



HAL
open science

Equilibres corrélés, jeux d'évolution et dynamique de populations.

Yannick Viossat

► **To cite this version:**

Yannick Viossat. Equilibres corrélés, jeux d'évolution et dynamique de populations.. Mathématiques [math]. Ecole polytechnique X, 2005. Français. NNT: . pastel-00002313

HAL Id: pastel-00002313

<https://pastel.hal.science/pastel-00002313>

Submitted on 29 Jul 2010

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Thèse de doctorat de l'École polytechnique
Domaine : mathématiques appliquées

Présentée par *Yannick Viossat*
sous la direction de *Sylvain Sorin*
pour obtenir le grade de Docteur de l'École polytechnique

**Equilibres corrélés, jeux d'évolution
et dynamique de populations**

Soutenue publiquement le 15 décembre 2005 à l'Institut Henri Poincaré, Paris

Devant le jury composé de :

Henri BERESTYCKI	Directeur d'études, EHESS	Examineur
Régis FERRIERE	Professeur, université Paris VI et ENS ULM	Examineur
Françoise FORGES	Professeur, université Paris Dauphine	Rapporteur
Josef HOFBAUER	Professeur, University College London	Rapporteur
Jacques ISTAS	Professeur, Grenoble II et Ecole polytechnique	Président
Sylvain SORIN	Professeur, Paris VI et Ecole polytechnique	Directeur

L'ECOLE POLYTECHNIQUE n'entend donner aucune approbation, ni improbation aux opinions émises dans les thèses ; ces opinions doivent être considérées comme propres à leurs auteurs.

A Andrew Bergel, champion du monde 2005 de Pierre-Papier-Ciseaux
(www.worldrps.com)

Remerciements

Je voudrais tout d'abord remercier profondément Sylvain Sorin. Nous sommes différents et nous ne nous sommes pas toujours compris, mais j'ai un grand respect pour lui, comme homme et comme scientifique. Il a subi d'innombrables manuscrits que les fautes de frappe ou les mélanges entre versions rendaient incompréhensibles, il m'a beaucoup appris et m'a beaucoup aidé, notamment lors des dernières années de ma thèse. Il m'a consacré de nombreuses heures et a toujours cherché à m'aider à trouver ma place dans le paysage scientifique. Merci.

Les travaux de Josef Hofbauer sont cités plus de cent fois dans ma thèse et je n'aurai jamais obtenu la plupart des résultats de la deuxième partie de ma thèse sans l'aide de son magnifique ouvrage, *Evolutionary Games and Population Dynamics*, co-écrit avec Karl Sigmund. Je ne pouvais rêver d'un meilleur rapporteur. Que certains de mes travaux ait pu l'intéresser est une grande source d'encouragement.

Françise Forges a jeté un regard bienveillant sur mes travaux à une époque où je doutais beaucoup et cela m'a soulagé et soutenu. Certaines des discussions que j'ai eu avec elle ont abouti directement au chapitre 4. Je lui suis particulièrement reconnaissant d'avoir accepté d'être un des rapporteurs de ma thèse, bien que je lui ai demandé extrêmement tardivement, et alors que son emploi du temps était déjà surchargé.

Je voudrais également remercier les autres membres du jury : Régis Ferrière, qui me suit depuis bien longtemps et dont le dynamisme et le sourire m'ont tant de fois redonné des forces, Jacques Istas, certainement un de mes professeurs les plus vivants et les plus sympathiques, et Henri Berestycki : tous trois très pris et qui m'ont fait la gentillesse de se libérer pour ma thèse.

Il y a d'autres professeurs et chercheurs que je voudrais saluer : mes premiers professeurs de mathématiques, Armand Jammot et Bertrand Renard, bien sûr, mais aussi Mme Mascarello, M. Bastianelli, et surtout Pierre Puchol, dont je n'oublie ni les cours ni qu'il nous traitait en adultes. Karl Sigmund a été mon premier professeur de jeux évolutionnistes, Richard Michod m'a convaincu que je pouvais peut-être faire quelque chose, however small, pour la recherche, et a co-écrit, avec Cristian Solari, Aurora Nedelcu et Mathilde Hurand, la troisième partie de cette thèse. Jean-Pierre Dupuy m'a montré ma première matrice de jeux. Jean Dalibard n'a rien à voir avec

la théorie des jeux mais était un professeur tellement exceptionnel. Pierre-Henri Gouyon, entre autres, m'a fait rencontrer Régis Ferrière. Bernhard von Stengel et Ehud Lehrer ont été parmi les premiers à s'intéresser à moi et m'ont toujours soutenu. Et puis Eilon Solan, Ken Binmore et Larry Samuelson, toute l'équipe du séminaire de théorie des jeux, et notamment les jeunes : Alexandre, Guillaume, qui m'ont bien aidé à Stony Brook; enfin les professeurs sous la responsabilité desquels j'ai enseigné, notamment Hervé Ledret, Joseph Oesterlé, Philippe Bich et Jérôme Renault.

Le laboratoire d'économétrie m'a accueilli pendant plus de quatre ans et je voudrais remercier son directeur, Jean-Pierre Ponsard, et tous ses membres. Une mention spéciale pour les théoriciens des jeux : Nicolas, Rida, Dinah (une thèse c'est dur, une thèse sans Dinah je ne veux pas savoir) ainsi qu'à Jean-François Laslier, Vincent Renard, Michel Balinski, et Damien, souvent providentiel.

J'ai passé l'essentiel de mon temps avec les doctorants : Thomas d'abord, avec qui j'ai partagé Sylvain, un bureau et des blagues débiles pendant l'essentiel de ma thèse, Alexia, Julien, Galit, amie chère, Erwann, qui m'a tout appris des cravates et des femmes, Jean-Pierre, Christophe, Thierry H. et Thierry L., Bilge, Jacques, François, Sylvie, Sami, Pierre, vieux camarade de prépa puis de DEA, Ignace, qui a souvent essayé d'aider Sudhir, Dip, aujourd'hui enterré sous deux tonnes de neige, Maximilien, Pascale, Blandine, Maïa, Tu et Marie-Anne.

J'ai croisé aussi des doctorants du CREA et du CRG, dont Sylvain, Corentin, Eric, et Lyse, aahhh Lyse... L'équipe d'X-Doc, les participants des Doctoriales et des Entreprenariales, dont Stéphane Fidanza, l'équipe de l'Ecole Doctorale, Fabrice, Audrey, Béatrice, Dominique, les participants à l'école d'été de Trento : Oxana, Martha, et les autres.

Un mot pour Chantal, Christine, Eliane, Lysa, Marie-Claude, qui s'est débattu pour me trouver une salle de soutenance, et Marie-Hélène, dont j'ai partagé le bureau et avec laquelle j'espère bien discuter encore de temps à autre.

A Tucson, outre Rick et Régis, bien sûr, Jeff, Cristian, Gabriella, Aurora, Denis, Bryan et Amanda ont rendu mon séjour fort agréable.

Depuis que je suis à Stockholm, Jörgen, Karolina, Chloé, Sven, Cédric, Marcus et Anneli ont bien pris soin de moi, et si la soutenance s'est bien passée je le dois en partie à Chloé et Cédric.

J'ai beaucoup douté pendant cette thèse, et souvent eu mal à d'autres

choses qu'à mes théorèmes. Beaucoup de gens m'ont aidé, beaucoup plus que je n'en peux citer. Jérôme m'accompagne depuis le DEA - nous nous suivons de si près qu'il a soutenu un jour avant moi, et avec un rapporteur commun. Il m'a fait connaître bon nombre de mes amis actuels : Alberto et Laurine, Julien, Zoheir, Charlotte, Cha, et Isabelle, Chafik, Laurence,... Ces années auraient été plus tristes sans les soirées avec Mathilde, Elie, Véro, Pacôme, Thomas, Jenlui, l'inconnu du CREA, Angela, Panu, Sanna, Yannick... Et puis Anne, ma plus ancienne amie, Laurence, qui a souvent senti mes mains sur ses paupières, Florence et Xavier, que je regrette de ne pas avoir vus davantage mais auxquels je pense quand ça ne va pas. Astrid a eu le bon goût de choisir un petit ami français, Stéphane, et de le choisir bien, ce qui m'a permis de les voir tous deux, et j'espère bien que nos vagabondages respectifs ne nous sépareront pas. Je n'oublie pas non plus qu'elle m'a fait connaître Susi, vers laquelle je me suis tourné aux moments les plus déchirants de ces années.

La Chorba : Khater, Houcine, Hacène, Momo, etc., ceux qu'on a perdu en route, les deux générations de Coup de Pouce : Philippe, François, Patrick, Stéphane, Etienne, Javier, Caro, Xavier, Amélie, Virginie, Adrien, Nicolas, Luc, Edouard, Isabelle, tant d'autres, les hommes et les femmes en galère que j'ai cotoyé, à Austerlitz, à Haxo, en maraude. J'avais besoin de vos sourires, de nos actions concrètes.

Sudhir a habité avec moi pendant une année particulièrement intense, puis chez mes parents, que je remercie profondément pour l'aide qu'ils nous ont apporté. Ce fut la grande aventure de ces années de thèse. Je ne peux pas ne pas repenser aussi à tous ceux que j'ai croisé au Népal, pendant les deux premiers mois de ma thèse, et notamment à la famille de Sudhir, sa femme, Shalini, seule pendant si longtemps, les gens de son village et de Janakpur.

Mes parents m'ont beaucoup aidé, en aidant Sudhir bien sûr, mais de bien d'autres façons. Ils m'ont toujours fait plus que bon accueil et ont adopté Romina. L'aide qu'ils m'ont apportée pour le pot n'est que la cerise sur le gâteau au chocolat. Cyrille est celui avec lequel j'ai grandi et les signes d'amitié qu'il m'a adressés depuis plusieurs années m'ont touché plus qu'il ne le sait. Raphaël, entre autres passements de jambes, a eu la gentillesse de me prévenir de ma calvitie. Par d'autres voies, Mouna a également beaucoup œuvré pour ralentir la dégradation de mon physique. Il y a aussi mes cousins les plus proches : Séverine, Jérémie, Armelle; mes tantes, Henri, qui sourit tout le temps, Inga et Loïc, la famille de Romina : Chiara, Andreina,

Franco, les cousins, tantes et oncles, qui m'ont accueilli. Gabriella et Giancarlo étaient même à mon pot de thèse. Et puis les amis de Romina : Chiara, Emi et bien d'autres. L'honnêteté m'oblige à avouer que la partie la plus originale de cette thèse a été écrite par Pistache, à petits coups de pattes sur le clavier.

Aline m'a accompagné au début de cette thèse et est aujourd'hui une amie, mère d'un joyeux Boubou dont les sourires me réjouissent.

Armelle m'a donné des moments de grâce, des moments que je n'espérais plus connaître. May she be showered with all the joy of Diwali.

Romina est une amie depuis mon arrivée au laboratoire d'économétrie. Nous vivons ensemble depuis plus de deux ans et j'espère pour longtemps. Elle a aidé Sudhir et m'a aidé. Elle a relu et même déposé la thèse, supporté mes indisponibilités, et pas seulement mes indisponibilités. Elle mérite plus que mes remerciements.

Ma grand-mère maternelle est morte au cours de ces années. Cette thèse est pour elle.

Table des matières / Contents

Présentation de la thèse	1
General Introduction	13
I Dual Reduction and Correlated Equilibria	15
1 Introduction to Part I	17
1.1 Definition and canonical representation	17
1.2 Why study correlated equilibria?	21
1.3 Contribution of this dissertation	22
2 Dual Reduction	25
2.1 Introduction	25
2.2 Notations and definitions	26
2.3 Existence of correlated equilibria	28
2.4 Dual reduction: definition and known results	29
2.5 Applications of dual reduction	35
2.6 Some results used in chapter 6	39
2.7 General properties of dual reduction	43
2.8 Properties in specific classes of games	50
2.9 The issue of uniqueness	55
2.10 Dual reduction as a refinement concept	59
2.11 Appendix	65
3 Games With a Unique Equilibrium	69
3.1 Definitions and main result	70
3.2 Proof	71
3.3 Remarks	75

4	Correlated Eq. in Zero-Sum Games	81
5	Elementary Games and Full Games	85
5.1	Introduction	85
5.2	Notations and definitions	86
5.3	Relation between elementary and full games	88
5.4	Characterizations	90
5.5	First appendix	92
5.6	Second appendix	95
6	Geometry, CE and Zero-Sum Games	97
6.1	Introduction	97
6.2	Notations	99
6.3	Definitions of tight and pre-tight games	100
6.4	Characterization of tight and pre-tight games	102
6.5	Links between tight and pre-tight games	105
6.6	Topology of tight and pre-tight games	107
6.7	Geometry of Nash and correlated equilibria	108
6.8	Two-player pre-tight games	110
II Evolutionary Dynamics and Strategic Concepts		117
7	Introduction to Part II	119
7.1	Basics of evolutionary game theory	119
7.2	Research topics in evolutionary game theory	125
7.3	Contribution of this dissertation	128
8	Elimination of Dominated Strategies	133
8.1	Introduction	133
8.2	Framework and classes of dynamics	134
8.3	Positive results	136
8.4	Negative results	142
8.5	Discrete-time dynamics	145
9	Dynamics and Equilibria: Low Dimension	149

Part A: Nash Equilibrium	151
9.1 Introduction	151
9.2 Replicator dynamics	152
9.3 Best-response dynamics	154
Part B: Correlated Equilibrium	157
9.4 Notations, definitions and main result	157
9.5 Proof	158
9.6 Extensions and comments	162
10 Elimination of Correlated Equilibria	165
Part A: Replicator dynamics	167
10.1 Notations and basic definitions	167
10.2 A reminder on Rock-Paper-Scissors	168
10.3 A family of 4×4 games	171
10.4 Decomposition of the dynamics	173
10.5 Main results	174
10.6 Extensions and variants	176
10.7 Discussion	180
Part B: More General Games and Other Dynamics	183
10.8 Replicator dynamics	183
10.9 Best-response dynamics	185
10.10 Brown-von Neumann-Nash dynamics	191
10.11 Monotonic and WSP Dynamics	193
10.12 Robustness to Addition of Mixed Strategies	197
10.13 First appendix	202
10.14 Second appendix	204
11 Elimination of Nash Equilibria	207
11.1 Introduction	207
11.2 Results	208
11.3 Proof	210
11.4 Other examples	213
11.5 Discussion	216
11.6 Appendix	216

12 Elimination for All Initial Conditions	219
12.1 Introduction	219
12.2 Best-response dynamics	221
12.3 Replicator dynamics	228
12.4 Discussion	243
III Origins of Multicellularity	245
13 Introduction to part III	247
14 Life-history Evolution and the Origin of Multicellularity	251
Bibliography of Part I and Part II	273

Présentation de la thèse

Les deux premières parties de la thèse se rattachent à la théorie des jeux et la troisième partie à la biologie théorique. La première partie étudie la notion d'équilibre corrélé : un des concepts centraux d'équilibre en théorie des jeux (Aumann, 1974, 1987); la deuxième partie relève de la théorie des jeux évolutionnistes et étudie le lien entre l'issue de dynamiques d'évolution et les concepts stratégiques statiques; la troisième partie, co-écrite, étudie les déterminants de la séparation entre lignée germinale et lignée somatique chez les algues vertes volvocales.

Notions de théorie des jeux. La théorie des jeux cherche à modéliser et analyser les situations d'interaction stratégique. Elle a de nombreuses applications, en sciences sociales, en écologie et en biologie de l'évolution, ainsi qu'en informatique théorique.

Jeu. Un jeu est donné par un triplet

$$G = \{I, (S_i)_{i \in I}, (U_i)_{i \in I}\}$$

où I est l'ensemble des *joueurs* (les agents qui interagissent), S_i l'ensemble des *stratégies pures* du joueur $i \in I$ (ce qu'il peut faire) et $U_i : S = \times_{i \in I} S_i \rightarrow \mathbb{R}$ la fonction de *paiement* du joueur i . Cette dernière associe à chaque *profil de stratégies* pure $s = (s_i)_{i \in I} \in S$ un paiement, qui est une mesure du gain ou de la perte que le joueur i tire de l'interaction. Dans le cadre de la thèse, l'ensemble des joueurs et les ensembles de stratégies pures de chacun des joueurs sont supposés finis.

Les agents peuvent avoir intérêt à agir de manière aléatoire, c'est à dire à choisir une probabilité sur leur ensemble de stratégies pure, appelé *stratégie mixte*. Lorsque les joueurs utilisent des stratégies mixtes, les fonctions de paiements sont étendues multi-linéairement. Enfin, la description du jeu peut être complété par l'introduction d'une structure d'information.

Équilibres. Un *équilibre de Nash* d'un jeu est un profil de stratégies mixtes tel que la stratégie de chaque joueur maximise son paiement, étant donné son information et la stratégie des autres joueurs. N'importe quel jeu fini à au moins un équilibre de Nash.

Un *équilibre corrélé* du jeu G est un équilibre de Nash d'un jeu où, avant de jouer le jeu G , les joueurs observent de manière privée des réalisations d'événement aléatoires sur lesquels ils peuvent conditionner leurs actions. Les observations des joueurs pouvant être corrélées entre elles, les joueurs peuvent ainsi corréliser leurs actions sans communication directe. Une *distribution d'équilibre corrélé* est une distribution de probabilité sur l'ensemble des profils de stratégies pures du jeu G induite par un équilibre corrélé. On peut montrer que l'ensemble des distributions d'équilibre corrélés d'un jeu est un polytope qui contient l'enveloppe convexe des équilibres de Nash.

Partie I : réduction duale et équilibre corrélé

La première partie de la thèse étudie les propriétés des équilibres corrélés ainsi que des techniques et classes de jeux liées. Par souci de concision, *on appellera par la suite "équilibres corrélés" les distributions d'équilibre corrélé.*

Chapitre 1 : Why Study Correlated Equilibria?

Cette introduction à la partie I rappelle quelques définitions et résultats fondamentaux sur les équilibres corrélés, explique l'intérêt de ce concept d'équilibre, et présente les contributions de la première partie de la thèse.

Chapitre 2 : Properties and Application of Dual Reduction

Cet article étudie la technique de réduction duale (Myerson, 1997) et en montre l'utilité pour l'étude des équilibres corrélés et des équilibres de Nash.

La technique de réduction duale permet de réduire un jeu en un jeu plus petit, dont les stratégies pures correspondent à des stratégies mixtes du jeu initial, invariantes sous certaines chaînes de Markov¹. Myerson (1997) a

¹Ces chaînes de Markov sont engendrées par les solutions du dual d'un programme linéaire dont les solutions sont les équilibres corrélés.

montré que tout équilibre corrélé du jeu réduit² induit un équilibre corrélé du jeu initial, et a caractérisé la classe des jeux qui ne peuvent être réduits strictement par réduction duale. Les autres propriétés de la réduction duale étaient pour l'essentiel inconnues.

L'article étudie ces propriétés. Il montre, par exemple, que le processus de réduction est invariant sous toute transformation affine positive des paiements, que les stratégies qui n'appartiennent au support d'aucun équilibre corrélé sont toujours éliminées, et que, pour les jeux à deux joueurs génériques, le processus de réduction est défini de manière unique. Il étudie également les propriétés de la réduction duale dans des classes particulières de jeux (jeux à somme nulle, symétriques, avec un seul équilibre, etc.) ainsi que le lien entre la réduction duale et des concepts de raffinement d'équilibre corrélé (Myerson, 1986; Dhillon et Mertens, 1996).

L'article s'attache d'autre part à montrer l'utilité et la puissance de la réduction duale. La réduction duale permet notamment de prouver qu'un équilibre corrélé unique est un équilibre de Nash quasi-strict³; l'existence dans certaines classes de jeux d'un équilibre de Nash à support plein; ou encore, sous des hypothèses supplémentaires, d'obtenir des informations sur la dimensionnalité du polytope des équilibres corrélés.

Chapitre 3 : Openness of the Set of Games with a Unique Correlated Equilibrium

Cette note montre notamment que l'ensemble des jeux ayant un unique équilibre corrélé est ouvert et que l'ensemble des jeux ayant un unique équilibre de Nash ne l'est pas.

Jansen (1981) a montré que l'ensemble des jeux à deux joueurs ayant un unique équilibre de Nash est ouvert. Il n'était pas sûr si ce résultat s'étendait aux jeux à plus de deux joueurs ou à des variantes de l'équilibre de Nash, comme les équilibres symétriques. De plus, pour la notion d'équilibre corrélé, aucun résultat n'avait été obtenu.

Le résultat principal de cette note est que, pour n'importe quel nombre

²En fait, un jeu peut admettre plusieurs réductions, et il faudrait donc dire : tout équilibre corrélé *de n'importe quel* jeu réduit...

³Un équilibre est quasi-strict si, pour tout joueur $i \in I$, le support de la stratégie mixte du joueur i est égal à l'ensemble des stratégies pures qui maximisent son paiement (face à la stratégie des autres joueurs).

de joueurs, l'ensemble des jeux ayant un unique équilibre corrélé est ouvert. De plus, un contre-exemple montre que pour les jeux à trois joueurs ou plus, l'ensemble des jeux avec un unique équilibre de Nash n'est pas ouvert.⁴

Il est montré également que, dans l'ensemble des jeux symétriques à deux joueurs, l'ensemble des jeux avec un unique équilibre de Nash symétrique⁵ n'est pas ouvert, mais que l'ensemble des jeux avec un unique équilibre de Nash symétrique et qui est quasi-strict est ouvert.

Les preuves utilisent des techniques de dualité linéaire, un résultat sur les polytopes, et la semi-continuité supérieure de la correspondance des équilibres.

Chapitre 4 : Correlated Equilibrium in Generic Zero-Sum Games

Cette très brève note montre que, pour les jeux à deux joueurs et à somme nulle génériques, tous les équilibres corrélés sont des équilibres de Nash et que, pour une notion plus forte de généricité, il y a un unique équilibre corrélé. La preuve apparaît à posteriori comme la réunion d'arguments de Forges (1990) et de Bonhenblust et al (1950), les résultats obtenus étant toutefois plus précis.

Chapitre 5 : Elementary Games and Games whose Correlated Equilibrium has Full Dimension

Cette note relie et caractérise deux classes de jeux. Un jeu est élémentaire (Myerson, 1997) s'il ne peut être réduit strictement par réduction duale. Un jeu est plein si le polytope des équilibres corrélés est de dimension maximale ($N-1$ où N est le nombre de profils de stratégies pures du jeu). Les équilibres de Nash se trouvent alors sur la frontière de ce polytope (Nau et al, 2004).

Il est montré qu'un jeu élémentaire est plein si et seulement si aucune des contraintes définissant les équilibres corrélés n'est vide. De plus des caractérisations des jeux élémentaires et des jeux pleins sont obtenues, à l'aide notamment de jeux auxiliaires similaires à celui utilisé par Hart et Schmeidler (1989) dans leur preuve d'existence des équilibres corrélés.

⁴En outre, j'ai montré récemment que pour tout entier $k \geq 2$ et n'importe quel nombre de joueurs, l'ensemble des jeux avec k équilibres n'est pas ouvert.

⁵Il peut y avoir plusieurs équilibres, mais un seul est symétrique.

Chapitre 6 : Geometry, Correlated Equilibrium and Zero-Sum Games

Cet article caractérise la classe des jeux dont le polytope des équilibres corrélés contient un équilibre de Nash dans son intérieur relatif (lorsque celui-ci est non vide) et montre que cette classe de jeux contient et généralise celle des jeux à somme nulle.

Comme expliqué plus haut, le polytope des équilibres corrélés contient l'enveloppe convexe des équilibres de Nash. Comprendre où sont situés les équilibres de Nash dans le polytope des équilibres corrélés permet de préciser la relation entre ces deux notions d'équilibre et pourrait déboucher sur de meilleurs algorithmes de calcul des équilibres de Nash (problème NP-dur même pour les jeux à deux joueurs; voir (Gilboa et Zemel, 1989)).

Comme indiqué précédemment, Nau et al (2004) ont montré que si le polytope des équilibres corrélés est de dimension maximale, alors les équilibres de Nash se trouvent sur la frontière de ce polytope. En revanche, si ce polytope n'est pas de dimension maximale, des équilibres de Nash peuvent se trouver dans son intérieur relatif. Nau et al (2004) donnent des conditions nécessaires pour que cela se produise, mais pas de conditions suffisantes.

Grâce à l'utilisation de la réduction duale, l'article montre que l'une des conditions nécessaires de Nau et al (2004) est en fait suffisante (le fait que, dans n'importe quel équilibre corrélé, un sous-ensemble des inégalités définissant les équilibres corrélés soient satisfaites avec égalité). De plus, il montre que la classe des jeux satisfaisant cette condition est conceptuellement importante. En particulier, dans le cas des jeux à deux joueurs, elle contient la classe des jeux à somme nulle, et la généralise, au sens où de nombreuses propriétés des jeux à somme nulle y sont vérifiées (les équilibres de Nash sont échangeables, les paiements d'équilibre corrélé coïncident avec les paiements d'équilibre de Nash, les marginales de tout équilibre corrélé induisent un équilibre de Nash, etc.)

Les preuves utilisent des propriétés de complémentarité forte ainsi que des arguments (élémentaires) d'analyse convexe.

Partie II : jeux d'évolution

Domaine de recherche

La deuxième partie de la thèse se rattache à la théorie des jeux évolutionniste (Maynard Smith, 1982; Hofbauer et Sigmund, 1998). Cette branche de la théorie des jeux analyse l'évolution de la fréquence des comportements au sein de populations d'agents interagissant stratégiquement et dotés d'une rationalité nulle ou limitée. En biologie, le paradigme le plus simple est celui d'interactions au sein d'une population infinie, entre individus dont les comportements sont codés dans les gènes, qui se reproduisent clonalement, et d'autant plus, en moyenne, que leur comportement obtient de bons paiements.

L'évolution de la fréquence des comportements dans la population est modélisée par une équation aux différences ou une équation différentielle comme la dynamique des répliqueurs:⁶

$$\dot{x}_i(t) = x_i(t) [\mathbf{A}\mathbf{x}(t) - \mathbf{x}(t) \cdot \mathbf{A}\mathbf{x}(t)]$$

Ici, $x_i(t)$ est la fréquence de la stratégie $i \in \{1, \dots, N\}$ à l'instant t , $\mathbf{x}(t) = (x_1(t), \dots, x_N(t))$ le vecteur des fréquences des stratégies (qu'on peut voir aussi comme la stratégie moyenne au sein de la population), et $\mathbf{A} = (a_{ij})_{1 \leq i, j \leq N}$ la matrice des paiements. I.e., a_{ij} est le paiement d'un individu jouant la stratégie i face à un individu jouant la stratégie j . Sous la dynamique des répliqueurs, le taux de croissance d'une stratégie est donc égal à la différence entre le paiement qu'elle obtient et le paiement moyen. La variable d'état est la stratégie moyenne, $\mathbf{x}(t)$, qui évolue dans un simplexe de dimension $N - 1$.

D'autres dynamiques sont très étudiées, comme la dynamique de meilleure réponse, d'inspiration économique :

$$\dot{\mathbf{x}}(t) \in BR(\mathbf{x}(t)) - \mathbf{x}(t)$$

Ici, $BR(\mathbf{x}(t))$ est l'ensemble des stratégie mixtes qui sont des meilleures réponses à la stratégie moyenne courante $\mathbf{x}(t)$. Notons que la dynamique

⁶Les dynamiques présentées ci-dessous décrivent l'évolution des comportements au sein d'une seule population, dont les membres interagissent. Pour représenter la co-évolution de plusieurs populations, on fait appel à des dynamiques analogues, mais où l'évolution de la fréquence des comportements dans une populations dépend de la fréquence des comportements courante dans les autres populations. Les chapitres 10 et 12 sont consacrés à l'étude de dynamiques à une population et les chapitres 8, 9 et 11 à l'étude de dynamiques à une ou plusieurs populations.

de meilleure réponse n'est pas donnée par une équation différentielle, mais par une inclusion différentielle : $\dot{\mathbf{x}}$ peut prendre plusieurs valeurs.

Une des questions centrales, dans la littérature comme dans la thèse, est de préciser le lien entre l'issue de ce type de dynamiques et les concepts stratégiques statiques, comme l'équilibre de Nash, l'équilibre corrélé, les stratégies évolutivement stables (un raffinement de l'équilibre de Nash due à Maynard-Smith et Price (1973)), ou l'élimination des stratégies dominées.⁷

Les dynamiques d'évolution ne convergent pas toujours (Hofbauer et Sigmund, 1998, section 8.6). Toutefois, cela ne signifie pas forcément qu'il n'y ait pas de lien entre l'issue de ces dynamiques et les équilibres. Ainsi, sous la dynamique des réplicateurs à une ou deux populations, il suffit qu'il existe un seuil $\epsilon > 0$ tel que toutes les stratégies soient toujours présentes en proportion au moins ϵ pour que la solution converge en moyenne temporelle vers l'ensemble des équilibres de Nash - même si la solution elle-même ne converge pas vers les équilibres (Hofbauer et Sigmund, 1998, theorem 7.6.4).

De plus, de nombreux résultats récents montrent qu'une classe de processus adaptatifs, appelés dynamiques de non-regret, convergent en moyenne temporelle vers l'ensemble des équilibres corrélés (voir Hart, 2005, pour une revue). Ces processus sont distincts des dynamiques d'évolution classiques, mais amènent à se demander s'il existe un lien général entre les dynamiques d'évolution et les équilibres corrélés.

L'apport majeur de la deuxième partie de la thèse est de montrer que ce n'est pas le cas, sauf en petite dimension : pour de grandes classes de dynamiques, toutes les stratégies appartenant au support d'au moins un équilibre corrélé peuvent être éliminées, si bien que seules les stratégies n'appartenant au support d'aucun équilibre survivent (voir le résumé du chapitre 10 ci-dessous). Pour les équilibres de Nash, les résultats sont encore plus forts : l'élimination de toutes les stratégies appartenant au support des équilibres de Nash peut se produire sous n'importe quelle dynamique d'adaptation myope régulière⁸ (chapitre 11) et, sous la dynamique des réplicateurs et la dynamique de meilleure réponse, à partir de presque toutes les condi-

⁷Une stratégie est (strictement) dominée s'il existe une stratégie qui obtient toujours (strictement) de meilleurs paiements, quelque soit le comportement de l'adversaire.

⁸Les dynamiques d'adaptation myopes sont celles dont le champ de vecteurs vérifie $\dot{\mathbf{x}} \cdot \mathbf{Ax} > 0$ pour tout \mathbf{x} qui n'est ni un équilibre ni un équilibre d'un des "sous-jeux" obtenus par élimination de certaines stratégies. Cela signifie intuitivement que la population évolue vers de meilleures réponses à la stratégie courante : une condition minimale d'adaptativité.

tions initiales (chapitre 12).

Contributions de la deuxième partie de la thèse

Chapitre 7 : Evolutionary Dynamics and Strategic Concepts

Cette introduction à la deuxième partie de la thèse donne un aperçu des questions, modèles et résultats essentiels de la théorie des jeux évolutionnistes, et présente les contributions de la thèse.

Chapitre 8 : Evolutionary Dynamics and Dominated Strategies

Cet article est essentiellement une revue de la littérature sur l'élimination - ou non - des stratégies dominées par les dynamiques d'évolution. Il traite tout d'abord des résultats positifs et négatifs concernant les dynamiques continues, puis de la manière dont ces résultats se transposent - ou non - pour les dynamiques discrètes. Quelques résultats nouveaux sont obtenus, comme le fait que, parmi une classe de dynamiques monotones, seules les dynamiques "concaves" éliminent les stratégies mixtes dominées par une stratégie pure.

Chapitre 9, partie A : Replicator Dynamics and Nash Equilibrium in Low Dimension

Cette note montre que, dans les jeux symétriques à trois stratégies, la dynamique des réplicateurs à une population élimine toutes les stratégies pures qui n'appartiennent au support d'aucun équilibre de Nash. Le même résultat est obtenu pour la dynamique de meilleure réponse. Les preuves exploitent le fait que les systèmes dynamiques correspondant évoluent dans une région compacte du plan.

Chapitre 9, partie B : Replicator Dynamics and Correlated Equilibrium

Cette note étudie toujours les jeux symétriques à trois stratégies, mais cette fois pour des dynamiques à deux populations, si bien que l'espace d'état a

maintenant dimension 4. Il est montré que la dynamique des réplicateurs, la dynamique de meilleure réponse, ainsi que n'importe quelle dynamique monotone convexe (Hofbauer et Weibull, 1996) élimine toutes les stratégies qui n'appartiennent au support d'aucun équilibre corrélé. La preuve repose sur la caractérisation des profils de stratégies non jouées en équilibre corrélé de Nau et McCardle (1990) et sur des arguments de réduction duale.

Chapitre 10, partie A : Replicator Dynamics do not Lead to Correlated Equilibrium

Cet article montre que pour la dynamique des réplicateurs à une population, il existe un ensemble ouvert de jeux tel que, pour un ensemble ouvert de conditions initiales, toutes les stratégies qui appartiennent au support des équilibres corrélés sont éliminées. Les mêmes résultats sont établis pour des versions perturbées ou discrétisées de la dynamique des réplicateurs. La preuve repose sur une méthode de construction de jeux ayant un unique équilibre, un principe de décomposition de la dynamique des réplicateurs dans les jeux considérés, qui permet d'analyser très précisément cette dynamique, et sur la construction d'une fonction de Lyapunov.

Chapitre 10, partie B : Evolutionary Dynamics do not Lead to Correlated Equilibrium

Cet article montre que les résultats obtenus dans la note précédente sont robustes, au sens où ils sont valables pour un grand nombre de dynamiques. En particulier, pour les dynamiques de meilleure réponse, de Brown-von Neumann-Nash, et pour n'importe quelle dynamique monotone⁹ qui dépend continûment des paiements, il existe un ensemble ouvert de jeux pour lesquels, pour un ensemble ouvert de conditions initiales, toutes les stratégies appartenant au support d'au moins un équilibre corrélé sont éliminées. De plus, pour la dynamique des réplicateurs et la dynamique de meilleure réponse, les résultats sont robustes à l'addition de stratégies mixtes comme nouvelles stratégies pures des jeux considérés.

⁹Une dynamique est monotone si le taux de croissance au temps t d'une stratégie est (strictement) plus grand que celui d'une autre stratégie si et seulement si son paiement au temps t est (strictement) plus élevé.

Les preuves utilisent notamment des fonctions de Lyapunov et les conditions de stabilité des cycles hétérocliniques sur le simplexe dues à Hofbauer (1994) et Brannath (1994).

Chapitre 11 : Elimination of All Strategies in the Support of Nash Equilibrium: a Universal Example

Cette note généralise les résultats obtenus dans l'article précédent à la (très grande) classe des dynamiques d'adaptation myope (Swinkels, 1993) et aux cas des dynamiques à plusieurs populations, mais pour la notion d'équilibre de Nash et non d'équilibre corrélé. Il est montré que pour n'importe quelle dynamique d'adaptation myope dépendant continûment des paiements, et à 1, 2 ou 3 populations, il existe un ensemble ouvert de jeux tels que, pour un ensemble ouvert de conditions initiales, toutes les stratégies jouées en équilibre de Nash sont éliminées. Ceci bien que les jeux considérés, comme ceux considérés au chapitre 10, aient un seul équilibre de Nash, qui est strict, et est donc évolutivement stable. La preuve repose notamment sur une version du théorème fondamental de la sélection naturelle de Fisher (1930).

Chapitre 12. Elimination of All Strategies in the Support of Nash Equilibrium from Almost All Initial Conditions

Cet article montre que, sous la dynamique des réplicateurs et la dynamique de meilleure réponse, toutes les stratégies appartenant au support des équilibres de Nash peuvent être éliminées non seulement pour un ensemble ouvert de conditions initiales, mais pour *presque toutes* les conditions initiales. Ceci pour un ensemble ouvert de jeux, tout du moins pour la dynamique de meilleure réponse. Les systèmes dynamiques considérés sont de dimension 5 ou 6, mais ils peuvent être précisément analysés grâce au principe d'amélioration ("Improvement Principle") de Monderer et Sela (1997), pour la dynamique de meilleure réponse, et à une décomposition de la dynamique, pour la dynamique des réplicateurs. Les preuves utilisent également des fonctions de Lyapunov et l'analyse du comportement de la moyenne temporelle de solutions convergeant vers un cycle hétéroclinique (Gaunersdorfer, 1992).

Partie III: différenciation germ-soma

La troisième partie de la thèse relève de la biologie théorique et comprend deux chapitres, comprend une unique contribution :

Chapitre 13 : Transition from Unicellular to Multicellular-Organisms and Gem-Soma Differentiation

Ce chapitre introduit et résume, à l'intention de non-biologistes, les questions abordées et les contributions apportées au chapitre 14.

14. Life-History Evolution and the Origin of Multicellularity, par R.E. Michod, Y. Viossat, C.A. Solari, M. Hurand et A. Nedelcu, à paraître dans Journal of Theoretical Biology

Cet article étudie les déterminants de la différenciation cellulaire, et plus particulièrement de la séparation entre lignée germinale et lignée somatique chez les algues vertes volvocales. Cette famille d'algues comprend de nombreuses espèces, qui varient en taille (de 1 à 50.000 cellules) et en degré de différenciation (absence de différenciation pour les colonies de petite taille, présence de cellules somatiques pour les colonies de taille intermédiaire, différenciation complète germ-soma pour les colonies de grande taille).

Nous proposons des modèles théoriques fondés sur le trade-off entre contribution à la fécondité et contribution à la viabilité de la colonie auxquels font face les cellules individuelles. Nous montrons que la différenciation germ-soma est favorisée si ce trade-off est convexe ou s'il existe un coût fixe à la reproduction (i.e. si les cellules qui se reproduisent, même rarement, contribuent nettement moins à la viabilité de la colonie que les cellules somatiques).

L'analogie d'un coût fixe de reproduction existe chez les algues vertes volvocales, à cause notamment d'un mode particulier de reproduction (palintomie) et de phénomènes hydrodynamiques. Compte tenu du fait que ce coût croît avec la taille de la colonie, les prédictions de nos modèles correspondent aux données biologiques rappelées ci-dessus (premier paragraphe).

General Introduction

This dissertation is divided in three parts. The first part groups contributions to the study of correlated equilibria. We focus on the properties and applications of the dual reduction (Myerson, 1997) and the geometry of Nash equilibria and correlated equilibria. The second part deals with evolutionary dynamics. We investigate the link between strategies belonging to the support of Nash or correlated equilibria and strategies surviving in the long-run. We find that many dynamics, including the replicator and best-response dynamics may eliminate all strategies in the support of correlated equilibria. Elimination of all strategies in the support of Nash equilibria is found to be even more universal, and may occur from almost all initial conditions. The third part consists of a single co-written article, which belongs to the field of theoretical biology. We study aspects of the transition from unicellular to multi-cellular organisms, in particular factors driving germ-soma specialization in volvocine green algae.

Lengthier introductions are given at the beginning of each part. The bibliography of part I and part II is disjoint from the bibliography of part III, and is given at the end of part II. Though interconnected, the chapters are essentially self-contained. In particular, the notations and some definitions are recalled each time. This accounts for some repetitions.

Part I

**Dual Reduction and Correlated
Equilibria**

Chapitre 1

Introduction to Part I: Why Study Correlated Equilibria?

We first recall the definition of correlated equilibria and the canonical representation theorem (any correlated equilibrium distribution is a canonical correlated equilibrium distribution). We then motivate the study of correlated equilibria (section 1.2) and present the contribution of this dissertation (section 1.3).

1.1 Definition and canonical representation

Correlated equilibrium. A correlated equilibrium (Aumann, 1974) is a Nash equilibrium of a game in which players can condition their action on payoff-irrelevant signals received before play. These signals may be correlated across players. Thus, even though players choose their action independently, they may correlate their actions through the signals. Formally, let

$$G = \{I, (S_i)_{i \in I}, (U_i)_{i \in I}\}$$

be a finite game. I is the set of players, S_i the set of pure strategies (or actions) of player i and $U_i : \times_{i \in I} S_i \rightarrow \mathbb{R}$ the payoff function of player i . For every i in I , let M_i be a nonempty finite set and let $M := \times_{i \in I} M_i$. Let ρ be a probability distribution over M . Consider the game Γ played as follows: first a profile $m = (m_i)_{i \in I} \in M$ is drawn at random according to the probability distribution ρ and player i privately observes m_i ; then G is played. That is, each player i chooses an action s_i in S_i and receives the payoff $U_i(s)$,

where $s = (s_i)_{i \in I}$. We say that Γ is a game based on G and extended by the *correlation device* (M, ρ) .

Definition. A correlated equilibrium of G is a Nash equilibrium of a game based on G and extended by a correlation device.

For any finite set Σ , let $\Delta(\Sigma)$ denote the set of probability distribution over Σ . A (behavioral) strategy of player i in the extended game Γ is a mapping

$$\begin{aligned} \phi_i : M_i &\rightarrow \Delta(S_i) \\ m_i &\rightarrow \phi_i(\cdot | m_i) \end{aligned}$$

from the set of messages that player i can receive to the set of mixed strategies of player i in G . In particular, Γ is a finite game. It follows that Γ has at least a Nash equilibrium hence G has at least a correlated equilibrium.¹ Direct proofs of existence of correlated equilibria (i.e. proofs that do not use the existence of Nash equilibria nor a fixed point theorem) have been given by Hart and Schmeidler (1989) and Nau and McCardle (1990). See chapter 2.

Let $S := \times_{i \in I} S_i$. A strategy profile $\phi = (\phi_i)_{i \in I}$ of a game extended by the correlation device (M, ρ) induces a probability distribution $\mu \in \Delta(S)$ given by:

$$\forall s \in S, \mu(s) = \sum_{m \in M} \rho(m) \prod_{i \in I} \phi_i(s_i | m_i)$$

Definition. A correlated equilibrium distribution is a probability distribution over the set S of pure strategy profiles of G induced by a correlated equilibrium.

Canonical representation. Assimilate a correlated equilibrium with a triplet (M, ρ, ϕ) , where $M = \prod_i M_i$ and $\rho \in \Delta(M)$ define the extended game Γ and $\phi = (\phi_i)_{i \in I}$ is a Nash equilibrium of Γ .

Definition. A correlated equilibrium (M, ρ, ϕ) is *canonical* if for every player i , $M_i = S_i$ and ϕ_i is the identity mapping².

¹Of course, existence of correlated equilibria follows also from existence of Nash equilibria in G ; see the end of section 1.1.

²That is, $\phi_i(s_i | s_i) = 1$ for every i in I and every s_i in S_i .

Definition. A canonical correlated equilibrium distribution is a correlated equilibrium distribution induced by a canonical correlated equilibrium.

The importance of canonical correlated equilibria stems from the following *canonical representation theorem*:

Theorem. Any correlated equilibrium distribution is a canonical correlated equilibrium distribution.

The proof is an instance of the revelation principle (see, e.g., (Myerson, 1994)).

Aumann's (1974) original approach to correlated equilibria is descriptive. The issue is to clarify the type of behaviour compatible with the assumption that the agents are rational (and with or without additional assumptions, such that the assumption that the agents have a common prior on the states of the world; see Aumann (1974, 1987)). The signals m_i received by the agents are then interpreted as realizations of any kind of random events on which the agents could condition their behaviour (e.g., the fact that it rained the day before or not). Another approach is in term of mechanism-design. The focus is now on the behaviour that rational agents may be induced to have, and in particular on the improvement of equilibrium payoffs which may result from the intervention of a referee (other correlation mechanism are also of interest, see, e.g., Forges (1986)). The signals are then sent by this referee, and the canonical representation theorem then means that the referee may restrict himself to send *recommendations* of pure strategies without loss of generality.

More precisely, let $\mu \in \Delta(S)$. Assume that a mediator draws at random a strategy profile s in S according to μ and then privately recommends the pure strategy s_i (the i^{th} component of s) to player i . In light of the canonical representation theorem, μ is a correlated equilibrium distribution if and only if, assuming that all the other players obey the mediator's recommendation, player i has no incentive to "deviate" from the mediator's recommendation. This is equivalent to:

$$\forall i \in I, \forall s_i, t_i \in S_i, \sum_{s_{-i} \in S_{-i}} \mu(s_i, s_{-i}) [U_i(t_i, s_{-i}) - U_i(s_i, s_{-i})] \leq 0 \quad (1.1.1)$$

where $S_{-i} := \prod_{j \neq i} S_j$. Note that these inequalities, which we call *incentive constraints*, are linear in μ . Since the conditions defining probability

distributions over S :

$$\forall s \in S, \mu(s) \geq 0$$

$$\sum_{s \in S} \mu(s) = 1$$

are also linear, it follows that the set of correlated equilibrium distributions is a convex polytope in $\Delta(S)$. In sharp contrast, the set of Nash equilibria of a finite game may be disconnected and its connected components need not be convex.

Nash Equilibria as Correlated Equilibria: we assimilate throughout $\times_{i \in I} \Delta(S_i)$ to a subset of $\Delta(S) = \Delta(\times_{i \in I} S_i)$. That is, we identify a mixed strategy profile σ in $\times_{i \in I} \Delta(S_i)$ with the product distribution in $\Delta(S)$, denoted also σ , and given by:

$$\forall s \in S, \sigma(s) = \prod_{i \in I} \sigma_i(s_i)$$

Nash equilibria correspond exactly to the correlated equilibrium distributions with a product distribution.

Abuse of vocabulary: in the remainder of this dissertation, *we often write correlated equilibrium for correlated equilibrium distribution*. This is lighter and without ambiguity since we focused throughout on correlated equilibrium distributions and never consider actual correlated equilibria (Nash equilibria of extended games) unless explicitly mentioned otherwise. In particular, the difference between correlated equilibria and canonical correlated equilibria will not appear.

However, this difference is conceptually important. For instance, Dhillon and Mertens (1996) call a correlated equilibrium (M, ρ, ϕ) *perfect* if ϕ is a perfect Nash equilibrium of the game extended by the correlation device (M, ρ) , and *perfect direct* if furthermore, for all i , $M_i = S_i$ and ϕ_i is the identity. They show that for this notion, the revelation principle fails. That is, the set of perfect correlated equilibrium distributions is in general larger than the set of perfect direct correlated equilibrium distributions. See also Chatterji and Govindan (2006, to appear).

1.2 Why study correlated equilibria?

In many situations it is difficult to exclude that agents condition their behaviours on (essentially) payoff-irrelevant signals. Furthermore it may be difficult to determine the signals available to each agent (in modelling terms, it may be difficult to determine the extended game being played). In such situations, correlated equilibrium seems a more natural solution concept than Nash equilibrium. This is the original and most straightforward argument to motivate the study of correlated equilibria. There is more however:

1. Aumann (1987) shows that correlated equilibrium is the natural outcome of common knowledge of Bayesian rationality.
2. Nau and McCardle (1990) show that strategy profiles that have zero probability in all correlated equilibrium distributions are precisely those that expose the players as a group to arbitrage from an outside observer. See chapter 2, proposition 2.5.5 and section 2.11.2 for the mathematical expression of this characterization.
3. The set of correlated equilibrium distributions is mathematically simpler than the set of Nash equilibria. While this does not imply that correlated equilibrium is a more interesting concept than Nash equilibrium, this is nonetheless an argument to study correlated equilibria, as whenever two concepts are as interesting, we should first investigate the simpler one.
4. Correlated equilibria are better related than Nash equilibria to the long-run behavior generated by classes of learning processes such as no-regret procedures. See (Hart, 2005) for a recent survey.
5. In ecology and evolutionary biology, the correlated equilibrium concept proved useful to interpret the outcome of asymmetric conflicts.³ Similarly, the correlated equilibrium concept is useful to interpret the outcome of games with local interactions. See Mailath et al (1997).

³For instance, in an animal contest for a territory, the fact that one of the animals is the owner of the territory is typically payoff-irrelevant but may act as a cue to settle the conflict. This results in a correlated equilibrium. See Maynard Smith and Parker (1976) and Selten (1980) for the introduction and study of such conditional strategies; see Cripps (1991) and Shmida and Peleg (1997) for the link with the correlated equilibrium concept.

6. Understanding the location of Nash equilibria within the polytope of correlated equilibrium distributions may lead to better algorithms to compute Nash equilibria.
7. More generally, studying correlated equilibria may lead to a better understanding of Nash equilibria. Examples will be given in the next section and in chapters 2 to 6.

1.3 Contribution of this dissertation

Properties and Applications of Dual Reduction. We begin by studying dual reduction (Myerson, 1997): a technique to reduce finite games into games with fewer strategies while selecting among correlated equilibria. More precisely, chapter 2 serves three purposes: first, we survey and unify the elementary proofs of existence of correlated equilibria (Hart and Schmeidler, 1989; Nau and McCardle, 1990), which form the mathematical foundations of dual reduction. Second, we argue that dual reduction is a useful tool to investigate properties of correlated equilibria and of Nash equilibria, and illustrate this point by a few examples. For instance, we give a direct proof of the fact that if a game has a unique correlated equilibrium, then this is a Nash equilibrium⁴, and a quasi-strict one. Third, we investigate systematically the properties of dual reduction (Is the reduction process sensitive to rescalings of the utility functions ? Is it uniquely defined ? Which strategies and equilibria are eliminated, which are not ? How does the reduction process operates in some important classes of games such as zero-sum games or symmetric games ? Etc.).

The geometry and topology of Nash equilibria and correlated equilibria. In the next four chapters, we focus on the geometry and the topology of correlated equilibria, and on geometrical relations between Nash equilibria and correlated equilibria.

The central result of chapter 3 is that the set of finite games with a unique correlated equilibrium is open. This generalizes a result of Nitzan (2005) and will prove crucial in part II, to show that certain properties of evolutionary dynamics are robust to perturbation of the game. Related results are also discussed (for instance, we show that the set of games with a

⁴The proof is direct in that it does not use the existence of a Nash equilibrium.

unique Nash equilibrium and the set of symmetric bimatrix games with a unique symmetric Nash equilibrium are not open, etc.)

Chapter 4 is a very brief note showing that in generic zero-sum games all correlated equilibria are Nash equilibria and that actually, still in generic zero-sum games but with a slightly more demanding notion of genericity, there is a unique correlated equilibrium distribution.⁵

The last two chapters are related to a result of Nau et al (2004). They show that when the correlated equilibrium distribution polytope has full dimension (i.e. the same dimension as the simplex of pure strategy profiles), then all Nash equilibria belong to the boundary of this polytope. In contrast, when the correlated equilibrium polytope does not have full dimension, then Nash equilibria may lie in its relative interior (Nau et al, 2004, section 6).

This brings attention to the class of games whose correlated equilibrium has full dimension. In chapter 5, we characterize this class of games and relate it to the class of elementary games (Myerson, 1997). In chapter 6, improving on Nau et al (2004), we characterize the class of games for which there exists a Nash equilibrium in the relative interior of the correlated equilibrium polytope (unless this polytope is a singleton). It is defined by requiring that, in every correlated equilibrium, all incentives constraints that stipulate not to deviate to a strategy that has positive probability in some correlated equilibrium be tight. We call these games “pre-tight”, consistent with Nitzan’s (2005) definition of tight games (which are pre-tight games in which all strategies have positive marginal probability in some correlated equilibrium).

Our initial motivation to study pre-tight games was the hope that a better understanding of the location of Nash equilibria within the correlated equilibrium polytope would eventually help improving existing algorithm for computing Nash equilibria. It turns out that pre-tight games are interesting in some other respect too. Indeed, we show that, in the two-player case, they include and generalize games that are best-response equivalent to a zero-sum game. In particular, Nash equilibria are exchangeable, any correlated equilibrium payoff is a Nash equilibrium payoff, the marginals of correlated equilibria yield Nash equilibria, and there are no good correlated equilibria in the sense of Rosenthal (1974). This is, up to our knowledge, the largest known class of two-person games in which Nash equilibria are exchangeable.

⁵As further discussed in chapter 4, this follows from a simple combination of arguments of Forges (1990) and of Bonhenblust et al (1950).

Chapitre 2

Dual Reduction

Abstract

Dual reduction (Myerson, 1997) is a technique to reduce finite games in a way that selects among correlated equilibria. This technique is shown to be a useful tool to study Nash equilibria and correlated equilibria and its properties are investigated.

2.1 Introduction

This chapter deals with dual reduction: a technique introduced by Myerson (1997), which allows to reduce finite games into games with fewer strategies, while selecting among correlated equilibria. The reduction operates by elimination of some pure strategies (e.g. dominated strategies) and replacement of sets of pure strategies by mixed strategies with support in these sets. The pure strategies of a reduced game are pure or mixed strategies of the original game.¹ It follows that any probability distribution on the set of strategy profiles of a reduced game induces a probability distribution on the set of strategy profiles of the initial game. In this sense, any correlated equilibrium of a reduced game induces a correlated equilibrium of the original game (Myerson, 1997). As we will see, the same result holds for Nash equilibria.

The material is organized as follows: we first survey the linear duality proofs of existence of correlated equilibria (Hart and Schmeidler, 1989; Nau and McCardle, 1990, Myerson, 1997), which form the mathematical foundation of dual reduction, and we recall Myerson's (1997) main results (sections

¹Some games may be reduced in several ways, hence the expression *a* reduced game instead of *the* reduced game.

2.3, 2.4 and appendix). We try throughout to synthesize the approaches of Hart and Schmeidler, Nau and McCardle and Myerson.

Second (sections 2.5 and 2.6), dual reduction is shown to be a useful tool to study properties of Nash equilibria and correlated equilibria. For instance, it allows to give a direct proof (not using the existence of a Nash equilibrium) of the fact that if a correlated equilibrium is the unique correlated equilibrium of a game, then it is a Nash equilibrium, and a quasi-strict one.² It also allows to show that if all pure strategies of a game are undominated, then there are certain dimensions that the correlated equilibrium polytope cannot have or that if in every correlated equilibrium, all incentive constraints are tight, then the game has a completely mixed Nash equilibrium.

Third, properties of dual reduction are investigated: Section 2.7 presents general properties (which strategies and equilibria are eliminated, which are not? If we rescale the payoffs of the game, how does it affect the reduction? Etc.). Section 2.8 deals with properties in specific classes of games, zero-sum games and symmetric games in particular. Section 2.9 shows that in almost all two-player games, the reduction process is uniquely defined. In section 2.10, equilibria remaining after reduction of the game are compared to acceptable correlated equilibria (Myerson, 1986) and to perfect correlated equilibrium distributions (Dhillon and Mertens, 1996); finally, the relevance of dual reduction as inducing a refinement concept is discussed.

2.2 Notations and definitions

We recall notations and definitions introduced in chapter 1, and introduce some new vocabulary.

2.2.1 Notations

Let

$$G = \{I, (S_i)_{i \in I}, (U_i)_{i \in I}\}$$

denote a finite game in strategic form: I is the nonempty finite set of players, S_i the nonempty finite set of pure strategies of player i and $U_i : S = \times_{i \in I} S_i \rightarrow \mathbb{R}$ the utility function of player i . As usual, $-i$ refers to the players other than i and we let $S_{-i} := \times_{j \in I \setminus \{i\}} S_j$. Pure strategies of player i (resp. strategy

²Quasi-strict equilibria are defined in definition 2.5.6.

profiles; strategy profiles of the players other than i) are denoted s_i or t_i (resp. s ; s_{-i}). We write (s_{-i}, t_i) to denote the strategy profile that differs from s only in that its i^{th} component is t_i . Finally, for any finite set Σ , $\Delta(\Sigma)$ denotes the set of probability distributions over Σ . In particular, $\Delta(S_i)$ denotes the set of mixed strategies of player i .

2.2.2 Correlated equilibria and deviation vectors

A *correlated strategy* of the players in I is a probability distribution over the set S of pure strategy profiles. Thus $\mu = (\mu(s))_{s \in S}$ is a correlated strategy if:

$$\mu(s) \geq 0 \quad \forall s \in S \quad (2.2.1)$$

$$\sum_{s \in S} \mu(s) = 1 \quad (2.2.2)$$

A correlated strategy is a *correlated equilibrium* (Aumann, 1974) if it satisfies the following *incentive constraints*:

$$\sum_{s_{-i} \in S_{-i}} \mu(s) [U_i(s) - U_i(s_{-i}, t_i)] \geq 0 \quad \forall i \in I, \forall s_i \in S_i, \forall t_i \in S_i \quad (2.2.3)$$

The following interpretation and vocabulary will be useful for the next sections. Let μ be a correlated strategy and consider the following extended game G_μ , based on G : before G is played, a strategy profile s in S is drawn at random with probability $\mu(s)$ by some mediator; then the mediator privately recommends s_i to player i ; finally, G is played. The players can thus condition their strategy in G on their private signal. A strategy of player i in this extended game is a mapping $\alpha_i : S_i \rightarrow \Delta(S_i)$ (a transition probability), which we call a *deviation plan*. Denoting by $\alpha_i(t_i|s_i)$ the probability that player i will play t_i when recommended s_i , we have:

$$\alpha_i(t_i|s_i) \geq 0 \quad \forall s_i \in S_i, \forall t_i \in S_i, \forall i \in I \quad (2.2.4)$$

$$\sum_{t_i \in S_i} \alpha_i(t_i|s_i) = 1 \quad \forall s_i \in S_i, \forall i \in I \quad (2.2.5)$$

A strategy profile is a *deviation vector*, i.e. a vector $\alpha = (\alpha_i)_{i \in I}$ of deviation plans. Such a deviation vector is *trivial* if, for all i in I , α_i is the identity.

³ The incentive constraints (2.2.3) mean that μ is a correlated equilibrium

³That is, $\alpha_i(s_i|s_i) = 1$ for all i and all s_i .

of G if and only if the trivial deviation vector (i.e. following the mediator's recommendations) is a Nash equilibrium of G_μ .

2.3 Existence of correlated equilibria

This section presents a direct proof of existence of correlated equilibria, based on those given by Hart and Schmeidler (1989), Nau and McCardle (1990) and Myerson (1997). These proofs are the mathematical roots of dual reduction and the notions introduced below will be used throughout the chapter.

Consider the following two-player, zero-sum auxiliary game Γ : the maximizer chooses a correlated strategy μ in $\Delta(S)$; the minimizer chooses a deviation vector α . The payoff is:

$$g(\mu, \alpha) := \sum_{i \in I} [U_i(\mu) - U_i(\alpha_i * \mu)] \quad (2.3.1)$$

where

$$U_i(\mu) := \sum_{s \in S} \mu(s) U_i(s)$$

is the average payoff of player i if μ is implemented, and $\alpha_i * \mu$ is the distribution on S that results if a mediator tries to implement μ but player i deviates according to α_i :

$$\alpha_i * \mu(s_{-i}, t_i) = \sum_{s_i \in S_i} \alpha_i(t_i | s_i) \mu_i(s) \quad \forall t_i \in S_i, \forall s_{-i} \in S_{-i}$$

Thus, the quantity $[U_i(\mu) - U_i(\alpha_i * \mu)]$ is the relative gain for player i that results from sticking to μ instead of deviating according to the deviation plan α_i . The payoff $g(\mu, \alpha)$ is the sum of these gains. It follows from the definition of correlated equilibria that the correlated strategy μ guarantees 0 to the maximizer (i.e. $g(\mu, \alpha) \geq 0$ for every deviation vector α) if and only if μ is a correlated equilibrium of G . It follows that G has a correlated equilibrium if and only if the value of the auxiliary game Γ is nonnegative. We have:

Theorem 2.3.1. *The value of the auxiliary game Γ is zero.*⁴

⁴The fact that the value is not only nonnegative but exactly 0 will be used later on.

We first need a definition: A deviation plan $\alpha_i : S_i \rightarrow \Delta(S_i)$ induces a Markov chain on S_i . This Markov chain maps the probability distribution $\sigma_i \in \Delta(S_i)$ to the probability distribution $\alpha_i * \sigma_i$ given by:

$$\alpha_i * \sigma_i(t_i) = \sum_{s_i \in S_i} \alpha_i(t_i | s_i) \sigma_i(s_i) \quad \forall t_i \in S_i$$

Definition 2.3.2. Let $\alpha = (\alpha_i)_{i \in I}$ be a deviation vector. A mixed strategy $\sigma_i \in \Delta(S_i)$ is α_i -invariant if $\alpha_i * \sigma_i = \sigma_i$. A correlated strategy $\mu \in \Delta(S)$ is α -invariant if $\alpha_i * \mu = \mu$ for all i in I .

We now prove the theorem:

Proof. Let α denote a deviation vector. It follows from the basic theory of Markov chains that, for each i in I , there exists at least one α_i -invariant strategy $\sigma_i \in \Delta(S_i)$. As is easily checked, the product distribution $\sigma := \prod_{i \in I} \sigma_i$ is α -invariant. This implies that $g(\sigma, \alpha) = 0$. Therefore the maximizer can defend 0. This shows that the value of the auxiliary game is nonnegative. Furthermore, the minimizer can guarantee 0 by choosing the trivial deviation vector. Therefore the value is actually 0. \square

2.4 Dual reduction: definition and known results

Unless stated otherwise, all definitions and results of this section are due to Myerson (1997).

2.4.1 Definition

The Markov chain on S_i induced by the deviation plan α_i partitions S_i into transient states and disjoint minimal absorbing sets⁵. For any minimal absorbing set B_i , there exists a unique α_i -invariant strategy with support in B_i (actually its support is exactly B_i). Let S_i/α_i denote the set of (randomized) α_i -invariant strategies with support in some minimal α_i -absorbing set. It may be shown that the set of α_i -invariant strategies is the set of random mixture of the strategies in S_i/α_i ; that is, the simplex $\Delta(S_i/\alpha_i)$.

⁵A subset B_i of S_i is α_i -absorbing if $\alpha(t_i | s_i) = 0$ for all s_i in B_i and all t_i in $S_i \setminus B_i$. An α_i -absorbing set is minimal if it contains no proper α_i -absorbing subset.

Let $\alpha = (\alpha_i)_{i \in I}$ be a deviation vector. The α -reduced game

$$G/\alpha = \{I, (S_i/\alpha_i)_{i \in I}, (U_i)_{i \in I}\}$$

is the game obtained from G by restricting the players to α -invariant strategies. That is, the set of players and the payoff functions are the same as in G but, for all i in I , the pure strategy set of player i is now S_i/α_i .⁶

Before turning to dual reduction and its properties, let us specify our vocabulary: let $s_i, t_i \in S_i$ and $s \in S$. The pure strategy s_i (resp. pure strategy profile s) is *eliminated* by the reduction process if $\sigma_i(s_i) = 0$ for all σ_i in S_i/α_i (resp. if $\sigma(s) = 0$ for all σ in S/α). Thus s_i (resp. s) is eliminated if and only if (resp. if and only if for some i in I) s_i is transient under α_i . The strategies s_i and t_i are *grouped together* if there exists σ_i in S_i/α_i such that both s_i and t_i belong to the support of σ_i . Thus, s_i and t_i are grouped together if and only if they are recurrent under α_i and belong to the same minimal α_i -absorbing set.⁷

Definition 2.4.1. *A dual vector is an optimal strategy of the minimizer in the auxiliary game of section 2.3. Since the value of this auxiliary game is zero, a deviation vector α is a dual vector if*

$$\forall s \in S, -g(s, \alpha) \geq 0 \tag{2.4.1}$$

where

$$-g(s, \alpha) = \sum_{i \in I} [U_i(\alpha_i * s) - U_i(s)] = \sum_{i \in I} \sum_{t_i \in S_i} \alpha_i(t_i | s_i) [U_i(s_{-i}, t_i) - U_i(s)] \tag{2.4.2}$$

Remark 2.4.2. *If s has positive probability in some correlated equilibrium, then for every dual vector α , equation (2.4.1) holds with equality ($g(s, \alpha) = 0$).*

Indeed, if μ is a correlated equilibrium and α a dual vector, then (μ, α) is a profile of optimal strategies of the auxiliary game of section 2.3, hence for every strategy profile s in the support of μ , $g(s, \alpha)$ is equal to the value of this auxiliary game, i.e. to 0.

⁶Strictly speaking the payoff functions of the reduced game are the functions *induced* by the original payoff functions on the reduced strategy space.

⁷A pure strategy is α_i -*recurrent*, or *recurrent under α_i* , if it belongs to a minimal absorbing set, and *transient* otherwise.

Definition 2.4.3. A dual reduction of G is an α -reduced game G/α where α is a dual vector. An iterative dual reduction of G is a reduced game $G/\alpha^1/\alpha^2/\dots/\alpha^m$, where m is a positive integer and, for all k in $\{1, 2, \dots, m\}$, α^k is a dual vector of $G/\alpha^1/\alpha^2/\dots/\alpha^{k-1}$.

Note that, depending on the context, “dual reduction” may refer either to the reduction process or to a game obtained by applying this reduction process. Myerson (1997, section 6) provides many instructive examples of dual vectors and dual reductions. Henceforth, unless stated otherwise, α is a dual vector.

2.4.2 Some intuition on dual vectors

This section aims at providing some intuition on the peculiarities of dual vectors and α -invariant strategies. The main point is that if α is a dual vector and if every player j other than i uses an α_j -invariant strategy, then there exists a best-response to the strategies of the players other than i which is α_i -invariant. More precisely, assume that for every j in $I \setminus \{i\}$ the mixed strategy σ_j is α_j -invariant and let σ_i be any mixed strategy of player i . Assimilate throughout a mixed strategy profile and the probability distribution it induces on $\Delta(S)$. It follows from (2.3.1) and the definition of dual vectors that

$$-g(\sigma, \alpha) = \sum_{j \in I} [U_j(\sigma_{-j}, \alpha_j * \sigma_j) - U_j(\sigma)] \geq 0 \quad (2.4.3)$$

Since for every j in I , $\alpha_j * \sigma = (\sigma_{-j}, \alpha_j * \sigma_j)$ and since for every $j \neq i$, $\alpha_j * \sigma_j = \sigma_j$, (2.4.3) boils down to:

$$U_i(\sigma_{-i}, \alpha_i * \sigma_i) \geq U_i(\sigma) \quad (2.4.4)$$

As a particular case, we obtain:

$$U_i(\sigma_{-i}, \alpha_i * s_i) \geq U_i(\sigma_{-i}, s_i), \quad \forall s_i \in S_i \quad (2.4.5)$$

Now fix s_i in S_i and let $\alpha_i^k * s_i$ denote the mixed strategy $\alpha_i * (\alpha_i * \dots * (\alpha_i * s_i))$, where α_i appears k times. It may be shown that the Cesaro average of the sequence $(\alpha_i^k * s_i)_{k \in \mathbb{N}}$ converges and that its limit, which we denote by σ_{s_i} , is

α_i -invariant.⁸ Furthermore, repeated applications of (2.4.5) yield

$$U_i(\sigma_{-i}, \sigma_{s_i}) \geq U_i(\sigma_{-i}, s_i)^9 \quad (2.4.6)$$

In particular, if s_i is a best-response to σ_{-i} , then so is σ_{s_i} , which shows that there exists an α_i -invariant best-response to σ_{-i} .

2.4.3 Main properties

First, dual reduction generalizes elimination of weakly dominated strategies in the following sense:

Proposition 2.4.4. *Let $s_i \in S_i$; assume that there exists a mixed strategy σ_i in $\Delta(S_i)$, $\sigma_i \neq s_i$, such that $U_i(s_{-i}, \sigma_i) \geq U_i(s_{-i}, s_i)$ for every s_{-i} in S_{-i} . Then there exists a dual vector α such that $S_i/\alpha_i = S_i \setminus \{s_i\}$ and $S_j/\alpha_j = S_j$ for every j in $I \setminus \{i\}$.*

Proof. Take for α : $\alpha_i * s_i = \sigma_i$, $\alpha_i * t_i = t_i$ if $t_i \in S_i \setminus \{s_i\}$, and $\alpha_j * s_j = s_j$ for every s_j in S_j if $j \neq i$. \square

The main property of dual reduction is that it selects among correlated equilibria: let G/α denote a dual reduction of G ; let $S/\alpha = \times_{i \in I} S_i/\alpha_i$ denote the set of pure strategy profiles of G/α . Let $\mu \in \Delta(S/\alpha)$ be a correlated strategy of the reduced game G/α ; the G -equivalent correlated strategy $\bar{\mu}$ is the distribution on S induced by μ :

$$\bar{\mu}(s) = \sum_{\sigma \in S/\alpha} \mu(\sigma) \left(\prod_{i \in I} \sigma_i(s_i) \right), \quad \forall s \in S \quad (2.4.7)$$

Theorem 2.4.5. *If μ is a correlated equilibrium of G/α , then $\bar{\mu}$ is a correlated equilibrium of G .*

It follows from (2.4.7) that if μ is a product distribution, then so is $\bar{\mu}$. Therefore we also have:

⁸The proof relies on the fact that every limit point of the Cesaro average of $(\alpha_i^k * s_i)_{k \in \mathbb{N}}$ is α_i -invariant, and on the decomposition of S_i in transient states and minimal absorbing sets.

⁹Note that σ_{s_i} does not depend on σ_{-i} ; therefore, provided that every player j other than i uses only α_j -invariant strategies, then, whatever the precise choice of strategies players other than i , player i is *always* (weakly) better off playing the α_i -invariant strategy σ_{s_i} than playing s_i .

Corollary 2.4.6. *If μ is a Nash equilibrium of G/α , then $\bar{\mu}$ is a Nash equilibrium of G .^{10,11}*

Let σ and τ be two pure strategy profile of the reduced game G/α . By definition of S/α , if $\sigma \neq \tau$, then σ and τ correspond to mixed strategy profiles of the original game *with disjoint support*. It follows that the application $\mu \rightarrow \bar{\mu}$ defined by (2.4.7) is injective, hence that *distinct* correlated equilibria of G/α induce *distinct* correlated equilibria of G . Thus, if G has a unique correlated equilibrium, then G/α has a unique correlated equilibrium too: a fact which will be used in section 2.5.

By induction, theorem 2.4.5 extends to iterative dual reductions. That is, any correlated equilibrium of an iterative dual reduction of G induces on $\Delta(S)$ a correlated equilibrium of G . Finally, as a corollary of theorem 2.4.5, Myerson shows that against any strategy of the other players in the reduced game, player i is indifferent between strategies belonging to the same minimal absorbing set:

Proposition 2.4.7. *Let B_i denote a minimal α_i -absorbing set. For every $j \neq i$, let $\sigma_j \in S_j/\alpha_j$ and let $\sigma_{-i} := \times_{j \in I \setminus \{i\}} \sigma_j$. For any pure strategies s_i and t_i in B_i , $U_i(\sigma_{-i}, s_i) = U_i(\sigma_{-i}, t_i)$.*

2.4.4 Jeopardization and Elementary Games

Let us call a dual vector *trivial* if this is the trivial deviation vector. It follows from the basic theory of Markov chains (see for instance Karr, 1990) that a game may be (strictly) reduced if and only if there exists a nontrivial dual vector. This leads to the question: when do nontrivial dual vectors exist? A first step to answer this question is to introduce the notions of jeopardization and elementary games:

Definition 2.4.8. *Let $s_i, t_i \in S_i$. The strategy t_i jeopardizes s_i if for all correlated equilibria μ :*

$$\sum_{s_{-i} \in S_{-i}} \mu(s) [U_i(s) - U_i(s_{-i}, t_i)] = 0$$

That is, in all correlated equilibria in which s_i is played, t_i is an alternative best response to the conditional probabilities on S_{-i} given s_i . Note that if

¹⁰We assimilate throughout a Nash equilibrium and its (product) distribution.

¹¹This property is not stated in (Myerson, 1997) but in (Myerson, 2003).

the pure strategy s_i has zero probability in all correlated equilibria, then any pure strategy t_i of player i jeopardizes s_i . Using complementary slackness properties allows to prove that:

Proposition 2.4.9. *The strategy t_i jeopardizes s_i if and only if there exists a dual vector α such that $\alpha_i(t_i|s_i) > 0$.*

Thus, there exists a nontrivial dual vector if and only if some strategy is jeopardized by some other strategy.

Definition 2.4.10. *A correlated equilibrium μ is strict if*

$$\mu(s_i \times S_{-i}) > 0 \Rightarrow \sum_{s_{-i} \in S_{-i}} \mu(s) [U_i(s) - U_i(s_{-i}, t_i)] > 0 \quad \forall i \in I, \forall s_i \in S_i, \forall t_i \neq s_i$$

A game is *elementary* if it has a strict correlated equilibrium with full support, or equivalently, if it has a correlated equilibrium μ satisfying with strict inequality all incentives constraints in the sense that:

$$\sum_{s_{-i} \in S_{-i}} \mu(s) [U_i(s) - U_i(s_{-i}, t_i)] > 0 \quad \forall i \in I, \forall s_i \in S_i, \forall t_i \neq s_i \quad (2.4.8)$$

Since¹² the set of correlated equilibria is convex, it follows from definition 2.4.8 that a game is elementary if and only if there exist no i , s_i and $t_i \neq s_i$ such that t_i jeopardizes s_i . Thus proposition 2.4.9 implies that:

Corollary 2.4.11. *A game may be reduced if and only if it is not elementary. By iterative dual reduction, any game is eventually reduced to an elementary game.*

2.4.5 Full dual reduction

Let us say that two dual reductions G/α and G/β of the same game are different if $S/\alpha \neq S/\beta$. A game may admit different dual reductions (for instance, if several strategies are weakly dominated). A tentative way to restore uniqueness is to consider only reductions by some special dual vectors, which minimize the number of pure strategies remaining in the reduced game:

¹²A strict correlated equilibrium with full support satisfies (2.4.8). Conversely, if μ satisfies (2.4.8), then every pure strategy has positive marginal probability in μ and, by perturbing μ , one obtains a correlated equilibrium with full support.

Definition 2.4.12. A dual vector α is full if $\alpha_i(t_i|s_i) > 0$ for all i in I , and all pure strategies s_i and t_i in S_i such that t_i jeopardizes s_i .

Full dual vectors always exist. This follows from proposition 2.4.9 and from the convexity of the set of dual vectors.

Definition 2.4.13. A full dual reduction of G is an α -reduced game G/α where α is a full dual vector. An iterative full dual reduction of depth m of G is a game $G/\alpha^1/\alpha^2/\dots/\alpha^m$ where m is a positive integer and, for all k in $\{1, 2, \dots, m\}$, α^k is a full dual vector of $G/\alpha^1/\alpha^2/\dots/\alpha^{k-1}$.

All full dual vectors α define the same minimal α_i -absorbing sets. Thus in all full dual reductions, the same strategies are eliminated and the same strategies are grouped together. A game may nonetheless admit different full dual reductions, because the way strategies are grouped together may differ quantitatively. We will return to this point in section 2.9.

2.5 Some applications of dual reduction

This section and the next illustrate the usefulness of dual reduction to study the sets of Nash equilibria and correlated equilibria. First, consider the following well-known fact:

Proposition 2.5.1. If a (finite) game has a unique correlated equilibrium, then this correlated equilibrium is a Nash equilibrium.

The standard proof invokes the existence of Nash equilibrium in finite games, hence relies implicitly on a powerful fixed point theorem such as Brouwer's or Kakutani's. Dual reduction yields an alternate proof relying solely on linear duality:

Proof. Let G be a game with a unique correlated equilibrium. By corollary 2.4.11, G has (at least one) elementary iterative reduction G^e . Since G^e is elementary (i.e. has a strict correlated equilibrium with full support), it follows that either G^e has an infinity of correlated equilibria, or G^e has a unique strategy profile. Since G has a unique correlated equilibrium and since different correlated equilibria of G^e induce different correlated equilibria of G , the first case is ruled out. Therefore, G^e has a unique strategy profile, hence trivially a Nash equilibrium. By corollary 2.4.6, this implies that G

has a Nash equilibrium hence that the unique correlated equilibrium of G is a Nash equilibrium. \square

Note that the above proof relies on: a) the definition of dual reduction, which requires the Minimax theorem and existence of invariant distributions for finite Markov chains; and b) corollary 2.4.11, which is proved by Myerson (1997) through the strong complementary property of linear programs. Since the existence of invariant distributions for finite Markov chains can be deduced from the Minimax theorem (see appendix A.1), it follows that our proof of proposition 2.5.1 relies solely on linear duality. In particular no fixed point theorem is used.

Dual¹³ reduction is also useful to prove existence of Nash equilibria with special properties:

Proposition 2.5.2. *In any finite game, there exists a Nash equilibrium σ such that, for every player i and every pure strategy s_i of player i that has marginal probability zero in all correlated equilibria, s_i is not a best response to σ_{-i} .*

To prove proposition 2.5.2, we first need to introduce some special dual vectors. Note that the set of dual vectors of a game is bounded and defined by linear inequalities, hence is a polytope.

Definition 2.5.3. *A dual vector α is interior if it is the unique dual vector of the game or if it belongs to the relative interior of the set of dual vectors.*

It follows from the definitions of dual vectors and full dual vectors that a strictly convex combination of a full dual vector with any dual vector is a full dual vector. This implies that:

Remark 2.5.4. *Any interior dual vector is full.*

Furthermore,

¹³The above proof relies on: a) the definition of dual reduction, which requires the Minimax theorem and existence of invariant distributions for finite Markov chains; and b) corollary 2.4.11, which is proved by Myerson (1997) through the strong complementary property of linear programs. Since the existence of invariant distributions for finite Markov chains can be deduced from the Minimax theorem (see appendix A.1), it follows that the above proof relies solely on linear duality.

Proposition 2.5.5. *If α is an interior dual vector, then for every strategy profile s that has probability zero in all correlated equilibria,*

$$-g(s, \alpha) = \sum_{i \in I} [U_i(\alpha_i * s) - U_i(s)] > 0$$

(The above equality merely repeats the definition of $g(s, \alpha)$.)

Proof. Call *strong* the dual vectors satisfying the above property. Existence of strong dual vectors follows from Nau and McCardle's (1990) proof of existence of correlated equilibria (see section 2.11.2 for details).¹⁴ Furthermore, any strictly convex combination of a dual vector with a strong dual vector is a strong dual vector. The result follows. \square

We are now in a position to prove proposition 2.5.2:

Proof. Let α be an interior dual vector. Let σ be a mixed strategy profile of G and let μ denote the product distribution induced by (σ_{-i}, s_i) . If s_i has marginal probability zero in all correlated equilibria, then it follows from proposition 2.5.5 that $g(\mu, \alpha) < 0$. Proceeding as in the derivation of 2.4.5, we obtain that if for every $j \neq i$, the mixed strategy σ_j is α_j -invariant, then

$$U_i(\sigma_{-i}, \alpha_i * s_i) - U_i(\sigma_{-i}, s_i) > 0$$

Therefore, s_i is not a best-response to σ_{-i} . It follows that if σ is a Nash equilibrium of the reduced game G/α (hence also a Nash equilibrium of G) then, for every player i and every pure strategy s_i that has probability zero in all correlated equilibria, s_i is not a best-response to σ_{-i} \square

Definition 2.5.6. *A Nash equilibrium σ is quasi-strict if for every player i in I any pure best-response to σ_{-i} belongs to the support of σ_i .*

¹⁴Existence of strong dual vectors can also be shown by applying the "equalizer theorem" to the auxiliary game Γ of section 2.3 (the equalizer theorem (see Raghavan, 1994) states that in a zero-sum game, a pure strategy of the maximizer belongs to the support of an optimal strategy if and only if it is a best-response to all optimal strategies of the minimizer; this is a version of the strong complementary property of linear programs). Indeed, in the auxiliary game Γ the optimal strategies of the maximizer (resp. minimizer) are the correlated equilibria (resp. dual vectors) of G . Thus, if the strategy profile s has probability zero in all correlated equilibria, then there exists a dual vector α such that $g(s, \alpha)$ is negative. Since the set of dual vectors is convex, this implies that there exists a dual vector such that $g(s, \alpha)$ is negative for every s with probability zero in all correlated equilibria.

It follows from proposition 2.5.2 that:

Corollary 2.5.7. *For any finite game, if a Nash equilibrium is the unique correlated equilibrium of the game, then it is quasi-strict.*

Proof. Let σ be a Nash equilibrium and the unique correlated equilibrium of a finite game. If the pure strategy s_i does not belong to the support of σ_i then it has marginal probability zero in all correlated equilibria; by proposition 2.5.2 (and because σ is the unique Nash equilibrium), this implies that s_i is not a best-response to σ_{-i} . \square

Recall that the set of correlated equilibria is a convex polytope (i.e. it is bounded and defined by a finite number of linear inequalities). Call it the *correlated equilibrium polytope*. The next proposition shows that if all pure strategies are undominated, then there are certain dimensions that this polytope cannot have:

Proposition 2.5.8. *Let $|S|$ denote the cardinal of the set of pure strategy profiles S . If no pure strategy is dominated in the sense that:*

$$\forall i \in I, \forall s_i \in S_i, \forall \sigma_i \in \Delta(S_i), \sigma_i \neq s_i \Rightarrow \exists s_{-i} \in S_{-i}, U_i(s) > U_i(s_{-i}, \sigma_i) \quad (2.5.1)$$

then the correlated equilibrium polytope does not have dimension $|S| - 2$.

(This shows for instance that, as is well known, the set of correlated equilibria of a 2×2 game cannot have dimension 2, unless some pure strategy is dominated.)

Proof. We first need a lemma:

Lemma 2.5.9. *Let $s_i \in S_i$; assume that there exists a dual vector α such that $s_i \notin S_i/\alpha_i$ and $S_j/\alpha_j = S_j$ for all j in $I \setminus \{i\}$. Then there exists a mixed strategy σ_i in $\Delta(S_i)$ such that $\sigma_i \neq s_i$ and $U_i(s_{-i}, \sigma_i) \geq U_i(s)$ for all s_{-i} in S_{-i} .*

Proof. Let $\sigma_i = \alpha_i * s_i$. For all $j \neq i$, every pure strategy in S_j is α_j -invariant. Therefore (2.4.1) yields:

$$U_i(s_{-i}, \sigma_i) \geq U_i(s) \quad \forall s_{-i} \in S_{-i}$$

Furthermore $s_i \notin S_i/\alpha_i$ hence s_i cannot be α_i -invariant. Therefore $\sigma_i \neq s_i$ \square

We¹⁵ now prove the proposition: if the game is elementary, then the correlated equilibrium polytope has dimension $|S| - 1$. Otherwise, there exists i in I , s_i in S_i and t_i in S_i such that t_i jeopardizes s_i . That is, for every correlated equilibrium μ ,

$$\sum_{s_{-i} \in S_{-i}} \mu(s) [U_i(s) - U_i(s_{-i}, t_i)] = 0 \quad (2.5.2)$$

Therefore, by proposition 2.4.9, there exists a dual vector α such that $s_i \notin S_i/\alpha_i$. Since by assumption s_i is undominated in the sense of (2.5.1), it follows from lemma 2.5.9 that there exists j in $I \setminus \{i\}$ and s_j in S_j such that $s_j \notin S_j/\alpha_j$. This implies that s_j is jeopardized by some strategy $t_j \in S_j \setminus \{s_j\}$. That is, for every correlated equilibrium μ ,

$$\sum_{s_{-j} \in S_{-j}} \mu(s) [U_j(s) - U_j(s_{-j}, t_j)] = 0 \quad (2.5.3)$$

Condition (2.5.1) implies that neither (2.5.2) nor (2.5.3) is satisfied by all μ in \mathbb{R}^S and that (2.5.2) and (2.5.3) are not equivalent. As an intersection of two non identical hyperplanes, the set of points of \mathbb{R}^S satisfying (2.5.2) and (2.5.3) is a vector space of dimension $|S| - 2$. Its intersection with the simplex $\Delta(S)$ has at most dimension $|S| - 3$ and includes the correlated equilibrium polytope. Therefore, this polytope has at most dimension $|S| - 3$. \square

2.6 Some results used in chapter 6

This section groups results which will be used in chapter 6. It is more convenient to derive these results here as they rely on dual reduction. Some definitions are needed:

Definition 2.6.1. *A game is tight (Nitzan, 2005) if in every correlated equilibrium μ , all incentives constraints are tight:*

$$\sum_{s_{-i} \in S_{-i}} \mu(s) [U_i(s) - U_i(s_{-i}, t_i)] = 0 \quad \forall i \in I, \forall s_i \in S_i, \forall t_i \in S_i$$

¹⁵As a side remark, lemma 2.5.9 implies the converse of proposition 2.4.4. That is, if a dual reduction simply consists in eliminating a pure strategy, then this pure strategy is dominated (in the sense of proposition 2.5.8, i.e. allowing for the “dominating” and the “dominated” strategies to always yield the same payoff).

Definition 2.6.2. A pure strategy is coherent if it has positive marginal probability in at least one correlated equilibrium.

Notation: S_i^c denotes the set of coherent pure strategies of player i .

Definition 2.6.3. A game is pre-tight if in every correlated equilibrium μ all incentive constraints that stipulate not to deviate to a coherent pure strategy are tight:

$$\sum_{s_{-i} \in S_{-i}} \mu(s) [U_i(s) - U_i(s_{-i}, t_i)] = 0 \quad \forall i \in I, \forall s_i \in S_i, \forall t_i \in S_i^c$$

The classes of tight and pre-tight games will be studied in chapter 6. The reason why dual reduction is useful to study these classes of games is that the definitions of tight and pre-tight games may be rephrased in terms of jeopardization:

Remark 2.6.4. A game is tight (resp. pre-tight) if and only if for every player i , every pure strategy s_i of player i , any (resp. any coherent) pure strategy of player i jeopardizes s_i .

