

**Evolutionary Behavioral Economics:
Essays on Adaptive Rationality in Complex Environments**



Stefano Benincasa

Supervisor: Prof. Marco Faillo

Department of Economics and Management
University of Trento

Dissertation submitted in fulfilment of the requirements for the degree of
Doctor of Philosophy in Economics and Management

June 2020

I would like to dedicate this dissertation to my parents,
without whose loving support the journey leading to this work would not have been possible.

Declaration

I hereby declare that the contents of this dissertation are original and have not been submitted in whole or in part for consideration in fulfillment of any other degree or qualification in this or any other university. This dissertation is my own work and contains nothing that is the work of others, except as specified in the text and the Acknowledgements. This dissertation contains fewer than 65,000 words including appendices, bibliography, tables, equations, and has fewer than 150 figures.

Stefano Benincasa
June 2020

Acknowledgements

My greatest debt is to my supervisor at the University of Trento, Marco Faillo, whose proficient guidance navigated me through various aspects of this dissertation. He has been a superlative mentor and our discussions have been a constant source of inspiration over the years. I wish to thank the technician of the Cognitive and Experimental Economics Laboratory (CEEL), Marco Tecilla, for offering me seamless assistance in the development of the experimental software, its ongoing upgrade, and the arrangement of the relevant sessions. I would like to thank the Doctoral School of Social Sciences for hosting me and generously funding my research. I am particularly indebted to my external advisor at the University of Southern Denmark, Stephan Billinger, for disclosing many critical details concerning the experimental framework and delivering several insightful comments. I also would like to thank the whole Strategic Organization Design (SOD) unit for welcoming my research visit, as well as all the people who made such treasured experience feasible. Last but not least, I am grateful to the referees, Luigi Marengo and Dimitry Rtischev, for contributing their time and expertise to critically reading the text and advancing thoughtful observations that definitely helped improving the quality of the work.

Abstract

Against the theoretical background of evolutionary behavioral economics, this project analyzes bounded rationality and adaptive behaviour in organizational settings characterized by complexity and persistent uncertainty. In particular, drawing upon the standard NK model, two laboratory experiments investigate individual and collective decision-making in combinatorial problems of resource allocation featuring multiple dimensions and various levels of complexity. In the first study, investment horizons of different length are employed to induce a near or distant future temporal orientation, in order to assess the effects of complexity and time horizon on performance and search behaviour, examine the presence of a temporal midpoint heuristic, and inspect the moderating effects of deadline proximity on the performance-risk relationship. This is relevant for organizational science because the passage of time is essential to articulate many strategic practices, such as assessing progress, scheduling and coordinating task-related activities, discerning the processual dynamics of how these activities emerge, develop, and terminate, or interpreting retrospected, current, and anticipated events. A greater or lesser amount of time reflects then a greater or lesser provision of resources, thereby representing a constraint that can greatly affect the ability to maintain a competitive advantage or ensure organizational survival. In the second study, the accuracy of the imitative process is varied to induce a flawless or flawed information diffusion system and, congruently, an efficient or inefficient communication network, in order to assess the effects of complexity and parallel problem-solving on autonomous search behaviour, clarify the core drivers of imitative behaviour, control for the degree of strategic diversity under different communication networks, and evaluate individual as well as collective performance conditional to the interaction between the levels of complexity and the modalities of parallel problem-solving. This is relevant for organizational science because imitating the practices of high-performing actors is one of the key strategies employed by organizations to solve complex problems and improve their performance, thereby representing a major part of the competitive process. The project is intended to contribute grounding individual and collective behaviour in a more psychologically and socially informed decision-making, with a view to further the research agenda of behavioral strategy and sustain the paradigm shift towards an evolutionary-complexity approach to real economic structures.

Contents

List of Figures	xiii
List of Tables	xv
1 Introduction	1
1.1 Theoretical Background	1
1.1.1 The As-If Hypothesis	1
1.1.2 Three Perspectives on Bounded Rationality	2
1.1.3 Towards an Evolutionary Behavioral Economics	5
1.2 The Project	9
1.2.1 Research Focus	9
1.2.2 Method	10
1.2.3 Aims and Objectives	10
1.2.4 Value	11
2 Time Horizon Effects in Adaptive Search on Rugged Landscapes	13
2.1 Introduction	14
2.2 Theory and Hypotheses	17
2.2.1 Complex Problems as Rugged Landscapes	17
2.2.2 Time-related Constructs in Organizational Research	18
2.2.3 Time and Individual Learning	20
2.2.4 Temporal Midpoint Heuristic	22
2.2.5 Moderating Effects of Deadline Proximity	23
2.3 Method	25
2.3.1 Task	25
2.3.2 Design	25
2.3.3 Participants	26
2.3.4 Procedure	26
2.4 Results	28

2.4.1	Descriptive Statistics	28
2.4.2	Regression Analyses	34
2.5	Discussion	42
2.6	Conclusions	44
3	Parallel Problem-Solving in Networked Groups	47
3.1	Introduction	48
3.2	Theory and Hypotheses	50
3.2.1	Complex Problems as Rugged Landscapes	51
3.2.2	Parallel Problem-Solving	52
3.2.3	Organizations and Imitative Behaviour	54
3.2.4	Cognitive Foundations of Imitative Behaviour	55
3.3	Method	57
3.3.1	Task	57
3.3.2	Design	58
3.3.3	Participants	59
3.3.4	Procedure	59
3.3.5	Descriptive Statistics	62
3.3.6	Regression Analyses	70
3.4	Discussion	80
3.5	Conclusions	83
4	Conclusions	85
4.1	Future Directions	88
4.2	Final Remarks	89
	References	91
	Appendix A	125
	Appendix B	131

List of Figures

2.1	Performance	28
2.2	Search Distance	30
2.3	Search Distance (Alternative Measure)	32
2.4	Adaptive Search	38
3.1	Strategic Diversity	63
3.2	Performance	64
3.3	Search Distance	66
3.4	Imitative Behavior	69
3.5	Adaptive Search	74

List of Tables

2.1	Frequency Distribution of Search Distances	31
2.2	Frequency Distribution of Search Distances (Alternative Measure)	33
2.3	Main Variables	34
2.4	OLS Models with Performance as Dependent Variable	35
2.5	Poisson Models with Search Distance as Dependent Variable	36
2.6	Probit Models with Feedback as Dependent Variable	38
2.7	Stepwise Poisson Models with Search Distance as Dependent Variable to Analyze the Temporal Midpoint Heuristic	40
2.8	Poisson Models with Search Distance as Dependent Variable to Analyze the Moderating Effects of Deadline	41
3.1	Frequency Distribution of Search Distances	68
3.2	Main Variables	70
3.3	Poisson Model with Strategic Diversity as Dependent Variable and OLS Model with Performance as Dependent Variable	71
3.4	Poisson Models with Search Distance as Dependent Variable (Independent Problem-Solving)	72
3.5	Poisson Models with Search Distance as Dependent Variable (Parallel Problem-Solving)	75
3.6	Probit Models with Imitation as Dependent Variable	78
3.7	Probit Models with Individual and Social Feedback as Dependent Variable	80

1

Introduction

“The goal of science is to make the wonderful and complex understandable and simple - but not less wonderful.”

H. A. Simon, *The Sciences of the Artificial*

1.1 Theoretical Background

Real economic structures are composed of multiple, interacting individuals who operate on various spatial and temporal scales. Of necessity, a critical issue is: how are we to understand these systems and the relative processes?

1.1.1 The As-If Hypothesis

Since the late nineteenth century, the standard economic approach - pioneered by talented figures such as Léon Walras (e.g., 1874/1954), Wilfred Pareto (e.g., 1906/1971), John Hicks (e.g., 1939), Paul Samuelson (e.g., 1947), Kenneth Arrow (e.g., 1954), Gérard Debreu (e.g., 1954), Milton Friedman (e.g., 1953), and others, including Leonard Savage (e.g., 1954) - has based its efforts on the theory of general equilibrium. Abstracting from the number of its alternative definitions, the general equilibrium posits that, based on their knowledge, all individuals make plans that are optimal. Provided that these beliefs are mutually consistent, the relevant plans can then be simultaneously implemented with no adjustment being required - or even desired - as time unfolds, except in response to occasional exogenous shocks.

To outline individual decision-making and the resulting aggregate outcomes, the standard approach resorted to a fictitious entity commonly referred as homo economicus, which is assumed to be endowed with perfect information on all alternatives, exact expectations about future events, faultless memory, and limitless computational capabilities. In line with this “Olympian” form of rationality (Simon, 1983, p. 19), the homo economicus is deemed able

to represent, evaluate, and compare all the possible alternatives that lie within her feasible choice set and select the best affordable option. Regardless of whether the favorite paradigm to model the decision-making process coincides with the one based on rational - that is, complete and transitive - preferences or the one based on choices that are presumed to reveal such preferences (Mas-Colell, 1995), the decisions of the homo economicus can be interpreted as-if she were a self-regarding utility maximizer - or expected utility maximizer, if the alternatives are associated with probabilities (Savage & Friedman, 1948; Friedman, 1953).

The creators of the standard, neoclassical economic approach did not intend the model to be a faithful reflection of human nature. Rather, their objective was inherently practical, as they aimed to overcome the arbitrariness and lack of focus afflicting classical economics through a shift of emphasis towards the logical rigor and theoretical clarity of the decision-making process. Thence, the rational choice theory was purposefully devised as a mathematical simplification able to narrow down the set of outcomes that could arise from economic systems and to make the problem under consideration tractable, thereby yielding a remarkable analytical power. In the parlance of Friedman (1953), even though individuals did not possess the formal tools - or the cognitive resources - to calculate the optimum, they still behave as-if they did, just as much as bicycle riders keep themselves in equilibrium even though they are unaware of the equations behind the dynamics of motion. Moving from this premises, Friedman concluded that there is no need for the rules in decision processes to be descriptively realistic, since predictions drawn upon largely simplified assumptions about individual choice are nonetheless accurate in the aggregate.

De facto, variants of the general equilibrium theory, such as the Arrow-Debreu model, the Solow growth model, and dynamic stochastic general equilibrium models, have been employed with modest success to understand the past, explain the present, and forecast the future with the purpose of guiding policy choices (see Epstein, 2008 for a list of applications). However, “modest success implies a degree of failure” (Page, 2016, p. 320), and neoclassical economics can be held responsible for some notable ones, such as the failure to conclusively solve - much less to predict - the most recent global economic collapse or to prevent the current hazard to the sustainability of the planet.

1.1.2 Three Perspectives on Bounded Rationality

The neoclassical program was soon questioned in the second half of the twentieth century, as it became evident that human behaviour systematically deviates from the assumptions of rational choice theory under a variety of circumstances, such as decision under risk (e.g., Allais, 1953; Ellsberg, 1961; Tversky & Kahneman, 1979), other-regarding preferences

(e.g., Rabin, 1993; Fehr & Schmidt, 1999), and intertemporal choices (e.g., Thaler, 1981; Loewenstein & Prelec, 1992; Liberman & Trope, 1998).

In a string of seminal contributions that laid the foundations for the behavioral revolution, Herbert Simon (e.g., 1956, 1978) epitomized the need to radically overhaul economic theory proceeding from an appropriate understanding of the cognitive limitations intrinsic to the decision-maker or, as he termed it, from her bounded rationality. In Simon's view, an appropriate understanding of human bounded rationality should satisfy three essential criteria. First, it should account for the real processes that underlie the decision of individuals and organizations, thereby moving beyond as-if theories of maximizing utility or expected utility. Second, it should address situations in which the preconditions for rationality established by the neoclassical model are not fulfilled (Simon, 1989). More thoroughly, it should extend to situations where an agent does not have all the possible alternatives and probabilities, but must generate them through a lengthy and costly search process, and where she cannot make an optimal choice, but must instead satisfice - which is to say, choose an option that is "good enough" (Simon, 1957) to meet some endogenous aspiration level. Third, it should explicate the ecological structures of the environment to which the agent must adapt in order to survive. Ultimately, the assumption of bounded rationality can be summarized using a theoretical analogy. As Simon magisterially put it: "Human rational behavior is shaped by a scissors whose two blades are the structure of the task environments and the computational capabilities of the actor" (Simon, 1990, p. 7).

Notably, Simon did not believe he had solved the problem of human decision-making and, indeed, never advanced a complete theory of bounded rationality (Gigerenzer, 2004). Rather, he advocated for the need of such theory, inspiring with his work the generations of researchers to come who, over time, have incorporated the notion of bounded rationality in several models along at least three major lines of inquiry.

The first is the reductionist line, which is predominant among neoclassical economists. According to the reductionist interpretation, bounded rationality is nothing subversive. It simply consents to model fully optimal procedures that accommodate costs in terms of time, computation, money, or any other resources constraint (e.g., Sargent, 1993; Rabin, 1998; Arrow, 2004). To this end, insights from psychology, the neurosciences, or biology are assimilated in the standard apparatus by introducing additional features, such as the value and weighting functions in decision under risk, the iniquity aversion or reciprocity parameters in other-regarding preferences, or the "beta" hyperbolic discounting factor in intertemporal choices. At last, this makes the behavioral approach a mere "neoclassical repair shop" (Güth, 2002), whose finality is to increase the empirical plausibility and explanatory power of rational choice theory (Camerer & Loewenstein, 2004). Simon rejected this interpretation

without reserves, openly observing that: "bounded rationality is not the study of optimization in relation to task environments" (Simon, 1991, p. 35).

The second line of inquiry is represented by the heuristics-and-biases program launched by the much-celebrated work of Kahneman and Tversky (e.g., Tversky & Kahneman, 1974; Kahneman et al., 1982; Kahneman, 2011). This approach relies on a research strategy based on a simple three-step *modus operandi*. First, take a principle from logic, probability theory, statistics, or decision theory that is universally recognized as a normative standard of rationality. Second, determine whether the decisions of individuals deviate from the principle in exam. Third, if a deviation is found, explain behaviour either in terms of heuristics - i.e., cognitive shortcuts - or in terms of biases - i.e., systematical errors in judgment and decision-making. Following this protocol, the heuristics-and-biases program has produced an impressive inventory of deviations from traditional rational choice theory (see Conlisk, 1996; Krueger & Funder, 2004), such as the representativeness heuristic, the availability heuristic, the anchoring-and-adjusting heuristic, loss aversion, the framing effect, the illusion of control, the confirmation bias, the in-group bias, and the present bias, to name just a few. The interpreters of the heuristics-and-biases program attributed these putative "fallacies" - or "mental illusions" - to the automatic and intuitive rather than to the controlled and analytical portion of the Janus-faced cognitive system that characterizes human beings. It is worth noting that, although these heuristics and biases posed a keen challenge to the normative standards of rational choice theory, they were still appreciated as evolutionary residues that have an instrumental role in supporting individuals with limited cognitive resources to confront a complex and uncertain world. For this reason, some authors have seen in the map of systematic and predictable differences between normative and descriptive behavior captured by the heuristics-and-biases program the achievement of Simon's theory of bounded rationality (e.g., Thaler, 1991; Ariely, 2009). However, this interpretation tends to demote psychology to the ancillary function of carrying out a clinical diagnosis of those instances that mislead human judgment and decision-making (Hertwig & Pedersen, 2016). What is more, Simon himself pointed out that "bounded rationality is not irrationality" (Simon 1985, p. 297), thereby suggesting that a theory of bounded rationality should be something more than a dry catalog of systematic flaws that burden human decision-making.

The third and last line of inquiry rests upon an ecological orientation, which has its roots in the work conducted by Gigerenzer, Todd, and the ABC group (e.g., Gigerenzer & Goldstein, 1996; Gigerenzer & Todd, 1999; Gigerenzer, 2000; Gigerenzer & Selten, 2002; Goldstein & Gigerenzer, 2002; Hertwig & Erev, 2009; Hertwig & Herzog, 2009; Todd et al., 2012; Hertwig et al., 2013; Hertwig & Engel, 2016; Hertwig & Mata, 2019). The ecological program resumes and emphasizes Simon's remark that, for a theory of

human rationality to be appropriate, it must be able to describe and predict not only the cognitive mechanisms underlying the behavior of individuals and organizations, but also the relationship between these mechanisms and the environment in which they are carried out. Given this, in stark contrast with the map of systematic biases proposed by the heuristics-and-biases program, Gigerenzer, Todd, and ABC group built a systematic model of ecological rationality, conceptualized as an “adaptive toolbox” to implement simple strategies able to exploit the informational structure of the natural and social environments. Notably, these strategies - usually referred as fast and frugal heuristics - are not unconditionally successful. Rather, their performance depends on whether cognition matches the environment (Brunswik, 1957), and whether the heuristic in place possesses the key informational structures of its specific domain (Sperber, 1994). Ultimately, rationality is no longer measured against the logical principles of consistency, transitivity, or compliance to some normative standard, but rather against the correspondence principles of speed, frugality, and accuracy of outcomes in the process of adaptation to the environment (Hammond, 1996). The research program on ecological rationality singled out an ensemble of domain-specific strategies that humans employ to make inferences (e.g., the take-the-best heuristic, the recognition heuristic, the fluency heuristic), choices (e.g., the priority heuristic), allocations in games against nature or against others (e.g., the social-circle rule, the average rule, the majority rule, the successful-member rule, the equity rule), or learn about the consequences of their decisions and the associated probabilities (e.g., the description-experience gap). To further highlight Simon’s tenet that “the mind is an adaptive system whose biological function is to enable the organism to behave effectively and, hence, to survive in a complex, changing, and often unpredictable environment” (Langley & Simon, 1981, p. 362), the notion of ecological rationality has been subsequently rephrased as adaptive rationality (e.g., Kenrick et al., 2012; Hertwig & Pedersen, 2016).

1.1.3 Towards an Evolutionary Behavioral Economics

Most recently, an assorted group of researchers gathered at the 2015 Ernst Strungmann Forum to probe whether concepts and methods from evolutionary theory and complexity theory could be integrated in order to promote a paradigm shift in the study of human economic systems. As will be seen, such evolutionary-complexity approach to economics, also referred as evolutionary behavioral economics (Burnham et al., 2016), represents an ideal framework to position the notion of adaptive rationality and study its implications for individual and organizational decision-making within a broader context.

Importantly, the forum was not meant to offer a ready-made and accomplished theory, but to review the state-of-art of research and herald future directions of inquiry. According to

its promoters, the creation of an evolutionary behavioral economics is a work in progress, which “will take many years and the concerted effort of many scholars” (Burnham et al., 2016, p. 114). However, different groups conveyed different positions, thereby holding divergent opinions about how to proceed. The reason is that, taken in the singular, the expressions ‘evolutionary theory’ and ‘complexity theory’ are oxymorons (Padgett, 2016). As a matter of fact, contemporary research on evolutionary processes cannot be merely equated to the Modern Synthesis of the Darwinian approach occurred throughout the 1930s and 1940s, now commonly referred as population genetic (e.g., Fisher, 1930; Wright, 1932; Dobzhansky, 1937; Ford, 1940; Huxley, 1942), but covers a range of domains that spans from the evolution of development (e.g., Gould, 1977; Jacob, 1977; Lewis, 1978; Nüsslein-Volhard & Wieschaus, 1980), to epigenesis (e.g., Waddington, 1942; Holiday, 1990; Riggs, 1996), neutral drift (e.g., Kimura, 1968; King & Jukes, 1969), autocatalysis in the origin of life (e.g., Kauffman, 1995; Dawkins, 2004), niche construction in the coevolution of genetic and culture (Boyd & Richerson, 1985; Richerson & Boyd, 2004; Gintis, 2011), and multilevel selection (e.g., Wilson & Sober, 1994; Nowak et al., 2010). Likewise, contemporary research on complex systems involves an array of domains as diverse as information entropy (e.g., Shannon, 1948), dissipative systems (e.g., Prigogine, 1962), nearly decomposable systems (e.g., Simon, 1969), spin-glass models (e.g., Anderson, 1994), genetic algorithms (e.g., Holland, 1975), NK models of genetic networks (e.g., Kauffman & Weinberger, 1989; Kaufmann, 1993), self-organized criticality (e.g., Bak, 1996), small-worlds networks (Watts, 1999), power laws (e.g., Barabási, 2002; Ijiri & Simon, 1977), and chaos theory (e.g., Birkhoff, 1927; Cartwright & Littlewood, 1945; Kolmogorov, 1954; Lorenz, 1963, 1972; Li & York, 1975; Libchaber et al., 1982). Thence, how evolutionary theory and complexity theory are to be integrated depends on a large extent upon which elements out of these loose clusters one decides to draw (Padgett, 2016).

Notwithstanding, an evolutionary behavioral economics can still be organized around a set of overarching principles. More in detail, in the evolutionary-complexity approach, real economies are conceptualized as complex adaptive systems that are frequently out of equilibrium (Kirman & Sethi, 2016; Wilson, 2016). By definition, for a system to be complex, it requires a number of components that interact with one another directly or indirectly, so that it is analytically impossible to determine *ex ante* what the aggregate behaviour of the overall system will be simply focusing on the components in question. However, for a complex system to be adaptive, it further requires that these components must be capable of adaptation to one another as well as to the environment (Axtel et al., 2016). In such systems, the components correspond to agents - which is to say, individuals or organizations - that make fitness-enhancing choices based on simple internal rules, past actions, and new inputs

(Holland, 2006; Miller & Page, 2007; Page, 2007; Mayfield, 2016). Simple rules specify how the agents interact with each other as well as how they interact with the environment. Past actions reflect path-dependence as determined by selective pressure, in the absence of which choices are no different than random mutations that, as such, may very well induce the agent to mismatch the environment and be no longer fitness-enhancing. The attainment of new inputs captures instead the processes of individual and social learning, in line to which preferences evolve, strategies evolve, and the external environment evolves, thereby changing fundamentally the nature of the system. In the last instance, individual choices based on simple rules can nonetheless determine the emergence of a dazzling variety of unstable aggregate patterns and non-linear systemic dynamics.

However, depending on the level at which adaptation and selection take place, complex adaptive systems can be of two types (Wilson, 2016). A complex adaptive system of type 1 (CAS1) is adaptive as a whole, so that it represents by itself the unit of selection (e.g., a market and its rules). In contrast, a complex adaptive system of type 2 (CAS2) is composed by individuals that employ adaptive strategies, in which case it is the single component that represents the unit of selection (e.g., the strategies of the agents that operate in a market). Clearly, human economic systems are more sophisticated than that. Strictly speaking, they are made of a multigroup population within a multi-tier organizational hierarchy, wherein adaptation and selection occur on multiple levels. To function well as a group, agents must provide services one to each other undertaking some cost in terms of time, energy, and risk of exploitation by eventual free-riders. In principle, groups of prosocial - that is, other-regarding - individuals have an advantage compared to groups of more antisocial - that is, self-regarding - individuals, but only provided that between-group selection does prevail against within-group selection (Wilson & Wilson, 2007; Wilson, 2015). The same latent conflict applies to all levels of the multi-tier hierarchy. What is good for me may be bad for my social circle, what is good for my social circle may be bad for my local community, what is good for my local community may be bad for my country, and so forth. Therefore, a multilevel selection posits that adaptation at any level of a multi-tier hierarchy involves a process of selection at that level and tends to be undermined by selection at lower levels (Wilson, 2015). Notably, the balance between levels of selection is not static, but evolves over time. Sometimes, higher-level selection can evolve to dominate lower-level selection, thereby prompting what has become known as a major transition (Szathmáry & Maynard Smith, 1995; Maynard Smith & Szathmáry, 1999). While in biological systems major transitions are understood as functional changes in the evolution from genes to cells, from cells to multicellular organisms, and from solitary organisms to societies, in human economic systems they are conceived as

large-scale technological innovations like those introduced by industrial revolutions (Axtel et al., 2016).

Ultimately, in the evolutionary-complexity approach to economics, individuals and organizations are the result of a slower-paced natural evolution of cognitive resources and a faster-paced cultural evolution of knowledge and technology (Wilson et al., 2016). Rather than a homogeneous - or “representative” - rational agent in possess of complete information and able to optimize a set of exogenous, fixed preferences subject to no internal constraints or contracting costs, the evolutionary-complexity approach postulates a population of heterogeneous purposive agents with mutable preferences that satisfice endogenous aspiration levels in accordance to internal constraints (Wilson et al., 2016; Gowdy et al., 2016; Colander, 2016; Axtel et al., 2016; Burnham et al., 2016; Padgett, 2016). Moving beyond the static allocations of goods and services that characterize efficient markets in a state of equilibrium, individuals adapt to changes in their natural and social environments by making incremental improvements to their strategies. In turn, these fitness-enhancing adjustments alter the external environment(s), thereby leading to further dynamic adjustments. The resulting trajectories might eventually converge to one out of several possible equilibria, in which case they function as an equilibrium selection device, although this is not necessary by any means (Kirman & Sethi, 2016). Furthermore, in contrast to a self-organizing economy resulting from the naïve group selection led by an invisible hand (Smith, 1759), which imagines that the unregulated pursue of self-interest - commonly conceptualized as financial wealth - provides robust benefits to the common good, the evolutionary-complexity approach argues for a multilevel group selection running across a multi-tier economy of mostly other-regarding individuals, in which selection at each level of the hierarchy is undermined by selection at lower levels. Under this perspective, a revised and more subtle conception of invisible hand is still admissible, as long as self-organization emerges from a process of adaptation and selection that occur at the level of the whole system (Wilson, 2015, 2016).

There are two core arguments in favor of a paradigm shift in the study of economic systems based on the integration of elements from evolutionary theory and complexity theory. In the first place, the evolutionary-complexity approach holds the potential to enhance the economic model of individual behaviour. In line with the programmatic research conducted in experimental and behavioral economics, this comprehends - but is not limited to - "grounding individual behaviour in a more psychologically and socially informed decision-making" (Burnham et al., 2016, p. 112). The rationale is that insights from psychology, neurosciences, and behavioral economics can be understood in a greater depth incorporating evolutionary theory and complexity theory into the levels of analysis. Secondly, the evolutionary-complexity approach holds the potential to aid in mapping not simply the

short-term, but most outstandingly the long-term causes and consequences (evolution) and connections (complexity) of human behaviours in their natural and social environments. On the one hand, this is expected to enhance our understanding of the dangers as well as the opportunities of behavioral interventions in a view of improving their effectiveness (Burnham et al., 2016). On the other, this would allow public policy to shape the coevolution of market and government by guiding the direction of innovation, resource use, and distribution through formation mechanisms able to modify the structures in which individuals and organizations happen to operate (Colander, 2016; Mazzucato, 2016; Gowdy et al., 2016).

1.2 The Project

1.2.1 Research Focus

This project fits into the research agenda of behavioral strategy. By definition, breaking out of the equilibrium mindset, behavioral strategy merges cognitive and social psychology with strategic management theory and practice, with the purpose to bring realistic assumptions about human cognition and social behavior to the strategic management of organizations and, thence, to enrich strategy theory, empirical research, and real-world practice (Powell et al., 2011). Its agenda identifies four core research problems: (1) scaling individual cognition to collective behavior, (2) defining the psychological underpinnings of strategy theory, (3) understanding complex problem-solving in organizations, and (4) improving the psychological architecture of the firm (Powell et al., 2011). Accordingly, I believe it is fair to argue that behavioral strategy occupies a central place in the evolutionary-complexity approach to economics.

In particular, the present project focuses on a laboratory investigation of bounded rationality and adaptive behaviour in combinatorial problems that, drawing upon the NK model, are operationalized as performance landscapes featuring multiple dimensions and various levels of complexity. Originally developed in the field of evolutionary biology (Kauffman & Weinberger, 1989; Kaufmann, 1993), the NK model has proven to be a powerful tool to analyze organizational adaptation processes (Levinthal, 1997) across a number of complex strategic activities, such as individual decision-making (e.g., Gavetti & Levinthal, 2000), organizational decision-making (e.g., Knudsen & Levinthal, 2007), new product development (e.g., Mihm et al., 2003), product modularity (e.g., Marengo & Dosi, 2005), open innovation (e.g., Almirall & Casadesus-Masanell, 2010), organizational design (e.g., Rivkin & Siggelkow, 2003), industry dynamics (e.g., Lenox et al., 2007), and more (see Ganco & Hoetker, 2009 for a comprehensive review). The implications of the NK model have also been taken into

account to inform empirical research on technology and product development (Fleming & Sorenson, 2004; Frenken, 2006), firm boundaries (Sorenson, 2003), the administration of start-up companies (Sommer et al., 2009), and industry profitability (Lenox et al., 2010).

1.2.2 Method

Thus far, most of the published work on strategic organization appealing to the NK model has been explicitly theoretical. Correspondingly, researchers have largely tapped into the behavioral theory of the firm (Cyert & March, 1963) to elaborate stylized behavioral rules that are expected to subtend the decision-making process of individuals and organizations in NK scenarios. As a result, despite the rich body of research, we still miss an appropriate understanding of how human beings actually behave in similar settings. In the interest of addressing this shortcoming, Billinger, Stieglitz & Schumacher (2014) have recently introduced an experimental framework to assess human behaviour on performance landscapes, thereby paving the way for laboratory implementations of the NK model. The studies conducted in pursuance of the present project build on such framework.

There are several reasons for employing laboratory experiments to investigate human behaviour in complex, combinatorial tasks. First, a laboratory experiment provides full control over the environment, which consents to rule out extraneous influences that may introduce more or less severe forms of endogeneity. Second, in the absence of empirical data from the field, a laboratory experiment provides a convenient device to weight the empirical plausibility of behavioral rules and, thence, to test for model validity (Hey, 1982; Lave & March, 1975; Winter, 1982; Sterman, 1987; Lam, 2010). Third, a laboratory experiment provides results that may foster theoretical enhancements and stimulate empirical research (Arthur, 1991). In a nutshell, the problem of adaptive behaviour as exhibited by real people requires an empirical foundation, and laboratory experiments are a primary source of empirical evidence (Selten, 1995). Undoubtedly, data from the field are also important, but they are more difficult to gain and harder to analyze (Selten, 1998).

1.2.3 Aims and Objectives

In the first study, investment time horizons of different length are implemented to induce a near or distant future temporal orientation. The aim is to appraise how available time affects individual adaptation on performance landscape with multiple dimension and various levels of complexity. More in detail, the study has three objectives: (1) assessing the effects of complexity and time horizon on performance and search behaviour; (2) examining the

presence of a temporal midpoint heuristic; (3) inspecting the moderating effects of deadline proximity on the performance-risk relationship.

