

Experiential Learning under Ambiguity

Thorsten Wahle

Institute of Management and Organization

Faculty of Economics

Università della Svizzera italiana

A dissertation submitted to the
Faculty of Economics of the Università della Svizzera italiana
in partial fulfillment of the requirements for the degree of
Doctor of Philosophy
under the direction of Prof. Dirk Martignoni

Lugano, March, 2020

ACKNOWLEDGMENT

I want to thank the Institute of Management and Organization at USI for the supportive and scholarly environment. I am indebted to everyone there for their feedback and friendship.

I thank my advisor, Professor Dirk Martignoni, for his guidance. I learned a lot from him in the past five years - especially in being more concise in my writing and efficient in approaching new tasks. I also thank Nikolaus Beck, Claus Rerup and Zur Shapira for agreeing to serve on my dissertation committee.

Further, I am very grateful for the opportunity to visit NYU – my time there has given me countless insights on the research process and truly made me enthusiastic about the prospect of academia. I'm particularly thankful to Professor Filippo Wezel for supporting my application, to Professor Zur Shapira for hosting me and working with me at NYU, and to Professor J.P. Eggers for becoming my co-author. I continue to learn from them, and am inspired by their insights, and way of working. Through the many interactions with Professor Shapira in particular, I gained much clarity on my research identity, the state of the field and the people involved in it.

Lastly, I want to thank my wife, Sarah Edris, and my parents for their unconditional love, support and advice during these last years!

ACKNOWLEDGMENT.....	2
LIST OF TABLES	4
LIST OF FIGURES	5
1. INTRODUCTION.....	7
2. THE SOCIAL COMPARISON TRAP	10
INTRODUCTION.....	10
LITERATURE REVIEW	14
METHOD.....	17
ANALYSIS AND RESULTS.....	22
DISCUSSION.....	31
3. LEARNING FROM OMISSION ERRORS.....	40
INTRODUCTION.....	40
THEORY AND HYPOTHESES	44
METHOD AND DATA.....	49
ANALYSIS.....	54
RESULTS.....	55
DISCUSSION.....	61
4. THE BOILED FROG IN ORGANIZATIONAL LEARNING.....	65
INTRODUCTION.....	65
THEORY.....	68
GENTRIFICATION AS ENVIRONMENTAL CHANGE.....	73
METHOD.....	75
RESULTS.....	79
SURVIVAL ANALYSIS	85
SIMULATION EXPERIMENT	89
DISCUSSION.....	101
5. DISCUSSION	105
DISSERTATION SUMMARY	105
CONTRIBUTION TO THEORY	106
LIMITATIONS AND FUTURE RESEARCH.....	108
6. REFERENCES.....	110
7. APPENDICES	121
APPENDIX CHAPTER 2 A	121
APPENDIX CHAPTER 2 B - EXPERIMENT	125
APPENDIX CHAPTER 2 C.....	141
APPENDIX CHAPTER 3 A	144
APPENDIX CHAPTER 3 B.....	145

List of Tables

Table 2-1: Treatment conditions chapter 2	20
Table 2-2: Performance implications of search errors and percentage of participants committing them	28
Table 3-1: Main results (full complexity)	55
Table 3-2: Reduced complexity results	57
Table 3-3: Split samples based on social aspirations	58
Table 3-4: Attention to competitor	60
Table 4-1: Gentrification and restaurant lifespan.....	80
Table 4-1a: Gentrification and restaurant lifespan.....	83
Table 4-2: Mediation analysis (lifespan and adaptation)	85
Table 4-3: Significance test between survival rates in Figure 4-7	87
Table 4-4: Mediation analysis	89
Table 7-1: Car features	144

List of Figures

Figure 2-1: Performance effect of social comparisons.....	22
Figure 2-2: Social comparison effects on search behavior	23
Figure 2-3: Level of exploration for 5 alternatives (left panel) and 100 alternatives (right panel).....	25
Figure 2-4: Share of participants never searching (left panel) never stopping to search (right panel).....	26
Figure 2-5: Share of participants searching beyond the best decision alternative	27
Figure 2-6: Level of exploration with a fixed sequence of alternatives.....	30
Figure 2-7: ANOVA of performance differences across the four conditions	31
Figure 3-1: Example decision screen	52
Figure 4-1: Number of active restaurants over the years	76
Figure 4-2: Number of active restaurants across boroughs.....	76
Figure 4-3: Distribution of menu change rates.....	78
Figure 4-4: Inverted U-shape of the effect of entry rate of bars on restaurant lifespan	81
Figure 4-5: Inverted U-shape of the effect of entry rate of cafés on restaurant lifespan	81
Figure 4-6: Inverted U-shape of the effect of the rate of eviction filings on restaurant lifespan	82
Figure 4-7: Differences in survival rates between restaurant exposed to gentrification at fast (strata=0) versus slow rates (strata =1)	87
Figure 4-8: Differences in survival rates between restaurant in relatively declining neighborhoods	88
Figure 4-9: Sensing accuracy as function of degree of continuous change	93
Figure 4-10: Sensing accuracy moderated by sensing threshold	94
Figure 4-11: Sensing accuracy moderated by learning rate	95
Figure 4-12: Performance implications of different degrees of continuous change	96
Figure 4-13: Performance deviation from discontinuity with increasing level of continuous change.....	96
Figure 4-14: Performance comparison with averaging heuristic	98
Figure 4-15: Performance comparison with regression heuristic.....	98
Figure 4-16: Performance comparison with minimum heuristic	99
Figure 4-17: Performance comparison with averaging-threshold.....	99
Figure 4-18: Performance comparison with minimum-threshold heuristic	100
Figure 4-19: Performance comparison with peak-end-threshold heuristic	101
Figure 4-20: Performance comparison with peak-end heuristic	101
Figure 7-1: Optimal task behavior	122
Figure 7-2: Optimal behavior using a 2/3 target heuristic	124
Figure 7-3: Instruction screen 1	125
Figure 7-4: Instruction screen 2	126
Figure 7-5: Instruction screen 3	127
Figure 7-6: Decision screen first period.....	128
Figure 7-7: Decision screen mid-task.....	129
Figure 7-8: Attention check.....	130
Figure 7-9: Final task screen	131
Figure 7-10: Comprehension questions 1.....	132
Figure 7-11: Comprehension questions 2.....	133
Figure 7-12: Additional questions 1	134
Figure 7-13: Additional questions 2.....	135
Figure 7-14: Changed introduction treatment 1	136

Figure 7-15: Changed decision screen treatment 1	137
Figure 7-16: Manipulation check treatment 1	138
Figure 7-17: Changed decision screen treatment 2	139
Figure 7-18: Changed decision screen treatment 3	140
Figure 7-19: Performance comparison between managers and non-managers.....	141
Figure 7-20: Performance comparison between age groups	142
Figure 7-21: Performance comparison between levels of education	143
Figure 7-22: Performance comparison between levels of income	143
Figure 7-23: Instructions omissions experiment	145
Figure 7-24: Instructions omissions experiment (cont'd).....	146
Figure 7-25: Omissions post-experiment screens	147
Figure 7-26: Omissions post-experiment questionnaire.....	149
Figure 7-27: Omissions treatment information – competitor treatment.....	150
Figure 7-28: Omissions treatment information – full information treatment.....	151

1. INTRODUCTION

Learning from experience is crucial for adaptation and to improve performance. Experiential learning takes place by making choices, observing the outcomes, and altering future choices accordingly (Levitt and March 1988 March and Olsen 1975). Over time, such learning is expected to improve choices and outcomes (Lant and Mezias 1992), but only when decision makers can classify outcomes as either successes or failures. This dichotomous outcome has clear behavioral implications - change after failure or persist after success (Greve 2002).

Such classification may be complicated when experience is ambiguous, i.e. a choice may not always produce an outcome that can be readily classified as success or failure (Rerup 2006). In particular, outcomes are ambiguous when they allow for classification into more than one category (Repenning and Sterman 2002). Prior literature has examined how learning from ambiguous outcomes may occur through attention and mindfulness (Levinthal and Rerup 2006). However, different choice contexts are likely to differ in the type of ambiguous experience they create (Weick and Sutcliffe 2006). It is therefore important to understand when learning under ambiguity is possible. Such choice contexts include situations when (1) decision makers do not have enough prior experience to determine a sensible threshold between success and failure (i.e. they do not have enough data to formulate historical aspirations), (2) choices do not produce observable outcomes, or (3) it is unclear if the decision maker's choice or external factors are causing the observed outcome. In this dissertation, I aim to study these impediments to experiential learning, particularly when decision makers perceive outcomes as ambiguous.

I explore three choice contexts of learning under ambiguity using experiments and a large dataset on restaurants in New York affected by gentrification as well as short simulation experiments.

In chapter 2 (study 1, co-authored with João Duarte and Dirk Martignoni) we aim to identify when social comparison improves or reduces search efficacy in contexts of insufficient prior experience. Social comparison is common in organizations and takes the form of pay transparency, ratings and rankings, among others.

In our study, we seek to enhance our understanding of the implication of this kind of information for search and, in turn, performance. Using a series of online experiments, we identify conditions under which peer performance information may have a positive or negative effect on performance. We identify the number of potential choice alternatives as a novel moderator affecting the performance implications of peer performance information. Our study also sheds new light on the mechanism through which peer performance information may affect (search) performance: positive performance effects only arise if peer performance suppresses rather than induces search.

In chapter 3 (study 2, co-authored with J.P. Eggers) we turn to examining choices which do not produce observable outcomes. Specifically, we are interested in learning from omission errors. Within organizations, the ability to learn from mistakes in general is central to performance improvement and adaptation. But different types of errors importantly produce different levels of observable feedback – commission errors generally produce direct feedback, but omission errors often do not. Thus, in order for managers to learn from omissions they must have the ability to know the outcomes of choices not pursued. One key source of such information comes from observing competitors, but attending to competitors may in itself affect learning and adaptation. In a series of experiments, we study whether decision makers learn from feedback provided by observing competitors. We argue and find that decision makers' ability to learn from their decision errors depends on their position relative to the competitor. Specifically, leaders tend to learn from their omission errors because they actually pay attention to competitor decisions, but laggards tend to learn only

from commission errors. This has implications for our understanding of learning from failure and organizations' attempts to learn from competitors.

In chapter 4 (study 3, co-authored with Zur Shapira) we explore how learning is inhibited by noisy feedback on outcomes and specifically, how the pace at which the environment is changing impacts organizational adaptation and survival. Such change is often viewed as discontinuous. We introduce the concept of continuous radical change by which old beliefs are rendered fallacious at an incremental pace. We theorize that incremental change can be as harmful as discontinuous change because it can remain undetected for too long by incumbent organizations. Specifically, we argue for a curvilinear relationship between the pace of environmental change and the organizational lifespan. We test our predictions using data from some 31,000 restaurants in the New York City Metropolitan area between the years 2007-2018. We use a fixed effects regression model to test for differences in the lifespan of restaurants located in areas of slow but radical gentrification versus those located in areas with accelerated and radical gentrification. We find support for an inverted U-shape between the pace of gentrification and restaurant lifespan. We explore whether this relationship is driven by differences in restaurants' adaptation behavior.

Overall, my dissertation makes two main contributions that distinguish it from prior research. Theoretically, I identify choice contexts that produce ambiguous outcomes, and discuss under which conditions decision makers should attempt to learn from such outcomes. Empirically, I develop a novel experimental design to measure search and learning. I complement the experiments with a unique dataset of gentrification and restaurants in New York that allows for the study of exogenous environmental change.

2. THE SOCIAL COMPARISON TRAP

Introduction

Social comparison among employees is common across all kinds of organizations (Adams 1963, Festinger 1954, Nickerson and Zenger 2008). Social comparisons often arise naturally whenever employees can observe their peers' performance (Kacperczyk et al. 2015, Nickerson and Zenger 2008). At the same time, organizations may also seek to promote social comparisons through, for example, providing their employees with performance rankings and ratings (Greve 2003b, Greve and Gaba 2017). Manipulating the availability of this peer performance information is an important aspect in designing the "social architecture" of organizations (Lee and Puranam 2017, Nickerson and Zenger 2008). In so doing, organizations hope to improve the efficacy of their employees' search processes (Baumann et al. 2018, Greve and Gaba 2017). Most importantly, providing information about better performing peers may trigger search, thereby helping to overcome problems of inertia and overexploitation (Chen and Miller 2007, Greve 2003a, Greve and Gaba 2017).

Though intuitively appealing, empirical research has produced mixed findings on the performance implications of social comparisons. Some studies point to positive performance effects (Blanes i Vidal and Nossol 2011, Stark and Hyll 2011) while others find negative performance effects of social comparisons (Nickerson and Zenger 2008, Obloj and Zenger 2017, Pfeffer and Langton 1993). Together, these findings suggest that the link between social comparison and performance outcomes is under-theorized, in particular that potential moderating effects are neglected.

In our study, we seek to refine our understanding of how and when social (performance) comparisons may affect the efficacy of search processes and, in turn,

performance¹. Specifically, we conducted a series of repeated-choice experiments in which we manipulated the availability of information about how peers performed in these experiments. We find that social comparison may improve or hamper the efficacy of search, depending on the type of search problem they face. Interestingly, and in contradiction to previous arguments (Chatterjee et al. 2003, Greve 2003b), positive effects of social comparison arise if the availability of peer performance information suppresses (rather than spurs) search and exploration. Negative effects, in contrast, arise if peer performance information induces exploration, specifically in contexts where overexploitation seems to be the key problem, i.e. when they face a multitude of decision alternatives. In sum, in refining and extending the scope of existing search-based theories, we seek to explain both the bright and the dark sides of social comparison. We identify boundary conditions for the positive and negative effects of social comparison (induced through peer performance information) in terms of the type of problem searched and in particular, the number of decision alternatives.

Our study contributes to the existing literature in several ways. First, we refine existing search-based theories to explain the emergence of both positive and negative effects of social comparison and identify boundary conditions for these opposing findings. While we are not the first to point to both positive and negative effects of social comparison, our theory allows explaining both positives and negatives within one theoretical framework, without resorting to adding other mechanisms such as envy (Mui 1995, Nickerson and Zenger 2008), effort (Blanes i Vidal and Nossol 2011, Levine 1993, Stark and Hyll 2011), or motivation (Gubler et al. 2006). For example, while it is often argued that social comparisons may have positive effects on search, they may also lead to envy among employees (Nickerson and Zenger 2008, Van de Ven 2017) and other asocial behavior (Charness et al. 2013, Chan et al.

¹ In our study, any differential effects of social comparison on performance are driven entirely by differences in search behavior rather than psychological factors such as envy or motivation.

2014) as well as a reduced willingness to cooperate (Colella et al. 2007, Dunn et al. 2012, Nickerson and Zenger 2008). As a result, the overall effect of social comparisons may be negative. In our study, we point to the fact that even if we focus on the implications for search, social comparisons may have positive or negative implications, depending on the size of the choice set decision makers face. When facing a large choice set of many alternatives, decision makers are tempted to search for too long in their attempts to explore at least some share of the set of the available alternatives. Such attempts require searching more choices compared to smaller choice sets (i.e. fewer decision alternatives). As a result, they may search for too long as the optimal stopping point for search is only a function of the time horizon rather than the size of the choice set. Thus, when facing large choice set and provided with peer performance information, decision makers are lured into searching for too long.

These findings also contribute to social aspiration theory (Cyert and March 1963, Greve 1998, Greve 2003b, March 1994) and stretch goal theories (Hamel and Prahalad 1993, Latham 2004, Kerr and Landauer 2004, Sitkin et al. 2011), intellectual offspring of social comparison theory. According to these theories, the benefits of social comparisons stem from the fact that observing better performing peers induces change (Audia et al. 2000, Lant 1992, Miller and Chen 1994), exploration (Baum and Dahlin 2007), and risk taking (Bromiley 1991). After having observed better-performing peers, the status quo appears no longer satisfactory (Simon 1956) and (problemistic) search processes are triggered. Our study points to a novel mechanism of how social comparison may affect search processes. Peer performance information may serve as a substitute for exploration and may help organizations put into perspective the performance of alternatives it already tried and knows. Thus, social comparison may both trigger exploration when performance can improve and suppress exploration when performance is satisfactory. It can only do so, however, when decision alternatives are few. When decision alternatives are many, observing high peer

performance may spur exploration when performance can improve but fail to curb exploration when performance is satisfactory.

Second, we discuss a novel mechanism by which social comparison affects performance outcomes. Prior research in the tradition of the behavioral theory of the firm highlights the role of social comparison as a means to overcome problems of inertia and under-exploration (Audia et al. 2000, Rhee and Kim 2014). Accordingly, social comparisons are thought to mitigate these issues by inducing search (Chen and Miller 2007, Greve 2003a, Greve and Gaba 2017). Our study points to a novel mechanism through which social comparisons may improve the efficacy of search processes - not by inducing but by suppressing search. If peer performance information is not available, decision makers tend to search for some time to determine what might constitute a satisfactory performance outcome. Imagine a situation in which a decision maker runs into the best alternative in his first try. Without any peer performance information (or information about the distribution of the other alternatives' values), she has to continue searching other alternatives in order to establish the superiority of this first alternative. Now, imagine a situation in which peer performance information is available. If peer performance information is available, it is much easier to recognize her first choice as satisfactory and thus she can avoid the (opportunity) costs of further exploration. In more abstract terms, in searching among a set of alternatives, decision makers seek to acquire two different kinds of information: first, information about what may constitute a better alternative; second, information about the distribution of performance and what may constitute a satisfactory level of performance. Peer performance information is particularly helpful in overcoming the second information problem.

Lastly, we contribute to the organization design literature (Puranam et al. 2012, Nickerson and Zenger 2008). In recent years, new forms of organizations have emerged, in particular meta-organizations that are characterized by open boundaries (Gulati et al. 2012).

Likewise, organizations face pressure to become more open, granting employees greater participation and transparency for example through pay transparency (Gartenberg and Wulf 2017). For employees, this new openness means both more opportunities to engage in social comparison (because outcomes of peers are now more available) and more decision alternatives to choose from (because employees are involved at a higher level of problem solving). Our study considers the co-existence of social comparison and a large number of decision alternatives. Our findings suggest that while social comparison may in theory be beneficial to decision makers, in open organizations with many decision alternatives, it may lead to excessive search for new solutions, thus ultimately hurting performance.

The remainder of our study is structured as follows. We first review the theoretical building blocks of our core argument. We then discuss the design of our experiment, followed by the presentation of our results. We conclude by discussing the implications of our results for existing research.

Literature Review

Social comparison in organizations

Social comparison is a ubiquitous phenomenon in social life (Festinger 1954). Social comparison is commonly understood as comparing own performance to the performance of a peer (Obloj and Zenger 2017). In organizations, employees often compare their performance to those of their fellow colleagues (Nickerson and Zenger 2008).² Organizations may seek to encourage social comparison by making peer performance information available (Nickerson and Zenger 2008, Kacperczyk et al. 2015, Gartenberg and Wulf 2017). For example, within

² However, social comparison is not limited to outcome comparisons among individuals (Festinger, 1954; Nickerson and Zenger, 2008). Business units compare themselves to other business units (Baumann et al., 2018; Birkinshaw and Lingblad, 2005; Galunic and Eisenhardt, 1996) and organizations compare their performance against other organizations (Greve, 1998; Greve 2003b; Joseph and Gaba, 2015).

(and across) banks, there is often complete transparency how fellow fund managers and traders performed (Castellaneta et al. 2017, Kacperczyk et al. 2015, Uribe 2017); call center agents are made very well aware of their performance compared to their peers (Neckermann et al. 2009); in many firms such as GE, rankings and ratings become an important part of their culture (Short and Palmer 2003). Finally, many firms are adopted policies of pay transparency (with pay as a proxy for performance) (Gartenberg and Wulf 2017).

While all these means provide more or less accurate information about how peers performed, they often provide far less information about how they achieved this performance, i.e., their peers' behavior and actions. This makes social comparison different from imitation and isomorphism, which requires observing not only peers' performance but also their behavior (Mizruchi and Fein 1999).

Performance Effects of Social Comparisons

Organizations often encourage social comparisons, hoping that such comparisons improve performance (Lee and Puranam 2017, Nickerson and Zenger 2008). As a consequence of social comparison, lower performing employees will realize that they could improve their performance (Festinger 1954) and thus, start searching for new solutions (Jordan and Audia 2012, Simon 1947). This follows the logic of problemistic search (Cyert and March 1963): Decision makers first engage in social comparison and take their peer's performance as a performance target. They then observe the outcome of their current solution and evaluate it against this target. When they fall short of reaching the performance target, they search for new solutions to improve performance. Exceeding the target, on the other hand, leads to reduced search efforts as performance is deemed satisfactory and opportunity costs would arise from continued search (Levinthal and March 1993, March 1991). Since organizations are often thought to be inert, i.e., their members search too little (Audia et al. 2000, Rhee and

Kim 2014), social comparison is assumed to help overcome inertia by prompting individuals in organizations to search more.

While intuitively appealing, there are several empirical studies that point to negative performance implications of social comparison. Some empirical studies find that social comparisons (encouraged by pay transparency or tournaments) may have negative performance implications (Pfeffer and Langton 1993, Obloj and Zenger 2017). This finding is confirmed by both experimental (Gino and Pierce 2009, 2010, Dunn et al. 2012) and anecdotal evidence (Nickerson and Zenger 2008).

In sum, the mixed theoretical and empirical evidence on the implications of social comparisons suggests that the link between social comparison and performance may still be under-theorized and that important moderating effects may be neglected. Below, we discuss one important moderator variable – the number of decision alternatives a decision maker faces.

The number of decision alternatives

The number of alternatives is an important characteristic of any decision problem (Simon 1955). Decision-makers face problems with varying numbers of decision alternatives (Payne 1976). Some decision problems are characterized by few decision alternatives. Take, for example, the procurement decision of an airline. When it comes to the decision from whom to buy a new wide body airplane, there are often only two choices – either Airbus or Boeing. For other procurement decisions, there are many more decision alternatives. For example, for its choice of supplier for aircraft seats, an airline can choose from +100 different manufacturers.

Organizations may further seek to manipulate the size of their employees' choice set as one element of designing an organization's "social architecture" (Nickerson and Zenger 2004). For example, organizations may seek to increase their set of choices by engaging in

open innovation (Terwiesch and Yi 2008) and crowdsourcing (Dahlander et al. 2018). Organizational rules and practices, in contrast, often reduce the size of the choice set (Cohen 1991). As a result, organizational structures “consider only limited decision alternatives” (Cyert and March 1963, p. 35).

Recently, the number of decision alternatives a decision maker within an organization faces, has increased as a consequence of greater organizational openness (Alexy et al. 2017, West 2003). This is apparent in new forms of organizations like meta-organizations characterized by open boundaries (Gulati et al. 2012). In such organizations, openness means a greater involvement of employees at a higher level of problem solving in the organization (Mack and Szulanski 2017). An organization solves big problems with many alternative solutions by dividing problems into smaller sub-problems with fewer alternative solutions that can be solved at a lower hierarchy level (Nickerson and Zenger 2004). In more open organizations, where employees are involved in the problem-solving process at a higher level, this thus means that employees search through a greater numbers of decision alternatives. It is thus the same force that drives social comparison (because open organizations are inherently more transparent) that increases the number of decision alternatives.

Method

Experimental framing

In our experiments, participants face a simple repeated (five periods) choice problem in which they face a set of five alternatives with different values. These values were not known to the participants. Only by choosing a particular alternative, they could learn about its value: in each period, after they have chosen a particular alternative, its value is revealed. The value of the alternative was drawn at random from a set of five values. In the first

period, participants have to explore one alternative at random since they had no prior information on the value of any alternative ahead of the task. From period two to period five participants could either choose an alternative whose value is known to the participants because they explored it in a past period or explore an alternative of unknown value. Feedback on the values was given without delay or any deception, i.e., the value revealed corresponds to the alternative's true value.³ We framed the experiment as a managerial choice task in which the participants (in the role of the manager) have to maximize their payoffs in US dollars (\$).

Treatment 1 (“Peer Performance Information”). First, to allow for social comparison, we provided them with information about how a peer performed in the very same task (“Treatment 1: Peer Performance Information”): *“A consulting study concludes that, with the right strategy, your company’s profits could be as high as \$121 million a year”*. We deliberately chose to frame it not as peer performance information (*“Your best performing competitor generated profits of \$121 million in his best year”*) for several reasons. Our pre-tests indicated that participants often continued searching for even better alternatives even after they had discovered what we told them to be the highest value achievable in this task (here, \$121 million). Likely reasons for this type of behavior are first, that participants assumed that their peer’s performance was a function of both their choices and effort, i.e. they would have continued searching because they thought they can do even better than their peer. And secondly, participants may have assumed that if several people simultaneously chose the very same alternative, its value or payoff will have to be split between them. To

³ In each period (after the first period), participants had to decide whether to explore a new alternative of unknown value or exploit an alternative of known value. To avoid confounding effects of, for example, memory constraints, we narrowed down the actual choices participants had to make between choosing “exploiting the best alternative explored so far” and “exploring an alternative never explored so far”. With such a setup, we avoid, that participants exploit inefficiently (i.e. they do not exploit the best-explored alternative) or re-explore alternatives that they had already explored in the past. In pre-tests of this experiment, we also tested a setup in which the choice was between 5 alternatives in each period, i.e. 5 alternatives were shown in a line and one of them would have to be picked. This led to more random choices of the participants, but had no effect on our findings.

avoid these confounding effects and isolate the informational effect of peer performance information, we presented the treatment information in a neutral way that would not invoke any feelings of rivalry (Kilduff et al. 2010).

Treatment 2 (“Large Choice Set”).

In this treatment group, participants faced a much larger choice set, i.e., 100 rather than just five alternatives. We adapted the instructions accordingly (“You will have to [...] choose among 100 different strategies”). This information was reinforced in every period while participants were informed that “You have 100 [99,98,97,96] alternatives left to explore”. Every time a new alternative was searched, its payoff was drawn at random from a set of 100 alternatives. Since our participants always only have five periods to search within their choice set; they could never search exhaustively. Thus, in our experiment, we only indicate that they face a larger choice set without changing the experimental task they face.

Treatment 3 (“Peer Performance Information & Large Choice Set”).

Finally, our last treatment combines treatment one and two. Participants had access to peer performance information and their choice set was large. This means, participants were shown the peer performance information (“*A consulting study concludes that, with the right strategy, your company’s profits could be as high as \$121 million a year*”) as well as informed that they face a choice set of 100 alternatives (“*You will have to [...] choose among 100 different strategies*”).

The study was conducted in a between-subject design. This means each participant was randomly assigned to one of the 4 conditions: control, peer performance information, large choice set, peer performance information and large choice set. In Table 2-1, we show an overview of all four conditions.

Table 1. Overview of treatment conditions			
		<i>Size of Choice Set</i>	
		<i>Small</i>	<i>Large</i>
<i>Peer Performance Information</i>	<i>Absent</i>	Control	Treatment 2
	<i>Available</i>	Treatment 1	Treatment 3

Table 2-1: Treatment conditions chapter 2

Experimental setup

Incentives. We relied on the induced-value approach (Smith 1976), i.e., we provided monetary incentives to the participants such that better choices translated into higher payoffs to participants. In keeping with the tradition of economic experiments, all information provided to the participants was truthful (Ortmann and Hertwig 2002). Before starting the experiment, participants were told that the experiment would take them around 5 minutes to complete and that their payment will be based on the payoffs they accumulate in this task. Each choice was associated with a particular payoff (e.g. “Period 1: The chosen alternative generated \$55m”). We chose payoffs in this experimental currency to provide higher incentives to participants. In particular, prior studies have found that participants consider the nominal (the value given in experimental currency) rather than the real payoff value (Davis and Holt 1993) and thus, behave more like they would in a real organizational setting. After each choice, the choice’s payoff was revealed. We also always showed them the profits they had accumulated so far.

We also informed our participants that upon completion, there would be a certain payment of \$0.10 and up to \$0.20 depending on their performance in the task, i.e. the cumulative profits associated with their choices. One might be tempted to conclude that such low stakes may result in some confounding effects such as excessive risk taking (Lefebvre et al. 2010, Weber and Chapman (2005), search, or lack of effort (Smith and Walker 1993). We do not observe such confounding effects in our experiments: First, if risk taking was higher due to small stakes, we would expect to see all treatment conditions affected equally.

Furthermore, we would expect to see very high levels of search across all treatments. We do not observe such effects – there are no differences in search between control and treatments and we also do not observe excessive search behavior. Also, prior research does not indicate a significant change in behavior if stakes are higher (Camerer and Hogarth 1999, Post et al. 2008). For instance, even if stakes are very high, participants commit decision errors to a similar extent as in low stake contexts (Post et al. 2008). Fourth, low payments are often thought to result in lower effort. In our context, however, effort does not matter. It is a pure choice task and all aspects that require some effort are designed in the way to minimize cognitive efforts by participants. For example, we always provide participants with a complete history of their past choices and their payoffs; we also always highlight the best alternative they have identified so far and make sure that inefficient exploitation is impossible. Generally, we would expect to see extremely monotone behavior, i.e., many participants only exploring or only exploiting throughout the task if they would generally exhibit a lack of effort. We find no such effect. Further, we included a manipulation check to control for a lack of (cognitive) effort. After participants finished the task, we asked them to report back the treatment information. Our findings are robust to the exclusion of those participants that did not accurately recall the peer performance information. In any case, the treatments effects found including those that did not accurately recall the peer performance information, provide a more conservative estimate of the effect.⁴

Respondents. Participants were recruited on Amazon MTurk. All treatments included, 857 subjects completed the experiment. 61.8 percent of the participants were women, the average age was 37 years and 27.8 percent claimed to hold a managerial or supervisory position at their workplace. Participation was restricted to the United States to avoid language difficulties. On average (median duration), participants completed the

⁴ Around 50% of participants accurately recalled the peer performance information.

experiment in 4 minutes and 38 seconds and spent 10 seconds on each of the four individual choices. Those who did not spend any time on the instruction pages or failed our attention check were excluded from the sample. All participants were randomly assigned to one of the 4 treatment groups (control, peer performance; each for a small and large choice set).

Analysis and Results

The performance implications of social comparison

In our first experiment, we are interested in how the availability of peer performance information may affect our participants search behavior and, in turn, performance. In Figure 2-1, we report the average accumulated (over all five periods) payoffs of our control and treatment groups.