The following proposition provides a criterion allowing to show that a game is tight (or pre-tight) without computing its correlated equilibria:

Proposition 2.6.5. (1). A game is tight if and only if there exists a dual vector α such that, for every player i in I and every pure strategy s_i in S_i , the mixed strategy $\alpha_i * s_i$ is completely mixed.

(2). A game is pre-tight if and only if there exists a dual vector α , and, for every player i in I , a subset $S'_i \subseteq S_i$ of pure strategies such that:

- (a) For every player i in I and every pure strategy s_i in S'_i , the mixed strategy $\alpha_i * s_i$ has support S'_i .
- (b) For every pure strategy profile s in S that does not belong to $S' := \times_{i \in I} S'_i$, we have $g(\alpha, s) < 0$

In that case, S'_i is the set of coherent pure strategies of player i . That is, $S'_i = S_i^c$.

Proof. We first prove point (1): If the game is tight, then it follows from the definition 2.4.12 of dual vectors and remark 2.6.4 that any full dual vector satisfies the desired property. Conversely, if there exists a dual vector α such

that, for all i and all s_i , $\alpha_i * s_i$ is completely mixed, then it follows from proposition 2.4.9 and remark 2.6.4 that the game is tight.

The proof of point (2) is more involved and requires two additional lemmas:

Lemma 2.6.6. *Let α be a dual vector.*

(i) *Let $s \in S$. If $g(s, \alpha) < 0$ then for every σ in S/α we have $\sigma(s) = 0$.*

(ii) *Let $s_i \in S_i$. If for every $s_{-i} \in S_{-i}$ we have $g(s, \alpha) < 0$ (with $s = (s_{-i}, s_i)$), then for every σ_i in S_i/α_i we have $\sigma_i(s_i) = 0$.*

Proof. Proof of (i): Assume that $g(s, \alpha) < 0$. Since $g(t, \alpha) \leq 0$ for all t in S , it follows that if s has positive probability in some correlated strategy μ then $g(\mu, \alpha) < 0$. Moreover, if σ belongs to S/α , then σ is α -invariant, hence $g(\sigma, \alpha) = 0$ by (2.3.1). Therefore s cannot have positive probability in σ .

Proof of (ii): Let $s_i \in S_i$. It follows from (i) that if for every $s_{-i} \in S_{-i}$, $g(s, \alpha) < 0$, then for every $s_{-i} \in S_{-i}$ and every σ in S/α , $\sigma(s) = 0$. This implies that $\sigma_i(s_i) = 0$ for every σ_i in S_i/α_i . \square

Lemma 2.6.7. *Let α be a dual vector. For every i in I , there exists a coherent pure strategy s_i which is recurrent under α_i .*

Proof. Let σ be a Nash equilibrium of G/α , hence also of G . Any pure strategy $s_i \in S_i$ in the support of σ_i is both coherent and recurrent under α_i . \square

We now prove point (2) of proposition 2.6.5: consider a pre-tight game. Let $S'_i = S_i^c$ and let α be an interior dual vector. It follows from proposition 2.5.5 that condition (b) is satisfied. We now prove that condition (a) is satisfied. Let $s_i \in S_i^c$. The dual vector α is interior hence full. Therefore, since every coherent pure strategy of player i jeopardizes s_i , it follows that the support of $\alpha_i * s_i$ contains S_i^c . It follows that the coherent pure strategies of player i are either all transient or all recurrent under α_i . The former is ruled out by lemma 2.6.7 hence all coherent strategies of player i are recurrent, in particular s_i is recurrent. Therefore if $\alpha_i(t_i|s_i) > 0$ then t_i is recurrent too. But it follows from lemma 2.6.6, condition (b) and $S'_i = S_i^c$, that every $t_i \notin S_i^c$ is transient. Therefore, $\alpha_i * s_i$ has exactly support S_i^c , hence condition (a) is satisfied .

Conversely, assume that there exists a dual vector α and, for every player i in I , a subset S'_i of S_i such that conditions (a) and (b) are checked. Assume

first that $S'_i = S_i^c$ for every i in I . In view of proposition 2.4.9, condition (a) then implies that the game is pre-tight. Thus, it suffices to prove that $S'_i = S_i^c$.

If $s_i \in S_i \setminus S'_i$, then for every strategy profile s_{-i} of the players other than i , the pure strategy profile $s = (s_i, s_{-i})$ does not belong to $\times_{i \in I} S'_i$. Therefore, $g(s, \alpha) < 0$ by condition (b). By remark 2.4.2, this implies that the strategy profile s has probability zero in every correlated equilibrium. Since this holds for every s_{-i} in S_{-i} , it follows that the strategy s_i is not coherent. Hence,

$$S_i^c \subseteq S'_i \tag{2.6.1}$$

It remains to prove the reverse inclusion. Condition (a) implies that S'_i is a minimal α_i -absorbing set. Furthermore, it follows from lemma 2.6.6 and condition (b) that the pure strategies of player i that do not belong to S'_i are transient under α_i . Therefore, S'_i is the unique minimal α_i -absorbing set. This implies that there exists a unique α_i -invariant strategy σ_i and that its support is S'_i . In the reduced game G/α , the corresponding strategy profile $\sigma = (\sigma_i)_{i \in I}$ is the unique strategy profile, hence, trivially, a Nash equilibrium. This implies that σ is a Nash equilibrium of G . Therefore, any pure strategy in the support of σ_i is also coherent, i.e. $S'_i \subseteq S_i^c$. Together with (2.6.1), this shows that $S'_i = S_i^c$. \square

We now show that tight and pre-tight games have special Nash equilibria:

Proposition 2.6.8. (i) *Every tight game has a completely mixed Nash equilibrium.* (ii) *Every pre-tight game has a quasi-strict Nash equilibrium with support $S^c = \times_i S_i^c$.*

Proof. It follows from proposition 2.6.5 that if a game is tight then it is pre-tight and $S_i^c = S_i$. Therefore, (ii) implies (i). We now prove (ii): Let G be a pre-tight game and let α be an interior dual vector. It follows from the proof of proposition 2.6.5, point (2), that in the reduced game G/α there is a unique strategy profile σ and that σ is a Nash equilibrium of G with support S^c . Furthermore, the proof of proposition 2.5.2 shows that if $s_i \notin S_i^c$ then s_i is not a best-response to σ_{-i} , hence that σ is quasi-strict. \square

We conclude this section by mentioning applications of dual reduction in later chapters: in chapter 3, corollary 2.5.7 is used to show that, for any number of players, the set of games with a unique correlated equilibrium

is open. In chapter 6, proposition 2.6.8 is used to characterize the set of games whose correlated equilibrium polytope contains a Nash equilibrium in its relative interior. In chapter 9, dual reduction is used to classify 3×3 symmetric games, and to prove that in every 3×3 symmetric game, every pure strategy that has probability zero in all correlated equilibria is eliminated by any convex monotonic dynamics (Hofbauer and Weibull, 1996) and by the two-population best-response dynamics. Finally, in chapter 10, our proof of the fact that general Rock-Paper-Scissors games have a unique correlated equilibrium relies implicitly on a dual reduction argument.

These applications of dual reduction suggests that this technique has great potentialities as a tool to study correlated equilibria. This leads us to investigate its properties.

2.7 General properties of dual reduction

A basic desirable property for a decision-theoretic concept is to be independent of the specific (von Neumann-Morgenstern) utility functions chosen to represent the preferences of the agents. We begin by showing that dual reduction meets this requirement; that is, the ways in which a game may be reduced are unaffected by positive affine transformations of the utility functions. We then show that theorem 2.4.5 extends to the equalizing correlated equilibrium concept, introduced by Sorin (1998)¹⁶ and point out that its converse holds: if a correlated strategy μ of a reduced game induces an equilibrium in the original game, then μ is an equilibrium of the reduced game. Elimination of strategies and equilibria is then investigated. We show that strategies that are weakly dominated (resp. are never played in correlated equilibria; have positive probability in some strict correlated equilibrium) need not be (resp. are always; are never) eliminated in full dual reductions. Finally, we show that if a pure strategy is redundant (in a sense to be made precise), then in a full dual reduction, it is either eliminated or grouped with the strategies making it redundant.

¹⁶Sorin uses the name “distribution equilibrium”.

2.7.1 Independence from the choice of utility functions

In this section G and G' are two games with the same sets of players and strategies and U'_i is the utility function of player i in G' .

Definition 2.7.1. *The games G and G' are best-response equivalent if for every player i in I , for every pure strategy s_i in S_i , and every correlated strategy μ_{-i} in $\Delta(S_{-i})$ the pure strategy s_i is a best-response to μ_{-i} in G if and only if it is a best-response to μ_{-i} in G' .¹⁷*

Since the fact that a strategy jeopardizes another strategy depends only on the best-response correspondence, it follows that games which are best-response equivalent are reduced similarly by dual reduction:

Proposition 2.7.2. *Let G and G' be best-response equivalent. Let s_i and t_i be pure strategies of player i . (i) The pure strategy t_i jeopardizes s_i in G if and only if it jeopardizes s_i in G' . (ii) The same strategies are grouped together (resp. eliminated) in full dual reductions of G and in full dual reductions of G' .*

Proof. Point (i) follows from definitions 2.7.1 and 2.4.8. Point (ii) follows from (i) and the definition of dual vectors (definition 2.4.12). \square

Definition 2.7.3. *The game G' is a rescaling of G if for every player i in I , there exists a positive constant a_i and a function $f_i : S_{-i} \rightarrow \mathbb{R}$ such that:*

$$U'_i(s) = a_i \cdot U_i(s) + f_i(s_{-i}), \quad \forall s \in S$$

A game and its rescalings need not have the same dual vectors, as the following example demonstrates:

Example 2.7.4.

	x_2	y_2		x_2	y_2
x_1	1, -1	0, 0		x_1	2, -1
y_1	0, 0	1, -1		y_1	0, 0

¹⁷In games with three or more players, there is another reasonable notion of best-response equivalence. It is obtained by requiring only that the best-responses to *products* of mixed strategies of the other players be the same in G and in G' , i.e. by replacing $\mu_{-i} \in \Delta(S_{-i})$ by $\sigma_{-i} \in \times_{j \in I \setminus \{i\}} \Delta(S_j)$ in definition 2.7.1. Originally, i.e. in the article of Rosenthal (1974), best-response equivalence was defined for two-player games, where the two notions coincide. The notion of definition 2.7.1 is more appropriate when dealing with correlated equilibria.

Let G denote the game on the left (a version of Matching Pennies) and G' its rescaling on the right. Define the deviation vector α by $\alpha_i * x_i = \alpha_i * y_i = \frac{1}{2}x_i + \frac{1}{2}y_i$ for every i in $\{1, 2\}$. The deviation vector α is a dual vector for G but not for G' .

The next proposition shows however that a game and its rescalings have the same dual reductions. We first need a lemma, which shows that different deviation vectors may induce the same reduced game:

Lemma 2.7.5. *Let α_i (resp. α_i^{id}) be a (resp. the trivial) deviation plan for player i . For any $0 \leq \epsilon \leq 1$, let $\alpha_i^\epsilon = \epsilon\alpha_i + (1 - \epsilon)\alpha_i^{id}$. If ϵ is positive then $S_i/\alpha_i = S_i/\alpha_i^\epsilon$.*

Proof. For any mixed strategy σ_i in $\Delta(S_i)$, $\alpha_i^\epsilon * \sigma_i - \sigma_i = \epsilon(\alpha_i * \sigma_i - \sigma_i)$. Therefore, α_i and α_i^ϵ induce the same invariant strategies. \square

Proposition 2.7.6. *Let G' be a rescaling of G . If α is a dual vector of G then there exists a dual vector α' of G' such that $G'/\alpha' = G/\alpha$.*

Proof. The proof consists in showing that dual vectors of G may be “rescaled” into dual vectors of G' : Let α be a dual vector of G . Let $a_k = \min_{i \in I} a_i$ and, for each i in I , let $\epsilon_i = a_k/a_i$ (the constants a_i are those of definition 2.7.3). Let α' denote the deviation vector whose i^{th} component is $\alpha_i^{\epsilon_i}$, (defined in lemma 2.7.5), for every i in I . Let g and g' denote the payoff functions in the auxiliary zero-sum games associated respectively to G and G' . We have:

$$g'(s, \alpha') = a_k \times g(s, \alpha) \quad \forall s \in S \quad (2.7.1)$$

Since α is a dual vector of G , it follows from 2.7.1 that α' is a dual vector of G' . Furthermore lemma 2.7.5 implies that $G'/\alpha' = G'/\alpha$. This completes the proof. \square

2.7.2 An extension of theorem 2.4.5

In this section, we present an equilibrium concept due to Sorin (1998) and show that theorem 2.4.5 extends to this concept. We then point out a converse of theorem 2.4.5.

Consider a correlated strategy μ in $\Delta(S)$ and a pure strategy s_i of player i . If $\mu(s_i \times S_{-i}) > 0$, let $\mu(\cdot|s_i)$ denote the conditional probability on S_{-i} given s_i :

$$\mu(s_{-i}|s_i) = \mu(s_{-i}, s_i)/\mu(s_i \times S_{-i})$$

Definition 2.7.7. *The correlated strategy μ in $\Delta(S)$ is an equalizing distribution if*

$$\mu(s_i \times S_{-i}) > 0 \Rightarrow \sum_{s_{-i} \in S_{-i}} \mu(s_{-i}|s_i) U_i(s) = U_i(\mu) \quad \forall i \in I, \forall s_i \in S_i,$$

That is, in an equalizing distribution, the expected payoff given the recommendation of a pure strategy is independent of this strategy.

Definition 2.7.8. *The correlated strategy μ in $\Delta(S)$ is an equalizing correlated equilibrium (Sorin¹⁸, 1998) if μ is both an equalizing distribution and a correlated equilibrium.*

Clearly, any Nash equilibrium is an equalizing correlated equilibrium, but the converse is false (see example 2.7.10).

Proposition 2.7.9. *Let μ be a correlated strategy of an iterative dual reduction G^r of G . If μ is an equalizing distribution (resp. equalizing correlated equilibrium) of G^r then the G -equivalent correlated strategy $\bar{\mu}$ is an equalizing distribution (resp. equalizing correlated equilibrium) of G .*

Proof. Let s_i be a pure strategy of player i that has positive marginal probability in $\bar{\mu}$. There exists a minimal α_i -absorbing set B_i such that s_i belongs to B_i . Let σ_{s_i} be the α_i -invariant strategy with support B_i . Since $\bar{\mu}(s_i \times S_{-i})$ is positive, it follows that $\mu(\sigma_{s_i} \times (S/\alpha)_{-i})$ is positive too. Furthermore,

(i) $U_i(\mu) = U_i(\bar{\mu})$;

(ii) the conditional probability $\bar{\mu}(\cdot|s_i) \in \Delta((S/\alpha)_{-i})$ is the conditional probability induced on S_{-i} by $\mu(\cdot|\sigma_{s_i})$. That is, if s_j is α_j -recurrent for all j in $I \setminus \{i\}$, then:

$$\bar{\mu}(s_{-i}|s_i) = \mu(\sigma_{s_{-i}}|\sigma_{s_i}) \left(\prod_{j \in I \setminus \{i\}} \sigma_{s_j}(s_j) \right) \quad \text{where } \sigma_{s_{-i}} = \times_{j \in I \setminus \{i\}} \sigma_{s_j}$$

Otherwise, i.e. if s_j is α_j -transient for some j in $I \setminus \{i\}$, then $\bar{\mu}(s_{-i}|s_i) = 0$.

Using (i) and (ii) we get:

$$\sum_{\sigma_{-i} \in (S/\alpha)_{-i}} \mu(\sigma_{-i}|\sigma_{s_i}) U_i(\sigma_{-i}, \sigma_{s_i}) = U_i(\mu) \Rightarrow \sum_{s_{-i} \in S_{-i}} \bar{\mu}(s_{-i}|s_i) U_i(s) = U_i(\bar{\mu})$$

Thus if μ is an equalizing distribution, then so is $\bar{\mu}$. Together with theorem 2.4.5, this implies that if μ is both an equalizing distribution and a correlated equilibrium, then so is $\bar{\mu}$. \square

¹⁸Sorin uses the expression *distribution equilibrium*.

The following example illustrates proposition 2.7.9:

Example 2.7.10.

	x_2	y_2	z_2			
x_1	2, 0	0, 2	0, -3			
y_1	0, 1	1, 0	0, 0			
z_1	-3, 0	0, 0	1, 1			

	σ_{B_2}	z_2
σ_{B_1}	2/3, 2/3	0, -1
z_1	-1, 0	1, 1

Let G denote the game on the left. Consider the deviation vector α such that for $i = 1, 2$:

$$\alpha_i(x_i|x_i) = 2/3, \alpha_i(y_i|x_i) = 1/3; \alpha_i(x_i|y_i) = 1/6, \alpha_i(y_i|y_i) = 5/6; \alpha_i(z_i|z_i) = 1,$$

and all other $\alpha_i(t_i|s_i)$ are zero. It is easily checked that α is a dual vector. The minimal α_i -absorbing sets are $B_i = \{x_i, y_i\}$ and $B'_i = \{z_i\}$. The α -reduced game G/α is the game on the right, where the α_i -invariant strategy σ_{B_i} is $(\frac{1}{3}; \frac{2}{3}; 0)$. Consider the distribution μ on S/α (below, right)¹⁹. This is an equalizing correlated equilibrium of G/α . Therefore, the G -equivalent distribution $\bar{\mu}$ (below, left) is an equalizing correlated equilibrium of G .

$\bar{\mu}$	1/24	1/12	1/24
	1/12	1/6	1/12
	1/24	1/12	3/8

μ	3/8	1/8
	1/8	3/8

We conclude this section by noting that theorem 2.4.5 has the following converse: Let α be a dual vector and $\mu \in \Delta(S/\alpha)$. If μ is not a correlated (resp. Nash) equilibrium of G/α , then the distribution on S induced by μ is not a correlated (resp. Nash) equilibrium of G . This is because the players have more options for deviating in G than in G/α . Similarly, it is easy to show that if μ is not an equalizing distribution (resp. equalizing correlated equilibrium) of G/α , then the distribution on S induced by μ is not an equalizing distribution (resp. equalizing correlated equilibrium) of G .

2.7.3 Elimination of strategies and equilibria

This section identifies classes of strategies or of equilibria which are always (or never) eliminated in full dual reductions.

¹⁹We represent correlated strategies in tables. For instance, $\mu(\sigma_{B_1}, z_2) = 1/8$.

First, note that if a strategy is weakly dominated, it is eliminated in some dual reductions (proposition 2.4.4), but not necessarily in full dual reductions:

Example 2.7.11.

	x_2	y_2
x_1	1, 1	1, 0
y_1	1, 0	0, 0

In the above game, μ is a correlated equilibrium if and only if y_2 is not played in μ . That is, $\mu(x_1, y_2) = \mu(y_1, y_2) = 0$. Therefore y_1 jeopardizes x_1 , and reciprocally. It follows that, in all full dual reductions, x_1 and y_1 are grouped together hence y_1 is not eliminated.

This raises the following question: except strictly dominated strategies, are there other classes of strategies that are always eliminated in full dual reductions or in iterative dual reductions? The answer is positive: strategies that are not played in correlated equilibrium are eliminated. Formally,

Proposition 2.7.12. (i) *Assume that the strategy profile $s \in S$ has probability zero in all correlated equilibria. Then s is eliminated in all full (resp. elementary iterative) dual reductions.*

(ii) *Assume that the pure strategy s_i of player i has marginal probability zero in all correlated equilibria. Then s_i is eliminated in all full (resp. elementary iterative) dual reductions.*

Proof. First, let α be an interior dual vector of the game G . It follows from proposition 2.5.5 and lemma 2.6.6 that in G/α all strategies and strategy profiles with probability zero in all correlated equilibria have been eliminated. Since an interior dual vector is full, and since in all full dual reductions the same strategies and strategy profiles are eliminated, the result for full dual reductions follows.

Second, if G^e is an elementary iterative dual reduction of G , it has a correlated equilibrium with full support, which induces a correlated equilibrium of G . Therefore all pure strategy profiles (resp. pure strategies) of G that belong to the support of a pure strategy profile (resp. pure strategy) of G^e have positive probability in some correlated equilibrium²⁰. \square

²⁰I owe my understanding of this point to Roger Myerson.

Let G^* denote the game obtained from G by deleting all pure strategies that have marginal probability zero in all correlated equilibria. Proposition 2.7.12 suggests that G and G^* have the same full dual reductions, but this is not so:

Example 2.7.13.

$$\begin{array}{cc} & \begin{array}{cc} x_2 & y_2 \end{array} \\ \begin{array}{c} x_1 \\ y_1 \end{array} & \begin{array}{cc} 1, 1 & 0, 1 \\ 0, 1 & 1, 0 \end{array} \end{array} \qquad \begin{array}{cc} & \begin{array}{cc} x_2 & y_2 \end{array} \\ \begin{array}{c} x_1 \\ y_1 \end{array} & \begin{array}{cc} 1, 1 & 0, 1 \\ 1, 1 & 0, 1 \end{array} \end{array}$$

Let G denote the left game. Then G^* is the game on the right. In G^* any mixed strategy profile is a Nash equilibrium. In G , a mixed strategy profile σ is a Nash equilibrium if and only if $\sigma_1(y_1) = 0$ and $\sigma_2(y_2) \leq 1/2$. Thus G^* has more Nash equilibria than G .

In any full dual reduction of G or G^* there is a single strategy profile, which corresponds to a Nash equilibrium σ such that $\sigma_2(x_2)$ and $\sigma_2(y_2)$ are both positive. Conversely, if σ is a Nash equilibrium of G (resp. G^*) such that $\sigma_2(x_2)$ and $\sigma_2(y_2)$ are both positive, then there exists a full dual vector α of G (resp. G^*) such that $S/\alpha = \{\sigma\}$ (resp. $S^*/\alpha = \{\sigma\}$). The set of full dual reductions of G is thus strictly included in the set of full dual reductions of G^* .

The following proposition shows that if a pure strategy of player i is redundant (from the point of view of player i) then it is either eliminated or grouped with the strategies making it redundant:

Proposition 2.7.14. *Assume that there exist i in I , s_i in S_i and σ_i in $\Delta(S_i)$ such that*

$$U_i(s_i, s_{-i}) \leq U_i(\sigma_i, s_{-i}) \quad \forall s_{-i} \in S_{-i} \quad (2.7.2)$$

Then in any full dual reduction, either s_i is eliminated or it is grouped with the strategies in the support of σ_i .

Proof. Equation (2.7.2) implies that every pure strategy in the support of σ_i jeopardizes s_i . The result follows. \square

Now consider elimination of equilibria. It is shown in section 2.10 that even completely mixed, hence proper Nash equilibria may be eliminated in full dual reductions. By contrast:

Proposition 2.7.15. *Strict correlated equilibria cannot be eliminated, not even in an iterative dual reduction.*

Proof. If μ is a strict correlated equilibrium, a strategy that has positive marginal probability in μ cannot be jeopardized by another strategy. It follows that in any dual reduction of G , all pure strategies of G used in μ remain as pure strategies. Furthermore, as the player's options for deviating are more limited in the reduced game than in G , μ is a fortiori a strict correlated equilibrium of the reduced game. Inductively, in any iterative dual reduction $G/\alpha^1/\dots/\alpha^m$ of G , all strategies used in μ are available and μ is still a strict correlated equilibrium. \square

The proof shows that a pure strategy that has positive marginal probability in some strict correlated equilibrium can never be eliminated nor grouped with other strategies. This generalizes the fact that elementary games cannot be reduced.

2.8 Properties in specific classes of games

In this section the additional properties of dual reduction in several classes of games are studied. It is shown that games that are best-response equivalent to zero-sum games and games with a unique correlated equilibrium are reduced in games with a single strategy profile by full dual reduction. We also show that if a game has some symmetry (e.g. cyclic symmetry), then there exist full dual reductions which preserve this symmetry. Finally, generic 2×2 games are analysed.

2.8.1 Two-player zero-sum games

It is easy to show that any dual reduction of a zero-sum game is a zero-sum game with the same value. Furthermore, dual vectors may be built easily from optimal strategies:

Proposition 2.8.1. *Let G denote a two-player zero-sum game and α a deviation vector. If for all $i = 1, 2$ and for all s_i in S_i , the mixed strategy $\alpha_i * s_i$ is an optimal strategy, then α is a dual vector.*

Proof. Let s be a pure strategy profile. Since $\alpha_1 * s_1$ is optimal, it follows that $U_1(\alpha_1 * s_1, s_2) \geq v$, where v is the value of the game. Similarly, $U_2(s_1, \alpha_2 * s_2) \geq$

$-v$. Therefore, $U_1(\alpha_1 * s_1, s_2) + U_2(s_1, \alpha_2 * s_2) \geq 0 = U_1(s) + U_2(s)$. Since this holds for all s in S , it follows that α is a dual vector. \square

This implies that:

Corollary 2.8.2. *For every Nash equilibrium σ of a zero-sum game, there exists a dual reduction in which the reduced set of strategy profiles is the singleton $\{\sigma\}$.*

Proof. In the particular case of proposition 2.8.1 where $\alpha_i * s_i = \sigma_i$ for every player i and every s_i in S_i , the only α_i -invariant strategy is σ_i . Therefore, the reduced set S_i/α_i of pure strategies of player i is the singleton $\{\sigma_i\}$. \square

If we restrict attention to full dual reduction, then we also obtain that in the reduced game there is a unique strategy profile; but this strategy profile now corresponds to a profile of optimal strategies *with maximal support*:²¹

Proposition 2.8.3. *If G is best-response equivalent to a two-player zero-sum game then in all full dual reductions of G all the strategies of player i with positive probability in some correlated equilibrium are grouped together, and his other pure strategies are eliminated.*

Proof. It is shown in chapter 6 that a game best-response equivalent to a two-player zero-sum game is pre-tight. Therefore it suffices to prove that the result holds for pre-tight games, which we now do: It follows from proposition 2.7.12 that, in a full dual reduction, all strategies with probability zero in all correlated equilibria must be eliminated. Furthermore, in a pre-tight game, all other pure strategies (of a given player) jeopardize each other; therefore, in a full dual reduction, they are all grouped together in a single mixed strategy. The result follows. \square

The following example shows that if G is best-response equivalent to a zero sum game, then the payoffs in a full dual reduction may depend on the reduction:

Example 2.8.4.

	x_2	y_2	z_2		x_2	y_2	z_2
x_1	0, 0	0, 0	0, 0		x_1	1, 1	0, 1
y_1	0, 0	1, -1	-1, 1		y_1	1, 0	1, -1
z_1	0, 0	-1, 1	1, -1		z_1	1, 0	-1, 1

²¹Conversely, it is easily checked that any profile of optimal strategies with maximal support may be obtained as a result of a full dual reduction.

Let G (resp. G') denote the game on the left (resp. right). G is zero-sum and G' is best-response equivalent to G . For $0 \leq \epsilon \leq 1$, let σ_i^ϵ denote the optimal strategy of player i such that: $\sigma_i^\epsilon(x_i) = \epsilon$ and $\sigma_i^\epsilon(y_i) = \sigma_i^\epsilon(z_i) = (1 - \epsilon)/2$. Let $\alpha^{\epsilon, \eta}$ denote the deviation vector such that: $\alpha_1 * x_1 = \alpha_1 * y_1 = \alpha_1 * z_1 = \sigma_1^\epsilon$ and $\alpha_2 * x_2 = \alpha_2 * y_2 = \alpha_2 * z_2 = \sigma_2^\eta$. It is easy to check that α is a dual vector of both G and G' . If $0 < \epsilon < 1$ and $0 < \eta < 1$, then α is full, the reduced strategy space $S'/\alpha^{\epsilon, \eta}$ is the singleton $(\sigma_1^\epsilon, \sigma_2^\eta)$ and the associated payoff is (η, ϵ) .

To conclude this section, we show that, within the set of games with a completely mixed Nash equilibrium, some rescalings of zero-sum games are characterized by their dual vectors. We first need a definition: let G and G' be two games with the same sets of players and strategies.

Definition 2.8.5. G' is a partial rescaling of G if for every player i in I , there exists a function $f_i : S_{-i} \rightarrow \mathbb{R}$ such that:

$$U'_i(s) = U_i(s) + f_i(s_{-i}), \quad \forall s \in S$$

(The difference with definition 2.7.3 is that we require $a_i = 1$ for all i in I .)

Let G be a two-player game with a completely mixed Nash equilibrium σ . Define the deviation vector α by $\alpha_i * s_i = \sigma_i$ for every player i in $\{1, 2\}$ and every pure strategy s_i in S_i .

Proposition 2.8.6. The deviation vector α is a dual vector if and only if G is a partially rescaled zero-sum game.

Proof. If G is a partial rescaling of \tilde{G} , then G and \tilde{G} have the same dual vectors. Therefore, it follows from proposition 2.8.1 that if G is a partially rescaled zero-sum game, then α is a dual vector.

Conversely, assume that α is a dual vector. This implies that for every mixed strategy profile τ ,

$$[U_1(\sigma_1, \tau_2) - U_1(\tau)] + [U_2(\tau_1, \sigma_2) - U_2(\tau)] \geq 0$$

Therefore (σ_1, σ_2) is a Nash-Pareto pair in the sense of Hofbauer and Sigmund (1998, section 11.4, page 135). By theorem 11.4.2 of Hofbauer and Sigmund (1998), this implies that G is a rescaled zero-sum game.

That is, there exists a zero-sum game \tilde{G} , and, for every player i in $\{1, 2\}$ a positive constant a_i and a function $h_i : S_{-i} \rightarrow \mathbb{R}$ such that:

$$U_i(s) = a_i \cdot \tilde{U}_i(s) + h_i(s_{-i}), \quad \forall s \in S$$

Without loss of generality, assume $a_2 = 1$ and $a_1 \geq 1$. If furthermore $a_1 = 1$, then G is a partially rescaled zero-sum game. So we may also assume $a_1 \neq 1$.

Let s be a pure strategy profile. Since G has a completely mixed Nash equilibrium, it follows that

$$g(s, \alpha) = \tilde{g}(s, \alpha) = 0 \quad (2.8.1)$$

(where \tilde{g} is the payoff function in the auxiliary game associated to \tilde{G})

Moreover, since $\alpha_i * s_i = \sigma_i$ for every i in $\{1, 2\}$, it follows that $g(s, \alpha) = [U_1(s) - U_1(\sigma_1, s_2)] + [U_2(s) - U_1(s_1, \sigma_2)] = a_1[\tilde{U}_1(s) - \tilde{U}_1(\sigma_1, s_2)] + [\tilde{U}_2(s) - \tilde{U}_1(s_1, \sigma_2)]$. Since $\tilde{g}(s, \alpha) = 0$ this yields:

$$g(s, \alpha) = (a_1 - 1)[\tilde{U}_1(s) - \tilde{U}_1(\sigma_1, s_2)] \quad (2.8.2)$$

Since σ is a Nash equilibrium of G , hence of \tilde{G} , and since s_2 belongs to the support of σ_2 , it follows that $\tilde{U}_1(\sigma_1, s_2) = v$ where v is the value of \tilde{G} . Since, by assumption, $a_1 \neq 1$, it follows from (2.8.1) and (2.8.2) that $\tilde{U}_1(s) = v$. Since this holds for every s in S , it follows that:

$$U_1(s) = \tilde{U}_1(s) + [h_1(s_2) + (a_1 - 1).v] \quad \forall s \in S$$

Therefore, G is a partial rescaling of \tilde{G} . This completes the proof. \square

2.8.2 Games with a unique correlated equilibrium

If G has a unique Nash equilibrium σ , then any iterative dual reduction of G has a unique Nash equilibrium, which induces σ in G ; but the strategy space need not be reducible to σ . In particular, it may be that a (nontrivial) game has a unique, pure Nash equilibrium but is nevertheless elementary, hence cannot be reduced. See (Nau and McCardle, 1990, example 4). By contrast,

Proposition 2.8.7. *Assume that G has a unique correlated equilibrium σ . Then σ is a Nash equilibrium, hence it may be seen as a mixed strategy profile. Let G^r be the reduced game in which the only strategy profile is σ and the payoff for player i is $U_i(\sigma)$. Any full (resp. elementary iterative) dual reduction of G is equal to G^r . In particular, G has a unique full dual reduction.*

Proof. The part of the proposition concerning elementary iterative dual reductions follows from the proof of proposition 2.5.1. The part of the proposition concerning full dual reduction follows from the fact that a game with a unique correlated equilibrium is pre-tight, and from the proof of proposition 2.8.3. \square

2.8.3 Symmetric Games

This section shows that if a game is symmetric with respect to some permutations of the set of players, then it may be reduced in a way respecting these symmetries. Some definitions are needed. Let P be a set of permutations of the set of players. Let G be a game for which, for every player i in I and every permutation p in P , player i and player $p(i)$ have the same number of pure strategies. Label the strategies of each player with integers so that the set of pure strategies of player i is $S_i = \{1, 2, \dots, m_i\}$, where $m_i = \#S_i$ (cardinal of S_i).

If $s = (k_1, k_2, \dots, k_n)$ and p a permutation of the set of players, let $p(s) = (k'_1, k'_2, \dots, k'_n)$ denote the strategy profile such that, for all i in I , $k'_{p(i)} = k_i$.

Definition 2.8.8. *The game is p -symmetric if*

$$\forall i \in I, U_{p(i)}(p(s)) = U_i(s)$$

The²² game is P -symmetric if it is p -symmetric for every p in P .

Proposition 2.8.9. *If the game G is P -symmetric then there exists an interior dual vector α such that G/α is P -symmetric.*

Proof. Without loss of generality, assume that P is the largest set of permutations such that G is P -symmetric. This implies that P is stable by internal permutation.²³ That is,

$$\forall p \in P, \{p \circ \tilde{p}, \tilde{p} \in P\} = P \tag{2.8.3}$$

Let²⁴ α be a dual vector. For every permutation p in P , let α^p denote the deviation vector such that:

$$\forall i \in I, \forall j \in S_i, \forall k \in S_i, \alpha^p_{p(i)}(k|j) = \alpha_i(k|j)$$

Define a deviation vector α to be p -symmetric if $\alpha = \alpha^p$ and to be P -symmetric if it is p -symmetric for every p in P . Let $\#P$ denote the cardinal

²²This means that if for every i in I , player $p(i)$ plays tomorrow as player i played today, then the payoff of player $p(i)$ tomorrow is the same as the payoff of player i today.

²³ P is a subgroup of the group of permutations of I .

²⁴To prove (2.8.3) The point is that if a game is both p -symmetric and \tilde{p} -symmetric then it is $p \circ \tilde{p}$ -symmetric.

of P . It follows from (2.8.3) that for any deviation vector α , the deviation vector

$$\bar{\alpha} = \frac{\sum_{p \in P} \alpha^p}{\#P}$$

is P -symmetric. Furthermore, if α is an (interior) dual vector then so is α^p for every p in P , hence so is $\bar{\alpha}$. Finally, it is easily checked that if α is P -symmetric then G/α is P -symmetric too. The result follows. \square

Note that nonsymmetric²⁵ games may also have symmetric full dual reductions, even if all strategies are undominated (take any nonsymmetric zero-sum game with value 0; see also example 2.7.10). Moreover, a symmetric game may have nonsymmetric full dual reductions (take any symmetric game in which every player is a dummy, i.e. has no influence on its own payoffs).

2.8.4 Generic 2×2 games

Proposition 2.8.10. *Let G be a 2×2 game such that a player is never indifferent between two different strategy profiles. That is,*

$$s \neq s' \Rightarrow U_i(s) \neq U_i(s'), \quad \forall s \in S, \forall s' \in S, \forall i \in \{1, 2\}$$

Then either G is elementary or G has a unique correlated equilibrium (in which case proposition 2.8.7 apply).

Proof. This follows from straightforward computations. The first case corresponds to games with three Nash equilibria: two pure and one completely mixed; the second case to games with either a strictly dominant strategy (for at least one of the players) or a unique, completely mixed Nash equilibrium. \square

2.9 The issue of uniqueness

In this section, we first show that, even if only full dual reductions are used, there might be multiple ways to reduce a game. We then show that generic two-player games have a unique sequence of iterative full dual reductions.

²⁵A game is *symmetric* if it is p -symmetric for any permutation p , and *nonsymmetric* otherwise.

Example 2.8.4 shows that a game may have several full dual reductions. This arises systematically when a player is indifferent between some of his strategies, and when these strategies are not eliminated:

Proposition 2.9.1. *Assume that player i is indifferent between s_i and t_i , i.e. $U_i(s) = U_i(s_{-i}, t_i)$ for all s_{-i} in S_{-i} . Then (i) for any $0 \leq \epsilon \leq 1$ there exists a dual reduction that simply consists in grouping s_i and t_i in the strategy σ_i such that $\sigma_i(s_i) = \epsilon$ and $\sigma_i(t_i) = 1 - \epsilon$; (ii) if s_i is not eliminated in full dual reductions, then there exists an infinity of full dual reductions.*

Proof. To prove (i) take as dual vector α : $\alpha_i(s_i|s_i) = \alpha_i(s_i|t_i) = \epsilon$, $\alpha_i(t_i|s_i) = \alpha_i(t_i|t_i) = 1 - \epsilon$ and all the other $\alpha_j(t_j|s_j)$ as in the trivial deviation vector. We now prove (ii): Assume that s_i is not eliminated in full dual reductions and let α be a full dual vector. For $0 < \lambda \leq 1$, define the dual vector α^λ by: $\alpha_i^\lambda(s_i|s_i) = \lambda\alpha_i(s_i|s_i)$, $\alpha_i^\lambda(t_i|s_i) = \alpha_i(t_i|s_i) + (1 - \lambda)\alpha_i(s_i|s_i)$ and all other $\alpha_j^\lambda(t_j|s_j)$ as in α . Since α is full and α and α^λ are positive in the same components, α^λ is full too. Therefore, there exists an α_i^λ -invariant strategy σ_i^λ such that $\sigma_i^\lambda(s_i) > 0$. We claim that if $\lambda' \neq \lambda$, σ_i^λ is not $\alpha_i^{\lambda'}$ -invariant (proof below). This implies that if $\lambda' \neq \lambda$, α^λ and $\alpha^{\lambda'}$ induce different full dual reductions. Therefore there exists an infinity of different full dual reductions.

To prove the claim, note that if σ_i^λ is $\alpha_i^{\lambda'}$ -invariant, then

$$\sum_{r_i \in S_i \setminus \{s_i\}} \alpha_i^{\lambda'}(s_i|r_i)\sigma_i^\lambda(r_i) = [1 - \alpha_i^{\lambda'}(s_i|s_i)]\sigma_i^\lambda(s_i)$$

But if $\lambda' \neq \lambda$:

$$\begin{aligned} \sum_{r_i \in S_i \setminus \{s_i\}} \alpha_i^{\lambda'}(s_i|r_i)\sigma_i^\lambda(r_i) &= \sum_{r_i \in S_i \setminus \{s_i\}} \alpha_i^\lambda(s_i|r_i)\sigma_i^\lambda(r_i) \\ &= [1 - \alpha_i^\lambda(s_i|s_i)]\sigma_i^\lambda(s_i) \neq [1 - \alpha_i^{\lambda'}(s_i|s_i)]\sigma_i^\lambda(s_i) \end{aligned}$$

□

Multiplicity of dual reductions may also arise if a player is indifferent between a pure and a mixed strategy (example 2.8.4) or if a player *becomes* indifferent between some of his strategies, after strategies of some other player have been eliminated (example 2.7.13). These are non-generic phenomena. We prove in this section that, for any positive integer m , two-player games generically have a unique iterative full dual reduction of depth m . We first

show that there are severe restrictions on the ways strategies may be grouped together in dual reductions:

Notation: for all i in I , let $B_i \subseteq S_i$ and let $B = \times_{i \in I} B_i$. We denote by $G_B = (I, (B_i)_{i \in I}, (U_i)_{i \in I})$ the game obtained from G by reducing the pure strategy set of player i to B_i , for all i in I .

Proposition 2.9.2. *Let α be a dual vector. For each i in I , let $B_i \subseteq S_i$ denote a minimal α_i -absorbing set and $B = \times_{i \in I} B_i$. Let σ_{B_i} denote the unique α_i -invariant strategy of player i with support in B_i and $\sigma_B = (\sigma_{B_i})_{i \in I}$. We have: σ_B is a completely mixed Nash equilibrium of G_B .*

Proof. By minimality of B_i , the support of σ_{B_i} is exactly B_i so σ_B is completely mixed. Moreover, let $\sigma_{B_{-i}} = \times_{j \in I \setminus \{i\}} \sigma_{B_j}$. Against $\sigma_{B_{-i}}$, player i is indifferent between the strategies of the minimal absorbing set B_i (proposition 2.4.7). Therefore, if player i is restricted to the strategies in B_i , σ_{B_i} is a best response to $\sigma_{B_{-i}}$ \square

Define α and σ_{B_i} as in proposition 2.9.2 and assume α full. If G_B has a unique completely mixed Nash equilibrium, then for any full dual vector β , the β_i -invariant strategy with support in B_i must be σ_{B_i} . So proposition 2.9.2 has the following corollary:

Corollary 2.9.3. *If for every product $B = \times_{i \in I} B_i$ of subsets B_i of S_i , G_B has at most one completely mixed Nash equilibrium, then there exists a unique full dual reduction.*

In the remainder of this section, G is a two-player game. To show that, generically, two-player games have a unique sequence of iterative full dual reductions, some suitable notions of genericity are needed:²⁶

Definition 2.9.4. *G is generic if for all Nash equilibria σ the supports of σ_1 and σ_2 have same cardinal²⁷. G is locally generic if G is generic and if any game obtained from G by deleting some pure strategies is generic.*

The following notion is quite ad hoc:

²⁶A property is sometimes called generic when it occurs for an open set of games. What we mean here by a generic property, is that it holds for a set of games with Lebesgue measure 1.

²⁷Any game which is nondegenerate in the sense of von Stengel (2002, def. 2.6 and thm 2.10) is generic in this sense.

Definition 2.9.5. *G is 2-generic if for any subset B_1 of S_1 of cardinal two or more, and for any disjoint subsets B_2 and B'_2 of S_2 , the following holds: if σ and σ' are completely mixed Nash equilibria of, respectively, $G_{B_1 \times B_2}$ and $G_{B_1 \times B'_2}$ then $\sigma_1 \neq \sigma'_1$ (the same mixed strategy cannot be a completely mixed Nash equilibrium strategy of player 1 both on $B_1 \times B_2$ and on $B_1 \times B'_2$). The notion of 1-genericity is defined similarly. A bimatrix game is $*$ -generic if it is both 1-generic and 2-generic.*

A bimatrix game in which players 1 and 2 have respectively p and q pure strategies is given by two $p \times q$ payoff matrices, thus it may be viewed as a point in $\mathbb{R}^{pq} \times \mathbb{R}^{pq}$. It may be shown that the set of $p \times q$ bimatrix games which are both locally generic and $*$ -generic contains an open and dense subset of $\mathbb{R}^{pq} \times \mathbb{R}^{pq}$. The two next propositions follow from proposition 2.9.2:

Proposition 2.9.6. *A locally generic bimatrix game has a unique full dual reduction.*

Proof. Locally generic bimatrix games check the conditions of corollary 2.9.3 □

Proposition 2.9.7. *If G is both locally generic and $*$ -generic, there are only three possibilities:*

1. *G is elementary*
2. *In all dual reductions of G , some strategies are eliminated, but no strategies are grouped together.*
3. *In any full dual reduction of G the reduced strategy space S/α is a singleton.*

Proof. Assume that G is not elementary and let α be a nontrivial dual vector. Assume that some strategies of player 1 (for instance) are grouped together. That is, there exists a minimal α_1 -absorbing set B_1 with at least two elements. Let B_2 and B'_2 be minimal α_2 -absorbing sets. Let σ_{B_1} denote the α_1 -invariant strategy with support in B_1 . Define σ_{B_2} and $\sigma_{B'_2}$ similarly. By proposition 2.9.2, σ_{B_1} is a Nash equilibrium strategy both of $G_{B_1 \times B_2}$ and of $G_{B_1 \times B'_2}$. Since G is $*$ -generic, this implies $B_2 = B'_2$. Therefore, there is a unique minimal α_2 -absorbing set, B_2 . That is, S_2/α_2 is a singleton. Moreover, since G is locally generic, B_1 and B_2 have same cardinal. Thus B_2 has at least two elements. Therefore, by the above reasoning, the strategy set of player 1 in G/α is also a singleton. This completes the proof. □

As an immediate corollary of proposition 2.9.7 and definitions 2.9.4 and 2.9.5 we get:

Corollary 2.9.8. *If G is both locally generic and $*$ -generic then any dual reduction of G is both locally generic and $*$ -generic.*

As an immediate corollary of proposition 2.9.6 and corollary 2.9.8 we obtain:

Theorem 2.9.9. *If G is both locally generic and $*$ -generic, then for any positive integer m , G has a unique iterative full dual reduction of depth m .*

2.10 Dual reduction as a refinement concept

When a game is reduced by dual reduction, not all equilibria of the original game remain (in the sense of being induced by an equilibrium of the reduced game), and in available examples of dual reduction, those equilibria that remain tend to have better stability properties than those that are eliminated. This suggests a refinement concept, both for correlated equilibria and Nash equilibria: retaining as solutions of the game only those equilibria that remain after dual reduction or iterative dual reduction of the game. In this section, correlated equilibria remaining after dual reduction are compared to acceptable correlated equilibria (Myerson, 1986) and also to perfect correlated equilibrium distributions (Dhillon and Mertens, 1996). The relevance of the refinement concept induced by dual reduction is then discussed.

2.10.1 Dual reduction and elimination of unacceptable pure strategies

There are similarities in the ways dual reduction and elimination of acceptable pure strategies are defined: in particular, the *aggregate incentive value* of s for the set of players I defined in (Myerson, 1986, p.141, equation (3.3)) is exactly the payoff $g(s, \alpha)$ defined in section 2.3. This motivates our comparison between the two concepts. We show by means of example that none of these refinement concepts is more stringent than the other and discuss further some of the differences between these concepts in the next section. Some notations and definitions are needed [unless stated otherwise, all definitions are from Myerson (1986); most of the phrasing is borrowed to Myerson (1986) and Dhillon and Mertens (1996); see also (Myerson, 1991)]:

Let $J \subseteq I$. If J is nonempty we let

$$S_J = \times_{i \in J} S_i$$

(so $S_I = S$), and we let $S_\emptyset = \{\emptyset\}$. If s is in S and t_J in S_J then (s_{-J}, t_J) denotes the strategy profile in which player i plays t_i if i is in J and s_i otherwise.

Definition 2.10.1. An ϵ -correlated strategy η is a lottery choosing a vector of “recommended” pure strategies (i.e. a point in S), a coalition J of trembling players, and a vector of trembles (i.e. a point in S_J) for those players (hence, formally, it is a probability distribution over $S \times (\cup_{J \subseteq I} J)$) such that:

(a) Given any vector of recommendations, the conditional probability of every coalition of trembling players and every vector of trembles for these players is strictly positive.

(b) Given any vector of recommendations s , any subset J of players not including player i and any vector of trembles t_J for those players: given that the coalition of trembling players is either J or $J \cup \{i\}$ and that the players of J tremble to t_J , the conditional probability of i also trembling is at most ϵ .

Let η be an ϵ -correlated strategy. Consider the extended game in which each player is first informed of his recommended action; next the non-trembling players are asked to move - while the trembling players are forced to move using the selected trembles. The ϵ -correlated strategy η is an ϵ -correlated equilibrium if, in this extended game, the obedient strategies form a Nash equilibrium.

A correlated strategy $\mu \in \Delta(S)$ is an *acceptable correlated equilibrium* if it is a limit ($\epsilon \rightarrow 0$) of distributions (i.e. marginal distributions on S) of ϵ -correlated equilibria. That is, if for all positive ϵ there exists some ϵ -correlated equilibrium η^ϵ such that for all s in S : $\lim_{\epsilon \rightarrow 0} \eta^\epsilon(s, \emptyset) = \mu(s)$, where $\eta^\epsilon(s, \emptyset)$ is the probability that s is recommended and that no player trembles. Acceptable correlated equilibria are correlated equilibria (Myerson, 1986, thm 1).

A pure strategy s_i is *acceptable* if, for every $\epsilon > 0$, there exists some ϵ -correlated equilibrium η such that

$$\sum_{s_{-i} \in S_{-i}} \eta(s, \emptyset) > 0$$

(that is, in Myerson’s (1986) terms, “if s_i can be rationally used when the probabilities of trembling are infinitesimal”).

The *acceptable residue* $R(G)$ of a game G is the game obtained from G by eliminating all the unacceptable pure strategies. Myerson (1986, theorems 2 and 4) shows that the acceptable correlated equilibria are exactly the correlated equilibria of the acceptable residue $R(G)$.

As dual reduction, elimination of unacceptable pure strategies may be iterated. A pure strategy is *predominant* if it remains after iterative elimination of unacceptable pure strategies, and correlated equilibria in which only predominant strategies are played are called *predominant*.

We now compare dual reduction and elimination of unacceptable pure strategies. We first need a lemma:

Lemma 2.10.2. *If there exists a correlated equilibrium with full support then all pure strategies are acceptable and predominant.*

Proof. Assume that there exists a correlated equilibrium μ with full support. By theorem 2 of Myerson (1986), if μ is acceptable, then any pure strategy is acceptable, hence any pure strategy is predominant. Therefore, it suffices to show that μ is acceptable. The trick is that, because μ has full support, it is possible to find trembles mimicking μ , so that whoever the players trembling, a non-trembling player always faces the same conditional probabilities given his signal as in μ .

More precisely, assume that there exists some ϵ -correlated strategy η such that:

$$\forall J \subseteq I, \forall t_J \in S_J, \forall s \in S, \eta(s, t_J) = K(J, \epsilon) \mu(s_{-J}, t_J) \quad (2.10.1)$$

where K is a positive constant that depends only on J and on ϵ (but not on s_{-J}). Thus, given any coalition J of trembling players, any vector t_J of trembles assigned to J , and any strategy profile s , the probability in η that (s_{-J}, t_J) will be played as a result of the players being recommended s , the players of $I \setminus J$ not trembling, and the players of J trembling to t_J , is proportional to the probability of (s_{-J}, t_J) in μ . The total probability in η that J and only J trembles and that (s_{-J}, t_J) is played is:

$$\sum_{r_J \in S_J} \eta((s_{-J}, r_J), t_J) = K'(J, \epsilon) \mu(s_{-J}, t_J)$$

where K' is a positive constant which depends only on J and on ϵ . It follows that for every player $i \notin J$ and every pure strategy s_i in S_i , the expected strategy of the other players in η , given s_i and given that J and only J

trembles, is equal to $\mu(\cdot|s_i)$ (the expected strategy of the other players in μ , given s_i). A fortiori, the expected strategy in η given s_i and given that player i does not tremble is equal to $\mu(\cdot|s_i)$, to which s_i is a best response. It follows that η is an ϵ -correlated equilibrium.

It remains to show that it is possible to find a sequence of ϵ -correlated strategy checking (2.10.1) and such that $\eta(s, \emptyset)$ tends to $\mu(s)$ as ϵ goes to zero. Such a sequence may be built by letting, for all s in S and for some suitable positive normalization constant A ,

$$\eta(s, \emptyset) = A \times \mu(s)$$

and, inductively, if the cardinal of $J \subseteq I$ is $m + 1$:

$$\eta(s, r_J) = \frac{\epsilon}{1 - \epsilon} A_m \times \mu(s_{-J}, r_J)$$

with

$$A_m = \min_{t \in S} \min_{L \in J: \text{Card } L = m} \min_{r_L \in S_L} \eta(t, r_L)$$

□

Lemma 2.10.2 implies that the class of games in which all pure strategies are acceptable is strictly larger than the class of elementary games. Consider for instance the following game of coordination where, moreover, player 2 has an outside option:

Example 2.10.3.

	x_2	y_2	z_2
y_1	0, 0	1, 1	-1, -1
z_1	0, 0	-1, -1	1, 1

In this game, playing each strategy with equal probability is a completely mixed Nash equilibrium. Therefore, by lemma 2.10.2, all strategies are acceptable and predominant. However, x_2 is eliminated in any nontrivial dual reduction. (To prove this, note that x_2 is equivalent to $\frac{1}{2}y_2 + \frac{1}{2}z_2$; this implies that y_2 and z_2 jeopardize x_2 . Furthermore y_i and z_i must be invariant under any dual vector because they have positive probability in some strict correlated equilibrium. So there is a unique dual reduction, which consists in eliminating x_2 .)

This example shows that dual reduction may eliminate acceptable and even predominant pure strategies. It also shows that dual reduction can

eliminate completely mixed, hence perfect Nash equilibria. Since any perfect Nash equilibrium is a perfect direct correlated equilibrium (Dhillon and Mertens, 1996), it shows that dual reduction may eliminate perfect direct correlated equilibria.

The next example shows that there may be unacceptable pure strategies that no dual reduction eliminates: let G denote the following three-player game, where player 1 chooses the matrix (x_1 or y_1), player 2 the row, and player 3 the column:

Example 2.10.4 (taken from (Myerson, 1986)).

	x_1				y_1		
	x_3	y_3	z_3		x_3	y_3	z_3
x_2	2, 1, 1	0, 2, 0	0, 2, 0	x_2	1, 3, 3	1, 3, 3	1, 3, 3
y_2	0, 0, 2	0, 3, 0	0, 0, 3	y_2	1, 3, 3	1, 3, 3	1, 3, 3
z_2	0, 0, 2	0, 0, 3	0, 3, 0	z_2	1, 3, 3	1, 3, 3	1, 3, 3

Myerson (1986) shows that for every player i in $\{1, 2, 3\}$, the only acceptable strategy is x_i . However, y_1 cannot be eliminated by dual reduction. Indeed, let $s = (y_1, y_2, y_3)$ and α be a dual vector; by definition 2.4.1, $\sum_{i \in I} [U_i(\alpha_i * s) - U_i(s)] \geq 0$; since s is a Nash equilibrium and all unilateral deviations from s by player 1 are strictly detrimental for him, this implies that y_1 is invariant under α .

Note that y_1 may be eliminated by *iterative* dual reduction. Indeed, to prove that y_2, z_2, y_3, z_3 and y_1 are unacceptable, Myerson uses the *codomination system*²⁸ (α^1, α^2) where α^1 and α^2 are the deviation vectors such that:

$$\alpha_i^1(x_i|y_i) = \alpha_i^1(x_i|z_i) = 1 \quad \forall i \in \{2, 3\},$$

$\alpha_1^2(x_1|y_1) = 1$, and all other $\alpha_i^k(t_i|s_i)$ are as in the corresponding trivial deviation vectors. It is easy to check that α^1 is a dual vector of G and α^2 a dual vector of G/α^1 . The only strategy profile remaining in $G/\alpha^1/\alpha^2$ is the strict Nash equilibrium (x_1, x_2, x_3) , thus y_1 has been eliminated. Whether some unacceptable (or non predominant) pure strategies cannot be eliminated by any iterative dual reduction is open.

²⁸For a definition of codomination systems, see (Myerson, 1986).

2.10.2 Discussion

Dual reduction, as inducing a refinement concept, has a number of satisfactory properties: the reduction process is unaffected by positive affine transformations of the utility functions; iteratively strictly dominated strategies are eliminated (since they are not played in correlated equilibrium); redundant strategies are either eliminated or grouped with the strategies making them redundant, strict correlated equilibria are never eliminated; etc. Other properties may seem less satisfactory; for instance, the fact that weakly dominated strategies need not be eliminated in full dual reductions or that different full dual reductions may select different outcomes. But these properties need not be weaknesses, as they are in accordance with the behaviour of evolutionary dynamics. Consider for instance the following two-player symmetric game:

$$\begin{array}{rcccl}
 & x_2 & y_2 & z_2 & \\
 x_1 & 1, 1 & 1, 1 & 1, 0 & \\
 y_1 & 1, 1 & 1, 1 & 0, 0 & \\
 z_1 & 0, 1 & 0, 0 & 0, 0 &
 \end{array} \tag{2.10.2}$$

The Nash equilibria of this game are exactly the strategy profiles σ such that $\sigma_i(z_i) = 0$ for $i = 1, 2$. The Nash equilibria which are limit of an interior solution of the two-population replicator dynamics are the Nash equilibria σ such that, for $i = 1, 2$, both $\sigma_i(x_i)$ and $\sigma_i(y_i)$ are positive²⁹. These are exactly the Nash equilibria which are selected by full dual reduction, while standard refinement concepts would select the equilibrium (x_1, x_2) . In this example, the outcomes selected by dual reduction could not coincide with the outcomes selected by the replicator dynamics if weakly dominated strategies were eliminated by dual reduction³⁰ or if dual reduction was to pinpoint a single equilibrium. Similarly, the fact that, contrary to most refinement concepts, dual reduction may eliminate completely mixed Nash equilibria need not be a weakness: in example 2.10.3, eliminating strategy x_2 seems more reasonable than keeping it in the game, since it is redundant and since the completely mixed Nash equilibria in which strategy x_2 has positive probability are unstable under most evolutionary dynamics.

²⁹The behaviour of the single-population replicator dynamics in this game is described by Weibull (1995); the extension to the two-population replicator dynamics is straightforward.

³⁰Note also that, after full dual reduction, only those weakly dominated strategies with positive probability in some correlated equilibrium may remain in the game. This follows from proposition 2.7.12. In this sense, dual reduction selects between “good” and “bad” weakly dominated strategies.

Dual reduction as a refinement concept suffers however from a severe drawback: it not clearly motivated. For this reason, we believe that dual reduction is primarily interesting as a tool to study correlated equilibria, and not as inducing a refinement concept. If dual reduction is seen, not as inducing a refinement concept, but as a way to simplify a game while retaining its most important elements, then another aspect emerges: dual reduction is too centered on equilibria to capture important out of equilibrium features of some games, such as best-response cycles. For instance, in chapter 10, we study a game in which a unique pure strategy is played in correlated equilibrium but a best-response cycle exists, which corresponds to a robust attractor for the replicator dynamics and many other dynamics. In this game, any nontrivial dual reduction eliminates all strategies that are not played in correlated equilibrium, missing the best-response cycle and the corresponding attractor as a solution of the game.

2.11 Appendix: Direct proofs of existence of correlated equilibria

In this appendix, we review and connect the proofs of existence of correlated equilibria given by Hart and Schmeidler (1989), Nau and McCardle (1990) and Myerson (1997), on which section 2.3 is based.

2.11.1 Hart and Schmeidler's proof

Consider the following two-player, zero-sum, auxiliary game Γ_{HS} (HS is for Hart and Schmeidler): the maximizer chooses a strategy profile $s = (s_1, \dots, s_n)$ in S ; the minimizer chooses a player i in I and a couple of pure strategies (s'_i, t_i) in $S_i \times S_i$. The payoff is $U_i(s) - U_i(s_{-i}, t_i)$ if $s'_i = s_i$ and 0 otherwise. In mixed strategies the maximizer chooses a correlated strategy μ in $\Delta(S)$ and the minimizer a probability distribution ν on triples $(i, s_i, t_i) \in I \times S_i \times S_i$; the expected payoff is then:

$$g_{hs}(\mu, \nu) = \sum_{s \in S} \mu(s) \sum_{i \in I} \sum_{t_i \in S_i} \nu(i, s_i, t_i) [U_i(s) - U_i(s_{-i}, t_i)] \quad (2.11.1)$$

Note that the payoff of the auxiliary game of section 2.3 may be written:

$$g(\mu, \alpha) = \sum_{s \in S} \mu(s) \sum_{i \in I} \sum_{t_i \in S_i} \alpha_i(t_i | s_i) [U_i(s) - U_i(s_{-i}, t_i)]$$

The only difference is that the normalization constraint for ν is:

$$\sum_{i \in I} \sum_{s_i \in S_i} \sum_{t_i \in S_i} \nu(i, s_i, t_i) = 1$$

while the normalization constraint for α is:

$$\sum_{t_i \in S_i} \alpha_i(t_i | s_i) = 1, \quad \forall i \in I, \forall s_i \in S_i$$

In Γ_{HS} , as in the auxiliary game of section 2.3, the correlated strategy μ guarantees 0 if and only if μ is a correlated equilibrium of the original game. Thus, to prove existence of correlated equilibria, it suffices to show that the value of Γ_{HS} is nonnegative. To do so, Hart and Schmeidler could have used the existence of invariant distributions for finite Markov chains (as we did in section 2.3):³¹

Lemma 2.11.1. *Let M be a $m \times m$ stochastic matrix (i.e. nonnegative with columns summing to unity); there exists a probability vector $x = (x_j)_{j=1, \dots, m}$ such that $Mx = x$.*

Instead, they used the following lemma:

Lemma 2.11.2 (Hart and Schmeidler). *Let $(a_{jk})_{1 \leq j, k \leq m}$ be nonnegative numbers. There exists a probability vector $x = (x_j)_{j=1, \dots, m}$ such that, for any vector $u = (u_j)_{j=1, \dots, m}$,*

$$\sum_{j=1}^m x_j \sum_{k=1}^m a_{jk} (u_j - u_k) = 0 \quad (2.11.2)$$

It turns out that:

Proposition 2.11.3. *Lemmas 2.11.1 and 2.11.2 are equivalent (in the sense that one is easily deduced from the other).*

³¹Indeed, let Γ denote the auxiliary game of section 2.3, λ a positive constant and n the number of players. If λ is small enough, any strategy of the minimizer in Γ can be emulated in Γ_{HS} , up to the scaling factor λ , by letting: $\nu(i, s_i, t_i) = \lambda \alpha_i(t_i | s_i) / n$ if $t_i \neq s_i$, and giving any value (up to normalization of ν) to $\nu(i, s_i, s_i)$. Conversely, any strategy ν of the minimizer in Γ_{HS} can be emulated in Γ by letting $\alpha_i(t_i | s_i) = \nu(i, s_i, t_i)$ if $s_i \neq t_i$ and $\alpha_i(s_i | s_i)$ be the complement to 1; it follows that the value of Γ is nonnegative if and only if the value of Γ_{HS} is nonnegative. Therefore the fact that the value of Γ_{HS} is nonnegative may be proved by the method of section 2.3.

Proof. First, in (2.11.2) we may assume $\sum_j a_{jk} = 1$ without loss of generality (indeed, one may increase arbitrarily the coefficients a_{kk} to ensure that each row sums to some positive constant and then divide all coefficients by this constant to normalize); second, by linearity, (2.11.2) holds for all vectors u if and only if it holds for all basis vectors (i.e. with one component equal to 1 and all the others zero); third, (2.11.2) holds for all basis vectors if and only if $\sum_j x_j a_{ji} = x_i$ ($= \sum_j a_{ji} x_i$) for all i ; that is, if and only if $A^T x = x$ where A^T denote the $m \times m$ square matrix whose (i, j) entry is a_{ji} . Thus lemma 2.11.2 boils down to lemma 2.11.1 applied to $M = A^T$. Conversely, lemma 2.11.1 is a special case of lemma 2.11.2. \square

Incidentally, Hart and Schmeidler prove lemma 2.11.2 through the Minimax theorem; so proposition 2.11.3 yields a game-theoretic proof of the existence of invariant distributions for finite Markov chains.³²

2.11.2 Other proofs

Nau and McCardle's proof is very similar. They also introduce (implicitly) the payoff matrix of Γ_{HS} . A pure strategy profile s is defined to be *jointly coherent* if $g(s, \alpha) = 0$ for all (unnormalized) dual vectors α . Nau and McCardle show through lemma 2.11.1, and essentially as in section 2.3, that there exists a jointly coherent strategy profile. Finally, they prove through a variant of Farkas lemma that a strategy profile is jointly coherent if and only if it has positive probability in some correlated equilibrium.³³ It follows that there exists a correlated equilibrium.

Myerson's proof is very close to the proof of section 2.3, but instead of introducing an auxiliary zero-sum game, Myerson introduces an auxiliary linear program. Deviation vectors appear as vectors of dual variables, hence the terms dual vector and dual reduction. Myerson's linear program corresponds to the maximisation's program of the maximizer in the auxiliary game of section 2.3.

³²I owe this remark to B. von Stengel, who first showed me a proof of lemma 2.11.1 based on linear duality. Such a proof can also be found in (Mertens et al, 1994, ex. 9, p. 41).

³³This implies that if s is not jointly coherent, then there exists a dual vector α such that $g(s, \alpha) < 0$. Since the set of dual vectors is convex, this implies that there exists a dual vector satisfying $g(s, \alpha) < 0$ for every non jointly coherent s , as used in the proof of proposition 2.5.5.

Chapitre 3

Openness of the Set of Games with a Unique Correlated Equilibrium¹

Abstract

This chapter shows that the set of n -player games with a unique correlated equilibrium is open, which is not true for Nash equilibrium. Related results are studied. For instance, we show that, even though the set of bimatrix games with a unique Nash equilibrium is open, the set of symmetric bimatrix games with a unique symmetric Nash equilibrium is not.

Keywords: correlated equilibrium, open set, quasi-strict equilibrium

The practical relevance of a phenomenon arising in a game often hinges upon this phenomenon being robust to perturbation of the game. To establish such robustness results typically requires proving that some of the properties of the game we initially considered are themselves robust; that is, that the set of games having these properties is open. We investigate here whether the set of games with a unique equilibrium is open, both for Nash equilibrium and correlated equilibrium.

A first result, due to Jansen (1981), is that the set of *bimatrix* games with a unique Nash equilibrium is open. However, this does not extend to three-player games. A counter-example due to Eilon Solan (personal communication) is given in section 3.3. Our main result is that, by contrast, for any

number of players n , the set of n -player finite games with a unique correlated equilibrium is open.² This generalizes an earlier result of Nitzan (2005).³ We also establish a number of related results. For instance, the fact that the set of games with a unique and strict Nash equilibrium is open, or the fact that, in two-player games, the set of symmetric games with a unique symmetric Nash equilibrium is not open, but the set of symmetric games with a unique and *quasi-strict* symmetric Nash equilibrium is.

The material is organized as follows. Definitions and notations are introduced in section 1. The fact that the set of games with a unique correlated equilibrium is open proved in section 2. Related results are stated and proved in section 3.

3.1 Definitions and main result

We first recall the notations: $I = \{1, 2, \dots, n\}$ is the set of players, S^i the set of pure strategies of player i and $S^{-i} := \times_{j \in I \setminus \{i\}} S^j$. The utility function of player i is $U^i : S = \times_{i \in I} S^i \rightarrow \mathbb{R}$. A pure strategy profile is denoted by $s = (s^i, s^{-i})$. Finally, given any finite set T , the simplex of probability distributions over T is denoted by $\Delta(T)$.

A *correlated strategy* of the players in I is a probability distribution over the set S of pure strategy profiles. Thus $\mu = (\mu(s))_{s \in S}$ is a correlated strategy if:

$$\mu(s) \geq 0 \quad \forall s \in S \quad (3.1.1)$$

$$\sum_{s \in S} \mu(s) = 1 \quad (3.1.2)$$

Henceforth, the conditions in (3.1.1) will be called *nonnegativity constraints*. For s^i, t^i in S^i and μ in $\Delta(S)$, let

$$h^{s^i, t^i}(\mu) := \sum_{s^{-i} \in S^{-i}} \mu(s) [U^i(s) - U^i(t^i, s^{-i})]$$

²This will be used in chapter 10 to show that elimination of all strategies used in correlated equilibria by a wide class of evolutionary dynamics occurs for an open set of games.

³She showed that, for two-player games with the same number of pure strategies for both players, the set of games with a unique correlated equilibrium and such that this correlated equilibrium has full support is open.

where, as throughout, $s = (s^i, s^{-i})$.

Definition. A correlated strategy μ is a correlated equilibrium (Aumann, 1974) if it satisfies the following incentive constraints:

$$h^{s^i, t^i}(\mu) \geq 0, \quad \forall i \in I, \forall s^i \in S^i, \forall t^i \in S^i \quad (3.1.3)$$

A n -player finite game has size $m_1 \times m_2 \times \dots \times m_n$ if, for every i in I , the pure strategy set of player i has cardinal m_i . Assimilating a game and its payoff matrices, a n -player game of size $m_1 \times m_2 \times \dots \times m_n$ may be seen as a point in $(\mathbb{R}^{m_1 m_2 \dots m_n})^n$, hence the notions of a neighborhood of a game and of an open set of games. The main result of this chapter is that:

Proposition 3.1.1. *The set of n -player games of size $m_1 \times m_2 \times \dots \times m_n$ with a unique correlated equilibrium is a nonempty, open subset of the set of games of size $m_1 \times m_2 \times \dots \times m_n$. Furthermore, if a n -player finite game has a unique correlated equilibrium σ , then the (unique) correlated equilibrium of every nearby game has the same support as σ .*

3.2 Proof

The fact that the set of $m_1 \times m_2 \times \dots \times m_n$ games with a unique correlated equilibrium is nonempty is obvious: any dominance solvable game has a unique correlated equilibrium⁴. Thus it suffices to prove that this set is open. Equivalently, letting G be a game with a unique correlated equilibrium and (G_n) a sequence of games converging towards G , we need to show that, for n large enough, the game G_n has a unique correlated equilibrium. The proof runs as follows: Let σ denote the unique correlated equilibrium of G . A dual reduction argument shows that σ is a quasi-strict Nash equilibrium (lemma 3.2.2). Together with the upper semi-continuity of the Nash equilibrium correspondence this implies that, for n large enough, G_n has a quasi-strict Nash equilibrium with the same support as σ (lemma 3.2.3). Since two quasi-strict Nash equilibria with the same support satisfy the same nonnegativity and incentive constraints with strict inequality (lemma 3.2.4), it follows that, for n large enough, G_n has a correlated equilibrium satisfying with strict inequality the same constraints as σ . Due to a general result on polytopes (lemma

⁴Note also that, as will be shown in the next chapter, generic two-player zero-sum games have a unique correlated equilibrium.

3.2.1), this implies that, for n sufficiently large, the correlated equilibrium polytope of G_n is a singleton. This completes the proof.

We begin with the result on polytopes: Let (\mathbf{A}_n) be a sequence of $p \times q$ real matrices and (\mathbf{a}_n) a sequence of column vectors of size p (we use bold letters for vectors and matrices). Assume that (\mathbf{A}_n) and (\mathbf{a}_n) converge respectively towards the matrix \mathbf{A} and the vector \mathbf{a} . Let

$$C_n = \{\mathbf{x} \in \mathbb{R}^q, \mathbf{A}_n \mathbf{x} \geq \mathbf{a}_n\}$$

and⁵ $C = \{\mathbf{x} \in \mathbb{R}^q, \mathbf{A} \mathbf{x} \geq \mathbf{a}\}$. Assume that C_n is uniformly bounded:

$$\exists M \in \mathbb{R}, \forall n \in \mathbb{N}, \forall \mathbf{x} \in C_n, \max_{1 \leq i \leq p} |(\mathbf{A}_n \mathbf{x})_i| \leq M \quad (3.2.1)$$

Assume finally that C is a singleton: $C = \{\bar{\mathbf{x}}\}$. Let $a_{n,i}$ denote the i^{th} component of \mathbf{a}_n . We have:

Lemma 3.2.1. *If there exists $N \in \mathbb{N}$ such that, for all $n \geq N$,*

$$\exists x_n \in C_n, \forall i \in \{1, \dots, p\}, (\mathbf{A}_n \mathbf{x}_n)_i > a_{n,i} \Rightarrow (\mathbf{A} \bar{\mathbf{x}})_i > a_i \quad (3.2.2)$$

then for n large enough, C_n is a singleton.

Proof. The proof is by contradiction. For notational simplicity, assume $\mathbf{a}_n = 0$ for all n (hence $\mathbf{a} = 0$). The proof in the general case is the same.⁶ If lemma 3.2.1 does not hold, then, up to considering a subsequence, we may assume that for every n in \mathbb{N} , condition (3.2.2) holds but C_n is not a singleton. As will later be proved, this implies that for every n in \mathbb{N} :

$$\exists \mathbf{z}_n \in C_n, \exists i \in \{1, \dots, p\}, (\mathbf{A}_n \mathbf{z}_n)_i = 0, (\mathbf{A} \bar{\mathbf{x}})_i > 0 \quad (3.2.3)$$

Due to (3.2.1), the sequence (\mathbf{z}_n) is bounded. Furthermore, since $\mathbf{A}_n \rightarrow \mathbf{A}$ and $\mathbf{A}_n \mathbf{z}_n \geq 0$, any accumulation point $\bar{\mathbf{z}}$ of (\mathbf{z}_n) satisfies $\mathbf{A} \bar{\mathbf{z}} \geq 0$, i.e. $\bar{\mathbf{z}} \in C$. Since $C = \{\bar{\mathbf{x}}\}$, it follows that (\mathbf{z}_n) converges to $\bar{\mathbf{x}}$. Therefore, for n sufficiently large,

$$\forall i \in \{1, \dots, p\}, (\mathbf{A} \bar{\mathbf{x}})_i > 0 \Rightarrow (\mathbf{A}_n \mathbf{z}_n)_i > 0$$

contradicting (3.2.3).

⁵ $\mathbf{A}_n \mathbf{x} \geq \mathbf{a}_n$ means that the inequality holds for each coordinate i in $\{1, \dots, p\}$.

⁶ Up to replacement of the scalar 0 by a_i or $a_{n,i}$ and of the vector 0 by \mathbf{a} or \mathbf{a}_n , depending on the context.

It only remains to prove (3.2.3). By assumption, there exists $\mathbf{y}_n \in C_n$, $\mathbf{y}_n \neq \mathbf{x}_n$. By convexity of C_n , for every λ in $[0, 1]$, $\lambda\mathbf{y}_n + (1 - \lambda)\mathbf{x}_n \in C_n$. Now extend the segment $[\mathbf{x}_n, \mathbf{y}_n]$ in the direction of \mathbf{y}_n : since $\mathbf{y}_n \neq \mathbf{x}_n$ and since C_n is bounded, there exists a maximal value of λ (≥ 1) such that $\lambda\mathbf{y}_n + (1 - \lambda)\mathbf{x}_n \in C_n$. Call this value λ_{max} and let

$$\mathbf{z}_n := \lambda_{max}\mathbf{y}_n + (1 - \lambda_{max})\mathbf{x}_n = \mathbf{x}_n + \lambda_{max}(\mathbf{y}_n - \mathbf{x}_n)$$

By definition of λ_{max} , there exists i in $\{1, \dots, p\}$ such that

$$(\mathbf{A}_n\mathbf{z}_n)_i = (\mathbf{A}_n\mathbf{x}_n)_i + \lambda_{max} [(\mathbf{A}_n\mathbf{y}_n)_i - (\mathbf{A}_n\mathbf{x}_n)_i] = 0 \quad (3.2.4)$$

and

$$\forall \lambda \in \mathbb{R}, \lambda > \lambda_{max} \Rightarrow (\mathbf{A}_n\mathbf{x}_n)_i + \lambda [(\mathbf{A}_n\mathbf{y}_n)_i - (\mathbf{A}_n\mathbf{x}_n)_i] < 0$$

It follows that $(\mathbf{A}_n\mathbf{y}_n)_i - (\mathbf{A}_n\mathbf{x}_n)_i < 0$. Since $(\mathbf{A}_n\mathbf{y}_n)_i \geq 0$, this implies $(\mathbf{A}_n\mathbf{x}_n)_i > 0$, hence by (3.2.2), $(\mathbf{A}\bar{\mathbf{x}})_i > 0$. Together with (3.2.4), this proves (3.2.3) and completes the proof of lemma 3.2.1. \square

Definition. A Nash equilibrium τ is quasi-strict if, for every player i and every pure strategy s^i in S^i , if s^i does not belong to the support of τ^i , then $U^i(s^i, \tau^{-i}) < U^i(\tau)$.

Lemma 3.2.2. *If a n -player finite game has a unique correlated equilibrium, then this correlated equilibrium is a quasi-strict Nash equilibrium.*

Proof. This was proved as corollary 2.5.7 in chapter 2. \square

Lemma 3.2.3. *If the n -player game G has a unique Nash equilibrium σ and that this Nash equilibrium is quasi-strict, then there exists a neighbourhood Ω_G of G such that, for every game \hat{G} in Ω_G and every Nash equilibrium $\hat{\sigma}$ of \hat{G} , the support of $\hat{\sigma}$ is equal to the support of σ and $\hat{\sigma}$ is quasi-strict.*

Proof. Let $(G_n)_{n \in \mathbb{N}}$ be a sequence of games converging to G and σ_n a Nash equilibrium of G_n . To prove lemma 3.2.3, it is enough to show that, for n large enough, the support of σ_n is equal to the support of σ and σ_n is quasi-strict. Since the Nash equilibrium correspondence is upper semi-continuous and since G has a unique Nash equilibrium, it follows that σ_n converges to σ . Therefore, if $\sigma^i(s^i)$ is positive then, for n sufficiently large, $\sigma_n^i(s^i)$ is positive too. It follows that there exists a constant N' such that, for all $n \geq N'$, the support of σ_n includes the support of σ .

Furthermore, the Nash equilibrium σ is quasi-strict. Therefore, if s^i does not belong to the support of σ^i , then $U^i(s^i, \sigma^{-i}) < U^i(\sigma)$. Since σ_n converges to σ and U_n^i to U^i , this⁷ implies that, for n large enough, $U_n^i(s^i, \sigma_n^{-i}) < U_n^i(\sigma_n)$. Therefore, there exists a constant N'' such that

$$\forall n \geq N'', \forall i \in I, \forall s^i \in S^i, \sigma^i(s^i) = 0 \Rightarrow U_n^i(s^i, \sigma_n^{-i}) < U_n^i(\sigma_n) \quad (3.2.5)$$

Since $U_n^i(s^i, \sigma_n^{-i}) < U_n^i(\sigma_n) \Rightarrow \sigma_n^i(s^i) = 0$, it follows that, for $n \geq N''$, the support of σ includes the support of σ_n . For $n \geq \max(N', N'')$, σ and σ_n have the same support; we can thus replace $\sigma^i(s^i) > 0$ by $\sigma_n^i(s^i) > 0$ in (3.2.5), which shows that σ_n is quasi-strict. This completes the proof. \square

Lemma 3.2.4. *Let G and \hat{G} be two games with the same set of players and strategies. Let σ and $\hat{\sigma}$ be Nash equilibria of, respectively, G and \hat{G} . Assume that σ and $\hat{\sigma}$ have the same support and are both quasi-strict. Then, among the nonnegativity and incentive constraints defining correlated equilibria, σ and $\hat{\sigma}$ satisfy the same constraints with strict inequality.*

Proof. Since, by assumption, σ and $\hat{\sigma}$ have the same support, they satisfy with strict inequality the same nonnegativity constraints. We now show that they satisfy with strict inequality the same incentive constraints. Since σ is a product distribution, it follows that

$$h^{s^i, t^i}(\sigma) = \sigma^i(s^i) [U^i(s^i, \sigma^{-i}) - U^i(t^i, \sigma^{-i})] \quad \forall i, \forall s^i, \forall t^i$$

Let $\tilde{S} = \times_i \tilde{S}^i$ denote the support of both σ and $\hat{\sigma}$. If $s^i \notin \tilde{S}^i$, then $\sigma^i(s^i) = 0$ hence $h^{s^i, t^i}(\sigma) = 0$ for every t^i in S^i . If $s^i \in \tilde{S}^i$ and $t^i \in \tilde{S}^i$, then, since σ is a Nash equilibrium, $U^i(s^i, \sigma^{-i}) = U^i(t^i, \sigma^{-i})$ hence $h^{s^i, t^i}(\sigma) = 0$. Finally, if $s^i \in \tilde{S}^i$ and $t^i \notin \tilde{S}^i$ then $\sigma^i(s^i) > 0$ and, since σ is quasi-strict, $U^i(s^i, \sigma^{-i}) - U^i(t^i, \sigma^{-i}) > 0$. Therefore, $h^{s^i, t^i}(\sigma) > 0$. Grouping these observations we obtain that $h^{s^i, t^i}(\sigma) > 0$ if and only if $s^i \in \tilde{S}^i$ and $t^i \notin \tilde{S}^i$. The same result holds for $\hat{\sigma}$ so that, letting $(\hat{h}^{s^i, t^i})_{s^i \in S^i, t^i \in S^i}$ denote the linear forms associated with the correlated equilibrium incentive constraints of \hat{G} , we have:

$$h^{s^i, t^i}(\sigma) > 0 \Leftrightarrow \hat{h}^{s^i, t^i}(\hat{\sigma}) > 0$$

This completes the proof. \square

⁷ U_n^i denotes the utility function of player i in the game G_n .

We now conclude. Let G be a game with a unique correlated equilibrium σ and (G_n) be a sequence of games converging towards G . Let C_n be the correlated equilibrium polytope of G_n . Combining lemmata 3.2.2, 3.2.3 and 3.2.4, we obtain that, for n large enough, G_n has a correlated equilibrium σ_n satisfying with strict inequality the same constraints as σ . By lemma 3.2.1, this implies that for n sufficiently large, C_n is a singleton. This completes the proof of proposition 3.1.1.

3.3 Remarks

1. Nitzan (2005) proved independently and earlier a weaker version of proposition 3.1.1. More precisely, she proved that if a two-player $m \times m$ game has a unique correlated equilibrium and that this correlated equilibrium has full support, then every nearby game has a unique correlated equilibrium and this correlated equilibrium has full support. To prove this result with our method, it suffices to note that if a game has a unique and completely mixed Nash equilibrium, then every nearby game has a completely mixed Nash equilibrium, and then to apply lemma 3.2.1. This illustrates a difference between our arguments and Nitzan's: while she uses a theorem of the alternative, we do not need any theorem of the alternative to prove her results.⁸

2. The set of 3-player games with a unique Nash equilibrium is not open. The following $2 \times 2 \times 2$ counter-example is adapted from (Flesch et al, 1997) and was provided by Eilon Solan (personal communication).

$$\begin{pmatrix} 1, 1, 1 & 0, 1, 1 \\ 1, 1, 0 & 1, 0, 1 \end{pmatrix} \quad \begin{pmatrix} 1, 0, 1 - \epsilon & 1, 1, 0 \\ 0, 1, 1 & 0, 0, 0 \end{pmatrix} \quad (3.3.1)$$

Player 1 chooses a row (Top or Bottom), player 2 a column (Left or Right) and player 3 a matrix (West or East). For $\epsilon = 0$, there is a unique Nash equilibrium, in which all players play their first strategy (this will be proved below). However, for $\epsilon > 0$, there is a continuum of Nash equilibria. Indeed, every (partially) mixed strategy profile in which player 1 plays Bottom with

⁸We do however use a theorem of the alternative to prove proposition 3.1.1. Indeed, the proof of lemma 3.2.2 uses Nau and McCardle's (1990) characterization of strategy profiles with positive probability in at least one correlated equilibrium, which itself relies on a theorem of the alternative.

probability less than $\epsilon/(1 + \epsilon)$ and player 2 and 3 stick to their first strategy is a Nash equilibrium. Thus, in 3-player games, there are sequences of games with a continuum of Nash equilibria converging towards a game with a unique Nash equilibrium.

The game (3.3.1) with $\epsilon = 0$ also provides an example of a game with a unique Nash equilibrium that is not quasi-strict. This cannot happen in two-player games, since bimatrix games always have a quasi-strict Nash equilibrium (Norde, 1999).⁹

Proof that the game (3.3.1) with $\epsilon = 0$ has a unique Nash equilibrium: for $\epsilon = 0$, the game (3.3.1) may be described as follows: player $i + 1$ (counted modulo 3) wants to mismatch player i , except if all players play their first strategy. Thus, in an hypothetical equilibrium different from Top-Left-West, if i plays in pure strategy, then $i + 1$ must mismatch i , $i + 2$ mismatch $i + 1$ and $i + 3 = i$ mismatch $i + 2$; therefore, i must mismatch itself, a contradiction. It follows that the only equilibrium in which one of the players plays in pure strategy is Top-Left-West.

It remains to show that there are no completely mixed Nash equilibria. By contradiction, let $x \in]0, 1[$ (resp. y, z) be the probability of Bottom (resp. Right, East) in an hypothetical completely mixed Nash equilibrium. Since player 1 is indifferent between Top and Bottom, we have $y(1 - z) = (1 - y)z + yz = z$, hence $y > z$. Since the game is cyclically symmetric, it follows that $y > z > x > y$, which cannot be. This completes the proof. ■

3. A corollary of lemma 3.2.3 is that:

Corollary 3.3.1. *If a finite game has a unique and strict Nash equilibrium σ , then σ is also the unique Nash equilibrium of every nearby game.*

Indeed, by lemma 3.2.3, every Nash equilibrium of every nearby game has the same support as σ , hence is equal to σ as σ is pure. Thus, the set of games with a unique and strict Nash equilibrium is open. We do not know whether the set of games with a unique and quasi-strict Nash equilibrium is

⁹It is well known that 3-player games need not have a quasi-strict equilibrium. However, the counter-examples I found in the literature, e.g. (Raghavan, 2002), are of games with several Nash equilibria. Thus, up to my knowledge, whether a unique Nash equilibrium is necessarily quasi-strict was still open.

open.