In the second study, the experimental framework is extended to allow for the diffusion of social information through imitation and, therewith, for parallel problem-solving. In particular, the accuracy of the imitative process is varied to induce a flawless or flawed information diffusion system and, congruently, an efficient or inefficient communication network. The aim is to appraise how parallel problem-solving in networked groups affects collective adaptation on performance landscape with multiple dimension and various levels of complexity. More precisely, the study has three objectives: (1) evaluating the degree of strategic diversity of networked groups and their subsequent performance conditional on the interaction between the levels of complexity and the modalities of parallel problem-solving; (2) appraising whether prior results on autonomous search behavior under independent problem-solving extends to autonomous search behavior under parallel problem-solving; (3) shedding light on the determinants of imitative behavior, with a special attention to crossover effects, social feedback, and strategic distance.

1.2.4 Value

The findings from the project are intended to provide empirical evidence conducive to improve our understanding of the psychological and social basis of individual and collective decision-making in organizational settings characterized by complexity and persistent uncertainty. In line with the prior work of Arthur (1991), Edmonds (2001), and Billinger et al. (2014), the findings are expected to have also theoretical implications for the development of empirically grounded search algorithms to simulate stylized behavioral rules within the computational modelling of adaptive processes. In so doing, I hope they will aid grounding human behaviour in more realistic assumptions, thereby contributing to the advancement of the behavioral strategy agenda and, by extension, to the construction of a behavioral evolutionary economics.

2

Time Horizon Effects in Adaptive Search on Rugged Landscapes

With Marco Faillo and Stephan Billinger

Abstract

This chapter presents findings from a laboratory experiment in which investment horizons of different length are employed to induce a near or distant future temporal orientation in combinatorial tasks with multiple dimensions and various levels of complexity crafted using the NK model. We find that, independently from the level of complexity, a shorter time frame introduces a form of pressure in the investment of scarce resources, thereby inducing subjects to explore the problem space on a larger scale than they would do with a longer time frame. In terms of performance, this is beneficial when the task is simple, while it tends to be detrimental when the task is complex. No evidence is found in support of a temporal midpoint heuristic. However, deadline proximity acts as moderator for the performance-risk relationship, so that underperforming subjects shift the focus of attention from performance to survival, thereby opting for risk-averse exploitation, while outperforming subjects shift the focus of attention from performance to experimenting with slack resources, thereby opting for risk-seeking exploration. We discuss the implications of our findings for the calibration of a temporal lens in organizational decision-making against the backdrop of behavioral strategy.

2.1 Introduction

Time and human experience are inextricably intertwined, even more so in the working life (Murray, 1938; Lewin, 1943). De facto, virtually all organizations operate within a time frame to deal with research and development, the design of strategies, and the solution of problems. As a result, the passage of time is of the essence when assessing progress (e.g., Gersick, 1988; Gevers et al., 2006, 2013; Lehman et al., 2011), scheduling and coordinating task-related activities (e.g., McGrath & Rotchford, 1983; Schriber & Gutek, 1987; McGrath & Tschan, 2004; Quintens & Matthysens, 2010), discerning the processual dynamics of how these activities emerge, develop, and terminate (Langley et al., 2013; Pitariu & Ployhart, 2010; Ployhart & Vandenberg, 2010; Roe, 2008; Zaheer et al., 1999), as well as interpreting retrospected, current, and anticipated events (e.g., Johns, 2006; Rousseau & Fried, 2001; Shipp & Jansen, 2011). Under this perspective, a greater or lesser amount of time reflects a greater or lesser provision of resources, thereby representing a constraint that can greatly affect the advancement of a project, the resulting outcome and, ultimately, the ability to gain a competitive advantage or ensure organizational survival (e.g., Hofer & Schendel, 1978; Gersick, 1988; Gersick & Hackman, 1990; Rajagopalan & Spreitzer, 1997; Eisenhardt & Martin, 2000; Pettigrew et al., 2001). Furthermore, as organizations aim to acquire assets or develop capabilities that yield future benefits exceeding upfront costs (Quirin & Wiginton, 1981), it may well be claimed that the very notion of investment relies on time (Reilly et al., 2016).

On these grounds, it is manifest how the time horizon has a critical role in the process of resource allocation. So there is no wonder if researchers have advocated on various occasions that adopting a temporal lens is indispensable for the advancement of organizational science (e.g., McGrath & Rotchford, 1983; Ancona et al., 2001a, 2001b; Bluedorn, 2002; George & Jones, 2000; Roe, 2008; Sonnentag, 2012; Shipp & Cole, 2015; Reilly et al., 2016; Kunisch et al., 2017). Indeed, multiple streams of research in the organizational literature are incorporating assumptions concerning the time horizon, while interest in temporal issues is gaining momentum and increasing exponentially (see Shipp & Cole, 2015; Kunisch et al., 2017). Although temporal research in organizational science is at a point of inflection (Shipp & Cole, 2015), with the exception of a handful of experimental studies (Gersick, 1989; Woolley, 1998; Okhuysen & Waller, 2002; Okhuysen et al., 2003; Joireman et al., 2006), conclusions are often grounded in plausible deduction moving from backwards-looking measures mainly distilled from archival data (Reilly et al., 2016). In this regard, some scholars have argued that such a deficit of empirical evidence on the effects of time horizon in organizational practices derives from the difficulty to actually measure temporal patterns and tie them with specific investment choices, thereby calling for more precise and nuanced

measurements of this hard-to-capture but necessary aspect of resource allocation (e.g., Devers et al., 2008; Souder & Shaver, 2010).

The present research aims to address this shortcoming by explicitly manipulating the time horizon in a laboratory experiment featuring a complex, combinatorial choice structure, which is particularly suitable to describe the problem of resource allocation in strategic organization activities (Schumpeter, 1934; Simon, 1962; Page, 1996; Levinthal, 1997). As it happens, an organizational strategy typically involves a number of decision variables, such as extent of vertical integration, product design, pricing policy, channels of distribution, research and development program, and more (Porter, 1991; Ghemawat & Levinthal, 2008). The conundrum of gaining a competitive advantage or ensuring the survival of the organization may then be conceptualized as a search for successful combinations based on alternative resource allocations over these decision variables (Rivkin, 2000). Strictly speaking, the problem of identifying a successful strategy is regarded as complex if the value of a combination depends on the interactions between the relevant decision variables (Simon, 1962; Page, 1996; Levinthal, 1997).

Resorting to a laboratory experiment to investigate search behaviour in complex, combinatorial problems entails several advantages. First, a laboratory experiment provides full control over the environment, thereby consenting to rule out extraneous influences that may introduce more or less severe forms of endogeneity. Second, a laboratory experiment offers a convenient device to weight the empirical plausibility of behavioral rules and, thence, to test for model validity, especially when empirical data from the field are scant (Hey, 1982; Lave & March, 1975; Winter, 1982; Sterman, 1987; Lam, 2010). Third, a laboratory experiment produces results that may foster theoretical enhancements and stimulate empirical research (Arthur, 1991). In a nutshell, the study of human behaviour requires an empirical foundation, and laboratory experiments are a primary source of empirical evidence (Selten, 1995). Clearly, data from the field are also important, but usually they are more difficult to gain and harder to analyze (Selten, 1998).

Over the last decades, a rich body of experimental works has studied the impact of time horizon in a wide range of settings, such as various game-theoretic scenarios (e.g., Roth & Murnighan, 1978; Holt, 1985; Camerer & Weigelt, 1993; Palfrey & Rosenthal, 1994; Engle-Warnick & Slonim, 2004; Dal Bó, 2005; Charness & Genicot, 2009; Blonski et al., 2011; Dal Bó & Frechette, 2011; Offerman et al., 2011), patch-leaving assignments (Hutchinson et al., 2008), multi-armed bandit policies (Lee et al., 2011; Wu et al., 2019), the secretary problem (Seale & Rapaport, 2000), negotiation of agreements concerning benefits and burdens (Okhuysen et al., 2003), and organizational citizenship behaviour (Joireman et al., 2006). Notwithstanding, prior experimental research failed to appraise the impact of time

horizon within a complex, combinatorial choice structure able to instantiate the problem of resource allocation in organizations.

In this laboratory experiment, investment horizons of different length are employed to induce a near or distant future temporal orientation in combinatorial tasks featuring multiple dimensions and various levels of complexity. The tasks were built using the NK model, which provides an algorithm to generate a combinatorial problem space characterized by a more or less rugged performance landscape. Parameter N represents the number of decision variables, while parameter K regulates the interactions among them. The key feature of the model is that, as K increases, the ruggedness of a performance landscape changes from a single-peaked landscape to a multi-peaked landscape (Kauffman & Weinberger, 1989; Kaufmann, 1993). For these properties, the NK model has established itself as a canonical approach to analyze behavioral rules across a wide range of complex, strategic organization settings.

Proceeding from this outline, the experimental investigation has three objectives: (1) assessing the effects of complexity and time horizon on performance and search behaviour; (2) examining the presence of a temporal midpoint heuristic; (3) inspecting the moderating effects of deadline proximity on the performance-risk relationship.

The main empirical findings suggest that, independently from the level of complexity, a shorter time frame introduces a form of pressure in the investment of scarce resources, such that human agents are prone to explore the problem space on a larger scale than they would do with a longer time frame. Although this seems to be somewhat beneficial in terms of performance when the task is simple, the same attitude tends to be particularly detrimental when the task is complex. Search behaviour does not obey to the temporal midpoint heuristic. However, deadline proximity does act as moderator for the performance-risk relationship, so that underperforming subjects tend to shift the focus of attention from the reference point coinciding with one's own best performance in prior trials to survival, thereby opting for risk-averse exploitation, while outperforming subjects are prone to shift the focus of attention from the same reference point to experimenting with slack resources, thereby opting for risk-seeking exploration.

The remainder of the chapter is organized as follows: section 2.2 reviews the related literature and introduces our main research hypotheses; section 2.3 illustrates the experimental design and procedures; section 2.4 reports our key results; section 2.5 discusses their implications for behavioral strategy; section 2.6 concludes.

2.2 Theory and Hypotheses

In what follows, we review prior organizational research on complex strategic activities, time-related constructs, individual learning, the temporal midpoint heuristic, and the moderating effects of deadline proximity on the performance-risk relationship. Contextually, we proceed to delineate our four primary hypotheses.

2.2.1 Complex Problems as Rugged Landscapes

A problem can be regarded as complex if the possible solutions involve the combination of multiple decision variables that are interdependent with respect to their impact on performance (Simon, 1962; Page, 1996; Levinthal, 1997). To derive a stylized representation of a complex problem, much of the recent theoretical work in organizational science has relied on the NK model. Originally developed in the field of evolutionary biology (Kauffman & Weinberger, 1989; Kaufmann, 1993), the NK model has proven to be a powerful device to analyze organizational adaptation processes (Levinthal, 1997) across a number of complex, strategic organization domains, such as individual decision-making (e.g., Gavetti & Levinthal, 2000), organizational decision-making (e.g., Knudsen & Levinthal, 2007), new product development (e.g., Mihm et al., 2003), product modularity (e.g., Marengo & Dosi, 2005), open innovation (e.g., Almirall & Casadesus-Masanell, 2010), organizational design (e.g., Rivkin & Siggelkow, 2003), industry dynamics (e.g., Lenox et al., 2007), and more (see Ganco & Hoetker, 2009 for a comprehensive review). The implications of the NK model have also been taken into account to inform empirical research on technology and product development (Fleming & Sorenson, 2004; Frenken, 2006), firm boundaries (Sorenson, 2003), the administration of start-up companies (Sommer et al., 2009), and industry profitability (Lenox et al., 2010).

The NK model specifies an algorithm to generate performance landscapes based on N attributes with K interactions, which determines their relative degree of interdependence. Accordingly, a problem can be expressed as a space of alternatives, where each alternative consists of N - in its basic version binary - attributes. Since each attribute can take on only two possible values, the landscape consists of 2^N alternatives. Nonetheless, depending on K , the contribution of each attribute may be influenced by the interactions with other attributes. In particular, if K is equal to 0, attributes are fully independent, and their contribution depends just on their own value. At the other extreme, if K is equal to $N-1$, attributes are fully interdependent, and their contribution depends also on the value of all other attributes. By modulating the degree of interdependence among the attributes, K affects the ultimate complexity of the landscape. Specifically, when attributes are fully independent, the landscape

appears smooth and it is characterized by exactly one peak. However, as K increases, local peaks augment and the landscape becomes progressively more rugged.

Adaptation on rugged landscapes is then assumed to unfold throughout two basic processes. The first is local search, in which the agent looks in the immediate neighborhood by changing one attribute at the time (i.e., one bit mutation or incremental search). This process reflects political and organizational routines in which time and budget constraints bind search to local improvements. The second is global search, in which the agent makes “long jumps” (i.e., random, distant, or non-incremental search) to visit more distant alternatives. This reflects a different sort of routines, leading to failure-induced radical reorientations. The key implication is that in non-complex - that is, single-peaked - landscapes, local search allows to reach the global maximum. However, in complex - that is, multi-peaked - landscapes, local search is likely to get trapped in local optima, thereby prompting a reduction in average performance.

Billinger et al. (2014) have recently introduced an experimental framework to assess individual search behaviour on rugged landscapes, thereby paving the way for laboratory implementations of the NK model. The results indicated that search behavior in the NK search task gradually adapts to performance feedback, encoded as success or failure in discovering a better-performing combination relative to a subjective reference point coinciding with the best-performing combination identified in prior trials. In particular, a positive feedback (success) induces subjects to exploit the region of the landscape where they have experienced the performance increase, while a negative feedback (failure) induces them to explore more distant regions. The evidence suggested also that complexity does not have a direct effect on search behavior. Rather, it affects the feedback conditions that guide success-induced exploitation and failure-induced exploration. Nonetheless, since human subjects gradually adapts to performance feedback, it can be concluded that complexity has an indirect impact on human search behaviour.

In the present investigation, we aim at replicating these results in the long-term condition, whose time horizon approximates that of the standard task in Billinger et al. (2014). Furthermore, we are interested in assessing whether their explanatory power holds up in the short-term condition.

2.2.2 Time-related Constructs in Organizational Research

A wide array of research traditions in organizational science - including agency theory, real options and the behavioral theory of the firm - have implicitly or explicitly incorporated some form of temporal component into their theorizing. Although interest is gaining momentum and increasing exponentially, thus far organizational research on temporal issues has been

unsystematic, haphazardly ranging from individual-level factors (see Ship & Cole, 2015 for a comprehensive review) to organizational-level strategic change (see Kunisch et al., 2017 for a comprehensive review). The result is that no dominant paradigm exists and the temporal lens of organizational science is still blurred. Nonetheless, circumscribing the focus of attention to those studies that have been specifically interested in addressing the role of time within the context of resource allocation, despite the fragmentary nature of research, an integrative summary can still be advanced to articulate clear definitions and subsume a number of dispersed insights under few cardinal constructs (Reilly et al., 2016).

The first relevant construct is temporal orientation, conceived as a “future time perspective” capturing the variance across individuals “in terms of the relative cognitive dominance of the near versus the distant future” (Das, 1987, p. 203), thereby influencing the allocation of resources between options with more immediate or more deferred payoffs. While Das formulated temporal orientation at the level of individual differences, the construct has been extended to account for the prevailing collective preference of the firm conditional on the personal preferences of current managers and their understanding of the relevant historical patterns (Marginson & McAulay, 2008; Souder & Bromiley, 2012), thereby reflecting a firm-specific investment policy (Bower, 1970; Bromiley, 1986; Maritan, 2001).

Another relevant construct is investment horizon, understood as the *ex ante* managerial expectation about the duration of time over which potential investments will generate productive returns (Reilly et al., 2016). It is straight-off evident how investment horizon is equivalent to payback (Connelly et al., 2010) or payoff horizon (Souder & Shaver, 2010). Similarly to temporal orientation, the concept of investment horizon is applicable at multiple levels of analysis. At the project-level, the investment horizon is indispensable to clock expectations for a single investment decision, thereby representing the condition *sine qua non* for the execution of a net present value (NPV) analysis (Brealey & Myers, 1996) functional to direct the allocation of resources. At the firm-level, instead, the investment horizon connotes in aggregate a portfolio of investments. Firms that routinely select short-horizon projects can be portrayed as having a short investment horizon, whereas firms that routinely select long-horizon project can be portrayed as having a long investment horizon. Clearly, due to the constrain stemming from the inherent uncertainty in estimating future returns, in reality firms may adopt heterogeneous approaches based on a synthesis of these two extremes.

Temporal orientation and investment horizon flow into two broader constructs, namely short-termism and long-termism. Short-termism represents the extent to which investment decisions aim at desirable near-term outcomes at the expenses of later-term outcomes (Lavery, 1996) or, alternatively, the propensity by decision-makers to forego good, long-term opportunities to pursue short-term targets (Marginson & McAulay, 2008). The same notion has

been conveyed using a number of terms, such as present focus (Cojuharenco et al., 2011), hyperbolic discounting (Frederick et al., 2002; Dasgupta & Maskin, 2005; Plambeck & Wang, 2013), and temporal myopia (Jacobs, 1991; Levinthal & March, 1993; Miller (2002). Real-world examples of short-termism are readily provided by the tendency exhibited by managers to emphasize near-term investments with the purpose to maximize their career value (Bower, 1970; Thakor, 1990), as well as by the narrow framing of decisions and outcomes displayed by stock market investors (Shleifer & Vishny, 1990; Stein, 1988, 1989; Thaler et al., 1997). In contrast, long-termism represents the extent to which decision-makers are prone to myopically overemphasize later-term prospects in place of near-term objectives. The idea of long-term myopia has received substantial attention in the theoretical literature (Stein, 1989; Bebchuk & Steal, 1993; Levinthal & March, 1993; Miller, 2002) and is presumed to be widespread among family-owned businesses (Chua et al., 2009; Lumpkin & Brigham, 2011).

Two laboratory experiments within organizational research have explicitly manipulated the time horizon inducing controlled shifts in the attention of individuals towards a shorter or longer temporal length. Specifically, Okhuysen et al. (2003) inspected the effects of time horizon on the efficiency of negotiated agreements, finding that a long-term horizon, particularized as a delay in the realization of outcomes, increases the efficiency of the negotiated agreements. In addition, the type of resource, namely burden or benefit, moderates this relationship, so that the salutary discounting effects are greater for burdens. Joireman et al. (2006) examined the effects of time horizon on organizational citizenship behaviors (OCBs), conceived as a social dilemma in which short-term employee sacrifice leads to long-term organizational benefits, finding that a long-term horizon is associated with higher levels of OCBs.

2.2.3 Time and Individual Learning

Nearly all of the research on time in organization science presumes that longer horizons foster better behavior and performance outcomes, whereas shorter horizons have negative consequences not only for the profitability of organizations, but also for the natural and social systems in which they operate (Hayes & Abernathy, 1980; Porter, 1992; Gross & Lewis, 2007). The rationale is to be found in the association between time and individual learning. When it comes to executive functions, in fact, the amount of time available for search and problem-solving has a key role in the formation and eventual revision of schemata (Billinger & Srikanth, in preparation), understood as organized pieces of information forming a coherent system of knowledge that underlies a particular capability (Rumelhart & Norman, 1978; Rumelhart, 1980; Ericsson & Staszewski, 1988).

More precisely, a longer time frame put subjects in the condition to carry out more experiments, identify better alternatives, and gradually refine the gathered knowledge, thereby increasing not only the volume, but also the worth of their schemata. From a cognitive standpoint, this ability relies on the progressive improvement of long-term memory and learning from recall (Karpicke & Roediger, 2008), and is exceptionally valuable as subjects must face complex phenomena requiring more complex mental models (White & Frederiksen, 1986; Clancey, 1987). Several authors have also argued that the benefits deriving from a longer time frame are intrinsically related with an incubation effect (Wallas, 1926) that enables subjects to forget misleading cues (Smith & Blankenship, 1991; Seifert et al., 1995), recover from mental fatigue to better take into account news cues (Seifert et al., 1995; Dorfman et al., 1996), and subconsciously refine the information stored in the recess of long-term memory (Yaniv & Meyer, 1987; Seifert et al., 1995).

In the wake of these arguments, it is easy to see how, conversely, a shorter time frame has instead a detrimental effect on human learning, especially when the required body of knowledge is unfamiliar and complex (White & Frederiksen, 1986; Clancey, 1987; Winter, 2000). Indeed, since subjects are unable to efficiently retain and employ the freshly acquired knowledge, a shorter time frame tends to hamper the level of practice and incubation, thereby leading to weak schemata, unfit learning and, ultimately, inferior performance (Wallas, 1926; Logan, 1988; Yaniv & Meyer, 1987; Smith & Blankenship, 1991; Seifert et al., 1995; Dorfman et al., 1996).

At an experimental level, the relevant effect has been documented in a wide range of tasks, such as attractiveness judgments relative to gain and losses (Svenson & Benson, 1993a) or positive and negative information (Edland, 1993), the exploration and exploitation of alternatives in order to appraise the resources that are available and those that are needed for solving a problem (Svenson & Benson, 1993b), but also criteria-dependent binary choices (Wallsten, 1993), sequential binary choices (Busemeyer, 1993), creativity in organizations (Amabile et al., 2002), cognitive lock-ups in the completion of an assignment (Schreuder & Mioch, 2011), binary bets framed in terms of gains and losses (Young et al., 2012), queuing (Conte et al., 2016), and University exams (De Paola & Gioia, 2016).

In line with epistemic lay theory, further studies have also highlighted how subjects exposed to shorter time frames tend to exhibit cognitive closure and epistemic freezing. Accordingly, they are prone to draw hasty conclusions based on insufficient experimentation and scarce cognitive plasticity, thus freezing upon solutions that they are reluctant to alter thereafter (Kruglanski & Freund, 1983; Mayseless & Kruglanski, 1987). Notably, such behaviour can readily be interpreted as further evidence in support of a diminished attitude to exploration corresponding to a shrink in the time frame.

(H1) *A longer time horizon is associated with higher average performance*

(H2) *A longer time horizon is associated with greater exploration*

2.2.4 Temporal Midpoint Heuristic

A different strand of research in the field of organizational behaviour has noticed that, when the development of a project is provided with a well-defined horizon, rather than looking at absolute levels of task progress, individuals pace their work using temporal milestones (Gersick, 1988, 1989). By definition, temporal milestones are reference points in time that prompt individuals to break off inertial patterns and change the way they approach a task (Gersick & Hackman, 1990). Although individuals may generate endogenous milestones associated with any arbitrary point in time (Gersick, 1994), the empirical evidence suggested that subjects are prone to follow an heuristic approach by placing the key temporal milestone in correspondence of the midpoint of the task, which acts as a turning point that splits work pacing into two major phases (Gersick, 1988, 1989). In the first phase, subjects focus on collecting information and producing new ideas. In the second, they restrain from introducing further inputs and concentrate on integrating the gathered knowledge to finalize the task. Accordingly, the temporal midpoint of the task marks the transition from exploration to exploitation.

In his theory-building paper, Gersick (1988) made use of a field study in which task forces were brought together to conduct a project in organizational settings as varied as banks, hospitals, fundraising groups, and universities. Likewise, Ericksen and Dyer (2004) employed a field-based approach centered on the analysis of task forces responsible for carrying out non-routinary projects in undisclosed organizations. Shifting empirical method, Gersick (1989) conducted a laboratory experiment in which participants acted as professional advertising writers who were asked to come up with a pilot for a client whose preferences they were made privy. Woolley (1998) implemented a laboratory study in which group discussions were held to solve a restaurant problem. In a similar fashion, Okhuysen and Waller (2002) resorted to a laboratory study focused on group discussions to manufacture ludic structures. Based on the arrangement of a well-defined time frame, all studies provided evidence in support of the midpoint heuristic and the associated behavioral pattern of exploration and exploitation.

Building on this framework, Ford and Sullivan (2004) extended the theoretical model posing that exploration is especially beneficial in the first phase of work pacing, whereas it is likely to disrupt performance and induce frustration in the second phase. Further investigating the role of feelings, Knight (2015) proposed that the mood of subjects around the midpoint of

the task is crucial for the unraveling of the exploration-exploitation pattern after the relative transition. In particular, he maintained that a positive mood is associated with increasing exploitation, while a negative mood is associated with sustained exploration. A contextual test relying upon longitudinal data on groups preparing for a military competition appeared to support these predictions.

(H3) *The midpoint of the task marks the transition from exploration to exploitation*

2.2.5 Moderating Effects of Deadline Proximity

Lehman et al. (2011) emphasized the importance of deadline proximity for moderating the performance-risk relationship in organizational decision-making. Drawing on the shifting-focus-of-attention model (March & Shapira, 1992) and its subsequent extensions (Chen & Miller, 2007; Iyer & Miller, 2008), they argued that risk-taking in organizational practices depends on three main factors: performance relative to an aspiration level, the focus of organizational attention, and the availability of time as determined by the proximity of a deadline.

Within their paradigm, the shifting-focus-of-attention model (March & Shapira, 1992; Chen & Miller, 2007; Iyer & Miller, 2008) served as foundation for understanding the former two factors. Strictly speaking, the aspiration level as a reference point represents a cornerstone of the classic behavioral theory of the firm (Cyert & March, 1963), which describes organizations as goal-oriented entities pursuing performance targets that can be based on an organization's own past performance or on the performance of other organizations (March, 1988; March & Shapira, 1987; Greve, 1998; Bromiley, 2004; Knudsen, 2008). On top of that, the relative distance rule (March & Shapira, 1992) holds that organizational attention will be focused on whatever reference point is closest. Correspondingly, the closer the performance of an organization is to its reference point, the greater the likelihood that attention will be focused on it. The further the performance of an organization is from its reference point, the greater the likelihood that attention will be focused on survival or slack resources.

Based on these conditions, Lehman and colleagues derived a combined risk function, which posits an inverted U-shaped performance-risk relationship for performance levels below the reference point, and a positive performance-risk relationship for performance levels above the reference point. More intently, they reasoned that the availability of time before the end of a given period - through which organizations measure and monitor their performance relative to a reference point (Carp, 2003) - represents a unique type of resource able to trigger shifts in the focus of organizational attention. When more time is available,

attention is likely to be directed at attaining and maintaining an aspired level of performance. Conversely, as less time is available and the deadline draws nearer, attention is likely to be redirected at ensuring survival in underperforming firms and at experimenting with slack resources in outperforming firms.

From a behavioral standpoint, it can be conjectured that, contingent on their prior performance, decision-makers form a “cognitive image” of the future, which tends to be perceived as less uncertain as the end of the period approximates. Specifically, decision-makers in underperforming organizations are more likely to visualize their end-of-period performance as underperforming. They become prey of learned helplessness (Abramson et al., 1980), presenting greater rigidity and the inability to generate novel alternatives (Staw et al., 1981). Otherwise, they may act in compliance of the principle imposing not to put the organization in peril (March & Shapira, 1987), exhibiting then a keen aversion to taking risks (Lopes, 1987). In both cases, decision-makers in underperforming organizations are expected to focus their attention on ensuring the survival of the organization by engaging in behaviours that are perceived as less risky (Audia & Greve, 2006; Miller & Chen, 2004).

On the other hand, decision-makers in outperforming organizations are more likely to visualize their end-of-period performance as outperforming. As the end of the period approach, with a level of performance above aspirations and fewer chances that unforeseen events will damage it, the same excess of resources previously interpreted as a “success” (Cyert & March, 1963) and regarded as a deterrent to any short-term pitfall (Chen & Miller 2007; Iyer & Miller 2008) will now be seen as an “opportunity” for additional experimentation (Baum et al., 2005). Moreover, outperformers are likely to perceive a greater control over aspiration attainments, building growing self-confidence in their ability to manage risks (Langer, 1975; March & Shapira, 1987). Ultimately, decision-makers in outperforming organizations are expected to believe they are in the position to indulge in examining novel alternatives that might further enhance the success of the organization (March, 1991; March & Shapira, 1992), thereby focusing on the investment of slack resources by engaging in behaviors previously considered too risky (Forlani, 2002; Chen & Miller 2007; Iyer & Miller 2008; Nohria & Gulati, 1996).

The authors tested their model using data from the American National Football League (NFL). In particular, they examined how available time before the end of a game affects the relationship between a team’s score and its willingness to take risks, finding overall support for their predictions.

(H4) As the end of the task approximates, underperformers should resort to more risk-averse exploitation, while outperformers should resort to more risk-seeking exploration.

2.3 Method

2.3.1 Task

Building on the laboratory framework introduced by Billinger et al. (2014), the core experimental task consisted of a combinatorial problem based on the NK model, in which multiple components were to be assembled to generate different possible configurations of a product. The specific number of components depended on parameter N , which was set equal to 10. Since any of the ten components was susceptible of a binary choice - that is, on or off - the resulting problem space comprised 2^{10} (1,024) combinations. Each combination was associated with a payment generated by a standard NK algorithm, thereby recreating a performance landscape that the subjects explored through experiential search. To limit the impact of foreknowledge and prevent subjects from tapping into some cognitive prior able to guide search (Gavetti & Levinthal, 2000), the task was framed as "the space traveler game", wherein participants were required to design a number of images by combining ten abstract symbols. The products were then to be sold to the space traveler, whose willingness to pay a different price for different products reflected the structure of the fitness landscape.

2.3.2 Design

The experiment revolved around a 2x2 mixed factorial design, with the 'Level of Complexity' (None, Intermediate) varied within subjects, and the 'Investment Time Horizon' (Short-term, Long-term) varied between subjects.

The level of complexity depended on parameter K , which regulates the ruggedness of a landscape. More thoroughly, the smooth, non-complex landscape was randomly generated setting $K = 0$, whereas the rugged, complex landscape was randomly generated setting $K = 5$. To specify the relevant interdependencies, we employed a random interaction matrix. The landscapes were normalized using two alternative multipliers to hide mean and maximum performance, so as to prevent the formation of reference points across the tasks. Although randomly generated, the landscapes had a representative number of peaks, that is: exactly 1 for $K = 0$, and 32 for $K = 5$.

Turning to the investment horizon, available time was manipulated to capture short-termism and long-termism by inducing a near or distant future temporal orientation, respectively. Both conditions involved a well-defined number of trials. In the short-term condition, subjects were informed that each space traveler was going to buy 11 images. In the long-term condition, they were informed that each space traveler was going to buy 21 images. Since the first image was sold automatically at the beginning of each encounter, subjects produced

10 images per encounter in the short-term condition and 20 per encounter in the long-term condition. Notably, the design was tuned so that the investment time horizon in the short-term condition was exactly half the span of the long-term condition.

The design accommodated two sources of persistent uncertainty. The first concerned payments, which remained unknown until an image was sold and, therefore, could only be progressively discovered over the course of the search process. The second concerned the interdependencies between components as specified by the random interaction matrix underpinning the complex task, owing to which the contribution of each component was impinged by K further components that could not be clearly identified.