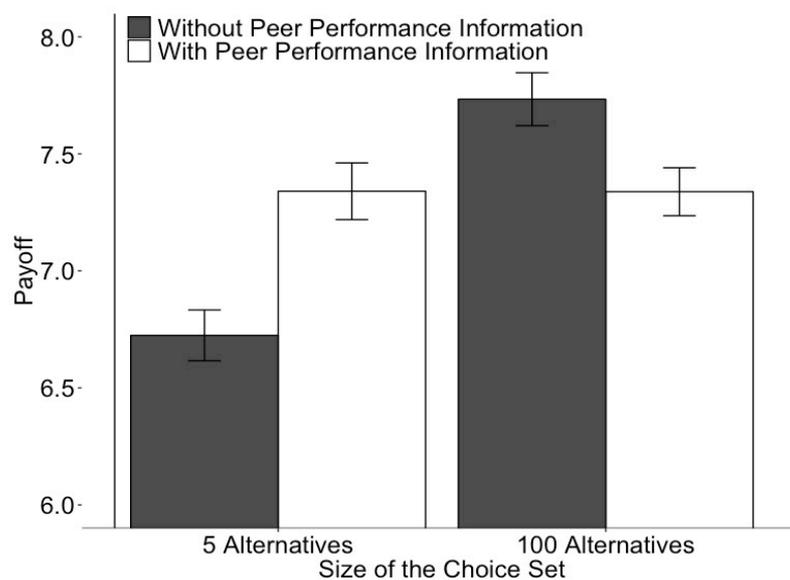
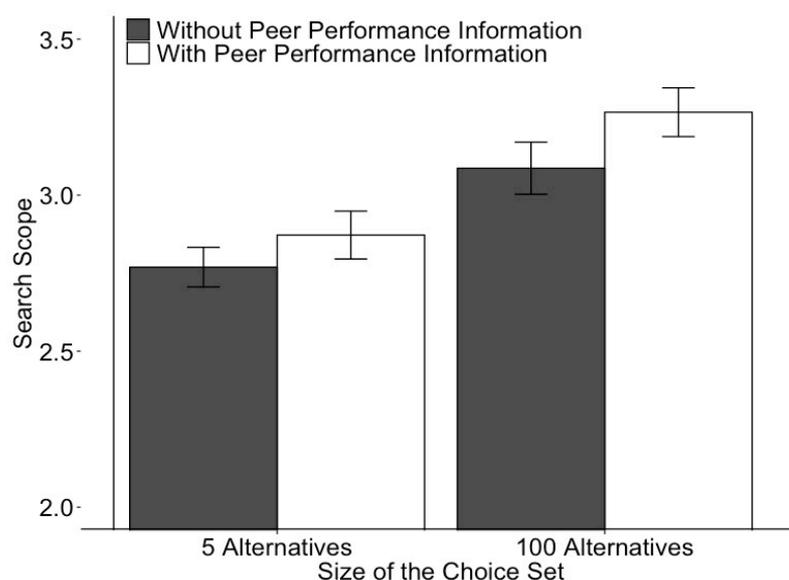


Figure 2-1: Performance effect of social comparisons

If the choice set is small, the availability of peer performance information has a significantly positive performance effect (MD=0.62, p=0.000). In terms of the effect size, this meant that participants earned a 9.16 percent higher payoff in the treatment condition

compared with control. If, however, the choice set is large, we observe the opposite effect, i.e., peer performance becomes costly (MD=-0.40, p=0.010). If peer performance information was available, participants earned a 5.12 percent lower payoff in the treatment condition when compared with the control condition. In other words, depending on the size of the choice set, peer performance information may have a positive or negative performance effect. It is important to note that these differential effects are not driven by differences in order in which alternatives were searched. The order in which participants could search alternatives is random and each order is equally likely.

In search problems, underperformance is often attributed to insufficient levels of exploration (Baum and Dahlin 2007, March 1991) and social comparisons are considered a remedy, i.e., they induce search (Cyert and March 1963, Greve 1998). Thus, we would expect that the positive performance effects of peer performance information for small choice sets are driven by higher levels of exploration while the negative effects for larger choice sets may be driven lower levels of exploration. In Figure 2-2, we compare the effect of peer performance information on how many alternatives participants explored in small and large choice sets.



Note. Search scope is measured by periods of exploration.

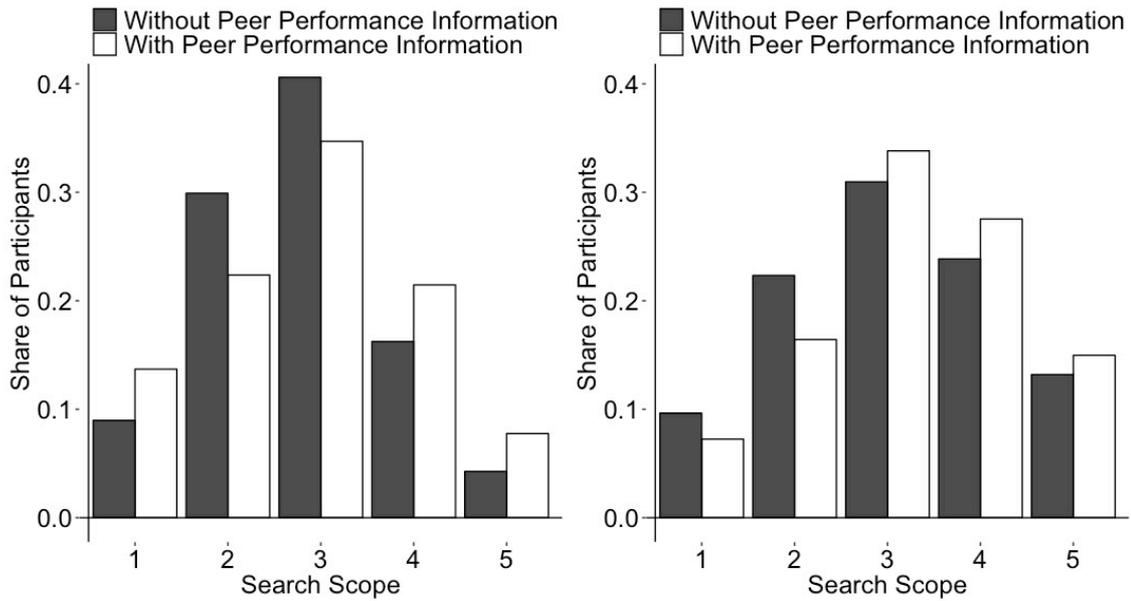
Figure 2-2: Social comparison effects on search behavior

With both small and large choice sets, peer performance information induces search: When facing a small choice set, the availability of peer performance information increases search by 3.7% but this increase is not significant (MD = 0.10, $p = 0.302$). For large choice sets, the increase (by 5.8%) becomes pronounced and more significant (MD = 0.18, $p = 0.118$). Thus, inducing search through social comparison cannot explain both the positive and negative performance effects of peer performance information.

It is sometimes overlooked that peer performance information may not only *induce* search but may also *suppress* further search. Imagine that, by chance, a participant discovers the best alternative with her first choice. Without peer performance information (and thus, without knowing that it is already the best alternative), the optimal behavior is to continue searching. With peer performance information, however, the participant knows that her current choice is already the best alternative and there is no point in searching for better alternatives. She can abort her search.

Using peer performance information to abort search once the best solution is found is always performance enhancing. Yet, in larger choice sets, it is much less likely that any benefits from aborting search can be reaped because it is much less likely that the best alternative is encountered within the first 1-2 periods. With a smaller choice set, in contrast, the probability is high that within the first 1-2 periods, the best alternative is encountered and search can be aborted.

In Figure 2-3, we seek to provide some evidence for the different ways (induce or suppress) peer performance information affects our participants search behavior. Specifically, we disentangle the average search scope and report the share of participants exploring 1, 2, 3, 4, or 5 alternatives. In the left panel, we show the results for small choice set; in the right panel, we show the results for large choice sets.



Note. The measure is the share of participants stopping exploration after trying out X alternatives in percent.

Figure 2-3: Level of exploration for 5 alternatives (left panel) and 100 alternatives (right panel)

With small choice sets, peer performance information both suppresses and induces search: the share of participants aborting search after the first period (presumably because – by chance (20%) – they run into the best alternative in in $t=1$) increases from 8.9% to 13.7%. At the same time, the share of participants continue searching (presumably because they did not run into the best alternative until period 5) for all 5 periods increases form 4.3% to 7.8% (see Figure 2-4).

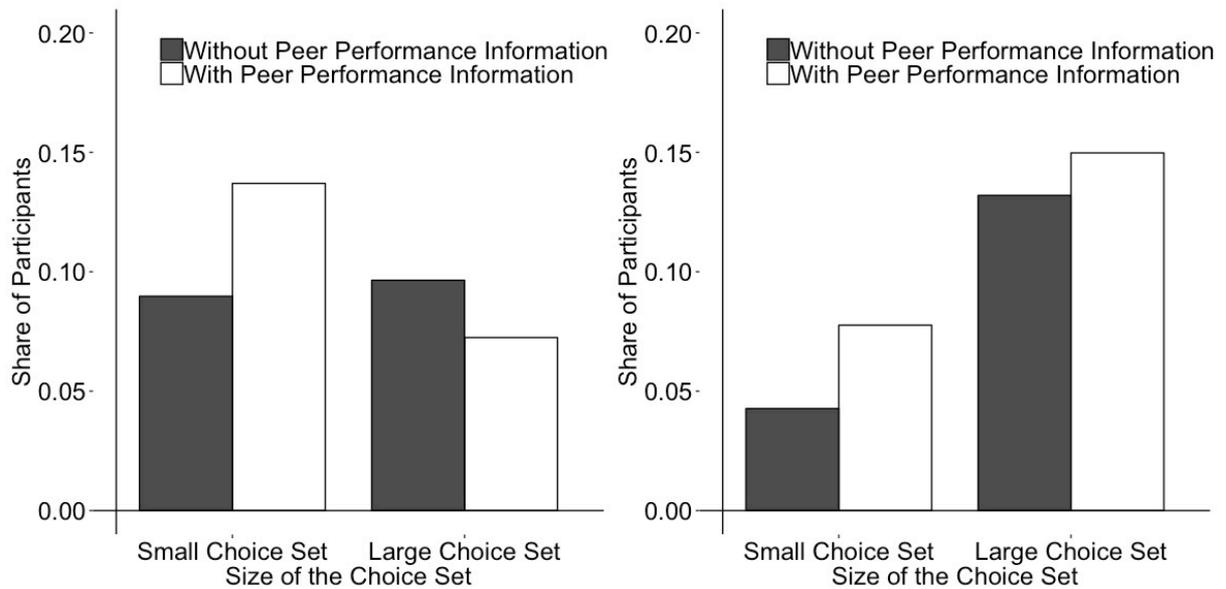


Figure 2-4: Share of participants never searching (left panel) never stopping to search (right panel)

With larger choice sets, the probability of running into the best alternative in the first two periods is much smaller (only 2% compared with 40% in small choice sets). As a result, peer performance information does not help much to suppress exploration. In fact, the share of participants aborting search after the first period decreases from 9.6% without peer performance information to 7.2% with peer performance information. At the same time, the share of participants never aborting search still increases with the availability of peer performance information (from 13.2% to 14.9%). Instead, there is a shift towards exploring more alternatives, even to the extent that it becomes very costly (search scope=5) (as can be seen in Figure 2-4).

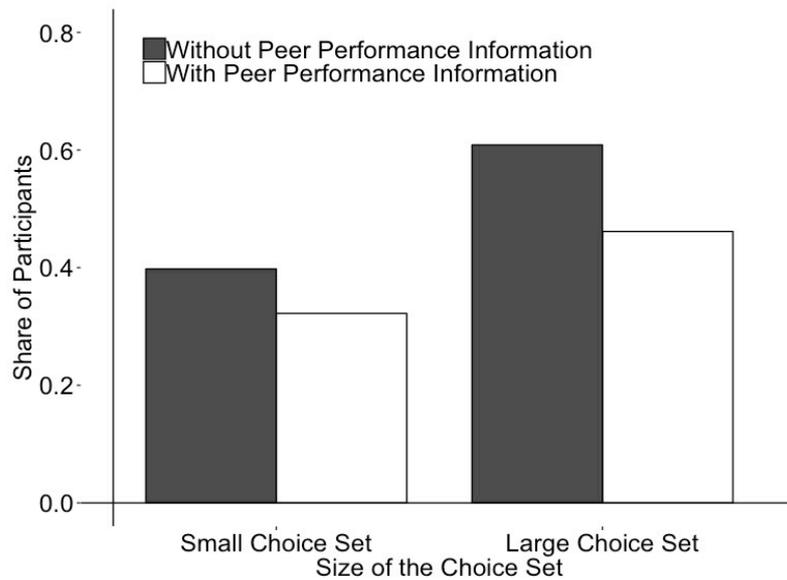


Figure 2-5: Share of participants searching beyond the best decision alternative

Inefficient search behavior and peer performance information

In our analyses above, we focused on the question of how our treatments affect *how many* alternatives (or their search scope) our participants search. These treatments, however, may also affect *when* participants search. This, in turn, may affect their performance in our task. In our experiments, the optimal search behavior is to first explore different alternatives (for 2-3 periods) and then switch to exploitation (i.e., play the best alternative found in the exploration phase). Our participants deviate from this optimal search behavior in two ways: First, in a behavior we call oscillation, they switch repeatedly between exploration and exploitation. For example, they may first explore two periods, then exploit one period, only to explore for the remaining two periods. Such a search behavior is suboptimal and costly, even compared to participants who search the same number of alternatives: a participant who explores for the first four periods can then exploit a much better alternative in the last period. Second, in what we call delayed exploration, participants may get the order wrong, i.e., they start with exploitation and then switch to exploration. Again, this is costly compared to the optimal search behavior because superior alternatives found later in the task cannot be exploited for

as long as would have been possible had they been found early in the task. Take, for example, a taxi driver. She may begin her tenure by randomly choosing an area in which she believes many potential customers are located and then service this area before finally, towards the end of her tenure, exploring other areas. This is clearly inferior to a strategy in which she tries out different areas at the beginning of her tenure before settling to service the area that turned out to have the most customers.

In Table 2-2, we report the frequency and performance implications of these two types of deviations from optimal search behavior.

		Search Error	
		<i>Delayed Exploration</i>	<i>Oscillation</i>
<i>Control</i>	<i>Share</i>	67.9%	56.8%
	<i>Performance</i>	~	~
<i>Peer Performance Information</i>	<i>Share</i>	46.6%	39.3%
	<i>Performance</i>	-	-
<i>Large Choice Set</i>	<i>Share</i>	63.9%	47.7%
	<i>Performance</i>	+	+
<i>Peer Performance Information & Large Choice Set</i>	<i>Share</i>	53.6%	46.4%
	<i>Performance</i>	+	+

Note. Search late indicates that participants first exploit a known alternative before searching for new ones. Alternate between Search/Exploit indicates that participants switch more than once between exploiting a known alternative and searching for new ones.

Table 2-2: Performance implications of search errors and percentage of participants committing them

A large share of participants switched repeatedly (rather than just once) between exploration and exploitation. In the control condition, 56.8% of participants exhibited this oscillation behavior. One participant remarked "*I made my decision based on what I would do in this actual event. I would play it safe then go for a goal. Once that is met I would then use that method again. Rather than playing it safe the entire time or just going all out. There needs to be an even ground somewhere.*" This would suggest that participants preferred to oscillate between exploration and exploitation out of a desire to pursue a performance goal while at

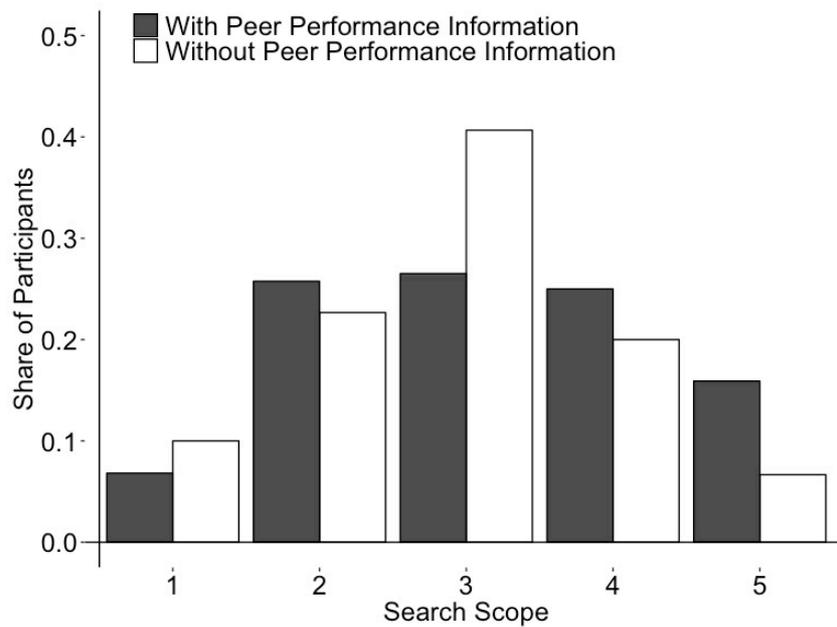
the same time minimizing risk. When peer performance information was available, the share of participants following an oscillation approach dropped to 39.3%. Interestingly, with larger choice sets, we do not observe such a drop if peer performance information is available (47.7% versus 46.4%). This makes sense in light of the statement above – there is no need to risk moving away from the fallback option for further exploration once the best option has already been identified, which is more likely to happen with peer performance information on small choice sets.

In the experiment, we observe, that many participants engage in delayed exploration, i.e. they only start searching after first exploiting. In the control condition 67.9% of participants fell back on known alternatives before searching for new alternatives again. When endowed with peer performance information, this proportion fell to 46.6%. Likewise, on large choice sets, 63.9% of participants exploited known alternatives before searching. When peer performance information was available, only 53.6% of participants searching on large choice sets showed this behavior. The universal drop in delayed exploration when peer performance information is available may be explained by the naïve assumption of some participants that whatever alternative they first come across is sufficiently good. For example, one participant remarked: *"Take the easy way or try something new. But, hey if the company already pulls in 76 million in profit why change."* When, however, peer performance information was available, participants seemed to question their initial assumption about the quality of the first alternative and start the task by exploring. Thus, overall, our results suggest that peer performance information helps to reduce decision errors for both small and large choice sets.

Robustness Checks

As the sequence of alternative was random in the experiments, a critic may argue that some of the performance effects are due to chance. We therefore run the peer performance

treatment again for the case of large choice sets, but this time use a fixed sequence of alternatives. This new round of experiments featured 415 participants across the US, 47 percent of which were women. We again find a significant performance difference between control and peer performance treatment (MD = -0.05; p = 0.030). We also find that these differences were borne by differences in search behavior and that those in the peer performance treatment often over-explored (see also Figure 2-6).



Note. The numbers on the X-axis reflect the stopping point, i.e. the number of alternatives explored.

Figure 2-6: Level of exploration with a fixed sequence of alternatives

Secondly, to address criticisms about the sample on the Mechanical Turk platform, we collected information on the participants' demographic characteristics. We then split the sample population by age (above or below 30 years of age), level of education (college educated versus not college educated) and income (an hourly salary of more versus less than \$20). Our results are robust to these different demographic groups. They also hold when we divide the sample of participants into those holding a managerial position versus those that do not (a detailed report is provided in the Appendix 2 C).

Critics may also argue that over-exploration is not driven by the prospect of finding the best alternative but rather by handling downward risk: Assuming a uniform distribution, it is more likely to receive a lower payoff in the next draw when choice sets are large. To avoid ending up with a payoff lower than the initial one, participants may just explore more. If that were the case, however, the number of searched alternatives below the initial draw would be higher in the treatment conditions with large choice sets. However, we find no difference in the number of draws with a lower value than the initial draw between conditions with large choice sets and small choice sets (MD = - 0.013, $p = 0.86$).

Our main effect also holds true if we use an ANOVA to conduct the analysis. We report the results in Figure 2-7. In particular, the analysis confirms that there is a significant effect of peer performance information on large choice sets on performance ($p = 0.000$).

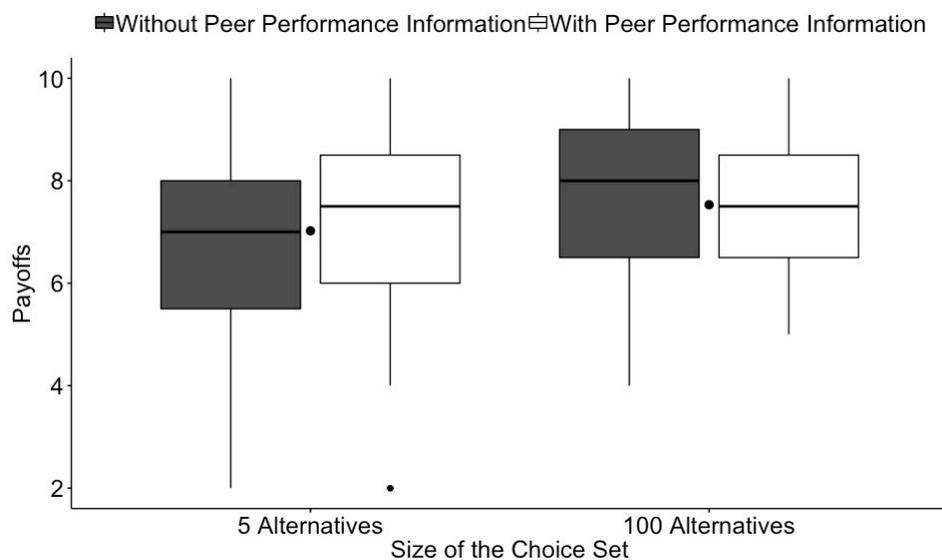


Figure 2-7: ANOVA of performance differences across the four conditions

Discussion

In our study, we investigated the performance implications of social comparisons in a simple search task. Practitioners and management scholars frequently assume that social comparison

improves search efficacy (and as a consequence, performance). However, this assumption has remained untested empirically. In an effort to fill this gap, we ran an experiment and found that while social comparison may improve performance when there are few decision alternatives, it may turn into a liability when there is a large number of decision alternatives. It is then that social comparison may lead to over-exploration and consequently, cause a performance penalty.

Contributions

Common assumptions in the extant literature are that (1) decision-makers generally search too little and (2) upward social comparison induces search. Thus, social comparison may increase performance by increasing search and overcoming inertia (Audia et al. 2000, Greve 1998, 2003a, Greve and Gaba 2017). However, this argument does not account for the reported negative performance effects of social comparison (Nickerson and Zenger 2008, Obloj and Zenger 2017). We extend the search-based line of reasoning by setting forth conditions under which upward social comparison may have negative performance consequences by altering search behavior. Concretely, we offer a model of search and learning which produces positive or negative performance effects depending on the number of decision alternatives. We contend that social comparison informs the decision-maker on how to improve the efficacy of search but may also lead to costly over-exploration when many decision alternatives are available (Baumann and Siggelkow 2013, Billinger et al. 2014, Clement and Puranam 2017).

Previous research has focused on either positive (Blanes i Vidal and Nossol 2011, Stark and Hyll 2011) or negative performance effects (Colella et al. 2007, Dunn et al. 2012, Gino and Pierce 2009, 2010, Nickerson and Zenger 2008, Obloj and Zenger 2017, Pfeffer and Langton 1993, Van de Ven 2017, Zenger 2016). In framing these conflicting findings, scholars have used distinct theories to explain such effects. Positive effects have often been

explained by economic theories of incentives while negative effects have been explained by theories rooted in the psychology literature such as theories of envy (Nickerson and Zenger 2008).

Our search model reconciles the conflicting finding in one theory. In so doing, we show that beyond traits and competitive dynamics, performance effects of social comparison can be thought of as properties of individual search and learning. We thus expand the scope of the theory of search and learning (Denrell and March 2001). What is more, we identify boundary conditions and key contingency factors of a search-based theory of social comparisons. We show that the key contingency of the benefits of social comparison is the size of the choice set, i.e. the number of decision alternatives. When upward social comparison is taken as information on a global optimum, it may prevent costly over-exploration once decision-makers find an already optimal solution (Levinthal 1997). Here, social comparison may regulate search efforts. At the same time, social comparison may mislead search efforts such that facing many decision alternatives can in fact cause over-exploration rather than regulation of search. Lastly, our findings add to social aspiration (Cyert and March 1963, Greve 1998 Greve 2003b, March 1994) and stretch goal theory (Hamel and Prahalad 1993, Latham 2004, Kerr and Landauer 2004, Sitkin et al. 2011) which are in part derived from social comparison theory. A common notion of social aspirations is that they are beneficial because they induce change (Audia et al. 2000, Lant 1992, Miller and Chen 1994), exploration (Baum and Dahlin 2007), and risk taking (Bromiley 1991). Observing better performing others leads to reevaluating current choices and thus triggers (problemistic) search (Cyert and March 1963). The challenge in this predominant view is to fight over-exploitation.

Since March (1991) introduced the exploration-exploitation dilemma into the organizational learning literature, scholars have focused on the cost of too little exploration

(Audia et al. 2000, Rhee and Kim 2014). This has been attributed to environmental factors such as delays in feedback (Repenning and Sterman 2002) as well as to individual factors such as overconfidence (Audia et al. 2000). We show that there can also be a problem of over-exploration. In particular, we observe that decision-makers search too much in the absence of social comparison. When social comparison is then possible, participants often search less than they otherwise would. Thus, our results indicate that social comparison often *decreases* rather than increases the level of search. This is so because information on one's relative standing to a top performer enables decision-makers to put their own performance in perspective. This way, next to learning about the performance of an alternative their peer is pursuing, decision-makers learn about the performance of their currently pursued alternative. In this way, social comparison is not a trigger but a substitute of search. This is particularly true for the case of few decision alternatives when a decision-maker has a high probability of coming across an alternative that is as good as their peer's. This finding adds an important contingency factor to Greve's (2003b) model of search. The direction of the effect of an attainment discrepancy on the level of search may be true for large choice sets but it does not hold for small choice sets. This contingency may explain why occasionally very hierarchical organizations manage innovation more successfully than organizations with flatter hierarchies (think Amazon): The hierarchy may be an attempt to limit the number of decision alternatives available to decision-makers thus regulating search within the organization.

Our findings also add to the organization design literature (Joseph and Ocasio 2012, Nickerson and Zenger 2008). Managers promote or demote social comparison as one pillar of the social architecture of an organization (Nickerson and Zenger 2008, Zenger 1992). Tools to encourage social comparison are pay transparency (Gartenberg and Wulf 2017), rankings and ratings (Greve and Gaba 2017), or internal benchmarking (Gaba and Joseph 2013). Tools to suppress social comparison include putting restrictions on information sharing and placing

offices farther apart (Kulik and Ambrose 1992, Nickerson and Zenger 2008). We point to the number of decision alternatives as a contingency factor of when social comparison should be promoted or discouraged. This applies in a particular way to managing innovation: While systems of open innovations do not present the decision maker with pre-structured problems, systems of closed innovation pre-structure the problem at the organizational level before passing parts of the problem on to employees in the organization (Felin and Zenger 2014). As a consequence, decision makers in systems of closed innovation face fewer decision alternatives than decision makers in systems of open innovation. In this paper, we suggest that in particular systems of open innovation may be harmed by encouraging social comparison, while systems of closed innovation may be benefitting from more social comparison.

Lastly, we contribute to the behavioral strategy literature by introducing a new contingency variable: the number of decision alternatives (i.e. the size of the choice set). The number of decision alternatives is a central dimension to task complexity (March and Simon 1958, Payne 1976, Simon 1962). Decision-makers often face problems of different degrees of complexity and by extension, varying numbers of decision alternatives. For instance, hierarchies filter information and suggest limited courses of action, thereby reducing the number of alternatives (Nickerson and Zenger 2004). Likewise, at the organizational level, the number of decision alternatives depends on the market structure in markets for sales and supply. For example, decision-makers may sometimes only have few suppliers available to choose from (think airlines buying either Airbus or Boeing planes) and other times have the choice between many alternative suppliers (think airlines choosing a company to supply the plane seats). The size of the choice set has in the past also been identified as an important decision-making factor in consumer choice (Iyengar and Lepper 2000, Schwartz 2004, Shin and Ariely 2004) in what has been dubbed the “paradox of choice” (Schwartz 2004). The

paradox of choice is that while more options should theoretically lead to better decision-making outcomes, it in fact often leads to worse outcomes. It should lead to better outcomes because, as rational choice theory predicts, more options imply a greater chance of optimally matching an alternative with one's preferences (Arrow 1963, Rieskamp et al. 2006). At odds with this prediction, studies of consumer choice find negative outcome effects of increased choice (Iyengar and Lepper 2000, Schwartz 2004). This is partly because consumers are deterred from making a decision at all because they anticipate regret over what is likely not an optimal choice (Iyengar and Lepper 2000, Schwartz 2004). Alternatively, scholars believe that consumers do not even consider a larger number of decision alternatives because their evaluation may be too costly (Hauser and Wernerfelt 1990) or not cognitively possible (Malhotra 1982, O'Reilly 1980, Reutskaja and Hogarth 2009).

We point to a different hazard of large choice sets that stems from making *worse choices* rather than *not choosing*: Larger choice sets may prompt greater levels of search, while the optimal level of search is first and foremost determined by the time horizon and less by the size of the choice set. To be clear, we loan the concept of the number of decision alternatives from consumer choice to introduce it to strategic decision-making. However, while the consumer choice literature examine whether or not consumers choose to begin with, we examine search efficacy, which is more relevant for decision-makers who must choose between engaging in exploitation vs exploration, e.g. deciding on R&D projects (Mudambi and Swift 2014), managing alliances (Yang et al. 2014) or coordinating the internal organization of the firm (Stettner and Lavie 2014).

Giving greater importance to the number of decision alternatives challenges our view on how upward social comparison affects adaptation (Levinthal 1997, Rivkin 2000). A common assumption is that knowing what is possible provides decision-makers with an incentive to persist in the face of adversarial situations (Nelson and Winter 1982). Consistent

with the view of “routines as targets”, social comparison with a successful competitor might trigger innovation and global search, even when it is impossible to imitate the competitor’s routines (Nelson and Winter 1982, p. 112). Meanwhile, information on a global optimum may prevent over-exploration once decision-makers find an already optimal solution (Levinthal 1997). Hence, social comparison may regulate search efforts. At the same time, attempting to imitate similar others may trigger exploration even after a global optimum has been found (Czasar and Siggelkow 2010). Our findings point to another downside of social comparison: social comparison may mislead search efforts. We show that the benefits of upward social comparison depend on the size of the choice set and that facing many decision alternatives can in fact cause over-exploration rather than regulation of search.

Limitations

Critics of our experimental approach will argue that our findings might not hold in actual organizational contexts. While we believe that there is some external validity to our findings, the approach is certainly more geared towards establishing the mechanisms through which social comparisons hurt or benefit performance. However, our setting has been proven to be a valid abstraction of many real organizational decision-making scenarios (Aggarwal et al. 2011, Cappelli and Hamori 2013, Puranam and Swamy 2016). Particularly because our setting is abstract enough to capture central features of a multitude of organizational settings, it is likely to yield externally valid results (Cook et al. 2002).