4. In bimatrix games, both the set of games with a unique Nash equilibrium and the set of games with a unique correlated equilibrium are open. Since there are games with a unique Nash equilibrium but many correlated equilibria, the latter set is included in the former. The following examples show that on the relative boundary of the set of bimatrix games with a unique correlated equilibrium, there are games with a continuum of Nash equilibria, games with a finite number (> 1) of Nash equilibria, and games with a unique Nash equilibrium:

$$\left(\begin{array}{c} 0 \\ \epsilon \end{array} \right) \quad \left(\begin{array}{cc} 1, 1 & 0, 0 \\ 0, 0 & -\epsilon, -\epsilon \end{array} \right) \quad \left(\begin{array}{ccc|c} 0, 0 & 2, 1 & 1, 2 & -1, x \\ 1, 2 & 0, 0 & 2, 1 & -1, x \\ 2, 1 & 1, 2 & 0, 0 & -1, x \\ \hline x, -1 & x, -1 & x, -1 & 0, 0 \end{array} \right)$$

The left game is a one-person game. For $\epsilon > 0$ it has a unique correlated equilibrium. For $\epsilon = 0$ it has a continuum of Nash equilibria. The middle game has a unique correlated equilibrium (Top-Left) for $\epsilon > 0$, but two Nash equilibria for $\epsilon = 0$. The game on the right is adapted from (Nau and McCardle, 1990, example 4). The 3×3 game in the top-left corner is due to Moulin and Vial (1978). This 3×3 game has a unique Nash equilibrium: $(1/3, 1/3, 1/3)$ for both players, with payoff 1; but putting probability $1/6$ on every off-diagonal square yields a correlated equilibrium with payoff $3/2$. Now consider the whole 4×4 game. For any value of x , $(4, 4)$ is a Nash equilibrium. For $1 < x \leq 3/2$, this is the unique Nash equilibrium, but not the unique correlated equilibrium (the correlated equilibrium with payoff $3/2$ of the 3×3 top-left game induces a correlated equilibrium of the whole game). For $x > 3/2$, this is the unique correlated equilibrium.

5. The following example shows that, within the set of two-person symmetric games, the set of games with a unique symmetric Nash equilibrium is not open:

$$\left(\begin{array}{ccc} -\epsilon, -\epsilon & 1, 0 & 1, 0 \\ 0, 1 & 0, 0 & -1, -1 \\ 0, 1 & -1, -1 & 0, 0 \end{array} \right)$$

For $\epsilon = 0$, this game has a unique symmetric Nash equilibrium: Top-Left. For $\epsilon > 0$, it has 3 symmetric Nash equilibria: $(\frac{1}{1+\epsilon}, \frac{\epsilon}{1+\epsilon}, 0)$, $(\frac{1}{1+\epsilon}, 0, \frac{\epsilon}{1+\epsilon})$,

$(\frac{3}{3+2\epsilon}, \frac{\epsilon}{3+2\epsilon}, \frac{\epsilon}{3+2\epsilon})$. This is linked to the fact that, for $\epsilon = 0$, the unique symmetric Nash equilibrium is not quasi-strict. Indeed, the openness of the set of bimatrix games with a unique Nash equilibrium has the following analogue for symmetric games:

Proposition 3.3.2. *Within the set of two-person symmetric games, the set of two-person symmetric games with a unique symmetric Nash equilibrium¹⁰ and such that this Nash equilibrium is quasi-strict is open.*

Proof. Let G be a two-person game with a unique symmetric Nash equilibrium, with support $\tilde{S} = \tilde{S}^1 \times \tilde{S}^2$, and such that this Nash equilibrium is quasi-strict. It follows from a variant of lemma 3.2.3 that, within the set of two-person symmetric games, there exists a neighborhood Ω_G of G such that, for any game G' in Ω_G , any symmetric Nash equilibrium of G' has support \tilde{S} and is quasi-strict.

Fix G' in Ω_G . Since G' is symmetric, it has a symmetric Nash equilibrium σ . To establish proposition 3.3.2, it is enough to show that G' has no other symmetric Nash equilibria. By contradiction, assume that G' has a symmetric Nash equilibrium $\tau \neq \sigma$. For every λ in \mathbb{R} , define the symmetric mixed strategy profile σ_λ by $\sigma_\lambda^i = \lambda\tau^i + (1 - \lambda)\sigma^i$, for $i = 1, 2$.

There are five types of incentive and nonnegativity constraints that σ_λ must satisfy in order to be a (symmetric) Nash equilibrium:

- (i) $\sigma_\lambda(s) \geq 0, s \in \tilde{S}$;
- (ii) $\sigma_\lambda(s) \geq 0, s \notin \tilde{S}$;
- (iii) $h^{s^i, t^i}(\sigma_\lambda) \geq 0, s^i \in \tilde{S}^i, t^i \notin \tilde{S}^i, i = 1, 2$;
- (iv) $h^{s^i, t^i}(\sigma_\lambda) \geq 0, s^i \in \tilde{S}^i, t^i \in \tilde{S}^i, i = 1, 2$;
- (v) $h^{s^i, t^i}(\sigma_\lambda) \geq 0, s^i \notin \tilde{S}^i, t^i \in S^i, i = 1, 2$.

Using the fact that both σ and τ are Nash equilibria with support \tilde{S} , it is easily checked that for every λ in \mathbb{R} , σ_λ satisfies (with equality) all constraints of types (ii), (iv) and (v).

Moreover, since $\sigma \neq \tau$ and since the set of Nash equilibria is compact, it follows that there exists a maximal value of λ such that σ_λ is a (symmetric) Nash equilibrium. Call this value λ_{max} . Since all symmetric Nash equilibria of G' have support \tilde{S} and are quasi-strict, they all satisfy with strict inequality all constraints of types (i) and (iii), hence so does $\sigma_{\lambda_{max}}$. Therefore, there exists $\lambda > \lambda_{max}$ such that σ_λ satisfies all constraints of type (i) and (iii).

¹⁰There might be other, asymmetric Nash equilibria.

Since, as mentioned in the previous paragraph, σ_λ also satisfies all other constraints, it follows that σ_λ is a (symmetric) Nash equilibrium, contradicting the maximality of λ_{max} . \square

Finally, a variant of the proof of corollary 3.3.1 shows that the set of n -player symmetric games with a unique and strict symmetric Nash equilibrium is open.

6. In order to prove the existence of correlated equilibria without using a fixed point theorem, Hart and Schmeidler (1989) associate to every finite game G an auxiliary zero-sum game whose size depends only on the size of G and whose payoff matrix depends continuously on the payoff matrices of G (see chapter 2, sections 2.3 and 2.11). In this auxiliary zero-sum game, the optimal strategies of the maximizer correspond exactly to the correlated equilibria of G , so that G has a unique correlated equilibrium if and only if, in the auxiliary game, the maximizer has a unique optimal strategy. It follows that, in order to get an alternative proof of the fact that the set of games with a unique correlated equilibrium is open, it would have been enough to prove that:

If in a two-player zero-sum game, one of the players has a unique optimal strategy, then in every nearby zero-sum game (in the space of payoff matrices), this player has a unique optimal strategy.

However, this turns out to be false: let G_ϵ be the two-player zero-sum game with payoff matrix for the row player

$$\begin{array}{c} \\ T \\ M \\ B \end{array} \begin{array}{cc} L & R \\ \left(\begin{array}{cc} -\epsilon & 0 \\ 0 & -1 \\ 0 & -1 \end{array} \right) \end{array}$$

For $\epsilon = 0$, the row player has a unique optimal strategy (playing T). But for $\epsilon > 0$, the row player has an infinite number of optimal strategies: playing T with probability $1/(1 + \epsilon)$ and playing M and B with any probabilities summing to $\epsilon/(1 + \epsilon)$.

Chapitre 4

Correlated Equilibrium in Generic Two-Person Zero-Sum Games

Abstract

This chapter combines arguments of Forges (1990) and of Bonhenblust et al (1950) to show that, in generic two-person zero-sum games, all correlated equilibria are Nash equilibria.¹

Forges (1990) showed that a correlated equilibrium of a two-person zero-sum game need not be a convex combination of Nash equilibria. This chapter establishes that, nevertheless, in generic two-person zero-sum games all correlated equilibria are Nash equilibria, and that with a more demanding notion of genericity, there is a unique correlated equilibrium. While this seems to have gone unnoticed, the arguments are hardly new: as noted by Forges (1990), if, in a zero-sum game, one of the players has a unique optimal strategy, then every correlated equilibrium is a Nash equilibrium. Since Bonhenblust et al (1950) have shown that in a generic zero-sum game, there is a unique pair of optimal strategies, it follows that in a generic zero-sum game there is a unique correlated equilibrium, hence that every correlated equilibrium is a Nash equilibrium.² The merit of this chapter is thus only to gather these two results, and to use a weaker notion of genericity than

¹This chapter originated with conversations with Françoise Forges.

²We thank Sylvain Sorin for making us aware of Bonhenblust et al's (1950) result.

Bonhenblust et al (1950).³

Notations (the notations are different from the notations used in previous chapters): G is a two-person zero-sum game with payoff matrix $A = (a_{ij})_{1 \leq i \leq n, 1 \leq j \leq p}$ and value v . The pure strategy sets of the players are denoted respectively I (for player 1) and J (for player 2). If μ is a probability distribution over $I \times J$, then $\mu(i) = \sum_{j \in J} \mu(i, j)$ is the marginal probability of the pure strategy i in μ and $I_\mu = \{i \in I, \mu(i) > 0\}$ the set of pure strategies of player 1 to which μ assigns a positive marginal probability. Similarly, $J_\mu = \{j \in J, \mu(j) > 0\}$. Also, for all $i \in I_\mu$, $\mu(j|i) = \mu(j, i)/\mu(i)$ is the conditional probability of j given i . A probability distribution μ over $I \times J$ is a correlated equilibrium (Aumann, 1974) if for every i in I_μ , the pure strategy i is a best-response to $\mu(\cdot|i)$, and if the symmetric condition for player 2 holds as well.

Let C be the set of all correlated equilibria of G and $I_C = \cup_{\mu \in C} I_\mu$ the set of pure strategies of player 1 which have positive marginal probability in at least one correlated equilibrium. Similarly, let $J_C = \cup_{\mu \in C} J_\mu$. Finally, let A_C be the $|I_C| \times |J_C|$ submatrix of A obtained by considering only the i 's in I_C and the j 's in J_C . This chapter establishes that:

Proposition *If A_C has full rank (i.e. if $\text{rank} A_C = \min(|I_C|, |J_C|)$), then all correlated equilibria of G are Nash equilibria.*

The proof relies on the following lemma:

Lemma (Forges, 1990) *If μ is a correlated equilibrium of G and $\mu(i) > 0$ then the conditional probability $\mu(\cdot|i)$ over J is an optimal strategy of player 2.*

Proof. Let $i \in I_\mu$. Since μ is a correlated equilibrium, i is a best-response to $\mu(\cdot|i)$. Therefore, $\sum_{j \in J} \mu(j|i) a_{ij} \geq v$, with strict inequality if $\mu(\cdot|i)$ is not optimal. Therefore, if for some i in I_μ , $\mu(\cdot|i)$ is not optimal, then the average payoff of player 1 in μ : $\sum_{i \in I, j \in J} \mu(i, j) a_{ij}$, is strictly greater than v ; hence the average payoff of player 2 is strictly less than v , i.e. less than what he can guarantee: a contradiction. \square

³The arguments of this chapter were developed “independently”, in that we were not aware of Bonhenblust et al’s (1950) result before completing the first draft of this chapter.

We now prove the proposition:

Step 1. Assume that G has a Nash equilibrium σ with support $I_C \times J_C$. Assume without loss of generality $|I_C| \geq |J_C|$. Let τ_2 be an optimal strategy of player 2. The mixed strategy profile (σ_1, τ_2) is a Nash equilibrium with payoff v , as a product of optimal strategies. Therefore, for all $i \in I_C$, $(A_C \tau_2)_i = \sum_{j \in J_C} \tau_2(j) a_{ij} = v$. This defines a linear system of $|I_C|$ equations with $|J_C|$ variables (the $\tau(j), j \in J_C$). Assuming that A_C has full rank (hence rank $|J_C|$ in our case) this linear system has at most one solution, hence player 2 has a unique optimal strategy.

Step 2. Assume now that player 2 has a unique optimal strategy and let μ be a correlated equilibrium of G . For all $i \in I_\mu$, $\mu(\cdot|i)$ is an optimal strategy of player 2. Since player 2 has a unique optimal strategy, $\mu(\cdot|i)$ is independent of $i \in I_\mu$. This implies that μ is a product distribution, hence a Nash equilibrium.⁴

Step 3. To complete the proof we need to show that G has indeed a Nash equilibrium with support $I_C \times J_C$. By convexity of the set of correlated equilibria, there exists a correlated equilibrium μ such that $\mu(i)\mu(j) > 0$ for all $(i, j) \in I_C \times J_C$. Recall that for all $i \in I_C$, $\mu(\cdot|i)$ is an optimal strategy of player 2. Hence, by convexity of the set of optimal strategies of player 2, $\sigma_2 = \sum_{i \in I_C} \mu(i)\mu(\cdot|i)$ is also optimal. Since $\sigma_2(j) = \mu(j)$ for all $j \in J_C$, it follows that σ_2 has support J_C . Symmetrically, player 1 has an optimal strategy σ_1 with support I_C . The mixed strategy profile (σ_1, σ_2) provides a Nash equilibrium with support $I_C \times J_C$. ■

Remarks:

1. Generically⁵, A and all its submatrices have full rank. Therefore, in generic two-person zero-sum games, A_C has full rank and all correlated equilibria are Nash.

2. In a similar spirit the following more precise results can be proved : if μ is a correlated equilibrium and $\text{rank}(A_C) \geq \min(|I_\mu|, |J_\mu|)$, then μ is a Nash equilibrium. In particular, if the submatrix of A obtained by considering only the lines in I_μ and the columns in J_μ has full rank, then μ is a Nash

⁴As already mentioned, the fact that if a player has a unique optimal strategy, then every correlated equilibrium is a Nash equilibrium had been noted by Forges (1990).

⁵That is, for an open and dense set of matrices.

equilibrium distribution.

3. The above arguments are essentially the arguments used by Bohnenblust et al (1950) to show that, with a more demanding notion of genericity, generic two-person zero-sum games have a unique pair of optimal strategies (hence, by the second step of the proof of the proposition, a unique correlated equilibrium).

Chapitre 5

Elementary Games and Games Whose Correlated Equilibrium Polytope has Full Dimension

Abstract

A game is elementary (Myerson, 1997) if it has a strict correlated equilibrium with full support. A game is “full” if its correlated equilibrium polytope has the dimension of the simplex of pure strategy profiles. This chapter relates and characterizes these classes of games.

5.1 Introduction

As mentioned in chapter 2, Myerson (1997) defines a game to be *elementary* if it has correlated equilibria that satisfy all nonnegativity and incentive constraints with strict inequality. This class of games satisfies two interesting properties. First, as Myerson (1997, p. 186) points out: “For such elementary games, any player can be motivated to choose any pure strategy with no indifference problems” so that, for elementary games, “correlated equilibrium refinements that generalize Selten’s perfectness concept should be unnecessary.” Furthermore, as already mentioned in chapter 2, proposition 2.4.11, a game is elementary if and only if it cannot be (strictly) reduced by dual reduction (Myerson, 1997). It follows that by iterative dual reduction, any finite game is eventually reduced to an elementary game.¹

¹Myerson (1997) writes that this suggests that “game theorists may study elementary games without loss of generality”, a view we do not endorse.

A slightly larger, closely related class of games is the class of games whose correlated equilibrium polytope C has full dimension², henceforth called *full games*. Nau et al (2004) proved that if a game G is full and nontrivial³ then there is no Nash equilibrium in the relative interior of C , and that this is not true for general games.

This chapter relates and characterizes these two classes of games. The material is organized as follows: in the next section, the main definitions and notations are introduced. Section 5.3 shows that a game is elementary if and only if it is full and all linear forms corresponding to the incentives constraints defining correlated equilibria are nonzero. The last section and the second appendix (section 5.6) are devoted to characterizations of full games. These can also be used to characterize elementary games. Finally, a method to build full but nonelementary games is explained in the first appendix (section 5.5).

5.2 Notations and definitions

5.2.1 Notations

Let $G = \{I, (S_i)_{i \in I}, (U_i)_{i \in I}\}$ denote a finite game in strategic form: I is the nonempty finite set of players, S_i the nonempty finite set of pure strategies of player i and $U_i : \times_{i \in I} S_i \rightarrow \mathbb{R}$ the utility function of player i . The set of (pure) strategy profiles is $S = \times_{i \in I} S_i$; the set of strategy profiles for the players other than i is $S_{-i} = \times_{j \in I-i} S_j$. Pure strategies of player i (resp. pure strategy profiles; pure strategy profiles of the players other than i) are denoted s_i or t_i (resp. s ; s_{-i}). We may write (t_i, s_{-i}) to denote the strategy profile that differs from s only in that its i^{th} component is t_i . Finally, N denotes the cardinal of S and $\Delta(S)$ the set of probability distribution over S .

5.2.2 Correlated equilibrium

We recall the definition of a correlated equilibrium and *introduce some new notations* for incentive constraints. The set $\Delta(S)$ is an $N - 1$ dimensional

²That is, dimension $N - 1$ where N is the number of pure strategy profiles in the game. See section 5.2.

³A game is trivial if the payoffs of the players are independent of their own strategy.

simplex, henceforth called *the simplex*. A *correlated strategy* of the players in I is an element of the simplex. Thus $\mu = (\mu(s))_{s \in S}$ is a correlated strategy if:

$$\text{(nonnegativity constraints)} \quad \mu(s) \geq 0 \quad \forall s \in S \quad (5.2.1)$$

$$\text{(normalization constraint)} \quad \sum_{s \in S} \mu(s) = 1 \quad (5.2.2)$$

For $(i, s_i, t_i) \in I \times S_i \times S_i$, let h_{s_i, t_i} denote the linear form on \mathbb{R}^S which maps $x = (x(s))_{s \in S}$ to

$$h_{s_i, t_i}(x) = \sum_{s_{-i} \in S_{-i}} x(s) [U_i(s) - U_i(t_i, s_{-i})]$$

A correlated strategy μ is a *correlated equilibrium* (Aumann, 1974) if:

$$\text{(incentive constraints)} \quad h_{s_i, t_i}(\mu) \geq 0 \quad \forall i \in I, \forall s_i \in S_i, \forall t_i \in S_i \setminus \{s_i\} \quad (5.2.3)$$

The set of correlated equilibria is a polytope, which we denote by C . Since C is a subset of the simplex, it has at most dimension $N - 1$.

Definition 5.2.1. *The polytope C has full dimension if it has dimension $N - 1$.*

5.2.3 Full games

Definition 5.2.2. *G is a full game if C has full dimension.*

To state more precisely the result of Nau et al (2004) mentioned in the introduction, we need some definitions:

Definition 5.2.3. *Let $(i, s_i, t_i) \in I \times S_i \times S_i$, with $s_i \neq t_i$. The incentive constraint $h_{s_i, t_i}(\cdot) \geq 0$ is vacuous if $h_{s_i, t_i} = 0$. That is, if $U_i(s_i, \cdot) = U_i(t_i, \cdot)$.*

Definition 5.2.4. *A game is nontrivial if at least one of the incentive constraints is nonvacuous: $\exists i \in I, \exists s_i \in S_i, \exists t_i \neq s_i, U_i(s_i, \cdot) \neq U_i(t_i, \cdot)$.*

Nau et al (2004) proved that if G is nontrivial, then all Nash equilibria lie on the boundary of C ⁴. If furthermore C has full dimension, its boundary coincides with its relative boundary, hence all Nash equilibria lie on its

⁴We could see C as a subset of \mathbb{R}^N , in which case C (and $\Delta(S)$) would always have an empty interior. Rather, we see C as a subset of the hyperplane containing the simplex. Therefore, a correlated equilibrium belongs to the *boundary* of C if and only if it belongs to a face of C whose dimension is at most $N - 2$.

relative boundary. In contrast, if C has less than full dimension, it consists entirely of boundary; the above result is then void and there are examples of nontrivial games with Nash equilibria in the relative interior of C (Nau et al, 2004).

5.2.4 Elementary games

Definition 5.2.5. *A game is elementary (Myerson, 1997) if it has a correlated equilibrium μ which satisfies all incentive constraints (5.2.3) with strict inequality. That is,*

$$\forall i \in I, \forall s_i \in S_i, \forall t_i \in S_i \setminus \{s_i\}, h_{s_i, t_i}(\mu) > 0 \quad (5.2.4)$$

A game is elementary if and only if it has a strict correlated equilibrium with full support.⁵ Indeed, if μ is a strict correlated equilibrium with full support, then it satisfies (5.2.4). Conversely, by definition of h_{s_i, t_i} , if s_i has marginal probability zero in μ , then $h_{s_i, t_i}(\mu) = 0$ for all t_i in S_i . Therefore if μ satisfies (5.2.4), then every pure strategy of every player has positive marginal probability in μ and, by perturbing μ , one obtains a strict correlated equilibrium with full support.

Finally, if some player i is indifferent between two pure strategies s_i and $t_i \neq s_i$ (that is, if $U_i(s_i, \cdot) = U_i(t_i, \cdot)$) then $h_{s_i, t_i}(\mu) = 0$ for all μ in $\Delta(S)$, and (5.2.4) cannot be satisfied. It follows that:

Remark 5.2.6. *If a game is elementary, then all incentive constraints are nonvacuous: $\forall i \in I, \forall s_i \in S_i, \forall t_i \in S_i \setminus \{s_i\}, U_i(s_i, \cdot) \neq U_i(t_i, \cdot)$.*

5.3 The relation between elementary games and full games

In this section, we first give necessary and sufficient conditions for a game to be full. We then precise the link between elementary games and full games.

Proposition 5.3.1. *The following properties are equivalent:*

- (i) C has full dimension

⁵That is, a correlated equilibrium satisfying all nonnegativity and incentive constraints with strict inequality.

(ii) *There exists a correlated equilibrium that satisfies all nonvacuous incentive constraints with strict inequality. Formally,*

$$\exists \mu \in C, \forall i \in I, \forall s_i \in S_i, \forall t_i \in S_i \setminus \{s_i\}, h_{s_i, t_i} \neq 0 \Rightarrow h_{s_i, t_i}(\mu) > 0 \quad (5.3.1)$$

(iii) *There exists a correlated equilibrium that satisfies all nonnegativity constraints and all nonvacuous incentive constraints with strict inequality.*

Proof. To prove (ii) \Rightarrow (iii) and (iii) \Rightarrow (i), note that by perturbing a correlated equilibrium satisfying all nonvacuous incentive constraints with strict inequality, one obtains a correlated equilibrium satisfying all nonnegativity constraints and all nonvacuous incentive constraints with strict inequality and that any correlated strategy in a sufficiently small neighborhood of this correlated equilibrium is a correlated equilibrium. We now prove (i) \Rightarrow (ii) by contraposition. By convexity of the set of correlated equilibria C , (ii) is equivalent to:

$$\forall i \in I, \forall s_i \in S_i, \forall t_i \in S_i \setminus \{s_i\}, h_{s_i, t_i} \neq 0 \Rightarrow (\exists \mu \in C, h_{s_i, t_i}(\mu) > 0)$$

Therefore, if (ii) does not hold, then there exists a nonvacuous incentive constraint that is binding in all correlated equilibria; this constraint defines an hyperplane whose intersection with the simplex has at most dimension $N - 2$ and includes C ; therefore C has at most dimension $N - 2$, contradicting (i). \square

Corollary 5.3.2. *G is elementary if and only if (a) none of the incentive constraints is vacuous and (b) C has full dimension.*

Proof. Clear from definition 5.2.5, remark 5.2.6 and the equivalence of (i) and (ii) in proposition 5.3.1. \square

Any trivial game (in the sense of definition 5.2.4) is a full, nonelementary game. A more subtle example of such a game is the following:

Example 5.3.3.

	L	R
T	1, 1	0, 0
B	1, 0	0, 1

There are four incentive constraints. Two of them are vacuous, hence this game is not elementary. However, the correlated strategy assigning probability 1/2 to both TL and BR checks the two nonvacuous incentive constraints with strict inequality, so, by proposition 5.3.1, this game is full.

A general method to build full, nonelementary games is given in the first appendix (section 5.5).

5.4 Characterizations of full games and elementary games

In this section, we provide criteria to determine whether C has full dimension. By corollary 5.3.2, these criteria can also be used to determine whether a game is elementary. We conclude with examples of applications of these criteria.

5.4.1 Characterizations

The following proposition is based on (Hart and Schmeidler, 1989), (Owen, 1995, p.186) and (Myerson, 1997). Let G be nontrivial. Consider the following two-player, zero-sum, auxiliary game Γ : the maximizer chooses a strategy profile s in S ; the minimizer chooses a player i in N and a couple of strategy (s'_i, t_i) in $S_i \times S_i$, such that $U_i(s'_i, \cdot) \neq U_i(t_i, \cdot)$.⁶ The payoff for the maximizer is $U_i(s) - U_i(t_i, s_{-i})$ if $s'_i = s_i$ and 0 otherwise.

Proposition 5.4.1. *C has full dimension if and only if the value of (the mixed extension of) Γ is positive.*

Proof. A mixed strategy of the maximizer is a correlated strategy μ of G ; the payoff if the minimizer chooses (s'_i, t_i) is $h_{s'_i, t_i}(\mu)$. Thus, μ guarantees a positive payoff if and only if μ checks all nonvacuous incentive constraints with strict inequality (and if it does $\mu \in C$). Then apply proposition 5.3.1. \square

The following propositions apply only to games with a correlated equilibrium with full support (for instance, a completely mixed Nash equilibrium). Let m be a positive integer and h_1, \dots, h_m denote the linear forms associated with the nonvacuous incentive constraints.

Proposition 5.4.2. *Assume that G has a correlated equilibrium with full support. If h_1, \dots, h_m are independent, then C has full dimension.*

⁶Such a triplet (i, s'_i, t_i) with $U_i(s'_i, \cdot) \neq U_i(t_i, \cdot)$ must exist, because G is nontrivial.

Proof. Given in the second appendix (section 5.6). \square

If h_1, \dots, h_m are not independent, let B be a basis of the linear span of $\{h_1, \dots, h_m\}$. Without loss of generality, assume that $B = (h_1, \dots, h_q)$ with $1 \leq q < m$. Let $A = (a_{kl})_{1 \leq k \leq q, q+1 \leq l \leq m}$ be the matrix of (h_{q+1}, \dots, h_m) in the basis B ; that is, for all $q+1 \leq l \leq m$,

$$h_l = \sum_{1 \leq k \leq q} a_{kl} h_k$$

Let Γ' denote the two-player, zero-sum, auxiliary game, whose payoff matrix for the maximizer is A ; that is the maximizer chooses k in $\{1, \dots, q\}$, the minimizer chooses l in $\{q+1, \dots, m\}$ and the payoff for the maximizer is a_{kl} .

Proposition 5.4.3. *Assume that G has a correlated equilibrium with full support. If h_1, \dots, h_m are not independent, C has full dimension if and only if the value of (the mixed extension of) Γ' is positive.*

Proof. Given in the second appendix (section 5.6). \square

5.4.2 Examples

Example 5.4.4. *An elementary game with linearly dependent incentive constraints.*

The following 3-player, $2 \times 2 \times 2$ game is studied by Nau et al (2004):

Up:		Left	Right
	Top	$\left(0, 0, 2$	$0, 3, 0 \right)$
	Bottom	$\left(3, 0, 0$	$0, 0, 0 \right)$
Down:			
	Top	$\left(1, 1, 0$	$0, 0, 0 \right)$
	Bottom	$\left(0, 0, 0$	$0, 0, 3 \right)$

This game has a completely mixed Nash equilibrium. There are only five distinct incentive constraints (the constraint stipulating that Row should not deviate from Top to Bottom is the same as the constraint stipulating that Column should not deviate from Left to Right). These five incentive constraints are linearly independent. It follows that the payoff matrix of the auxiliary game Γ' is a 5×1 column matrix whose entries are four 0 and a 1. Therefore, the value of Γ' is positive and, by proposition 5.4.3, C has dimension 7. Furthermore, none of the incentive constraints is vacuous. Therefore, by corollary 5.3.2, this game is elementary.

Example 5.4.5. *A nonelementary game:*

$$\begin{array}{cc} & \begin{array}{cc} L & R \end{array} \\ \begin{array}{c} T \\ B \end{array} & \left(\begin{array}{cc} 2, -1 & 0, 0 \\ 0, 0 & 1, -2 \end{array} \right) \end{array}$$

This game has a completely mixed Nash equilibrium. Any incentive constraint is a nonpositive linear combination of the three other incentive constraints, which are linearly independent. It follows that the payoff matrix of Γ' is a 3×1 , nonpositive column matrix, hence that the value of Γ' is nonpositive. By proposition 5.4.3, this implies that C has less than dimension 3. In particular, this game is not elementary.⁷

5.5 First appendix: A method to build full, nonelementary games

We first need a definition:

Definition 5.5.1. *Let $G = (I, (S_i)_{i \in I}, (U_i)_{i \in I})$ and $G' = (I', (S'_i)_{i \in I'}, (U'_i)_{i \in I'})$ be two finite games. G' is built on G by adding a semi-duplicate to player i if:*

- $I' = I$
- $S'_j = S_j \quad \forall j \neq i$
- $\exists t'_i \in S'_i, S'_i = S_i \cup \{t'_i\}$
- $U'_k(s) = U_k(s) \quad \forall s \in S, \forall k \in I$
- $\exists t_i \in S_i, \forall s_{-i} \in S_{-i}, U'_i(t'_i, s_{-i}) = U_i(t_i, s_{-i})$ ⁸

⁷Of course, this can also be proved by other methods, for instance by noting that this game is best-response equivalent to a two-player zero-sum game (Rosenthal, 1974) and that in a non trivial zero-sum game, the correlated equilibrium polytope cannot have full dimension.

⁸In words, in G' the set of players is the same as in G and the pure strategy sets are the same for all players but i , who has an additional pure strategy t'_i ; when player i does not use his additional strategy the payoffs in G' are the same as in G ; furthermore player i is indifferent between his additional strategy and a strategy t_i that was already available in G .

Example 5.5.2.

$$G_1 = (1, 1 \quad 0, 0) \quad G'_1 = \begin{pmatrix} 1, 1 & 0, 0 \\ 1, 0 & 0, 1 \end{pmatrix} \quad G''_1 = \begin{pmatrix} 1, 1 & 0, 0 & 0, 1 \\ 1, 0 & 0, 1 & 1, 0 \end{pmatrix}$$

G'_1 is built on G_1 by adding a semi-duplicate to the row player and G''_1 is built on G'_1 by adding a semi-duplicate to the column player.

We can now provide the method:

Proposition 5.5.3. *Let G be elementary and G' be built on G by adding a semi-duplicate to some player. Then G' is full and nonelementary.*

Proof. G' is clearly nonelementary, so we only have to prove that G' is full. Let μ in $\Delta(S)$ check all incentive constraints of G with strict inequality (in the sense of (5.2.4)). Define μ' and ν' in $\Delta(S')$ by:

$$\begin{aligned} \mu'(s) &= \mu(s) \quad \forall s \in S ; \quad \mu'(t'_i, s_{-i}) = 0 \quad \forall s_{-i} \in S_{-i} \\ \nu'(s) &= 0 \quad \forall s \in S ; \quad \nu'(t'_i, s_{-i}) = \frac{1}{\mu(t_i \times S_{-i})} \mu(t_i, s_{-i}) \quad \forall s_{-i} \in S_{-i} \end{aligned}$$

where $\mu(t_i \times S_{-i}) = \sum_{s_{-i} \in S_{-i}} \mu(t_i, s_{-i})$. For $\epsilon > 0$ small enough, $\mu'_\epsilon := (1 - \epsilon)\mu' + \epsilon\nu'$ is a correlated equilibrium of G' that satisfies all its nonvacuous incentive constraints with strict inequality. By proposition 5.3.1, it follows that G' is full. \square

Full, nonelementary games cannot all be built by adding semi-duplicates to an elementary game. For instance, G'_1 is full and nonelementary (see example 5.3.3) but cannot be built in this way. Note also that if G is full but not elementary, then adding a semi-duplicate to G need not yield a full game. For instance, G''_1 is not full.⁹ The point is that adding a new strategy to some player may lift the indifference of some other player between two of her strategies. This shall be clear from proposition 5.5.5, which generalizes proposition 5.5.3. We first need a definition:

Definition 5.5.4. *Let G' be a game built on G by adding a semi-duplicate to player i . G' preserves indifference in G if for all $j \neq i$ and all s_j, t_j in S_j :*

$$U_j(s_j, \cdot) = U_j(t_j, \cdot) \Rightarrow U'_j(s_j, \cdot) = U'_j(t_j, \cdot)$$

⁹This follows from proposition 5.3.1 and from the fact that the first strategy (Top) of player 1 is weakly dominated.

That is, if player j was indifferent between s_j and t_j in G , she is still indifferent between s_j and t_j in G' .

Proposition 5.5.5. *Let G be full and G' be built on G by adding a semi-duplicate t'_i to player i . If G' preserves indifference in G , then G' is full.¹⁰ If G is a two-player game, then the converse holds, so that G' is full if and only if G' preserves indifference in G .*

Proof. In G' , there are three kinds of incentive constraints: constraints of type

- (i) $h'_{s_j, t_j}(\cdot) \geq 0$ with $j \neq i$ or, if $j = i$, $s_i \neq t'_i$ and $t_i \neq t'_i$;
- (ii) $h'_{s_i, t'_i}(\cdot) \geq 0$ with $s_i \in S_i$;
- (iii) $h'_{t'_i, s_i}(\cdot) \geq 0$ with $s_i \in S_i$.

(The prime in h' indicates that we consider incentive constraints of G' .)

Since G is full, there exists a correlated strategy μ that checks all the nonvacuous incentive constraints of G with strict inequality. Define μ' , ν' and μ'_ϵ as in the proof of proposition 5.5.3. Assuming that G' preserves indifference in G , we now show that for ϵ small enough, μ'_ϵ satisfies with strict inequality all the nonvacuous incentive constraints of G' . By proposition 5.3.1, this implies that G' is full.

First, for ϵ small enough, μ'_ϵ satisfies with strict inequality all the incentive constraints of type (i) corresponding to incentive constraints of G satisfied by μ with strict inequality. Since G' preserves indifference in G , the other incentive constraints of type (i) are vacuous. Since for all $s_i \in S_i$, $h'_{s_i, t'_i} = h'_{s_i, t_i}$, the above argument also takes care of constraints of type (ii). Finally, the conditional probabilities on S_{-i} given t'_i in μ'_ϵ are the same as the conditional probabilities given t_i in μ . Since $U'_i(t'_i, \cdot) = U_i(t_i, \cdot)$, it follows that μ'_ϵ satisfies with strict inequality all the nonvacuous incentive constraints of type (iii).

Now assume that G is a 2-player game and that $i = 2$. Let t'_2 be the strategy added to player 2 in G' . If G' does not preserve indifference in G , then there exists $s_1, t_1 \in S_1$ such that player 1 is indifferent between s_1 and t_1 in G but not in G' : $U_1(s_1, s_2) \neq U_1(t_1, s_2)$ for all s_2 in S_2 but $U_1(s_1, t'_2) \neq U_1(t_1, t'_2)$. Assume w.l.o.g. $U_1(s_1, t'_2) > U_1(t_1, t'_2)$; then, in G' , s_1 weakly dominates t_1 . So the incentive constraint $h'_{t_1, s_1}(\cdot) \geq 0$, which is

¹⁰If G is elementary, then there are no strategies s_j, t_j such that $U_j(s_j, \cdot) = U_j(t_j, \cdot)$ hence G' necessarily preserves indifference in G . The first part of proposition 5.5.5 then reduces to proposition 5.5.3.

nonvacuous, cannot be satisfied with strict inequality. Therefore G' cannot be full. \square

5.6 Second appendix: Proof of propositions 5.4.2 and 5.4.3

We begin with a claim:

Claim 5.6.1. *C has full dimension if and only if (α) there exists a correlated equilibrium μ with full support and (β) there exists x in \mathbb{R}^S such that x satisfies all nonvacuous incentive constraints with strict inequality.*

Proof. If C has full dimension, then it follows from proposition 5.3.1 that (α) and (β) are satisfied. Conversely, assuming that (α) and (β) hold, let $\nu = (1 - \epsilon)\mu + \epsilon x$. For ϵ positive small enough, normalizing ν yields a correlated equilibrium which satisfies all nonvacuous incentive constraints with strict inequality. Therefore, by proposition 5.3.1, C has full dimension. \square

Claim 5.6.1 implies that if there exists a correlated equilibrium with full support, then C has full dimension if and only if (β) holds. We now show that the condition required on top of (α) in proposition 5.4.2 (resp. proposition 5.4.3) imply (resp. is equivalent to) condition (β) . We will use the following standard result:

Lemma 5.6.2. *Let E be a finite dimensional real vector space, q a positive integer, and f_1, \dots, f_q linear forms on E . Then f_1, \dots, f_q are linearly independent if and only if for any y in \mathbb{R}^q there exists x in E such that $y = (f_1(x), \dots, f_q(x))$.*

The notations below are taken from section 5.4.1. Furthermore, if y is a vector then $y > 0$ (resp. $y \geq 0$) means that every coordinate of y is positive (resp. nonnegative). Assume that h_1, \dots, h_m are linearly independent; lemma 5.6.2 then implies that (β) holds, proving proposition 5.4.2. Assume now that $B = (h_1, \dots, h_q)$ is a basis of the linear span of $\{h_1, \dots, h_m\}$, for some $1 \leq q < m$. The value of the auxiliary game of proposition 5.4.3 is positive if and only if

$$\exists y \in \mathbb{R}^q, y \geq 0, \sum_{k=1}^q y_k = 1, yA > 0 \quad (5.6.1)$$

For x in \mathbb{R}^S , let $y(x) = (h_1(x), \dots, h_q(x))$. By definition of the matrix A , we have $(h_{q+1}(x), \dots, h_m(x)) = y(x)A$. Therefore (β) holds if and only if there exists x in \mathbb{R}^S such that $y(x) > 0$ and $y(x)A > 0$. But, by lemma 5.6.2, $y(x)$ may be given any value in \mathbb{R}^q by an appropriate choice of x . Therefore (β) is equivalent to:

$$\exists y \in \mathbb{R}^q, y > 0, yA > 0 \tag{5.6.2}$$

It is easy to see that (5.6.2) is equivalent to (5.6.1), completing the proof of proposition 5.4.3.

Chapitre 6

Geometry, Correlated Equilibrium and Zero-Sum Games

Abstract

We characterize the class of games whose correlated equilibrium polytope contains a Nash equilibrium in its relative interior (unless this relative interior is empty). It is defined by requiring that, in every correlated equilibrium, all incentives constraints stipulating not to deviate to a strategy played with positive probability in at least one correlated equilibrium be tight. This class of games, though not defined by some antagonistic property, is shown to include and generalize two-player zero-sum games.

6.1 Introduction

The set of correlated equilibria of a finite game is a convex polytope which contains the Nash equilibria. Understanding the location of the Nash equilibria within this polytope may allow to get a better understanding of the connections between Nash equilibria and correlated equilibria and to design more efficient algorithm for computing Nash equilibria.¹

¹The idea is to compute Nash equilibria as special correlated equilibria. It is motivated by the fact that the set of correlated equilibria has a much simpler structure than the set

Progress has been made in the last decade: Evangelista and Raghavan (1996) showed that in bimatrix games, extreme Nash equilibria are extreme point of the correlated equilibrium polytope (see also, Gomez Canovas et al (2000) for an alternate proof). More recently, Nau et al (2004) showed that in any nontrivial² n -player game, if the correlated equilibrium polytope has “full” dimension³ then *all* Nash equilibria belong to the relative boundary of this polytope. Improving on this result, we characterize in this chapter the class of games for which the correlated equilibrium polytope contains a Nash equilibrium in its relative interior (unless there is a unique correlated equilibrium, in which case this relative interior is empty). It is defined by requiring that in every correlated equilibrium, all incentive constraints that stipulate not to deviate towards a pure strategy that has positive probability in at least one correlated equilibrium be tight.

This class of games, which we call pre-tight, turns out to be interesting in some other respect: we show that in the two-player case, it includes and generalizes games that are best-response equivalent to a zero-sum game. In particular, Nash equilibria are exchangeable, Nash equilibrium payoffs and correlated equilibrium payoffs coincide, and profiles of marginals of correlated equilibria are Nash equilibria. Up to our knowledge, this is the largest known class of games in which Nash equilibria are exchangeable.

The material is organized as follows: the next section is devoted to basic notations and definitions. In section 6.3, we recall the definition of tight games (Nitzan, 2005) and introduce the class of pre-tight games. These classes of games are defined by special properties of correlated equilibria. However, section 6.4 shows that whether a game is tight (resp. pre-tight) or not may be checked without computing its correlated equilibria. The link between tight and pre-tight games is made precise in section 6.5. Topological properties of the sets of tight and pre-tight games are studied in section 6.6. In section 6.7, we show that the relative interior of the correlated equilibrium polytope contains a Nash equilibrium if and only if the game is tight and does not have a unique correlated equilibrium. Finally, in section 6.8, we show that in the two-player case, pre-tight games include and generalize zero-sum

of Nash equilibria and that efficient (polynomial-time) algorithm to compute a correlated equilibrium are available. See Gilboa and Zemel (1989) for results on computation of Nash equilibria and correlated equilibria.

²As in the previous chapter, a game is called trivial if the payoffs of the players are independent of their own move.

³That is, dimension $N - 1$ where N is the number of pure strategy profiles of the game.

games.

6.2 Notations

The notations are the same as in previous chapters:

$$G = \{I, (S_i)_{i \in I}, (U_i)_{i \in I}\}$$

denotes a finite game in strategic form; I is the nonempty finite set of players, S_i the nonempty finite set of pure strategies of player i and $U_i : \times_{i \in I} S_i \rightarrow \mathbb{R}$ the utility function of player i . We let $S := \times_{i \in I} S_i$ and $S_{-i} := \times_{j \in I-i} S_j$. For any finite set Σ , $\Delta(\Sigma)$ denotes the set of probability distributions over Σ . Finally, N denotes the cardinal of S .

6.2.1 Correlated equilibrium

We recall the definition of correlated equilibria and the notation for incentive constraints introduced in chapter 5. The set $\Delta(S)$ of probability distributions over S is an $N - 1$ dimensional simplex, henceforth called *the simplex*. A *correlated strategy* of the players in I is an element of the simplex. Thus $\mu = (\mu(s))_{s \in S}$ is a correlated strategy if:

$$\text{(nonnegativity constraints)} \quad \mu(s) \geq 0 \quad \forall s \in S \quad (6.2.1)$$

$$\text{(normalization constraint)} \quad \sum_{s \in S} \mu(s) = 1 \quad (6.2.2)$$

For $(i, s_i, t_i) \in I \times S_i \times S_i$, let h_{s_i, t_i} denote the linear form on \mathbb{R}^S which maps $x = (x(s))_{s \in S}$ to

$$h_{s_i, t_i}(x) = \sum_{s_{-i} \in S_{-i}} x(s) [U_i(s) - U_i(t_i, s_{-i})]$$

A correlated strategy μ is a *correlated equilibrium* (Aumann, 1974) if:

$$\text{(incentive constraints)} \quad h_{s_i, t_i}(\mu) \geq 0 \quad \forall i \in I, \forall s_i \in S_i, \forall t_i \in S_i \quad (6.2.3)$$

The set of correlated equilibria is a polytope, which we denote by C .

More notations and vocabulary: Let $s_i \in S_i$, $s \in S$ and $\mu \in \Delta(S)$. The strategy s_i (resp. strategy profile s) is *played* in the correlated strategy μ if

$\mu(s_i \times S_{-i}) := \sum_{s_{-i} \in S_{-i}} \mu(s_i, s_{-i}) > 0$ (resp. $\mu(s) > 0$). In that case, we let $\mu(\cdot|s_i) \in \Delta(S_{-i})$ denote the correlated strategy given s_i of the players other than i :

$$\forall s_{-i} \in S_{-i}, \mu(s_{-i}|s_i) = \frac{\mu(s_i, s_{-i})}{\mu(s_i \times S_{-i})}$$

Finally, the *average payoff of player i in μ* is

$$U_i(\mu) := \sum_{s \in S} \mu(s) U_i(s)$$

6.3 Definitions of tight and pre-tight games

In this section, we define the two classes of games studied in this chapter: tight and pre-tight games.

6.3.1 Tight games

Definition 6.3.1. *A game is tight (Nitzan, 2005) if in any correlated equilibrium all the incentive constraints are tight. Formally,*

$$\forall \mu \in C, \forall i \in I, \forall s_i \in S_i, \forall t_i \in S_i, h_{s_i, t_i}(\mu) = 0 \quad (6.3.1)$$

This means that whenever a pure strategy s_i is played in a correlated equilibrium μ , then every pure strategy of player i is an alternative best-response to $\mu(\cdot|s_i)$.

Example 6.3.2.

$$G_1 = \begin{pmatrix} 1, -1 & -1, 1 \\ -1, 1 & 1, -1 \end{pmatrix} \quad G_2 = \begin{pmatrix} 1, -1 & -1, 1 & 0, -1 \\ -1, 1 & 1, -1 & 0, -1 \end{pmatrix}$$

The game G_1 (i.e. Matching Pennies) is tight. Indeed, it has a unique correlated equilibrium: the Nash equilibrium σ in which both players play $(1/2, 1/2)$. Therefore, G_1 is tight if, in σ , all incentive constraints are tight. Since σ is a completely mixed Nash equilibrium, this is indeed the case.

In contrast, the game G_2 is not tight. Indeed, the mixed strategy profile in which the row player plays $(\frac{1}{2}, \frac{1}{2})$ and the column player $(\frac{1}{2}, \frac{1}{2}, 0)$ is a Nash equilibrium, hence a correlated equilibrium. However, against $(\frac{1}{2}, \frac{1}{2})$, player 2 has a strict incentive not to play her third strategy. More examples will be given in section 6.4.2.

6.3.2 Pretight games

Following Nau et al (2004), let us define a pure strategy (resp. a pure strategy profile) to be *coherent* if it is played in some correlated equilibrium. Formally,

Definition 6.3.3. *The pure strategy s_i (resp. the pure strategy profile s) is coherent if there exists a correlated equilibrium μ such that $\mu(s_i \times S_{-i}) > 0$ (resp. $\mu(s) > 0$).*

Denote by S_i^c the set of coherent pure strategies of player i .

Definition 6.3.4. *A game is pre-tight if in any correlated equilibrium all the incentive constraints stipulating not to “deviate” to a coherent strategy are tight. Formally,*

$$\forall \mu \in C, \forall i \in I, \forall s_i \in S_i, \forall t_i \in S_i^c, h_{s_i, t_i}(\mu) = 0$$

(The above condition is equivalent to

$$\forall \mu \in C, \forall i \in I, \forall s_i \in S_i^c, \forall t_i \in S_i^c, h_{s_i, t_i}(\mu) = 0$$

Indeed, if s_i is not coherent, then for any correlated equilibrium μ , we have $\mu(s_i \times S_{-i}) = 0$, hence $h_{s_i, t_i}(\mu) = 0$ for every t_i in S_i .)

This means that whenever a pure strategy s_i is played in a correlated equilibrium μ , every *coherent* pure strategy of player i is an alternative best-response to $\mu(\cdot | s_i)$. This does not imply that every coherent pure strategy is played in all correlated equilibria. For instance, the game G_3 (below, left) is pre-tight (see proposition 6.3.7). Furthermore, since the correlated strategy μ (below, center) is a completely mixed Nash equilibrium, it follows that every pure strategy is coherent. Nevertheless, the third column is not played in the Nash equilibrium ν (below, right).

$$G_3 = \begin{pmatrix} 1, -1 & -1, 1 & 0, 0 \\ -1, 1 & 1, -1 & 0, 0 \end{pmatrix} \quad \mu = \begin{array}{|c|c|c|} \hline 1/6 & 1/6 & 1/6 \\ \hline 1/6 & 1/6 & 1/6 \\ \hline \end{array} \quad \nu = \begin{array}{|c|c|c|} \hline 1/4 & 1/4 & 0 \\ \hline 1/4 & 1/4 & 0 \\ \hline \end{array}$$

Proposition 6.3.5. *Any tight game is pre-tight.*

Proof. This follows directly from the definitions of tight and pre-tight games. □

As already mentioned in chapter 2:

Proposition 6.3.6. *Any game with a unique correlated equilibrium is pre-tight.*

Proof. We recall the proof: if a game has a unique correlated equilibrium σ , then σ is a Nash equilibrium. Furthermore, the set of coherent strategies of player i is simply the support of σ_i . Therefore, by definition 6.3.4, the game is pre-tight if and only if: $U_i(s_i, \sigma_{-i}) = U_i(t_i, \sigma_{-i})$ for any player i in I and any pure strategies s_i and t_i in the support of σ_i . Since σ is a Nash equilibrium, this condition is satisfied. \square

Proposition 6.3.7. *Any two-player zero-sum game is pre-tight.*

This will be proved in section 6.8, proposition 6.8.1.

6.3.3 Best-Response Equivalence

We note here that whether a game is tight or not (resp. pre-tight or not) depends only on the best-response correspondence, and is in particular unaffected by positive affine transformation of the payoff functions. Consider two games G and G' with the same sets of players and strategies, but with different utility functions:

Definition 6.3.8. *The games G and G' are best-response equivalent if for every player i in I , every pure strategy s_i in S_i , and every correlated strategy μ_{-i} in $\Delta(S_{-i})$, the pure strategy s_i is a best-response to μ_{-i} in G if and only if s_i is a best-response to μ_{-i} in G' .*

Proposition 6.3.9. *If G is tight (resp. pre-tight) then any game that is best-response equivalent to G is tight (resp. pre-tight).*

Proof. In the definitions of tight and pre-tight games, the utility functions only intervene via best-responses to correlated strategies of the other players. The result follows. \square

6.4 Characterization of tight and pre-tight games

This section provides a criterion allowing to check that a game is tight (resp. pre-tight) without having to compute its correlated equilibria.

6.4.1 Statement of the results

We first recall the definition of a dual vector (chapter 2, definition 2.4.1): For each player i in I , let α_i be a transition probability over the set of pure strategies of player i :

$$\begin{aligned}\alpha_i : S_i &\rightarrow \Delta(S_i) \\ s_i &\rightarrow \alpha_i * s_i\end{aligned}$$

The vector $\alpha = (\alpha_i)_{i \in I}$ is a *dual vector* (Myerson, 1997) if

$$\forall s \in S, f(\alpha, s) := \sum_{i \in I} [U_i(\alpha_i * s_i, s_{-i}) - U_i(s)] \geq 0$$

The⁴ following proposition was proved in chapter 2 as proposition 2.6.5:

Proposition 6.4.1. (1) *The game G is tight if and only if there exists a dual vector α such that, for every player i in I and every pure strategy s_i in S_i , the mixed strategy $\alpha_i * s_i$ is completely mixed.*

(2) *The game G is pre-tight if and only if there exists a dual vector α , and, for every player i in I , a subset $S'_i \subseteq S_i$ of pure strategies such that:*

- (a) *For every player i in I and every pure strategy s_i in S'_i , the mixed strategy $\alpha_i * s_i$ has support S'_i .*
- (b) *For every pure strategy profile s in S that does not belong to $S' := \times_{i \in I} S'_i$, we have $f(\alpha, s) > 0$*

In that case, S'_i is the set of coherent pure strategies of player i . That is, $S'_i = S_i^c$.

6.4.2 Examples

This section illustrates the use of proposition 6.4.1 and provides more examples of tight and pre-tight games.

Example 6.4.2 (A pre-tight game).

Consider the following game, due to Bernheim (1984) and studied by Nau and McCardle (1990):

$$\begin{array}{c} \\ T \\ M \\ B\end{array} \quad \begin{array}{ccc} L & M' & R \\ \left(\begin{array}{ccc} 0, 7 & 2, 5 & 7, 0 \\ 5, 2 & 3, 3 & 5, 2 \\ 7, 0 & 2, 5 & 0, 7 \end{array} \right)\end{array}$$

⁴In chapter 2, $f(\alpha, s)$ is denoted $-g(\alpha, s)$.

A way to check that this game is pre-tight is to prove that it has a unique correlated equilibrium: the Nash equilibrium (M, M') , and to apply proposition 6.3.6. An alternative way is to use proposition 6.4.1: define α by $\alpha_1 * T = \alpha_1 * M = \alpha_1 * B = M$ and $\alpha_2 * L = \alpha_2 * M' = \alpha_2 * R = M'$. Let $S'_1 = \{M\}$ and $S'_2 = \{M'\}$. As noted by Nau and McCardle (1990, example 2), α is a dual vector.⁵ Furthermore, if $s_1 \neq M$ and $s_2 \neq M'$, then $f(\alpha, s) = 3$. If $s_1 \neq M$ or $s_2 \neq M'$ (but not both), then $f(\alpha, s) = 1$. Thus, in any case, if $s \notin S'_1 \times S'_2$ (i.e. $s \neq (M, M')$), then $f(\alpha, s) > 0$. By proposition 6.4.1 this implies that the game is pre-tight (and that it has a unique correlated equilibrium).

Example 6.4.3 (General Rock-Paper-Scissors games).

A Rock-Paper-Scissors game is a 3×3 symmetric game in which the second strategy (Paper) beats the first (Rock), the third (Scissors) beats the second, and the first beats the third. Up to normalization (i.e. putting zeros on the diagonal, which is without loss of generality) the payoff matrix of player 1 is of the form:

$$\begin{array}{c} \begin{array}{ccc} & 1 & 2 & 3 \\ \begin{array}{c} 1 \\ 2 \\ 3 \end{array} & \left(\begin{array}{ccc} 0 & -a_2 & b_3 \\ b_1 & 0 & -a_3 \\ -a_1 & b_2 & 0 \end{array} \right) & \text{with } a_k > 0, b_k > 0 \text{ for all } k = 1, 2, 3. \end{array} \end{array} \quad (6.4.1)$$

(Note that we consider general Rock-Paper-Scissors games and not only the zero-sum version. That is, we do not require $b_1 = a_2$ and so forth.)

Proposition 6.4.4. *Any Rock-Paper-Scissors game (6.4.1) is tight.*

Proof. Let $\{1,2,3\}$ denote the set of pure strategies of both players (recall that the game is symmetric). Assume without loss of generality that $a_k + b_k < 1$ for all k in $\{1, 2, 3\}$. Counting k modulo 3, define α_1 as follows: α_1 maps the pure strategy k on the mixed strategy consisting in playing $k+1$ with probability a_k , $k-1$ with probability b_k and k with the remaining probability $1 - a_k - b_k$. Let $\alpha_2 = \alpha_1$. It is clear that $\alpha_i * k$ is completely mixed, for every player i in $\{1, 2\}$ and every pure strategy k in $\{1, 2, 3\}$. Thus, in view of proposition 6.4.1, it only remains to check that α is a dual

⁵Of course, Nau and McCardle (1990) do not use the expression dual vector, which was only later coined by Myerson (1997).

vector. Due to the symmetry and cyclic symmetry of both α and the game, it is enough to check that $f(\alpha, s)$ is nonnegative for $s = (1, 1)$ and $s = (1, 2)$. For $s = (1, 1)$ we get:

$$f(\alpha, s) = 2(a_1[b_1] + b_1[-a_1]) = 0$$

For $s = (1, 2)$, we get

$$f(\alpha, s) = (a_1[a_2] + b_1[a_2 + b_2]) + (a_2[-b_1 - a_1] + b_2[-b_1]) = 0 \quad \blacksquare$$

Example 6.4.5 (An example with an arbitrary large number of players).

The following example (an n -player version of Matching Pennies) generalizes an example which appeared in an earlier version of (Nau et al, 2004). Consider an n -player game G^n in which every player has two pure strategies: K (*keep*) and R (*reverse*). The payoff of player $i \in \{1, 2, \dots, n\}$ is $(-1)^{i+r}$ where r is the number of players playing R .

Proposition 6.4.6. *For every positive integer n , the game G^n is tight*

Proof. If n is even, define α by $\alpha_i(K|R) = \alpha_i(R|K) = 1/2$ for every i in $\{1, 2, \dots, n\}$. If n is odd, hence $n = 2p+1$, define α_i by $\alpha_i(R|K) = \alpha_i(K|R) = \frac{p+1}{2p+1}$ if i is even and by $\alpha_i(R|K) = \alpha_i(K|R) = \frac{p}{2p+1}$ if i is odd. It is easily checked that α is a dual vector. Furthermore, $\alpha_i * s_i$ is completely mixed for every player i in $\{1, 2, \dots, n\}$ and every pure strategy s_i in $\{K, R\}$. By proposition 6.4.1, this implies that the game is tight. \square

6.5 Links between tight and pre-tight games

This section clarifies the links between tight and pre-tight games

Proposition 6.5.1. *A game is tight if and only if it is pre-tight and every pure strategy of every player is coherent.*

Proof. If a game is pre-tight and if all pure strategies are coherent, then it follows from the definitions of tight and pre-tight games that the game is tight. Conversely, if a game is tight, then it is pre-tight, as noted in proposition 6.3.5. Furthermore, it follows from proposition 6.4.1, point (1), that there exists a dual vector α such that, for every s_i in S_i , the mixed strategy $\alpha_i * s_i$ is completely mixed. By proposition 6.4.1, point (2), this implies that $S_i = S_i^c$. \square

To state the next result, we first need to introduce the game G^c obtained from G by restricting the players to their coherent strategies:

$$G^c = \{I, (S_i^c)_{i \in I}, (U_i)_{i \in I}\}$$

Proposition 6.5.2. *A game G is pre-tight if and only if the game G^c is tight.*

We first need a lemma: recall that a Nash equilibrium σ is *quasi-strict* if for every player i in I , any pure best-response to σ_{-i} belongs to the support of σ_i . We proved in chapter 2, proposition 2.6.8 that

Lemma 6.5.3. *Any pre-tight game has a quasi-strict Nash equilibrium with support $S^c = \times_{i \in I} S_i^c$.*

We now prove proposition 6.5.2:

Proof. First, denote by $C^c \subseteq \Delta(S^c)$ the set of correlated equilibria of G^c . Since any correlated equilibrium of G has support in S^c , the set of correlated equilibria of G may be seen as a subset of $\Delta(S^c)$. Since, in G^c , the players have less options than in G , it follows that any correlated equilibrium of G is a correlated equilibrium of G^c . That is, $C \subseteq C^c$.

Second, by definition 6.3.1, the game G^c is tight if and only if

$$\forall \mu \in C^c, h_{s_i, t_i}(\mu) = 0 \quad \forall i \in I, \forall s_i \in S_i^c, \forall t_i \in S_i^c \quad (6.5.1)$$

Similarly, by definition 6.3.4, G is pre-tight if and only if

$$\forall \mu \in C, h_{s_i, t_i}(\mu) = 0 \quad \forall i \in I, \forall s_i \in S_i^c, \forall t_i \in S_i^c \quad (6.5.2)$$

Since $C \subseteq C^c$, it follows that (6.5.1) implies (6.5.2).

Third, we show that (6.5.2) implies (6.5.1) by contraposition.⁶ Assume that (6.5.1) does not hold. Then:

$$\exists \mu \in C^c, \exists i \in I, \exists s_i^* \in S_i^c, \exists t_i^* \in S_i^c, h_{s_i, t_i}(\mu) > 0$$

By lemma 6.5.3, there exists a quasi-strict Nash equilibrium with support S^c , hence a correlated equilibrium σ such that:

$$\forall i \in I, \forall s_i \in S_i^c, \forall t_i \in S_i \setminus S_i^c, h_{s_i, t_i}(\sigma) > 0$$

For $\epsilon > 0$ small enough, $\mu_\epsilon = \epsilon\mu + (1 - \epsilon)\sigma$ is in C . But $h_{s_i^*, t_i^*}(\mu_\epsilon) > 0$. This contradicts (6.5.2). \square

⁶The reason why this implication is not trivial is that the inclusion $C \subseteq C^c$ may be strict, as shown by example 2.7.13 in chapter 2.

To conclude this section note that, as already proved in chapter 2:

Corollary 6.5.4. *Any tight game has a completely mixed Nash equilibrium*

Proof. This follows from lemma 6.5.3 and proposition 6.5.1. \square

6.6 Topology of tight and pre-tight games

In this section we first show that the set of tight (resp. pre-tight) games is neither closed nor open; we then study the size of the class of tight (resp. pre-tight) games.

Example 6.6.1. *Consider the following 2×2 games:*

$$\begin{array}{cc} & \begin{array}{cc} L & R \end{array} \\ \begin{array}{c} T \\ B \end{array} & \left(\begin{array}{cc} \epsilon, -\epsilon & 0, 0 \\ 0, 0 & \epsilon, -1 \end{array} \right) \end{array} \qquad \begin{array}{cc} & \begin{array}{cc} L & R \end{array} \\ \begin{array}{c} T \\ B \end{array} & \left(\begin{array}{cc} \epsilon, \epsilon & 0, 0 \\ 0, 0 & \epsilon, \epsilon \end{array} \right) \end{array}$$

For $\epsilon > 0$, the left game is tight (apply proposition 6.4.1 with α defined by: $\alpha_1(B|T) = \alpha_2(L|R) = \alpha_2(R|L) = \epsilon/2$ and $\alpha_1(T|B) = 1/2$). However, for $\epsilon = 0$, the left game is not even pre-tight, as in the Nash equilibrium (B, L) , player 2 has a strict incentive not to play R . This shows that the set of tight (resp. pre-tight) games is not closed. Furthermore, the game on the right is tight for $\epsilon = 0$, but for $\epsilon > 0$ it is not even pre-tight. This shows that the set of tight (resp. pre-tight) games is not open.

Another issue is the size of the class of tight (resp. pre-tight) games. Fix a positive integer n :

Proposition 6.6.2. *(i) Within the set of $n \times n$ bimatrix games, the set of tight games contains an open set. (ii) If $n \neq m$, then within the set of $n \times m$ bimatrix games, the set of tight games has Lebesgue measure 0.*

Proof. Proof of (i): Nitzan (2005) shows that the set of $n \times n$ bimatrix games with a unique correlated equilibrium and such that this correlated equilibrium is a completely mixed Nash equilibrium, is nonempty and open. It follows from proposition 6.3.6 and proposition 6.5.1 that such games are tight, hence the result.

Proof of (ii): It follows from (von Stengel, 2002, discussion following theorem 2.10) that, if $n \neq m$, the set of $n \times m$ games with a completely mixed Nash equilibrium has Lebesgue measure 0. Since by corollary 6.5.4, any tight game has a completely mixed Nash equilibrium, point (ii) follows. \square

In contrast with point (ii) of proposition 6.6.2, for any number of players n and any positive integers m_1, m_2, \dots, m_n :

Proposition 6.6.3. *The set of n -player pre-tight games of size $m_1 \times m_2 \times \dots \times m_n$ contains a nonempty, open subset of the set of all n -player games of size $m_1 \times m_2 \times \dots \times m_n$.*

Proof. We showed in chapter 3 that the set of n -player games of size $m_1 \times m_2 \times \dots \times m_n$ with a unique correlated equilibrium, is a nonempty, open subset of the set of all n -player games of size $m_1 \times m_2 \times \dots \times m_n$. Since any game with a unique correlated equilibrium is pre-tight, the result follows. \square

Thus, at least for non square bimatrix games, the set of pre-tight games is much bigger than the set of tight games.

6.7 The geometry of Nash equilibria and correlated equilibria

Nau et al (2004) proved the following:

Proposition 6.7.1. *If G has a Nash equilibrium σ in the relative interior of C , then:*

- (a) *The Nash equilibrium σ assigns positive probability to every coherent strategy of every player; that is, σ has support $S^c := \times_{i \in I} S_i^c$.*
- (b) *G is pre-tight.⁷*

For completeness, we recall the proof:

Proof. If (a) is not checked, then σ satisfies with equality some nonnegativity constraint which is not satisfied with equality by all correlated equilibria, hence σ belongs to the relative boundary of C . Assuming now that condition (a) is checked, σ renders indifferent every player among its coherent strategies; therefore σ satisfies with equality all incentive constraints of type $h_{s_i, t_i}(\cdot) \geq 0$, where s_i and t_i are coherent. If G is not pre-tight, at least one of these constraints is not satisfied with equality by all correlated equilibria, hence σ belongs to the relative boundary of C . \square

⁷While Nau et al (2004) introduce the condition defining pre-tight games, they do not give a name to this class of games.

This section proves a converse of this result:

Proposition 6.7.2. *If a game is pre-tight, then either C is a singleton or C contains a Nash equilibrium in its relative interior.*

Proof. We first need a lemma:

Lemma 6.7.3. *Let G be pre-tight and assume that C is not a singleton. A Nash equilibrium belongs to the relative interior of C if and only if it is quasi-strict.*

Proof. First note that if σ is a quasi-strict Nash equilibrium, then σ has necessarily support S^c . Indeed, if $s_i \in S_i^c$, then it follows from definition 6.3.4 that s_i is a best-response to σ_{-i} . Therefore either $\sigma_i(s_i) > 0$ or σ is not quasi-strict. The second sentence of lemma 6.7.3 may thus be rephrased as follows: a Nash equilibrium σ belongs to the relative interior of C if and only if it is quasi-strict and has support S^c . That is,

$$\forall s \in S^c, \sigma(s) > 0 \quad (6.7.1)$$

and

$$\forall i \in I, \forall s_i \in S_i^c, \forall t_i \in S_i \setminus S_i^c, h_{s_i, t_i}(\sigma) > 0 \quad (6.7.2)$$

Now, let σ be a Nash equilibrium of G . By lemma 6.5.3, G has a quasi-strict Nash equilibrium with support S^c , hence a correlated equilibrium checking (6.7.1) and (6.7.2). Therefore, if σ does not check (6.7.1) and (6.7.2) then there exists a nonnegativity constraint or an incentive constraint which is tight in σ but not in all correlated equilibria; hence σ belongs to a strict face of C .

Conversely, assume that σ satisfies (6.7.1) and (6.7.2). Since (6.7.1) and (6.7.2) are strict inequalities, there exists an neighbourhood Ω of σ in \mathbb{R}^S in which they are still satisfied. Let E denote the linear subspace of \mathbb{R}^S consisting of all points $x = (x(s))_{s \in S}$ such that:

(i) $\sum_{s \in S} x(s) = 1$; $\forall s \in S \setminus S^c, x(s) = 0$ and $\forall i \in I, \forall s_i \in S_i \setminus S_i^c, \forall t_i \in S_i, h_{s_i, t_i}(x) = 0$;

(ii) $\forall i \in I, \forall s_i \in S_i^c, \forall t_i \in S_i^c, h_{s_i, t_i}(x) = 0$. Since G is pre-tight, the affine span of C is a subset of E .⁸ Furthermore, $\Omega \cap E \subseteq C$. Finally, since C is not a singleton, E is not a singleton either. Therefore, σ belongs to the relative interior of C . \square

⁸Indeed, if μ is a correlated equilibrium of any game, μ satisfies the constraints in (i), and if the game is pre-tight, then μ satisfies the constraints in (ii).

We can now prove proposition 6.7.2: assume that G is pre-tight. By lemma 6.5.3, G has a quasi-strict Nash equilibrium. If furthermore C is not a singleton, lemma 6.7.3 implies that this Nash equilibrium belongs to the relative interior of C . \square

Corollary 6.7.4. *A game G is pre-tight if and only if C is a singleton or C contains a Nash equilibrium in its relative interior.*

Proof. If C is a singleton (resp. contains a Nash equilibrium in its relative interior) then it follows from proposition 6.3.6 (resp. proposition 6.7.1) that the game is tight. The converse is exactly proposition 6.7.2. \square

Thus, C contains a Nash equilibrium in its relative interior if and only if G is pre-tight and C is not a singleton.

6.8 Two-player pre-tight games

In this section we first show that two-player zero-sum games are pre-tight but that a pre-tight game need not be best-response equivalent to a zero-sum game. We then show that, nevertheless, some of the properties of the equilibria and equilibrium payoffs of zero-sum games extend to pre-tight games. We then discuss the interest and implications of these findings.

6.8.1 Pretight games and zero-sum games

Proposition 6.8.1. *A two-player game which is best-response equivalent to a zero-sum game is pre-tight.*

Proof. In view of proposition 6.3.9 we only need to prove the result for two-player zero-sum games. Consider a two-player zero-sum game with value v . As noted by Forges (1990), and recalled in chapter 4:

(i) If the pure strategy s_1 has positive marginal probability in the correlated equilibrium μ , then the conditional probability $\mu(\cdot|s_1)$ is an optimal strategy for player 2.

It follows that:

(ii) If a pure strategy of player 1 is coherent, then it is a best response to any optimal strategy of player 2.

Indeed, if t_1 is coherent then there exists μ in C and s_2 in S_2 such that $\mu(t_1|s_2)$ is positive. Assume that there exists an optimal strategy σ_2 of player

2 to which t_1 is not a best response. By playing σ_2 against $\mu(\cdot|s_2)$, player 2 would obtain strictly more than $-v$. Therefore $\mu(\cdot|s_2)$ is not an optimal strategy of player 1. This contradicts the analogue of (i) for player 2.

It follows from (i) and (ii) that in every correlated equilibrium μ and for every pure strategy s_i with positive marginal probability in μ , every coherent pure strategy of player 1 is a best-response to $\mu(\cdot|s_1)$. Together with the analogous result for player 2, this shows that the game is pre-tight. \square

The converse of proposition 6.8.1 is false:

Proposition 6.8.2. *A two-player tight game need not be best-response equivalent to a zero-sum game.*

Proof. Recall that every Rock-Paper-Scissors game is tight (proposition 6.4.4). We now show that Rock-Paper-Scissors games need not be best-response equivalent to a zero-sum game: In all bimatrix games that are best-response equivalent to a zero-sum game, fictitious play and its continuous time analog: the best-response dynamics, converge to the set of Nash equilibria (Robinson 1951, Hofbauer and Sorin, 2002). But, in Rock-Paper-Scissors games (6.4.1) such that $a_1a_2a_3 > b_1b_2b_3$, the best-response dynamics does not converge to the unique Nash equilibrium but to a triangle (see, for instance, Hofbauer and Sigmund, 1998). The result follows.⁹ \square

The next section shows that, nevertheless, some of the main properties of two-player zero-sum games extend to pre-tight games. Noticeably, in two-player pre-tight games, the Nash equilibria are exchangeable and any correlated equilibrium payoff is a Nash equilibrium payoff.

6.8.2 Equilibria of pre-tight games

Let us first introduce some notations: we denote by NE the set of Nash equilibria of G and by NE_i the set of Nash equilibrium strategies of player i . That is,

$$NE_i = \{\sigma_i \in \Delta(S_i), \exists \sigma_{-i} \in \times_{j \neq i} \Delta(S_j), (\sigma_i, \sigma_{-i}) \in NE\}$$

Proposition 6.8.3. *In a two-player pre-tight game:*

⁹It follows from this proof that, within the class of 3×3 symmetric games, there exists an open set of games which are tight (hence pre-tight) but are not best-response equivalent to a zero-sum game.

- (a) NE_1 and NE_2 are convex polytopes.
- (b) $NE = NE_1 \times NE_2$. That is, Nash equilibria are exchangeable.

We first need a lemma:

Lemma 6.8.4. *Let G be a two-player pre-tight game and let $\sigma_1 \in \Delta(S_1)$. The following assertions are equivalent:*

- (i) σ_1 is a Nash equilibrium strategy. That is, $\sigma_1 \in NE_1$.
- (ii) For some pure strategy s_2 of player 2, σ_1 is the conditional strategy of player 1 given s_2 in some correlated equilibrium. Formally, $\exists \mu \in C, \exists s_2 \in S_2, \mu(s_2 \times S_1) > 0$ and $\sigma_1 = \mu(\cdot | s_2)$.
- (iii) Every pure strategy in the support in σ_1 is coherent and all coherent strategies of player 2 are best responses to σ_1 .

(The symmetric results for σ_2 in $\Delta(S_2)$ hold obviously just as well.)

Proof. (i) trivially implies (ii) and (ii) implies (iii) by definition of pretight games (definition 6.3.4). So we only need to prove that (iii) implies (i). Let σ_1 check (iii) and let $\tau_2 \in NE_2$. Necessarily, any pure strategy played in τ_2 is coherent. Since any coherent strategy of player 2 is a best response to σ_1 , it follows that τ_2 is a best response to σ_1 . Similarly, by the analogue of (i) \Rightarrow (iii) for player 2, any coherent strategy of player 1 is a best response to τ_2 . Since all pure strategies played in σ_1 are coherent, σ_1 is a best response to τ_2 . Grouping these results, we get that (σ_1, τ_2) is a Nash equilibrium, hence $\sigma_1 \in NE_1$. \square

We now prove proposition 6.8.3: it follows from the proof of lemma 6.8.4 that if $\sigma_1 \in NE_1$, then for any $\tau_2 \in NE_2$, (σ_1, τ_2) is a Nash equilibrium. This implies that Nash equilibria are exchangeable (point (b)). Furthermore, from the equivalence of (i) and (iii) it follows that NE_1 can be defined by a finite number of linear inequalities. Therefore, NE_1 is a polytope, and so is NE_2 by symmetry (point (a)). \blacksquare

Our second result is that if μ is a correlated equilibrium, then the profile of its marginals is a Nash equilibrium:

Proposition 6.8.5. *Let μ be a correlated equilibrium of a two-player pre-tight game. Let $\sigma_1 \in \Delta(S_1)$ (resp. $\sigma_2 \in \Delta(S_2)$) denote the marginal probability distribution of μ on S_1 (resp. S_2). That is, $\forall s_1 \in S_1, \sigma_1(s_1) = \mu(s_1 \times S_2)$. Let $\sigma = (\sigma_1, \sigma_2)$ so that σ is the profile of the marginals of μ . We have:*

- (a) σ is a Nash equilibrium
- (b) The average payoff of the players is the same in σ and in μ . That is, $\forall i \in \{1, 2\}, U_i(\sigma) = U_i(\mu)$.

Proof. First note that σ_2 may be written:

$$\sigma_2 = \sum_{s_1 \in S_1 : \mu(s_1 \times S_2) > 0} \mu(s_1 \times S_2) \mu(\cdot | s_1) \quad (6.8.1)$$

Proof of (a): assume $\mu(s_1 \times S_2) > 0$; then by lemma 6.8.4, $\mu(\cdot | s_1) \in NE_2$. Therefore, by (6.8.1) and convexity of NE_2 , $\sigma_2 \in NE_2$. Similarly, $\sigma_1 \in NE_1$, so that, by proposition 6.8.3, $\sigma \in NE$.

Proof of (b): assume $\mu(s_1 \times S_2) > 0$; then s_1 is coherent and, by the analogue for player 2 of (ii) \Rightarrow (iii) in lemma 6.8.4, any coherent strategy of player 1 is a best response to $\mu(\cdot | s_1)$. Since σ is a Nash equilibrium, every pure strategy in the support of σ_1 is coherent, so that

$$U_1(\sigma_1, \mu(\cdot | s_1)) = U_1(s_1, \mu(\cdot | s_1)) \quad (6.8.2)$$

Using successively (6.8.1), (6.8.2) and a straightforward computation, we get

$$\begin{aligned} U_1(\sigma) &= \sum_{s_1 \in S_1 : \mu(s_1 \times S_2) > 0} \mu(s_1 \times S_2) U_1(\sigma_1, \mu(\cdot | s_1)) \\ &= \sum_{s_1 \in S_1 : \mu(s_1 \times S_2) > 0} \mu(s_1 \times S_2) U_1(s_1, \mu(\cdot | s_1)) = U_1(\mu) \end{aligned}$$

Similarly, $U_2(\sigma) = U_2(\mu)$, completing the proof. \square

Finally, as noted by Forges (1990) and already mentioned in chapter 4, a two-player zero-sum game has a unique Nash equilibrium if and only if it has a unique correlated equilibrium. Since Bohnenblust et al (1950) showed that almost all zero-sum games have a unique Nash equilibrium, this implies that almost all zero-sum games have a unique correlated equilibrium. The next two propositions extend these results to two-player pre-tight games:

Proposition 6.8.6. *A two-player pre-tight game has a unique Nash equilibrium if and only if it has a unique correlated equilibrium.*

(Proof below.) Furthermore,

Proposition 6.8.7. *Within the set of $p \times q$ bimatrix games, the set of pre-tight games which do not have a unique correlated equilibrium has Lebesgue measure 0.*¹⁰

Proof of propositions 6.8.6 and 6.8.7 If C is a singleton, then G has trivially a unique Nash equilibrium. Conversely, let G be a two-player pre-tight game such that C is not a singleton. By proposition 6.7.2, there exists a Nash equilibrium σ in the relative interior of C . Let τ be an extreme Nash equilibrium (in the sense of Evangelista and Raghavan (1996)). Since, in two-player games, an extreme Nash equilibrium is an extreme point of C (Evangelista and Raghavan, 1996), it follows that τ is an extreme point of C . Therefore $\tau \neq \sigma$. This proves proposition 6.8.6.¹¹ Furthermore, since τ does not belong to the relative interior of C , it follows from proposition 6.7.3 that τ is not quasi-strict. Since almost all games have only quasi-strict equilibria (Harsanyi, 1973), this implies proposition 6.8.7 ■

Together with corollary 6.7.4, proposition 6.8.7 shows that for almost all bimatrix games, all Nash equilibria belong to the relative boundary of C . The author does not know whether this extends to games with three or more players. The reason why the above proof does not go through is that in games with three or more players, there need not exist a Nash equilibrium that is an extreme point of C (Nau et al, 2004).

6.8.3 Equilibrium payoffs of pre-tight games

Let NEP (resp. NEP_i , CEP) denote the set of Nash equilibrium payoffs (resp. Nash equilibrium payoffs of player i , correlated equilibrium payoffs). That is,

$$NEP = \{g = (g_i)_{i \in I} \in \mathbb{R}^I : \exists \sigma \in NE, \forall i \in I, U_i(\sigma) = g_i\}$$

$$NEP_i = \{g_i \in \mathbb{R} : \exists \sigma \in NE, U_i(\sigma) = g_i\}$$

$$CEP = \{g = (g_i)_{i \in I} \in \mathbb{R}^I : \exists \mu \in C, \forall i \in I, U_i(\mu) = g_i\}$$

¹⁰By proposition 6.6.3 the set of $p \times q$ pre-tight games contains a nonempty, open subset of the set of $p \times q$ bimatrix games. Therefore, it indeed follows from proposition 6.8.7 that almost all pre-tight games have a unique correlated equilibrium.

¹¹Proposition 6.8.6 also follows, and more directly, from lemma 6.8.4; but the above argument is convenient to prove jointly propositions 6.8.6 and 6.8.7.

Two-player games which are best-response equivalent to zero-sum games may have an infinity of Nash equilibrium payoffs. So pre-tight games need not have a unique Nash equilibrium payoff. Nonetheless, some of the properties of equilibrium payoffs of zero-sum games are preserved. In particular, proposition 6.8.3 and proposition 6.8.5 imply respectively that:

Corollary 6.8.8. *In a two-player pre-tight game, NEP_1 and NEP_2 are convex and $NEP = NEP_1 \times NEP_2$*

Corollary 6.8.9. *In a two-player pre-tight game, $CEP = NEP$*

Thus, allowing for correlation is useless in two-player pre-tight games, in the sense that it cannot improve the payoffs of the players in equilibria. In particular, there are no “good” correlated equilibria in the sense of Rosenthal (1974). Furthermore:

Corollary 6.8.10. *In a two-player pre-tight game, any correlated equilibrium payoff of player i given his move is a Nash equilibrium payoff of player i :*

$$\forall \mu \in C, \forall i \in \{1, 2\}, \forall s_i \in S_i, \mu(s_i \times S_{-i}) > 0 \Rightarrow \sum_{s_{-i} \in S_{-i}} \mu(s_{-i} | s_i) U_i(s) \in NEP_i$$

Proof. For clarity we take $i = 1$. In equation (6.8.2), $(\sigma_1, \mu(\cdot | s_1))$ is a Nash equilibrium (by lemma 6.8.4, proposition 6.8.5(a) and proposition 6.8.3). Therefore, $U_1(s_1, \mu(\cdot | s_1)) = \sum_{s_2 \in S_2} \mu(s_2 | s_1) U_1(s) \in NEP_1$. \square

Finally, there exists a dominant Nash equilibrium. That is,

Proposition 6.8.11. *There exists a Nash equilibrium σ such that*

$$\forall i \in \{1, 2\}, U_i(\sigma) = \max NEP_i \quad (6.8.3)$$

Proof. Let τ, τ' be Nash equilibria such that $U_1(\tau) = \max NEP_1$ and $U_2(\tau') = \max NEP_2$. From exchangeability of equilibria, it follows that $\sigma = (\tau'_1, \tau_2)$ is a Nash equilibrium which satisfies (6.8.3). \square

6.8.4 Discussion

(a) Several classes of non-zero sum games in which some of the properties of two-player zero-sum games are satisfied have been studied. Most are defined in either of these three ways:

(i) by requiring some conflict in the preferences of the players over strategy profiles (“Strictly competitive games” (Aumann, 1961; Friedman, 1983), “Unilaterally competitive games” (Kats and Thisse, 1992));

(ii) by comparing the best or better response correspondence in G and in some zero-sum game (games “best-response equivalent” (Rosenthal, 1974) or “order-equivalent” (Shapley, 1964) to a zero-sum game; “strategically zero-sum games” (Moulin and Vial, 1978));

(iii) by comparing the Nash equilibria or Nash equilibrium payoffs of G and of some auxiliary game (“Almost strictly competitive games” (Aumann, 1961) and other classes of games studied by Beaud (2002)).

The definition of tight and pre-tight games do not fall in these categories; tight games however may be defined by comparing the *correlated equilibria* of G and of some auxiliary game. Indeed, let $-G$ be the game with the same sets of players and strategies as G but in which all the payoffs are multiplied by -1 :

$$-G = \{I, (S_i)_{i \in I}, (-U_i)_{i \in I}\}$$

It is easily checked that G is tight if and only if G and $-G$ have the same correlated equilibria.

(b) Lemma 6.8.4 implies that in two-player *tight* games, as in two-player zero-sum games, the Nash equilibrium strategies of the players can be computed independently, as solutions of linear programs that depend only on the payoffs of the *other* player. In two-player *pre-tight* games, the additional knowledge of the sets of individually coherent strategies is required (indeed the 1×2 games $(0, 1 \mid 0, 0)$ and $(0, 0 \mid 0, 1)$ are both pre-tight and in both games the payoffs of player 1 are the same; but the Nash equilibrium strategies of player 2 are not the same in both games).

(c) A wide range of dynamic procedures converge towards correlated equilibria in all games (for instance generalized no-regret procedures (Hart and Mas-Colell, 2001 and 2003a)). By proposition 6.8.5, suitably modified versions of these dynamics converge towards Nash equilibria in all two-player pre-tight games.

(d) In 3-player tight games, Nash equilibria are not exchangeable (for instance, the game in (Nau et al, 2004, section 6) is tight, since its correlated equilibrium polytope contains a completely mixed Nash equilibrium in its relative interior; but its Nash equilibria are not exchangeable). Up to our knowledge, whether the other properties of section 6.8 extend to n -player games is not known.

Part II

**Evolutionary Dynamics and
Strategic Concepts**

Chapitre 7

Introduction to part II: Evolutionary Dynamics and Strategic Concepts

This introduction is organized as follows: section 7.1 introduces basic notions and methodological issues. Section 7.2 surveys some of the main topics and results of evolutionary game theory. Section 7.3 presents the main contributions of the second part of this dissertation.¹

7.1 Basics of evolutionary game theory

7.1.1 What do we try to understand?

Evolutionary game theory studies the evolution of populations of agents interacting strategically. Applications were originally in biology, but then spread to economics and other social sciences. Examples of situations that evolutionary game theory helps to understand include animal conflicts (Maynard

¹The introduction to evolutionary game theory given below is very basic and almost completely ignores static evolutionary concepts. For more substantial introductions to evolutionary game theory, we refer to the books of Maynard Smith (1982), Weibull (1995), Vega-Redondo (1996), Samuelson (1997), Hofbauer and Sigmund (1998) and Cressman (2003), and to the surveys of Hammerstein and Selten (1994) and Hofbauer and Sigmund (2003). We also recommend, in French, the very clear introduction to the Ph'D dissertation of Philippe Rivière (1997).

Smith and Price, 1973), the evolution of dispersal (Hamilton and May, 1977), or the evolution of conventions (Young, 1993).

The interactions may be with individuals from the same population (males fighting other males, firms competing with similar firms), with individuals from other populations (males interacting with females, rabbits with wolves, or buyers with traders) or with individuals from the same population and individuals from other populations at the same time (plants competing for access to light with many other plants, including plants from their own species and from other species).² The individuals are typically bacteriae or animals in biology, humans or firms in economics.