2.3.3 Participants

72 students (51 % females, 49 % males) were randomly selected using the university recruitment system and received real monetary incentives. Subjects were on average 20 years old and came from a variety of educational backgrounds. The participants were randomly assigned to the 2 between-subjects conditions, each carried out across 2 sessions of 18 individuals, for a total of 36 subjects per condition. The 4 sessions were held in the computer-based Cognitive and Experimental Economics Laboratory (CEEL) at the University of Trento.

2.3.4 Procedure

At the beginning of every session, participants visualized treatment-specific instructions, which were read aloud by one of the experimenters. The instructions described the experimental task and the structure of the relevant session, emphasizing how each participant was going to confront two space travelers who came from two different planets and, consequently, had completely different tastes and preferences, thereby making clear that the knowledge acquired in the previous encounter would not be pertinent in the following one. Technically, the two space travelers reflected two different landscapes: one smooth and simple ($K = 0$), the other rugged and complex ($K = 5$). Participants were also informed about the exact number of trials.

As participants correctly answered all the understanding questions, they had access to the first encounter. Thereafter, as all of them had concluded the first encounter, they had access to the second and last encounter. To prevent the occurrence of carry-over effects, for each between-subject condition (Time Horizon) we counter-balanced the within-subject levels (Complexity) across the two dedicated sessions. The resulting sequences induced two alternative task positions for each landscape, thereby allowing to control for eventual learning effects.

In the first trial of each encounter, participants were provided with a preliminary combination that, unbeknownst to them, always coincided with the lowest-performing one. Starting from the second trial, each subject could modify none, some or all the attributes of the preliminary combination. After making a choice, she learnt the payment of the selected combination, thereby completing a trial. No limit was imposed to the number of component that could be set on or off, nor to the time that subjects could spend in producing a combination before submitting it. However, since the number of possible combinations (1024) far outweighed the number of available trials (10 or 20), each choice entailed a high opportunity cost. This captured the constraints in available resources to be invested in the search task.

To support the decision-making process, the software displayed the complete history of submitted combinations and their relative payments. Furthermore, it also displayed the current trial and the task-specific cumulative payment up to that trial. Notably, participants were also constantly reminded about available time and, then, the number of available attempts by having the current trial followed by the relevant upper bound (i.e., “Trial x of 11” in the short-term condition, and “Trial x of 21” in the long-term condition). Lastly, in line with theoretical assumptions, subjects were not apprised of key parameters, such as task complexity, average and peak performance, or the performance of other participants.

As all subjects completed the second and last encounter, participants were required to fill a brief survey to collect socio-demographics information. The final compensation involved a show-up fee of €5 plus an additional sum contingent on participants’ performance across the two landscapes. In particular, the three subjects with the best final performance within a session received a prize: the first prize was €15, the second €10, and the third €5. The prizes were announced at the beginning of the session, thereby instituting a compensation scheme based on a rank-order tournament system to motivate competitive behavior. Indeed, rank-order compensation schemes are pervasive across many real-world settings in which competing agents are required to allocate scarce resources, such as effort, money, or time (e.g., Tullock, 1967; Krueger, 1974; Gibbons, 1998; Lazear, 1999; Prendergast, 1999). This is most prominent within the labor market where, based on the relative performance of agents, rank-order tournaments often regulate intra-organizational competition between lower-level employees (Lazear & Rosen, 1981) as well as inter-organizational competition between higher-level executives (Gibbons & Murphy, 1990). Compensation schemes based on rank-order tournaments have been given careful consideration in laboratory research, which suggested that average effort levels in tournaments are consistent with theoretical equilibrium levels, despite presenting higher variance (Bull et al., 1987). Notably, this result has been replicated in a large number of experiments (see Dechenaux et al., 2015 for a comprehensive review). Further laboratory research suggested that, relative to compensation

schemes based on the plain cumulation of individual performance, rank-order tournaments do not alter the risk-taking attitude of participants, while consenting to minimize transaction costs and to prevent the occurrence of wealth effects or cross-task contaminations (see Charness et al., 2016 for a comprehensive review).

Sessions lasted about 60 minutes, including instructions and questionnaire. The experimental software was developed and implemented using oTree (Chen et al., 2016).

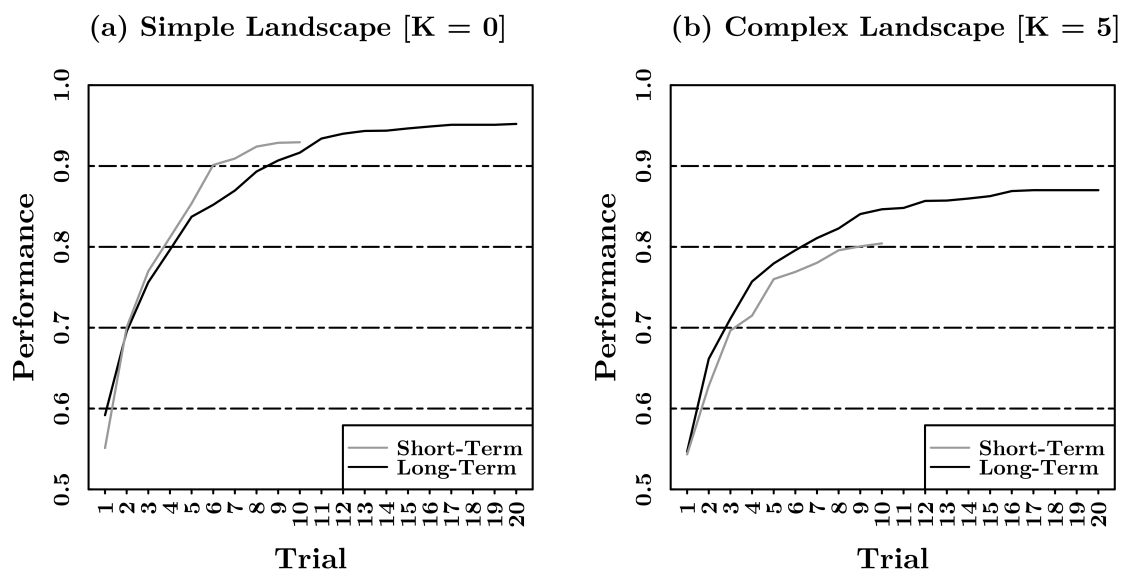
2.4 Results

2.4.1 Descriptive Statistics

We begin the report of the findings from the experiment presenting descriptive statistics concerning human performance and the underlying search behaviour, which are supplemented with hypothesis-based tests of significance. In particular, differences relative to within-subject conditions were assessed using the Wilcoxon signed-rank test, a non-parametric statistic commonly employed to compare two dependent samples. In contrast, differences relative to between-subjects conditions were assessed using the Mann-Whitney U test, a non-parametric statistic commonly employed to compare to independent samples.

Figure 2.1 displays the average performance of participants for the two levels of complexity in each of the two time horizons. To mirror the simulative results from the NK literature, where an agent adopts a new combination if and only if it provides a fitness improvement,

Figure 2.1 Performance




performance is measured based on the cumulative maximum, which is to say based on the best-performing combination identified by a subject up to that trial. Consistently with the temporal pattern suggested by simulations (Siggelkow & Levinthal, 2003), overall average performance increases in the number of trials, with marginal gains becoming progressively lower. More in detail, in the short-term horizon, average performance on the rugged, complex landscape ($M = 0.73$, $SD = 0.13$) is lower than average performance on the smooth, non-complex landscape ($M = 0.83$, $SD = 0.20$). Furthermore, in the long-term horizon, average performance on the rugged, complex landscape ($M = 0.81$, $SD = 0.12$) is lower than average performance on the smooth, non-complex landscape ($M = 0.88$, $SD = 0.16$). Two separate Wilcoxon tests suggested that differences in average performance are significant in both the short-term ($V = 51497$, $p < .001$) and the long-term condition ($V = 192291$, $p < .001$). As a matter of fact, 56.9% of participants reached the global maximum on the smooth landscape (33.3% in the long-term horizon, 26.6% in the short-term horizon), whereas nobody managed to reach it on the rugged landscape. Consistently with the last observation, on the smooth landscape, average performance in the long-term horizon is higher than average performance in the short-term horizon. Two separate Mann-Whitney U tests suggested that, indeed, differences in average performance are significant in both the smooth ($W = 111338$, $p < .001$) and the rugged condition ($W = 73080$, $p < .001$). However, looking at the brief period represented by the first 10 trials, *ceteris paribus* the short-term horizon is associated with higher average performance than the long-term horizon ($M = 0.81$, $SD = 0.19$) when the landscape is smooth, while it is associated with lower average performance than the long-term horizon ($M = 0.76$, $SD = 0.14$) when the landscape is rugged. In other terms, relative to the near future, the short-term horizon seems to provide an adaptive advantage if the task is simple, while it is detrimental if the task is complex. In this regard, two separate Mann-Whitney U tests suggested that, however, differences in average performance are significant on the smooth ($W = 86512$, $p < .001$) but not on the rugged landscape ($W = 86512$, $p = .078$).

Figure 2.2 displays the patterns of search behaviour for the two levels of complexity in each of the two time horizons. In keeping with Billinger et al. (2014), search behaviour is assessed based on the number of changed attributes compared with the best-performing combination identified by an individual in prior trials. The resulting measure is referred as search distance. It reflects the Hamming distance in information theory (Hamming, 1950), a metric to compare binary data of strings commonly employed in the organizational literature to model search behaviour in combinatorial tasks (Frenken, 2006; Levinthal, 1997; Rivkin, 2000). More precisely, search distance is measured based on discrete values ranging between 1 (change of one attribute) and 10 (change of all attributes). What immediately stands out is that average search distance on both the smooth ($M = 2.63$, $SD = 2.11$) and the rugged

landscape ($M = 2.53$, $SD = 2.13$) is higher than it would be with the pure incremental search strategy customarily assumed in computational models. The same applies to average search distance in the short-term horizon ($M = 2.85$, $SD = 2.23$) and the long-term horizon ($M = 1.64$, $SD = 2.03$). Specifically, on the smooth landscape, average search distance in the short-term horizon ($M = 2.86$, $SD = 2.26$) is higher than average search distance in the long-term horizon ($M = 2.42$, $SD = 2.06$). Furthermore, on the rugged landscape, average search distance in the short-term horizon ($M = 2.83$, $SD = 2.20$) is higher than average search distance in the long-term horizon ($M = 2.34$, $SD = 2.11$). Two separate Mann-Whitney tests indicated that differences in average search distance between the short-term and the long-term condition are significant on both the simple landscape ($W = 80321$, $p < .05$) and the complex landscape ($W = 83115$, $p < .001$). Nonetheless, two separate Wilcoxon tests suggested also that average search distance on the simple landscape is not significantly different than on the complex landscape neither in the short-term ($V = 9040$, $p > .1$) nor in the long-term condition ($V = 14122$, $p > .1$).

Overall, consistently with Billinger et al. (2014), subjects begin the search process resorting to distant search, thereby changing many attributes at once. Then, they gradually narrow down their search distance with the number of trials. Still in compliance with Billinger et al. (2014), the behavioral pattern appears to be different for the simple and the complex task. Specifically, in the long-term horizon search appears to broad again throughout the latter trials in the rugged landscape, but not in the smooth landscape. However, this is not

Figure 2.2 Search Distance



case in the short-term horizon, where search behaviour does not appear to present peculiar differences between the smooth and rugged landscape.

Table 2.1 displays the frequency distribution of search distances for the two levels of complexity in each of two time horizons, thereby providing additional insights into search behavior. Manifestly, local search represents the predominant strategy (38.9% in the short-term horizon, 50.2% in the long-term horizon, and 45.3% overall). Very distant search (changing 9 or 10 attributes) is quite infrequent and occurred mostly on the simple landscape in the short-term horizon, and on the complex landscape in the long-term horizon. Intermediate distant search (changing between 3 and 8 attributes) are more frequent in the short-term horizon than in the long-term horizon, although in both cases it occurs mainly on the simple landscape.


Table 2.1 Frequency Distribution of Search Distances

	Short-Term Horizon					
	K = 0		K = 5		Average	
	Percent	Cum. percent	Percent	Cum. percent	Percent	Cum. percent
1	39.0	39.0	38.7	38.7	38.9	38.9
2	18.8	57.9	21.2	59.9	20.0	58.9
3	11.6	69.5	9.9	69.9	10.8	69.7
4	9.6	79.1	7.9	77.8	8.8	78.5
5	9.2	88.4	8.9	86.8	9.1	87.5
6	3.1	91.4	6.0	92.7	4.5	92.1
7	2.7	94.2	3.0	95.7	2.9	94.9
8	2.4	96.6	1.7	97.4	2.0	97.0
9	0.7	97.3	0.3	97.7	0.5	97.5
10	2.7	100	2.3	100	2.5	100
Total	100	100	100	100	100	100

	Long-Term Horizon					
	K = 0		K = 5		Average	
	Percent	Cum. percent	Percent	Cum. percent	Percent	Cum. percent
1	47.3	47.3	53.3	53.3	50.2	50.2
2	18.0	65.3	18.4	71.7	18.2	68.4
3	9.0	74.3	8.2	79.9	8.6	77.0
4	10.2	84.4	6.3	86.3	8.3	85.3
5	6.6	91.0	5.5	91.8	6.1	91.4
6	4.2	95.2	1.9	93.7	3.1	94.5
7	1.6	96.8	1.3	94.9	1.4	95.9
8	1.2	98.0	1.9	96.8	1.5	97.4
9	0.4	98.4	0.8	97.7	0.6	98.0
10	1.6	100	2.3	100	2.0	100
Total	100	100	100	100	100	100

In a view of examining the presence of a temporal midpoint heuristic, we considered an alternative measure of search distance based on discrete values ranging between 0 (change of no attribute) and 10 (change of all attributes), rather than between 1 (change of one attribute) and 10 (change of all attributes). Notably, a search distance of 0 reflects the notion of pure exploitation, which captures the attempt to maximize one’s own cumulative payments through the reiteration of the best-performing combination identified in prior trials. Reintroducing a search distance of 0 consents then to inspect the trade-off between active search and the maintenance of the status quo, with the main objective to control whether the temporal pattern of exploration and exploitation obeys to the midpoint heuristic. Figure 2.3 displays the patterns that emerge from this alternative measure. Visibly, average search distance on both the smooth ($M = 1.93$, $SD = 2.15$) and the rugged landscape ($M = 1.81$, $SD = 2.15$), as well as average search distance in the short-term ($M = 2.35$, $SD = 2.30$) and the long-term horizon ($M = 1.64$, $SD = 2.03$), is still higher than it would be with a pure incremental search strategy. More thoroughly, on the smooth landscape, average search distance in the short-term horizon ($M = 2.32$, $SD = 2.33$) is higher than in the long-term horizon ($M = 1.73$, $SD = 2.03$). Furthermore, on the rugged landscape, average search distance in the short-term horizon ($M = 2.37$, $SD = 2.27$) is higher than in the long-term horizon ($M = 1.54$, $SD = 2.04$). Two separate Mann-Whitney tests indicated that differences in average search distance between the short-term and the long-term condition are significant on both the simple landscape ($W = 151715$, $p < .001$) and the complex landscape ($W = 164872$, $p < .001$). Moreover, two separate Wilcoxon tests suggested that differences in average search distance between the

Figure 2.3 Search Distance (Alternative Measure)



smooth and rugged landscape are not statistically significant in the short-term condition ($V = 15047$, $p. > .05$), although average search distance on the smooth landscape is significantly higher than on the rugged landscape in the long-term condition ($V = 45402$, $p. < 0.01$).

As before, subjects begin the search process resorting to distant search. Then, they gradually narrow down their search distance with the number of trials. However, in a departure from the previous scenario, search behaviour does not broad again in the latter search trials, but seems to steadily converge towards incremental search. Notably, this process is coherent with the exploration-exploitation pattern suggested by the standard models of reinforcement learning (Sutton & Barto, 1981, 1998; Kaelbling et al., 1996) and

Table 2.2 Frequency Distribution of Search Distances (Alternative Measure)

	Short-Term Horizon					
	K = 0		K = 5		Average	
	Percent	Cum. percent	Percent	Cum. percent	Percent	Cum. percent
0	18.9	18.9	16.1	16.1	17.5	17.5
1	31.7	50.6	32.5	48.6	32.1	49.6
2	15.3	65.8	17.8	66.4	16.5	66.1
3	9.4	75.3	8.3	74.7	8.9	75.0
4	7.8	83.1	6.7	81.4	7.2	82.2
5	7.5	90.6	7.5	88.9	7.5	89.7
6	2.5	93.1	5.0	93.9	3.8	93.5
7	2.2	95.3	2.5	96.4	2.4	95.8
8	1.9	97.2	1.4	97.8	1.7	97.5
9	0.6	97.8	0.3	98.1	0.4	97.9
10	2.2	100	1.9	100	2.1	100
Total	100	100	100	100	100	100

	Long-Term Horizon					
	K = 0		K = 5		Average	
	Percent	Cum. percent	Percent	Cum. percent	Percent	Cum. percent
0	30.4	30.4	34.3	34.3	32.4	32.4
1	32.9	63.3	35.0	69.3	34.0	66.3
2	12.5	75.8	12.1	81.4	12.3	78.6
3	6.2	82.1	5.4	86.8	5.8	84.4
4	7.1	89.2	4.2	91.0	5.6	90.1
5	4.6	93.8	3.6	94.6	4.1	94.2
6	2.9	96.7	1.2	95.8	2.1	96.2
7	1.1	97.8	0.8	96.7	1.0	97.2
8	0.8	98.6	1.2	97.9	1.0	98.3
9	0.3	98.9	0.6	98.5	0.4	98.7
10	1.1	100	1.5	100	1.3	100
Total	100	100	100	100	100	100

the experimental evidence on instance-based learning (Gonzalez et al., 2003; Gonzalez & Dutt, 2011).

Table 2.2 displays the frequency distribution of search distances for the two levels of complexity in each of the two time horizons. Again, local search clearly represents the predominant strategy (32.1% in the short-term horizon, 34.0% in the long-term horizon, and 33.3% overall), while pure exploitation (changing no attribute) is the second most frequent strategy (17.5% in the short-term horizon, 32.4% in the long-term horizon, and 27.4% overall). Very distant search (changing 9 or 10 attributes) is quite infrequent and occurred mostly in the long-term horizon independently from the level of complexity. Lastly, intermediate distant search (changing between 3 and 8 attributes) are considerably more frequent in the short-term horizon than in the long-term horizon, again independently from the level of complexity.

2.4.2 Regression Analyses

A regression analysis was conducted to assess the differential impact of time horizon on performance in the simple ($K = 0$) and the complex task ($K = 5$). In particular, since performance is a continuous variable measured based on the cumulative maximum up to that trial, we adopted the OLS model.

Table 2.4 displays the two OLS regressions employed to analyze the effect of time horizon on performance in the simple and the complex task, respectively. Both models control for the possible influence of the number of search trials and the position of the task in the session-specific sequence. As can be seen, Model 1 provides strong evidence in support

Table 2.3 Main Variables

Variable	Data	Min	Max	Mean	SD	Skewness	Kurtosis
Complexity	Any	0	5	-	-	-	-
Time horizon	Any	10	20	-	-	-	-
Task position	Any	1	2	-	-	-	-
Feedback	Overall	0	1	0.32	0.47	0.76	-1.43
Unsuccessful attempts	Overall	0	18	2.29	3.22	1.95	4.01
Prior search distance	Overall	0	10	2.24	2.2	1.49	1.99
Trial	Overall	1	20	-	-	-	-
Performance	Simple task	0.20	1	0.83	0.17	-1.45	1.61
Performance	Complex task	0.20	1	0.80	0.13	-1.05	0.96
Feedback	Short-term	0	1	0.38	0.48	0.46	-1.78
Trial	Short-term	1	10	-	-	-	-
End of the task proximity	Short-term	0	9	-	-	-	-
Feedback	Long-term	0	1	0.28	0.45	0.95	-1.09
Trial	Long-term	1	20	-	-	-	-
End of the task proximity	Long-term	0	19	-	-	-	-

Table 2.4 OLS Models with Performance as Dependent Variable

	Model 1 (K = 0)	Model 2 (K = 5)
Time horizon (Short)	0.032** (0.011)	-0.015† (0.008)
Trial	0.016*** (0.001)	0.014*** (0.001)
Task Position (2)	0.044*** (0.009)	0.051*** (0.006)
Constant	0.612 *** (0.017)	0.615*** (0.013)
R ²	0.258	0.381
Adjusted R ²	0.256	0.379
Residual Std. Error	0.149	0.107
F Statistic	124.512***	220.938***
Observations	1,080	1,080

Note: † p = 0.05; * p < 0.05; ** p < 0.01, *** p < 0.001

of the positive association between short-term horizon and performance in the simple task, while Model 2 provides weak - but clearly inconclusive - evidence in support of the negative association between short-term horizon and performance in the complex task ($p = 0.051$).

Another regression analysis was conducted to assess the determinants of search behaviour. The dependent variable at issue is the positive search distance of subjects from the first serviceable trial, which consists of discrete count data with values ranging between 1 and 10 conditional on the number of changed attributes relative to best-performing combination identified in prior trials. Thence, to assess search behaviour we adopted the Poisson model. Notably, since average search distance is non-normally distributed and positively skewed, we resorted to robust standard errors for the parameter estimates (Cameron & Trivedi, 2009), so as to control for the possible violation of the assumption that the variance should equal the mean. Moving from these premises, we ran four Poisson models. Model 1 assesses how the treatment conditions, namely complexity and time horizon, affect the search distance. Model 2-4 inspect the impact of individual fixed effects. In particular, Model 2 estimates the effect of performance feedback, while Model 3 and 4 evaluates the search history. The four models are structured to control for the possible influence of the number of search trials and the position of the task in the session-specific sequence.

Focusing on fixed effects, performance feedback is built as a binary variable, which takes value of 1 if in the last trial the subject managed to achieve a performance improvement relative to the best-performing combination identified in prior trials (success) and 0 otherwise (failure). The variable has first been proposed by Billinger et al. (2014) based on prior models of adaptive search (Levinthal & March, 1981) and individual risk attitudes (Kahneman &

Table 2.5 Poisson Models with Search Distance as Dependent Variable

	Model 1	Model 2	Model 3	Model 4
Complexity (K = 5)	-0.041 (0.041)	-0.067 (0.040)	-0.049 (0.040)	-0.043 (0.038)
Time horizon (Short)	0.099* (0.044)	0.100* (0.043)	0.101* (0.043)	0.041 (0.041)
Feedback (Positive)		-0.412*** (0.047)	-0.346*** (0.052)	-0.515*** (0.049)
Number of unsuccessful attempts			0.034*** (0.010)	0.012 (0.009)
Prior search distance				0.117*** (0.007)
Trial	-0.021*** (0.005)	-0.033*** (0.005)	-0.048*** (0.008)	-0.044*** (0.007)
Task position (2)	-0.064 (0.041)	-0.075 (0.040)	-0.082* (0.040)	-0.058 (0.038)
Constant	1.104*** (0.060)	1.319*** (0.063)	1.319*** (0.063)	1.318*** (0.071)
Deviance	2234.4	2104.5	2085.5	1767.4
Log Likelihood	-3,164.577	-3,099.614	-3,090.144	-2,931.101
Pseudo-R ²	0.0275	0.0841	0.0923	0.2307
Observations	1,568	1,568	1,568	1,568

Note: * p < 0.05; ** p < 0.01; *** p < 0.001

Tversky, 1979; March, 1988; March & Shapira, 1992) to capture the idea that individuals encode performance feedback as success or failure with respect to a subjective reference point (Markowitz, 1952; March, 1988; Bromiley, 1991). In particular, positive feedback (success) is expected to induce risk-averse exploitation, thereby reducing the search distance. Conversely, negative feedback (failure) is expected to induce risk-seeking exploration, thereby increasing the search distance. As regards the search history, it encompasses two variables. The first is the number of unsuccessful attempts, which counts the number of trials since the last improvement in performance. The second is the search distance in the prior trial, which has the function to test for path dependencies in search behavior.


Table 2.5 displays the four Poisson models of search behaviour. Model 1 suggests that complexity has no immediate effect on the search distance, while a short-term horizon appears to have a positive impact. Model 2 suggests that positive feedback reduces the search distance, thereby inducing subjects to exploit the gathered knowledge focusing on the region of the landscape where they experienced a performance increase. Model 3 suggests that the number of consecutive unsuccessful attempts is positively associated with search distance, hinting that subjects tend to become more impatient with the systematic failure to obtain performance improvements. Model 4 suggests that the prior search distance is positively

related with the search distance, thereby providing evidence in support of the presence of strong path dependence. Notably, in this model the effect of prior search distance seems to absorb that of time horizon and the number of unsuccessful attempts. Nonetheless, in all models the search distance tends to decrease with the number of trials, while the task position has no significant effect.

Taken together, the results suggest that search behaviour responds to the time horizon and gradually adapts to performance feedback. In particular, a shorter time horizon or temporal orientation and, thence, scarcer temporal resources promote greater exploration. Success in obtaining performance improvements induces more exploitation, while failure induces more exploration. Search behaviour responds also to the current interval without performance improvements and exhibits a strong path-dependence, such that a higher search distance in the prior trial brings about a higher search distance in the current trial. Last but not least, the search distance decreases with time, possibly as a result of the increasing difficulty to obtain performance improvement and the concurrent effort to maximize the sum of one's own rewards by resorting to more exploitation. Since the variance of the dependent variable is higher than the mean, we conducted a robustness check for overdispersion of the data by running a negative binomial regression obtaining the same qualitative results. Nonetheless, several pseudo- R^2 measures suggested that the goodness of the fit for the Poisson models was superior to that of the negative binomial models.

Figure 2.4 provides a compact summary of the effects contained in the Poisson regressions, displaying how prior search distance and performance feedback affect search behaviour in the current trial. Supplemented with the relationship between complexity and performance feedback depicted in Table 2.6, it outlines the basic tenets behind the behavioral model of adaptive search (Billinger et al., 2014), understood as a solid, stylized explanation of human search behaviour in diversely complex, combinatorial tasks. The main argument is that complexity does not have a direct effect on search distance. Rather, it affects feedback conditions in such a way that the rate of success decreases with the level of complexity. Thus, for human subjects gradually adapt to performance feedback, complexity has nonetheless an indirect impact on human search behaviour. In particular, subjects adapt to positive feedback by resorting to more incremental search and to negative feedback by gradually broadening their search distance. Since identifying better-performing combinations becomes more difficult as trials go by, this implies that subjects tend to become impatient with the lack of performance improvements, moving away from local search well before a widespread examination of the immediate neighborhood has been achieved.

To ensure that task complexity impacted performance feedback, we ran an additional regression analysis with feedback as the dependent variable. More precisely, since the

Figure 2.4 Adaptive Search

response is binary, we adopted the probit model. Table 2.6 displays the two probit regressions employed to analyze the effect of complexity on performance feedback in the short-term and long-term horizon, respectively. The evidence suggests that the likelihood of a positive feedback decreases with the level of complexity in the long-term but not in the short-term condition. Notwithstanding, in both cases the rate of success tends to decrease with the number of trials.

Table 2.6 Probit Models with Feedback as Dependent Variable

	Model 1 (Short-term)	Model 2 (Long-term)
Complexity (K = 5)	-0.198 (0.106)	-0.207* (0.090)
Trial	-0.093*** (0.019)	-0.104*** (0.010)
Task position (2)	-0.115 (0.106)	-0.045 (0.090)
Constant	0.431* (0.309)	0.378* (0.167)
Deviance	765.57	1033.9
Log Likelihood	-382.786	-516.948
Observations	594	974

Note: * $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$

A supplementary regression analysis was conducted to assess whether the temporal patterns of search behaviour comply with the midpoint heuristic (Gersick, 1988, 1989). The dependent variable corresponds to the alternative measure of search distance exhibited by searchers from the first serviceable trial that, conditional on the number of changed attributes relative to the best-performing combination identified in prior trials, consists of discrete count data with values ranging between 0 and 10, rather than 1 and 10. As mentioned, a search distance of 0 reflects the notion of pure exploitation, which is to say the attempt to maximize one's own cumulative payments through the reiteration of the best-performing combination assumed as reference point. As before, since the dependent variable is made out of non-normally distributed, positively skewed count data, we adopted a Poisson model with robust standard errors for parameter estimates (Cameron & Trivedi, 2009). Furthermore, because we were interested in determining whether the temporal midpoint acted as a breakpoint and, especially, whether the slope of the trends before and after the temporal midpoint were somehow different, we fitted a stepwise Poisson regression, whose coefficients were parametrized to represent the slope of the trend to the left and the right of the breakpoint located at the middle of the task. Lastly, we ruled out observations from the first serviceable trial. There reason is that the very first attempt is a form of random rather than directed exploration that, as such, is typically associated with very high distant search. Thence, maintaining the first serviceable trial within the data would have unduly increased the steepness of the directed exploration trend before the breakpoint. Table 2.7 displays the two stepwise Poisson models employed to analyze the midpoint heuristic in the short-term and long-term horizon, respectively. As can be seen, in the short-term horizon, the trend before the temporal midpoint is negative but not significant, while the trend after the temporal midpoint is negative and significant. Furthermore, in terms of logs of expected counts, the slope of the trend after the midpoint is steeper than before the midpoint. Instead, in the long-term horizon, the trends before and after the temporal midpoint are negative and significant. Moreover, the slope of the trend after the midpoint is steeper than before the midpoint. Nonetheless, two separate Wald-test-based comparisons of the two trends indicated that the change in the slope before and after the midpoint is not statistically significant neither in the short-term horizon (chi-square = 1.14, $p > .05$) nor in the long-term horizon (chi-square = 1.25, $p > .05$). The same results were obtained fitting a stepwise Poisson regression whose coefficients were parametrized to represent, respectively, the slope of the trend before the breakpoint and the change in the slope after the breakpoint. Further robustness checks indicated also that no breakpoint was present in the neighborhood of the temporal midpoint. Over and above the midpoint heuristic, the stepwise models offer complementary details concerning the impact of complexity and the role of the task position. In particular, they suggest that,

Table 2.7 Stepwise Poisson Models with Search Distance as Dependent Variable to Analyze the Temporal Midpoint Heuristic

	Model 1 (Short-term)	Model 2 (Long-term)
Complexity (K = 5)	0.039 (0.054)	-0.132* (0.045)
Slope before midpoint	-0.066 (0.029)	-0.054*** (0.009)
Slop after midpoint	-0.112*** (0.020)	-0.072*** (0.009)
Task position (2)	-0.213** (0.054)	-0.325*** (0.045)
Constant	0.954*** (0.063)	0.649*** (0.052)
Residuals	1189.4	2479.3
Log Likelihood	-1,291.097	-2,369.050
Pseudo-R ²	0.0763	0.1064
Observations	648	1,368

Note: * p <0.05; ** p<0.01; *** p<0.001

as hinted by preliminary tests of significance, the alternative measure of search distance is negatively associated with the level of complexity in the long-term but not in the short-term horizon. Moreover, contrary to the standard measure, the alternative measure of search distance tends to decrease with the task position in both the short-term and the long-term horizon. Taken in conjunction, these findings may suggest that subjects learn to resort more extensively to exploitation through the reiteration of one's own best-performing combination in a bid to maximize their cumulative payment, especially when the task is more complex and improvements more difficult to obtain.