While the decision scenario in our experimental task is a simplified version of reality, the way in which managers perceive the decisions they face can be depicted in a dramatically simplified mental model as well (Gavetti and Levinthal 2000, Halford, et al. 1994, Kelley 1973). In reality, there may also be situations where search is a coordinated team effort rather than conducted by an individual (Knudsen and Srikanth 2014). Teams in which decisions are based on such coordination tend to exhibit lower levels of search (Davis and Eisenhardt

2011) and as a consequence, the problem of over-exploration through social comparison may be less pronounced. This is because information sharing and coordination are costly and so teams are more likely to settle on strategies with more predictable outcomes, i.e. exploitation rather than exploration (Lavie et al. 2010). The present research would suggest that social comparison may be capable of breaking up patterns of over-exploitation in leadership teams, but more empirical research is needed to scrutinize the role of social comparisons in coordinated search.

We also do not claim to present a task that is a valid abstraction for all search tasks in real organizational life. Our experiment does not feature intercorrelated payoffs between alternatives; a situation that can be found for instance when a new region is explored as a sales market and neighboring regions' profitabilities are highly correlated. For these cases, a landscape model would be the better abstraction (e.g. Billinger et al. 2014, Levinthal 1997). Likewise, we do not model situations where payoffs are dynamically changing over time, which could be better accounted for with a restless bandit model (Loch and Kavadias 2002).

What is more, we deliberately excluded an effort component from our experimental task to focus on the choice problem. However, increased effort might partially offset the negative effect we found of social comparison with a highly successful peer when facing many decision alternatives. Increased effort might make decision-makers find solutions faster and thus partly balance the decreased probabilities of success.

Lastly, we cannot fully explain why decision-makers over-explore when there are many decision alternatives. The finding is, however, consistent with prior experimental studies (Billinger et al. 2014, MacLeod and Pingle 2005). One reasonable explanation for that behavior is that decision-makers are insensitive to changes in probabilities and more focused on the amount of outcomes (March and Shapira 1987).

Conclusion

Upward social comparison is often encouraged in organizational contexts to increase managers' effort in a given task. It may also yield valuable information in search tasks that allow managers to balance exploration and exploitation efficiently. However, this study suggests that it may constitute a previously understudied liability in organizational learning by causing over-exploration when faced with an abundance of decision alternatives.

3. LEARNING FROM OMISSION ERRORS

Introduction

Adaptive learning is the core of the behavioral theory of the firm (Cyert and March 1963, March and Olsen 1976). Decision makers choose a solution to a specific problem, and then use performance feedback on that solution to affect the search for future solutions: if the performance feedback is positive, they are expected to choose the successful solution again; if the performance feedback is negative, they are expected to avoid the solution in the future (Glynn et al. 1991). While decision makers have been found to learn from their successful choices, it is much less clear that they learn from their errors (Bennett and Snyder 2017, Kc et al. 2013, Haunschild and Sullivan 2002). Behavioral changes spurred by errors in previous choices can lead to improved performance in some circumstances, but not in others (Eggers and Suh 2019) suggesting that learning from mistakes is not a simple process.

One complication in learning from errors is that judgment errors take two different forms - errors of commission, where the decision maker does something that produces a negative result, and errors of omission, where the decision maker declines to follow an action that would have improved performance (Green and Swets 1966, Sah and Stiglitz 1986). Commission errors lead to specific penalties for behavior, while omission errors lead to foregone rewards. For example, a manager of an automaker may decide to launch a car that turns out to sell poorly or she may decide against the launch of a car that would have been a success with customers. Such errors in decision-making occur regularly, making learning from prior errors feasible and necessary (Csaszar 2012). While inferring the correct action to take in response to an error is always a challenge, the opportunity to learn from commission errors is always present – experiential learning processes provide feedback that directly affect future choices (Eggers 2012, Maslach 2016). Learning from omission errors, meanwhile, is more complex and challenging. The primary difficulty is that information on

errors of omission is not always available to decision makers. In most situations where information on errors is available, it comes from observing competitor choices on identical or similar projects (Maslach et al. 2018). Venture capital firms and record labels may be able to see the success of opportunities that they rejected, and pharmaceutical companies may be able to observe valuable drugs from competitors that represent paths they chose not to pursue. This suggests that observation of competitors may allow decision makers to learn from their omission errors.

A core challenge, however, is that attending to competitors affects not just the ability to learn from omission errors, but also affects social comparison processes that produce additional changes in behavior and challenges to learning that complicate the ability to learn and adapt to feedback (Kacperczyk et al. 2015). In this study, we seek to separate the effect of information on omission errors from the competitive effect of observing competitors through the use of an experiment (Di Stefano and Gutierrez 2019). This allows us to assess whether individuals learn from their commission and omission errors, as well as how learning differs across sources of performance feedback. We focus on a simple product approval task where participants repeatedly decide whether to accept or reject a product for development. Features of a product are repeated in multiple products over time, such that learning from rejecting a profitable product (*omission error*) or accepting an unprofitable product (*commission error*) is feasible. We manipulate what type of feedback participants receive: (a) just their own (historical) performance without information on their hypothetical performance in the case of product rejections, which should allow for learning from commissions but not omissions, (b) adding information on the hypothetical performance of product rejections, and (c) disclosing information about the value of omission errors through the behavior of a competitor. We theorize that the presence of a competitor introduces a

reference point which alters learning depending on whether the decision maker is performing above or below the competitor.

Our results demonstrate multiple aspects of the complexity of learning from omission errors. As expected, we find no evidence of learning from omissions without information on the value of rejected products. While the inclusion of this information facilitates learning from omissions, we find that such learning tends to crowd out learning from commission errors – there are limits on how much information participants can process. When feedback on omissions is provided through information about hypothetical competitors, we find that learning effects diminish in general. While information overload appears to be partly responsible for this diminishment, the results intriguingly show that the lack of learning from omission errors through competitive performance stems from the introduction of an aspiration level. Specifically, participants focus on learning from their commission errors when trailing the competitor, while their focus shifts to using their competitor’s experiences to learning from their omission errors when they are ahead of the competitor. These findings suggest that learning by observing competitors can be an effective way to improve the opportunity to learn from omission errors, but that the social comparison effects of such information clouds the ability to efficiently use all available data to learn from mistakes.

Our findings contribute to a stream of literature examining learning from failure (Eggers 2012, Maslach 2016). Prior research has begun to outline conditions under which learning from failure may be impaired (Audia and Greve 2006, Eggers and Suh 2019, Gaba and Joseph 2013, Maslach 2016). For example, Eggers and Suh (2019) find that learning from decision errors is possible when decision makers have the capability to determine which part of the decision caused it. We suggest that learning may even be impaired under situations of simple decision tasks and with a single source of performance feedback

(Haunschild and Beckman 1998) and that this is due to dynamics of social comparison. Specifically, the presence of a competitor induces a reference point above which learning differs from being below that reference point. While the presence of a competitor has been previously found to impact learning behavior (e.g. Baum and Dahlin 2007), we point to their differential effects on omission versus commission errors.

In addition, we contribute to social aspiration research (Cyert and March 1963, Greve 1998, Joseph and Gaba 2015) by extending our understanding of the behavioral consequences of social aspirations. Such aspirations may aid self-assessment (Audia et al. 2015), induce search for new solutions (Greve 1998, Baum et al. 2005), and increase effort and productivity (Blanes i Vidal and Nossol 2011). We point to another behavioral consequence – learning from omission errors may be suppressed, i.e. choosing a performance target to strive for based on competitor performance may lead decision makers to discard useful competitor information that would otherwise allow them to learn from their omissions. This may be particularly relevant when omission errors are costly, i.e. firms cannot miss too many opportunities to stay competitive (Csaszar 2012). Thus, ironically, focusing on a better performing competitor with the explicit goal of closing the performance gap to that competitor may lead to more omission errors and thus cement the position behind the leading competitor.

Lastly, we build on the literature of competition between individuals. Prior studies have found an effect of rivalry (Kilduff et al. 2010) as well as co-acting in a similar task on performance (Flynn and Amanatullah 2012). While prior work established performance effect, the mechanism remains unclear. Specifically, Flynn and Amanatullah (2012, p.412) have speculated that effects are “the result of learning rather than motivation”. Our study confirms an effect of observing someone else’s performance in the same task on learning. Specifically, our results suggest that co-acting in a similar task can inhibit either learning

from commission or learning from commission errors depending on whether the focal actor is trailing or leading in task performance.

Theory and Hypotheses

The ability to learn from mistakes and make improvements in behavior that improve future performance is a central aspect of both individual and organizational learning. From a behavioral perspective, learning typically emerges from a simple process of decision and feedback – a decision maker chooses a solution and then observes the performance feedback she receives from this choice (Cyert and March 1963). If the feedback is positive, she chooses this solution again in the future, if the feedback is negative, she will not choose this solution again (Glynn et al. 1991). Specifically, a solution that looks similar to a solution that worked in the past, will likely be implemented while a solution that looks similar to a solution that did not work in the past, will likely not be implemented again. This suggests that successes will lead to repetition, while failures will lead to changed behavior (Kacperczyk et al. 2015).

A central aspect of this simple learning dynamic is feedback, as it is feedback on performance that enables learning and adaptation (Greve 2003). The importance of feedback creates complications for learning from errors, as different types of errors produce different types of feedback (Sah and Stiglitz 1986, Klingebiel 2018). Errors of commission – engaging in behavior that produces undesirable results – typically produces clear feedback that can enable learning. Numerous studies have focused on learning from commission errors at the individual and organizational level, studying learning from unsuccessful new products launched by the firm (Eggers and Suh 2019), from accidents (Baum and Dahlin 2007, Madsen and Desai 2010), and failed R&D experimentation (Khanna et al. 2016). The overall finding is that learning from commission errors can take place (and exceed learning

from successes) to the extent that decision makers can adequately interpret the feedback in terms of which aspect(s) of the initial decision produced negative feedback. Learning is further dependent on number, timing, and importance of failures (Khanna et al. 2016).

By contrast, omission errors – decisions not to engage in behavior that would have produced appealing results – are often more complex in terms of learning. Some omission decisions produce clear and actionable feedback. For example, a physician who errs by not prescribing a useful treatment may see a clear deterioration in the health of their patient (though health may deteriorate even with proper treatment). Likewise, a city mayor who does not declare an evacuation for a hurricane, but the hurricane makes a direct hit, will clearly understand the gravity of their mistake (Dye et al. 2013). But in many cases, feedback on omission errors is not readily available. When Csaszar (2012) looked at omission errors by mutual fund managers, it was difficult as a researcher to identify specific omission errors, and would be even more difficult for the mutual fund manager herself. The difficulty is to have information on every decision (including the decision not to pursue a project) as well as on the quality of the projects that a decision was made on. The manager would therefore need to be privy to all past decisions (of which there will often be no record when the project was not pursued) as well as project quality. In many cases, decision makers may get little or no direct feedback on their omission errors. In such a case, learning from omissions would likely be difficult or impossible. Thus, we expect that:

Hypothesis 1. Decision makers will learn from their errors of commission but not their errors of omission in the absence of direct feedback on their hypothetical payoffs.

It is also conceivable that feedback on omission errors is directly available. For example, when IBM shut down its Scientific Data Systems project, some of the engineers involved in the project went on to leave IBM to form the new company SAP. The fast success of this new company gave the IBM executives direct feedback on their error of

omission. A venture capital partner, publishing house executive, or recording label producer who "passes" on a given potential deal to sign may be able to observe the eventual success or failure of the venture, writer, or musical act, respectively. Bessemer Venture Partners, for example, publishes a webpage with their "anti-portfolio" of missed investment opportunities, which they presumably try to learn from to improve future decision-making. In such cases, the availability of direct feedback on the missed performance of omission errors can help decision makers learn from their omission errors.

Information on omission errors being available, however, also may have an effect on learning from errors of commission. Generally, decision makers in most situations exhibit an omission bias, i.e. they put more weight on omissions than commissions in the evaluation of their errors (Baron and Ritov 2004). Beyond a greater emphasis on omissions when information on them is available, learning from omission errors also requires the opposite behavioral antecedent as learning from commission error – while one requires accepting more projects, the other requires an increase in rejections (Klingebiel 2018). Thus, we argue that:

Hypothesis 2. If direct feedback on omission errors is available, decision makers will learn from their errors of omission but not their errors of commission.

Learning about omission errors by observing a competitor

One of the most common sources of information about omission errors comes from competitors. To the extent that a decision maker's competitors face a similar set of decisions to make, the decision maker can learn about the performance of their omission errors by focusing attention on their competitor.

Decision makers frequently attempt to learn from better-performing competitors by imitating their practices (Haunschild and Miner 1997). Specifically, they imitate what they can observe and what they believe to be generalizable (Baum and Dahlin 2007). There are, however, a number of problems with attempting to learn in this way. First, since imitated

practices are not selected based on their effectiveness but on the overall performance of the competitor using them, those practices may be ineffective in the focal firm (Abrahamson 1991, DiMaggio and Powell 1983, Haunschild and Miner 1997). Secondly, very successful peers may employ riskier practices and just have gotten lucky to make it to the top by chance in which case imitating their practices would likely be detrimental to the imitator (Denrell 2003). Lastly, the focus on few successful peers' practices may lead to a narrow set of practices in the field (Porac et al. 1989) and thus hinder adaptability in the long run.

The potential challenges of focusing on competitors in order to obtain useful information are especially relevant because competitors are typically the most important source of information about omission errors available to decision makers (Haunschild and Miner 1997). For example, an automobile manufacturer may decide to cancel an R&D project at a developmental stage but later observe a similar project being pursued to completion by a competitor and, subsequently, observe the sales numbers of the competitor model. The publishing executives that failed to sign J.K. Rowling, for example, could observe her success with Bloomsbury publishing.

Social comparisons, however, may affect the behavior of decision makers and organizations in ways beyond learning and imitation. We particularly focus on the idea of competitors offering a social comparison reference point for performance (Bromiley and Harris 2014). Given that extensive research has explored how learning behavior is different when decision makers are above or below social comparison reference points (Baum and Dahlin 2007), we identify below how focusing attention on the behavior of competitors to learn from the decision maker's own omission errors may affect the learning process.

When faced with the choice between accepting and rejecting potential investments, the decision to accept is inherently risky to performance. Rejection always leads to status quo, while acceptance either increases or decreases performance. Thus, taking action is

inherently risk seeking, while avoiding commitment is risk averse⁵. When decision makers are performing below their competitive reference point, we expect that they will become increasingly risk-tolerant in an effort to catch up with the competitor (Bromiley et al. 2001, March and Shapira 1987, 1992). This means that decision makers will be more likely to make commission errors when performing below the reference point. This will affect learning in two ways. First, making more commission errors leads to more data that can improve the opportunity to learn. If decision makers do not make commission errors, it will be difficult to learn from their commission errors, by default. Small sample learning is exceptionally difficult (Lampel et al. 2009, March et al. 1991), so an increased sample of commission errors will provide more opportunity. Second, while decision makers below aspirations may become more aggressive in order to try and catch up, they will want to avoid potentially large mistakes that could impair this effort. Any opportunity they see as "on the margin" will likely be pursued, so the real challenge is avoiding significant commission errors. As a result, the decision maker will focus their own limited attention on trying to diagnose their own commission mistakes in the past, and will pay less attention to the competitor's specific decision history. This will manifest as a laggard decision maker focusing on "getting their own house in order" first, before paying any significant attention to the decisions of the competitor. Such a focus on the decision maker's own history will inherently limit their ability to learn from anything other than commission errors. As a result, we hypothesize that:

Hypothesis 3. Decision makers will learn from their errors of commission when performing below their reference point but not from their errors of omission.

Above the reference point, we expect decision makers to show the opposite behavior. First, they will be less willing to take risks, which translates to approving fewer projects.

⁵ This, of course, is only true in an environment in which performance rewards are stable, whereas in an environment with changing rewards, omissions may well be riskier (see Klingebiel 2018, Silcoff et al. 2013)

Consequently, they will commit errors of omission more often, which analogously provides additional data on omission errors from which to learn (using information from the competitor to understand their omission). Second, successful decision makers will have the opportunity to engage in more explorative learning. The success of being above a reference point makes decision makers more confident, leading to an increase in exploring new approaches (Cyert and March 1963, Levinthal and March 1981). Scholars have identified two different types of learning – explorative and exploitative. While performing below the reference point may cause both explorative and exploitative learning, performing above the reference point causes explorative learning (March 1991). Explorative learning implies seeking new types of data from which to learn, which will increase the salience of competitive performance data. For this reason, explorative learning is linked to learning from the experience of others (Baum and Dahlin 2007). The resulting behavior will be that higher performing decision makers who have limited reasons to study their own previous commission errors will instead focus on the opportunity to learn from their own omissions by focusing attention on the behavior of the competitor. This leads to our final hypothesis:

Hypothesis 4. Decision makers will learn from their errors of omission when performing above their reference point but not from their errors of commission.

Method and Data

This study explores how different available performance feedback mechanisms affect learning from both omission and commission errors. We do so by using an online experiment that allows for causal identification of the effect of performance feedback on learning in the absence of confounding elements like the simultaneous arrival of performance feedback from multiple sources (Busemeyer et al. 1986, Cook and Campbell 1979, Sterman 1987). This empirical approach also allows for easy replication to verify our findings (Bettis et al. 2016).

Experimental task

Our experimental task mirrors the approach used in many modeling papers on learning (e.g., Csaszar 2013), where a decision maker is asked to assess projects that arrive in sequence, where each project must be either approved or rejected before considering the next project. Each project is an aggregation of multiple characteristics, which affect project performance. The sequencing of projects allows for the potential for learning from previous choices. Each project has a true, but hidden, project performance value which is only revealed after a decision is made. As a result, each round the participant will experience one of four potential outcomes drawn from signal detection theory (Dye et al. 2014) – a true positive, a false positive (commission error), a true negative, or a false negative (omission error). This provides a simple opportunity to test our broader theory about learning from omission errors based on the availability of feedback.

The procedure was conducted through an online, 20 round experiment on Amazon's Mechanical Turk. In the experiment, participants are put in the shoes of a manager deciding repeatedly whether to introduce a new car to the market. A car in our task can be described by a unique combination of its three main features (engine, transmission, body style), each of which in turn can hold 4 values (e.g. “V engine” or “convertible”; full list of features and values in Appendix 1). A mix of features could for example be a convertible with a V engine and manual transmission, and each affected profit in a simple and independent manner. Thus, a manual transmission might always be positive for car performance across all potential cars considered by the subject. While many car features may actually affect performance in an interdependent manner, such a task would make learning significantly more difficult. For each of the three car features, the best performing value increased performance by three points (see more details on points below), while the worst decreased performance by three points and the others performed in between. Thus, a given car's

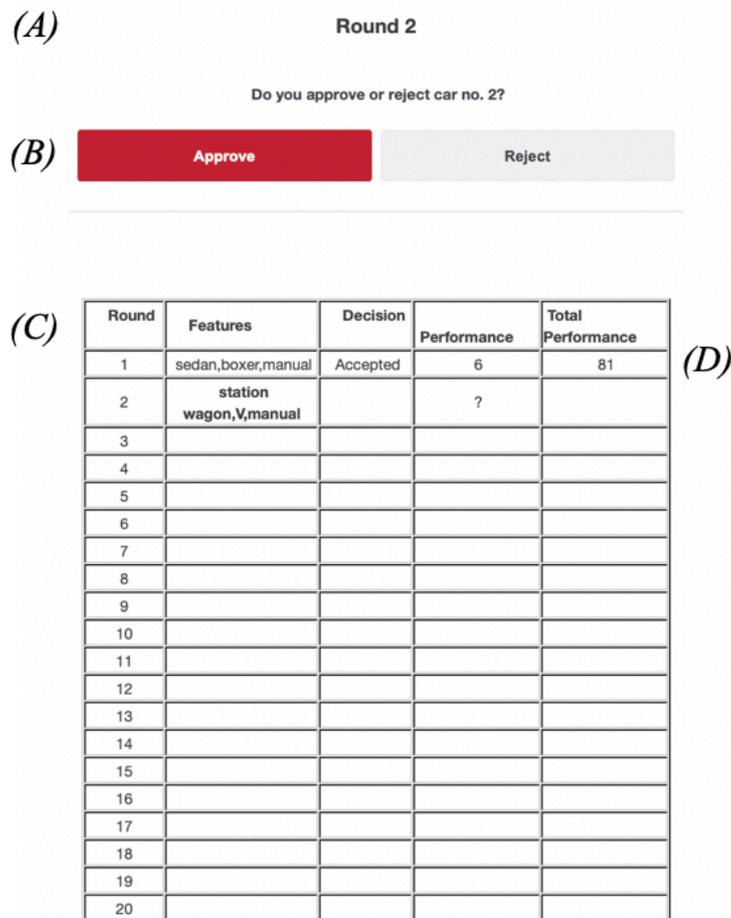
aggregate performance ranged from positive nine to negative nine. All cars which were launched by the participant affect profits (positively or negatively), while cars that are not launched do not affect profits. Thus, learning could take place about a mix of features by observing the performance feedback on other, similar cars. The performance feedback was immediate, unambiguous and non-stochastic. It was made clear to the participants that no knowledge of actual cars was required to perform well in the experiment.

Experimental design

Incentives. In providing monetary incentives to the participants, we applied the induced-value approach (Smith 1976): for the participants, more profitable product launches resulted in higher bonus payments. In line with prior experimental work, the participants were shown only correct information (Ortmann and Hertwig 2002). For 10 minutes of their time, participants were paid 25 US cents as a base payment plus an additional payment of up to one US dollar depending on their performance in the task. The monetary reward at stake in the experiment is commensurate with prior experimental work (Harmon et al. 2015) as well as with reservation wages of workers on Amazon Mechanical Turk (Burbano 2016, Horton et al. 2010). What is more, previous studies on monetary stakes in experiments have shown that behavior is largely invariant to varying stakes for choice (as opposed to effort) tasks (Camerer and Hogarth 1999, Smith and Walker 1993). Throughout the task, participants could accumulate performance points (an experimental currency) that would then later be translated into US dollars. At the beginning of the task, all participants were given 75 performance points and each round, participants could lose or win up to 9 points. One performance point translated into one US cent (\$0.01). We chose an experimental currency with an initial endowment rather than a bonus because accumulated profits during the task could otherwise plunge into the negative, thus taking away money from the participant. This

format ensured that participants always achieved at least a small bonus, even if they got everything wrong. Participants knew of the experimental currency / US dollar exchange rate.

Procedure. Upon starting the survey, participants were informed of the estimated length, the performance incentives, and the number of rounds of the exercise. In each of the 20 rounds, participants were shown the same decision screen (see Figure 3-1 below). On that screen, the current round was displayed (A) and under it the question of whether to accept or reject the current proposed car (B). The bottom of the screen contained a table (C) with information on the current car features as well as information from past rounds on features, decisions and performance. The table also noted the participant's current point total (D).



Note. Decision screen from the control treatment from round 2 including information on (A) current round, (B) question of whether to accept or reject the current proposed car (B), (C) information on the current car features as well as information from past rounds on features, decisions and performance, and (D) total points.

Figure 3-1: Example decision screen

After each decision, participants were shown their profits or losses, respectively, for last round's car. During the task, an attention test was shown between rounds eight and nine instructing participants to not click an accept or reject button to check they were engaging in the task mindfully. After task completion, participants were informed of their final score and presented a short questionnaire on demographic characteristics. Among the information polled in this questionnaire were a verbal account of the participants' strategy, their gender, age, motivation to participate, level of education, background in STEM subjects and average hourly wages on Amazon Mechanical Turk and offline. See Appendix 3 B for screen shots of the entire survey.

Treatments. The central manipulation was the type of performance feedback that was shown to participants. All comparisons between treatments were done between subjects, i.e. each participant only participated in one treatment. In the control condition, only feedback on the participant's own product launches was displayed. Importantly, no information was given on how the car would have fared in the market if it was not chosen to be launched by the participant. In the first treatment condition, we introduced a competitor. Participants were instructed that they had a counterpart at their main competitor's company tasked with similar duties and that they would see the competitor's success in the market and their competitor's car features. This treatment added additional columns to Figure 3-1, noting the competitor's current performance level as well as the cars shown to the competitor. In the second treatment condition, there was no competitor information but direct feedback was provided (see Appendix 3 B) on how cars not chosen to launch would have performed in the market.

Respondents. All participants were recruited on the Amazon Mechanical Turk platform. We used only data on participants that finished the task, passed the attention check, and more than 3 seconds on reading the instructions. A total of 123 both finished and passed the attention check and spent more than 3 seconds reading the instructions. All statements

about participants demographics refer to those 123 participants. We restricted the subject pool to residents of the United States to avoid difficulties with English as the language of instruction. Half (50%) of the participants were women, the average age was 37 years old, 49% had an academic background in a STEM subject, and 72% had at least a college degree. On average, participants earned 76 cents in bonus payment.

Analysis

To understand learning, we arranged our data with a single observation for each round of the experiment (20 rounds). Using fixed effects (at the individual level) allowed us to look within any given subject at their rate of learning, while controlling for any time-invariant factors that may affect performance. This fixed effect approach combined with randomization of the treatment assigned to any given participant allowed us to use a regression-based approach (since we have two independent variables), but we need employ no other control variables. As our dependent variable (see below) is binary, we used a fixed effects logit regression.

Dependent variable. The dependent variable captures success versus failure in a given round i . Success means making a correct choice, i.e. either rejecting a proposal that yields a negative reward or accepting a proposal that yields a positive reward. Thus, the dependent variable is a binary variable.

Independent variables. Given our focus on learning from omission and commission errors, we use the respondent's cumulative omission and commission errors as two independent variables. To avoid an induced slope effect in the cumulative errors, we focus only on errors committed in the previous four rounds (Bennett and Snyder 2017). Expanding the window of errors considered produces similar, and statistically stronger, results.

Results

We report findings of our baseline regression analysis in Table 3-1. For participants in the control condition (Model 1), we find a positive and highly significant effect for (0.360, $p < 0.001$) for the effect of prior commission errors on success. This coefficient (odds ratio = 1.43) suggests a 43% higher likelihood of making the right decision (i.e. not making a decision error) for each recent commission error. Meanwhile, there is no such effect for omission errors which is not surprising given that participants had no information on the hypothetical performance with cars they rejected. These findings are consistent with H1 and highlight how learning from mistakes is only feasible with adequate feedback on performance and potential outcomes of omitted decisions.

	<i>Model 1</i> <i>Control</i>	<i>Model 2</i> <i>Full Info</i>	<i>Model 3</i> <i>Competitor</i>
<i>Commissions</i>	0.360*** (0.095)	0.073 (0.086)	0.202* (0.119)
<i>Omissions</i>	-0.064 (0.129)	0.333** (0.145)	0.327* (0.192)
<i>Fixed Effects</i>	<included>	<included>	<included>
R ²	0.021	0.007	0.012
N	800	752	416

*Note: * $p < 0.10$; ** $p < 0.01$; *** $p < .001$*
Standard errors in parentheses

Table 3-1: Main results (full complexity)

Participants in the full information treatment receive complete performance feedback on both proposals they have accepted and rejected, and the results are shown in Model 2. The results suggest that participants demonstrate positive learning from omission errors ($p < 0.01$), but now no longer demonstrate learning from commission errors. In terms of the

effect size, this means that an additional error of omission is associated with a 39% higher likelihood of making the right decision subsequently. Meanwhile, participants in the third condition (Model 3) are aware of competitor choices and performance. These participants continue to show positive learning from omission errors (38% increase, $p < 0.1$), though the standard errors increase suggesting more variance. They do however demonstrate a small positive effect of learning from commission errors as well (22% increase, $p < 0.1$). Our second hypothesis (H2) had suggested that disclosing information about omission errors would help participants learn from omissions (supported in both Model 2 and Model 3), but doing so would crowd out learning from commission errors. The latter effect is supported in Model 2 (full information), but while the effects are weaker in Model 3 (competition), learning still manifests. This offers only limited support for H2, but the general pattern is broadly consistent with the idea that that learning from one type of error may crowd out the ability to learn from the other type of error.

To further explore the idea of "crowding out", we re-ran our experiment with a simpler task. Our supposition was that any impairment to learning that crowded out one type of learning could be driven by information overload. For example, rather than processing only own choices and feedback, participants in the competitor treatment also need to process the same information for the competitor. If information (or cognitive) overload is the reason for reduced learning, we would expect an increase in learning when the task becomes less complex. To simplify the task we reduced the number of features (e.g., engine) from three in the baseline condition to two in the reduced complexity condition. These results are reported in Table 3-2. In these results, we now find a consistent effect of learning from commission errors (35% in the control, 42% with full information, and 43% in competitor in Models 4-6, respectively), but we no longer find support for any learning from omission errors. Overall,

we still observe a lack of learning in the competitor treatment even when the information load is reduced.

	<i>Model 4</i> <i>Control</i>	<i>Model 5</i> <i>Full Info</i>	<i>Model 6</i> <i>Competitor</i>
<i>Commissions</i>	0.306** (0.142)	0.351*** (0.125)	0.363** (0.176)
<i>Omissions</i>	0.191 (0.223)	0.119** (0.191)	0.174* (0.254)
<i>Fixed Effects</i>	<included>	<included>	<included>
R ²	0.012	0.014	0.017
N	400	576	288

Note: * $p < 0.10$; ** $p < 0.01$; *** $p < .001$
Standard errors in parentheses

Table 3-2: Reduced complexity results

Our third and fourth hypotheses focused on the idea that participants would be more able to learn from commission errors when they were performing worse than their competitor (H3), and more able to learn from omission errors when they were performing above their competitor (H4). In Table 3-3, we report the results split by reference point for the full complexity task (Models 7 and 8) and the reduced complexity task (Models 9 and 10). In line with H3, we find consistent and significant positive learning effects from commission errors when the participant is trailing the competitor (38% increase in full complexity, $p < 0.1$; 157% increase in reduced complexity, $p < 0.05$). When the participant is above aspirations, the effect is absent in the full complexity task and appears but is much smaller than the learning from omissions in the reduced complexity task (51% increase, $p < 0.1$). This generally supports H3.

	<i>Full Complexity</i>		<i>Reduced Complexity</i>	
	<i>Model 7 Below</i>	<i>Model 8 Above</i>	<i>Model 9 Below</i>	<i>Model 10 Above</i>
<i>Commissions</i>	0.328* (0.171)	0.263 (0.183)	0.947** (0.417)	0.419* (0.219)
<i>Omissions</i>	0.385 (0.263)	0.680* (0.350)	0.560 (0.265)	0.971** (0.424)
<i>Fixed Effects</i>	<included>	<included>	<included>	<included>
R ²	0.025	0.022	0.102	0.038
N	205	211	62	226

*Note: *p<0.10; **p<0.01; ***p<.001
Standard errors in parentheses*

Table 3-3: Split samples based on social aspirations

In line with H4, we find that learning from omissions only manifests significantly when the participant is ahead of the competitor, both in the full complexity task (97% increase, $p < 0.1$) and the reduced complexity task (163% increase, $p < 0.05$). While the coefficients on omission errors are not significant when the participant is below aspirations (in either complexity condition), it is worth noting that the coefficients are positive but the standard errors are quite large. This suggests that learning from omissions for participants below aspiration may be very uneven. It is important to note that participants above aspirations in the reduced complexity task seem to be able to learn from both commission and omission errors, though in line with H4 the effect of learning from omission errors is much stronger. This does provide limited support for the notion that task complexity may limit the ability to learn from both types of errors.