The interaction is strategic in that there is something to be gain or lost, and that this gain or loss depends both on the behaviour of the agent and on the behaviour of the agents it is interacting with. Behaviours resulting in relatively large gains spread. In biology, this is because animals gaining access to better resources tend to reproduce more and that their behaviour is assumed to be hereditary.³ In economics and other social sciences, this is mostly because social agents tend to imitate currently successful behaviours.

As currently successful behaviours (strategies) spread, the distribution of behaviours in the population changes, so that new behaviours might become successful and spread, while those behaviours that were successful in the past might now decline.⁴ This leads to a potentially complex dynamics for the distribution of behaviours in the population, which we try to understand.

²The behaviour of agents adapting to an uncertain environment (single-player games) is also of interest but belongs to optimization and not to game theory.

³More precisely, it is usually assumed that the behavior of an agent is coded in its genes and that reproduction is clonal so that, unless a mutation arises, the offspring and the parent have the same behavior. The fact that reproduction is assumed clonal is one of the many problematic assumptions in deriving evolutionary game dynamics, since many applications concern sexually reproducing populations. Maynard-Smith (1982), among others, discusses this issue.

⁴In many models in social sciences, agents revise their strategy at some random times and then adopt a better or best-reply to the current population's behaviour. These agents are boundedly rational in that they do not take into account the later evolution of the population's behaviour.

7.1.2 Models

Games. The interaction is modeled by a game. That is, by a specification of the number and type of agents interacting (the players), of the possible and relevant behaviours for each player (his pure strategies), and of a payoff function, which maps every possible behaviour profile to a payoff for each agent: his gain or loss resulting from the interaction.

As most authors, we focus on finite games (a finite number of players, each having a finite set of pure strategies), with multi-linear payoff functions. We even usually focus on two-player games played within a single-population. The underlying paradigm is that there is a large population from which, repeatedly, individuals are randomly drawn to interact. This models the fact that they meet in a certain context.

This is appropriate to model some situations, such as animal conflicts over mates, food, or territories, with which evolutionary game theory originated. However, many biological (or economic) situations call for other models, with a continuous set of players and/or a continuous set of pure strategies. For instance, in the third part of this dissertation, we study the optimal investment into fecundity of some cells. It is then natural to assume that any level of investment into fecundity between two extreme levels is relevant, which we model by assuming that the pure strategy set of a cell is a continuum. Also, quite often in biology, e.g., when studying the evolution of the sex-ratio⁵, it is natural to assume that the individuals interact with the population as a whole, which we might want to model by a continuum of agents.

Dynamics. An evolutionary dynamics is a rule governing the evolution of the behaviours' distribution. We focus on deterministic differential equations (or differential inclusions) but discrete and stochastic dynamics are also of interest. To introduce specific examples, we first need some notations:

Consider a two-player game played within a single-population. Let $\{1, \dots, N\}$ denote the pure strategies available to a player. Let $x_i(t)$ denote the probability at time t that an individual drawn at random will play the pure strategy i . Let $\mathbf{x}(t) := (x_i(t))_{1 \leq i \leq N}$ denote the so-called *mean strategy* in the population (we use bold letters for vectors). Let $U(i, j)$ denote the payoff (utility, fitness gain) of a player playing the pure strategy i facing a player playing the

⁵The evolution of the sex-ratio, i.e., the respective proportions of males and females that an individual should produce, is a traditional topic of evolutionary biology.

pure strategy j . The dynamics we study are (usually) differential equations $\dot{x}(t) = f(\mathbf{x}(t))$ such that strategies which at time t obtain a good payoff tend to spread. For instance, the replicator dynamics (Taylor and Jonker, 1978) is given by:

$$\dot{x}_i(t) = x_i(t) [U(i, \mathbf{x}(t)) - U(\mathbf{x}(t), \mathbf{x}(t))] \quad (7.1.1)$$

where $U(i, \mathbf{x}) = \sum_j x_j U(i, j)$ is the expected payoff of a player playing strategy i when the mean-strategy is \mathbf{x} , and $U(\mathbf{x}, \mathbf{x}) = \sum_i x_i U(i, \mathbf{x})$ is the mean payoff in the population. It may be derived as follows (see Maynard Smith, 1982, appendix D): start with a generation model (time is discrete: $t, t+1, \dots$; individuals live only one period, then they reproduce and die). Let $X_i(t)$ denote the density of individuals playing strategy i at generation t (hence $x_i(t) = X_i(t) / \sum_i X_i(t)$). Assume that the expected number of offspring (or *fitness*) of an individual playing strategy i is

$$W_i := C + U(i, \mathbf{x}(t)) \quad (7.1.2)$$

where C is large enough for W_i to be positive for all i . The constant C represents the background fitness, i.e., the part of the reproduction rate which is not related to the game we study.

Neglecting stochasticities, we get: $X_i(t+1) = W_i X_i(t)$ which leads to

$$x_i(t+1) - x_i(t) = x_i(t) \frac{U(i, \mathbf{x}) - U(\mathbf{x}, \mathbf{x})}{C + U(\mathbf{x}, \mathbf{x})} \quad (7.1.3)$$

When C is large enough, this difference equation is well approximated (on compact intervals of time) by its continuous time counterpart

$$\dot{x}_i(t) = x_i(t) \frac{U(i, \mathbf{x}) - U(\mathbf{x}, \mathbf{x})}{C + U(\mathbf{x}, \mathbf{x})} \quad (7.1.4)$$

Since the denominator is independent of i , omitting it does not change the orbits. This leads to (7.1.1).

Another well studied dynamics is the best-response dynamics (Gilboa and Matsui, 1991), given by:

$$\dot{\mathbf{x}}(t) \in BR(\mathbf{x}(t)) - \mathbf{x}(t) \quad (7.1.5)$$

where⁶ $BR(\mathbf{x})$ is the set of mixed strategies that are best-responses to \mathbf{x} . The underlying model is that at each period of time, a fraction of the population

⁶Note that the best-response dynamics is a differential inclusion; that is, $\dot{\mathbf{x}}$ may be multi-valued.

revises its choice of strategy and chooses a best-response to the current state of the population.

These two dynamics are those that we will study most. The replicator dynamics, under which every strategy earning above average has a positive growth-rate, may be seen as the benchmark dynamics for applications in biology; the best-response dynamics, under which only strategies that are best-responses to the current state have a positive growth rate, is more adapted to economic applications, with myopic but otherwise rational, actively optimizing agents.

Crudeness of derivations and benchmark dynamics. We will not expand more on the derivations of these and other dynamics but to point out that they are usually rather crude. For instance, in the above derivation of the replicator dynamics, we ignored the intricacies of sexual reproduction, the age structure of the population, stochastic aspects, etc. This is appropriate as a first step: introducing additional layers of complexity is bound to be fruitless as long as we do not understand simple models, and focusing on very specific situations may yield results of limited applicability. However, as in all models, we should worry that aspects of real-life situations that are being ignored might alter our results.

This is not to say that studying extensively a particular dynamics, however crudely derived, is not of interest: this allows to go deeper into our understanding of this dynamics and to discover phenomena or properties that we can then investigate in larger classes of dynamics. But we should always wonder whether the properties that we have established are likely to hold for similar dynamics or whether they are really specific of the dynamics we studied.⁷

Lotka-Volterra dynamics and replicator dynamics. There is a close analogy between the evolution of the abundance of different species in an ecological system and the evolution of the relative abundance of different behaviours in a single species or population. This analogy is grounded mathematically. In particular, the ecological Lotka-Volterra dynamics is closely related - “mathematically equivalent” - to the game-theoretical replicator dynamics (see Hof-

⁷For instance, many properties of the time-average of the replicator dynamics do not hold for dynamics similar to the replicator dynamics. See, e.g., Gaunersdorfer and Hofbauer, 1995, section 4.

bauer and Sigmund, 1998). Evolutionary game dynamics are thus also relevant for the study of competition between species.⁸

Biological and economical dynamics. We conclude this section by discussing a difference between ecological and economical dynamics. In an ecological context, it is highly improbable that a large share of the population will suddenly mutate to some new strategy. Rather, mutations are rare, and even when a strategy is highly favoured, the speed at which it spreads is limited by the current number of agents playing this strategy. This is because the number of offspring that an agent can have, and a fortiori the effect of the game on the number of offspring that an agent can have, is bounded. For most biological applications, the role of mutations is confined to introducing novelty, and mutations terms may be omitted as in the replicator dynamics, as long as we focus on interior solutions; that is, solutions with all species present at all times.⁹

In an economic context, it might be that the speed at which new behaviours spread is bounded too. For instance, if agents need to encounter an agent playing a “mutant” strategy, to become aware of the possibility of playing this strategy. Such assumptions lead to dynamics resembling biological dynamics. But we might also consider that the agents are aware of all the strategies they could use, and periodically revise their strategy to a best (or good) reply to the current state of the population (as in the derivation of the best-response dynamics sketched above). In such situations, the growth \dot{x}_i of the share of the population playing strategy i is not limited by

⁸For instance, Bomze (1983) uses the phase portraits’ classification of the two-dimensional replicator dynamics to classify the phase portraits of the two-dimensional Lotka-Volterra dynamics.

⁹Even when focusing on interior solutions, mutations cannot always be ignored: for instance, in chapter 10, proposition 10.9.5, we give an example of an heteroclinic cycle which, under the replicator dynamics (7.1.1), is not asymptotically stable, but still attracts an open and - for some parameters - large set of solutions. If we want to understand the long-run behaviour of the system we model (and not only of the system of differential equations we model it with), it is crucial to note that in a finite population model *à la* Kandori et al (1993), a *single* mutation suffices to take the population out of the basin of attraction of this heteroclinic cycle. This issue is different from the consideration of *large* deviations to select between strict equilibria (Foster and Young, 1990; Kandori et al, 1993), where a *large* number of almost simultaneous mutations - a very rare event - is needed to switch from the basin of attraction of an equilibrium to the basin of attraction of another equilibrium.

the current value of x_i . This calls for *innovative dynamics* for which, even if stochasticities are neglected, new strategies may appear and spread quickly, so that the growth-rate of a strategy: \dot{x}_i/x_i , is unbounded.

We now present some of the main research topics in evolutionary game theory.

7.2 Research topics in evolutionary game theory

Main issues. There are two main questions: what is the outcome of evolution? What are the links between the outcome of evolution and static solution concepts (Nash equilibrium, correlated equilibrium, rationalizability, etc.)? The first question seems to us the most important one. But the second serves to orientate research, and is important per se, to decide whether results based on the paradigm of rational agents (typically, common knowledge of rationality) are still relevant once the paradigm is changed to populations of agents adapting to the environment through rules of thumb.

Convergence to and stability of equilibria. A first issue is whether the outcome of evolution tends to be a Nash equilibrium. The main positive result is the so-called folk theorem of evolutionary game theory. It states that, for a very wide class of dynamics: first, for every initial condition with all strategies present, if an interior trajectory converges to a point, then this point is a Nash equilibrium: second, every (Lyapunov) stable state is a Nash equilibrium.^{10,11} Furthermore, for several dynamics and classes of games (potential games, games with an interior ESS, dominance solvable games, etc.) global convergence to Nash equilibria has been established. This might be seen as a foundation for an “as-if” interpretation of Nash equilibrium play, and John Nash had actually an evolutionary interpretation in mind when conceiving his equilibrium concept:

¹⁰A point \mathbf{p} is Lyapunov stable (under some dynamics) if for every neighborhood N_1 of \mathbf{p} , there exists a neighborhood N_2 of \mathbf{p} such that if the solution starts in N_2 , then it remains in N_1 for ever.

¹¹This holds in particular for any (smooth) dynamics for which, whenever not all strategies earn the same payoff, the growth rate of at least one of the strategies earning strictly less than average is negative (or the growth rate of at least one of the strategies earning strictly more than average is positive). See Weibull, 1995, proposition 4.11.

“[Assume] that there is a population (...) of participants for each position of the game. Let us also assume that the “average playing” of the game involves n participants selected at random from the n populations, and that there is a stable average frequency with which each pure strategy is employed by the “average member” of the appropriate population. (...) The assumptions (...) made in this “mass-action” interpretation lead to the conclusion that the mixed strategies representing the average behavior in each of the populations form an equilibrium point.” (Nash, 1950, p.21-23)

However, Nash equilibria need not be (Lyapunov) stable.¹² For instance, in so called outward cycling Rock-Paper-Scissors games, the Nash equilibrium is unique, completely mixed, and passes all (non-evolutionary) equilibrium refinements¹³, but it is unstable under the replicator dynamics and many other evolutionary dynamics (see, e.g., section 7.7. in Hofbauer and Sigmund, 1998, or chapter 10, section 10.2 of this dissertation).¹⁴

With respect to stability, ESS (evolutionary stable strategies), a refinement of Nash equilibrium, fare better: ESS are asymptotically stable, and interior ESS globally asymptotically stable under many dynamics including the replicator and best-response dynamics (Hofbauer, 2000). But ESS may fail to exist, as in Rock-Paper-Scissors games.

Nonconvergence to Nash equilibrium is not a feature of only a few dynamics, but appears to be a universal phenomenon. In particular, Hofbauer and Swinkels (personal communication) showed that cyclic behaviour is “universal”, in that it occurs for essentially all adjustment dynamics¹⁵ of Swinkels (1993). See (Hofbauer and Sigmund, 1998, section 8.6). Related results have been obtained by Hart and Mas-Collel (2003b). Once acknowledged the possibility of nonconvergence, attention turns to general properties of evolutionary dynamics which do not depend on whether the solution converges or not.

¹²Paraphrasing the above quotation, there need not be a “stable average frequency with which each pure strategy is employed”.

¹³Except that it is not strict.

¹⁴For multi-population dynamics, the contrast between the refinement literature and the evolutionary games literature is even starker: while completely mixed equilibria pass most refinements in van Damme (1991), they are always unstable for wide classes of evolutionary dynamics. See, e.g., Ritzberger and Weibull, 1996.

¹⁵Adjustment dynamics are, essentially, those dynamics under which the mean strategy evolves towards better reply to the current state: a rather minimal condition of adaptivity.

Permanence. A very important issue, from an ecological perspective, is whether diversity will be maintained. That is, whether, with high probability, all behaviours initially present will still be observed in the long run, or whether some will go extinct. This issue has been addressed through the concept of permanence. Boldly said, a deterministic ecological dynamical system is permanent if there exists a positive threshold ϵ such that, for any initial condition with all species present, all species will eventually be present in a proportion greater than ϵ . Of course, even if the evolution of the distribution of behaviors in a population is well described in the short run by a permanent dynamical system, large stochastic fluctuations will drive some behaviours extinct; Hofbauer and Sigmund (1998, part III) give necessary conditions and sufficient conditions for some systems to be permanent, and study the related concept of *persistence*.

Classes of strategies that tend to survive and rationalizability. For non-permanent system, the main issue is to determine the strategies that survive and those that are eliminated. There is a parallel between trying to define classes of strategies that survive (or classes of strategies that are eliminated) and weak static solution concepts such as elimination of dominated strategies and rationalizability. In both cases, the basic tenet is that it would be too ambitious to try to predict a single outcome but that some outcomes may be discarded as unreasonable or selected against. Accordingly, many researchers studied whether evolutionary dynamics eliminate dominated strategies. We survey some of this literature in chapter 8. Suffice here to say that, for many dynamics, including all monotonic¹⁶ dynamics and the best-response dynamics, it has been found that for every game and every interior initial condition, all pure strategies that are iteratively strictly dominated are eliminated. A dual statement is that only rationalizable pure strategies survive.

Time-averages of dynamics and equilibria. Another avenue of research is to investigate the behaviour of time-averages of dynamics, and whether they are related to Nash equilibrium, correlated equilibrium or other standard concepts. For instance, for permanent systems (and actually along all solutions that are eventually bounded away from the boundary of the state-space), the

¹⁶A dynamics is *monotonic* if, on top of some smoothness assumptions, whenever strategy i earns more than strategy j , its growth-rate is higher than the growth-rate of strategy j .

time-average of the (single-population or two-population) replicator dynamics converges to a Nash equilibrium (Hofbauer and Sigmund, 1998). In any finite game, the time-average of the (multi-population) replicator dynamics converges to the set of weak correlated equilibria¹⁷ (Hofbauer, 2004). As another example of a link between the time-average of a dynamics and an equilibrium concept, consider the following $2 \times 2 \times 2$ game, a variant of Jordan's (1993) 3-player matching pennies game studied by Gaunersdorfer and Hofbauer (1995):

$$\begin{pmatrix} -1, -1, -1 & -1, +1, +1 \\ +1, +1, -1 & +1, -1, +1 \end{pmatrix} \quad \begin{pmatrix} +1, -1, +1 & +1, +1, -1 \\ -1, +1, +1 & -1, -1, -1 \end{pmatrix}$$

This game has a unique Nash equilibrium, in which every player plays $(1/2, 1/2)$. This gives a payoff of 0 (for every player). The probability distribution which puts probability 0 on the two squares with payoff -1 for all players and probability $1/6$ on each of the six other squares is a correlated equilibrium, with payoff $1/3$ (for every player). For almost all initial conditions, the time average of the three-population best-response dynamics converges to this correlated equilibrium. This follows from the analysis of the best-response dynamics in this game made by Gaunersdorfer and Hofbauer (1995, section 5, in particular p. 298). See also (Jordan, 1993).

7.3 Contribution of this dissertation

Subject and motivations. We investigate the link between strategies that tend to survive under evolutionary dynamics and strategies belonging to the support of a Nash equilibrium or of a correlated equilibrium. This is related to two of the main issues of evolutionary game theory: trying to define classes of strategies that tend to survive (or to be eliminated), and investigating the relevance of standard, static concepts for evolutionary dynamics.

There are three more specific motivations. First, we saw in the preceding section that even when a solution of an evolutionary dynamics does

¹⁷Let G be a finite game, S_i the set of pure strategies of player i and $U_i : S = \times_i S_i \rightarrow \mathbb{R}$ its payoff function. A weak correlated equilibrium (Moulin and Vial, 1978) is a probability distribution μ on the set of strategy profiles S such that, for every player i and every pure strategy t_i in S_i , $\sum_{s \in S} \mu(s)[U_i(s) - U_i(s_{-i}, t_i)] \geq 0$ (where (s_{-i}, t_i) is the strategy profile which differs from s only in that its i^{th} component is t_i). The set of weak correlated equilibria is also called the *Hannan set*.

not converge, there might still be a link between its long-run behavior and equilibrium concepts. Second, a number of adaptive procedures converge, in a time-average sense, towards the set of correlated equilibria (for a survey, see Hart, 2005). While these procedures, often based on a notion of regret, differ from the evolutionary dynamics that we study¹⁸, it is natural to ask whether correlated equilibria enjoy similar good properties under evolutionary dynamics; in particular, whether correlated equilibrium is a better concept than Nash equilibrium from an evolutionary perspective. Finally, as already mentioned in the first part of this dissertation, Nau and McCardle (1990) characterized the set of strategy profiles that do not belong to the support of any correlated equilibrium. They show that these strategy profiles are aggregatively bad, in the sense that they expose the group of players to arbitrage from an outside observer¹⁹. This suggests that strategies never played in correlated equilibrium might be eliminated under some evolutionary dynamics.

Small dimension: Elimination of all strategies not played in equilibrium For 2×2 games and 3×3 symmetric games, this intuition is correct. Indeed, for these games, Nau and McCardle's characterization implies that if a pure strategy is never played in correlated equilibrium, then certain domination relations hold, which imply that for any interior initial condition, this strategy is eliminated by the (two-population) replicator dynamics, the best-response dynamics and any convex monotonic dynamics. For the *single*-population replicator dynamics, the fact that, in 3×3 symmetric games, all strategies that are not used in correlated equilibrium are eliminated, follows from a more basic result: all strategies not used in Nash equilibrium are eliminated. More precisely, in a 3×3 symmetric game, either there exists a completely mixed Nash equilibrium, or for any interior initial condition, the replicator dynamics converges to the set of Nash equilibria. This may be seen by examining Bomze's (1983) classification of the single-population replicator dynamics in 3×3 symmetric games (see also Zeeman, 1980); we give a direct proof in chapter 9 and show that the same result holds for the best-response dynamics. The fundamental reason for this result, so we think, is that the

¹⁸In particular, in no-regret dynamics, the state variable is a time-average and through this state variable, current play depends on the whole past. In contrast, the evolutionary dynamics that we study here have no memory and the state variable is simply current play.

¹⁹Mathematically, this takes the form of proposition 2.5.5 in chapter 2

single-population dynamics in 3×3 games is a two-dimensional smooth dynamical system evolving in a compact region of the plane. Due to, e.g., the Poincaré-Bendixson theorem, there are severe constraints on the ways such a dynamical system can behave.²⁰

The general case: All strategies used in correlated equilibrium may be eliminated. Summing up, chapter 9 shows that for small dimensional games, strategies which do not belong to the support of any correlated equilibrium are eliminated by many dynamics. This does not extend to higher dimensions. Indeed, chapter 10 shows that in 4×4 symmetric games, all strategies played in correlated equilibrium may be eliminated, so that only strategies that do *not* belong to the support of any correlated equilibrium remain. The examples we build are very much related to Dekel and Scotchmer's (1992) example of a game for which a strictly dominated pure strategy survives under a discrete-time version of the replicator dynamics.

Elimination of all strategies in the support of correlated equilibria is very robust (it occurs for many dynamics, including all monotonic dynamics, and for each of these dynamics, for an open set of games and an open set of initial conditions; for the replicator dynamics, it is robust to perturbation of the vector field, occurs in continuous-time as well as in discrete-time, and for an arbitrarily large set of initial conditions).

Universal elimination of all strategies used in Nash equilibrium. For Nash equilibrium, even more can be shown: in particular, elimination of all strategies used in Nash equilibrium occurs for essentially all adjustment dynamics. See chapter 11. In our examples, there is a unique, strict Nash equilibrium. In particular, this equilibrium is an ESS and satisfies all refinements from the literature. Thus, failure to converge to Nash equilibrium does not come from a failure to coordinate on the right equilibrium or from an absence of perfectness.

Finally, we show that for the replicator and best-response dynamics, all strategies played in Nash equilibrium may be eliminated from almost all initial conditions (chapter 12). Here again, the Nash equilibrium is unique and

²⁰In contrast, the *two-population* replicator dynamics in 3×3 games is 4-dimensional, hence from this point of view, our result on elimination of all strategies unused in correlated equilibrium is non-trivial; this is still due to a small dimension effect, but not of the same type. See chapter 9.

passes most refinements (all refinements studied by van Damme (1991), but ESS and strictness).

Relevance of Nash and correlated equilibria for evolutionary dynamics. So what should we think of Nash and correlated equilibria? Are they relevant to evolutionary dynamics? Are they not? Our position is mixed. Much effort has been devoted to “justify” traditional concepts by finding dynamics converging to equilibria. This need not be the right approach: if we think of evolutionary game theory as a tool to describe the evolution of human or animal behaviour, then it seems more fruitful to try to design dynamics resembling this behaviour, rather than dynamics converging to equilibria. However, it cannot be ignored that, as previously mentioned, many dynamics do indeed converge to Nash equilibria in several important classes of games. Furthermore, even if we study disequilibrium behaviour, or games for which all strategies used in equilibrium are eliminated, Nash equilibria remain important as rest points of evolutionary dynamics.

Chapters 10, 11 and 12 should not be seen as an attempt to disqualify the Nash equilibrium concept or the correlated equilibrium concept. They suggest though that other, more dynamical solution concepts are of interest, such as limit cycles or invariant measures.

Chapitre 8

Elimination of Dominated Strategies

Abstract

We survey and unify results on elimination of dominated strategies by evolutionary dynamics.

8.1 Introduction

One of the main issues in evolutionary game theory is whether evolutionary dynamics lend support to strategic concepts. A most basic strategic concept is the elimination of dominated strategies and, accordingly, whether evolutionary dynamics wipe out dominated strategies has been studied by a number of authors (Akin, 1980; Nachbar, 1990; Dekel and Scotchmer, 1992; Samuelson and Zhang, 1992; Cabrales and Sobel, 1992; Hofbauer and Weibull, 1996; Berger and Hofbauer, 2005; and others). Many results concern monotonic dynamics and are proved in a very similar way. We unify these results, and prove some new results on concave monotonic dynamics (defined in the next section). We also discuss elimination of dominated strategies by the best-response dynamics and by the Brown-von Neumann-Nash dynamics. The material is organized as follows. In the next section we precise the framework and define several classes of dynamics. Section 8.3 is devoted to positive results and section 8.4 to negative results. Finally, section 8.5 considers discrete-time dynamics.

8.2 Framework and classes of dynamics

The setting we introduce is rather abstract in order to encompass a variety of situations, including 1-player decision problem, games within a population or games between individuals from different populations, which need not evolve according to the same dynamics.

We consider a large population of players from which individuals are repeatedly and randomly drawn to play a game against some opponent, which could be the weather, or (one or several) individuals from the same population and/or other populations. For instance, if the underlying game is a n -person game, then we focus on player 1 and consider players 2 to n as a single entity, which we call the opponent (whether players 2 to n can correlate their actions or not will not be relevant). In what follows, unless mentioned otherwise, “player” means player 1, “the population” means the population of players 1 and “a strategy” means a strategy of player 1.

We assume that the set $I := \{1, \dots, N\}$ of pure strategies of player 1 is finite and that the set S^{opp} of strategies of the opponent is compact. Let S_N denote the $N - 1$ dimensional simplex over I :

$$S_N = \left\{ \mathbf{x} \in \mathbb{R}_+^N, \sum_{i \in I} x_i = 1 \right\}$$

and $\text{int } S_N := \{\mathbf{x} \in S_N : \forall i \in I, x_i > 0\}$ its interior. Let $\mathbf{x}(t) := (x_i(t))_{i \in I} \in S_N$ denote the mean-strategy in the population at time t , and $\mathbf{y}(t)$ the opponent’s strategy at time t .¹

We assume that the population adapts to the opponent’s strategy so that the mean-strategy $\mathbf{x}(t)$ evolves over-time.² This is modeled by assuming that $\mathbf{x}(t)$ follows a differential equation or, possibly, a differential inclusion. We will mostly consider dynamics of the form:

$$\dot{x}_i(t) = x_i(t) \left[g_i(\mathbf{x}(t), \mathbf{y}(t)) - \sum_{k \in I} x_k(t) g_k(\mathbf{x}(t), \mathbf{y}(t)) \right] \quad (8.2.1)$$

¹For instance, if the game consists of an interaction with an individual of another population, then $\mathbf{y}(t)$ would be the mean strategy at time t in this other population; if the game consists of a symmetric interaction with an individual of the same population, then $\mathbf{y}(t) = \mathbf{x}(t)$.

²Evolution may take place by hereditary or cultural transmission of strategies to offspring and differential reproduction of individuals playing different strategies, or by imitation of currently successful strategies.

where the functions g_i are continuous. Note that S_N and its boundary faces are invariant under (8.2.1). The best-known example of such a dynamics is the replicator dynamics given by (time indices suppressed):

$$\dot{x}_i = x_i [U(i, \mathbf{y}) - U(\mathbf{x}, \mathbf{y})] \quad (8.2.2)$$

where $U(i, \mathbf{y})$ is the payoff of player 1 when playing the pure strategy i and facing strategy \mathbf{y} , and $U(\mathbf{x}, \mathbf{y}) := \sum_{i \in I} x_i U(i, \mathbf{y})$. Throughout, the payoff function $U(\cdot, \cdot)$ is assumed to be continuous.

We say that $g_i(\mathbf{x}(t), \mathbf{y}(t))$ is the unnormalized growth rate of strategy i at time t . By analogy, if \mathbf{p} is a mixed strategy, we say that

$$g_{\mathbf{p}}(\mathbf{x}, \mathbf{y}) := \sum_{i \in I} p_i g_i(\mathbf{x}, \mathbf{y}) \quad (8.2.3)$$

is the *unnormalized growth rate of the mixed strategy \mathbf{p}* . To understand this name note that in the replicator dynamics, $g_{\mathbf{p}}(\mathbf{x}, \mathbf{y}) = U(\mathbf{p}, \mathbf{y})$ so that

$$U(i, \mathbf{y}) = U(\mathbf{p}, \mathbf{y}) \Rightarrow g_{\mathbf{p}}(\mathbf{x}, \mathbf{y}) = g_i(\mathbf{x}, \mathbf{y}) \quad \forall (\mathbf{x}, \mathbf{y}) \quad (8.2.4)$$

We now introduce some classes of dynamics. Let i and j denote pure strategies and \mathbf{p} and \mathbf{q} mixed strategies.

Definition 8.2.1. *A dynamics (8.2.1) is monotonic³ if*

$$U(i, \mathbf{y}) > U(j, \mathbf{y}) \Rightarrow g_i(\mathbf{x}, \mathbf{y}) > g_j(\mathbf{x}, \mathbf{y}) \quad \forall (i, j, \mathbf{x}, \mathbf{y}) \quad (8.2.5)$$

It is convex monotonic (Hofbauer and Weibull, 1996) if furthermore

$$U(\mathbf{p}, \mathbf{y}) > U(i, \mathbf{y}) \Rightarrow g_{\mathbf{p}}(\mathbf{x}, \mathbf{y}) > g_i(\mathbf{x}, \mathbf{y}) \quad \forall (i, \mathbf{p}, \mathbf{x}, \mathbf{y}) \quad (8.2.6)$$

It is strictly convex monotonic if furthermore

$$\begin{aligned} [U(\mathbf{p}, \mathbf{y}) = U(i, \mathbf{y}) \quad \text{and} \quad \exists (k, l) \in I \times I, p_k p_l > 0, U(k, \mathbf{y}) \neq U(l, \mathbf{y})] \\ \Rightarrow g_{\mathbf{p}}(\mathbf{x}, \mathbf{y}) > g_i(\mathbf{x}, \mathbf{y}) \end{aligned} \quad (8.2.7)$$

The dynamics is concave monotonic if it is monotonic and

$$U(i, \mathbf{y}) > U(\mathbf{p}, \mathbf{y}) \Rightarrow g_i(\mathbf{x}, \mathbf{y}) > g_{\mathbf{p}}(\mathbf{x}, \mathbf{y}) \quad \forall (i, \mathbf{p}, \mathbf{x}, \mathbf{y}) \quad (8.2.8)$$

³This property is called *relative monotonicity* by Nachbar (1990), *order-compatibility of pre-dynamics* by Friedman (1991), *monotonicity* by Samuelson and Zhang (1992), which we follow, and *payoff monotonicity* by Hofbauer and Weibull (1996).

It is strictly concave monotonic if furthermore,

$$\begin{aligned} [U(i, \mathbf{y}) = U(\mathbf{p}, \mathbf{y}) \quad \text{and} \quad \exists(k, l) \in I \times I, p_k p_l > 0, U(k, \mathbf{y}) \neq U(l, \mathbf{y})] \\ \Rightarrow \quad g_i(\mathbf{x}, \mathbf{y}) > g_{\mathbf{p}}(\mathbf{x}, \mathbf{y}) \end{aligned} \quad (8.2.9)$$

Finally, the dynamics is aggregate monotonic (Samuelson and Zhang, 1992) if it is both convex and concave monotonic; that is, if

$$U(\mathbf{p}, \mathbf{y}) > U(\mathbf{q}, \mathbf{y}) \Leftrightarrow g_{\mathbf{p}}(\mathbf{x}, \mathbf{y}) > g_{\mathbf{q}}(\mathbf{x}, \mathbf{y}) \quad \forall(\mathbf{p}, \mathbf{q}, \mathbf{x}, \mathbf{y}) \quad (8.2.10)$$

An example of an aggregate monotonic dynamics is the replicator dynamics⁴ (8.2.2). To understand the name convex monotonic, note that if there exists an increasing function g such that

$$g_i(\mathbf{x}, \mathbf{y}) = g(U(i, \mathbf{y})) \quad \forall(i, \mathbf{x}, \mathbf{y})$$

then the dynamics is (strictly) convex monotonic if and only if g is (strictly) convex in \mathbf{x} , and (strictly) concave monotonic iff g is (strictly) concave in \mathbf{x} .

Any aggregate monotonic satisfies (8.2.4). In this sense, we may think of aggregate monotonic dynamics as being neutral between mixed and pure strategies. In contrast, in the sense of (8.2.7) and (8.2.9), strictly convex (resp. strictly concave) monotonic dynamics give an advantage to mixed (resp. pure) strategies. This explains that pure strategies are more easily eliminated by convex monotonic dynamics than by concave monotonic dynamics, as will be clear from the next section.⁵

8.3 Positive results

Definition 8.3.1. *The pure strategy $i \in I$ is eliminated if $x_i(t) \rightarrow 0$ as $t \rightarrow +\infty$. The mixed strategy $\mathbf{q} \in S_N$ is eliminated if $\min_{\{i \in I: q_i > 0\}} x_i(t) \rightarrow 0$ (or equivalently $\prod_{i \in I} x_i^{q_i}(t) \rightarrow 0$).*

⁴Actually, in the context of bimatrix games, Samuelson and Zhang (1992, theorem 3) have shown that for any aggregate monotonic dynamics, there exists a positive speed function λ such that $\dot{\mathbf{x}}(t) = \lambda(\mathbf{x}(t), \mathbf{y}(t)) \dot{\mathbf{x}}_{REP}(t)$, where $\dot{\mathbf{x}}_{REP} = x_i(U(i, \mathbf{y}) - U(\mathbf{x}, \mathbf{y}))$. For single-population dynamics ($\mathbf{y}(t) = \mathbf{x}(t) \forall t$), this implies that any aggregate monotonic dynamics has the same orbits than the replicator dynamics; for multi-population dynamics, this need not be so, because the speed function is population specific.

⁵Conversely, mixed strategies are more easily “eliminated” by concave than by convex monotonic dynamics; see next section.

Definition 8.3.2. *The mixed strategy \mathbf{q} is strictly dominated by the mixed strategy \mathbf{p} if*

$$U(\mathbf{p}, \mathbf{y}) > U(\mathbf{q}, \mathbf{y}) \quad \forall \mathbf{y} \in S^{opp} \quad (8.3.1)$$

The mixed strategy \mathbf{q} is weakly dominated by the mixed strategy \mathbf{p} if

$$U(\mathbf{p}, \mathbf{y}) \geq U(\mathbf{q}, \mathbf{y}) \quad \forall \mathbf{y} \in S^{opp} \quad (8.3.2)$$

and furthermore there exists $\mathbf{y} \in S^{opp}$ such that $U(\mathbf{p}, \mathbf{y}) > U(\mathbf{q}, \mathbf{y})$.

Note that if the mixed strategies \mathbf{q} and $\tilde{\mathbf{q}}$ have the same support and that \mathbf{q} is eliminated then so is $\tilde{\mathbf{q}}$. In this sense, whether a mixed strategy is eliminated depends only on its support.^{6,7}

Now to the results. We first focus on aggregate monotonic dynamics. Let \mathbf{p} and \mathbf{q} be two mixed strategies. The main result is the following:

Proposition 8.3.3. *Assume that $\mathbf{x}(0) \in \text{int } S_N$. Assume furthermore that*

$$\limsup_{t \rightarrow +\infty} (U(\mathbf{q}, \mathbf{y}(t)) - U(\mathbf{p}, \mathbf{y}(t))) < 0 \quad (8.3.3)$$

or that the following holds: There exists T in \mathbb{R} such that

$$\forall t \geq T, U(\mathbf{q}, \mathbf{y}(t)) \leq U(\mathbf{p}, \mathbf{y}(t)), \quad (8.3.4)$$

$\mathbf{y}(\cdot)$ is uniformly continuous, and

$$\liminf_{t \rightarrow +\infty} [U(\mathbf{q}, \mathbf{y}(t)) - U(\mathbf{p}, \mathbf{y}(t))] < 0 \quad (8.3.5)$$

Under any aggregate monotonic dynamics⁸, $\prod_{i \in I} x_i^{q_i - p_i}(t) \rightarrow 0$ hence strategy \mathbf{q} is eliminated.⁹

⁶It is also true that if \mathbf{q} is weakly (strictly) dominated by a mixed strategy, then $\tilde{\mathbf{q}}$ is weakly (strictly) dominated by a mixed strategy too, though not necessarily by the same mixed strategy.

⁷A way to see elimination of mixed strategies is as follows: Assume that individuals may play mixed strategies. If a mixed strategy is eliminated in the sense of definition 8.3.1, then it follows that the proportion of the population playing this mixed strategy goes to zero.

⁸What we mean is that $\mathbf{x}(t)$ evolves according to an aggregate monotonic dynamics; unless explicitly mentioned otherwise, we never make any assumption on $\mathbf{y}(t)$.

⁹Note that the fact that \mathbf{q} is eliminated is much weaker than the fact that $\prod_{i \in I} x_i^{q_i - p_i}(t) \rightarrow 0$; in particular, if $\mathbf{x}(t)$ converges to the boundary of S_N , which is common and the usual starting point to establish counter-intuitive properties of evolutionary dynamics, then any completely mixed strategy is eliminated in the sense of definition 8.3.1.

Proof. We begin with a basic result:

Lemma 8.3.4. *Let $w : \mathbb{R} \rightarrow \mathbb{R}$ be differentiable. If $\limsup_{t \rightarrow +\infty} \dot{w}(t) < 0$ or if*

$$\exists T \in \mathbb{R}, \forall t \geq T, \dot{w}(t) \leq 0,$$

\dot{w} is uniformly continuous and $\liminf_{t \rightarrow +\infty} \dot{w}(t) < 0$, then $w(t) \rightarrow -\infty$ as $t \rightarrow +\infty$

Proof. Both sets of assumptions imply $\int_0^t \dot{w}(\tau) d\tau \rightarrow -\infty$ as $t \rightarrow +\infty$. \square

We now prove proposition 8.3.3. Let $w(t) := \sum_{i \in I} (q_i - p_i) \ln x_i(t)$. We have:

$$\dot{w} = \sum_{i \in I} (q_i - p_i) \frac{\dot{x}_i}{x_i} = \sum_{i \in I} (q_i - p_i) g_i(\mathbf{x}, \mathbf{y}) = g_{\mathbf{q}}(\mathbf{x}, \mathbf{y}) - g_{\mathbf{p}}(\mathbf{x}, \mathbf{y}) \quad (8.3.6)$$

We show below that the assumptions of proposition 8.3.3 imply that the assumptions of lemma 8.3.4 are satisfied; it follows that $w(t) \rightarrow -\infty$ hence that $\prod_{i \in I} x_i^{q_i - p_i} \rightarrow 0$ as $t \rightarrow +\infty$.

First, it follows from (8.2.10) and (8.3.6) that

$$(8.3.4) \Rightarrow (\forall t \geq T, \dot{w} \leq 0)$$

Second, let $\epsilon > 0$ and $K_\epsilon := \{\mathbf{y} : U(\mathbf{q}, \mathbf{y}) - U(\mathbf{p}, \mathbf{y}) \leq -\epsilon\}$; it follows from (8.2.10), compactness of $K_\epsilon \times S_N$ and continuity of $g_{\mathbf{q}} - g_{\mathbf{p}}$ that if $K_\epsilon \neq \emptyset$ then

$$-\alpha := \max_{\mathbf{y} \in K_\epsilon, \mathbf{x} \in S_N} (g_{\mathbf{q}}(\mathbf{x}, \mathbf{y}) - g_{\mathbf{p}}(\mathbf{x}, \mathbf{y})) < 0$$

Together with (8.3.6), this implies that

$$\limsup [U(\mathbf{q}, \mathbf{y}(t)) - U(\mathbf{p}, \mathbf{y}(t))] < 0 \Rightarrow \limsup \dot{w} < 0$$

and similarly $\liminf [U(\mathbf{q}, \mathbf{y}(t)) - U(\mathbf{p}, \mathbf{y}(t))] < 0 \Rightarrow \liminf \dot{w} < 0$.

Third, the functions g_i are continuous hence uniformly continuous on the compact set $S_N \times S^{opp}$; therefore, if $\mathbf{y}(\cdot)$ is uniformly continuous then so is $\dot{w}(\cdot)$. This completes the proof. \square

Remark. In some contexts, the assumption that $\mathbf{y}(\cdot)$ is uniformly continuous might seem unnatural. This assumption and (8.3.5) might be replaced

by any assumption which ensures that the selection pressure against \mathbf{q} is not asymptotically released; that is,

$$\exists \epsilon > 0, \int_0^{+\infty} \mathbb{1}_{\mathbf{p} > \mathbf{q} + \epsilon}(t) = +\infty$$

where $\mathbb{1}_{\mathbf{p} > \mathbf{q} + \epsilon}(t) = 1$ if $U(\mathbf{p}, \mathbf{y}(t)) > U(\mathbf{q}, \mathbf{y}(t)) + \epsilon$ and 0 otherwise.

Proposition 8.3.3 implies two better known properties, due to Samuelson and Zhang (1992):

Corollary 8.3.5 (Strict domination). *Let \mathbf{p} and \mathbf{q} be two mixed strategies. If \mathbf{p} strictly dominates \mathbf{q} then, under any aggregate monotonic dynamics, \mathbf{q} is eliminated.*

Proof. Since S^{opp} is compact, it follows from (8.3.1) that there exists $\epsilon > 0$ such that

$$\forall \mathbf{y} \in S^{opp}, U(\mathbf{p}, \mathbf{y}) \geq U(\mathbf{q}, \mathbf{y}) + \epsilon$$

Therefore (8.3.3) holds. The result follows. \square

Corollary 8.3.6 (Weak domination). *Let \mathbf{p} and \mathbf{q} be two mixed strategies. Assume that $\mathbf{y}(t)$ is uniformly continuous and that \mathbf{p} weakly dominates \mathbf{q} . For any aggregate monotonic dynamics, we have:*

(a) *If \mathbf{q} is not eliminated then $U(\mathbf{q}, \mathbf{y}(t)) - U(\mathbf{p}, \mathbf{y}(t)) \rightarrow 0$.*

(b) *Assume that $S^{opp} = S_M$ (the $M - 1$ dimensional simplex) and that $U(i, \mathbf{y}) = \sum_{1 \leq j \leq M} y_j(t) U(i, j)$. If \mathbf{q} is not eliminated, then*

$$U(\mathbf{p}, j) > U(\mathbf{q}, j) \Rightarrow y_j(t) \rightarrow 0 \quad (8.3.7)$$

Proof. Since $U(\mathbf{q}, \mathbf{y}(t)) - U(\mathbf{p}, \mathbf{y}(t)) \leq 0$ for all t , it follows that if $U(\mathbf{q}, \mathbf{y}(t)) - U(\mathbf{p}, \mathbf{y}(t))$ does not converge to zero, then $\liminf (U(\mathbf{q}, \mathbf{y}(t)) - U(\mathbf{p}, \mathbf{y}(t))) < 0$. This being seen, (a) follows from proposition (8.3.3), and (b) follows from (a). \square

Remark. For n -player finite games, letting M_k be the number of pure strategies of player k and $M = \prod_{2 \leq k \leq n} M_k$, we may see S^{opp} as a subset of S_M . Corollary 8.3.6 then yields that if strategy \mathbf{q} is weakly dominated by the mixed strategy \mathbf{p} but not eliminated, then every *strategy profile* of players 2

to n to which \mathbf{p} is a strictly better response than \mathbf{q} is eliminated.

Proposition 8.3.3, corollary 8.3.5 and corollary 8.3.6 extend trivially (with suitable modifications) to the different classes of games introduced in definition 8.2.1:

Proposition 8.3.7.

- (a) *If \mathbf{q} is a pure strategy, proposition 8.3.3, corollary 8.3.5 and corollary 8.3.6 hold for any convex monotonic dynamics.*
- (b) *If \mathbf{p} is a pure strategy, the same results hold for any concave monotonic dynamics.*
- (c) *If both \mathbf{p} and \mathbf{q} are pure, then the same results hold for any monotonic dynamics.*

(Point (a) is due to Hofbauer and Weibull (1996) and point (c) to Samuelson and Zhang (1992); point (b) appears to be new.)

The proofs of (the suitably modified versions of) proposition 8.3.3, corollary 8.3.5 and corollary 8.3.6 for these other classes of dynamics is word by word the same as the proofs for aggregate monotonic dynamics given above, up to replacement of “aggregate monotonic” by “convex monotonic” (resp. “concave monotonic”, “monotonic”) and of equation (8.2.10) by equation (8.2.6) (resp. eq. (8.2.8), eq. (8.2.5)).

Best-response dynamics. Consider now the best-response dynamics (Gilboa and Matsui, 1991; Matsui 1992) given by:

$$\dot{\mathbf{x}}(t) \in BR(\mathbf{y}(t)) - \mathbf{x}(t) \quad (8.3.8)$$

for almost all t ; here, $BR(\mathbf{y})$ is the set of mixed best-responses to \mathbf{y} :

$$BR(\mathbf{y}) = \left\{ \mathbf{p} \in S_N, U(\mathbf{p}, \mathbf{y}) = \max_{\mathbf{q} \in S_N} U(\mathbf{q}, \mathbf{y}) \right\}$$

It follows from (8.3.8) that if from some time T on, the pure strategy i is never a best-response to $\mathbf{y}(t)$, then $x_i(t)$ decreases exponentially to zero. In particular, if the pure strategy i is strictly dominated or if there exists a mixed strategy \mathbf{p} such that $U(\mathbf{p}, \mathbf{y}(t)) > U(i, \mathbf{y}(t))$ for all $t \geq T$, then $x_i(t) \rightarrow 0$.¹⁰

¹⁰This is in particular the case if \mathbf{p} weakly dominates i , $S^{opp} = S_M$ for some integer M , and $\mathbf{y}(t) \in \text{int } S_M$ for $t \geq T$.

For these results, we do not need any assumption on $\mathbf{y}(\cdot)$, in particular, $\mathbf{y}(\cdot)$ need not be uniformly continuous.

Iterative elimination. We conclude this section with a standard result on elimination of iteratively strictly dominated strategies. Consider the special case of a two-player finite game, that is, $S^{opp} = S_M$ for some positive integer M and $U(i, \mathbf{y}) = \sum_{j \in J} y_j U(i, j)$ where $J = \{1, \dots, M\}$ denotes the set of pure strategies of player 2. Let $I^0 = I$, $J^0 = J$, and inductively let I^{k+1} (resp. J^{k+1}) denote the set of pure strategies $i \in I^k$ (resp. $j \in J^k$) that are not strictly dominated by a mixed strategy in the game restricted to $I^k \times J^k$. Similarly, let $\tilde{I}^0 = I$, $\tilde{J}^0 = J$, and let \tilde{I}^{k+1} (resp. \tilde{J}^{k+1}) denote the set of pure strategies $i \in \tilde{I}^k$ (resp. $j \in \tilde{J}^k$) that are not strictly dominated by a *pure* strategy in the game restricted to $\tilde{I}^k \times \tilde{J}^k$.

Definition 8.3.8. *A pure strategy $i \in I$ is iteratively strictly dominated (resp. iteratively strictly dominated by pure strategies) if there exists k in \mathbb{N} such that $i \notin I^k$ (resp. $i \notin \tilde{I}^k$).*

Proposition 8.3.9. *Assume that both $\mathbf{x}(t)$ and $\mathbf{y}(t)$ follow a convex monotonic (resp. monotonic) dynamics¹¹. Then for every interior initial condition, every pure strategy iteratively strictly dominated (resp. iteratively strictly dominated by pure strategies) is eliminated.*

Proof. We first prove the result on convex monotonic dynamics. It follows from the analogue of corollary 8.3.5 for convex monotonic dynamics that the pure strategies in $I^0 \setminus I^1$ and in $J^0 \setminus J^1$ are eliminated. By induction assume that for k in \mathbb{N} , the strategies in $I^0 \setminus I^k$ and in $J^0 \setminus J^k$ are eliminated, hence in particular

$$\forall \eta > 0, \exists T \in \mathbb{R}, \forall t \geq T, \max_{j \in J \setminus J^k} y_j(t) \leq \eta \quad (8.3.9)$$

Let $i \in I^k \setminus I^{k+1}$. Since strategy i is strictly dominated in the game restricted to $I^k \times J^k$, there exists $\epsilon > 0$, $\eta > 0$ and $\mathbf{p} \in S_N$ such that

$$\max_{j \in J \setminus J^k} y_j \leq \eta \Rightarrow U(i, \mathbf{y}(t)) < U(\mathbf{p}, \mathbf{y}(t)) + \epsilon$$

Therefore, it follows from (8.3.9) and the analogue of proposition (8.3.3) for convex monotonic dynamics that strategy i is eliminated. The result

¹¹Not necessarily the same.

follows. The proof of the result on monotonic dynamics is the same, up to replacement of “convex monotonic” by “monotonic”, “strictly dominated” by “strictly dominated by a pure strategy”, and “ $\mathbf{p} \in S_N$ ” by “ $i' \in I$ ”. \square

(The result on monotonic dynamics is due to Samuelson and Zhang (1992), the result on convex monotonic dynamics to Hofbauer and Weibull (1996).)

The extension to n -player games and to the best-response dynamics are immediate. Note that there are no results on iterative elimination of dominated *mixed* strategies, because when a mixed strategy is eliminated, this does not give rise to a well defined reduced game. We now turn to negative results.

8.4 Negative results

Consider a dynamics (8.2.1) in which the unnormalized growth rate of a pure strategy depends only on its payoff:

$$\dot{x}_i = x_i \left[g(U(i, \mathbf{y})) - \sum_{k \in I} x_k g(U(k, \mathbf{y})) \right] \quad (8.4.1)$$

for some continuous function g independent of $i \in I$. If g is increasing, then the dynamics is monotonic.

Proposition 8.4.1. (a). *If g is not convex, then (8.4.1) need not eliminate pure strategies dominated by mixed strategies.*¹²

(b). *If g is not concave, then (8.4.1) need not eliminate mixed strategies strictly dominated by pure strategies.*

(Point (a) is due to Hofbauer and Weibull (1996); point (b) appears to be new.)

Proof. The proofs of points (a) and (b) are very similar. We give a detailed proof of (b) below and sketch a proof of (a) in the appendix.

¹²More precisely, for any continuous and non-convex function g , there exist a 3×2 game and an opponent's strategy function $\mathbf{y}(\cdot) : \mathbb{R} \rightarrow S_2$ for which (8.4.1) does not eliminate a pure strategy strictly dominated by a mixed strategy. The next statement is to be understood in the same way. See the proofs.

If g is not concave, then there exist reals a and b such that

$$\frac{g(a) + g(b)}{2} > g\left(\frac{a+b}{2}\right)$$

By continuity of g , there exists $\epsilon > 0$ such that:

$$\alpha := \frac{g(a) + g(b)}{2} - g\left(\frac{a+b}{2} + \epsilon\right) > 0 \quad (8.4.2)$$

Assume that player 1 plays a 3×2 game with payoff matrix:

$$\begin{array}{c} \\ \\ \\ \end{array} \begin{array}{cc} L & R \\ \left(\begin{array}{cc} a & b \\ \frac{a+b}{2} + \epsilon & \frac{a+b}{2} + \epsilon \\ b & a \end{array} \right) \end{array}$$

Let $y_l(t)$ denote the probability at time t that the opponent chooses the left column. Fix some large positive constant T (more precisely, $T > (2C + 1)/\alpha + 1$ with $C = \max_{[a,b]} |g|$; this condition will appear later). Assume that the function $y_l(\cdot)$ is $2T$ -periodic with $y_l(t) = 1$ if $t \in [0, T - 1]$, $y_l(t) = 0$ if $t \in [T, 2T - 1]$, and linear variation on $[T - 1, T]$ and $[2T - 1, 2T]$.

Assume that $\mathbf{x}(0) \in \text{int}S_N$ and let

$$w(t) := \ln\left(\frac{x_2}{\sqrt{x_1 x_3}}\right)(t)$$

Letting $g_i(t) := g(U(i, \mathbf{y}(t)))$, we have:

$$\dot{w}(t) = g_2(t) - \frac{g_1(t) + g_3(t)}{2}$$

hence

$$w((k+1)T) - w(kT) = \int_{kT}^{(k+1)T} \left(g_2(t) - \frac{g_1(t) + g_3(t)}{2} \right) dt$$

Since for all t , $g_2(t) = g((a+b)/2 + \epsilon)$ and since

$$\forall t \in [kT, (k+1)T - 1], g_1(t) + g_3(t) = g(a) + g(b)$$

it follows from (8.4.2) that:

$$w((k+1)T) - w(kT) \leq -(T-1)\alpha + 2C$$

with $C = \max_{[a,b]} |g|$. Since we assumed $T > (2C + 1)/\alpha + 1$, it follows that

$$w((k + 1)T) \leq w(kT) - 1$$

Therefore $w(kT) \rightarrow -\infty$ as $k \rightarrow +\infty$. Since the variation of w between kT and $(k + 1)T$ is bounded (less than $2CT$), it follows that $w(t) \rightarrow -\infty$ hence $x_2(t) \rightarrow 0$ as $t \rightarrow +\infty$. Furthermore, it is easy to see by the same kind of computation that $x_1(t)/x_3(t)$ is $2T$ -periodic. It follows that $x_1 x_3 \not\rightarrow 0$. Actually, as is easily seen,

$$\inf_{t \in \mathbb{R}_+} x_1(t)x_3(t) = \min_{t \in [0, 2T]} x_1(t)x_3(t) > 0$$

Therefore, though $\mathbf{q} = (1/2, 0, 1/2)$ is strictly dominated by strategy 2, the mixed strategy \mathbf{q} is not eliminated.¹³ \square

The function $\mathbf{y}(\cdot)$ used in the proofs quite ad-hoc, and thus one may wonder whether we could obtain similar counter-examples with more realistic opponent's behaviours, e.g., with $\mathbf{y}(t) = \mathbf{x}(t)$. For point (a), the answer is positive: Hofbauer and Weibull (1996) have shown that for every C^1 function g which is increasing but not convex, there exists a 4×4 symmetric game for which strategy 4 is strictly dominated by a mixed strategy but survives along an open set of solutions of the single-population dynamics

$$\dot{x}_i = x_i \left[g(U(i, \mathbf{x})) - \sum_{k \in I} g(U(k, \mathbf{x})) \right] \quad (8.4.3)$$

For point (b), we strongly conjecture that the answer is positive but our proof is still incomplete.¹⁴

We conclude this section with two remarks. First, in chapter 10, we implicitly give an example of a 4×4 symmetric game for which the mixed strategy $(1/3, 1/3, 1/3, 0)$ is strictly dominated by the fourth pure strategy but survives along an open set of solutions of the single-population best-response dynamics. Thus, the best-response dynamics need not eliminate strictly dominated *mixed* strategies.

¹³Note also that for T sufficiently large or $x_1(0)$ sufficiently close to $x_3(0)$, there exists t in $[0, 2T]$ such that $x_1(t) = x_3(t)$, hence $x_1(t + 2kT) = x_3(t + 2kT)$ for all $k \in \mathbb{N}$. Together with $x_2(t) \rightarrow 0$, this implies that $\limsup_{t \rightarrow +\infty} x_1(t)x_3(t) = 1/4$.

¹⁴In the spirit of Hofbauer and Weibull (1996) and of chapter 10, it is easy to show that if g is increasing but nonconcave, then there exists a 4×4 symmetric game for which strategy 4 strictly dominates $\mathbf{q} = (1/3, 1/3, 1/3, 0)$ but $x_4(t) \prod_{i \in I} x_i^{-q_i}(t) \rightarrow 0$; but it might still be that $\prod_{i \in I} x_i^{q_i}(t) \rightarrow 0$ (as in game (10.3.1) in chapter 10).

Second, consider the Brown-von Neumann-Nash dynamics (Brown and von Neumann, 1950):

$$\dot{x}_i = k_i(\mathbf{x}, \mathbf{y}) - x_i \sum_{j \in I} k_j(\mathbf{x}, \mathbf{y}) \quad (8.4.4)$$

where $k_i(\mathbf{x}, \mathbf{y}) := \max \left[0, U(i, \mathbf{y}) - \sum_j x_j U(j, \mathbf{y}) \right]$ is the excess payoff of strategy i over the average payoff. This is not a dynamics of type (8.2.1) and it may be that a pure strategy has a higher payoff than another pure strategy yet a lower growth rate. Nevertheless, this dynamics is adaptive in that strategies earning less than average are selected against; that is, (8.4.4) is weakly sign-preserving in the sense of (Ritzberger and Weibull, 1995). This implies that (8.4.4) is an adjustment dynamics (Swinkels, 1993); i.e. $\mathbf{x}(t)$ evolves in the direction of strategies that are better replies to $\mathbf{y}(t)$ than $\mathbf{x}(t)$.¹⁵ Nevertheless, in the context of a game played within a single population ($\mathbf{y}(t) = \mathbf{x}(t)$), Berger and Hofbauer (2005) have shown that, under (8.4.4), a pure strategy strictly dominated by another *pure* strategy need not be eliminated. Generalizations of this result should appear in (Sandholm, 2006). Of course, this cannot happen with monotonic dynamics, as implied by corollary 8.3.5.

8.5 Discrete-time dynamics

Dekel and Scotchmer (1992) have shown that a discrete-time version of the replicator dynamics need not eliminate pure strategies strictly dominated by mixed strategies. This implies that corollary 8.3.5 does not extend straightforwardly to discrete-time dynamics. To understand the difficulty, consider the discrete-time dynamics:

$$x_i(n+1) = x_i(n) \frac{C + g_i(\mathbf{x}, \mathbf{y})}{C + \sum_k x_k g_k(\mathbf{x}, \mathbf{y})} \quad (8.5.1)$$

with $C > \max_{i, \mathbf{x}, \mathbf{y}} g_i(\mathbf{x}, \mathbf{y})$, $\mathbf{x} = \mathbf{x}(n)$ and $\mathbf{y} = \mathbf{y}(n)$. This is equivalent to

$$x_i(n+1) - x_i(n) = x_i(n) \frac{g_i(\mathbf{x}, \mathbf{y}) - \sum_k x_k g_k(\mathbf{x}, \mathbf{y})}{C + \sum_k x_k g_k(\mathbf{x}, \mathbf{y})} \quad (8.5.2)$$

¹⁵Formally, $U(\dot{\mathbf{x}}, \mathbf{y}) \geq 0$ with strict inequality whenever \mathbf{x} is not a best-response to \mathbf{y} .

Assume that the mixed strategy \mathbf{p} strictly dominates the mixed strategy \mathbf{q} and let $w(n) := \sum_{i \in I} (q_i - p_i) \ln x_i(n)$. A straightforward computation shows that

$$w(n+1) - w(n) = \sum_{i \in I} (q_i - p_i) \tilde{g}_i(\mathbf{x}, \mathbf{y}) \text{ with } \tilde{g}_i = \ln \left(1 + \frac{g_i}{C} \right) \quad (8.5.3)$$

This should be compared to (8.3.6), i.e.

$$\dot{w} = \sum_{i \in I} (q_i - p_i) g_i(\mathbf{x}, \mathbf{y})$$

Since the logarithm is concave, it is as if going from the continuous dynamics (8.2.1) to its discrete-counterpart (8.5.1) made the growth rates more concave. Thus, while the results on monotonic, concave monotonic, and “non-convex” dynamics carry over with no change to discrete-time dynamics (8.5.1), this is not true of the results on aggregate and convex monotonic dynamics. In particular, the analogue of point (b) in proposition 8.4.1 is:

Proposition 8.5.1. *If the function g is non-convex or linear, then the dynamics*

$$x_i(n+1) = x_i(n) \frac{C + g(U(i, \mathbf{y}))}{C + \sum_k x_k g(U(k, \mathbf{y}))}$$

need¹⁶ not eliminate pure strategies strictly dominated by mixed strategies.¹⁷

Proof. Similar to the proof of point (b) of proposition 8.4.1. □

Dekel and Scotchmer’s (1992) example might be seen as a particularly beautiful instance of this general result.

We now consider the limit $C \rightarrow +\infty$. If $C \gg \max_{(i, \mathbf{x}, \mathbf{y})} |g_i(\mathbf{x}, \mathbf{y})|$ then $\ln(1 + g_i/C)$ is approximately equal to g_i/C . From this remark we obtain:

Proposition 8.5.2. *Fix a game and functions g_i ; fix a mixed strategy \mathbf{q} strictly dominated by a mixed strategy \mathbf{p} . Assume that (8.2.10) holds, or that (8.2.6) holds and \mathbf{q} is pure. Then there exists \bar{C} in \mathbb{R} such that for all $C \geq \bar{C}$, the discrete dynamics (8.5.1) eliminates strategy \mathbf{q} .*

¹⁶As before, \mathbf{x} and \mathbf{y} are taken at time n in the right-hand side.

¹⁷Again, in the sense that for any continuous, non-convex or linear g , there exist a 3×2 game and an opponent’s strategy function $\mathbf{y}(\cdot) : \mathbb{N} \rightarrow S_2$ for which (8.4.1) does not eliminate a pure strategy strictly dominated by a mixed strategy.

Proof. There exists $\epsilon > 0$ such that

$$\forall(\mathbf{x}, \mathbf{y}), \sum_{i \in I} (p_i - q_i) U(i, \mathbf{y}) \geq \epsilon$$

Therefore, under the above assumptions, there exists $\alpha > 0$ such that

$$\forall(\mathbf{x}, \mathbf{y}), \sum_{i \in I} (p_i - q_i) g_i(\mathbf{x}, \mathbf{y}) \geq \alpha$$

For C large enough,

$$\forall(i, \mathbf{x}, \mathbf{y}), \left| \ln \left(1 + \frac{g_i(\mathbf{x}, \mathbf{y})}{C} \right) - \frac{g_i(\mathbf{x}, \mathbf{y})}{C} \right| < \frac{\alpha}{4C}$$

so that

$$\forall(\mathbf{x}, \mathbf{y}), \sum_{i \in I} (p_i - q_i) \tilde{g}_i(\mathbf{x}, \mathbf{y}) \geq \alpha/2C > 0$$

where $\tilde{g}_i = \ln \left(1 + \frac{g_i}{C} \right)$. The result follows. \square

Note that the constant \bar{C} depends not only on the game and the functions g_i , but also, through ϵ , on the strategies \mathbf{p} and \mathbf{q} . Compare proposition 8.5.2 and proposition 8.5.1. Note that for aggregate monotonic dynamics, the order of the quantifiers (whether we first fix the game, the functions g_i and the strategies \mathbf{p} and \mathbf{q} , or we first fix the constant C) is crucial.

To conclude, note that if the constant C depends on the step n :

$$x_i(n+1) = x_i(n) \frac{C_n + g_i(\mathbf{x}, \mathbf{y})}{C_n + \sum_k x_k g_k(\mathbf{x}, \mathbf{y})} \quad (8.5.4)$$

then we have:

Proposition 8.5.3. *Assume that $C_n \rightarrow +\infty$ as $n \rightarrow +\infty$ and that*

$$\sum_{n \in \mathbb{N}} \frac{1}{C_n} = +\infty, \quad (8.5.5)$$

Then for any game, any functions g_i and any mixed strategies \mathbf{p} and \mathbf{q} such that \mathbf{p} strictly dominates \mathbf{q} , if (8.2.10) holds, or if (8.2.6) holds and \mathbf{q} is pure, then the discrete dynamics (8.5.4) eliminates strategy \mathbf{q} .

Proof. This follows from proposition 8.5.2. The condition (8.5.5) is needed for the dynamics not to slow down too much and “stop”. \square

Many complements to this section may be found in (Cabrales and Sobel, 1992).

Appendix: Proof of proposition 8.4.1, point (a).

We proceed as in the proof of point (b). If g is continuous and not convex, then we can find reals a , b , and $\epsilon > 0$ such that

$$\alpha := g\left(\frac{a+b}{2} - \epsilon\right) - \frac{g(a) + g(b)}{2} > 0 \quad (8.5.6)$$

Assume that player 1 plays a 3×2 game with payoff matrix:

$$\begin{pmatrix} a & b \\ \frac{a+b}{2} - \epsilon & \frac{a+b}{2} - \epsilon \\ b & a \end{pmatrix}$$

Define $y_i(t)$ as in the proof of point (a), assume $\mathbf{x}(0) \in \text{int}S_N$ and let $w(t) := \ln(x_1(t)) - \ln(x_2(t))$. From $\dot{w}(t) = g(U(1, \mathbf{y}(t))) - g(U(2, \mathbf{y}(t)))$, the definition of $y_i(t)$ and (8.5.6) we obtain:

$$w((2k+2)T) - w(2kT) \leq -2(T-1)\alpha + 4C$$

with $C = \max_{[a,b]} |g|$. For T sufficiently large (which is without loss of generality), this implies $w((2k+2)T) \leq w(2kT) - 1$ hence $w(2kT) \rightarrow -\infty$. Since the variation of w between $2kT$ and $(2k+2)T$ is bounded (less than $4CT$), it follows that $w(t) \rightarrow -\infty$ hence $x_1(t) \rightarrow 0$ as $t \rightarrow +\infty$. Similarly, we obtain $x_3(t) \rightarrow 0$ so that $x_2(t) \rightarrow 1$ as $t \rightarrow +\infty$. Therefore, though strategy 2 is strictly dominated by the mixed strategy $(1/2, 0, 1/2)$, strategy 2 is the only strategy that survives.

Chapitre 9

Evolutionary Dynamics and Equilibria in Low Dimension: Elimination of All Strategies not Played in Equilibrium

In this chapter, we study evolutionary dynamics in 3×3 symmetric games. In the first part, we give a simple proof of the fact that, along all interior solutions, the *single population* replicator dynamics eliminates all strategies that do not belong to the support of any Nash equilibrium. We also show that the same result holds for the best-response dynamics.

In the second part, we show that, along all interior solutions, strategies that have zero (marginal) probability in all correlated equilibria are eliminated by the *two-population* best-response dynamics, and any *two-population* convex monotonic dynamics. The proof is based on the dual reduction technique, studied in chapter 2.

Part A: Nash Equilibrium

9.1 Introduction

We give a short proof of the fact that in every 3×3 symmetric game and along all interior solutions, the single-population replicator dynamics eliminates all strategies that are not used in Nash equilibrium. The proof is direct, in that it does not use Bomze's (1983) classification of 3×3 symmetric games. We also show that the same result holds for the best-response dynamics.

The material is organized as follows: we first introduce the notations and a lemma crucial to our proofs. The main results are proved in section 9.2 for the replicator dynamics and in section 9.3 for the best-response dynamics.

Notations. We consider finite, two-player symmetric games, played within a single population. Such a game is given by a set $I = \{1, \dots, N\}$ of pure strategies (the same for each player) and a payoff function $U : I \times I \rightarrow \mathbb{R}$. Thus, $U(i, j)$ is the payoff of a player playing strategy i and facing strategy j .

Let S_N denote the simplex of mixed strategies

$$S_N := \left\{ \mathbf{x} \in \mathbb{R}_+^I : \sum_{i \in I} x_i = 1 \right\}$$

and let $\text{int } S_N := \{\mathbf{x} \in S_N, \forall i \in I, x_i > 0\}$. Let \mathbf{e}_i denote the vertex of the simplex corresponding to the pure strategy i . The payoff of a player playing the mixed strategy \mathbf{x} and facing the mixed strategy \mathbf{y} will be denoted

$$U(\mathbf{x}, \mathbf{y}) := \sum_{i \in I, j \in I} x_i y_j U(i, j)$$

We will use the following lemma:

Lemma 9.1.1. *If a symmetric game does not have a completely mixed symmetric Nash equilibrium, then there exist mixed strategies \mathbf{p} and \mathbf{q} such that \mathbf{p} weakly dominates \mathbf{q} ; i.e. for all \mathbf{x} in S_N*

$$[x_i > 0 \text{ for all } i \text{ in } I] \Rightarrow [U(\mathbf{p}, \mathbf{x}) > U(\mathbf{q}, \mathbf{x})] \quad (9.1.1)$$

The above result is due to Akin (1980). Guidelines to the proof can also be found in (Hofbauer and Sigmund, 1998, exercice 7.6.3).

Two additional remarks will be useful: first, we may assume that \mathbf{p} and \mathbf{q} have disjoint support¹. Furthermore, it follows from (9.1.1) that $U(\mathbf{p}, k) \geq U(\mathbf{q}, k)$ for every k in $\{1, 2, 3\}$, with strict inequality for at least one k ; therefore, fixing k such that $U(\mathbf{p}, k) > U(\mathbf{q}, k)$, we have:

$$x_k > 0 \Rightarrow U(\mathbf{p}, \mathbf{x}) > U(\mathbf{q}, \mathbf{x}) \quad (9.1.2)$$

9.2 Replicator dynamics

Given some initial condition $\mathbf{x}(0)$ in S_N , the (single population) replicator dynamics is given by

$$\dot{x}_i(t) = x_i(t) [U(i, \mathbf{x}(t)) - U(\mathbf{x}(t), \mathbf{x}(t))] \quad (9.2.1)$$

For lightness of notations, we often write \dot{x}_i , x_i and \mathbf{x} for $\dot{x}_i(t)$, $x_i(t)$ and $\mathbf{x}(t)$.

Definition *An initial condition is interior if it belongs to $\text{int}S_N$.*

Definition *A pure strategy i in I is eliminated (along a solution of a given dynamics) if $x_i(t) \rightarrow 0$ as $t \rightarrow +\infty$.*

Proposition 9.2.1. *Under the replicator dynamics (9.2.1), in every 3×3 symmetric game and from every interior initial condition, every strategy that does not belong to the support of a symmetric Nash equilibrium is eliminated.*

We will actually show a slightly stronger result. Namely, if the game does not have a completely mixed symmetric Nash equilibrium, then every interior solution converges to a Nash equilibrium.

¹If $p_i q_i > 0$ for some i in $\{1, 2, 3\}$, then replace \mathbf{p} by $\tilde{\mathbf{p}} = (\mathbf{p} - \lambda \mathbf{e}_i)/(1 - \lambda)$ and \mathbf{q} by $\tilde{\mathbf{q}} = (\mathbf{q} - \lambda \mathbf{e}_i)/(1 - \lambda)$, with $\lambda = \min(p_i, q_i)$; note that since $\mathbf{p} \neq \mathbf{q}$, it follows that $\lambda < 1$. Clearly, (9.1.1) still holds for $\tilde{\mathbf{p}}$ and $\tilde{\mathbf{q}}$. If $\tilde{\mathbf{p}}$ and $\tilde{\mathbf{q}}$ do not have disjoint supports, iterate the process.

Proof. Consider a 3×3 symmetric game with no completely mixed, symmetric Nash equilibrium and fix an interior initial condition $\mathbf{x}(0)$. By lemma 9.1.1, there exist mixed strategies \mathbf{p} and \mathbf{q} with disjoint supports such that \mathbf{p} weakly dominates \mathbf{q} . By corollary 8.3.6 in chapter 8, this implies that if no pure strategy is eliminated then $\prod_{i:q_i>0} x_i(t) \rightarrow 0$. Since \mathbf{p} and \mathbf{q} have disjoint supports, it follows that there are at most two pure strategies, say i and j , in the support of \mathbf{q} . Therefore, in any case, there exist pure strategies i and j such that $x_i(t)x_j(t) \rightarrow 0$ as $t \rightarrow +\infty$.

Assume for concreteness that $x_1x_2 \rightarrow 0$. This means that the ω -limit of $\mathbf{x}(t)$, i.e.

$$\left\{ \mathbf{y} \in S_N, \exists(t_n) \in \mathbb{R}^{\mathbb{N}}, t_n \rightarrow +\infty, \mathbf{x}(t_n) \rightarrow \mathbf{y} \right\}$$

is included in the union of the edge $[\mathbf{e}_1, \mathbf{e}_3]$ (i.e. $\{\mathbf{x} : x_2 = 0\}$) and of the edge $[\mathbf{e}_2, \mathbf{e}_3]$.

If the solution $\mathbf{x}(t)$ converges to a point, then by the folk-theorem of evolutionary game theory (see, e.g., Weibull, 1995, proposition 4.11), this point is a Nash equilibrium and the statement of the proposition is satisfied. Otherwise, the ω -limit of $\mathbf{x}(t)$ contains at least two points \mathbf{x}' and \mathbf{x}'' . Since a bounded ω -limit set is always connected, we may assume that \mathbf{x}' and \mathbf{x}'' belong to the same edge, say $[\mathbf{e}_2, \mathbf{e}_3]$; furthermore, the whole segment $[\mathbf{x}', \mathbf{x}'']$ must belong to the ω -limit set of $\mathbf{x}(t)$.²

There are two cases: either (case 1) the whole edge consists of rest points or (case 2) there is at most one rest point in the relative interior of the edge $[\mathbf{e}_2, \mathbf{e}_3]$.³

In case 1, $U(2, i) = U(3, i)$ for $i = 2, 3$. Thus, if for instance $U(2, 1) \geq U(3, 1)$, then the payoff of 2 is always weakly greater than the payoff of 3. Since the function $w(t) = \ln(x_2(t)/x_3(t))$ satisfies

$$\dot{w} = U(2, \mathbf{x}) - U(3, \mathbf{x}) \tag{9.2.2}$$

it follows that the ratio x_2/x_3 is always weakly increasing. Since this ratio is

²Being connected, the ω -limit set of $\mathbf{x}(t)$ must include a connected path from \mathbf{x}' to \mathbf{x}'' ; therefore it must include the segment $[\mathbf{x}', \mathbf{x}'']$ if \mathbf{x}' and \mathbf{x}'' belong to the same edge, and the segments $[\mathbf{x}', \mathbf{e}_3]$ and $[\mathbf{e}_3, \mathbf{x}'']$ otherwise; in the latter case, replace \mathbf{x}' by \mathbf{e}_3 .

³The rest-points of the replicator dynamics within the edge $[\mathbf{e}_2, \mathbf{e}_3]$ are the vertices \mathbf{e}_2 , \mathbf{e}_3 and the completely mixed symmetric Nash equilibria of the 2×2 game obtained by eliminating strategy 1; in this 2×2 symmetric game, as is well known (see, e.g., Weibull, 1995), either every symmetric mixed strategy profile is a Nash equilibrium or there is at most one completely mixed symmetric Nash equilibrium.

different at \mathbf{x}' and at \mathbf{x}'' and since $\mathbf{x}(t)$ visits infinitely often arbitrarily small neighborhood of \mathbf{x}' and \mathbf{x}'' , this yields a contradiction.

In case 2, up to replacement of $[\mathbf{x}', \mathbf{x}'']$ by one of its subsegments, we may assume that there is no rest-point in the segment $[\mathbf{x}', \mathbf{x}'']$. Therefore, there exists a neighborhood V of this segment in which one of the strategies 2 and 3 earns strictly less than the other one. By (9.2.2), this implies that the ratio $\mathbf{x}_2/\mathbf{x}_3$ evolves monotonically in V . Note that this ratio is different at \mathbf{x}' and \mathbf{x}'' . Thus, without loss of generality, we may assume both that this ratio increases (weakly) in V and that it is strictly higher at \mathbf{x}' than at \mathbf{x}'' . This implies that there exist neighborhoods V' and V'' of, respectively, \mathbf{x}' and \mathbf{x}'' , such that $\mathbf{x}(t)$ cannot go from V' to V'' while staying in V . This yields a contradiction because, due to the structure of its ω -limit, $\mathbf{x}(t)$ must go infinitely often from V' to V'' while eventually staying in an arbitrarily small neighborhood of the union of the edge $[\mathbf{e}_1, \mathbf{e}_3]$ and of the edge $[\mathbf{e}_2, \mathbf{e}_3]$, hence must eventually go from V' to V'' while staying in V . \square

9.3 Best-response dynamics

The single-population best-response dynamics (Gilboa and Matsui, 1991) is given by the differential inclusion:

$$\dot{\mathbf{x}} \in BR(\mathbf{x}) - \mathbf{x} \quad (9.3.1)$$

where $BR(\mathbf{x})$ is the set of (mixed) best-responses to \mathbf{x} :

$$BR(\mathbf{x}) = \{\mathbf{p} \in S_N : U(\mathbf{p}, \mathbf{x}) = \max_{\mathbf{q} \in S_N} U(\mathbf{q}, \mathbf{x})\}$$

Proposition 9.3.1. *Consider a 3×3 symmetric game with no completely mixed symmetric Nash equilibrium. From every initial condition, every solution of the single-population best-response dynamics converges to the set of Nash equilibria.*

Proof. The proof is similar to the proof given for the replicator dynamics. A difference is that in general the best-response dynamics need not eliminate mixed strategies strictly dominated by pure strategies.⁴ However, consider a

⁴As shown in chapter 8, it holds if \mathbf{q} is a pure strategy; it also holds for 3×3 symmetric games, as will be shown in the proof of lemma 9.3.2 below; but it does not hold for 4×4 symmetric games, even if \mathbf{p} is a pure strategy: a counter-example is (implicitly) given in chapter 10.

3×3 symmetric game with no completely mixed symmetric Nash equilibrium and fix an initial condition $\mathbf{x}(0)$:

Lemma 9.3.2. *For every solution of the best-response dynamics, there exist i in $\{1, 2, 3\}$ and a time T such that $x_i(t) \rightarrow 0$ and, for all $t \geq T$, strategy i is not a strict best-response to $\mathbf{x}(t)$.*

We now prove proposition 9.3.1 (lemma 9.3.2 will be proved in the end): assume without loss of generality that $x_1(t) \rightarrow 0$ (so that $\mathbf{x}(t)$ converges to the edge $[\mathbf{e}_2, \mathbf{e}_3]$) and that for t sufficiently large, strategy 1 is not a strict best-response to $\mathbf{x}(t)$. If $\mathbf{x}(t)$ converges to a point, then this point is a Nash equilibrium and the statement of proposition 9.2.1 is satisfied. Otherwise, as in the proof of proposition 9.2.1, the ω -limit of $\mathbf{x}(t)$ contains two points \mathbf{x}' and \mathbf{x}'' in the edge $[\mathbf{e}_2, \mathbf{e}_3]$, and there are the same two cases: either (case 1) $U(2, i) = U(3, i)$ for $i = 2, 3$ or we may assume that (case 2) $U(2, \mathbf{x}) \neq U(3, \mathbf{x})$ for every \mathbf{x} in $[\mathbf{x}', \mathbf{x}'']$.

In case 1, if \mathbf{p} belongs to the ω -limit of $\mathbf{x}(t)$, then $p_1 = 0$ and $U(2, \mathbf{p}) = U(3, \mathbf{p}) \geq U(1, \mathbf{p})$; this implies that \mathbf{p} is a Nash equilibrium.

In case 2, there exists a neighborhood of the segment $[\mathbf{x}', \mathbf{x}'']$ in which one of the strategies 2 and 3 earns strictly less than the other one, hence this strategy is not a best-response. It follows that in the neighborhood of the segment $[\mathbf{x}', \mathbf{x}'']$, the ratio x_2/x_3 evolves monotonically (at least in a weak sense). As in the proof of proposition 9.2.1, this yields a contradiction. \square

Proof of lemma 9.3.2. Recall lemma 9.1.1. Fix mixed strategies \mathbf{p} and \mathbf{q} with disjoint support and a pure strategy k such that (9.1.2) is satisfied. If $x_k(t) = 0$ for all $t \geq 0$, then k is never a strict best-response to $\mathbf{x}(t)$ due to (9.3.1) and the result follows.

Otherwise, there exists T such that $x_k(T) > 0$. Due to (9.3.1), this implies that for every $t \geq T$, $x_k(t) \geq x_k(T) \exp(t - T) > 0$, hence $U(\mathbf{p}, \mathbf{x}(t)) > U(\mathbf{q}, \mathbf{x}(t))$.

Since \mathbf{p} and \mathbf{q} have disjoint support and since there are only three pure strategies, it follows that at least one of the strategies \mathbf{p} and \mathbf{q} is pure.

Case 1: If \mathbf{q} is a pure strategy, e.g., $\mathbf{q} = \mathbf{e}_1$. Then for all $t \geq T$, $U(\mathbf{p}, \mathbf{x}(t)) > U(1, \mathbf{x}(t))$, hence 1 is not a best-response to $\mathbf{x}(t)$. This implies that $x_1(t) \rightarrow 0$ and the result follows.

Case 2: If \mathbf{p} is a pure strategy, e.g., $\mathbf{p} = \mathbf{e}_1$. Then $q_1 = 0$ and

$$\forall t \geq T, \quad U(1, \mathbf{x}(t)) > q_2 U(2, \mathbf{x}(t)) + (1 - q_2) U(3, \mathbf{x}(t)) \quad (9.3.2)$$

so that

$$\forall t \geq T, \quad U(2, \mathbf{x}(t)) = U(3, \mathbf{x}(t)) \Rightarrow U(1, \mathbf{x}(t)) > U(3, \mathbf{x}(t)) \quad (9.3.3)$$

If, e.g., $U(2, \mathbf{x}(t)) > U(3, \mathbf{x}(t))$ for all $t \geq T$, then from time T on, strategy 3 is not a best-response to $\mathbf{x}(t)$, hence $x_3(t) \rightarrow 0$ and the result follows. Therefore we may assume that there exists $T' \geq T$ such that $U(2, \mathbf{x}(T')) = U(3, \mathbf{x}(T'))$. Now, by (9.3.3), at time T' as at any time $t \geq T$ such that $U(2, \mathbf{x}(t)) = U(3, \mathbf{x}(t))$, the unique best-response to $\mathbf{x}(t)$ is strategy 1, hence the solution locally points towards \mathbf{e}_1 .

Subcase 2.1: If $U(2, 1) = U(3, 1)$, then the equality $U(2, \mathbf{x}(t)) = U(3, \mathbf{x}(t))$ is maintained as $\mathbf{x}(t)$ points towards \mathbf{e}_1 . It follows that $\mathbf{x}(t)$ points towards \mathbf{e}_1 for all $t \geq T'$, hence converges to \mathbf{e}_1 and the result follows.

Subcase 2.2: If, e.g., $U(2, 1) > U(3, 1)$, then, as the solution points towards \mathbf{e}_1 , the difference $U(2, \mathbf{x}(t)) - U(3, \mathbf{x}(t))$ strictly increases. This implies that for $t \geq T'$, $U(2, \mathbf{x}(t)) - U(3, \mathbf{x}(t))$ is nonnegative, hence $U(3, \mathbf{x}(t)) < U(2, \mathbf{x}(t))$ or $U(3, \mathbf{x}(t)) = U(2, \mathbf{x}(t)) < U(1, \mathbf{x}(t))$. In any case, strategy 3 is not a best-response to $\mathbf{x}(t)$ and the result follows. This concludes the proof of lemma 9.3.2. ■

Part B: Correlated Equilibrium

9.4 Notations, definitions and main result

9.4.1 Notations

In this second part, we still focus on two-player finite symmetric games but we study two populations dynamics and so we need some new notations⁵: we let $\mathbf{y}(t) \in S_N$ denote the mean-strategy in the population of players 2. The payoff of a player playing the mixed strategy \mathbf{x} and facing the mixed strategy \mathbf{y} is:

$$U(\mathbf{x}, \mathbf{y}) := \sum_{i \in I, j \in I} x_i y_j U(i, j)$$

Given an initial condition $(\mathbf{x}(0), \mathbf{y}(0))$ in $S_N \times S_N$, the *two-population replicator dynamics* is given by:

$$\dot{x}_i = x_i [U(i, \mathbf{y}) - U(\mathbf{x}, \mathbf{y})] \quad \text{and} \quad \dot{y}_i = y_i [U(i, \mathbf{x}) - U(\mathbf{y}, \mathbf{x})] \quad (9.4.1)$$

(here \mathbf{x} (resp. \mathbf{y}) represents the mean strategy in the population of players 1 (resp. 2); the reason why the same payoff function U appears in both equations is that the game is symmetric)

Note that, for symmetric games and from a mathematical point of view, the single population replicator dynamics corresponds to the two-population replicator dynamics with symmetric initial conditions (that is, with $\mathbf{x}(0) = \mathbf{y}(0)$).

Definition *The pure strategy i of player 1 (resp. player 2) is eliminated by the two-population replicator dynamics (for some initial condition $(\mathbf{x}(0), \mathbf{y}(0))$)*

⁵Since we restrict attention to symmetric games, the pure strategy set and the payoff function are the same for both players; I (resp. U) denotes the pure strategy set (resp. payoff function) both of player 1 and of player 2.

if $x_i(t)$ (resp. $y_i(t)$) goes to zero as $t \rightarrow +\infty$.

Definition The pure strategy i in I (resp. the pure strategy profile (i, j) in $I \times I$) is used in correlated equilibrium⁶ if there exists a correlated equilibrium distribution μ such that $\sum_{j \in I} \mu(i, j) > 0$ (resp. $\mu(i, j) > 0$).

Remark 9.4.1. Due to the symmetry of the game, the existence of a correlated equilibrium μ such that $\sum_{j \in I} \mu(i, j) > 0$ is equivalent to the existence of a correlated equilibrium μ' such that $\sum_{j \in I} \mu'(j, i) > 0$.