One more regression analysis was conducted to assess the moderating effects of deadline proximity. To do that, we reprised the standard measure of search distance as dependent variable, which consists of discrete count data with values ranging between 1 and 10 conditional on the number of changed attributes relative to the best-performing combination identified in prior trials. As previously mentioned, the relevant variable is non-normally distributed and positively skewed. Therefore, we resorted again to a Poisson model with robust standard errors for parameter estimates (Cameron & Trivedi, 2009). As before, feedback conditions signaling whether performance is above or below the reference point are captured using a binary variable, which takes value of 1 if in the last trial the subject managed to achieve a performance improvement relative to the best-performing combination identified in prior trials (success), and 0 otherwise (failure). Accordingly, the subjects that systematically

Table 2.8 Poisson Models with Search Distance as Dependent Variable to Analyze the Moderating Effects of Deadline Proximity

	Model 1 (Short-term)	Model 2 (Long-term)
Complexity (K = 5)	-0.043 (0.061)	-0.088 (0.054)
Feedback (Negative)	0.860*** (0.171)	0.905*** (0.260)
End of task proximity	0.121*** (0.025)	0.058*** (0.016)
Feedback (Negative) : Proximity	-0.064* (0.029)	-0.038* (0.018)
Task position (2)	-0.079 (0.061)	-0.076 (0.054)
Constant	0.193 (0.178)	-0.015 (0.257)
Deviance	776.54	1292.5
Log Likelihood	-1192.039	-1889.84
Pseudo-R ²	0.1405	0.0552
Observations	594	974

Note: * p < 0.05; ** p < 0.01; *** p < 0.001

experienced a negative feedback are classified as underperformers, while those that systematically experienced a positive feedback are classified as outperformers. Thereafter, following Lehman et al. (2011), deadline proximity is defined as the number of trials that remain to be done at the time of the current trial. More in detail, this is coded as the difference between the last and the current trial, which is equal to $10 - t$ in the short-term horizon, and $20 - t$ in the long-term horizon, where t represents the current trial. Ultimately, the moderating effect of deadline proximity on the performance-risk relationships and, thence, on search behaviour can be tested creating an interaction term between deadline proximity and performance feedback. Table 2.8 displays the two Poisson models employed to analyze the effects of interest in the short-term and long-term horizon, respectively. In both cases, we set success as baseline group for the feedback variable, thereby showing the effect of performing below the reference point. As evident, the interaction terms between performance feedback and deadline proximity are significant and go in the expected direction. Specifically, they suggest that, as the deadline of the short-term or the long-term investment horizon approximates, underperforming subjects focus their attention on survival, thereby opting for risk-averse exploitation by lowering their search distance, while outperforming subjects focus their attention on slack resources, thereby opting for risk-seeking exploration by increasing their

search distance. Remarkably, a robustness check centering on the integration of the search history highlighted that, as in the case of the effect of the short-term horizon in the Poisson models of search behaviour, the prior search distance seems to absorb the effect of feedback, deadline proximity and the resulting interaction in the long-term but not in the short-term horizon. A feasible reason may be that, given enough time, path dependence becomes so strong to act as a potential negative confounder for weaker but still present effects.

2.5 Discussion

Consistently with theoretical predictions and prior empirical evidence, the findings from the experiment suggest that average performance in the smooth, non-complex landscape is significantly higher than the rugged, complex landscape. Still in line with the prior empirical evidence, the findings suggest also that the level of complexity does not have a direct effect on search behaviour. Besides, human agents gradually adapt to feedback conditions, so that success in discovering better-performing combinations induces more exploitation, whereas failure induces more exploration. Moreover, in the long-term horizon, positive feedback significantly decreases with the level of complexity. Since the level of complexity influences feedback conditions and, in turn, feedback conditions impact search behaviour, complexity can be regarded as having an indirect effect on search behaviour. However, in the short-term horizon, the negative association between complexity and positive feedback is not statistically significant. Since the rate of positive feedback decreases with time, this may be due to the fact that, in the first 10 trials of the task, subjects are still in the position to frequently discover better-performing combinations.

The principal contributions of this experiment arises from the investigation of the effects exerted by time horizon on human performance and the underlying search behaviour. In this regard, the findings suggest that the short-term horizon is associated with lower average performance than the long-term horizon both at the end of the task and overall. However, restricting the focus of attention to the brief period represented by the first 10 trials, *ceteris paribus* the short-term horizon is associated with higher average performance than the long-term horizon when the task is simple, whereas it is associated with lower average performance when the task is complex. That being the case, the empirical evidence supports H1, while also introducing more subtle nuances. Properly speaking, the detrimental effect of the short-term horizon on final and overall performance in the simple and complex task is in line with predictions. So it is the near future drawback relative to the long-term horizon in the presence of complexity. However, still with respect to the long-term horizon, the short-term horizon appears to produce unexpected near-future benefits that are in contrast

with standard assumptions. One possible explanation lies in the time horizon effect on the underlying search behavior. On this matter, the experimental findings suggest that, although average search distance in the simple task is not significantly different than in the complex tasks neither in the short-term nor the long-term condition, average search distance in the short-term horizon is significantly higher than in the long-term horizon. Otherwise stated, independently from the level of complexity, in our combinatorial task, the short-term horizon appears to be positively associated with search distance, thereby leading to the rejection of H2. From a cognitive standpoint, it may then be argued that a shorter time frame introduces a form of pressure in the investment of scarce resources that induces human subjects to take more risks and explore the problem space on a larger scale than they would do with a longer time frame, in which case exploration can be spread over a lengthier span. Notably, although this is somewhat beneficial in terms of performance when subjects cope with a simple task, the same attitude tends to be particularly detrimental when the task to be carried out is complex. A further inspection of the pattern of search behaviour based on the reintroduction of pure exploitation suggests also that the trend after the midpoint of the task is steeper than the trend before it. However, the change in the slope of the two trends turns out to be not significant, thereby failing to provide support for H3. A plausible interpretation for the trend after the midpoint being steeper is that, rather than responding to the midpoint of the task, the transition from exploration to pure exploitation emerges as a progressive process occurring along the whole time frame of the task. Notably, this phenomenon is coherent with the exploration-exploitation pattern suggested by the standard models of reinforcement learning and the experimental evidence on instance-based learning. Last but not least, the assessment of the moderating effects of deadline proximity on search behaviour suggests that underperforming subjects tend to shift the focus of attention from the reference point represented by the best-performing combination identified in prior trials to survival, thereby opting for risk-averse exploitation, while outperforming subjects are prone to shift the focus of attention from the same reference point to experimenting with slack resources, thereby opting for risk-seeking exploration. Correspondingly, the empirical evidence provides support for H4. When all is said and done, a final consideration should also be made on the impact of search history, especially when it comes to path dependence as expressed by the prior search distance. In this regard, some models pointed out how, in the long-term horizon, the prior search distance tends to absorb the effect of a number of other factors. On this account, it can be mooted that, given enough time, path dependence becomes so strong to act as a potential negative confounder for weaker but still present effects.

The experimental study has several limitations. In the first place, the negative association between short-term horizon and performance in the complex task is at the borderline of

conventional significance ($p = 0.051$). In a narrow sense, the result is inconclusive and cannot be safely received. However, since it is actually expected on the basis of the previous literature (e.g., White & Frederiksen, 1986; Clancey, 1987; Winter, 2000), we believe that the result in point cannot be readily dismissed as insignificant without facing a substantial risk of type-II error and, for this reason, are ultimately inclined to grant the weak evidence in its favor. Secondly, the study does not address the problem of uncertainty about the future. In the real world, in fact, agents are often not aware of the horizon constraining a choice sequence, in the sense that they do not know how often they will carry out a similar choice. Such a “shadow of the future” (Dal Bó, 2005) makes then extremely difficult for them weighting the costs and benefits of exploration and exploitation (Mehlhorn et al., 2015), thereby introducing a further complication to the process of resource allocation (Yaari, 1965; Hakansson, 1969, 1971; Merton, 1971; Richard, 1975; Karatzas & Wang, 2001). In this regard, a preliminary experimental manipulation has been attempted in the context of the present study based on the traditional termination rule. However, the manipulation led to unbalanced samples and unreliable data, thereby indicating the needful to consider a different experimental strategy. An additional problem left to further research is related to the moderating effect of deadline proximity in the presence of vicarious information, which would allow for the formation of an alternative reference point based on the performance of other subjects (e.g., Knudsen, 2008) rather than on one’s own history, thereby influencing search behaviour (e.g., Greve, 2003).

2.6 Conclusions

We carried out a laboratory experiment in which investment horizons of different length were employed to induce a near or distant future temporal orientation in combinatorial tasks crafted using the NK model and mirroring organizational settings characterized by multiple dimensions and various levels of complexity. This is relevant for organizational science because the passage of time is essential to articulate many strategic practices, such as assessing progress, scheduling and coordinating task-related activities, discerning the processual dynamics of how these activities emerge, develop, and terminate, or interpreting retrospectively, current, and anticipated events. A greater or lesser amount of time reflects then a greater or lesser provision of resources, thereby representing a constraint that can greatly affect the ability to maintain a competitive advantage or ensure organizational survival. Also, as organizations aim to acquire assets or develop capabilities that yield future benefits exceeding upfront costs, it may well be claimed that the very notion of investment relies on time. In the light of these considerations, adopting a temporal lens is crucial to deepen our

understanding of organizational phenomena (e.g., Ancona et al., 2001a, 2001b; Bluedorn 2002; George & Jones 2000; McGrath & Rotchford 1983; Roe 2008; Sonnentag, 2012; Shipp & Cole, 2015; Reilly et al., 2016; Kunisch et al., 2017).

Our findings provide empirically grounded insights that are instrumental in moving a first step towards the integration of the temporal dimension into the model of adaptive search (Billinger et al., 2014). Of particular importance is the shift of attention from performance to either survival or experimenting with slack resources, which results from the moderating effects of deadline proximity on the performance-risk relationship. The model of adaptive search describes how individuals encode performance feedback as success or failure with respect to a subjective reference point based on prior performance and adapt their search behavior accordingly, capturing the fundamental pattern of success-induced exploitation and failure-induced exploration. It has the foremost merit to synthesize kindred ideas previously scattered across a plurality of research traditions, such as the risk dynamics conditional on a reference point in prospect theory (Bromiley, 2010; Kahneman & Tversky, 1979), the attitudes contingent on a subjective aspiration level in the work on managerial risk-taking (Koop & Johnson, 2012; March, 1988; March & Shapira, 1992), and a similar relationship between aspirations and search expressed in the context of the behavioral theory of the firm (Cyert & March, 1963; Greve, 2003; Levinthal, 1997, Levinthal & March, 1981). Following Lehman et al. (2011), the findings from our experiment extend the model of adaptive search showing that time represents a unique type of resource able to trigger shifts in the focus of attention. In particular, when more time is available, attention is likely to be directed at attaining and maintaining an aspired level of performance. Conversely, as less time is available and the deadline draws nearer, attention is likely to be redirected at ensuring survival in underperforming subjects or at experimenting with slack resources in outperforming subjects. The rationale is that, besides encoding performance feedback as success or failure with respect to a reference point, subjects form a “cognitive image” of the future, which tends to be perceived as less uncertain as the end of the period approximates. Thence, underperformers - that is, subjects that systematically experienced a negative feedback - are more likely to visualize their end-of-period performance as underperforming, while outperformers - that is, subjects that systematically experienced a positive feedback - are more likely to visualize their end-of-period performance as outperforming. The subsequent performance-risk relationship is then U-shaped for the former and linear for the latter. Accordingly, as the end of the task approximates, underperformers are likely to redirect their attention from maintaining the aspired level of performance to ensuring survival, thereby opting for risk-averse exploitation, while outperformers are likely to redirect their attention from the same reference point to experimenting with slack resources, thereby opting for risk-seeking exploration. The

present experiment offers the first validation of this phenomenon under laboratory conditions. Consistently with the prior work of Arthur (1991), Edmonds (2001), and Billinger et al. (2014), the relevant findings have further spillovers for theoretical work in organizational research, in that they provide evidence in support of the development of empirically informed search algorithms to simulate stylized behavioral rules within the computational modelling of adaptive processes.

In closing, we hope that our investigation will contribute to calibrate a temporal lens for organizational science and improve our comprehension of the psychological foundations of individual decision-making in settings characterized by complexity and persistent uncertainty, in a view to advance the behavioral strategy agenda aiming to ground human behaviour in more realistic assumptions and, thereby, to promote strategy theory and empirical research.

3

Parallel Problem-Solving in Networked Groups

With Marco Faillo

Abstract

This chapter presents findings from a laboratory experiment in which the accuracy of the imitative process is manipulated to induce a flawless or flawed information diffusion system and, congruently, an efficient or inefficient communication network in combinatorial tasks with multiple dimensions and various levels of complexity shaped using the NK model. We find that networks that are more efficient at disseminating information quickly converge on a single solution and perform well over the short-run but poorly over the long-run. Conversely, networks that are less efficient at disseminating information maintain strategic diversity and perform poorly over the short-run but well over the long-run. However, the performance benefits produced by perfect as well as imperfect imitation decrease with the level of complexity. Furthermore, imitative behavior adapts to social feedback, so that subjects are inclined to follow the members of the reference group who are more successful, while they prefer to rely on their own search ability when there is nobody who outperforms them. Remarkably, the harmonization of targeted exploration through the imitation of more successful others with fine-grained exploitation through autonomous adaptive search tends to improve as subjects gain experience. Despite the potential benefits from crossover effects, within inefficient networks the attitude to imitate more successful others tends to decline as with the strategic distance and, thence, with the risk of an error intervening in the imitative process. The implications of our findings for decision-making in organizational settings are discussed in the light of behavioral strategy research.

3.1 Introduction

The world we live in is becoming increasingly connected for the everyday person as well as for groups and organizations, with more distant linkages rapidly broadcasting information from one corner of the globe to another. In such a reduced world, services that allow distant actors to access knowledge and technologies one from each other have proliferated, promoting to various extents collaboration in the solution of problems. This is especially gainful when the relevant actors parallelly confront similar problems that are sufficiently complex to elude more or less hasty individual solutions. Paradigmatic cases are organizational activities, which are structured so that a strategy usually comprises a plurality of decision variables, such as extent of vertical integration, product design, pricing policy, channels of distribution, research and development program, and more (Porter, 1991; Ghemawat & Levinthal, 2008). The challenge of gaining a competitive advantage or ensuring the survival of the organization may then be conceived as a search for successful combinations of these decision variables (Rivkin, 2000). By definition, the problem of identifying a successful strategy becomes complex if the value of a combination is dependent on the interaction between the relevant decision variables (Levinthal, 1997; Page, 1996; Simon, 1962).

In this scenario, imitating the practices of high-performing actors is indeed one of the key strategies employed by organizations to solve problems and improve their performance (Nutt, 1998; Argote, 1999; Csaszar & Siggelkow, 2010), thereby representing a major part of the competitive process (Lieberman & Asaba, 2006). Accordingly, it has found considerable attention in the theoretical as well as in the practitioner-oriented literature. Imitative behaviour may assume many forms. Organizations may invest ample resources to detect and imitate best practices, engage consultants or experts to gain access to ideas and practices that have worked well in other organizations, adhere to trade associations to share information, enter business incubators to get in contact with well-connected venture capitalists with the prospect to gain also access to better practices used by others. Moreover, they invest resources in capabilities that may put them in the condition to imitate others more quickly and more extensively (Csaszar & Siggelkow, 2010).

The prior corpus of theoretical research has widely investigated how the structure behind the communication network of an organization can affect the balance between autonomous search and the social diffusion of information through imitation, thereby causing the levels of individual and collective performance to vary considerably (e.g., Rogers, 1988; Watts & Strogatz, 1998; Watts, 2002; Bettencourt, 2003; Schilling & Phelps, 2006; Lazer & Friedman, 2007; Mason et al., 2008; Fang et al., 2010; Mason & Watts, 2012; Derex & Boyd, 2016; Barkoczi et al., 2016). Furthermore, behavioral scientists have conducted a number of laboratory experiments to submit the relevant mathematical predictions to an empirical test,

finding overall support to their benchmark model. However, thus far experimentation has been limited to the use of two-dimensional functions (Mason et al., 2008; Mesoudi, 2008; Mason & Watts, 2012) or simple combinatorial tasks involving pair-wise interactions in single-peaked problems (Wisdom et al., 2013). No empirical test has been conducted to assess parallel problem-solving in a task characterized by a combinatorial choice structure with multiple decision variables encompassing varying degree of interactions, which is particularly suitable to capture complex and, thence, tendentially multi-peaked organizational issues that exhibit the above-mentioned structure (Schumpeter, 1934; Simon, 1962; Page, 1996; Levinthal, 1997).

The present research aims to address this shortcoming manipulating the accuracy of the imitative process under laboratory conditions to induce a flawless or flawed information diffusion system and, congruently, an efficient or inefficient communication network in combinatorial tasks featuring multiple dimensions and various levels of complexity. The tasks were created using the NK model, which provides an algorithm to generate a combinatorial problem space characterized by a more or less rugged performance landscape. Parameter N represents the number of decision variables, while the parameter K regulates the interactions among them. The hallmark of the model is that, as K increases, the ruggedness of a performance landscape changes from a single-peaked landscape to a multi-peaked landscape (Kauffman & Weinberger, 1989; Kaufmann, 1993). Due to these properties, the NK model has established itself as a canonical approach to analyze behavioral rules across a wide range of complex, strategic organization settings.

Against this backdrop, the present experimental investigation had three objectives: (1) evaluating the degree of strategic diversity of networked groups and their subsequent performance conditional on the interaction between the levels of complexity and the modalities of parallel problem-solving; (2) appraising whether prior results on autonomous search behavior under independent problem-solving extends to autonomous search behavior under parallel problem-solving; (3) shedding light on the determinants of imitative behavior, with a special attention to crossover effects, social feedback, and strategic distance.

There are several reasons for employing laboratory experiments to investigate human behaviour in complex, combinatorial tasks. First, a laboratory experiment provides full control over the environment, which consents to dismiss extraneous influences that may produce more or less acute manifestations of endogeneity. Second, in the absence of empirical data from the field, a laboratory experiment provides a convenient device to evaluate the empirical plausibility of behavioral rules and, thence, to test for model validity (Hey, 1982; Lave & March, 1975; Winter, 1982; Sterman, 1987; Lam, 2010). Third, a laboratory experiment provides results that may promote theoretical refinements and encourage further

empirical tests (Arthur, 1991). In brief, the problem of adaptive behaviour as manifested by real people requires an empirical foundation, and laboratory experiments are a primary source of empirical evidence (Selten, 1995). Undoubtedly, data from the field are also important, but they are more difficult to gain and harder to analyze (Selten, 1998).

The main empirical findings suggested that perfect imitation - or better to say, the flawless information diffusion system that characterizes an efficient network - accelerates the selection of better-performing strategies, resulting in a keen strategic convergence on a single solution and, thence, the leveling of performance in networked groups. Nonetheless, it promotes a near-term adaptive advantage characterized by a higher average performance over the short-run. In contrast, imperfect imitation - or better to say, the flawed information diffusion system that characterizes an inefficient network - delays the propagation of better-performing combinations, leading to a reduction in average performance over the short-run. However, the increased degree of strategic diversity engendered by the crossover mechanism consents to explore the space of the problem more extensively, thereby preserving strategic diversity and promoting a long-term adaptive advantage that translates into a higher average performance over the long-run. Still, the performance benefits produced by perfect as well as imperfect imitation decrease with the level of complexity. The findings suggested also that imitative behavior adapts to social performance feedback. In particular, subjects are inclined to follow the members of the reference group who are more successful, while they prefer to rely on their own search ability when there is nobody who outperforms them. Remarkably, the harmonization of "wise long-jumps" through the imitation of more successful others with the incremental enhancements of performance through autonomous adaptive search tends to improve as subjects gain experience. Despite the potential benefits from crossover effects, within inefficient networks the attitude to imitate more successful others declines as with the strategic distance and, thence, with the risk of an error intervening in the imitative process.

The remainder of the chapter is organized as follows: section 3.2 reviews the related literature and introduces our main research hypotheses; section 3.3 illustrates the experimental design and procedures; section 3.4 reports our key results; section 3.5 discusses their implications for behavioral strategy; section 3.6 concludes.

3.2 Theory and Hypotheses

In what follows, we review prior organizational research on complex strategic activities, parallel problem-solving, organizations and imitative behavior, and the cognitive foundations of imitative behavior. Contextually, we proceed to outline our four primary hypotheses.

3.2.1 Complex Problems as Rugged Landscapes

A problem is assumed to be complex if the possible solutions involve the combination of multiple decision variables that are interdependent with respect to their impact on performance (Simon, 1962; Page, 1996; Levinthal, 1997). Much of the recent theoretical work in organizational science has drawn upon the NK model to produce a stylized representation of complex problems. Although primarily devised in the field of evolutionary biology (Kauffman & Weinberger, 1989; Kaufmann, 1993), the NK model has established itself as a compelling apparatus to probe organizational adaptation processes (Levinthal, 1997) across a number of complex, strategic organization areas, such as individual decision-making (e.g., Gavetti & Levinthal, 2000), organizational decision-making (e.g., Knudsen & Levinthal, 2007), new product development (e.g., Mihm et al., 2003), product modularity (e.g., Marengo & Dosi, 2005), open innovation (e.g., Almirall & Casadesus-Masanell, 2010), organizational design (e.g., Rivkin & Siggelkow, 2003), industry dynamics (e.g., Lenox et al., 2007), and more (see Ganco & Hoetker, 2009 for a comprehensive review). The insights from the NK model have also informed empirical research on technology and product development (Fleming & Sorenson, 2004; Frenken, 2006), firm boundaries (Sorenson, 2003), the management of start-up companies (Sommer et al., 2009), and industry profitability (Lenox et al., 2010).

The NK model delivers an algorithm to generate performance landscapes based on N attributes and K interactions between these attributes. Correspondingly, a problem can be expressed as a space of alternatives where each alternative consists of N - in its basic version binary - attributes. Since each attribute admits only two possible values, the landscape encompasses 2^N alternatives. However, depending on the value assigned to the K parameter, the contribution of each attribute may be influenced by the interactions with other attributes. Specifically, if K is set equal to 0, the attributes are fully independent, and their contribution depends uniquely on their own value. At the opposite end of the spectrum, if K is set equal to $N-1$, attributes are fully interdependent, which implies that their contribution depends on their own value as well as on the value of all other attributes. By varying the degree of interdependence among the attributes, K influences the complexity of the landscape. More In detail, when attributes are fully independent, the landscape is smooth and features precisely one peak. However, as K rises, the number of local peaks grows, thereby making the landscape progressively more rugged.

Adaptation on rugged landscapes is then assumed to develop throughout two fundamental processes. The first is local search, in line to which the agent examines the proximate neighborhood by changing one attribute at the time. This process - also referred as one bit mutation or incremental search - reflects political and organizational routines in which time and budget constraints bind the search activity to local improvements. The second is global

search, in line to which the agent undertakes “long jumps” to inspect more distant alternatives. This process - also referred as random, distant, or non-incremental search - reflects a different sort of routine in which the systematic failure in obtaining marginal improvements induces more radical reorientations. The major implication is that, in non-complex - that is, single-peaked - landscapes, local search consents to reach the global maximum. However, in complex - that is, multi-peaked - landscapes, local search tends to get trapped in local optima, prompting a decrease in average performance.

Billinger et al. (2014) have recently proposed an experimental framework to study individual search behaviour on rugged landscapes, thereby laying the foundations for laboratory implementations of the NK model. The findings from the experiment indicated that search behavior in the NK search task gradually adapts to individual performance feedback, encoded as success or failure in discovering a better-performing combination relative to a subjective reference point coinciding with the best-performing combination identified in prior trials. More precisely, a positive feedback (success) prompts subjects to exploit the region of the landscape where they have experienced the performance improvement, whereas a negative feedback (failure) prompts them to explore more distant regions. The evidence suggested also that complexity does not have an immediate effect on search behavior. Rather, it affects the feedback conditions behind success-induced exploitation and failure-induced exploration. However, since performance feedback guides search behaviour, complexity can still be regarded as having an indirect effect on search behaviour.

Besides aiming at replicating these results in our experimental baseline, which reproduces the framework of autonomous search under independent problem-solving from Billinger et al. (2014), in the present investigation we are interested in assessing whether their explanatory power further extends to autonomous search behaviour under parallel problem-solving.

3.2.2 Parallel Problem-Solving

Under parallel problem-solving, a set of subjects pursue the solution of the same problem, so that the performance of any one or subset of them has no direct effect on the performance of the others (Lazer & Friedman, 2007). What is of primary importance in parallel problem-solving, however, is not how subjects solve problems on their own, but rather how they solve problems collectively. De facto, even though the performance of each subject has no direct effect on the performance of others, the parallel setup of problem-solving is such that she indirectly influences collective performance through a network of peers, in which the solutions produced by any of them offer to the others valuable information on the best practices within the problem space. Otherwise stated, each subject can examine what the other members of the group are doing and how they are performing. If someone produced a

different solution that provides a performance superior to her, the subject will most likely replicate it. In the last instance, a similar convention generates a flow of information among subjects that affects the systemic performance of the network. This is particularly convenient in complex problems characterized by multiple peaks. In fact, while any incremental search strategy would force a subject to move through solutions with worse outcomes before ascending again to solutions superior to her status quo, the flow of information on the best available practices consents to make targeted non-incremental changes - or so to say, wise “long-jumps” - in order to improve one’s own performance without undergoing inferior solutions.

Ultimately, while independent problem-solving relies uniquely upon individual learning, so that subjects search the problem space as isolated monades and gather information on it autonomously, parallel problem-solving can be regarded as a blend of individual and social learning, whereby autonomous search is integrated with the spread of social information through imitation (Tarde, 1903; Miller & Dollard, 1941). In line with the computational work of March (1991) on organizational learning, both independent and parallel problem-solving are concerned with striking a balance between exploration and exploitation, which is to say between the attempt to introduce new information and the disposition to leverage existing knowledge for some productive end (Gupta et al., 2006). Nonetheless, under parallel problem-solving, the trade-off between exploration and exploitation acquires a wider scope. In the case of independent problem-solving, in fact, exploration relates to more distant search, reflecting the ability to probe uncharted regions within the problem space, while exploitation relates to more local search, capturing the ability to take advantage of the region within the problem space where the subject experienced a performance improvement (March, 1991; Levinthal, 1997; Gavetti & Levinthal, 2000; Siggelkow & Rivkin, 2006). Under parallel problem-solving, in contrast, imitation can be seen as targeted exploration of more distant region within the landscape, while the subsequent local searches can be interpreted as more fine-grained, incremental adjustments to further enhance performance (Csaszar & Siggelkow, 2010).

As evident, the idea of parallel problem-solving is intrinsically tied to that of information diffusion (e.g., DiMaggio & Powell, 1983; Granovetter, 1978; Mahajan et al., 1990; Carley, 1991; Krackhardt, 2001; Wejnert, 2002; Roger, 2003) and information cascade (e.g., Banerjee, 1992; Bikhchandani et al., 1992; Strang & Macy, 2001; Watts, 2002; Bettencourt, 2003), with which it shares the basic assumption that subjects earn information from their social environment. Notwithstanding, whereas information diffusion and information cascade restrain the focus of the analysis on the endogenous dissemination or aggregation of signals within a population as subjects observe the available solutions or the performance of others,

parallel problem-solving takes a more radical stance and argues that subjects actively seek out novel signals from the environment while these signals are propagating throughout the network (Lazer & Friedman, 2007).

3.2.3 Organizations and Imitative Behaviour

Imitating the practices of high-performing actors is indeed one of the key strategies employed by organizations to solve problems and improve their performance (Nutt, 1998; Argote, 1999; Csaszar & Siggelkow, 2010), thereby representing a major part of the competitive process (Lieberman & Asaba, 2006). The creation of explicit alliances among organizations (Kogut, 1988; Mowery & Sampat, 1999), the adhesion by start-ups to business incubators involving well-connected venture capitalists (Bøllingtoft & Ulhøi, 2005), or the hiring of information brokers allowing for the interchange between otherwise unconnected groups within the network (Burt, 1992) are only some of the mechanisms that organizations can adopt to access knowledge located outside their boundary.

Provided that organizations are able to successfully imitate one another, it is common assumption that imitation is largely beneficial, in that it consents to accelerate the development of organizational capabilities (Kogut, 1988; Mowery & Sampat, 1999) and, consequently, to attain a sharp adaptive advantage over the short-run (Lazer & Friedman, 2007). This type of mimetic isomorphism becomes especially useful when the environment is characterized by persistent uncertainty and requires the implementation of rapid measures (Lieberman & Asaba, 2006), as it allows organizations to stay in tune with the norms and conventions of their industry (DiMaggio & Powell, 1983) economizing on search and communication costs (Csaszar & Siggelkow, 2010). However, as organizations imitate more practices from each other, strategic convergence induces a progressive decline in the differentiation of solutions that, owing to increasing competition, results in profit erosion (Porter, 1996).

Networked groups based on perfect imitation and flawless information diffusion should

(H1a) quickly converge on a single solution

(H1b) perform well over the short-run but poorly over the long-run

A further body of work investigated the implications deriving from the assumption that organizations are not always able to successfully imitate one another. Under these circumstances, the network suffers by definition a loss of efficiency in terms of information diffusion and knowledge sharing. Nonetheless, as shown in computational models, errors in imitation can vastly expand the potential space examined by organizations because, in principle, any solution in the entire space between two strategies might be sampled during the

imitative process (Lazer & Friedman, 2007). This guarantees the maintenance of diversity among organizations, which represents a necessary condition for the seamless improvement of solutions over time (George, 1972; Janis, 1972; Nemeth, 1985; Florida, 2002; Sunstein, 2003; Page, 2007). Although in the beginning the diffusion of the most successful practices takes more time, thereby impairing the adaptive advantage provided by imitation over the short-run, the possibility of errors gives then rise to an adaptive advantage leading to superior systemic performance over the long-run (Lazer & Friedman, 2007). From a technical standpoint, this is a direct consequence of the crossover mechanism intervening in the imitative process, owing to which taking a piece of an organization's solution and mixing it with another organization's solution may result in something that is better than the two starting solution (Mitchell, 1996).

