



HAL
open science

Essays in Behavioral Strategy : Re-biased Search, Misconceived Complexity, and Cognitive Aliens

Aleksey Korniychuk

► **To cite this version:**

Aleksey Korniychuk. Essays in Behavioral Strategy : Re-biased Search, Misconceived Complexity, and Cognitive Aliens. Business administration. Université Paris Saclay (COmUE), 2017. English. NNT : 2017SACLH002 . tel-01569432

HAL Id: tel-01569432

<https://pastel.hal.science/tel-01569432v1>

Submitted on 26 Jul 2017

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.

THESE DE DOCTORAT
DE
L'UNIVERSITE PARIS-SACLAY
PREPAREE A "HEC PARIS"

Ecole Doctorale n° 578
Sciences de l'homme et de la société (SHS)

Spécialité de doctorat : Sciences de gestion

Par

Aleksey KORNIYCHUK

**Essays in Behavioral Strategy:
Re-biased Search, Misconceived Complexity, and Cognitive Aliens**

Composition du Jury :

Stefano BRUSONI	Professeur, ETH Zurich	Rapporteur
Arnaldo CAMUFFO	Professeur, Bocconi University	Examinateur
Giada DI STEFANO	Professeur Associé, HEC Paris	Examinatrice
Margaret KYLE	Professeur, Ecole des Mines	Rapporteuse
Philippe MONIN	Professeur, EM Lyon	Président du jury
Tomasz OBLOJ	Professeur Associé, HEC Paris	Co-directeur de thèse
Gonçalo PACHECO DE ALMEIDA	Professeur Associé, HEC Paris	Directeur de thèse
Eric Luis UHLMANN	Professeur Associé, INSEAD	Co-encadrant

L'université Paris-Saclay n'entend donner aucune approbation ou improbation aux opinions émises dans cette thèse. Ces opinions doivent être considérées comme propres à leur auteur.

Résumé

i. L'Objet

La rationalité humaine est limitée (Conlisk 1996, Fiori 2011, Simon 1955, 1957). Du point de vue de l'école Carnegie (néo-Carnegie), cela signifie que i) les décideurs recherchent des solutions offrant un minimum de satisfaction (à partir d'ici, ce processus est appelé le « satisficing consécutif » en anglais, on l'appelle « search ») plutôt qu'optimisent l'ensemble des options disponibles, et ii) la réflexion des décideurs est un processus complexe qui n'est pas parfaite, c'est-à-dire les décideurs pensent d'une manière typique de l'homme (voir Cyert et March 1963, Gavetti, Levinthal, Ocasio 2007, Kahneman 2003, Simon 1955, 1956). Bien que la littérature existante offre un aperçu important de ces sujets, l'analyse conjointe de (i) et (ii) reste rare (voir Gavetti et al., 2007, Powell, Lovallo, Fox 2011). Cette large lacune dans la littérature existante est au cœur de la présente dissertation.

L'objectif de ce travail est d'identifier (potentiellement utiles) les modèles dans la dynamique du satisficing consécutif qui survient en raison des idiosyncrasies cognitives des décideurs. Concentrons-nous séparément sur chaque élément de cet objectif. Premièrement, le satisficing consécutif est un moyen d'apprentissage et d'adaptation (Levinthal 1997). Bien qu'il existe plusieurs façons dont les organisations apprennent et s'adaptent (par exemple, l'échantillonnage séquentiel, l'apprentissage indirecte, l'écologie de la population, etc.), le processus d'adaptation particulier considéré tout au long de ce travail est le satisficing consécutif de recombinaison. Ce processus caractérise la reconfiguration des connaissances organisationnelles, du développement technologique ou plus largement de la résolution des problèmes (voir par exemple Fleming et Sorenson 2001, Nickerson, Silverman, Zenger 2007).

Essentiellement, cela représente la prise de décision consécutive où les agents font générer et évaluer des nouvelles combinaisons (Knudsen et Levinthal 2007). Deuxièmement, l'accent mis sur les idiosyncrasies cognitives implique la comparaison de différentes façons dont les décideurs pensent, comme avoir une préférence intuitive systématique contre ne pas l'avoir (essai I) ou avoir une plus complexe compréhension contre une plus simple (essai II). En reconnaissant la variance de la connaissance des décideurs (sur l'hétérogénéité dans la pensée humaine, voir Baron 2007), cette thèse s'engage à développer une compréhension plus nuancée de l'adaptation organisationnelle, qui permettrait d'informer mieux du potentiel de l'esprit humain, ainsi elle sert à améliorer l'efficacité de la gestion.

Cela définit le domaine académique auquel appartient le présent travail, c'est-à-dire la stratégie comportementale. La stratégie comportementale est un sous-domaine émergent de sciences de l'organisation qui cherche à découvrir des sources comportementales cognitives de création de valeur supérieure ou, plus largement, considère les phénomènes stratégiques à travers des lentilles comportementales (Powell et al., 2011). En tant que tel, ce sous-domaine de la connaissance repose sur la théorie comportementale de l'entreprise et plus large sur l'école Carnegie (voir Gavetti 2012). En effet, on peut considérer l'émancipation de la stratégie comportementale (en tant que domaine distinct) comme une réponse à l'appel pour « une perspective renouvelée, comportementalement plausible et centrée sur la décision sur les organisations » (traduit, Gavetti et al., 2007: p. 531).

Une propriété clé qui identifie la stratégie comportementale est l'accent mis sur l'utilisation (ou la mauvaise utilisation) de la psychologie des individus pour la création et la capture de la valeur dans les organisations (voir Gavetti 2012). Les travaux qui se rapportent

directement à ce sujet remontent au moins aux années 1980 (voir par exemple Schwenk 1984, 1986). Cependant, une enquête plus exhaustive sur ce sujet tombe sur les dernières années et concorde quelque peu avec les pièces qui définissent la notion de stratégie comportementale avec ses racines dans des hypothèses positives plutôt que sur des suppositions idéalisées sur la prise de décision (voir aussi Levinthal 2011, Powell et al. 2011).

Dans cette littérature, un sous-ensemble d'œuvres définit la logique de la présente dissertation. Les principaux développements dans cette logique synthétisent comme suit. Tout d'abord, en s'appuyant sur la théorie comportementale de l'entreprise ainsi que sur les travaux en matière d'adaptation et d'apprentissage organisationnels (par exemple, Cyert et March 1963, Levinthal 1997, Levitt et March 1988), Denrell et March illustrent qu'un biais contre les alternatives risquées peut être un résultat de l'apprentissage et de la sélection (2001). Cela pourrait avoir des conséquences préjudiciables pour une adaptation ultérieure car les organisations persistent avec des alternatives plus sûres (mais moins optimales). Reflétant cette conclusion dans une œuvre postérieure, Winter et ses collègues montrent comment l'obsession modérée d'une action particulière peut conduire à des performances supérieures à long terme (Winter, Cattani, Dorsch, 2007). Contrairement à l'aversion au risque, un certain niveau de surévaluation consciente de certaines options peut payer à l'avenir. La raison en est que les erreurs d'évaluation permettent un niveau d'exploration plus élevé (Knudsen et Levinthal 2007, Baumann et Martignoni 2011). Une analyse de l'interaction de la confiance et de la qualité de l'évaluation illustre en outre que la confiance peut compenser partiellement la compétence (Posen, Martignoni, Lang 2013). Bien que ces travaux démontrent comment les éléments du comportement humain peuvent être à la fois un sous-produit du satisficing consécutif et un déterminant du gain à l'adaptation, ils se concentrent sur les résultats

spécifiques de l'évaluation des alternatives sans tenir compte du processus cognitif qui sous-tend des erreurs de jugement (par exemple, imparfait modèles mentaux) ou précède le choix conscient (p. ex. réactions spontanées et affectives).

Un travail récent met en lumière la signification de ces éléments de la connaissance, en particulier des modèles mentaux imparfaits, pour le processus d'adaptation organisationnelle. Csaszar et Levinthal (2016) étudient le choix entre la recherche de représentations mentales et la recherche d'une stratégie efficace, compte tenu de la représentation mentale des décideurs du problème organisationnel. Ils démontrent que des modèles mentaux incomplets peuvent faciliter l'exploration et ainsi conduire à de meilleurs résultats organisationnels. Dans le même ordre d'idées, Martignoni et ses collègues (2016) analysent les conséquences de modèles mentaux sous ou trop précisés. Leur analyse indique que, compte tenu d'une combinaison de complexité environnementale et de contrôle organisationnel sur les choix, les modèles mentaux mal spécifiés peuvent conduire à des résultats supérieurs. En conséquence, ils mettent en relief l'importance de l'ajustement cognitif, un argument selon lequel les organisations devraient chercher à associer la connaissance des décideurs aux attributs organisationnels et environnementaux.

Les essais dans la présente dissertation continuent cette façon de pensée. Chaque essai représente une analyse autonome axée sur un attribut spécifique de la cognition managériale et son interaction avec le processus du satisficing consécutif (notez que les essais II et III partagent une base analytique commune). Les principaux attributs de la cognition managériale visés dans le présent travail sont a) la pensée spontanée, et spécifiquement l'heuristique d'affect (essai I), b) les modèles mentaux imparfaits ou la notion de la réalité mal perçue

(essais II et IV) et c) la complexité de la pensée humaine (essai II et III). Les quatre essais partagent un thème commun de traçage des effets de la cognition managériale et des biais correspondants au fur et à mesure que les organisations recherchent des stratégies efficaces. Le paragraphe suivant décrit chaque essai séparément.

ii. Résumé des essais

Essai I *Re-biaisé: la gestion des préférences intuitives au fil du temps* explore le processus et les conséquences du satisficing consécutif biaisé. Nous montrons que les biais peuvent avoir des effets variant dans le temps, ou des cycles de vie, qui apparaissent dans le processus d'adaptation organisationnelle. Plus précisément, nous constatons que le biais, opérationnalisé comme une préférence intuitive injustifiée, peut améliorer la performance à court terme, mais compromettre la performance à long terme dans des environnements complexes. Ce modèle de comportement est le résultat d'un mécanisme auquel nous nous référons comme récurrence générative: les préférences injustifiées empêchent une expérimentation excessive mais au détriment de l'accumulation de connaissances. Nous analysons plus avant les traitements comportementaux qui peuvent exploiter les avantages de ces effets dynamiques et constatons que l'élimination d'un biais existant (de-biaison) ne peut être bénéfique que pour les organisations opérant dans des environnements très stables. Dans des environnements turbulents caractérisés par une forte incertitude (sans doute les deux propriétés des environnements visés par la plupart des organisations modernes), une stratégie plus efficace peut en fait consister au re-biaison (c.-à-d. changer stratégiquement le biais à son contraire). Cela fournit une perspective nouvelle sur la gestion des biais, car les travaux antérieurs dans des contextes expérimentaux se sont concentrés presque exclusivement sur le de-biaison (par exemple, Wilson et Brekke 1994, Wilson et al., 2002). Les résultats actuels identifient le re-

biaison comme une stratégie non considérée mais potentiellement très efficace pour les entreprises. En plus, nos résultats parlent au domaine émergent de la stratégie comportementale en ce sens que nous montrons que les biais cognitifs peuvent être exploités comme un source de création de valeur (voir Gavetti 2012, Felin et Foss 2009, Levinthal 2011, Powell, Lovallo et Fox 2011).

Essai II *Un parti pris en son terme: la complexité méconnu et les deux aspirations* se concentrent sur le fait que les organisations formulent des problèmes stratégiques, définissent ses objectifs et recherchent des configurations supérieures. Une propriété importante de ce processus est que la compréhension de l'organisation est imparfaite et sujette à des erreurs systématiques en vue de la complexité (c.-à-d. on peut simplifier et on peut compliquer). Le chapitre montre que les organisations peuvent utiliser ces biais dans la formulation du problème pour poursuivre leurs objectifs de manière plus efficace. Plus précisément, le biais envers la complexité domine lorsque le temps ou la performance est plus important pour les organisations. Et inversement, le biais vers la simplicité facilite la réalisation des objectifs organisationnels lorsque le temps et la performance sont aussi importants dans les aspirations des entreprises. Dans l'ensemble, nos résultats suggèrent une règle pratique pour former une compréhension efficace pour guider le choix stratégique. S'il y a une forte préférence soit pour le temps soit pour la performance, il est préférable de compliquer ce choix. Si au contraire, le décideur doit s'inquiéter du temps et de la performance également, une stratégie comportementale plus adaptée c'est de le maintenir simple.

Essai III *La complexité conceptuelle des cadres supérieurs et les dynamiques de la performance organisationnelle* teste empiriquement un corollaire du deuxième chapitre. En

poursuite du travail sur la complexité cognitive, dans le troisième essai de ma dissertation sont utilisées les transcriptions des conférences téléphoniques à mesurer la tendance des dirigeants à reconnaître et intégrer les interdépendances multiples entre les choix organisationnels. Supposant que les individus cognitivement complexes sont moins enclins à commettre des erreurs d'omission lorsqu'ils définissent la structure des interdépendances, l'essai analyse la relation entre la perception erronée de la complexité et de la dynamique de performance. Les résultats préliminaires confirment partiellement mes prédictions, ce qui indique que le gain de la pensée complexe change à mesure que les organisations s'adaptent à un nouvel ensemble de conditions.

Essai IV *Aliens Cognitifs* intègre les primitives de la rationalité limitée avec les fondements de la concurrence et montre que les erreurs idiosyncrasiques dans la compréhension de la réalité par les décideurs peuvent être une source d'avantage concurrentiel. Plus précisément, il met en lumière une tension peu étudiée entre l'exactitude et la concurrence, démontrant que les organisations orientées vers le long terme pourraient vouloir demander aux décideurs une compréhension étrange (plutôt que précise) de la réalité. Ces résultats se rapportent directement au domaine émergent de la stratégie comportementale dans la mesure où ils montrent comment les idiosyncrasies du comportement humain, si elles sont correctement gérées, peuvent conduire à une création de valeur supérieure et à une capture de valeur.

Les décideurs avec un caractère étrange, excentrique et, du point de vue de la majorité, la compréhension absurde de la réalité évalue les alternatives différemment. Même s'ils errent et peut-être même plus souvent que tout le monde, ils voient de la valeur lorsque d'autres ne

voient aucun. Et peut être que c'est l'inexactitude de leur compréhension qui les amène à découvrir des configurations supérieures. En raison de cela, lorsque les concurrents sont nombreux, ce sont ceux qui pensent différents et ont en fait une idée moins précise de la réalité, qui guide les organisations vers des stratégies pour un avantage concurrentiel durable à long terme.

iii. Conclusion générale

a. Aperçu et contribution

Cette dissertation a débuté dans le but de découvrir des modèles de création de valeur dans le processus du satisficing consécutif qui émergent grâce aux idiosyncrasies de la pensée des décideurs. Une base fondamentale, c'est que la rationalité des décideurs est bornée, et les limites à la rationalité ne sont pas identiques pour tous les décideurs.

En reconnaissant cette prémisse de base, le présent travail suggère qu'il peut être rationnel pour les organisations de comprendre et d'utiliser les idiosyncrasies cognitives de leurs décideurs. Comme nous vivons dans le monde du deuxième meilleur, où on n'a pas l'accès à l'intégralité de l'information et des options disponibles, la gestion de l'imparfait peut surpasser l'effort pour (NB: pas possession) le parfait. Pour voir cela, considérez que si l'amélioration du jugement des décideurs est coûteuse (par exemple, le coût de l'augmentation de l'exactitude globale des modèles mentaux des décideurs), il peut y avoir un point où les efforts supplémentaires pour l'améliorer sont moins efficaces que la gestion de la pensée des concepteurs (par exemple, identifier et associer le style cognitif du décideur à une certaine tâche organisationnelle). Ainsi, un corollaire général important de ce travail est un compromis entre l'investissement dans la capacité des décideurs et la gestion de leurs idiosyncrasies.

À un niveau plus spécifique, la contribution de la présente dissertation est double. Tout d'abord, cela avance notre compréhension de l'adaptation organisationnelle en étendant la recherche classique, Simonian (voir essai I ou Fiori 2011) pour rendre compte plus en détail des attributs cognitifs de la prise de décision. La notion de rationalité limitée est liée au processus du satisficing consécutif. Bien que le satisficing consécutif représente effectivement une distinction fondamentale de l'optimisation instantanée, ce n'est pas le seul aspect dans lequel notre rationalité est limitée.

Outre les contraintes purement informatiques accentuées dans le satisficing consécutif Simonien, les décideurs sont également soumis aux effets des instruments cognitifs qui ont évolué pour maximiser notre persistance génétique (plutôt que le profit ou autre fonction objective au niveau individuel) compte tenu des contraintes computationnelles de notre esprit. Un résultat est une rationalité qui est limitée selon des dimensions multiples. Bien qu'il existe un large consensus sur la multidimensionnalité de la rationalité délimitée (voir Fiori 2011), les travaux qui explorent les conséquences organisationnelles du satisficing consécutif sont restés largement séparés des progrès de la psychologie cognitive qui mettent en lumière la complexité de la pensée humaine. Curieusement, cette séparation persiste malgré le fait que Simon lui-même a reconnu d'autres limites à la rationalité humaine (voir, par exemple Simon 1987). Le présent travail commence à combler cette division ad hoc en intégrant des éléments plus larges de la connaissance avec les primitives du satisficing consécutif.

Deuxièmement, ce travail contribue au développement du domaine émergent de la stratégie comportementale en identifiant le biais comme un levier important dans la gestion de la connaissance des décideurs. Une vision commune dans les sciences de l'organisation

considère le biais comme une déviation de la rationalité (voir Arnott 2006). Dans ce sens, un biais est strictement préjudiciable à la fonction objective de l'organisation. Cependant, en regardant le biais comme écart par rapport à un choix "impartial" (plutôt que rationnel), nous pouvons découvrir d'autres moyens pour valoriser la création. Dans cette perspective, le biais est effectivement une variation selon une certaine dimension, potentiellement (mais pas nécessairement) non pertinente. Considérons, par exemple que, compte tenu du degré d'exactitude en termes de reconnaissance des interdépendances, les modèles mentaux peuvent être plus ou moins complexes, c'est-à-dire qu'il peut y avoir une tendance à compliquer ou à simplifier (voir essai II). Dans cet exemple, la notion de partialité en soi ne porte pas directement sur la rationalité du choix du décideur. Cependant, comme le montre l'essai II de ce travail, il en résulte l'atteinte des objectifs organisationnels. Le biais n'est pas une déviation de la rationalité (voir Arnott 2006), mais plutôt un moyen d'efficacité organisationnelle. Le biais du décideur, par conséquent, peut être un instrument de stratégie comportementale.

b. Limites

Comme ce travail comporte les primitives du satisficing consécutif avec les éléments de la pensée humaine, les principales limitations classent en conséquence, c'est-à-dire les limitations liées au satisficing consécutif et les limitations liées à la cognition. Notez que les conditions aux limites suivantes s'appliquent à la dissertation complète, des contraintes spécifiques peuvent être trouvées dans les essais correspondants.

Limitations liées au satisficing consécutif. Une condition de frontière importante de ce travail est qu'il met l'accent sur le satisficing consécutif de recombinaison. Alors que le satisficing consécutif en soi est un comportement omniprésent, elle n'est pas uniforme dans la

manière dont elle se manifeste. Les spécialistes de l'organisation ont analysé de multiples formes de comportement, qui partagent les principales propriétés du satisficing consécutif. Il s'agit notamment du satisficing consécutif de recombinaison (par exemple, Levinthal 1997) et, plus généralement, de recherches sur des paysages accidentés (cf. Levinthal 1997 et Winter et al., 2007), échantillonnage séquentiel (p. ex. Posen et Levinthal 2012), problèmes d'affectation de crédit (par exemple, Denrell et al., 2004) etc. Chaque forme du satisficing consécutif s'applique à un type particulier de problèmes organisationnels. Par exemple, le processus d'échantillonnage séquentiel caractérise la façon dont les organisations découvrent des options supérieures parmi les alternatives. De même, le satisficing consécutif de recombinaison répond à un problème particulier, c'est-à-dire celui de trouver une configuration efficace d'éléments. Un contexte typique auquel cela s'applique est le développement technologique (voir Fleming et Sorenson 2001) ou la recherche d'une stratégie efficace (Rivkin 2000).

En outre, il est important de noter que ce travail considère les organisations comme des agents équivalents aux décideurs. Bien que cela s'applique aux entrepreneurs individuels, des organisations suffisamment petites ou des acteurs suffisamment puissants dans les entreprises, les processus de décision de groupe, les aspects politiques et structurels échappent à la présente thèse. Les processus du satisficing consécutif conjointe (par exemple Knudsen et Srikanth 2014) ou le satisficing consécutif organisationnelle avec plusieurs décideurs (par exemple, Rivkin et Siggelkow 2002, Siggelkow et Levinthal 2005) peuvent impliquer des interactions plus complexes que celles mises en évidence dans ce travail.

Limites liées à la connaissance. Bien que ce travail considère des aspects cognitifs différents (relativement) disjoints, cela ne représente naturellement qu'une fraction de ce qui

constitue la pensée humaine. L'analyse de tous les attributs pertinents de la prise de décision dans un seul travail est évidemment peu pratique. Néanmoins, une limite importante pour les analyses ci-dessus est qu'elles reflètent des interactions isolées avec le satisficing consécutif, c'est-à-dire des effets purs d'un attribut cognitif donné. Les éléments de la connaissance humaine, cependant, ne sont pas séparés les uns des autres. Le système 1 et le système 2, l'intuition et les modèles mentaux coexistent et se mêlent dans l'esprit humain (voir par exemple Alos-Ferrer et Strack 2014). En d'autres termes, les modèles rapportés représentent un esprit simplifié qui est soumis à a) des contraintes de calcul qui déclenchent le satisficing consécutif, et b) un attribut choisi de la pensée humaine.

Les éléments choisis de la connaissance humaine sont en outre modélisés pour capturer une propriété spécifique d'intérêt. Considérons, par exemple, le système 1 où la pensée intuitive est incontrôlée. Bien qu'il existe un large éventail de fonctions que le traitement automatique de l'information accomplit dans notre esprit (voir, par exemple, Kahneman 2011), l'analyse de cette thèse (essai I) réside essentiellement dans la plupart de ses rôles et se concentre sur une seule manifestation, l'heuristique d'affect.

En principe, une telle simplification est inhérente à la modélisation scientifique. Une question est de savoir si cette simplification nous permet de formuler des prévisions utiles. Ce travail rapporte des preuves empiriques cohérentes avec certaines des prédictions théoriques. Cependant, cela ne rejette pas le besoin de recherches supplémentaires pour bien comprendre les interactions de la pensée humaine avec le processus du satisficing consécutif.

c. Recherches futures

La présente dissertation développe une base théorique pour gérer les éléments de la connaissance des décideurs dans le temps, pendant que les organisations s'adaptent. Une question juste est de savoir comment les organisations pourraient le faire et si les tendances identifiés existent dans la réalité. Bien que les essais spéculent sur ces sujets, établir comment les organisations peuvent en pratique gérer la connaissance de leurs décideurs pour extraire une plus grande valeur du processus du satisficing consécutif représente un terrain fructueux pour des recherches futures.

Une avenue particulièrement intrigante consiste à identifier les leviers spécifiques que les organisations peuvent utiliser pour extraire la valeur des idiosyncrasies cognitives des décideurs. Les exemples incluent la rotation de la gestion (p. ex., le re-biaison, voir la section 2.4.3.), l'architecture de la prise de décision (p. ex. la complexité de la décision réalisée, voir la section 3.6.), les politiques de recrutement et de promotion (p. ex., pensées étranges, voir la section 5.6), etc.

De la même manière, les organisations peuvent potentiellement imiter certains des effets individuels spécifiques avec leurs éléments structurels. Ethiraj et Levinthal, par exemple, discutent des asymétries de performance de sous- et sur- modularité (2004). Leur analyse met en relief l'idée que les organisations devraient gérer les types d'erreurs qu'elles produisent pour spécifier certains aspects de la réalité objective. Ce genre de raisonnement peut être utile à la lumière des résultats signalés dans le présent travail. Considérons que les organisations ont des moyens multiples pour réguler la complexité relative de leurs processus internes, y compris l'articulation et la codification des connaissances (Zollo et Winter 2002),

ou le degré d'automatisation dans les processus de production (Camuffo et Volpato 1996) entre autres. En établissant si les effets sur la performance de ces choix et des choix similaires suivent, la dynamique identifiée dans les essais II et III peut révéler d'autres moyens de gérer l'adaptation organisationnelle.

La thèse selon laquelle la trajectoire adaptative de l'organisation dépend de la connaissance des principaux décideurs suggère que des organisations multiples et interdépendantes, à savoir les écosystèmes organisationnels, peuvent être façonnées par la synergie (ou la collision) des styles cognitifs. La littérature académique aussi bien que la littérature axée sur les praticiens porte l'accent sur la nature vitale des processus conjoints et coévolutifs dans les industries (voir par exemple Adner et Kapoor 2010, Iansity et Levien 2004, McKelvey, 1999). Dans cette perspective, une action d'une organisation définit les conditions pour une autre. Les industries, ou les écosystèmes d'organisation, peuvent donc différer en fonction de la façon dont leurs décideurs multiples pensent. De même, cela s'applique aux contextes intra-organisationnels où les décisions des différentes divisions et départements dépendent les uns des autres (voir par exemple Rivkin et Siggelkow 2002, 2003).

En outre, il existe un potentiel important pour analyser l'interaction entre les éléments de la connaissance humaine et les instruments d'architecture qui modifient la dynamique du satisficing consécutif. Par exemple, la modularité représente un sujet intrigant dans cette perspective (voir ci-dessus, Brusoni et Prencipe 2001, Brusoni et al., 2007, Ethiraj et Levinthal, 2004).

Une autre extension naturelle du présent travail est de lever les principales limitations et conditions aux limites. Tout d'abord, il peut être utile d'envisager d'autres manifestations du processus du satisficing consécutif. Bien que les résultats signalés dans ce travail soient fiables dans le cadre du problème de recombinaison, dans le cadre d'autres types de problèmes organisationnels (par exemple l'échantillonnage séquentiel), ils peuvent être différents. De même, l'élargissement du spectre des propriétés considérées de la pensée humaine peut révéler un potentiel supplémentaire d'imperfections dans la cognition managériale pour la création de valeur. Un chemin très prometteur dans cette direction est de rendre compte des émotions humaines et de leur interaction avec le processus du satisficing consécutif. Les émotions interfèrent avec la prise de décision (par exemple Lerner, Small, Loewenstein 2004, Maitlis et Ozcelik 2004, Simon 1987). Par extension, cela signifie que les émotions affectent le processus du satisficing consécutif. Cependant, comme le suggère le présent travail, la compréhension des résultats d'une telle influence peut nécessiter une analyse contrôlée. Enfin, il existe une possibilité d'interactions complexes entre les multiples éléments de la cognition humaine et leurs implications communes pour le processus du satisficing consécutif. Tout en étudiant ces divers voies d'interactions impliquent d'importants obstacles informatiques et analytiques, leur analyse permettra une meilleure approximation du potentiel de notre esprit, des stratégies plus efficaces et, en fin de compte, de meilleures prédictions.

Abstract

This work centers on the tenet that organizational rationality is bounded: decision makers search, satisfice, and think in a way that is typical (in its integrity) only of humans. The dissertation explores this interplay between search and decision maker's cognition and demonstrates how biases in characteristic aspects of our thinking can be instruments of behavioral strategy.

As a starting point, I take search, sequential generation and evaluation of alternatives, as the first primitive of bounded rationality and complement it with integral elements of human cognition, such as automatic, intuitive thinking, specifically affect heuristic, and imperfect mental representations of reality. With the help of computational models, I track the effects of the corresponding biases (systematic affective preferences and systematic errors in mental representations) over time as organizations adapt to complex environments. This allows me to identify life cycles of the elements of human cognition and show that organizations should manage (rather than eliminate) some biases over time. Finally, I derive predictions and empirically test a subset of my propositions.

In conclusion, this work aims to advance the emerging theory of behavioral strategy by jointly considering different primitives of bounded rationality and integrating them with the existing knowledge in organization sciences. A broad question that motivates this work is how organizations can manage the many bounds to human rationality.

Contents

Résumé	ii
Abstract	xvii
Contents	xviii
Acknowledgements	xxii
General Introduction	1
1.1. Subject matter	1
1.2. Essay abstracts	4
Essay I. Re-biased Search: Managing Intuitive Preferences over Time	7
2.1. Introduction	8
2.2. Search and Intuitive Preferences	11
2.3. A Model of Biased Search	16
2.3.1. Modeling Intuitive Bias	17
2.3.2. Model limitations	22
2.4. Results	25
2.4.1. The Dynamics of Biased Search	26
2.4.2. Mechanism	28
2.4.3. Re-Biased and De-Biased Search	32

2.4.4.	The Optimal Timing of Behavioral Treatments.....	36
2.5.	Conclusion and Discussion	39
Essay II. A Bias to an End: Misconceived Complexity and Dual Aspirations.....		44
3.1.	Introduction.....	45
3.2.	Search and aspiration levels.....	48
3.3.	Problem formulation and misspecified Interdependencies.....	51
3.4.	Model.....	55
3.4.1.	Objective strategy landscape.....	56
3.4.2.	Subjective understanding.....	57
3.4.3.	Search process.....	60
3.5.	Analysis.....	61
3.5.1.	Likelihood of attainment.....	62
3.5.2.	Effect of biases.....	64
3.5.3.	Underlying mechanisms.....	68
3.5.4.	Aspiration level adaptation.....	70
3.5.5.	Problem reformulation.....	73
3.6.	Intentionally biased.....	74
3.7.	Conclusion and Discussion.....	77
Essay III. CEO's Conceptual Complexity and Performance Dynamics.....		82
4.1.	Introduction.....	83

4.2. From model to data	84
4.2.1. Empirical challenges	84
4.2.2. Predictions and context	87
4.3. Main measures	88
4.3.1. Performance and time	88
4.3.2. Conceptual complexity	89
4.4. Method	91
4.5. Results and conclusion	96
4.5.1. Results	96
4.5.2. Conclusion	98
Essay IV. Cognitive Aliens	99
5.1. Introduction	100
5.2. Bounded rationality	102
5.2.1. Search. Evaluation of alternatives.....	103
5.2.2. Mental representations	104
5.2.3. Cognitive aliens: Premise.....	106
5.3. Competitive pressure	107
5.4. Model	108
5.4.1. Objective landscape	109
5.4.2. Cognitive aliens: A formal definition	110

5.4.3.	Performance and search	112
5.5.	Analysis	114
5.5.1.	Terminal strategies.....	115
5.5.2.	Performance differential	116
5.5.3.	Mechanism.....	118
5.5.4.	Lower accuracy	120
5.5.5.	Low accuracy as cognitive distinctness	123
5.5.6.	Imitation.....	125
5.5.7.	Robustness	128
5.6.	Conclusion and discussion.....	129
General Conclusion		134
6.1.	Overview and contribution	134
6.2.	Limitations.....	136
6.2.1.	Search-related limitations	136
6.2.2.	Cognition-related limitations	137
6.3.	Future research	138
Bibliography.....		141

Acknowledgements

As any expedition, this dissertation owes its greatest thanks to curiosity. It has persisted with me over the years at HEC. Here, in the quietude of Jouy-en-Josas, I have been fortunate to work with people who have taught me, among other things, to question my assumptions and never lock myself in the frames of my current understanding. To these people, I am grateful.

First, for a chance to find a subject that would truly interest me. Perhaps, if it were not for Gonçalo Pacheco de Almeida, my main advisor who encouraged me to search for an independent identity, I would not have explored broadly enough to complete or even begin this dissertation. Gonçalo's advice, continuous support since the very first days and belief in his student have shaped this work as well as parts of myself.

Second, for an example of a scholarly discussion. In our countless meetings and occasional walks with Tomasz Obloj, I have found a model of an intellectual conversation: receptive, straightforward, focused on meaning rather than form. Indeed, so open were our talks that my guess is that Tomasz may underrate the significance of this learning experience for his student. His training has forged my knowledge in virtually every aspect of research.

Third, for a lesson in integrity, professional and personal. During my early years at HEC, I was fortunate to meet Eric Luis Uhlmann. Little did I know how much Eric would mean for my research as well as for me. Besides his invaluable guidance in the vastness of psychology, always-available feedback and advice, he gave me to see many a principle of professional and personal conduct that I want to stand by.

Fourth, for a memory that will stay with me forever. Thanks to our breakfast talks with Giada Di Stefano, I now know that a week will be a ray brighter if started at 7 a.m. on Monday at a meeting with a very good person. To Giada I am indebted for a far better, however much remains for me to learn, understanding of the social aspects of decision-making, strategy and life.

Also deeply grateful, I am to those who shaped my knowledge as well as this work's progress with their courses, advice and encouragement: Olivier Chatain, Mark Desjardine, Rodolphe Durand, Pierre Dussauge, Gary Dushnitsky, Alfonso Gambardella, Giovanni Gavetti, Oliver Gottschalg, Samina Karim, Elie Matta, John Mawdsley, Corey Phelps, Joe Porac, Bertrand Quélin, PK Toh, Dmitry Volkov, Maurizio Zollo and Thomas Åstebro among others. To Stefano Brusoni, Arnaldo Camuffo, Margaret Kyle and Philippe Monin, my gratitude extends also for taking the roles on my committee. This urges me to thank this dissertation itself, for it has given me a chance to meet great people. Even a brief interaction with Ron Adner will inspire me long into the future, regardless of my work's matter.

In the course of this work, I also met friends: João Albino Pimentel, Moumita Das, Ashkan Faramarzi, Cédric Gutierrez Moreno, Pier Vittorio Mannucci, Timofey Shalpegin, Ali Shantia, Sara Rezaee Vessal, Jiachen Yang among others. Ali, Cédric, João, Pier and Sara not only shared with me the minutiae of our students' life but also helped form parts of this work. To Tim, I am thankful for our long tea debates that taught me to find counsel in a very different, sometimes opposite, point of view and every now and then enjoy being proven wrong.

Finally, my family supported me in ways that only those dear to us can. To my mother and father, I dedicate this work. Эту работу я посвящаю моим маме и папе.

My hope is that this dissertation only begins to reveal what I have learnt, in work and in life, from those who supported, mentored and encouraged me. Thank you.