Thus, when we say that some pure strategy i is used (or not used) in correlated equilibrium, it is unnecessary to specify whether we see this strategy as a strategy of player 1 or as a strategy of player 2. Furthermore, due to the symmetry of the game and to the convexity of the set of correlated equilibrium distributions, a pure strategy is used in correlated equilibrium if and only if it is used in some symmetric equilibrium (i.e. in a correlated equilibrium μ such that $\mu(k, l) = \mu(l, k)$ for every (k, l) in $I \times I$). Thus, we do not have to specify whether we are only interested in symmetric equilibria or not.

9.4.2 Main result

Definition An initial condition is interior if it belongs to $\text{int } S_N \times \text{int } S_N$

Proposition 9.4.2. Consider a 3×3 symmetric game. If the pure strategy i is not used in correlated equilibrium, then $x_i(t)$ and $y_i(t)$ both converge to 0 under the two-population replicator dynamics (9.4.1), for any interior initial condition $(\mathbf{x}(0), \mathbf{y}(0))$.

9.5 Proof

9.5.1 A property of correlated equilibria

From Nau and McCardle's (1990) characterization of strategy profiles never played in correlated equilibria and from the fact that dual vector of symmetric games may be symmetrized (proof of proposition 2.8.9), it follows that there

⁶For a definition of correlated equilibria see chapter 1.

exists a mapping $\alpha : I \rightarrow S_N$ with the following properties: For all (i, j) in $I \times I$,

$$[U(\alpha * i, j) - U(i, j)] + [U(\alpha * j, i) - U(j, i)] \geq 0 \quad (9.5.1)$$

where the mixed strategy $\alpha * i$ is the image of the pure strategy i by the mapping α . Furthermore⁷, (9.5.1) holds with strict inequality whenever (i, j) has probability zero in all correlated equilibria. In particular, if the pure strategy i has marginal probability zero in all correlated equilibria then, for all j in I ,

$$[U(\alpha * i, j) - U(i, j)] + [U(\alpha * j, i) - U(j, i)] > 0 \quad (9.5.2)$$

9.5.2 Properties of the replicator dynamics

The only properties of the replicator dynamics that will be used in the proof are the one given below. The fact that the replicator dynamics satisfies these properties was proved in chapter 8.

Let i, i' and \mathbf{p} denote respectively two pure strategies and a mixed strategy of player 1. Fix an interior initial condition $(\mathbf{x}(0), \mathbf{y}(0))$.

Property 9.5.1. *If there exist $\epsilon > 0$ and a time T in \mathbb{R} such that, for all $t \geq T$, $U(i, \mathbf{y}(t)) < U(i', \mathbf{y}(t)) - \epsilon$, then $\mathbf{x}_i(t) \xrightarrow[t \rightarrow +\infty]{} 0$.*

Property 9.5.2. *If \mathbf{p} weakly dominates i and if there exists a pure strategy j in $\{j \in I, U(\mathbf{p}, j) > U(i, j)\}$ such that $y_j(t)$ does not go to zero as time goes to infinity, then $\mathbf{x}_i(t) \xrightarrow[t \rightarrow +\infty]{} 0$.*

Property 9.5.3. *If a pure strategy is strictly dominated by a mixed strategy, then for every interior initial condition this pure strategy is eliminated by the two-population replicator dynamics.*

Of course, the symmetric properties (i.e. on elimination of strategies of player 2) hold as well.

⁷In chapter 2, a mapping from the set of pure strategy of player k to its set of mixed strategies is denoted α_k . Here, due to the symmetry of the game and of the dual vector we consider, we may skip the players' subscripts, and write α both for α_1 and for α_2 .

9.5.3 Proof of proposition 9.4.2

We are now ready to prove proposition 9.4.2. From now on, there are only three pure strategies: $S = \{1, 2, 3\}$, and strategy 3 is not used in correlated equilibrium. The aim is to show that strategy 3 is eliminated by the replicator dynamics. By symmetry, we only need to show that $x_3(t)$ converges to 0. We first exploit the inequations (9.5.1) and (9.5.2). These inequations are particularly interesting in two cases: first, taking $j = i$ in (9.5.1) yields

$$\forall i \in I, U(\alpha * i, i) \geq U(i, i) \quad (9.5.3)$$

If moreover strategy i is not used in correlated equilibrium then (9.5.2) yields:

$$U(\alpha * i, i) > U(i, i)$$

In particular,

$$U(\alpha * 3, 3) > U(3, 3) \quad (9.5.4)$$

Second, if j is α -invariant, i.e. if $\alpha * j = j$, then (9.5.1) yields

$$U(\alpha * i, j) \geq U(i, j) \quad (9.5.5)$$

If moreover strategy i is not used in correlated equilibrium then (9.5.2) yields:

$$U(\alpha * i, j) > U(i, j) \quad (9.5.6)$$

Now, distinguish the following cases:

Case 9.5.4. *If at least one of the strategies 1 and 2 is α -invariant*

Assume, for instance, that strategy 1 is α -invariant. Then, by (9.5.6),

$$U(\alpha * 3, 1) > U(3, 1) \quad (9.5.7)$$

and by (9.5.5)

$$U(\alpha * 2, 1) \geq U(2, 1) \quad (9.5.8)$$

Furthermore, taking $i = 3$ and $j = 2$ in (9.5.2) yields:

$$[U(\alpha * 3, 2) - U(3, 2)] + [U(\alpha * 2, 3) - U(2, 3)] > 0$$

Thus, at least one of the two brackets must be positive. If the first bracket is positive (subcase 1.1), i.e. if $U(\alpha * 3, 2) > U(3, 2)$, then (9.5.4) and (9.5.7)

imply that $\alpha * 3$ strictly dominates 3, hence $x_3(t)$ converges to 0 by property 9.5.3 and we are done.

Otherwise (subcase 1.2), the second bracket is positive, i.e. $U(\alpha * 2, 3) > U(2, 3)$. Together with (9.5.3) and (9.5.8) this implies that $\alpha * 2$ weakly dominates strategy 2, with strict inequality against strategy 3. Therefore, by property 9.5.2, if $x_3(t)$ does not converge to 0, then $y_2(t)$ does.

Now consider the 3×2 game obtained by eliminating the second strategy of player 2: it follows from 9.5.4) and (9.5.7) that, in this reduced game, the third strategy of player 1 is strictly dominated by $\alpha * 3$; since strategy 2 is weakly dominated, this implies that, in the reduced game, strategy 3 is strictly dominated by strategy 1. Therefore there exists $\epsilon > 0$ such that, once $y_2(t)$ is low enough, $U(3, \mathbf{y}) \leq U(1, \mathbf{y}) - \epsilon$. By property 9.5.1, this implies that $x_3(t)$ converges to 0.

Case 9.5.5. *If neither strategy 1 nor strategy 2 is α -invariant*

Consider the 2×2 game G_r obtained by elimination of the third strategy of both players. Since G_r is a 2×2 symmetric game, it may a priori be of three kinds:

Subcase 2.1 a coordination game, i.e. a game with two strict Nash equilibria and a completely mixed one;

Subcase 2.2 a game with a weakly or strictly dominated strategy;

Subcase 2.3 a trivial game, i.e. a game where the players have no influence on their own payoff.

Since α is an interior dual vector and since strategy 3 is not played in correlated equilibrium, it follows from proposition 2.7.12 that strategy 3 is transient under the Markov chain on I induced by α . Since moreover, strategy 1 is not α -invariant, this implies that the support of $\alpha * 1$ contains strategy 2 but not strategy 3; similarly, the support of $\alpha * 2$ contains strategy 1 but not strategy 3. This, in turn, implies two things:

First, the game G_r has no strict Nash equilibrium (indeed, if (i', j') is a strict Nash equilibrium of G_r then the inequality (9.5.1) for $i = i'$ and $j = j'$ cannot be satisfied). This rules out subcase 2.1.

Second, the Markov chain on I induced by α has a unique recurrent communicating set: $\{1, 2\}$. By the basic theory of dual reduction (Myerson, 1997), this implies that the game G may be reduced, in the sense of dual

reduction, into a game with a unique strategy profile, which corresponds to a mixed strategy profile of G with support $\{1, 2\} \times \{1, 2\}$. By corollary 2.4.6, this implies that G has a Nash equilibrium with support $\{1, 2\} \times \{1, 2\}$ and so, that G_r has a completely mixed Nash equilibrium. This rules out subcase 2.2. It follows that G_r is necessarily a trivial game.

Now, only two possibilities remain: first (subcase 2.3.1), it may be that $U(1, 3) = U(2, 3)$, so that $U(1, j) = U(2, j)$ for all j in I . Since for $i = 1, 2$, $\alpha * i$ has support in $\{1, 2\}$, this implies that $U(\alpha * i, j) = U(i, j)$ for every i in $\{1, 2\}$ and every j in I . Therefore, repeated applications of (9.5.2) show that strategy 3 is strictly dominated by $\alpha * 3$, which implies that $x_3(t)$ converges to 0.

Otherwise (subcase 2.3.2), $U(1, 3) \neq U(2, 3)$, so that we may assume for instance $U(1, 3) > U(2, 3)$. This implies that strategy 2 is weakly dominated by strategy 1, with strict inequality against strategy 3. Thus, if $x_3(t)$ does not converge to 0, then $y_2(t)$ does. But in the 3×2 game obtained by elimination of the second strategy of player 2, strategy 1 strictly dominates 3. Therefore, as in subcase 1.2, $x_3(t)$ converges to 0.

9.6 Extensions and comments

9.6.1 Other dynamics

The only properties of the replicator dynamics that are used in the proof of proposition 9.4.2 are properties 9.5.1, 9.5.2 and 9.5.3. It follows that proposition 9.4.2 extends to any dynamics satisfying these properties. As proved in chapter 8, this is the case in particular of the *best-response dynamics* of Gilboa and Matsui (1991) and of the *convex monotonic dynamics* of Hofbauer and Weibull (1996).

9.6.2 Asymmetric games

When considering multi-population dynamics, there is no compelling reason to focus on symmetric games. A more general result than proposition 9.4.2 would consist in proving that the two-population replicator dynamics eliminates all strategies that are not played in correlated equilibrium in every 3×3

game (and not only in symmetric ones).⁸ The author does not know whether this is true or not. However, it may be shown that in every 2×2 game, the two-population replicator dynamics eliminates all strategy *profiles* that are not used in Nash equilibrium in the following sense: if the pure strategy profile (i, j) has probability zero in all Nash equilibria, then $x_i(t)y_j(t) \rightarrow 0$ as $t \rightarrow +\infty$ for any interior initial condition $(\mathbf{x}(0), \mathbf{y}(0))$. This implies that in every 2×2 game, the two-population replicator dynamics eliminates all strategies that are not used in Nash equilibrium.

Here again, the proof relies solely on properties 9.5.1, 9.5.2 and 9.5.3, so that the same result holds for the two-population best-response dynamics and for any two-population convex monotonic dynamics.

9.6.3 Nash Equilibrium

In the first part of this chapter, we have shown that in every 3×3 symmetric games and from every interior initial condition, the *single-population* replicator dynamics and best-response dynamics eliminate strategies which are not used in Nash equilibrium. We do not know whether this extends to the two-population replicator dynamics (or to the two-population best-response dynamics).

9.6.4 Higher dimensional games

In both the first part and the second part of this chapter, the small dimensionality of the games was crucial to the proofs. In the first part, we studied dynamical systems evolving in a compact region of the plane. Due to, e.g., Poincaré-Bendixson theorem, there are severe constraints on the behavior of such dynamical systems;⁹ this is, so we think, the fundamental reason why proposition 9.2.1 holds.

In the second part, the dynamical systems we study are four-dimensional hence could, a priori, exhibit more complex behaviors. However, since we restrict attention to symmetric games, there are only a few pure strategy

⁸For two-player nonsymmetric games, the replicator dynamics is defined by taking \mathbf{x} in S_N , \mathbf{y} in S_M and by replacing U by U_1 (resp. U_2) in the first (resp. second) equation of (9.4.1).

⁹For this reason, we think that proposition 9.2.1 extends to all weakly payoff positive dynamics (for a definition of weakly payoff positive dynamics, see, e.g., Weibull, 1995, section 4.3.4).

payoffs in the game; together with Nau and McCardle's (1990) characterization of strategy profiles never used in correlated equilibrium, this allows to show that if a pure strategy is never played in correlated equilibrium, certain domination relation hold, which imply that strategies never played in correlated equilibrium must be eliminated. This would not be the case in, e.g., 4×4 symmetric games. Indeed, in the 4×4 symmetric games studied in the next chapter, a unique strategy is used in correlated equilibrium, but all pure strategies are undominated.

In this light, it should not come as a surprise that the results of this chapter do not extend to higher dimensions, as will be clear from the next chapters.

Chapitre 10

Elimination of All Strategies in the Support of Correlated Equilibria

This chapter is divided in two parts. In the first part, we show that there exists an open set of 4×4 symmetric games for which, under the replicator dynamics and from a large set of initial conditions, all strategies in the support of correlated equilibria are eliminated (hence only strategies that are NOT used in equilibrium remain). In the second part, we show that the same result holds for vast classes of dynamics, in particular, for the best-response dynamics (Gilboa and Matsui, 1991) and for every monotonic (Samuelson and Zhang, 1992) or weakly sign preserving (Ritzberger and Weibull, 1995) dynamics which depends continuously on the payoffs and for which pure strategies initially absent remain absent.

Part A: Replicator dynamics

The material is organized as follows: First, we introduce the notations and basic definitions, and recall some classical results on Rock-Paper-Scissors (RPS) games. In addition, we prove that these games have a unique correlated equilibrium. We then introduce a family of 4×4 symmetric games built by adding a strategy to a RPS game. We describe in details the orbits of the replicator dynamics in these games and show that, from an open set of initial conditions, all strategies used in correlated equilibrium are eliminated. We then show that elimination of all strategies used in correlated equilibrium is robust to perturbation of the vector field and occurs also for the discrete-time replicator dynamics. We conclude with a short discussion.

10.1 Notations and basic definitions

Consider a finite, two-player symmetric game played within a single population. Such a game is given by a set $I = \{1, \dots, N\}$ of pure strategies and a payoff matrix $\mathbf{U} = (u_{ij})_{1 \leq i, j \leq N}$. Here u_{ij} is the payoff of a player playing strategy i and facing strategy j . We use bold characters for vectors and matrices.

The proportion of the population playing strategy i at time t is denoted $x_i(t)$. Thus, the vector $\mathbf{x}(t) = (x_1(t), \dots, x_N(t))^T$ denotes the mean strategy at time t . It belongs to the $N - 1$ dimensional simplex over I

$$S_N := \left\{ \mathbf{x} \in \mathbb{R}_+^I : \sum_{i \in I} x_i = 1 \right\}$$

(henceforth, “the simplex”) whose vertices $\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_N$ correspond to the pure strategies of the game. We study the evolution of the mean strategy \mathbf{x} under the single-population replicator dynamics (Taylor and Jonker, 1978):

$$\dot{x}_i(t) = x_i(t) [(\mathbf{U}\mathbf{x}(t))_i - \mathbf{x}(t) \cdot \mathbf{U}\mathbf{x}(t)] \quad (10.1.1)$$

Remark: for notational simplicity, we usually write x_i and \mathbf{x} instead of $x_i(t)$ and $\mathbf{x}(t)$.

We now recall the definition of a correlated equilibrium: Consider a (non necessarily symmetric) bimatrix game with strategy set I (resp. J) for player 1 (resp. 2). Let $g_k(i, j)$ denote the payoff of player k when player 1 plays i and player 2 plays j . A correlated equilibrium (Aumann, 1974) is a probability distribution μ on the set $I \times J$ of pure strategy profiles (i.e. $\mu(i, j) \geq 0$ for all (i, j) in $I \times J$ and $\sum_{(i,j) \in I \times J} \mu(i, j) = 1$) which satisfies the following inequalities:

$$\sum_{j \in J} \mu(i, j) [g_1(i, j) - g_1(i', j)] \geq 0 \quad \forall i \in I, \forall i' \in I \quad (10.1.2)$$

and

$$\sum_{i \in I} \mu(i, j) [g_2(i, j) - g_2(i, j')] \geq 0 \quad \forall j \in J, \forall j' \in J \quad (10.1.3)$$

Though the above definition applies to general bimatrix games, from now on, we only consider *symmetric* bimatrix games. We use the same vocabulary as in the previous chapter:

Definition: the pure strategy i is *used in correlated equilibrium* if there exists a correlated equilibrium μ and a pure strategy j such that $\mu(i, j) > 0$.¹

Definition: the pure strategy i is *eliminated* (for some initial condition $\mathbf{x}(0)$) if $x_i(t)$ goes to zero as $t \rightarrow +\infty$.

10.2 A reminder on Rock-Paper-Scissors

A RPS (Rock-Paper-Scissors) game is a 3×3 symmetric game in which the second strategy (Paper) beats the first (Rock), the third (Scissors) beats the second, and the first beats the third. Up to normalization (i.e. putting zeros

¹As already noted in chapter 9, due to the symmetry of the game, if a strategy is used in a correlated equilibrium, it is also used in a symmetric correlated equilibrium. Thus it is immaterial whether we focus on symmetric correlated equilibria or not.

on the diagonal) the payoff matrix is of the form:

$$\begin{array}{c} \begin{array}{ccc} & 1 & 2 & 3 \\ \begin{array}{c} 1 \\ 2 \\ 3 \end{array} & \left(\begin{array}{ccc} 0 & -a_2 & b_3 \\ b_1 & 0 & -a_3 \\ -a_1 & b_2 & 0 \end{array} \right) & \text{with } a_i > 0, b_i > 0 \text{ for all } i = 1, 2, 3. \end{array} \end{array} \quad (10.2.1)$$

Any RPS game has a unique Nash equilibrium: (\mathbf{p}, \mathbf{p}) with

$$\mathbf{p} = \frac{1}{\Sigma} (a_2 a_3 + a_3 b_2 + b_2 b_3, a_1 a_3 + a_1 b_3 + b_3 b_1, a_1 a_2 + a_2 b_1 + b_1 b_2) \quad (10.2.2)$$

with $\Sigma > 0$ such that $\mathbf{p} \in S_4$ (see Zeeman, 1980; Gaunersdorfer and Hofbauer, 1995; or Hofbauer and Sigmund, 1998). Actually,

Notation: for $\mathbf{x} \in S_N$, $\mathbf{x} \otimes \mathbf{x}$ denotes the probability distribution on S_N induced by \mathbf{x} .

Proposition 10.2.1. *Any RPS game has a unique correlated equilibrium: $\mathbf{p} \otimes \mathbf{p}$.*

Proof. Let μ be a correlated equilibrium of (10.2.1). For $i = 1$ and, respectively, $i' = 2$ and $i' = 3$, the incentive constraint (10.1.2) reads:

$$\mu(1, 1)(-b_1) + \mu(1, 2)(-a_2) + \mu(1, 3)(a_3 + b_3) \geq 0 \quad (10.2.3)$$

$$\mu(1, 1)a_1 + \mu(1, 2)(-a_2 - b_2) + \mu(1, 3)b_3 \geq 0 \quad (10.2.4)$$

Add (10.2.3) multiplied by a_1 to (10.2.4) multiplied by b_1 . This gives

$$-\mu(1, 2)(a_1 a_2 + a_2 b_1 + b_1 b_2) + \mu(1, 3)(a_1 a_3 + a_1 b_3 + b_3 b_1) \geq 0$$

That is, recalling (10.2.2):

$$p_2 \mu(1, 3) \geq p_3 \mu(1, 2)$$

Every choice of a player and a strategy i yields a similar inequality. So we get six inequalities which together read:

$$p_2 \mu(1, 3) \geq p_3 \mu(1, 2) \geq p_1 \mu(3, 2) \geq p_2 \mu(3, 1) \geq p_3 \mu(2, 1) \geq p_1 \mu(2, 3) \geq p_2 \mu(1, 3)$$

Therefore all the above inequalities hold as equalities. Letting λ be such that the common value of the above expressions is $\lambda p_1 p_2 p_3$, we have: $\mu(i, j) = \lambda p_i p_j$ for every $j \neq i$. Together with (10.2.3) and (10.2.4), this implies that we also have $\mu(1, 1) = \lambda p_1^2$ (and by symmetry $\mu(i, i) = \lambda p_i^2$ for all i). Therefore $\lambda = 1$ and $\mu = \mathbf{p} \otimes \mathbf{p}$. \square

The behaviour of the replicator dynamics in RPS games has been totally analyzed by Zeeman (1980). In particular, letting $\partial S_3 := \{\mathbf{x} \in S_3 : x_1 x_2 x_3 = 0\}$ denote the boundary of the simplex:

Proposition 10.2.2 (Zeeman (1980)). *If $a_1 a_2 a_3 > b_1 b_2 b_3$, then for every initial condition $\mathbf{x}(0) \neq \mathbf{p}$, the solution $\mathbf{x}(t)$ converges to ∂S_3 as $t \rightarrow +\infty$.*

In the case of *cyclic symmetry* (i.e. $a_1 = a_2 = a_3$ and $b_1 = b_2 = b_3$) then the unique Nash equilibrium is $\mathbf{p} = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$. Furthermore, up to division of all payoffs by the common value of the a_i , the payoff matrix may be taken of the form:

$$\begin{pmatrix} 0 & -1 & \epsilon \\ \epsilon & 0 & -1 \\ -1 & \epsilon & 0 \end{pmatrix} \text{ with } \epsilon > 0 \quad (10.2.5)$$

The condition $a_1 a_2 a_3 > b_1 b_2 b_3$ then reduces to $\epsilon < 1$ and in this case proposition 10.2.2 may be proved as follows: for $\epsilon < 1$, the Nash equilibrium \mathbf{p} is globally inferior in the sense that:

$$\forall \mathbf{x} \in S_3, \mathbf{x} \neq \mathbf{p} \Rightarrow \mathbf{p} \cdot \mathbf{U}\mathbf{x} < \mathbf{x} \cdot \mathbf{U}\mathbf{x} \quad (10.2.6)$$

More precisely,

$$\mathbf{p} \cdot \mathbf{U}\mathbf{x} - \mathbf{x} \cdot \mathbf{U}\mathbf{x} = -(\mathbf{p} - \mathbf{x}) \cdot \mathbf{U}(\mathbf{p} - \mathbf{x}) = -\left(\frac{1-\epsilon}{2}\right) \sum_{1 \leq i \leq 3} (p_i - x_i)^2 \quad (10.2.7)$$

Now, let $\hat{V}(\mathbf{x}) := (x_1 x_2 x_3)^{1/3}$. Note that the function \hat{V} takes its minimal value 0 on ∂S_3 and its maximal value $1/3$ at \mathbf{p} . Letting $\hat{v}(t) := \hat{V}(\mathbf{x}(t))$ we get:

$$\dot{\hat{v}}(t) = (\mathbf{p} \cdot \mathbf{U}\mathbf{x} - \mathbf{x} \cdot \mathbf{U}\mathbf{x}) \hat{v}(t) = -\hat{v}(t) \left(\frac{1-\epsilon}{2}\right) \sum_{1 \leq i \leq 3} (p_i - x_i)^2 \quad (10.2.8)$$

where $\mathbf{x} = \mathbf{x}(t)$. The above expression is negative whenever $\hat{v}(t) \neq 0$ and $\mathbf{x} \neq \mathbf{p}$. It follows that for every initial condition $\mathbf{x}(0) \neq \mathbf{p}$, $\hat{v}(t)$ decreases to zero hence $\mathbf{x}(t)$ converges to the boundary.

10.3 A family of 4×4 games

Fix ϵ in $]0, 1[$, $\alpha \geq 0$, and consider the following 4×4 symmetric game which is built by adding a strategy to a RPS game:

$$\begin{array}{c}
 \begin{array}{c} 1 \\ 2 \\ 3 \\ 4 \end{array}
 \left(
 \begin{array}{cccc|c}
 & 1 & 2 & 3 & 4 & \\
 \hline
 1 & 0 & -1 & \epsilon & & -\alpha \\
 2 & \epsilon & 0 & -1 & & -\alpha \\
 3 & -1 & \epsilon & 0 & & -\alpha \\
 4 & \frac{-1+\epsilon}{3} + \alpha & \frac{-1+\epsilon}{3} + \alpha & \frac{-1+\epsilon}{3} + \alpha & & 0
 \end{array}
 \right)
 \end{array}
 \quad (10.3.1)$$

For $0 < \alpha < (1 - \epsilon)/3$, the interesting case, this game is very similar to the example used by Dekel and Scotchmer (1992) to show that a discrete version of the replicator dynamics need not eliminate all strictly dominated strategies.² We now describe the main features of the above game.

Let $\mathbf{n}_{123} = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3}, 0)$ denote the rest-point of the replicator dynamics corresponding to the Nash equilibrium of the underlying RPS game. Let \mathbf{U}_α denote the payoff matrix (10.3.1).

The case $\alpha = 0$. The strategies \mathbf{n}_{123} and \mathbf{e}_4 always earn the same payoff:

$$\mathbf{n}_{123} \cdot \mathbf{U}_0 \mathbf{x} = \mathbf{e}_4 \cdot \mathbf{U}_0 \mathbf{x} \quad \forall \mathbf{x} \in S_4 \quad (10.3.2)$$

Furthermore, against \mathbf{e}_4 , as against \mathbf{n}_{123} , all strategies earn the same payoff:

$$(\mathbf{x} - \mathbf{x}') \cdot \mathbf{U}_0 \mathbf{e}_4 = (\mathbf{x} - \mathbf{x}') \cdot \mathbf{U}_0 \mathbf{n}_{123} = 0 \quad \forall \mathbf{x} \in S_4, \forall \mathbf{x}' \in S_4 \quad (10.3.3)$$

The set of symmetric Nash equilibria is the segment $E_0 = [\mathbf{n}_{123}, \mathbf{e}_4]$.³ This shall be clear from the proof of proposition 10.3.2 below. A key property is that whenever the mean strategy \mathbf{x} does not belong to the segment of equilibria E_0 , every strategy in E_0 earns less than the mean payoff. Formally,

$$\forall \mathbf{x} \notin E_0, \forall \mathbf{p} \in E_0, \mathbf{p} \cdot \mathbf{U}_0 \mathbf{x} < \mathbf{x} \cdot \mathbf{U}_0 \mathbf{x}$$

²More precisely, the game obtained from (10.3.1) by multiplying all payoffs by -1 belongs to the family of games *à la* Dekel and Scotchmer (1992) considered by Hofbauer and Weibull (1996). In particular, figure 1 of (Hofbauer and Weibull, 1996, p.570) describes the dynamics on the boundary of the simplex in game (10.3.1), up to reversal of all arrows and permutation of strategies 2 and 3.

³That is, (\mathbf{x}, \mathbf{x}) is a Nash equilibrium if and only if $\mathbf{x} \in E_0$.

More precisely, for $\mathbf{x} \neq \mathbf{e}_4$, define \hat{x}_i as the share of the population that plays i relative to the share of the population that plays 1, 2 or 3. Formally,

$$\hat{x}_i = x_i / (x_1 + x_2 + x_3) \quad (10.3.4)$$

Lemma 10.3.1. *For every \mathbf{p} in E_0 and every $\mathbf{x} \neq \mathbf{e}_4$,*

$$\mathbf{p} \cdot \mathbf{U}_0 \mathbf{x} - \mathbf{x} \cdot \mathbf{U}_0 \mathbf{x} = -\frac{(1-\epsilon)}{2} (1-x_4)^2 \sum_{1 \leq i \leq 3} (\hat{x}_i - 1/3)^2 \quad (10.3.5)$$

Proof. Let $K = \mathbf{p} \cdot \mathbf{U}_0 \mathbf{x} - \mathbf{x} \cdot \mathbf{U}_0 \mathbf{x} = (\mathbf{p} - \mathbf{x}) \cdot \mathbf{U}_0 \mathbf{x}$. By (10.3.2), $\mathbf{p} \cdot \mathbf{U}_0 \mathbf{x} = \mathbf{n}_{123} \cdot \mathbf{U}_0 \mathbf{x}$ so that $K = (\mathbf{n}_{123} - \mathbf{x}) \cdot \mathbf{U}_0 \mathbf{x}$. Now let $\mathbf{y} = (\hat{x}_1, \hat{x}_2, \hat{x}_3, 0)$. Using (10.3.3) we get:

$$K = (\mathbf{n}_{123} - \mathbf{x}) \cdot \mathbf{U}_0 [(1-x_4)\mathbf{y} + x_4\mathbf{e}_4] = (1-x_4)(\mathbf{n}_{123} - \mathbf{x}) \cdot \mathbf{U}_0 \mathbf{y}$$

Noting that $\mathbf{n}_{123} - \mathbf{x} = (1-x_4)(\mathbf{n}_{123} - \mathbf{y}) + x_4(\mathbf{n}_{123} - \mathbf{e}_4)$ and using (10.3.2), we get: $K = (1-x_4)^2 (\mathbf{n}_{123} - \mathbf{y}) \cdot \mathbf{U}_0 \mathbf{y}$. Now apply (10.2.7). This gives (10.3.5) and concludes the proof. \square

The case $\alpha > 0$. The mixed strategy \mathbf{n}_{123} is no longer an equilibrium. Actually:

Proposition 10.3.2. *For every $\alpha > 0$, there exists a neighborhood of the game with payoffs (10.3.1), in which every game has a unique correlated equilibrium: $\mathbf{e}_4 \otimes \mathbf{e}_4$.*

Proof. We proved in chapter 3 that if a game has a unique correlated equilibrium, then every nearby game has a unique correlated equilibrium, and with the same support. Therefore, it suffices to show that $\mathbf{e}_4 \otimes \mathbf{e}_4$ is the unique correlated equilibrium of the game with payoffs (10.3.1).

Assume, by contradiction, that there exists a correlated equilibrium μ different from $\mathbf{e}_4 \otimes \mathbf{e}_4$. Since \mathbf{e}_4 is a strict Nash equilibrium, there exists $1 \leq i, j \leq 3$ such that $\mu(i, j) > 0$. Define the correlated distribution $\hat{\mu} \in \Delta(S_3)$ of the underlying RPS game (10.2.5) by:

$$\hat{\mu}(i, j) = \frac{\mu(i, j)}{K} \quad 1 \leq i, j \leq 3$$

with $K = \sum_{1 \leq i, j \leq 3} \mu(i, j) > 0$. For $1 \leq i, i' \leq 3$, we have $u_{i4} = u_{i'4} (= -\alpha)$, so that:

$$\sum_{j=1}^3 \hat{\mu}(i, j) [u_{ij} - u_{i'j}] = \sum_{j=1}^3 \frac{\mu(i, j)}{K} [u_{ij} - u_{i'j}] = \frac{1}{K} \sum_{j=1}^3 \mu(i, j) [u_{ij} - u_{i'j}] \geq 0$$

(The latter inequality holds because μ is a correlated equilibrium)

Together with symmetric inequalities, this implies that $\hat{\mu}$ is a correlated equilibrium of (10.2.5). By proposition 10.2.1, this implies that for every i in $\{1, 2, 3\}$, we have $\hat{\mu}(i, j) = 1/9$ hence $\mu(i, j) = K/9$. From this and the fact that strategy 4 is a best-response to itself, it follows that for any $1 \leq i, j \leq 3$

$$\sum_{1 \leq j \leq 4} \mu(i, j) [u_{ij} - u_{4j}] \leq \sum_{1 \leq j \leq 3} \mu(i, j) [u_{ij} - u_{4j}] = -\frac{K\alpha}{3} < 0$$

This contradicts the fact that μ is a correlated equilibrium. \square

Nevertheless, for $\alpha < (1 - \epsilon)/3$, the above game has a best-response cycle: $\mathbf{e}_1 \rightarrow \mathbf{e}_2 \rightarrow \mathbf{e}_3 \rightarrow \mathbf{e}_1$. We will show that for $\alpha > 0$ small enough, the corresponding set

$$\Gamma := \{\mathbf{x} \in S_4, x_4 = 0 \text{ and } x_1 x_2 x_3 = 0\} \quad (10.3.6)$$

attracts all nearby orbits. We first show that the (replicator) dynamics in the interior of S_4 may be decomposed in two parts: an increase or decrease in x_4 , and an outward spiralling movement around the segment $E_0 = [\mathbf{n}_{123}, \mathbf{e}_4]$.

10.4 Decomposition of the dynamics

First, note that for every \mathbf{x} in E_0 , we have: $(\mathbf{U}\mathbf{x})_1 = (\mathbf{U}\mathbf{x})_2 = (\mathbf{U}\mathbf{x})_3$. This implies that the segment E_0 is globally invariant. Second, recall the definition (10.3.4) of \hat{x}_i . For $\mathbf{x} \neq \mathbf{e}_4$, let $\hat{\mathbf{x}} = (\hat{x}_1, \hat{x}_2, \hat{x}_3)$. Let $\hat{\mathbf{U}}$ denote the payoff matrix (10.2.5) of the underlying RPS game.

Lemma 10.4.1. *Let $\mathbf{x}(\cdot)$ be a solution of the replicator dynamics (10.1.1) with $\mathbf{x}(0) \neq \mathbf{e}_4$. For every i in $\{1, 2, 3\}$,*

$$\dot{\hat{x}}_i = (1 - x_4) \hat{x}_i [(\hat{\mathbf{U}}\hat{\mathbf{x}})_i - \hat{\mathbf{x}} \cdot \hat{\mathbf{U}}\hat{\mathbf{x}}] \quad (10.4.1)$$

Proof. Let i in $\{1, 2, 3\}$. If $x_i = 0$, then (10.4.1) holds trivially. Otherwise, for every j in $\{1, 2, 3\}$ such that x_j is positive,

$$\frac{\dot{\hat{x}}_i}{\hat{x}_i} - \frac{\dot{\hat{x}}_j}{\hat{x}_j} = \frac{d}{dt} \ln \left(\frac{\hat{x}_i}{\hat{x}_j} \right) = \frac{d}{dt} \ln \left(\frac{x_i}{x_j} \right) = (\mathbf{U}\mathbf{x})_i - (\mathbf{U}\mathbf{x})_j = (1 - x_4) [(\hat{\mathbf{U}}\hat{\mathbf{x}})_i - (\hat{\mathbf{U}}\hat{\mathbf{x}})_j]$$

Multiplying the above equality by \hat{x}_j and summing over all j such that $x_j > 0$ yields (10.4.1). \square

The lemma means that, up to a change of velocity, $\hat{\mathbf{x}}$ follows the replicator dynamics for the game with payoff matrix $\hat{\mathbf{U}}$ (hence spirals towards the boundary).⁴ Now, recall equation (10.2.8) and the definition of \hat{V} . For $\mathbf{x} \neq \mathbf{e}_4$, let $V(\mathbf{x}) := \hat{V}(\hat{\mathbf{x}})$. That is,

$$V(\mathbf{x}) = (\hat{x}_1 \hat{x}_2 \hat{x}_3)^{1/3} = \frac{(x_1 x_2 x_3)^{1/3}}{x_1 + x_2 + x_3}$$

Corollary 10.4.2. *Let $\mathbf{x}(\cdot)$ be a solution of (10.1.1) with $\mathbf{x}(0) \neq \mathbf{e}_4$. The function $v(t) := V(\mathbf{x}(t))$ satisfies:*

$$\dot{v}(t) = v(t)f(\mathbf{x}(t)) \text{ with } f(\mathbf{x}) = -(1 - x_4) \left(\frac{1 - \epsilon}{2} \right) \sum_{1 \leq i \leq 3} (\hat{x}_i - 1/3)^2 \quad (10.4.2)$$

Proof. We have: $v(t) = V(\mathbf{x}(t)) = \hat{V}(\hat{\mathbf{x}}(t))$. Therefore $\dot{v} = \vec{\text{grad}} \hat{V} \cdot \dot{\hat{\mathbf{x}}}$, with $\vec{\text{grad}} \hat{V} = (\partial \hat{V} / \partial \hat{x}_i)_{1 \leq i \leq 3}$. Applying lemma 10.4.1 and equation (10.2.8) yields (10.4.2). \square

Note that $v(t)$ is nonnegative and that the function f is negative everywhere but on the interval $[\mathbf{n}_{123}, \mathbf{e}_4[$, where V attains its maximal value $1/3$. Therefore, it follows from (10.4.2) that V decreases along all interior trajectories, except the ones starting (hence remaining) in the interval $]\mathbf{n}_{123}, \mathbf{e}_4[$. We now exploit this fact to build a Lyapunov function⁵ for the set Γ defined in (10.3.6).

10.5 Main results

Let $W(\mathbf{x}) := \max(x_4, 3V(\mathbf{x}))$ for $\mathbf{x} \neq \mathbf{e}_4$ and $W(\mathbf{e}_4) = 1$, so that W is continuous on S_4 . Note that W takes its maximal value 1 on the segment $E_0 = [\mathbf{n}_{123}, \mathbf{e}_4]$ and its minimal value 0 on Γ .

For $\delta \geq 0$, let K_δ denote the compact set:

$$K_\delta := \{\mathbf{x} \in S_N, W(\mathbf{x}) \leq \delta\}$$

⁴The fact that when the $N - 1$ first strategies earn the same payoff against the N^{th} (and last) strategy, the dynamics may be decomposed as in lemma 10.4.1 was known to Josef Hofbauer (personal communication). This results from a combination of theorem 7.5.1 and of exercise 7.5.2 in (Hofbauer and Sigmund, 1998). The author rediscovered it independently.

⁵For an introduction to Lyapunov functions, see, e.g., Bhatia and Szegö, 1970.

Recall that \mathbf{U}_0 denotes the payoff matrix (10.3.1) with $\alpha = 0$. Finally, on the set of 4×4 payoff matrices \mathbf{U} , define the norm $\|\mathbf{U}\| = \max_{1 \leq i, j \leq 4} |u_{ij}|$.

Proposition 10.5.1. *Let $0 < \delta < 1$. There exists $\gamma > 0$ such that for every 4×4 symmetric game with payoff matrix \mathbf{U} such that $\|\mathbf{U} - \mathbf{U}_0\| \leq \gamma$ and every initial condition $\mathbf{x}(0)$ in K_δ ,*

$$W(\mathbf{x}(t)) \leq W(\mathbf{x}(0)) \exp(-\gamma t) \quad \forall t \geq 0$$

In particular, $x_4(t) \rightarrow 0$ as $t \rightarrow +\infty$.

Proof. Since $\delta < 1$, the set K_δ is disjoint from E_0 . Therefore, it follows from (10.3.5) that for every \mathbf{x} in K_δ , the quantity $(\mathbf{U}_0\mathbf{x})_4 - \mathbf{x} \cdot \mathbf{U}_0\mathbf{x}$ is negative. Similarly, it follows from the definition of the function f in (10.4.2) that for every \mathbf{x} in K_δ , $f(\mathbf{x})$ is negative. Therefore, by compactness of K_δ , there exists a positive constant γ such that

$$\max_{\mathbf{x} \in K_\delta} ((\mathbf{U}_0\mathbf{x})_4 - \mathbf{x} \cdot \mathbf{U}_0\mathbf{x}, f(\mathbf{x})) \leq -3\gamma < 0 \quad (10.5.1)$$

Fix a 4×4 payoff matrix \mathbf{U} such that $\|\mathbf{U} - \mathbf{U}_0\| \leq \gamma$ and let $\mathbf{x}(\cdot, \mathbf{U})$ denote a solution of the replicator dynamics for the game with payoff matrix \mathbf{U} , with initial condition different from \mathbf{e}_4 . Since $(\mathbf{U}\mathbf{x})_4 - \mathbf{x} \cdot \mathbf{U}\mathbf{x}$ is Lipschitz continuous in \mathbf{U} with Lipschitz constant 2 (as the difference of two 1-Lipschitz terms), it follows from (10.5.1) that

$$\forall \mathbf{x} \in K_\delta, (\mathbf{U}\mathbf{x})_4 - \mathbf{x} \cdot \mathbf{U}\mathbf{x} \leq -\gamma$$

Since $(\mathbf{U}\mathbf{x})_4 - \mathbf{x} \cdot \mathbf{U}\mathbf{x}$ is the growth rate of strategy 4, this implies that:

$$\mathbf{x}(t, \mathbf{U}) \in K_\delta \Rightarrow \frac{d}{dt} x_4(t, \mathbf{U}) \leq -\gamma x_4(t, \mathbf{U}) \quad (10.5.2)$$

Let $v_U(t) := V(\mathbf{x}(t, \mathbf{U}))$. We claim that the function v_U satisfies:

$$\dot{v}_U(t) = v_U(t) f_U(\mathbf{x}(t, \mathbf{U})) \text{ with } f_U(\mathbf{x}) = \frac{1}{3} \sum_{1 \leq i, j \leq 3} \hat{x}_j [(\mathbf{U}\mathbf{x})_i - (\mathbf{U}\mathbf{x})_j] \quad (10.5.3)$$

This will be proved in the end. The above expression for f_U is Lipschitz continuous in \mathbf{U} with Lipschitz constant 2. Since⁶ the function f in (10.5.1) is

⁶ f_U is actually Lipschitz in \mathbf{U} with constant $4/3$ but we do not need it.

equal to f_{U_0} (see (10.4.2)), it follows from (10.5.1) that $\max_{\mathbf{x} \in K_\delta} f_U(\mathbf{x}) \leq -\gamma$. Together with (10.5.3) this implies that

$$\mathbf{x}(t, \mathbf{U}) \in K_\delta \Rightarrow \dot{v}_U(t) \leq -\gamma v_U(t) \quad (10.5.4)$$

Let $w_U(t) = W(\mathbf{x}(t, \mathbf{U}))$. Equations (10.5.2) and (10.5.4) imply that:

$$\mathbf{x}(t, \mathbf{U}) \in K_\delta \Rightarrow \dot{w}_U(t) \leq -\gamma w_U(t) \quad (10.5.5)$$

Thus, if $\mathbf{x}(t, \mathbf{U})$ is in K_δ (i.e. $w_U(t) \leq \delta$) then w_U decreases weakly. This implies that K_δ is forward invariant. It follows that for every initial condition $\mathbf{x}(0)$ in K_δ and every time $t \geq 0$, $\dot{w}_U(t) \leq -\gamma w_U(t)$, hence $w_U(t) \leq w_U(0) \exp(-\gamma t)$.

It only remains to prove the claim: proceeding as in the proof of lemma 10.4.1, we get (with $\hat{x}_i = \hat{x}_i(t, \mathbf{U})$):

$$\dot{\hat{x}}_i = \hat{x}_i \sum_{1 \leq j \leq 3} \hat{x}_j [(\mathbf{U}\mathbf{x})_i - (\mathbf{U}\mathbf{x})_j] \quad (10.5.6)$$

Furthermore, by definition of v_U , we have $\dot{v}_U = 0$ when $v_U = 0$ and otherwise:

$$\dot{v}_U(t) = v_U(t) \left(\frac{1}{3} \sum_{1 \leq i \leq 3} \frac{\dot{\hat{x}}_i}{\hat{x}_i}(t) \right)$$

Replacing $\dot{\hat{x}}_i$ by the expression in (10.5.6) leads to (10.5.3). This completes the proof. \square

It follows from propositions 10.3.2 and 10.5.1 that if $\alpha > 0$ is small enough⁷, then there exists a neighborhood of the game with payoffs (10.3.1) such that: for any game in this neighborhood, the unique strategy used in correlated equilibrium is strategy 4, but $x_4(t) \rightarrow 0$ from an open set of initial conditions.

10.6 Extensions and variants

10.6.1 Robustness to perturbation of the vector field

We have shown that convergence to Γ does not only occur in the game with payoffs (10.3.1), but is robust to perturbation of the game. We now show

⁷As will be clear from proposition 10.8.1 in part 2, it suffices that α be smaller than $(1 - \epsilon)/3$.

that, more generally, convergence to Γ is robust to perturbation of the vector field.⁸ Let $0 < \delta < 1$. Fix γ as in proposition 10.5.1. Let $0 < \alpha \leq \gamma$ and let \mathbf{U}_α denote the payoff matrix (10.3.1).

Let $f(\mathbf{x}, t) = (f_i(\mathbf{x}, t))_{i \in I}$ be a (time dependent) locally Lipschitz vector field such that the perturbed dynamics

$$\dot{x}_i = x_i [(\mathbf{U}_\alpha \mathbf{x})_i - \mathbf{x} \cdot \mathbf{U}_\alpha \mathbf{x}] + f_i(\mathbf{x}, t) \quad (10.6.1)$$

leaves the simplex S_N forward invariant. Finally, let

$$g(t) = \max_{i \in I, \mathbf{x} \in K_\delta} |f_i(\mathbf{x}, t)| \text{ and } h(t) = \sup_{i \in I, \mathbf{x} \in K_\delta, x_i > 0} |f_i(\mathbf{x}, t)/x_i|$$

(g measures the intensity of the perturbation in absolute terms, h in terms of perturbation of the growth-rates; note that a priori h may take the value $+\infty$).

Corollary 10.6.1. *If*

$$\forall t \geq 0, g(t) < \gamma \delta^3 (1 - \delta) \text{ and } \lim_{t \rightarrow +\infty} g(t) = 0 \quad (10.6.2)$$

or if

$$\forall t \geq 0, h(t) < \gamma(1 - \delta) \quad (10.6.3)$$

then for every initial condition $\mathbf{x}(0)$ in K_δ , the solution $\mathbf{x}(t)$ of (10.6.1) converges to Γ .

Proof. Let \mathbf{x} belong to K_δ and to the interior of the simplex. Recall that $V(\mathbf{x}) = (x_1 x_2 x_3)^{\frac{1}{3}} / (1 - x_4)$. Therefore, letting $v(t) = V(\mathbf{x}(t))$, we have:

$$\frac{\dot{v}}{v} = \frac{1}{3} \sum_{1 \leq i \leq 3} \frac{\dot{x}_i}{x_i} + \frac{\dot{x}_4}{1 - x_4} = \frac{1}{3} \sum_{1 \leq i \leq 3} \frac{\dot{x}_i}{x_i} + \frac{\dot{x}_4}{x_4} \frac{x_4}{1 - x_4} \quad (10.6.4)$$

Since this expression is less than $-\gamma$ under the unperturbed dynamics, we get:

$$\frac{\dot{v}}{v} + \gamma \leq h(t) \left(1 + \frac{x_4}{1 - x_4} \right) \leq \frac{h(t)}{1 - x_4} \leq \frac{h(t)}{1 - \delta}$$

(for the last inequality, recall that $\mathbf{x} \in K_\delta$, so that $x_4 \leq \delta$). Furthermore, under the perturbed dynamics, $\dot{x}_4 \leq x_4(-\gamma + h(t))$. It follows that if (10.6.3)

⁸Perturbing the payoffs amounts to a certain type of perturbation of the vector field. In this sense the results of this section are more general than those of the previous section.

holds then both V and x_4 decrease (and actually decrease exponentially) along the trajectories starting in K_δ .

Similarly, from the first equality in (10.6.4), we get:

$$\begin{aligned} \frac{\dot{V}}{V} + \gamma &\leq \frac{g(t)}{3} \left(\sum_{1 \leq i \leq 3} \frac{1}{x_i} + \frac{3}{x_1 + x_2 + x_3} \right) \\ &\leq \frac{4g(t)}{9} \sum_{1 \leq i \leq 3} \frac{1}{x_i} = \frac{4}{3(1-x_4)} \sum_{1 \leq i \leq 3} \frac{1}{\hat{x}_i} \end{aligned}$$

(for the second inequality, use the convexity of $y \rightarrow 1/y$). But

$$\sum_{1 \leq i \leq 3} \frac{1}{\hat{x}_i} = \frac{1}{V^3} \sum_{1 \leq i < j \leq 3} \hat{x}_i \hat{x}_j = \frac{1}{V^3} \left[\frac{1}{3} - \frac{1}{6} \sum_{1 \leq i < j \leq 3} (\hat{x}_i - \hat{x}_j)^2 \right] \leq \frac{1}{3V^3}$$

so⁹ that finally for all \mathbf{x} in K_δ

$$\dot{V} \leq -\gamma V + \frac{4g(t)}{27(1-x_4)V^2} \leq -\gamma V + \frac{4g(t)}{27(1-\delta)V^2} \quad (10.6.5)$$

Furthermore, $\dot{x}_4 \leq -\gamma x_4 + g(t)$. Together with (10.6.5), this implies that if $g(t) \leq \gamma \delta^3(1-\delta)$, then both V and x_4 decrease when equal to δ (hence K_δ is forward invariant) and that if furthermore $g(t) \rightarrow 0$, then both V and x_4 converge to 0. \square

10.6.2 Discrete-time replicator dynamics

Consider the discrete-time version of the replicator dynamics (Maynard-Smith, 1982):

$$x'_i = x_i \frac{C + (\mathbf{U}\mathbf{x})_i}{C + \mathbf{x} \cdot \mathbf{U}\mathbf{x}} \quad (10.6.6)$$

where \mathbf{x} (resp. \mathbf{x}') is the value of the mean-strategy at time t (resp. $t+1$), and C is a large positive constant (larger than all the payoffs in the game). This is equivalent to

$$x'_i - x_i = x_i \frac{(\mathbf{U}\mathbf{x})_i - \mathbf{x} \cdot \mathbf{U}\mathbf{x}}{C + \mathbf{x} \cdot \mathbf{U}\mathbf{x}} \quad (10.6.7)$$

We will show that

⁹For the second equality, note that $1 = (\sum_i \hat{x}_i)^2 = \sum_i \hat{x}_i^2 + 2 \sum_{i < j} \hat{x}_i \hat{x}_j$ and that $(\sum_{i < j} (\hat{x}_i - \hat{x}_j)^2) / 2 = \sum_i \hat{x}_i^2 - \sum_{i < j} \hat{x}_i \hat{x}_j$ and then subtract the latter equality to the former.

Proposition 10.6.2. *If $0 < \alpha < (1 - \epsilon)/12$ then for the game with pay-offs (10.3.1) and from an open set of initial conditions, the mean-strategy converges to Γ . In particular, strategy 4 is eliminated.*

Proof. Recall that for the continuous-time version of the replicator dynamics, we showed that the function $W(\mathbf{x}) = \max(x_4, 3\hat{x}_1\hat{x}_2\hat{x}_3)$ decreases exponentially in the neighborhood of Γ . We show that the analogous property holds:

$$\exists \delta > 0, \exists k < 1, W(\mathbf{x}) < \delta \Rightarrow W(\mathbf{x}') \leq kW(\mathbf{x}) \quad (10.6.8)$$

The result clearly follows.

It is easy to see that the mean-payoff on Γ is always greater than $(-1 + \epsilon)/4$ (this value is attained in the middle of each edge $[\mathbf{e}_i, \mathbf{e}_{i+1}]$). Therefore, the assumption $\alpha < (1 - \epsilon)/12$ implies that there exists $\delta > 0$ such that if $W(\mathbf{x}) \leq \delta$ then $(\mathbf{U}\mathbf{x})_4 < \mathbf{x} \cdot \mathbf{U}\mathbf{x}$. Therefore, by compactness of $K_\delta := \{\mathbf{x} : W(\mathbf{x}) \leq \delta\}$,

$$\max_{\mathbf{x} \in K_\delta} \frac{C + (\mathbf{U}\mathbf{x})_4}{C + \mathbf{x} \cdot \mathbf{U}\mathbf{x}} < 1$$

Together with (10.6.6), this implies

$$\exists k_1 < 1, W(\mathbf{x}) < \delta \Rightarrow x'_4 \leq k_1 x_4 \quad (10.6.9)$$

Furthermore, as its continuous-time counterpart, the dynamics (10.6.6) may be decomposed. Precisely (derivation omitted), for every i in $\{1, 2, 3\}$,

$$\hat{x}'_i = \hat{x}_i \frac{C_t + (\hat{\mathbf{U}}\hat{\mathbf{x}})_i}{C_t + \hat{\mathbf{x}} \cdot \hat{\mathbf{U}}\hat{\mathbf{x}}} \quad \text{with } C_t = \frac{C - \alpha x_4}{1 - x_4} \quad (10.6.10)$$

To compare with (10.4.1), note that by subtracting \hat{x}'_i to both sides, we obtain:

$$\hat{x}'_i - \hat{x}_i = (1 - x_4) \hat{x}_i \frac{(\hat{\mathbf{U}}\hat{\mathbf{x}})_i - \hat{\mathbf{x}} \cdot \hat{\mathbf{U}}\hat{\mathbf{x}}}{C - \alpha x_4 + \hat{\mathbf{x}} \cdot \hat{\mathbf{U}}\hat{\mathbf{x}}} \quad (10.6.11)$$

It¹⁰ follows from (10.6.10) that:

$$A(\mathbf{x}) := \prod_{1 \leq i \leq 3} \frac{\hat{x}'_i}{\hat{x}_i} = \prod_{1 \leq i \leq 3} \frac{C_t + (\hat{\mathbf{U}}\hat{\mathbf{x}})_i}{C_t + \hat{\mathbf{x}} \cdot \hat{\mathbf{U}}\hat{\mathbf{x}}} \quad (10.6.12)$$

¹⁰As is well known, (10.1.1) is obtained from (10.6.7) by replacing $x'_i - x_i$ by dx_i/dt , rescaling the time from t to $\tau = t/C$ and finally letting C go to infinity. As should be, the same operations applied to (10.6.11) yield (10.4.1)

We claim that for every \mathbf{x} in K_δ , $A(\mathbf{x}) < 1$ (proof delayed). By compactness of K_δ , this implies that

$$\exists k_2 < 1, W(\mathbf{x}) \leq \delta \Rightarrow [\hat{x}'_1 \hat{x}'_2 \hat{x}'_3 \leq k_2 \hat{x}_1 \hat{x}_2 \hat{x}_3] \quad (10.6.13)$$

Equation (10.6.8) follows from (10.6.9) and (10.6.13). Therefore, it only remains to prove that $A(\mathbf{x}) < 1$ for every \mathbf{x} in K_δ . Let $\mathbf{x} \in K_\delta$. We have:

$$\ln A(\mathbf{x}) = \left(\sum_{1 \leq i \leq 3} \ln \left(C_t + (\hat{\mathbf{U}}\hat{\mathbf{x}})_i \right) \right) - 3 \ln \left(C_t + \hat{\mathbf{x}} \cdot \hat{\mathbf{U}}\hat{\mathbf{x}} \right)$$

Let $\mathbf{p} = (1/3, 1/3, 1/3)$. By concavity of the logarithm and by (10.2.6),

$$\sum_{1 \leq i \leq 3} \frac{1}{3} \ln \left(C_t + (\hat{\mathbf{U}}\hat{\mathbf{x}})_i \right) \leq \ln \left(C_t + \mathbf{p} \cdot \hat{\mathbf{U}}\hat{\mathbf{x}} \right) < \ln \left(C_t + \hat{\mathbf{x}} \cdot \hat{\mathbf{U}}\hat{\mathbf{x}} \right)$$

It follows that $\ln A(\mathbf{x})$ is negative. This completes the proof. \square

Since we assumed that the constant C is larger than all the coefficients of the game, it follows from (10.6.6) that $V(\mathbf{x}')/V(\mathbf{x})$ and x'_4/x_4 are (locally) Lipschitz in the payoff matrix \mathbf{U} . This¹¹ implies that (10.6.8) is robust to perturbation of the game, hence that elimination of the unique strategy used in correlated equilibrium occurs for an open set of games.

10.7 Discussion

1. Propositions 10.3.2 and 10.5.1 also imply that the *two-population* replicator dynamics may eliminate all strategies used in correlated equilibrium along interior solutions. See the remark in (Hofbauer and Weibull, 1996, p.571).

2. The basic idea is that if an attractor is disjoint from the set of equilibria, then it is likely that we may add a strategy in a way that strongly affects the set of equilibria but does not perturb much the dynamics in the neighborhood of the attractor.

3. We showed that elimination of all strategies used in correlated equilibrium occurs for an open set of games. In the second part of this chapter, we will see that it also occurs for vast classes of dynamics. This robustness

¹¹That is, there exists a neighborhood of (10.3.1) on which $V(\mathbf{x}')/V(\mathbf{x})$ and x'_4/x_4 are Lipschitz in \mathbf{U} .

is crucial for the practical relevancy of our results. Indeed, in practical situations, we are unlikely to have an exact knowledge of the payoffs or of the dynamics followed by the agents.

4. Proposition 10.5.1 shows much more than nonconvergence to correlated equilibrium: all strategies used in correlated equilibrium are wiped out. In particular, no kind of time-average of the replicator dynamics can converge to the set of correlated equilibria. In contrast, Hofbauer (2004) shows that, in all n -player finite games and along all interior solutions, the time-average of the (n -population) replicator dynamics converges to the Hannan set.

5. For α small, in the game with payoffs (10.3.1), there are two asymptotically stable attractors: the set Γ and the equilibrium \mathbf{e}_4 . So we face a problem of selection between attractors, similar to the standard problem of selection between two strict Nash equilibria. For α small enough, the basin of attraction of Γ is much larger than the basin of attraction of \mathbf{e}_4 (actually, as $\alpha \rightarrow 0$, the basin of attraction of Γ increases and converges towards $S_4 \setminus E_0$ where E_0 denotes the segment $[\mathbf{e}_4, \mathbf{n}_{123}]$). In particular, in models *à la* Kandori et al (1993), for α small enough, Γ would be the only stochastically stable attractor.¹²

¹²Let $i \in \{1, 2, 3\}$. As α gets small, a rough measure of the number of mutations needed to go from Γ to the basin of attraction of \mathbf{e}_4 is the distance between \mathbf{e}_i and \mathbf{n}_{123} ; a rough measure of the number of mutations needed to go from \mathbf{e}_4 to the basin of attraction of Γ is the distance between \mathbf{e}_4 and the rest-point of the replicator dynamics in $] \mathbf{e}_4, \mathbf{e}_i [$, which is of the order of α .

Part B: More General Games and Other Dynamics

In the first part of this chapter, we showed that there exists an open set of games for which, under the single-population replicator dynamics and from an open set of initial conditions, all strategies used in correlated equilibrium are eliminated. In this second part, we first give another proof of this result, which applies to more general games (section 10.8). We then show that the same result holds for a wide class of dynamics. Namely, the best-response dynamics (section 10.9), the Brown-von Neumann-Nash dynamics (section 10.10) and any monotonic or weakly-sign preserving dynamics that depends continuously on the payoffs and for which no new strategy appears by mutation (section 10.11). We also show that for the replicator and best-response dynamics, and in a sense to be made precise, elimination of all strategies used in correlated equilibria is robust to the addition of mixed strategies as new pure strategies of the game (section 10.12). Finally some proofs are gathered in sections 10.13 and 10.14.

The notations are the same as in the first part and, as in the first part, we focus on single-population dynamics in 4×4 symmetric games.

10.8 Replicator dynamics

Any game in a sufficiently small neighborhood of the game with payoffs (10.3.1) studied in part 1 is a particular case of a 4×4 symmetric game with payoffs

$$\left(\begin{array}{ccc|c} a_1 & b_2 & c_3 & d_1 \\ c_1 & a_2 & b_3 & d_2 \\ b_1 & c_2 & a_3 & d_3 \\ \hline f_1 & f_2 & f_3 & a_4 \end{array} \right) \quad (10.8.1)$$

satisfying

$$b_i < a_i < c_i \text{ for } i = 1, 2, 3, \quad \prod_{i=1}^3 (a_i - b_i) > \prod_{i=1}^3 (c_i - a_i) \quad (10.8.2)$$

and

$$f_i < a_i \text{ for } i = 1, 2, 3 \quad (10.8.3)$$

The conditions (10.8.2) mean that the game restricted to $\{1, 2, 3\} \times \{1, 2, 3\}$ is an outward cycling RPS (Rock-Paper-Scissors) game. The conditions in (10.8.3) imply that near the vertices $\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3$ of the simplex, strategy 4 earns strictly less than the mean payoff. Together with (10.8.2), this implies that there is a best-response cycle from \mathbf{e}_1 to \mathbf{e}_2 to \mathbf{e}_3 and back to \mathbf{e}_1 . As in part 1, we let Γ denote the corresponding set (see (10.3.6)).

Definition Let C be a closed subset of S_4 . The set C is *asymptotically stable* if it is:

- (a) invariant: $\mathbf{x}(0) \in C \Rightarrow (\forall t \in \mathbb{R}, \mathbf{x}(t) \in C)$
- (b) Lyapunov stable: for every neighborhood N_1 of C , there exists a neighborhood N_2 of C such that, for every initial solution $\mathbf{x}(0)$ in N_2 , $\mathbf{x}(t) \in N_1$ for all $t \geq 0$.
- (c) locally attracting: there exists a neighborhood N of C such that, for every initial condition $\mathbf{x}(0)$ in N , $\min_{c \in C} \|\mathbf{x}(t) - c\| \rightarrow_{t \rightarrow +\infty} 0$ (where $\|\cdot\|$ is any norm on \mathbb{R}^I .)

Proposition 10.8.1. *For every game with payoffs (10.8.1) satisfying conditions (10.8.2) and (10.8.3), the set Γ is asymptotically stable.*

Proof. The proof consists in checking that the stability criteria for heteroclinic cycles developed by Hofbauer (1994) are satisfied. As these criteria will be introduced in section 10.11, the proof is made in the appendix (section 10.13). \square

It follows from proposition 10.3.2 that there exists an open set of games with payoffs (10.8.1) satisfying conditions (10.8.2) and (10.8.3) for which $\mathbf{e}_4 \otimes \mathbf{e}_4$ is the unique correlated equilibrium. Together with proposition 10.8.1, this provides an alternate proof of the fact that there exists an open set of 4×4 symmetric games for which, from an open set of initial conditions, the unique strategy used in correlated equilibrium is eliminated.

Note that not all games satisfying (10.8.2) and (10.8.3) have $\mathbf{e}_4 \otimes \mathbf{e}_4$ as unique correlated or Nash equilibrium. Actually, proposition 10.8.1 provides an example of a family of games with a common attractor but very different sets of Nash equilibria¹³. The point is that this attractor, Γ , is bounded away from the set of equilibria and that its asymptotic stability only depends on the payoffs in its neighbourhood. This explains that the stability of Γ be in a large part unrelated to the structure of the equilibrium set.

10.9 Best-response dynamics

The best-response dynamics (Gilboa and Matsui, 1991) is given by the differential inclusion:

$$\dot{\mathbf{x}}(t) \in BR(\mathbf{x}(t)) - \mathbf{x}(t) \quad (10.9.1)$$

where $BR(\mathbf{x})$ is the set of best-responses to \mathbf{x} :

$$BR(\mathbf{x}) = \{\mathbf{y} \in S_N : \mathbf{y} \cdot \mathbf{U}\mathbf{x} = \max_{\mathbf{z} \in S_N} \mathbf{z} \cdot \mathbf{U}\mathbf{x}\}$$

A solution $\mathbf{x}(\cdot)$ of the best-response dynamics is an absolutely continuous function satisfying (10.9.1) for almost every t . In general, there might be several solutions starting from the same initial condition. However, for the games and the initial conditions that we will consider, there will be a unique solution starting from each initial condition.¹⁴

Consider a 4×4 symmetric game with payoff matrix (10.8.1) satisfying (10.8.2) and (10.8.3). Let

$$V(\mathbf{x}) := \max_{1 \leq i \leq 3} \left[(\mathbf{U}\mathbf{x})_i - \sum_{1 \leq i \leq 4} a_i x_i \right] \text{ and } W(\mathbf{x}) := \max(x_4, |V(\mathbf{x})|)$$

¹³For instance, assuming throughout that (10.8.2) and (10.8.3) are satisfied: if the f_i are low enough, then the Nash equilibrium of the underlying RPS game induces a Nash equilibrium of (10.8.1). If $d_1 = d_2 = d_3 < a_4$ [resp. $> a_4$] and if the f_i are high enough, then there is a unique correlated equilibrium: $\mathbf{e}_4 \otimes \mathbf{e}_4$ [resp. a convex combination of $\mathbf{e}_4 \otimes \mathbf{e}_4$ and of the Nash equilibrium of the underlying RPS game]; to see this, mimick the proof of proposition 10.3.2 in part 1. If $d_i < a_4$ for all i and if d_2 and d_3 are low enough (with respect to d_1 and f_1), then there is a Nash equilibrium in the interior of the edge $[\mathbf{e}_1, \mathbf{e}_4]$, etc.

¹⁴We focus on forward time and so will never study whether a solution is uniquely defined in backward time.

The set

$$ST := \{\mathbf{x} \in S_4 : W(\mathbf{x}) = 0\} \quad (10.9.2)$$

is a triangle, which, following Gaunersdorfer and Hofbauer (1995), we call the Shapley triangle. Gaunersdorfer and Hofbauer (1995) show that in the underlying RPS game, this triangle attracts all solutions of (10.9.1) except the one starting and remaining at the Nash equilibrium. Here, we show that in the full game (10.8.1), this triangle still attracts all solutions from an open set of initial conditions.

Proposition 10.9.1. *If strategy 4 is not a best-response to $\mathbf{x}(0)$ and if strategies 1, 2 and 3 are not all best-responses to $\mathbf{x}(0)$, then for all $t \geq 0$, $\mathbf{x}(t)$ is uniquely defined, and $\mathbf{x}(t)$ converges to the Shapley triangle (10.9.2) as $t \rightarrow +\infty$.*

Proof. We begin with a lemma, which is the continuous time version of the improvement principle of Monderer and Sela (1997):

Lemma 10.9.2. *Let $t_1 < t_2$, let $\mathbf{b}, \mathbf{b}' \in S_4$. Assume that $\dot{\mathbf{x}} = \mathbf{b} - \mathbf{x}$ for all t in $]t_1, t_2[$. If \mathbf{b}' is a best-response to $\mathbf{x}(t_2)$ then $\mathbf{b}' \cdot \mathbf{U}\mathbf{b} \geq \mathbf{b} \cdot \mathbf{U}\mathbf{b}$, with strict inequality if \mathbf{b}' is not a best-response to $\mathbf{x}(t_1)$.*

Proof of lemma 10.9.2. There exists λ in $]0, 1[$ such that

$$\mathbf{x}(t_2) = \lambda\mathbf{x}(t_1) + (1 - \lambda)\mathbf{b} \quad (10.9.3)$$

If \mathbf{b}' is a best-response to $\mathbf{x}(t_2)$ then $(\mathbf{b}' - \mathbf{b}) \cdot \mathbf{U}\mathbf{x}(t_2) \geq 0$ so that, using (10.9.3),

$$(1 - \lambda)(\mathbf{b}' - \mathbf{b}) \cdot \mathbf{U}\mathbf{b} \geq \lambda(\mathbf{b} - \mathbf{b}') \cdot \mathbf{U}\mathbf{x}(t_1)$$

Since the right hand side is nonnegative, and positive if \mathbf{b}' is not a best-response to $\mathbf{x}(t_1)$, the result follows. ■

Using lemma 10.9.2, it is easy to see that, at least for some time, the solution $\mathbf{x}(t)$ is uniquely defined and has the following behavior: assume for concreteness that at some time t , strategy 1 is the unique best-response to $\mathbf{x}(t)$. The solution will then point towards \mathbf{e}_1 till some time $t' > t$ when some other pure strategy becomes a best-response. Due to the improvement principle, this strategy can only be strategy 2. Thus, the solution must then point towards the edge $[\mathbf{e}_1, \mathbf{e}_2]$. Since strategy 2 strictly dominates strategy 1 in the game restricted to $\{1, 2\} \times \{1, 2\}$, it follows that immediately after time T , strategy 2 becomes the unique best-response; therefore the

solution will actually point towards \mathbf{e}_2 . The solution keeps pointing towards \mathbf{e}_2 till some other pure strategy becomes a best-response; due to - again - the improvement principle, this strategy must be strategy 3. The solution then changes direction again, and points towards \mathbf{e}_3 till 1 becomes a best-response again, and so on.

A priori, it could be that the times when the direction of the trajectory changes accumulate, as happens in inward cycling RPS games (see, e.g., Gaunersdorfer and Hofbauer, 1995). But we claim that:

Claim 10.9.3. *The time-length between two successive times when the direction of $\mathbf{x}(t)$ changes is bounded away from zero.*

This implies that the above description of the dynamics holds for all positive times. Now when $\mathbf{x}(t)$ points towards \mathbf{e}_i (with $i \in \{1, 2, 3\}$), letting $v(t) := V(\mathbf{x}(t))$, we have:

$$\dot{v} = (\mathbf{U}\dot{\mathbf{x}})_i - \sum_{1 \leq i \leq 4} a_i \dot{x}_i = (\mathbf{U}(\mathbf{e}_i - \mathbf{x}))_i - \left(a_i - \sum_{1 \leq i \leq 4} a_i x_i \right) = -v \quad (10.9.4)$$

and we also have $\dot{x}_4 = -x_4$. Therefore, letting $w(t) := W(\mathbf{x}(t))$,

$$\dot{w}(t) \leq -w(t) \text{ for almost all } t \geq 0$$

It follows that $w(t)$ decreases exponentially to 0, hence that $\mathbf{x}(t)$ converges to the Shapley triangle.

To complete the proof, we still need to prove claim 10.9.3:

Proof of claim 10.9.3: In what follows $i \in \{1, 2, 3\}$ and $i + 1$ is counted modulo 3. Fix an initial condition and let

$$g(t) := \max_{1 \leq i, j \leq 3} [(\mathbf{U}\mathbf{x}(t))_i - (\mathbf{U}\mathbf{x}(t))_j]$$

Let t_i^k denote the k^{th} time at which strategy i becomes a best-response and choose i such that $t_i^k < t_{i+1}^k$. Let $\mathbf{x} = \mathbf{x}(t_i^k)$, $g = g(t_i^k)$ and $\mathbf{x}' = \mathbf{x}(t_{i+1}^k)$, $g' = g(t_{i+1}^k)$. We now compute g' as a function of g .

Between t_i^k and t_{i+1}^k , the solution points towards \mathbf{e}_i . Therefore, there exists λ in $]0, 1[$ such that

$$\mathbf{x}' = \lambda \mathbf{e}_i + (1 - \lambda) \mathbf{x}$$

Furthermore, by definition of t_i^k and t_{i+1}^k , we have:

$$(\mathbf{U}\mathbf{x})_{i-1} = (\mathbf{U}\mathbf{x})_i = (\mathbf{U}\mathbf{x})_{i+1} + g \quad \text{and} \quad (\mathbf{U}\mathbf{x}')_i = (\mathbf{U}\mathbf{x}')_{i+1} = (\mathbf{U}\mathbf{x}')_{i-1} + g'$$

Therefore

$$0 = (\mathbf{e}_{i+1} - \mathbf{e}_i) \cdot \mathbf{U}\mathbf{x}' = (\mathbf{e}_{i+1} - \mathbf{e}_i) \cdot \mathbf{U}(\lambda\mathbf{e}_i + (1 - \lambda)\mathbf{x}) = \lambda(c_i - a_i) - (1 - \lambda)g$$

and

$$g' = (\mathbf{e}_i - \mathbf{e}_{i-1}) \cdot \mathbf{U}\mathbf{x}' = (\mathbf{e}_i - \mathbf{e}_{i-1}) \cdot \mathbf{U}(\lambda\mathbf{e}_i + (1 - \lambda)\mathbf{x}) = \lambda(a_i - b_i)$$

Solving for g' we get $g'/g = \alpha_i/(g + \beta_i)$ with $\alpha_i = a_i - b_i$ and $\beta_i = c_i - a_i$. Iterating this argument, we obtain the return map:

$$g(t_i^{k+1}) = \frac{\alpha_1\alpha_2\alpha_3}{\beta_1\beta_2\beta_3 + g(t_i^k)(\alpha_1\alpha_2 + \alpha_1\beta_3 + \beta_2\beta_3)}g(t_i^k)$$

Since, by condition (10.8.2), $\alpha_1\alpha_2\alpha_3 > \beta_1\beta_2\beta_3$, it follows that for small $g(t_i^k)$, we have $g(t_i^{k+1}) > g(t_i^k)$, hence that $g(t_i^k)$ is bounded away from zero. Now, since $(\mathbf{U}\mathbf{x}(t))_i - (\mathbf{U}\mathbf{x}(t))_{i+1}$ decreases from $g(t_i^k)$ to 0 between t_i^k and t_{i+1}^k , and since the speed is bounded, it follows that $t_{i+1}^k - t_i^k$ is bounded away from zero. This proves claim 10.9.3 and completes the proof of proposition 10.9.1. \square

Together with proposition 10.3.2, proposition 10.9.1 implies that there exists an open set of 4×4 symmetric games for which the unique strategy used in correlated equilibrium is strategy 4, but, from an open set of initial conditions, the solution $\mathbf{x}(\cdot)$ of (10.9.1) is uniquely defined and $x_4(t) \rightarrow 0$.

10.9.1 Remarks

Remark 1. For every $\eta > 0$, we may set the parameters of (10.8.1) so that the set $\{\mathbf{x} \in S_4 : \mathbf{e}_4 \in BR(\mathbf{x}) \text{ or } \{\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3\} \subseteq BR(\mathbf{x})\}$ has Lebesgue measure less than η . In this sense, the basin of attraction of the Shapley triangle may be made arbitrarily large.

Remark 2. Recall the definition of $\hat{\mathbf{x}}$ and the decomposition of the replicator dynamics in section 10.4. Consider a game of kind (10.8.1) with the additional assumption that $d_1 = d_2 = d_3$. This implies that, provided that 4 is not a best-response to \mathbf{x} , the strategies i in $\{1, 2, 3\}$ which are best-responses to \mathbf{x} are exactly those which are best-responses to $\hat{\mathbf{x}}$ in the underlying RPS game \hat{G} . It follows that, up to a change of velocity, $\hat{\mathbf{x}}$ follows the

best-response dynamics in \hat{G} . More precisely, straightforward computations show that if 4 is not a best-response to \mathbf{x} then:

$$\dot{\mathbf{x}} \in BR(\mathbf{x}) - \mathbf{x} \Rightarrow (1 - x_4)\dot{\hat{\mathbf{x}}} \in BR(\hat{\mathbf{x}}) - \hat{\mathbf{x}}$$

Interestingly, the change in velocity is the inverse of the change in velocity found for the replicator dynamics (see (10.4.1)).

Remark 3. The proof of proposition 10.9.1 uses condition (10.8.3), i.e., $f_i < a_i$ for $i = 1, 2, 3$. Since $(\mathbf{U}\mathbf{x})_4 - \sum_{1 \leq i \leq 4} a_i x_i = \sum_{1 \leq i \leq 3} (f_i - a_i)x_i$ is linear in \mathbf{x} , condition (10.8.3) means that $(\mathbf{U}\mathbf{x})_4 - \sum_{1 \leq i \leq 4} a_i x_i$ is negative on the face of the simplex spanned by $\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3$. If instead of requiring (10.8.3), we only require that $(\mathbf{U}\mathbf{x})_4 - \sum_{1 \leq i \leq 4} a_i x_i$ be negative on the Shapley triangle:

$$W(\mathbf{x}) = 0 \Rightarrow (\mathbf{U}\mathbf{x})_4 - \sum_{1 \leq i \leq 4} a_i x_i < 0 \quad (10.9.5)$$

then proposition 10.9.1 does not hold. For instance, if $f_1 > a_1$, $f_3 < a_3$ and $d_1 = d_2 = d_3 < a_4$, then from every initial condition sufficiently close to the mixed strategy $\mathbf{x} \in [\mathbf{e}_3, \mathbf{e}_4]$ to which strategies 1 and 4 are both best-responses, including initial conditions to which strategy 1 is the unique best-response, every solution of (10.9.1) converges to \mathbf{e}_4 . However,

Proposition 10.9.4. *If (10.9.5) holds, then there exists $\gamma > 0$ such that from every initial condition in $N_\gamma := \{\mathbf{x} \in S_4 : W(\mathbf{x}) < \gamma\}$, there is a unique solution to (10.9.1), and it converges to the Shapley triangle.*

Proof. If we can find $\gamma > 0$ such that on N_γ strategy 4 is never a best-response, then the proof of proposition 10.9.1 implies that, as long as $\mathbf{x}(t) \in N_\gamma$, the solution is unique and W decreases exponentially. The later implies that N_γ is forward invariant and that W goes to zero, hence the result.

Now, if (10.9.5) holds, then there exists an open neighborhood Ω of the Shapley triangle on which 4 is not a best-response. Since W is positive on the compact set $S_4 \setminus \Omega$, it follows that $\gamma := \min_{\mathbf{x} \notin \Omega} W(\mathbf{x})$ is positive. Furthermore, the definition of γ implies that $N_\gamma \subseteq \Omega$; hence, on N_γ , strategy 4 is never a best-response and the result follows. \square

Following Gaunersdorfer and Hofbauer (1995), it is interesting to compare the behavior of the best-response dynamics and of the time-average of the replicator dynamics. If $f_i > a_i$ for some $i \in \{1, 2, 3\}$, then under the replicator dynamics, the heteroclinic cycle Γ is not stable. But

Proposition 10.9.5. *If (10.9.5) holds, then Γ attracts an open set of orbits, along which the time-average converges to the Shapley triangle.*

Proof. This follows from proposition 3.1 of Brannath (1994). For completeness we give a sketch of proof in the appendix (section 10.14). \square

Remark 4. In order to understand whether elimination of all strategies used in Nash equilibrium is a common phenomenon for, say, the replicator dynamics, it is natural to try to characterize all games for which this phenomenon occurs. While such a characterization seems out of reach in the general case, it should not be too difficult to obtain for 4×4 symmetric games. More precisely, say that a game is of type 1 if under the replicator dynamics and for an open set of initial conditions, all strategies used in Nash equilibrium are eliminated. We just showed that 4×4 symmetric games satisfying (10.8.2) and (10.9.5) are of type 1. We conjecture that, conversely:

Conjecture 10.9.6. *Up to permutation of the strategies, almost all 4×4 symmetric games of type 1 satisfy (10.8.2) and (10.9.5).*

We now motivate the conjecture. All statements concern symmetric games and the replicator dynamics. Proceeding as in chapter 9, proof of proposition 9.2.1, it is easy to show that if at most two pure strategies survive, then the solution converges to the set of Nash equilibria. Therefore a 4×4 game of type 1 must have a unique, pure Nash equilibrium. By uniqueness and since the game is symmetric, this equilibrium must be symmetric, and we may assume that this is $(\mathbf{e}_4, \mathbf{e}_4)$. Furthermore, as a unique and pure Nash equilibrium of a bimatrix game, this equilibrium is strict.¹⁵

Now focus on the three strategies not used in Nash equilibrium. By the folk theorem of evolutionary game theory, if an interior solution converges to a point, then this point is a Nash equilibrium. Therefore, in a game of type 1, we must have an open set of nonconverging solutions. Together with Bomze's (1983) classification of 3×3 symmetric games, this suggests that the three strategies not used in Nash equilibrium form an outward cycling Rock-Paper-Scissors game; more precisely, that the inequalities (10.8.2) hold at least as weak inequalities. Finally, if (10.8.2) holds and if (10.9.5) does not hold at least with a weak inequality, i.e. if there exists $\mathbf{y} \in ST$ to which

¹⁵In a bimatrix game, there always exists a quasi-strict Nash equilibrium (Norde, 1999); it follows that a unique Nash equilibrium is quasi-strict, and that a unique and pure Nash equilibrium is strict.

strategy 4 is a strict best-response, then a Poincaré map argument shows that the heteroclinic cycle Γ cannot attract an interior solution.

The reason why the conjecture concerns almost all games of type 1, and not all, is that we want to exclude the possibility that (10.8.2) and (10.9.5) hold as weak inequalities but not as strict inequalities: a degenerate case.

10.10 Brown-von Neumann-Nash dynamics

The Brown-von Neumann-Nash dynamics (henceforth BNN) is given by:

$$\dot{x}_i = k_i(\mathbf{x}) - x_i \sum_{j \in I} k_j(\mathbf{x}) \quad (10.10.1)$$

where $k_i(\mathbf{x}) := \max(0, (\mathbf{U}\mathbf{x})_i - \mathbf{x} \cdot \mathbf{U}\mathbf{x})$ is the excess payoff of strategy i over the average payoff. As in the best-response dynamics, strategies that are initially absent may appear, the proportion of every strategy earning less than average decreases and the rest-points are exactly the Nash equilibria of the game.¹⁶ Furthermore, since the right-hand side of (10.10.1) is Lipschitz continuous, BNN has a unique solution from each initial condition. We refer to (Hofbauer, 2000; Berger and Hofbauer, 2005) and references therein for a motivation of and results on BNN.

Recall the following notations introduced in part 1: G_0 denotes the game (10.3.1) with $\alpha = 0$ and \mathbf{U}_0 its payoff matrix; $\mathbf{n}_{123} = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3}, 0)$ denotes the mixed strategy corresponding to the unique Nash equilibrium of the underlying RPS game, and $E_0 = [\mathbf{n}_{123}, \mathbf{e}_4]$ the set of symmetric Nash equilibria of G_0 . This section is devoted to a proof of the following proposition:

Proposition 10.10.1. *If C is a closed subset of S_4 disjoint from E_0 , then there exists a neighborhood of G_0 such that, for every game in this neighborhood and every initial condition in C , $x_4(t) \rightarrow 0$ as $t \rightarrow +\infty$.*

Proposition 10.3.2 implies that any neighborhood of the game G_0 contains an open set of games for which the unique correlated equilibrium is $\mathbf{e}_4 \otimes \mathbf{e}_4$. Together with proposition 10.10.1, this implies that there exists an open set of games for which, under BNN, the unique strategy played in correlated equilibrium is eliminated from an open set of initial conditions. We first show that in G_0 the segment E_0 is locally repelling.

¹⁶The *symmetric* Nash equilibria, for the single-population version presented here.

The function

$$V_0(\mathbf{x}) := \frac{1}{2} \sum_{i \in I} k_i^2 = \frac{1}{2} \sum_{i \in I} [\max(0, (\mathbf{U}_0 \mathbf{x})_i - \mathbf{x} \cdot \mathbf{U}_0 \mathbf{x})]^2$$

is continuous, nonnegative and equals 0 exactly on E_0 so that $V_0(\mathbf{x})$ may be seen as a distance from \mathbf{x} to E_0 . Fix an initial condition and let $v_0(t) := V_0(\mathbf{x}(t))$.

Lemma 10.10.2. *There exists an open neighborhood N_{eq} of E_0 such that, under BNN in the game G_0 , $\dot{v}_0(t) > 0$ whenever $\mathbf{x}(t) \in N_{eq} \setminus E_0$.*

Proof. We will need the following equations (derived in part 1 as equations (10.3.2), (10.3.3) and (10.3.5), respectively):

$$\mathbf{n}_{123} \cdot \mathbf{U}_0 \mathbf{x} = \mathbf{e}_4 \cdot \mathbf{U}_0 \mathbf{x} \quad \forall \mathbf{x} \in S_4 \quad (10.10.2)$$

$$(\mathbf{x} - \mathbf{x}') \cdot \mathbf{U}_0 \mathbf{e}_4 = (\mathbf{x} - \mathbf{x}') \cdot \mathbf{U}_0 \mathbf{n}_{123} = 0 \quad \forall \mathbf{x} \in S_4, \forall \mathbf{x}' \in S_4 \quad (10.10.3)$$

and, for every \mathbf{p} in E_0 and every $\mathbf{x} \neq E_0$,

$$(\mathbf{x} - \mathbf{p}) \cdot \mathbf{U}_0 \mathbf{x} = (\mathbf{x} - \mathbf{p}) \cdot \mathbf{U}_0 (\mathbf{x} - \mathbf{p}) = \frac{1-\epsilon}{2} \sum_{1 \leq i \leq 3} \left(x_i - \frac{1-x_4}{3} \right)^2 > 0 \quad (10.10.4)$$

Hofbauer (2000) shows that the function v_0 satisfies

$$\dot{v}_0 = \bar{k}^2 [(\mathbf{q} - \mathbf{x}) \cdot \mathbf{U}_0 (\mathbf{q} - \mathbf{x}) - (\mathbf{q} - \mathbf{x}) \cdot \mathbf{U}_0 \mathbf{x}] \quad (10.10.5)$$

with $\mathbf{x} = \mathbf{x}(t)$, $\bar{k} = \sum_i k_i$ and $q_i = k_i / \bar{k}$. It follows from equation (10.10.3) that if $\mathbf{p} \in E_0$, then against \mathbf{p} all strategies earn the same payoff. Therefore, the second term $(\mathbf{q} - \mathbf{x}) \cdot \mathbf{U}_0 \mathbf{x}$ goes to 0 as \mathbf{x} approaches E_0 . Thus, to prove lemma 10.10.2, it suffices to show that as \mathbf{x} approaches E_0 , the first term $(\mathbf{q} - \mathbf{x}) \cdot \mathbf{U}_0 (\mathbf{q} - \mathbf{x})$ is positive and bounded away from 0. But it follows from (10.10.2) and (10.10.4) that for $\mathbf{x} \notin E_0$,

$$\min_{1 \leq i \leq 3} (\mathbf{U}_0 \mathbf{x})_i \leq \mathbf{n}_{123} \cdot \mathbf{U}_0 \mathbf{x} = (\mathbf{U}_0 \mathbf{x})_4 < \mathbf{x} \cdot \mathbf{U}_0 \mathbf{x} \quad (10.10.6)$$

This implies that $q_4 = 0$ and $q_i = 0$ for some i in $\{1, 2, 3\}$. Therefore, it follows from (10.10.4) that for every \mathbf{p} in E_0 ,

$$(\mathbf{q} - \mathbf{p}) \cdot \mathbf{U}_0 (\mathbf{q} - \mathbf{p}) = \frac{1-\epsilon}{2} \sum_{1 \leq i \leq 3} \left(q_i - \frac{1}{3} \right)^2 \geq \frac{1-\epsilon}{18}$$

This completes the proof. \square

We now prove proposition 10.10.1. Consider first the BNN dynamics in the game G_0 . Recall lemma 10.10.2 and let $0 < \delta < \min_{\mathbf{x} \in S_4 \setminus N_{eq}} V_0(x)$ (the latter is positive by compactness of $S_4 \setminus N_{eq}$). By lemma 10.10.2 and definition of δ , the right hand side of (10.10.5) is positive for every \mathbf{x} in $V_0^{-1}(\delta) := \{\mathbf{x} \in S_4 : V_0(\mathbf{x}) = \delta\}$. Therefore, by compactness of $V_0^{-1}(\delta)$,

$$\exists \gamma > 0, \forall \mathbf{x} \in S_4, (v_0(t) = \delta) \Rightarrow (\dot{v}_0 \geq \gamma > 0) \quad (10.10.7)$$

Let $C_\delta := \{\mathbf{x} : V_0(\mathbf{x}) \geq \delta\}$. If $V_0(\mathbf{x}) \geq \delta$, then $\mathbf{x} \notin E_0$, so that by (10.10.4), $(\mathbf{U}_0 \mathbf{x})_4 - \mathbf{x} \cdot \mathbf{U}_0 \mathbf{x}$ is negative. Since C_δ is compact, it follows that there exists $\gamma' > 0$ such that

$$\mathbf{x} \in C_\delta \Rightarrow (\mathbf{U}_0 \mathbf{x})_4 - \mathbf{x} \cdot \mathbf{U}_0 \mathbf{x} \leq -\gamma' < 0 \quad (10.10.8)$$

so that x_4 decreases as long as $\mathbf{x} \in C_\delta$. Since by (10.10.7), the set C_δ is forward invariant, it follows that if $\mathbf{x}(0) \in C_\delta$, then $x_4(t) \rightarrow 0$ as $t \rightarrow +\infty$.