Our theory behind H3 and H4 suggested that only participants who were above aspirations would really pay attention to the competitor and try to learn from their omission errors. In this case, the core mechanism is attention to the competitor. We wanted to explore

this mechanism directly, and subsequently created an additional experiment to test attention. Specifically, we displayed stars (“*”) next to the features of the competitor cars. The number of stars displayed varied between two and six. After completing the task, we asked the participants to recall the minimum number of stars that appeared in any round. We then computed the distance of this measure from the true value. We normalized the measure between zero and one where zero meant a low focus on the competitor and one meant a high focus (i.e. an absolute accurate recollection of the minimum of stars). For the following analysis, the sample was split by those focused on the competitor (attention value above the median) and those not focused on the competitor (attention value below the median). This allows us to have a direct and unobtrusive measure of attention. We further asked participants directly how much attention they had paid to the competitor on a scale from zero (all attention focused on own performance feedback) to 100 (attention exclusively focused on the competitor).

We use this measure as an outcome variable in Model 11 in Table 3-4, where we check whether the number of periods the participant spent above social aspirations predicts the likelihood of paying attention to the competitor. As shown in the results, for each additional round above the reference point the participant was significantly more likely to pay attention to the competitor (4.5% increase, $p < 0.01$). This supports the core mechanism from our theory that the reason that participants above and below aspirations learn differently is because of attention to the competitor.

	<i>DV=Attention</i>	<i>DV=Performance</i>	
	<i>Model 11</i>	<i>Model 12</i> <i>Focus on</i> <i>Competitor</i>	<i>Model 13</i> <i>Not</i> <i>Focused on</i> <i>Competitor</i>
<i>Rounds Above Aspirations</i>	0.767** (0.374)		
<i>Commissions</i>		0.369*** (0.084)	0.180 (0.112)
<i>Omissions</i>		0.289** (0.132)	0.076 (0.161)
<i>Fixed Effects Constant</i>	<included>	<include>	<included>
R ²	0.046	0.022	0.006
N	90	972	463

*Note: *p<0.10; **p<0.01; ***p<.001*
Standard errors in parentheses

Table 3-4: Attention to competitor

To complete the test for the mechanism, we explored whether attention to the competitor affected the way in which participants learned. To do this, we split the sample based on those who paid attention to the competitor (i.e. those who remembered the minimum number of stars accurately) (Model 12) versus those that did not (Model 13) in Table 3-4. We find that participants only learn significantly from both types of errors when focused on the competitor. Specifically, an increase of one commission error is associated with a 33% higher likelihood of making the right decision ($p < 0.01$), while an increase of one omission error is associated with a 28% increase ($p < 0.05$). These results are somewhat surprising, suggesting that attention to the competitor is indeed the mechanism by which learning from omissions emerges, but it also seems to affect learning from commission

errors, as well. This may be driven by the idea that participants need motivation to learn in order to engage in active learning, and this motivation may emerge from attending to competitor performance.

Discussion

Previous research often does not differentiate between omission and commission errors when discussing learning from failure (e.g. Eggers 2012, Maslach 2016). But learning from omission errors is critical yet difficult as performance feedback on omissions is often absent. Decision makers may observe their competitors to learn from omissions, but we argue that such competitor focus may affect behavior and learning beyond simply providing insight into omission errors. This suggests that learning from decision errors is importantly influenced by context. Using a series of online experiments, we found no evidence for learning from omission errors unless feedback on the value of rejected projects was provided. When direct feedback on omissions was available, however, learning from omission errors crowded out learning from commission errors. When feedback was given by observing a competitor, we found overall diminished learning effects for both commission and omission errors. Specifically, we observed that learning effects were contingent on the position relative to the competitor – leaders were learning from their errors of omission but not their commission errors while laggards were learning from their errors of commission but not their omission errors. The core of our theory revolves around when and how leaders versus laggards will attend differently to competitors based on both their own incentives and availability of data. While laggards seem to learn mostly from commission errors, leaders focus on learning from their omissions. This is consistent with theories of performance feedback and risk taking (Bromiley et al. 2001, March and Shapira 1987, 1992) but has been ignored in previous studies on learning from decision errors.

Our study provides three primary contributions. First, this study contributes to the literature of learning from failure (Eggers 2012, Maslach 2016). Previous studies have pointed out that learning may be impaired when it is unclear which part of the decision caused the error (Eggers and Suh 2019) or when performance feedback is coming from multiple sources (Haunschild and Beckman 1998). We identify an additional constraint to learning from failure. Concretely, we find that competitive dynamics may not only spur or suppress learning (Baum and Dahlin 2007) but also determine what errors are learned from. This suggests that research on learning from failure (or learning more broadly) needs to focus not only on the firm itself, but also the competitive context that may directly affect the ability to learn.

Second, our findings shed light on the behavioral consequences of social aspirations. Prior research has found that setting a performance target based on competitors' performance may aid self-assessment (Audia et al. 2015), induce search for new solutions (Greve 1998, Baum et al. 2005), and increase productivity and effort (Blanes i Vidal and Nossol 2011). Our results indicate that there may be another side effect of social aspirations. They may curb learning from omission errors. This may hurt organizational performance in particular when avoiding omission errors -i.e. not missing important opportunities- is crucial to gaining a competitive advantage (Csaszar 2012).

Third, this research adds to the literature of competition between individuals (Deutsch 1949, Deci et al. 1981, Beersma et al. 2003). The literature of interindividual competition spans mere co-acting (Flynn and Amanatullah 2012) to intense rivalry (Kilduff et al. 2010). Specific competition between individuals has been shown to have consequences for individuals' performance. For example, Kilduff et al. (2010) found a significant effect of rivalry on team performance in NCAA basketball competitions. Nonetheless, the mechanisms mediating such performance effects are not fully explained. While most studies

point to motivation mediators (e.g. Kilduff et al. 2010), some (e.g. Flynn and Amanatullah 2012) have suspected that learning may explain part of the performance effects. Our findings suggest that observing a competitor may curb learning from commission errors when leading against the competitor and curb learning from omission errors when trailing the competitor.

More empirical work is needed to further explore how competitive dynamics influence learning from decision errors. Specifically, it is important to understand whether learning from omission errors above reference points and learning from commission errors below reference points replicates in an organizational context. In particular, do less profitable businesses focus too much on commissions but miss the next big opportunity when they are trailing their competitor(s)? This could be tested in the context of R&D activity in the Pharma industry. Omissions could be measured by patenting activity related to previously passed up projects and it would be easy to approximate historical and social reference points as done by previous studies (Lungeanu et al. 2016). Further, we focused on a single competitor's behavior as a source of learning about omission errors. It would be interesting to study the effect of observing groups of competitors. For example, firms increasingly have access to public repositories of other firms' failures, such as the MAUDE database for failures in medical devices (Maslach et al. 2018). This may constitute a boundary condition for our findings since the formation of reference points may be decoupled from the opportunity to learn about omission errors.

Importantly, our task represents a dramatically simplified version of actual organizational decision makers. For example, a real R&D manager deciding on which projects to invest in receives performance feedback on many dimensions and from multiple sources, think customer feedback, sales data etc. However, managers need to form a mental model in order to make any decision that is a simplification of reality, too (Gavetti and

Levinthal 2000, Halford et al. 1994, Kelley 1973). Similarly, our experiments are different from many organizational decisions as the stakes involved are much lower. A failed car rollout may cost a manager his job, whereas in the experiment he is risking merely cents. Prior research has identified behavior to change in two ways for lower (monetary) stakes. First, individuals typically invest less effort (Smith and Walker 1993) and secondly, they take greater risks (Lefebvre et al. 2010, Weber and Chapman 2005). Our experiment, however, is not an effort task. Risk may play a role in the sense that approving a project always features a greater outcome variance than rejecting a project (outcome is always 0). This should affect all our treatment conditions in the same way and thus, is not expected to bias out results.

Another way in which the experimental task varies from a number of decision problems is that the likelihood of making an error of commission versus an error omission is balanced. For many managerial decision problems, on the contrary, the likelihood of making an omission error is much more likely than making a commission error (Klingebiel 2018). This is the case because the decision space is large and thus, many potentially beneficial projects are foregone at each point in time. Assuming a reality where omission errors are much more likely than commission errors, would make it even more difficult to learn about omission errors by gathering better data and all the more important to rely on similar competitors.

4. THE BOILED FROG IN ORGANIZATIONAL LEARNING

Introduction

“There is a quasi-scientific fable that if you can get a frog to sit quietly in a saucepan of cold water, and if you then raise the temperature of the water very slowly and smoothly so that there is no moment marked to be the moment at which the frog should jump, he will never jump. He will get boiled” (Bateson 1979, p. 98)

Organizational survival is periodically challenged by exogenous environmental change, i.e. a shift in the organization’s external reality which makes its current beliefs useless for identifying satisfactory solutions (March 1991). Thus, to survive, organizations must be able to sense the change in the environment and adapt (March 1991, Teece 2007). While some adapt to change and thrive, others perish. This heterogeneity in organizational responses to exogenous change has been attributed to differences in routines (Benner 2009), strategy (Hannan and Freeman 1984), leadership (Tushman and Romanelli 1985) and generally, the balance between exploration and exploitation (March 1991, Posen and Levinthal 2012). In this paper, we argue that characteristics of the external change itself widely determine organizational survival. Specifically, we propose that the pace of environmental change affects an organization’s capability to sense, adapt to, and survive exogenous environmental change.

Various scholars have described the consequences of the inability to sense radical environmental changes (Anderson and Tushman 1990, Kaplan et al. 2003), while others have described the organization’s ability to adapt, conditioned on having been able to sense the change (Keister 2002, Kiesler and Sproull 1982, Levinthal and March 1981). The ability to sense is an important antecedent of the adapting process, i.e. the ability to learn from experience about an environment that has changed. This ability depends on the pace of environmental change, i.e. whether the environment changes in a more continuous or a discontinuous manner. We argue that it is harder to sense continuous change for reasons

analogous to our initial quote on the boiled frog, i.e. an organization's ability to sense continuous change is limited.

In developing our arguments, we draw on two strands of literature. First, we use arguments from the organizational learning and adaptation literature to flesh out the relationship between environmental change and sensing by experiential learning (Posen and Levinthal, 2012, Stieglitz et al. 2016). Second, we draw upon the literature on environmental change to examine the precipitousness of environmental change (Asgari et al. 2017, Maula et al. 2013 Tushman and Anderson 1986). Specifically, we define environmental change as more or less continuous or discontinuous. By borrowing arguments from the learning and adaptation literature and applying them to the change literature, we claim that sensing change by experiential learning implies the ability to differentiate a signal of change from mere noise (Levinthal and Posen 2007). Naturally, this may require a performance shortfall that exceeds the level of noise in a stable environment (Posen and Levinthal 2012). If change, however, occurs continuously, the performance shortfall would rarely exceed the level of noise at any particular time. Thus, we argue that continuous change may be as harmful as its discontinuous cousin, as continuous change often goes undetected.

We test this claim analyzing the lifespans of restaurants facing gentrification. Specifically, we analyze how the pace of gentrification affects a restaurant's operating duration. We find that both very a fast and very slow pace of gentrification is associated with a reduced life expectancy. For faster changing restaurant environments, we argue that a lower life expectancy stems from the inability to adapt on time. Conversely, for slowly changing restaurant environments, we theorize that a lower life expectancy stems from the inability to sense the change on time. This is because organizations experiencing a continuous performance decline face the difficulty of telling apart (i) an altered environment from (ii) random noise in an unchanged environment. In the first instance, organizations would ignore

the change, whereas in the second, organizations would unlearn old beliefs and start again. Consider, for example, a restaurant owner offering Bologna sandwiches. If she finds that fewer people are coming into her store this could be due to the fact that many regular customers happen to be on vacation or it could be because the neighborhood around her has changed such that its new inhabitants have less appreciation for cold cuts and demand different food choices like salad wraps. In the former case, waiting until customers return from vacation while leaving the menu unchanged will do the trick. In the latter case of a changed neighborhood, however, she would need to forget about Bologna sandwiches and learn a new dish like the aforementioned salad wraps. In actuality, this mechanism could for example be observed in the US automobile industry's response to the rise in imports of foreign cars over the past decades. Virtually insignificant in the 1970s, Toyota alone increased its market share in the US from three to 17 percent. Quite unable to tell that a permanent shift in customer preferences had occurred, US automakers' market share fell dramatically. While the big three US manufacturers of cars (GM, Ford, and Chrysler) made had a combined market share of more than 80 percent in the late 1980s, this share fell through 2009 to less than 50 percent.

Given that the main source of learning is performance feedback, organizations need to determine a threshold between actual performance and expected performance beyond which it would perceive a performance shortfall as a change. It may be feasible to determine such a threshold when the level of noise is low, and changes are sudden. When however, change is gradual, it is harder to detect change.

We primarily add to the literature on radical environmental change (Anderson and Tushman 1990, Sosa 2011, Tushman and Anderson 1986) by expanding the concept of radical environmental change beyond cases of discontinuous change. Indeed, we show that continuous change may alter the value of strategic choices at least as revolutionarily as

discontinuous change. Continuous change may even be more harmful than discontinuous change because it could go unnoticed, which would thereby lead organizations to act on outdated beliefs and reduce the ability for future learning (Levinthal 1997) and ultimately, performance. Contrary to much of the change literature, however, we analyze the impact of exogenous environmental change as opposed to technological change brought about by one of the players (e.g. Benner 2010). We also contribute to the literature on organizational learning and adaptation by examining the capability to sense rather than just adapt to environmental change. While some recent work has looked at the sensing of change (Chakravarti et al. 2011, Green and Shapira 2018), a systematic analysis of conditions under which sensing is possible as well as an empirical test thereof, is absent from the literature.

Theory

In this section, we apply arguments from the literature on organizational learning and adaptation to the realm of radical environmental change. First, we try to parse out the relationship between the ability to sense change and environmental dynamism. We then use the insights on such relationship to differentiate between two types of radical environmental change (continuous and discontinuous) and flesh out the neglected perils of continuous radical change (in the simulation section, we explicate the sensing mechanism theorized upon in a simulation experiment).

Sensing change in the environment by experiential learning

Experiential learning leads to a better fit of an organization with its current environment. More experience lets organizations repeat strategies that are beneficial under current circumstances while discouraging the use of strategies that are less beneficial (Fang 2012, Levinthal 1997). As a consequence, however, organizations are less apt to perform well in a different context (Miles and Snow 1984). Thus, a lower performance can typically be

expected when a fundamental shift in the environment goes unnoticed (Haveman 1992, Henderson 1993, Nickerson and Silverman 2003, Tushman and Romanelli 1985). To prevent this from happening, organizations engage in sensing (Kiesler and Sproull 1982). Sensing change in an organization's environment is thought to happen by contrasting a stimulus against some level of aspirations (Kiesler and Sproull 1982). In concrete terms, this means that an organization compares its current performance with its performance target, and detects a problem when the current performance is below the target (Cyert and March 1963). The perceived performance shortfall could then be attributed to factors internal or external to the organization such as to a radically changed environment (Siggelkow 2001). Both conceptual and experimental work have suggested that decision makers can only sense environmental change if the current performance deviates by some magnitude from their own beliefs about the environment (Chakravarti et al. 2011, Kiesler and Sproull 1982). The threshold for the difference between beliefs and state of the environment to trigger sensing differs across organizations. Beliefs about the environment in turn are formed by experiential learning. Learning takes place "by encoding inferences from history into routines that guide behavior" where routines refer to -among other things- the "structure of beliefs" (Levitt and March 1988, p. 320). This implies that learning is about forming beliefs about the value of solutions available to the organization. Beliefs are refined by continued experiential learning. Thus, experiential learning is a prerequisite to sensing radical environmental change, because only when the organization holds narrow beliefs about what performance it can expect, can it reliably sense change from a sudden performance shortfall. Its ability to sense also depends on the threshold it sets to differentiate between noise and a signal for a changed environment: A low threshold will lead to noticing change more often but also to a greater rate of mistaking noise for change. Likewise, a high threshold will lead to noticing change less often but also to a lower rate of mistaking noise for change.

Limitations of the ability to sense change in the environment lie in the “signal-to-noise ratio (Green and Swets 1966, Kiesler and Sproull 1982, Peterson et al. 1954) where the signal is the performance shortfall due to change, and the noise is random variation in a solution’s payoff in a stable environment. The more noise characterizes the environment, the harder it is to sense the genuine signal of change. This is so, because it is difficult for a decision maker to attribute a drop in performance to either radical environmental change or random noise in a stable environment. Beyond a level of noise equal to the signal of change, attempts of sensing is merely guesswork. Hence, it follows that, when the level of noise exceeds the amount by which the solution exploited by an organization decreases in value, an organization cannot sense the change in its environment.

Further, the ability to sense may be impaired by organizational memory loss, i.e. the fact that lessons from experiential learning (Argyris and Schon 1978) may not always be retained or may be forgotten over time (Haunschild et al. 2015, Walsh and Ungson 1991).

Unlearning in response to sensing environmental change

Sensing in itself is not easy to identify, because we cannot access the mental representation decision makers have of their environment (Green and Shapira 2018, Teece 2007). Prior research, however, suggests that often times sensing is accompanied by a distinct change in learning behavior: unlearning, which allows for the subsequent relearning of old beliefs. This is so because the same experiential learning that makes possible the sensing of change also hinders relearning of the value of old solutions in the changed environment. Previous work has proposed that outdated beliefs may be a threat to adaptation in the changed environment (Levinthal 1997) and may best be erased from the organization’s memory, i.e. unlearned (Starbuck 2017). In sum, when radical change is identified, unlearning of old beliefs appears to be the best way to deal with the altered environment as otherwise the burden of inaccurate beliefs would hinder adaptation to the new environment (Asmussen et al. 2016, Cirillo et al.

2014, Holan and Phillips 2004, Kim 1998, Starbuck and Hedberg 1977). Hence, “changing situations” require, among other things, the “unlearning of prior premises” (Garud et al. 1997, p. 298). Overall, an organization will set a threshold between expected performance and observed performance based on the level of noise at which its performance usually varies. When an organization experiences a drop in performance exceeding this threshold, the organization will attempt to unlearn old beliefs.

Types of radical environmental change

Fatal disruptions, arising in an incumbent organization’s environment, are commonly linked to discontinuous change (Tushman and Anderson 1986) or even equated with radical change (Tushman and O’Reilly 1996). These forms of radical change are seen as competence destroying and believed to alter the values attached to a firm’s set of choices dramatically and instantaneously (Tushman and Anderson 1986). Overall, the attribute often highlighted about radical environmental change is its suddenness (Tushman and Anderson 1986). This means the change is assumed to take place in a rather short period of time (Asgari et al. 2017, Maula et al. 2013). This aspect is also reflected in the fact that the change is often referred to as a *discontinuity* (e.g. Sosa 2011).

Meanwhile, radical change that happens continuously is not often considered. In Bateson's (1979) boiled frog analogy, he explains that if you were to put a frog in boiling water, it would sense the danger and jump out. However, if you were to place it in colder water and then gradually turn up the heat, it would not be able to sense the threat and die. Likewise, an organization for which the environment is changing completely but over an extended period of time would remain unaware of the radical change and not be able to adapt to the new situation in time. We accordingly define *continuous radical change* as change that alters the valuations of all the solutions available to an organization completely, but which happens over an extended period of time. There are many recent examples of situations

where radical change took place in a continuous way. For example, the US postal services lost most of their mail business to electronic mail providers over a period of the last 20 years while the losses only started to exceed self-set performance targets since 2011 (Pociask 2016). This resulted in a radically changed situation for them that, however, did not happen right after the new email technology arrived. Another example can be found among firms in the grocery industry, where consumer preferences have radically shifted from home cooked to prepared meals (Yoon 2017). This change has taken place over the course of almost 50 years. To our baseline argument above, this type of change clearly provides a challenge: Even for moderate levels of uncertainty it is hard to discern radical change from environmental noise if the change is taking places in such incremental steps. Hence, we propose

Hypothesis 1. When an organization experiences continuous radical change, sensing becomes more challenging even for low levels of noise. Therefore, the more continuous the nature of the change, the slower an organization will unlearn.

Implications for organizational survival

Prior research has often linked threats to an organization's survival with the failure to sense radical environmental changes (Siggelkow 2001, Tripsas and Gavetti 2000). This line of reasoning is as follows: Organizations that fail to sense change in their environment act on outdated beliefs (Levinthal 1997). From engaging in learning about the environment prior to the change, the organization has learned to pursue solutions that had proven to be profitable while ignoring others that had proven to be unprofitable (Denrell and March 2001). Once the environment has changed, however, the previously profitable solutions may now be inferior while the previously unprofitable solutions might earn the organizations great profits. Acting on outdated beliefs thus means consistently realizing inferior profits by exploiting the wrong

solutions. Consistent under-performing relative to those organizations that do sense change, however, leads to a higher likelihood of demise (Klepper 2002). It follows that

Hypothesis 2. An organization experiencing continuous radical change will have lowered survival chances all else being equal.

At the same time, discontinuous environmental change may also be associated with lower organizational chances of survival. An immediately altered environment makes timely implementation of new solutions necessary (Benner 2010). In particular, it means replacing old solutions with new solutions. Since organizations tend to get inert in their ways over time, they often fail in replacing old solutions with new ones in time (Sosa 2011). As more continuous change as well as sudden, discontinuous change results in lower survival chances, we argue that

Hypothesis 3. There will be an inverted U-shape between the pace of environmental change and organizational survival.

Gentrification as environmental change

The Merriam-Webster dictionary introduces gentrification as the “the process of repairing and rebuilding homes and businesses in a deteriorating area (such as an urban neighborhood) accompanied by an influx of middle-class or affluent people and that often results in the displacement of earlier, usually poorer residents”. In New York, gentrification is also called “brownstoning” as it started by members of the middle class buying and renovating old brownstone houses in poorer neighborhoods. The process of gentrification often follows the same steps. To first qualify for later gentrification, richer inhabitants need to trade in their homes in a formerly affluent neighborhood for suburban homes. The abandoned neighborhood attracts low income residents involving a decrease in investments and a decrease of the quality of neighborhood infrastructure. When now inner city living becomes more attractive and house prices are low, gentrification sets in. At first, middle class residents

buy homes in the run-down neighborhood which they renovate themselves. Those early stage gentrifiers are not primarily driven by the potential increase in house prices but rather by living in an urban community and by the historic value of the houses they are buying. This type of home ownership is also called sweat equity. As a consequence of the efforts of the first wave of gentrifiers, quality of life improves in their communities and others start noticing the potential for real estate in gentrifying neighborhoods as an investment. In a second step, investors start buying real estate and market it to wealthy buyers. At the same time, businesses which are catering to the new tenants start pouring in, for example, vegan restaurants, bars, or Starbucks cafes. To make room for wealthier tenants, investors also engage in active displacement, i.e. they try to evict poorer members of the community to be able to renovate and sell properties to a more affluent group of tenants.

In sum, the process of gentrification is characterized by an increase in home prices, evictions of poorer tenants and an influx of bars and cafes catering to the new middle class.

Take, for example, the history of the Mott Haven neighborhood in the South Bronx. After Jordan Mott's ironworks had set up shop in the 19th century, rows of brownstone houses were constructed in the present-day Mott Haven historic District to house the managerial elite. This process aided by commuter trains beginning to service the Bronx. Post-World War II, following an influx of Puerto Rican immigrants and a greater accessibility of the suburbs by highways, many of the white residents moved out of Mott Haven and into the suburbs. From the 1970s onwards, Mott Haven was ridden with gang violence, poverty and poor living conditions. In recent years, however, gentrification has taken hold of Mott Haven. First, local investments and efforts turned vacant spaces into community gardens and low-income co-ops. Next, investors bought waterfront property and developed them into luxury lofts. The investment in housing is met also with an increase in fancy coffee shops, bars and

other businesses. At the same time, eviction rates are increasing - proof of the ongoing displacement of low-income tenants.

Method

Context

The restaurant industry is ideal to test the effect of radical environmental change on organizational performance because it provides a sizable sample with a large variance of firm survival rates as well as a frequently changing environment by virtue of demographically changing city neighborhoods (Kalnins et al. 2006). Specifically, as survival rates vary greatly across organizations, restaurants offer the unique advantage of providing large number of organizations exposed to the same environmental changes. Further, restaurants usually have shorter lifespans than other types of organizations making the necessary time window shorter (Hannan and Freeman 1983). The setting also constitutes one of the most important industries in the US and in New York City. Nationwide, restaurant sales gross at around 785 billion US dollars while food establishments employ around 10 percent of the country's workforce. In New York City alone, eateries sell food worth around 43 billion US dollars and employ more than 800,000 people¹.

Data

The first source of data is the permit database of the New York City Department of Health and Mental Hygiene (DOHMH). It contains information on when an operating permit (Food Service Establishment Permits) was first granted and when a permit ultimately expired. Such permits are a necessary prerequisite to operating any business in New York City that serves food. Thus, it gives an accurate overview of the lifespan of restaurants in New York City.

The cleaned data set contains 30,392 entities. In each given year, there were between 10,000 and 14,000 restaurants active (Figure 4-1), most of them located in Manhattan (Figure 4-2).

The earliest permit in the data was granted July 2007 and the last one in August 2018. The average time a restaurant was in business is 662 days. The DOHMH data set further contains information on cuisine type, permit violations, owner, and exact location of the restaurant. In the following, we leverage this location information to assign each restaurant to a gentrification rate in its neighborhood.

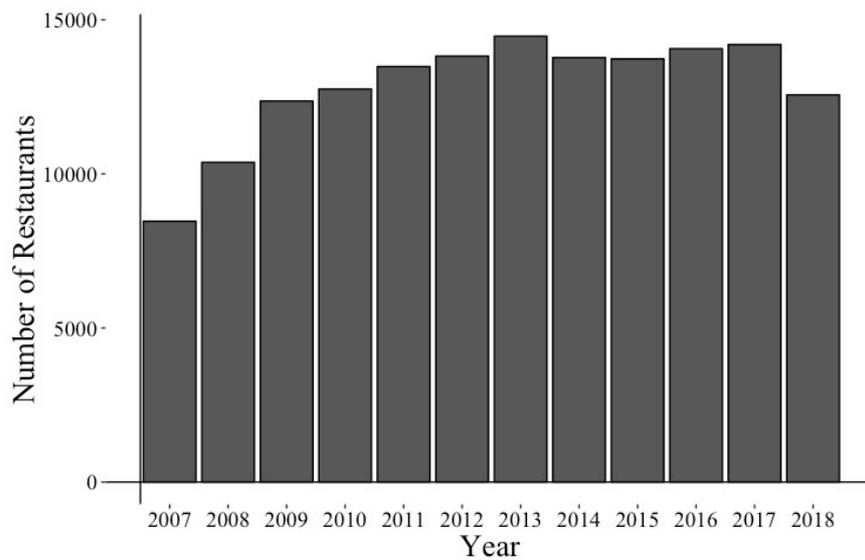


Figure 4-1: Number of active restaurants over the years

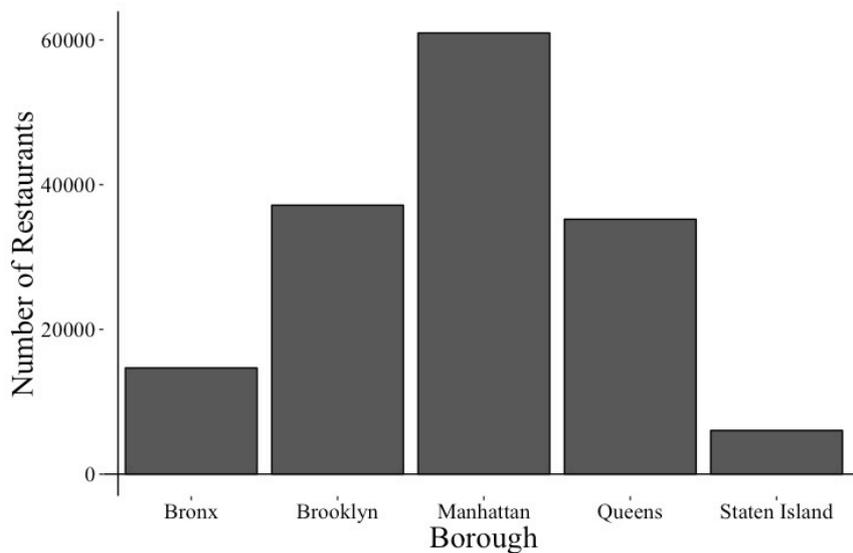
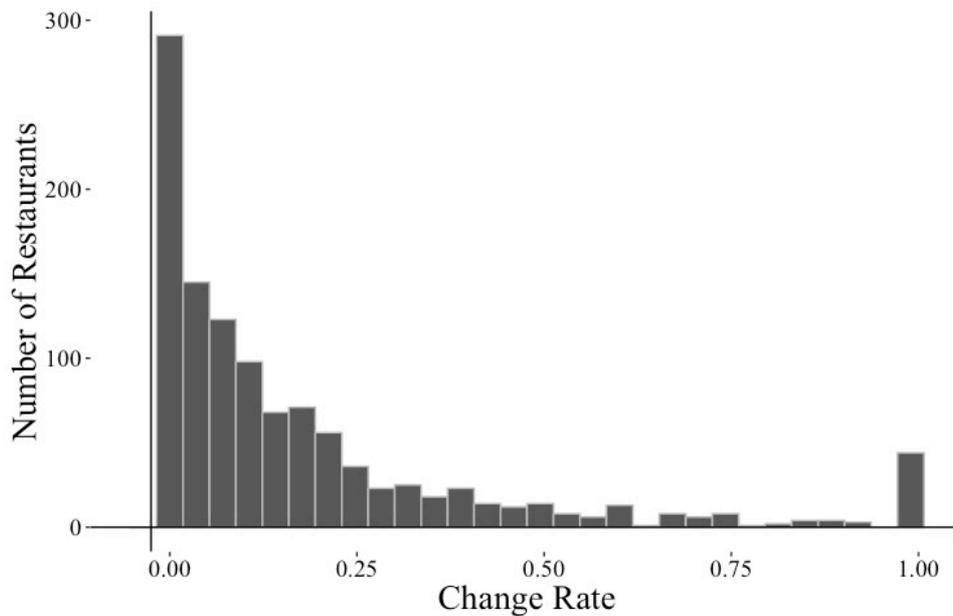


Figure 4-2: Number of active restaurants across boroughs

We use different data to approximate the pace of gentrification. First, we use evictions. We extract data on the number of evictions filings for each New York neighborhood from data provided by the NYC Department of Investigation. Eviction filings are closely associated with gentrification as landlords have an incentive to replace space for poorer tenants with offers catered to incoming richer tenants. Secondly, we use the entry rate of new cafes and bars into a given neighborhood to approximate gentrification. The influx of cafes and bars has been shown to be a hallmark of gentrification (Glaeser et al. 2018). We estimate entry rates of cafes and bars using the same data from the DOHMH which contains additional details on the type of a restaurant which allows us to distinguish cafes and bars from other types of eateries.