General Introduction

1.1. Subject matter

Human rationality is bounded (Conlisk 1996, Fiori 2011, Simon 1955, 1957). From the perspective of the Carnegie (neo-Carnegie) School, this means i) that organizations search, *i.e.* sequentially generate and evaluate alternatives, and satisfice rather than optimize over the entirety of available options, and ii) that decision makers' thinking is a complex process that is not perfect (see Cyert and March 1963, Gavetti, Levinthal, Ocasio 2007, Kahneman 2003, Simon 1955, 1956). While the existing literature offers a sizeable insight into these subjects, the joint analysis of (i) and (ii), remains scarce (see Gavetti *et al.* 2007, Powell, Lovallo, Fox 2011). This broad gap in the existing literature is at the core of the present dissertation.

The objective of this work is to identify (potentially useful) patterns in the dynamics of search that arise due to decision makers' cognitive idiosyncrasies. Let us focus on each element of this objective separately. First, search is a means to organizational learning and adaptation (Levinthal 1997). While there are multiple ways in which organizations learn and adapt (*e.g.* sequential sampling, vicarious learning, population ecology, *etc.*), the particular process of adaptation considered throughout this work is recombination search. This process characterizes reconfiguration of organizational knowledge, technological development or more broadly problem solving (see for example Fleming and Sorenson 2001, Nickerson, Silverman, Zenger 2007). In essence, it represents consecutive decision making wherein agents generate options or new combinations and evaluate them (Knudsen and Levinthal 2007). Second, the focus on cognitive idiosyncrasies implies comparison of different ways in

which decision makers think, such as having vs. not having a systematic intuitive preference (essay I) or having a more complex vs. more simple understanding of reality (essay II). By recognizing the variance in decision makers' cognition (on heterogeneity in human thinking see Baron 2007), this dissertation undertakes to develop a more nuanced understanding of organizational adaptation, one that would better inform organizations on the limit as well as potential of the human mind, and thereby serve to advance management efficiency.

This defines the academic domain to which the present work belongs, *i.e.* behavioral strategy. Behavioral strategy is an emerging subfield of organization sciences that seeks to uncover cognitive-behavioral sources of superior value creation, or, more broadly, regards strategy phenomena through behavioral lenses (Powell *et al.* 2011). As such, this subfield of knowledge is grounded in the behavioral theory of the firm, and the broader Carnegie School (see Gavetti 2012). Indeed, one may view the emancipation of behavioral strategy (as a distinct domain) as a response to the call for "a renewed *behaviorally plausible, decision-centered* perspective on organizations" (Gavetti *et al.* 2007: p. 531).

A key property that arguably identifies behavioral strategy is the focus on the use (or misuse) of individuals' psychology for value creation and capture in organizations (see Gavetti 2012). Works that directly relate to this subject date back to at least 1980s (see for example Schwenk 1984, 1986). However, a more inclusive investigation of this subject falls on the later years and somewhat concurs with the pieces that define the notion of behavioral strategy with its roots in positive rather than idealized assumptions about decision-making (see also Levinthal 2011, Powell *et al.* 2011).

Within this literature, a subset of works sets the logic for the present dissertation. The key developments in this logic synthesize as follows. First, drawing on the behavioral theory of the firm as well as earlier work in organizational adaptation and learning (*e.g.* Cyert and March 1963, Levinthal 1997, Levitt and March 1988), Denrell and March illustrate that a bias against risky alternatives can be an outcome of learning and selection (2001). This in turn may have detrimental consequences for subsequent adaptation as organizations persist with safer (but suboptimal) alternatives. Mirroring this conclusion in a later work, Winter and colleagues show how moderate obsession with a particular course of action can lead to superior performance in the long run (Winter, Cattani, Dorsch 2007). Unlike risk aversion, some level of conscious overvaluation of certain options can pay off in the future. The reason is that evaluation errors allow for higher level of exploration (Knudsen and Levinthal 2007, Baumann and Martignoni 2011). An analysis of the interaction of confidence and evaluation quality further illustrates that confidence can partially compensate for competence (Posen, Martignoni, Lang 2013). While these works demonstrate how elements of human behavior can be both a by-product of search and a determinant of the payoff to adaptation, they focus on the specific outcomes of alternatives evaluation without accounting for the cognitive processing that underlies errors in judgement (*e.g.* imperfect mental models) or precede conscious choice (*e.g.* spontaneous, affective reactions).

A recent stream of work brings to light the significance of such elements of cognition, in particular imperfect mental models, for the process of organizational adaptation. Csaszar and Levinthal (2016) study the choice between searching mental representations and searching for an efficient strategy given decision maker's mental representation of the organizational problem. They demonstrate that incomplete mental models can facilitate

exploration and thereby lead to better organizational outcomes. In a similar vein, Martignoni and colleagues (2016) study the consequences of under- or overspecified mental models. Their analysis indicates that given some combination of environmental complexity and organizational control over choices, misspecified mental models can lead to superior outcomes. Accordingly, they emphasize the importance of cognitive fit, an argument that organizations should seek to match decision makers' cognition with organizational and environmental attributes.

The essays in the present dissertation continue this line of thought. Each essay represents a standalone analysis focused on a specific attribute of managerial cognition and its interaction with the process of search (note that essays II and III share a common analytical basis). The main attributes of managerial cognition viewed in the present work are a) spontaneous thinking, and specifically affect heuristic (essay I), b) imperfect mental models or the notion of misperceived reality (essays II and IV), and c) complexity of human thought (essay II and III). The four essays share a common theme of tracing the effects of managerial cognition and corresponding biases *over time* as organizations search for efficient strategies. The following subsection describes each essay separately.

1.2. Essay abstracts

Essay I *Re-biased Search: Managing Intuitive Preferences over Time* extends the model of organizational adaptation on rugged landscapes and examines the dynamics of biased search. Drawing from work on the psychology of human decision-making, the essay construes biases as unjustified preferences that arise due to automatic, spontaneous thinking. This property of

decision-making gives rise to a mechanism labeled generative recurrence. In the presence of this mechanism, unjustified intuitive preferences produce two opposing effects on organizational adaptation: they curb excessive experimentation, but at the expense of knowledge accumulation. In the context of organizational search, these regularities allow leveraging the value of biases. Specifically, the results suggest that re-biasing (adopting the opposite bias) often dominates both de-biasing (eliminating the bias) as well as consistently unbiased search. The paper, therefore, provides evidence that managing rather than eliminating intuitive biases can be an effective instrument of behavioral strategy.

Essay II *A Bias to an End: Misconceived Complexity and Dual Aspirations* focuses on the fact that organizations formulate strategic problems, set objectives and search for superior configurations. An important property of this process is that the organizational understanding is imperfect and prone to systematic errors in seeing complexity (*i.e.*, oversimplifying and overcomplicating). The chapter shows that organizations can use such biases in problem formulation to pursue their objectives more efficiently. Specifically, bias toward complexity dominates when either time or performance is more important for organizations. And inversely, bias toward simplicity facilitates the achievement of organizational goals when both time and performance are equally important in the firms' aspirations. Overall, this chapter suggests that when formulating strategic problems—and in particular when conceiving of complexity—organizations may want to be intentionally biased rather than unbiased.

Essay III *CEO's Conceptual Complexity and Performance Dynamics* empirically tests a corollary from the second chapter. Drawing from the work on cognitive complexity, the third

essay of my dissertation uses conference call transcripts to measure executives' tendency to recognize and integrate multiple interdependencies among organizational choices. Assuming that cognitively complex individuals are less prone to make errors of omission when defining the structure of interdependencies, the essay analyzes the relationship between executives' misperception of complexity and performance dynamics. Preliminary results partially support my predictions, indicating that the payoff to complex thinking changes as organizations adapt to a new set of conditions.

Essay IV *Cognitive Aliens* integrates the primitives of bounded rationality (search on rugged landscapes and imperfect mental models) with the fundamentals of competition and shows that idiosyncratic errors in decision makers' understanding of reality can be a source of competitive advantage. More specifically, it brings to light an understudied tradeoff between accuracy and competition, demonstrating that organizations oriented toward the long term may want to seek decision makers with an odd (rather than accurate) understanding of reality. These findings speak directly to the emerging field of behavioral strategy in that they show how idiosyncrasies of human behavior, if managed properly, may lead to superior value creation and value capture.

The conclusion to this dissertation highlights the dissertation-wide limitations, contributions, and directions for future research. The present work aims to advance the emerging theory of behavioral strategy by jointly considering different primitives of bounded rationality and integrating them with the existing knowledge in organization sciences. A broad question that motivates this work is how organizations can manage the many bounds to human rationality.

**Essay I. Re-biased Search: Managing Intuitive
Preferences over Time**

2.1. Introduction

Decision making is prone to the effects of intuitive thinking, most notably biases (*e.g.* Kahneman 2003, Tversky and Kahneman 1974). While existing work in organizational science and psychology considers strategies that allow de-biasing or, under some conditions, biasing of organizational choices (Christensen and Knudsen 2010, Schwenk 1986, Wilson and Brekke 1994, Wilson, Centerbar, Brekke 2002, Winter, Cattani, and Dorsch 2007), we show that dynamically re-biasing search—that is, reversing biases during the search process—is often a strictly dominant behavioral strategy. To do so, we integrate a first primitive of biased decision-making—namely, the presence of spontaneous, intuitive thinking and inherent heuristic preferences—with a standard model of search. This allows a more cognitively nuanced representation of organizational choice and the identification of *generative recurrence*, a behavioral mechanism at the interplay of organizational search and intuitive thinking that determines the value of biases over time.

Our work builds on and contributes to the behavioral theory of the firm that aims at integrating the positive principles of human behavior into the study of organizations (Cyert and March 1963, Gavetti, Levinthal, Ocasio 2007). Unlike the hyper-rational firms populating neoclassical models, the organization seen through the lenses of behavioral theory does not instantaneously optimize, and instead must engage in a lengthy process of search (*e.g.* Levinthal 2011, Simon 1955). This process of organizational search, in turn, is influenced by the fundamentals of human decision making (Knudsen and Levinthal 2007). An important behavioral attribute of decision making is that it is a product of both deliberate and intuitive

responses (*e.g.* Sloman 1996, Gawronski and Bodenhausen 2006, Greenwald and Banaji 1995, Wilson, Lindsey, Schooler 2000, Pacini and Epstein 1999). This former, intuitive, component represents an important element of human judgment—decision makers routinely call on their intuition or “gut feeling” when making both day-to-day and strategic choices (Khatri and Ng 2000, Miller and Ireland 2005). Despite its significance, the intuitive side of decision making remains understudied by the scholars of organizations as well as in the organizational sciences more generally (Dane and Pratt 2007). Yet, without explicitly accounting for this fundamental component of cognition, we can only partially understand the implications of bounded rationality for organizational processes and performance.

From the behavioral perspective, intuitive thinking often results in unjustified preferences that shape organizational choices before conscious reasoning is engaged (Inbar, Cone, Gilovich 2010, Kahneman 2003). We examine the effects of such biases in the process of organizational adaptation, focusing on their dynamic rather than static consequences. While existing work has pointed to the possibly adaptive role of biases (*e.g.* Gigerenzer, Todd and the ABC Research Group 1999, Johnson and Fowler 2011), we analyze the lifecycles of biases and demonstrate that time is an important factor in *managing* them. Indeed, organizations—unlike individuals—possess instruments to calibrate and manipulate biases, such as changing decision-making processes, redesigning organizational structures, or replacing key decision makers (see for example Christensen and Knudsen 2010, Csaszar 2012). That is, organizations have structural and contextual means to alter the effective biasedness of their decisions, and therefore can proactively manage their effects.

Our work specifically assesses the effectiveness of two basic strategies that organizations can use to manipulate biases: *de-biasing*, or entirely eliminating a bias, and *re-biasing*, or adopting the exact opposite unjustified preference, as well as their optimal timing. Although de-biasing interventions have been widely attempted in experimental laboratory research with mixed success (Wilson and Brekke 1994, Wilson *et al.* 2002), to our knowledge, re-biasing has not been previously considered as a strategic option at either the individual or organizational level. Further highlighting the moderating role of time, we show that the effectiveness of both behavioral treatments hinges critically on the stage of organizational search. These observations contribute to an emerging behavioral research stream seeking to discover how organizations "can act with intelligence and efficacy in strategic contexts," where decisions are products of human rather than machine judgment (Levinthal 2011: 1521, cf. Ethiraj and Levinthal 2009, Fiori 2011, Kahneman 2003, Simon 1990).

Overall, our analysis provides an account of a well-documented phenomenon of intuitive thinking and associated biases in a parsimonious model of search. While existing work has considered properties of generation of alternatives that can be partially explained by intuitive decision making (such as local search), we are not aware of any prior work that provides a direct analysis of the consequences of dual thinking in the process of organizational choice. We begin to fill this gap by addressing the following research questions: How do intuitive biases affect the process and effectiveness of organizational adaptation over time? Can organizations intervene to manage such biases to create value? And if so, when and how?

2.2. Search and Intuitive Preferences

Organizational search is a sequential decision-making process by which firms seek to improve their performance over time (Levinthal 1997, Siggelkow 2001, 2002). While there are numerous properties that characterize this process, its fundamental attribute is that organizations cannot fully and faultlessly comprehend the entire space of alternatives and therefore need to experiment by trying new options (Gavetti and Levinthal 2000, Knudsen and Levinthal 2007).

Originating from Herbert Simon's seminal works, this view of organizational decision making replaces "the global rationality of economic man with a kind of rational behavior that is compatible with the access to information and the computational capacities that are actually possessed by organisms, including man" (Simon 1955, p. 99). Full rationality, from this perspective, is not attainable because the human mind cannot process the entirety of available information, let alone compare all alternatives and perfectly evaluate their mapping to the objective performance function (Levinthal 2011). The decision-making process, therefore, is akin to that of artificial intelligence—a system that encodes and interprets information in terms of symbols (Simon 1990, Newell and Simon 1972). Since processing of a coded signal requires time and cognitive effort, complex problems may actually be prohibitively costly to optimize and as a result "intelligent systems must use approximate methods to handle most tasks" (Simon 1990, p. 6). Specifically, human behavior—and, by extension, organizational decision making—is more likely governed by a search for continual improvement rather than by an unconstrained optimization process (Simon 1955).

Although this view of algorithmic search significantly advances our understanding of organizations from the perspective of the behavioral theory of the firm, it still omits some fundamental elements of human psychology. In particular, the existing specifications of search come short of fully embracing the largely intuitive nature of decision making. Individuals form spontaneous judgments that predetermine choice and lead to systematic and predictable biases (Kahneman 2003).

Organizational scholars have stressed the importance of biases since the early works on organizations. Cyert and March, for example, state that “[...] search is biased. The way in which the environment is viewed and the communications about the environment that are processed through the organization reflect variations in training, experience and goals of the participants in the organization” (1963, p. 121). Similarly, Isaack (1978) and later Dane and Pratt (2007) stress the centrality of human intuition for organizational decision making, while Schwenk (1984, 1986) points to the crucial role of cognitive processes and the accompanying biases in the formation of strategy. However, despite the importance of intuition and concomitant biases for human behavior and firms’ strategies, they have received only incidental treatment in the analysis of organizational adaptation (notable exceptions include Winter, Cattani, Dorsch 2007, Lounamaa and March 1987). The few works that explicitly examine idiosyncrasies of human judgment in the process of search essentially equate biases with imperfections of evaluation (see Winter *et al.* 2007, Knudsen and Levinthal 2007).

Indeed, in the existing analyses, the effect of biases tends to be reduced to that of a systematic error in evaluation. When search is sufficiently local, moderate levels of evaluation error can support better organizational performance (Baumann and Martignoni 2011, Knudsen

and Levinthal 2007, Winter *et al.* 2007). Similarly, Denrell and March show that inaccuracies in learning can diminish the hot stove effect—that is, decelerate "stabilization of knowledge concerning alternatives that have resulted in early failures" and thereby bolster organizational performance (2001, p. 534). Further, erring in evaluating the probability of failure fuels innovation as overconfident decision makers are more likely to pursue risky projects (Galasso and Simcoe 2011). Although these works jointly indicate that the role of biases in the process of organizational adaptation is not unidimensional and not necessarily negative, they largely underemphasize the intuitive nature of biases.

Biases often arise due to a reliance on heuristics—or cognitive simplification mechanisms—that link one's choice to a certain immediately observable attribute of options (Gilovich and Griffin 2002, Tversky and Kahneman 1974). Inasmuch as this attribute is uncorrelated with the underlying value of alternatives (*i.e.* it is unjustified), heuristics and associated biases that occur in the process of spontaneous intuitive reasoning may alter the process of generation of alternatives leading to stable choice outcomes that are qualitatively distinct from evaluation errors. We focus on this aspect of heuristics and biases and analyze the general search implications of having unjustified intuitive preferences. The existing literature identifies a wide spectrum of intuitive biases or spontaneous "response[s] because of mental processing that is unconscious or uncontrollable" (Wilson and Brekke 1994, p. 117). These biases systematically contaminate decision making, often without the person's awareness of their influence. Indeed, such blindness to the rationale behind one's own choices reflects the complexity of human thought (Greenwald and Banaji 1995, Haidt 2001, Kahneman, Lovallo, Sibony 2011, Nisbett and Wilson 1977). Extensive research in psychology indicates that human cognition involves the simultaneous functioning of two

systems. One is spontaneous, intuitive, uncontrolled and fast—this system (System 1) is based on the law of association. The other is deliberate, effortful and relatively slow—this system (System 2) can be said to rely on the law of logic (Stanovich and West 2000).

The responses of these two Systems are not always aligned. In situations in which System 1 dominates System 2 (*e.g.* limited time, high cognitive load, or when the judgment is closer to perception than to deliberate assessment), the decision maker's judgment is especially likely to deviate from the rules of logic. Therefore, we cannot assume that intuitive preferences will be "reasonable by the cooler criteria of reflective reasoning. In other words, the preferences of System 1 are not necessarily consistent with preferences of System 2" (Kahneman 2003, p. 1463).

One reason is precisely that the functioning of System 1 is closely linked to the use of heuristics, or mental shortcuts that allow for speeded and simplified decision making (Frederick and Kahneman 2002, Stanovich and West 2000). While such intuitive rules of thumb are variegated (Hutchinson and Gigerenzer 2005), in the context of choice, heuristics can be fundamentally reduced to a form of preference for a certain—immediately observable or perceivable—attribute of options. That is, "pick A, if A is" more readily accessible, more representative of a category, implies lesser losses, *etc.* Such a simplified representation of a heuristic can be directly linked to automatic evaluations that are computed by System 1 upon perception (Duckworth *et al.* 2002, Fazio 2001, Fazio, Sanbonmatsu, Powell, Kardes 1986, Slovic, Finucane, Peters, MacGregor 2002, Zajonc 1980). This basic intuitive reaction can substitute deliberate evaluation in any task that requires a favorable or unfavorable response, or simply a choice. Accordingly, Slovic *et al.* (2002) used the notion of the affect heuristic and

showed that liking or disliking—or equivalently having a preference—is the underlying heuristic attribute for numerous target choices (*e.g.*, assessments of technologies, safe levels in the use of chemicals, *etc.*).

Therefore, an unjustified intuitive preference represents a primitive heuristic that permeates human decision making and, at the level of a population of choices, leads to an observable bias. To see this consider that if System 1 at least at times dominates System 2—that is if the judgment is not fully deliberate—there will be instances, where the decision is made based on the unjustified intuitive preference. Choices that satisfy the intuitive and automatic preference will therefore be overrepresented in the population of decisions.

Importantly, the presence of biases is not uniform over all stages of the decision-making processes. Specifically, the greater the involvement of System 1, the more liable to bias the choice is. This happens because intuitive judgments originate "between the automatic parallel operations of perception and the controlled serial operations of reasoning" (Frederick and Kahneman 2002, p. 50). The effect of intuitive preference, therefore, may be consequential for organizational search, as this process involves both deliberate choice and generation of alternatives, which is closer to perception.

A basic premise behind our analyses is that individuals vary in their proclivity to rely on spontaneous, intuitive judgments as well as in the direction of the bias they display (Cacioppo, Petty, Feinstein, Jarvis 1996, Epstein 1990, Frederick 2005, Kruglanski and Webster 1996, Nosek, Banaji, Greenwald 2002, Pacini and Epstein 1999). In effect, people can exhibit no bias or even a bias that is opposite to that prevalent in the population (Baron 2007). Therefore, while the impact of intuitive preferences for the process of search is likely

substantial at the population level, it differs from one individual to another. For organization this implies that the psychology of its decision makers—in particular, the extent to which they are prone to rely on an intuitive bias when constructing the set of considered alternatives—potentially represents a key consideration for behavioral strategy. In the following section we construct a model of organizational search to analyze whether this is indeed so.

2.3. A Model of Biased Search

Consistent with the existing work on organizational adaptation, we think of search as a recombination process occurring in the space of interrelated decisions. The mapping from organizational configurations to performance is defined using the NK model (Kauffman 1993, Levinthal 1997). In short, the NK model allows for a representation of an organizational state as an N -large sequence of binary decisions that can take the value of either 1 or 0. Each of these configurations (forms, alternatives, options or courses of action) corresponds to a certain level of performance—or, in terms of organizational adaptation, a position on the fitness landscape. By construction, the geometry of the landscape depends on the extent and pattern of interactions among the N decisions (Rivkin and Siggelkow 2007). Specifically, when the amount of interdependencies, $K \in \{0, 1, \dots, N - 1\}$, is large ($K \rightarrow N - 1$), complexity is high and the performance landscape is rugged (Levinthal 1997). On the contrary, when decisions tend to be independent ($K \rightarrow 0$), complexity is low and local maxima are rare—and absent when K is at the minimum. Therefore, K is the critical parameter that determines the efficacy of local search (Levinthal 1997). A more detailed description of the technical apparatus of the

baseline NK model can be found in subsection 5.4.1 (for intuition, see also subsection 3.4.1) or elsewhere (see, for example, Ganco and Hoetker 2009).

A long tradition of theoretical work emphasizes the tendency of organizations to search locally, that is, to generate alternatives within the immediate vicinity of their status quo and at least partially rely on online experimentation (*e.g.* Cyert and March 1963, Levinthal 1997, Stuart and Podolny 1996). Accordingly, in the model, we follow this line of work and consider strictly local experiential search. In this process, organizations generate alternatives by randomly flipping one choice in their current configuration, and accept only those alternatives that have not been tried in the past. Should the new alternative worsen performance, organizations return to the previous configuration in the period following the trial. While in the interest of clarity we only report our observations from this simplified model, our general results are robust across different specifications, including a combination of local and distant search. A detailed discussion of the distance of search, the possibility of offline evaluation, and other attributes of the search process are included in the Appendix.


2.3.1. Modeling Intuitive Bias

The described sequential process characterizes an organization as an algorithmic, purely computational system myopically searching for better configurations. Although partially embracing the computational constraints of the human mind and its inability to optimize (along with organizational rigidities), this model of organization does not fully capture the effect of automatic intuitive preferences on human decision making (Gawronski and Bodenhausen 2006, Greenwald and Banaji 1995, Wilson *et al.* 2000). Our model accounts for

these intuitive attributes of human reasoning by extending prior standard models of organizational search, as follows.

Unlike most agents in prior models—whose judgment is unaffected by spontaneous responses—our agents perceive not one but *many* alternatives. Immediately upon discovering these alternatives, our agents intuitively select one of them to be considered more thoroughly (see Figure 2.1). The basic rationale is that organizations simply cannot carefully evaluate all ideas that come to mind, and hence, decision makers automatically discard some of these alternatives. As a result, there is a temporal and qualitative distinction between intuitive and deliberate judgment. [continued on the next page]

Figure 2.1. Process of Biased Search



Notes. (a) end of intuitive phase; (b) organizations deliberately assess, *i.e.* compare to previous trials, one alternative per period; (c) an extended model described in the Appendix further accounts for the formal evaluation stage as well as for imperfections in memory.

This representation is consonant with research in psychology, which demonstrates that human intuitive reactions do not require conscious reasoning and can occur automatically even when the stimuli are novel (Duckworth *et al.* 2002, Fazio 2001, Fazio *et al.* 1986). Our characterization is therefore consistent with the co-existence of System 1 and System 2 cognition. Further, intuitive, spontaneous processing can precede and operate independently of more conscious processes (Greenwald and Banaji 1995, Wilson *et al.* 2000, Zajonc 1980). Careful, deliberate evaluation of alternatives, therefore, is likely to be distinct from the more intuitive reactions of decision makers, which typically take place well before full information regarding the alternatives' expected value is available (Duckworth *et al.* 2002, Gawronski and Bodenhausen 2006, Wilson *et al.* 2000, Zajonc 1980). As modeled in our paper, intuitive responses take place between the automatic processes of perception and the more deliberate processes of reasoning (Kahneman 2003, Zajonc 1980).

Any ranking of options during intuitive, spontaneous phase is based on purely subjective interpretations, associations, and preferences—and, therefore, may be especially prone to biases. Accordingly, we operationalize bias as a preference for certain (randomly chosen *ex ante*) organizational configurations. Note that, while arbitrary and random *ex ante*, the underlying bias is consistent over time. That is, biased agents exhibit the same unjustified intuitive preference over the entire process of search. We assume that preferences are random to show that our results hold for any type of underlying intuitive preference.

For simplicity, and without loss of generality, we take the number of perceived or generated alternatives to be equal to two (our results remain qualitatively unchanged for a

greater number of generated alternatives, as shown in the Appendix). In our model, only one of the two generated alternatives passes to the stage of deliberate, conscious evaluation, whereas the other is intuitively dismissed as less promising. An important property of this intuitive filter is that it can be skewed by biases.¹ We capture this effect as follows: for a biased decision maker, if one of the two generated alternatives conforms to the agent's intuitive preference whereas the other alternative does not, the latter is eliminated.² For an unbiased decision maker no such preference exists and thus, for comparison purposes, we assume that the choice between the two alternatives is random (see Figure 2.1).³ Finally, for reasons of simplicity, we modeled the deliberate evaluation stage as a simple comparison between the selected alternative and previous trials. This comparison leads to the acceptance of the selected alternative if, and only if, it has not been tried in the past. In the appendix, we

¹ The idea of organizational filters has surfaced in organizational studies before, although in a different context. The notion of dominant logics has been conceived of as an information filter that prevents organizations from acting on certain types of decisions (Bettis and Prahalad 1995).

² While biased spontaneous responses that we study closely align with an intuitive selection, we note that they could also be modeled as a sampling bias. This would require a different specification but would be functionally equivalent to our approach: mechanically, both result in a non-uniform distribution governing the probability that different alternatives reach the evaluation stage. That is, in our model, for each intuitive preference operating after a random generation of options (conditional on local search) there exists a corresponding non-uniform probability distribution that would produce qualitatively similar results. Mathematically, in our model, the probability that an option reaches the evaluation stage is $\frac{g}{2}(1 - b)$ if the option does not conform to the bias and $g\left(1 - \frac{b}{2}\right)$ otherwise, where g is the uniform probability that any option has of being generated under local search and b is the probability that the second generated option conforms to the bias.

³ Unbiased agents select randomly between the two generated alternatives because deliberate reasoning is not yet engaged at this stage of the search process and, thus, random selection is the only plausible unbiased comparison baseline. The concept of rationality is illusive in the context of our model because organizational search assumes, by default, that agents are boundedly rational. It is conceivable that a fully rational agent would not need to search and would instead instantaneously optimize or use backward induction to identify the optimal mix of biased and unbiased search to maximize performance. Therefore, in this paper we do not equate unbiasedness with rationality and focus on the effect of intuitive biases on the search patterns of boundedly rational agents.

generalize this simple decision algorithm to show that our results remain qualitatively similar when we include a formal evaluation stage in our model.

To see how such biases operate in the space of organizational configurations defined as N -dimensional binary strings of the NK model, consider the status quo of 000—corresponding to values of three possible choices. Further, consider an unjustified preference for configuration 100. If, for example, 001 and 100 represent the generated pair of options, then when bias is present, only the second alternative will pass to the stage of deliberate evaluation. When the generated pair contains no option that satisfies the intuitive preference or both options are preferred, then the choice is random. For example, from the pair of 001 and 010, both alternatives have equal chances of being evaluated.⁴ When, on the contrary, search is unbiased, there is no arbitrary discriminating criterion and regardless of the composition of the generated pair, any alternative has the same probability of passing to the stage of evaluation. In the process of search, unjustified intuitive preferences, therefore, narrow the set of consciously considered alternatives.

2.3.2. Model limitations

The described model of search explicitly factors in intuitive thinking and associated biases. As any other model, it entwines with the existing knowledge as well as necessarily restricts the real-world complexity (Levinthal 2011). Accordingly, before reporting our observations, we lay down the key bounds on how the specified model furthers our understanding of intuitive biases in the process of search.

⁴ Note that we allow for the possibility that the generated pair contains only one alternative.

In stylized terms, intuition originates closer to the automatic processes of perception and before the deliberate processes of reasoning (Kahneman 2003). As a result, intuitive biases are likely to tamper with the process of generation of alternatives, changing the likelihood that a given option will reach the stage of conscious evaluation. This being said, the non-uniformity of generation of alternatives is a commonly accepted tenet. In particular, there is a broad consensus that organizations tend to generate alternatives within the immediate vicinity of the current state. Although such a tendency can be seen as reflective of a bias (*i.e.* status quo bias, see Kahneman, Knetsch, Thaler 1991, Samuelson and Zeckhauser 1988), by itself it does not fully embrace the generally unjustified nature of biases nor informs on their consequences conditional on generating local alternatives.

Local search is driven by a multitude of factors at the individual and organizational levels, many of them dependent on deliberate or System 2 processes (Stuart and Podolny 1996). Due to following organizational routines, abiding by social norms, or extracting value from the firm's current capabilities organizations can gravitate toward local search (see for example Kogut and Zander 1992, Nelson and Winter 1982). More generally, incompleteness of knowledge, awareness of organizational rigidities and other non-intuitive factors may underlie generation of alternatives within the immediate vicinity of the current state. As a result, the propensity to generate local alternatives is characteristic of both biased and unbiased decision makers.

Similarly, other non-uniformities in the generation of alternatives currently present in the literature are largely grounded in the processes of System 2. In particular, forward-looking search, where certain choices are pre-specified *ex ante*, relies on deliberate reasoning and

careful assessment of the regions on the organizational landscape (see Gavetti and Levinthal 2000, Ghemawat and Levinthal 2008). In the same vein, choosing an optimal pattern of chunky search, where organizations sequentially address their attention to different parts of the landscape, arguably requires an informed and weighted reasoning (see Baumann and Siggelkow 2013).

Including all such non-intuitive constraints into the model is, of course, technically and conceptually impractical. However, since it is possible that local search partially originates in intuitive processes, we derive our observations for strictly local generation of alternatives (for an extension with non-local search see Appendix). Although arguably confining our claims (note that we assume that localness of search is orthogonal to intuitive biases), such an approach has clear analytical benefits. First, it allows us to illustrate a qualitatively distinct mechanism of generative recurrence. This mechanism arises on fully cognitive priors and holds for both local and distant generation of alternatives. Second and key from the perspective of behavioral strategy, our approach allows a general conceptualization of an equivalent bias (*i.e.* an exact opposite intuitive preference). The analytical import of this is best seen by comparison with local search. Consider that the converse of local is distant. Albeit it is conceivable that some individuals intuitively prefer distant alternatives, that they can as precisely implement them is not certain (see Rivkin 2000). In other words, there is an *ex post* qualitative difference between a bias toward local alternatives and that toward distant.⁵ Third and from a more technical perspective, the consequences of local search are well-

⁵ Note also that if non-intuitive factors lead to strictly local generation of alternatives, an intuitive preference toward distant options is meaningless.

established in the existing literature. Accordingly, we construct a simple model of search that assumes local generation of alternatives and specifically focuses on unjustified, affective preferences that arise in the functioning of System 1.

Although this model captures some essential aspects of the decision making under a joint influence of Systems 1 and 2, the full spectrum of cognitive complexities is naturally beyond the scope of the present paper. In particular, we abstract away from any interactions or simultaneity between intuitive and deliberate processes and focus on first principles of each system. Although complex higher-order interactions can occur between intuitive and deliberative mental processes (Gawronski and Bodenhausen 2006, Wilson et al 2000) for the purposes of simplicity we model the performance of intuitively biased vs. unbiased decision makers over time. So long as there is a non-zero main effect of intuitive biases on human decision making (Kahneman 2003), our analysis has important implications for organizational search. Including complex interactions into a single model would inevitably undermine tractability and cloud the mechanisms. The following section illustrates that even a simple parsimonious representation that we use suffices to unveil some potentially far-reaching implications of biased search for organizations.

2.4. Results


Since bias has no bearing on the mapping from configurations to performance, the comparison of behavioral differences in search is only meaningful for ante-limit time spans ($T < \infty$). By definition, the performance of biased and unbiased firms converge in the limit ($T \rightarrow \infty$), as sooner or later both behaviors lead to the discovery of (on average) a similar local optimum.

Accordingly, we focus on comparative dynamics rather than a static analysis of outcomes. In the following simulations, we generate random landscapes of $N = 10$ (our observations hold for greater dimensionality of the organizational space, *i.e.* $N = 15$, $N = 20$) and let the agents search for $T = 10^3$ periods. Our results are averaged over 10^7 runs.

2.4.1. The Dynamics of Biased Search

Interdependencies among decisions give rise to the ruggedness of the performance landscape, which in turn affects how fast organizations gain fit and progress to the global optimum (Levinthal 1997, Rivkin and Siggelkow 2007). Accordingly, we start by analyzing how the presence of unjustified intuitive preferences affects search in environments of varied complexity. An essential aspect of this analysis is that it considers temporal consequences of decision makers' intuitive responses and associated biases. Figure 2.2 shows the dynamics of relative performance of the biased organization. We measure relative performance at each point in time as the difference between the average fit of the biased firm minus that of the unbiased firm.

Figure 2.2. Performance of Biased Search over Time



Notes. The curves show the dynamics of relative performance of biased search. Positive (negative) values indicate that at the given moment in time, biased search has an advantage (disadvantage) over unbiased search in terms of the observed fit. The value of zero means that biased and unbiased organizations tend to have exactly the same level of fit at that moment in time. The different curves illustrate how the relative dynamics of biased search changes with complexity. For example, at $t = 20$ in the low- K environment ($K = 0$) biased search tends to result in a level of fit that is lower by approximately $8 \cdot 10^{-3}$ units, whereas in the higher- K environment ($K = 5$) it leads to an average level of fit that is higher than that of an unbiased search by approximately $3 \cdot 10^{-3}$ units.