Moreover, V_0 is Lipschitz in \mathbf{x} , and $\dot{\mathbf{x}}$ is Lipschitz in the payoff matrix \mathbf{U} . Therefore, for \mathbf{U} close enough to \mathbf{U}_0 , we still have $v_0(t) = \delta \Rightarrow \dot{v}_0 > 0$ under the perturbed dynamics, and also $\mathbf{x} \in C_\delta \Rightarrow (\mathbf{U} \mathbf{x})_4 - \mathbf{x} \cdot \mathbf{U} \mathbf{x} < 0$. Therefore, the above reasoning applies and for every initial condition in C_δ , we have $x_4(t) \rightarrow 0$ as $t \rightarrow +\infty$.

Finally, the set C is compact and disjoint from the compact E_0 , hence bounded away from E_0 . Therefore, for δ small enough, $C \subset C_\delta$. This completes the proof ■

10.11 Monotonic and weakly sign preserving dynamics

Consider a dynamics of the form

$$\dot{x}_i = x_i g_i(\mathbf{x}) \quad (10.11.1)$$

where the C^1 functions g_i have the property that $\sum_{i \in I} x_i g_i(\mathbf{x}) = 0$ for all \mathbf{x} in S_4 , so that S_4 and its boundary faces are invariant.

Such a dynamics is *monotonic* if the growth rates of the different strategies

are ranked according to their payoffs¹⁷:

$$g_i(\mathbf{x}) > g_j(\mathbf{x}) \Leftrightarrow (\mathbf{U}\mathbf{x})_i > (\mathbf{U}\mathbf{x})_j \quad \forall i \in I, \forall j \in I \quad (10.11.2)$$

A dynamics of type (10.11.1) is *weakly sign preserving* (WSP) (Ritzberger and Weibull, 1995) if whenever a strategy earns below average, its growth rate is negative:

$$[(\mathbf{U}\mathbf{x})_i < \mathbf{x} \cdot \mathbf{U}\mathbf{x}] \Rightarrow g_i(\mathbf{x}) < 0 \quad (10.11.3)$$

Before¹⁸ stating the result, a definition is still needed: implicitly, dynamics of type (10.11.1) depend on the payoff matrix \mathbf{U} . Thus, a more correct writing of (10.11.1) would be:

$$\dot{x}_i = x_i g_i(\mathbf{x}, \mathbf{U})$$

where $\mathbf{U} \in \mathbb{R}^{N \times N}$. Such a dynamics *depends continuously on the payoff matrix* if, for every i in I , the functions g_i are defined for an open set of payoff matrices and depend continuously on \mathbf{U} . We now state the result: fix a monotonic or WSP dynamics (10.11.1) that depends continuously on the payoff matrix.

Proposition 10.11.1. *For every α in $]0, 1/3[$, there exists $\epsilon > 0$ such that for every game in the neighborhood of the game with payoffs (10.3.1), the set Γ defined by (10.3.6) is asymptotically stable.*

Proof. For every monotonic or WSP dynamics (10.11.1), and for every game in the neighborhood of (10.3.1), the set Γ is a so-called heteroclinic cycle. That is, a globally invariant set, consisting of saddle rest points and of the saddle orbits connecting these rest points. Thus we may use the asymptotic stability's criteria for heteroclinic cycles developed by Hofbauer (1994) (a more accessible reference for this result is theorem 17.5.1 in Hofbauer and Sigmund, 1998). Specifically, associate with the heteroclinic cycle Γ its so-called characteristic matrix. That is, the 3×4 matrix whose entry in row

¹⁷As already mentioned in chapter 8, this property goes under various names in the literature: *relative monotonicity* in (Nachbar, 1990), *order-compatibility* of pre-dynamics in (Friedman, 1991), *monotonicity* in (Samuelson and Zhang, 1992) and *payoff monotonicity* in (Hofbauer and Weibull, 1996).

¹⁸Instead of dynamics of type (10.11.1), Ritzberger and Weibull (1995) consider dynamics of the more general type $\dot{x}_i = h_i(\mathbf{x})$, that need not leave the faces of the simplex positively invariant. Thus, we only consider a subclass of their WSP dynamics.

i and column j is $g_j(e_i)$ (i.e. the eigenvalue in the direction of \mathbf{e}_j of the linearization of the vector field at \mathbf{e}_i):

	1	2	3	4
e_1	0	$g_2(e_1)$	$g_3(e_1)$	$g_4(e_1)$
e_2	$g_1(e_2)$	0	$g_3(e_2)$	$g_4(e_2)$
e_3	$g_1(e_3)$	$g_2(e_3)$	0	$g_4(e_3)$

(Note that $g_i(e_i) = 0$ because \mathbf{e}_i is a rest point of (10.11.1).)

Call \mathbf{C} this matrix. If \mathbf{p} is a real vector, let $\mathbf{p} < 0$ (resp. $\mathbf{p} > 0$) mean that all coordinates of \mathbf{p} are negative (resp. positive). Hofbauer (1994) shows that if the following conditions are satisfied:

$$\Gamma \text{ is asymptotically stable within the boundary of } S_4 \text{ }^{19} \quad (10.11.4)$$

$$\text{There exists a vector } \mathbf{p} \text{ in } \mathbb{R}^4 \text{ such that } \mathbf{p} > 0 \text{ and } \mathbf{C}\mathbf{p} < 0 \quad (10.11.5)$$

then Γ is asymptotically stable. Therefore, it is enough to show that these conditions hold for small positive ϵ . We begin with a lemma. In the remainder of this section, $i \in \{1, 2, 3\}$ and $i - 1$ and $i + 1$ are counted modulo 3.

Lemma 10.11.2. *For every $0 < \alpha < 1/3$, there exists $\epsilon > 0$ such that in the game with payoffs (10.3.1) and for every i in $\{1, 2, 3\}$,*

$$g_4(\mathbf{e}_i) < 0 \text{ and } 0 < g_{i+1}(\mathbf{e}_i) < -g_{i-1}(\mathbf{e}_i) \quad (10.11.6)$$

Proof of lemma 10.11.2 for monotonic dynamics. For $\epsilon = 0$, at the vertex \mathbf{e}_i , the payoff of strategy 4 (resp. $i + 1$) is strictly smaller (greater) than the payoff of strategy i . Since the growth rate of strategy i at \mathbf{e}_i is 0, this implies by monotonicity $g_4(\mathbf{e}_i) < 0$ (resp. $g_{i+1}(\mathbf{e}_i) > 0$). Moreover, for $\epsilon = 0$, we have: $(\mathbf{U}\mathbf{e}_i)_i = (\mathbf{U}\mathbf{e}_i)_{i+1} > (\mathbf{U}\mathbf{e}_i)_{i-1}$ so that $0 = g_{i+1}(\mathbf{e}_i) < -g_{i-1}(\mathbf{e}_i)$. Since the dynamics depends continuously on the payoff matrix, the latter strict inequality still holds for small $\epsilon > 0$. ■

Proof of lemma 10.11.2 for WSP dynamics. For concreteness, set $i = 2$. At \mathbf{e}_2 , strategy 4 earns less than average. Therefore $g_4(\mathbf{e}_2) < 0$. Now consider the case $\epsilon = 0$: At every point \mathbf{x} in the (relative) interior of the edge $[\mathbf{e}_1, \mathbf{e}_2]$, strategy 3 earns strictly less than average hence its growth rate is negative.

¹⁹That is, for each proper face (subsimplex) F of S_4 , if $\Gamma \cap F$ is nonempty, then it is asymptotically stable for the dynamics restricted to F .

By continuity at \mathbf{e}_2 this implies $g_3(\mathbf{e}_2) \leq 0$. Since at \mathbf{e}_2 , strategy 1 earns strictly less than average, it follows that $g_1(\mathbf{e}_2) < 0$, hence $0 < -g_1(\mathbf{e}_2)$, hence $g_3(\mathbf{e}_2) < -g_1(\mathbf{e}_2)$. Since the dynamics depends continuously on the payoff matrix, the latter strict inequality still holds for small $\epsilon > 0$.

To establish (10.11.6), we still need to show that $g_3(\mathbf{e}_2)$ is positive for small positive ϵ . Let $\epsilon > 0$. If $\lambda > 0$ is sufficiently small then, for all $\mu > 0$ small enough, the unique strategy which earns above average at $\mathbf{x} = (\lambda\mu, 1 - \mu - \lambda\mu, \mu, 0)$ is strategy 3. Since $\sum_i x_i g_i(\mathbf{x}) = 0$, it follows that $x_1 g_1(\mathbf{x}) + x_3 g_3(\mathbf{x}) = \lambda\mu g_1(\mathbf{x}) + \mu g_3(\mathbf{x}) > 0$, hence $g_3(\mathbf{x}) > -\lambda g_1(\mathbf{x})$. By letting μ go to zero, we obtain $g_3(\mathbf{e}_2) \geq -\lambda g_1(\mathbf{e}_2) > 0$. ■

We now prove proposition 10.11.1. Fix α and ϵ as in lemma 10.11.2. Note that since the dynamics we consider depends continuously on the payoff matrix, there exists a neighborhood of the game with payoffs (10.3.1) in which the strict inequalities (10.11.6) still hold. Fix a game for which (10.11.6) holds.

Proof that condition (10.11.5) holds. It follows from (10.11.6) that $g_4(e_i)$ is negative for all i in $\{1, 2, 3\}$. This implies that condition 2 holds (fix $p_1 = p_2 = p_3 = -1$ and take a very high p_4).

Proof that condition (10.11.4) holds. To prove asymptotic stability on the boundary, we use again characteristic matrices. Let $\hat{\mathbf{C}}$ denote the 3×3 matrix obtained from \mathbf{C} by eliminating the fourth column. This corresponds to the characteristic matrix of Γ , when viewed as an heteroclinic cycle of the underlying 3×3 RPS game. In this RPS game, the set Γ is trivially asymptotically stable on the relative boundary of S_3 (Γ is the relative boundary!). Furthermore, for $\hat{\mathbf{p}} = (1/3, 1/3, 1/3) > 0$, the second inequation in (10.11.6) implies that $\hat{\mathbf{C}}\hat{\mathbf{p}} < 0$. Therefore, it follows from theorem 1 of Hofbauer (1994) than, in the 4×4 initial game, Γ is asymptotically stable on the face spanned by $\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3$. Asymptotic stability on the face spanned by $\mathbf{e}_i, \mathbf{e}_{i+1}, \mathbf{e}_4$ is easy. This concludes the proof. □

10.12 Robustness to the addition of mixed strategies as new pure strategies

We showed that for many dynamics, there exists an open set of symmetric 4×4 games for which, from an open set of initial conditions, the unique strategy used in correlated equilibrium is eliminated. Since we might not want to rule out the possibility that some players use mixed strategies, and that mixed strategies be heritable, it is important to check whether our results change if we explicitly introduce mixed strategies as new pure strategies of the game. We tackle this question for the best-response dynamics and for the replicator dynamics. We first precise the issue we have in mind and introduce some notations.

Let G be a finite game with strategy set $I = \{1, \dots, N\}$ and payoff matrix \mathbf{U} . A finite game G' is said to be *built on G by adding mixed strategies as new pure strategies* if:

First, letting $I' = \{1, \dots, N, N+1, \dots, N'\}$ be the set of pure strategies of G' and \mathbf{U}' its payoff matrix, we may associate to each pure strategy i in I' a mixed strategy \mathbf{p}^i in S_N in such a way that:

$$\forall i \in I', \forall j \in I', \mathbf{e}'_i \cdot \mathbf{U}' \mathbf{e}'_j = \mathbf{p}^i \cdot \mathbf{U} \mathbf{p}^j \quad (10.12.1)$$

where \mathbf{e}'_i is the unit vector in $S_{N'}$ corresponding to the pure strategy i .

Second, if $1 \leq i \leq N$, the pure strategy i in the game G' corresponds to the pure strategy i in the base game G :

$$1 \leq i \leq N \Rightarrow \mathbf{p}^i = \mathbf{e}_i \quad (10.12.2)$$

If $\mu' = (\mu(k, l))_{1 \leq k, l \leq N'}$ is a probability distribution over $I' \times I'$, then it induces the probability distribution μ on $I \times I$ given by:

$$\mu(i, j) = \sum_{1 \leq k, l \leq N'} \mu'(k, l) p_i^k p_j^l \quad \forall (i, j) \in I \times I,$$

It follows from a version of the revelation principle (see Myerson, 1994) that, if G' is built on G by adding mixed strategies as new pure strategies, then for any correlated equilibrium μ' of G' , the induced probability distribution on $I \times I$ is a correlated equilibrium of G . In particular, if G is a 4×4 symmetric game with $\mathbf{e}_4 \otimes \mathbf{e}_4$ as unique correlated equilibrium, then μ' is a correlated equilibrium of G' if and only if, for every k, l in I' such that $\mu'(k, l)$ is positive,

$\mathbf{p}^k = \mathbf{p}^l = \mathbf{e}_4$. Thus, the unique strategy of G used in correlated equilibria of G' is strategy 4.

The question investigated below is whether there exists an open set of 4×4 symmetric games with $\mathbf{e}_4 \otimes \mathbf{e}_4$ as unique correlated equilibrium such that, for any game G in this set and any game G' built on G by adding mixed strategies as new pure strategies, we have: from an open set of initial conditions, every pure strategy k in I' such that $p_4^k > 0$ is eliminated.

We show that this is the case for the best-response dynamics and for the replicator dynamics.²⁰ For the best-response dynamics, we actually prove a much more general result: for any finite game and in a sense made precise in the next section, adding mixed strategies as new pure strategies does not modify the behavior of the best-response dynamics.

10.12.1 Best-response dynamics

Let G be a finite game and let G' be built on G by adding mixed strategies as new pure strategies. We want to relate the behaviour of the best-response dynamics in the game G' to its behaviour in the base game G . For this purpose, associate to each mixed strategy \mathbf{x}' in $S_{N'}$ the induced mixed strategy \mathbf{x} in S_N defined by:

$$\mathbf{x} := \sum_{k=1}^{N'} x'_k \mathbf{p}^k \quad (10.12.3)$$

Let $\mathbf{x}'(\cdot)$ be a solution of the best-response dynamics in G' and $\mathbf{x}(\cdot)$ the induced mapping from \mathbb{R}_+ to S_N .

Proposition 10.12.1. *$\mathbf{x}(\cdot)$ is a solution of the best-response dynamics in G .*

Before proving proposition 10.12.1, note that together with propositions 10.3.2 and 10.9.1, proposition 10.12.1 implies that there indeed exists an open set of 4×4 symmetric games such that, for every game G in this set and every game G' built on G by adding mixed strategies as new pure strategies, we have: from an open set of initial conditions, every pure strategy k such that $p_4^k > 0$ is eliminated. We now prove proposition 10.12.1.

²⁰For BNN, the condition $p_4^k > 0$ does not seem sufficient to ensure elimination. This is because, as noted in (Berger and Hofbauer, 2005), BNN tends to equalize the shares of pure strategies that earn nearly the same payoff. Thus, if $p^k = (1 - p_4^k)\mathbf{e}_1 + p_4^k\mathbf{e}_4$ with p_4^k very low, then it seems likely that if strategy 1 survives, then so will strategy k .

Proof. For almost all $t \geq 0$, there exists a vector $\mathbf{b}' \in BR(\mathbf{x}'(t))$ such that $\dot{\mathbf{x}}'(t) = \mathbf{b}' - \mathbf{x}'(t)$. Let $\mathbf{b} := \sum_{k \in I'} b'_k \mathbf{p}^k \in S_N$. It follows from (10.12.3) that:

$$\dot{\mathbf{x}}(t) = \sum_{k=1}^{N'} (b'_k - x'_k) p^k = \mathbf{b} - \mathbf{x}(t) \quad (10.12.4)$$

Furthermore, since \mathbf{b}' is a best-response to $\mathbf{x}'(t)$ it follows from (10.12.1) and (10.12.2) that \mathbf{b} is a best-response to $\mathbf{x}(t)$ (otherwise, letting $i \in \{1, \dots, N\}$ be a best-response to \mathbf{x} , we have: $\mathbf{b}' \cdot \mathbf{U}' \mathbf{x}' = \mathbf{b} \cdot \mathbf{U} \mathbf{x} < \mathbf{e}_i \cdot \mathbf{U} \mathbf{x} = \mathbf{e}'_i \cdot \mathbf{U}' \mathbf{x}'$, hence \mathbf{b}' is not a best-response to \mathbf{x}' , a contradiction). Together with (10.12.4), this implies that, for almost all t , $\dot{\mathbf{x}} \in BR(\mathbf{x}) - \mathbf{x}$. The result follows. \square

We have shown that any solution of the best-response dynamics in the augmented game G' induces in the natural way a solution of the best-response dynamics in the base game G . Since there may be several solutions to the best-response dynamics with the same initial condition, it is also interesting to check that for any \mathbf{x}' in $S_{N'}$, any solution in the base game with initial condition \mathbf{x} (induced by \mathbf{x}') is induced by some solution in the augmented game with initial condition \mathbf{x}' . This is our next result:²¹

Let $\mathbf{x}' \in S_{N'}$. Define \mathbf{x} as in (10.12.3). Let $\mathbf{x}(\cdot)$ be a solution to (10.9.1) in G with initial condition $\mathbf{x}(0) = \mathbf{x}$.

Proposition 10.12.2. *There exists a solution $\mathbf{x}'(\cdot)$ to (10.9.1) in G' with initial condition $\mathbf{x}'(0) = \mathbf{x}'$ such that, for all $t \geq 0$, $\mathbf{x}'(t)$ induces $\mathbf{x}(t)$ in the sense of (10.12.3).*

Proof. Let $t \geq 0$. For all k in $\{N+1, \dots, N'\}$ let $x'_k(t) = x'_k(0) \exp(-t)$ and choose $x'_1(t), \dots, x'_N(t)$ such that $\mathbf{x}'(t)$ induces $\mathbf{x}(t)$. To any vector \mathbf{b} in S_N associate the vector \mathbf{b}' in $S_{N'}$ such that: $b'_k = b_k$ if $1 \leq k \leq N$ and $b'_k = 0$ otherwise. A simple computation shows that if $\dot{\mathbf{x}} = \mathbf{b} - \mathbf{x}$ and $\mathbf{b} \in BR(\mathbf{x})$ then $\dot{\mathbf{x}}' = \mathbf{b}' - \mathbf{x}'$ and $\mathbf{b}' \in BR(\mathbf{x}')$. It follows that $\mathbf{x}'(\cdot)$ is a solution of the best-response dynamics in G' . \square

10.12.2 Replicator dynamics

Recall that \mathbf{U}_0 denote the payoff matrix (10.3.1) with $\alpha = 0$ and for \mathbf{x} in S_4 let

$$V(\mathbf{x}) := 3 \frac{(x_1 x_2 x_3)^{1/3}}{x_1 + x_2 + x_3} \quad (10.12.5)$$

²¹This second result is unrelated to the issue of elimination of all strategies used in correlated equilibrium, but we think that it is interesting as a general result.

Fix a 4×4 matrix \mathbf{U} and a solution $\mathbf{x}(\cdot)$ of the replicator dynamics with payoff matrix \mathbf{U} , and let $v(t) := V(\mathbf{x}(t))$. Fix δ in $]0, 1[$. Proceeding as in the proof of proposition 10.5.1, we get that there exists a positive constant γ such that if $\|\mathbf{U} - \mathbf{U}_0\| < \gamma$ then:

$$[x_1 + x_2 + x_3 = 1 \text{ and } V(\mathbf{x}) \leq \delta] \Rightarrow [(\mathbf{U}\mathbf{x})_4 - \mathbf{x} \cdot \mathbf{U}\mathbf{x} \leq -\gamma] \quad (10.12.6)$$

and

$$[x_1(t) + x_2(t) + x_3(t) = 1 \text{ and } v(t) = \delta] \Rightarrow \dot{v}(t) \leq -\gamma \quad (10.12.7)$$

Fix a game G with payoff matrix \mathbf{U} such that $\|\mathbf{U} - \mathbf{U}_0\| < \gamma$. Let G' be a game built on G by adding mixed strategies of G as new pure strategies, and let \mathbf{U}' be its payoff matrix. Recall 10.12.5 and for \mathbf{x}' in $S_{N'}$, let

$$V'(\mathbf{x}') := 3 \frac{(x'_1 x'_2 x'_3)^{1/3}}{x'_1 + x'_2 + x'_3}$$

Consider a solution $\mathbf{x}'(\cdot)$ of the replicator dynamics in G' and let $v'(t) = V'(\mathbf{x}'(t))$. It follows from (10.12.6) and (10.12.7) that there exists positive constant η and γ' such that

$$\left[\max_{k \in \{4, \dots, N'\}} x'_k \leq \eta \text{ and } V'(\mathbf{x}') \leq \delta \right] \Rightarrow (\mathbf{U}'\mathbf{x}')_4 \leq \mathbf{x}' \cdot \mathbf{U}'\mathbf{x}' - \gamma' \quad (10.12.8)$$

and

$$\left[\max_{k \in \{4, \dots, N'\}} x'_k \leq \eta \text{ and } v'(t) = \delta \right] \Rightarrow \dot{v}'(t) \leq -\gamma' \quad (10.12.9)$$

Fix $\mathbf{y}' \in S_{N'}$ such that $y'_1 + y'_2 + y'_3 = 1$, $V(\mathbf{y}') < \delta$ and $C := \min_{1 \leq i \leq 3} y'_i > 0$. There exists a open neighborhood Ω of \mathbf{y}' in $S_{N'}$ such that for every \mathbf{x}' in Ω ,

$$\min_{1 \leq i \leq 3} x'_i > C/2, \quad \max_{k \in \{4, \dots, N'\}} x'_k < C\eta/2, \quad \text{and } V'(\mathbf{x}') < \delta$$

Consider an interior solution $\mathbf{x}'(\cdot)$ of the replicator dynamics in G' with initial condition in Ω . Recall that \mathbf{p}^k denote the mixed strategy of G associated with the pure strategy k of G' . We will show that:

Proposition 10.12.3. *For all k in $\{4, \dots, N'\}$ such that $p_4^k > 0$, $x'_k(t) \rightarrow_{t \rightarrow +\infty} 0$*

Proof. We begin with two lemmas:²²

Lemma 10.12.4. *Let $T > 0$ and $k \in \{4, \dots, N'\}$. If $x'_4(T) \leq x'_4(0)$ then $x'_k(T) < \eta$.*

Proof. By definition of \mathbf{p}^k ,

$$(\mathbf{U}'\mathbf{x}')_k = \sum_{1 \leq i \leq 4} p_i^k (\mathbf{U}'\mathbf{x}')_i$$

Therefore, it follows from the definition of the replicator dynamics that:

$$\frac{\dot{x}'_k}{x'_k} = \sum_{1 \leq i \leq 4} p_i^k \frac{\dot{x}'_i}{x'_i}$$

Integrating between 0 and T and taking the exponential of both members leads to:

$$x'_k(T) = x'_k(0) \prod_{1 \leq i \leq 4} \left(\frac{x'_i(T)}{x'_i(0)} \right)^{p_i^k} \quad (10.12.10)$$

Noting that for $1 \leq i \leq 3$, we have $x'_i(T) \leq 1$ and $1 \leq 1/x'_i(0) \leq 2/C$, we get:

$$\prod_{1 \leq i \leq 3} \left(\frac{x'_i(T)}{x'_i(0)} \right)^{p_i^k} \leq \prod_{1 \leq i \leq 3} \left(\frac{2}{C} \right)^{p_i^k} = \left(\frac{2}{C} \right)^{1-p_4^k} \leq \frac{2}{C} \quad (10.12.11)$$

Since furthermore $x'_k(0) < C\eta/2$, we obtain from (10.12.10) and (10.12.11):

$$x'_k(T) = \frac{C\eta}{2} \frac{2}{C} \left(\frac{x'_4(T)}{x'_4(0)} \right)^{p_4^k} < \eta \left(\frac{x'_4(T)}{x'_4(0)} \right)^{p_4^k} \quad (10.12.12)$$

The result follows. □

Lemma 10.12.5. *For all $t > 0$, $\max_{k \in \{4, \dots, N'\}} x'_k(t) < \eta$ and $v'(t) < \delta$*

²²The heteroclinic cycle $\Gamma' : \mathbf{e}'_1 \rightarrow \mathbf{e}'_2 \rightarrow \mathbf{e}'_3 \rightarrow \mathbf{e}'_1$ need not be asymptotically stable on the boundary of $S_{N'}$. This is because there might be strategies $k \neq i+1$ (counted modulo 3) such that p^k is a better response to \mathbf{e}_i than \mathbf{e}_i (e.g. $p^k = (1-\lambda)\mathbf{e}_{i+1} + \lambda\mathbf{e}_4$, with λ small). Therefore we cannot use theorem 1 of (Hofbauer, 1994) as in our first proof of proposition 10.8.1.

Proof. Otherwise there exists a first time $T > 0$ such that $\max_{k \in \{4, \dots, N'\}} x'_k(T) = \eta$ or $v'(T) = \delta$ (or both). It follows from (10.12.8) and the definition of the replicator dynamics that if $0 \leq t \leq T$ then $\dot{x}'_4(t) \leq -\gamma' < 0$. Therefore $x'_4(T) \leq x'_4(0)$. By lemma 10.12.4, this implies that $\max_{k \in \{4, \dots, N'\}} x'_k(T) < \eta$. Therefore, $v'(T) = \delta$. Due to (10.12.9) this implies that $\dot{v}'(T) < 0$. Therefore, there exists a time $T_1 < T$ such that $v(T_1) \geq \delta$, contradicting the minimality of T . \square

We now conclude: it follows from lemma 10.12.5, equation (10.12.8) and the definition of the replicator dynamics that for all $t \geq 0$, $x'_4(t) \leq \exp(-\gamma't)x'_4(0)$. By (10.12.12) this implies that for every k in $\{4, \dots, N'\}$,

$$\forall t \geq 0, x_k(t) < \eta \exp(-\gamma' p_4^k t)$$

Therefore, if $p_4^k > 0$ then $x'_k(t) \rightarrow 0$. \square

10.13 First appendix: Proof of proposition 10.8.1

We provide two proofs (which we think provide different insights). The first one, in the spirit of section 10.11, consists in checking that the sufficient conditions for asymptotic stability of heteroclinic cycles given by Hofbauer (1994) are satisfied. The second proof, in the spirit of proposition 10.5.1 exhibits an (average) Lyapunov function.²³ In both proofs, $i \in \{1, 2, 3\}$ and $i + 1$ and $i - 1$ are counted modulo 3.

Proof 1. We use the tools introduced at the beginning of the proof of proposition 10.11.1 (up to lemma 10.11.2): The heteroclinic cycle Γ is asymptotically stable on the boundary of S_4 : asymptotic stability on the face spanned by $\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3$ follows from condition (10.8.2), as shown by Zeeman (1980); asymptotic stability on the face spanned by $\mathbf{e}_i, \mathbf{e}_{i+1}, \mathbf{e}_4$ is easy. Furthermore, under the replicator dynamics, the characteristic matrix \mathbf{C} of

²³Our first proof relies on theorem 1 of Hofbauer (1994), which itself relies on the construction of an (average) Lyapunov function; however this average Lyapunov function is more abstract than the one we exhibit in our second proof.

Γ has the sign structure

	1	2	3	4
e_1	0	-	+	-
e_2	+	0	-	-
e_3	-	+	0	-

It follows that there exists a vector \mathbf{p} in \mathbb{R}^4 such that $\mathbf{p} > 0$ and $\mathbf{C}\mathbf{p} < 0$ (fix $p_1 = p_2 = p_3 = 1$ and take a very high p_4). By theorem 1 of Hofbauer (1994), this implies that Γ is asymptotically stable. ■

Proof 2 (sketch). Applying lemma 7 from Zeeman (1980) in the spirit of (Hofbauer and Sigmund, 1998, proof of theorem 7.7.2), we may assume without loss of generality that there exists a positive constant c such that $b_i - a_{i+1} = c$ for all $i = 1, 2, 3$. Let $\mathbf{p} \in S_3$ denote the Nash equilibrium of the underlying RPS game and let $V(\mathbf{x}) = \prod_{1 \leq i \leq 3} \hat{\mathbf{x}}_i^{p_i}$ (where $\hat{\mathbf{x}}_i = x_i / (1 - x_4)$). The function \dot{V}/V extends to a continuous function on S_4 which is strictly negative on Γ (more precisely, if $x_4 = 0$, then $\dot{V} = -cV \sum_{1 \leq i \leq 3} (x_i - p_i)^2$; see Hofbauer and Sigmund, 1998, proof of theorem 7.7.2). This implies that V decreases exponentially in the neighborhood of Γ . The only difference with the proof of proposition 10.5.1 in part 1 is then that $W(\mathbf{x}) = \max(x_4, V(\mathbf{x}))$ is no longer a local Lyapunov function (because x_4 need not decrease everywhere in the neighbourhood of Γ) but a local average Lyapunov function (x_4 decreases in average over an approximate cycle). We only give the heuristic argument: there exists a neighbourhood N_i of e_i in which strategy 4 earns strictly less than the mean payoff, so that x_4 decreases. As long as $\mathbf{x}(t)$ is close enough to Γ , it describes a cycling movement from N_1 to N_2 to N_3 and back to N_1 . During this (approximate) cycle, most of the time²⁴ is spent in the N_i , so that x_4 decreases over the cycle. This allows to show that for every $\delta > 0$, there exists $\delta' > 0$ such that if $W(\mathbf{x}(0)) \leq \delta'$ then $W(\mathbf{x}(t)) \leq \delta$ for all $t \geq 0$ and $W(\mathbf{x}(t)) \rightarrow 0$ as $t \rightarrow +\infty$.

²⁴A proportion of the time which can be made arbitrarily large by requiring that the solution starts close enough to Γ .

10.14 Second appendix: (Sketch of) proof of proposition 10.9.5

It is convenient and without loss of generality to normalize the payoff matrix to have zeros on the diagonal. Consider thus a 4×4 symmetric game with payoff matrix:

$$\left(\begin{array}{ccc|c} 0 & -\alpha_2 & \beta_3 & \delta_1 \\ \beta_1 & 0 & -\alpha_3 & \delta_2 \\ -\alpha_1 & \beta_2 & 0 & \delta_3 \\ \hline \gamma_1 & \gamma_2 & \gamma_3 & 0 \end{array} \right) \quad (10.14.1)$$

such that all the α_i and β_i are positive, and assume furthermore that

$$\prod_{i=1}^3 \alpha_i > \prod_{i=1}^3 \beta_i \quad (10.14.2)$$

and that (10.9.5) holds. That is, $\mathbf{x} \in ST \Rightarrow (\mathbf{U}\mathbf{x})_4 < 0$, where ST is the Shapley triangle defined in (10.9.2). Throughout, $i \in \{1, 2, 3\}$ and $i + 1$ and $i - 1$ are counted modulo 3. Let $V(\mathbf{x}) := \min_{1 \leq i \leq 3} x_i$ and $W(\mathbf{x}) = \max(V(\mathbf{x}), x_4)$ so that $W(\mathbf{x}) = 0$ if and only if $\mathbf{x} \in \Gamma$. Fix a small positive constant η . Let

$$N_i = \{\mathbf{x} \in S_4, x_i > 1 - 2\eta, x_4 < \eta\}$$

(recall that $\hat{x}_i = x_i/(x_1 + x_2 + x_3)$). Thus, N_i is a small neighborhood of \mathbf{e}_i . Let S_i^- and S_i^+ denote the sections “before” and “after” \mathbf{e}_i :

$$S_i^- := \{\mathbf{x} \in N_i, x_{i-1} = \eta\}$$

$$S_i^+ := \{\mathbf{x} \in N_i, x_{i+1} = \eta\}$$

Consider an interior solution $\mathbf{x}(\cdot)$ with initial condition close to Γ . If

$$\forall t \geq 0, W(\mathbf{x}(t)) < \eta/2 \quad (10.14.3)$$

then $\mathbf{x}(t)$ describes cycles

$$S_i^- \rightarrow S_i^+ \rightarrow S_{i+1}^- \rightarrow S_{i+1}^+ \rightarrow S_{i+2}^- = S_{i-1}^- \rightarrow S_{i-1}^+ \rightarrow S_i^- \quad (10.14.4)$$

Let $\mathbf{x}(t_i^-) \in S_i^-$ and let t_i^+ be the first time after time t_i^- when $\mathbf{x}(t)$ hits S_i^+ :

$$t_i^+ := \min\{t \geq t_i^- : \mathbf{x}(t) \in S_i^+\}$$

Setting $y_j := x_j/\eta$ and computing the approximate Poincaré map from S_i^- to S_i^+ we get that (approximately)

$$y_{i-i}(t_i^+) = [y_{i+1}(t_i^-)]^{\alpha_i/\beta_i} \quad (10.14.5)$$

$$x_4(t_i^+) = x_4(t_i^-) [y_{i+1}(t_i^-)]^{-\gamma_i/\beta_i} \quad (10.14.6)$$

Iterating the argument²⁵, we may compute the approximate return map $S_i^- \rightarrow S_i^-$. We obtain that if \tilde{t}_i^- is the first time after t_i^- when the solution comes back to S_i^- , then (approximately)

$$y_{i-i}(\tilde{t}_i^-) = (y_{i-1}(t_i^-))^{\alpha_1\alpha_2\alpha_3/\beta_1\beta_2\beta_3} \quad (10.14.7)$$

$$x_4(\tilde{t}_i^-) = x_4(t_i^-) (y_{i+1})^{\theta_i} \quad (10.14.8)$$

with

$$-\theta_i = \frac{\gamma_i\beta_{i-1}\beta_i + \gamma_{i+1}\alpha_i\beta_{i-1} + \gamma_{i+2}\alpha_i\alpha_{i+1}}{\beta_1\beta_2\beta_3} \quad (10.14.9)$$

Let $\rho := (\alpha_1\alpha_2\alpha_3/\beta_1\beta_2\beta_3) - 1$. It follows from (10.14.2) that ρ is positive and from (10.14.7) that

$$y_{i-i}(\tilde{t}_i^-) \leq k_1 y_{i-1}(t_i^-) \text{ with } k_1 := \eta^\rho < 1 \quad (10.14.10)$$

Furthermore, a simple computation²⁶ shows that the numerator of the right-hand side of (10.14.9) is equal to $(\mathbf{U}\mathbf{q})_4$ where \mathbf{q} is the vertex of the Shapley triangle such that $(\mathbf{U}\mathbf{q})_{i-1} = (\mathbf{U}\mathbf{q})_i = 0$. Therefore, (10.9.5) implies that θ_i is positive. Since by assumption $y_{i+1}(t_i^-) \leq \eta$, it follows from (10.14.8) that

$$x_4(\tilde{t}_i^-) \leq k_2 x_4(t_i^-) \text{ with } k_2 := \eta^{\theta_i} < 1 \quad (10.14.11)$$

Fix k_3 such that $\max(k_1, k_2) < k_3 < 1$. It follows from (10.14.10) and (10.14.11) that for $W(\mathbf{x}(t_i^-))$ sufficiently small²⁷

$$W(\mathbf{x}(\tilde{t}_i^-)) \leq k_3 W(\mathbf{x}(t_i^-)) \quad (10.14.12)$$

Furthermore, if $\gamma_i < 0$ and $\gamma_{i-1} > 0$, then during a cycle (10.14.4), x_4 first decreases in the neighborhood of \mathbf{e}_i then increases or decreases near \mathbf{e}_{i+1} and

²⁵If $W(\mathbf{x}(t_i^-))$ is sufficiently small, we may neglect the transit time between S_i^+ and S_{i+1}^- .

²⁶See (Gaunersdorfer and Hofbauer, 1995) or (Hofbauer and Sigmund, 1998, p.82 and p.96).

²⁷We need $\ln x_{i+1}(t_i^-) \ll \ln \eta$ to control the error made when neglecting the transit times between the S_j^+ and the S_{j+1}^- .

finally increases near \mathbf{e}_{i-1} . In any case, x_4 starts by decreasing and then increases till \mathbf{x} reaches S_i^- . This implies that (approximately)

$$\max_{t \in [\tilde{t}_i^-, \tilde{t}_i^+]} x_4(t) = \max(x_4(t_i^-), x_4(t_i^+)) = x_4(t_i^-) \quad (10.14.13)$$

(the last equality follows from (10.14.11)).

We are now in a position to prove that if $\gamma_i < 0$ and $\gamma_{i-1} > 0$ then:

Lemma 10.14.1. *Assume that $\mathbf{x}(t) \in S_i^-$. There exist positive constants a , b , c and k , with $k < 1$ such that, if $W(\mathbf{x}(t)) < a$, then there exists $\tilde{t} > t$ such that $\mathbf{x}(\tilde{t}) \in S_i^-$, and*

$$W(\mathbf{x}(\tilde{t})) \leq kW(\mathbf{x}(t)) \quad (10.14.14)$$

$$W(\mathbf{x}(t')) \leq b[W(\mathbf{x}(t))]^c \quad \forall t' \in [t, \tilde{t}] \quad (10.14.15)$$

Proof. Let \tilde{t} be the first time ($> t$) such that either $W(\mathbf{x}(\tilde{t})) = \eta$ or $\mathbf{x}(\tilde{t}) \in S_i^-$.

Between t and \tilde{t} , the function W is less than η , so that the above estimates apply. Choose c such that $0 < \sqrt[3]{c} < \min(1, \alpha_1/\beta_1, \alpha_2/\beta_2, \alpha_3/\beta_3)$. It follows from (10.14.5) that there exists $b > 1$ and $a > 0$ such that if $W(\mathbf{x}(t)) < a$ then $V(\mathbf{x}(t')) < bV(\mathbf{x}(t))^c$ for every t' in $[t, \tilde{t}]$. Henceforth, we assume $W(\mathbf{x}(t)) < a$. It follows from (10.14.13) that we have $x_4(t') \leq x_4(t) \leq bx_4(t)^c$. Thus, assuming without loss of generality $ba^c < \eta$, we have:

$$\forall t' \in [t, \tilde{t}], W(\mathbf{x}(t')) \leq b[W(\mathbf{x}(t))]^c \leq ab^c < \eta$$

In particular, $W(\mathbf{x}(\tilde{t})) < \eta$. By definition \tilde{t} , this implies $\mathbf{x}(\tilde{t}) \in S_i^-$. Finally, it follows from (10.14.12) that (10.14.14) holds provided that a is sufficiently small. This concludes the proof. \square

It follows from lemma 10.14.1 that there exist positive constants a , b , c and k , with $k < 1$ such that, if $\mathbf{x}(0) \in S_i^-$ and $W(\mathbf{x}(0)) < a$ then

$$\limsup_{t \rightarrow +\infty} W(\mathbf{x}(t)) \leq \limsup_{n \rightarrow +\infty} k^n ba^c = 0$$

hence $\mathbf{x}(t) \rightarrow \Gamma$. This implies that for an open set of initial conditions $\mathbf{x}(t) \rightarrow \Gamma$ and²⁸ concludes the proof of proposition 10.9.5.

²⁸Let $\xi(\mathbf{x}, t)$ denote the value at time t of the solution of the replicator dynamics with initial condition \mathbf{x} . The set $\Omega := \{\xi(\mathbf{x}, t) | t \in \mathbb{R}, \mathbf{x} \in S_i^-, W(\mathbf{x}) < a\}$ contains an open set and for every \mathbf{x} in Ω , $\xi(\mathbf{x}, t) \rightarrow \Gamma$ as $t \rightarrow +\infty$.

Chapitre 11

Elimination of All Strategies in the Support of Nash Equilibria: a Universal Example

Abstract

For every Myopic Adjustment Dynamics that depends continuously on the payoffs, there exists an open set of games for which, from an open set of initial conditions, all strategies used in Nash equilibrium are (nearly) eliminated. This holds both for single-population dynamics and multi-population dynamics.

11.1 Introduction

Chapter 10 showed that under many single-population dynamics, all strategies played in correlated equilibria may be eliminated. This chapter provides examples of games for which, for an even wider class of dynamics, all strategies played in Nash equilibria are eliminated (or nearly eliminated). As in chapter 10, the examples we give are of games with a unique, strict Nash equilibrium. Thus, elimination of all strategies played in Nash equilibria does not come from a problem of coordination on the right equilibrium.¹

¹When does elimination of all strategies played in Nash equilibria occur is more difficult to clarify. Our results rely on the fact that, in the games we study, the unique Nash equilibrium is not Pareto-efficient and the mean-payoff of the players tend to increase under evolutionary dynamics.

The material is organized as follows: the main results are stated in section 11.2 and proved in section 11.3. Complementary examples are given in section 11.4. Section 11.5 concludes. Finally, an auxiliary result is proved in the appendix (section 11.6).

11.2 Framework and statement of the results

Let G be a two-player bimatrix game. Let I (resp. J) be the set of pure strategies of player 1 (resp. 2). Let $\mathbf{A} = (a_{ij})$ (resp. $\mathbf{B} = (b_{ji})$) denote the payoff matrix of player 1 (resp. 2). Here a_{ij} (resp. b_{ji}) denote the payoff of player 1 (resp. 2) when player 1 plays i and player 2 plays j . Note that the matrix \mathbf{B} has the same number of rows and columns than \mathbf{A}^T . Let $\Delta(I)$ (resp. $\Delta(J)$) denote the simplex over I (resp. J). That is,

$$\Delta(I) = \left\{ \mathbf{x} \in \mathbb{R}_+^I : \sum_{i \in I} x_i = 1 \right\}$$

The class of dynamics we consider is a variant of Swinkels'(1993) *myopic adjustment dynamics* (MAD)²: Consider two-population dynamics of the form

$$(\dot{\mathbf{x}}, \dot{\mathbf{y}}) = (f_1(\mathbf{x}, \mathbf{y}, A, B), f_2(\mathbf{x}, \mathbf{y}, A, B)) \quad (11.2.1)$$

where \mathbf{x} and \mathbf{y} are the mean strategies in, respectively, population 1 and population 2, and where the function $f = (f_1, f_2) : \mathbb{R}^{I \times J} \rightarrow \mathbb{R}^{I \times J}$ is Lipschitz continuous and such that the state space $\Delta(I) \times \Delta(J)$ is positively invariant. Such dynamics may be innovative (initially absent pure strategies may appear). A definition is now needed:

Definition. (\mathbf{x}, \mathbf{y}) is a *selection equilibrium* if $x_i > 0 \Rightarrow (\mathbf{A}\mathbf{y})_i = \mathbf{x} \cdot \mathbf{A}\mathbf{y}$ and $y_j > 0 \Rightarrow (\mathbf{B}\mathbf{x})_j = \mathbf{y} \cdot \mathbf{B}\mathbf{x}$.³

The dynamics we consider are required to be adaptive in the sense that, for all $(\mathbf{x}, \mathbf{y}) \in \Delta(I) \times \Delta(J)$,

$$\dot{\mathbf{x}} \cdot \mathbf{A}\mathbf{y} \geq 0, \dot{\mathbf{y}} \cdot \mathbf{B}\mathbf{x} \geq 0 \quad (11.2.2)$$

²The class of myopic adjustment dynamics is itself a variant of Friedman's (1991) weak compatible dynamics.

³Thus, (\mathbf{x}, \mathbf{y}) is a selection equilibrium if and only if it is a Nash equilibrium or a Nash equilibrium of one of the "subgames" obtained by eliminating some of the pure strategies; the selection equilibria are exactly the rest-points of the replicator dynamics.

with at least one strict inequality whenever (\mathbf{x}, \mathbf{y}) is not a selection equilibrium. The condition $\dot{\mathbf{x}} \cdot \mathbf{A}\mathbf{y} \geq 0$ means that population 1 evolves towards (weakly) better replies to the current mean strategy of players 2.

Finally, and this is the main difference with (Swinkels, 1993), the dynamics is required to evolve smoothly with the game. That is, the functions f_1, f_2 in (11.2.1) are jointly continuous in $(\mathbf{x}, \mathbf{y}, \mathbf{A}, \mathbf{B})$.

Definition A dynamics (11.2.1) is a *myopic adjustment dynamics** (henceforth, MAD*) if it satisfies the above requirements.

An important subclass of dynamics satisfying (11.2.2) is the class of weakly sign preserving (WSP) dynamics (Ritzberger and Weibull, 1995). A dynamics is WSP if the proportion of each strategy earning less than average decreases:

$$(x_i > 0 \text{ and } (\mathbf{A}\mathbf{y})_i < \mathbf{x} \cdot \mathbf{A}\mathbf{y}) \Rightarrow \dot{x}_i < 0; \quad (y_j > 0 \text{ and } (\mathbf{B}\mathbf{x})_j < \mathbf{y} \cdot \mathbf{B}\mathbf{x}) \Rightarrow \dot{y}_j < 0$$

Before stating the results, more definitions are needed: let $I_N \subset I$ denote the set of pure strategies of player 1 belonging to the support of at least one Nash equilibrium:

$$I_N = \{i \in I, \exists NE(\mathbf{x}, \mathbf{y}), x_i > 0\}$$

(where “ $\exists NE$ ” means “there exists a Nash equilibrium”). Let $N(\mathbf{x}) = \sum_{i \in I_N} x_i$. Define $J_N \subset J$ and $N(\mathbf{y})$ similarly. Our main result is the following:

Theorem 11.2.1. *For any MAD* and every $\alpha > 0$, there exists an open set of (bimatrix) games for which, from an open set of initial conditions, $\limsup[N(\mathbf{x}(t)) + N(\mathbf{y}(t))] \leq \alpha$, (and $\lim[N(\mathbf{x}(t)) + N(\mathbf{y}(t))] = 0$ if the dynamics is WSP).*

The analogue of theorem 11.2.1 for single-population dynamics is also obtained: For symmetric games ($I = J$ and $\mathbf{A} = \mathbf{B}$), we may consider single-population dynamics of the form $\dot{\mathbf{x}} = f(\mathbf{x}, A)$, call MAD* the single-population dynamics satisfying the analogue of the above requirements, and let $N(\mathbf{x}) = \sum_{i \in I_N} x_i$ denote the proportion of strategies used in Nash equilibrium. We then have:

Theorem 11.2.2. *Within the set of two-player symmetric games, for any single-population MAD* and every $\alpha > 0$, there exists an open set of games for which, from an open set of initial conditions, $\limsup N(\mathbf{x}(t)) \leq \alpha$ (and $\lim N(\mathbf{x}(t)) = 0$ if the dynamics is WSP).*

11.3 Proof

The proof combines ideas of (Hofbauer and Sigmund, 1998, section 8.6) showing that cycling behaviour occurs in every MAD* with the method used in chapter 10 to build games in which all strategies used in equilibria are eliminated. Crucial to our proof is the fact that, in games with identical interests ($A = B^T$) and under any MAD, the mean payoff increases along the trajectory:

Lemma 11.3.1. *If $B = A^T$ then, for any MAD*, the mean-payoff $P(\mathbf{x}, \mathbf{y}) = \mathbf{x} \cdot \mathbf{A}\mathbf{y} = \mathbf{y} \cdot \mathbf{B}\mathbf{x}$ satisfies $\frac{d}{dt}P(\mathbf{x}(t), \mathbf{y}(t)) \geq 0$, with strict inequality whenever $(\mathbf{x}(t), \mathbf{y}(t))$ is not a selection equilibrium.*

Proof. Let $p(t) := \mathbf{x}(t) \cdot \mathbf{A}\mathbf{y}(t)$. We have: $\dot{p} = \dot{\mathbf{x}} \cdot \mathbf{A}\mathbf{y} + \mathbf{x} \cdot \mathbf{A}\dot{\mathbf{y}} = \dot{\mathbf{x}} \cdot \mathbf{A}\mathbf{y} + \dot{\mathbf{y}} \cdot \mathbf{A}^T \mathbf{x}$. Since $\mathbf{B} = \mathbf{A}^T$ it follows that $\dot{p} = \dot{\mathbf{x}} \cdot \mathbf{A}\mathbf{y} + \dot{\mathbf{y}} \cdot \mathbf{B}\mathbf{x}$. The result now follows from (11.2.2). \square

The sole use of condition (11.2.2) is to derive lemma 11.3.1. Thus, the above results also hold for any dynamics that does not satisfy (11.2.2) but for which lemma 11.3.1 holds.⁴

Consider the family of 6×6 symmetric games with payoff matrix:

$$\left(\begin{array}{ccccc|c} 0 & \epsilon & -1 & -1 & 0 & -2/5 \\ 0 & 0 & \epsilon & -1 & -1 & -2/5 \\ -1 & 0 & 0 & \epsilon & -1 & -2/5 \\ -1 & -1 & 0 & 0 & \epsilon & -2/5 \\ \epsilon & -1 & -1 & 0 & 0 & -2/5 \\ \hline -2/5 + \epsilon & -2/5 + \epsilon & -2/5 + \epsilon & -2/5 + \epsilon & -2/5 + \epsilon & -2/5 + \epsilon \end{array} \right) \quad (11.3.1)$$

Denote this game by G_ϵ and its payoff matrix (11.3.1) by \mathbf{A}_ϵ . Since the game is symmetric, \mathbf{A}_ϵ is the payoff matrix of both players. The 5×5 game obtained

⁴This would be the case for instance if, instead of requiring the nonnegativity of both $\dot{\mathbf{x}} \cdot \mathbf{A}\mathbf{y}$ and $\dot{\mathbf{y}} \cdot \mathbf{B}\mathbf{x}$, we required a form of ‘‘aggregate adaptiveness’’: that the sum $\dot{\mathbf{x}} \cdot \mathbf{A}\mathbf{y} + \dot{\mathbf{y}} \cdot \mathbf{B}\mathbf{x}$ be nonnegative (and positive if (\mathbf{x}, \mathbf{y}) is not a selection equilibrium).

by eliminating the sixth strategy is a variant of game (8.24) in (Hofbauer and Sigmund, 1998). For $\epsilon > 0$, it has a unique Nash equilibrium^{5,6}: the symmetric equilibrium in which both players play $(\frac{1}{5}, \frac{1}{5}, \frac{1}{5}, \frac{1}{5}, \frac{1}{5})$, with payoff $(-2 + \epsilon)/5$. Together with proposition 11.6.1 in the appendix, this implies that for $\epsilon > 0$, game (11.3.1) has a unique Nash equilibrium: the symmetric, strict Nash equilibrium (6,6). Note that there are no nonsymmetric Nash equilibria.

For $\epsilon = 0$, the game (11.3.1) is a game with identical interests. Let

$$P_0(\mathbf{x}, \mathbf{y}) = \mathbf{x} \cdot \mathbf{A}_0 \mathbf{y}$$

denote the mean payoff. The maximal value of P_0 is 0. Furthermore (proof omitted),

Lemma 11.3.2. *There are no selection equilibria (\mathbf{x}, \mathbf{y}) with $-\frac{2}{5} < P_0(\mathbf{x}, \mathbf{y}) < 0$.*

Together with lemma 11.3.1, this implies that if the mean payoff is initially greater than $-2/5$ then it increases monotonically to 0. Therefore, the solution $(\mathbf{x}(t), \mathbf{y}(t))$ converges to $\{(\mathbf{x}, \mathbf{y}) : P_0(\mathbf{x}, \mathbf{y}) = 0\}$. It follows that $x_6(t) + y_6(t) \rightarrow 0$. If we perturb the game, then $x_6(t) + y_6(t)$ need not go to zero, but it will still get very low from an open set of initial conditions: let $\|\mathbf{A}\| = \max_{i,j} |a_{ij}|$.

Proposition 11.3.3. *For every $\alpha > 0$, there exists $\bar{\epsilon} > 0$ such that, for all $(6 \times 6$ bimatrix) games with payoff matrices \mathbf{A}, \mathbf{B} such that $\|\mathbf{A} - \mathbf{A}_0\| < \bar{\epsilon}$, $\|\mathbf{B} - \mathbf{A}_0\| < \bar{\epsilon}$, and for every initial condition $(\mathbf{x}(0), \mathbf{y}(0))$ such that $P_0(\mathbf{x}(0), \mathbf{y}(0)) > -\frac{2}{5} + \alpha$,*

$$\limsup(x_6(t) + y_6(t)) \leq \alpha$$

and $\lim(x_6(t) + y_6(t)) = 0$ if the dynamics is WSP.

Proof. This follows from the fact that for games sufficiently close to G_0 , the function P_0 is still a Lyapunov function outside a neighbourhood of the

⁵In contrast, the game (8.24) in (Hofbauer and Sigmund, 1998) had a unique symmetric equilibrium but several Nash equilibria, including Pareto-efficient ones, which is why we need to modify it.

⁶We omit the proof, as we found no other than checking all the different possible supports.

selection equilibria of G_0 . More precisely, for $\alpha > 0$, let K_α denote the compact set:

$$K_\alpha = \left\{ (\mathbf{x}, \mathbf{y}) \in \Delta(I) \times \Delta(J) : -\frac{2}{5} + \alpha \leq P_0(\mathbf{x}, \mathbf{y}) \leq -\frac{\alpha}{10} \right\}$$

In the game G_0 , at every point in K_α , the mean payoff $P_0(\mathbf{x}, \mathbf{y})$ is strictly increasing:

$$\forall (\mathbf{x}, \mathbf{y}) \in K_\alpha, f_1(\mathbf{x}, \mathbf{y}, \mathbf{A}_0, \mathbf{A}_0) \cdot \mathbf{A}_0 \mathbf{y} + f_2(\mathbf{x}, \mathbf{y}, \mathbf{A}_0, \mathbf{A}_0) \cdot \mathbf{A}_0 \mathbf{x} > 0$$

By compactness of K_α and continuity of f_1 and f_2 , we actually have:

$$\exists \beta > 0, \forall (\mathbf{x}, \mathbf{y}) \in K_\alpha, f_1(\mathbf{x}, \mathbf{y}, \mathbf{A}_0, \mathbf{A}_0) \cdot \mathbf{A}_0 \mathbf{y} + f_2(\mathbf{x}, \mathbf{y}, \mathbf{A}_0, \mathbf{A}_0) \cdot \mathbf{A}_0 \mathbf{x} > \beta$$

Since f_1 and f_2 are jointly continuous⁷ in $(\mathbf{x}, \mathbf{y}, \mathbf{A}, \mathbf{B})$, it follows that there exists $\bar{\epsilon} > 0$ such that for all payoff matrices \mathbf{A}, \mathbf{B} with $\|\mathbf{A} - \mathbf{A}_0\| < \bar{\epsilon}$, $\|\mathbf{B} - \mathbf{A}_0\| < \bar{\epsilon}$,

$$\forall (\mathbf{x}, \mathbf{y}) \in K_\alpha, f_1(\mathbf{x}, \mathbf{y}, \mathbf{A}, \mathbf{B}) \cdot \mathbf{A}_0 \mathbf{y} + f_2(\mathbf{x}, \mathbf{y}, \mathbf{A}, \mathbf{B}) \cdot \mathbf{A}_0 \mathbf{x} > 0 \quad (11.3.2)$$

That is, as long as the solution is in K_α , $P_0(\mathbf{x}, \mathbf{y})$ strictly increases along solutions of the perturbed dynamics. This implies that for every initial condition (\mathbf{x}, \mathbf{y}) such that $P_0(\mathbf{x}, \mathbf{y}) > -2/5 + \alpha$, the solution $(\mathbf{x}(t), \mathbf{y}(t))$ enters the set $\{(\mathbf{x}, \mathbf{y}) : P_0(\mathbf{x}, \mathbf{y}) > -\alpha/10\}$ and stays there forever.

Fix payoff matrices \mathbf{A} and \mathbf{B} such that $\|\mathbf{A} - \mathbf{A}_0\| < \bar{\epsilon}$ and $\|\mathbf{B} - \mathbf{A}_0\| < \bar{\epsilon}$. We have:

$$P_0(\mathbf{x}, \mathbf{y}) \leq \bar{\epsilon} - \frac{1}{5}(x_6 + y_6)$$

Therefore, provided that $\bar{\epsilon} \leq \alpha/10$,

$$\left(P_0(\mathbf{x}, \mathbf{y}) \geq -\frac{\alpha}{10} \right) \Rightarrow (\mathbf{x}_6 + \mathbf{y}_6 \leq \alpha)$$

It follows that if $P_0(\mathbf{x}(0), \mathbf{y}(0)) > -2/5 + \alpha$ then $\limsup(x_6(t) + y_6(t)) \leq \alpha$.

Finally, note that if $P_0(\mathbf{x}, \mathbf{y}) > -\alpha/10$ then, (provided that α and $\bar{\epsilon}$ are chosen small enough), the sixth strategy of player $k \in \{1, 2\}$ earns strictly less than the mean payoff in population k . It follows that if the dynamics is WSP and $P_0(\mathbf{x}(0), \mathbf{y}(0)) > -2/5 + \alpha$ then $x_6(t) + y_6(t) \rightarrow 0$. \square

⁷This is the only use of the joint continuity of f_1 and f_2 .

We are now in a position to prove theorems 11.2.1 and 11.2.2:

Proof of theorem 11.2.1. Fix $\alpha > 0$. Recall that for $\epsilon > 0$, the game G_ϵ has a unique Nash equilibrium: $(6, 6)$. Since the set of bimatrix games with a unique Nash equilibrium is open (Jansen, 1981) and since $(6, 6)$ is a strict equilibrium, it follows that there is a neighborhood⁸ of G_ϵ in which every game has a unique Nash equilibrium: $(6, 6)$. Since every neighborhood of G_0 contains a neighborhood of some G_ϵ , it follows from proposition 11.3.3 that there exists an open set of games for which the unique Nash equilibrium is $(6, 6)$ but, from every initial condition (\mathbf{x}, \mathbf{y}) with $P_0(\mathbf{x}, \mathbf{y}) > -2/5 + \alpha$, $\limsup(x_6(t) + y_6(t)) \leq \alpha$ (and $\lim(x_6(t) + y_6(t)) = 0$ if the dynamics is WSP). ■

Proof of theorem 11.2.2: the orbits of the single-population dynamics $\dot{\mathbf{x}} = f(\mathbf{x}, \mathbf{A})$ correspond to the orbits of the two-population dynamics $(\dot{\mathbf{x}}, \dot{\mathbf{y}}) = (f(\mathbf{y}, \mathbf{A}), f(\mathbf{x}, \mathbf{A}))$ for symmetric initial conditions $(\mathbf{x}(0) = \mathbf{y}(0))$. Therefore, since the game G_ϵ is symmetric, it follows from the above proof that there exists an open set of symmetric games for which the unique Nash equilibrium is $(6, 6)$ but, from every initial condition \mathbf{x} with $P_0(\mathbf{x}, \mathbf{x}) = \mathbf{x} \cdot \mathbf{A}_0 \mathbf{x} > -2/5 + \alpha$, $\limsup x_6(t) \leq \alpha$ (and $\lim x_6(t) = 0$ if the dynamics is WSP). ■

11.4 Other examples

In the previous section, we considered a game with a relatively high number of strategies to prove theorem 11.2.1 and theorem 11.2.2 simultaneously. Here we provide three smaller dimensional examples: the first, based on game (8.24) of (Hofbauer and Sigmund, 1998), is adapted to single-population dynamics; the second and third ones are based on games designed by Josef Hofbauer and Jeroen Swinkels (personal communication). They are respectively adapted to two-population and three-population dynamics. All proofs are similar to the proof of theorem 11.2.1 and therefore omitted.

⁸A set of (6×6) games is a neighborhood of the game (A, B) if, for some $\alpha > 0$, this set contains every game (A', B') with $\|A - A'\| < \alpha$, $\|B - B'\| < \alpha$

11.4.1 Single-population dynamics

Consider the following symmetric two-player game G_ϵ :

$$\left(\begin{array}{cccc|c} 0 & 0 & -1 & \epsilon & -1/4 \\ \epsilon & 0 & 0 & -1 & -1/4 \\ -1 & \epsilon & 0 & 0 & -1/4 \\ 0 & -1 & \epsilon & 0 & -1/4 \\ \hline -1/4 + \epsilon & -1/4 + \epsilon & -1/4 + \epsilon & -1/4 + \epsilon & -1/4 + \epsilon \end{array} \right) \quad (11.4.1)$$

For $\epsilon > 0$, this game has a unique symmetric Nash equilibrium: the strict Nash equilibrium $(5, 5)$. This follows from proposition 11.6.1 in the appendix and from the fact that the underlying 4×4 game has a unique symmetric Nash equilibrium: $(\frac{1}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4})$, whose payoff $\frac{-1+\epsilon}{4}$ is less than $-\frac{1}{4} + \epsilon$ (see Hofbauer and Sigmund, 1998, section 8.6, for a more detailed presentation of this 4×4 game). The mean payoff in the game G_0 is

$$P_0(\mathbf{x}) := -2(x_1x_3 + x_2x_4) - (2x_6 - x_6^2)/4$$

Proposition 11.4.1. *For every single-population MAD* and every $\alpha > 0$, there exist $\epsilon > 0$ and a neighborhood of G_ϵ such that, for every game in this neighborhood, the unique symmetric Nash equilibrium is $(5, 5)$ but, for every initial condition \mathbf{x} with $P_0(\mathbf{x}) \geq -\frac{1}{4} + \alpha$, we have $\limsup x_5 \leq \alpha$, and $\lim x_5 = 0$ if the dynamics is WSP.*

11.4.2 Two-population dynamics

Consider the following two-player game G_ϵ :

$$\left(\begin{array}{ccc|c} -1, -1 & \epsilon, 0 & 0, \epsilon & -1/3, -1/3 + \epsilon \\ 0, \epsilon & -1, -1 & \epsilon, 0 & -1/3, -1/3 + \epsilon \\ \epsilon, 0 & 0, \epsilon & -1, -1 & -1/3, -1/3 + \epsilon \\ \hline -1/3 + \epsilon, -1/3 & -1/3 + \epsilon, -1/3 & -1/3 + \epsilon, -1/3 & -1/3 + \epsilon, -1/3 + \epsilon \end{array} \right) \quad (11.4.2)$$

For $\epsilon > 0$, this game has a unique Nash equilibrium: the strict Nash equilibrium $(4, 4)$. This follows from proposition 11.6.1 in the appendix and from the fact that the underlying 3×3 game has a unique Nash equilibrium: $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$, whose payoff $\frac{-1+\epsilon}{3}$ is less than $-\frac{1}{3} + \epsilon$.

The mean payoff in the game G_0 is

$$P_0(\mathbf{x}, \mathbf{y}) = -(x_1y_1 + x_2y_2 + x_3y_3) - \frac{1}{3}(x_6 + y_6 - x_6y_6)$$

Proposition 11.4.2. *For every MAD* and every $\alpha > 0$, there exists $\epsilon > 0$ and a neighborhood of G_ϵ such that, for every game in this neighborhood, the unique Nash equilibrium is (4, 4) but, for every initial condition (\mathbf{x}, \mathbf{y}) with $P_0(\mathbf{x}, \mathbf{y}) \leq -\frac{1}{3} + \alpha$, we have $\limsup(x_4 + y_4) \leq \alpha$, and $\lim(x_4 + y_4) = 0$ if the dynamics is WSP.*

Note that the game (11.4.2) is symmetric. However, for our purposes, it is not adapted to single-population dynamics as there is no mixed strategy \mathbf{x} such that $P_0(\mathbf{x}, \mathbf{x}) > -1/3$. Actually, in the game G_ϵ , under any sign-preserving dynamics (Ritzberger and Weibull, 1995), the unique Nash equilibrium (4, 4) is globally asymptotically stable. This is because the payoff of the fourth strategy is always greater than the mean payoff.

11.4.3 A three-player game

$$\begin{pmatrix} -1, -1, -1 & 0, 0, \epsilon \\ 0, \epsilon, 0 & \epsilon, 0, 0 \end{pmatrix} \quad \begin{pmatrix} \epsilon, 0, 0 & 0, \epsilon, 0 \\ 0, 0, \epsilon & -1, -1, -1 \end{pmatrix} \quad (11.4.3)$$

The $2 \times 2 \times 2$ game (11.4.3) has a unique Nash equilibrium, in which each player plays (1/2, 1/2), with payoff $-(1 + \epsilon)/4$. Consider the $3 \times 3 \times 3$ game G_ϵ with the following payoffs (for pure strategies): if every player plays one of his two first strategies, then the payoffs are as in (11.4.3); if one of the players plays his third strategy, then the payoff of each player is $(-1/4) + m\epsilon$ where m is the number of players playing strategy 3. Let $(\mathbf{x}, \mathbf{y}, \mathbf{z})$ denote the vector of mean strategies and $P_0(\mathbf{x}, \mathbf{y}, \mathbf{z})$ the mean payoff in G_0 . Define three-population MAD* by replacing (11.2.2) by its analogue for three-player games (see equation (1.1) in (Swinkels, 1993)).

Proposition 11.4.3. *For every MAD* and every $\alpha > 0$, there exists $\epsilon > 0$ and a neighborhood of the game G_ϵ such that, for every game in this neighborhood, the unique Nash equilibrium is (3, 3, 3) but for every initial condition $(\mathbf{x}, \mathbf{y}, \mathbf{z})$ with $P_0(\mathbf{x}, \mathbf{y}, \mathbf{z}) > -\frac{1}{4} + \alpha$, we have $\limsup(x_4 + y_4) \leq \alpha$, and $\lim(x_4 + y_4) = 0$ if the dynamics is WSP.⁹*

⁹The proof of theorem 11.2.1 relies on the openness of the set of bimatrix games with a unique Nash equilibrium. This is to be replaced here by the fact that the set of n -player games with a unique, strict Nash equilibrium is open. This follows from the upper semi-continuity of the Nash equilibrium correspondence, as shown in chapter 3, corollary 3.3.1.

11.5 Discussion

Chapter 10 dealt with elimination of all strategies used in *correlated* equilibria, while this chapter only studies elimination of all strategies used in *Nash* equilibria. From this point of view, the results of this chapter are weaker. On the other hand, the dynamics considered in this chapter are more general: In chapter 10, the results on monotonic or WSP dynamics is obtained under the additional assumption that initially absent strategies do not appear: an assumption which, as discussed by Swinkels (1993), is hardly satisfactory in an economical context. Here this assumption is relaxed. Furthermore, the results of this chapter are not restricted to single-population dynamics but hold for multi-population dynamics as well¹⁰.

Up to our knowledge the following problems remain open:

1. whether for *all* MAD* (or all WSP dynamics), there exists *some* games for which all strategies used in correlated equilibrium are (almost) eliminated along *some* interior solutions;
2. whether for *some* two-population MAD*, there exists *an open set* of games for which, from *an open set* of initial conditions, all strategies used in correlated equilibrium are eliminated.

11.6 Appendix: A class of games with a unique equilibrium

In this appendix, we prove a general result implying that the games G_ϵ considered in this chapter have a unique Nash equilibrium (symmetric Nash equilibrium, in the case of game (11.4.1)). Consider a n -player finite game $G = \{I, (S_i), (U_i)\}$: I is the set of players, S_i the set of pure strategies of player i and $U_i : S = \times_{k \in I} S_k \rightarrow \mathbb{R}$ its utility function. As usual, let $S_{-i} := \times_{j \neq i} S_j$.

For each player i , fix t_i in S_i and assume that $\hat{S}_i := S_i \setminus \{t_i\}$ is nonempty. Let $\hat{S} = \times_{i \in I} \hat{S}_i$ and let \hat{G} denote the game obtained by restricting the players to their strategies in \hat{S} . For instance, if G is the two-player game (11.3.1) and

¹⁰The results of chapter 10 show that, for a large class of two-population dynamics, there exist an open set of *symmetric* games for which, for an open set of *symmetric* initial conditions; but it does not show that this occurs on a open set of games nor from an open set of initial conditions.

if, for $i = 1, 2$, the strategy t_i is the sixth strategy, then \hat{G} is the underlying 5×5 game.

Proposition 11.6.1. (i) *Assume that no Nash equilibrium of \hat{G} induces a Nash equilibrium of G , and that if at least one player j plays t_j (with probability 1), then every player $i \neq j$ is indifferent between his strategies in \hat{S}_i and has a strict incentive to play t_i . Formally, letting $NE(\hat{G})$ denote the set of Nash equilibria of \hat{G} :*

$$\forall \hat{\sigma} \in NE(\hat{G}), \exists i \in I, U_i(t_i, \hat{\sigma}_{-i}) > U_i(\sigma) \quad (11.6.1)$$

$$\forall i, \forall s_{-i} \in S_{-i} \setminus \hat{S}_{-i}, \forall (s_i, s'_i) \in \hat{S}_i \times \hat{S}_i, U_i(s) = U_i(s'_i, s_{-i}) < U_i(t_i, s_{-i}) \quad (11.6.2)$$

Then G has a unique Nash equilibrium: the strict equilibrium $t = (t_i)_{i \in I}$.

(ii) *For symmetric games, the same result holds if we replace everywhere Nash equilibrium by symmetric Nash equilibrium.*

Proof. We prove (i). The proof of (ii) is similar. First, the fact that t is a strict Nash equilibrium follows from (11.6.2). Second, let σ be a Nash equilibrium of G . If for some j , $\sigma_j = t_j$, then it follows from (11.6.2) and from the fact that σ is a Nash equilibrium that $\sigma = t$. Otherwise, for every i , $K_i := \sum_{s_i \in \hat{S}_i} \sigma_i(s_i)$ is positive. Let $\hat{\sigma}_i$ be the mixed strategy of player i in the game \hat{G} defined by

$$\forall s_i \in \hat{S}_i, \hat{\sigma}_i(s_i) = \sigma_i(s_i) / K_i$$

Let $K_{-i} := \prod_{j \neq i} K_j$. Fix a couple of pure strategies (s_i, s'_i) in $\hat{S}_i \times \hat{S}_i$. We have:

$$U_i(s_i, \hat{\sigma}_{-i}) - U_i(s'_i, \hat{\sigma}_{-i}) = \sum_{s_{-i} \in \hat{S}_{-i}} \frac{\sigma_{-i}(s_{-i})}{K_{-i}} [U_i(s) - U_i(s'_i, s_{-i})] \quad (11.6.3)$$

Furthermore, by (11.6.2), $U_i(s) = U_i(s'_i, s_{-i})$ for all $s_{-i} \notin \hat{S}_{-i}$. It follows that

$$U_i(s_i, \sigma_{-i}) - U_i(s'_i, \sigma_{-i}) = \sum_{s_{-i} \in \hat{S}_{-i}} \sigma_{-i}(s_{-i}) [U_i(s) - U_i(s'_i, s_{-i})] \quad (11.6.4)$$

If s_i belongs to the support of $\hat{\sigma}_i$, then it belongs to the support of σ_i so that $U_i(s_i, \sigma_{-i}) - U_i(s'_i, \sigma_{-i}) \geq 0$ (because σ is a Nash equilibrium). Together with (11.6.3) and (11.6.4), this implies that if $\hat{\sigma}_i(s_i) > 0$ then

$$U_i(s_i, \hat{\sigma}_{-i}) - U_i(s'_i, \hat{\sigma}_{-i}) = K_{-i} [U_i(s_i, \sigma_{-i}) - U_i(s'_i, \sigma_{-i})] \geq 0$$

It follows that $\hat{\sigma} := (\hat{\sigma}_i)_{i \in I}$ is a Nash equilibrium of \hat{G} . Therefore, by (11.6.1), there exists a player i such that $U_i(t_i, \hat{\sigma}_{-i}) > U_i(\hat{\sigma})$ hence

$$U_i(t_i, \hat{\sigma}_{-i}) > U_i(\sigma_i, \hat{\sigma}_{-i}) \quad (11.6.5)$$

Let

$$C := \sum_{s_i \in S_i} \sum_{s_{-i} \notin \hat{S}_{-i}} \sigma(s) [U_i(t_i, s_{-i}) - U_i(s)]$$

Note that C is nonnegative by (11.6.2). Therefore,

$$U_i(t_i, \sigma_{-i}) - U_i(\sigma) = K_{-i} [U_i(t_i, \hat{\sigma}_{-i}) - U_i(\sigma_i, \hat{\sigma}_{-i})] + C > C \geq 0$$

where the strict inequality follows from (11.6.5). This implies $U_i(t_i, \sigma_{-i}) - U_i(\sigma) > 0$, contradicting the assumption that σ is a Nash equilibrium. This completes the proof. \square

Chapitre 12

Elimination of All Strategies in the Support of Nash Equilibria From Almost All Initial Conditions

Abstract

We show that under the replicator dynamics and the best-response dynamics, all strategies in the support of at least one Nash equilibrium may be eliminated from almost all initial conditions. For the best-response dynamics, this holds for an open set of games.

12.1 Introduction

In the previous chapters, examples were given of games for which, for many dynamics and from an open set of initial conditions, all strategies in the support of the unique Nash equilibrium are eliminated. However, in these examples, the Nash equilibrium is strict and thus asymptotically stable under most dynamics.

This leads to the following question: are there examples of games for which all strategies played in Nash equilibrium distributions are eliminated *for almost all initial conditions* ? This chapter shows that the answer is positive, at least for the single-population replicator dynamics (REP) and the best response dynamics (BR). For BR, we even exhibit *an open set* of

games for which all strategies played in Nash equilibrium are eliminated from almost all initial conditions.

Our examples are relatively high dimensional: 6×6 for BR and 7×7 for REP. The reason why, for the replicator dynamics, we need an extra-dimension, and study only specific games instead of an open set of games, seems purely technical: we think that our examples for the best-response dynamics work as well for the replicator dynamics, but this does not seem so easy to prove, as the replicator dynamics is more difficult to “control” than the best-response dynamics.

The reason why our games are relatively high dimensional is deeper: though we have no formal proof of this, it seems that for the single-population replicator (or best-response) dynamics, examples of games for which all strategies played in Nash equilibrium are eliminated from almost all initial conditions cannot be found in games with less than 5 strategies.

The reason is twofold: first, by the folk-theorem of evolutionary game theory (see, e.g., Weibull, 1995, prop. 4.11), if an interior trajectory of REP of BR converges to a point, then this point is a Nash equilibrium. Thus, we need nonconvergent trajectories, and along which, asymptotically, only strategies that do not belong to the support of a Nash equilibrium have positive probability. This seems to require at least three strategies not played in Nash equilibrium. Second, the only solution for having a unique strategy played in Nash equilibrium is to have a unique, pure Nash equilibrium. But such a Nash equilibrium would be strict. Indeed, as already mentioned in chapter 10:

Proposition 12.1.1. *In a bimatrix game, a unique and pure Nash equilibrium is strict.*

Proof. Every bimatrix game has a quasi-strict Nash equilibrium¹ (Norde, 1999). Therefore, if a Nash equilibrium is unique, it is quasi-strict; if it is unique and pure, it is quasi-strict and pure, hence strict. \square

We thus need at least two strategies played in Nash equilibrium, which together with the three strategies not played in Nash equilibrium which seems to be needed, make at least five strategies. Our examples for the

¹With standard notations, a Nash equilibrium σ is quasi-strict if for every player i , the pure strategies in the support of σ_i are the only pure best-responses to σ_{-i} .

best-response dynamics are 6×6 games; thus, there might be room for improvement, but not much.²

The main results are proved in section 12.2 for the best-response dynamics and in section 12.3 for the replicator dynamics. Section 12.4 concludes.

Framework and notations. We study single-population dynamics in two-player, finite symmetric games. The set of pure strategies is $I = \{1, 2, \dots, N\}$ and the payoff matrix is $\mathbf{U} = (u_{ij})_{1 \leq i, j \leq N}$. We let S_N denotes the simplex of mixed strategies. Its vertices \mathbf{e}_i , $1 \leq i \leq N$, correspond to the pure strategies of the game. The mean strategy at time t is denoted $\mathbf{x}(t) = (x_i(t))_{1 \leq i \leq N}$.

12.2 Best-response dynamics

The best-response dynamics (Gilboa and Matsui, 1991) is given by the differential inclusion:

$$\dot{\mathbf{x}}(t) \in BR(\mathbf{x}(t)) - \mathbf{x}(t) \quad (12.2.1)$$

where $BR(\mathbf{x})$ is the set of best-responses to \mathbf{x} :

$$BR(\mathbf{x}) = \{\mathbf{y} \in S_N : \mathbf{y} \cdot \mathbf{U}\mathbf{x} = \max_{\mathbf{z} \in S_N} \mathbf{z} \cdot \mathbf{U}\mathbf{x}\}$$

A solution $\mathbf{x}(\cdot)$ of the best-response dynamics is an absolutely continuous function satisfying (12.2.1) for almost every t . In the games considered below, there might be several solutions with the same initial condition. This explains the way our first result is phrased:

Theorem 12.2.1. *There exists an open set of 6×6 bimatrix games with a unique Nash equilibrium and such that, for almost every initial condition \mathbf{x} in S_6 , there exists a solution $\mathbf{x}(t)$ of the best-response dynamics with $\mathbf{x}(0) = \mathbf{x}$, along which all strategies in the support of the Nash equilibrium are eliminated.*

²It might be possible to design an example whose restriction to the set of strategies played in Nash equilibrium is a 2×2 coordination game, and whose unique Nash equilibrium corresponds to the mixed, unstable equilibrium of this coordination game.

12.2.1 Intuition and steps of the proofs

The games we study are based on outward cycling Rock-Paper-Scissors (RPS) games, i.e. 3×3 symmetric games

$$\begin{pmatrix} a_1 & b_2 & c_3 \\ c_1 & a_2 & b_3 \\ b_1 & c_2 & a_3 \end{pmatrix} \quad (12.2.2)$$

satisfying

$$b_i < a_i < c_i \text{ for } i = 1, 2, 3 \quad \text{and} \quad \prod_{i=1}^3 (a_i - b_i) > \prod_{i=1}^3 (c_i - a_i) \quad (12.2.3)$$

As already mentioned in chapter 10, these games have a unique, completely mixed Nash equilibrium. Furthermore, for every initial condition \mathbf{x} different from the Nash equilibrium, there is a unique solution $\mathbf{x}(t)$ to (12.2.1) with initial condition \mathbf{x} and it converges to the so called Shapley triangle:

$$ST = \{\mathbf{x} \in S_3 : \left[\max_{1 \leq i \leq 3} (\mathbf{U}\mathbf{x})_i \right] - \sum_{1 \leq i \leq 3} a_i x_i = 0\}$$

See (Gaunersdorfer and Hofbauer, 1995).

Consider the following 6×6 symmetric game G :

$$\left(\begin{array}{ccc|ccc} 0 & -3 & 1 & -1 & -1 & -1 \\ 1 & 0 & -3 & -1 & -1 & -1 \\ -3 & 1 & 0 & -1 & -1 & -1 \\ \hline -4 & -4 & 3 & 0 & -5 & 1 \\ -1 & -1 & -3 & 1 & 0 & -5 \\ -1 & -1 & -3 & -5 & 1 & 0 \end{array} \right) \quad (12.2.4)$$

Let G_{123} (resp. G_{456}) denote the 3×3 game obtained from G by restricting the players to their three first (resp. last) strategies. Both G_{123} and G_{456} are outward cycling RPS games with cyclic symmetry. Their unique Nash equilibrium correspond in the whole game to:

$$\mathbf{n}_{123} = \left(\frac{1}{3}, \frac{1}{3}, \frac{1}{3}, 0, 0, 0 \right)$$

for G_{123} and to

$$\mathbf{n}_{456} = \left(0, 0, 0, \frac{1}{3}, \frac{1}{3}, \frac{1}{3} \right)$$

for G_{456} . The payoffs are chosen so that $(\mathbf{n}_{123}, \mathbf{n}_{123})$ be a Nash equilibrium of (12.2.4) but not $(\mathbf{n}_{456}, \mathbf{n}_{456})$. Actually,

Proposition 12.2.2. *The game (12.2.4) has a unique Nash equilibrium: $(\mathbf{n}_{123}, \mathbf{n}_{123})$.*

Proof. See next section. □

Proposition 12.2.2 does not only state that $(\mathbf{n}_{123}, \mathbf{n}_{123})$ is the unique symmetric Nash equilibrium, but also that there are no asymmetric Nash equilibria. Nevertheless, we will show that from almost all initial conditions, strategies 1, 2 and 3 are eliminated. More precisely, let ST_{456} denote the Shapley triangle

$$ST_{456} = \left\{ \mathbf{x} : x_4 + x_5 + x_6 = 1 \text{ and } \max_{4 \leq i \leq 6} (\mathbf{U}\mathbf{x})_i = 0 \right\} \quad (12.2.5)$$

Proposition 12.2.3. *For every mixed strategy \mathbf{x} in S_6 to which there is a unique best-response, there exists a solution $\mathbf{x}(t)$ of the best response-dynamics (12.2.1) with initial condition $\mathbf{x}(0) = \mathbf{x}$ which converges to the Shapley triangle ST_{456} ; in particular, $x_1(t) + x_2(t) + x_3(t) \rightarrow 0$. Furthermore, if there is not time T such that strategies 1, 3 and 4 are all best-responses to $\mathbf{x}(T)$, then this is the unique solution to (12.2.1) with initial condition \mathbf{x} .*

This will be proved in section 12.2.3. The intuition is that, while the best-response cycle of G_{456} : $4 \rightarrow 5 \rightarrow 6 \rightarrow 4$ is still a best-response cycle in the whole game, the best-response cycle of G_{123} : $1 \rightarrow 2 \rightarrow 3 \rightarrow 1$ is perturbed by the fact that, in the whole game, the unique best-response to strategy 3 is strategy 4. More precisely, let

$$ST_{123} = \left\{ \mathbf{x} : x_1 + x_2 + x_3 = 1 \text{ and } \max_{1 \leq i \leq 3} (\mathbf{U}\mathbf{x})_i = 0 \right\} \quad (12.2.6)$$

denote the Shapley triangle of the RPS game G_{123} . Let \mathbf{q} denote the vertex of this triangle at which strategies 1 and 3 earn the same payoff, i.e.

$$\mathbf{q} = (1/13, 3/13, 9/13, 0, 0, 0)$$

The reason why the Shapley triangle ST_{123} does not attract any trajectory is that, in the neighborhood of \mathbf{q} , the unique best-response is strategy 4.

Taken together, propositions 12.2.2 and 12.2.3 show that in the game with payoffs (12.2.4), from almost every initial condition, there exists a solution of the best-response dynamics along which all strategies used in Nash equilibrium are eliminated. To prove theorem 12.2.1, it only remains to show that this result is robust to perturbations of the game:

Proposition 12.2.4. *There exists a neighborhood of (12.2.4) in which every game has a unique Nash equilibrium and its support is $\{1, 2, 3\} \times \{1, 2, 3\}$.*

Proof. This follows directly from proposition 12.2.2 and lemma 3.2.3 in chapter 3. \square

Proposition 12.2.5. *There exists a neighborhood of the game (12.2.4) such that proposition 12.2.3 holds for every game in this neighborhood, up to replacement of $(\mathbf{U}\mathbf{x})_i$ by $(\mathbf{U}\mathbf{x})_i - \sum_{i \in I} u_{ii}x_i$ in the definition of the Shapley triangle (12.2.5).*

Proof. See section 12.2.3. \square

12.2.2 Proof of proposition 12.2.2

We begin with a general result related to proposition 11.6.1 of chapter 11 and which will be used several times: let $I' \subset I$. For \mathbf{x} in S_N , define $\mathbf{x}' \in \mathbb{R}^N$ by $x'_i = x_i$ if $i \in I'$ and $x'_i = 0$ otherwise. Let $\mathbf{x}'' = \mathbf{x} - \mathbf{x}'$. Finally, let $x(I') = \sum_{i \in I'} x_i$. Define \mathbf{y}' , \mathbf{y}'' and $y(I')$ similarly.