We use data from the Federal Housing Finance Agency as well as zillow.com on the development of home prices. We are particularly interested in the 10-year annualized home value index. The agency's house price index is based on the development of single-family house price per year and is a "weighted, repeat-sales index, meaning that it measures average price changes in repeat sales or re-financings on the same properties" (*Federal Housing Finance Agency*, 2018). The Zillow index is based on a machine learning estimate of all home values in a zip code sold or not. We use the FHFA index as a control.

Additionally, we collected data on changes in restaurants' menus during our observation period from 2007 to 2018. Concretely, we collected menu texts from restaurant websites and yelp.com. We then used the *wayback machine* to find older versions of the same menu and compare the texts to see to what degree it had changed. All in all, we gather information on 1395 restaurants in this way. The distribution of change rates is depicted in Figure 4-3.



Note. Change refers to changes in items of sections on a menu.

Figure 4-3: Distribution of menu change rates

Estimation

To test whether very slow and very fast environmental change have more negative performance effects than change rates in between the extremes, we run a simple OLS model with fixed effects for the type of restaurant cuisine. The results are robust to other estimations such as a propensity score matching. We exclude extraordinarily high values of evictions as those will likely be the consequence of a high poverty neighborhood as opposed to gentrification (Desmond and Gershenson 2017).

Dependent variable

The dependent variable is the restaurant lifespan. The information for this variable was extracted using the DOHMH operating permits. Lifespan is measured in days of having an operating permit.

Independent variables

The central independent variable is the entry rate of new bars. The information was taken from the same DOHMH database counting the number of new bars entries over the observation window from December 2007 to April 2018. In the same way, we used the entry

rate of new cafés as a secondary explanatory variable. Lastly, we used the rate of eviction filings as another way to approximate the pace of gentrification. In particular, we used the average number of evictions filed per year for each neighborhood. The information was taken from the New York City Marshals. The entry of new cafes or bars is a good indicator of gentrification as this type of businesses is typically associated with an influx of a younger, more affluent demographic (Glaeser et al. 2018). The eviction filings measure the other side of the same coin, i.e. the displacement of poorer residents (Lee 2003).

Control variables

We control for the number of employees as a proxy for firm size. Further, we control for changes in rent prices. The latter is important to separate effects driven by a change in the exogenous environment - i.e. the customer base has changed - from a change in costs. Through including cuisine fixed effects we also exclude that effects are driven by cuisine specific trends.

Results

The effect of gentrification pace on restaurant lifespan

To find the effect of gentrification pace on restaurant lifespan, we regress lifespan – i.e. how many days a restaurant survived for - on the entry rate of new bars (see Table 4-1, model (1)). It is important to note that we consider only restaurant lifespans for restaurants that are not bars to avoid including the same observations in both the left and the right-hand side of the regression equation. We further exclude neighborhoods with an entry rate of less than 1 bars per year this would not constitute radical change. Since we hypothesized the effect to take the form of an inverted U-shape, we include also the squared effect of the entry rate. In line with our argument we find that the coefficient for entry rate is positive and significant ($p = 0.035$) while the coefficient for the squared term is negative and significant ($p = 0.035$). We plot the

curvilinear effect in Figure 4-4. The effects also hold when using the entry of cafes rather than the entry of bars. Besides our main effect, both restaurant size ($p < 0.000$) and change in rental prices ($p = 0.004$) have a positive and significant effect on restaurant lifespan.

<i>DV: Lifespan</i>			
<i>Variable</i>	<i>(1)</i>	<i>(2)</i>	<i>(3)</i>
<i>Entry Rate Bars</i>	20.31** (7.86)		
<i>Entry Rate Bars²</i>	-0.24** (2.18)		
<i>Entry Rate Cafes</i>		10.28 (9.09)	
<i>Entry Rate Cafes²</i>		-0.11 (0.10)	
<i>Avg. Evict. Filings p.a.</i>			3.66* (1.75)
<i>Avg. Evict. Filings p.a.²</i>			-0.005. (0.003)
<i>Size (# staff)</i>	20.81*** (0.91)	20.95*** (0.77)	14.74*** (1.22)
<i>ΔRent Prices</i>	2284.00** (812.73)	100.10 (1000.18)	9346.70** (848.97)
<i>Fixed Effects</i>	<included>	<included>	<included>
<i>R²</i>	0.15	0.15	0.09
<i>Adj. R²</i>	0.13	0.14	0.05
<i>Note: *p<0.05, **p<0.01, ***p=0.001</i>			
<i>Standard errors in parentheses</i>			

Table 4-1: Gentrification and restaurant lifespan

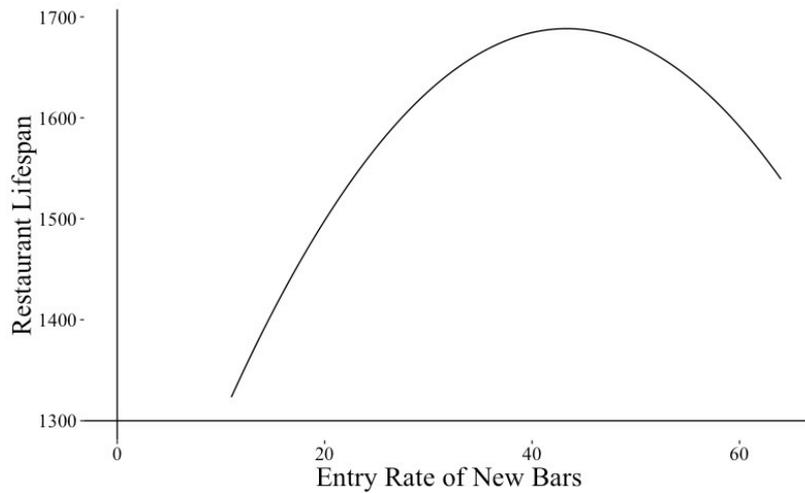


Figure 4-4: Inverted U-shape of the effect of entry rate of bars on restaurant lifespan

Robustness check: entry of cafes as a measure of gentrification

As an additional check, we conduct the same analysis with the entry of new cafes rather than bars (see Table 4-1, model (2)). Results are consistent with findings using the entry of bars. The coefficient associated with the entry rate of cafes is positive and significant ($p = 0.024$) while the coefficient associated with the quadratic term is negative and significant ($p = 0.011$). Notably, an increase in rent prices is no longer significantly associated with a change in restaurant lifespan. The inverted U-shape is depicted in Figure 4-5.

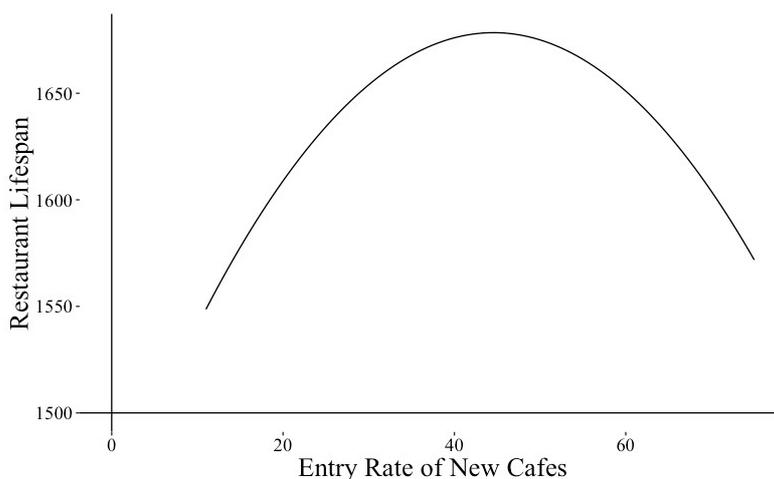


Figure 4-5: Inverted U-shape of the effect of entry rate of cafés on restaurant lifespan

Robustness check: eviction filings as a measure of gentrification

We repeat the analysis by replacing the independent variable with the number of eviction filings (see Table 4-1, model (3)). Coefficients take the same direction as with the entry rates of cafes. There is a positive effect of eviction filings on restaurant lifespan and a negative effect of the squared term. However, neither the effect of filings ($p = 0.137$) nor the effect of the square term ($p = 0.292$) become as significant as the effects of entry rates. We plot the inverted U-shaped relationship between eviction filings and restaurant longevity in Figure 4-6.

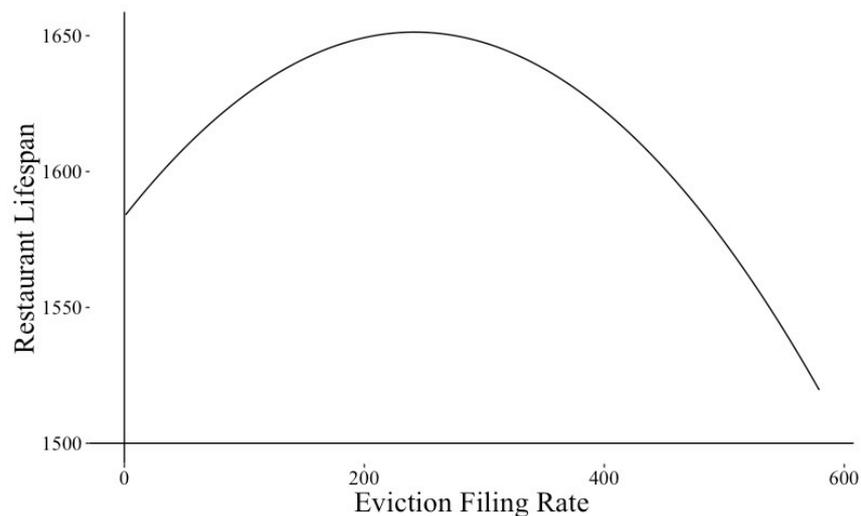


Figure 4-6: Inverted U-shape of the effect of the rate of eviction filings on restaurant lifespan

Robustness check: repeating the analysis with only those restaurants founded in the observation window

To check that the effects were not driven by events in the history of the restaurants prior to our observation period, we repeated the main analysis leaving out those restaurants that were founded prior to December 2007. The main curvilinear effect of gentrification pace on restaurant lifespan remains unchanged. In particular, the coefficient for entry rates (of new

bars) is positive and significant ($p = 0.044$) and the coefficient of its squared term is negative and significant ($p = 0.034$). The effect size is stronger than when leaving all restaurants independent of their founding date in the observations. Results are shown in Table 4-1a. In line with our argument we find that the coefficient for entry rate is positive and significant ($p = 0.035$) while the coefficient for the squared term is negative and significant ($p = 0.035$).

<i>DV: Lifespan</i>	
<i>Variable</i>	<i>(1)</i>
<i>Entry Rate Bars</i>	43.93* (21.80)
<i>Entry Rate Bars²</i>	-0.70* (0.33)
<i>Size (# staff)</i>	16.17*** (1.04)
<i>ΔRent Prices</i>	714.79 (812.73)
<i>Fixed Effects</i>	<included>
<i>R²</i>	0.11
<i>Adj. R²</i>	0.08
<i>Note: *p<0.05, **p<0.01, ***p=0.001</i>	
<i>Standard errors in parentheses</i>	

Table 4-1a: Gentrification and restaurant lifespan

Alternative explanation: the effect of fast gentrification on lifespan is driven by increased competition

One could argue that the hypothesized inverted U-shaped relationship between pace of gentrification and restaurant lifespan is driven by increased competition for fast gentrification and poverty in neighborhoods of slow gentrification. Following this logic, fast gentrification brings about more competitors and thus it is not timely adaptation but dealing with increased competition which drives reduced lifespans. Effectively this would mean that competition

would mediate the relationship between pace and lifespan. In a first step, for that to be true, gentrification pace would need to predict the level of competition. To test this, we construct competition as a density measure capturing the number of restaurants in a given neighborhood per 1,000 inhabitants. We then regress the level of competition on entry rates of new bars. This yields a positive and significant effect suggesting that competition actually may work through the gentrification rate. However, to truly mediate the effect of entry rate on lifespan, the former effect would have to subside when including both terms into the regression. But, the effect of entry rate on lifespan stays highly significant ($p = 0.007$).

Alternative explanation: the effect of slow gentrification on lifespan is driven by particularly poor neighborhoods

Following the same logic, slow gentrification may not drive reduced lifespans through a lack of sensing but in fact through high poverty which just does not generate enough of a market for entering restaurants. In practical terms, this would imply that the effect of gentrification on lifespan is driven by poverty. We therefore use the NYC.gov poverty index (i.e. what percentage of the population lives below the poverty threshold of ~\$33k a year). Again, we find no support that the effect of pace of gentrification becomes any weaker once a measure of poverty is included into the regression ($p = 0.009$).

Mediation analysis: How the inverted U-shaped relationship is driven by changes in adaptation

As previously discussed, we operationalize restaurant adaptation by changes in menus. In particular, we want to know if the degree of changes in menus mediates the effect of gentrification on restaurant lifespan. Results of the mediation analysis can be found in Table 4-2. We follow the mediation approach suggested by Barron and Kenny (1986). If adaptation truly mediates the relationship between gentrification pace and restaurant lifespan, we would expect to see a significant effect of gentrification pace (here: entry rates of bars) on restaurant

lifespan which we do observe (Table 4-2, model (1)). Further, we would expect to see a significant and positive effect of gentrification pace on menu changes where instead we find an insignificant effect (Table 4-2, model (2)). Lastly, we would expect the effect of gentrification pace on lifespan to become insignificant once the adaptation variable is included as an additional explanatory variable which we do find (Table 4-2, model (3)). Thus, overall we find limited support for a mediating effect of adaptation.

<i>Variable</i>	<i>DV: Lifespan</i>	<i>DV: Menu Changes</i>	<i>DV: Lifespan</i>
	(1)	(2)	(3)
Entry Rate	-3.29** (0.10)	-0.002 (0.002)	-9.29 (7.87)
<i>Menu Changes</i>			-52.97 (874.41)

Table 4-2: Mediation analysis (lifespan and adaptation)

Survival Analysis

Additional data

A second source of data is Census data. From the US Census Bureau, we extract household income data necessary to construct our gentrification measure. In particular, we were interested at the median household income in New York for the beginning of our observation period in 2007. Additional income data is collected from the American Community Survey (ACS) which contains information on zip code level household income data for the aggregate period of 2006-2010.

Operationalization

The dependent variable is the duration of a restaurant's stay in business or its age. The main independent variable is the change rate in a restaurant's zip code. The change rate is

operationalized by the gentrification rate. This gentrification rate we construct following an established method to measure gentrification (e.g. Herrine et al. 2016). In particular, we consider a zip code gentrified if (i) by the beginning of our observation period the average household income was in the bottom 40% of the city, and (ii) over the entire observation period, the home values increased by an above median amount. Note, that we observe different rates of increasing home values which we take as a measure for different degrees of continuous change. We are left with 13 gentrifying neighborhoods.

Results

In a first step, we gather evidence consistent with our claim that continuous radical change can be more harmful to organizational survival than discontinuous radical change. To do so, we compare survival chances of restaurants in faster versus slower gentrifying neighborhoods.

Since we have a substantial amount of right censored observations in our data, i.e. restaurants still in business at the end of our observation period, event history analysis is the appropriate method for our empirical test (Park and Russo 1996). The event on record is the case of a restaurant losing its operating permit.

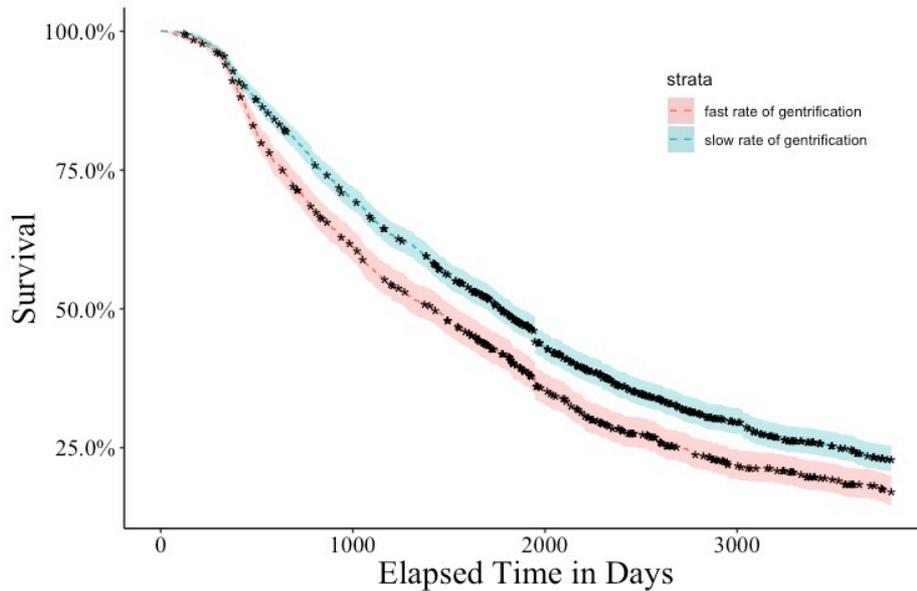
We use a parametric hazard rate model and in particular a piecewise exponential model that allows for different hazard rates for different time periods (Blossfeld et al. 1989).

$$r(t) = \exp(\alpha_i + A * \alpha_i),$$

where α will be estimated, A is the matrix of covariates and α_i is the constant element in period i .

Utilizing the official home price index of the FHFA as a proxy for change rates, we compute a binary variable reflecting whether a restaurant is in a neighborhood which is changing faster than the median gentrification rate or slower than the median gentrification rate. In Figure 4-7, we plot the survival rates for restaurants in slow versus fast gentrifying

neighborhoods. Restaurants in neighborhoods that gentrifying at a slower rate have lower survival chances. To test the significance of the difference in survival chances, we use a G-Rho test and find that a restaurant in a slowly changing neighborhood can be expected to survive more than a 100 days less than a restaurant in a fast changing neighborhood ($p < 0.001$). We show the test results in Table 4-3.



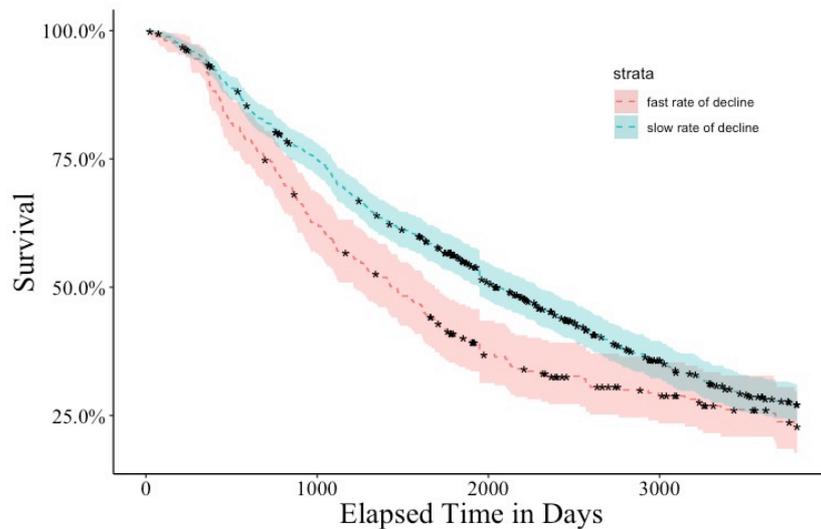
Note. From 13 NYC neighborhoods.

Figure 4-7: Differences in survival rates between restaurant exposed to gentrification at fast ($strata=0$) versus slow rates ($strata =1$)

	<i>N</i>	<i>Observed</i>	<i>Expected</i>	$(O-E)^2/E$	$(O-E)^2/V$
<i>Fast Gentrification</i>	1321	776	885	13.4	28.9
<i>Slow Gentrification</i>	1311	878	769	15.4	28.9
Chi ²	28.9				
Df	1				
p	0.00				

Table 4-3: Significance test between survival rates in Figure 4-7

Similar results can be found for the analysis of relatively declining rather than gentrifying neighborhoods (Figure 4-8).



Note. From 19 NYC neighborhoods.

Figure 4-8: Differences in survival rates between restaurant in relatively declining neighborhoods

The results from the survival analysis do not contradict the predictions of the simulation analysis but they do not show the mechanism of sensing and subsequent adaptation behavior as a cause of differences in organizational survival. To establish that firms' adaptation behavior indeed accounts for differences in survival, we analyze changes in restaurants' menus. We make use of the mediation framework of Barron and Kenny (1986). In first step, we thus regress survival chances on the gentrification rate to show that gentrification does drive restaurant survival. As expected, we find that a higher gentrification rate is associated with a lower likelihood of restaurant death (Table 4-4, model (1)). In other words, a slower change rate implies a shorter life expectancy.

In a second step, we regress the mediator (the change in restaurant menus) on the gentrification rate. We find no significant effect (Table 4-4, model (2)). Next, we regress menu restaurant survival on menu changes (Table 4-5, model (3)). We find that the more the menu changes, the less likely the restaurant is to vanish ($p=0.1$). This, too, is in line with our argument that it is a lack of adaptation that harms firm survival chances.

Lastly, we regress restaurant survival on both gentrification rate and menu change rate (Table 4-4, model (4)). Both variables negatively impact the likelihood of restaurant death. The effect of gentrification is now weaker than in our first regression albeit still highly significant ($p=0.001$). The effect of menu changes is weakly significant ($p=0.09$). In conclusion, the mediation analysis suggests that there is limited evidence for a mediating effect of changes in menus. We argue, however, that such limited support is to be expected due to (1) a small sample size and (2) the fact that changing menus is only one of several possible measures of adaptation.

<i>Variable</i>	<i>DV: Survival</i>	<i>DV: Menu Changes</i>	<i>DV: Survival</i>	<i>DV: Survival</i>
	(1)	(2)	(3)	(4)
<i>Rent Prices</i>	-0.29** (0.10)	-0.01 (0.05)		-0.32** (0.10)
<i>Menu Changes</i>			-0.19 (0.12)	-0.20. (0.11)
R^2	0.02	0.0003	0.008	0.04
Adj. R^2	0.02	-0.003	0.005	0.03

*Note: * $p < 0.05$, ** $p < 0.01$, *** $p = 0.001$*
Standard errors in parentheses

Note. The event of restaurant death is the dependent variable.

Table 4-4: Mediation analysis

Simulation Experiment

A model

Since we cannot empirically observe sensing in our restaurants, we conduct a simulation experiment for the purpose of exploring if our theoretical arguments are plausible. Further, using a simulation model we can test if the assumptions we are making on sensing behavior can lead to outcomes consistent with our empirical findings.

In the simulation experiments we model a simple choice task under radical environmental change. In most learning environments, decision makers have to sequentially decide which solution to implement thereby ignoring all other potential solutions. Specifically, they need to choose between exploring an unknown solution or exploiting a known one. Learning in such a scenario only occurs by trial and error. We are particularly interested in the effect of radical environmental change the values of available solutions are turned upside down. Underlying the simulation is the multi-armed bandit model (Holland 1975, March 1996, Robbins 1952). It is widely used to describe learning scenarios and has even been called the canonical model to describe the tradeoff between exploration and exploitation (March 1996, Posen and Levinthal 2012).

Task environment

The decision maker faces an environment of 10 potential solutions. All solutions have a fixed and a random payoff component. The fixed component (v_i) is drawn from a normal distribution with mean 0.5 and standard deviation 0.5. It stays constant throughout the task. The random component (ε) is drawn from a normal distribution with mean 0 and a variable standard deviation ranging from 0 to 1. This variable reflects the noise parameter in the environment. Thus, the payoff of solution i in a stable environment is

$$p_i = v_i + \varepsilon \tag{1}$$

Radical environmental change enters into the situation through redrawing solutions' payoffs from the initial distribution of payoffs. Thus, the change is stationary in the sense that the set of payoffs is unchanged but the assignment of payoffs to solutions differs before and after the change. We look at 2 types of change. First, there is discontinuous radical change. In this case, payoffs change from the old to the redrawn values from one period to the next. Second, there is continuous radical change. Here, payoffs change from the old to the redrawn values

in small increments. In particular, for a fraction of the 1,400 total periods, the payoffs are generated by taking values from a linear graph connecting the old and the redrawn values.

Beliefs updating

Much like decision makers in organizations decide for a strategy at the expense of other possible strategies in each fiscal year, the agents in this simulation choose a solution among a n potential solutions in each period. Initially the agent's beliefs are flat, i.e. she believes all solutions to pay 0.5. Beliefs are subsequently updated by experiential learning. Concretely, we rely on an incremental updating rule that has frequently been used in simulations modelling learning from performance feedback (Sutton and Barto 1998). We define the belief $B_{i,t+1}$ as the average of the belief in period t and the current payoff ($p_{i,t}$) weighted by a factor α that determines how much weight is put on the most recent versus more distant experiences.

$$B_{i,t+1} = (1 - \alpha) * B_{i,t} + \alpha * p_{i,t} \quad (2)$$

Choice rule

A common assumption about decision makers is that they choose what has had the highest payoffs in the past -i.e. best practices- in most occasions, but sometimes deviate from this behavior and choose a random unknown solution (Posen and Levinthal 2012). A simple rule capturing this type of behavior is an ϵ -greedy choice rule (Sutton and Barto 1998). Therein, in $1 - \epsilon$ cases, the decision maker chooses the solution she believes to have the highest payoff. In ϵ cases, she chooses a previously unexplored solution a random.

Sensing and unlearning

Beliefs built up prior to radical environmental change often hinders learning in post change times (Levinthal 1997). In dealing with this threat, theoretical and empirical accounts have maintained that unlearning old beliefs facilitates relearning in the altered environment (Cirillo et al. 2014, Hedberg 1981, Nystrom and Starbuck 1984) and ultimately, yields a higher

performance compared with continued regular belief updating. In the present model, an agent unlearns old beliefs when the observed payoff of solution i deviates by more than a threshold θ from her belief about this solution's payoff. In the next section, we discuss several alternative threshold levels. Once the difference between belief and observed payoff exceeds θ , beliefs are set back to the initial flat beliefs of 0.5.

$$B_{i,t+1} = \begin{cases} 0.5, & \text{for } B_{i,t} - p_{i,t} > \theta \\ (1 - \alpha) * B_{i,t} + \alpha * p_{i,t}, & \text{otherwise} \end{cases}$$

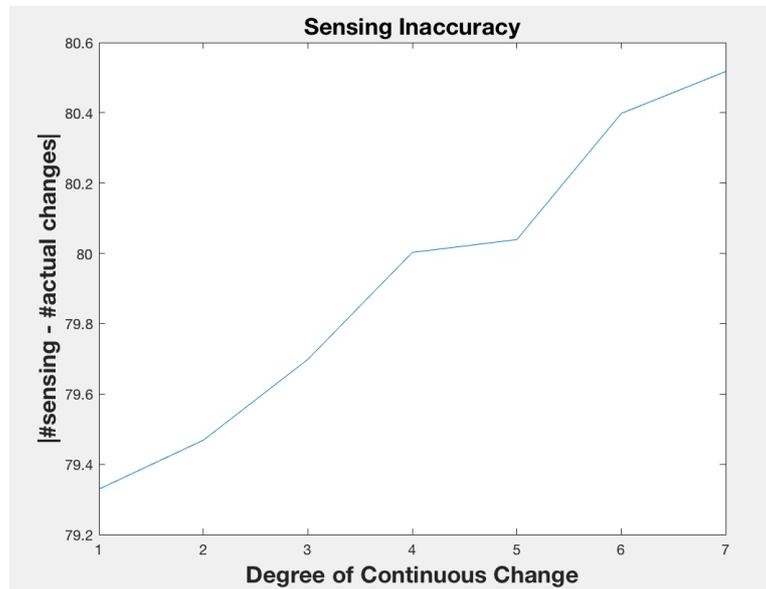
Simulation results

The purpose of the simulation experiments is twofold. First, we investigate the effect of the degree of continuous change on sensing accuracy. Second, we analyze performance implications of different types of radical environmental change, i.e. discontinuous and varying degrees of continuous change. As the performance feedback in our model has a noise component, we average over the results of 10,000 simulated organizations.

Sensing accuracy and degree of continuous change

We start our analysis by looking at the effect of the degree of continuous environmental change on an organization's ability to sense that change. Recall that to sense a change, the organization sets a value determining how far performance can drop below beliefs. This value also informs the organization when to reset those beliefs, enabling it to relearn the value of available solutions in the changed environment. Consequently, we measure sensing accuracy as the difference between the number of times an organization senses change, and the number of times change actually happens. In Figure 4-9, we show the effect of continuous change on sensing accuracy for a baseline model of a medium threshold between performance and beliefs (the simulation runs for 100 rounds, uncertainty = 0.01, learning rate $\alpha=0.7$, threshold=10%). Here, the degree of continuous change goes from a discontinuity (change happens from period to the next) to change that happens over the entire remaining

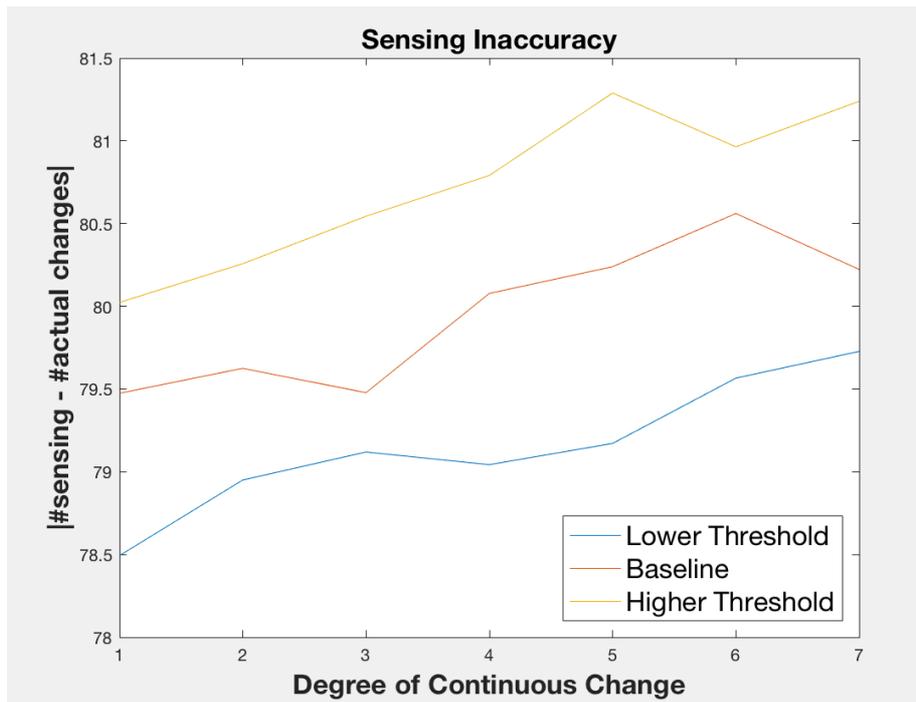
periods. As can be seen, the more continuous the change, the less accurate the sensing becomes.