Three key observations are immediate: first, the effect of bias is time-variant; second, being unbiased is not always better; third, the payoff to biased search depends on the degree of interdependencies among decisions. Specifically, when the complexity of the environment is relatively high (as is arguably the case with most modern organizations), it pays off to be


biased in the short-term and hurts in the long-term. In contrast, when complexity of the environment is low, this pattern does not hold. Contrary to the characterizations of intuitive biases as leading to errors and suboptimal outcomes (*e.g.*, Gilbert 2011), we find that unjustified preferences may actually prove advantageous at certain stages of the process of search. Importantly, this indicates that the effect of spontaneous, intuitive thinking is not necessarily constant (and neither strictly detrimental nor strictly beneficial) but rather accumulates and varies in a systematic way. In the following subsection, we characterize the underlying processes that give rise to the reported pattern of performance.

2.4.2. Mechanism

The dynamics identified above result from the interplay between two forces: (a) accumulation of knowledge and (b) the implicit cost of experimentation. The two panels in Figure 2.3 isolate and illustrate separately each of these two mechanisms. The y-axis in Figure 2.3 measures the relative performance of biased search—that is, the observed differential between biased and unbiased search. The left panel in Figure 2.3 shows the main disadvantage of biases, as reflected in the negative relative value of accumulated knowledge of the biased search. The downside of an intuitive bias is that it constrains search and, thereby, hinders the accumulation of knowledge about regions of the landscape that lie outside the scope of the bias. This creates a pattern in which organizations focus disproportionately more on a narrower subset of options of the landscape. Intuitively focusing on an arbitrary subset of alternatives increases the diminishing returns to search by raising the probability of redundancies. The organization learns less because it stumbles upon the same configurations more often. In contrast, unbiased search allows for broader exploration and less redundancies in learning, which leads to greater

awareness about the whole landscape—and a faster discovery of local optima. The left panel in Figure 2.3 shows that the value of the best known state is strictly higher in the case of unbiased search.

Figure 2.3. Mechanism



Notes. The curves show relative values, where zero means that there is no difference between biased and unbiased search in moderately complex environment ($K = 5$). The left panel shows the relative performance of biased search in terms of accumulated knowledge. We measure accumulated knowledge as the fit of the best known state. For example, a point with coordinates approximately $(20, -9 \cdot 10^{-3})$ means that the best known alternative that the biased organization has tried by the time $t = 20$ has fit on average $-9 \cdot 10^{-3}$ units lower than that known by the unbiased organization. The right panel shows the relative performance in terms of value from stability. We measure value from stability as the inverse of the implicit cost of experimentation (*i.e.* probability of trying a new alternative multiplied by the expected reduction in fit). This means, for example, that at the time $t = 20$ the implicit cost of trying new alternatives is lower for the biased search by approximately $11 \cdot 10^{-3}$ units.

Conversely, the right panel in Figure 2.3 illustrates the second mechanism and a source of the main advantage of intuitive biases in organizational adaptation, as evidenced by the positive relative value from stability. The upside of biased search is that it may provide

stability by regulating excessive experimentation. Since unjustified preferences increase trials in the same subset of configurations, organizations draw more frequently on options that were already experienced in previous periods. This regularity of a biased decision maker—that we refer to as generative recurrence—is biases’ primary advantage.⁶ Past experience with a given configuration helps resolve uncertainty about the alternative: organizations know that the option they are facing is inferior to their status quo and therefore need not try it. Biased search therefore skews organizational learning in a way that allows organizations to recognize and avoid inferior alternatives and, thereby, reduces the implicit cost of experimentation. A real world analog of this mechanism might include situations in which the new alternative is not exactly identical to previously tested configurations, but is sufficiently similar to warrant meaningful comparisons and support accurate inferences about its intrinsic value.

Following the work on the use of analogy in organizational decision making (see Gavetti, Levinthal, Rivkin 2005), we assume that conscious effort facilitates managers’ ability to accurately recollect experiences from the past and thereby avoid making similar mistakes repeatedly. Although memory operates both intuitively and deliberately (Jacoby 1991, Roediger 1990, Schacter, Chiu, Ochsner 1993, Smith and DeCoster 2000), there is evidence that conscious recall adds value above-and-beyond implicit or intuitive recall alone. For instance, empirical works indicate that conscious associations exhibit more rapid updating in light of new experiences than automatic associations do (Gregg, Seibt, Banaji 2006, see Wilson *et al.* 2000, for a review). Conscious, deliberate reasoning improves recall and

⁶ Recall that, from the standpoint of conscious evaluation, intuitive selection is functionally equivalent to a sampling bias in generation of alternatives (footnote 4).

therefore adds to the agent's ability to effectively use experiences from the past (see Smith and DeCoster 2000). Arguments in favor of this can be further invoked at the organizational level. Although both organization and its members are liable to forgetfulness, organizational memory—for example through a process of knowledge codification—is likely to further add to the memory of its decision makers (see Levitt and March 1988, Walsh 1995, Walsh and Ungson 1991). A direct corollary is that in organizations, memory at the stage of generation of alternatives tends to be inferior to memory at the evaluation stage, when deliberate recall regarding past experiences with the considered alternatives is more likely to be engaged. Although, for simplicity of the baseline model, we assume no intuitive memory and perfect deliberate memory, in the Appendix we show that our main qualitative observations hold as long as (a) organizational memory is at least moderate, and (b) there is value-added from conscious recall above-and-beyond any intuitive recollections about the alternatives that have been generated. Therefore, the main mechanism behind our results holds for search processes where generation is at least partially repetitive and environmental uncertainty warrants online experimentation. A lean startup—or a venture development approach that is based on iterative product experimentation—is a real-world modern example of an adaptive process in which intuitive preferences are likely to result in the described performance dynamics. In such contexts, where highly experiential organizations operate in turbulent environments, our findings suggest that it should be possible to improve organizational performance by managing the process of biased search.

2.4.3. Re-Biased and De-Biased Search

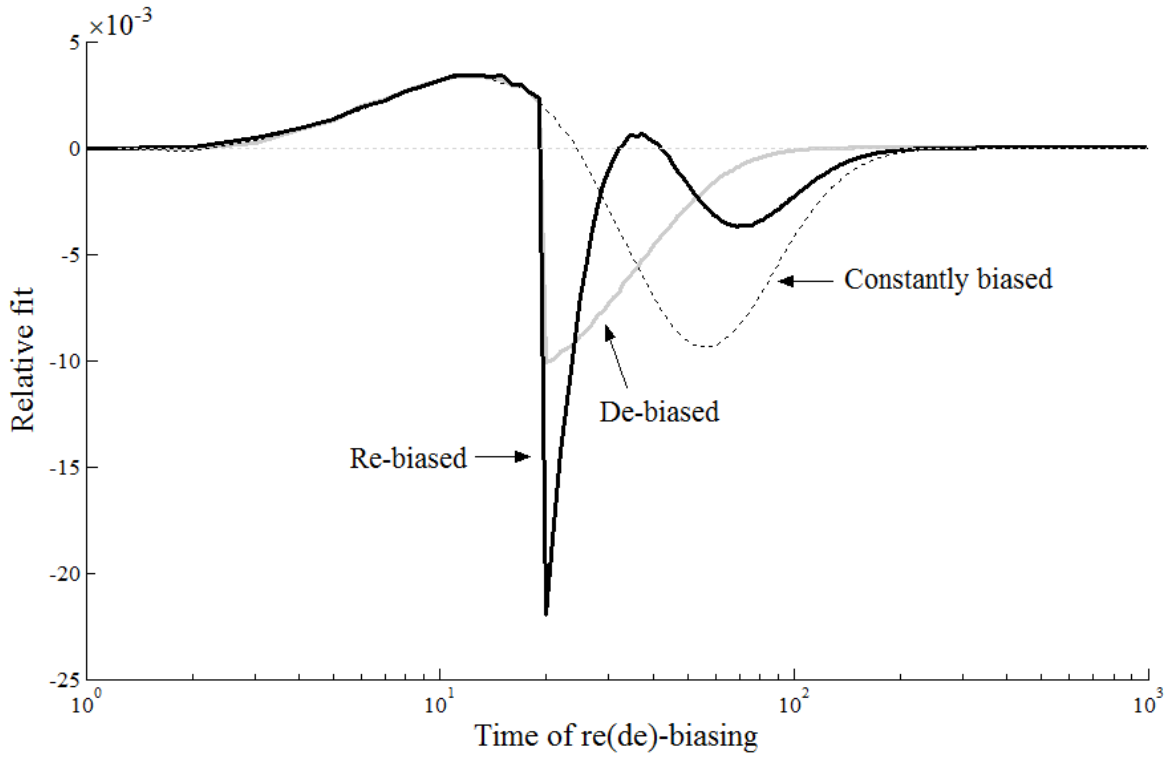
In the computational experiments reported in previous sections, we treated intuitive bias as a constant behavioral tendency that is either present or absent over the entire period of organizational adaptation. However, the magnitude, direction, and way in which organizations are biased is not necessarily a stable attribute as organizations can experience drastic, often exogenous, changes in their biases. Consider an example of core decisions in football. From the perspective of the coach, finding the right composition of the team is a typical recombination problem that requires trial and error. While searching for an efficient solution to this problem, the coach may well be biased. That is, he or she may intuitively reject those alternatives that do not favor players with whom the coach has friendly relationships. Swapping the positions of two favored players is considered as an option, whereas putting one of them on the bench may not be. Should this coach retire, his or her successor is likely to form a different pattern of liking and disliking towards the players, and thereby change the way organizational search is biased. A similar logic applies to a more general problem of finding an efficient correspondence between organizational members and positions. A change of the key decision maker, therefore, represents a basic instrument that can lead to the inversion of organizational bias, or re-biasing.

Scholars in psychology as well as industry practitioners have also discussed an array of techniques that can abate the effect of biases, or de-bias, decision makers (see for example, Kahneman *et al.* 2011). Similarly, the literature in management has shown that organizations have structural means to manipulate and attempt to reduce bias in organizational decision making (see Christensen and Knudsen 2010), and case studies highlight instances in which

companies have changed management teams and completely reversed their previous management practice orientations (see for example, Maddux, Williams, Swaab, Betania 2014).

Given that organizations often experience either exogenous or intended transformation in their biases, here we analyze the consequences of such changes. In the following analyses, we are agnostic as to the exact levers that organizations use to manipulate biases—whether it involves replacement of the key decision makers or implementing other management practices—and focus solely on the outcomes of such strategic interventions. Our starting condition is that of the biased firm and its performance dynamics. Subsequently, we examine the temporal implications of re-biasing (changing the bias to its opposite), and de-biasing (eliminating the bias entirely). Specifically, we operationalize re-biasing as adopting the exact opposite of the initial intuitive preference. For example, if *ex ante* an organization is biased towards alternatives that have decision 1 in state 1, *ex post* it will be biased towards alternatives that have decision 1 in state 0. De-biasing means that *ex post*, the agent is unbiased and no longer discriminates among alternatives. We compare the effects of adopting these behavioral strategies (re-biasing, de-biasing) to the performance of continuously unbiased search. We focus on managing biases in environments with some complexity, as these are more likely to reflect the environments faced by real firms. [continued on the next page]

Figure 2.4. Dynamics of Re-Biased and De-Biased Search



Notes. The curves show relative performance of (de)re-biased search (cf. Figure 2.1) in a moderately complex environment ($K = 5$), such that the value of zero means that the difference between unbiased and (de)re-biased search is nil. We assume that initially biased search is exogenously de-biased or re-biased at $t = 19$, *i.e.* close to the exhaustion of the initial advantage of biased search. That is, starting from period $t = 19$ organizations exhibit either no preference (de-bias) or the opposite preference (re-bias) for certain options.

The dynamics of biased search in complex environments may suggest that eliminating the bias when it no longer creates value—where there are no more rents to generative recurrence—should prevent organizations from experiencing the long-term costs of unjustified, intuitive preferences. Although in the long run this is indeed the case, we find that immediately after de-biasing, organizational performance suffers a sharp decline relative to that of continuously unbiased search (see Figure 2.4). Since the set of alternatives that used to

be intuitively discarded as less promising remains comparatively unknown by the biased organization, elimination of the unjustified preference results in an increased likelihood of trying new options. By equilibrating their focus between the preferred and the underexplored subsets of alternatives, de-biased organizations are more likely to sample novel, untried options. The corresponding increase in experimentation raises the implicit cost of search (*i.e.* the probability of trying a new option multiplied by the expected reduction in performance associated with online experimentation). However, since a large part of the available alternatives is already encoded in organizational memory, such an increase in experimentation does not provide a commensurate improvement in the value of the best-known state. This is consequent to the fact that the maximum is a concave function of the number of trials. The resulting imbalance between increased experimentation and the reduced increment in the value of the best-known state produces an immediate drop in relative performance. Gradually, as organizations progress to superior configurations, this initial shock of de-biasing fades out and the performance of the de-biased organization converges to that of the continuously unbiased firm.

In contrast to de-biasing (eliminating the bias), re-biasing (adopting the opposite bias) actually leads to a second-order advantage. After an initial drop in performance, re-biasing produces a temporary, but significant improvement in performance. The mechanism behind this is essentially the same as the one explaining the short-term dominance of the biased search (see section 5.2). Because of the heightened focus on the underexplored part of the landscape—after a more pronounced initial shock—re-biased search allows for a speeded accumulation of knowledge, which soon approaches that of the continuously unbiased search. As this occurs, the implicit relative cost of re-biased search declines and the organization takes

advantage of the opposite bias. We call this a second-order advantage because it builds on the asymmetries in knowledge accumulation that were generated in the course of exercising the initial bias.


2.4.4. The Optimal Timing of Behavioral Treatments

Thus far, our analysis has shown that, while in terms of accumulated knowledge unbiased search is strictly dominant, biased search can create more value at certain time points. This is because it provides stability and lowers the burden from excessive experimentation. However, comparing performance differentials at specific points in time does not reflect the total relative value created or destroyed by search in the periods leading up to and following the specific moment of a behavioral treatment. Therefore, in this section we analyze the optimality of stylized behavioral strategies over an interval of organizational adaptation, that is, we compute the total cumulative performance of these strategies, relative to unbiased search. Specifically, we define three generic forms of behavior—continuously unbiased search, biased search with a re-biased treatment at a certain moment in time t , and biased search with a de-biased treatment at a certain moment in time t —and, then, we measure these interventions' relative performance over a given time span.

The definition of the given time span of analysis is particularly important. By limiting this time period during which organizations search over an unchanged landscape, we follow previous work that examines the effects of turbulence or exogenous shocks on the organizational environment (*e.g.* Eisenhardt 1989, Posen and Levinthal 2012, Siggelkow and Rivkin 2005). Siggelkow and Rivkin suggest that "an environment is turbulent, dynamic, *etc.*, if the mapping from firm actions to performance outcomes changes frequently, profoundly,

and in ways that are difficult to predict" (2005, p. 103). The question that we address is what form of behavioral strategy allows greater value creation for a given level of turbulence or a given interval of time when the organizational landscape remains unchanged. Such a period of environmental stability can be thought of as an interim between two consecutive Schumpeterian shocks. During this time, organizations adapt to the given set of technological, strategic, or demand conditions. After the next shock that redefines the mapping from organizational configurations to performance, the knowledge that organizations accumulated about the previous environment becomes obsolete and the search starts anew. The critical determinant of such punctuated process, therefore, is the amount of time that organizations have on a given landscape. Intensifying environmental turbulence—that is, shortening the time on a stable landscape—should render more attractive strategies that allow for a faster increase in fit (Siggelkow and Rivkin 2005). Conversely, the lower the turbulence, the more critical an organization's ability to discover superior configurations becomes.

Figure 2.5. Optimal Time of Re-Biasing



Notes. Curves plot the relative cumulative performance of a given search behavior in moderately complex environments ($K = 5$), such that the value of zero indicates that the average accumulated levels of fit of the unbiased and re(de)-biased searches are equal. For example, a point on the solid black line (left panel) that coordinates approximately (6, 0.03) means that re-biasing in a highly turbulent environment (where organizations have $T = 20$ periods to search a stable landscape) at $t_{re(de)-bias} = 6$ leads to the overall gain of approximately 0.03 units over the entire period of adaptation to the given landscape.

Figure 2.5 shows how re-biasing (adopting the opposite bias) or de-biasing (eliminating the bias) at a given point in time $t_{re(de)-bias}$ affect the cumulative value created relative to continuously unbiased search. It illustrates that, counter to what one may expect, changing rather than eliminating bias can be a superior form of behavioral strategy. In turbulent and moderately stable environments (e.g. $T = 20$ or $T = 50$)—arguably the case with most modern industries—organizations can actually benefit from reversing their decision makers' biases. Observe that as turbulence declines ($T = 20 \rightarrow T = 50$) organizations are better off delaying the actions that can re-bias their decision making. Specifically, in a turbulent

environment, where organizations have $T = 20$ periods to search a stable landscape, the optimal timing of re-biasing is $t_{\text{re(de)-bias}} = 2$. In a more stable environment ($T = 50$), however, optimal timing of re-biasing increases to $t_{\text{re(de)-bias}} = 8$. At the extremes (not shown here), in very turbulent (*e.g.* $T = 10$) or very stable (*e.g.* $T = 100$) contexts it is optimal to be constantly biased or constantly unbiased respectively (cf. Figure 2.1).

Strikingly, although de-biasing occasionally outperforms re-biasing, it is never the dominant form of behavior. Specifically, it is always dominated either by continuously unbiased search or by re-biased search. In contrast, re-biasing emerges as a superior behavioral strategy in less turbulent environments. Contrary to characterizations of biases as undesirable, our findings indicate that completely eliminating bias in the process of organizational adaptation is not only an inferior behavioral strategy, but can also be dominated by strategic changes in the direction of the bias.

2.5. Conclusion and Discussion

Our work explores the process and consequences of biased search. We show that biases may have time-variant effects, or lifecycles, that emerge in the process of organizational adaptation. Specifically, we find that bias, operationalized as an unjustified intuitive preference, may improve short-term performance but undermine long-term levels of fit in complex organizational environments. This behavioral pattern is an outcome of a mechanism we refer to as generative recurrence: unjustified preferences curb excessive experimentation but at the expense of knowledge accumulation. We further analyze behavioral treatments that can exploit the benefits from these dynamic effects and find that eliminating an existing bias

(de-biasing) may only be beneficial for organizations operating in very stable environments. In turbulent environments characterized by high uncertainty (arguably two properties of the environments faced by most modern organizations), a more effective strategy can in fact consist of re-biasing (*i.e.*, strategically changing the bias to its opposite). This provides a novel perspective on managing biases as previous work in experimental settings has focused almost exclusively on de-biasing (*e.g.*, Wilson and Brekke 1994, Wilson *et al.* 2002). The present results identify re-biasing as an often unconsidered but potentially highly effective strategy for firms. More broadly, our results speak to the emerging field of behavioral strategy in that we show that cognitive biases can be strategically exploited as a source of value creation (see Gavetti 2012, Felin and Foss 2009, Levinthal 2011, Powell, Lovallo, and Fox 2011).

The value of biases in the form of unjustified intuitive preferences stems from stability and decreased cost of experimentation. The flipside of a biased search, however, is that it lowers the pace of knowledge accumulation. This double-edged effect points to the important role biases play in balancing the critical organizational tradeoff between exploration and exploitation (see March 1991). Contrary to what may be expected, having an unjustified intuitive preference suppresses exploration, but allows for a more efficient exploitation (*cf.* March 2006). Since the seminal work of March (1991), management scholars have scrutinized the organizational levers that balance the tension between exploration and exploitation (Gupta, Smith, Shalley 2006). An important role in resolving this tradeoff has been attributed to organizational architecture. Designing efficient organizations requires that temporal aspects of the interplay between search and stability be taken into account (Rivkin and Siggelkow 2003, Siggelkow and Levinthal 2005). Some organizational architectures provide greater emphasis on exploiting the status quo, while others facilitate experimentation. Our analysis points to the critical importance of taking into account the psychological dispositions of organizational

decision makers when developing efficient organizational designs. More broadly, we concur with the recent work emphasizing the role of individual psychology in organizational processes (see Reitzig and Sorenson 2013, Larkin, Pierce, Gino 2012, Powell *et al.* 2011).

From a broader theoretical standpoint, the present research begins to fill the gap between fundamental advances in psychology of human decision making and research and theory on organizational adaptation. In particular, there is an important distinction between the Simonian perspective on decision making—which focuses on the computational limits of the mind and largely underlies the current view of organizational adaptation—and that of Kahneman and Tversky, which further embraces the notion of intuitive bias (see Fiori 2011). Human reasoning, unlike that of artificial intelligence, is not a monolithic process that operates on a single algorithm of logic. Rather, it functions as a function of both intuitive and deliberative cognition (Kahneman and Frederick 2002, Sloman 1996, Stanovich and West 2000). As a consequence, human rationality is different from that of artificial intelligence—and a key element is the presence of automatic and intuitive responses that may give rise to systematic and predictable biases. We show that complementing this element of decision making with the fundamentals of Simonian bounded rationality and analyzing the consequences of biased judgment for organizational search can help us better understand the available levers to manage our behavior to create value.

Our point of comparison is unbiased search as we effectively compare having an intuitive preference to not having any preference. Although from the psychological perspective not having any intuitive preference may appear unrealistic, we choose such an idealized form of search to focus on differences in kind rather than degree. Our observations also apply to a comparison between agents who are more prone to exhibit intuitive preferences

against those who are less prone to have such preferences. Importantly, the preferences we consider are fully random, *i.e.* they embrace those biases that have no bearing on the organization's location on the performance landscape. Therefore, our analysis shows that even position-independent non-uniformities in generation (cf. local search) can be useful for organizational adaptation.

Our findings are consistent with prior research and theory suggesting that biases are not necessarily a liability (*e.g.* Denrell and March 2001, Galasso and Simcoe 2011, Miller, Zhao, Calantone 2006, Winter *et al.* 2007, Xu, Liu, Liu 2014). For instance, prior work in psychology points to an important ecological benefit of some biases. Specifically, biased decision-making rules can lead to fast and frugal choices and ultimately more adaptive behavior in ecological contexts (*e.g.* Bernardo and Welch 2001, Johnson and Fowler 2011, Gigerenzer and Goldstein 1996). Although this work jointly suggests a complex—and not necessarily negative—role of biases in human and organizational decision making, they underemphasize the temporal aspects of relying on unjustified intuitive preferences and the potential of managing them over time.

We extend this line of thought by developing a model of organizational adaptation that accounts for spontaneous managerial choices and focuses on the effects of biased search over time. Among our core findings is that biases often improve organizational learning across the lifespan of the organization. While most experimental laboratory research documents the consequences of biases in the context of a single, well-defined choice, we tract temporal effects of biases in the context of multiple (and highly interdependent) choices. Counter to some of the more negative characterizations of biases in the psychological literature, we show

that they may lead to superior outcomes at key stages in the process of organizational adaptation.

Because of the reported temporal effects, biases contain strategic tradeoffs. A particularly important tension that we identify is that between reduced experimentation and accumulated knowledge. Importantly, we show that it may be reasonable to manage such tradeoffs not only *ex ante* (e.g. by appointing a top executive with or without a particular bias), but also during the process of organizational adaptation (e.g. by replacing the top executive with another who has the opposite bias). Specifically, because of the lifecycle pattern described above, in a complex environment it may be reasonable to re-bias, or change the bias to its opposite. Although re-biasing results in an immediate short-term decline in fit, it can substantially improve future organizational performance over time in complex and turbulent environments. Overall, our results indicate that research on elements of organizational design that can be used to manage biased decision making over time may uncover further means of value creation.

**Essay II. A Bias to an End: Misconceived Complexity
and Dual Aspirations**

3.1. Introduction

Business strategy derives from decision-makers' understanding of the complex reality, an understanding that is inherently imperfect (Levinthal 2011). Due to cognitive constraints, decision-makers often err in assigning values to choices or seeing what the choices are, and, more generally, in understanding how complex the reality is (Martignoni, Menon, and Siggelkow 2016). Such a misconception of business complexity affects the dynamics of boundedly rational search and thereby alters the likelihood that organizations will meet their aspirations. We focus on this relationship and study how decision-makers should conceive of complex reality to meet their aspirations. By doing so, we derive an explicit connection between mental models, micro-foundations of problem formulation and the attainment of organizational aspirations.

The present-day business reality is a complex plexus of interdependent factors, ill-specified alternatives, and ambiguous tradeoffs that prevent global optimization and hamper organizational adaptation (Levinthal 1997, Lyles and Mitroff 1980, Rivkin 2000, Sargut and McGrath 2011). In such an environment, the limits to decision-makers' understanding become ever more evident. Increasingly, organizations fail because they cannot fully comprehend the complex problems they face (*e.g.* Tripsas and Gavetti, 2000). At the same time, the competition squeezes every bit of decision-makers' ability to process and interpret information, leaving little, if any, room for a better understanding (see Gary and Wood 2011, Helfat and Peteraf 2015). As a result, organizations compete at the frontiers of human cognitive abilities, striving to reach their objectives in an environment where decision-makers'

mental models of the complex reality are anything but perfect (Csaszar and Levinthal 2016, Gavetti 2012, Gavetti and Levinthal 2000).

Accordingly, our premise here is twofold. First, we assume that agents form an idea or a mental representation of the complex structure of interdependencies, which is *stochastically imperfect*. We explicitly model this process and regard it as a basic primitive of problem formulation. From this perspective, mental representations are (in part) an outcome of deliberate problem formulation, or problem framing (see Levinthal 2011, cf. Baer, Dirks, Nickerson 2013).⁷ Second, we consider that individuals and organizations differ, and can intentionally differ, in how they misconceive the objective complexity, as well as in their aspirations.

Given an unknown structure of the choice setting, organizations routinely follow behavioral strategies that increase their chances of meeting aspirations rather than maximize performance (Camerer, Babcock, Loewenstein, Thaler 1997, March and Shapira 1992). Although broadly recognizing this in principle, the existing analytical work has mainly focused on performance as the primary objective function. We depart from this approach and treat meeting aspirations (*i.e.* the probability of attaining a certain outcome) as a measure of success. We conceptualize aspirations as compounds of two basic dimensions: time and performance. Such a duality of aspirations (albeit not explicit in the existing work on

⁷ The other part is intuitive, uncontrollable, or implicit. That is, intuitive, associative linkages beyond one's conscious understanding guide decision-makers' choices as well. In the interest of clarity, our narration here centers on the notion of problem formulation, and more deliberate processes. The empirical supplement, however, illustrates how our analysis applies to cognitive constructs that arguably also encapsulate less conscious antecedents of one's choice.

organizational adaptation) closely aligns with human behavior in general and goal setting in particular (see Locke and Latham 2002).⁸

Using attainment of such dual aspirations as an objective function for boundedly rational agents, we show that, conditional on errors in specifying individual interdependencies, perceiving complexity at the correct overall level can be suboptimal. In other words, organizations should strive, *i.e.* have bias, for either complexity or simplicity depending on their aspirations. To illustrate this, we derive an explicit mapping from such biases in problem formulation to the likelihood of attaining different aspirations and show that overcomplicating or oversimplifying can allow organizations to pursue their objectives more efficiently. For the behavioral strategy, this means that organizations may want to manage their approaches to formulating strategic problems. In particular, it may be optimal to select decision-makers or define procedures in ways that promote rather than eliminate biases in understanding business complexity. From a broader theoretical perspective, our results indicate that the value or detriment of a bias (in problem formulation or problem understanding) is a function of the decision-maker's aspirations. In other words, systematic rather than random errors in problem formulation can be instrumental for organizational success.

In the supplement, we bring our model to data and empirically assess a corollary of our predictions. In particular, we illustrate that time and clockspeed—*i.e.* the pace of decision-making and organizational change—likely moderate the relationship between perceiving

⁸ Note that we follow previous work in considering aspirations, objectives, or goals as essentially synonymous (see for example Hu, Blettner, and Bettis 2011, Lant 1992, Mezias, Chen, Murphy 2002, Shinkle 2012, Washburn and Bromiley 2012).

complexity and organizational performance. Our work, therefore, points to the importance of managing the unavoidable misconceptions of business complexity over time. The following sections characterize the priors of our analysis in detail as well as specify the technical apparatus we use to arrive at conclusions.

3.2. Search and aspiration levels

A first principle of bounded rationality and a fundamental of human behavior in sufficiently complex problems is search (Cyert and March 1963, Simon 1955, 1990). Individuals—and by extension organizations—cannot comprehend the entirety of available information, and therefore must rely on consecutive generation of alternatives and satisficing (Knudsen and Levinthal 2007, Simon 1955, 1956). In particular, organizations tend to consider options near the immediate vicinity of their current state and try those supposed to improve performance (Cyert and March 1963).

A natural consequent of such a stepwise process is that organizations spend a non-trivial time at suboptimal configurations, or strategies (see Levinthal 1997, Rivkin 2000). More importantly, since the entire space of alternatives is unknowable, the extent to which a given strategy is suboptimal cannot be reliably assessed; that is, decision-makers cannot say how well- or underperforming their strategy is on an absolute scale. Instead, organizations use cues from past experience and social comparison to construct their aspiration levels (Cyert and March 1963, Washburn and Bromiley 2012).

Aspiration levels—"levels of outcomes that will satisfy the individual or organization" (Washburn and Bromiley 2012: 896)—serve boundedly rational decision-makers to distinguish success from failure (Lant and Shapira 2008, March and Simon 1958). The origin of these levels being a separate subfield in organizational and behavioral sciences, their presence in choice settings (albeit not always explicit) of both individuals and organizations is nearly a universal. Individuals use conscious goal-setting as a means to motivation (*e.g.* Locke 1996, Locke and Latham 2002, Locke, Shaw, Saari, Latham 1981). Similarly, aspiration levels guide organizations in choosing strategies, accepting risks and implementing the chosen courses of action (Bromiley 1991, Greve 1998, Lant and Mezias 1992, March and Shapira 1992).

Although recognizing the multidimensional nature of organizational objectives, the existing work often views aspirations on a given criterion or dimension as integrating only performance on that criterion (*e.g.* profit; for organizational search with multiple criteria see Ethiraj and Levinthal 2009). Reflective of Simon's search—where an aspiration level "defines a satisfactory alternative" (1955: 111)—such a view of aspirations is a useful and plausible conception for many settings. In particular, it allows a clear-cut proxy of the organizational choice process. Organizations accept the first alternative that satisfices along all essential criteria, dimensions, or all fundamental constraints (Simon 1955). However parsimonious, this view yet markedly underplays the intricacies of organizational planning, which often involves prioritizing time and performance (see for example Frederick, Loewenstein, O'Donoghue 2002, Pacheco-de-Almeida and Zemsky 2007). In other words, from the current perspective on organizational search, the idea that objectives reflect organizational preferences with respect to the time-performance tradeoff remains largely underexplored.

Consider, for example, the following objectives from the 2015 Annual Reports for Daimler, Microsoft, and Adidas Group, respectively:

Daimler Trucks occupies a very good position in the competitive field. On this basis, we continue to target sales of 700,000 units in the year 2020.

With all of this new innovation hitting the market, we are on track to reach two important goals for fiscal year 2018: a commercial cloud revenue annual run rate of \$20 billion and 1 billion Windows 10 active devices per month.

We have set ourselves the goal to become the first fast sports company by 2020. This means that it is our ambition to increase the share of so-called 'speed-enabled products' to 50% of our net sales by 2020 from 15% of net sales in 2015.

These examples suggest that organizational aspirations are often set in a multi-dimensional space of various objectives (such as sales, growth, market share, profitability *etc.*) and a well-defined time horizon by which these metrics are to reach certain levels. Importantly, to the extent that organizations can actively change and observe performance feedback in the meantime (*e.g.* annual or daily results), such aspirations are qualitatively distinct from satisfactory alternatives (cf. Gaba and Bhattacharya 2012, Schneider 1992). Observe, for example, that a target of 700K trucks in five years, with current sales of about 502K and global market size of around 30M units, is an ambitious (albeit reasonable) goal, which may require complex real-time adjustment of positioning, technology or general strategy. This means that organizations will make choices and accept alternatives whose performance feedback will be observable before the time horizon in their aspirations. In the example above, Daimler Trucks may need to change its positioning, enter new markets, renegotiate its agreements with suppliers, or consider other choices that may aid or impede its chances of

reaching the target of 700K trucks in the year 2020. From this perspective, organizations should incrementally adapt to reach their objectives, *i.e.* they should search toward aspirations.⁹

3.3. Problem formulation and misspecified Interdependencies

As organizations search and adapt to the complex environment, they make choices that in and of themselves tend to be imperfect (Knudsen and Levinthal 2007). In particular, an [intendedly] "rational choice process requires an *a priori* act of problem framing and representation" (Levinthal 2011: 1517). That is, to make a decision, one needs an understanding of whether a given choice is instrumental for attaining one's objectives. While existing models of recombination search tend to depict a picture wherein organizations act on the direct (albeit possibly noisy) signals from the objective reality, a broad literature posits that managerial understanding is in fact imperfect and organizations act on their decision-makers' mental representations of reality (*e.g.* Anderson and Paine 1975, Porac, Thomas 1990, Porac, Thomas, Baden-Fuller 1989, Walsh 1995).

Research in psychology and management provides ample evidence that decision-makers' mental models are heterogeneous and their attributes consequential for organizational strategies and performance (*e.g.* Barr, Stimpert, Huff 1992, Gavetti 2005, Hodgkinson, Bown, Maule, Glaister, and Pearman 1999, Nadkarni and Narayanan 2007). The way individuals

⁹ Note that scholars have considered search above aspiration levels (see for example, March and Shapira 1992). While we fully concur that continuous search is inherent in all adaptive organisms, including firms, we focus on how boundedly rational organizations pursue their objectives.

understand relationships among different factors determines their decision rules, alters their awareness of plausible actions and ultimately affects strategic choices (Gary and Wood 2011). Yet, irrespective of one's decision processes, it is the true state of the world that determines choice outcomes. Indeed, the work on sequential sampling (n-armed bandit models) centers on precisely the idea that organizations act on their mental models as they learn about the true state of the world (*e.g.* Denrell and March 2001, Posen and Levinthal 2012, for the distinction between organizational learning and adaptation see Fiol and Lyles 1985). In other words, while organizations experience and adapt to the objective landscape, they act on their subjective understanding thereof.