Lemma 12.2.6. *Let (\mathbf{x}, \mathbf{y}) be a Nash equilibrium. If $x(I')y(I') > 0$ and if*

$$\forall i \in I', \forall j \in I', [(\mathbf{U}\mathbf{y}'')_i = (\mathbf{U}\mathbf{y}'')_j \text{ and } (\mathbf{U}\mathbf{x}'')_i = (\mathbf{U}\mathbf{x}'')_j] \quad (12.2.7)$$

then $(\mathbf{x}', \mathbf{y}')$ induces an (unnormalized) Nash equilibrium of the game restricted to $I' \times I'$.

Proof. Let $i \in I'$. For every j in I' , by (12.2.7), $(\mathbf{U}\mathbf{y})_i - (\mathbf{U}\mathbf{y})_j = (\mathbf{U}\mathbf{y}')_i - (\mathbf{U}\mathbf{y}')_j$. If $x'_i > 0$, then $x_i > 0$, hence, since \mathbf{x} is a best-response to \mathbf{y} , $(\mathbf{U}\mathbf{y})_i \geq (\mathbf{U}\mathbf{y})_j$ for every j in I ; therefore, for every j in I' , $(\mathbf{U}\mathbf{y}')_i \geq (\mathbf{U}\mathbf{y}')_j$. Similarly, if $y'_i > 0$, then for every j in I' , $(\mathbf{U}\mathbf{x}')_i \geq (\mathbf{U}\mathbf{x}')_j$. Therefore, $(\mathbf{x}', \mathbf{y}')$ induces an (unnormalized) Nash equilibrium of the game restricted to $I' \times I'$. \square

We are now ready to prove proposition 12.2.2: let (\mathbf{x}, \mathbf{y}) be a Nash equilibrium of (12.2.4) and note that:

(i) \mathbf{n}_{456} is strictly dominated by \mathbf{n}_{123}

(ii) $x_4x_5x_6 = 0$ and by symmetry $y_4y_5y_6 = 0$.

Indeed, if $x_4x_5x_6 > 0$, then strategies 4, 5 and 6 are all best-responses to \mathbf{y} , hence so is \mathbf{n}_{456} . But this cannot be due to (i).

(iii) If $y_1 = y_2 = y_3 = 0$, then $x_1 + x_2 + x_3 > 0$

Indeed, if for every i in $\{1, 2, 3\}$, $x_i = y_i = 0$, then \mathbf{x} and \mathbf{y} have support in $\{4, 5, 6\}$, hence must induce a Nash equilibrium of the game restricted to $\{4, 5, 6\} \times \{4, 5, 6\}$. This implies that $\mathbf{x} = \mathbf{y} = \mathbf{n}_{456}$, which contradicts (ii).

(iv) $y_1 + y_2 + y_3 > 0$ and by symmetry $x_1 + x_2 + x_3 > 0$

Assume by contradiction that $y_1 = y_2 = y_3 = 0$. It follows that

$$\forall i \in \{1, 2, 3\}, (\mathbf{U}\mathbf{y})_i = -1 < 0 \quad (12.2.8)$$

Furthermore, due to (ii), \mathbf{y} has support in $\{4, 5\}$, $\{5, 6\}$ or $\{4, 6\}$. In any case, there exists i in $\{4, 5, 6\}$ such that $(\mathbf{U}\mathbf{y})_i \geq 0$. Together with (12.2.8), this implies that strategies 1, 2 and 3 are not best-responses to \mathbf{y} , hence $x_1 = x_2 = x_3 = 0$. But since $y_1 = y_2 = y_3 = 0$, this contradicts (iii).

(v) $x_1 = x_2 = x_3$ and $y_1 = y_2 = y_3$.

Let $\mathbf{x}' = (x_1, x_2, x_3, 0, 0, 0)$ and $\mathbf{x}'' = \mathbf{x} - \mathbf{x}' = (0, 0, 0, x_4, x_5, x_6)$. Define \mathbf{y}' and \mathbf{y}'' symmetrically. For every i and j in $\{1, 2, 3\}$, we have $(\mathbf{U}\mathbf{x}'')_i = (\mathbf{U}\mathbf{x}'')_j$ and $(\mathbf{U}\mathbf{y}'')_i = (\mathbf{U}\mathbf{y}'')_j$. Therefore, it follows from (iv) and from lemma 12.2.6 that $(\mathbf{x}', \mathbf{y}')$ is an unnormalized Nash equilibrium of the game restricted to $\{1, 2, 3\} \times \{1, 2, 3\}$. Therefore \mathbf{x} and \mathbf{y} are both proportional to \mathbf{n}_{123} . The result follows.

(vi) If $x_4 + x_5 + x_6 > 0$ then $y_4 + y_5 + y_6 = 0$.

Let $\mathbf{x}' = (0, 0, 0, x_4, x_5, x_6)$ and $\mathbf{x}'' = \mathbf{x} - \mathbf{x}' = (x_1, x_2, x_3, 0, 0, 0)$ (in order to be consistent with the notations of lemma 12.2.6, the vector denoted \mathbf{x}'' (resp. \mathbf{x}') in the proof of (v) is now denoted \mathbf{x}' (resp. \mathbf{x}'')). Similarly, let $\mathbf{y}' = (0, 0, 0, y_4, y_5, y_6)$ and $\mathbf{y}'' = \mathbf{y} - \mathbf{y}'$. Against \mathbf{n}_{123} , every strategy i in $\{4, 5, 6\}$ earns the same payoff: $-5/3$. Therefore it follows from (v) that for every i and j in $\{4, 5, 6\}$, we have $(\mathbf{U}\mathbf{x}'')_i = (\mathbf{U}\mathbf{x}'')_j$ and $(\mathbf{U}\mathbf{y}'')_i = (\mathbf{U}\mathbf{y}'')_j$. Therefore, by lemma 12.2.6, if $x_4 + x_5 + x_6 > 0$ and $y_4 + y_5 + y_6 > 0$ then \mathbf{x}' and \mathbf{y}' are proportional to \mathbf{n}_{456} , hence $x_4x_5x_6 > 0$, which contradicts (ii).

(vii) $x_4 + x_5 + x_6 = 0$

Otherwise, by (vi), $y_4 + y_5 + y_6 = 0$. But then, by (v), $\mathbf{y} = \mathbf{n}_{123}$. Therefore,

strategies 4, 5 and 6 are not best-responses to \mathbf{y} , hence $x_4 = x_5 = x_6 = 0$.

We now conclude: it follows from (vii) that $x_4 = x_5 = x_6 = 0$ and, by symmetry, $y_4 = y_5 = y_6 = 0$. Therefore, by (v), $\mathbf{x} = \mathbf{y} = \mathbf{n}_{123}$. It follows that the only Nash equilibrium of (12.2.4) is $(\mathbf{n}_{123}, \mathbf{n}_{123})$. ■

12.2.3 Proof of propositions 12.2.3 and 12.2.5

It suffices to prove proposition 12.2.5 as it is more general than proposition 12.2.3. The proof relies on the following properties, which hold for any game in a sufficiently small neighborhood of the game with payoffs (12.2.4):

- (a) for every $i \in \{1, 2, 4, 5\}$ and every $j \notin \{i, i+1\}$, $u_{ji} < u_{ii} < u_{i+1,i}$; for every $j \notin \{6, 4\}$, $u_{j6} < u_{66} < u_{46}$; for every $j \notin \{1, 3, 4\}$, $u_{j3} < u_{33} < u_{13} < u_{43}$
- (b) For any i in I and any j such that $u_{ji} > u_{ii}$, in the game restricted to $\{i, j\} \times \{i, j\}$, strategy i is strictly dominated.³
- (c) The Rock-Paper-Scissors games on $\{1, 2, 3\} \times \{1, 2, 3\}$ and on $\{4, 5, 6\} \times \{4, 5, 6\}$ are outward cycling.
- (d) Let ST_{123} denote the Shapley triangle

$$ST_{123} = \left\{ \mathbf{x} : x_1 + x_2 + x_3 = 1 \text{ and } \max_{1 \leq i \leq 3} [(\mathbf{U}\mathbf{x})_i] - \sum_{i \in I} u_{ii}x_i = 0 \right\} \quad (12.2.9)$$

denote the Shapley triangle of the RPS game G_{123} . Let \mathbf{q} denote the vertex of this triangle at which strategies 1 and 3 earn the same payoff:

$$q_1 + q_2 + q_3 = 1 \text{ and } (\mathbf{U}\mathbf{q})_1 = (\mathbf{U}\mathbf{q})_3 = \sum_{i \in I} u_{ii}q_i \quad (12.2.10)$$

The unique best-reply to \mathbf{q} is strategy 4.

Finally, we will use the following version of the improvement principle (Monderer and Sela, 1997):

Lemma 12.2.7 (Improvement principle). *Assume that $BR(\mathbf{x}(T)) = \{\mathbf{e}_i\}$, and let T' be the first time greater than T such that strategy i is not the only best-response to $\mathbf{x}(t)$; if strategy j is a best-reply to $\mathbf{x}(T')$ then $u_{ji} > u_{ii}$.*

Proof. See chapter 10, lemma 10.9.2. □

³This is actually a consequence of (a).

We are now in a position to prove proposition 12.2.5. We proved much as in chapter 10, proof of proposition 10.9.1. Let $PBR(\mathbf{x}) := \{i \in I, \mathbf{e}_i \in BR(\mathbf{x})\}$ denote the set of pure best-responses to \mathbf{x} . Recall that we only consider initial conditions to which there is a unique best-reply. There are two cases:

Case 1: if $PBR(\mathbf{x}(0)) \in \{4, 5, 6\}$. Let us assume for concreteness that $PBR(\mathbf{x}(0)) = \{4\}$. It follows that $\mathbf{x}(t)$ must point towards \mathbf{e}_4 for some time. Let T be the first positive time at which strategy 4 is not the unique best-response (T is finite as strategy 4 is not a best-response to itself). Due to the improvement principle and to property (a), the pure best-replies to $\mathbf{x}(T)$ must be strategies 4 and 5. Thus, the solution will then point towards the edge $[\mathbf{e}_4, \mathbf{e}_5]$. But, by property (b), strategy 4 is strictly dominated by strategy 5 in the game restricted to $\{4, 5\} \times \{4, 5\}$. It follows that, immediately after time T , $\mathbf{x}(t)$ must point towards \mathbf{e}_5 , so that strategy 5 becomes the unique best-reply.⁴

Iterating the argument, we see that the solution will point towards \mathbf{e}_5 till 6 becomes a best-response, then it will point towards \mathbf{e}_6 , till 4 becomes a best-response, then it will point towards \mathbf{e}_4 , so that 4 becomes the unique best-response again, etc. To make sure that this cycling movement goes on for ever, we need to check that the times when the direction of $\mathbf{x}(t)$ change do not accumulate. This can be done exactly as in chapter 10, part 2, proof of proposition 10.9.1; the proof relies on the fact that the RPS game on $\{4, 5, 6\} \times \{4, 5, 6\}$ is outward cycling.

It follows from the above description of the dynamics that, if $PBR(\mathbf{x}(0)) \in \{4, 5, 6\}$, then there is a unique solution to (12.2.1) with initial condition $\mathbf{x}(0)$ and that strategies 1, 2 and 3 never become best-responses; therefore, x_1 , x_2 and x_3 decrease exponentially to zero. Furthermore, for almost all $t \geq 0$, there exists i in $\{4, 5, 6\}$ such that $BR(\mathbf{x}(t)) = \{\mathbf{e}_i\}$ and $\dot{\mathbf{x}}(t) = \mathbf{e}_i - \mathbf{x}(t)$. The function

$$v(t) := \max_{4 \leq i \leq 6} \left[(\mathbf{U}\mathbf{x}(t))_i - \sum_{i \in I} u_{ii}x_i \right]$$

then satisfies

$$\dot{v} = -v$$

(the computation is detailed in chapter 10, equation (10.9.4))

⁴Here and in what follows, we say that a property holds immediately after (before) some time T if it holds on an interval $]T, T + \tau[$ ($]T - \tau, T[$) with $\tau > 0$.

It follows that $v(t)$ converges to zero, hence that $\mathbf{x}(t)$ converges to the Shapley triangle ST_{456} .

Case 2: if $PBR(\mathbf{x}(0)) \in \{1, 2, 3\}$. As long as none of the strategies 4, 5 and 6 becomes a best-response, the dynamics may be described as above. That is, assuming for concreteness that $PBR(\mathbf{x}(0)) = \{1\}$, the solution will first points towards \mathbf{e}_1 , then towards \mathbf{e}_2 , then towards \mathbf{e}_3 , etc. If none of the strategies 4, 5 and 6 ever becomes a best-response, then $\mathbf{x}(t)$ converges to the Shapley triangle ST_{123} and visits arbitrarily small neighborhood of the vertex \mathbf{q} defined in the statement of property (d). This cannot be, because, by property (d), in the neighborhood of \mathbf{q} the unique best-reply is strategy 4, hence 4 would become a best-response (contradicting the assumption that none of the strategies 4, 5 and 6 ever becomes a best-response).

Therefore, there exists a time at which one of the strategies 4, 5 or 6 becomes a best-response. Let T be the smallest such time. Due to the improvement principle and to the behavior of the dynamics before time T , it must be that immediately before time T the unique best-response is strategy 3, and that the pure best-response to $\mathbf{x}(T)$ are either (subcase 2.1) strategies 3 and 4, or (subcase 2.2) strategies 1, 3 and 4.

In the first subcase, since in the game restricted to $\{3, 4\} \times \{3, 4\}$, strategy 3 is strictly dominated, it follows that immediately after time T the solution points towards \mathbf{e}_4 so that strategy 4 becomes the unique best-response. We are then back to case 1. In the second subcase, the solution to (12.2.1) ceases to be uniquely defined; however, since \mathbf{e}_4 induces a Nash equilibrium of the game restricted to $\{1, 3, 4\} \times \{1, 3, 4\}$, it follows that there exists a solution to (12.2.1) which points to \mathbf{e}_4 immediately after time T , so that strategy 4 becomes the unique best-response. For this solution, we are back to case 1 and thus this solution converges to ST_{567} . ■

12.3 Replicator dynamics

We showed in the previous section that, for the best-response dynamics, there exists games for which, from almost all initial conditions, all strategies used in Nash equilibrium are eliminated. We now show that this is also true of the replicator dynamics:

$$\dot{x}_i(t) = x_i(t) [(U\mathbf{x}(t))_i - \mathbf{x}(t) \cdot U\mathbf{x}(t)] \quad (12.3.1)$$

12.3.1 Example and statement of the result

Consider for instance the following 7×7 symmetric game:

$$\left(\begin{array}{ccc|ccc} 0 & -1 & \epsilon & -10 & -1/3 + \epsilon & -1/3 + \epsilon & -1/3 + \epsilon \\ \epsilon & 0 & -1 & -10 & -1/3 + \epsilon & -1/3 + \epsilon & -1/3 + \epsilon \\ -1 & \epsilon & 0 & -10 & -1/3 + \epsilon & -1/3 + \epsilon & -1/3 + \epsilon \\ \hline -2 & -2 & 2 & 0 & -1/3 & -1/3 & -1/3 \\ -1/3 & -1/3 & -1/3 & 10 & 0 & -1 & \epsilon \\ -1/3 & -1/3 & -1/3 & 10 & \epsilon & 0 & -1 \\ -1/3 & -1/3 & -1/3 & 10 & -1 & \epsilon & 0 \end{array} \right) \quad (12.3.2)$$

with $0 < \epsilon < 1/48$. This game is similar to the game with payoffs (12.2.4). In particular, both the three first and the three last strategies form, when alone, a Rock-Paper-Scissors game with cyclic symmetry. Each of these Rock-Paper-Scissors games has a unique Nash equilibrium. The corresponding rest-points of the replicator dynamics are:

$$\mathbf{n}_{123} = \left(\frac{1}{3}, \frac{1}{3}, \frac{1}{3}, 0, 0, 0, 0 \right)$$

and

$$\mathbf{n}_{567} = \left(0, 0, 0, 0, \frac{1}{3}, \frac{1}{3}, \frac{1}{3} \right)$$

The payoffs are chosen so that $(\mathbf{n}_{123}, \mathbf{n}_{123})$ be a Nash equilibrium of (12.3.2). Actually,

Proposition 12.3.1. *$(\mathbf{n}_{123}, \mathbf{n}_{123})$ is the unique Nash equilibrium of the game with payoffs (12.3.2)⁵*

Proof. See next section. □

However, we will show that for almost all initial conditions, strategies 1, 2 and 3 are eliminated. The intuition is as follows: the behavior of the replicator dynamics in RPS games is well known (see, e.g., Hofbauer and Sigmund, 1998). In particular, for any initial condition $\mathbf{x} \in S_7$ such that $x_1 + x_2 + x_3 = 1$ and $(x_1, x_2, x_3) \neq (1/3, 1/3, 1/3)$, $\mathbf{x}(t)$ converges to

$$\Gamma_{123} := \{ \mathbf{x} \in S_7, x_1 + x_2 + x_3 = 1 \text{ and } x_1 x_2 x_3 = 0 \} \quad (12.3.3)$$

⁵There are no asymmetric Nash equilibria.

However, in the whole game, Γ_{123} is not asymptotically stable because strategy 4 destabilizes it in the neighbourhood of \mathbf{e}_3 .⁶ In contrast, in the neighborhood of the vertices \mathbf{e}_5 , \mathbf{e}_6 and \mathbf{e}_7 , the payoffs of strategies 1 to 4 are negative, hence less than the mean payoff. It follows (using, e.g., the tools in (Hofbauer and Sigmund, 1998, chapter 17)) that the heteroclinic cycle⁷ $5 \rightarrow 6 \rightarrow 7 \rightarrow 5$, i.e.

$$\Gamma_{567} := \{\mathbf{x} \in S_7, x_5 + x_6 + x_7 = 1 \text{ and } x_5x_6x_7 = 0\} \quad (12.3.4)$$

is asymptotically stable. Actually,

Proposition 12.3.2. *For any interior initial condition $\mathbf{x} = \mathbf{x}(0)$ such that⁸ neither $x_1 = x_2 = x_3$ nor $x_5 = x_6 = x_7$, the solution $\mathbf{x}(t)$ of the replicator dynamics converges to Γ_{567} .*

Proof. See section 12.3.3. □

The same result holds for the best-response dynamics, up to replacement of the heteroclinic cycle Γ_{567} by the corresponding Shapley triangle:

$$ST_{567} := \{\mathbf{x} \in S_7 : x_5 + x_6 + x_7 = 1 \text{ and } \max_{5 \leq i \leq 7} (\mathbf{U}\mathbf{x})_i = 0\}$$

More precisely,

Proposition 12.3.3. *If $0 < \epsilon \leq 2/11$ then, for any initial condition \mathbf{x} such that neither $x_1 = x_2 = x_3$ nor $x_5 = x_6 = x_7$, ALL solutions of the best-response dynamics converge to the Shapley triangle ST_{567} .*

Proof. See section 12.3.4. □

12.3.2 Proof of proposition 12.3.1

Recall the definition of \mathbf{n}_{123} and \mathbf{n}_{567} :

$$\mathbf{n}_{123} = \left(\frac{1}{3}, \frac{1}{3}, \frac{1}{3}, 0, 0, 0, 0 \right)$$

⁶While not asymptotically stable, Γ_{123} could still attract an open set of orbits (see Brannath, 1994; see also chapter 10, proposition 10.9.5); the reason why this is not so is that at the point \mathbf{q} defined as in (12.2.10), strategy 4 is the unique best-response.

⁷Recall that a heteroclinic cycle is a globally invariant set consisting of saddle rest-points and saddle orbits connecting these rest-points.

⁸An initial condition \mathbf{x} is interior if $x_i > 0$ for every i in I

and

$$\mathbf{n}_{567} = \left(0, 0, 0, 0, \frac{1}{3}, \frac{1}{3}, \frac{1}{3}\right)$$

Let (\mathbf{x}, \mathbf{y}) be a Nash equilibrium of the game with payoffs (12.3.2). It follows from lemma 12.2.6 that

Lemma 12.3.4. *If*

$$x_1 + x_2 + x_3 > 0 \text{ and } y_1 + y_2 + y_3 > 0 \quad (12.3.5)$$

then $\bar{\mathbf{x}} = \bar{\mathbf{y}} = (1/3, 1/3, 1/3)$ and if

$$x_5 + x_6 + x_7 > 0 \text{ and } y_5 + y_6 + y_7 > 0 \quad (12.3.6)$$

then $\hat{\mathbf{x}} = \hat{\mathbf{y}} = (1/3, 1/3, 1/3)$

We now formally prove proposition 12.3.1.

Case 1: assume that (12.3.5) holds. Then, by lemma 12.3.4, $y_1 = y_2 = y_3$. Therefore $(\mathbf{n}_{567} - \mathbf{e}_4) \cdot \mathbf{U}\mathbf{y} > 0$, hence $x_4 = 0$. By symmetry, $y_4 = 0$.

Subcase 1.1. If furthermore (12.3.6) holds, then by lemma 12.3.4, $y_5 = y_6 = y_7$. Since we proved $y_4 = 0$ and $y_1 = y_2 = y_3$, it follows that \mathbf{y} is a convex combination of \mathbf{n}_{567} and \mathbf{n}_{123} . Since \mathbf{n}_{123} is a strictly better response than 5, 6 and 7 both against \mathbf{n}_{567} and against \mathbf{n}_{123} ,⁹

it follows that none of the strategies 5, 6, and 7 is a best-response to \mathbf{y} . Therefore, $x_5 + x_6 + x_7 = 0$, contradicting (12.3.6).

Subcase 1.2. If (12.3.6) does not hold. Without loss of generality, assume that $y_5 + y_6 + y_7 = 0$. Since $y_4 = 0$ and $y_1 = y_2 = y_3$, this implies that $\mathbf{y} = \mathbf{n}_{123}$. Therefore, as above, none of the strategies 5, 6 and 7 is a best-response to \mathbf{y} . Therefore $x_5 + x_6 + x_7 = 0$ which by the same argument implies $\mathbf{x} = \mathbf{n}_{123}$. Therefore, $\mathbf{x} = \mathbf{y} = \mathbf{n}_{123}$.

Case 2: assume that (12.3.5) does not hold. Without loss of generality, assume $x_1 + x_2 + x_3 = 0$. This implies that \mathbf{n}_{567} is a strictly better response to \mathbf{x} than strategy 4. Thus, $y_4 = 0$.

Subcase 2.1. If furthermore (12.3.6) holds, then \mathbf{y} is a convex combination of \mathbf{n}_{567} and strategies 1, 2, 3. This implies that \mathbf{n}_{123} is a strictly

⁹That is, $\mathbf{n}_{123} \cdot \mathbf{U}\mathbf{n}_{567} > \mathbf{e}_i \cdot \mathbf{U}\mathbf{n}_{567}$ and $\mathbf{n}_{123} \cdot \mathbf{U}\mathbf{n}_{123} > \mathbf{e}_i \cdot \mathbf{U}\mathbf{n}_{123}$ for any strategy i in $\{5, 6, 7\}$.

better response to \mathbf{y} than either 5, 6 or 7. Therefore, $x_5 = x_6 = x_7 = 0$, contradicting (12.3.6).

Subcase 2.2. If (12.3.6) does not hold, then $x_5 + x_6 + x_7 = 0$ or $y_5 + y_6 + y_7 = 0$. In the latter case, since $y_4 = 0$, it follows that \mathbf{y} has support in $\{1, 2, 3\}$, hence that \mathbf{n}_{123} is a strictly better response to \mathbf{y} than either 5, 6, or 7; therefore, in any case, $x_5 + x_6 + x_7 = 0$. Since we assumed $x_1 + x_2 + x_3 = 0$, it follows that $\mathbf{x} = \mathbf{e}_4$. As the best-responses to 4 are strategies 5, 6 and 7, \mathbf{y} must have support in $\{5, 6, 7\}$. This implies that strategies 1, 2, 3 are strictly better responses to \mathbf{y} than strategy 4, hence $x_4 = 0$, contradicting $\mathbf{x} = \mathbf{e}_4$. This completes the proof. ■

12.3.3 Proof of proposition 12.3.2

We first introduce some notations and explain the main features of the dynamics. We then prove proposition 12.3.2.

Notations and intuition We proceed much as in chapter 10. Let $\mathbf{x}(\cdot)$ be an interior solution of the replicator dynamics; that is, $\mathbf{x}(0) \in \text{int } S_7$, so that $\mathbf{x}(t) \in \text{int } S_7$ for all t . For each pure strategy i in $\{1, 2, 3\}$, define $\bar{x}_i(t)$ as the share of strategy i at time t relative to the total share of strategies 1, 2 and 3:

$$\bar{x}_i := \frac{x_i}{x_1 + x_2 + x_3} \quad (12.3.7)$$

and let $\bar{\mathbf{x}} = (\bar{x}_1, \bar{x}_2, \bar{x}_3)$. For $i \in \{5, 6, 7\}$, define similarly:

$$\hat{x}_i := \frac{x_i}{x_5 + x_6 + x_7} \quad (12.3.8)$$

and let $\hat{\mathbf{x}} = (\hat{x}_5, \hat{x}_6, \hat{x}_7)$. Finally, let $\lambda(t) := x_1(t) + x_2(t) + x_3(t)$ (resp. $\mu(t) = x_5(t) + x_6(t) + x_7(t)$) denote the total share of the three first (resp. last) strategies at time t .

The evolution of \mathbf{x} is fully described by the joint evolution of $\bar{\mathbf{x}}$, $\hat{\mathbf{x}}$, λ and μ . The interest of such a description is that as in chapter 10, lemma 10.4.1, the dynamics may be decomposed. That is, up to a change in velocity, $\bar{\mathbf{x}}$ (resp. $\hat{\mathbf{x}}$) follows the replicator dynamics of the Rock-Paper-Scissors game formed by the three first (resp. last) strategies. More precisely, letting $\bar{\mathbf{U}}$ (resp. $\hat{\mathbf{U}}$) denote the matrix obtained from (12.3.2) by restricting the players to their three first (resp. last) strategies:¹⁰

¹⁰Of course, with our choice of payoffs for the game (12.3.2), it turns out that $\bar{\mathbf{U}} = \hat{\mathbf{U}}$,

Lemma 12.3.5. *Let $\mathbf{x}(\cdot)$ be an interior solution to (12.3.1). We have (time indices suppressed):*

$$\dot{\bar{x}}_i = \lambda \bar{x}_i [(\bar{\mathbf{U}}\bar{\mathbf{x}})_i - \bar{\mathbf{x}} \cdot \bar{\mathbf{U}}\bar{\mathbf{x}}] \quad \forall i = 1, 2, 3 \quad (12.3.9)$$

and

$$\dot{\hat{x}}_i = \mu \hat{x}_i [(\hat{\mathbf{U}}\hat{\mathbf{x}})_i - \hat{\mathbf{x}} \cdot \hat{\mathbf{U}}\hat{\mathbf{x}}] \quad \forall i = 5, 6, 7 \quad (12.3.10)$$

Proof. The proof is the same as the proof of lemma 10.4.1 in chapter 10. \square

Define $\mathbf{y}(t) = (y_1(t), y_2(t), y_3(t))$ as the solution of the replicator dynamics:

$$\dot{y}_i = y_i ((\bar{\mathbf{U}}\mathbf{y})_i - \mathbf{y} \cdot \bar{\mathbf{U}}\mathbf{y}) \quad \forall i = 1, 2, 3$$

with initial condition $\mathbf{y}(0) = \bar{\mathbf{x}}(0)$. It follows from (12.3.9) that:

Lemma 12.3.6.

$$\forall t \geq 0, \quad \bar{\mathbf{x}}(t) = \mathbf{y}(\tau(t))$$

where $\tau(t)$ is the rescaled time:

$$\tau(t) := \int_0^t \lambda(s) ds$$

Proof. By (12.3.9), $\mathbf{y}(\tau(t))$ and $\bar{\mathbf{x}}(t)$ are solutions of the same differential equation, which admits a unique solution through each initial condition. \square

It follows from lemma 12.3.6 and from the behavior of the replicator dynamics in RPS games (see, e.g., Hofbauer and Sigmund (1998)) that:

Lemma 12.3.7. *If $\bar{\mathbf{x}}(0) \neq (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ and $\int_0^{+\infty} \lambda(t) dt = +\infty$ then we have: $\bar{x}_1(t)\bar{x}_2(t)\bar{x}_3(t) \rightarrow 0$ as $t \rightarrow +\infty$. Similarly, if $\hat{\mathbf{x}}(0) \neq (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ and $\int_0^{+\infty} \mu(t) dt = +\infty$, then $\hat{x}_5(t)\hat{x}_6(t)\hat{x}_7(t) \rightarrow 0$ as $t \rightarrow +\infty$.*

Before proving proposition 12.3.2, we describe informally the behaviour of the replicator dynamics in the game with payoffs (12.3.2). Let $\bar{\Gamma}$ (resp. $\hat{\Gamma}$) denote the heteroclinic cycle of the rock-scissors-paper game with payoff matrix $\bar{\mathbf{U}}$ (resp. $\hat{\mathbf{U}}$). Fix $\mathbf{x}(0) \in \text{int } S_7$ such that $\bar{\mathbf{x}}(0)$ and $\hat{\mathbf{x}}(0)$ are both different from $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$. Assume for concreteness that $\mathbf{x}(0)$ is close from the unique Nash equilibrium of the game. Initially, the payoff of strategies 4 to 7 is less than the mean payoff, hence λ increases. However, at the same time,

but this is coincidental.

$\bar{\mathbf{x}}$ spirals outwards towards $\bar{\Gamma}$. Eventually, $\bar{\mathbf{x}}$ stays long enough close to \mathbf{e}_3 for x_4 to increase substantially. Since strategy 4 is a better reply to 4 than 1, 2 and 3, it follows that x_4 keeps increasing as long as μ is low. When x_4 is sufficiently high and as long as λ stays low, μ increases so that $\hat{\mathbf{x}}$ spirals towards the heteroclinic cycle $\hat{\Gamma}$ (recall that by (12.3.10) the speed at which $\hat{\mathbf{x}}$ spirals outwards is proportional to μ). If initially $\hat{\mathbf{x}}$ was very close to $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$, it may be that, at this point, λ increases again and that the same process starts again: \mathbf{x} first approaches the cycle $1 \rightarrow 2 \rightarrow 3 \rightarrow 1$, but eventually x_4 increases, triggering an increase in μ , etc. However, each time this process runs, $\hat{\mathbf{x}}$ gets closer to the boundary (and the maximal value of μ tends to get higher). Eventually, $\hat{\mathbf{x}}$ will be sufficiently close to the boundary and μ sufficiently high for \mathbf{x} to be in the basin of attraction of the (asymptotically stable) heteroclinic cycle $5 \rightarrow 6 \rightarrow 7 \rightarrow 5$ and thus to converge to this heteroclinic cycle. In particular, λ will converge to zero, hence all strategies played in Nash equilibrium will be eliminated.

We now make these points formal.

Proof of proposition 12.3.2. Fix $\mathbf{x}(0) \in \text{int } S_7$ such that $\bar{\mathbf{x}}(0)$ and $\hat{\mathbf{x}}(0)$ are both different from $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$. We must show that $\mathbf{x}(t)$ converges to Γ_{567} (defined in (12.3.4)). It is enough to show that $\lambda(t)$ goes to zero as t goes to infinity. Indeed, in the absence of the three first strategies, the fourth strategy is strictly dominated by \mathbf{n}_{567} . Thus, if λ goes to zero, then there exists $\alpha > 0$ and a time T such that, for every $t \geq T$, $(\mathbf{n}_{567} - \mathbf{e}_4) \cdot \mathbf{U}\mathbf{x} > \alpha > 0$. Therefore, by proposition 8.3.3 in chapter 8, $x_4(t) \rightarrow 0$, hence $\mu(t) \rightarrow 1$. By lemma 12.3.7, this implies that $\mathbf{x}(t)$ converges to Γ_{567} . We now show that λ converges to zero.

The proof is by contradiction. Assume thus that λ does not converge to zero. Computing $\dot{\lambda}$, \dot{x}_4 and $\dot{\mu}$ from (12.3.1), we get (time indices suppressed):

$$\frac{\dot{\lambda}}{\lambda} = \lambda \bar{\mathbf{x}} \cdot \bar{\mathbf{U}}\bar{\mathbf{x}} + \left(\epsilon - \frac{1}{3} \right) \mu - 10x_4 - \mathbf{x} \cdot \mathbf{U}\mathbf{x} \quad (12.3.11)$$

$$\frac{\dot{x}_4}{x_4} = 4x_3 - 2\lambda - \frac{\mu}{3} - \mathbf{x} \cdot \mathbf{U}\mathbf{x} \quad (12.3.12)$$

$$\frac{\dot{\mu}}{\mu} = \mu \hat{\mathbf{x}} \cdot \hat{\mathbf{U}}\hat{\mathbf{x}} - \frac{\lambda}{3} + 10x_4 - \mathbf{x} \cdot \mathbf{U}\mathbf{x} \quad (12.3.13)$$

From (12.3.11) and (12.3.12) we get:

$$\frac{d}{dt} \ln \left(\frac{x_4}{\lambda} \right) = 4x_3 - 2\lambda - \lambda \bar{\mathbf{x}} \cdot \bar{\mathbf{U}}\bar{\mathbf{x}} - \epsilon\mu + 10x_4 \quad (12.3.14)$$

Similarly, from (12.3.11) and (12.3.13) we get:

$$\frac{d}{dt} \ln \left(\frac{\mu}{\lambda} \right) = \mu \left(\hat{\mathbf{x}} \cdot \hat{\mathbf{U}} \hat{\mathbf{x}} + \frac{1}{3} - \epsilon \right) - \lambda \cdot \left(\frac{1}{3} + \bar{\mathbf{x}} \cdot \bar{\mathbf{U}} \bar{\mathbf{x}} \right) + 20x_4 \quad (12.3.15)$$

A lemma is now needed:

Lemma 12.3.8. $\limsup_{t \rightarrow +\infty} \mu(t) \geq \frac{1}{1 + \epsilon}$

Proof. Since $\epsilon \leq 1$, it follows that $\bar{\mathbf{x}} \cdot \bar{\mathbf{U}} \bar{\mathbf{x}}$ is nonpositive. Therefore (12.3.14) implies:

$$\frac{d}{dt} \ln \left(\frac{x_4}{\lambda} \right) \geq 4x_3 - 2\lambda - \epsilon\mu + 10x_4 \quad (12.3.16)$$

If $\limsup_{t \rightarrow +\infty} \mu(t) < \frac{1}{1 + \epsilon}$, then there exists a time $T > 0$ such that for all $t \geq T$, $(1 + \epsilon)\mu(t) < 1 = \mu(t) + \lambda(t) + x_4(t)$ hence $\epsilon\mu(t) \leq \lambda(t) + x_4(t)$. Together with (12.3.16), this implies that for $t \geq T$:

$$\frac{d}{dt} \ln \left(\frac{x_4}{\lambda} \right) \geq 4x_3 - 3\lambda + 9x_4 \quad (12.3.17)$$

Therefore, for $t' \geq T$,

$$\ln \left(\frac{x_4}{\lambda} \right) (t') \geq \ln \left(\frac{x_4}{\lambda} \right) (T) + \int_T^{t'} (4x_3 - 3\lambda) dt \quad (12.3.18)$$

Since, by assumption, λ does not converge to zero, it follows from (12.3.18) that:

Claim 12.3.9. $\limsup \left(\frac{x_4}{\lambda} \right) = +\infty$

Proof. Recall the definition of $\mathbf{y}(t)$ and $\tau(t)$ and the fact that, by lemma 12.3.6, $y(\tau(t)) = \mathbf{x}(t)$. It follows from lemma 12.3.6 that if $t' \geq 0$ and $\tau' = \tau(t')$, then

$$\int_0^{t'} [4x_3(t) - 3\lambda(t)] dt = \int_0^{\tau'} \lambda(t) [4\bar{x}_3(t) - 3] dt = \int_0^{\tau(t')} [4y_3(\tau) - 3] d\tau$$

As, by assumption, λ does not converge to zero, $\tau(t) \rightarrow +\infty$ as $t \rightarrow +\infty$. Thus, in order to prove claim 12.3.9, it suffices to show that

$$\limsup_{\tau' \rightarrow +\infty} I(\tau') = +\infty \text{ with } I(\tau') = \int_0^{\tau'} [4y_3(\tau) - 3] d\tau \quad (12.3.19)$$

It follows from Gaunersdorfer and Hofbauer (1995, proof of theorem 2, p. 287-288) that the time-average

$$\mathbf{z}(\tau') := \frac{1}{\tau'} \int_0^{\tau'} \mathbf{y}(\tau) d\tau$$

converges to the Shapley triangle

$$\bar{S}T := \left\{ \mathbf{y} \in S_3 : \max_{1 \leq i \leq 3} (\bar{\mathbf{U}}\mathbf{y})_i = 0 \right\}$$

and that there exists a sequence $\tau'_n \rightarrow +\infty$ such that $\mathbf{z}(\tau'_n)$ converges to its vertex $\hat{\mathbf{q}}$ given by:

$$\hat{\mathbf{q}} = \frac{1}{1 + \epsilon + \epsilon^2} (\epsilon^2, \epsilon, 1)$$

Therefore,

$$I(\tau'_n) \sim \tau'_n \left[\frac{4}{1 + \epsilon + \epsilon^2} - 3 \right] \sim \tau'_n \left[\frac{1 - 3\epsilon - 3\epsilon^2}{1 + \epsilon + \epsilon^2} \right]$$

Since we assumed $\epsilon > 1/48$, it follows that $1 - 3\epsilon - 3\epsilon^2$ is positive, hence that $I(\tau'_n) \rightarrow +\infty$ as $n \rightarrow +\infty$. This concludes the proof of claim 12.3.9. \square

It follows from claim 12.3.9 that there exists a time $T' > T$ at which the ratio x_4/λ is greater than 1. But it follows from (12.3.17) that once the ratio x_4/λ is greater than $1/3$, then this ratio keeps increasing. Therefore, at any time $t \geq T'$, the ratio x_4/λ is greater than 1, i.e. $x_4(t) \geq \lambda(t)$. Together with (12.3.17), this implies

$$\frac{d}{dt} \ln \left(\frac{x_4}{\lambda} \right) \geq 6\lambda$$

Hence for $t \geq T'$,

$$\ln \left(\frac{x_4}{\lambda} \right) (t) \geq \ln \left(\frac{x_4}{\lambda} \right) (T') + \int_{T'}^t 6\lambda(s) ds \quad (12.3.20)$$

Since λ is continuous and, by assumption, does not go to zero, this implies that the ratio x_4/λ goes to $+\infty$, hence that λ goes to zero, a contradiction. Therefore, $\limsup_{t \rightarrow +\infty} \mu \geq \frac{1}{1+\epsilon}$. This completes the proof of lemma 12.3.8. \square

A corollary of lemma 12.3.8 is that $\int_0^{+\infty} \mu(t) dt = +\infty$. This implies that $\hat{\mathbf{x}}$ converges to the heteroclinic cycle $\hat{\Gamma}$. Since along this cycle, the mean payoff is always greater than $-\frac{1}{4}$, this implies:

$$\exists T_1 \geq 0, \forall t \geq T_1, \hat{\mathbf{x}}(t) \cdot \hat{\mathbf{U}}\hat{\mathbf{x}}(t) \geq -\frac{1}{4} - \epsilon \quad (12.3.21)$$

Assuming $\epsilon \leq \frac{1}{48}$, equations (12.3.21) and (12.3.15) jointly imply that for $t \geq T_1$:

$$\frac{d}{dt} \ln \left(\frac{\mu}{\lambda} \right) \geq \frac{\mu}{24} - \frac{\lambda}{3} \quad (12.3.22)$$

It follows from lemma 12.3.8 that there exists a time $T_2 \geq T_1$ at which the ratio μ/λ is greater than 16. Since (12.3.22) implies that the ratio μ/λ keeps increasing once greater than 8, it follows that at any time $t \geq T_2$, the ratio μ/λ is still greater than 16. Together with (12.3.22), this implies that:

$$\forall t \geq T_2, \quad \frac{d}{dt} \ln \left(\frac{\mu}{\lambda} \right) (t) \geq \frac{16\lambda}{24} - \frac{\lambda}{3} \geq \frac{\lambda(t)}{3} \quad (12.3.23)$$

Since λ is continuous, this implies that λ goes to zero, a final contradiction.

■

12.3.4 Proof of proposition 12.3.3

We first explain the main difference between proposition 12.3.3 and proposition 12.2.3. Consider an initial condition \mathbf{x} such that neither $x_1 = x_2 = x_3$ nor $x_5 = x_6 = x_7$. It is easy to show along the lines of the proof of 12.2.3 that there exists a time at which a strategy in $\{4, 5, 6, 7\}$ becomes a best-response, and that this implies that there exists a solution with initial condition \mathbf{x} which converges to the Shapley triangle ST_{567} . The difficulty is to prove that *all* solutions with initial condition \mathbf{x} converge to ST_{567} .

To see why this difficulty arises, consider a time t such that the pure best-responses to $\mathbf{x}(t)$ are strategies 1, 3 and 4. After time t , the solution is not uniquely defined: it might point towards \mathbf{e}_4 , but it might also point towards \mathbf{e}_1 or towards the mixed Nash equilibrium of the game restricted to $\{\mathbf{e}_1, \mathbf{e}_4\} \times \{\mathbf{e}_1, \mathbf{e}_4\}$; each of these three possibilities gives rise to a valid solution to (12.2.1).

In proposition 12.2.3, we avoided this difficulty by focusing on a particular solution. In game (12.3.2), it may be shown that, up to a change in velocity, $\bar{\mathbf{x}}$ and $\hat{\mathbf{x}}$ follow a Rock-Paper-Scissors best-response dynamics. This allows to prove asymptotic results valid for all solutions without having to follow precisely each solution.

We now turn to the proof. We use the notations introduced in section 12.3.1 and at the beginning of section 12.3.3 (that is, $\bar{\mathbf{x}}$, $\hat{\mathbf{x}}$, λ , μ , $\bar{\mathbf{U}}$, $\hat{\mathbf{U}}$ and \mathbf{n}_{567}). Let \bar{G} (resp. \hat{G}) denote the RPS game with payoff matrix $\bar{\mathbf{U}}$ (resp. $\hat{\mathbf{U}}$).

We first show that:

Lemma 12.3.10. *If there exists a time T such that*

$$\text{none of the strategies 1, 2 and 3 is a best-response to } \mathbf{x}(T) \quad (12.3.24)$$

and

$$x_5(T) = x_6(T) = x_7(T) \text{ does not hold} \quad (12.3.25)$$

then $\mathbf{x}(t)$ converges to the Shapley triangle ST_{567} .

We then show that, up to a change of velocity, $\bar{\mathbf{x}}$ and $\hat{\mathbf{x}}$ follow the best-response dynamics in respectively \bar{G} and \hat{G} . More precisely,

Lemma 12.3.11. *For almost all time t , if $\lambda(t) > 0$ then (time indices suppressed):*

$$\dot{\bar{\mathbf{x}}} \in \left(1 + \frac{\dot{\lambda}}{\lambda}\right) (BR(\bar{\mathbf{x}}) - \bar{\mathbf{x}}) \quad (12.3.26)$$

where $BR(\cdot)$ is the best-response correspondence in \bar{G} . Similarly, if $\mu(t) > 0$ then

$$\dot{\hat{\mathbf{x}}} \in \left(1 + \frac{\dot{\mu}}{\mu}\right) (BR(\hat{\mathbf{x}}) - \hat{\mathbf{x}}) \quad (12.3.27)$$

where $BR(\cdot)$ is the best-response correspondence in \hat{G} .

This allows to show that:

Lemma 12.3.12. *For any initial condition with neither $x_1 = x_2 = x_3$ nor $x_5 = x_6 = x_7$ there exists a time T such that none of the strategies 1, 2 and 3 is a best-response to $\mathbf{x}(T)$ and such that $x_5(T) = x_6(T) = x_7(T)$ does not hold.*

Proposition 12.3.3 follows from the combination of lemmas 12.3.10 and 12.3.12.

Proof of lemma 12.3.10. Let $PBR(\mathbf{x}) := \{i \in I, \mathbf{e}_i \in BR(\mathbf{x})\}$ denote the set of pure best-responses to \mathbf{x} . Let T check the assumptions of lemma 12.3.10. Since none of the strategies 1, 2 and 3 is a best-response to $\mathbf{x}(T)$ and since strategies 5, 6 and 7 are cyclically symmetric, we may assume that we are in one of the following cases:

- (a) $\{5, 6, 7\} \subseteq PBR(\mathbf{x}(T))$
- (b) $PBR(\mathbf{x}(T)) \subseteq \{5, 6\}$
- (c) $PBR(\mathbf{x}(T)) = \{4, 6\}$ or $PBR(\mathbf{x}(T)) = \{4, 5, 6\}$
- (d) $PBR(\mathbf{x}(T)) = \{4\}$

In case (a), $(\mathbf{U}\mathbf{x})_5 = (\mathbf{U}\mathbf{x})_6 = (\mathbf{U}\mathbf{x})_7$ for $\mathbf{x} = \mathbf{x}(T)$. But as explained in the proof of proposition 12.3.2, the differences between the payoffs of strategy 5, 6 and 7 are only due to the values of x_5 , x_6 and x_7 . In particular, if $(\mathbf{U}\mathbf{x})_5 = (\mathbf{U}\mathbf{x})_6 = (\mathbf{U}\mathbf{x})_7$, then either $x_5 = x_6 = x_7 = 0$ or $(\hat{\mathbf{U}}\hat{\mathbf{x}})_5 = (\hat{\mathbf{U}}\hat{\mathbf{x}})_6 = (\hat{\mathbf{U}}\hat{\mathbf{x}})_7$; the latter implies that $\hat{\mathbf{x}}$ is a symmetric Nash equilibrium of the game restricted to $\{5, 6, 7\} \times \{5, 6, 7\}$, hence $\hat{\mathbf{x}} = (1/3, 1/3, 1/3)$. Thus, in any case, $x_5 = x_6 = x_7$. Therefore case (a) is ruled out by the assumption that $x_5(T) = x_6(T) = x_7(T)$ does not hold.

In case (b), an adaptation of the proof of proposition 12.2.5 shows that $\mathbf{x}(t)$ converges towards ST_{567} .

In case (c), immediately after time T , the solution must point towards the convex hull of \mathbf{e}_4 , \mathbf{e}_5 and \mathbf{e}_6 . Since in the game restricted to $\{4, 5, 6\} \times \{4, 5, 6\}$, strategy 4 is strictly dominated by strategy 6, it follows that the solution actually points towards the edge $[\mathbf{e}_5, \mathbf{e}_6]$; since in the game restricted to $\{5, 6\} \times \{5, 6\}$, strategy 5 is strictly dominated by strategy 6, this implies that the solution actually points towards \mathbf{e}_6 , so that strategy 6 becomes the unique best-response. We are then back to case (b).

In case (d), the solution will point towards \mathbf{e}_4 till some time T' when strategy 4 ceases to be the unique best-response. It follows from the improvement principle that none of the strategies 1, 2, and 3 is a best-response to $\mathbf{x}(T')$. Furthermore, between time T and time T' , the shares of strategies 5, 6 and 7 decrease exponentially, so that:

$$\forall i \in \{5, 6, 7\}, x_i(T') = x_i(T) \exp(-(T' - T))$$

Together with (12.3.25) this implies that at time T' the condition $x_5 = x_6 = x_7$ does not hold. Therefore, we are back to case (c). This concludes the proof. ■

Proof of lemma 12.3.11. We only prove (12.3.26). The proof of (12.3.27) is exactly the same. Consider a time t and a vector $\mathbf{b} \in BR(\mathbf{x}(t))$ such that

$\dot{\mathbf{x}}(t) = \mathbf{b} - \mathbf{x}(t)$, or equivalently $\mathbf{x}(t) + \dot{\mathbf{x}} = \mathbf{b}$. We then have

$$\lambda + \dot{\lambda} = b_1 + b_2 + b_3 \quad (12.3.28)$$

Furthermore, differentiating the equality $\lambda \bar{x}_i = x_i$ for i in $\{1, 2, 3\}$ we get:

$$\dot{\lambda} \bar{x}_i + \lambda \dot{\bar{x}}_i = \dot{x}_i = b_i - x_i = b_i - \lambda \bar{x}_i$$

hence

$$\dot{\bar{x}}_i = \frac{b_i - (\lambda + \dot{\lambda}) \bar{x}_i}{\lambda} \quad (12.3.29)$$

If $b_1 = b_2 = b_3 = 0$, then by (12.3.28), $\lambda + \dot{\lambda} = 0$ (hence $1 + \frac{\dot{\lambda}}{\lambda} = 0$), and by (12.3.29), $\dot{\bar{x}} = 0$. Therefore (12.3.26) holds trivially.

If $b_1 + b_2 + b_3 > 0$ define $\bar{\mathbf{b}} \in S_3$ by $\bar{b}_i = b_i / (b_1 + b_2 + b_3)$ for $i = 1, 2, 3$. It follows from (12.3.28) that $\mathbf{b} = (\lambda + \dot{\lambda}) \bar{\mathbf{b}}$. Therefore, by (12.3.29):

$$\dot{\bar{\mathbf{x}}} = \frac{\lambda + \dot{\lambda}}{\lambda} (\bar{\mathbf{b}} - \bar{\mathbf{x}}) = \left(1 + \frac{\dot{\lambda}}{\lambda}\right) (\bar{\mathbf{b}} - \bar{\mathbf{x}}) \quad (12.3.30)$$

To complete the proof, it suffices to show that $\bar{\mathbf{b}}$ is a best-response to $\bar{\mathbf{x}}(t)$. This is where we use the special structure of the game. Note that for any i and j in $\{1, 2, 3\}$,

$$(\mathbf{U}\mathbf{x})_i \geq (\mathbf{U}\mathbf{x})_j \Leftrightarrow (\bar{\mathbf{U}}\bar{\mathbf{x}})_i \geq (\bar{\mathbf{U}}\bar{\mathbf{x}})_j$$

This implies that if strategy $i \in \{1, 2, 3\}$ is a best-response to \mathbf{x} then it is a best-response to $\bar{\mathbf{x}}$. Now, if $\bar{b}_i > 0$ then $b_i > 0$ hence strategy i is a best-response to $\mathbf{x}(t)$; therefore, as we just saw, strategy i is a best-response to $\bar{\mathbf{x}}(t)$. It follows that $\bar{\mathbf{b}} \in BR(\bar{\mathbf{x}}(t))$, which completes the proof. ■

Proof of lemma 12.3.12. For simplicity, we assume that $\lambda(0)$ and $\mu(0)$ are both positive (otherwise, $\bar{\mathbf{x}}$ and/or $\hat{\mathbf{x}}$ are not initially defined and this complicates some arguments). It will be clear ex-post that this is without loss of generality. Let $\mathbf{y}(\cdot)$ be the (unique) solution of the best-response dynamics in the game \bar{G} with initial condition $\mathbf{y}(0) = \bar{\mathbf{x}}(0)$. Let $\tau(t)$ denote the rescaled time:

$$\tau(t) := \int_0^t \left(1 + \frac{\dot{\lambda}}{\lambda}\right)(s) ds = t + \ln \left(\frac{\lambda(t)}{\lambda(0)}\right) \quad (12.3.31)$$

Note that $\tau(t)$ is nondecreasing as, due to (12.3.28), $\dot{\lambda} + \lambda$ is always nonnegative. It follows from (12.3.26) that

$$\forall t \geq 0, \quad \bar{\mathbf{x}}(t) = \mathbf{y}(\tau(t)) \tag{12.3.32}$$

Similarly, let $\mathbf{y}'(\cdot)$ be the (unique) solution of the best-response dynamics in \hat{G} with initial condition $\mathbf{y}'(0) = \hat{\mathbf{x}}(0)$ and let:

$$\tau'(t) := t + \ln \left(\frac{\mu(t)}{\mu(0)} \right) \tag{12.3.33}$$

It follows from (12.3.27) that

$$\forall t \geq 0, \quad \hat{\mathbf{x}}(t) = \mathbf{y}'(\tau'(t)) \tag{12.3.34}$$

Let $\bar{V}(\bar{\mathbf{x}}) = \max_{1 \leq i \leq 3} (\bar{\mathbf{U}}\bar{\mathbf{x}})_i$. Let

$$\mathbf{q} = \left(\frac{\epsilon^2}{1 + \epsilon + \epsilon^2}, \frac{\epsilon}{1 + \epsilon + \epsilon^2}, \frac{1}{1 + \epsilon + \epsilon^2} \right)$$

be the vertex of the Shapley triangle $\bar{S}T := \{\mathbf{y} \in S_3, \bar{V}(\mathbf{y}) = 0\}$ such that $(\bar{\mathbf{U}}\mathbf{q})_3 = (\bar{\mathbf{U}}\mathbf{q})_1 = 0$. Since $\bar{\mathbf{x}}(0) \neq (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$, it follows from (12.3.32) that:

Claim 12.3.13. *If $\lambda(t)$ does not converge to 0 then $\bar{V}(\bar{\mathbf{x}}(t)) \rightarrow 0$ and there exists a sequence (t_n) such that $t_n \rightarrow +\infty$ and $\bar{\mathbf{x}}(t_n) \rightarrow \mathbf{q}$.*

Proof. In view of (12.3.32) and of the behavior of the best-response dynamics in RPS games (see chapter 10 or, e.g., Gaunersdorfer and Hofbauer, 1995), it suffices to show that if $\lambda \not\rightarrow 0$ then $\tau(t) \rightarrow +\infty$. Since $\tau(t)$ is nondecreasing, it actually suffices to show that if $\lambda \not\rightarrow 0$ then $\limsup_{t \rightarrow +\infty} \tau(t) = +\infty$. This is clear from (12.3.31). \square

Similarly, let $\hat{V}(\hat{\mathbf{x}}) = \max_{4 \leq i \leq 6} (\hat{\mathbf{U}}\hat{\mathbf{x}})_i$. It follows from (12.3.34) that:

Claim 12.3.14. *If $\mu(t)$ does not converge to 0 then $\hat{V}(\hat{\mathbf{x}}(t)) \rightarrow 0$.*

Another consequence of (12.3.34) is that there is no time $T \geq 0$ such that $x_5(T) = x_6(T) = x_7(T)$. Therefore, to establish lemma 12.3.12, it suffices to show that there exists a time $T \geq 0$ such that none of the strategies 1, 2, 3 is a best-response to $\mathbf{x}(T)$. Assume by contradiction that:

$$\forall t \geq 0, \quad \left[\max_{1 \leq i \leq 3} (\mathbf{U}\mathbf{x})_i \right] - (\mathbf{U}\mathbf{x})_4 \geq 0 \tag{12.3.35}$$

$$\forall t \geq 0, \left[\max_{1 \leq i \leq 3} (\mathbf{U}\mathbf{x})_i \right] - \left[\max_{5 \leq j \leq 7} (\mathbf{U}\mathbf{x})_j \right] \geq 0 \quad (12.3.36)$$

Straightforward computations show that:

$$\left[\max_{1 \leq i \leq 3} (\mathbf{U}\mathbf{x})_i \right] - (\mathbf{U}\mathbf{x})_4 = \lambda [\bar{V}(\bar{\mathbf{x}}) - (\mathbf{U}\bar{\mathbf{x}})_4] - 10x_4 + \epsilon\mu \quad (12.3.37)$$

(($\mathbf{U}\bar{\mathbf{x}}$) = $-2\bar{x}_1 - 2\bar{x}_2 + 2\bar{x}_3$ is defined by identifying $\bar{\mathbf{x}}$ and $(\bar{x}_1, \bar{x}_2, \bar{x}_3, 0, 0, 0, 0)$)

Furthermore,

$$\left[\max_{1 \leq i \leq 3} (\mathbf{U}\mathbf{x})_i \right] - \left[\max_{5 \leq j \leq 7} (\mathbf{U}\mathbf{x})_j \right] = \lambda [\bar{V}(\bar{\mathbf{x}}) + 1/3] - 20x_4 - \mu [\hat{V}(\hat{\mathbf{x}}) + 1/3 - \epsilon] \quad (12.3.38)$$

Equations (12.3.35) and (12.3.37) imply:

$$\frac{\mu}{\lambda} \geq \frac{(\mathbf{U}\bar{\mathbf{x}})_4 - \bar{V}(\bar{\mathbf{x}})}{\epsilon} \geq 0 \quad (12.3.39)$$

Equations (12.3.36) and (12.3.38) imply:

$$\lambda [\bar{V}(\bar{\mathbf{x}}) + 1/3] - \mu [\hat{V}(\hat{\mathbf{x}}) + 1/3 - \epsilon] \geq 20x_4 \geq 0 \quad (12.3.40)$$

and in particular, if $\hat{V}(\hat{\mathbf{x}}) > -1/3 + \epsilon$,

$$\frac{\mu}{\lambda} \leq \frac{\bar{V}(\bar{\mathbf{x}}) + 1/3}{\hat{V}(\hat{\mathbf{x}}) + 1/3 - \epsilon} \quad (12.3.41)$$

If $\lambda \rightarrow 0$ then there exists a time T_1 such that, for all $t \geq T_1$, $(\mathbf{n}_{567} - \mathbf{e}_4) \cdot \mathbf{U}\mathbf{x}(t) > 0$ so that strategy 4 is not a best-response to $\mathbf{x}(t)$. Therefore, $x_4 \rightarrow 0$ hence $\mu \rightarrow 1$. By claim (12.3.14), this implies that $\hat{V}(\hat{\mathbf{x}}) \rightarrow 0$. Therefore the left-hand side of (12.3.40) converges to $-1/3 + \epsilon$, which is negative as, by assumption, $\epsilon \leq 2/11$. This contradicts (12.3.40).

Thus, we may assume that λ does not converge to 0. By claim 12.3.13 this implies that

$$\limsup \frac{\mu}{\lambda} \geq \limsup \left(\frac{(\mathbf{U}\bar{\mathbf{x}})_4 - \bar{V}(\bar{\mathbf{x}})}{\epsilon} \right) \geq \frac{(\mathbf{U}\mathbf{q})_4}{\epsilon}$$

hence

$$\limsup \frac{\mu}{\lambda} \geq \frac{1}{\epsilon} \left[\frac{2(1 - \epsilon - \epsilon^2)}{1 + \epsilon + \epsilon^2} \right] \quad (12.3.42)$$

The first inequality in (12.3.40) implies that $\limsup x_4 < 1$. Therefore, it follows from (12.3.42) and the assumption $\epsilon \leq 2/11$ that μ does not converge to zero. By claim 12.3.14 and (12.3.39), this implies that $\hat{V}(\hat{x}) \rightarrow 0$. Furthermore, we have seen that λ does not converge to zero and that this implies that $\bar{V}(\bar{x}) \rightarrow 0$. Therefore, it follows from (12.3.41) that

$$\limsup \frac{\mu}{\lambda} \leq \frac{1/3}{(1/3) - \epsilon} = \frac{1}{1 - 3\epsilon} \quad (12.3.43)$$

For $\epsilon \leq 2/11$, equations (12.3.42) and (12.3.43) are contradictory. Therefore, there exists a time T such that none of the strategies 1, 2 and 3 is a best-response to $\mathbf{x}(T)$. This completes the proof of lemma 12.3.12 and of proposition 12.3.3. ■

12.4 Discussion

In the games we presented, the Nash equilibrium is unique and quasi-strict, and therefore persistent, regular¹¹, hence strongly stable, essential, strictly proper, strictly perfect, etc. Thus, from the traditional, rationalistic point of view, it is the unambiguous solution of the game. However, under the replicator dynamics and/or best-response dynamics, all strategies in the support of this Nash equilibrium are eliminated from almost all initial conditions. This indicates an even wider gap between strategic and evolutionary considerations that had been noted before.

It is the author's view that elimination of all strategies in the support of Nash equilibria from almost all initial conditions likely occurs for many dynamics, and is not a specific feature of the replicator dynamics and of the best-response dynamics. However, as mentioned in the introduction, this phenomenon can only occur in relatively large games and these games might prove difficult to analyse for dynamics more complicated or less studied than the best-response or replicator dynamics.

¹¹In a bimatrix game, an isolated and quasi-strict Nash equilibrium is regular, see Van Damme, 1991

Part III

Origins of Multicellularity

Germ-Soma Differentiation in Volvocine Green Algae

Chapitre 13

Introduction to part III: Transition from Unicellular to Multicellular Organisms and Germ-Soma Separation

This part consists of a single article, co-written with Richard. E. Michod, Cristian Solari, Mathilde Hurand, and Aurora M. Nedelcu, and accepted for publication in *Journal of Theoretical Biology*. We study aspects of the transition from unicellular to multicellular organisms, and in particular factors driving germ-soma separation. While the points made in the formal models are fairly general, we focus on specific organisms: volvocine green algae. The purpose of this introduction is both to introduce some basic issues related to the article and to sum up our contributions. A more technical introduction, including extensive references, will be given in the article.

The living world is organized hierarchically: genes inside chromosomes, chromosomes inside genomes, genomes inside cells, cells inside multi-cellular organisms, and multi-cellular organisms inside animal societies. This is not how the living world originally looked like. Rather, the evolution of such a hierarchical organization required a series of evolutionary transitions between levels of organization. These are key-events in evolution and the study of these transitions is accordingly one of the major topics in evolutionary biology (Buss, 1987; Maynard Smith and Szathmary, 1995; Michod, 1999).

We focus on the transition between unicellular and multicellular organisms. A fundamental aspect of this transition is that in a multicellular or-

ganism, cells may differentiate. In particular, some cells (called germ cells) may specialize in reproduction while other cells (somatic cells) specialize in vegetative functions (they help the reproductive cells to survive and reproduce). This division of labor is not an option open to solitary cells, since a solitary cell which would not reproduce would simply not transmit its genes. In contrast, in multicellular organism, a somatic cell may transmit its genes even though it does not reproduce itself, provided that its genes are identical or related to the reproductive cells' genes.

While cell differentiation opens up immense possibilities of evolution and complexification to multicellular organisms, cell differentiation, and germ-soma separation in particular, is not present in all multicellular organisms. It may well be that multicellular organisms evolve, for instance because group-living enhances viability due to predation avoidance or ability to catch bigger preys, or because it provides a buffered environment, etc., but that cells stay undifferentiated, because differentiation does not provide an immediate selective advantage. We try to understand the factors driving the evolution of germ-soma differentiation, both in general and in the special case of volvocine green algae.

The special interest of these algae for understanding the transition from unicellular to multicellular organism and aspects of their biology relevant to our analysis will be explained in the text. However, it may help at this point to know that there are many different (but closely related) species of volvocine green algae, which range from unicellular forms to colonies involving more than 50,000 cells. Boldly said (the article will be more precise), the smaller colonies are undifferentiated, colonies of intermediate size are partially differentiated (some cells specialize in somatic functions and the other cells are generalists) and the larger colonies are totally differentiated in somatic cells and germ cells. Furthermore, the proportion of somatic cells increases with the size of the colony. These are the facts we try to understand.

A basic and general point is that cells should both survive and reproduce (or, in a group setting, contribute to the survival of the group and contribute to its reproduction). Since resources are limited, the resources affected to one of these two basic functions is likely to detract from the resources affected to the other one. This results in a trade-off between fecundity and viability, embodied in a trade-off function. More precisely, let v and b be the contribution of a cell to the viability and fecundity of the colony.¹ We assume that

¹Assume for now that the contributions of the cells to the viability and fecundity and

there exists an intrinsic relationship between v and b , so that $v = f(b)$ for some decreasing trade-off function f .

We argue that the evolution of the curvature (concavity/convexity) of this trade-off function is crucial to germ-soma separation. The essence of the argument is as follows: division of labor is favored if there are increasing returns on investment in viability and/or fecundity (specialists are more efficient than generalists); increasing (decreasing) returns on investment corresponds to a convex (concave) trade-off function; therefore a convex (concave) trade-off selects for (against) specialization.

A factor that tends to make the trade-off function convex like, hence selects for germ-soma specialization, is the existence of a fixed cost to reproduction (a minimal investment in reproduction needed for there to be any reproduction at all).² In the presence of such a fixed cost of reproduction³, a cell specializing in somatic functions would contribute much more to the viability of the group than a cell performing viability functions during most of its life cycle but also reproducing.

Recall that, as the size of volvocine green algae colonies increases, these colonies tend to be more and more differentiated. In light of the previous paragraphs, this suggests that the viability-fecundity trade-off becomes increasingly convex as size increases. Why should it be so? This seems to be linked for a good part to the reproduction mode of these algae: rather than growing, dividing in two, growing, dividing in two, etc. reproductive cells of a 2^N cell colony first grow a lot and then divide N times. Furthermore, in large colonies ($N > 5$), a reproductive cell cannot both divide and contribute to the motility of the colony.⁴ It follows that, in large colonies, a reproductive cell must go through a long phase during which it cannot contribute to the motility of the colony.

Now imagine a large colony of undifferentiated cells. Assume for simplic-

colony are additive. See the article for more discussion and relaxation of this assumption.

²We do not mean that such a fixed cost necessarily exists, but that when it exists, it selects for specialization.

³In the article, we use the expression “initial cost of reproduction to survival” to stress that we focus on the cost resulting from going to no reproduction to a little reproduction, and that this cost is in term of reduced contribution to survival.

⁴More precisely, cells have a flagella, which serves several purposes: it helps to move the colony but to increase the exchange of chemical with the environment due to a vortex effect. For simplicity, we focus on motility.

ity that all cells reproduce at the same time.⁵ Since there are no somatic cells, the colony undergoes a long reproduction phase during which it has no motility at all, a phase which is increasingly long as the size of the colony increases. Since motility is an important component of survival, such a colony is likely to have low survival capabilities. If some cells were to specialize in vegetative functions, then they could provide at least some motility during this phase, and this would likely lead to a substantial increase in survival capabilities of the colony. We argue that this is one of the factors leading to germ-soma specialization in large colonies. Additional factors will be discussed in the article.

The fact that the proportion of somatic cells increases as the size of the colonies increases may be understood as follows: in a large colony, reproductive cells, just before dividing, are much larger and heavier compared to somatic cells than in a small colony (this is due to the particular reproduction mode of volvocine green algae, explained above). Accordingly, a much higher proportion of somatic cells is needed to move the colony during the reproduction phase in a large colony than in a small one. Other factors are also important, in particular, hydrodynamical factors: as the size of the colony increases, drag increases superlinearly, hence it is very important for a large colony to have a compact organization, which requires specialization of a high proportion of cells in soma.

References

- Buss, L. W. (1987), *The Evolution of Individuality*, Princeton University Press, Princeton, NJ.
- Maynard-Smith J. and E. Szathmary (1995), *The Major Transitions in Evolution*, W.H.Freeman Press, New York
- Michod, R. E. (1999), *Darwinian Dynamics, Evolutionary Transitions in Fitness and Individuality*, Princeton University Press, Princeton, NJ.

⁵This is what is observed in extant volvocine green algae. If the cells were reproducing at different times or at different speeds, then this would likely disrupt the individuality and functionality of the colony.

Chapitre 14

Life-history Evolution and the Origin of Multicellularity



Life-history evolution and the origin of multicellularity

Richard E. Michod^{a,*}, Yannick Viossat^{b,c}, Cristian A. Solari^a,
Mathilde Hurand^d, Aurora M. Nedelcu^e

^a*Department of Ecology and Evolutionary Biology, University of Arizona, Tucson, AZ 85721, USA*

^b*Laboratoire d'Économétrie, Ecole Polytechnique, 1, rue Descartes, 75005 Paris, France*

^c*Department of Economics, Stockholm School of Economics, Box 6501, SE-113 83 Stockholm, Sweden*

^d*Laboratoire d'Informatique, Ecole Polytechnique, 91128 Palaiseau, France*

^e*Department of Biology, University of New Brunswick, Fredericton, NB, Canada E3B 6E1*

Received 30 January 2005; accepted 30 July 2005

Abstract

The fitness of an evolutionary individual can be understood in terms of its two basic components: survival and reproduction. As embodied in current theory, trade-offs between these fitness components drive the evolution of life-history traits in extant multicellular organisms. Here, we argue that the evolution of germ–soma specialization and the emergence of individuality at a new higher level during the transition from unicellular to multicellular organisms are also consequences of trade-offs between the two components of fitness—survival and reproduction. The models presented here explore fitness trade-offs at both the cell and group levels during the unicellular–multicellular transition. When the two components of fitness negatively covary at the lower level there is an enhanced fitness at the group level equal to the covariance of components at the lower level. We show that the group fitness trade-offs are initially determined by the cell level trade-offs. However, as the transition proceeds to multicellularity, the group level trade-offs depart from the cell level ones, because certain fitness advantages of cell specialization may be realized only by the group. The curvature of the trade-off between fitness components is a basic issue in life-history theory and we predict that this curvature is concave in single-celled organisms but becomes increasingly convex as group size increases in multicellular organisms. We argue that the increasingly convex curvature of the trade-off function is driven by the initial cost of reproduction to survival which increases as group size increases. To illustrate the principles and conclusions of the model, we consider aspects of the biology of the volvocine green algae, which contain both unicellular and multicellular members.

© 2005 Elsevier Ltd. All rights reserved.

Keywords: Cost of reproduction; Germ–soma differentiation; Fitness; Body size; Evolutionary transitions; Life-history evolution; Volvox

1. Introduction

Fitness can be understood in terms of its two basic components: survival (viability) and reproduction (fecundity). Investment in one component often detracts from the other, leading to trade-offs in fitness components. A wide body of work shows that fitness trade-offs underlie the evolution of diverse life-history traits in extant organisms

(Stearns, 1992; Charlesworth, 1980). We show here that trade-offs between survival and reproduction have special significance during evolutionary transitions; in particular, they may drive the evolution of individuality during the transition from unicellular to multicellular organisms.

The emergence of individuality during the unicellular–multicellular transition is based on the evolution of cells that differentiate and specialize in reproductive and survival-enhancing vegetative functions. In unicellular individuals, the same cell must contribute to both fitness components, these contributions typically being separated in time. In multicellular groups, cells may specialize during development in either component, leading to the differentiation and specialization in reproductive (germ) and

*Corresponding author. Tel.: +1 520 621 7509.

E-mail addresses: michod@u.arizona.edu (R.E. Michod),
yannick.viossat@polytechnique.org (Y. Viossat),
casolari@email.arizona.edu (C.A. Solari),
mathilde.hurand@polytechnique.org (M. Hurand),
anedelcu@unb.ca (A.M. Nedelcu).

vegetative survival-enhancing functions (soma)—what we term germ–soma or “G–S” specialization. As cells specialize in these different but essential fitness components, they relinquish their autonomy in favor of the group and, as a result, fitness and individuality are transferred from the cell level to the group level. We argue here that the evolution of G–S separation and the emergence of individuality at the new higher level are consequences of fitness trade-offs among life-history components—in short, that life-history evolution is a fundamental factor in evolutionary transitions. We first present an overview of the volvocine green algae, which are the organisms we had in mind when constructing the models. Although we discuss the models with regard to the volvocine algae, we have kept the assumptions of the models general so that they will apply to other groups.

2. The volvocine green algae

The evolution of multicellular organisms from unicellular and colonial ancestors is the premier example of the integration of lower level units into a new, higher level individual. Unfortunately, for the major multicellular lineages, the factors underlying their origin lie hidden deep in their evolutionary past, obscured by hundreds of millions of years of evolution. In contrast, according to one estimate (Rausch et al., 1989), the colonial volvocine algae (Fig. 1) diverged from a unicellular ancestor just 35 million years ago, providing a unique window into this major transition.

Volvocine algae are flagellated photosynthetic organisms that range from unicellular (i.e. *Chlamydomonas*) and multicellular forms with no cell differentiation (e.g., *Gonium* and *Eudorina*; 8–32 cells) or incomplete G–S differentiation (*Pleodorina*; 64–128 cells) to multicellular

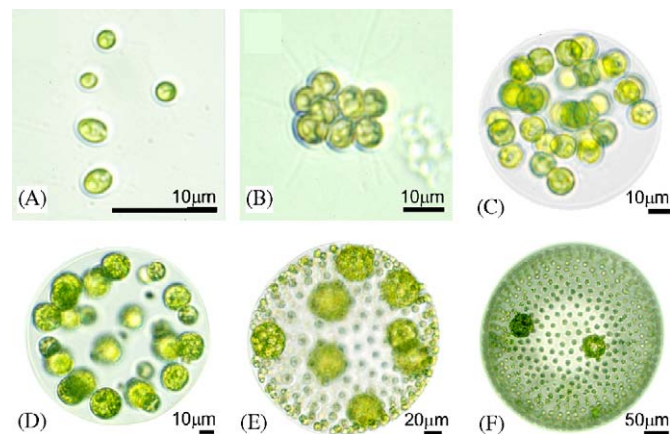


Fig. 1. Subset of volvocine species which shows an increase in complexity, cell number, volume of extracellular matrix, division of labor between somatic and reproductive cells, and proportion of somatic cells. A: *Chlamydomonas reinhardtii*; B: *Gonium pectorale*; C: *Eudorina elegans*; D: *Pleodorina californica*; E: *Volvox carteri*; F: *Volvox aureus*. Where two cell types are present (D, E and F), the smaller cells are the vegetative sterile somatic cells, whereas the larger cells are the reproductive germ cells. Picture credit: C. Solari.

forms with complete G–S separation (i.e. *Volvox*; 500–50,000 cells) (Kirk, 1998). In multicellular volvocine colonies the number of cells is determined by the number of cleavage divisions that take place during their initial formation, and cell number is not augmented by additional cell divisions (Kirk, 1997). In colonies without G–S separation (i.e., *Gonium*, *Eudorina*), each cell gives rise to a daughter colony. The life cycle corresponds to one of discrete generations as the parent colony dies as soon as the daughter colonies hatch.

It is believed that all multicellular volvocine algae have evolved from a common ancestor similar to the extant *Chlamydomonas reinhardtii* (Coleman, 1999; Larson et al., 1992). Within this closely-related monophyletic group (Buchheim et al., 1994; Coleman, 1999; Larson et al., 1992; Nozaki et al., 2000, 2002, 2003; Nozaki, 2003), significant evolutionary transitions have occurred repeatedly within a relatively short period of time (possibly as short as 35 million years (Rausch et al., 1989), as already mentioned), suggesting strong selective pressures driving the evolution of multicellularity and G–S specialization.

Although several model systems have been used to investigate the origins of multicellularity, including choanoflagellates (King and Carroll, 2001), cellular slime molds (Strassmann et al., 2000; Foster et al., 2002; Queller et al., 2003) and myxobacteria (Velicer et al., 2000; Shimkets, 1990), volvocine algae exhibit a number of features that make them especially suitable for our work (see the Volvocales Information Project at www.unbf.ca/vip). Like most familiar multicellular forms, and unlike other model experimental systems such as slime molds or myxobacteria, multicellular volvocine algae develop from a single cell, so the cells in the group are related. They can easily be obtained from nature (where uni- and multicellular forms coexist) and maintained in the lab under realistic conditions that allow for an eco-physiological framework. Many aspects of their biology have been studied (Kirk, 1998) (cytology, biochemistry, development, genetics, physiology, natural history, ecology and life-history). The ‘social’ genes necessary for group living and fitness reorganization have been identified in *V. carteri* (Kirk et al., 1999; Miller and Kirk, 1999), indicating that the underlying genetics of cellular differentiation and G–S specialization is likely simple and may not involve many genetic steps (Kirk, 1997, 1998).

3. Overview of models

The models studied below focus on the trade-offs between survival and reproduction and on how these trade-offs change as group size increases and cells specialize in reproductive and vegetative functions. The models are based on three general assumptions: (i) there are both advantages and disadvantages associated with increasing group size, (ii) generations are discrete, so that fitness is the product of viability and fecundity, and (iii) variation in

fitness exists primarily at the group or colony level; within-group variation is assumed negligible.

Larger group size may be beneficial for survival (for example, in terms of predation avoidance, ability to catch bigger prey, or a buffered environment within a group), as well as for reproduction (for example, in terms of a higher number or quality of offspring). Reduced predation is likely to be especially important in the volvocine algae (Morgan, 1980; Pentecost, 1983; Porter, 1977; Reynolds, 1984; Shikano et al., 1990). However, we do not explicitly model or discuss further this assumed advantage of larger groups.

Increasing group size may also detract from fitness, because of the increasing need for local resources, less effective movement within the environment, and longer generation time. In volvocine algae, these disadvantages of larger size are the result of (i) the ‘flagellation constraint’ which impedes motility in dividing cells (Koufopanou, 1994) and (ii) the ‘enlargement constraint’ which refers to the transport and hydrodynamic problems associated with the metabolism and translocation of an increasingly larger colony (Solari et al., 2005a, b). Bell (1985) has also discussed with respect to the volvocines the effect of increased colony size on the increased generation time and the resources needed.

The flagellation constraint impedes motility, and thus viability, during cell division (Koufopanou, 1994), and is a consequence of the coherent glycoprotein cell wall that does not allow the flagellar basal bodies to move laterally and take the expected position of centrioles in cell division while still attached to the flagella (as they do in naked green flagellates). This constraint sets an upper limit of five for the number of times a cell can divide while still maintaining an active flagellum, and thus becomes critical at about the 32-cell stage.

The enlargement constraint stems from the particular way in which volvocine algae reproduce. Because post-embryonic cell divisions are not possible (although the young colonies do increase in size after their release from the mother colony through an increase in cell size and volume of extracellular matrix), the embryo contains all the cells present in the adult. Consequently, the larger the number of cells in the colony, the larger the embryo that develops and must be supported by the swimming mother colony. And, the larger the colony, the larger the investment needed for there to be any reproduction at all. This initial cost of reproduction is especially acute in species in which cells do not double in size and then undergo binary fission, but grow about $N = 2^d$ fold in size and then undergo a rapid synchronous series of d divisions (under the mother cell wall). This type of cell division, which is considered the ancestral developmental program in this lineage (Desnitski, 1995), is known as “palintomy” and is thought to have predisposed these algae to multicellularity (Kirk, 1998). It occurs in the smaller species (including *Chlamydomonas*, *Gonium*, *Eudorina* and *Pleodorina*) and in some of the G–S specialized *Volvox* species (e.g., *V. carteri*) (Fig. 1).

The assumption of selection at the group level (assumption (iii)) is likely to hold in volvocine algae because of their mode of reproduction and colony formation, in which all cells in the group are derived clonally from a single cell after a specific number of cell divisions, d ($d = 3$ for *Gonium*, $d = 5$ for *Eudorina*, $d = 6–7$ for *Pleodorina* and $d = 8–16$ for species of *Volvox*). We have previously studied the conditions under which multilevel selection may select for systems of conflict mediation that enhance selection at the group level (Michod, 1996, 1997, 1999; Michod and Roze, 1997, 1999; Michod et al., 2003). Another factor favoring selection at the group level is “parental control” of the cell phenotype, in which the behavioral phenotype (i.e., the cell fate) is determined during development by the “mother” cell. This is the case in *V. carteri*, as the cell fate (somatic or germ) is established early in development through a series of asymmetric cell divisions of the anterior blastomeres (for discussion see Michod et al. (2003)). It is well known that it is easier for cooperation to be maintained in a group under parental control than under offspring control (in which the phenotype is determined by the genotype of the cell), because the sacrifice of cooperation is spread over the different genotypes present in the cell group (see, for example, Michod (1982)).

Here we consider a group of cells and seek to understand the selective pressures that mold the allocation of energy and resources at the cell level to the two fitness components of the group, survival and reproduction. We present two models. In the fitness isocline model we consider whether, as groups increase in size, when a single new cell is added, it could increase the fitness of the group by changing its reproductive effort from what the existing cells in the group have been doing. In the full optimization model we consider whether a small change in behavior of one or several cells could increase the fitness of the group. The optimization model is clearly more general, but we begin with the fitness isocline model for heuristic reasons. Throughout, we seek qualitative results that are independent of the specific functions involved (so long as these functions meet the general assumptions stated: differentiability, concavity, convexity, etc.).

4. Fitness isocline model

Consider a group of $N - 1$ cells with a group viability V_{N-1} and fecundity B_{N-1} . We would like to predict the allocation of energy and resources to reproduction, e , and survival-enhancing vegetative functions, $1 - e$, for the N th cell, resulting in $b(e)$ and $v(1 - e)$ contributions to fecundity and viability, respectively. The variable e is the familiar reproductive effort variable of life-history theory (Stearns, 1992; Charlesworth, 1980). Since both b and v are assumed to be monotonic functions of e , we follow precedence in this area and generally work in terms of b and v directly (instead of in terms of e). We assume a simple additive model of fitness at the group level (termed “group selection

I'' by Damuth and Heisler (1988)) so that the fitness components of the group are the sum of the contributions of the cells, or considering the additional N th cell, the fitness of the group is $W = (V + v)(B + b)$, where $V = V_{N-1}$ and $B = B_{N-1}$. We suppress the group size subscript here and in what follows for notational simplicity. Additivity of cell contributions to viability and fecundity might apply, for example, to the simpler forms of volvocine algae considered in Fig. 1, in which cells stay together after cell division.

The new fitness of the group with the additional cell is then given by

$$W = bv + bV + vB + BV. \tag{1}$$

We would like to maximize the fitness, W_i , contributed by the new cell given by

$$W_i = bv + bV + vB. \tag{2}$$

For fixed W (fixed V and B), the fitness contributed by the new cell, W_i , is a function of two variables, b and v . We can plot isoclines for W_i by using Eqs. (3) and (4) to plot v as a function of b as done in Fig. 2:

$$v = \frac{W_i - bV}{b + B}, \tag{3}$$

$$v'(b) = -\frac{W_i + BV}{(b + B)^2} < 0 \quad \text{and}$$

$$v''(b) = 2\frac{W_i + BV}{(b + B)^3} > 0. \tag{4}$$

We note a few points about Fig. 2 that will be useful below. The isoclines are convex functions (first derivative increasing) which do not overlap and, for increasing fitness return, W_i , they occur increasingly farther from the origin. For any particular W_i , W_i/B and W_i/V are the maximal fitness that could be attained at the group level for viability and




Fig. 2. Fitness isoclines for the contribution of the new cell to viability, v , and fecundity, b , at the group level. Four isoclines are shown, the heavy solid line is the isocline of interest, the others are dashed. Tangents to the isocline are shown at the maximal contributions possible: W_i/V for reproduction and W_i/B for viability.

fecundity, respectively. Using Eq. (4), the slopes of tangents to the isocline at these points are indicated in Fig. 2.

In addition to the fitness relations at the group level graphed by the isoclines in Fig. 2, we assume there is an intrinsic relation that links b and v within the cell because of cell physiology and/or other constraints. We refer to this intrinsic relation as the “trade-off function”, as it embodies life-history trade-offs between the two fitness components at the cell level. During the origin of multicellularity, we expect these trade-offs to depend upon the size of the group that the cells must create (investigated below), as well as a host of other factors; but, for the moment, we consider the implications of the simple linear relation

$$v = v_{\max} - \alpha b. \tag{5}$$

As illustrated in Fig. 3A with a linear intrinsic function, the cell will likely invest in both reproductive and viability-related functions. Indeed, a simple inductive argument given in Appendix A shows that for a linear intrinsic function (Eq. (5)), cell groups have no incentive to specialize. No matter how large the group is, provided that the $N-1$ first cells exert intermediate reproductive effort at $b = b_{\max}/2$ and $v = v_{\max}/2$, it is optimal that the N th cell exerts the same effort (this yields the best unspecialized group; some specialized groups may achieve the same level of fitness, but not a higher one). In the case of a linear trade-off, the ratio of viability to fecundity at the group level is determined directly by the trade-offs at the cell level (as represented by α) and is given by

$$V = \alpha B. \tag{6}$$

Since α governs the basic relationship between survival and reproduction at the cell level, it imposes severe constraints on fitness components V and B at the group level (Eq. (6)). Indeed, as there is yet nothing else in the model that might change the relationship between viability and fecundity at the group level, we may expect that Eq. (6) will hold as the group increases in size, so long as we assume the linear constraint at the cell level (Eq. (5)). Below we consider a cost of reproduction to survival that increases as the size of the group produced by the cell increases. This cost changes the basic relationship between survival and reproduction at the group level from that given in Eq. (6), because certain fitness advantages of cell specialization may be realized by the group, but not the cell.

From the graphs in Fig. 3 (B and C), we may anticipate a central result of the model. Note that as the intrinsic curve becomes convex (meaning its derivative increases with b), the cell will specialize in viability (panel B) or fecundity (panel C) functions to attain the maximum fitness gains allowed at the group level. Such specialization in viability or fecundity functions is tantamount to the evolution of soma (panel B) and germ (panel C). In what follows, we approximate a convex intrinsic function in a piecewise linear fashion.