Networked groups based on imperfect imitation and flawed information diffusion should (H2a) preserve strategic diversity

(H2b) perform poorly in the short-run but well in the long-run

An alternative strand of research underscored the threat deriving from the level of complexity. Specifically, a higher level of complexity is assumed to reduce the diffusion of practices in or from a large organization, thereby exerting a detrimental effect on the imitation of high-performing actors (Ounjian & Carne, 1987). Plausibly, this is due to the increased time required to successfully reproduce multifaced innovations (Galbraith, 1990). The relevant assumption has been validated at a formal level using computational models, which showed how interactions among multiple dimensions of an organization produce severe barriers to imitation (Rivkin, 2000).

(H3) The performance benefits provided by imitation decrease with the level of complexity

3.2.4 Cognitive Foundations of Imitative Behaviour

Following Tomasello's (1999) classic definition, to imitate "means reproducing an instrumental act understood intentionally, that is reproducing not just the behavioral means but also the intentional end for which the behavioral means was formulated" (p. 512). The cognitive faculty of human beings to discern within the actions of others the underlying goal and the alternative means that can be employed to achieve it has also been labeled no-trial learning (Bandura, 1965), because it is even faster than the one-trial learning exhibited by animals with a strong tendency to form cause-and-effect associations. Ultimately, this peculiar talent at the core of - no-trial - social learning allows human imitators to introduce novel behaviors in

their repertory avoiding the costs imposed by - trial-and-error - individual learning (Wisdom et al., 2013).

Notwithstanding, making use of an evolutionary simulation, Rogers (1988) found that unconditionally reproducing the solutions of others in a dynamic environment provided no adaptive advantage to the relevant population, since the benefits from social learning were eventually outweighed by the cost of using obsolete information. In the wake of this result, other researchers noted that, for social learning to provide an overall adaptive advantage to collective performance, subjects must be able to imitate conditionally (Laland, 2004), so that a synergy between individual and social learning can be established leading to the progressive refinement of information and, consequently, of solutions (Boyd & Richerson, 1995; Kameda & Nakanishi, 2002).

One of the most effective social learning conditions to determine who should be imitated is represented by the follow-the-successful - or imitate-the-best(s) - heuristic, which posits that subjects are inclined to imitate the behavior of the most successful individual(s) within the reference group (Boyd & Richerson, 2005). This cognitive short-cut is largely employed to address information diffusion relative not only to strategic practices within the organizational (e.g. Lazer & Friedman, 2007; Barkoczi et al., 2016) or economic domain (Ostrom, 1991; Ellickson, 1994; Samuelson, 1997; Binmore, 1998), but also relative to foraging attitudes, cultural norms, and symbolic artifacts at a broader anthropological level (Katz, 1974; Harris, 1979; Jensen & Erickson, 1979; McEvoy & Land, 1981; Tomasello, 1999). Notably, the relevant heuristic appears to be especially valuable when problems present an increasing level of difficulty conditional to environmental variables (Boyd & Richerson, 1985, 1988, 1989, 1995, 1996).

Several authors have eventually argued that differences between solutions tend to promote learning opportunities (Mill, 1848; Burt, 1992) and collective intelligence of a group (Hutchins, 1995; Page, 2007). From an organizational perspective, this implies that strategic distance can be regarded as the very seed of innovation. Nonetheless, the empirical evidence hinted that the success of imitation is sturdily influenced by the degree of similarity between the knowledge bases of organizations (Helfat, 1998; Lane & Lubatkin, 1998; Lieberman & Asaba, 2006). Once more, this result has been verified at a formal level using a computational model, which showed how the value of transferring knowledge or technologies from one context to another tends to decrease as the similarity between the relevant contexts decreases (Gavetti, 2005).

(H4) Within inefficient networks, a greater strategic distance entails a greater risk of error and is associated with less imitation of more successful others.

3.3 Method

3.3.1 Task

Capitalizing on the laboratory framework introduced by Billinger et al. (2014), the core experimental task consisted of a combinatorial problem based on the NK model, in which multiple components were to be assembled to generate different possible configurations of a product. The specific number of components was contingent on parameter N , which was set equal to 10. Since any of these ten components could take on only two possible values - namely, on or off - the resulting problem space was constituted of 2^{10} (1,024) combinations. Each combination was associated with a payment generated by a standard NK algorithm, thereby recreating a performance landscape that the subjects explored through experiential search. To cut off the impact of foreknowledge and prevent subjects from tapping into some cognitive prior able to guide search (Gavetti & Levinthal, 2000), the task was framed as "the space traveler game", wherein participants were required to design a number of images by combining ten abstract symbols. The products were then to be sold to the space traveler, whose willingness to pay a different price for different products reflected the structure of the fitness landscape.

We extended the original task to accommodate the introduction of network dynamics and, thus, the diffusion of information, so as to allow not only for individual but also for social learning. To reach this objective, in each trial subjects visualized the payments obtained in the prior trial by two further, randomly drawn participants. Based on this information, every subject could then decide whether pursuing autonomous search, thereby assembling a combination on his own, or whether resorting to imitation, thereby adopting the combination submitted in the prior trial by one of the two further, randomly drawn participants with whom she was put in communication in the trial of interest. It is worth noting how the choice to have subjects being able to visualize the payments of two further participants - but not the actual combinations behind these payments - was consistent with common assumptions in the theoretical literature, according to which worse-performing agents are expected to imitate those who exhibit a better performance. Furthermore, this was instrumental in ruling out possible extraneous influences that might have undermined the internal validity of the experiment. Strictly speaking, if subjects were provided with actual combinations of other participants, they could decide to reproduce the relevant combination manually rather than using the imitation button, even if this was not time efficient. Also, they could decide to implement some form of local or distant search moving from the knowledge of the combination submitted by another participant and its relative payment. In both cases, this

would end up making impossible to properly disentangle autonomous search from imitation, thereby impairing the correct measurement of the two variables.

3.3.2 Design

The experiment hinged on a 3x3 mixed factorial design, with the level of ‘Complexity’ (None, Moderate, High) varied within subjects, and the modality of ‘Problem-Solving’ (Independent, Parallel Perfect, Parallel Imperfect) varied between subjects.

The level of complexity depended on parameter K , which regulates the ruggedness of a landscape. More thoroughly, the smooth landscape was randomly generated setting $K = 0$ to capture the absence of complexity, whereas the rugged landscapes were randomly generated setting $K = 3$ to capture moderate complexity, and $K = 9$ to capture high complexity. Notably, a random interaction matrix governed the interdependencies between components in place on the two rugged landscapes. To prevent the formation of reference points across the three tasks, we hid mean and maximum performance normalizing the three landscapes based on as many multipliers. Although randomly generated, the landscapes had a representative number of peaks, that is: exactly 1 for $K = 0$, 32 for $K = 5$, and 93 for $K = 9$.

Turning to the modality of problem-solving, it was contingent on the manipulation of the accuracy of imitation, the subsequent flow of information and, ultimately, the emerging network dynamics. In the independent problem-solving condition, which replicated the original experiment by Billinger et al. (2014), the subjects involved in a session searched the problem space as isolated monads, relying uniquely on individual learning. In the two parallel problem-solving conditions, on the other hand, the subjects involved in a session functioned as a fully interconnected network and were able to search the problem space collectively, relying simultaneously on individual and social learning. The essential difference between the two parallel problem-solving conditions dwelled in the quality of information diffusion and knowledge sharing within the network, which was conditional on the accuracy of the imitative process. In the case of flawless information diffusion, the imitation process entailed no error, so that subjects were always able to reproduce the intended combination successfully, thereby incurring in no mistake. Conversely, in the case of flawed information diffusion, the imitation process did entail the possibility of error, so that subjects were not always able to reproduce the intended combination successfully, but faced the eventuality of one or more mistakes. Specifically, to introduce the possibility of error in the imitative process, we employed the Lazer-Friedman mechanism (Lazer & Friedman, 2007). Consistently with the relevant mechanism, the arise of an error depended on a parameter e affecting only the components within the last combination of a subject that were dissimilar to the components within the combination of the participant she intended to imitate. In other terms, when a

subject imitated another, provided that two components in their respective combinations were dissimilar - namely one on and the other off, or vice-versa - there was a probability $1 - e$ that at the end of the imitative process the component of interest was replaced with the corresponding component, and a probability e that it remained in its original state, meaning that an error had occurred. Ultimately, the likelihood of no error occurring in the imitative process was $(1 - e)^d$, with d representing the distance between the last combination of a subject and the combination of the other participant she intended to imitate as given by the number of dissimilar components. Notably, the mechanism in point captured the assumption that the success of imitation depends on the degree of similarity between the knowledge bases of organizations (Helfat, 1998; Lane & Lubatkin, 1998; Gavetti, 2005; Lieberman & Asaba, 2006). In keeping with Lazer and Friedman (2007), we set $e = 0.20$.

The design allowed for the introduction of two sources of persistent uncertainty. The first pertained payments, which remained unknown until an image was sold and, thence, could only be progressively discovered through search activity. The second pertained the interdependencies between the different components as stipulated by the random interaction matrix at the root of the moderately complex and highly complex tasks, due to which the contribution of each component depended upon K further components impossible to single out.

3.3.3 Participants

135 students (58% females, 42% males) were randomly selected using the university recruitment system and received real monetary incentives. Subjects were on average 21 years old and came from a variety of educational backgrounds. The participants were randomly assigned to the 3 between-subject conditions, each carried out across 3 sessions of 15 individuals, for a total of 45 subjects per condition. The 9 sessions were held in the computer-based Cognitive and Experimental Economics Laboratory (CEEL) at the University of Trento. An inconvenient with the early version of the software caused two computer stations to freeze in the first out of three sessions of the independent problem-solving condition. This forced the students associated with the two stations to leave the laboratory prematurely, thereby lowering the sample size of the independent condition to 43 subjects.

3.3.4 Procedure

At the beginning of every session, participants visualized treatment-specific instructions, which were read aloud by one of the experimenters. The instructions described the experimental task and the structure of the relevant session, emphasizing how each participant was

going to confront three space travelers who came from three different planets and, therefore, had completely different tastes and preferences, so as to make clear that the knowledge acquired in the previous encounter would not be pertinent in the following one. Technically, the three space travelers reflected three different landscapes: one smooth and simple ($K = 0$), one rugged and moderately complex ($K = 3$), and the other rugged but highly complex ($K = 9$). Participants were also informed about the exact number of trials.

Focusing on treatment-specific instructions, in the independent problem-solving condition, subjects were informed that they were going to receive no information about the combination produced by other participants or their relative performance. In the parallel problem-solving condition with flawless information diffusion, subjects were informed that, in each trial, they were going to visualize the payments obtained in the prior trial by two further, randomly drawn participants, but not the actual combinations behind these payments. Moreover, subjects were made aware that, if retained suitable, there was the opportunity to imitate one of two participants they were put in communication with by using a specific button. If they decided to do so, in the subsequent trial the combination behind the payment of the participant who was imitated would have been explicitly visible as last produced combination. In the parallel problem-solving condition with flawed information diffusion, subjects received similar information to the condition with flawless information diffusion, but with two essential additions. First, they knew that, in each trial, along with the payments obtained by two further, randomly drawn participants in the prior trial, they would have also visualized the distance between each of these combinations and one's own last combination, with the distance being equal to the number of dissimilar components. Second, they were made aware that the imitative process entailed a margin of error that might have led to some inaccuracy in the reproduction of the combination behind the payment of the participant who was intended to be imitated, with the likelihood of error being proportional to the distance.

As participants correctly answered all the understanding questions, they had access to the first encounter. Thereafter, as all of them had concluded the first encounter, they had access to the second encounter, and so on. To prevent the occurrence of carry-over effects, for each between-subject condition (Problem-Solving) we partially counter-balanced the within-subject levels (Complexity) to ensure that each of them appeared equally often in each position. The resulting sequences induced three alternative task positions for each landscape, thereby allowing to control for eventual learning effects.

In the first trial of each encounter, participants were provided with a preliminary combination that, unbeknownst to them, always coincided with the lowest-performing one. Starting from the second trial, each subject could modify none, some or all the attributes of the preliminary combination or resort to imitation. After making a choice, she learnt the payment of

the selected combination, thereby completing a trial. No limit was imposed to the number of component that could be set on or off, nor to the time that subjects could spend in producing a combination before submitting it. However, since the number of possible combinations (1024) greatly exceeded the number of available trials (24), each choice entailed a high opportunity cost. This captured the constraints in available resources to be invested in the search task.

To support the decision-making process, the software displayed the complete history of submitted combinations and their relative payments. Furthermore, it also displayed the current trial and the task-specific cumulative payment up to that point. Nonetheless, in line with theoretical assumptions, subjects were not made privy to key parameters, such as task complexity, average and peak performance, or the specific probability of error employed for the mechanism at basis of imperfect imitation and the resulting flaws in information diffusion.

As all subjects completed the third and last encounter, participants were required to fill a brief survey to collect socio-demographics details. The final compensation included a show-up fee of €5 added to an extra amount conditional on participants' performance across the three landscapes. In particular, the three subjects with the best final performance within a session received a prize: the first prize was €15, the second €10, and third €5. The prizes were announced at the beginning of the session, with the purpose to institute a compensation scheme based on a rank-order tournament system that could motivate competitive behavior. Actually, rank-order compensation schemes permeate numerous real-world settings in which competing agents are demanded to allocate scarce resources, such as effort, money, or time (e.g., Tullock, 1967; Krueger, 1974; Gibbons, 1998; Lazear, 1999; Prendergast, 1999). This is particularly salient in the labor market where, based on the relative performance of agents, rank-order tournaments frequently modulate intra-organizational competition between lower-level employees (Lazear & Rosen, 1981) as well as inter-organizational competition between higher-level executives (Gibbons & Murphy, 1990). Compensation schemes based on rank-order tournaments have been attentively scrutinized in laboratory research, which highlighted how average effort levels in tournaments are coherent with theoretical equilibrium levels, even though they exhibit higher variance (Bull et al., 1987). Signally, the relevant result has been replicated in a wide number of experiments (see Dechenaux et al., 2015 for a comprehensive review). Further laboratory research has also highlighted how, compared to compensation schemes based on the bare cumulation of individual performance, rank-order tournaments do not modify the risk-taking attitude of participants, while enabling to minimize transaction costs and to counteract the arise of wealth effects or cross-task contaminations (see Charness et al., 2016 for a comprehensive review).

Including instructions and questionnaire, the sessions dedicated to independent problem-solving lasted about 60 minutes, while those dedicated to parallel problem-solving lasted about 90 minutes. The experimental software was developed and implemented using oTree (Chen et al., 2016).

3.3.5 Descriptive Statistics

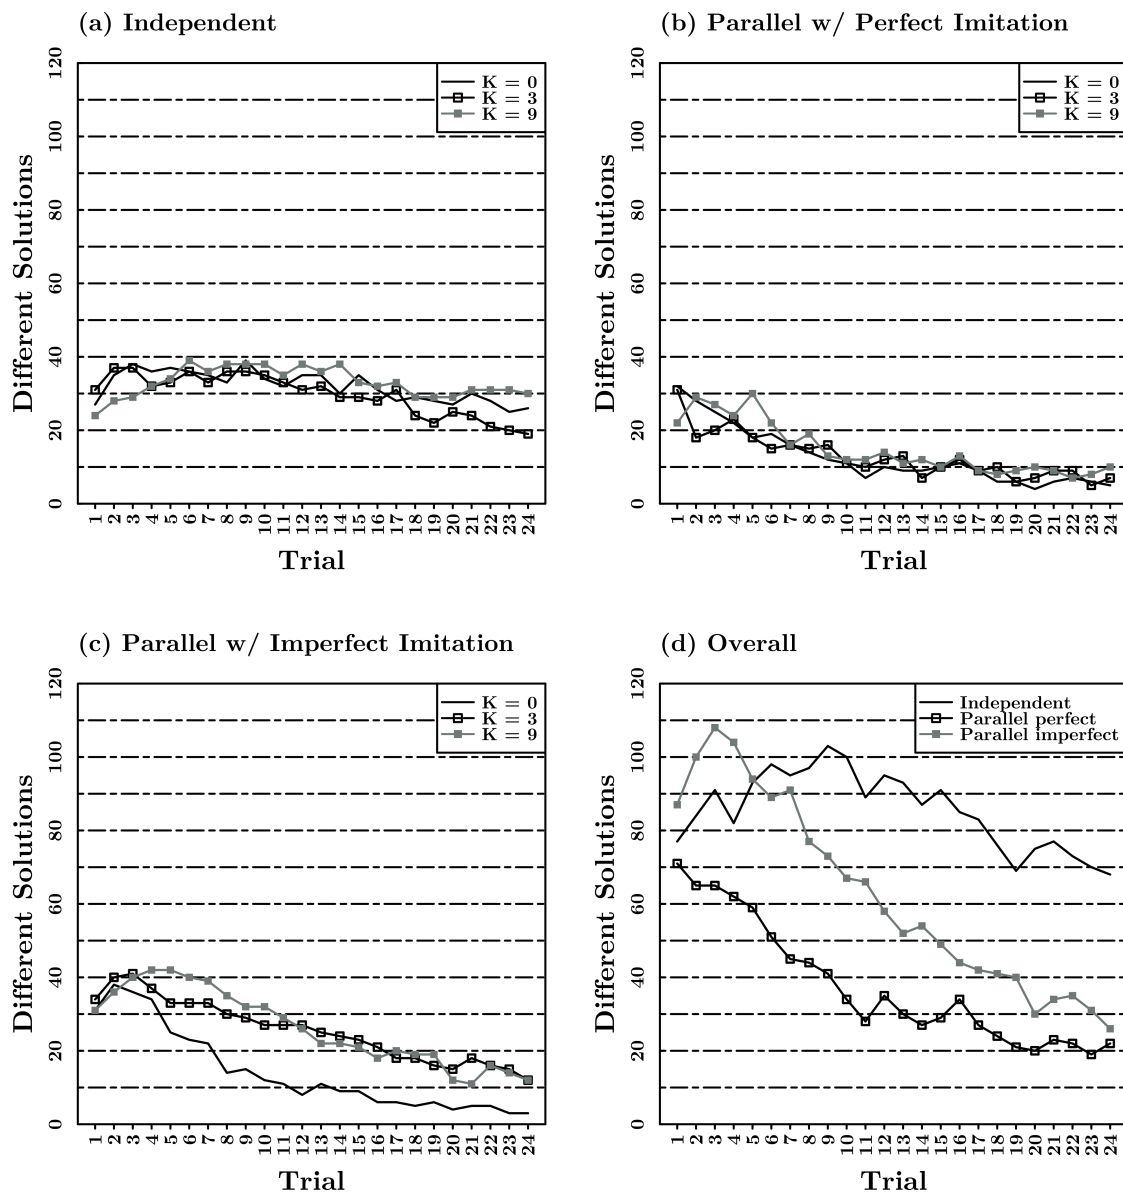
We begin the report of the findings from the experiment presenting descriptive statistics concerning strategic diversity, human performance, autonomous search behavior, and imitative behavior. These statistics are paralleled with hypothesis-based tests of significance. More precisely, differences relative to within-subject conditions were assessed using the Wilcoxon signed-rank test, a non-parametric test commonly employed to compare two dependent samples. On the other hand, differences relative to between-subjects conditions were assessed using the Mann-Whitney U test, a non-parametric test commonly employed to compare to independent samples.

Figure 3.1 displays strategic diversity for the three levels of complexity in each of the three modalities of problem-solving and overall. Strictly speaking, strategic diversity captures the number of distinct combinations per trial. Overall, the greatest degree of strategic diversity occurs under independent problem-solving ($M = 85.45$, $SD = 10.45$). However, focusing on parallel problem-solving, strategic diversity under the flawed information diffusion system instantiated by imperfect imitation ($M = 62.16$, $SD = 26.15$) is greater than under the flawless information diffusion system ensured by perfect imitation ($M = 37.41$, $SD = 16.48$), especially when the landscape is moderately or highly complex. In fact, in the simple task, average strategic diversity is 14.20 under imperfect imitation ($SD = 7.70$) and 12.58 under perfect imitation ($SD = 11.22$). However, on the moderately complex landscape, average strategic diversity is 25.50 under imperfect imitation ($SD = 8.45$) and 12.87 under perfect imitation ($SD = 6.08$). Lastly, on the highly complex landscapes, average strategic diversity is 26.25 under imperfect imitation ($SD = 10.37$) and 14.83 under perfect imitation ($SD = 7.09$). A Mann-Whitney test indicated that the scores from the imperfect imitation condition were significantly different than scores from the perfect imitation condition ($W = 118$, $p < .001$).

Figure 3.2 displays the average performance of participants for the three levels of complexity in each of the three modalities of problem-solving and overall. To mirror the simulative results from the NK literature, where an agent adopts a new combination if and only if it provides a fitness improvement, performance is measured based on the cumulative maximum, which is to say based on the best-performing combination identified by a subject up to that trial. Overall, average performance increases in the number of trials, with marginal gains becoming progressively lower. As evident, parallel problem-solving with perfect imitation

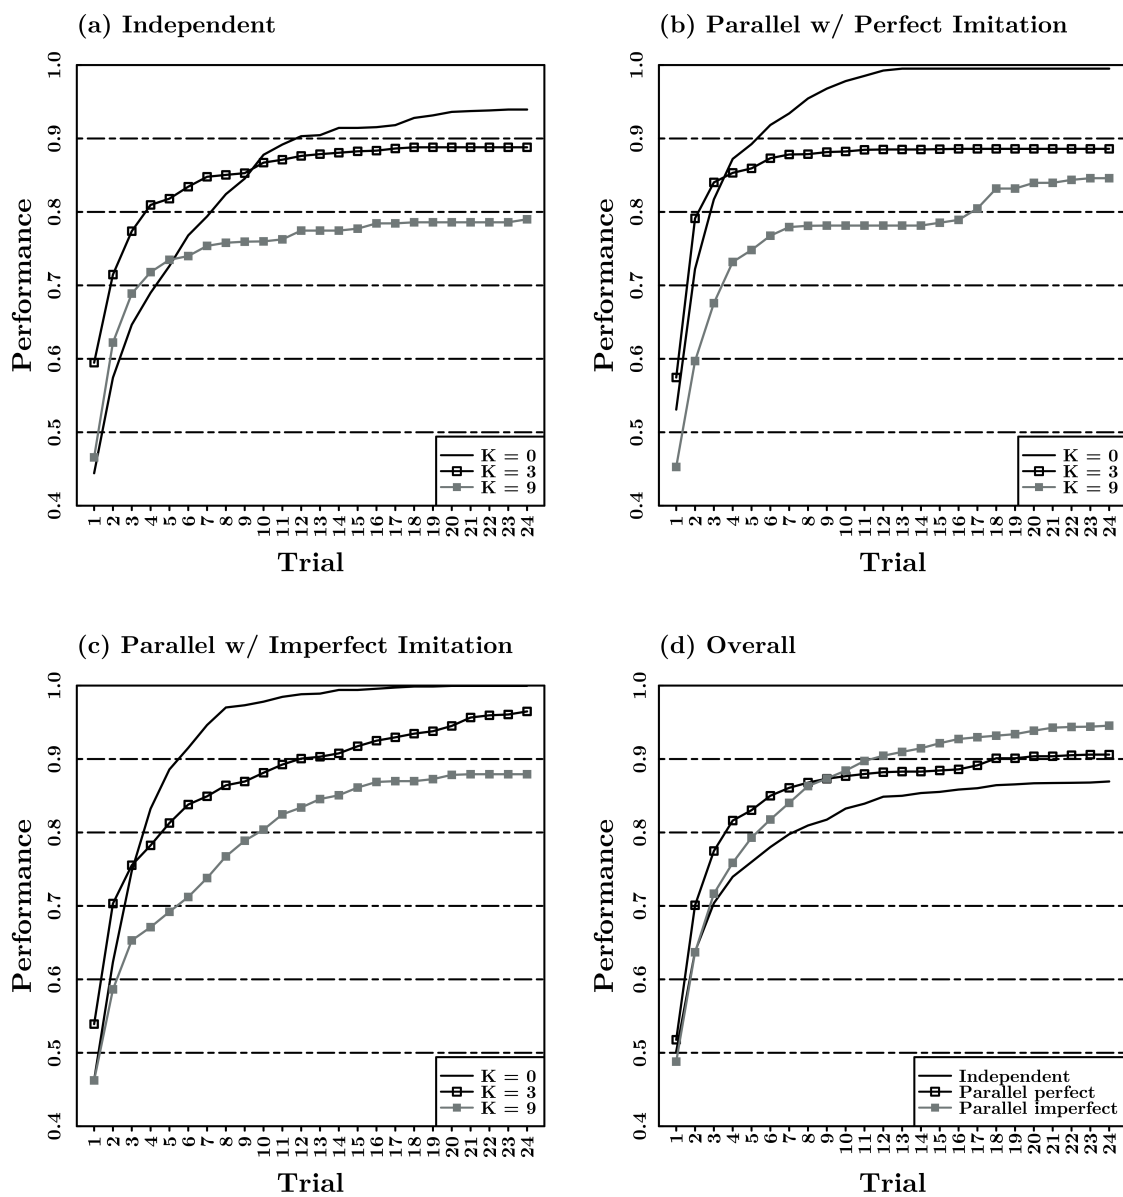
provides an adaptive advantage over the short-run, while parallel problem-solving with imperfect imitation provides an adaptive advantage over the long-run. These trends are consistent with the temporal pattern suggested by simulations (Siggelkow & Levinthal, 2003; Lazer & Friedman, 2007). In detail, under independent problem-solving, average performance in the simple landscape ($M = 0.83$, $SD = 0.10$) is approximately the same as average performance in the moderately complex landscape ($M = 0.84$, $SD = 0.17$), although in both cases it is higher than average performance in the highly complex landscape ($M = 0.74$, $SD = 0.12$). Three

Figure 3.1 Strategic Diversity



separate Wilcoxon tests indicated that scores from the simple landscape are not significantly different than scores from the moderately complex landscape, while scores from the highly complex landscape are significantly different than scores from both the simple ($V = 419196$, $p < .001$) as well as the moderately complex landscape ($V = 468743$, $p < .001$). In contrast, under parallel problem-solving with perfect and imperfect imitation, average performance on the moderately complex and highly complex landscape is lower than the simple landscape. More thoroughly, with perfect imitation, average performance on the simple landscape (M

Figure 3.2 Performance



= 0.93, SD = 0.13) is lower than the moderately complex (M = 0.86, SD = 0.08) and the highly complex landscape (M = 0.76, SD = 0.10). Three further Wilcoxon tests indicated that the scores from the simple landscape are significantly different than scores from the moderately complex (V = 467056, $p < .001$) and the highly complex landscape (V = 562664, $p < .001$), while scores from the moderately complex landscape are significantly different than scores from the highly complex landscape (V = 555930, $p < .001$). Likewise, with imperfect imitation, average performance on the simple landscape (M = 0.92, SD = 0.16) is lower than the moderately complex (M = 0.87, SD = 0.12) and the highly complex landscape (M = 0.79, SD = 0.14). Three additional Wilcoxon tests indicated that scores from the simple landscape are significantly different than scores from the moderately complex (V = 483508, $p < .001$) and the highly complex landscape (V = 538888, $p < .001$), while scores from the moderately complex landscape are significantly different than scores from the highly complex landscape (V = 496437, $p < .001$). Comparing the three modalities of problem-solving, average performance in the independent condition (M = 0.80, SD = 0.13) is lower than the parallel condition with perfect imitation (M = 0.85, SD = 0.11) that, in turn, is slightly lower than the parallel condition with imperfect imitation (M = 0.86, SD = 0.14). Three separate Mann-Whitney tests indicated that scores from the independent condition are significantly different than scores from the parallel condition with perfect (W = 3957308, $p < .001$) and imperfect imitation (W = 3314222, $p < .001$), while scores from the parallel condition with perfect imitation are significantly different than scores from the parallel condition with imperfect imitation (W = 4393585, $p < .001$). More insights derives from the region of the performance landscapes where subjects ended the search process. In the simple landscape, 39.5% of participants reached the global maximum under independent problem-solving, 97% managed to reach it in the under parallel problem-solving with perfect imitation, and 97% under parallel problem-solving with imperfect imitation. In the moderately complex landscape, nobody reached the global maximum under independent problem-solving or parallel problem-solving with perfect imitation, while 23.2% of participants managed to reach it under parallel problem-solving with imperfect imitation. Lastly, in the highly complex landscape, nobody managed to reach the global maximum neither under independent problem-solving nor under parallel problem-solving with perfect or imperfect imitation.