Recent work in behavioral strategy and organizational adaptation has begun bringing to light this element of decision making, focusing on the search consequences of imperfections in mental models (see Csaszar and Levinthal 2016, Gavetti and Levinthal 2000, Martignoni *et al.* 2016). In particular, and related to our work, Martignoni *et al.* show that errors of over- and underspecifying interdependencies are not equivalent and may be, under some conditions, instrumental (2016). While indicating a non-trivial link between performance and deficiencies in understanding, this work contrasts errors against a fully accurate representation. Mental models of boundedly rational agents, however, are stochastically inaccurate. In essence, when it comes to the complex objective reality, we all err. We also err differently, and can differ in our propensity to commit certain types of errors. Accordingly, to determine the consequences of misconceiving reality more fully, it is essential to consider biases—rather than additional

errors—in inherently imperfect mental models, as well as the process of their construction, or problem formulation.¹⁰

Strategic problem formulation is a process whereby decision-makers define an organization's objective function and its constraints, or simply define "*what* problem is solved" (Baer *et al.* 2013: 198). This includes forming an understanding of relationships among the existing policy choices and their effects on organizational performance. In other words, when formulating the problem—explicitly or implicitly—decision-makers construct a mental image of the objective problem, an image or set of rules that guide strategic decision-making (see Gary and Wood 2011, Levinthal 2011). More broadly, "problem formulation [...] is embedded in the firm's norms for organizational decision-making" (Lyles and Thomas 1988: p. 139).

Despite the long-recognized importance of strategic problem formulation (see for example, Shrivastava and Grant 1985, Volkema 1983), its exact search implications remain largely underexplored. In part, this stems from the fact that research on organizations in general—and the work on organizational adaptation in particular—has remained focused on the organizational ability to *resolve* complex problems, rather than the ability to identify or *formulate* such problems (Baer *et al.* 2013).¹¹

¹⁰ Note that we do not equate problem formulation and construction of mental models, as the two can have profound differences in the underlying cognitive processes (*i.e.* deliberate, associative thinking, learning and socialization, *etc.*). Rather we observe that from the perspective of problem solving, the two processes are functionally equivalent.

¹¹ Problem formulation represents a broad concept that also embraces identification of a problem, or search for a problem (see Lyles and Mitroff 1980, Nickerson *et al.* 2007). While we recognize that our work does not fully

At the same time, understanding the problem is arguably as important (if not more) for the outcome as solving it (Mintzberg, Raisinghani, Theoret 1976, Nickerson, Silverman, Zenger 2007). Even if the solution to a formulated problem is perfect, any mismatch between the perceived problem and the reality may render the outcome suboptimal. Consequent to seeing a different strategic problem, organizations can stall at spurious or imagined maxima. The reason is that unlike white noise in evaluation, imperfections in mental models lead to systematic errors in assessing the potential of different configurations (see Csaszar and Levinthal 2016). As a result, superior solutions may be inaccessible because they do not appear as such, whereas inferior solutions can seem acceptable, thereby causing undue experimentation.

Systematic errors in the understanding and evaluation of alternatives can result from such factors as overconfidence or obsession (Winter, Cattani, Dorsch 2007), incentives design (Ethiraj and Levinthal 2009, Rivkin and Siggelkow 2002, 2003), or internal politics (Cyert and March 1963). From the perspective of strategy formulation, a central cause of systematic errors in strategic choices is misperception of interdependencies. Indeed, Ghemawat and Levinthal argue that "choice settings are strategic to the extent that they entail cross-sectional or intertemporal linkages" (2008: 1638). The reason is that interdependencies essentially underlie strategic tradeoffs. Failure to specify the linkages among organizational choices correctly results in decision-makers acting on spurious tradeoffs or omitting the essential ones.

cover the entirety of problem formulation, we use this term to establish a broader conceptual link between problem formulation and aspiration attainment.

Omission of interdependencies or simplified representation of reality results from a decision-maker's inability to account for the entirety of available information and the ensuing necessity for reduction (*e.g.* Heath and Staudenmayer 2000, Nelson 2008, Walsh 1995). Similarly, errors of commission in specifying interdependencies are consequent to the fact that cognitively constrained individuals cannot always discriminate between informative and noisy signals, thereby constructing illusory relationships and solving spurious tradeoffs (*e.g.* Denrell, Fang, Levinthal 2004, Einhorn and Hogarth 1986).

Importantly, while organizations are prone to making both types of errors, they can be biased (and as argued below, intentionally biased) towards seeing less or more interdependencies. Such biases in the way organizations formulate complex problems can affect the efficacy of search and attainment of organizational aspirations. The following section describes a model that seeks to uncover precisely these relationships.

3.4. Model

Following the neo-computational tradition in organizational sciences, we think of search as a recombination process (Levinthal 1997). Specifically, we define the mapping from strategic configurations to organizational performance using the technical apparatus of the NK model (Kauffman 1993, Levinthal 1997, for a detailed description of the standard NK model see Ganco and Hoetker 2009). In our model, there are three basic elements: objective strategy landscape, subjective understanding thereof, and the process of search.

3.4.1. Objective strategy landscape

The space of organizational strategies is characterized by an N -dimensional hypercube, where each of the N dimensions can take two values (0 and 1) that describe a certain policy choice.¹² Such a space of configurations or strategies has the cardinality of 2^N , and each alternative corresponds to a certain level of fit.

The mapping between organizational strategies and fit reflects the structure of interdependencies among the N elements. Specifically, the value of any configuration is the average of the contribution functions of the activity choices, wherein each contribution function depends on the state of the activity choice itself and those of other elements with which it interacts. Note that we allow K_i to vary across elements, setting the total number of interactions for all N elements (see Rivkin and Siggelkow 2007). For each element there are 2^{K_i+1} unique contribution functions (where K_i is the number of elements with which choice i interacts), which are defined by a random value drawn from a uniform distribution between 0 and 1. If an element does not depend on any other choice, its contribution function remains the same regardless of the state of other elements. For any given choice, therefore, the effective organizational configuration can be characterized as a Hadamard product of the objective organizational form and an adjacency vector (see Ghemawat and Levinthal 2008). The latter is an N -large vector of ones (depends) and zeros (does not depend) indicating the relevance of all elements for the contribution function of the focal one. Therefore, from the perspective of a given choice, an organizational form is different if and only if there is a change in the relevant

¹² Following Adner, Csaszar, and Zemsky, we refer to policies as “all the choices that affect the firm’s performance” (2014: 2796).

element, *i.e.* an element that has the value of 1 in the adjacency vector. When an element does not depend on a given choice, it will have the same contribution function regardless of the state of that decision.

As an example, consider $N = 3$. In this case, the space of possible organizational strategies is defined by the following set: $\{[0\ 0\ 0], [0\ 0\ 1], [0\ 1\ 1], [1\ 1\ 1], [1\ 0\ 0], [1\ 1\ 0], [1\ 0\ 1], [0\ 1\ 0]\}$. Assume that there is only a unique interaction between elements 1 and 3, so that the latter affects the payoff of the former. Then, the contribution functions $c(1, \mathbf{f})$, where \mathbf{f} is the vector of organizational configuration, for element 1 satisfy the following: $c(1, [0\ 0\ 0]) \equiv c(1, [0\ 1\ 0]) \sim U(0, 1)$; $c(1, [1\ 0\ 0]) \equiv c(1, [1\ 1\ 0]) \sim U(0, 1)$; $c(1, [1\ 0\ 1]) \equiv c(1, [1\ 1\ 1]) \sim U(0, 1)$; $c(1, [0\ 0\ 1]) \equiv c(1, [0\ 1\ 1]) \sim U(0, 1)$. For each of elements 2 and 3 there are only two unique values of the contribution functions, which correspond to their states.

3.4.2. Subjective understanding

A key element of our model is how organizations form their perception of the underlying adjacency matrixes. We assume that the process of problem formulation is inherently imperfect in that the probability that agents correctly specify any given interdependence (or lack thereof) is less than one. Guided by this premise, we further consider that problem formulation can be biased. The probability that agents accurately perceive the relationship between choices may depend on the kind of relationship they consider, *i.e.* dependence or independence. Formally, we model problem formulation as the following random process:

$$\Pr(\hat{k}_{ij} = k_{ij}) = \begin{cases} \gamma + \kappa, & \text{if } k_{ij} = 1 \\ \gamma + \omega \cdot \kappa, & \text{if } k_{ij} = 0 \end{cases}, \quad (3.1)$$

where k_{ij} and \hat{k}_{ij} represent, respectively, objective and perceived relationships between choices $i \neq j$ and j (1 for interdependence and 0 for lack thereof), $\gamma \in [0, 1]$ is the overall level of accuracy, κ denotes the bias in defining the interdependence between i and j , and ω is the correction to maintain constant level of accuracy (note that it is negative so that $\omega \cdot \kappa$ differs in sign from κ).¹³ The terms κ and γ are central in our model. While the γ specifies to what extent decision-makers err on average, the κ denotes a bias, and a key lever of behavioral strategy, *i.e.* propensity to make either an error of omission or that of commission in seeing an interdependence. The logic behind the correction term is as follows: if κ increases the accuracy for all $k_{ij} = 1$ becomes greater; hence to preserve overall accuracy constant, the accuracy for all $k_{ij} = 0$ must become correspondingly lower. In other words, while organizations cannot escape from making errors in complex problem formulation, they can choose (or differ in) the type of errors they tend to make. An apt analogy is that of a statistician who is bound by a sample size but can change inference strategies to decrease the probability of false negatives against false positives. We further assume that agents are able to recognize that the performance of the

¹³ Specifically, $\omega = \frac{K}{K + N - N^2}$ and $\kappa \in \left[\max \left\{ -\gamma, \frac{(1-\gamma) \cdot (K + N - N^2)}{K} \right\}, \min \left\{ 1 - \gamma, \frac{-\gamma \cdot (K + N - N^2)}{K} \right\} \right]$, where

N is the number of decisions, K is total number of interdependencies. Observe that the limits of possible bias and the correction term depend on the parameters of the space, N and K . Consider that if the amount of interdependencies K equals $N^2 - N$ there can be no bias in problem formulation as the underlying structure of interdependencies can be only underspecified.

choice depends on its own state—that is, $\Pr(\hat{k}_{ij} = k_{ij}) = 1$ for all $i = j$ (our main observations remain qualitatively similar if we relax this assumption).

Figure 3.1. Misperception of interdependencies: example




Figure 3.1 illustrates how a process of problem formulation may alter the perception of the structure of interdependencies. When $\kappa = 0$ there is no bias, and organizations are equally likely to commit the error of commission and that of omission in specifying the relationship between choices. However, when $\kappa > 0$ organizations are more likely to perceive interdependencies where they exist not (*i.e.*, they overcomplicate); and inversely, when $\kappa < 0$ organizations have a greater tendency to overlook existing interdependencies (*i.e.*, they oversimplify). As a result, the space of alternatives is characterized not by a single landscape but rather by a pair of landscapes. One describes the true nature of the environment, whereas the other reflects how the organization perceives it (cf. Csaszar and Levinthal 2016).¹⁴

¹⁴ Note that mental representations of reality are generally considered to be simplified, abridged or crude models of the underlying problem (Gavetti and Levinthal 2000, Knudsen and Srikanth 2014, Levinthal 2011). We fully

In our model, the difference between the subjective and the objective landscapes is driven solely by the discrepancies in the actual and perceived adjacency matrixes. Specifically, for each of the N organizational elements we generate 2^N would-be contribution functions. We then construct the two landscapes using this set of contribution functions and the respective adjacency matrixes. The actual and the perceived landscapes share the same pool of values but differ (albeit intersect) in the subsets of contribution functions used to construct the mapping from configurations to performance.

3.4.3. Search process

Given a mental representation of the strategic problem, organizations search for solutions. A defining behavioral element of this process is the tendency to generate alternatives within the immediate vicinity of the current state. While there is empirical evidence as well as broad theoretical support for the prevalence of local search (*e.g.* Stuart and Podolny 1996), organizational scholars have also repeatedly stressed that occasionally organizations consider distant alternatives (*e.g.* Gavetti 2012, Levinthal 1997, Rosenkopf and Almeida 2003). Accordingly, we model a process whereby organizations tend to generate local alternatives, but sometimes come up with distant courses of action. Specifically, we assume that in every period organizations generate distant alternatives with some low probability.¹⁵

concur with this view and observe that our conclusions hold for a reduced representation of the strategy landscape, one where perceived N' is lower than the actual N . In the interest of clarity, however, we leave this analysis beyond the scope of the present work.

¹⁵ In the results reported below, this probability is set at 0.1. However, our observations remain qualitatively similar for other levels of non-strictly local search. We have also considered specifications where the likelihood of generating an alternative monotonically declines with the distance and report that such formulations lead to similar conclusions.

In the same vein, since organizations experience objective reality but make decisions based on their subjective understanding thereof, we model search as partially cognitive, partially experiential (see Gavetti and Levinthal 2000). Specifically, agents evaluate alternatives and make decisions based on their mental representations but reject alternatives on receiving negative feedback from the objective reality and may return to the previous state. Since such a backward-looking correction is possible if and only if organizations experience the actual (adverse) performance feedback, cognitive evaluation on misperceived landscapes can produce persistent blind spots (cf. Zajac and Bazerman 1991). These blind spots, in combination with potential stumbles upon acceptance of inferior alternatives, decelerate and may preclude attainment of organizational objectives, or aspirations.

3.5. Analysis

We regard aspiration level as a vector $\mathbf{v} = (v_1, v_2)$, where v_1 denotes target performance and v_2 stands for target time. Such conception of aspiration levels closely aligns with organizational goals in that we consider aspiration level to reflect a desired level of performance within a given time horizon. Observe that annual performance goals represent a special case of our approach, where v_2 corresponds to one year (note that this is not to say that $v_2 = 1$). In this conception, organizations satisfice by accepting whatever alternatives they believe will reduce the discrepancy between the aspiration level and the current state.

Note that technically our specification directly corresponds to that in the existing models (see Ganco and Hoetker 2009 and references therein). However, we deviate from previous modeling approaches conceptually in that we think of aspiration levels as not

necessarily synonymous with acceptable alternatives. In our conception, an alternative below the aspiration level can be accepted, but if (and only if) it is perceived to bring the organization closer to the desired performance. This applies to contexts where acceptance of one alternative does not in principle preclude acceptance of another, which is the case of strategic recombination problem (cf. marriage problem). Note further that for the mechanics of adaptation, both views are equivalent: organizations satisfice by accepting the first alternative expected to add value over and above the status quo.

3.5.1. Likelihood of attainment

Given that in our conceptualization aspiration level internalizes both performance and time, a natural objective function is

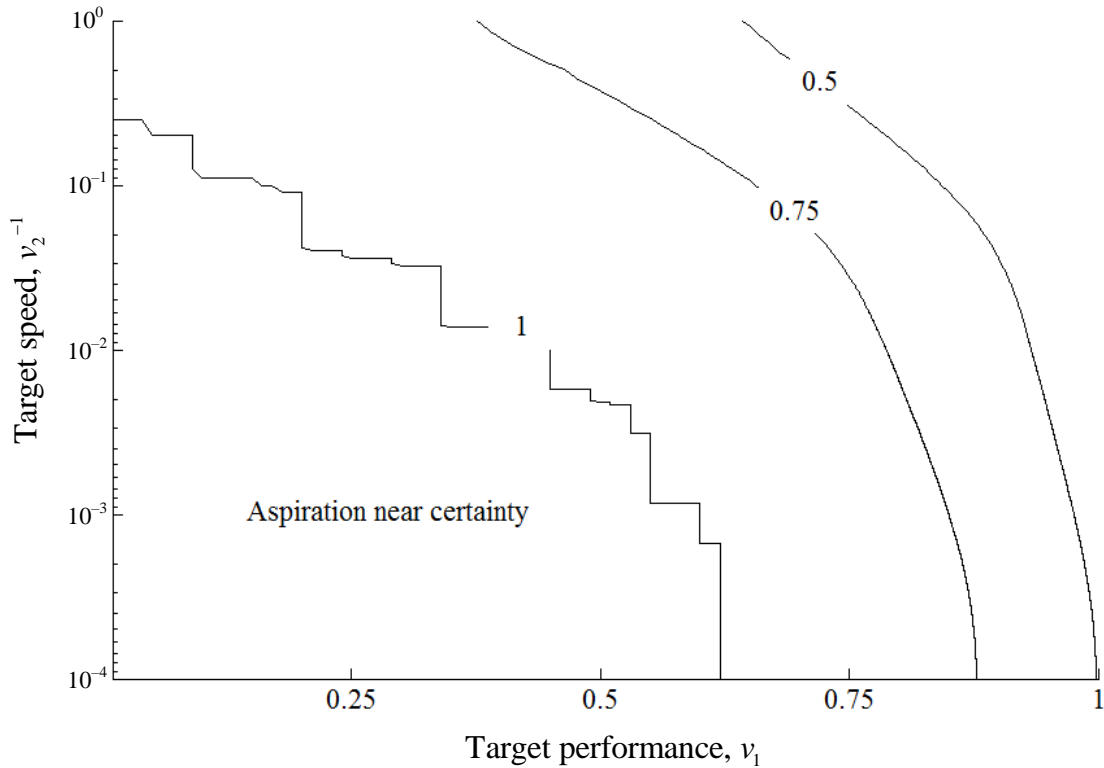
$$p = \Pr(\pi_{v_2} \geq v_1), \quad (3.2)$$

where $\pi_{(.)}$ is performance at time $(.)$, and p denotes the probability that aspiration level \mathbf{v} is reached. This function is declining in v_1 and v_2^{-1} (note that we invert time to account for increasing difficulty associated with greater speed of desired performance increment). Accordingly, Figure 3.2 shows that iso-probability curves—that is, curves that connect aspirations with equal likelihood of attainment—are concave (these and the following results are derived for at least 10^5 simulation runs).¹⁶ For a given p , an increase in either v_1 or v_2 requires an increase in the other. Higher target performance renders aspiration level harder to

¹⁶ The iso-probability curves are analogous to production-possibility frontiers in that they indicate combinations of time and performance attainable with a given level of certainty.

reach. Inversely, more time supports greater likelihood of attaining a given level of performance.

Figure 3.2. Iso-probability curves

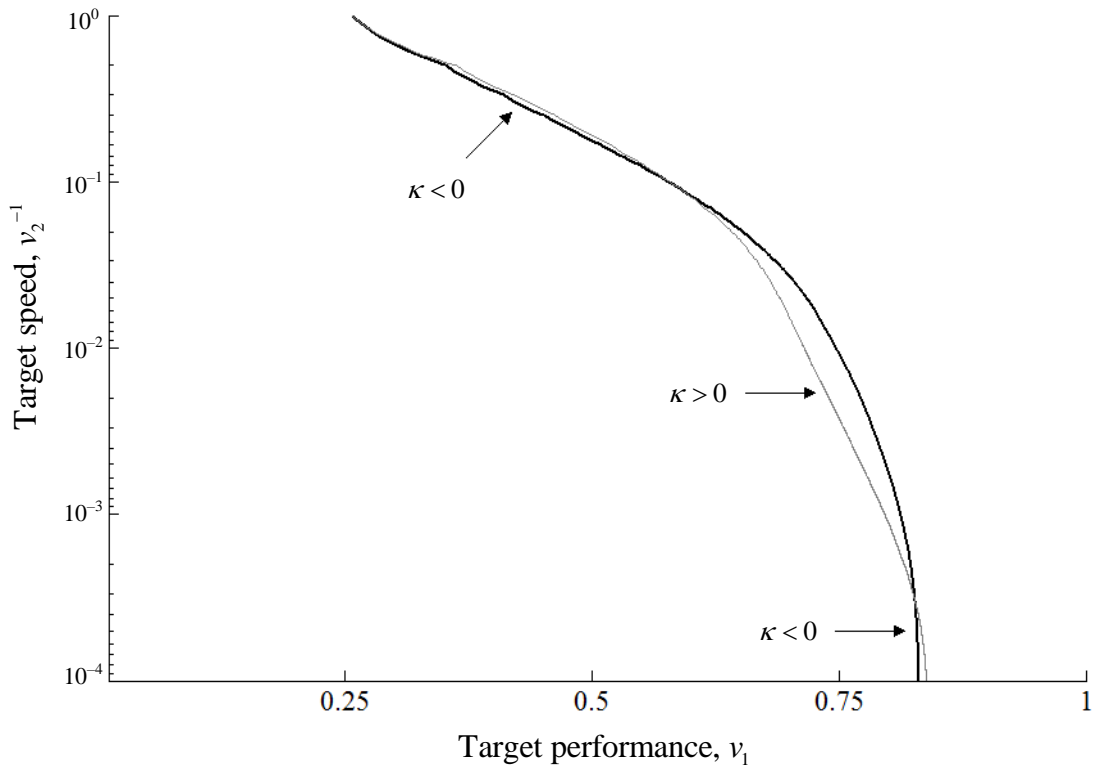


Errors in problem formulation tend to reduce maximum achievable level of performance. Specifically, iso-probability curves for an accurate representation (not displayed for clarity) are strictly dominant. However, for a given time and accuracy there is performance level a boundedly rational agent will reach with $p \rightarrow 1$. All \mathbf{v} for which this holds form what we refer to as aspiration near certainty. In simple terms, objectives characterized by aspiration near certainty correspond to the lowest ambition possible.

3.5.2. Effect of biases

While similar iso-probability curves hold for all misspecifications, we find that their form depends on the bias in problem formulation. Specifically, when $\kappa > 0$ (bias toward complexity), iso-probability curves—in the plane $v_1 \rightarrow v_2^{-1}$ —bend outward for low and high values of v_1 , and inward for medium values of v_1 (see Figure 3.3). An inverse effect is true for $\kappa < 0$ (bias toward simplicity).

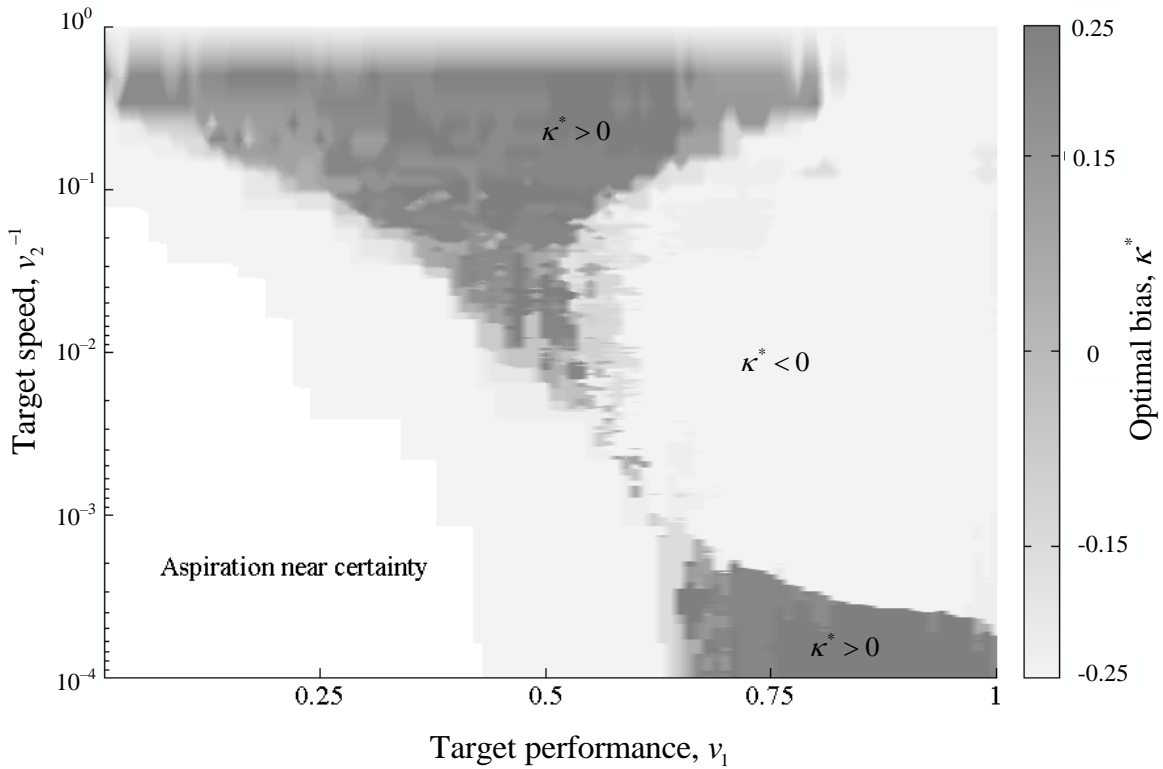
Figure 3.3. Shift in iso-probability curves



These observations have direct implications for optimal behavior with respect to formulating problems. Given some convex preference between speed and performance, the maximum utility from search can be a function of biases in problem formulation. To illustrate

this, Figure 3.4 shows how the described change in the curvature of the iso-probability curves reflects on the optimal bias in the process of problem formulation. Counter to what may be expected, biased problem formulation is almost universally superior to the unbiased. In other words, for nearly all \mathbf{v} there exists a bias $\kappa \neq 0$ that leads to greater probability of attainment than unbiased problem formulation, $\kappa = 0$. Recall that the actual size of the bias has no bearing on the average count of errors in specifying interdependencies, *i.e.* bias only changes the relative prevalence of one type or error over the other. These observations suggest that from the perspective of efficiency, the question that matters is not *whether* boundedly rational organizations should be biased when specifying the structure of interdependencies, but rather *how* they should be biased. Our analyses point to the following general pattern: bias toward complexity dominates when either time or performance is prioritized, whereas bias toward simplicity is superior when both time and performance are important. [continued on the next page]

Figure 3.4. Optimal bias (maximizing p)



In the above analysis, we have considered aspiration levels exogenous in that the pair \mathbf{v} was independent of p . However, in reality there is an (efficient) correspondence between uncertainty and aspiration level adaptation (see Greve 2002). In other words, decision-makers may also take into account the underlying (perceived) risk of not reaching the desired performance at the desired time. Specifically, organizations may, for example, want to seek minimum time (ν_2) such that performance (ν_1) is reached with a given level of certainty.

Figure 3.5. Optimal bias (minimizing $t \mid p = 0.9$)




Figure 3.5 shows the relationship between the target performance and the optimal bias in problem formulation that minimizes the time needed to reach this performance holding the probability of attainment constant.¹⁷ Specifically, if aspirational performance is medium to relatively high, a bias toward the error of omission when specifying the structure of interdependencies minimizes the time needed to meet an aspiration. In contrast, if organizations seek low or very high performance, organizations are better off erring on the side of seeing greater rather than lower business complexity.

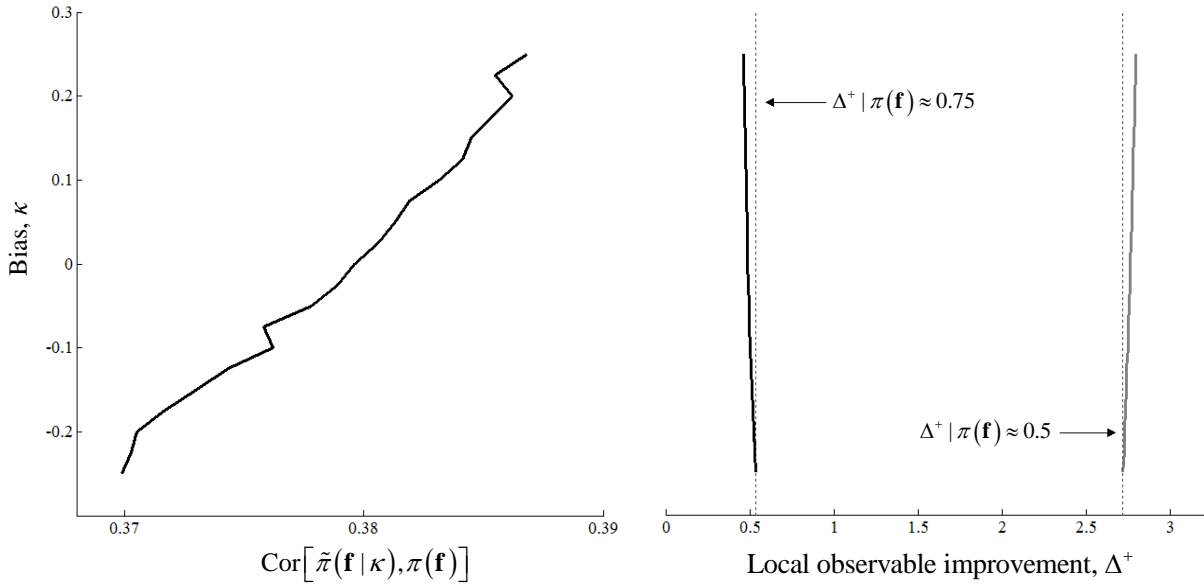
¹⁷ Specifically, in the interest of better exposition, in the figure we illustrate the results for $p = 0.9$. Note also that for clarity here we omit interim degrees of bias.

From the perspective of behavioral strategy, this implies that depending on the context (the relative importance of aspirational time and performance) organizations may choose to strategically alter their behavior in relation to formulating problems. Alternatively, given decision-makers' propensity to see greater or lower complexity in business problems, organizations may choose to compete in contexts that are more aligned with the way their key decision-makers think.

3.5.3. Underlying mechanisms

Errors in defining interdependencies change the perceived geometry of the organizational landscape. By ignoring the interactions among organizational elements, decision makers construct a more internally correlated mental image of reality, one that contains fewer local maxima as compared to the objective landscape. Given such an understanding, the perceived optimal configuration is more readily within the reach of search. In contrast, by assuming nonexistent interactions, decision-makers mentally construct a more complex strategic problem, where the actual peaks are complemented by spurious local maxima. This decelerates adaptation as organizations stall at configurations thought to offer no immediate improvement. [continued on the next page]

Figure 3.6. Properties of perceived landscapes



Note. $\pi(\mathbf{f})$ and $\tilde{\pi}(\mathbf{f}|\kappa)$ denote respectively objective and perceived (given bias κ) performance of configuration \mathbf{f} , \mathbf{f}_a stands for an alternative within immediate vicinity of \mathbf{f} , and local observable improvement, for brevity, defined formally as $\Delta^+ = |\{\mathbf{f}_a : \tilde{\pi}(\mathbf{f}_a|\kappa) > \pi(\mathbf{f}), \pi(\mathbf{f}_a) > \pi(\mathbf{f})\}|$.

Such changes in the perceived geometry of the landscape allow those who oversimplify ($\kappa < 0$) to continue active search even after having accrued relatively high performance.¹⁸ Figure 3.6 shows that when fit is in the vicinity of 0.75 (we take the reported level ± 0.05), local observable improvement is highest for $\kappa = -0.25$ (see right panel). That is, errors of omission in defining interdependencies allow for greater likelihood of discovering a better configuration near a well-performing status quo. However, this does not hold for all levels of fit. In particular, Figure 3.6 shows that when performance is relatively low (*e.g.* in

¹⁸ Observe that while we compare performance in terms of accumulated knowledge, accounting for the implicit cost of trial—that is, measuring the spot performance—does not change our qualitative conclusions.

the vicinity of 0.5) it is $\kappa = 0.25$ that provides highest local observable improvement. While $\kappa > 0$ increases the probability of freezing at a spurious local peak, when not on a peak it affords more paths to improvement. Consider a case where the objective reality offers three better alternatives near the current state, one leading to the global optimum and two to local peaks. At the extreme, oversimplifying will entirely conceal those not on the way to the perceived maximum, whereas overcomplicating will create false positives but is less likely to hide true local optima. As a result, the chance of discovering local improvement is higher when $\kappa > 0$.

Given local search, these effects characterize the relationship between κ and performance. However, relaxing the assumption of strictly local generation of alternatives reveals that $\kappa > 0$ tends to dominate for aspirations with sufficiently long time horizons, or equivalently with sufficiently low speed. Figure 3.6 shows that the average correlation between the objective reality and the subjective understanding thereof is increasing in κ (see left panel). This means that seeing a more complex picture of reality reduces the loss of information about the true landscape, a property that is beneficial when organizations have found ways to walk away from local peaks or before they have reached them. As a result, $\kappa > 0$ increases the chances of success for aspirations that allow for either long or short period of search.

3.5.4. Aspiration level adaptation

Although in our conception aspirations integrate both performance and time, this is not to say that they remain unchanged. A long tradition in organizational sciences points to an adaptive

nature of aspiration levels (Cyert and March 1963, Simon 1955). In broad terms, organizational objectives tend to change in response to the observed performance feedback.


While there is a wealth of models describing the dynamics of organizational objectives (see for example, Greve 2003, Lant 1992, Levinthal and March 1981, Mezas *et al.* 2002), empirical evidence suggests that attainment discrepancy is key to understanding the formation of goals (Lant 1992, Mezas, Chen, Murphy 2002). Specifically, an aspiration level at a given moment in time is some function of the difference between the observed performance and the prior aspiration level.¹⁹ Our conception is logically consistent with such a dynamic in that attainment discrepancy at any moment in time is a critical determinant of the targeted performance. To see this, consider how organizations can respond to the observed performance feedback.

Figure 3.7 illustrates the implications of both negative and positive attainment discrepancy (left panel). When the observed performance at time $t < v_2$ is greater than v_1 (positive attainment discrepancy), there is some probability that performance at v_2 will be even better. Accordingly, organizations can raise their aspirations by adjusting v_1 upwards by a factor greater than one. The response, however, is not as simplex when performance at time $t < v_2$ is lower than v_1 .

¹⁹ Note that scholars have identified a swath of factors that affect the dynamics of aspiration levels, including, industry performance, social expectations, *etc.* (see Washburn and Bromiley 2012). While we agree that internal and external antecedents affect organizational objectives, in the interest of tractability and focus, our model abstracts away from these factors.

A simple reaction to negative attainment discrepancy is depression of the desired performance at v_2 . If the probability of reaching v_1 at v_2 conditional on the observed performance is lower than that implicit in \mathbf{v} , organizations can lower their expectations of v_1 and thereby maintain constant level of attainment. Equivalently, decision-makers can extend their estimate of the time needed to reach v_1 , *i.e.* increase v_2 . While similar in effect on v_1 (observe that in both cases organizations will reduce their expected performance at the initially planned moment in time), the latter response is distinct behaviorally. To see this, consider a project rather than an ongoing concern; for example, a car manufacturer launching a new electric model. Suppose that the initially planned lifetime period is ten years. If after five years, the figures show mediocre sales, the manufacturer can either lower its outlooks for the cash flow from the project or extend its lifetime. By doing so, it will shift the probability of reaching a certain level of performance upwards (see Figure 3.7, left panel). [continued on the next page]

Figure 3.7. Conditional probability of attainment



Note. Consider aspiration $\mathbf{v} = (v_1, v_2)$, where $v_1 = 0.8$ and $v_2 = 5 \cdot 10^2$. For such a \mathbf{v} it is optimal that $\kappa = -0.25$, *i.e.* oversimplify. Assume that at moment $t = 2.5 \cdot 10^2$ organizations observe performance π_t and reformulate strategic problem following process (3.1) where $\kappa = \kappa_r$. The following curves show the probability of observing π_{v_2} at $t = v_2$ conditional on: (a) $\pi_t = 0.7$, (b) $v_2 = 2.5 \cdot 10^3$, (c) $\pi_t = 0.85$, (d) $\kappa_r = -0.25$ (e) $\kappa_r = 0.25$.