Fig. 3. Optimal investment strategy determined by the intrinsic functions and the fitness isocline. For a linear intrinsic curve (A), the new cell will perform a mix of viability and fecundity functions. For convex intrinsic curves, the new cell will more likely specialize, for example, in survival (B) or reproduction (C).




Fig. 4. Initial cost of reproduction. The piecewise convex curve in panel (C) is formed out of the functions in panels (A) and (B) and approximates the convex curves in Fig. 3. In panel (A) the reproductive effort e_N is the initial (or fixed) cost of reproduction. In panel (C) the quantity $v_{\max} - v_{\max}(1 - e_N)$ is the bonus to viability of soma specialization. This bonus is realized only in groups. The negative of the bonus may be referred to as the initial cost of reproduction to survival. See text for further explanation.

An initial investment is often necessary to get any reproduction. For example, growing the embryo inside the mother colony in the case of the volvocine algae takes time, energy and resources away from other functions (or a mating display, producing a flower, etc.). These initial reproductive costs tend to create a convex relationship between reproductive effort, e , and fecundity, $b(e)$, as depicted in Fig. 4A in a piecewise linear way. We assume that this initial cost of reproduction detracts from survival and so we term it “initial cost of reproduction to survival” or sometimes just “initial cost of reproduction”. This initial cost will depend on the group size N which the cell must produce, e_N . Combining this initial cost of reproduction (Fig. 4A) with a linear intrinsic function for viability (Fig. 4B) and using the construction given in Fig. 4, we obtain the piecewise linear intrinsic function relating v and b given in Eq. (7) and plotted in Fig. 4C. By varying the initial cost of reproduction, the piecewise linear curve in Eq. (7) (Fig. 4B) can approximate the convex curves graphed in Fig. 3 (panels B and C).

$$v = v_{\max}(1 - e_N) \left(1 - \frac{b}{b_{\max}}\right), \quad v \leq v_{\max}(1 - e_N),$$

$$v = v_{\max}(1 - e), \quad v > v_{\max}(1 - e_N). \quad (7)$$

In Fig. 4, the initial cost of reproduction to survival is the vertical portion of the intrinsic curve running along the v -axis from $v_{\max}(1 - e_N)$ up to v_{\max} . The modulus of this initial cost also equals the benefit to viability of soma specialization stemming from not having to pay the initial

cost of reproduction. For the volvocine green algae with palintomic development (Fig. 1), the initial cost of reproduction, e_N in Fig. 4, is directly related to the group size N which the reproductive cell must produce, and thus to the cell size the reproductive cell must attain before initiating the rapid series of embryonic divisions to create the daughter colony.

We note three points about this benefit of soma specialization, the “bonus” diagrammed in Fig. 4C. First, this bonus is only obtainable through group living and is only expressed at the group level, it is not an option open to solitary cells. Second, it changes the basic relationship that governs the fitness components at the cell level into a new relationship between viability and reproduction at the group level. Third, the benefit will likely change with the size and organization of the group. For example, if there are already many somatic cells in the group, the benefit of a new somatic cell may be small.

If we assume that at the colony size at which the initial cost of reproduction becomes operative, the ratio of survival to reproduction at the group level is α (as the linear constraint predicts), $V = \alpha B$, then a straightforward analysis is possible. In this case, we must have $Bv_{\max} > Vb_{\max}$, so germ specialization never pays as a first step from undifferentiated cells. Assuming $V = \alpha B$, a critical value for the initial cost of reproduction, e_N , can be derived to determine whether the new cell will specialize in somatic functions or remain undifferentiated. This critical cost is obtained by investigating the conditions when $Bv_{\max} > W_i^*$, where W_i^* is the maximal value of added

fitness obtained for optimal intermediate allocation to reproduction, b^* , and survival, v^* . The critical value of the initial cost of reproduction is given by

$$e_N^{crit} = \frac{b_{max}}{4B + b_{max}}. \quad (8)$$

If $e_N > e_N^{crit}$, soma specialization pays. Otherwise, the cell continues to allocate resources to both survival and reproduction.

By inspecting Eq. (8), we can see that the larger the group fecundity (B) is, the smaller the initial cost of reproduction may be for soma specialization to evolve. In other words, in colonies with larger fecundity (and, all things being equal, this means larger colonies) it is easier for a specialized and sterile soma to evolve. This may be explained as follows. The difference between the added fitness brought about by a cell specializing in soma and the added fitness brought about by a cell having the same fecundity, $b = b_{max}/2$, as the first $N-1$ cells is $(v_{max} - v_0)B - v_0 b_{max}/4$. The first term is the advantage of specialization linked to the initial cost of reproduction. The second term is the loss linked to the fact that specialization disrupts the balance between viability and fecundity (see Section 5.4.1 for more discussion). While this loss is independent of colony size, the advantage increases with colony fecundity. Therefore, it is more likely for the advantage of specialization to exceed the loss due to specialization for larger and already more fecund colonies.

The significance of the loss due to specialization is a result of the assumption that only one cell changes. In particular, if the cell reduces its reproductive effort, there must be a loss to fecundity that must be overcome for this specialization to pay in the overall group fitness. In the optimization model considered next, we allow two (or more) cells to simultaneously change their allocation strategy. If one cell increases and another cell decreases their respective reproductive efforts by the same amount, the total fecundity will remain the same (hence, there is no fecundity loss to the group), but gains in viability are possible under convex curvature.

5. Optimization model

5.1. Overview

We now apply optimization theory to the cell group, so as to consider all the cells simultaneously and study strategies in which cells jointly increase or decrease their reproductive effort so as to maximize the fitness of the group. In the fitness isocline model, we considered how a single new cell could maximize its fitness contribution to the group. In the optimization model we test whether small deviations by two or more cells could increase the fitness of the group. The stability conditions of the optimization model include, and are more general than, the stability conditions of the fitness isocline model.

5.2. The model

Consider groups of N cells, with cells indexed $i = 1, 2, \dots, N$. Let e_1, e_2, \dots, e_N be the reproductive effort for each cell, and let b_1, b_2, \dots, b_N be the resulting contribution to the fecundity of the group. As we did above, we assume the contribution to fecundity is an increasing function of reproductive effort; therefore, we can work in terms of fecundity, instead of reproductive effort. Let v_1, v_2, \dots, v_N be the vegetative, viability-enhancing capabilities of each cell. As more effort is put into reproduction, less is available for vegetative functions, resulting in a trade-off between the contributions to the fitness components of the group. We assume that if $b = b_{max}$ then $v = 0$, and if $b = 0$ then $v = v_{max}$. As above, for simplicity, we assume that the viability and fecundity of the group, V and B , respectively, are simple additive functions of the cell properties given by

$$B = \sum_{i=1}^N b_i \quad \text{and} \quad V = \sum_{i=1}^N v_i. \quad (9)$$

Note that while in the fitness isocline model V and B denoted the contribution to viability and fecundity of the first $N-1$ cells, here they denote the viability and fecundity of the whole colony. While it seems biologically reasonable to assume additivity of the contributions to fecundity of the group, additivity of the contributions to viability is more questionable. We have in mind a trait-like flagellar motility (or mixing) as a proxy for viability and assume there is a simple linear relationship between the effort or time a cell invests in flagellar action and the overall motility of the group. While this assumption may hold over a limited range, it would likely fail as the group gets larger and more integrated. We show in Appendix B.3 that we may dispense with the additivity assumption as it applies to viability so long as we maintain it for fecundity and still reach the same qualitative conclusions.

For our purposes, it is not necessary to normalize fitness, since the analysis of optimal behavior in the optimization model or in the fitness isocline model would not change. Normalizing fitness means multiplying the fitness we have by a coefficient which depends only on the size of the colony. When we ask, for a colony of a particular size, should cells specialize or remain generalists, normalization would not change the answer, because the maxima for the normalized or not-normalized fitness functions would be the same.

We assume that group fitness, W , is the product of viability and fecundity. This is appropriate for a life cycle involving discrete generations as is the case with the volvocine green algae.

$$W = VB. \quad (10)$$

Although the multiplicative decomposition of fitness into viability and fecundity assumed in Eq. (10) applies when generations are discrete, most of the qualitative points

made in the following sections would still hold were fitness a more general function $W(V, B)$ which was nonnegative, zero if and only if $V = 0$ or $B = 0$, and strictly increasing in both arguments whenever V and B are both positive. In particular, the fundamental point that cell specialization allows the group to increase fitness under conditions of convexity holds for this more general fitness function (because specialization can retain the same group fecundity while increasing viability).

As already mentioned, the additivity assumed in Eq. (9) is an example of group selection of type I as discussed by Damuth and Heisler (1988). However, there are interesting implications of combining the fitness components at the group level after first summing the cell contributions (as assumed in Eqs. (9) and (10)). Most important (and critical to our analysis below) is the fact that, if one cell has a high fecundity (and hence a low viability, so that it would have a low fitness by itself), this may be compensated for if another cell has a high viability (and hence low fecundity). Consequently, even though each of these cells by itself would have a low fitness, together they can bring a high fitness to the group (especially under conditions of convexity of the trade-off). This kind of joint effect is a first step towards integration of the group, and would not be possible if we used as group fitness the average cell fitness, $(1/N)\sum_{i=1}^N v_i b_i$.

More formally, the normalized fitness, VB/N^2 , is greater than the average cell fitness by the negative of the covariance between the two fitness components. Since in our case the covariance is negative, the normalized fitness associated with Eq. (10), VB/N^2 , is greater than the average cell fitness, $(1/N)\sum_{i=1}^N v_i b_i$, by the magnitude of the covariance between fitness components. This covariance effect at the group level appears to be quite general. Its contribution to a property like fitness depends on the property being a multiplicative function (or some other function requiring a strong balance) of two components (e.g., viability and fecundity) which themselves covary so that higher values of one component bring lower values of the other (the trade-off principle). Of course, if there is no variance in these components among the lower level units (cells) then there is no covariance and no effect at the group level. What factors might produce variance among the lower level units? We can think of two factors: noise, and the curvature of the trade-off function being convex.

5.3. Implications of different curvatures of the trade-off function

When investigating the implications of the different possible curvatures of the trade-off function, $v(b)$, we will repeatedly make use of the definitions of convex and concave functions. For a strictly convex (concave) function $v(b)$, if we take a particular point, say b^* , and two points equidistant below and above b^* , say b^- and b^+ , respectively, then $v(b^-) + v(b^+) > (<) 2 v(b^*)$. If b is fecundity and $v(b)$ viability, then convexity of v implies that there is an

advantage to specializing in the two components of fitness, while concavity implies there are diminishing returns on an investment in either component. We first assume that there is no initial cost of reproduction.

5.3.1. Concave trade-off

If the function $v(b)$ is strictly concave, then the cell group should remain unspecialized. More precisely, all cells should exert the effort b^* that maximizes the product $bv(b)$. The key to this result is to observe that if two cells, i and j , have different reproductive efforts, $b_i \neq b_j$, then they could both change their fecundities to $(b_i + b_j)/2$. This change in reproductive effort would not change the overall fecundity of the group but would (by definition of concavity) increase group viability, and hence increase fitness. Indeed, the change in group viability would be $\Delta V = 2v((b_i + b_j)/2) - v(b_i) - v(b_j)$, which is positive because $v(b)$ is a strictly concave function. This shows that all cells should exert the same effort. If this common effort is b then the viability is $V = Nv(b)$, the fecundity is $B = Nb$ and the fitness is $W = N^2bv(b)$. Thus, independently of the number of cells in the colony, the optimal value of the cell fecundity common to all cells is the one that maximizes the product $bv(b)$. The result that the optimal fecundity for each cell does not depend on the number of cells is likely not robust and depends crucially on the assumptions that viability is additive and that the trade-off function $v(b)$ does not depend on the size of the colony. In contrast, the result that a concave trade-off function selects against specialization is robust, as we show in Appendix B.

5.3.2. Linear trade-off

If the function $v(b)$ is linear ($v(b) = v_{\max} - \alpha b$, as in Eq. (5)) then the group viability only depends on the group fecundity (and not on the particular values of the component cell fecundities). Indeed, we have $V = \sum_i v_i = \sum_i (v_{\max} - \alpha b_i) = Nv_{\max} - \alpha B$.

Thus, any values of the fecundities b_1, \dots, b_N leading to the same global fecundity B yield the same fitness $W = (Nv_{\max} - \alpha B)B$. The possible values of B range from 0 to Nv_{\max}/α . The maximum fitness is obtained for $B = Nb_{\max}/2$ (hence $V = Nv_{\max}/2$) and is equal to $W^* = N^2 b_{\max} v_{\max}/4$. Any arrangement of the fecundities b_1, \dots, b_N such that $\sum_i b_i = Nb_{\max}/2$ is optimal and these are the only optimal choices of the fecundities. In particular, assuming that the $N-1$ first cells have a fecundity $b = b_{\max}/2$, then it is optimal for the N th cell to exert the same effort, which yields the first result of the fitness isocline model. Formally, the group of cells behaves as if there was just one cell. There is no incentive to specialize and so no individuality at this stage.

5.3.3. Convex trade-off

If the function $v(b)$ is strictly convex, then the vast majority of cells will specialize (some in soma, some in germ). It may be that, at most, one cell remains unspecialized if, for example, there is an odd number of

cells in the group. If there is an even number of cells, then half should specialize in germ and half in soma. Indeed, this yields a fitness of $W = \frac{1}{4}N^2v_{\max}b_{\max}$, which is the same fitness obtained in the more favorable linear case considered above with the same values of b_{\max} and v_{\max} . The linear trade-off is more “favorable” in the sense that, for any value of the contribution to fecundity b , the contribution to viability $v(b)$ is higher or equal in the linear case than in the convex case. As a result, for a convex trade-off, complete specialization in equal proportions must be optimal since it attains this highest possible fitness. If there is an odd number of cells, it may be that one cell remains unspecialized (for instance, when there is just one cell), but at most one cell may remain unspecialized. Indeed, assume that two cells i and j have an intermediate fecundity; without loss of generality assume $b_i \leq b_j$. Let δ be positive and smaller than both b_i and $b_{\max} - b_j$. If cell i decreases its fecundity while cell j increases its fecundity by the same amount δ , then the global fecundity of the colony does not change. However, the viability increases, hence the fitness increases. Indeed, the change in viability is $\Delta V = v(b_i - \delta) + v(b_j + \delta) - (v(b_i) + v(b_j))$, which is positive due to the strict convexity of the function $v(b)$, as shown in Appendix B.1.

5.3.4. Neither convex nor concave

It might be that the function $v(b)$ is neither concave nor convex. In that case, in the absence of additional information, whether specialization pays cannot be decided. Some partial results may be obtained though. For instance, assuming that $v(b)$ is differentiable, if at a fitness maximum a cell i has an intermediate fecundity b_i , then we must have $\partial v / \partial b(b_i) = -V/B$ (taking the derivative of fitness with respect to fecundity of cell i and setting it to zero using Eqs. (10) and (9)). Also, if $v(b)$ is twice differentiable, then at a fitness maximum at most one cell may have an intermediate fecundity b' such that the trade-off function is locally strictly convex at b' (that is, $\partial^2 v / \partial^2 b(b') > 0$). This generalizes the above result on convex trade-offs. The proof of this result (omitted for brevity) consists in differentiating the fitness function and investigating the standard first- and second-order optimality conditions.

5.4. Initial cost of reproduction

We now investigate the effect of an initial cost of reproduction. Formally, letting v_{\max} be the contribution to viability of a completely specialized somatic cell ($b = 0$) and letting v_0 be the limit of $v(b)$ when the fecundity $b > 0$ tends to zero, we assume that v_{\max} is greater than v_0 and study how this modifies the results of the preceding section. Recall that $v_{\max} - v_0$ is the bonus discussed in Fig. 4. Intuitively, an initial cost of reproduction makes the trade-off function more “convex-like”, and thus tends to select for specialization.

5.4.1. Convex or linear trade-off

If the trade-off function $v(b)$ is strictly convex, then adding an initial cost of reproduction only reinforces the conclusion that specialization should be favored. More interesting are the cases of a linear trade-off (considered now) or a concave trade-off (considered below). Recall that in the linear case, without an initial cost of reproduction, the maximal fitness may be obtained with or without specialization. Thus, in the previous case, the colony is indifferent to specialization or no specialization. An initial cost of reproduction, however small, tilts the balance in favor of specialization. Indeed, assume that for $b > 0$, $v(b) = v_0 - \alpha b$ with $v_0 = \alpha b_{\max} < v_{\max}$. If no cell specializes, then, as discussed in the preceding section, the best fitness that the colony can obtain is $W^* = N^2 b_{\max} v_0 / 4$. This will be the case when the group fecundity is $B^* = N b_{\max} / 2$, and the group viability is $V^* = N v_0 / 2$. We now apply the argument above concerning a pair of cells and whether they might jointly specialize by considering a colony in which $N-2$ cells have fecundity $b = b_{\max} / 2$, and in which the two other cells are specialized, one in soma and one in germ. The overall fecundity of the colony is still B^* , but the viability is now $V = (N-2)v_0/2 + v_{\max} = N(v_0/2) + v_{\max} - v_0 > V^*$; hence, the fitness is greater than W^* . Therefore, if some of the cells, specialize, the colony can obtain a greater fitness than if all cells are generalists. If there is an even number of cells then, as in the case of a convex trade-off and no initial cost of reproduction, half of the cells should specialize in soma and half in germ.

Note that, with a linear trade-off, specialization occurs as soon as there is any initial cost of reproduction. This contrasts with the fitness isocline model, in which specialization requires that the initial cost of reproduction be greater than a critical value (Eq. (8)). To understand this difference, note that in the fitness isocline model we assumed that the behavior of the $N-1$ first cells was fixed. Thus, if the N th cell specializes in soma, this yields a benefit to viability (the bonus to specialization linked to the initial cost of reproduction), but disrupts the balance between fecundity and viability. For specialization to be optimal in the fitness isocline model, the benefits must outweigh the costs. In the more general optimization model considered here, if a cell specializes in soma, the other cells may increase their reproductive effort in order to compensate for the corresponding loss of fecundity. Thus, the group may obtain the benefits of specialization without having to pay for a disruption in the balance between fecundity and viability.

5.4.2. Concave trade-off

We assume now that the trade-off function $v(b)$ is strictly concave ($0 < b \leq b_{\max}$), but that there is an initial cost of reproduction, so that $v_{\max} > v_0$ (Fig. 5). As in the preceding section, if two cells i and j which are not specialized in soma have a different fecundity, e.g. $b_i > b_j > 0$, then by changing their fecundity to $(b_i + b_j)/2$, they would retain the same overall contribution to group fecundity while increasing

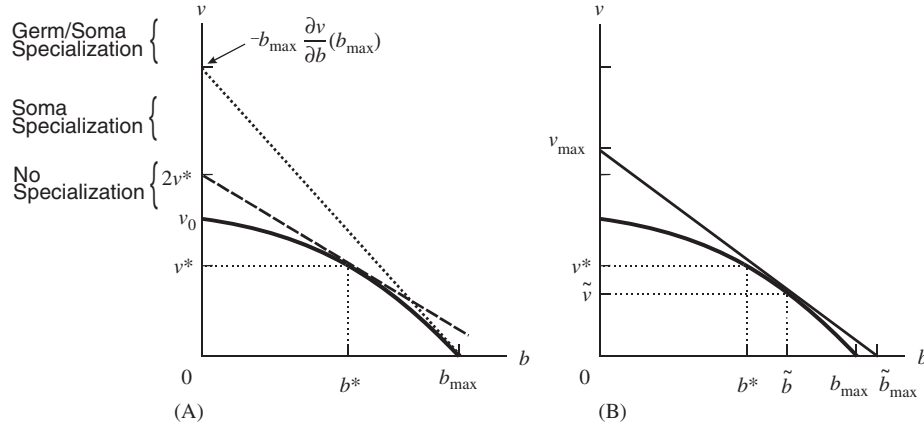


Fig. 5. Evolution of soma and germ with concave trade-off and cost of reproduction. The kinds of colonies predicted by the model are given along the ordinate in panel (A) for different regions of v_{\max} . The quantities b^* and v^* are the fecundity and viability, respectively, of the cells in the optimal unspecialized colony. The solid concave curve (identical in both panels) is the trade-off between viability and fecundity. Three tangents relevant to the analysis are drawn. In panel (A), the dotted and dashed lines are the tangents to the trade-off curve at $b = b_{\max}$ and $b = b^*$, respectively. In panel (B), the dashed line is the tangent to the curve at $b = \tilde{b}$. The significance of these points is discussed in the text and Appendix B.2. The quantity $2v^* - v_0$ is the detriment to soma specialization due to the concavity of the trade-off. If $v_{\max} > 2v^*$, then specialization in soma pays. See text for further explanation.

their contribution to group viability. It follows that, at a fitness maximum, all cells which are not specialized in soma must have the same fecundity. An immediate consequence is that specialization in germ without specialization in soma cannot pay as a first step. The question is whether specialization of some cells in soma allows for an increase in fitness.

Assuming that the function $v(b)$ is twice differentiable, an analysis given in Appendix B.2 gives the conditions under which soma specialization will evolve, and, if soma evolves, whether germ specialization will also evolve. We now summarize those conditions referring to Fig. 5. Let b^* denote the fecundity of the cells in the best unspecialized colony (i.e. b^* is the fecundity which maximizes the product $bv(b)$). If Eq. (11) holds, specialization does not pay (the best colony is the one in which all cells have intermediate fecundity b^*).

$$v_{\max} \leq 2v^*. \quad (11)$$

Since b^* maximizes $bv(b)$, we have $v^* + b^*(\partial v/\partial b)(b^*) = 0$. It follows that Eq. (11) is equivalent to

$$v_{\max} \leq v^* - b^* \frac{\partial v}{\partial b}(b^*). \quad (12)$$

Eq. (12) means that the tangent to the trade-off curve at $b = b^*$ crosses the line $b = 0$ below v_{\max} . On the other hand, if the reverse of Eq. (12) (given in Eq. (13)) holds, then some cells should specialize in soma.

$$v^* - b^* \frac{\partial v}{\partial b}(b^*) = 2v^* < v_{\max}. \quad (13)$$

Furthermore, if, in addition to satisfying Eq. (13), $v_{\max} < -b^*(\partial v/\partial b)(b^*)$, then the reproductive cells should have an intermediate fecundity; more precisely, they should have fecundity \tilde{b} such that the tangent to the trade-off curve at $b = \tilde{b}$ crosses the line $b = 0$ precisely at v_{\max} (Fig. 5

panel (B)) (that is $v_{\max} = \tilde{v} - \tilde{b}(\partial v/\partial b)(\tilde{b})$, see Appendix B.2). The proportion of somatic cells should be such that the mean viability is $v_{\max}/2$. The mean fecundity is then $\tilde{b}_{\max}/2$, $\tilde{b}_{\max} = -v_{\max}/[(\partial v/\partial b)(\tilde{b})]$ and the fitness $W = N^2 v_{\max} \tilde{b}_{\max}/4$, which is the highest obtainable fitness for a colony facing a linear trade-off $v = v_{\max} - \alpha b$ with $\alpha = -\partial v/\partial b(\tilde{b})$.

Finally, if

$$v_{\max} \geq -b_{\max} \frac{\partial v}{\partial b}(b_{\max}), \quad (14)$$

then half of the cells should specialize in soma ($b = 0$) and the other half in germ ($b = b_{\max}$). This yields the fitness $W = N^2 v_{\max} b_{\max}/4$.

The above analysis shows that there are two threshold values for v_{\max} (in the region $v_{\max} > v_0$ (see Fig. 5)). The first threshold, $v_{\max} = 2v^* = v^* - b^*(\partial v/\partial b)(b^*)$, concerns specialization in soma. The second threshold, $v_{\max} = -b_{\max}(\partial v/\partial b)(b_{\max})$, concerns specialization in germ. In the case of a linear trade-off, both thresholds are equal to v_0 . Thus, while $v_{\max} - v_0$ is the bonus of soma specialization linked to the cost of reproduction, we may see $2v^* - v_0$ as the detriment of soma specialization linked to the concavity of the trade-off. Eq. (11) (or, equivalently, Eq. (12)) expresses the condition that the bonus of soma specialization is smaller than its detriment.

As detailed in Appendix B.2, if we take into account the fact that the proportion of somatic cells must be a multiple of $1/N$, then the condition for soma specialization is slightly more demanding than Eq. (13) and is more easily satisfied if the number of cells in the colony is large. This effect of colony size in facilitating the evolution of soma specialization may be interpreted as follows: there is a tension between being efficient (that is, having high ratios of viability/(resources allocated to viability) and fecundity/(resources allocated to fecundity)) and keeping a balance

between viability and fecundity. If a cell specializes in soma, then in order to keep a balance between fecundity and viability, the other cells must increase their fecundity. If there are only a few other cells in the colony, then they will have to increase their fecundity a lot, possibly moving to inefficient functioning points. This is likely to be the case when there are decreasing returns on efforts (that is, when the fecundity/(resources to fecundity) ratio decreases as the resources allotted to fecundity increase, and similarly for viability), which corresponds to a concave trade-off. In contrast, in a large colony where there are more cells available, a balance between viability and fecundity can more easily be maintained since the increase in reproductive effort by the rest of the colony can be divided among more cells, which would then not greatly effect each cell's efficiency. Thus, it is possible to reap the advantage of soma specialization linked to the cost of reproduction while maintaining both a balance between viability and fecundity and an efficient functioning of all cells.

6. Discussion

We have concluded that during the origin of multicellularity convex trade-offs between survival and reproduction, such as those created by a significant initial cost of reproduction to survival, select for specialization in the two fitness components as colonies increase in size. As a result of this specialization, the individuality of the cell group is enhanced. The conclusion about the role of convexity in specialization is very general; in particular, we have shown it holds for more general fitness functions than multiplicative (Eq. (10)) and for non-additive viabilities. The main point is that when the trade-off is convex, specialization allows for the increase of one component of fitness (we focused on viability) without any decrease in the other component (fecundity).

The conclusion that convexity favors specialization resembles the standard results of life-history theory, which state that convex fitness trade-offs select for specialization in reproductive function as organisms increase in age. That is, convex fitness trade-offs select for semelparity or “big bang” reproduction in which there is no reproduction until the last stage of the life cycle (Schaffer, 1974; Charlesworth and Leon, 1976). Big bang reproduction is analogous to cell specialization in the sense that age classes specialize in either no reproduction or complete reproduction (for the last class).

However, the life-history problem of optimization of reproductive effort over the lifespan of an individual is different in important ways from the problem of optimization of the reproductive effort of cells in a group. There is, most fundamentally, the very question we wish to answer: is the individual the cell or the cell group? While we do assume selection at the group level, without cell specialization there is no property that would make the group indivisible and hence a true individual. In answering this question, we are not concerned with how the reproductive

effort at the group level changes, but rather with whether there is specialization at the lower level among cells. Indeed, in our argument for cell specialization in the optimization model, the average reproductive effort at the group level does not change at all. However, how this effort is distributed among cells can be critical for the group in terms of its viability and individuality.

Consider, for example, the following question in life-history theory which may seem similar to the one studied here. How will the trade-off between viability and fecundity evolve as an external parameter, such as the quality of the environment, varies (Kisdi et al., 1998)? Here we also investigate the evolution of the trade-off between viability and fecundity as another parameter, the size of the group, varies. However, the question we investigate is not how the investment in fecundity will evolve as this additional parameter varies, as in Kisdi et al. (1998), but whether some cells will specialize (again, the overall investment in fecundity staying more or less the same).

On a qualitative level, what we have studied is how the relative changes in viability and fecundity linked to cell specialization evolve as colony size increases. Due to the assumed multiplicative nature of fitness (Eqs. (1) and (10)), $W = VB$, what matters when a cell changes its reproductive effort are the relative changes in viability and fecundity for the group. Formally, if a cell specializes in soma, leading to a decrease in group fecundity of δB and an increase in group viability of δV , then the change in group fitness is $\Delta W = -V\delta B + B\delta V - \delta B\delta V$. If we neglect the last term (which is a second-order term), then we see that the condition for fitness to increase is that the relative increase in viability $\delta V/V$ be greater than the relative decrease in fecundity $-\delta B/B$. This is more likely to be the case if viability is low, as will occur if colonies increase in size without specializing in somatic functions.

What is the fecundity viability trade-off curve like in single-celled organisms? The multiplicative nature of fitness requires that single-cell organisms be generalists and have intermediate efforts at both reproduction and viability, regardless of the curvature of the trade-off curve. Nevertheless, the curvature of the trade-off determines whether the unicellular habit will be stable to two- (or greater) cell groups. This will be the case when the trade-off curve is strictly concave. Since, in nature, smaller groups are not specialized (Fig. 1), our model suggests that the trade-off is concave rather than convex in single-celled species. Furthermore, a concave trade-off seems more natural for small groups, as it expresses a law of decreasing return on efforts.

The curvature of the function describing the relationship between the two main fitness components' reproduction and survival is a basic issue in life-history theory (Benkman, 1993; Michod, 1978; Schaffer, 1974; Benson and Stephens, 1996; Blows et al., 2004; Carriere and Roff, 1995; Kisdi, 2001; Reznick, 1985; Roff, 2002; Rueffler et al., 2004; Sato, 2002; Strohm and Linsenmair, 2000; Takada and Nakajima, 1996; Levins, 1968; Stearns, 1992). Despite

the central relevance of this issue to life-history theory, a recent review (Rueffler et al., 2004) of the data concerning the curvature of the trade-off curve states: “Unfortunately, there is no study known to us which has revealed the details of this curvature for any life-history trade-off in a specific organism. However, these curvatures are central in life-history theory which indicates a major gap between theory and empirical knowledge...”.

Our analysis predicts that a large initial cost of reproduction to survival is sufficient to select for G–S specialization. Measuring this initial survival cost of reproduction is empirically more practical than measuring the complete curvature of the trade-off curve between survival and reproduction, the latter having been studied in a variety of organisms with no clear results (if the above quotation is accepted). We return now to the volvocine green algae (Fig. 1) and consider this central prediction of our model.

As already discussed when introducing the volvocine algae (Fig. 1), the investment of the parent colony in reproductive cell growth illustrates an initial cost of reproduction to survival, which increases with organism size. Besides using more resources, a larger embryo increases the volume, mass and drag of the mother colony, as has been quantified in Fig. 4 of Solari et al. (2005b). Solari et al. (2005b) show that these initial survival costs increase with colony size, requiring more swimming force as well as more flagellar mixing (for nutrient acquisition and removal of waste) per embryo.

We believe the need to pay this initial cost of reproduction to survival accounts for the observed increase in the somatic/reproductive (S/R) cell ratio as colony size increases in the volvocine algae (see Table 3 of Solari et al., 2005b). We think that the evolution of soma (as well as the evolution of increased S/R ratios) provides the benefits that compensate for the increasing initial costs of reproduction in colonies of increasing size. There are also direct costs of germ and soma specialization which must be overcome by these benefits, as germ specialization reduces the number of cells available for vegetative functions and soma specialization reduces the number of reproducing cells.

The benefits of soma specialization include: (i) colony motility while reproducing (overcoming the flagellation constraint discussed in Section 3), (ii) motility while large (overcoming the enlargement constraint discussed in Section 3), (iii) increased resource uptake due to the ‘source-sink’ effect (in which somatic cells transfer resources to germ cells which act as a sink) (Bell, 1985; Koufopanou and Bell, 1993; Solari et al., 2005a), and (iv) enhanced uptake of resources and removal of waste by flagellar beating (Niklas, 1994, 2000; Solari et al., 2005a).

In addition, soma specialization reduces the detriment to viability of germ specialization. Once larger colonies invest in a high proportion of somatic cells, non-somatic cells can focus on reproduction rather than contribute to vegetative functions which are sufficiently dealt with by somatic cells. When soma separation is complete, germ specialization can

provide additional benefits, such as decreased generation time, increased productivity by specialization at photosynthesis, and hydrodynamic advantages stemming from the location of germ. Since specialized germ cells are non-flagellated and do not contribute to motility, they are located in the interior of the colony, making the colony spheroid smaller and lowering drag (Solari et al., 2005b).

Single gene mutations in life-history traits can be a powerful approach to understanding the cost of reproduction and trade-offs between life history traits, both longstanding topics of considerable interest (Reznick, 1985; Roff, 2000, 2002). Various *V. carteri* developmental mutants are known (Kirk, 1998), which differ in the basic factors hypothesized in our models for the origin of multicellularity: group size, S/R ratio, type and timing of G–S specialization, and motility; yet they differ in just one or a few genes. These mutants include *lag*[−] (germ cells perform motility functions before reproducing; these mutant colonies are similar to *Volvox* species such as *V. aureus* and *V. rousellei*), *regA*[−] (somatic cells regenerate to become reproductive), and *glsA*[−]/*regA*[−] (all cells perform vegetative functions first and then become reproductive; this mutant is similar to *Eudorina*; see Fig. 1).

These mutants are especially useful for studying fitness decomposition at the cell and group levels, because a certain known number of cells (or amount of tissue) have changed their reproductive effort. We can measure the consequences of this change at the colony level, and in this way estimate the contribution to the group fitness of the changed effort at the cell level as is required by our model in Fig. 4. In the *regA*[−] mutants, ~235 cells have changed their phenotype from somatic to unspecialized; in *lag*[−] ~9 cells have changed their phenotype from germ to unspecialized; and in *glsA*[−] *regA*[−] there are ~561 unspecialized cells — similar to a *Eudorina* colony, but larger.

As a result of these changes in reproductive effort at the cell level, the size, productivity and motility of the group change (Solari et al., 2005b; Solari, 2005). For example, in colonies with the *regA*[−] mutation, as once-specialized somatic cells ($b = 0$ in Fig. 4) begin exerting reproductive effort ($b > 0$), there is not only a large decrease in colony motility, but also a large decrease in the motility contributed by a single changed cell. Specifically, the average force exerted for group motility by a single motile cell is about half in the *regA*[−] mutant and a quarter in the *glsA*[−] *regA*[−] mutant compared to the wild type (Solari et al., 2005b). The initial cost of reproduction to survival that underlies the convex nature of the fitness trade-offs (Fig. 4) is real and directly measurable in these organisms, and attributable to a change in the effort exerted by single cells within the cell group as required by the models considered above.

Acknowledgments

This research was supported by the NSF and the University of Arizona College of Science. The authors

appreciate the thoughtful advice of two anonymous reviewers.

Appendix A

We use Eq. (5) and build the group one cell at a time, assuming that after a cell is added it does not change its strategy. For the first cell, $V = B = 0$ and the isoclines are simply $v = W_i/b$. Maximizing W_i (Eq. (3)) subject to Eq. (5) gives the optimum strategy for the cell defined by

$$b^* = \frac{v_{\max}}{2\alpha} = \frac{b_{\max}}{2} \quad \text{and} \quad v^* = \frac{v_{\max}}{2}. \quad (\text{A.1})$$

The additional fitness (Eq. (3)) is now given by

$$W_i^* = \frac{b_{\max}v_{\max}}{4}. \quad (\text{A.2})$$

We now add the second cell. The isocline (Eq. (2)) now has $V_1 = v_{\max}/2$ and $B_1 = b_{\max}/2$. Note that $V_1 = \alpha B_1$. For the second cell, maximizing the added fitness (Eq. (3)) subject to the linear constraint (Eq. (5)) gives

$$b^* = \frac{b_{\max}}{2} \quad \text{and} \quad v^* = \frac{v_{\max}}{2}, \quad (\text{A.3})$$

with added fitness given by

$$W_i^* = \frac{3}{4} V_1 B_1. \quad (\text{A.4})$$

The new 2-cell group has $V_2 = 2v_{\max}/2 = v_{\max}$ and $B_2 = 2b_{\max}/2 = b_{\max}$ and Eq. (A.4) holds since, of course, $v_{\max} = \alpha b_{\max}$. If we now consider the $N+1$ cell and maximize W_i subject to the linear constraint (Eq. (5)), and that for the N cell group $V_N = \alpha B_N$ (Eq. (A.2)), we find again Eq. (A.1) and $B_{N+1} = B_N + b_{\max}/2 = (N+1)b_{\max}/2$ and $V_{N+1} = V_N + v_{\max}/2 = (N+1)v_{\max}/2$. The optimal value of intermediate reproductive effort can be obtained by maximizing W_i subject to the linear intrinsic constraint. For use in the text in deriving Eq. (8), we assume $V = \alpha B$. When there is a cost of reproduction, using the piecewise linear curve defined in Fig. 4, we obtain the optimal intermediate values to be

$$b^* = \frac{b_{\max}}{2} \quad \text{and} \quad v^* = \frac{(1 - e_N)v_{\max}}{2}. \quad (\text{A.5})$$

Using the values in Eq. (A.5), we obtain as the maximal added fitness for intermediate strategies

$$W_i^* = \frac{(4B + b_{\max})(1 - e_N)v_{\max}}{4}. \quad (\text{A.6})$$

Appendix B

B.1. Proof of a result on convex trade-offs

Assume that $v(b)$ is strictly convex. Let $b_i < b_j$ be intermediate fecundities and let $\delta > 0$ be smaller than b_i and $b_{\max} - b_j$. From the mean-value theorem, it follows that there exist fecundities b' in $[b_i - \delta, b_i]$ and b''

in $[b_j, b_j + \delta]$ such that

$$v(b_i) = v(b_i - \delta) + \delta \frac{\partial v}{\partial b}(b') \quad \text{and}$$

$$v(b_j + \delta) = v(b_j) + \delta \frac{\partial v}{\partial b}(b'').$$

It follows that

$$\begin{aligned} v(b_i - \delta) + v(b_j + \delta) - (v(b_i) + v(b_j)) \\ = \delta \left[\frac{\partial v}{\partial b}(b'') - \frac{\partial v}{\partial b}(b') \right]. \end{aligned} \quad (\text{B.1})$$

Since the function v is strictly convex, its first derivative is strictly increasing. Therefore, since $b'' > b'$, Eq. (B.1) is positive, as claimed in Section 5.3.3. The result still holds if the function v is not differentiable (proof omitted).

B.2. Proof of results on concave trade-offs with an initial cost of reproduction

We compute here the optimal behavior of a colony facing a strictly concave trade-off with an initial cost of reproduction. We first assume, as an approximation, that the proportion of somatic cells can take any value. We then discuss how taking into account the fact that the proportion of somatic cells must be a multiple of $1/N$ changes the results. Throughout, we assume that the contributions to viability are additive (the case of non-additive viabilities is treated in Appendix B.3). We also assume for simplicity that the function $v(b)$ is twice differentiable, so that the fitness function is twice differentiable. This allows us to use the standard first- and second-order optimality conditions of optimization theory (see below). However, this assumption is not necessary: as discussed at the end of this section, a graphical analysis shows that the same conclusions may be reached if the function v is not differentiable.

The fitness of a colony with a proportion p of somatic cells and fecundity b for all other cells is

$$W(p, b) = N^2 [p v_{\max} + (1 - p)v(b)] \cdot (1 - p)b. \quad (\text{B.2})$$

As explained in Section 5.4.2, in the best colony, all non-somatic cells have the same fecundity. Thus, to find the best colony, we only need to find the values of p and b that maximize Eq. (B.2). Let \tilde{p} and \tilde{b} be such optimal values and let $\tilde{v} = v(\tilde{b})$. Specialization of some cells in soma is optimal if $\tilde{p} > 0$. Specialization of the non-somatic cells in germ is optimal if $\tilde{b} = b_{\max}$. Recall that b^* denotes the optimal value of the fecundity for a non-specialized colony. If specialization in soma does not pay ($\tilde{p} = 0$), then $\tilde{b} = b^*$; but, if specialization in soma pays, then we expect (and we will prove) that the optimal fecundity of the non-somatic cells \tilde{b} is greater than b^* .

Since the values \tilde{p} and \tilde{b} are optimal, we have

$$\frac{\partial W}{\partial b}(\tilde{p}, \tilde{b}) = 0 \quad (\text{B.3})$$

if $\tilde{b} < b_{\max}$,

$$\frac{\partial W}{\partial b}(\tilde{p}, \tilde{b}) \geq 0 \quad (\text{B.4})$$

if $\tilde{b} = b_{\max}$, and finally, if $\tilde{p} > 0$,

$$\frac{\partial W}{\partial p}(\tilde{p}, \tilde{b}) = 0. \quad (\text{B.5})$$

To use these conditions, we need to compute the partial derivatives of W . We get

$$\frac{\partial W}{\partial p}(p, b) = N^2 b [(1 - 2p)v_{\max} - 2(1 - p)v(b)], \quad (\text{B.6})$$

$$\begin{aligned} \frac{\partial W}{\partial b}(p, b) &= N^2(1 - p) \left[(1 - p) \left[b \frac{\partial v}{\partial b}(b) + v(b) \right] + p v_{\max} \right]. \end{aligned} \quad (\text{B.7})$$

From Eqs. (B.5) and (B.6) we obtain

$$\tilde{p} = \frac{v_{\max} - 2\tilde{v}}{2[v_{\max} - \tilde{v}]} \quad (\text{B.8})$$

and

$$\tilde{p} v_{\max} = (1 - \tilde{p})(v_{\max} - 2\tilde{v}). \quad (\text{B.9})$$

Furthermore, plugging Eq. (B.9) into Eq. (B.7) we get

$$\frac{\partial W}{\partial b}(\tilde{p}, \tilde{b}) = N^2(1 - \tilde{p})^2 \left[\tilde{b} \frac{\partial v}{\partial b}(\tilde{b}) + (v_{\max} - \tilde{v}) \right]. \quad (\text{B.10})$$

Therefore, Eqs. (B.3) and (B.4) imply that

$$v_{\max} \geq \tilde{v} - \tilde{b} \frac{\partial v}{\partial b}(\tilde{b}), \quad (\text{B.11})$$

with equality if $\tilde{b} < b_{\max}$.

We now distinguish three cases which correspond, respectively, to no specialization, specialization in soma but not in germ, and specialization in soma and in germ. Let $v^* = v(b^*)$.

Case 1: $v_{\max} \leq 2v^*$. It follows from Eq. (B.8) that for specialization to be optimal, i.e. $\tilde{p} > 0$, we must have

$$v_{\max} > 2\tilde{v}. \quad (\text{B.12})$$

Since $v_{\max} \leq 2v^*$, this implies $v^* > \tilde{v}$, hence $\tilde{b} > b^*$.

Since the function v is concave, it follows that

$$\tilde{v} - \tilde{b} \frac{\partial v}{\partial b}(\tilde{b}) > v^* - b^* \frac{\partial v}{\partial b}(b^*).$$

Together with Eq. (B.11), this implies that

$$v_{\max} > v^* - b^* \frac{\partial v}{\partial b}(b^*) = 2v^*$$

(the latter equality is proved in Section 5.4.2). This contradicts the assumption $v_{\max} \leq 2v^*$. It follows that if $v_{\max} \leq 2v^*$, then specialization is not optimal.

Case 2: $2v^* < v_{\max} < -b_{\max}(\partial v/\partial b)(b_{\max})$. Applying Eq. (B.6) at $p = 0$ and $b = b^*$ we obtain

$$\frac{\partial W}{\partial p}(0, b^*) = N^2 b^* (v_{\max} - 2v^*) > 0. \quad (\text{B.13})$$

Thus a colony with some somatic cells and fecundity b^* for the non-somatic cells would have a higher fitness than the colony in which all cells have fecundity b^* (i.e. the best unspecialized colony). This implies that at a fitness maximum, the proportion of somatic cells is positive: $\tilde{p} > 0$. Furthermore, since

$$v_{\max} < -b_{\max} \frac{\partial v}{\partial b}(b_{\max}) = v(b_{\max}) - b_{\max} \frac{\partial v}{\partial b}(b_{\max})$$

it follows that if $\tilde{b} = b_{\max}$, then Eq. (B.10) is negative, contradicting Eq. (B.4). Therefore, $\tilde{b} < b_{\max}$, i.e. the non-somatic cells should be generalists. This implies that Eq. (B.11) holds with equality. That is,

$$v_{\max} = \tilde{v} - \tilde{b} \frac{\partial v}{\partial b}(\tilde{b}). \quad (\text{B.14})$$

Since the equation of the tangent to the trade-off curve at \tilde{b} is

$$v = \tilde{v} + (b - \tilde{b}) \frac{\partial v}{\partial b}(\tilde{b}),$$

Eq. (B.14) means that this tangent crosses the line $b = 0$ exactly at $v = v_{\max}$, as depicted in Fig. 5. This implies that $\tilde{v} > v^*$.

Using Eq. (B.8), we compute the viability, fecundity and fitness of the colony to obtain $V = Nv_{\max}/2$, $B = Nv_{\max}/2\alpha$, and $W = N^2 v_{\max}^2/4\alpha$, with $\alpha = -(\partial v/\partial b)(\tilde{b})$.

Case 3: $v_{\max} \geq -b_{\max}(\partial v/\partial b)(b_{\max})$. As in case 2, specialization in soma pays: $\tilde{p} > 0$, but now $\tilde{b} = b_{\max}$ (otherwise Eq. (B.10) would be positive, contradicting Eq. (B.3)). That is, the non-somatic cells should be germ (and not generalist). The above formulae for p , V , B and the fitness W still hold. In particular, Eq. (B.8) gives $p = \frac{1}{2}$; thus, half of the cells should specialize in germ and half in soma.

Note that our results do not require that the function $v(b)$ stays the same. In particular, the trade-off function could change as the number of cells increases ($v = v(b, N)$), in which case the optimal fecundity of the cells in an unspecialized colony may depend on the size of the colony. In addition, it is not necessary that the function v be concave (it could be, e.g., neither convex nor concave). A graphical analysis (included as supplementary material in the online version of this article) shows that provided that there exists a fecundity \tilde{b} which is greater than the mean fecundity in the best unspecialized colony and such that the line joining the points $(0, v_{\max})$ and $(\tilde{b}, v(\tilde{b}))$ is above the graph of v , then specialization will be favored.

Taking into account the fact that the proportion of somatic cells is a multiple of $1/N$: We now discuss how the results change if we take into account the fact that the proportion of somatic cells cannot vary continuously but must be a multiple of $1/N$. Consider a colony in which a cell specializes in soma and the other cells increase their fecundities from b^* to $(N/(N-1))b^*$. This colony would have the same fecundity $B^* = Nb^*$ as the best unspecialized

colony, but its viability would be

$$V = v_{\max} + (N - 1)v \left(\frac{Nb^*}{N - 1} \right). \quad (\text{B.15})$$

Thus, a sufficient condition for specialization to be favored is that $V > V^* = Nv^*$, or equivalently

$$v_{\max} > Nv^* - (N - 1)v \left(\frac{Nb^*}{N - 1} \right). \quad (\text{B.16})$$

Noting that $Nb^*/(N - 1) = b^*(1 + 1/(N - 1))$ and approximating the right-hand side by a Taylor expansion including up to second-order terms, Eq. (B.16) becomes

$$v_{\max} > v^* - b^* \frac{\partial v}{\partial b}(b^*) - \frac{(b^*)^2}{2(N - 1)} \frac{\partial^2 v}{\partial b^2}(b^*). \quad (\text{B.17})$$

This is exactly Eq. (13) when the trade-off is linear, but is more demanding for a concave trade-off as the second derivative of v is then negative. Furthermore, for a concave trade-off, Eq. (B.17) reduces to Eq. (13) in the limit of a very large number of cells, but may be significantly more demanding when N is small.

Intuitively, if N is small and v_{\max} is only slightly greater than $v^* - b^*(\partial v/\partial b)(b^*) = 2v^*$, then if a cell specializes in soma, the resulting proportion of somatic cells, $1/N$, might be much higher than the optimal proportion of somatic cells given in Eq. (B.8) and specialization in soma need not be favored even though Eq. (13) is satisfied.

B.3. Non-additive viabilities

Up to now, we assumed for simplicity that the cells contributions to the viability of the group were additive. We show here that we may dispense with this assumption. The assumptions we keep are that the fecundities of the cells are additive, $B = \sum_i b_i$, and (for some results) that the viability V of the group is a symmetric function of the fecundities; that is, the cells are interchangeable in the sense that if cell i and cell j exchange their fecundities, then the viability of the group does not change. We first consider trade-offs with no initial cost of reproduction and show that a convex (concave) trade-off selects for (against) specialization.

Convex trade-off: Assume that the function V is strictly convex. Then in an optimal group, at most one cell may have an intermediate fecundity. Indeed, assume by contradiction that two cells, say cells 1 and 2, have an intermediate fecundity. If cell 1 increases its fecundity by some small quantity x and cell 2 simultaneously decreases its fecundity by the same quantity, then the fecundity of the group does not change but the viability becomes

$$f(x) = V(b_1 + x, b_2 - x, b_3, \dots, b_N).$$

It follows from the strict convexity of V that the function f is strictly convex. Therefore, for $x > 0$, $f(x) + f(-x) > 2f(0)$. It follows that at least one of the quantities $f(x)$ and $f(-x)$ is strictly greater than $f(0)$. Without loss of generality, assume $f(x) > f(0)$. This means that, while the

fecundity of a colony with fecundities $b_1 + x, b_2 - x, b_3, \dots, b_N$ is the same as the fecundity of the initial colony, its viability, hence its fitness, is higher. Therefore, the initial colony was not optimal.

Linear trade-off: In this case, assuming that V is symmetric in the fecundities, then viabilities are additive and we are back to the model of Section 5.3.2. Indeed, if the trade-off is linear then there exist constants $V_{\max}, \alpha_1, \dots, \alpha_N$ such that

$$V(b_1, \dots, b_n) = V_{\max} - \left(\sum_i \alpha_i b_i \right).$$

If V is symmetric, then the constants α_i are all equal. Letting α be the common value of the α_i and $v_{\max} = V_{\max}/N$, we get

$$V(b_1, \dots, b_n) = \sum_i (v_{\max} - \alpha b_i)$$

as in the case of additive viabilities.

Concave trade-off: If the viability V is strictly concave and symmetric, then in an optimal colony, all cells have the same fecundity. Indeed, consider a colony with fecundities b_1, \dots, b_N and assume that two cells, say cells 1 and 2, have different fecundities. Let

$$g(x) = V \left(\frac{b_1 + b_2}{2} + x, \frac{b_1 + b_2}{2} - x, b_3, \dots, b_N \right)$$

so that $V(b_1, \dots, b_N) = g([b_1 - b_2]/2)$. If V is strictly concave, then so is g , so that for $x = [b_1 - b_2]/2 \neq 0$

$$g(x) + g(-x) < 2g(0).$$

Furthermore, if V is symmetric, then

$$g(-x) = V(b_2, b_1, b_3, \dots, b_N) = V(b_1, b_2, b_3, \dots, b_N) = g(x)$$

so that $g(x) < g(0)$. It follows that a colony with fecundities $(b_1 + b_2)/2, (b_1 + b_2)/2, b_3, \dots, b_N$ would have the same group fecundity but a higher group viability than a colony with fecundities $b_1, b_2, b_3, \dots, b_N$.

Initial cost of reproduction: We now consider the effect of an initial cost of reproduction. By an initial cost of reproduction, we mean that if cell i specializes in soma, the viability is substantially higher than if it provides a little fecundity. Formally,

$$V(b_1, \dots, b_{i-1}, 0, b_{i+1}, \dots, b_N) > V(b_1, \dots, b_{i-1}, 0^+, b_{i+1}, \dots, b_N)$$

with

$$V(b_1, \dots, b_{i-1}, 0^+, b_{i+1}, \dots, b_N) = \lim_{\varepsilon \rightarrow 0, \varepsilon > 0} V(b_1, \dots, b_{i-1}, \varepsilon, b_{i+1}, \dots, b_N).$$

When the fecundities are additive, then the difference

$$V(b_1, \dots, b_{i-1}, 0, b_{i+1}, \dots, b_N) - V(b_1, \dots, b_{i-1}, 0^+, b_{i+1}, \dots, b_N)$$

is simply the difference between v_{\max} and v_0 (see Fig. 5).

We focus on the case of a concave trade-off (the case of a convex or linear trade-off is easily dealt with as in Section 5.4.1.). Recall the above argument showing that if the function V is strictly concave and symmetric and if there is no initial cost of reproduction, then all cells should have the same fecundity. The same argument shows that, when there is a cost of reproduction, all non-somatic cells should have the same fecundity.

Let b^* denote the fecundity of the cells in the best unspecialized colony and $B^* = Nb^*$. If, starting from the best unspecialized colony, one cell specializes in soma and the other cells increase their fecundities to $b' = Nb^*/(N-1)$, then the global fecundity does not change but the viability goes from $V(b^*, \dots, b^*)$ to $V(0, b', \dots, b')$. Thus, a sufficient condition for fitness to increase is that

$$V(0, b', \dots, b') > V(b^*, \dots, b^*) \quad \text{with } b' = \frac{Nb^*}{N-1}. \quad (\text{B.18})$$

In the case of a linear (and symmetric) trade-off, the quantity

$$V(0^+, b', \dots, b') - V(b^*, \dots, b^*) \quad (\text{B.19})$$

is zero. Thus we may see this quantity as a detriment to soma specialization due to the concavity of the trade-off. Eq. (B.18) expresses that if the bonus to soma specialization

$$V(0, b', \dots, b') - V(0^+, b', \dots, b')$$

is greater than the detriment in Eq. (B.19), then specialization in soma is favored.

Another perspective is as follows: let $V(p, b)$ and $W(p, b)$ denote, respectively, the viability and fitness of a colony with a proportion p of somatic cells and fecundity b for the non-somatic cells. Assume for simplicity that p can vary continuously and let

$$f(p) = V(p, b^*/(1-p))$$

denote the viability of a colony with a proportion p of somatic cells and fecundity $b^*/(1-p)$ for the other cells, so that the global fecundity equals B^* . A sufficient condition for specialization is that

$$\frac{\partial f}{\partial p}(0) > 0. \quad (\text{B.20})$$

This expresses the fact that by having some cells specialized in soma, a colony can retain the same fecundity as the best unspecialized colony but increase its viability.

Simple computations show that Eq. (B.20) is equivalent to

$$\frac{1}{V} \frac{\partial V}{\partial p}(0, b^*) > 1 \quad (\text{B.21})$$

which means that the relative increase in viability provided by specialization in soma should be greater than a certain threshold (when the viabilities are additive, Eq. (B.21) boils down to $(v_{\max} - v^*)/v^* > 1$ or equivalently $v_{\max} > 2v^*$). An effect of an increasing initial cost of reproduction is that

the relative increase in viability provided by specialization in soma increases with the size of the colony (In volvocine algae, this is essentially because, in the absence of somatic cells, the viability gets lower and lower, hence $1/V$ increases). Thus, the higher the initial cost of reproduction, the more likely it is that Eq. (B.21) will be satisfied, hence specialization favored.

Appendix C. Supplementary data

Supplementary data associated with this article can be found in the online version at [doi:10.1016/j.jtbi.2005.08.043](https://doi.org/10.1016/j.jtbi.2005.08.043).

References

- Bell, G., 1985. The origin and early evolution of germ cells as illustrated by the Volvocales. In: Halvorson, H.O., Monroy, A. (Eds.), *The Origin and Evolution of Sex*. Alan R. Liss, Inc., New York, pp. 221–256.
- Benkman, C.W., 1993. Adaptation to single resources and the evolution of crossbill (*Loxia*) diversity. *Ecol. Monogr.* 63, 305–325.
- Benson, K.E., Stephens, D.W., 1996. Interruptions, tradeoffs, and temporal discounting. *Am. Zool.* 36, 506–517.
- Blows, M.W., Chenoweth, S.F., Hine, E., 2004. Orientation of the genetic variance-covariance matrix and the fitness surface for multiple male sexually selected traits. *Am. Nat.* 163, E329–E340.
- Buchheim, M.A., McAuley, M.A., Zimmer, E.A., Theriot, E.C., Chapman, R.L., 1994. Multiple origins of colonial green flagellates from unicells: evidence from molecular and organismal characters. *Mol. Phylogene. Evol.* 3, 322–343.
- Carriere, Y., Roff, D.A., 1995. The evolution of offspring size and number—a test of the Smith–Fretwell model in 3 species of crickets. *Oecologia (Berlin)* 102, 389–396.
- Charlesworth, B., 1980. *Evolution in Age-Structured Populations*. Cambridge University Press, Cambridge.
- Charlesworth, B., Leon, J.A., 1976. The relation of reproductive effort to age. *Am. Nat.* 110, 449–459.
- Coleman, A.W., 1999. Phylogenetic analysis of “volvocaceae” for comparative genetic studies. *Proc. Natl Acad. Sci. USA* 96, 13892–13897.
- Damuth, J., Heisler, I.L., 1988. Alternative Formulations of Multilevel Selection. *Biol. Philos.* 3, 407–430.
- Desnitski, A.G., 1995. A review on the evolution of development in *Volvox*—morphological and physiological aspects. *Eur. J. Protistol.* 31, 241–247.
- Foster, K.R., Fortunato, A., Strassmann, J.E., Queller, D.C., 2002. The costs and benefits of being a chimera. *Proc. R. Soc. Lond B Biol. Sci.* 269, 2357–2362.
- King, N., Carroll, S.B., 2001. A receptor tyrosine kinase from choanoflagellates: molecular insights into early animal evolution. *Proc. Natl Acad. Sci. USA* 98, 15032–15037.
- Kirk, D.L., 1997. The genetic program for germ-soma differentiation in *Volvox*. *Annu. Rev. Genet.* 31, 359–380.
- Kirk, D.L., 1998. *Volvox: Molecular-Genetic Origins of Multicellularity and Cellular Differentiation*. Cambridge University Press, Cambridge.
- Kirk, M.M., Stark, K., Miller, S.M., Muller, W., Taillon, B.E., Gruber, H., Schmitt, R., Kirk, D.L., 1999. RegA, a *Volvox* gene that plays a central role in germ-soma differentiation, encodes a novel regulatory protein. *Development* 126, 639–647.
- Kisdi, E., 2001. Long-term adaptive diversity in Levene-type models. *Evol. Ecol. Res.* 3, 721–727.
- Kisdi, E., Meszéna, G., Pásztor, L., 1998. Individual optimization: mechanisms shaping the optimal reaction norm. *Evol. Ecol.* 12, 211–221.

- Koufopanou, V., 1994. The evolution of soma in the Volvocales. *Am. Nat.* 143, 907–931.
- Koufopanou, V., Bell, G., 1993. Soma and germ—an experimental approach using *Volvox*. *Proc. R. Soc. London B* 254, 107–113.
- Larson, A., Kirk, M.M., Kirk, D.L., 1992. Molecular phylogeny of the volvocine flagellates. *Mol. Biol. Evol.* 9, 85–105.
- Levins, R., 1968. *Evolution in Changing Environments: Some Theoretical Explorations*. Princeton University Press, Princeton, NJ.
- Michod, R.E., 1978. Evolution of life histories in response to age-specific mortality factors. *Am. Nat.* 113, 531–550.
- Michod, R.E., 1982. The theory of kin selection. *Annu. Rev. Ecol. Syst.* 13, 23–55.
- Michod, R.E., 1996. Cooperation and conflict in the evolution of individuality. II. conflict mediation. *Proc. R. Soc. London B* 263, 813–822.
- Michod, R.E., 1997. Cooperation and conflict in the evolution of individuality I. Multi-level selection of the organism. *Am. Nat.* 149, 607–645.
- Michod, R.E., 1999. *Darwinian Dynamics, Evolutionary Transitions in Fitness and Individuality*. Princeton University Press, Princeton, NJ.
- Michod, R.E., Roze, D., 1997. Transitions in individuality. *Proc. R. Soc. London B* 264, 853–857.
- Michod, R.E., Roze, D., 1999. Cooperation and conflict in the evolution of individuality. III. transitions in the unit of fitness. In: Nehaniv, C.L. (Ed.), *Mathematical and Computational Biology: Computational Morphogenesis, Hierarchical Complexity, and Digital Evolution*. American Mathematical Society, Providence, RI, pp. 47–92.
- Michod, R.E., Nedelcu, A.M., Roze, D., 2003. Cooperation and conflict in the evolution of individuality IV. Conflict mediation and evolvability in *Volvox carteri*. *BioSystems* 69, 95–114.
- Miller, S.M., Kirk, D.L., 1999. *glsA*, a *Volvox* gene required for asymmetric division and germ cell specification, encodes a chaperone-like protein. *Development* 126, 649–658.
- Morgan, N.C., 1980. Secondary production. In: Le Cren, E.D., Lowe-McConnell, R.H. (Eds.), *The Functioning of Freshwater Ecosystems*, IBP 22. Cambridge University Press, Cambridge, pp. 247–340.
- Niklas, K.J., 1994. *Plant allometry: The Scaling of Form and Process*. University of Chicago Press, Chicago, IL.
- Niklas, K.J., 2000. The evolution of plant body plans—a biomechanical perspective. *Ann. Bot.* 85, 411–438.
- Nozaki, H., 2003. Origin and evolution of the genera *Pleodorina* and *Volvox* (Volvocales). *Biologia* 58, 425–431.
- Nozaki, H., Misawa, K., Kajita, T., Kato, M., Nohara, S., Watanabe, M., 2000. Origin and evolution of the colonial Volvocales (Chlorophyceae) as inferred from multiple, chloroplast gene sequences. *Mol. Phylogenet. Evol.* 17, 256–268.
- Nozaki, H., Takahara, M., Nakazawa, A., Kita, Y., Yamada, T., Takano, H., Kawano, S., Kato, M., 2002. Evolution of *rbcL* group IA introns and intron open reading frames within the colonial Volvocales (Chlorophyceae). *Mol. Phylogenet. Evol.* 23, 326–338.
- Nozaki, H., Misumi, O., Kuroiwa, T., 2003. Phylogeny of the quadri-flagellate Volvocales (Chlorophyceae) based on chloroplast multigene sequences. *Mol. Phylogenet. Evol.* 29, 58–66.
- Pentecost, A., 1983. The distribution of daughter colonies and cell numbers in a natural population of *Volvox aureus* Ehrenb. *Ann. Bot.* 52, 769–776.
- Porter, K.G., 1977. Plant-animal interface in freshwater ecosystems. *Am. Sci.* 65, 159–170.
- Queller, D.C., Ponte, E., Bozzaro, S., Strassmann, J.E., 2003. Single-gene greenbeard effects in the social amoeba *Dictyostelium discoideum*. *Science* 299, 105–106.
- Rausch, H., Larsen, N., Schmitt, R., 1989. Phylogenetic relationships of the green alga *Volvox carteri* deduced from small-subunit ribosomal RNA comparisons. *J. Mol. Evol.* 29, 255–265.
- Reynolds, C.S., 1984. *The Ecology of Freshwater Phytoplankton*. Cambridge University Press, Cambridge, UK.
- Reznick, D., 1985. Costs of reproduction—an evaluation of the empirical evidence. *Oikos* 44, 257–267.
- Roff, D.A., 2000. Trade-offs between growth and reproduction: an analysis of the quantitative genetic evidence. *J. Evol. Biol.* 13, 434–445.
- Roff, D.A., 2002. *Life History Evolution*. Sinauer Association, Sunderland, MA.
- Rueffler, C., Van Dooren, T.J.M., Metz, J.A.J., 2004. Adaptive walks on changing landscapes: Levins' approach extended. *Theor. Popul. Biol.* 65, 165–178.
- Sato, H., 2002. Invasion of unisexuals in hermaphrodite populations of animal-pollinated plants: effects of pollination ecology and floral size-number trade-offs. *Evolution* 56, 2374–2382.
- Schaffer, W.M., 1974. Selection for optimal life histories: the effects of age structure. *Ecology* 55, 291–303.
- Shikano, S., Luckinbill, L.S., Kurihara, Y., 1990. Changes of traits in a bacterial population associated with protozoal predation. *Microb. Ecol.* 20, 75–84.
- Shimkets, L.J., 1990. Social and developmental biology of the myxobacteria. *Microbiol. Rev.* 54, 473–501.
- Solari, C.A., A hydrodynamics approach to the evolution of multicellularity: flagellar motility and the evolution of germ–soma differentiation in volvoclean green algae. Thesis/Dissertation, University of Arizona, 2005.
- Solari, C.A., Ganguly, S., Kessler, J.O., Michod, R.E., Goldstein, R.E., 2005a. Multicellularity and the functional interdependence of motility and molecular transport. *Proc. Natl. Acad. Sci. USA*, in revision.
- Solari, C.A., Kessler, J.O., Michod, R.E., 2005b. A hydrodynamics approach to the evolution of multicellularity: flagellar motility and germ–soma differentiation in volvoclean green algae. *Am. Nat.*, in revision.
- Stearns, S.C., 1992. *The Evolution of Life Histories*. Oxford University Press, Oxford.
- Strassmann, J.E., Zhu, Y., Queller, D.C., 2000. Altruism and social cheating in the social amoeba *Dictyostelium discoideum*. *Nature (London)* 408, 965–967.
- Strohm, E., Linsenmair, K.E., 2000. Allocation of parental investment among individual offspring in the European beewolf *Philanthus triangulum* F. (Hymenoptera : Sphecidae). *Biol. J. Linnean Soc.* 69, 173–192.
- Takada, T., Nakajima, H., 1996. The optimal allocation for seed reproduction and vegetative reproduction in perennial plants: an application to the density-dependent transition matrix model. *J. Theor. Biol.* 182, 179–191.
- Velicer, G.J., Kroos, L., Lenski, R.E., 2000. Developmental cheating in the social bacterium *Myxococcus xanthus*. *Nature (London)* 404, 598–601.

Appendix C. (Supplementary material for the online version of the article)

In this appendix, we generalize some of the results given in appendix B. We focus on the case of additive viabilities and fecundities. That is, we assume that the viability and fecundity of the colony are given by the sum of the contributions of the cells, as given in equation 9. Furthermore, we assume (counterfactually) that the proportion of cells with a given fecundity need not be a multiple of $1/n$ but may take any value in $[0, 1]$.

We consider a nonnegative trade-off function $v(b)$ defined on an interval $[0, b_{max}]$ (the possible contributions of a cell to the fecundity of the colony). No other assumption the trade-off function is needed. In particular, we do not assume that this function is differentiable, nor that it is convex or concave, nor even that it is decreasing.

Let $b^* \in [0, b_{max}]$ be an optimal value of the fecundity of all cells in an unspecialized colony. That is, b^* maximizes the product $bv(b)$. (Since we make no particular assumption on the curvature of the trade-off, there may be multiple values of b maximizing the product $bv(b)$; in this case, we choose as b^* one of these values.) Let $v^* = v(b^*)$ and for all b in $[0, b_{max}]$ let

$$\bar{v}(b) = v^* (2 - b/b^*)$$

That is, \bar{v} represents a linear trade-off with no initial cost of reproduction and such that $\bar{v}(0) = 2v^*$ and $\bar{v}(b^*) = v^*$. Note that if for all b , $v(b) = \bar{v}(b)$, then, as proved in section 5.3.2, any arrangement of the fecundities such that the average fecundity is b^* is optimal. In particular, the undifferentiated colony in which all cells have fecundity b^* is an optimal colony.

More generally, we prove below that:

Proposition. *The undifferentiated colony in which all cells have fecundity b^* is an optimal colony if and only if*

$$\forall b \in [0, b_{max}], v(b) \leq \bar{v}(b) \tag{C.1}$$

Indeed, let

$$W(b_1, \dots, b_N) = \left(\sum_{1 \leq i \leq N} v(b_i) \right) \left(\sum_{1 \leq i \leq N} b_i \right)$$

and

$$\bar{W}(b_1, \dots, b_N) = \left(\sum_{1 \leq i \leq N} \bar{v}(b_i) \right) \left(\sum_{1 \leq i \leq N} b_i \right)$$

Let $W^* = W(b^*, \dots, b^*)$ and $\bar{W}^* = \bar{W}(b^*, \dots, b^*)$

Since $\bar{v}(b^*) = v(b^*)$, we have

$$W^* = \bar{W}^*$$

Furthermore, it follows from the definition of \bar{v} and the analysis in section 5.3.2. that, for any values b_1, \dots, b_n of the fecundities,

$$\bar{W}^* \geq \bar{W}(b_1, \dots, b_n)$$

Finally, if Eq. (C.1) is satisfied then, for any b_1, \dots, b_n ,

$$\bar{W}(b_1, \dots, b_n) \geq W(b_1, \dots, b_n)$$

Combining these three equations, we obtain that for any values b_1, \dots, b_n of the fecundities,

$$W^* = \bar{W}^* \geq \bar{W}(b_1, \dots, b_n) \geq W(b_1, \dots, b_n)$$

Therefore, the undifferentiated colony in which all cells have fecundity b^* is an optimal colony.

Conversely, if Eq. (C.1) is not satisfied then there exists a fecundity b' in $[0, b_{max}]$ such that

$$v(b') > \bar{v}(b') \tag{C.2}$$

For p in $[0, 1]$, let $W(p)$ denote the fitness of a colony with a proportion p of cells having fecundity b' and a proportion $1 - p$ of cells having fecundity b^* :

$$W(p) = N^2 (pv' + (1 - p)v^*) (pb' + (1 - p)b^*)$$

Thus, $W(0) = W^*$. A straightforward computation shows that due to Eq. (C.2), we have

$$\frac{\partial W}{\partial p}(0) > 0 \tag{C.3}$$

(actually, Eqs. (C.2) and (C.3) are equivalent).

Eq. (C.3) means that if, departing from the best undifferentiated colony, a small proportion of cells adopt the fecundity b' , then the fitness of the colony increases. It follows that the best undifferentiated colony is not an optimal colony, hence that specialization is favored. This completes the proof of the proposition.

We conclude with some remarks: first, if Eq. (C.1) is satisfied then, for any b' and \tilde{b} such that $b' < b^* < \tilde{b}$, the line joining the points $(b, v(b))$ and $(\tilde{b}, v(\tilde{b}))$ crosses the vertical line $b = b^*$ below v^* . It follows that, as mentioned in appendix B.2, if there exists a fecundity $\tilde{b} > b^*$ such that the line joining the points $(0, v(0))$ and $(\tilde{b}, v(\tilde{b}))$ is strictly above the graph of v at $b = b^*$, then condition (C.1) is not satisfied, hence specialization is favored.

Second, if the trade-off is concave (strictly convex) then Eq. (C.1) is (is not) satisfied so that specialization is not (is) favored, as mentioned in appendix B.2.

Finally, a complement to the above proposition is as follows (proof omitted): If Eq. (C.1) holds then the optimal colonies are exactly those such that the average fecundity is v^* and such that, for every cell i , $v(b_i) = \bar{v}(b_i)$. It follows that the undifferentiated colony in which every cell has fecundity b^* is the unique optimal colony if and only if, on top of condition (C.1), at least one of the two following conditions holds:

$$\forall b \in [0, b^*[, v(b) < \bar{v}(b) \tag{C.4}$$

$$\forall b \in]b^*, b_{max}], v(b) < \bar{v}(b) \tag{C.5}$$

Bibliography of parts I and II

- [1] Akin, E. (1980), “Domination or Equilibrium”, *Mathematical Biosciences* **50**, 239-250
- [2] Aumann R.J. (1961), “Almost Strictly Competitive Games”, *Journal of the SIAM* **9**, 544-550
- [3] Aumann, R. (1974), “Subjectivity and Correlation in Randomized Strategies”, *Journal of Mathematical Economics* **1**, 67-96
- [4] Aumann, R.J. (1987), “Correlated Equilibria as an Expression of Bayesian Rationality”, *Econometrica* **55**, 67-96
- [5] Beaud, J.P. (2002), *Contributions à la théorie des jeux : jeux antagonistes et modèles de jeux répétés*, Ph'D thesis, Université Paris VI.
- [6] Berger, U. and J. Hofbauer (2005), “Irrational Behavior in the Brown-von Neumann-Nash Dynamics”, *Games and Economic Behavior*, forthcoming
- [7] Bernheim D.G. (1984), “Rationalizable Strategic Behavior”, *Econometrica* **52**, 1007-1028
- [8] Bhatia, N.P. and G.P. Szegö (1970), *Stability Theory of Dynamical Systems*, Springer-Verlag
- [9] Bohnenblust, H.F., S. Karlin and L.S. Shapley (1950), “Solutions of Discrete, Two-Person Games”, in H.W. Kuhn and A.W. Tucker, Eds. *Contributions to the Theory of Games I*, 51-72. Annals of Mathematical Studies **24**. Princeton University Press.
- [10] Bomze, I. (1983) “Lotka-Volterra equations and replicator dynamics: a two-dimensional classification”, *Biological Cybernetics* **48**, 201-11

- [11] Brannath W. (1994), "Heteroclinic networks on the tetrahedron", *Non-linearity* **7**, 1367-1384
- [12] Brown, G.W. and J. von Neumann (1950), "Solutions of Games by Differential Equations", in H.W. Kuhn and A.W. Tucker, Eds. *Contributions to the Theory of Games I*, 73-79. Annals of Mathematical Studies **24**. Princeton University Press.
- [13] Cabrales, A. and J.Sobel (1992), "On the Limit Points of Discrete Selection Dynamics", *Journal of Economic Theory* **57**, 407-420
- [14] Chatterji, S. and S. Govindan (2006), "Message Spaces for Perfect Correlated Equilibria", to appear in *Economic Theory*, **28**, 475-479
- [15] Cressman, R. (2003), *Evolutionary Dynamics and Extensive Form Games*, MIT Press
- [16] Cripps, M. (1991), "Correlated Equilibria and Evolutionary Stability", *Journal of Economic Theory* **55**, 428-434
- [17] Dhillon, A. and J.F. Mertens (1996), "Perfect Correlated Equilibria", *Journal of Economic Theory* **68**, 279-302
- [18] Dekel, E. and S. Scotchmer (1992), "On the Evolution of Optimizing Behavior", *Journal of Economic Theory* **57**, 392-406
- [19] Evangelista, F. and T.E.S. Raghavan (1996), "A note on correlated equilibrium" *International Journal of Game Theory* **25**, 35-41
- [20] Fisher, R.A. (1930), *The Genetical Theory of Natural Selection*, Oxford: Clarendon Press.
- [21] Flesch, J., F. Thuisjman and K. Vrieze (1997), "Cyclic Markov Equilibria in Stochastic Games", *International Journal of Game Theory* **26**, 303-314
- [22] Forges, F. (1986) "An Approach to Communication Equilibria", *Econometrica* **54**, 1375-1385
- [23] Forges, F. (1990), "Correlated Equilibrium in Two-Person Zero-Sum Games", *Econometrica* **58**, 515

- [24] Foster D. and H.P. Young (1990), "Stochastic Evolutionary Game Dynamics", *Journal of Theoretical Biology* **38**, 219-232
- [25] Friedman J. (1983), "On Characterizing Equilibrium Points in Two-Person Strictly Competitive Games", *International Journal of Game Theory* **12**, 245-247.
- [26] Friedman D. (1991), "Evolutionary Games in Economics", *Econometrica* **59**, 637-666
- [27] Gaunersdorfer A. and J. Hofbauer (1995), "Fictitious Play, Shapley Polygons, and the Replicator Equation", *Games and Economic Behavior* **11**, 279-303
- [28] Gilboa, I. and A. Matsui (1991), "Social Stability and Equilibrium", *Econometrica* **59**, 859-867
- [29] Gilboa I. and E. Zemel (1989), "Nash and Correlated Equilibria: Some Complexity Considerations", *Games and Economic Behavior* **1**, 80-93
- [30] Gomez Canovas, S., P. Hansen, and B. Jaumard (1999), "Nash Equilibria from the Correlated Equilibria Viewpoint", *International Game Theory Review* **1**, 33-44
- [31] Hamilton, W.D. and R.M. May (1977), "Dispersal in stable habitats", *Nature* **269**, 578-581
- [32] Hammerstein, P. and R. Selten (1994), "Game Theory and Evolutionary Biology", in *Handbook of Game Theory with Economic Applications*, R.J. Aumann and S. Hart Eds, Elsevier, vol. 2, 931-993
- [33] Harsanyi, J.C. (1973), "Oddness of the Number of Equilibrium Points: A New Proof", *International Journal of Game Theory* **2**, 235-250
- [34] Hart, S. (2005), "Adaptive Heuristics", *Econometrica* **73**, 5, 1401-1430
- [35] Hart, S. and A. Mas-Colell (2001), "A General Class of Adaptive Strategies", *Journal of Economic Theory* **98**, 26-54
- [36] Hart, S. and A. Mas-Colell (2003a), "Regret-Based Continuous-Time Dynamics", *Games and Economic Behavior* **45**, 375-394

- [37] Hart, S. and A. Mas-Colell (2003b), “Uncoupled Dynamics Do Not Lead to Nash Equilibrium”, *American Economic Review* **93**, 1830-1836
- [38] Hart, S. and D. Schmeidler (1989), “Existence of Correlated Equilibria”, *Mathematics Of Operations Research* **14**, 18-25
- [39] Hofbauer, J. (1994), “Heteroclinic Cycles In Ecological Differential Equations”, *Tatra Mountains Mathematical Publications* **4**, 105-116
- [40] Hofbauer, J. (2000), “From Nash and Brown to Maynard-Smith: Equilibria, Dynamics and ESS”, *Selection* **1**, 81-88
- [41] Hofbauer, J. (2004), “Time Averages of the Replicator Dynamics and Correlated Equilibria”, mimeo.
- [42] Hofbauer, J. and K. Sigmund (1998), *Evolutionary Games and Population Dynamics*, Cambridge University Press
- [43] Hofbauer, J. and K. Sigmund (2003), “Evolutionary Game Dynamics”, *Bulletin (New Series) of the American Mathematical Society* **40**, 479-519
- [44] Hofbauer, J. and S. Sorin (2002), “Best-Response Dynamics for Continuous Zero-Sum Games”, cahier du laboratoire d'économétrie 2002-28, Ecole polytechnique, Paris
- [45] Hofbauer, J. and J.W. Weibull (1996), “Evolutionary Selection against Dominated Strategies”, *Journal of Economic Theory* **71**, 558-573
- [46] Jansen, M.J.M (1981) “Regularity and Stability of Equilibrium Points of Bimatrix Games”, *Mathematics Of Operations Research* **6**, 18-25
- [47] Jordan, J.S. (1993), “Three Problems in Learning Mixed-Strategy Equilibria”, *Games and Economic Behavior* **5**, 368-386
- [48] Kandori, M., G.J. Mailath and R. Rob (1993), “Learning, Mutation and Long-Run Equilibria in Games”, *Econometrica* **61**, 29-56
- [49] Karr A.F. (1990), “Markov Processes” in *Handbooks in Operations Research and Management Sciences*, D.P. Heyman, M.J. Sobel Eds, Elsevier Science Publishers (North Holland), vol. 2, chap. 2, 95-123

- [50] Kats, A. and J.-F. Thisse (1992), “Unilaterally Competitive Games”, *International Journal of Game Theory* **21**, 291-299
- [51] Mailath G.J., L. Samuelson and A. Shaked (1997), “Correlated Equilibria and Local Interactions”, *Economic Theory* **9**, 551-556
- [52] Matsui, A. (1992), “Best-Response Dynamics and Socially Stable Strategies”, *Journal of Economic Theory* **57**, 343-362
- [53] Maynard Smith, J. and G.R. Price (1973), “The Logic of Animal Conflict”, *Nature* **246**, 15-18
- [54] Maynard Smith, J., and G.A. Parker (1976), “The Logic of Asymmetric Contests”, *Animal Behavior* **24**, 159-175
- [55] Maynard Smith, J. (1982), *Evolution and the Theory of Games*, Cambridge University Press
- [56] Monderer, D. and A. Sela (1997), “Fictitious-Play and No-Cycling Condition”, SFB 504 Discussion Paper 97-12, Universität Mannheim
- [57] Moulin, H and J-P. Vial (1978), “Strategically Zero-Sum Games: The Class of Games Whose Completely Mixed Equilibria Cannot be Improved Upon”, *International Journal of Game Theory* **7**, 201-221
- [58] Mertens, J.F., S. Sorin and S. Zamir (1994) “Repeated games, Part A, Background material”, CORE discussion paper 9402, Université Catholique de Louvain
- [59] Myerson, R.B. (1986), “Acceptable and Predominant Correlated Equilibria”, *International Journal of Game Theory* **15**, 133-154
- [60] Myerson, R.B. (1991), *Game Theory: Analysis of Conflict*, Harvard University Press, Cambridge, MA.
- [61] Myerson, R.B. (1994), “Communication, Correlated Equilibria and Incentive Compatibility” in *Handbook of Game Theory*, R.J. Aumann and S. Hart Eds, Elsevier Science Publishers (North Holland), vol. 2, chap. 24, 827-848
- [62] Myerson, R.B. (1997), “Dual Reduction and Elementary Games”, *Games and Economic Behavior* **21**, 183-202

- [63] Myerson, R.B. (2003), "Incentive-Dual Methods in Game Theory: Virtual Utility and Dual Reduction Revisited", mimeo.
- [64] Nachbar, J. (1990) " 'Evolutionary' Selection Dynamics in Games: Convergence and Limit Properties", *International Journal of Game Theory* **19**, 59-89
- [65] Nash, J. (1950), *Non-Cooperative Games*, Ph'D Dissertation, Princeton University
- [66] Nau R.F., S. Gomez Canovas, and P. Hansen (2004), "On the Geometry of Nash Equilibria and Correlated Equilibria", *International Journal of Game Theory* **32**, 443-453
- [67] Nau, R.F. and K.F. McCardle (1990), "Coherent Behavior in Noncooperative Games", *Journal of Economic Theory* **50**, 424-444
- [68] Norde, H. (1999), "Bimatrix games have quasi-strict equilibria", *Mathematical Programming* **85**, 35-49
- [69] Nitzan, N. (2005), "Tight Correlated Equilibrium", discussion paper #394, Center for the Study of Rationality, the Hebrew University of Jerusalem
- [70] Owen, G. (1995), *Game Theory*, Academic Press
- [71] Raghavan, T.E.S. (1994), "Zero-Sum Two-Person Games" in *Handbook of Game Theory*, R.J. Aumann and S. Hart Eds, Elsevier Science Publishers (North Holland), vol. 2, chap. 20, 735-768
- [72] Raghavan, T.E.S. (2002), "Non-Zero Sum Two-Person Games" in *Handbook of Game Theory*, R.J. Aumann and S. Hart Eds, Elsevier Science Publishers (North Holland), vol. 3, chap. 44, 1687-1721
- [73] Rivière, P. (1997), *Quelques modèles de jeux d'évolution*, Ph'D dissertation, Université Paris 6
- [74] Ritzberger K. and J.W. Weibull (1995), "Evolutionary Selection in Normal-Form Games", *Econometrica* **63**, 1371-1399
- [75] Robinson, J. (1951). "An iterative method of solving a game", *Annals of Mathematics* **54**, 296-301

- [76] Rosenthal, R.W. (1974), "Correlated Equilibria in Some Classes of Two-Person Games", *International Journal of Game Theory* **3**, 119-128
- [77] Samuelson, L. (1997), "Evolutionary Games and Equilibrium Selection", MIT Press
- [78] Samuelson, L. and J. Zhang (1992), "Evolutionary stability in asymmetric games", *Journal of Economic Theory* **57**, 363-391
- [79] Sandholm, W.H. (2006), *Population Games and Evolutionary Dynamics*, unpublished manuscript to appear in 2006, University of Wisconsin
- [80] Selten, R. (1980), "A Note on Evolutionary Stable Strategies in Asymmetric Animal Conflicts", *Journal of Theoretical Biology* **84**, 93-101
- [81] Shapley, L. (1964), "Some Topics in Two-Person Games", *Annals of Mathematical Studies* **52**, Princeton, 1-28
- [82] Shmida A. and B. Peleg (1997), "Strict and Symmetric Correlated Equilibria are the Distributions of the ESS's of Biological Conflicts with Assymmetric Roles", in *Understanding Strategic Interaction, Essays in Honor of R. Selten*, W. Albers, W. Guth, P. Hammerstein, B. Moldovanu and E. van Damme Eds, Springer-Verlag, 149-170
- [83] Sorin, S. (1998), "Distribution Equilibrium I: Definition and Examples", cahier du THEMA 9835, Université de Paris X - Nanterre
- [84] Swinkels, J. (1993), "Adjustment Dynamics and Rational Play in Games", *Games and Economic Behavior* **5**, 455-484
- [85] Taylor, P.D., and L. Jonker (1978), "Evolutionary Stable Strategies and Game Dynamics", *Mathematical Biosciences* **40**, 145-156
- [86] van Damme, E. (1991), *Stability and Perfection of Nash Equilibria*, Springer-Verlag
- [87] Vega-Redondo, F (1996), *Evolution, Games, and Economic Behaviour*, Oxford University Press

- [88] von Stengel, B. (2002), “Computing Equilibria for Two-Person Games”, in *Handbook of Game Theory*, R.J. Aumann and S. Hart Eds, Elsevier Science Publishers (North Holland), vol. 3, chap. 45, 1723-1759
- [89] Weibull, J.W. (1995), *Evolutionary Game Theory*, Cambridge, MA: MIT Press
- [90] Young, H.P. (1993), “The Evolution of Conventions”, *Econometrica* **61**, 57-84
- [91] Zeeman, E.C. (1980), “Population Dynamics From Game Theory”, in *Global Theory of Dynamical Systems*, Lecture Notes in Mathematics **819**, Nitecki, A. and C. Robinson Eds, Springer, New York.