Note. Lowest degree of continuous change is a discontinuity, highest degree means change is stretched out over increments over the entire remaining periods.

Figure 4-9: Sensing accuracy as function of degree of continuous change

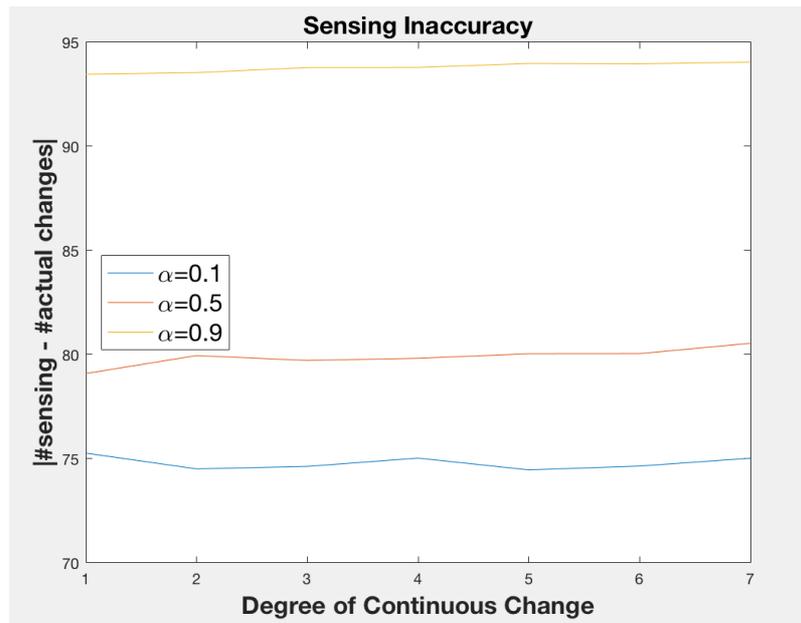
Next, we explore what happens when we lower or increase the sensing threshold (the simulation runs for 100 rounds, uncertainty = 0.01, learning rate $\alpha=0.5$, threshold lower=9%, threshold increased=11%). Intuitively, we will assume that when the threshold is low, change in the environment is sensed, but that random variation is also mistaken for such change. When the threshold is high, those mistakes happen less often, but the likelihood of sensing the actual change in the environment is low. As can be seen in Figure 4-10, sensing accuracy decreases most from increasing the sensing threshold. Put differently, it is beneficial for the organization to set a low sensing threshold than a high threshold – i.e. in effort to avoid mistaking noise for change.



Note. Lowest degree of continuous change is a discontinuity, highest degree means change is stretched out over increments over the entire remaining periods.

Figure 4-10: Sensing accuracy moderated by sensing threshold

Next, we analyze the moderating effect of the learning rate on sensing accuracy. We define the learning rate (α) as the weight an organization puts on recent versus more distant experiences. In Figure 4-11, accuracy curves are depicted for an organization focused mostly on recent experiences ($\alpha=0.9$), a balanced learning strategy ($\alpha=0.9$), and a focus on distant experiences ($\alpha=0.1$). Put differently, organizations which are fast learners (i.e. place *more* value on recent signals from their environment) suffer greater sensing inaccuracy than organizations which are slow learners (i.e. place *less* value on recent signals from their environment).



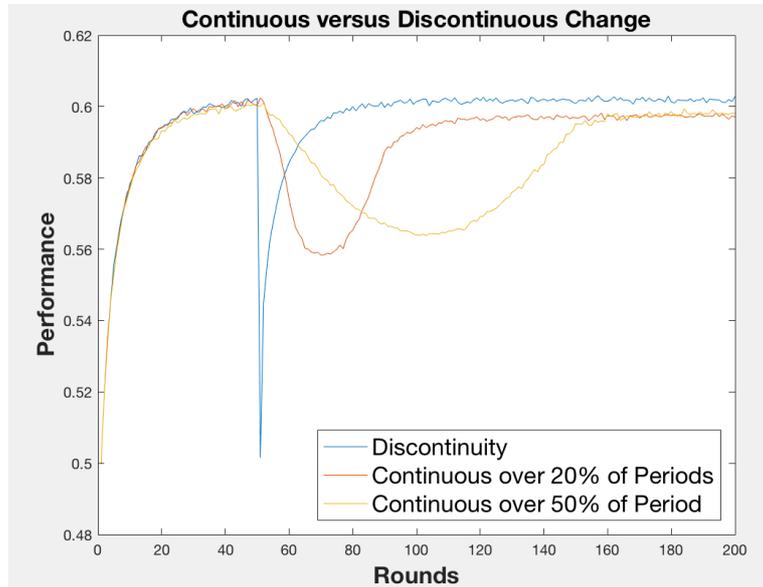
Note. Learning rate given by α in $B_{i,t+1} = (1 - \alpha) * B_{i,t} + \alpha * p_{i,t}$

Figure 4-11: Sensing accuracy moderated by learning rate

Performance implications of different degrees of continuous change

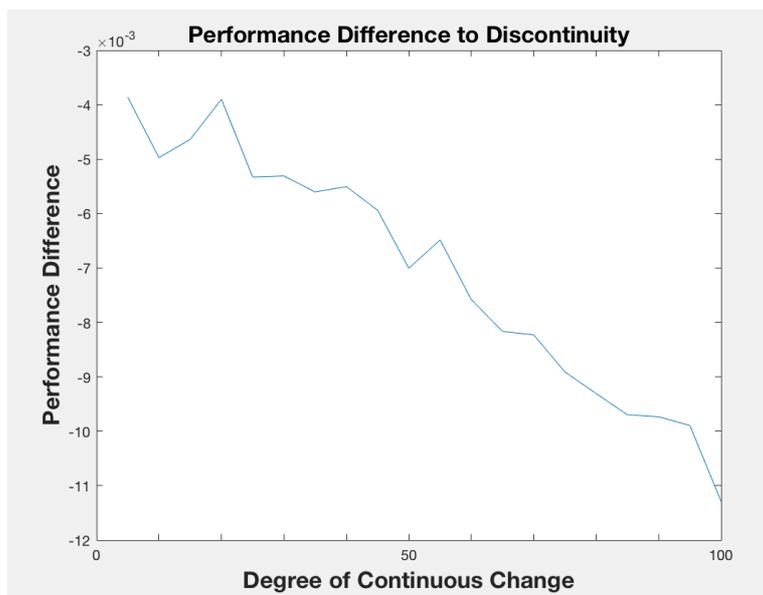
In the following, we turn to performance implications of different rates of environmental change. We begin by plotting organizational performance over time (the simulation runs for 200 rounds, uncertainty = 0.01, learning rate $\alpha=0.5$, threshold=10%). We plot the performance effects over 200 rounds instead of a 100 as before to make clear the equilibrium or final level of performance organizations return to after the change has happened. The graphs in Figure 4-12 reflect discontinuous change, continuous change over 20 percent of the total rounds and continuous change over 50 percent of the total rounds. As can be seen, a discontinuity results in a drastic performance loss that the organization would then be able to recover from quickly. The more continuous the change in the environment, the more prolonged the performance loss – i.e. the longer it takes to recover to old performance heights. Overall, total performance remains on average lower when change happens continuously. To corroborate this insight, we plot performance deviations from the performance in the discontinuity case as a function of the degree of continuous change. The

degree of continuous change is operationalized as the percentage of all periods over which radical change is stretched out. As becomes apparent in Figure 4-13, performance decreases with a greater degree of continuous change.



Note. Discontinuous change means solutions' payoffs are redrawn from one period to the next, continuous change means values in each period are drawn from a linear function connecting the old payoffs and the redrawn ones. The length of that linear function is 20% of all periods for the orange line and 50% of all periods for the yellow line.

Figure 4-12: Performance implications of different degrees of continuous change



Note. Degree of continuous change in percent of total remaining rounds after initial change over which change is stretched out.

Figure 4-13: Performance deviation from discontinuity with increasing level of continuous change

Alternative sensing by looking for trends

An alternative way of sensing change may be to average performance feedback over multiple periods to detect a trend. This could be done by either averaging over the changes from period to period over a number of periods or running a linear regression model. In the following we compare the results for these two approaches. In the first case, instead of comparing the performance feedback from a choice against expectations thereof, the agent averages the changes performance of the last 25 periods. Subsequently, she senses change and resets her beliefs when the average change in performance over the last periods is negative. Performance implications of this behavior are depicted in Figure 4-14. While the relationship between continuous and discontinuous change stays the same, overall agents using the averaging heuristic perform worse than in our prior threshold model, because it is now more likely to trigger un- and re-learning. However, as Kahneman et al. (1997), among others, have pointed out, decision makers rarely actually apply an averaging heuristic. Therefore, in the second case, we the agent runs a simple linear regression over the performance feedback on a decision alternative she receives over the last 25 periods. As can be seen in Figure 4-15, performance implications are similar to those of our original simulation. This is so because while sensing may become more accurate, the time it takes to “collect” the necessary data to detect change is now longer and hinders immediate sensing.

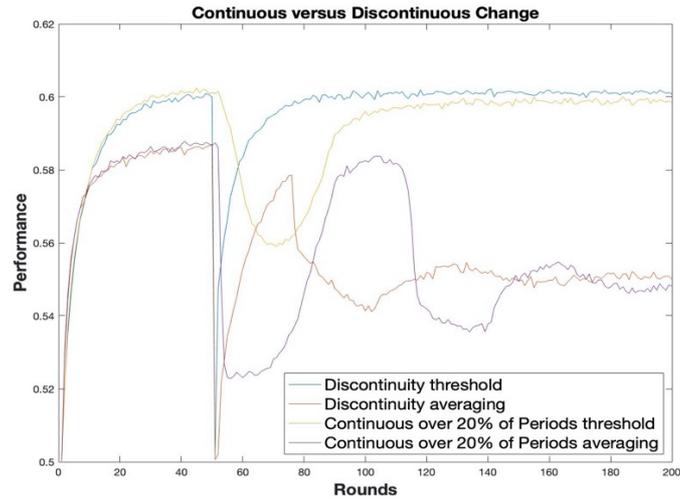


Figure 4-14: Performance comparison with averaging heuristic

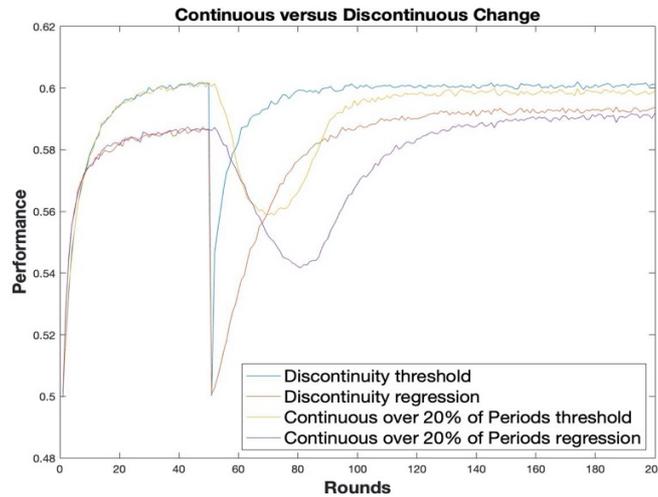


Figure 4-15: Performance comparison with regression heuristic

Kahneman et al. (1997) also suggest yet another heuristic, decision makers may use: They may recall only the worst experience in their recent performance history. In Figure 4-16, results of operationalizing such a minimum heuristic are displayed. In concrete terms, the agent compares the worst performance she experienced in the last 25 rounds to the worst performance of the 25 rounds before that. Performance implications are similar than for the averaging heuristic. Here, too, accidental triggering of change is frequent.

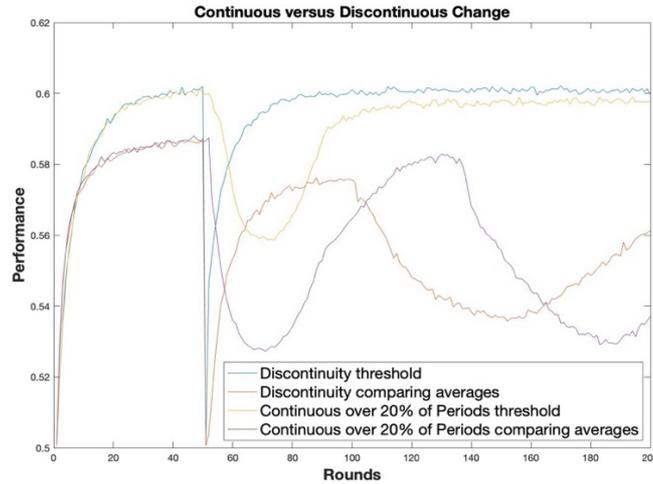


Figure 4-16: Performance comparison with minimum heuristic

Alternatively, decision makers may consider a mix of a threshold model and analyzing the recent performance trend. In Figures 4-17 and 4-18, the averaging and minimum heuristic are being used but with an additional threshold of 10%. This means, change is detected when the current minimum is at least 10% worse than the prior minimum, or the current average is at least 10% worse than the prior average. Results suggest that performance improves overall by adopting such a modification of the original trend analysis but the relationship between continuous and discontinuous change stays the same.

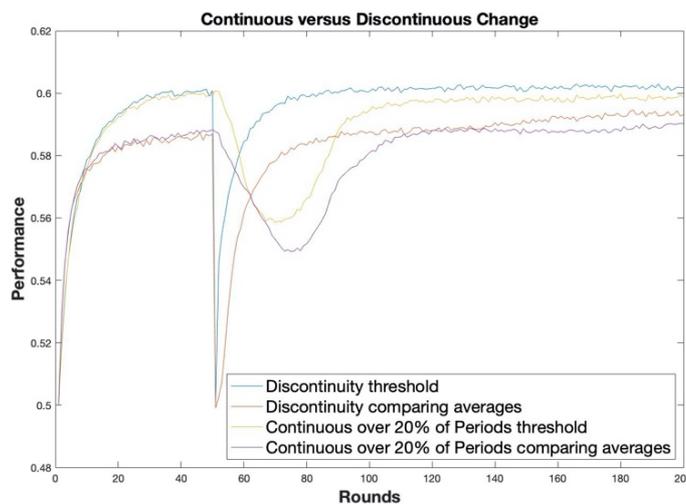


Figure 4-17: Performance comparison with averaging-threshold

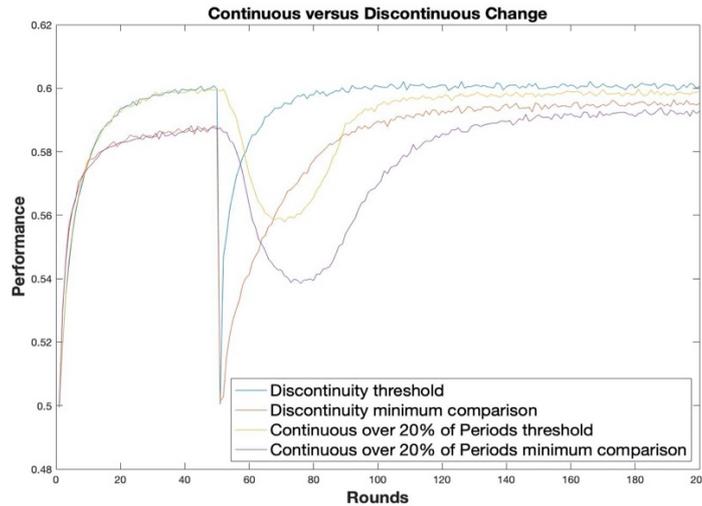


Figure 4-18: Performance comparison with minimum-threshold heuristic

Lastly, we implement a heuristic closest to Redelmeier and Kahneman’s (1996) peak-end rule. Here, the decision maker senses change when the last experience and the worst experience within a 25 period interval is worse than the respective peak and end of the previous 25 periods. In Figure 4-19, we show the performance implications of this heuristic including a threshold rule. Specifically, this means change was sensed when peak and end of the last 25 periods was at least 10 percent worse than the peak and end of the previous 25 periods. In Figure 4-20, we show the performance implications without a threshold. This means that change was sensed when peak and end of the last 25 periods was worse than the peak and end of the previous 25 periods. The results do not yield different performance consequences than previously discussed trend heuristics.

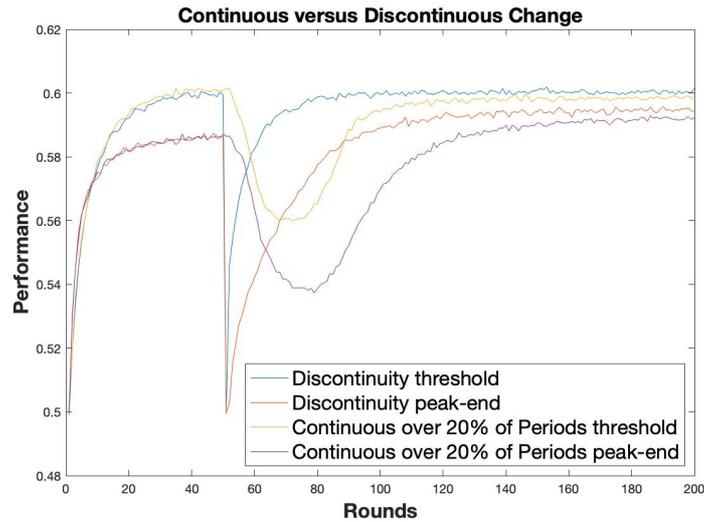


Figure 4-19: Performance comparison with peak-end-threshold heuristic

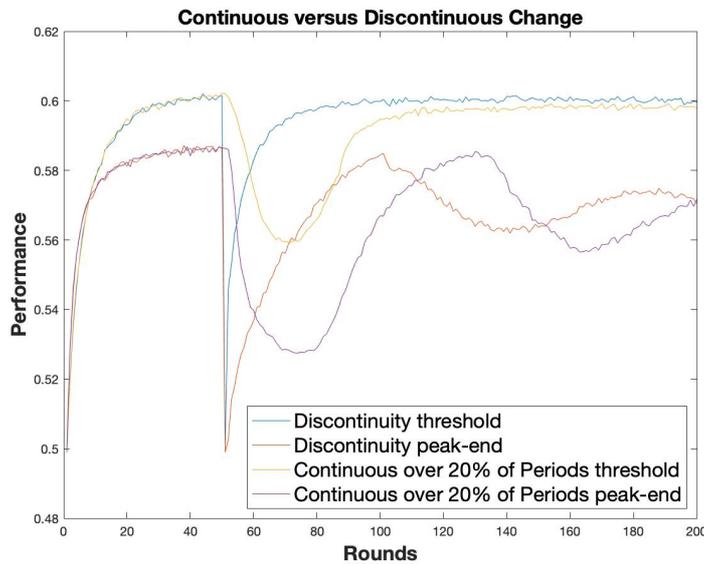


Figure 4-20: Performance comparison with peak-end heuristic

Discussion

Conclusion

In the analysis presented here, we show that organizations may suffer from both very slow and very fast exogenous change in their environment. While very fast change may make it impossible to adapt on time, very slow change may make it difficult to sense that the change is happening in the first place. We have presented evidence that this is true for restaurants in New York City experiencing gentrification.

Contributions to existing literature

Our primary contribution is to the literature on radical environmental change (Anderson and Tushman 1990, Sosa 2011, Tushman and Anderson 1986). While most prior work has focused on discontinuous change, we point to the fact that disruptions sometimes take a long time to manifest. When they do -a scenario we call continuous radical change- organizations may be harmed even more than when change is fast as it goes unnoticed for too long to react to it on time. This adds to literature on environmental change in two ways. First, we answer the call of Tushman and Anderson (1986) to study the patterns of change. We find that in particular the pace of change plays a crucial role in the performance implications of environmental change for incumbent firms. Further, while prior work has focused on innovation as a source of change, we study changing demand in a market of many small players as opposed to an oligopoly situation of few big players that endogenously bring about change. This applies to many other contexts for example in the service industry that have not been covered by prior work.

Further, our findings speak to the debate on organizational learning and adaptation. Much of prior work on learning and adaptation has looked at radical environmental change. For example, Posen and Levinthal (2012) model environmental change and conclude that adaptation sometimes means to exploit known solutions because benefits from learning cannot be accrued as fast as they diminish as the environment is changing again. Relatedly, Benner (2009) emphasizes that routines must be adapted to fit changing environments. Overall, most prior studies take a normative approach to adaptation in changing environments, i.e. they model how optimal adaptation should look like. On the contrary, we emphasize under which condition adaptation is possible. Specifically, we argued that adaptation can only take place when organizations sense the change in their environments which in turn depends on the pace

at which the environment is changing. This adds to recent work on conditions under which sensing is feasible (Chakravarti et al. 2011, Green and Shapira 2018).

By extension, we also address the dynamic capability literature. Teece (2007) argued that dynamic capabilities, i.e. the capability to continuously update an organization's core of assets to meet the current environment's requirements, consist of (a) sensing and (b) seizing opportunities. Subsequent work, however, has focused on the seizing, i.e. the creation or altering of organizational resources (e.g. Helfat et al. 2007). Our findings draw attention to the sensing part of dynamic capabilities and suggest that no changes in an organization's resource endowment can be made without first accurately sensing the environmental change.

Lastly, we also add to the debate on which types of organizations are more likely to survive what kind of environmental change originating in theories of population ecology. For example, Hannan and Freeman (1983) distinguished between specialist and generalist organizations. We, on the contrary, focus on the pace of change as a factor for organizational survival and conclude that survival chances might be following a pattern of an inverted U-shaped such that survival chances are lowest in very fast and very slow changing environments.

Limitations

Despite its contributions, this study has a number of limitations. First, we purposefully study a situation where change is exogenous to the firm. Often times, however, change may be endogenously induced by one of the big players in an industry. In such an industry with few big players, we may observe that change is induced from fewer directions making it easier to sense.

Further, vicarious learning is not central to our model but may be more important in more concentrated industries. Prior behavioral work has uncovered that organizational decision makers belonging to the group of "reciprocators", i.e. those that respond primarily to social

norms, may sense environmental change much faster because they sample from more than just their own organization's experience (Bridoux et al. 2017).

Another limitation is the lack of a direct measure of sensing. Sensing -an inherently cognitive construct- can be approximated by observed behavior such as unlearning or adaptation but is hard to observe directly. We also have a limited amount of data on the adaptation behavior of restaurants. Further studies should explore to study the boiled frog phenomenon in other contexts with a greater availability of adaptation data.

Future research

It is hoped that our finding will serve as a starting point for more empirical research which tests the causal chain of environmental change, sensing and adaptation. In particular, more work is needed on measuring the process of adaptation in the face of very slow and fast changing environments. One way to measure adaptation in exogenously changing environments would be to use patent data and the advent of new subclasses. New subclasses emerge frequently and grow at various rates. Growth could be measured by citations per subclass while a firm patenting in that subclass could represent adaptation. However, neither this nor our approach captures actual sensing. For that, more qualitative work is required.

Another interesting avenue for future research is to study which types of organizations fall prey to the boiled frog effect. Recent work has made headway into this by exploring the role different hierarchies play in sensing (Green and Shapira 2018). Future work could ask specifically what governance type is more or less likely to be the boiled frog.

5. DISCUSSION

Dissertation Summary

In this dissertation, I set out to examine learning under ambiguity in different choice contexts. Specifically, I explored which contexts lead to experiencing ambiguous outcomes and when it is possible for decision makers to learn from such experience. An experience is considered to be ambiguous when it cannot be readily classified into success or failure, i.e. it can be construed as both a success or a failure (Repenning and Sterman 2002, Rerup 2006). While prior research has focused on strategies (such as mindfulness) to cope with ambiguity, I focused on the choice contexts that produce ambiguous outcomes and explore for which of those learning is inhibited or facilitated. I pointed out three choice contexts: (1) a lack of prior experience to determine a sensible threshold between success and failure (i.e. they do not have enough data to formulate historical aspirations), (2) omission choices which do not produce observable outcomes, and (3) confounding exogenous change and outcomes produced by own choices.

The first study dealt with ambiguous experience generated by a lack of historical comparisons. Categorizing outcomes implies that there is a standard of comparison that determines success or failure. In the absence of such a benchmark, experience remains ambiguous. We then explored the role of social comparison (via outcome comparison) as a means to form a standard of comparison. Using an experimental search task, we have found that social comparison may help to improve behavior by informing decision makers of when and when not to search. We have further identified the number of decision alternatives as a moderating factor – the positive effect on search behavior only exists for a small number of decision alternatives while the effect turns negative for many decision alternatives.

The second study explored choices that are in fact omissions and thus, do not produce direct observable feedback. The outcome is in this way ambiguous as it remains unclear if the

omission was positive (i.e. the outcome from action would have been negative) or negative (i.e. the outcome from action would have been positive). We proposed that to learn from errors of omission, decision makers turn to observing competitor outcomes from choices similar to their own. Utilizing a product approval task in an experimental setting, we have found that decision makers do learn from their omission errors by observing competitors but only when they are performing better than the competitor. At the same time, our findings suggest that laggards do not learn from their omission errors and instead focus only on their errors of commission. We have argued that this is due to (1) leaders' broadening of attention and (2) increased data on omission errors as a consequence of leaders becoming more risk averse and thus, more likely to commit an error of omission than an error of commission.

Lastly, the third study focused on ambiguity produced by an unclear cause-effect relationship. Specifically, it dealt with the challenge of telling exogenous environmental change from the outcomes of own choices. We have identified the pace of environmental change as a critical factor in recognizing and responding to change. We have argued that slow change in the environment may be harmful because change cannot be sensed and so, adaptation is never initiated. Very fast change, however, may also be harmful because it does not allow for adaptation at all when there is no time to implement new alternatives. In sum, we have proposed and found an inverted U-shape between pace of environmental change and organizational survival chances and laid out evidence that suggests that this effect is driven by differences in adaptation behavior. Arguments were tested in the context of the New York restaurant industry where we used gentrification to operationalize change.

Contribution to Theory

Generally, this dissertation contributes to developing a theory of learning under ambiguity. Prior literature has treated ambiguity as a given and focused on how decision makers can learn in such contexts. However, for a theory of learning under ambiguity it is important to

first understand which choice contexts produce ambiguous outcomes and how these facilitate or hinder learning. I find that contrary to what prior essays on learning under ambiguity suggest, the ability to learn from ambiguous outcomes varies greatly across choice contexts. In some contexts, it may be lack of prior experience as opposed to lack of immediate feedback that poses challenges for learning. Here, it is difficult to form a standard of comparison against which to evaluate success or failure. It may be possible to construct such standard of comparison by looking at comparable others. In such cases, I find that observing peers' outcomes may trigger over-exploration depending on the size of one's choice set. In other choice contexts, there may be ample prior experience but a lack of immediate feedback that poses challenges for learning. For the context of learning from omissions, the key challenge is to conceive the consequences of a hypothetical action. While this may be possible by observing competitors pursue similar actions, I find that this introduces new behavioral effects on learning and that the ability to learn from competitors about omission errors may be dependent on relative performance to the competitor. Lastly, there are conceivably choice contexts where choices produce feedback and there is enough prior experience to form a standard of comparison, but it is not obvious whether the choice or an exogenous change in the environment caused the outcome. I find that the ability to clearly categorize feedback blurred by exogenous change critically depends on the pace of the environmental change. In conclusion, each of these choice contexts poses a distinct challenge for learning thus, has to be understood separately.

In exploring the challenges different choice contexts that produce ambiguity pose for learning, I built on March and Olsen's (1975) groundwork of a theory of organizational learning under ambiguity. March and Olsen (1975) argued that learning can be described in the following cycle: (1) Individuals have models about the environment, (2) they engage in a choice situation, (3) this leads to organizational action, which (3) produces responses from

the environment which in turn (1) alters individuals' models of the environment, and so forth. They went on to suggest that this learning cycle may be incomplete at several places. Individuals' beliefs may not lead to choices because individuals are constrained in their roles. Their choices may also not lead to organizational action for example because coalitions to pursue one organizational course of actions cannot be formed. Further, the relation between organizational action and environmental response may be severed, which may be described as superstitious learning. Lastly, learning under ambiguity occurs when individuals cannot update their models of the environment because they are unable to classify environmental responses in a meaningful way. In this dissertation, I study cases of "breaking the cycle in different ways under different circumstances", thus contributing to the development of a "theory of a full cycle of organizational choice" (March and Olsen 1975, p. 169).

This dissertation also contributes to social aspiration theory (Cyert and March 1963, Greve 1998, Greve 2003b, March 1994). In study 1 and 2, I examine two instances where decision makers construct social aspiration levels in their learning process. In study 1, I show how decision makers can utilize social aspirations not only to induce change (Audia et al. 2000, Lant 1992, Miller and Chen 1994), exploration (Baum and Dahlin 2007), and risk taking (Bromiley 1991), but also to inform them of when to start and when to stop searching. In study 2, I demonstrate that social aspirations can also be inadvertently constructed in an attempt to learn from competitor experiences and thus, may alter decision makers' learning behavior. Overall, this suggests a broader look at social aspiration that takes into account different ways in which aspirations can be utilized and analyzes whether they are constructed deliberately or as a byproduct of learning.

Limitations and Future Research

While this dissertation is meant to be a step in the direction of a theory of learning under ambiguity, it falls short of providing a complete theory. In future research, it would be

necessary to identify other choice contexts that produce ambiguous outcomes along with the specific challenges each context imposes on learning behavior and then to integrate those into one typology of learning under ambiguity.

Further, I tested my arguments in a narrow empirical setting. The experiments in study 1 and 2 allowed for an exact identification of the mechanism of peer information on learning. They did not, however, allow for identifying all possible ways in which ambiguous experience from omissions or lack of experience can affect learning behavior. Future research may be needed to examine additional mechanisms. Additionally, the empirical setting of the restaurant industry in study 3 does not allow for showing the causal chain of pace of environmental change on survival via sensing and adaptation. I have limited data on restaurants' adaptation behavior and no data on their sensing activities. Further studies are required to better measure sensing and adaptation directly.