Figure 3.3 displays the patterns of autonomous search behavior for the three levels of complexity in each of the three modalities of problem-solving and overall. Following Billinger et al. (2014), autonomous search behavior is assessed based on the number of changed attributes compared with the best-performing combination identified by a subject in prior trials. The resulting measure is referred as search distance. It reflects the Hamming distance in information theory (Hamming, 1950), a metric to compare binary data of strings traditionally

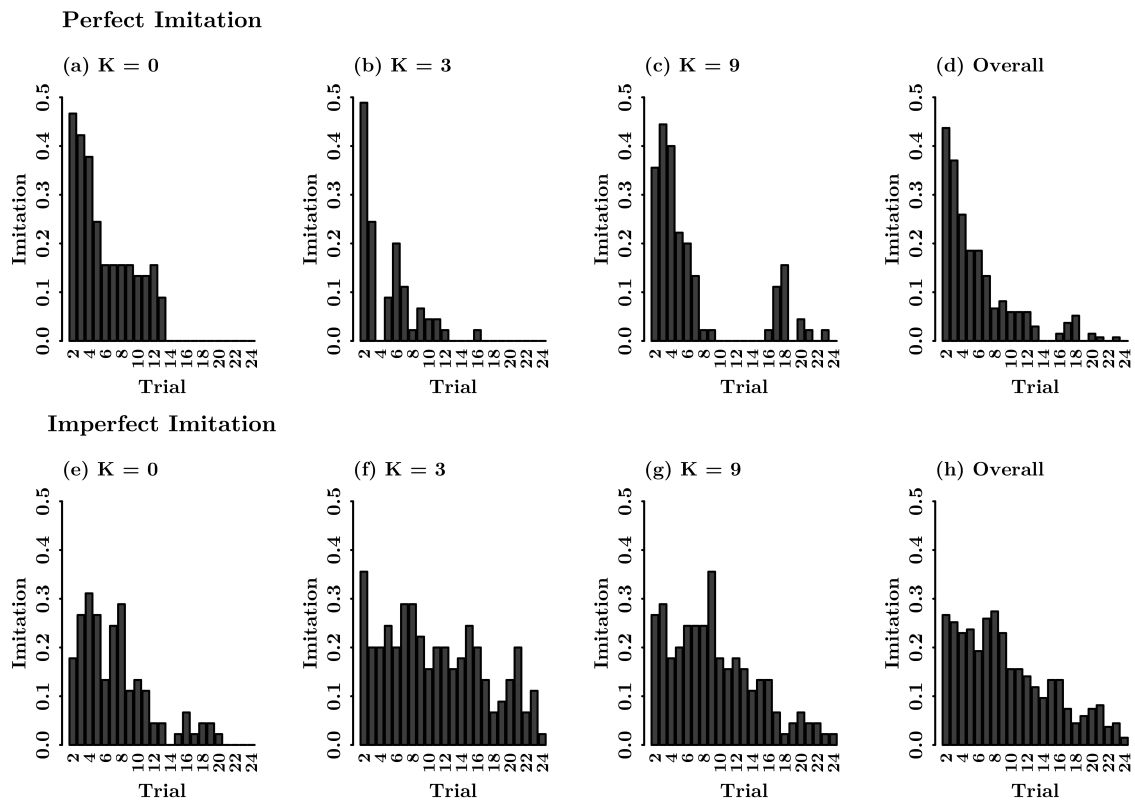
independent problem-solving, average search on the simple landscape ($M = 2.28$, $SD = 1.81$) is higher than the moderately complex landscape ($M = 2.25$, $SD = 1.88$) and lower than the highly complex landscape ($M = 2.33$, $SD = 1.84$). Three separate Wilcoxon tests indicated that the scores from the three conditions are not significantly different. Under parallel problem-solving with perfect imitation, average search on the simple landscape ($M = 3.19$, $SD = 2.16$) is higher than the moderately complex landscape ($M = 2.97$, $SD = 2.49$) and the highly complex landscape ($M = 2.94$, $SD = 2.22$). Three further Wilcoxon tests indicated that, again, the scores from the three conditions are not significantly different. Likewise, under parallel problem-solving with imperfect imitation, average search on the simple landscape ($M = 2.94$, $SD = 2.23$) is higher than the moderately complex landscape ($M = 2.84$, $SD = 2.34$) and the highly complex landscape ($M = 2.91$, $SD = 2.18$). Three additional Wilcoxon tests indicated that, once more, the scores from the three conditions are not significantly different.

Comparing the three modalities of problem-solving, average search distance under independent problem-solving ($M = 2.28$, $SD = 1.18$) is lower than under parallel problem-solving with perfect ($M = 3.03$, $SD = 2.31$) and imperfect imitation ($M = 2.89$, $SD = 2.26$). Three separate Mann-Whitney tests indicated that scores from the independent condition are significantly different than scores from the parallel condition with perfect ($W = 647702$, $p < .001$) and imperfect imitation ($W = 811642$, $p < .001$), while scores from the parallel condition with perfect imitation are not significantly different than scores from the parallel condition with imperfect imitation. Across all conditions average search distance is systematically higher than it would be with the pure local search strategy ordinarily presupposed in computational models. De facto, consistently with Billinger et al. (2014), subjects begin the search process resorting to distant search, thereby changing many attributes at once. Afterwards, they gradually narrow down their search distance, broadening it again as the end of the task approaches. The behavioral pattern appears to be different for the smooth and the rugged landscapes, with the increase of search distance in the latter trials being especially marked when the task is more complex.

Table 3.1 displays the frequency distribution of search distances for the three levels of complexity in each of the three modalities of problem-solving, thereby providing additional insights into autonomous search behavior. It is apparent that local search represents the predominant strategy (50.4% under independent problem-solving, 50.2% under parallel problem-solving with perfect imitation, and 37.9% under parallel problem-solving with imperfect imitation). Very distant search (changing 9 or 10 attributes) is quite infrequent and occurred mostly under parallel problem-solving, particularly in complex tasks. Lastly,

which takes value of 1 if a subject imitated another participant, and 0 otherwise. We only consider observations in which the the decision to imitate results in a search distance ranging between 1 (change of one attribute) and 10 (change of all attributes). Accordingly, we drop the first serviceable trial, in which all subject shared the same combination and the imitative process necessarily resulted in a search distance of 0 - therefore being meaningless. Overall, the proportion of perfect imitation ($M = 0.28$, $SD = 0.45$) is lower than that of imperfect imitation ($M = 0.33$, $SD = 0.47$). As can be seen from the histograms, the relative frequency of both variables is higher at the beginning of the search process and then decreases in the number of trials. Nonetheless, their respective patterns are fundamentally different. In fact, perfect imitation tends to concentrate in the first half of the task, whereas imperfect imitation is more sustained over time. Furthermore, while perfect imitation is greater when the task is simple ($M = 0.35$, $SD = 0.48$) than when it is moderately complex ($M = 0.19$, $SD = 0.39$) or highly complex ($M = 0.29$, $SD = 0.46$), imperfect imitation is roughly the same regardless whether the task is simple simple ($M = 0.33$, $SD = 0.47$), moderately complex ($M = 0.34$, $SD = 0.47$) or highly complex ($M = 0.31$, $SD = 0.46$). A chi-square test for equality of proportions with continuity correction indicated that scores from the

Figure 3.4 Imitative Behavior



perfect imitation condition are significantly different than scores from the imperfect imitation condition (chi-square = 6.87, $df = 1$, $p < .01$).

3.3.6 Regression Analyses

A regression analysis was conducted to evaluate the impact of the level of complexity and the modality of problem-solving on strategic diversity and human performance, respectively. Strategic diversity consists of count data on the number of distinct combinations per trial, and was assessed using a Poisson regression. Performance is a continuous variable measured with respect to the cumulative maximum identified in prior trials, and was assessed using an OLS regression. Table 3.3 displays the Poisson model of strategic diversity and the OLS model of human performance. Both models control for the possible influence of the number of search trials and the position of the task in the session-specific sequence. The evidence suggests that strategic diversity increases with the level of complexity, while it decreases under parallel problem-solving as well as with the number of trials. More precisely, the logs of expected counts suggest that, relative to a flawless information diffusion system, a flawed information diffusion system does preserve strategic diversity. Turning to performance, the evidence suggests that it decreases with the level of complexity, while it increases under parallel problem-solving, the number of trials, and task position. Furthermore, the comparison of the estimated effects exerted by the two modalities of problem-solving suggests that the improvement in performance under a flawed information diffusion system is greater than under a flawless information diffusion system.

Table 3.2 Main Variables

Variable	Data	Min	Max	Mean	SD	Skewness	Kurtosis
Complexity	Any	0	9	-	-	-	-
Task position	Any	1	3	-	-	-	-
Information diffusion	Overall	0	2	-	-	-	-
Cumulative maximum	Overall	0.19	1	0.84	0.14	-2.06	5.22
Strategic diversity	Overall	3	42	21.98	11.17	-0.07	-1.37
Trial	Overall	1	24	-	-	-	-
Individual feedback	Independent problem-solving	0	1	0.26	0.44	1.04	-0.93
Unsuccessful attempts	Independent problem-solving	0	22	3.70	4.61	1.50	1.74
Prior search distance	Independent problem-solving	0	10	1.98	1.94	1.50	1.90
Trial	Independent problem-solving	1	24	-	-	-	-
Information diffusion	Parallel problem-solving	1	2	-	-	-	-
Individual feedback	Parallel problem-solving	0	1	0.24	0.43	1.17	-0.63
Unsuccessful attempts	Parallel problem-solving	0	20	3.06	4.10	1.67	2.36
Prior search distance	Parallel problem-solving	0	10	2.26	2.36	1.09	0.49
Social feedback	Parallel problem-solving	0	1	0.49	0.50	0.04	-1.99
Strategic distance (PCA)	Parallel problem-solving	-1.24	3.77	0.01	0.50	1.31	-1.14
Trial	Parallel problem-solving	2	24	-	-	-	-

Table 3.3 Poisson Model with Strategic Diversity as Dependent Variable and OLS Model with Performance as Dependent Variable

	Strategic Diversity	Performance
Complexity	0.037*** (0.006)	-0.004*** (0.001)
Information diffusion (Flawless)	-0.903*** (0.039)	0.045*** (0.003)
Information diffusion (Flawed)	-0.472*** (0.038)	0.052*** (0.003)
Trial	-0.042*** (0.002)	0.011*** (0.001)
Task position		0.006*** (0.001)
Constant	3.889*** (0.032)	0.680*** (0.004)
Residuals	333.38	
Log Likelihood	-680.963	
Akaike Inf. Crit.	1,371.925	
R ²	0.7534	0.336
Adjusted R ²		0.336
Residual Std. Error		0.110
F Statistic		970.313***
Observations	216	9,576

Note: *p<0.05; **p<0.01; ***p<0.001

A further regression analysis was conducted to assess the determinants of autonomous search behavior under independent problem-solving. The dependent variable of interest is the positive search distance of subjects from the first serviceable trial, which consists of discrete count data with values ranging between 1 and 10 conditional on the number of changed attributes relative to the best-performing combination identified in prior trials. The average search distance is non-normally distributed and positively skewed. Therefore, we adopted a Poisson regression. In particular, to control for the violation of the assumption that the variance should equal the mean, we had recourse to robust standard errors for the parameter estimates (Cameron & Trivedi, 2009). Moving from these premises, we ran four Poisson models. Model 1 assesses how the within-subject treatment conditions, namely the levels of complexity, affect search distance. Model 2-4 inspect the impact of individual fixed effects. More precisely, Model 2 estimates the effect of individual performance feedback, while Model 3 and 4 evaluates the search history. All models control for the possible influence of the number of search trials and the position of the task in the session-specific sequence.

Focusing on fixed effects, individual performance feedback is defined as a binary variable, which takes value of 1 if in the last trial the subject managed to achieve a performance

**Table 3.4 Poisson Models with Search Distance as Dependent Variable
(Independent Problem-Solving)**

	Model 1	Model 2	Model 3	Model 4
Complexity (K = 3)	-0.006 (0.045)	-0.054 (0.045)	-0.077 (0.044)	-0.059 (0.039)
Complexity (K = 9)	0.033 (0.045)	-0.038 (0.046)	-0.100* (0.046)	-0.058 (0.039)
Individual feedback (Positive)		-0.420*** (0.047)	-0.310*** (0.049)	-0.422*** (0.047)
Number unsuccessful attempts			0.047*** (0.006)	0.026*** (0.005)
Prior search distance				0.153*** (0.007)
Trial	0.008* (0.003)	-0.002 (0.003)	-0.027*** (0.005)	-0.024*** (0.004)
Task position (2)	-0.029 (0.044)	-0.021 (0.042)	-0.019 (0.041)	-0.034 (0.037)
Task position (3)	-0.117* (0.047)	-0.108* (0.046)	-0.093* (0.047)	-0.073 (0.041)
Constant	0.792*** (0.052)	1.015*** (0.058)	1.061*** (0.058)	0.771*** (0.061)
Residuals	2281.8	2169.5	2080.5	1588.6
Log Likelihood	-3,529.495	-3,473.346	-3,428.869	-3,182.883
Pseudo-R ²	0.0096	0.0583	0.0969	0.3105
Observations	1,893	1,893	1,893	1,893

Note: * p < 0.05; ** p < 0.005; *** p < 0.001

improvement relative to the best-performing combination identified in prior trials (success), and 0 otherwise (failure). The variable has been introduced by Billinger et al. (2014) building on prior models of adaptive search (Levinthal & March, 1981) and individual risk attitudes (Kahneman & Tversky, 1979; March, 1988; March & Shapira, 1992) to capture the idea that individuals encode performance feedback as success or failure with respect to a subjective reference point (Markowitz, 1952; March, 1988; Bromiley, 1991). Specifically, positive individual feedback (success) is expected to prompt risk-averse exploitation, thereby lowering the search distance. In contrast, negative feedback (failure) is expected to prompt risk-seeking exploration, thereby increasing the search distance. As regards the search history, it encompasses two variables. The first is the number of unsuccessful attempts, which counts the number of trials since the last performance improvement. The second is the search distance in the prior trial, which allows to control for path dependencies in search behavior.

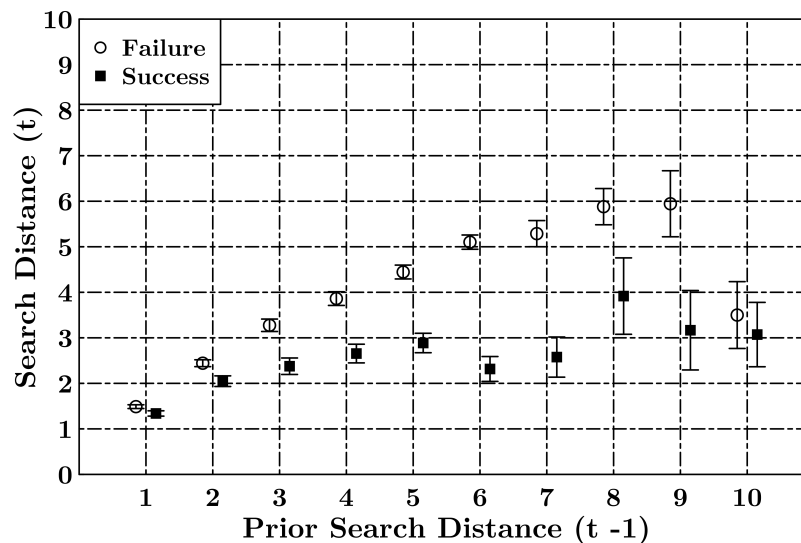
Table 3.4 displays the four Poisson models of autonomous search behavior under independent problem-solving. Model 1 suggests that moderate and high complexity have no immediate effect on search distance. Model 2 suggests that positive individual feedback

reduces the search distance, thereby prompting subjects to exploit the gathered knowledge focusing on the region of the landscape where they experienced a performance improvement. Model 3 suggests that the number of consecutive unsuccessful attempts is positively related with the search distance, indicating that subjects tend to become more impatient with the systematic failure to obtain performance improvements. Model 4 suggests that the prior search distance is positively associated with the search distance, thereby providing evidence in support of the presence of strong path dependence. In the specification with the highest goodness of the fit, Model 4, the search distance decreases with the number of trials, while the task position has no significant effect.

Taken together, this results suggest that autonomous search behavior under independent problem-solving gradually adapts to individual performance feedback. In particular, success in obtaining performance improvements prompts more exploitation, while failure prompts more exploration. Search behavior responds also to the current interval without performance improvements and exhibits a strong path-dependence, such that a higher search distance in the prior trial elicits a higher search distance in the current trial. Eventually, search distance decreases with time, possibly as a result of the progressive difficulty to obtain performance improvement and the concurrent endeavor to maximize the sum of one's own rewards by resorting to more exploitation. Since the variance of the dependent variable is higher than the mean, we conducted a robustness check for overdispersion of the data by using a negative binomial regression obtaining the same qualitative results. Nonetheless, several pseudo- R^2 measures suggested that the goodness of the fit for the Poisson models was superior to the negative binomial models.

A separate regression analyses was conducted to assess the determinants of autonomous search behavior under parallel problem-solving. Compared to the previous case, the Poisson models present one essential addition, which is represented by the categorical variable capturing the effect of the inefficient information diffusion system arising from imperfect imitation relative to the efficient information diffusion system ensured by perfect imitation. Accordingly, Model 1 assesses how the within-subject and between-subject treatment conditions, that is the levels of complexity and the information diffusion system entailed by the modality of parallel-problem-solving, affect search distance. Model 2 examines the effect of individual performance feedback, while Model 3 and 4 evaluates the search history based on the number of unsuccessful attempts and the prior search distance, respectively. As before, all four models control for the possible influence of the number of search trials and the position of the task in the session-specific sequence.

Table 3.5 displays the four additional Poisson models of autonomous search behavior under parallel problem-solving. Model 1 suggests that, again, moderate and high complexity

Figure 3.5 Adaptive Search

have no immediate effect on search distance, while a flawed information diffusion has a positive impact. Model 2 suggests that positive individual feedback still reduces the search distance, prompting subjects to leverage the acquired knowledge focusing on the region of the landscape where they experienced a performance increase. Model 3 suggests that the number of consecutive unsuccessful attempts continues to be positively related with the search distance, so that subjects tend to become more impatient with the systematic failure to obtain performance improvements. Model 4 suggests that the prior search distance maintains a positive association with search distance, providing more evidence in support of strong path dependence. According to the specification with the highest goodness of the fit, Model 4, search distance decreases with the number of trial. However, in a departure from independent problem-solving, under parallel problem-solving the task position appears to have a significant negative effect on search distance. As it will be seen in more detail over the course of the analysis, this may suggest that, when provided with the opportunity to select between autonomous search and imitation, subjects learn to harmonize incremental search with the diffusion of social information. Again, we conducted a robustness check for overdispersion of the data using a negative binomial regression. As before, we obtained the same qualitative results, although several pseudo- R^2 measures suggested that the goodness of the fit for the Poisson models was superior to the negative binomial models.

Figure 3.5 presents a compact summary of the effects contained in the Poisson regressions for autonomous search under independent and parallel problem-solving. It displays how search behavior in the current trial changes conditional to prior search distance and

**Table 3.5 Poisson Models with Search Distance as Dependent Variable
(Parallel Problem-Solving)**

	Model 1	Model 2	Model 3	Model 4
Complexity (K = 3)	-0.062 (0.045)	-0.078 (0.044)	-0.070 (0.044)	-0.087* (0.044)
Complexity (K = 9)	-0.035 (0.044)	-0.051 (0.043)	-0.054 (0.043)	-0.062 (0.042)
Information diffusion (Flawed)	-0.046 (0.036)	-0.034 (0.036)	-0.006 (0.036)	-0.024 (0.035)
Individual feedback (Positive)		-0.314*** (0.044)	-0.245*** (0.047)	-0.458*** (0.050)
Number unsuccessful attempts			0.038*** (0.008)	0.027*** (0.008)
Prior search distance				0.086*** (0.007)
Trial	-0.005 (0.003)	-0.011*** (0.003)	-0.029*** (0.006)	-0.031*** (0.005)
Task position (2)	-0.140** (0.043)	-0.149*** (0.042)	-0.186*** (0.042)	-0.140*** (0.042)
Task position (3)	-0.091* (0.044)	-0.113** (0.044)	-0.157*** (0.044)	-0.111* (0.044)
Constant	1.254*** (0.049)	1.389*** (0.051)	1.413*** (0.052)	1.293*** (0.056)
Residuals	2906.0	2821.8	2776.2	2556.5
Log Likelihood	-3,992.063	-3,950.002	-3,927.181	-3,817.352
Pseudo-R ²	0.0098	0.0384	0.0540	0.1288
Observations	1,857	1,857	1,857	1,857

Note: * $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$

performance feedback, thereby providing an outline of the behavioral model of adaptive search (Billinger et al., 2014). As evident, subjects respond to positive feedback by resorting to more incremental search and to negative feedback by gradually broadening their search distance. Since identifying better-performing combinations becomes increasingly difficult in later trials, this suggests that subjects tend to become impatient with the lack of performance improvements, departing from local search well before a widespread examination of the immediate neighborhood has been achieved.

In order to provide a full account of parallel problem-solving, we complemented the latter survey of autonomous search with a further regression analysis to assess the determinants of imitative behavior. The dependent variable of interest coincides with the decision of subjects to imitate other members of the reference groups after the first serviceable trial. The relevant decision is coded as a binary variable, which takes value of 1 if a subject decided to imitate one out of the two subjects with whom she was put in communication in the current trial,

and 0 otherwise. Thence, to assess the determinants of imitative behavior, we adopted a probit regression. In particular, we ran four probit models. Model 1 assesses how the within- and between-subject treatment conditions, which is to say the levels of complexity and the information diffusion system entailed by the modality of parallel-problem-solving, affect imitation. Model 2-4 inspect the impact of individual fixed effects. In particular, Model 2 estimates the effect of social performance feedback, while Model 3 and 4 evaluates the role of strategic distance and its interaction with social feedback. All four models are structured to control for the possible influence of the number of search trials and the position of the task in the session-specific sequence.

As concerns the latter fixed effects, social performance feedback aims to extend the notion of individual performance feedback previously introduced by Billinger et al. (2014). In detail, social feedback is defined as a binary variable, which takes value of 1 if at least one out of the two participants with whom a subject is put in communication in the current trial exhibits a higher performance relative to the best-performing combination that the subject of interest managed to identify in prior trials, and 0 otherwise. Since a subject visualizes the performance obtained by two randomly drawn participants in the previous trial, this means that she receives a positive social feedback if it exists another participant who has been more successful than she has ever been, while she receives a negative feedback if it does not. The variable retains the idea that individuals encode feedback as success or failure with respect to a subjective reference point (Markowitz, 1952; March, 1988; Bromiley, 1991) and integrates it with the dynamic postulated by the follow-the-successful heuristic (Boyd & Richerson, 2005), according to which human beings are naturally inclined to imitate more successful others. Turning to strategic distance, it captures the number of dissimilar components between the last combination of a subject and the combinations of the two subjects she is put in communication. More in detail, it is obtained integrating the distance with respect to the combination of the first randomly drawn participant and the distance with respect to the combination of the second randomly drawn participant, each measured based on discrete values ranging from 0 to 10. Since these two variables presented a high correlation - as suggested by the Cronbach's alpha statistics (0.86) and confirmed by a Pearson's correlation test ($r = 0.75$, $p < .001$) - the choice to integrate them was mainly motivated by the necessity to prevent multicollinearity, beyond the intention to safeguard the parsimony of the models. Technically, the two variables were integrated building an index based upon a principal component analysis (PCA). Notably, the PCA index showed also a perfect correlation with an alternative index obtained based upon the summing procedure. The preference for the PCA index depended on the superior performance of the resulting model in estimating social feedback when the interaction with strategic distance was present,

since the estimate in point appeared to be less inflated and more stable than the case in which the additive index was used. Last but not least, it should be noted that a value of 0 in the two underlying distance measures reflects the impossibility of an error to occur, and was used to code strategic distance under parallel problem-solving with perfect imitation, where the number of differences is irrelevant, as well as strategic distance under parallel problem-solving with imperfect imitation, where the number of differences does is relevant. The rationale is that the differential effect of the two 0s is caught by the estimator of the dummy used to code the perfect and imperfect condition, which takes value of 0 for the former and of 1 for the latter.

Table 3.6 displays the four probit models of imitative behavior under parallel problem-solving. Model 1 suggests that moderate complexity has a negative effect on imitation. However, it must be immediately noted that the relevant effect dissolves in the following models, which present a higher goodness of the fit. Model 1 suggests also that high complexity has no significant effect on imitation, while the possibility of error and, thence, of a beneficial crossover effect has a positive impact on the decision to imitate others. Model 2 suggests that a positive social feedback is associated with a higher likelihood of imitating others, thereby inducing subjects to follow who has been more successful. Model 3 suggests that the strategic distance is negatively related with imitation, hinting that, although the potential benefits from crossover effects are attractive, a higher risk of incurring in an error results in a lower likelihood of imitating others. Model 4 provides further support to this insight, as it suggests that, in the presence of positive social feedback, subjects are more likely to refrain from imitation as the strategic distance increases. The specification with the highest goodness of the fit, Model 4, suggests also that the likelihood of imitation is not affected by the number of trials or the task position.

Taken together, the results suggest that imitative behavior adapts to social performance feedback. This implies that, when faced with combinatorial problems, subjects are inclined to follow the members of the group who are more successful. Conversely, they prefer to rely on their own search ability when there is nobody who outperforms them. Notably, the absence of a significant association with the task position implies that the negative effect of the latter variable on autonomous search is not due to a decrease in the rate of exploration through imitation that reverberates onto the search distance, but rather to a lower degree of exploration through random, distant search - or the other way around, to a higher degree of exploitation through local search. Plainly speaking, it implies that subjects learn to harmonize the "wise long-jumps" made possible by the imitation of more successful others with the incremental enhancements of performance through autonomous adaptive search. Imitative behavior responds also to the potential benefits that may derive from crossover, although the

Table 3.6 Probit Models with Imitation as Dependent Variable

	Model 1	Model 2	Model 3	Model 4
Complexity (K = 3)	-0.141* (0.069)	-0.116 (0.071)	-0.077 (0.071)	-0.095 (0.072)
Complexity (K = 9)	-0.089 (0.069)	-0.069 (0.071)	-0.024 (0.071)	-0.058 (0.072)
Information diffusion (Flawed)	0.185** (0.056)	0.082 (0.058)	0.508*** (0.100)	0.462*** (0.101)
Social feedback (Positive)		0.632*** (0.057)	0.671*** (0.058)	0.670*** (0.058)
Strategic distance			-0.198*** (0.038)	-0.029 (0.047)
Social feedback (Positive) : Distance				-0.261*** (0.045)
Trial	-0.028*** (0.005)	-0.020*** (0.005)	-0.024*** (0.005)	-0.021*** (0.005)
Task position (2)	0.054 (0.066)	0.074 (0.067)	0.019 (0.068)	0.052 (0.069)
Task position (3)	0.126 (0.068)	0.118 (0.069)	0.088 (0.070)	0.130 (0.071)
Constant	-0.330*** (0.077)	-0.698*** (0.085)	-0.932*** (0.097)	-0.905*** (0.097)
Residuals	2847.6	2723.2	2696.5	2663.3
Log Likelihood	-1,423.811	-1,361.583	-1,348.246	-1,331.668
Akaike Inf. Crit.	2,861.622	2,739.165	2,714.492	2,683.335
Observations	2,352	2,352	2,352	2,352

Note: * $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$

effect is not so straightforward. In fact, although the benefits in point clearly appeal subjects, an increase in the perceived risk of error conveyed by a higher strategic distance tends to reduce the likelihood of imitating others.

We conducted several robustness checks. First, we built a model in which, rather than a unique social feedback predictor, we used a variable for the feedback received from the first randomly drawn participant and a different variable for the feedback received from the second randomly drawn participant. The estimates for the two feedback variables were again positively associated with imitation, but the model presented a higher residual deviance and a lower goodness of the fit. Further robustness checks were conducted to assess social feedback and strategic distance using exclusively the subset of observations from the parallel problem-solving condition with imperfect imitation. Again, we obtained the same qualitative results, but the models presented an inferior goodness of the fit. Also, we ran a model in which we introduced an index capturing the difference between the performances of the two randomly drawn participants with whom a subject was put in communication and the

performance provided by the best combination she identified in prior trials. The estimate of the performance difference was not significant and ended up reducing the goodness of the fit. The reiteration of the same model using two different measures of performance difference, one for each of the two randomly drawn participants, produced identical results. This led us to keep performance difference out of the final model. Last but not least, given that both randomly drawn participants within a reference group simultaneously provided a positive social feedback, we controlled to what extent subjects decided to imitate the one exhibiting the best performance. Consistently with predictions, the evidence suggested that 69% of participants did follow the most successful when imitation entailed no error, while the proportion decreased to 56% when imitation entailed a possibility of error proportional to the strategic distance.

We ran supplementary regressions to probe whether - and eventually how - individual and social feedback respond to changes in the level of complexity and the modality of problem-solving. According to the model of adaptive search under independent problem-solving (Billinger et al., 2014), in fact, individual feedback has an essential role in mediating the influence of the level of complexity on search behavior. More precisely, the model holds that, although complexity has no direct effect on search behavior, the rate of individual positive feedback decreases with the level of complexity and human search behavior adapts to individual performance feedback. Therefore, complexity still has an indirect effect on human search behavior through individual feedback. In the heels of this finding, we were interested in determining whether, in the context of parallel problem-solving, individual and social feedback might have a similar role relative to autonomous search and imitative behavior, respectively. Since the relevant feedback conditions are coded as binary variables, we adopted a probit model. Table 3.7 displays the three probit models used to analyze, respectively, individual feedback under independent problem-solving and individual or social feedback under parallel problem-solving. The evidence suggests that, consistently with the model of adaptive search, the likelihood of individual positive feedback under independent problem-solving decreases with the level of complexity and the number of trials. On the other hand, individual positive feedback under parallel problem-solving still decreases with number of trials, but the negative association with the level of complexity ceases to be significant. Rather, relative to a flawless information diffusion system, the likelihood of individual positive feedback under parallel problem-solving increases in the presence of a flawed information diffusion system. The same applies to social feedback. Specifically, the level of complexity does not affect the likelihood of positive social feedback that, rather, tends to increase with a flawed information diffusion and to decrease with the number of trials. Task position negatively affects individual feedback under parallel problem-solving,

Table 3.7 Probit Models with Individual and Social Feedback as Dependent Variable

	Independent problem-solving	Parallel problem-solving	
	Individual feedback	Individual feedback	Social feedback
Complexity	-0.350*** (0.042)	-0.079 (0.043)	-0.048 (0.034)
Information diffusion (Flawed)		0.168* (0.068)	0.474*** (0.054)
Trial	-0.100*** (0.007)	-0.087*** (0.007)	-0.034*** (0.004)
Task position	0.041 (0.042)	-0.132** (0.042)	0.012 (0.033)
Constant	0.808*** (0.132)	0.263 (0.136)	0.092 (0.106)
Residuals	1858.1	1854.7	3125.1
Log Likelihood	-929.043	-927.342	-1,562.571
Akaike Inf. Crit.	1,866.086	1,864.684	3,135.143
Observations	1,893	1,857	2,352

Note: * $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$

but not individual feedback under independent problem-solving or social feedback under parallel problem-solving. It is also worth noting that, since performance tends to increase with task position and, thence, with the level of gained experience, the negative association between task position and individual - but not social - feedback may reflect an increased level of efficiency in the trade-off between exploration through imitation and exploitation through autonomous adaptive search, in line to which the achievement of quicker and larger improvements at the outset of subsequent tasks ends up reducing the overall rate of success in autonomously discovering better-performing combinations, thereby making the occurrence of a positive individual feedback somewhat less frequent.