3.5.5. Problem reformulation

An essential element of organizational behavior, a downward aspiration level adaptation represents a passive response. Yet organizations can also take active measures to enhance their chances of reaching the desired performance at the desired time. An effective means of doing so can be a periodic reconsideration of a problem at hand, or the construction of an alternative mental representation. While existing works have pointed to this mode of the adaptive search (see Csaszar and Levinthal 2016, Gavetti and Levinthal 2000), we further find that by

managing systematic errors in problem *reformulation*, organizations can raise the efficacy of search.

Counterintuitively, optimal biases in reformulation do not necessarily align with those in the initial problem formulation. Figure 3.7 (right panel) shows that for a \mathbf{v} that requires $\kappa < 0$ in the initial problem formulation, $\kappa_r > 0$ in reformulation can result in higher conditional probability of attainment (note, however, that we do not observe an inverse pattern). That is, it may be useful to change the bias in one's understanding "midway" in the process of search toward organizational objectives.

These observations suggest that organizations should not only manage how they formulate strategic problems *ex ante*, but also consider reassessing their understanding *ex intra*. Biases in such reassessments can further enhance the chances of attaining organizational objectives.

3.6. Intentionally biased

Our analysis suggests that unbiased understanding of business complexity can be suboptimal for the attainment of organizational objectives. Organizations may therefore strive to intentionally bias the decision-making processes to increase the chances of attaining their aspirations. While a full exploration of this issue is beyond the scope of this study, here we discuss potential instruments (not restricted to problem formulation) that organizations can use to be intentionally biased.

Consider that organizations differ in the architecture of their decision-making. In particular, organizations can be classified along a continuum between hierarchies and polyarchies (*e.g.* Christensen and Knudsen 2010, Knudsen and Levinthal 2007, Sah and Stiglitz 1986). In a pure form of a hierarchy, decisions are made based on the rule of unanimity. Conversely, in a pure form of a polyarchy, a single member's approval suffices for the decision to be accepted. If judgments are noisy and independent, hierarchies mitigate the error of commission and polyarchies that of omission (Csaszar 2012). These properties of the two architectural extremes (as well as intermediate hybrids) are often viewed as directly applying to the choice outcome, *i.e.* acceptance or rejection of alternatives (see Knudsen and Levinthal, 2007). Yet in complex strategic choices, decisions are preceded by discussions whereby the understandings of the problems at hand are socially constructed (*e.g.* Cyert and March 1963, Eisenhardt and Bourgeois 1988, Smith *et al.* 1994). In other words, problem formulation is a social process that can be regulated by formal or informal structure (Baer *et al.* 2013). The properties of hierarchies and polyarchies, therefore, apply equally to the social construction of the perceived complexity being resolved as they do to the process of evaluation. Specifically, given no (or few) interdependencies as the initial condition, hierarchies and polyarchies will result in simplified and complicated representations respectively.

Modularity is potentially another instrument that organizations can use to manage the extent to which they complicate or simplify certain problems, particularly that of searching for efficient product or technology solutions (Brusoni and Prencipe 2001). As Adner and Kapoor note, "modularity does reduce interdependence" (2010: 312). Hence, by mismatching forms of modularity, organizations can reproduce some of the effects related to misperception of

complexity (see Ethiraj and Levinthal 2004). However, it is important to note that conceptually modularity is distinct from choice interdependence (*ibid.*). Therefore, misspecified modularity does not fully account for misconception of the complexity of a given organizational problem (see also Brusoni, Marengo, Prencipe, Valente 2007).

More broadly, organizations may rest on different kinds of decision-making routines: more complex, where a focal member consults multiple departments, subsidiaries, controllers, or managers; or more simplex or decentralized, where members autonomously solve their problems (see Nelson and Winter 1982). From this perspective, organizations are complex adaptive systems that misalign their internal processes with the complexity of business environment (on organizations as complex adaptive systems, see for example, Anderson 1999). By architecting more inclusive, interactive processes, organizations can create conditions in which members consider more factors and perspectives than objectively warranted (cf. Ethiraj and Levinthal 2004).

In a similar vein, organizations can control the complexity of members' decision processes by regulating the content and format of the coded knowledge. Manuals or decision-support systems can be designed to invoke more or less complex patterns of reasoning. More importantly, however, the very presence of articulated knowledge (as contrasted with absence thereof) can serve as a mechanism to promote a simplified understanding of the objective complexity. In particular, explicit articulation resolves part of the *ex-ante* ambiguity as to the mapping between actions and outcomes (Zollo and Winter 2002). As a result, decision-makers more readily accept choices determined only by the articulated contingencies; contextual or unarticulated interdependencies are effectively omitted. Further, organizations can rely on the

use of deliberately simplified decision-making procedures guided by heuristics or basic rules of thumb (Eisenhardt and Sull 2001, Sull and Eisenhardt 2012).

Finally, at a more fundamental level, an organization's choices reflect the personal characteristics and behavioral attributes, and ultimately the mental models, of its managers (*e.g.* Gary and Wood 2011, Hambrick and Mason 1984). Central from this perspective is one's cognitive complexity, or propensity to take into account multiple perspectives, dimensions, or relationships (*e.g.* Streufert, Suedfeld, Driver 1965, Suedfeld, Tetlock, Streufert 1992, Tetlock, 1983). More cognitively complex individuals are arguably more susceptible to the error of commission when specifying interdependencies among organizational choices, and vice versa. To the extent that decision-makers are replaceable, organizations can intentionally bias problem formulation, *e.g.* by choosing those with needed behavioral patterns. Inversely, individuals may want to self-select into projects, ventures, or environments that align with their (mis)perceptions of reality.

3.7. Conclusion and Discussion

Our analysis points to an important and understudied relationship between problem formulation and attainment of organizational objectives. Specifically, our findings suggest that organizations can strategically use biases when defining the structure of interdependencies among decisions. Depending on the aspirations, organizations can benefit from either overcomplicating or oversimplifying when formulating strategic problems. From the microfoundations perspective (see Baer *et al.* 2013), this points to a need to further explore the link between a group's problem formulation and the complexity of the resultant representations

that define organizational choice. A direct implication of the reported regularities is a possible mapping between group processes of problem formulation and attainment of organizational objectives.

At the conceptual level, our analysis underscores the importance of viewing aspirations as compounds of performance and time. While in no conflict with the existing knowledge on aspirations (recall that time can be implicit), such a conception may allow a more inclusive analysis of aspiration levels in organizations, as well as their dynamics and consequences.

An important aspect of aspirations is that they are in part a product of social comparison (Cyert and March 1963). Such social nature of aspiration levels—albeit not explicit in our analysis—is fully consistent with our observations. Consider that the intensity and kind of competition in the industry as well as strategies of other firms may well define the marginal rate of substitution between performance and time. By extending our argument, the optimality of bias in problem formulation can be, at least to some extent, socially constructed.

Our findings directly speak to the emerging field of behavioral strategy. One of the core aims within this research stream is to uncover psychological sources of value creation (see Gavetti 2012, Levinthal 2011, Powell *et al.* 2011). From this perspective, our work points to a correspondence between organizational aspirations and the cognitive styles of the key decision-makers. Our results also indicate that sticking points on the performance landscape (see Rivkin and Siggelkow 2002) can arise even in the absence of any incentive distortions. The reason for such non-incentive-based sticking points lies in the imperfections of human mind, which cannot faultlessly define the true structure of the interdependencies among organizational decisions. Faulty representations can create specific behavioral failures that

prevent organizations from seeing superior configurations (see Gavetti 2012). Importantly, however, we show that it is possible to partially overcome these failures by a timely and directed reformulation of the strategic problem.

Our work further contributes to the ongoing debate on how decision makers should think for organizations to cope with the challenges of the modern business environment. Nadkarni and Narayanan observe that the existing literature offers arguments for both complex and simple decision processes (2007). On the one hand, the managerial cognition literature advocates the use of complex and varied knowledge to support a more efficient organizational change (*e.g.* Lant, Milliken, Batra 1992, Nadkarni and Narayanan 2007, Wally and Baum 1994). On the other hand, prior work has also suggested that simple choice rules promote more active experimentation and thereby greater learning and superior outcomes (*e.g.* Eisenhardt and Martin 2000, Eisenhardt and Sull 2001). We show that both perspectives may in fact hold true, the appropriateness of decision-making complexity in a given setting depends on how organizations prioritize time and performance in their aspirations.

These observations derive from a formal computational model. Although the technical apparatus we use—NK model—finds support across different research domains, a note on the main boundary conditions is nonetheless in order. First, our findings apply to ambiguous problems that require active search, or trial and error. Observe that when the problem is objectively well-defined, there is no misspecification of interdependencies. Second, we assume that the ability to search is independent of performance. An implicit assumption in most work on organizational adaptation (for an exception see the works on survival, *e.g.* Levinthal and Posen 2007), this may alter the mapping between aspirations and optimal

biases. Consider that if greater performance affords more active search, the initial advantage of overcomplicating will expand. However, to the extent that the magnitude of such an addition in experimentation is not extreme, our qualitative conclusions will hold. Third, our analysis applies to aspirations on new landscapes, *i.e.* new ventures, start-ups, projects, products, market entries, post-shock performance, *etc.* The relationship between misconception of reality and incremental aspirations of ongoing concerns represents an intriguing subject for further study. Our analyses of aspiration level adaptation and problem reformulation may serve as a useful starting point for such a study. Fourth, the optimality of biases is derived for an abstract clock speed, where one unit of time comprises one instance of search. How this model clock speed translates into real-world time is an empirical question, the answer to which is likely to differ across contexts.

While at the conceptual level, our work develops the first explicit link between problem formulation and aspiration attainment, in its subject matter it relates to the recent work by Martignoni *et al.* (2016). Besides the conceptual argument, however, several specific elements advanced in our work are worth separate attention. To start with, while we draw from Martignoni and colleagues, we use a different (and arguably more broadly validated) analytical apparatus of NK model. This allows direct parallels with the previous work, but more importantly makes clear how misconception of reality alters the outcomes of recombination search in organizations. In particular, distinct from Martignoni and colleagues, we find that if such a general search is not prohibitively local (Levinthal 1997, cf. Knudsen and Levinthal 2007), misconception of reality is strictly detrimental for organizations. In other words, organizations should strive to have an accurate representation. However, in the world of bounded rationality—where complete knowledge is inaccessible—an accurate

representation of the business problem is often not a fitting benchmark. Accordingly, as the benchmark we take a stochastically imperfect (and unbiased) specification of interdependencies rather than accurate problem formulation. A result is a picture, where an unbiased problem formulation is nearly always dominated by a bias toward either simplifying or complicating. Unlike previously reported observations, our pattern is remarkably robust, holding for essentially all levels of complexity (provided a bias is possible, see equation (3.1) and footnote 13), reduced representations of the landscape, and different non-extreme distances of search. In contrast with the previous work, therefore, we report a general effect that holds for a wide range of conditions.

Overall, our results suggest a practical rule of thumb for forming an efficient understanding to guide strategic choices. If one has a strong preference for either time or performance, she is better off overcomplicating. If on the contrary, the decision-maker must care about time and performance equally, a more apt behavioral strategy for her is to keep it simple.

**Essay III. CEO's Conceptual Complexity and
Performance Dynamics**


4.1. Introduction

Essay III draws on the same conceptual and analytical apparatus as essay II. Indeed, this chapter essentially provides an empirical test of a corollary of the effects identified in the previous essay. Accordingly, to suppress repetition, this essay proceeds directly to the analytical part.

While the previous essay focuses on the attainment of dual aspirations, the results of the model (see section 3.4) imply a direct and testable relationship between systematic errors in perceiving interdependencies among organizational choices and the dynamics of organizational performance. Figure 4.1 illustrates such theoretical performance dynamics for the base case parameters in the model in essay II. In this figure, the horizontal axis measures model time t and the vertical axis shows the difference in performance between overcomplicating and oversimplifying agents. The figure shows a highly non-linear relationship in model time.²⁰ Below we provide empirical evidence elements of this relationship may hold in reality. [continued on the next page]

²⁰ To see the connection to the analytical part, consider the difference between two iso-probability curves in Figure 3.3 as a function of time.

Figure 4.1. Relative performance dynamics



Note. Horizontal axis measures time on logarithmic scale. Vertical axis depicts the relative performance of seeing more complex reality, or formally $\pi_t(\mathbf{f} | \kappa > 0) - \pi_t(\mathbf{f} | \kappa < 0)$.

4.2. From model to data

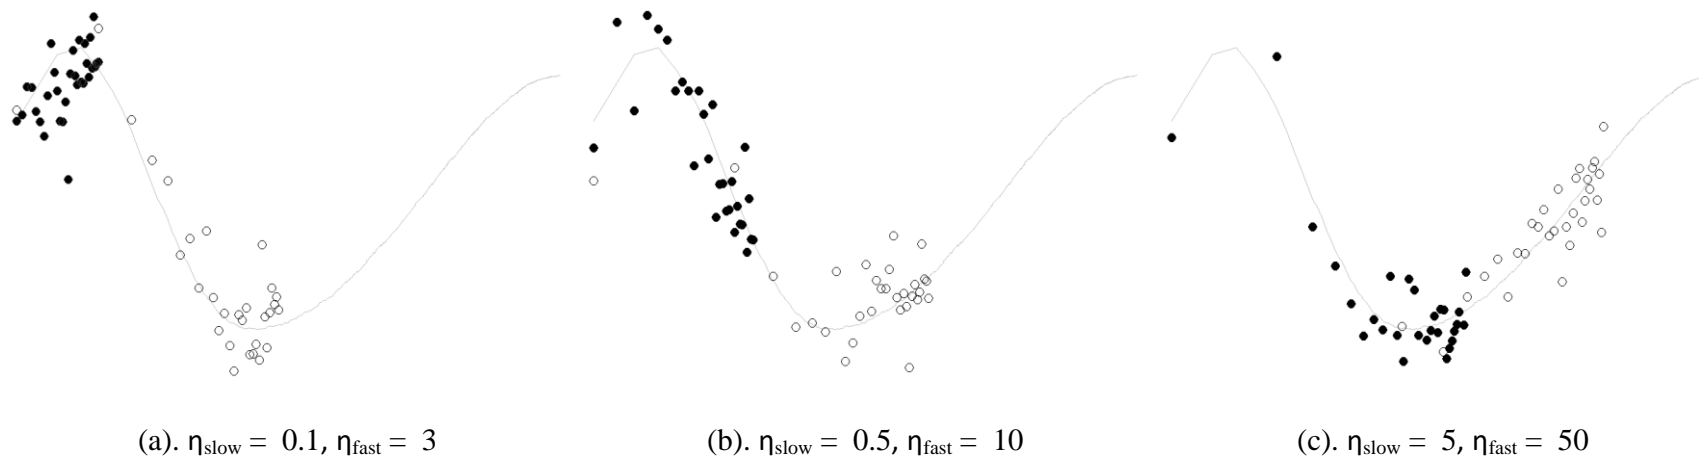
4.2.1. Empirical challenges

The first key challenge in bringing our (or any similar) model to data is that the reported pattern of performance dynamics applies to the process of organizational search on a new landscape. In other words, moment $t = 0$ corresponds to the beginning of organizational adaptation to a novel set of conditions. In theory, this need not complicate the empirical

analyses, provided that the dependent variable is a strictly monotone function of model time. However, as our model's predictions are highly non-linear in model time, to test such effects the observation period should follow an exogenous shock that transformed the landscape: *e.g.* venture creation, major economic crisis, or economically substantial political event (such as the collapse of the Soviet Union or, possibly, Brexit).

The second key challenge is that the mapping from model time to real-world time cannot be known *ex ante*. Consider that in our model, one period corresponds to a single instance of generation and evaluation of alternatives. In reality, however, organizations can make multiple choices per unit of time (*e.g.* year, quarter, month, *etc.*). Since performance is measured discretely (*e.g.* quarterly or annual results), the mapping from the model's predictions to the observed performance dynamics will depend on the pace with which organizations generate and evaluate alternatives, or clock speed. Figure 4.2 illustrates how the predictions of our model change as a function of the clock speed (panels illustrate different clock speeds). Note that by clock speed η we mean the amount of alternatives, *i.e.* potential actions, implicitly generated and evaluated per empirical unit of time (see organizational clock speed in Fines 1998). That is, if $\eta = 1$ then over 10^4 units of time (months, quarters, or years), we can expect to observe the relationship as shown in Figure 4.1, corresponding to 10^4 cycles of generation and evaluation of alternatives. However, in reality, it is unlikely that organizations make a single choice per observable unit of time, nor can we observe performance for such long periods. Accordingly, Figure 4.2 demonstrates possible empirical patterns of observations for different η and observation period of 30 units of real-world time (*e.g.* 30 quarters or years).

Figure 4.2. Stylized empirical patterns



Note. Solid dots show possible observations in an environment with slow clockspeed, η_{slow} , whereas void ones exemplify fast clockspeed, η_{fast} . Horizontal axis indicates model time t (logarithmic scale), relative rank of dots from left to right corresponds to the real-world time (*e.g.* first, second, third quarter, *etc.*), and the vertical axis is relative performance of overcomplicating.

4.2.2. Predictions and context

As a result, multiple patterns of performance dynamics under systematic errors in understanding complex business reality are consistent with our model. At the fundamental level, however, our general prediction is that misconceiving complexity changes the performance trajectory on a new landscape. In other words, time moderates the relative payoff to errors of omission and commission (or their correlates) in specifying interdependencies among organizational choices. Further, one can also expect that such a relationship depends on the relative clock speed. In particular, in contexts with greater clock speed, each real-world moment of time should correspond to more distant periods in the model (see Figure 4.2).

We test these predictions in the context of two sectors of the US economy. Specifically, we consider constituents of S&P 500 (as of October 12, 2015) that belong to two basic sectors in the Global Industry Classification Standard (2-digit): industrials and information technology. These sectors are comparable in size but arguably different in the underlying organizational pace of decision-making, a property that may allow us to observe whether the temporal effect of misunderstanding reality varies across clock speeds and whether this change is consistent with our model. Further, this context agrees with the key boundary conditions of our analysis. First, organizational setting represents an ambiguous environment where organizations learn from trial and error, and therefore the mechanisms identified in our analysis should be at play (*e.g.* Lant and Montgomery 1987). Second, the companies in our sample are sufficiently large to ensure that the ability to search, *i.e.* make strategic decisions, is not critically dependent on performance (*e.g.* Sharfman, Wolf, Chase,

Tansik 1988). And finally, to partially address the necessity of observing search from its starting point, we can use the event of the 2008 financial crises as an exogenous shock that should, in principle, have substantially changed the organizational landscape (see for example, Ghemawat 2010, Lewin, Long, Carroll 1999).

4.3. Main measures

The key variables of interest in our analysis are time on a new landscape, performance, and a correlate of systematic errors in misconceiving complex reality.

4.3.1. Performance and time

As a measure of time on a new landscape, we take the count of quarters since January 1, 2009—the bottom of the latest global financial crisis—where the first quarter of 2009 is counted as 0. We use quarters rather than years to allow for the possibility that the correspondence between model time and real-world time is substantially shifted to the right or left, *i.e.* η is very high or very low.

We use standard return on assets as a measure of performance. Although organizational fit (performance in the model) cannot be negative, while ROA certainly can, the premise is that low levels of return on assets correspond to low levels of fit. Importantly, the use of return on assets, a commonly accepted measure of performance, allows for a direct link of our analysis to the existing body of work on the antecedents of organizational success.

4.3.2. Conceptual complexity

As a correlate of seeing, perceiving, or implicitly formulating a more complex picture of reality, we consider cognitive complexity. Cognitive complexity is one's ability to discern multiple dimensions in a unit of information and to interact or integrate them (Suedfeld and Coren, 1992). Together with its near alternates, such as integrative or conceptual complexity, this personality attribute is central in psychology and has been shown to be consequential in competitive settings, such as international conflicts (Suedfeld and Bluck 1988, Suedfeld and Tetlock 1977). Similarly, individuals' cognitive complexity plays a role in different organizational phenomena, as, for example, employees' pay satisfaction, organizational learning, corporate social performance, or organizational scope (see Carraher and Buckley 1996, Calori, Johnson, Sarnin 2006, Hayes and Allinson 1998, Wong, Ormiston, Tetlock 2011). Here we argue that cognitive complexity can further serve to assess the temporal consequences of misconceiving complex business reality.

Our basis is that cognitively complex individuals, all else being equal, are more prone to the error of commission when understanding interdependencies, and vice versa. Consider that individuals who perceive, integrate, and interact multiple dimensions in units of information will be more likely to entertain spurious (or negligible) relationships among organizational elements. While situational factors affect this tendency, existing work indicates that cognitive complexity is also an important personality trait (*e.g.* Suedfeld and Bluck 1993, Suedfeld and Coren 1992, Tetlock 1985, Tetlock, Peterson, Berry 1993, Thoemmes and Conway 2007, cf. for example, Suedfeld and Bluck 1988, Porter and Suedfeld 1981).

Therefore, organizations with cognitively complex decision-makers will tend to adapt along a pattern consistent with that of a bias toward seeing greater complexity, *i.e.* $\kappa > 0$.²¹

We measure cognitive complexity using an automated multiple-pass procedure developed by Young and Hermann (2014). This procedure allows unobtrusive and consistent measurement of cognitive or more specifically conceptual complexity, which we use as a proxy for seeing the objective reality as less or more complex. While this measure is based on a complex multiple-pass text analysis (see Young and Hermann 2014), the underlying idea is that conceptually complex individuals see their environment as more ambiguous, involving multiple tradeoffs, whereas conceptually simple individuals see a clearer picture of the problem (Ishiyama and Backstrom 2011). In particular, the usage of words and phrases like “probably”, “perhaps”, “tends to”, *etc.* suggests higher levels of cognitive complexity, whereas words like “always”, “definitely”, “absolutely”, *etc.* indicate lower levels of complexity in the thought process (*e.g.* Hermann 1980).

To apply this measure, we use the following procedure. First, for each company in our sample, we identify a chief executive officer (CEO) who headed the company after the crisis, *i.e.* she or he was the CEO in the first quarter of 2009. Next, we find the earliest year (excluding 2008), for which this CEO's communication in the earnings conference calls of the given company is available on Factiva. A result is a sample of 68 CEOs, 39 industrial companies and 29 IT companies. We collect all available earnings conference call transcripts for this year and combine all paragraphs that open with the CEO's name to estimate conceptual

²¹ Note that here we draw on the upper echelon literature (see Hambrick and Mason 1984). In particular, we consider that personal attributes of decision-makers will reflect upon organizational choices.

complexity in these decision-makers' communication in a given quarter. The resulting measure is likely to be a compound of the CEOs' innate cognitive complexity and temporal or contextual effects. Accordingly, we use a two-stage procedure where we first isolate the time-invariant component of decision-makers' cognitive complexity, and second, test our model predictions. Note also that given that our measure is at the CEO level, our observation period is bounded by either the quarter when a given CEO left the company or availability of data for all measures in the model, *i.e.* five years or 20 quarters past the crisis. In the interest of clarity, from here on we will use *conceptual complexity in communication* to denote the measure from the decision-maker's speech and innate *cognitive complexity* to refer to the time- and context-invariant component of one's complexity of thought.

4.4. Method

As mentioned above, an inherent property of rhetoric is that it is affected by situational factors. This means that organizational performance, size, competitive environment, and other contextual elements alter decision-makers' communication. To test our predictions, however, we need a metric that reflects cognitive complexity as a stable trait of a core decision-maker, rather than a time- or performance-contingent construct. The structure of our data allows us to isolate such a stable component and measure one's innate cognitive complexity. To this end, we first estimate the following model:

$$\text{conceptual complexity in communication}_{iq} = \alpha_0 + \mathbf{D}_i \boldsymbol{\beta} + \boldsymbol{\psi} \circ \boldsymbol{\Omega}_{iq} + \nu_{iq},$$

where the dependent variable is the measured conceptual complexity in decision-maker i 's communication in quarter q (prior to 2008); β is a vector of interest that corresponds to the coefficients (*i.e.* the CEO's innate cognitive complexity) estimated on CEO fixed effects (\mathbf{D}); Ω is a vector of control variables. For controls, Ω , we include the annual quarter dummies; logarithm of total assets, logarithm of quarterly revenues as proxies of firm size, average consumer confidence index in OECD countries as a proxy of general economic conditions, Herfindahl-Hirschmann index of market concentration computed for text-based network industry classification, total competitive similarity of the focal firm's product portfolio to other firms in the industry as measured from the 10-K product descriptions (see Hoberg and Phillips, 2016), and quarterly return on assets together with its increment to partially account for a possible reverse causality and a change in communication in response to the observed performance; finally, we control for CEO age.

In the second stage, we use the vector of the estimated coefficients $\hat{\beta}$, that is, executives' innate cognitive complexity, as a key independent variable in the following model:

$$ROA_{it} = \gamma_0 + \gamma_1 \cdot time_{it} + \gamma_2 \cdot \hat{\beta}_i \cdot time_{it} + \gamma_3 \cdot \hat{\beta}_i \cdot time_{it} \cdot sector_i + \lambda \circ \Phi_{it} + u_i + e_{it},$$

where $time$ denotes the number of quarters after the crisis (\ln); $\hat{\beta}_i \in \hat{\beta}$ is the decision-maker's innate cognitive complexity represented by standardized coefficients from the first-stage model, $sector$ is a dummy variable, such that $sector = 0$ for industrials and 1 for information technologies; Φ_{it} is a vector of control variables (see Tables 4.1 and 4.2). Finally, u_i is a vector of time-invariant firm fixed effects.

Table 4.1. Descriptive statistics and correlation table

	Mean	S.D.	Min.	Max.	1	2	3	4	5	6	7
1. Return on assets	0.033	0.020	-0.041	0.139	1						
2. ln(Time)	1.947	0.823	0.000	2.944	0.1591	1					
3. Interaction: Cognitive complexity, ln(Time)	0.191	2.151	-8.717	12.124	0.2404	0.0344	1				
4. Interaction: Cognitive complexity, Sector, ln(Time)	0.358	1.493	-3.410	12.124	0.0434	0.0863	0.7125	1			
5. Interaction: Sector, ln(Time)	0.798	1.094	0.000	2.944	0.2635	0.3107	0.2767	0.3168	1		
6. ln(Revenues)	7.513	1.196	5.038	10.629	-0.0987	0.0649	-0.7096	-0.5005	-0.2539	1	
7. ln(Total assets)	9.182	1.290	6.891	13.577	-0.295	0.0488	-0.7411	-0.4319	-0.1482	0.8905	1
8. Debt/Equity ratio	0.243	8.501	-232.126	88.489	-0.0434	0.0447	-0.0106	0.0045	0.0202	0.0176	0.0261
9. Consumer Confidence Index	98.365	0.629	96.860	99.550	0.0312	0.5118	0.0198	0.0216	0.1593	0.0182	0.0269
10. Decision maker's age	57.917	5.569	40.000	69.000	-0.003	0.1367	0.0352	-0.1671	-0.2009	0.1252	0.0748
11. Market competition (HHI)	0.207	0.175	0.047	1.000	-0.1585	0.0326	-0.1386	-0.0314	-0.2756	0.1037	0.0895
12. Competitive similarity	2.732	1.970	1.000	9.124	0.1923	0.1018	0.1174	0.1069	0.6993	-0.1209	-0.0255

	Mean	S.D.	Min.	Max.	8	9	10	11	12
8. Debt/Equity ratio	0.243	8.501	-232.126	88.489	1				
9. Consumer Confidence Index	98.365	0.629	96.860	99.550	0.019	1			
10. Decision maker's age	57.917	5.569	40.000	69.000	0.0069	0.0963	1		
11. Market competition (HHI)	0.207	0.175	0.047	1.000	0.0113	0.0561	0.1033	1	
12. Competitive similarity	2.732	1.970	1.000	9.124	-0.0124	0.0172	-0.0766	-0.4514	1

Notes. Correlations estimated for the observation period in stage 2. Categorical variables in model (5), i.e. industry and education level, not included. Extreme values of Debt/Equity ratio are due to significant amount of convertible notes on the balance sheet of "Linear Technology Corp." prior to 2011; descriptive statistics for Debt/Equity ratio excluding "Linear Technology Corp." are as follows: Mean = 0.298, S.D. = 0.3384, Min. = 0, Max. = 1.981.

Table 4.1 reports the summary statistics and pairwise correlations of the variables in the second stage of our analysis. Note that we observe significant negative correlation between the estimated innate cognitive complexity and firm's size as measured by both total assets and sales. Interestingly, this correlation is only present for the innate cognitive complexity estimated in the first stage; we report no such correlation between a CEO's conceptual complexity in communication and a firm's size. Here we mark this regularity, pointing to the possibility of negative association between one's innate propensity to think complexly and organizational size. It is also worth noting that the relationship between a firm's size and the conceptual complexity in communication in the first stage is positive but not significant. These observations need, of course, further analysis, but the intuition can be that individuals who are less cognitively complex may self-select into large businesses, whereas greater size of their companies increases the complexity of their thinking and communication. Besides this association, we also expectedly report high positive correlation between consumer confidence index and time since the bottom of the crisis as well as between total assets and sales.

To account for unobserved correlates of cognitive complexity at individual or organizational levels, we use fixed effects model as the main specification. Note that this also allows us to keep multicollinearity within commonly accepted limits. However, we further consider a random effects model that includes cognitive complexity as a standalone independent variable (see model 5 in Figure 4.2). We rely on this specification to see if there is evidence conflicting with our predictions. Consider, for example, that a case where the main effect of cognitive complexity is positive in a fast-paced setting but negative when the pace of decision-making is slow, while declining in both contexts, is inconsistent with our model predictions.

Table 4.2. Estimated models

Dependent variable: ROA	(1)	(2)	(3)	(4)	(5)
ln(Time)	0.00282 ^{***} 0.001	0.00246 ^{**} 0.001	0.00253 ^{**} 0.001	0.00180 [*] 0.001	0.00143 ^{***} 0.001
Interaction: Cognitive complexity, ln(Time)	0.00137 ^{***} 0.000	0.00204 ^{***} 0.001	0.00198 ^{***} 0.001	0.00260 ^{**} 0.001	0.00198 ^{***} 0.001
Interaction: Cognitive complexity, Sector, ln(Time)	8.50e-05 0.000	-0.00233 ^{***} 0.000	-0.00203 ^{***} 0.000	-0.00476 ^{***} 0.001	-0.00351 ^{**} 0.002
Interaction: Sector, ln(Time)	0.00240 0.002	0.00118 0.002	0.000553 0.002	0.00340 0.002	0.00326 ^{***} 0.001
ln(Revenues)		0.0707 ^{***} 0.005	0.0712 ^{***} 0.004	0.0761 ^{***} 0.005	0.0561 ^{***} 0.019
ln(Total assets)		-0.0642 ^{***} 0.007	-0.0633 ^{***} 0.007	-0.0668 ^{***} 0.008	-0.0476 ^{**} 0.023
Debt/Equity ratio		-2.68e-05 0.000	-2.50e-05 0.000	-2.87e-05 0.000	-3.48e-05 ^{***} 0.000
Decision maker's age		-0.000685 0.001	-0.00103 0.001	-0.000977 0.001	-8.73e-05 0.001
Consumer Confidence Index			0.000507 0.000	0.000871 ^{**} 0.000	-0.000334 0.001
Market competition (HHI)			0.00573 ^{**} 0.003	0.00579 0.004	0.00414 [*] 0.003
Competitive similarity			0.00127 0.001	0.00138 [*] 0.001	0.00108 0.001
Cognitive complexity					0.0146 0.009
Sector					0.0201 ^{***} 0.006
Interaction: Cognitive complexity, Sector					-0.00553 0.010
Education					-0.00264 0.008
Firm fixed effects	yes	yes	yes	yes	
Constant	0.0255 ^{***} 0.002	0.125 ^{***} 0.022	0.0793 ^{**} 0.036	0.0337 0.041	0.0672 0.101
R-squared	0.088	0.563	0.566	0.600	0.589
Only significant estimates from 1-stage (p-value < 0.1)				yes	yes
Robust standard errors					yes
Driscoll-Kraay standard errors	yes	yes	yes	yes	
Observations	923	923	923	721	721
Firms	65	65	65	51	51

Note. Standard errors below coefficients. R-squared within. In models 1 through 4 fixed effects absorb the main effect of cognitive complexity.

*** p<0.01, ** p<0.05, * p<0.1

4.5. Results and conclusion

4.5.1. Results

The pooled ordinary least squares regression in stage 1 is significant at 0.0001 level, adjusted R-squared is 0.314. Our observations from the second stage are summarized in Table 4.2. Specification 1 illustrates the model with our key variables only; specifications 2 and 3 progressively add organization-specific and environmental covariates; specification 4 replicates specification 3 but considers only significant estimates from the first stage; specification 5 extends specification 4 by lifting fixed effects and adding the main effect of cognitive complexity with time-invariant controls. Since there is evidence of both heteroscedasticity and autocorrelation (Baltagi-Wu LBI = 0.8757; Breusch-Pagan test for heteroscedasticity significant at 0.0001 level), in the fixed effects specifications we use Driscoll-Kraay standard errors, whereas in the random effects specification we cluster errors by the corresponding sector. We also consider that the estimates of cognitive complexity derived in the first stage vary in their accuracy. Accordingly, we also report our results for a subsample of estimates from the first stage that are significant at 0.1 level (see Table 4.2, specifications 4 and 5).

Overall, the data indicates that a relationship between decision-makers' innate cognitive complexity and performance dynamics is highly probable. In particular, in all specifications we see that the decision-makers' cognitive complexity moderates the performance dynamics and this moderation differs across contexts with varying clock speeds, *i.e.* the key interactions are at least marginally significant.