Lastly, in study 1 and 2 I focus on the individual decision maker while in study 3 I study the decisions of organizations as a whole. However, more research is needed to determine how individual level learning behavior translates into organizational learning. This is also an important part of March and Olsen's (1975) theory of organizational learning: In their cycle of learning, individual action may not always lead to organizational action, a phenomenon they call *audience experiential learning*. Here, individuals learn but the organization does not adapt, for example because of organizational politics. This again, may introduce ambiguity as now there are multiple standards of comparison from different organizational coalitions against which decision makers evaluate success or failure. The combination of audience experiential learning and learning under ambiguity therefore deserves further attention.

6. REFERENCES

- Abrahamson E (1991) Managerial fads and fashions: The diffusion and rejection of innovations. *Acad. Management Rev.* 16(3): 586-612.
- Adams JS (1963) Towards an understanding of inequity. *J. Abnormal Soc. Psych.* 67(5): 422-436.
- Aggarwal VA, Siggelkow N, Singh H (2011) Governing collaborative activity: interdependence and the impact of coordination and exploration. *Strategic Management J.* 32(7):705-730.
- Anderson P, Tushman ML (1990) Technological Discontinuities and Dominant Designs: A Cyclical Model of Technological Change. *Admin. Sci. Quart.* 35(4): 604-633.
- Argote L (1999) *Organizational Learning: Creating, Retaining and Transferring*. Kluwer, Boston, MA.
- Argote L, Ingram P (2000) Knowledge transfer: A basis for competitive advantage in firms. *Organ. Behav. Human Decision Processes* 82(1): 150-169.
- Argyris C, Schön D (1978) *Organizational learning: A theory of action perspective*, Reading, MA: Addison-Wesley.
- Arkes HR, Dawes RM, Christensen C (1986) Factors influencing the use of a decision rule in a probabilistic task. *Organ. Behav. Human Decision Processes* 37(1): 93-110.
- Arrow KJ (1963) Uncertainty and the welfare economics of medical care. *American Econom. Rev.* 53(5): 941-973.
- Asgari N, Singh K, Mitchell W (2017) Alliance Portfolio Reconfiguration Following a Technological Discontinuity. *Strategic Management J.* 38(5): 1062-1081.
- Ashton RH (1990) Pressure and performance in accounting decision settings: Paradoxical effects of incentives, feedback, and justification. *J. Accounting Research* 28: 148-180.
- Asmussen CG, Larsen MM, Pedersen T (2016) Organizational Adaptation in Offshoring: The Relative Performance of Home- and Host-Based Learning Strategies. *Organ. Sci.* 27(4): 911-928.
- Audia PG, Brion S, Greve HR (2015) Self-assessment, self-enhancement, and the choice of comparison organizations for evaluating organizational performance. Gavetti G, Ocasio W, eds. *Cognition and Strategy*, Vol. 32 (Emerald Group Publishing Ltd., Bingley, UK), 89-118.
- Audia PG, Greve HR (2006) Less likely to fail: Low performance, firm size, and factory expansion in the shipbuilding industry. *Management Sci.* 52(1): 83-94.
- Audia PG, Locke EA, Smith KG (2000) The paradox of success: An archival and a laboratory study of strategic persistence following radical environmental change. *Acad. Management J.* 43(5): 837-853.
- Baron J, Ritov I (2004) Omission bias, individual differences, and normality. *Organ. Behav. Human Decision Processes* 94(2): 74-85.
- Baum JAC, Dahlin KB (2007) Aspiration Performance and Railroads' Patterns of Learning from Train Wrecks and Crashes. *Organ. Sci.* 18(3): 368-385.
- Baum JA, Rowley TJ, Shipilov AV, Chuang, YT (2005) Dancing with strangers: Aspiration performance and the search for underwriting syndicate partners. *Admin. Sci. Quart.* 50(4): 536-575.
- Bateson G (1979). *Mind and nature: A necessary unity* (Vol. 255). New York: Bantam Books.
- Baumann O, Eggers JP, Stieglitz N (2018) Colleagues and Competitors: How Internal Social Comparisons Shape Organizational Search and Adaptation. *Admin. Sci. Quart.* 63(1): 1-35.

- Berinsky AJ, Huber GA, Lenz GZ (2012) Evaluating online labor markets for experimental research: Amazon.com's Mechanical Turk. *Political Anal.* 20(3): 351-368.
- Bendor JB, Kumar S, Siegel DA (2009) Satisficing: A 'pretty good' heuristic. *BE J. Theoretical Econom.* 9(1): 1-36.
- Bennett VM, Snyder J (2017) The empirics of learning from failure. *Strategy Sci.* 2(1): 1-12.
- Benner MJ (2009) Dynamic or static capabilities? Process management practices and response to technological change. *J. Product Innovation Management* 26(5): 473-486.
- Benner MJ (2010) Securities analysts and incumbent response to radical technological change: Evidence from digital photography and internet telephony. *Organ. Sci.* 21(1):42-62.
- Beersma B, Hollenbeck JR, Humphrey SE, Moon H, Conlon DE, Ilgen DR (2003) Cooperation, competition, and team performance: Toward a contingency approach. *Acad. Management J.* 46(5): 572-590.
- Bennett VM, Snyder J (2017) The empirics of learning from failure. *Strategy Sci.* 2(1): 1-12.
- Bettis RA, Helfat CE, Shaver JM (2009) The necessity, logic and form of replication. *Strategic Management J.* 37(11): 2193-2203.
- Billinger S, Stieglitz N, Schumacher TR (2014) Search on Rugged Landscapes: An Experimental Study. *Organ. Sci.* 25(1): 93-108.
- Blanes i Vidal J, Nossol M (2011) Tournaments without prizes: Evidence from personnel records. *Management Sci.* 57(10): 1721-1736.
- Botti S, Iyengar SS (2006) The dark side of choice: When choice impairs social welfare. *J. Public Policy & Marketing* 25(1): 24-38.
- Bridoux F, Coeurderoy R, Durand R (2017) Heterogeneous social motives and interactions: The three predictable paths of capability development. *Strategic Management J.* 38(9): 1755-1773.
- Bromiley P (1991) Testing a Causal Model of Corporate Risk Taking and Performance. *Acad. Management J.* 34(1): 37-59.
- Bromiley P, Harris JD (2014) A comparison of alternative measures of organizational aspirations. *Strategic Management J.* 35(3): 338-357.
- Bromiley P, Miller KD, Rau D (2001) Risk in strategic management research. *The Blackwell Handbook of Strategic Management* 259-288.
- Brown DJ, Ferris DL, Heller D, Keeping LM (2007) Antecedents and consequences of the frequency of upward and downward social comparisons at work. *Organ. Behav. Human Decision Processes* 102(1): 59-75.
- Brynjolfsson E, Hu Y, Smith MD (2003) Consumer surplus in the digital economy: Estimating the value of increased product variety at online booksellers. *Management Sci.* 49(11): 1580-1596.
- Burbano VC (2016) Social responsibility messages and worker wage requirements: Field experimental evidence from online labor marketplaces. *Organ. Sci.* 27(4): 1010-1028.
- Busemeyer JR, Swenson KN, Lazarte A (1986) An adaptive approach to resource allocation. *Organ. Behav. Human Decision Processes* 38(3): 318-341.
- Camerer CF, Hogarth RM (1999) The effects of financial incentives in experiments: A review and capital-labor-production framework. *J. Risk Uncertainty* 19(1-3): 7-42.
- Cirillo B, Brusoni S, Valentini G (2014) The Rejuvenation of Inventors Through Corporate Spinouts. *Organ. Sci.* 25(6): 1764-1784.

- Chakravarti A, Fang C, Shapira Z (2011) Detecting and Reacting to Change: The Effect of Exposure to Narrow Categorizations. *J. Experimental Psychology: Learning Memory and Cognition* 37(6): 1563–1570.
- Cappelli P, Hamori M (2013) Understanding executive job search. *Organ. Sci.* 25(5): 1511-1529.
- Chatterjee S, Wiseman RM, Fiegenbaum A, Devers CE (2003) Integrating behavioural and economic concepts of risk into strategic management: The twain shall meet. *Long Range Planning* 36(1): 61-79.
- Chen WR, Miller KD (2007) Situational and institutional determinants of firms' R&D search intensity. *Strategic Management J.* 28(4): 369-381.
- Clement J, Puranam P (2017) Searching for structure: Formal organization design as a guide to network evolution. *Management Sci.* 64(8): 3879–3895.
- Colella A, Paetzold RL, Zardkoohi A, Wesson MJ (2007) Exposing pay secrecy. *Acad. Management Rev.* 32(1): 55-71.
- Cook TD, Campbell DT (1979) *Quasi-experimentation: Design and analysis for field settings* (Vol. 3). Rand McNally: Chicago, IL.
- Cook TD, Campbell DT, Shadish W (2002) *Experimental and quasi-experimental designs for generalized causal inference*. Houghton Mifflin. Boston, MA.
- Croson R, Anand J, Agarwal R (2007) Using experiments in corporate strategy research. *European Management Rev.* 4(3): 173-181.
- Cyert RM, March JG (1963) *A behavioral theory of the firm* (2nd ed.). Englewood Cliffs: New York, NY.
- Csaszar FA (2012) Organizational structure as a determinant of performance: Evidence from mutual funds. *Strategic Management J.* 33(6): 611-632.
- Csaszar FA (2013) An efficient frontier in organization design: Organizational structure as a determinant of exploration and exploitation. *Organ. Sci.* 24(4): 1083-1101.
- Csaszar FA, Siggelkow N (2010) How much to copy? Determinants of effective imitation breadth. *Organ. Sci.* 21(3): 661-676.
- Davis JP, Eisenhardt KM, Bingham CB (2009) Optimal structure, market dynamism, and the strategy of simple rules. *Admin. Sci. Quart.* 54(3): 413-452.
- Davis JP, Eisenhardt KM (2011) Rotating leadership and collaborative innovation: Recombination processes in symbiotic relationships. *Admin. Sci. Quart.* 56(2): 159-201.
- Deci EL, Betley G, Kahle J, Abrams L, Porac J (1981) When trying to win: Competition and intrinsic motivation. *Personality Social Psych. Bulletin* 7(1): 79-83.
- Denrell J, March JG (2001) Adaptation as information restriction: The hot stove effect. *Organ. Sci.* 12(5): 523-538.
- Denrell J (2003) Vicarious learning, undersampling of failure, and the myths of management. *Organ. Sci.* 14(3): 227-243.
- Desmond M, Gershenson C (2017) Who gets evicted? Assessing individual, neighborhood, and network factors. *Soc. Sci. Res.* 62: 362-377.
- Deutsch M (1949) A theory of co-operation and competition. *Human Relations* 2(2): 129-152.
- DiMaggio PJ, Powell WW (1983) The iron cage revisited: Institutional isomorphism and collective rationality in organizational fields. *Amer. Sociol. Rev.* 48(April): 147-160.
- Di Stefano G, Gutierrez C (2018) Under a magnifying glass: On the use of experiments in strategy research. *Strategic Organ.* 17(4): 1-11.
- Dunn J, Ruedy NE, Schweitzer ME (2012) It hurts both ways: How social comparisons harm affective and cognitive trust. *Organ. Behav. Human Decision Processes* 117(1): 2-14.
- Dye KC, Eggers JP, Shapira Z (2014) Trade-offs in a tempest: Stakeholder influence on hurricane evacuation decisions. *Organ. Sci.* 25(4): 1009-1025.

- Eggers JP (2012) Falling flat: Failed technologies and investment under uncertainty. *Admin. Sci. Quart.* 57(1): 47-80.
- Eggers JP, Suh JH (2019) Experience and Behavior: How Negative Feedback in New Versus Experienced Domains Affects Firm Action and Subsequent Performance. *Acad. Management J.* 62(2): 309-334.
- Fang C (2012) Organizational learning as credit assignment: A model and two experiments. *Organ. Sci.* 23(6): 1717-1732.
- Festinger L (1954) A theory of social comparison processes. *Human Relations* 7(2): 117-140.
- Fleming L, Sorenson O (2004) Science as a map in technological search. *Strategic Management J.* 25(8-9): 909-928.
- Flynn FJ, Amanatullah ET (2012) Psyched up or psyched out? The influence of coactor status on individual performance. *Organ. Sci.* 23(2): 402-415.
- Gaba V, Joseph J (2013) Corporate structure and performance feedback: Aspirations and adaptation in M-form firms. *Organ. Sci.* 24(4): 1102-1119.
- Gartenberg C, Wulf J (2017) Pay harmony? Social comparison and performance compensation in multibusiness firms. *Organ. Sci.* 28(1): 39-55.
- Garud R, Nayyar PR, Shapira ZB (1997) *Technological innovation: Oversights and foresights*. Cambridge University Press.
- Gavetti G, and Levinthal D (2000) Looking Forward and Looking Backward: Cognitive and Experiential Search. *Admin. Sci. Quart.* 45(1): 113-137.
- Gavetti G (2012) PERSPECTIVE—Toward a behavioral theory of strategy. *Organ. Sci.* 23(1): 267-285.
- Gino F, Pierce L (2009) Dishonesty in the name of equity. *Psych. Sci.* 20(9): 1153-1160.
- Gino F, Pierce L (2010) Lying to level the playing field: Why people may dishonestly help or hurt others to create equity. *J. Bus. Ethics* 95(1): 89-103.
- Glaeser EL, Kim H, Luca M (2018) Nowcasting gentrification: using yelp data to quantify neighborhood change. *AEA Papers and Proceedings* (Vol. 108, pp. 77-82).
- Glynn MA, Lant TK, Mezias SJ (1991) Incrementalism, learning, and ambiguity: An experimental study of aspiration level adaptation. *1991 Acad. Management Best Paper Proceedings*: 384-388.
- Green E, Shapira Z (2018) Hierarchical Sensing and Strategic Decision-Making. In *Behavioral Strategy in Perspective* (pp. 123-138). Emerald Publishing Limited.
- Green DM, Swets JA (1966) *Signal detection theory and psychophysics* (Vol. 1). Wiley: New York, NY.
- Greve HR (1998) Performance, Aspirations, and Risky Organizational Change. *Admin. Sci. Quart.* 43(1): 58-86.
- Greve HR (2002) Sticky aspirations: Organizational time perspective and competitiveness. *Organ. Sci.* 13(1): 1-17.
- Greve HR (2003a) A Behavioral Theory of R&D Expenditures and Innovations: Evidence from Shipbuilding. *Acad. Management J.* 46(6): 685-702.
- Greve HR (2003b) *Organizational Learning from Performance Feedback*. Cambridge University Press, Cambridge: UK.
- Greve HR, Gaba V (2017) *Performance feedback in organizations and groups: Common themes*. Oxford University Press, Oxford, UK.
- Halford GS, Wilson WH, Guo J, Gayler RW, Wiles J, Stewart JEM (1994) Connectionist implications for processing capacity limitations in analogies. *Advances in connectionist and neural computation theory*, 2: 363-415.
- Hamel G, Prahalad CK (1993) Strategy as stretch and leverage. *Harvard Bus. Rev.* 71(2): 75-84.

- Freeman J, Hannan MT (1983) Niche width and the dynamics of organizational populations. *Amer. J Sociol.* 88(6): 1116-1145.
- Halford GS, Wilson WH, Guo J, Gayler RW, Wiles J, Stewart JEM (1994) Connectionist implications for processing capacity limitations in analogies. *Advances in Connectionist and Neural Computation Theory 2*: 363-415.
- Hannan MT, Freeman J (1984) Structural inertia and organizational change. *Amer. Sociological Rev.* 49(2): 149-164.
- Harmon DJ, Kim PH, Mayer KJ (2015) Breaking the letter vs. spirit of the law: How the interpretation of contract violations affects trust and the management of relationships. *Strategic Management J.* 36(4): 497-517.
- Haunschild PR, Beckman CM (1998) When do interlocks matter?: Alternate sources of information and interlock influence. *Admin. Sci. Quart.* 44(4): 815-844.
- Haunschild PR, Miner AS (1997) Modes of interorganizational imitation: The effects of outcome salience and uncertainty. *Admin. Sci. Quart.* 42(3): 472-500.
- Haunschild PR, Polidoro F Jr, Chandler D (2015) Organizational oscillation between learning and forgetting: The dual role of serious errors. *Organ. Sci.* 26(6): 1682-1701.
- Haunschild PR, Sullivan BN (2002) Learning from complexity: Effects of prior accidents and incidents on airlines' learning. *Admin. Sci. Quart.* 47(4): 609-643.
- Hauser JR, Wernerfelt B (1990) An evaluation cost model of consideration sets. *J. Consumer Res.* 16(4): 393-408.
- Haveman H (1992) Between a rock and a hard place: Organizational change and performance under conditions of fundamental environmental transformation. *Admin. Sci. Quart.* 37(1): 48-75.
- Henderson R (1993) Underinvestment and incompetence as responses to radical innovation: Evidence from the photolithographic alignment equipment industry. *RAND J. Econom.* 24: 248-270.
- 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 J.* 20(10): 977-985.
- Hogarth RM, Gibbs BJ, McKenzie CRM, Marquis, MA (1991) Learning from feedback: exactingness and incentives. *J Experiment. Psych.: Learn., Memory, Cognition* 17(4): 734-752.
- Holan PM de, Phillips N (2004) Remembrance of Things Past? The Dynamics of Organizational Forgetting. *Management Sci.* 50(11): 1603-1613.
- Holt CA. (1995) Industrial organization: A survey of laboratory research. *The Handbook of Experimental Economics* 349: 402-03.
- Horton JJ, Rand DG, Zeckhauser RJ (2011) The online laboratory: Conducting experiments in a real labor market. *Experiment. Econom.* 14(3): 399-425.
- Iyengar SS, Lepper MR (2000) When choice is demotivating: Can one desire too much of a good thing? *J. Personality Soc. Psych.* 79(6): 995-1006.
- Jemison B (1981) The importance of an integrative approach to strategic management research. *Acad. Management Rev.* 6(4): 601-608.
- Joseph J, Gaba V (2015) The fog of feedback: Ambiguity and firm responses to multiple aspiration levels. *Strategic Management J.* 36(13): 1960-1978.
- Kacperczyk A, Beckman CM, Moliterno TP (2015) Disentangling risk and change: Internal and external social comparison in the mutual fund industry. *Admin. Sci. Quart.* 60(2): 228-262.
- Kahneman D, Tversky A (1979) Prospect Theory: An Analysis of Decision Under Risk. *Econometrica* 47: 263-291.

- Kalnins A, Swaminathan A, Mitchell W (2006) Turnover events, vicarious information, and the reduced likelihood of outlet-level exit among small multiunit organizations. *Organ. Sci.* 17(1): 118-131.
- Kaplan S, Murray F, Henderson R (2003) Discontinuities and senior management: assessing the role of recognition in pharmaceutical firm response to biotechnology. *Indust. Corporate Change* 12(2): 203-233.
- Kc D, Staats BR, Gino F (2013) Learning from my success and from others' failure: Evidence from minimally invasive cardiac surgery. *Management Sci.* 59(11): 2435-2449.
- Keister LA (2002) Adapting to Radical Change: Strategy and Environment in Piece-Rate Adoption During China's Transition. *Organ. Sci.* 13(5): 459-474.
- Kelley HH (1973) The processes of causal attribution. *Amer. Psych.* 28(2): 107-128.
- Kerr S, Landauer S (2004) Using stretch goals to promote organizational effectiveness and personal growth: General Electric and Goldman Sachs. *Acad. Management Exec.* 18(4): 134-138.
- Kiesler S, Sproull L (1982) Managerial Response to Changing Environments: Perspectives on Problem Sensing from Social Cognition. *Admin. Sci. Quart.* 27(4): 548-570.
- Kilduff GJ, Elfenbein HA, Staw BM (2010) The psychology of rivalry: A relationally dependent analysis of competition. *Acad. Management J.* 53(5): 943-969.
- Kim L (1998) Crisis Construction and Organizational Learning: Capability Building in Catching-up at Hyundai Motor. *Organ. Sci.* 9(4): 506-521.
- Kim H, Hoskisson RE, Lee SH (2015) Why strategic factor markets matter: "New" multinationals' geographic diversification and firm profitability. *Strategic Management J.* 36(4): 518-536.
- Khanna R, Guler I, Nerkar A (2016) Fail often, fail big, and fail fast? Learning from small failures and R&D performance in the pharmaceutical industry. *Acad. of Management J.* 59(2): 436-459.
- Klepper S (2002) Firm Survival and the Evolution of Oligopoly. *RAND J. Econom.* 33(1): 37-61.
- Klingebliel R (2018) Risk-type preference shifts in response to performance feedback. *Strategic Organ.* 16(2): 141-166.
- Knudsen T, Srikanth K (2014) Coordinated exploration: Organizing joint search by multiple specialists to overcome mutual confusion and joint myopia. *Admin. Sci. Quart.* 59(3): 409-441.
- Lamin A (2013) Business groups as information resource. *Strategic Management J.* 56(1): 1486-1509.
- Lampel J, Shamsie J, Shapira Z (2009) Experiencing the improbable: Rare events and organizational learning. *Organ. Sci.* 20(5): 835-845.
- Lant TK (1992) Aspiration level adaptation: An empirical exploration. *Management Sci.* 38(5): 623-644.
- Lant TK, Mezias SJ (1992) An organizational learning model of convergence and reorientation. *Organ. Sci.* 3(1): 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 J.* 13(8): 585-608.
- Latham GP (2004) The motivational benefits of goal-setting. *Acad. Management Perspectives* 18(4): 126-129.
- Laureiro-Martínez D, Brusoni S, Canessa N, Zollo M (2015) Understanding the exploration-exploitation dilemma: An fMRI study of attention control and decision-making performance. *Strategic Management J.* 36(3): 319-338.

- Lavie D, Stettner, U, Tushman ML (2010) Exploration and Exploitation within and across organizations. *Acad. Management Annals* 4(1): 109-155.
- Lees L (2003) Super-gentrification: The case of Brooklyn heights, New York city. *Urban Studies* 40(12): 2487-2509.
- Lefebvre M, Vieider FM, Villeval MC (2010) Incentive effects on risk attitude in small probability prospects. *Econom. Letters* 109(2): 115-120.
- Levine SS, Bernard M, Nagel R (2017) Strategic intelligence: The cognitive capability to anticipate competitor behavior. *Strategic Management J.* 38(12): 2390-2423.
- Levinthal DA (1997) Adaptation on rugged landscapes. *Management Sci.* 43(7): 934-950.
- Levinthal DA (2011) A behavioral approach to strategy—what's the alternative? *Strategic Management J.* 32(13): 1517-1523.
- Levinthal D, March JG (1981) A model of adaptive organizational search. *J. Econom. Behav. Organ.* 2(4): 307–333.
- Levinthal DA, March JG (1993) The myopia of learning. *Strategic Management J.* 14(S2): 95-112.
- Levinthal D, Posen HE (2007) Myopia of Selection : Does Organizational Adaptation Limit the Efficacy of Population Selection? *Admin. Sci. Quart.* 52(4): 586–620.
- Levinthal D, Rerup C (2006) Crossing an apparent chasm: Bridging mindful and less-mindful perspectives on organizational learning. *Organ. Sci.* 17(4): 502-513.
- Levitt B, March JG (1988) Organizational learning. *Annual Rev. Sociol.* 14(1): 319-338.
- Loch CH, Kavadias S (2002) Dynamic portfolio selection of NPD programs using marginal returns. *Management Sci.* 48(10): 1227-1241.
- Long X, Nasiry J (2018) Wage transparency and social comparison in sales force compensation. Working paper, University of Wisconsin.
- Lungeanu R, Stern I, Zajac EJ (2016) When do firms change technology-sourcing vehicles? The role of poor innovative performance and financial slack. *Strategic Management J.* 37(5): 855-869.
- MacLeod WB, Pingle M (2005) Aspiration uncertainty: Its impact on decision performance and process. *J. Econom. Behav. Organ.* 56(4): 617-629.
- Madsen PM, Desai V (2010) Failing to learn? The effects of failure and success on organizational learning in the global orbital launch vehicle industry. *Acad. Management J.* 53(3): 451-476.
- Malhotra NK (1982) Information load and consumer decision making. *J. Consumer Res.* 8(4): 419-430.
- March JG (1991) Exploration and exploitation in organizational learning. *Organ. Sci.* 2(1): 71–87.
- March JG (1994) *Primer on decision making: How decisions happen*. Simon and Schuster: New York.
- March JG (1996) Continuity and change in theories of organizational action. *Admin. Sci. Quart.* 41(2): 278-287.
- March JG, Olsen JP (1975) Uncertainty of past—Organizational learning under ambiguity. *Eur. J. Political Res.* 3(2):147–171.
- March JG, Olsen JP (1976) Organizational choice under ambiguity. *Ambiguity and choice in organizations.* 2: 10-23.
- March JG, Shapira Z (1987) Managerial perspectives on risk and risk taking. *Management Sci.* 33(11): 1404-1418.
- March JG, Shapira Z (1992) Variable risk preferences and the focus of attention. *Psych. Rev.* 99(1): 172-183.
- March JG, Simon HA (1958) *Organizations*. Wiley: New York, NY.

- March JG, Sproull LS, Tamuz M (1991) Learning from samples of one or fewer. *Organ. Sci.* 2(1): 1-13.
- Maslach D (2016) Change and persistence with failed technological innovation. *Strategic Management J.* 37(4): 714-723.
- Maslach D, Branzei O, Rerup C, Zbaracki MJ (2018) Noise as signal in learning from rare events. *Organ. Sci.* 29(2): 225-246.
- Maula MVJ, Keil T, Zahra SA (2013) Top Management's Attention to Discontinuous Technological Change: Corporate Venture Capital as an Alert Mechanism. *Organ. Sci.* 24(3): 926-947.
- Mezias SJ, Chen YR, Murphy PR (2002) Aspiration-level adaptation in an American financial services organization: A field study. *Management Sci.* 48(10): 1285-1300.
- Miles RE, Snow CC (1984) Fit, failure and the hall of fame. *California Management Rev.* 26(3): 10-28.
- Miller D, Chen MJ (1994) Sources and consequences of competitive inertia: a study of the U. S. airline industry. *Admin. Sci. Quart.* 39(1): 1-23.
- Moliterno TP, Beck N, Beckman CM, Meyer M (2014) Knowing Your Place: Social Performance Feedback in Good Times and Bad Times. *Organ. Sci.* 25(6): 1684-1702.
- Mudambi R, Swift T (2014) Knowing when to leap: Transitioning between exploitative and explorative R&D. *Strategic Management J.* 35(1): 126-145.
- Myerson RB (1979) Incentive compatibility and the bargaining problem. *Econometrica* 47(1): 61-73.
- Nelson RR, Winter SG (1982) *An Evolutionary Theory of Economic Change*. Harvard University Press: Cambridge, MA.
- Nickerson JA, Silverman BS (2003) Why firms want to organize efficiently and what keeps them from doing so: Inappropriate governance, performance, and adaptation in deregulated industry. *Admin. Sci. Quart.* 48(3): 433-465.
- Nickerson JA, Zenger TR (2004) A knowledge-based theory of the firm—The problem-solving perspective. *Organ. Sci.* 15(6): 617-632.
- Nickerson JA, Zenger TR (2008) Envy, comparison costs, and the economic theory of the firm. *Strategic Management J.* 29(13): 1429-1449.
- Obloj T, Zenger T (2017) Organization design, proximity, and productivity responses to upward social comparison. *Organ. Sci.* 28(1): 1-18.
- O'Reilly III CA (1980) Individuals and information overload in organizations: is more necessarily better? *Acad. Management J.* 23(4): 684-696.
- Ortmann A, Hertwig R (2002) The costs of deception: Evidence from psychology. *Experiment. Econom.* 5(2): 111-131.
- Paolacci G, Chandler J, Ipeirotis, PG (2010) Running experiments on amazon mechanical Turk. *Judgment Decision Making* 5(5): 411-419.
- Payne JW (1976) Task complexity and contingent processing in decision making: An information search and protocol analysis. *Organ. Behav. Performance* 16(2): 366-387.
- Peterson WWTG, Birdsall T, Fox W (1954) The theory of signal detectability. *Transactions of the IRE professional group on information theory* 4(4): 171-212.
- Pfeffer J, Langton N (1993) The effect of wage dispersion on satisfaction, productivity, and working collaboratively: Evidence from college and university faculty. *Admin. Sci. Quart.* 38(3): 382-407.
- Pociask S (2016) Post Office Performance Is "Out Of Service." Retrieved from <https://www.forbes.com/sites/stevepociask/2016/04/08/post-office-performance-is-out-of-service%0Da0d928636f2>
- Porac JF, Thomas H, Baden-Fuller C (1989) Competitive groups as cognitive communities: The case of Scottish knitwear manufacturers. *J. Management Stud.* 26(4): 397-416.

- Porter ME (1996) What is strategy. *Harvard Bus. Rev.* November-December.
- Posen HE, Levinthal DA (2012) Chasing a Moving Target: Exploitation and Exploration in Dynamic Environments. *Management Sci.* 58(3): 587–601.
- Post T, Van den Assem MJ, Baltussen G, Thaler RH (2008) Deal or no deal? Decision making under risk in a large-payoff game show. *Amer. Econom. Rev.* 98(1): 38-71.
- Puranam P, Swamy M. 2016. How initial representations shape coupled learning processes. *Organ. Sci.* 27(2): 323-335.
- Rapoport A, Tversky A (1970) Choice behavior in an optional stopping task. *Organ. Behav. Performance* 5(2): 105-120.
- Ref O, Shapira Z (2017) Entering new markets: The effect of performance feedback near aspiration and well below and above it. *Strategic Management J.* 38(7): 1416-1434.
- Repenning NP, Sterman JD (2002) Capability traps and self-confirming attribution errors in the dynamics of process improvement. *Admin. Sci. Quart.* 47(2): 265-295.
- Rerup C (2006) Success, failure and the gray zone: how organizations learn or don't from ambiguous experience. Paper presented at the Academy of Management Conference, Atlanta, GA.
- Reutskaja E, Hogarth RM (2009) Satisfaction in choice as a function of the number of alternatives: When “goods satiate”. *Psych. Marketing* 26(3): 197-203.
- Rieskamp J, Busemeyer JR, Mellers BA (2006) Extending the bounds of rationality: Evidence and theories of preferential choice. *J. Econom. Literature* 44(3): 631-661.
- Rivkin JW (2000) Imitation of complex strategies. *Management Sci.* 46(6): 824-844.
- Rhee M, Kim T (2014) Great vessels take a long time to mature: Early success traps and competences in exploitation and exploration. *Organ. Sci.* 26(1): 180–197.
- Sah RK, Stiglitz JE (1986) The architecture of economic systems: Hierarchies and polyarchies. *Amer. Econom. Rev.* 76(4): 716-727.
- Schwartz B (2004) The paradox of choice: Why more is less. New York: Ecco.
- Schwenk CR (1984) Cognitive simplification processes in strategic decision-making. *Strategic Management J.* 5(2): 111-128.
- Shapira Z, Venezia I (1981) Optional stopping on nonstationary series. *Organ. Behav. Performance* 27(1): 32-49.
- Shin J, Ariely D (2004) Keeping doors open: The effect of unavailability on incentives to keep options viable. *Management Sci.* 50(5): 575-586.
- Shugan SM (1980) The cost of thinking. *J. Consumer Res.* 7(2): 99-111.
- Siggelkow N (2001) Change in the presence of fit: The rise, the fall, and the renaissance of Liz Claiborne. *Acad. Management J.* 44(4): 838-857.
- Siggelkow N, Levinthal DA (2005) Escaping real (non-benign) competency traps: Linking the dynamics of organizational structure to the dynamics of search. *Strategic Organ.* 3(1): 85-115.
- Silcoff S, McNish J, Ladurantaye S (2013) Inside the fall of BlackBerry: How the smartphone inventor failed to adapt. *The Globe and Mail* (September 27), <https://www.theglobeandmail.com/report-on-business/the-inside-story-of-why-blackberry-is-failing/article14563602/>.
- Simon HA (1955) A behavioral model of rational choice. *Quart. J. Econom.* 69(1): 99-118.
- Simon HA (1956) Rational choice and the structure of the environment. *Psych. Rev.* 63(2): 129-138.
- Simon HA (1962) New Developments in the Theory of the Firm. *Amer. Econom. Rev.* 52(2): 1-15.
- Sitkin SB, See KE, Miller CC, Lawless MW, Carton AM (2011) The paradox of stretch goals: Organizations in pursuit of the seemingly impossible. *Acad. Management Rev.* 36(3): 544-566.