3.4 Discussion

Consistently with theoretical predictions and prior empirical evidence, the findings from the experiment suggest that average performance on the simple landscape ($K = 0$) is significantly higher than on the highly complex landscape ($K = 9$). Furthermore, while Billinger et al. (2014) found that average performance on the medially complex landscape ($K = 5$) was not significantly different than on the highly complex landscape, we found that average performance on the moderately complex landscape ($K = 3$) was not significantly different than on the simple landscape. Nonetheless, a regression analysis suggested that average performance does tend to decrease with the level of complexity. Still consistently with prior results from Billinger et al. (2014), we found that the level of complexity does not exert an immediate effect on search distance under independent problem-solving. Moreover, it does

not exert an immediate effect on search distance or imitation under parallel problem-solving either. We found also that human agents gradually adapt to individual feedback, so that success in discovering better-performing combinations prompts more exploitation, whereas failure prompts more exploration. Again, this applies to both the independent and the parallel problem-solving condition, thereby extending the validity of prior results to a different modality of problem-solving. In addition, the findings from the experiment highlighted that, under independent problem-solving, positive individual feedback significantly decreases with the level of complexity. Since the level of complexity influences individual feedback and, in turn, individual feedback impacts autonomous search behavior, complexity can be regarded as having an indirect effect on autonomous search behavior. However, the findings highlighted also that the negative association between complexity and individual feedback ceases to be significant under parallel problem-solving. Plausibly, the reason is to be found in the diffusion of better-performing combinations across networked groups that, in principle, is able to contrast the increasingly persistent uncertainty introduced by higher levels of complexity. The ability of information diffusion to contrast the negative association between complexity and individual feedback is especially apparent in the presence of potential benefits from crossover effects arising within the flawed information diffusion system that, relative to its flawless counterpart, has a significant positive impact on both individual and social feedback, thereby providing a long-run adaptive advantage. Taken in conjunction, these findings provide support to prior results only limited to the independent problem-solving condition, while adding further nuances to the model of adaptive search when the solution of problems unfolds in parallel.

The principal contribution of this experimental study stems from the investigation of the effects of parallel problem-solving on strategic diversity, collective performance, and imitative behavior in networked groups of human subjects. In this regard, consistently with H1a and H1b, we found that perfect imitation - or better to say, the flawless information diffusion system that characterizes an efficient network - accelerates the selection of better-performing strategies. Although resulting in a neat strategic convergence on a single solution, this still promotes a near-term adaptive advantage characterized by higher average performance over the short-run. On another note, consistently with H2a and H2b, we found that the crossover mechanisms behind imperfect imitation - or better to say, the flawed information diffusion system that characterizes an inefficient network - delays the propagation of better-performing combinations, which impairs the adaptive advantage over the short-run but allows to explore the problem space more extensively. In so doing, it preserves strategic diversity, a necessary condition for the seamless improvement of solutions, thereby promoting a long-term adaptive advantage that translates into a higher average performance over the long-run.

Moreover, within both the efficient and the inefficient networked group, we found that average performance on the simple landscape is significantly lower than on the moderately complex landscape that, in turn, is significantly lower than on the highly complex landscape. These findings suggest that, ultimately, consistently with H3, the performance benefits produced by perfect and imperfect imitation decrease with the level of complexity. Last but not least, we observed that imitative behavior does adapt to social feedback. In particular, subjects are inclined to imitate the members of the reference group who are more successful, while they prefer to rely on their own search ability if nobody outperforms them. Remarkably, the harmonization of "wise long-jumps" through the imitation of more successful others with the incremental enhancements of performance through autonomous adaptive search tends to improve as subjects gain experience. On top of that, consistently with H4, we found that, despite the potential benefits from crossover effects, within inefficient networks the attitude to imitate more successful others declines as with the strategic distance and, thence, with the risk of an error intervening in the imitative process. This phenomenon reflects the tension between the increased learning opportunities entailed by more differences and the subsequent decline in the success of imitation.

The study has several limitations. First, in the present design we only considered fully interconnected networks in which, potentially, each subject was connected with everybody else. However, a number of alternative communication structures - e.g., lattice, random, small-world - may be considered in order to assess their relative impact on the balance between individual and social learning and, consequently, on the performance of singles and collectives (Lazer & Friedman, 2007). Moreover, in the experiment subjects received information from the social environment in every trial. In contrast, the velocity of information diffusion may be manipulated decreasing the frequency with which subjects observe other participants to appreciate how this would affect, again, parallel problem-solving and the resulting individual and collective performance (Lazer & Friedman, 2007). In addition, our subjects received information concerning the performance of other participants and, in one condition, the strategic distance between themselves and the participants with whom they were put in communication, but not the relevant combinations. Nonetheless, the study of imitative behavior on rugged landscape may be further advanced considering not only more social learning conditions related to who should be imitated, such as the average rule and the majority rule, but also related to when imitation should happen or which element should be imitated (Laland, 2004). The when conditions are relevant if subjects have to evaluate any risk associated with a dynamic environment or any cost to access social information, while the which conditions are important if the disclosed solutions of those who can be imitated are similar to one's own or similar to each other (e.g., Laughlin & Ellis, 1986; Baron et al.,

1996; Rogers, 2003; Rubinstein, 2003; Knudsen, 2008; McElreath et al., 2008; Gureckis & Goldstone, 2009; Rendell et al., 2010; Wisdom et al., 2013).

3.5 Conclusions

We conducted a laboratory experiment in which the accuracy of the imitative process was manipulated to induce a flawless or flawed information diffusion system and, congruently, an efficient or inefficient communication network in combinatorial tasks shaped using the NK model and reflecting organizational settings featuring multiple dimensions and various levels of complexity. This is relevant for organizational science because imitating the practices of high-performing actors is one of the key strategies employed by organizations to solve complex problems and improve their performance (Nutt, 1998; Argote, 1999; Csaszar & Siggelkow, 2010), thereby representing a major part of the competitive process (Lieberman & Asaba, 2006). Ultimately, the question of under what circumstances social learning provides adaptive advantages over individual learning - that is, what their appropriate contexts are in the social world - goes to the very heart of adaptive rationality (Todd et al., 2012).

Our experiment provides empirically grounded insights that enrich the prior literature along different lines. To begin, the findings from the experiment extend the validity of the model of adaptive search (Billinger et al., 2014) from independent to parallel problem-solving. In this regard, they suggest that, when the solution of problems unfolds in parallel, autonomous search behavior still adapts to performance feedback. However, in a departure from the case in which subjects search the problem space as isolated monads, individual feedback ceases to be negatively associated with the level of complexity, while it becomes positively associated with the potential benefits from crossover effects arising within a flawed information diffusion system. Arguably, this is inherently tied with the diffusion of better-performing combinations across networked groups that, especially in the presence of crossover effects, is able to contrast the increasingly persistent uncertainty introduced by higher levels of complexity. What is more, our experiment provides the first laboratory validation of the causal chain behind the effects of efficient and inefficient networks in complex, multi-peaked problems as specified within the model of parallel problem-solving on rugged landscapes (Lazer & Friedman, 2007). Consistently with predictions, the findings from the experiment indicated that networks that are more efficient at disseminating information quickly converge on a single solution and perform well over the short-run but poorly over the long-run. Conversely, networks that are less efficient at disseminating information maintain strategic diversity and perform poorly over the short-run but well over the long-run. Moreover, the relevant findings extend the original model of parallel problem-solving on

multiple planes. In fact, besides providing support for the proclivity of human beings to act in compliance of the follow-the-successful heuristic (Boyd & Richerson, 2005), they show that the potential benefits from crossover effects have a positive influence on imitation. Most importantly, they expose the moderating effects of strategic distance within inefficient networks, according to which the attitude to imitate more successful others tends to decline as the number of differences with the targeted solution and, thence, with the risk of an error intervening in the imitative process. This finding is particularly meaningful in that it captures all the tension between the potential gains entailed by the capacity of differences to promote learning opportunities (Mill, 1848; Burt, 1992) as well as collective intelligence of a group (Hutchins, 1995; Page, 2007), thereby sowing the seeds for innovation, and the potential losses associated with the subsequent decline in the success of imitation, which is heavily influenced by the degree of similarity between the knowledge bases of organizations (Helfat, 1998; Lane & Lubatkin, 1998; Lieberman & Asaba, 2006). Notably, in line with the prior work of Arthur (1991), Edmonds (2001), and Billinger et al. (2014), the findings from our experiment have also value for further theoretical work in organizational research, since they offer evidence in favor of the calibration of empirically grounded search algorithms to simulate stylized behavioral rules within the computational modelling of adaptive processes.

In closing, we hope that this investigation will contribute to improve our understanding of the psychological and social mechanisms behind individual and collective decision-making in settings characterized by complexity and persistent uncertainty, thereby furthering the behavioral strategy agenda interested in bringing realistic assumptions into human behavior and, therewith, to foster the theoretical and empirical advancement of strategic organization.

4

Conclusions

Against the theoretical background of evolutionary behavioral economics, this project has investigated bounded rationality and adaptive behaviour in organizational settings characterized by complexity and persistent uncertainty. Embracing the principle that the development of theories on human rationality and behaviour requires an empirical foundation (e.g., Simon, 1956, 1978; Selten, 1995, 1998; Powell et al., 2011), two laboratory experiments were conducted to shed light on the determinants of individual and collective decision-making in combinatorial problems that, drawing upon the NK model, were operationalized as performance landscapes featuring multiple dimensions and various levels of complexity. The relevant choice structure is particularly apt to capture complex strategic activities as diverse as individual decision-making (e.g., Gavetti & Levinthal, 2000), organizational decision-making (e.g., Knudsen & Levinthal, 2007), new product development (e.g., Mihm et al., 2003), product modularity (e.g., Marengo & Dosi, 2005), open innovation (e.g., Almirall & Casadesus-Masanell, 2010), organizational design (e.g., Rivkin & Siggelkow, 2003), industry dynamics (e.g., Lenox et al., 2007), and more (see Ganco & Hoetker, 2009 for a comprehensive review).

In the first study, investment time horizons of different length were implemented to induce a near or distant future temporal orientation. It had three objectives: (1) assessing the effects of complexity and time horizon on performance and search behaviour; (2) examining the presence of a temporal midpoint heuristic; (3) inspecting the moderating effects of deadline proximity on the performance-risk relationship.

Besides providing further evidence in support of the behavioral model of adaptive search (Billinger et al., 2014), according to which a positive performance feedback (success) induces more risk-averse exploitation, while a negative performance feedback (failure) induces more risk-seeking exploration, the main empirical findings suggested that, independently from the level of complexity, a shorter time frame introduces a form of pressure in the investment of scarce resources, such that human agents are prone to explore the problem space more

extensively than they would do with a longer time frame. Although this is somewhat beneficial in terms of performance when the task is simple, the same attitude tends to be particularly detrimental when the task is complex. Search behaviour does not obey to the temporal midpoint heuristic, but is rather sensitive to the approaching of the end of the task. In particular, deadline proximity moderates the dynamics of the performance-risk relationship, so that underperforming subjects tend to shift the focus of attention from the reference point coinciding with one's own best performance in prior trials to survival, thereby opting for risk-averse exploitation, while outperforming subjects are prone to shift the focus of attention from the same reference point to experimenting with slack resources, thereby opting for risk-seeking exploration.

These findings provide empirically grounded insights that are instrumental to move a first step towards the integration of the temporal dimension into the model of adaptive search (Billinger et al., 2014). In particular, they highlight how time represents a unique type of resource able not only to induce more or less pressure in the process of resource allocation and, with this, to elicit a different trade-off between exploration and exploitation, but especially to trigger shifts in the focus of attention from the subjective reference point based on performance to either survival or experimenting with slack resources. Accordingly, they extend the model of adaptive search suggesting that, besides encoding performance feedback as success or failure relative to the reference point in question, subjects form a "cognitive image" of the future, which tends to be perceived as less uncertain as the end of the period approximates. Thence, underperformers - that is, subjects that systematically experienced a negative feedback - are more likely to visualize their end-of-period performance as underperforming, while outperformers - that is, subjects that systematically experienced a positive feedback - are more likely to visualize their end-of-period performance as outperforming. In line with the predictions from Lehman et al. (2011), the subsequent performance-risk relationship is U-shaped for the former and linear for the latter, thereby leading to the distinguishing dynamics described above. The present experiment offers the first validation of such phenomenon under laboratory conditions.

In the second study, the experimental framework allowed for parallel problem-solving based on the diffusion of social information through imitation. In particular, the accuracy of the imitative process was varied to induce a flawless or flawed information diffusion system and, congruently, an efficient or inefficient communication network. It had three objectives: (1) evaluating the degree of strategic diversity of networked groups and their subsequent performance conditional on the interaction between the levels of complexity and the modalities of parallel problem-solving; (2) appraising whether prior results on autonomous search behavior under independent problem-solving extends to autonomous

search behavior under parallel problem-solving; (3) shedding light on the determinants of imitative behavior, with a special attention to crossover effects, social feedback, and strategic distance.

The main empirical findings suggested that perfect imitation - or better to say, the flawless information diffusion system that characterizes an efficient network - accelerates the selection of better-performing strategies, resulting in a keen strategic convergence on a single solution and, thence, the leveling of performance in networked groups. Nonetheless, it promotes a near-term adaptive advantage characterized by a higher average performance over the short-run. In contrast, imperfect imitation - or better to say, the flawed information diffusion system that characterizes an inefficient network - delays the propagation of better-performing combinations, leading to a reduction in average performance over the short-run. However, the increased degree of strategic diversity engendered by the crossover mechanism consents to explore the space of the problem more extensively, thereby preserving strategic diversity and promoting a long-term adaptive advantage that translates into a higher average performance over the long-run. Still, the performance benefits produced by perfect as well as imperfect imitation decrease with the level of complexity. The findings suggested also that imitative behavior adapts to social performance feedback. In particular, subjects are inclined to follow the members of the reference group who are more successful, while they prefer to rely on their own search ability when there is nobody who outperforms them. Remarkably, the harmonization of "wise long-jumps" through the imitation of more successful others with the incremental enhancements of performance through autonomous adaptive search tends to improve as subjects gain experience. Despite the potential benefits from crossover effects, within inefficient networks the attitude to imitate more successful others declines as the strategic distance and, thence, the risk of an error intervening in the imitative process escalates.

These findings have several merits. In the first place, they extend the validity of the adaptive search model (Billinger et al., 2014) from independent to parallel problem-solving. In this regard, they suggest that, when the solution of problems unfolds in parallel, autonomous search behaviour still adapts to performance feedback. However, in departure from the case in which subjects search the problem space as isolated monades, individual feedback ceases to be negatively associated with the level of complexity, while it becomes positively associated with the potential benefits from crossover effects arising within a flawed information diffusion system. Arguably, this is inherently tied with the diffusion of better-performing combinations across networked groups that, especially in the presence of crossover effects, is able to contrast the increasingly persistent uncertainty introduced by higher levels of complexity. What is more, our findings provides first-time laboratory evidence in support of the causal

chain behind efficient and inefficient networks in complex, multi-peaked problems according to the model parallel problem-solving model (Lazer & Friedman, 2007). Moreover, they extend the original model of parallel problem-solving on multiple planes. In fact, besides providing empirical support for the proclivity of human beings to act in compliance of the follow-the-successful heuristic (Boyd & Richerson, 2005) as postulated by the model, they show that the potential benefits from crossover have a positive influence on imitation. Most importantly, they expose the moderating effects of strategic distance. This finding is particularly meaningful in that it captures all the tension between the possible gains entailed by the capacity of differences to promote learning opportunities (Mill 1848; Burt, 1992) as well collective intelligence of a group (Hutchins, 1995; Page, 2007), thereby sowing the seeds for innovation, and the possible losses associated with the subsequent decline in the success of imitation, which is heavily influenced by the degree of similarity between the knowledge bases of organizations (Helfat, 1998; Lane & Lubatkin, 1998; Lieberman & Asaba, 2006).

4.1 Future Directions

The study of time horizon effects is susceptible of several enhancements. The first regards the role of uncertainty about the future. In the real world, in fact, individuals and organizations are often not aware of the horizon constraining a choice sequence, in the sense that they do not know how often they will carry out a similar choice. Such a “shadow of the future” (Dal Bó, 2005) makes then extremely difficult for them weighting the costs and benefits of exploration and exploitation (Mehlhorn et al., 2015), thereby introducing a further complication to the process of resources allocation (Yaari, 1965; Hakansson 1969, 1971; Merton, 1971; Richard, 1975; Karatzas & Wang, 2001). Another possible extension is related to the moderating effect of deadline proximity in the presence of vicarious information, which would allow for the formation of an alternative reference point based on the performance of other subjects rather than on one’s own best-performing combination in prior trials (e.g., Knudsen, 2008), thereby influencing search behaviour (e.g., Greve, 2003).

A number of enhancements are also possible for the study of parallel problem-solving and information diffusion in networked groups. First, many different communication structures - e.g., lattice, random, small-world - may be considered in the interest of assessing their relative impact on the balance between individual and social learning and, consequently, on the performance of singles and collectives (Lazer & Friedman, 2007). Moreover, the velocity of information diffusion may be manipulated by decreasing the frequency with which subjects observe other participants to appreciate how this would affect, again, parallel problem-solving and the resulting individual and collective performance (Lazer & Friedman, 2007). The

study on imitative behaviour on rugged landscape may be further advanced considering not only more social learning conditions related to who should be imitated, but also related to when imitation should happen or which element should be imitated. The when conditions are relevant if subjects have to evaluate any risk associated with a dynamic environment or any cost to access social information, while the which conditions are important if the disclosed solutions of those who can be imitated are similar to one's own or similar to each other (e.g., Laughlin & Ellis, 1986; Baron et al., 1996; Rogers, 2003; Rubinstein, 2003; Laland, 2004; McElreath et al., 2008; Gureckis & Goldstone, 2009; Rendell et al., 2010; Wisdom et al., 2013).

Far beyond enhancing the present studies, the experimental investigation of individual and collective adaptation in complex environments holds the potential to open an array of future directions. Prime examples are the study of memory, mental representation of the external environment, neural correlates of random and direct exploration, as well as more conventional aspects such as modularity, vertical integration, or the implications of dynamic environments.

4.2 Final Remarks

This project fits into the research agenda of behavioral strategy, which aims to bring realistic assumptions about human cognition and social behavior to the strategic management of organizations, in a view to enhance strategy theory, empirical research, and real-world practice (Powell et al., 2011). Despite the rich body of research, most of the published work on strategic organization in complex environments has been explicitly theoretical, so that we still miss an appropriate understanding of how human beings actually behave in similar settings.

To address this shortcoming, the laboratory studies conducted in pursuance of the present project built on the laboratory framework introduced by Billinger et al. (2014) to shed light on several aspects of bounded rationality and adaptive behaviour in complex, combinatorial problems. Relative to other empirical methods, laboratory experiments provide full control over the environment, thereby consenting to rule out possible sources of endogeneity. As a result, they represent a privileged location to assess the empirical plausibility of behavioral rules and to test for model validity (Hey, 1982; Lave & March, 1975; Winter, 1982; Sterman, 1987; Selten, 1995, 1998; Lam, 2010).

The findings from the project are intended to provide empirical insights to improve our understanding of the psychological and social basis of individual and collective decision-making in settings characterized by complexity and persistent uncertainty. Furthermore, they

are expected to have theoretical implications for the development of empirically grounded search algorithms to simulate stylized behavioral rules within the computational modelling of adaptive processes. Ultimately, I hope the findings from the project will aid grounding human behaviour in more realistic assumptions, thereby contributing to furthering the behavioral strategy agenda and, over and above, to the establishment of a behavioral evolutionary economics.

References

- Abramson, L.Y., Garber, J., & Seligman, M.E. (1980). Learned helplessness in humans: An attributional analysis. In Garber, J., & Seligman, M.E. (Eds.), *Human helplessness: Theory and application* (pp. 3-34). New York: Academic Press.
- Allais, P. M. (1953). Behavior of rational man in risk situations: Critique of American school postulates and axioms. *Econometrica*, 21(4), 503-546.
- Almirall, E., & Casadesus-Masanell, R. (2010). Open versus closed innovation: A model of discovery and divergence. *Academy of Management Review*, 35(1), 27-47.
- Amabile, T.M., Mueller, J.S., Simpson, W.B., Hadley, C.N., Kramer, S.J., & Fleming, L. (2002). Time pressure and creativity in organizations: a longitudinal field study. HBS Working Paper 02-073.
- Ancona, D. G., Goodman, P. S., Lawrence, B. S., & Tushman, M. L. (2001a). Time: A new research lens. *Academy of Management Review*, 26(4), 645-663.
- Ancona, D. G., Okhuysen, G. A., & Perlow, L. A. (2001b). Taking time to integrate temporal research. *Academy of Management Review*, 26(4), 512-529.
- Anderson, P. W. (1994). *A career in Theoretical Physics*. Singapore: World Scientific.
- Argote, L. (1999). *Organizational Learning: Creating, Retaining, and Transferring Knowledge*. Boston: Kluwer Academic.
- Ariely, D. (2009). *Predictably Irrational*, Revised edition. London: Harper Collings Publishers.
- Arrow, K. J. (2004). Is bounded rationality unboundedly rational? Some ruminations. In M. Augier & J. G. March (Eds.), *Models of a Man: Essays in Memory of Herbert A. Simon* (pp. 47-55). Cambridge: MIT Press.

- Arrow, K. J., & Debreu, G. (1954). Existence of an equilibrium for a competitive economy. *Econometrica*, 22(3), 265-290.
- Arthur, W. B. (1991). Designing economic agents that act like human agents: A behavioral approach to bounded rationality. *The American Economic Review*, 81(2), 353-359.
- Audia, P. G., & Greve, H. R. (2006). Less likely to fail: Low performance, firm size, and factory expansion in the shipbuilding industry. *Management Science*, 52(1), 83-94.
- Axelrod, R. (1980). More effective choice in the prisoner's dilemma. *Journal of Conflict Resolution*, 24(3), 379-403.
- Axtel, R., Kirman, A., Couzin, I. D., Fricke, D., Hens, T., Hochberg, M. E., Mayfield, P. S., Schuster, P., & Sethi, R. (2016). Challenges of integrating complexity and evolution into economics. In D. S. Wilson & A. Kirman (Eds.), *Complexity and Evolution: Toward a New Synthesis for Economics* (pp. 65-81). Cambridge: MIT Press.
- Bak, P. (1996). *How Nature Works: The Science of Self-organized Criticality*. New York: Copernicus.
- Bandura, A. (1965). Vicarious processes: A case of no-trial learning. In L. Berkowitz (Ed.), *Advances in Experimental Social Psychology* (Vol. 2, pp. 1-55). New York: Academic Press.
- Banerjee, A. V. (1992). A simple model of herd behavior. *The Quarterly Journal of Economics*, 107(3), 797-817.
- Barabási, A. L., & Frangos, J. (2002). *Linked: The New Science of Networks Science of Networks*. New York: Perseus.
- Barkoczi, D., Analytis, P. P., & Wu, C. (2016). Collective search on rugged landscapes: A cross-environmental analysis. In Proceedings of the 38th annual conference of the cognitive science society (pp. 918-923). Cognitive Science Society.
- Baron, R. S., Vandello, J. A., & Brunsman, B. (1996). The forgotten variable in conformity research: The impact of task performance on social influence. *Journal of Personality & Social Psychology*, 71(5), 915-927.

-
- Baum, J. A., Rowley, T. J., Shipilov, A.V., & Chuang, Y. T. (2005). Dancing with strangers: Aspiration performance and the search for underwriting syndicate partners. *Administrative Science Quarterly*, 50(4), 536-575.
- Bebchuk, L. A., & Stole, L. A. (1993). Do short-term objectives lead to under-or overinvestment in long-term projects?. *The Journal of Finance*, 48(2), 719-729.
- Bettencourt, L. (2003). Tipping the balances of a small world. arXiv preprint cond-mat/0304321.
- Bikhchandani, S., Hirshleifer, D., & Welch, I. (1992). A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of Political Economy*, 100(5), 992-1026.
- Billinger, S., Srikanth, K. (in preparation). The effect of time compression on capability learning.
- Billinger, S., Stieglitz, N., & Schumacher, T. R. (2013). Search on rugged landscapes: An experimental study. *Organization Science*, 25(1), 93-108.
- Binmore, K. G. (1998). *Just playing: Game theory and the Social Contract*. Cambridge: MIT Press.
- Birkhoff G. D. (1927). *Dynamical Systems*. American Mathematical Society Colloquium Publications (Vol. 9). Providence: American Mathematical Society.
- Blonski, M., Ockenfels, P., & Spagnolo, G. (2011). Equilibrium selection in the repeated prisoner's dilemma: axiomatic approach and experimental evidence. *American Economic Journal, Microeconomics*, 3(3), 164-192.
- Bluedorn, A. C. (2002). *The Human Organization of Time: Temporal Realities and Experience*. Stanford: Stanford University Press.
- Bøllingtoft, A., & Ulhøi, J. P. (2005). The networked business incubator - leveraging entrepreneurial agency?. *Journal of Business Venturing*, 20(2), 265-290.
- Bower, J. L. (1970). Planning within the firm. *The American Economic Review*, 60(2), 186-194.

- Boyd, R., & Richerson, P. J. (1985). *Culture and the Evolutionary Process*. Chicago: University of Chicago Press.
- Boyd, R., & Richerson, P. J. (1988). An evolutionary model of social learning: The effects of spatial and temporal variation. In T. Zentall, & B. G. Galef (Eds.), *Social learning: Psychological and biological perspectives* (pp. 29-48). Hillsdale: Lawrence Erlbaum.
- Boyd, R., & Richerson, P. J. (1989). Social learning as an adaptation. *Lectures on Mathematics in the Life Sciences*, 20, 1-26.
- Boyd, R., & Richerson, P. J. (1995). Why does culture increase human adaptability? *Ethology and Sociobiology*, 16(2), 125-143.
- Boyd, R., & Richerson, P. J. (1996). Why culture is common, but cultural evolution is rare. *Proceedings of the British Academy*, 88, 77-93.
- Boyd, R., & Richerson, P. J. (2005). *The Origin and Evolution of Cultures*. New York: Oxford University Press.
- Brealey, R., & Myers, S. (1996). *Principles of Financial Management*. New York: McGraw-Hill.
- Bromiley, P. (1986). *Corporate capital investment: A behavioral approach*. Cambridge: Cambridge University Press.
- Bromiley, P. (1991). Testing a causal model of corporate risk taking and performance. *Academy of Management Journal*, 34(1), 37-59.
- Bromiley, P. (2004). *The Behavioral Foundations of Strategic Management*. New York: Wiley.
- Bromiley, P. (2010). Looking at prospect theory. *Strategic Management Journal*, 31(12), 1357-1370.
- Brunswik, E. (1957). Scope and aspects of the cognitive problem. In H. Gruber, K. R. Hammond, & R. Jessor (Eds.), *Contemporary Approaches to Cognition* (pp. 5-31). Cambridge: Harvard University Press.

-
- Bull, C., Schotter, A., & Weigelt, K. (1987). Tournaments and piece rates: An experimental study. *Journal of Political Economy*, 95(1), 1-33.
- Burnham, T. C., Lea, S. E. G., Bell, A. V., Gintis, H., Glimcher, P. W., Kurzban, R., Lades, L., McCabe, K., Panchanathan, K., Teschl, M., & Witt, U. (2016). Evolutionary behavioural economics. In D. S. Wilson & A. Kirman (Eds.), *Complexity and Evolution: Toward a New Synthesis for Economics* (pp. 113-144). Cambridge: MIT Press.
- Burt, R. S. (1992). *Structural Holes: The Social Structure of Competition*. Cambridge: Harvard University Press.
- Busemeyer, J. R. (1993). Violations of the speed-accuracy tradeoff relation. In O. Svenson & A.J. Maule (Eds.), *Time Pressure and Stress in Human Judgment and Decision Making* (pp. 181-193). Boston: Springer.
- Cagan, P. (1956), The monetary dynamics of hyperinflation. In M. Friedman (Ed.), *Studies in the Quantity Theory of Money* (pp. 25-117). Chicago: University of Chicago Press.
- Camerer, C., & Weigelt, K. (1993). Convergence in experimental double auctions for stochastically lived assets. In D. Friedman & J. Rust (Eds.), *The Double Auction Market: Theories, Institutions and Experimental Evaluations* (pp. 355-396). Redwood City: Addison-Wesley.
- Camerer, Colin F. & Loewenstein, George (2003). Behavioral economics: Past, present, future. In C. Camerer, G.Loewenstein, & M. Rabin (Eds.), *Advances in Behavioral Economics* (pp. 3-51). Princeton: Princeton University Press.
- Cameron, A. C., & Trivedi, P. K. (2009). *Microeconometrics Using Stata*. College Station: Stata press.
- Carley, K. (1991). A theory of group stability. *American Sociological Review*, 56(3), 331-354.
- Carp, S. (2003). Risk behavior and performance benchmarks. Unpublished doctoral dissertation, Stern School of Business, New York University, New York.
- Cartwright, M. L., & Littlewood, J. E. (1945). On Non-Linear Differential Equations of the Second Order: I. *Journal of the London Mathematical Society*, 1(3), 180-189.

- Charness, G., & Gencicot, G. (2009). Informal risk sharing in an infinite-horizon experiment. *The Economic Journal*, 119(537), 796-825.
- Charness, G., Gneezy, U., & Halladay, B. (2016). Experimental methods: Pay one or pay all. *Journal of Economic Behavior & Organization*, 131(PA), 141-150.
- Chen, D. L., Schonger, M., & Wickens, C. (2016). oTree - An open-source platform for laboratory, online, and field experiments. *Journal of Behavioral and Experimental Finance*, 9(1), 88-97.
- Chen, W. R., & Miller, K. D. (2007). Situational and institutional determinants of firms' R&D search intensity. *Strategic Management Journal*, 28(4), 369-381.
- Chua, J. H., Chrisman, J. J., & Bergiel, E. B. (2009). An agency theoretic analysis of the professionalized family firm. *Entrepreneurship Theory and Practice*, 33(2), 355-372.
- Clancey, W. J. (1987). The Knowledge Engineer as Student: Metacognitive Bases for Asking Good Questions (No. STAN-CS-87-1183). Stanford University, Dept of Computer Science.
- Cojuharenco, I., Patient, D., & Bashshur, M. R. (2011). Seeing the "forest" or the "trees" of organizational justice: Effects of temporal perspective on employee concerns about unfair treatment at work. *Organizational Behavior and Human Decision Processes*, 116(1), 17-31.
- Colander, D. (2016). Complexity, economics and workaday economic policy. In D. S. Wilson & A. Kirman (Eds.), *Complexity and Evolution: Toward a New Synthesis for Economics* (pp. 285-298). Cambridge: MIT Press.
- Conlisk, J. (1996). Why bounded rationality?. *Journal of Economic Literature*, 34(2), 669-700.
- Connelly, B. L., Tihanyi, L., Certo, S. T., & Hitt, M. A. (2010). Marching to the beat of different drummers: The influence of institutional owners on competitive actions. *Academy of Management Journal*, 53(4), 723-742.
- Conte, A., Scarsini, M., & Sürücü, O. (2016). The impact of time limitation: Insights from a queueing experiment. *Judgment and Decision Making*, 11(3), 260-274.