In the slow-clock-speed context (industrials), cognitive complexity positively moderates the relationship between performance and time. Specifically, after five years (recall that time represents the logarithm of the number of quarters since the bottom of the crisis) in the slow setting, one standard deviation in conceptual complexity adds more than 0.5 percent to the ROA.²² In the fast-clockspeed setting, however, after the same five years, conceptual complexity reduces ROA by approximately 0.1 percent. Note further that we observe a positive association between time and performance. This means that organizations exhibit behavior of intendedly rational agents. While a firm's performance does not always grow, all else constant, organizations show systematic tendency to improve performance over time (albeit at a declining rate).²³

These observations are consistent with our analytical model. Specification 5 (see Figure 4.2) further shows that considering the main effect of cognitive complexity does not contradict our predictions. In particular, this specification suggests that the effect of seeing greater complexity in the business problem leads to superior outcomes within the analyzed timeframe when the clock speed is slow. Inversely, when the clock speed is fast, cognitively complex decision-makers discover superior strategies in the short run but bear the cost of long-term underperformance. Recall that our observation period is five years. Accordingly, the distinction between short and long term is bound by this time span. These results roughly correspond to the case presented in panel (a) of Figure 4.2. That is, the implicit

²² Note also that ROA in the model is measured in absolute terms.

²³ It is important to note that this observation has no bearing on the dynamics of superior performance (cf. Waring 1996, Wiggins and Ruefli 2002). Rather it refers to the tendency of average returns to increase on a new landscape to a certain normal level, so long as the competitive forces can be considered constant.

correspondence between model time and real-world time, η_{slow} , is approximately 0.1 in slow-clock-speed industries and several times greater in fast-clock-speed settings (possibly, η_{fast} is around 1).

4.5.2. Conclusion

Our empirical analysis, therefore, indicates that time moderates the relationship between cognitive complexity and performance. Importantly, the data provides no substantive evidence against this moderation being in agreement with our model predictions. On the contrary, the patterns in the data align with those predicted by our theory. Hence, we conclude that the relationships brought to light in our analytical model are consequential for organizational search. The way decision-makers understand the business problem, not merely the accuracy but also the conception of the overall complexity, alters the dynamics of performance and hence the chances of attaining organizational aspirations (see essay II).

Essay IV. Cognitive Aliens

Talent hits a target no one else can hit; Genius hits a target no one else can see.

Arthur Schopenhauer

5.1. Introduction

Organizational strategies reflect their decision makers' understandings of reality (Hambrick and Mason 1984, Marcel, Barr, Duhaime 2011, Tripsas and Gavetti 2000). CEOs as well as other members of the top management team imprint their thoughts, ideas, and beliefs onto organizational acts (Nadkarni and Barr 2008, Nadkarni and Narayanan 2007). Greater managerial insight, from this perspective, improves decision-making and helps discover value-creating configurations or more efficient routines and capabilities, and thereby leads to superior performance (Gary and Wood 2011, Gavetti 2012, Helfat and Peteraf 2015). Yet, in reality, even outstanding organizations are not always headed by the greatest intellectuals among us.

Indeed, casual observation suggests that many of the most successful CEOs think in distinct and not always accurate ways. They connect unconnected and focus on business aspects that others consider unimportant or simply irrelevant (*e.g.* Dyer, Gregersen, Christensen 2009, Kanter 2011, Roger 2007). Their organizations welcome unconventional reasoning and communicate this attitude to their customers. A known advertising campaign from the late 90s—*Think different*—is a spotlight example of such a philosophy. Headed by late Steve Jobs—arguably, the most discussed icon of avant-garde understanding in contemporary management—Apple Inc. skyrocketed, powered by their explicit emphasis on the difference in thinking. Similarly, in a highly competitive industry of low-cost airlines,

Ryanair is unfailingly demonstrating outstanding business results. A correlate of this company's notable performance is Michael O'Leary's, Ryanair's CEO's, often-criticized behavior and an irregular business acumen. Another prominent example of a leader with unconventional vision, whose atypical reading of the organizational environment reflects in his peculiar manner, is Steve Ballmer. Like many other *cognitive aliens*—that is, people whose mental representations of reality are distinct from that shared by the many others—this eccentric CEO of Microsoft constructed an organization that persistently reports superior performance.²⁴

Puzzled by these observations (albeit recognizing that the relationship can be spurious), we ask a simple question of whether an odd understanding can in fact lead to a sustained competitive advantage. Note that by odd we mean distinct and objectively less accurate. To answer this question we draw on two adjacent fields of knowledge. One relates to the notion of organizational search (*e.g.* Cyert and March 1963, Levinthal 1997), the other centers on imperfections in decision makers' mental representations (*e.g.* Gary and Wood 2011, Daft and Weick 1984, Walsh 1995). Although these concepts—search and imperfect mental representation—belong to the same domain of bounded rationality and therefore are inherently linked, there is a stark dearth of research analyzing the interplay between them (Gavetti, Levinthal, Ocasio 2007). From the perspective of our question, however, both are essential. Accordingly, we construct a model of search wherein all agents misperceive the structure of

²⁴ We borrow the term from psychology, where children are considered cognitive aliens because their perception and information-processing profoundly differ from those of adults (see for example, Forgas, Burnham, Trimboli 1988).

interdependencies among organizational choices—that is, they err in understanding the underlying strategic problem—but some err differently.

We then show that competition gives an advantage to those whose understanding is distinct and even odd. In line with the preface to our work, cognitive aliens discover valuable configurations that others do not see. Isolated from competitors' attention *ex ante*, such configurations yield greater performance *ex post*. When a common understanding is shared by a sufficient number of competitors, in the long run such an advantage can be strong enough to outweigh the reduction in performance associated with lower accuracy. What is more, we explicitly show that inferior acumen is a form of cognitive distinctness. As a result, those whose judgment is similar but less accurate than that shared by others, can have superior performance in the long run. Grounded in human cognition, such performance differential persists in the face of both substitution and imitation. We, therefore, conclude that in the realm of boundedly rational decision making, where there are competitive forces, thinking odd can in fact underlie sustainable competitive advantage. Although quasi-monopolistic positions exist even in markets with rational agents (*e.g.* Besanko, Dranove, Shanley 1996, Hotelling 1929, Sraffa 1926), when decision makers err, it is those who err different that appropriate greater value in the long run. In the closing section, we discuss how these observations add to a broad spectrum of knowledge on how organizations come to discover superior value.

5.2. Bounded rationality

That human rationality is bounded is a generally accepted fact and a central tenet of the behavioral theory of the firm (*e.g.* Eisenhardt and Zbaracki 1992, Gavetti *et al.* 2007,

Kahneman 2003). Unlike hyper-rational firms in the neoclassical economics, real-world decision makers have cognitive constraints and therefore cannot optimize. Rather, organizations search and satisfice (Simon 1955, 1956).

5.2.1. Search. Evaluation of alternatives

Essentially a process of consecutive decision making, search comprises two primitive stages: generation and evaluation of alternatives (Knudsen and Levinthal 2007). While the former (*i.e.* generation of alternatives) describes sequential sampling from the space of possible options, the latter (*i.e.* evaluation of alternatives) characterizes how organizations decide whether a given alternative should be accepted.

Extensive research in organizational sciences focuses on the attributes of the first stage in the process of search. In particular, existing works point to a pattern, whereby organizations tend to generate alternatives within their immediate vicinity but sometimes discover distant courses of action (Levinthal 1997, Rosenkopf and Almeida 2003, Rosenkopf and Nerkar 2001, Stuart and Podolny 1996). This propensity to search locally is a principal known constraint that prevents organizations from deriving optimal strategies. Although global search improves long-term performance, whether seeing beyond the immediate vicinity of the current state is sufficient to warrant discovery of the optimal strategy requires further analysis of how organizations evaluate alternatives.

Although evaluation accounts for a substantial part of the decision makers' time (see Mintzberg, Raisinghani, Theoret 1976), it has received only cursory attention in the work on organizational adaptation. In fact, because of primary focus on generation, "the mechanism by

which [...] alternatives are evaluated are less clearly developed" (Knudsen and Levinthal 2007: 39). Specifically, in the existing analyses evaluation is either equated with blind trial and error or considered to mechanically inform of the underlying performance with some level of noise (*e.g.* Levinthal 1997, Winter, Cattani, Dorsch 2007). How evaluation occurs—that is, how decision makers derive the exact estimate of the potential increment in performance—remains largely exogenous to the process of search. In effect, evaluation tends to be reduced to a black box that somehow produces an answer as to whether a given alternative should be accepted or rejected.

Errors in the existing works tend to be consequent to a random (albeit possibly biased) noise, which is not consistent over time (*cf.* Martignoni, Menon, Siggelkow 2016). Specifically, inaccurate decisions occur without any pattern and do not reflect the noise idiosyncratic to a given choice (*see* Knudsen and Levinthal 2007). However, given no possibility of vicarious learning, a decision maker that mistakenly rejects an alternative is likely to do so again the next time the same alternative is generated. A firm that, for example, decides not to enter a certain market has its own subjective reasons which will guide its choice so long as the conditions remain unchanged. Such subjective motives behind the firm's choice characterize its understanding of the relevant factors. Consequently, errors in evaluation are not fully random but rather reflect the decision makers' understandings of reality (Gavetti and Levinthal 2000).

5.2.2. Mental representations

A vast array of works indicates that organizations interact with the environment through the prism of their decision makers' mental representations (*see for example* Barr, Stimpert, Huff

1992, Kaplan and Tripsas 2008, Tripsas and Gavetti 2000).²⁵ Broadly construed as knowledge structures about reality, such mental representations essentially reflect managers' understanding of "how the business environment works" (Gary and Wood 2011: 569). A basic attribute of these representations is that they provide a mental image of the cause-effect relationships among choices (Rehder 2003).

While existing research addresses the implications of mental representations from various perspectives, a common theme is that human understanding of reality is imperfect. In particular, the way individuals interpret information and form deliberate judgment is based on their *subjective understanding* of the relationships among relevant factors, rather than on the objective structure thereof (Csaszar and Levinthal 2016, Martignoni *et al.* 2016). This means that the mental image of interdependencies does not perfectly map onto the factual state of the world.

A natural product of errors in perceiving reality is heterogeneity of mental representations. Indeed, prior research provides ample support for the idea that decision makers often understand the same objective environment differently (*e.g.* Barr *et al.* 1992, Bourgeois, 1985, Helfat and Peteraf 2015, Miller, Burke, Glick 1998, Tripsas and Gavetti 2000). Further, various attributes of mental representations—such as completeness, complexity, *etc.*—have multifaceted implications for organizational and social outcomes (*e.g.* Calori, Johnson, Sarnin 1994, Nadkarni and Barr 2008, Nadkarni and Narayanan 2007).

²⁵ Note that prior works have used a variety of terms to describe the decision makers' mental image of reality, such as: mental models, cognitive schemata, cognitive maps, dominant logics, *etc.* (see for example Axelrod 1976, Bettis and Prahalad 1995, Eden 1992, Giola and Manz 1985).

Differences in decision makers' understanding of the strategic problem, therefore, are not only present but also consequential.

Presence of variance in mental representations, however, does not imply that all individuals form entirely unique ideas of what the reality is. On the contrary, there is also evidence that people often come to share similar beliefs and at least partially converge in their understandings of the organizational landscape (*e.g.* Daft and Weick 1984, Langfield-Smith, 1992, Porac, Thomas, Wilson, Paton, Kanfer 1995, Smircich and Stubbart 1985). Common cultural and educational background, repetitive interactions, and other antecedents of socialization create a push toward similar mental representations (*e.g.* Knight *et al.* 1999, Rentsch and Klimoski 2001).

5.2.3. Cognitive aliens: Premise

Existing literature, therefore, portrays a picture wherein individuals' mental representations are inaccurate and diverse, but at the same time convergent on certain aspects. That is, decision makers have partially shared imperfect understanding of organizational environment. While prior works tend to focus either on the attributes of what is different about mental representations (*e.g.* Nadkarni and Narayanan 2007) or on the implications of convergence in beliefs (*e.g.* Porac *et al.* 1995), our analysis centers on the deviation from what is common in the way most people understand reality, *i.e.* cognitive distinctness.

Whereas all individuals have their idiosyncratic image of reality, cognitively distinct are those whose interpretation can differ from what is commonly assumed. The premise is simple: if one's thought process is unlike that of others, he or she can mentally construct a

similar or a different representation. On the contrary, if one's cognition is akin to that of other people, his or her understanding, evaluation of alternatives, and ultimately decision making will tend to align with the standard. Note that mental representation is a reflection of how an individual thinks rather than the thought process per se (Eden 1992). Therefore, individuals who have uncommon mental representations of reality tend to reason in a way that differs from that shared by others. We dub these individuals cognitive aliens because, perhaps, such is the thinking that we can expect from the intelligence that has never interacted with ours. By definition, such a phenomenon can only occur when there are many individuals (*i.e.* number of players is greater or equal to three).

5.3. Competitive pressure

Where there are multiple players, there are competitive forces. While search and competition have been considered jointly before (see Adner, Csaszar, Zemsky 2014, Lenox, Rockart, Lewin 2006, 2007), existing analyses assume independence between the two. That is, the objective function of search, *i.e.* landscape, does not change in the presence of competition. The rationale is that such an objective function—albeit possibly complex itself—represents a simple input in the economic decision, such as cost, quality, *etc.*

If, on the contrary, the searched landscape is that of strategy as such, the shape of the objective function should sag under competitive pressure. The rationale behind this can be seen from both economic and management perspectives. In economic terms, if firms are exactly the same they find themselves in the situation of perfect competition, where economic rents are null (see Besanko *et al.* 1996). In management terms, firms with sufficiently close

strategies compete for the same resources (Barney 1986, Carroll 1985, Peteraf 1993). To the extent that resources are finite, such competition will erode performance of shared strategies (Baum and Mezias 1992, Baum and Singh 1994).

Present such diseconomies of having identical strategies, the objective function of search becomes dynamic. Specifically, the performance of a given configuration differs as a function of the number of players occupying it. As the following analysis illustrates, accounting for these effects may further our knowledge about organizational consequences of bounded rationality in general, and cognitive distinctness in particular.

5.4. Model

Following the neo-computational tradition in organizational sciences, we think of organizational strategy as an N -dimensional binary string, where every element can take on two values: zero and one (*e.g.* Posen, Lee, Yi 2013). Each of the 2^N possible configurations of such a strategy corresponds to a particular level of fit.²⁶ To define the mapping from configurations to fit, we use the technical apparatus of the NK model (Kauffman 1993). Detailed description of the mechanics of this model can be found elsewhere (see for example Ganco and Hoetker 2009). Here we focus on the main elements that are essential for the construction of the NK landscape (*i.e.* mapping from configurations to performance, and the mapping from mental representations to landscapes) and our analysis.

²⁶ Note that the following terms refer to the N -dimensional binary string of organizational strategy: configurations, alternatives, options, courses of action, states.

5.4.1. Objective landscape

A central property of the NK model is that it accounts for the structure of interdependencies among choices. Technically, this is accomplished by means of altering the input of every decision conditional on its relationship with other decisions. Specifically, the contribution function of every choice (element in the N -dimensional binary string) changes depending on (i) its own state and (ii) states of other choices, with which it interacts. The fit of the entire configuration is a simple average of the contribution functions of the N decisions.

The structure of relationships among organizational choices can be represented as an adjacency matrix, $\mathbf{A} \in \{0,1\}^{N \times N}$ (Ghemawat and Levinthal 2008). In such a matrix, every element a_{ij} indicates the relationship between choices i and j . Specifically, if $a_{ij} = 1$, the state of choice j affects the contribution function of choice i . That is, holding the remainder of the configuration constant, the input of decision i will differ when there is a change in decision j , even if the state of i itself remains unchanged. When $a_{ij} = 0$, the state of choice j has no effect on choice i , and the contribution function of decision i given that choice $j = 1$ is the same as that when choice $j = 0$, the remainder of the configuration being constant. Thus, the fitness landscape is fully determined by the adjacency matrix \mathbf{A} and the matrix of contribution functions, $\mathbf{C} \in \mathbb{R}^{2^N \times N}$:

$$\xi(\mathbf{f}) = \frac{\sum_{i=1}^N c_{\varphi(\mathbf{f} \circ \mathbf{a}_i) i}}{N}, \quad (5.1)$$

where $\xi(\mathbf{f})$ denotes fit of strategy $\mathbf{f} \in \{0,1\}^{1 \times N}$, function $\varphi(\mathbf{f})$ provides one-to-one mapping from configurations to row index in \mathbf{C} , $c_{ki} \sim U(0,1)$ —an element in \mathbf{C} —is a contribution

function of decision i given configuration $\varphi^{-1}(k)$, \mathbf{a}_i is an i 's row in \mathbf{A} , and \circ stands for Hadamard product. These technical elements allow for an explicit definition of cognitive distinctness.

5.4.2. Cognitive aliens: A formal definition

Objective landscape characterizes the actual mapping from configurations to fit. An agent whose understanding is flawless will act on precisely this true landscape. However, if an agent misperceives any element in \mathbf{A} , its actions will be guided by a different landscape. When $\tilde{a}_{ij} \neq a_{ij}$ —where \tilde{a}_{ij} is the perceived relationship between choices i and j —the decision maker will either neglect a relevant change in the contribution function of decision i (if $a_{ij} = 1$) or conceive an alteration that does not in fact exist (if $a_{ij} = 0$). As a result, for some configurations, perceived fit $\tilde{\xi}(\mathbf{f})$ will differ from the true fit $\xi(\mathbf{f})$.²⁷

What decision makers think about interdependencies, therefore, affects how they perceive the landscape. Organizations with the same understanding of the relationships among choices see the landscape similarly and vice versa. In other words, distinct mental representations reflect that individuals connect choices in unusual ways, employ different decision rules and essentially think different (Gary and Wood 2011). This basic property allows a parsimonious formal definition of cognitive aliens, or, in broad terms, those who think unlike everybody else.

²⁷ Perceived fit is calculated using equation (5.1), where \mathbf{a}_i is replaced with agent's perception of each decision's interdependencies, $\tilde{\mathbf{a}}_i$.

Given \mathbf{A} , \mathbf{C} , and set of agents $S = \{1, 2, \dots\}$, agent $w \neq 1$ is a cognitive alien if:

$$\begin{cases} p(\tilde{a}_{ij,s} = a_{ij,s}) = \theta, & \forall i, j, s \\ \exists i \text{ s.t. } \tilde{a}_{ij,s} = \tilde{a}_{ij,1}, & \forall j, s \neq w \end{cases}, \quad (5.2)$$

where $\tilde{a}_{ij,s}$ is agent s 's perceived interdependence between choices i and j , $s \in S$, and $\theta \in [0,1]$ limits the probability that a given interdependence is perceived accurately. What this definition means is that for a fragment of a strategic problem all but one share the same θ -accurate mental representation, whereas that of cognitive alien is also θ -accurate but not necessarily the same. In other words, cognitive aliens tend to have equally good but in some respect unique understanding. The suggested definition implies a strong form of uniqueness in that cognitive alien is lone. A weaker form would imply presence of unequally-sized and differently thinking groups. In the interest of clarity and without loss of generality, however, we consider the strong form of cognitive alien defined by condition (5.2).

Figure 5.1. Mental representations: An example

$$\begin{array}{cccc} \begin{pmatrix} 1 & 0 & \dots & 0 & 0 \\ 1 & 1 & \dots & 1 & 1 \\ \vdots & \ddots & \ddots & \vdots & \vdots \\ 1 & 0 & \dots & 1 & 1 \\ 1 & 0 & \dots & 0 & 1 \end{pmatrix} & \begin{pmatrix} 1 & 0 & \dots & 0 & \mathbf{1} \\ 1 & 1 & \dots & 1 & 1 \\ \vdots & \ddots & \ddots & \vdots & \vdots \\ \mathbf{0} & \mathbf{1} & \dots & 1 & 1 \\ 1 & 0 & \dots & 0 & 1 \end{pmatrix} & \begin{pmatrix} 1 & 0 & \dots & 0 & 0 \\ \mathbf{0} & 1 & \dots & 1 & 1 \\ \vdots & \ddots & \ddots & \vdots & \vdots \\ \mathbf{0} & \mathbf{1} & \dots & 1 & 1 \\ 1 & 0 & \dots & 0 & 1 \end{pmatrix} & \begin{pmatrix} 1 & \mathbf{1} & \dots & 0 & 0 \\ 1 & 1 & \dots & 1 & 1 \\ \vdots & \ddots & \ddots & \vdots & \vdots \\ \mathbf{0} & 0 & \dots & 1 & \mathbf{0} \\ 1 & 0 & \dots & 0 & 1 \end{pmatrix} \\ \mathbf{A} & \text{Perceived 1} & \text{Perceived 2} & \text{Perceived 3} \end{array}$$

Figure 5.1 illustrates an example of how cognitive distinctness can manifest itself in different mental representations. Given \mathbf{A} —an accurate representation—Perceived 1, Perceived 2, and Perceived 3 are equally good: all have three errors (highlighted in bold). As

Perceived 1 and Perceived 2 share the same understanding of decision $i = N - 1$ (recall that $\mathbf{A} \in \{0,1\}^{N \times N}$) whereas Perceived 3 has a different representation of this choice's dependencies, the latter decision maker is a revealed cognitive alien. Note, however, that our definition allows that occasionally cognitive aliens derive the same mental representation as that shared by everybody else.

Imperfect ability to specify the structure of interdependencies (*i.e.* $\theta < 1$) is necessary to observe cognitive distinctness (note that otherwise $\tilde{a}_{ij,s} = a_{ij}, \forall i, j, s$). Importantly however, regardless of how agents perceive the structure of interdependencies, the underlying fit remains unchanged. Accordingly, while organizations act on their subjective understanding, they all—cognitive aliens or not—experience the true, objective reality. As a result, search is an interplay between cognition and performance feedback.

5.4.3. Performance and search

To capture competitive pressure, performance in our model is a function of fit of a given strategy and the number of organizations sharing this exact configuration.²⁸ Specifically, performance is strictly increasing in fit and monotonically decreasing in the number of clones. While there is an infinite set of functions that can be used to formalize these relationships, without loss of generality we assume the following:

²⁸ Note that with a few exceptions (*e.g.* Lenox *et al.* 2006, 2007) prior works essentially equate the notions of fit and performance. While we agree that such an equation is often justified, we believe that in the context of our study, it is practical to separate the two concepts. Consider that fit characterizes the extent to which organizational strategy aligns with the market conditions other than strategies of other players, whereas performance integrates all.

$$\pi(\mathbf{f}) = (1-\gamma)^r \cdot \xi(\mathbf{f}), \quad (5.3)$$

where $\pi(\mathbf{f})$ denotes performance of strategy \mathbf{f} given r , the number of *other* players with strategy \mathbf{f} , and $\gamma \in [0,1]$ sets competitive pressure, so that when $\gamma = 0$ performance stays constant regardless of the number of competitors with strategy \mathbf{f} , whereas when $\gamma > 0$ performance deflates when multiple competitors share the same strategy. Since the nature of competitive pressure can vary (*e.g.* use of the same resources, contraction of monopolistic rents, *etc.*), we parametrize the discount of having identical strategies, rather than model its antecedents. This further allows for a more general analysis as industries may well differ in the magnitude of γ .

Whatever the environment, however, organizations search for strategies that lead to better performance. As for our analysis the process of generation is only tangential (note that the effect of cognitive distinctness is qualitatively unchanged for both local and global search), we impose no restrictions on how organizations derive alternatives. This assumption is not technically driven, rather it emphasizes a key contribution of our analysis. In particular, by allowing distant search we abstract away from recombination constraints (organizations can change multiple decisions simultaneously), and focus on the cognitive antecedents of how organizations decide whether an alternative should be accepted.

The basic premise behind evaluation in our model is that organizations make judgments based on their subjective understanding of the landscape. In every period an arbitrary organization compares a random alternative to its current state, following a basic rule: accept if the alternative's perceived performance is greater than the actual performance in the status quo, reject otherwise. If upon acceptance, it turns out that the alternative leads to

inferior performance, the organization returns to the previous state. This means that an alternative $\mathbf{f}_{al.}$ becomes a new status quo if it in fact performs better than the previous status quo $\mathbf{f}_{s.q.}$, *i.e.* if $\tilde{\pi}(\mathbf{f}_{al.}) > \pi(\mathbf{f}_{s.q.})$ and $\pi(\mathbf{f}_{al.}) > \pi(\mathbf{f}_{s.q.})$, $\tilde{\pi}$ denotes perceived performance. Organizations, therefore, are intendedly rational in that they do not quit superior strategies (see Simon 1957).

5.5. Analysis


Unlike random noise in evaluation, inaccuracies in defining the structure of interdependencies produce systematic errors. An important property of such errors is that they are idiosyncratic to alternatives. That is, perceived fit of a given alternative is consistently under- or overestimated. This, in turn, leads to a consistent variance in actions determined by discrepancies in understanding of the links among organizational choices, and ultimately by the way people think (Martignoni *et al.* 2015).

In the following computational experiments we examine the consequences of these effects. Unless otherwise specified we set $N = 8$, total number of interdependencies equal to 28 (*i.e.* $K \approx 3.5$), number of competitors $|S| = 5$, $\theta = 0.75$, and average the results over 10^5 runs. For computational reasons, we use these parameter levels as the base case; in the robustness section we consider alternative specifications. Note also that following the previous work to facilitate comparison, we standardize fit on all landscapes between zero and one. To define the basic properties of search present imperfect mental representations, we begin by discussing general consequences of misperceiving the structure of interdependencies.

5.5.1. Terminal strategies

Systematic, alternative-specific errors in evaluation (that occur when \mathbf{A} is misperceived) may preclude organizational discovery of superior strategies even when there are no constraints on the distance of generated alternatives, *i.e.* if search is global. To see why, consider that if an organizations discovers $\mathbf{f}_{p,max}$ such that $\tilde{\xi}(\mathbf{f}_{p,max}) \geq \tilde{\xi}(\mathbf{f})$ for all \mathbf{f} , search discontinues and as long as there is an \mathbf{f} for which $\xi(\mathbf{f}_{p,max}) < \xi(\mathbf{f})$, there is a strategy or strategies that can lead to better performance. A close analogy of such mentally inaccessible strategies is that of "blind spots" (see Zajac and Bazerman, 1991). Figure 5.2 shows typical cases of the correspondence between perceived and actual fit, and highlights possible terminal strategies.

Figure 5.2. Terminal strategies: An example



Note. Each panel depicts a single simulation, where x axis is $\tilde{\xi}(\mathbf{f})$ and y axis is $\xi(\mathbf{f})$.

More generally, search terminates when organization discovers either $\mathbf{f}_{p,max}$ or an $\mathbf{f}_{terminal}$ such that $\xi(\mathbf{f}_{terminal}) > \xi(\mathbf{f})$ for all \mathbf{f} for which $\tilde{\xi}(\mathbf{f}) > \tilde{\xi}(\mathbf{f}_{terminal})$. In Figure 5.2 $\mathbf{f}_{p,max}$

and $\mathbf{f}_{\text{terminal}}$ are illustrated by filled black and empty circles respectively. Note that as θ increases, correlation between \tilde{a}_{ij} and a_{ij} goes up, the perceived fit becomes more descriptive of the underlying landscape, and $\xi(\mathbf{f}_{\text{p.max}}) \rightarrow \max_{\mathbf{f}} \xi(\mathbf{f})$. In the limit, when organizations perceive the structure of interdependencies perfectly, perceived maximum is equivalent to the true maximum. For any $\theta < 1$, however, the actual fit at the perceived global peak can take on any value between zero and one. Importantly, this value as well as the underlying configuration vary depending on the way the organization perceives the structure of interdependencies. That is, two agents with unlike mental representations will tend to terminate search with different strategies.

5.5.2. Performance differential

The fact that there are non-unique (and idiosyncratic to mental representation) terminal strategies implies that performance will vary depending on the initial conditions and chance. Yet, given a mental representation, number of competitors and the magnitude of competitive pressure, the expected performance is definite. Consistent with existing experimental studies, the pattern is that better understanding leads to better outcomes (Gary and Wood 2011). However, we further document that holding accuracy constant, cognitive distinctness supports superior performance.

Figure 5.3. Cognitive alien's advantage

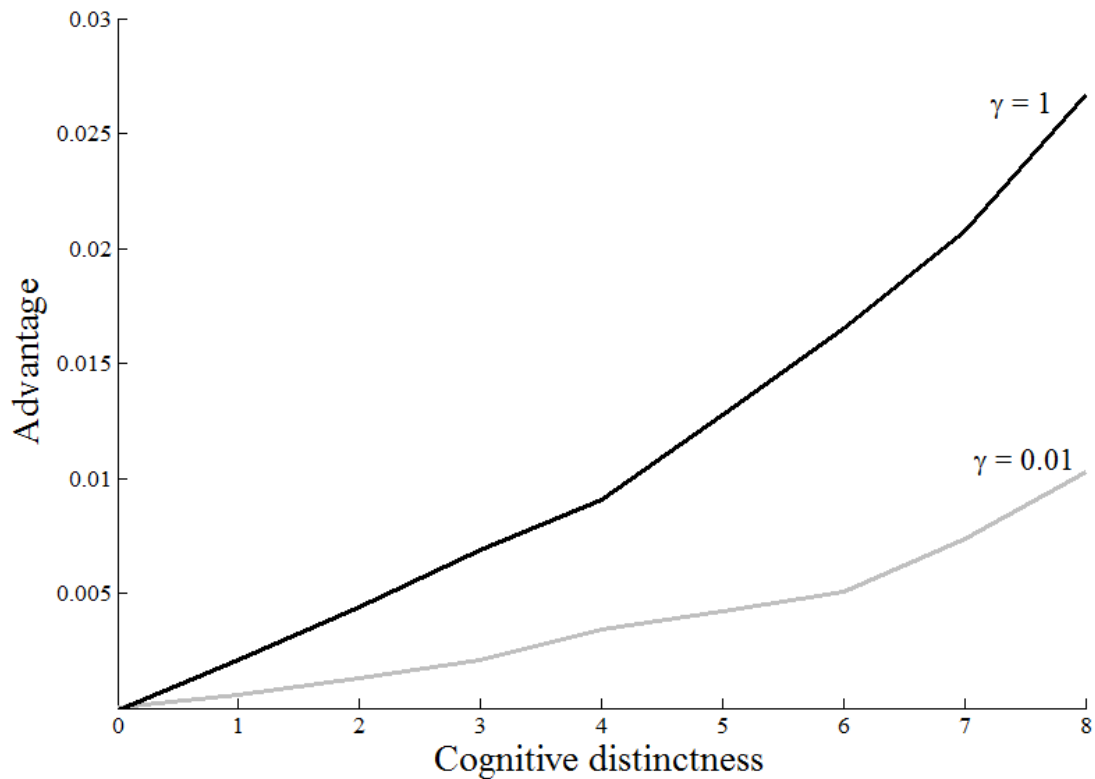


Figure 5.3 shows that this regularity holds for both low and high magnitude of competitive pressure, γ (unless specified otherwise, to focus on the difference of the kind rather than degree, in the following analyses we take $\gamma = 1$). It further illustrates that the more organizational understanding misaligns with that shared by competitors, the greater the advantage. We measure cognitive distinctness (abscissa) as the number of decisions for which cognitive alien's understanding differs from that shared by others. That is, we take the number of decisions for which condition (5.2) holds (unless specified otherwise, in the following analyses we take this number equal to N). Advantage (ordinate) represents the steady-state (after 10^4 periods) difference between performance of the cognitive alien and the mean

performance of others. Note further that while Figure 5.3 reports the long-run advantage of cognitive distinctness, short-term observations are qualitatively similar (the shorter the time, the lower the magnitude of advantage).

Thinking unlike everybody else, therefore, can indeed lead to better performance. Decision makers with unique—and from the standpoint of the mainstream, eccentric—understanding of the structure of interdependencies alter organizational search in a way that allows for systematic superior value creation.

5.5.3. Mechanism

To see the mechanism behind the reported performance differential, consider a simplified case where there are only three players who ignore any feedback from the objective reality. In other words, their decision rule is reduced to the following: accept if $\tilde{\pi}(\mathbf{f}_{al.}) > \tilde{\pi}(\mathbf{f})$, reject otherwise.

Given such reduced evaluation, organization's first best is to discover $\mathbf{f}_{p.max}$. As a result, for agents with identical mental representations there is an implicit race to find $\mathbf{f}_{p.max}$. The second to do so will have to be content with a suboptimal strategy (recall that this applies only to the simplified case where competitive pressure is high and search ignores any feedback from the objective reality). Contrary to those who share a similar mental representation, cognitive aliens are not necessarily participants in this race to a certain thought-to-be superior strategy. Insofar as terminal strategies are idiosyncratic to mental representations, decision makers with distinct understanding will configure organizations differently. Importantly, so

long as their accuracy on average is the same, cognitive aliens will tend to search for unique but equally fit strategies.

Figure 5.4. Possible realizations in the simplified case





Figure 5.4 illustrates all possible combinations of matches and mismatches between three thought-to-be superior strategies perceived by two players with shared understanding and those of the cognitive alien. For example, symbol $\begin{matrix} \blacksquare & \square \\ \blacksquare & \square \\ \blacksquare & \square \end{matrix}$ means that $\mathbf{f}_{p,\max}$ is the same for all agents, but with respect to second- and third-best strategies cognitive alien holds a different idea of where they are. Accordingly, *ex ante* an agent's private probability that it will be the first one to reach $\mathbf{f}_{p,\max}$ is the same for all, *i.e.* 1/3. However, cognitive alien's private probability (columns denoted p_w) that it will terminate at its second best is 2/3, whereas that for the other two players (columns denoted $p_{w \neq 3}$) is 1/3. Since the expected fit of strategies with the same perceived ranking (columns denoted ξ) is asymptotically equivalent for all, difference in private probabilities results in a performance differential. The average of performance differentials for each combination weighted by the respective probability (indicated under each combination) is exactly the magnitude of the overall advantage.

While this details the calculus behind the reported performance differential, the intuition is that cognitive aliens discover strategies that have the same fit but at the same time are less contested. A similar rationale applies to the case when organizations learn from the objective reality, although the explanation is compounded by the presence of terminal strategies other than $\mathbf{f}_{p,\max}$. The ability to think in a way that differs from that of the majority, therefore, allows organizations to detect value concealed from others by their common mental boundaries.

5.5.4. Lower accuracy


So far, we have relied on the assumption that cognitive aliens tend to hold a distinct but equally accurate mental representation. However, existing works find that consensus is often indicative of better judgment and combined opinion of many is highly descriptive of reality (*e.g.* Ashton 1985, Bourgeois 1980, Galton 1907, Yi, Steyvers, Lee, and Dry 2012). An implication is that those who think independently—and therefore, on average differently—may underutilize shared knowledge and thereby construct less accurate mental representations. Consider, for example, that while a disciplined education (MBA, Ph.D., *etc.*) improves reasoning, it also likely imprints common analytical patterns that may lead to convergence in understanding. Here we explicitly model such a tradeoff between cognitive distinctness and accuracy.