- Smith VL (1976) Experimental economics: Induced value theory. *Amer. Econom. Rev.* 66(2): 274-279.
- Smith VL, Walker JM (1993) Monetary rewards and decision cost in experimental economics. *Econom. Inquiry* 31(2): 245-261.
- Sosa ML (2011) From Old Competence Destruction to New Competence Access: Evidence from the Comparison of Two Discontinuities in Anticancer Drug Discovery. *Organ. Sci.* 22(6): 1500–1516.
- Starbuck WH (2017) Organizational learning and unlearning. *Learning Organ.* 24(1): 30-38.
- Starbuck WH, Hedberg B (1977) Saving an organization from a stagnating environment. Hans B. Thorelli, ed. *Strategy+Structure=Performance: The Strategic Imperative* Indiana University Press, Bloomington 249-258.
- Stark O, Hyll W (2011) On the economic architecture of the workplace: Repercussions of social comparisons among heterogeneous workers. *J. Labor Econom.* 29(2): 349-375.
- Sterman JD (1987) Testing behavioral simulation models by direct experiment. *Management Sci.* 33(12): 1572-1592.
- Stettner U, Lavie D (2014) Ambidexterity under scrutiny: Exploration and exploitation via internal organization, alliances, and acquisitions. *Strategic Management J.* 35(13): 1903-1929.
- Stieglitz N, Knudsen T, Becker M (2016) Adaptation and Inertia in Dynamic Environments. *Strategic Management J.* 37(9): 1854–1864.
- Teece DJ (2007) Explicating dynamic capabilities: the nature and microfoundations of (sustainable) enterprise performance. *Strategic Management J.* 28(13): 1319-1350.
- The Economist* (2009) The importance of irrelevant alternatives. 22 May. <https://www.economist.com/democracy-in-america/2009/05/22/the-importance-of-irrelevant-alternatives>[20 April 2019].
- Tripsas M, Gavetti G (2000) Capabilities, Cognition, and Inertia: Evidence from Digital Imaging. *Strategic Management J.* 21(10/11): 1147–1161.
- Tushman ML, Anderson P (1986) Technological Discontinuities and Organizational Environments. *Admin. Sci. Quart.* 31(3): 439–465.
- Tushman ML, Romanelli E (1985) Organizational evolution: A metamorphosis model of convergence and reorientation. *Res. Organ. Behav.* 7, 171–222.
- Tushman ML, O'Reilly CA (1996) Ambidextrous Organizations: Managing Evolutionary and Radical Change. *California Management Rev.* 38(4): 8–30.
- Van de Ven N (2017) Envy and admiration: Emotion and motivation following upward social comparison. *Cogn. Emotion* 31(1): 193-200.
- Walsh JP, Ungson GR (1991) Organizational memory. *Acad. Management Rev.* 16(1): 57-91.
- Weber BJ, Chapman GB (2005) Playing for peanuts: Why is risk seeking more common for low stakes gambles? *Organ. Behav. Human Decision Processes* 97(1): 31-46.
- Weick KE, Sutcliffe KM (2006) Mindfulness and the quality of organizational attention. *Organ. Sci.* 17(4): 514-524.
- Winter SG (2000) The satisficing principle in capability learning. *Strategic Management J.* 21(10-11): 981-996.
- Yang H, Zheng Y, Zhao X (2014) Exploration or exploitation? Small firms' alliance strategies with large firms. *Strategic Management J.* 35(1): 146-157.
- Yoon E (2017) The Grocery Industry Confronts a New Problem: Only 10 Percent of Americans Love Cooking. Retrieved from <https://hbr.org/2017/09/the-grocery-industry-confronts-a-new-problem-only-10-of-americans-love-cooking%0D>
- Zenger T (2016) *Beyond competitive advantage: How to solve the puzzle of sustaining growth while creating value.* Harvard Business Review Press. Cambridge, MA.

Zenger (2012) We Wait Too Long to Train Our Leaders. 17 December.
<https://hbr.org/2012/12/why-do-we-wait-so-long-to-trai>[20 April 2019].

7. APPENDICES

Appendix Chapter 2 A

Optimal task behavior

It is worthwhile to contrast the observed behavior from the experiment with simulated behavior using different stylized decision strategies. Participants face a simple task of searching for several periods among alternatives of unknown value. Put differently, given the opportunity costs that would otherwise arise, the optimal behavior is to first explore a few alternatives and then exploit the best of the explored alternatives. Recall that the participants' task is to maximize the sum of the alternatives' values they played over the four periods of the task. In exploring a new alternative, participants face (opportunity) costs of failing to exploit the best-known alternatives. In exploiting a known alternative, participants face (opportunity) costs of failing to discover an even better alternative. Below, we discuss in more detail how the possibility for social comparison affects optimal search behavior.

Control condition (no social comparison). If the social comparison is not possible (control condition), a simple stopping rule is the optimal search strategy: first, explore for X periods and then exploit the best alternative found in these first X periods for the remaining $(4-X)$ periods. The optimal stopping point is independent of the number of alternatives; it is only a function of the number of periods available for exploitation and exploration. In our context, the optimal stopping point is also unaffected by the distribution of the alternatives' payoffs: in Figure 7-1 (dashed line), we plot the average payoffs (averaged over all 4 periods) for agents stopping exploration after 1 to 4 periods. Ideally, participants should explore 2 or 3 alternatives before settling on the best among the explored. Exploring fewer or more alternatives has a negative impact on average performance.

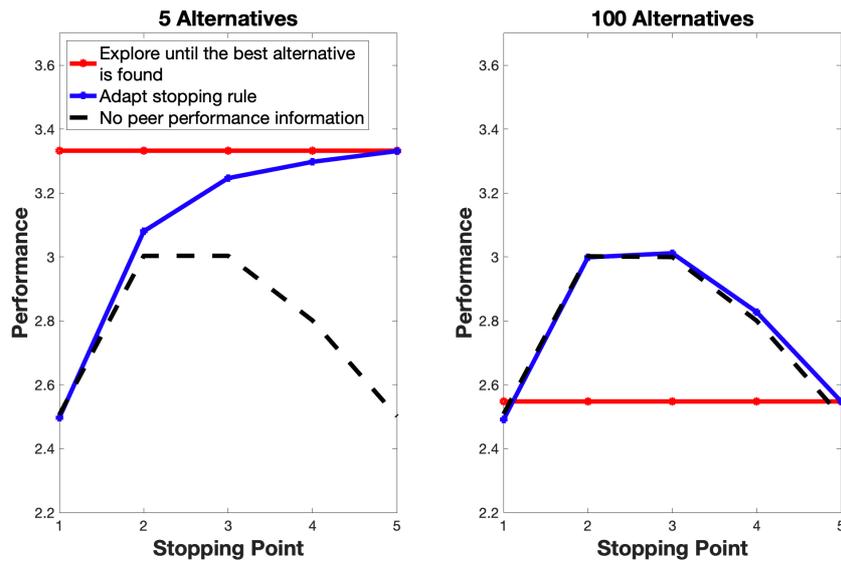


Figure 7-1: Optimal task behavior

Treatment (social comparison available). In the treatments, participants are informed about the payoff of the best alternatives. Broadly speaking, participants may use this information in two ways. First, they may explore until they find the best alternative. We call this strategy the exploration strategy. Second, they may adapt their stopping rule, i.e., explore for X periods unless they find the best alternative earlier than that. We call this second strategy the adaption strategy. By tendency, using the exploration strategy to utilize social comparison leads to increased exploration whereas the adaption strategy curbs exploration. The probability that the alternative corresponding to the treatment information is found within four periods is obviously a function of the number of decision alternatives in the choice set: if there are only very few alternatives, the probability is very high. If, however, there are many alternatives, this probability is lower. In our setups, the probability of finding the maximum of alternatives is 100 percent in the case of few decision alternatives (5 alternatives) and drops to less than 10 percent for many decision alternatives (100 alternatives). This, in turn, is consequential for the right strategy, exploration or adaption. Consider knowing the payoff of the best alternative. While it is beneficial to explore until it the maximum is found, particularly when the probability of finding this maximum is 100 percent, it is likely to induce over-exploration

when the probability of finding the maximum is low. Our simulation confirms that the exploration strategy is the most beneficial for limited choice, whereas the adaption strategy is the most beneficial for abundant choice, when the exploration strategy becomes a liability. In combining adaption on large choice sets and exploration on small ones, the peer performance treatment information always carries positive value.

Note that participants could also interpret the treatment information differently. While so far, we have treated the payoff of the best alternative as a performance target, one could also conceivably set the performance target to a fraction of this value. For example, a decision-maker confronted with his best peer performance may try to reach only a certain percentage of the maximum payoff. When allowing for these heuristics, some pitfalls of social comparison can be mitigated. In particular, they are a way to prevent over-exploration in the peer performance treatment.

In Figure 7-2, we show the performance implications of setting two-thirds of the maximum as the target. Analogously to utilizing the maximum, there are two possible strategies. First, participants could explore until an alternative at least as good as two-thirds of the maximum performance has been found. When there are only few decision alternatives, this yields similar performance results as aiming for the maximum. However, when faced with many decision alternatives, a two-third target yields almost as good results as in a situation of few decision alternatives, and significantly outperforms the control condition. The reason for this is simple: While it was less likely to find the maximum when there are more alternatives, the likelihood of finding a sufficiently close alternative is almost as high for small as for large choice sets. The same insight holds for the second possible strategy: Participants could adapt their stopping rule, i.e. explore X times unless an alternative at least two-thirds the value of the maximum can be found before. Performance effects for limited choice are similar to those featured in the case the 100 percent target. However, what we now

observe for abundant choice is a pattern similar to the pattern for limited choice. Thus, overall, setting a two-thirds of the maximum performance target can be superior to shooting for the maximum, particularly for a large choice set.

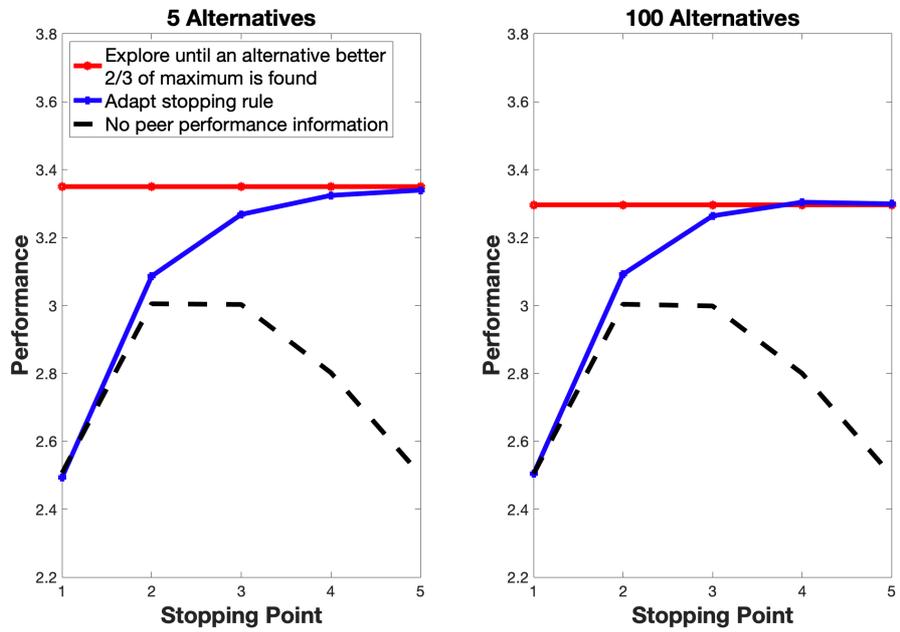


Figure 7-2: Optimal behavior using a 2/3 target heuristic

Appendix Chapter 2 B - Experiment

Dear mTurk worker,

You are now taking part in an experiment on decision making and learning.

It will take you about **5 minutes** to complete this experiment. Your base payment will be **5 cents** and you can earn a **bonus of up to 10 cents**, depending on your decisions (and, to a lesser extent, on luck). You may also choose to earn an additional **bonus of up to 5 cents** by answering a few questions about yourself and your decision-making approach, during this experiment.

Everything that you need to know to perform well in this experiment is explained on the following screens. If you have any questions, please drop us an email at empiricist1740@gmail.com

Figure 7-3: Instruction screen 1

In this experiment, you will take on the role of a manager of a company. There will be 4 periods in which you will have to decide about your company's strategy, i.e., choose among **5 different strategies**. The profits of these different strategies are initially unknown to you; only by trying a strategy, you will learn about its profits.

At the end of each period, you will learn what profits your chosen strategy generated. In the next period, you can decide whether to stick with the currently best performing strategy ("**exploit**") or try a different strategy ("**explore**").

Figure 7-4: Instruction screen 2

You will now play the experiment. The more profits you generate, the higher your bonus will be.

You are the manager of company X in industry Y.

Proceed to the next screen to know about the value of your company's current strategy.

Figure 7-5: Instruction screen 3

Period 1

The current strategy of your company generates **profits of \$76 million**.

Do you want to

- exploit** this strategy to earn a profit of \$86 million
- explore** an alternative strategy with an unknown profit

Period	Profits (in Million of \$)
1	
2	
3	
4	
5	

Figure 7-6: Decision screen first period

Period 3

The current strategy of your company generates **profits of \$86 million**.

In total, your company has now earned profits of \$162 million and explored 2 out of 5 possible strategies.

Do you want to

- exploit** the best strategy you tried so far to earn a profit of \$86 million
- explore** an alternative strategy with an unknown profit

Period	Profits (in Million of \$)
1	76
2	86
3	
4	
5	

Figure 7-7: Decision screen mid-task

Period 4

The current strategy of your company generates **profits of \$86 million**.
This is an attention check. If you are paying attention please proceed to the next screen without choosing any option.

Do you want to

- exploit** the best strategy you tried so far to earn a profit of \$27 million
- explore** an alternative strategy with an unknown profit

Period	Profits (in Million of \$)
1	
2	
3	
4	
5	

Figure 7-8: Attention check

Thank you for completing this part of the experiment.

In total, your company earned profits of \$347 million.

Please click below to proceed.

Figure 7-9: Final task screen

Finally, we will ask you four short questions on the understanding of the experiment.

To get an above average bonus in the end, you needed to choose a strategy with an above average profit in every period.

- True
- False

It was ideal to explore a few times and then only exploit the best found strategy for the remaining periods.

- True
- False

Figure 7-10: Comprehension questions 1

It was ideal to explore until the best available strategy has been found.

- True
- False

The bonus for the first part of the experiment was based on the sum of the profits you earned for your company.

- True
- False

Figure 7-11: Comprehension questions 2

Gender

- Male
- Female

Age

Why are you participating in this study?

- MTurk experiments are my primary source of income.
- To earn some additional money.
- For entertainment.
- To kill time.
- Out of curiosity.
- I am a researcher and I like to spy on other projects.

What is your highest level of education?

- No formal degree
- High school
- Associate's
- Bachelor's
- Master's
- J.D.
- M.D.
- Ph.D.

Have you ever heard of the multi-armed bandit ("n-armed bandit") model?

- Yes
- No

Have you ever heard about the "exploration / exploitation dilemma"?

- Yes
- No

Have you ever heard about "optimal search" theory?

- Yes
- No

Have you ever heard about the term "ambidexterity"?

- Yes
- No

Do you work in an organization with more than 10 employees?

- Yes
- No

If so, are you working in a managerial role?

- Yes
- No

Figure 7-12: Additional questions 1

Do you have a background in one of these fields?

- Mathematics
- Natural Sciences
- Computer Science
- Economics or Finance
- Statistics
- Engineering
- No I do not have a background in any of those fields

How much do you make, as an mTurk worker, per day on average?

- Less than \$5.
- Between \$5 and \$10
- Between 10 and \$20
- Between \$20 and \$30
- Between \$30 and \$40
- More than \$40.
- Prefer not to say.

What hourly wage did your last or current job pay?

- Less than \$5.
- Between \$5 and \$10
- Between 10 and \$20
- Between \$20 and \$30
- Between \$30 and \$40
- More than \$40.
- Prefer not to say.

Figure 7-13: Additional questions 2

You will now play the experiment. The more profits you generate, the higher your bonus will be.

You are the manager of company X in industry Y.

A consulting study concludes that, with the right strategy, your company's profits could be as high as \$121 million.

Proceed to the next screen to know about the value of your company's current strategy.

Figure 7-14: Changed introduction treatment 1

Period 3

The current strategy of your company generates **profits of \$86 million**.

In total, your company has now earned profits of \$162 million and explored 2 out of 5 possible strategies.

Remember that, with the right strategy, your company's profits could be as high as \$121 million.

Do you want to

- exploit** the best strategy you tried so far to earn \$86 million
- explore** an alternative strategy with an unknown profit

Period	Profits (in Million of \$)
1	76
2	86
3	
4	
5	

Figure 7-15: Changed decision screen treatment 1

What was the profit of the best available strategy in the previous experiment?

Million of \$:

Figure 7-16: Manipulation check treatment 1

Period 3

The current strategy of your company generates **profits of \$86 million**.
In total, your company has now earned profits of \$162 million and explored 2 out of 100 possible strategies.

Do you want to

- exploit** the best strategy you tried so far to earn \$86 million
- explore** an alternative strategy with an unknown profit

Period	Profits (in Million of \$)
1	76
2	86
3	
4	
5	

Figure 7-17: Changed decision screen treatment 2

Period 3

The current strategy of your company generates **profits of \$86 million**.

In total, your company has now earned profits of \$162 million and explored 2 out of 100 possible strategies.

Remember that, with the right strategy, your company's profits could be as high as \$121 million.

Do you want to

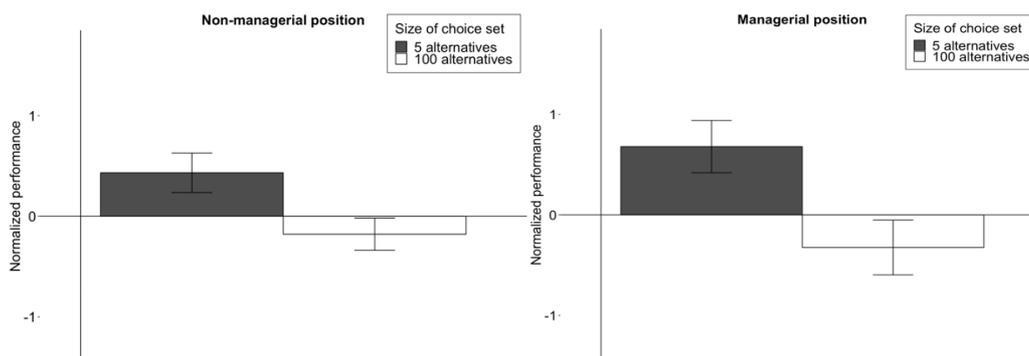
- exploit** the best strategy you tried so far to earn \$86 million
- explore** an alternative strategy with an unknown profit

Period	Profits (in Million of \$)
1	76
2	86
3	
4	
5	

Figure 7-18: Changed decision screen treatment 3

Appendix Chapter 2 C

To address the concern that our sample differs from the managers commonly being the object of investigation in the strategic management literature, we utilize the demographic information available about our participants. Concretely, we split the sample population by participants' characteristics that are common among managers. First, we asked participants if they held managerial positions at an organization with at least 10 employees. In Figure 7-19, we compare the performance effects across managers and non-managers. As a consequence of a much smaller sample of managers ($n = 160$) compared to non-managers ($n = 407$), performance effects become less pronounced. However, there are no significant performance differences between managers and non-managers.

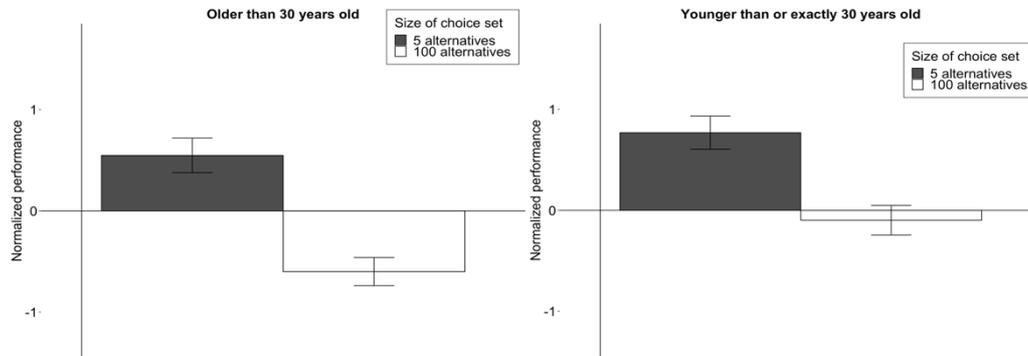


Note. Separation by managerial position is based on self-reported occupations of participants. Performance is normalized by subtracting performance in the control treatment.

Figure 7-19: Performance comparison between managers and non-managers

Second, we address the criticism that managers tend to be older than the average participant in online experiments. The average age of American CEOs upon appointment is 50. The average age of first managerial appointment is 30 years (Zenger, 2012). In our sample, the average age is at 37 years already close to the age group of managers. To provide additional evidence that the results are not driven by participants being considerably younger than managers, we split the sample by age. Specifically, we look at performance effects for

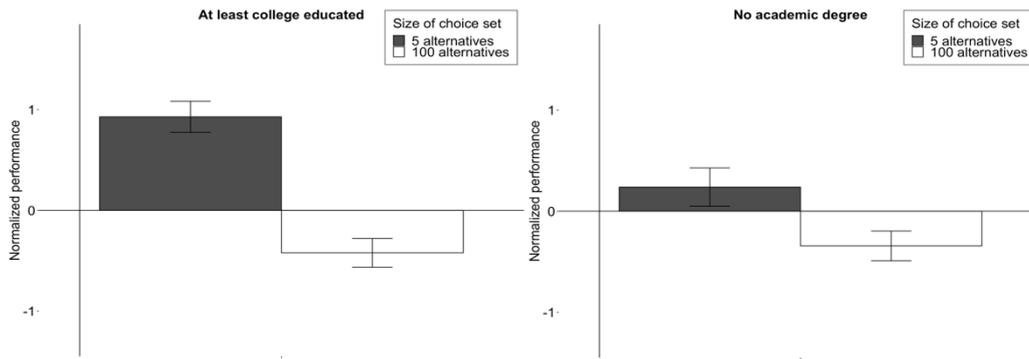
below versus above 30-year-old participants. As can be seen in Figure 7-20, no significant differences in performance exist between the two age groups.



Note. We chose 30 years as a cutoff point to split our sample by age because the vast majority of employees in their twenties are not occupying a managerial position with a large increase in that number in their thirties.

Figure 7-20: Performance comparison between age groups

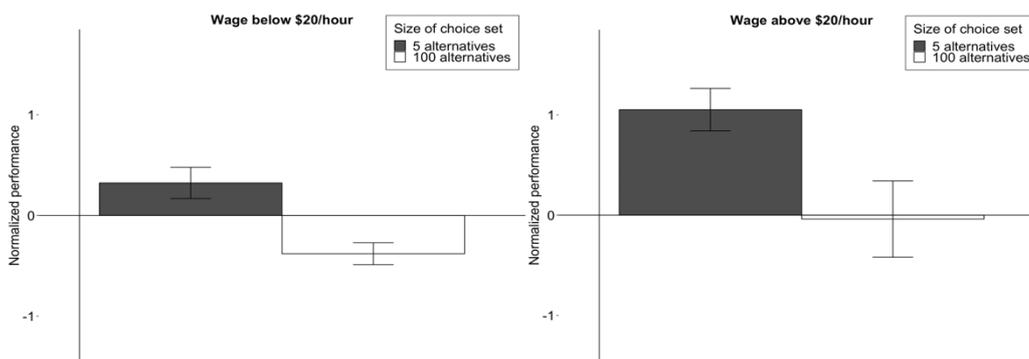
A third difference between mTurk workers and managers may be the level of education. While managers are often reported to have at least a college degree but often times also an MBA, a JD or comparable graduate-level degree, our sample features college educated and participants without college degrees in about equal shares. To establish that our results are not driven by these less educated part of the sample, we again conduct a split sample analysis contrasting performance effects for participants who hold at least a 4-year college degree with those who do not. As we see in Figure 7-21, there are no significant differences between the two levels of education.



Note. The cutoff point here is having completed at least a 4-year college education.

Figure 7-21: Performance comparison between levels of education

Lastly, managers ordinarily command a higher income than mTurk workers. To establish that our results are not in fact driven by a specific low-income category of participants, we conduct a split sample analysis for those who report earning more versus less than \$20 per hour, respectively (see Figure 7-22). Again, the direction of the effects is the same for both samples, even if participants with higher real-life salaries perform better and at a higher variance among themselves.



Note. We chose to divide the sample by a \$20-dollar threshold because that is roughly the median hourly wage for American workers (according to the Bureau of Labor Statistics).

Figure 7-22: Performance comparison between levels of income

Note. Sample choice screen from a first period of the peer performance treatment.

Appendix Chapter 3 A

<i>Values of Features</i>	<i>Car Features</i>		
	<i>Engine</i>	<i>Transmission</i>	<i>Body</i>
	V	Manual	Sedan
	Boxer	Semi-automatic	Convertible
	Inline	Automatic	Station wagon
	Wankel	Dual clutch	Hatchback

Table 7-1: Car features

Appendix Chapter 3 B

Dear mTurk worker,

You are participating in an experiment on decision making.

It will take you about **10 minutes** to complete this experiment. Your base payment will be **25 cents** and you can earn a **bonus of up to 1 dollar**, depending on your decisions (and, to a lesser extent, on luck). At the end, we would like you to answer a few questions about yourself.

Everything that you need to know to perform well in this experiment is explained on the following screens. If you have any questions, please drop us an email at empiricist1740@gmail.com



In the following experiment, you will have the opportunity to accumulate **performance points**.

You start out with 75 performance points.

In each round, you can win or lose performance points.

At the end of the experiment, your performance points are translated into US dollars. Each performance point will become 1 cent.



Figure 7-23: Instructions omissions experiment

In this experiment, you will assume the role of a manager of a car manufacturer.

You play for 20 rounds.

In each round, your chief technology officer (CTO) will present you a new car.

You can then decide that it may be introduced to the market. In this case, between -9 and 9 performance points will be added to your account.

If you reject the engineers' proposal, you will earn no performance points and no score of what you would have earned if you had accepted the proposal will be shown to you.



A car has three different features that your company can modify: an **engine (V, inline, boxer, wankel)**, **transmission system (manual, semi-automatic, automatic, dual clutch)** and its **body style (sedan, station wagon, convertible, hatchback)**.

Before you make a decision on a car, you will be informed about the composition of those features.

Each feature contributes to its performance points.

This means, one engine will be more profitable than another engine and one body style will be giving you more performance points than another one.

Please note, that these the value of he cars do in no way correspond to market prices of actual cars. Thus, you don't need any knowledge what so ever about cars to perform well in the following task.



Figure 7-24: Instructions omissions experiment (cont'd)

Your final balance is 79 performance points!



Please explain your strategy in the experimental task:

A large, empty rectangular text input field with a thin gray border, intended for the user to write their strategy.

Figure 7-25: Omissions post-experiment screens

Gender

- Male
- Female

Age

Why are you participating in this study?

- MTurk experiments are my primary source of income.
- To earn some additional money.
- For entertainment.
- To kill time.
- Out of curiosity.
- I am a researcher and I like to spy on other projects.

What is your highest level of education?

- No formal degree
- High school
- Associate's
- Bachelor's
- Master's
- J.D.
- M.D
- Ph.D.

Do you have a background in one of these fields?

- Mathematics
- Natural Sciences
- Computer Science
- Economics or Finance
- Statistics
- Engineering
- No, I do not have a background in any of those fields.

What hourly wage did your last or current job pay?

- Less than \$5.
- Between \$5 and \$10.
- Between \$10 and \$20.
- Between \$20 and \$30.
- Between \$30 and \$40.
- More than \$40.
- Prefer not to say.

How much do you make, per day, as an mTurk worker on average?

- Less than \$5.
- Between \$5 and \$10.
- Between \$10 and \$20.

- Between \$20 and \$30.
- Between \$30 and \$40.
- More than \$40.
- Prefer not to say.

Figure 7-26: Omissions post-experiment questionnaire

Your main competitor company employs a decision maker who has the same duties as you.

She gets the same proposals featuring cars with the exact same elements as the ones you are confronted with.

In every round, you will see her performance as well as the proposals she received.



Round 2

Do you approve or reject car no. 2?

			Own Performance		Competitor Performance		
Round No.	Features	Decision	Round	Total	Features	Round	Total
1	hatchback,wankel,manual	Accepted	7	82	hatchback,wankel,manual	7	82
2	station wagon,V,automatic		?				

The car earned you 3 performance points in the market!

Figure 7-27: Omissions treatment information – competitor treatment

Round 3

Do you approve or reject car no. 3?

Approve

Reject

Round	Features	Decision	Performance	Total Performance
1	hatchback,inline,automatic	Accepted	1	76
2	station wagon,boxer,dual clutch	Accepted	5	81
3	hatchback,boxer,automatic		?	
4				

You rejected the proposed car and hence, earn no performance points in this round. If you would have accepted the proposal you would have earned -3 performance points.

Figure 7-28: Omissions treatment information – full information treatment