Figure 5.5. Lower accuracy: dynamics



We first observe that time and presence of competitive pressure are critical in resolving this tradeoff. Figure 5.5 shows that in the absence of diseconomies of having identical strategies, lower accuracy results in strictly inferior performance (note that in this case, $\xi \equiv \pi$). However, present competition, the benefit of cognitive distinctness may outweigh the cost of inferior understanding in the long run, when performance reaches the steady state ($T = 10^4$).

Figure 5.6. Lower accuracy: Number of competitors



Lesser accuracy of mental representations reduces the pace of performance improvement (because more of the potentially better strategies become concealed by systematic errors in evaluation) but at the same time furthers distinctness of cognitive alien's understanding. The latter effect expands the long-term dividend to thinking different, which in turn can fuel advantage if competition is strong. Figure 5.6 shows that for a given level of accuracy of the cognitive alien θ_w , a sufficient number of same-minded competitors flips the performance differential from negative to positive.

5.5.5. Low accuracy as cognitive distinctness

An important tenet in the existing works on behavioral strategy is that decision makers with better mental representations of organizational environment attain greater performance (Gary and Wood 2011, Gavetti 2012, Gavetti and Levinthal 2000). Yet, the above analyses suggest that an exact inverse of this statement can also be true. Specifically, an inferior but unique mental representation can lead to superior performance. The main factor behind this effect is cognitive distinctness—that is, an ability to think different. But this is not yet to say that inferior acumen per se can improve performance. In reality, those who have unlike mental representations do not necessarily think different, they may simply have inferior knowledge. Accordingly, here we analyze if lower accuracy can result in better performance even when the decision maker's thought process is the same (albeit less correct). Such a decision maker misspecifies the same interdependencies as others, plus some additional ones. [continued on the next page]

Figure 5.7. Lower accuracy: Error as distinctness

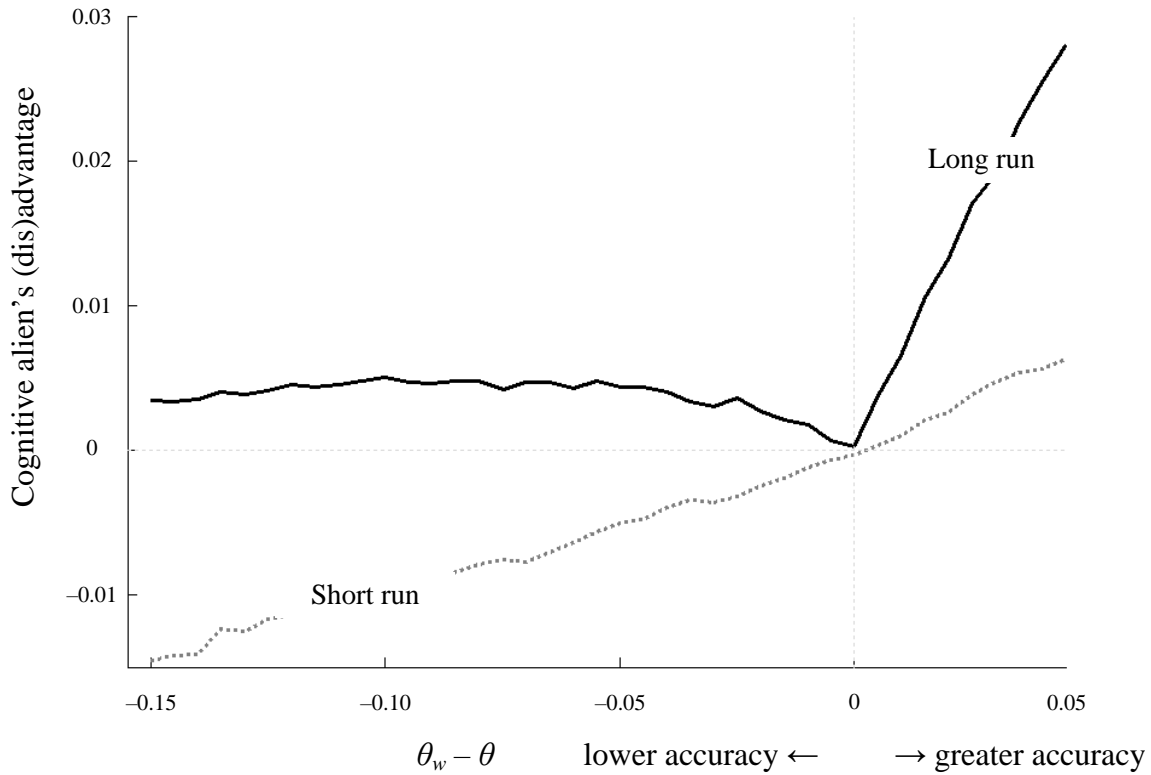


Figure 5.7 illustrates that when multiple competitors share a common understanding of the business environment (note: number of competitors is 10), in the long run the effect of lower accuracy is in fact similar to that of greater accuracy. Despite that errors in specifying interdependencies reduce the expected level of fit, they steer organizations toward less contested regions on the organizational landscape, thereby protecting performance from the erosive effects of competition. Note that in this analysis we assume no cognitive distinctness (save lower accuracy) in that there are no mistakes that the more accurate have and the less accurate does not. When $\theta_w < \theta$, this holds for the focal firm; when $\theta < \theta_w$, this holds for the majority.

An important implication of this effect is that there may be an implicit tradeoff between greater accuracy of one's understanding and lower intensity of competition. To the extent that one's acumen derives from education, socialization, *etc.*, an increase in accuracy is likely to reflect convergence to a frontier representation inherent in the accumulated social knowledge. By investing in learning from the shared knowledge, individuals bring their mental representations closer to those of other learned minds and thereby engage in a competition on a similarly misperceived landscape. In contrast, decision makers who draw less from the shared knowledge preserve uniqueness of their mental representations, thereby on average see inferior options but at the same time experience less competition.²⁹ Insofar as organizations headed by such individuals can live through the period of initial underperformance, in the long run they discover strategies that on average outperform those of more accurate decision makers (see Figure 5.5).

5.5.6. Imitation

Our analysis indicates that cognitive distinctness—even as an artifact of lesser acumen—may allow for creation of value that is not fully understood by others. A question is whether competitors can appropriate this value without knowing its origin, *i.e.* by imitating the observed attributes. The answer is they can, but only partially.

Imitation and substitution are the two main forces that cause erosion of superior performance (Dierickx and Cool 1989). While accurate global search implies both substitution

²⁹ Note that without competitive forces, on average such individuals would underperform compared to those who have more accurate understanding.

(as organizations generate new alternatives) and imitation (as organizations also generate alternatives already used by others), allowing for such a search in the presence of misperceptions does not fully capture the force of imitation. Consider that if a firm exhibits higher performance than another one, the latter can copy all observable attributes of the former even if it thinks that such a move should worsen its state. The premise is that organizations have greater trust in what they observe than in what they believe. Accordingly, we allow for such an imitation by assuming that if there is a firm with strategy \mathbf{f} , then $\tilde{\pi}(\mathbf{f}) = \pi(\mathbf{f})$ for all agents, that is, all agents can learn vicariously. By doing so, we effectively abstract away from imperfections in imitation, which have been argued to both obstruct and accelerate performance improvement (Posen *et al.* 2013, Rivkin 2000). We find that these effects naturally deflate the advantage of cognitive alien but do not eliminate it.

Figure 5.8. Effect of imitation

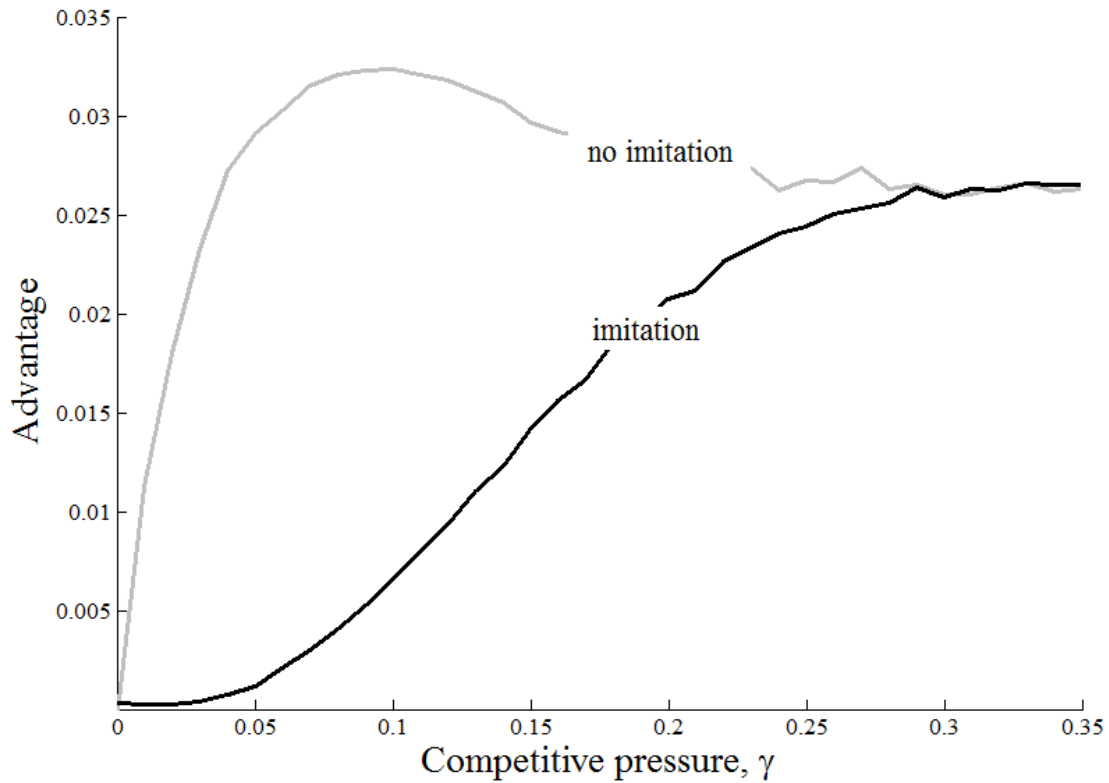


Figure 5.8 shows that when competitive pressure γ is low, competitors appropriate a large part of the value discovered by the cognitive alien: advantage with imitation is substantially lower than that without imitation. However, as γ increases, imitation becomes uneconomical and cognitive aliens benefit from the full possible advantage: performance differentials with and without imitation are near equivalent.³⁰ For any $\gamma > 0$, however, superior

³⁰ Note that the mild curvilinearity in the relationship between competitive pressure γ and the magnitude of advantage is a consequence of "greedy" behavior of laggards. When γ is not high there occur instances where laggards act in their own interest but at the expense of industry performance, hence inflating relative advantage of cognitive aliens. To see this, consider the following example: $\pi(\mathbf{f}_{s.q.}) = 0.8$, $\xi(\mathbf{f}_{al.}) = \bar{\xi}(\mathbf{f}_{al.}) = 0.86$, $\gamma = 0.04$, and

strategies are at least partially inimitable because there are instances when copying the leader is against the follower's interest (see Pacheco-de-Almeida and Zemsky 2007).

5.5.7. Robustness

An unusual feature of our work is that we use NK model without explicit interest in the underlying complexity. Rather, we take advantage of the model's power in providing a mapping from a set of possible adjacency matrixes—or mental representations—to a set of landscapes. When constructing such a mapping we use several methodological assumptions. Here we discuss robustness of our results to changes in these assumptions.

First, to keep the time for the population of firms to reach the steady state within reasonable limits (10^4 periods) we took $N = 8$. We find that the effect of cognitive distinctness remains qualitatively similar for higher ($N = 10$) and lower ($N = 6$) dimensionality of organizational space. Second, we assumed that the objective landscape is based on an \mathbf{A} with 28 interdependencies off the main diagonal. The rationale is that in this case agents are equally likely to make errors of omission and errors of commission when specifying the structure of interdependencies. Performance of cognitive distinctness, however, remains qualitatively unchanged for any number of interdependencies in \mathbf{A} (specifically, we tested $K = 0$ and $K = N - 1$). Third, we considered \mathbf{C} —a matrix of real and perceived contribution functions—to be the same for all agents in a given run. The rationale is that c_{ij} is a constituent of understanding similar to a_{ij} : if agents differ in the way they perceive c_{ij} , they also differ in the way they see

there is already one competitor at $\mathbf{f}_{al.}$, which means that $\pi(\mathbf{f}_{al.}) = \tilde{\pi}(\mathbf{f}_{al.}) = 0.82$. In this case the organization will accept the alternative, thereby causing the average performance of same-minded competitors to go down from 0.83 to 0.82. When on the contrary γ is sufficiently high, such behavior does not happen.

the landscape (see Ethiraj and Levinthal, 2009). Yet this is not to say that agents must have the exact same estimates of c_{ij} . Accordingly, we add agent-specific noise in the estimation of contribution functions such that $c_{ij,s} = c_{ij} + \varepsilon_{ij,s}$, where $\varepsilon_{ij,s} \sim N(0, 0.25)$ is noise, and $c_{ij,s}$ denotes agent s 's estimate of c_{ij} . Although variance of 0.25 is extremely high—recall that $c_{ij} \sim U(0, 1)$ —we nevertheless observe a similar advantage of cognitive distinctness.

5.6. Conclusion and discussion

This paper studies performance outcomes of seeing the reality in an unusual or even odd way. We construe such cognitive distinctness as an aptitude to form mental representations that differ from those shared by majority. Present competitive forces that give rise to diseconomies of having identical strategies, this trait can support superior performance. Such an advantage persists despite the erosive effects of competitors' efforts to substitute or imitate the strategy of the cognitive alien, a firm with an unlike understanding of the strategic problem.

A core premise in our analysis is that extreme similarity of strategies erodes performance. While a sizeable body of work provides ample support for this principle (*e.g.* Baum and Mezias 1992, Baum and Singh 1994, Hotelling 1929), we point to an important and starkly understudied antecedent of a value-creating difference. Uncommon errors in perceiving the structure of interdependencies can steer organizations toward less contested regions on the fitness landscape. Without such errors—*holding recombination capability constant*—superior performance is a mere outcome of luck.

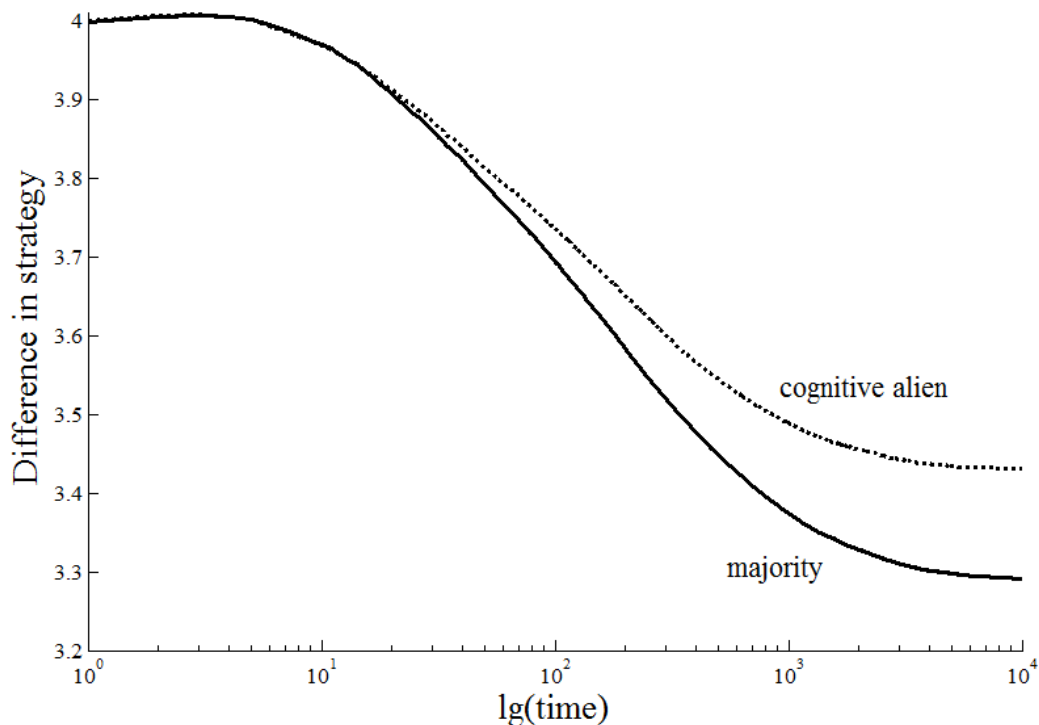
These observations speak directly to the emerging field of behavioral strategy (see Gavetti 2012, Levinthal 2011, Powell, Lovallo, Fox 2011). Broadly defined, behavioral strategy seeks to uncover psychological roots of superior performance (Gavetti 2012). A central construct in this view is cognitive distance. Organizations that manage to see opportunities beyond what is immediately available should tend to earn greater profits. In keeping with this argument we show that cognitive aliens discover strategies that are *ex ante* concealed from others by their common mental boundaries. Importantly, however, we demonstrate that to perceive such strategies organizations do not need to have a better representation—in fact they can have a worse one—but they do need to have a representation that is distinct.

Mental distinctness is a central element in our analysis. Aligned with recent work in behavioral strategy, we show that imperfections in understanding reality can lead to superior performance (see Csaszar and Levinthal 2016, Martignoni *et al.* 2016). Different from these analyses, however, we show that an important key to furthering our knowledge about the search outcomes—without a limiting assumption of strictly or nearly local search—is integrating this process with the socio-cognitive aspects.

From this perspective, our analyses can also be of interest for the institutional theory (see for example DiMaggio and Powell 1983, Scott 1987). Figure 5.9 shows that our model—despite explicitly assuming diseconomies of similarity—does not negate isomorphism. In fact, we observe that over time all organizations, including cognitive aliens, become increasingly alike. In Figure 5.9, difference in strategy is measured as the average Hamming distance between a given firm and all other players. As organizations search, certain practices become

common. Those, however, who have distinct mental representations do not necessarily adopt all of the practices widely perceived to be superior. While present institutional forces the efficient degree of strategic difference will be lower for all organizations, our analysis indicates that for those who think different the balance between distinctness and similarity will be shifted toward the former (see Deephouse 1999, Uzzi, Mukherjee, Stringer, Jones 2013).

Figure 5.9. Isomorphism



This, however, does not mean that organizations should seek to skew the balance between strategic difference and strategic similarity. Observe that given one's mental representation, attempting to be different—*i.e.* by constraining generation and evaluation of alternatives—will impede discovery of superior strategies. Firm's behavioral strategy, therefore, should consist in finding and promoting individuals with unusual understanding of

the business environment. Organizations with better capabilities in doing so will create value inaccessible to others because of their common mental boundaries rather than economic constraints.

Thus, in line with recent modeling approaches (see for example Lenox *et al.* 2006, 2007), our work further highlights the use of integrating the primitives of bounded rationality with fundamental economic laws. Although the very idea of bounded rationality has an inherent antilogy with neoclassical decision making, it is not incompatible with the basic principles of economic interaction.

Another broad stream of research that may potentially relate to our analysis explores individual antecedents of creativity and innovation (see for example Amabile 1983, Mumford 2003). Specifically, interested researchers may draw parallels between the key constructs in this paper and the main notions in the work on creativity. Consider that our analytical approach may allow a clear formal separation of novelty (uniqueness of configuration \mathbf{f}), usefulness (the underlying fit ζ) and performance (π as a function of both fit and uniqueness). From this angle, our analysis provides a formal link between idiosyncrasy of cognition and performance. Specifically, we explicitly show why creation of superior value requires deviation from common knowledge structures (see Mumford and Gustafson 1988). Different from the existing studies, however, our analysis focuses not so much on the attributes of mental representations, but rather on the frictions that exist between one's understanding and reality. This allows us to show that lower accuracy can in fact be instrumental to better performance. Overall, our work—albeit also exploring the consequences of unlike reasoning—centers not on the ability to recombine different elements or divergent thinking per

se (organizations can generate any idea) but rather on the implications of cognitive distinctness in evaluation of alternatives.

Decision makers with odd, eccentric and—from the standpoint of the same-minded majority—absurd understanding of organizational reality evaluate alternatives differently. Although they also err and perhaps even more often than everybody else, they see value where others see none. And it could be that it is the very inaccuracy of their understanding that leads them to discover superior configurations. Because of this, when competitors are many, it is those who think different, *err different* and have in fact a less accurate idea of what the real is, that guide organizations toward strategies for a long-term sustainable competitive advantage.

General Conclusion

6.1. Overview and contribution

This dissertation started with an aim of discovering value-creating patterns in the process of search that arise from the idiosyncrasies of the decision makers' thinking. A fundamental premise it stands on is that decision makers' rationality is bounded, and the bounds to rationality are not identical for all decision makers.

By recognizing this basic premise, the present work suggests that *it may be rational* for organizations *to understand and use cognitive idiosyncrasies* of their decision makers. As we live in the world of the second best, where we do not have access to the entirety of information and available options, managing the imperfect can outperform striving for (NB: not having) the perfect. To see this, consider that if improving decision makers' judgment is costly (*e.g.* the cost of raising the overall accuracy of decision makers' mental models), there can be a point at which further efforts to enhance decision quality are less expedient than managing decision makers' cognition (*e.g.* identifying and matching decision maker's cognitive style to a certain organizational task). Thus, an important general corollary of this work is a tradeoff between investing in decision makers' ability and managing their idiosyncrasies.

At a more specific level, the contribution of the present dissertation is twofold. First, it furthers our understanding of organizational adaptation by extending the classical, Simonian (see essay I or Fiori 2011) search to account more fully for the cognitive attributes of decision-making. The notion of bounded rationality is often equated with the process of search. While search does indeed represent a fundamental distinction from instantaneous optimizing, an

inherent attribute of the full rationality, it is by no means the sole aspect in which our rationality is bounded.

Besides purely computational constraints emphasized in Simonian search, decision makers are also subject to the effects of cognitive instruments that have evolved to maximize our genetic persistence (rather than profit or other individual-level objective function) given computational constraints of our mind. A result is an intended rationality that is bounded along multiple dimensions. Although there is a broad consensus on multidimensionality of bounded rationality (see Fiori 2011), the works that explore the organizational consequences of search have remained largely separate from the advances in cognitive psychology that bring to light the intricacy of human thinking. Curiously enough, this separation persisted despite the fact that Simon himself recognized other bounds to human rationality (see for example Simon 1987). The present work begins to bridge this *ad hoc* divide by integrating broader elements of cognition with the primitives of boundedly rational search.

Second, this work contributes to the development of the emerging field of behavioral strategy by identifying bias as an important lever in managing decision makers' cognition. A common view in organization sciences regards bias as a deviation from rationality (see Arnott 2006). In this sense, a bias is strictly detrimental for organizational objective function. However, by viewing bias as a deviation from an "unbiased" (rather than rational) choice, we can discover further means to value creation. From this perspective, bias is effectively a variation along a certain, potentially but not necessarily irrelevant, dimension. Consider, for example, that given some level of accuracy in terms of recognizing interdependencies, mental models can be more or less complex, *i.e.* there can be a bias toward complicating or

simplifying (see essay II). In this example, the notion of bias per se does not directly bear on the rationality of the decision maker's choice; however, as essay II of this work demonstrates, it is consequential for the attainment of organizational objectives. Here bias is not a deviation from rationality (cf. Arnott 2006), but rather a means to organizational efficiency. Decision maker' bias, therefore, may be an instrument of behavioral strategy that serves organizational purpose.

6.2. Limitations

As this work weds the primitives of boundedly rational search with the elements of human thinking, the main limitations categorize accordingly, *i.e.* search-related limitations and cognition-related limitations. Note that the following boundary conditions apply to the whole dissertation, specific constraints can be found in the corresponding essays.

6.2.1. Search-related limitations

An important boundary condition of this work is that it focuses on the recombination search. While search per se is a ubiquitous behavior, it is not uniform in the way it manifests itself. Organizational scholars have analyzed multiple forms of behavior, which share the main properties of search. These include recombination search (*e.g.* Levinthal 1997) and more generally search on rugged landscapes (*c.f.* Levinthal 1997 and Winter *et al.* 2007), sequential sampling (*e.g.* Posen and Levinthal 2012), credit assignment problems (*e.g.* Denrell *et al.* 2004), *etc.* Each form of search applies to a particular kind of organizational problems. For example, the process of sequential sampling characterizes how organizations discover superior

options from among noisy alternatives. Similarly, recombination search addresses a particular problem, *i.e.* that of finding an efficient configuration of elements. A typical context to which this applies is technological development (see Fleming and Sorenson 2001) or search for an efficient strategy (Rivkin 2000).

Further, it is important to observe that this work considers organizations as agents equivalent to decision makers. While this applies to individual entrepreneurs, sufficiently small organizations or sufficiently powerful actors in firms, group decision processes, political and structural aspects are beyond the scope of the present dissertation. Processes of joint search (*e.g.* Knudsen and Srikanth 2014) or organizational search with multiple decision makers (*e.g.* Rivkin and Siggelkow 2002, Siggelkow and Levinthal 2005) may involve more complex interactions than those brought to light in this work.

6.2.2. Cognition-related limitations

Although this work considers different and (relatively) disjoint cognitive aspects, it naturally accounts for only a fraction of what constitutes human thinking. Analyzing all relevant attributes of decision making in a single work is, of course, impractical. Nonetheless, an important boundary condition to the above analyses is that they reflect isolated interactions with the process of search, *i.e.* pure effects of a given cognitive attribute. The elements of human cognition, however, are not separate from one another. System 1 and System 2, intuition and mental models coexist and intermingle in the human mind (see for example Alos-Ferrer and Strack 2014). In other words, the reported models depict a simplified mind that is subject to a) computational constraints that trigger search, and b) a select attribute of human thinking.

The select elements of human cognition are further modeled to capture a specific property of interest. Consider, for example, System 1 or uncontrolled, intuitive thinking. While there is a broad spectrum of functions that the automatic information processing performs in our mind (see for example, Kahneman 2011), the analysis in this dissertation (essay I) essentially abstracts away from the most of its roles and focuses on a single manifestation, affect heuristic.

In principle, such simplification is inherent in scientific modeling. A question is whether suppressing possible interactions and depicting elements of human thinking in a simplified manner allows us to formulate useful predictions. This work does report empirical evidence which is consistent with some of the theoretical predictions. However, this by no means dismisses the need in further research to fully understand the interactions of human thinking with the process of boundedly rational search.

6.3. Future research

The present dissertation develops a theoretical basis for managing elements of decision makers' cognition over time, as organizations adapt to a given set of conditions. A just question is how organizations could do so and whether the reported tendencies hold in reality. While the essays do speculate on these subjects, establishing how organizations can in practice manage their decision-makers' cognition to extract greater value from the process of search represents a fruitful ground for future research.

A particularly intriguing avenue lies in identifying specific levers that organizations can use to extract value from decision makers' cognitive idiosyncrasies. Examples include

managerial rotation (*e.g.* re-biasing, see section 2.4.3.), architecture of decision-making (*e.g.* complexity of realized decision, see section 3.6.), hiring and promotion policies (*e.g.* odd thinking, see section 5.6.), *etc.*

In a similar vein, organizations can potentially mimic some of the individual-specific effects with their structural elements. Ethiraj and Levinthal, for example, discuss performance asymmetries of under- and overmodularity (2004). Their analysis points to the idea that organizations should manage the types of errors they make in specifying certain aspects of objective reality. This kind of reasoning can be useful in light of the regularities reported in the present work. Consider that organizations have multiple means to regulate the relative complexity of their internal processes, including knowledge articulation and codification (Zollo and Winter 2002) or the degree of automation in the production processes (Camuffo and Volpato 1996) among others. Establishing whether performance effects of these and similar choices follow the dynamics identified in essays II and III may reveal further means to managing organizational adaptation.

Thesis that organization's adaptive trajectory depends on the cognition of the key decision makers suggests that multiple, interdependent organizations, *i.e.* organizational ecosystems, may be shaped by synergy (or collision) of cognitive styles. Both practitioner-oriented and academic literature points to the vital nature of joint, coevolutionary processes in industries (see for example, Adner and Kapoor 2010, Iansity and Levien 2004, McKelvey 1999). From this perspective, an action of one organization defines the landscape for another. Industries, or organizational ecosystems, therefore, may fare differently depending on how their multiple decision makers think. Equally, this applies to intraorganizational settings where

decisions of various divisions and departments depend on one another (see for example, Rivkin and Siggelkow 2002, 2003).

Further, there is a substantial potential for analyzing the interplay between elements of human cognition and architectural instruments that alter the dynamics of search. For example, in view of its relevance for some aspects of the phenomena studied here and virtually ubiquitous presence, modularity represent an intriguing subject along this way (see above, Brusoni and Prencipe 2001, Brusoni *et al.* 2007, Ethiraj and Levinthal 2004).

Another natural extension of the present work is to lift the main limitations and boundary conditions. First, it can be useful to consider other manifestations of the process of search. While the regularities reported in this work are generally robust for alternative model specifications, different conceptual settings (*i.e.* other types of organizational problems), for example sequential sampling, may not necessarily show identical patterns. Similarly, broadening the spectrum of considered properties of human thinking may reveal further potential of imperfections in managerial cognition for value creation. A very promising path in this direction is to account for human emotions and their interaction with the process of search. Emotions interfere with decision-making (*e.g.* Lerner, Small, Loewenstein 2004, Maitlis and Ozcelik 2004, Simon 1987). By extension, this means that emotions affect the process of search. However, as the present work suggests, understanding the outcomes of such influence may require a controlled analysis. Finally, there is the possibility of complex interactions among multiple elements of human cognition and their joint implications for the process of search. While studying such multi-way interactions entails substantial computational and analytical hurdles, their analysis will allow better approximation of our mind's potential, more efficient strategies and ultimately better predictions.

Bibliography

- Adner R, Csaszar F, Zemsky P. 2014. Positioning on a multiattribute landscape. *Management Science* **60**: 2794–2815.
- Adner R, Kapoor R. 2010. Value creation in innovation ecosystems: How the structure of technological interdependence affects firm performance in new technology generations. *Strategic Management Journal* **31**: 306–333.
- Alos-Ferrer C, Strack F. 2014. From dual processes to multiple selves: Implications for economic behavior. *Journal of Economic Psychology* **41**: 1–11.
- Amabile TM. 1983. The social psychology of creativity: A componential conceptualization. *Journal of Personality and Social Psychology* **45**: 357–376.
- Anderson C, Paine F. 1975. Managerial perceptions and strategic behavior. *Academy of Management Journal* **18**: 811–823.
- Anderson P. 1999. Complexity theory and organization science. *Organization Science* **10**: 216–232.
- Arnott D. 2006. Cognitive biases and decision support systems development: A design science approach. *Information Systems Journal* **16**: 55–78.
- Ashton AH. 1985. Does consensus imply accuracy in accounting studies of decision making? *The Accounting Review* **60**: 173–185.
- Axelrod RM. 1976. *Structure of Decision: The Cognitive Maps of Political Elites*. Princeton University Press: Princeton, NJ.
- Baer M, Dirks KT, Nickerson JA. 2013. Microfoundations of strategic problem formulation. *Organization Science* **34**: 197–214.
- Barney JB. 1986. Strategic factor markets: Expectations, Luck, and Business Strategy. *Management Science* **32**: 1231–1241.
- Baron J. 2007. *Thinking and Deciding* 4th ed. Cambridge Univ. Press: New York, NY.

- Barr PS, Stimpert JL, Huff AS. 1992. Cognitive change, strategic action and organizational renewal. *Strategic Management Journal* **13**: 15–36.
- Baum JAC, Mezias SJ. 1992. Localized competition and organizational failure in the Manhattan hotel industry, 1898–1990. *Administrative Science Quarterly* **37**: 580–604.
- Baum JAC, Singh JV. 1994. Organizational niche overlap and the dynamics of organizational mortality. *American Journal of Sociology* **100**: 346–380.
- Baumann O, Martignoni D. 2011. Evaluating the new: The contingent value of a pro-innovation bias. *Schmalenbach Business Review* **63**: 393–415.
- Baumann O, Siggelkow N. 2013. Dealing with complexity: Integrated vs. chunky search process. *Organization Science* **24**: 116–132.
- Bernardo AE, Welch I. 2001. On the evolution of overconfidence and entrepreneurs. *Journal of Economics and Management Strategy* **10**: 301–330.
- Besanko D, Dranove D, Shanley M. 1996. *Economics of Strategy*. Wiley: New York.
- Bettis RA, Prahalad CK. 1995. The dominant logic: Retrospective and extension. *Strategic Management Journal* **16**: 5–14.
- Bourgeois LJ. 1980. Performance and consensus. *Strategic Management Journal* **1**: 227–248.
- Bourgeois LJ. 1985. Strategic goals, perceived uncertainty, and economic performance in volatile environments. *Academy of Management Journal* **28**: 548–573.
- Bromiley P. 1991. Testing a causal model of corporate risk taking and performance. *Academy of Management Journal* **34**: 37–59.
- Brusoni S, Marengo L, Prencipe A, Valente M. 2007. The value and costs of modularity: A problem-solving perspective. *European Management Review* **4**: 121–132.
- Brusoni S, Prencipe A. 2001. Unpacking the black box of modularity: Technologies, products and organizations. *Industrial and Corporate Change* **10**: 179–205.
- Cacioppo JT, Petty RE, Feinstein J, Jarvis WBG. 1996. Dispositional differences in cognitive motivation: The life and times of individuals varying in need for cognition. *Psychological Bulletin* **119**: 197–253.

- Calori R, Johnson J, Sarnin P. 1994. CEOs' cognitive maps and the scope of the organization. *Strategic Management Journal* **15**: 437–457.
- Camerer C, Babcock L, Loewenstein G, Thaler R. 1997. Labor supply of New York City cab drivers: One day at a time. *Quarterly Journal of Economics* **112**: 407–441.
- Camuffo A, Volpato G. 1996. Dynamic capabilities and manufacturing automation: Organizational learning in the Italian automobile industry. *Industrial and Corporate Change* **5**: 813–838.
- Carraher SM, Buckley MR. 1996. Cognitive complexity and the perceived dimensionality of pay satisfaction. *Journal of Applied Psychology* **81**: 102–109.
- Carroll GR. 1985. Concentration and specialization: Dynamic of niche width in populations of organizations. *American Journal of Sociology* **90**: 1262–1283.
- Christensen M, Knudsen T. 2010. Design of decision-making organizations. *Management Science* **56**: 71–89.
- Conlisk J. 1996. Why bounded rationality? *Journal of Economic Literature* **34**: 669–700.
- Csaszar FA, Levinthal D. 2016. Mental representations and the discovery of new strategies. *Strategic Management Journal* **37**: 2031–2049.
- Csaszar FA. 2012. Organizational structure as a determinant of performance: Evidence from mutual funds. *Strategic Management Journal* **33**: 611–632.
- Cyert R, March J. 1963. *A Behavioral Theory of the Firm*. Prentice-Hall: Englewood Cliffs, NJ.
- Dane E, Pratt MG. 2007. Exploring intuition and its role in managerial decision making. *Academy of Management Review* **32**: 33–54.
- Deepphouse DL. 1999. To be different, or to be the same? It's a question of (and theory) of strategic balance. *Strategic Management Journal* **20**: 147–166.
- Denrell J, Fang C, Levinthal DA. 2004. From T-mazes to labyrinths: Learning from model-based feedback. *Management Science* **50**: 1366–1378.

- Denrell J, March JG. 2001. Adaptation as information restriction: The hot stove effect. *Organization Science* **12**: 523–538.
- Dierickx I, Cool K. 1989. Asset stock accumulation and sustainability of competitive advantage. *Management Science* **35**: 1504–1511.
- DiMaggio PJ, Powell WW. 1983. The iron cage revisited: Institutional isomorphism and collective rationality in organizational fields. *American Sociological Review* **48**: 147–160.
- Duckworth KL, Bargh JA, Chaiken MG, Chaiken S. 2002. The automatic evaluation of novel stimuli. *Psychological Science* **13**: 513–519.
- Dyer JH, Gregersen HB, Christensen, C. M. 2009. The innovator's DNA. *Harvard Business Review* **87**: 60–67.
- Eden C. 1992. On the nature of cognitive maps. *Journal of Management Studies* **29**: 261–265.
- Einhorn HJ, Hogarth RM. 1986. Judging probable cause. *Psychological Bulletin* **99**: 3–19.
- Eisenhardt K, Sull D. 2001. What is strategy in the new economy? *Harvard Business Review* **79**: 107–116.
- Eisenhardt KM, Bourgeois LJ III. 1988. Politics of strategic decision making in high-velocity environments: Toward a midrange theory. *Academy of Management Journal* **31**: 737–770.
- Eisenhardt KM, Martin JA. 2000. Dynamic capabilities: what are they? *Strategic Management Journal* **21**: 1105–1121.
- Eisenhardt KM, Zbaracki MJ. 1992. Strategic decision making. *Strategic Management Journal* **13**: 17–37.
- Eisenhardt KM. 1989. Making fast strategic decisions in high velocity environments. *Academy of Management Journal* **32**: 543–576.
- Epstein S. 1990. Cognitive-experiential Self-theory. Pervin LA, John OP, eds. *Handbook of personality theory and research: Theory and research*. Guilford Publications Inc.: New York.

- Ethiraj SK, Levinthal D. 2009. Hoping for A to Z while rewarding only A: Complex organizations and multiple goals. *Organization Science* **20**: 4–21.
- Ethiraj SK, Levinthal DA. 2004. Modularity and innovation in complex systems. *Management Science* **50**: 159–173.
- Fazio RH, Sanbonmatsu DM, Powell MC, Kardes FR. 1986. On the automatic activation of attitudes. *Journal of Personality and Social Psychology* **50**: 229–238.
- Fazio RH. 2001. On the automatic activation of associated evaluations: An overview. *Cognition and Emotion* **15**: 115–141.
- Felin T, Foss NJ. 2009. Strategic organization: A field in search of micro-foundations. *Strategic Organization* **3**: 441–455.
- Fines CH. 1998. *Clockspeed: Winning Industry Control in the Age of Temporary Advantage*. Perseus: Reading, MA.
- Fiol CM, Lyles MA. 1985. Organizational learning. *Academy of Management Review* **10**: 803–813.
- Fiori S. 2011. Forms of bounded rationality: The reception and redefinition of Herbert A. Simon's perspective. *Review of Political Economy* **23**: 587–612.
- Fleming L, Sorenson O. 2001. Technology as a complex adaptive system: Evidence from patent data. *Research Policy* **30**: 1019–1039.
- Forgas JP, Burnham DK, Trimboli C. 1988. Mood, memory, and social judgments in children. *Journal of Personality and Social Psychology* **54**: 697–703.
- Frederick S, Kahneman D. 2002. Representativeness revisited: Attribute substitution in intuitive judgment. Gilovich T, Griffin D, Kahneman D, eds. *Heuristics and biases: The psychology of intuitive judgment*. Cambridge University Press: New York.
- Frederick S, Loewenstein G, O'Donoghue T. 2002. Time discounting and time preference: A critical review. *Journal of Economic Literature* **40**: 351–401.
- Frederick S. 2005. Cognitive reflection and decision making. *Journal of Economic Perspectives* **19**: 25–42.

- Gaba V, Bhattacharya S. 2012. Aspirations, innovation, and corporate venture capital: A behavioral perspective. *Strategic Entrepreneurship Journal* **6**: 178–199.
- Galasso A, Simcoe TS. 2011. CEO overconfidence and innovation. *Management Science* **57**: 1469–1484.
- Ganco M, Hoetker G. 2009. NK modeling methodology in the strategy literature: Bounded search on a rugged landscape. In *Research Methodology in Strategy and Management*, Vol. **5**, Bergh DD, Ketchen DJ Jr. Eds. Emerald Group Publishing: Bingley, UK: 237–268.
- Gary MS, Wood RE. 2011. Mental models, decision rules, and performance heterogeneity. *Strategic Management Journal* **32**: 569–594.
- Gavetti G, Levinthal D, Ocasio W. 2007. Neo-Carnegie: The Carnegie school's past, present, and reconstructing for the future. *Organization Science* **18**: 523–536.
- Gavetti G, Levinthal D. 2000. Looking forward and looking backward: Cognitive and experiential search. *Administrative Science Quarterly* **45**: 113–137.
- Gavetti G, Levinthal DA, Rivkin JW. 2005. Strategy making in novel and complex worlds: The power of analogy. *Strategic Management Journal* **26**: 691–712.
- Gavetti G, Rivkin JW. 2007. On the origin of strategy: Action and cognition over time. *Organization Science* **18**: 420–439.
- Gavetti G. 2005. Cognition and hierarchy: rethinking the microfoundations of capabilities' development. *Organization Science* **16**: 599–617.
- Gavetti G. 2012. Toward a behavioral theory of strategy. *Organization Science* **23**: 267–285.
- Gawronski B, Bodenhausen GV. 2006. Associative and propositional processes in evaluation: An integrative review of implicit and explicit attitude change. *Psychological Bulletin* **132**: 692–731.
- Ghemawat P, Levinthal D. 2008. Choice interactions and business strategy. *Management Science* **54**: 1638–1651.
- Ghemawat P. 2010. Finding your strategy in the new landscape. *Harvard Business Review* **88**: 54–60.

- Gigerenzer G, Goldstein DG. 1996. Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review* **103**: 650–669.
- Gigerenzer G, Todd PM, the ABC Research Group. 1999. *Simple Heuristics That Make Us Smart*. Oxford University Press: New York.
- Gilbert DT. 2011. Buried by bad decisions. *Nature* **474**: 276–277.
- Gilovich T, Griffin D. 2002. Introduction—heuristics and biases: Then and now. Gilovich T, Griffin D, Kahneman D, eds. *Heuristics and biases: The psychology of intuitive judgment*. Cambridge University Press: New York.
- Giola DA, Manz CC. 1985. Linking cognition and behavior: A script processing interpretation of vicarious learning. *Academy of Management Review* **10**: 527–539.
- Greenwald AG, Banaji MR. 1995. Implicit social cognition: Attitudes, self-esteem, and stereotypes. *Psychological Review* **102**: 4–27.
- Gregg AP, Seibt B, Banaji MR. 2006. Easier done than undone: asymmetry in the malleability of implicit preferences. *Journal of Personality and Social Psychology* **90**: 1–20.
- Greve HR. 1998. Performance, aspirations, and risky organizational change. *Administrative Science Quarterly* **43**: 58–86.
- Greve HR. 2002. Sticky aspirations: Organizational time perspective and competitiveness. *Organization Science* **13**: 1–17.
- Greve HR. 2003. A behavioral theory of R&D expenditures and innovations: Evidence from shipbuilding. *Academy of Management Journal* **46**: 685–702.
- Gupta AK, Smith KG, Shalley CE. 2006. The interplay between exploration and exploitation. *Academy of Management Journal* **49**: 693–706.
- Haidt J. 2001. The emotional dog and its rational tail: A social intuitionist approach to moral judgment. *Psychological Review* **108**: 814–834.
- Hambrick DC, Mason PA. 1984. Upper echelons: The organization as a reflection of its top managers. *Academy of Management Review* **9**: 193–206.

- Hayes J, Allinson CW. 1998. Cognitive style and the theory and practice of individual and collective learning in organizations. *Human Relations* **51**: 847–871.
- Heath C, Staudenmayer N. 2000. Coordination neglect: How lay theories of organizing complicate coordination in organizations. *Research in Organizational Behavior* **22**: 155–193.
- Helfat CE, Peteraf MA. 2015. Managerial cognitive capabilities and the microfoundations of dynamic capabilities. *Strategic Management Journal* **36**: 831–850.
- Hermann MG. 1980. Explaining foreign policy behavior using the personal characteristics of political leaders. *International Studies Quarterly* **24**: 7–46.
- Hoberg G, Phillips G. 2016. Text-based network industries and endogenous product differentiation. *Journal of Political Economy* **124**: 1423–1465.
- Hodgkinson GP, Bown NJ, Maule AJ, Glaister KW, Pearman AD. 1999. Breaking the frame: an analysis of strategic cognition and decision making under uncertainty. *Strategic Management Journal* **20**: 977–985.
- Hotelling H. 1929. Stability in competition. *Economic Journal* **39**: 41–57.
- Hu S, Blettner D, and Bettis RA. 2011. Adaptive aspirations: Performance consequences of risk preferences at extremes and alternative reference groups. *Strategic Management Journal* **32**: 1426–1436.
- Hutchinson JMC, Gigerenzer G. 2005. Simple heuristics and rules of thumb: Where psychologists and behavioral biologists might meet. *Behavioral Processes* **69**: 97–124.
- Iansity M, Levien R. 2004. Strategy as ecology. *Harvard Business Review* **82**: 68–81.
- Inbar Y, Cone J, Gilovich T. 2010. People's intuitions about intuitive insight and intuitive choice. *Journal of Personality and Social Psychology* **99**: 232–247.
- Isaack TS. 1978. Intuition: An ignored dimension of management. *Academy of Management Review* **3**: 917–922.
- Ishiyama J, Backstrom J. 2011. Reconciliation and conceptual complexity: The case of post conflict Kenya. *African and Asian Studies* **10**: 366–386.

- Jacoby LL. 1991. A process dissociation framework: Separating automatic from intentional uses of memory. *Journal of Memory and Language* **30**: 513–541.
- Johnson DDP, Fowler JH. 2011. The evolution of overconfidence. *Nature* **477**: 317–320.
- Kahneman D, Knetsch JL, Thaler RH. 1991. Anomalies: The endowment effect, loss aversion, and status quo bias. *Journal of Economic Perspectives* **5**:193–206.
- Kahneman D, Lovallo D, Sibony O. 2011. Before you make that big decision... *Harvard Business Review* **89**: 50–60.
- Kahneman D, Tversky A. 1996. On the reality of cognitive illusions. *Psychological Review* **103**: 582–591.
- Kahneman D. 2003. A perspective on judgment and choice: Mapping bounded rationality. *American Psychologist* **58**: 697–720.
- Kahneman D. 2011. *Thinking, Fast and Slow*. Farrar, Straus and Giroux, NJ.
- Kanter RM. 2011. How great companies think differently. *Harvard Business Review* **89**: 66–78.
- Kaplan S, Tripsas M. 2008. Thinking about technology: Applying a cognitive lens to technical change. *Research Policy* **37**: 790–805.
- Kauffman SA. 1993. *The Origins of Order: Self-Organization and Selection in Evolution*. Oxford University Press: New York.
- Khatri N, Ng HA. 2000. The role of intuition in strategic decision making. *Human Relations* **53**: 57–86.
- Knight D, Pearce CL, Smith KG, Olian JD, Sims HP, Smith KA, Flood P. 1999. Top management team diversity, group process and strategic consensus. *Strategic Management Journal* **20**: 445–465.
- Knudsen T, Levinthal DA. 2007. Two faces of search: Alternative generation and alternative evaluation. *Organization Science* **18**: 39–54.

- Knudsen T, Srikanth K. 2014. Coordinated exploration: Organizing joint search by multiple specialists to overcome mutual confusion and joint myopia. *Administrative Science Quarterly* **59**: 409–441.
- Kogut B, Zander U. 1992. Knowledge of the firm, combinative capabilities, and the replication of technology. *Organization Science* **3**: 383–397.
- Kruglanski AW, Webster DM. 1996. Motivated closing of the mind: "Seizing" and "freezing." *Psychological Review* **103**: 263–283.
- Langfield-Smith K. 1992. Exploring the need for a shared cognitive map. *Journal of Management Studies* **29**: 349–368.
- Lant T, Shapira Z. 2008. Managerial reasoning about aspirations and expectations. *Journal of Economic Behavior and Organization* **66**: 60–73.
- Lant TK, Mezias SJ. 1992. An organizational learning model of convergence and reorientation. *Organization Science* **3**: 47–71.
- Lant TK, Milliken FJ, Batra B. 1992. The role of managerial learning and interpretation in strategic persistence and reorientation: an empirical exploration. *Strategic Management Journal* **13**: 585–608.
- Lant TK, Montgomery DB. 1987. Learning from strategic success and failure. *Journal of Business Research* **15**: 503–517.
- Lant TK. 1992. Aspiration level adaptation: An empirical exploration. *Management Science* **38**: 623–644.
- Larkin I, Pierce L, Gino F. 2012. The psychological costs of pay-for-performance: Implications for the strategic compensation of employees. *Strategic Management Journal* **33**: 1194–1214.
- Lenox MJ, Rockart SF, Lewin AY. 2006. Interdependency, competition, and the distribution of firm and industry profits. *Management Science* **52**: 757–772.
- Lenox MJ, Rockart SF, Lewin AY. 2007. Interdependency, competition, and industry dynamics. *Management Science* **53**: 599–615.

- Lerner JS, Small DA, Loewenstein G. 2004. Heart strings and purse strings: Carryover effects of emotions on economic decisions. *Psychological Science* **15**: 337–341.
- Levinthal DA, March JG. 1981. A model of adaptive organizational search. *Journal of Economic Behavior & Organization* **2**: 307–333.
- Levinthal DA, Posen HE. 2007. Myopia of selection: Does organizational adaptation limit the efficacy of population selection? *Administrative Science Quarterly* **52**: 586–620.
- Levinthal DA. 1997. Adaptation on rugged landscapes. *Management Science* **43**: 934–950.
- Levinthal DA. 2011. A behavioral approach to strategy - what's the alternative? *Strategic Management Journal* **32**: 1517–1523.
- Levitt B, March JG. 1988. Organizational learning. *Annual Review of Sociology* **14**: 319–340.
- Lewin AY, Long CP, Carroll TN. 1999. The coevolution of new organizational forms. *Organization Science* **10**: 535–550.
- Locke EA, Latham GP. 2002. Building a practically useful theory of goal setting and task motivation. *American Psychologist* **57**: 705–717.
- Locke EA, Shaw KN, Saari LM, Latham JP. 1981. Goal setting and task performance: 1969—1980. *Psychological Bulletin* **90**: 125–152.
- Locke EA. 1996. Motivation through conscious goal setting. *Applied and Preventive Psychology* **5**: 117–124.
- Lounamaa PH, March JG. 1987. Adaptive coordination of a learning team. *Management Science* **33**: 107–123.
- Lyles MA, Mitroff II. 1980. Organizational problem formulation: An empirical study. *Administrative Science Quarterly* **25**: 102–119.
- Lyles MA, Thomas H. 1988. Strategic problem formulation: Biases and assumptions embedded in alternative decision-making models. *Journal of Management Studies* **25**: 131–145.
- Maddux WW, Williams E, Swaab R, Betania T. 2014. *Ricardo Semler: A Revolutionary Model of Leadership*. Case Study. Harvard Business Publishing: Boston, MA.

- Maitlis S., Ozcelik H. 2004. Toxic decision processes: A study of emotion and organizational decision making. *Organization Science* **15**: 375–393.
- Marcel JJ, Barr PS, Duhaime IM. 2011. The influence of executive cognition on competitive dynamics. *Strategic Management Journal* **32**: 115–138.
- March JG, Shapira Z. 1992. Variable risk preference and the focus of attention. *Psychological Review* **99**: 172–183.
- March JG, Simon HA. 1958. *Organizations*. John Wiley & Sons: New York.
- March JG. 1991. Exploration and exploitation in organizational learning. *Organization Science* **2**: 71–87.
- March JG. 2006. Rationality, foolishness and adaptive intelligence. *Strategic Management Journal* **27**: 201–214.
- Martignoni D, Menon A, Siggelkow N. 2016. Consequences of misspecified mental models: Contrasting effects and the role of cognitive fit. *Strategic Management Journal* **37**: 2545–2568.
- McKelvey B. 1999. Avoiding complexity catastrophe in coevolutionary pockets: Strategies for rugged landscapes. *Organization Science* **10**: 294–321.
- Mezias SJ, Chen Y, Murphy PR. 2002. Aspiration-level adaptation in an american financial services organization: A field study. *Management Science* **48**: 1285–1300.
- Miller CC, Burke LM, Glick WH. 1998. Cognitive diversity among upper-echelon executives: Implications for strategic decision processes. *Strategic Management Journal* **19**: 39–58.
- Miller CC, Ireland RD. 2005. Intuition in strategic decision making: friend or foe in the fast-paced 21st century? *Academy of Management Executive* **19**: 19–30.
- Miller KD, Zhao M, Calantone RJ. 2006. Adding interpersonal learning and tacit knowledge to March's exploration-exploitation model. *Academy of Management Journal* **49**: 709–722.
- Mintzberg H, Raisinghani D, Theoret A. 1976. The structure of "unstructured" decision processes. *Administrative Science Quarterly* **21**: 246–275.

- Mumford MD, Gustafson SB. 1988. Creativity syndrome: Integration, application and innovation. *Psychological Bulletin* **103**: 27–43.
- Mumford MD. 2003. Where have we been, where are we going? Taking stock in creativity research. *Creativity Research Journal* **15**: 107–120.
- Nadkarni S, Barr PS. 2008. Environmental context, managerial cognition, and strategic action: An integrated view. *Strategic Management Journal* **29**: 1395–1427.
- Nadkarni S, Narayanan VK. 2007. Strategic schema, strategic flexibility, and firm performance: The moderating role of industry clockspeed. *Strategic Management Journal* **28**: 243–270.
- Nelson RR, Winter SG. 1982. *An Evolutionary Theory of Economic Change*. Harvard University Press: Cambridge, MA.
- Nelson RR. 2008. Bounded rationality, cognitive maps, and trial and error learning. *Journal of Economic Behavior and Organization* **67**: 78–89.
- Newell A, Simon HA. 1972. *Human Problem Solving*. Prentice-Hall: Englewood Cliffs, NJ.
- Nickerson JA, Silverman BS, Zenger TR. 2007. The 'problem' of creating and capturing value. *Strategic Organization* **5**: 211–225.
- Nisbett RE, Wilson TD. 1977. Telling more than we can know: Verbal reports on mental processes. *Psychological Review* **84**: 231–259.
- Nosek BA, Banaji MR, Greenwald AG. 2002. Harvesting implicit group attitudes and beliefs from a demonstration website. *Group Dynamics* **6**: 101–115.
- Pacheco-de-Almeida G, Zemsky P. 2007. The timing of resource development and sustainable competitive advantage. *Management Science* **53**: 651–666.
- Pacini R, Epstein S. 1999. The relation of rational and experiential information processing styles to personality, basic beliefs, and the ratio-bias phenomenon. *Journal of Personality and Social Psychology* **76**: 972–987.
- Peteraf MA. 1993. The cornerstones of competitive advantage: A resource-based view. *Strategic Management Journal* **14**: 179–191.

- Porac JF, Thomas H, Baden-Fuller C. 1989. Competitive groups as cognitive communities: The case of Scottish knitwear manufacturers. *Journal of Management Studies* **26**: 397–416.
- Porac JF, Thomas H. 1990. Taxonomic mental models in competitor definition. *Academy of Management Review* **15**: 224–240.
- Porter CA, Suedfeld P. 1981. Integrative complexity in the corresponding of literary figures: Effects of personal and societal stress. *Journal of Personality and Social Psychology* **40**: 321–330.
- Posen HE, Lee J, Yi S. 2013. The power of imperfect imitation. *Strategic Management Journal* **34**: 149–164.
- Posen HE, Levinthal DA. 2012. Chasing a moving target: Exploration and exploitation in dynamic environments. *Management Science* **58**: 587–601.
- Posen HE, Martignoni D, Lang M. 2013. *Fortune favors fools: How confidence can compensate for competence in learning*. Manuscript presented at the Annual Meeting of the Academy of Management.
- Powell TC, Lovallo D, Fox CR. 2011. Behavioral strategy. *Strategic Management Journal* **32**: 1369–1386.
- Rehder B. 2003. Categorization as causal reasoning. *Cognitive Science* **27**: 709–748.
- Reitzig M, Sorenson O. 2013. Biases in the selection stage of bottom-up strategy formulation. *Strategic Management Journal* **34**: 782–799.
- Rentsch JR, Klimoski RJ. 2002. Why do 'great minds' think alike?: Antecedents of team member schema agreement. *Journal of Organizational Behavior* **22**: 107–120.
- Rivkin JW, Siggelkow N. 2002. Organizational sticking points on NK landscapes. *Complexity* **7**: 31–43.
- Rivkin JW, Siggelkow N. 2003. Balancing search and stability: Interdependencies among elements of organizational design. *Management Science* **49**: 290–311.
- Rivkin JW, Siggelkow N. 2007. Patterned interactions in complex systems: Implications for exploration. *Management Science* **53**: 1068–1085.

- Rivkin JW. 2000. Imitation of complex strategies. *Management Science* **46**: 824–844.
- Roediger HL. 1990. Implicit memory: Retention without remembering. *American Psychologist* **45**: 1043–1056.
- Roger M. 2007. How successful leaders think. *Harvard Business Review* **85**: 60–67.
- Rosenkopf L, Almeida P. 2003. Overcoming local search through alliances and mobility. *Management Science* **49**: 751–766.
- Rosenkopf L, Nerkar A. 2001. Beyond local search: Boundary-spanning, exploration and impact in the optical disk industry. *Strategic Management Journal* **22**: 287–306.
- Sah R, Stiglitz J. 1986. Committees, hierarchies, and polyarchies. *Economic Journal* **98**: 451–470.
- Samuelson W, Zeckhauser R. 1988. Status quo bias in decision making. *Journal of Risk and Uncertainty* **1**: 7–59.
- Sargut G, McGrath RG. 2011. Learning to live with complexity. *Harvard Business Review* **89**: 68–76.
- Schacter DL, Chiu CYP, Ochsner KN. 1993. Implicit memory: A selective review. *Annual Review of Neuroscience* **16**: 159–182.
- Schneider SL. 1992. Framing and conflict: Aspiration level contingency, status quo, and current theories of risky choice. *Journal of Experimental Psychology: Learning, Memory, and Cognition* **18**: 1040–1057.
- Schwenk CR. 1984. Cognitive simplification processes in strategic decision-making. *Strategic Management Journal* **5**: 111–128.
- Schwenk CR. 1986. Information, cognitive biases, and commitment to a course of action. *Academy of Management Review* **11**: 298–310.
- Scott WR. 1987. The adolescence of institutional theory. *Administrative Science Quarterly* **32**: 493–511.
- Sharfman MP, Wolf G, Chase RB, Tansik DA. 1988. Antecedents of organizational slack. *Academy of Management Review* **13**: 601–614.

- Shinkle GA. 2012. Organizational aspirations, reference points, and goals: building on the past and aiming for the future. *Journal of Management* **38**: 415–455.
- Shrivastava P, Grant JH. Empirically derived models of strategic decision-making processes. *Strategic Management Journal* **6**: 97–113.
- Siggelkow N, Levinthal DA. 2005. Escaping real (non-benign) competency traps: Linking the dynamics of organizational structure to the dynamics of search. *Strategic Organization* **3**: 85–115.
- Siggelkow N, Rivkin JW. 2005. Speed and search: Designing organizations for turbulence and complexity. *Organization Science* **16**: 101–122.
- Siggelkow N. 2001. Change in the presence of fit: The rise, the fall, and the renaissance of Liz Claiborne. *Academy of Management Journal* **44**: 838–857.
- Siggelkow N. 2002. Evolution toward fit. *Administrative Science Quarterly* **47**: 125–159.
- Simon HA. 1955. A behavioral model of rational choice. *Quarterly Journal of Economics* **69**: 99–118.
- Simon HA. 1956. Rational choice and the structure of the environment. *Psychological Review* **63**: 129–138.
- Simon HA. 1957. *Models of Man: Social and Rational*. John Wiley and Sons: New York.
- Simon HA. 1987. Making management decisions: The role of intuition and emotion. *Academy of Management Executives* (1987–1989) **1**: 57–64.
- Simon HA. 1990. Invariants of human behavior. *Annual Review of Psychology* **41**: 1–19.
- Sloman SA. 1996. The empirical case for two systems of reasoning. *Psychological Bulletin* **119**: 3–22.
- Slovic P, Finucane M, Peters E, MacGregor DG. 2002. Rational actors or rational fools: Implications of the affect heuristic for behavioral economics. *Journal Socio-Economics* **31**: 329–342.
- Smircich L, Stubbart C. 1985. Strategic management in an enacted world. *Academy of Management Review* **10**: 724–736.

- Smith ER, DeCoster J. 2000. Dual-process models in social and cognitive psychology: Conceptual integration and links to underlying memory systems. *Personality and Social Psychology Review* **4**: 108–131.
- Smith KG, Smith KA, Olian JD, Sims HP Jr., O'Bannon DP, and Scully JA. 1994. Top management team demography and process: The role of social integration and communication. *Administrative Science Quarterly* **39**: 412–438.
- Sraffa P. 1926. The laws of returns under competitive conditions. *Economic Journal* **36**: 535–550.
- Stanovich KE, West RF. 2000. Individual differences in reasoning: Implications for the rationality debate. *Behavioral and Brain Sciences* **23**: 645–726.
- Streufert S, Suedfeld P, and Driver MJ. 1965. Conceptual structure, information search, and information utilization. *Journal of Personality and Social Psychology* **2**: 736–740.
- Stuart TE, Podolny JM. 1996. Local search and the evolution of technological capabilities. *Strategic Management Journal* **17**: 21–38.
- Suedfeld P, Bluck S. 1988. Changes in integrative complexity prior to surprise attacks. *Journal of Conflict Resolution* **32**: 626–635.
- Suedfeld P, Bluck S. 1993. Changes in integrative complexity accompanying significant life events: Historical evidence. *Journal of Personality and Social Psychology* **64**: 124–130.
- Suedfeld P, Coren S. 1992. Cognitive correlates of conceptual complexity. *Personality and Individual Differences* **13**: 1193–1199.
- Suedfeld P, Tetlock P. 1977. Integrative complexity of communications in international crises. *Journal of Conflict Resolution* **21**: 169–184.
- Suedfeld P, Tetlock PE, and Streufert S. 1992. Conceptual/integrative complexity. In CP Smith (Ed.), *Motivation and personality: Handbook of thematic content analysis*: 393–400. Cambridge University Press: New York.
- Sull D, Eisenhardt KM. 2012. Simple rules for a complex world. *Harvard Business Review* **90**: 68–74.

- Tetlock PE, Peterson RS, Berry JM. 1993. Flattering and unflattering personality portraits of integratively simple and complex managers. *Journal of Personality and Social Psychology* **64**: 500–511.
- Tetlock PE. 1983. Accountability and complexity of thought. *Journal of Personality and Social Psychology* **45**: 74–83.
- Tetlock PE. 1985. Integrative complexity of American and Soviet foreign policy rhetoric: A time-series analysis. *Journal of Personality and Social Psychology* **49**: 1565–1585.
- Thoemmes F, Conway LG III. 2007. Integrative complexity of 41 U.S. presidents. *Political Psychology* **28**: 193–226.
- Tripsas M, Gavetti G. 2000. Capabilities, cognition, and inertia: Evidence from digital imaging. *Strategic Management Journal* **21**: 1147–1161.
- Tversky A, Kahneman D. 1974. Judgment under uncertainty: Heuristics and biases. *Science* **185**: 1124–1131.
- Uzzi B, Mukherjee S, Stringer M, Jones B. 2013. Atypical combinations and scientific impact. *Science* **342**: 468–472.
- Volkema RJ. 1983. Problem formulation in planning and design. *Management Science* **29**: 639–652.
- Wally S, Baum JR. 1994. Personal and structural determinants of the pace of strategic decision making. *Academy of Management Journal* **37**: 932–956.
- Walsh JP, Ungson GR. 1991. Organizational memory. *Academy of Management Review* **16**: 57–91.
- Walsh JP. 1995. Managerial and organizational cognition: Notes from a trip down memory lane. *Organization Science* **6**: 280–321.
- Waring GF. 1996. Industry difference in the persistence of firm-specific returns. *American Economic Review* **86**: 1253–1265.
- Washburn M, Bromiley P. 2012. Comparing aspiration models: The role of selective attention. *Journal of Management Studies* **49**: 896–917.

- Wiggins RR, Ruefli TW. 2002. Sustained competitive advantage: Temporal dynamics and the incidence and persistence of superior economic performance. *Organization Science* **13**: 81–105.
- Wilson TD, Brekke N. 1994. Mental contamination and mental correction: Unwanted influences on judgments and evaluation. *Psychological Bulletin* **116**: 117–142.
- Wilson TD, Centerbar DB, Brekke N. 2002. Mental contamination and the debiasing problem. Gilovich T, Griffin D, Kahneman D, eds. *Heuristics and biases: The psychology of intuitive judgment*. Cambridge University Press: New York.
- Wilson TD, Lindsey S, Schooler TY. 2000. A model of dual attitudes. *Psychological Review* **107**: 101–126.
- Winter SG, Cattani G, Dorsch A. 2007. The value of moderate obsession: Insights from a new model of organizational search. *Organization Science* **18**: 403–419.
- Wong EM, Ormiston ME, Tetlock PE. 2011. The effects of top management team integrative complexity and decentralized decision making on corporate social performance. *Academy of Management Journal* **54**: 1207–1228.
- Xu B, Liu R, Liu W. 2014. Individual bias and organizational objectivity: An agent-based simulation. *Journal of Artificial Societies and Social Simulation* **17**:2.
- Yi SKM, Steyvers M, Lee MD, Dry JM. 2012. The wisdom of the crowd in combinatorial problems. *Cognitive Science* **36**: 452–470.
- Young MD, Hermann MG. 2014. Increased Complexity Has Its Benefits. *Political Psychology* **35**: 635–645.
- Zajac EJ, Bazerman MH. 1991. Blind spots in industry and competitor analysis: Implications of interfirm (mis)perceptions for strategic decisions. *Academy of Management Review* **16**: 37–56.
- Zajonc RB. 1980. Feeling and thinking: Preferences need no inferences. *American Psychologist* **35**: 151–175.
- Zollo M, Winter SG. 2002. Deliberate learning and the evolution of dynamic capabilities. *Organization Science* **13**: 339–351.

Titre : Essais en Stratégie Comportementale: L'Adaptation Re-biaisée, Complexité Mal Comprise, et Aliens Cognitifs

Mots clés : Stratégie comportementale, cognition, adaptation organisationnelle

Résumé : Ce travail se concentre sur l'idée que la rationalité organisationnelle est limitée: les décideurs recherchent des solutions offrant un minimum de satisfaction et pensent d'une manière qui est typique pour l'homme. La thèse explore cette interaction entre le processus de recherche organisationnelle et la cognition des décideurs et démontre que certains biais (distorsions) dans les aspects caractéristiques de notre pensée peuvent être des instruments de stratégie comportementale.

Comme le point de départ, je complète le premier primitif de la rationalité limitée, i.e. la génération et l'évaluation d'alternatives, avec des éléments intégrés de la cognition humaine, tels que la pensée intuitive, spécifiquement l'heuristique d'affect, et des représentations mentales imparfaites.

À l'aide de modèles de calcul, j'étudie les effets des biais correspondants à ces éléments de la cognition humaine (préférences affectives et erreurs systématiques dans les représentations mentales) dans le temps lorsque des organisations s'adaptent à des environnements complexes. Cela me permet d'identifier les cycles de vie des éléments de la cognition humaine et de montrer que les organisations devraient gérer (plutôt que d'éliminer) certains biais. Enfin, je fais des propositions et je teste empiriquement un sous-ensemble de mes prédictions.

En conclusion, ce travail vise à faire progresser la théorie émergente de la stratégie comportementale en considérant conjointement différents primitifs de la rationalité limitée et en les intégrant aux connaissances existantes en sciences organisationnelles. Une question générale qui motive ce travail est la façon dont les organisations peuvent gérer les nombreuses limites de la rationalité humaine.

Title: Essays in Behavioral Strategy: Re-biased Search, Misconceived Complexity, and Cognitive Aliens

Keywords: Behavioral strategy, cognition, organizational adaptation

Abstract: This work centers on the tenet that organizational rationality is bounded: decision makers search, satisfice, and think in a way that is typical (in its integrity) only of humans. The dissertation explores this interplay between search and decision maker's cognition and demonstrates how biases in characteristic aspects of our thinking can be instruments of behavioral strategy.

As a starting point, I take search, sequential generation and evaluation of alternatives, as the first primitive of bounded rationality and complement it with integral elements of human cognition, such as automatic, intuitive thinking, specifically affect heuristic, and imperfect mental representations of reality.

With the help of computational models, I track the effects of the corresponding biases (systematic affective preferences and systematic errors in mental representations) over time as organizations adapt to complex environments. This allows me to identify life cycles of the elements of human cognition and show that organizations should manage (rather than eliminate) some biases over time. Finally, I derive predictions and empirically test a subset of my propositions.

In conclusion, this work aims to advance the emerging theory of behavioral strategy by jointly considering different primitives of bounded rationality and integrating them with the existing knowledge in organization sciences. A broad question that motivates this work is how organizations can manage the many bounds to human rationality.