-
- Csaszar, F. A., & Siggelkow, N. (2010). How much to copy? Determinants of effective imitation breadth. *Organization Science*, 21(3), 661-676.
- Cyert, R. M., March, J. G. (1963). *A Behavioral Theory of the Firm*. Englewood Cliffs: Prentice-Hall.
- Dal Bó, P. D. (2005). Cooperation under the shadow of the future: experimental evidence from infinitely repeated games. *American Economic Review*, 95(5), 1591-1604.
- Dal Bó, P., & Frechette, G. (2011). The evolution of cooperation in infinitely repeated games: experimental evidence. *American Economic Review*, 101(1), 411-429.
- Das, T. K. (1987). Strategic planning and individual temporal orientation. *Strategic Management Journal*, 8(2), 203-209.
- Dasgupta, P., & Maskin, E. (2005). Uncertainty and hyperbolic discounting. *American Economic Review*, 95(4), 1290-1299.
- Dawkins, R. (2004). *The Ancestor's Tale*. London: Phoenix.
- De Paola, M., & Gioia, F. (2016). Who performs better under time pressure? Results from a field experiment. *Journal of Economic Psychology*, 53(53c), 37-53.
- Dechenaux, E., Kovenock, D., & Sheremeta, R. M. (2015). A survey of experimental research on contests, all-pay auctions and tournaments. *Experimental Economics*, 18(4), 609-669.
- Derex, M., & Boyd, R. (2016). Partial connectivity increases cultural accumulation within groups. *Proceedings of the National Academy of Sciences*, 113(11), 2982-2987.
- Devers, C. E., McNamara, G., Wiseman, R. M., & Arrfelt, M. (2008). Moving closer to the action: Examining compensation design effects on firm risk. *Organization Science*, 19(4), 548-566.
- DiMaggio, P. J., & Powell, W. W. (1983). The iron cage revisited: Institutional isomorphism and collective rationality in organizational fields. *American Sociological Review*, 48(2), 147-160.

- Dobzhansky, T. (1937). *Genetics and the Origin of Species*. New York: Columbia University Press.
- Dorfman, J., Shames, V. A., & Kihlstrom, J. F. (1996). Intuition, incubation, and insight: Implicit cognition. In G.D.M. Underwood (Ed.), *Implicit Cognition* (pp. 257-296). Oxford: Oxford University Press.
- Edland, A. (1993). The effects of time pressure on choices and judgments of candidates to a university program. In O. Svenson & A.J. Maule (Eds.), *Time Pressure and Stress in Human Judgment and Decision Making* (pp. 145-156). Boston: Springer.
- Edmonds, B. (2001). Towards a descriptive model of agent strategy search. *Computational Economics*, 18(1), 111-133.
- Eisenhardt, K. M., & Martin, J. A. (2000). Dynamic capabilities: what are they?. *Strategic Management Journal*, 21(10-11), 1105-1121.
- Ellickson, R. C. (1994). *Order Without Law: How Neighbors Settle Disputes*. Cambridge: Harvard University Press.
- Ellsberg, D. (1961). Risk, ambiguity, and the Savage axioms. *The Quarterly Journal of Economics*, 75(3), 643-669.
- Engle-Warnick, J., & Slonim, R. L. (2004). The evolution of strategies in a repeated trust game. *Journal of Economic Behavior & Organization*, 55(4), 553-573.
- Epstein, J. M. (2008). Why model?. *Journal of Artificial Societies and Social Simulation*, 11(4), 12.
- Ericksen, J., & Dyer, L. (2004). Right from the start: Exploring the effects of early team events on subsequent project team development and performance. *Administrative Science Quarterly*, 49(3), 438-471.
- Ericsson K. A, Staszewski J. (1988). Skilled memory and expertise: mechanisms of exceptional performance. In D.K.K. Kotovsky (Ed.), *Complex Information processing: The impact of Herbert A. Simon*. Hillsdale: Erlbaum.

-
- Fang, C., Lee, J., & Schilling, M. A. (2010). Balancing exploration and exploitation through structural design: The isolation of subgroups and organizational learning. *Organization Science*, 21(3), 625-642.
- Fehr, E., & Schmidt, K. M. (1999). A theory of fairness, competition, and cooperation. *The Quarterly Journal of Economics*, 114(3), 817-868.
- Feinberg, R. M., & Husted, T. A. (1993). An experimental test of discount-rate effects on collusive behaviour in duopoly markets. *The Journal of Industrial Economics*, 41(2), 153-160.
- Fisher, R.A. (1930). *The Genetical Theory of Natural Selection*. Oxford: Clarendon Press.
- Fleming, L., & Sorenson, O. (2004). Science as a map in technological search. *Strategic Management Journal*, 25(8-9), 909-928.
- Florida, R. (2002). *The Rise of the Creative Class: And How It's Transforming Work, Leisure, Community and Everyday Life*. New York: Basic Books.
- Ford E. B. (1940). Polymorphism and taxonomy. In J. Huxley (Ed.), *The New Systematics*. London: Oxford University Press.
- Ford, C., & Sullivan, D. M. (2004). A time for everything: How the timing of novel contributions influences project team outcomes. *Journal of Organizational Behavior*, 25(2), 279-292.
- Forlani, D. (2002). Risk and rationality: the influence of decision domain and perceived outcome control on managers' high-risk decisions. *Journal of Behavioral Decision Making*, 15(2), 125-140.
- Fouraker, L.E., & Siegel, S. (1963). *Bargaining Behavior*. McGraw-Hill: New York.
- Frederick, S., Loewenstein, G., & O'donoghue, T. (2002). Time discounting and time preference: A critical review. *Journal of Economic Literature*, 40(2), 351-401.
- Frenken, K. (2006). *Innovation, Evolution and Complexity Theory*. Cheltenham: Edward Elgar Publishing.

- Friedman, M., & Friedman, M. (1953). *Essays in Positive Economics*. Chicago: University of Chicago Press.
- Friedman, M., & Savage, L. J. (1948). The utility analysis of choices involving risk. *Journal of Political Economy*, 56(4), 279-304.
- Galbraith, C. S. (1990). Transferring core manufacturing technologies in high-technology firms. *California Management Review*, 32(4), 56-70.
- Ganco M., & Hoetker G. (2009). NK modeling methodology in the strategy literature: Bounded search on a rugged landscape. In D. D. Bergh, D. J. Ketchen (Ed.), *Research Methodology in Strategy and Management*, (Vol. 5, pp. 237-268). Bingley: Emerald Group Publishing.
- Gavetti, G. (2005). Cognition and hierarchy: Rethinking the microfoundations of capabilities' development. *Organization Science*, 16(6), 599-617.
- Gavetti, G., & Levinthal, D. (2000). Looking forward and looking backward: Cognitive and experiential search. *Administrative Science Quarterly*, 45(1), 113-137.
- George, A. L. (1972). The case for multiple advocacy in making foreign policy. *American Political Science Review*, 66(3), 751-785.
- George, J. M., & Jones, G. R. (2000). The role of time in theory and theory building. *Journal of Management*, 26(4), 657-684.
- Gersick, C. J. (1988). Time and transition in work teams: Toward a new model of group development. *Academy of Management Journal*, 31(1), 9-41.
- Gersick, C. J. (1989). Marking time: Predictable transitions in task groups. *Academy of Management Journal*, 32(2), 274-309.
- Gersick, C.J., & Hackman, J.R. (1990). Habitual routines in task-performing groups. *Organizational Behavior and Human Decision Processes*, 47(1), 65-97.
- Gersick, C. J. (1994). Pacing strategic change: The case of a new venture. *Academy of Management Journal*, 37(1), 9-45.

-
- Gevers, J. M., Rutte, C. G., & Van Eerde, W. (2006). Meeting deadlines in work groups: Implicit and explicit mechanisms. *Applied Psychology*, 55(1), 52-72.
- Gevers, J., Mohammed, S., & Baytalskaya, N. (2015). The conceptualisation and measurement of pacing styles. *Applied Psychology*, 64(3), 499-540.
- Ghemawat, P., & Levinthal, D. (2008). Choice interactions and business strategy. *Management Science*, 54(9), 1638-1651.
- Gibbons, R., (1998). Incentives in organizations. *Journal of Economic Perspectives*, 12(4), 115-132.
- Gibbons, R., & Murphy, K. J. (1990). Relative performance evaluation for chief executive officers. *Industrial and Labor Relations Review*, 43(3), 30-51.
- Gigerenzer, G. (2000). *Adaptive Thinking: Rationality in the Real World*. Oxford: Oxford University Press.
- Gigerenzer, G. (2004). Striking a blow for sanity in theories of rationality. In M. Augier & J. G. March (Eds.), *Models of a Man: Essays in Memory of Herbert A. Simon* (pp. 389-409). Cambridge: MIT Press.
- Gigerenzer, G., & Goldstein, D. G. (1996). Reasoning the fast and frugal way: models of bounded rationality. *Psychological Review*, 103(4), 650.
- Gigerenzer, G., & Selten, R. (2002). *Bounded Rationality: The Adaptive Toolbox*. Cambridge: MIT press.
- Gigerenzer, G., & Todd, P. M. (1999). *Simple Heuristics that Make Us Smart*. New York: Oxford University Press.
- Gintis, H. (2011). Gene-culture coevolution and the nature of human sociality. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 366(1566), 878-888.
- Goldstein, D. G., & Gigerenzer, G. (2002). Models of ecological rationality: the recognition heuristic. *Psychological Review*, 109(1), 75-90.

- Gonzalez, C., & Dutt, V. (2011). Instance-based learning: Integrating sampling and repeated decisions from experience. *Psychological Review*, 118(4), 523-551.
- Gonzalez, C., Lerch, J.F., & Lebiere, C. (2003). Instance-based learning in dynamic decision making. *Cognitive Science*, 27(4), 591-635.
- Gould, S. J. (1977). *Ontogeny and Phylogeny*. Cambridge: Harvard University Press.
- Gowdy, J. Mazzucato, M., van den Bergh, J. C. (2016). Shaping the evolution of complex societies. In D. S. Wilson & A. Kirman (Eds.), *Complexity and Evolution: Toward a New Synthesis for Economics* (pp. 327-350). Cambridge: MIT Press.
- Granovetter, M. (1978). Threshold models of collective behavior. *American Journal of Sociology*, 83(6), 1420-1443.
- Greve, H. R. (1998). Performance, aspirations, and risky organizational change. *Administrative Science Quarterly*, 43(1), 58-86.
- Greve H.R. (2003). *Organizational Learning from Performance Feedback*. Cambridge: Cambridge University Press.
- Gross, P.W., & Lewis, W.W. (2007). Built to last: Focusing corporations on long-term performance. Arlington: Research and Policy Committee of the Committee for Economic Development.
- Gupta, A. K., Smith, K. G., & Shalley, C. E. (2006). The interplay between exploration and exploitation. *Academy of Management Journal*, 49(4), 693-706.
- Gureckis, T. M., & Goldstone, R. L. (2009). How you named your child: Understanding the relationship between individual decision making and collective outcomes. *Topics in Cognitive Science*, 1(4), 651-674.
- Güth, W. (2002). Interview with Werner Güth. *Experimental Economics: Financial Markets, Auctions, and Decision Making*, 19-28.
- Hackenberg, T. D., & Axtell, S. A. (1993). Human's choices in situations of time-based diminishing returns. *Journal of the Experimental Analysis of Behavior*, 59(3), 445-470.

-
- Hakansson, N. H. (1969). Optimal investment and consumption strategies under risk, an uncertain lifetime, and insurance. *International Economic Review*, 10(3), 443-466.
- Hakansson, N. H. (1971). Optimal entrepreneurial decisions in a completely stochastic environment. *Management Science*, 17(7), 427-449.
- Hamming, R.W. (1950). Error detecting and error correcting codes. *The Bell System Technical Journal*, 29(2), 147-160.
- Hammond, K. R. (1996). *Human Judgment and Social Policy: Irreducible Uncertainty, Inevitable Error, Unavoidable Injustice*. New York: Oxford University Press.
- Harris, M. (1979). *Cultural Materialism: The Struggle for a Science of Culture*. New York: Random House.
- Hayes, R. H., & Abernathy, W. J. (1980). Managing our way to economic decline. *Harvard Business Review*, 58(4), 66-77.
- Helfat, C. E. (1998). Simple indicators of adaptation versus rigidity in history-dependent firm activities and decision rules. *Industrial and Corporate Change*, 7(1), 49-75.
- Hertwig, R., & Engel, C. (2016). Homo ignorans: Deliberately choosing not to know. *Perspectives on Psychological Science*, 11(3), 359-372.
- Hertwig, R., & Erev, I. (2009). The description-experience gap in risky choice. *Trends in Cognitive Sciences*, 13(12), 517-523.
- Hertwig, R., & Herzog, S. M. (2009). Fast and frugal heuristics: Tools of social rationality. *Social Cognition*, 27(5), 661-698.
- Hertwig, R., Hoffrage, U. E., & ABC Group (2013). *Simple Heuristics in a Social World*. New York: Oxford University Press.
- Hertwig, R., & Pedersen, A. P. (2016). Finding Foundations for Bounded and Adaptive Rationality. *Minds & Machines*, 1(26), 1-8.

- Hertwig, R., Wulff, D. U., & Mata, R. (2019). Three gaps and what they may mean for risk preference. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 374(1766), 20180140.
- Hey J. D. (1987). Still searching. *Journal of Economic Behavioral Organization*, 8(1), 137-144.
- Hicks, J. R. (1939). The foundations of welfare economics. *The Economic Journal*, 49(196), 696-712.
- Hofer, C. W. & Schendel, D., (1978). *Strategy Formulation: Analytical Concepts*. St. Paul: West Publishing.
- Holland, J. (1975). *Adaptation in Natural and Artificial Systems*. Ann Arbor: University of Michigan Press.
- Holland, J. H. (2006). Studying complex adaptive systems. *Journal of Systems Science and Complexity*, 19(1), 1-8.
- Holliday, R. (1990). DNA methylation and epigenetic inheritance. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 326(1235), 329-338.
- Holt, C. A. (1985). An Experimental Test of the Consistent-Conjectures Hypothesis. *American Economic Review*, 75(3), 314-25.
- Hutchins, E. (1995). *Cognition in the Wild*. Cambridge: MIT Press.
- Hutchinson, J.M., Wilke, A., & Todd, P.M. (2008). Patch leaving in humans: can a generalist adapt its rules to dispersal of items across patches?. *Animal Behaviour*, 75(4), 133-1349.
- Huxley, J. S. (1942). *Evolution: The Modern Synthesis*. London: Allen and Unwin.
- Ijiri, Y., & Simon, H. A. (1977). *Skew Distributions and the Sizes of Business Firms*. Amsterdam: North Holland.
- Iyer, D. N., & Miller, K. D. (2008). Performance feedback, slack, and the timing of acquisitions. *Academy of Management Journal*, 51(4), 808-822.

-
- Jacob, F. (1977). Evolution and tinkering. *Science*, 196(4295), 1161-1166.
- Jacobs, M.T. (1991). Short-term America: The causes and cures of our business myopia. Harvard Business School Proceedings.
- Janis, I. L. (1972). *Victims of Groupthink*. Boston: Houghton Mifflin.
- Jensen, G. F., & Erickson, M. L. (1979). The religious factor and delinquency: Another look at the hellfire hypothesis. In R. Wenthow (Ed.), *The religious dimension: New directions in quantitative research* (pp. 157-177). New York: Academic Press.
- Johns, G. (2006). The essential impact of context on organizational behavior. *Academy of Management Review*, 31(2), 386-408.
- Joireman, J., Kamdar, D., Daniels, D., & Duell, B. (2006). Good citizens to the end? It depends: empathy and concern with future consequences moderate the impact of a short-term time horizon on organizational citizenship behaviors. *Journal of Applied Psychology*, 91(6), 1307-1320.
- Kaelbling, L. P., Littman, M. L., & Moore, A.W. (1996). Reinforcement learning: A survey. *Journal of Artificial Intelligence Research*, 4(11,28), 237-285.
- Kahneman, D. (2011). *Thinking, Fast and Slow*. New York: Farrar, Strauss and Giroux.
- Kahneman, D., Slovic, P., & Tversky, A. (1982). *Judgment Under Uncertainty: Heuristics and Biases*. New York: Cambridge University Press.
- Kameda, T., Tsukasaki, T., Hastie, R., & Berg, N. (2011). Democracy under uncertainty: The wisdom of crowds and the free-rider problem in group decision making. *Psychological Review*, 118(1), 76-96.
- Karatzas, I., & Wang, H. (2000). Utility maximization with discretionary stopping. *SIAM Journal on Control and Optimization*, 39(1), 306-329.
- Karpicke, J. D., & Roediger, H. L. (2008). The critical importance of retrieval for learning. *Science*, 319(5865), 966-968.

- Katz, S. H., M. L. Hediger, & L. A. Valleroy. (1974). Traditional maize processing techniques in the New World. *Science*, 184(4138), 765-773.
- Kauffman, S. A. (1993). *The origins of order: Self-organization and selection in evolution*. New York: Oxford University Press.
- Kauffman, S. A., & Weinberger, E. D. (1989). The NK model of rugged fitness landscapes and its application to maturation of the immune response. *Journal of Theoretical Biology*, 141(2), 211-245.
- Kauffman, S. A. (1995). *At Home in the Universe: The Search for the Laws of Self-Organization and Complexity*. New York: Oxford University Press.
- Kenrick, D. T., Li, Y. J., White, A. E., & Neuberg, S. L. (2012). Economic subselves: Fundamental motives and deep rationality. In J. P. Forgas, K. Fiedler, C. Sedikides (Eds.), *Social Thinking and Interpersonal Behavior* (pp. 23-44). London & New York: Taylor and Francis.
- Kimura, M. (1968). Evolutionary rate at the molecular level. *Nature*, 217(5129), 624-626.
- King, J. L., & Jukes, T. H. (1969). Non-darwinian evolution. *Science*, 164(3881), 788-798.
- Kirman, A., & Sethi, R. (2016). Disequilibrium adjustment and economic outcomes. In D. S. Wilson & A. Kirman (Eds.), *Complexity and Evolution: Toward a New Synthesis for Economics* (pp. 15-30). Cambridge: MIT Press.
- Knight, A.P. (2015). Mood at the midpoint: Affect and change in exploratory search over time in teams that face a deadline. *Organization Science*, 26(1), 99-118.
- Knudsen, T. (2008). Reference groups and variable risk strategies. *Journal of Economic Behavior & Organization*, 66(1), 22-36.
- Knudsen, T., & Levinthal, D. A. (2007). Two faces of search: Alternative generation and alternative evaluation. *Organization Science*, 18(1), 39-54.
- Kogut, B. (1988). Joint ventures: Theoretical and empirical perspectives. *Strategic Management Journal*, 9(4), 319-332.

-
- Kolmogorov, A. N. (1954). On Conservation of Conditionally Periodic Motions for a Small Change in Hamilton's Function. In Dokl. Akad. Nauk SSSR, 98, 527-530.
- Koop, G. J., & Johnson, J. G. (2012). The use of multiple reference points in risky decision making. *Journal of Behavioral Decision Making*, 25(1), 49-62.
- Krackhardt, D. (2001). Viscosity models and the diffusion of controversial innovation. *Dynamics of Organizations: Computational Modeling and Organization Theories*, 22(2), 268-273.
- Krueger, A. O. (1974). The political economy of the rent-seeking society. *American Economic Review*, 64(3), 291-303.
- Krueger, J. I., & Funder, D. C. (2004). Toward a balanced social psychology: Causes, consequences, and cures for the problem-seeking approach to social behavior and cognition. *Behavioral and Brain Sciences*, 27(3), 313-327.
- Kruglanski, A. W., & Freund, T. (1983). The freezing and unfreezing of lay-inferences: Effects on impression primacy, ethnic stereotyping, and numerical anchoring. *Journal of Experimental Social Psychology*, 19(5), 448-468.
- Kunisch, S., Bartunek, J. M., Mueller, J., & Huy, Q. N. (2017). Time in strategic change research. *Academy of Management Annals*, 11(2), 1005-1064.
- Laland, K. N. (2004). Social learning strategies. *Learning & Behavior*, 32(1), 4-14.
- Lam, S. Y. (2010). What kind of assumptions need to be realistic and how to test them: A response to Tsang (2006). *Strategic Management Journal*, 31(6), 679-687.
- Lane, P. J., & Lubatkin, M. (1998). Relative absorptive capacity and interorganizational learning. *Strategic Management Journal*, 19(5), 461-477.
- Langer, E. J. (1975). The illusion of control. *Journal of Personality and Social Psychology*, 32(2), 311.
- Langley, A. N. N., Smallman, C., Tsoukas, H., & Van de Ven, A. H. (2013). Process studies of change in organization and management: Unveiling temporality, activity, and flow. *Academy of Management Journal*, 56(1), 1-13.

- Langley, P., & Simon, H. A. (1981). The central role of learning in cognition. In J. Anderson (Ed.), *Cognitive Skills and Their Acquisition* (pp. 361-380). Hillsdale: Lawrence Erlbaum Associates.
- Laughlin, P. R., & Ellis, A. L. (1986). Demonstrability and social combination processes on mathematical intellectual tasks. *Journal of Experimental Social Psychology*, 22(3), 177-189.
- Lave C. A., & March J. G. (1975). *An Introduction to Models in the Social Sciences*. New York: Harper & Row.
- Laverty, K. J. (1996). Economic “short-termism”: The debate, the unresolved issues, and the implications for management practice and research. *Academy of Management Review*, 21(3), 825-860.
- Lazear, E. P. (1999). Personnel economics: Past lessons and future directions. *Journal of Labor Economics*, 17(2), 199-236.
- Lazear, E. P., & Rosen, S. (1981). Rank-order tournaments as optimum labor contracts. *Journal of Political Economy*, 89(5), 841-864.
- Lazer, D., & Friedman, A. (2007). The network structure of exploration and exploitation. *Administrative science quarterly*, 52(4), 667-694.
- Lee, M.D., Zhang, S., Munro, M., & Steyvers, M. (2011). Psychological models of human and optimal performance in bandit problems. *Cognitive Systems Research*, 12(2), 164-174.
- Lehman, D.W., Hahn, J., Ramanujam, R., & Alge, B. J. (2011). The dynamics of the performance-risk relationship within a performance period: The moderating role of deadline proximity. *Organization Science*, 22(6), 1613-1630.
- Lei, V., & Noussair, C.N. (2002). An experimental test of an optimal growth model. *American Economic Review*, 92(3), 549-570.
- Lenox, M. J., Rockart, S. F., & Lewin, A. Y. (2007). Interdependency, competition, and industry dynamics. *Management Science*, 53(4), 599-615.

-
- Lenox, M. J., Rockart, S. F., & Lewin, A. Y. (2010). Does interdependency affect firm and industry profitability? An empirical test. *Strategic Management Journal*, 31(2), 121-139.
- Levinthal, D. A. (1997). Adaptation on rugged landscapes. *Management Science*, 43(7), 934-950.
- Levinthal, D. A., & March, J. G. (1981). A model of adaptive organizational search. *Journal of Economic Behavior & Organization*, 2(4), 307-333.
- Levinthal, D. A., & March, J. G. (1993). The myopia of learning. *Strategic Management Journal*, 14(S2), 95-112.
- Lewin, K. (1943). Defining the 'field at a given time'. *Psychological Review*, 50(3), 292-310.
- Lewis, E. B. (1978). A gene complex controlling segmentation in *Drosophila*. In H.D. Lipshitz (Ed.), *Genes, Development and Cancer* (pp. 205-217). Boston: Springer.
- Li, T.Y.; Yorke, J.A. (1975). Period three implies chaos. *American Mathematical Monthly*. 82(10), 985-92.
- Libchaber, A., Laroche, C., & Fauve, S. (1982). Period doubling cascade in mercury, a quantitative measurement. *Journal de Physique Lettres*, 43(7), 211-216.
- Liberman, N., & Trope, Y. (1998). The role of feasibility and desirability considerations in near and distant future decisions: A test of temporal construal theory. *Journal of Personality and Social Psychology*, 75(1), 5.
- Lieberman, M. B., & Asaba, S. (2006). Why do firms imitate each other?. *Academy of Management Review*, 31(2), 366-385.
- Loewenstein, G., & Prelec, D. (1992). Anomalies in intertemporal choice: Evidence and an interpretation. *The Quarterly Journal of Economics*, 107(2), 573-597.
- Logan, G. D. (1988). Toward an instance theory of automatization. *Psychological Review*, 95(4), 492-527.

- Lopes, L. L. (1987). Between hope and fear: The psychology of risk. In L. Berkowitz (Ed.), *Advances in Experimental Social Psychology* (Vol. 20, pp. 255-295). San Diego: Academic Press.
- Lorenz, E. N. (1963). Deterministic nonperiodic flow. *Journal of the Atmospheric Sciences*, 20(2), 130-141.
- Lorenz, E. N. (1972). Predictability: Does the Flap of a Butterfly's Wings in Brazil Set Off a Tornado in Texas?. 139th Annual Meeting of the American Association for the Advancement of Science, in E.N. Lorenz (Ed.), *Essence of Chaos* (Appendix 1).
- Lumpkin, G. T., & Brigham, K. H. (2011). Long-term orientation and intertemporal choice in family firms. *Entrepreneurship Theory and Practice*, 35(6), 1149-1169.
- Mahajan, V., Muller, E., & Bass, F. M. (1990). New product diffusion models in marketing: A review and directions for research. *Journal of Marketing*, 54(1), 1-26.
- March, J. G. (1988). Variable risk preferences and adaptive aspirations. *Journal of Economic Behavior & Organization*, 9(1), 5-24.
- March, J. G. (1991). Exploration and exploitation in organizational learning. *Organization Science*, 2(1), 71-87.
- March, J. G., & Shapira, Z. (1987). Managerial perspectives on risk and risk taking. *Management Science*, 33(11), 1404-1418.
- March, J. G., & Shapira, Z. (1992). Variable risk preferences and the focus of attention. *Psychological Review*, 99(1), 172.
- Marengo, L., & Dosi, G. (2005). Division of labor, organizational coordination and market mechanism in collective problem-solving. *Journal of Economic Behavior and Organization*, 58(2), 303-326.
- Marginson, D., & McAulay, L. (2008). Exploring the debate on short-termism: A theoretical and empirical analysis. *Strategic Management Journal*, 29(3), 273-292.
- Marimon, R., and Sunder, S. (1993). Indeterminacy of Equilibria in a Hyperinflationary World: Experimental evidence, *Econometrica*, 61(5), 1073-1107.

-
- Maritan, C. A. (2001). Capital investment as investing in organizational capabilities: An empirically grounded process model. *Academy of Management Journal*, 44(3), 513-531.
- Markowitz, H. (1952). The utility of wealth. *Journal of Political Economy*, 60(2), 151-158.
- Mas-Colell, A., Whinston, M. D., & Green, J. R. (1995). *Microeconomic Theory*. New York: Oxford University Press.
- Mason, W. A., Jones, A., & Goldstone, R. L. (2008). Propagation of innovations in networked groups. *Journal of Experimental Psychology: General*, 137(3), 422-433.
- Mason, W., & Watts, D. J. (2012). Collaborative learning in networks. *Proceedings of the National Academy of Sciences*, 109(3), 764-769.
- Mayfield, J. E. (2016). Evolution and market complexity. In D. S. Wilson & A. Kirman (Eds.), *Complexity and Evolution: Toward a New Synthesis for Economics* (pp. 47-63). Cambridge: MIT Press.
- Maynard Smith, J., & Szathmáry, E. (1999). *The Origins of Life: From the Birth of Life to the Origin of Language*. Oxford: Oxford University Press.
- Mayseless, O., & Kruglanski, A. W. (1987). What makes you so sure? Effects of epistemic motivations on judgmental confidence. *Organizational Behavior and Human Decision Processes*, 39(2), 162-183.
- Mazzuccato, M. (2016). Innovation policy as creating markets, not only fixing them: Implications for complexity theory. In D. S. Wilson & A. Kirman (Eds.), *Complexity and Evolution: Toward a New Synthesis for Economics* (pp. 271-184). Cambridge: MIT Press.
- McElreath, R., Lubell, M., Richerson, P., Waring, T., Baum, W., Edsten, E., Efferson, C., & Paciotti, B. (2005). Applying evolutionary models to the laboratory study of social learning. *Evolution and Human Behavior*, 26(6), 483-508.
- McEvoy, L., & Land, G. (1981). Life-style and death patterns of the Missouri RLDS church members. *American Journal of Public Health*, 71(12), 1350-1357.

- McGrath, J. E., & Rotchford, N. L. (1983). Time and behavior in organizations. In B. Staw & L. Cummings (Eds.), *Research in Organizational Behavior* (Vol. 5, pp. 57-101). Greenwich: JAI Press.
- McGrath J. E., & Tschan F. (2004). *Temporal Matters in Social Psychology: Examining the Role of Time in the Lives of Groups and Individuals*. Washington: American Psychology Association.
- Mehlhorn, K., Newell, B.R., Todd, P.M., Lee, M.D., Morgan, K., Braithwaite, V.A., Hausmann, D., Fiedler, K., & Gonzalez, C. (2015). Unpacking the exploration-exploitation tradeoff: A synthesis of human and animal literatures. *Decision*, 2(3), 191-215.
- Merton, R.C. (1971). Optimal consumption and portfolio rules in a continuous-time model, *Journal of Economic Theory*, 3(4), 373-413.
- Mesoudi, A. (2008). An experimental simulation of the “copy-successful-individuals” cultural learning strategy: Adaptive landscapes, producer–scrounger dynamics, and informational access costs. *Evolution and Human Behavior*, 29(5), 350-363.
- Mihm, J., Loch, C., & Huchzermeier, A. (2003). Problem-solving oscillations in complex engineering projects. *Management Science*, 49(6), 733-750.
- Mill, J. S. (2004). Principles of political economy. Prometheus Books, Amherst, NY. (Original work published 1848)
- Miller, K. D. (2002). Knowledge inventories and managerial myopia. *Strategic Management Journal*, 23(8), 689-706.
- Miller, K. D., & Chen, W. R. (2004). Variable organizational risk preferences: Tests of the March-Shapira model. *Academy of Management Journal*, 47(1), 105-115.
- Miller, J. H., & Page, S. E. (2007). *Complex Adaptive Systems: An Introduction to Computational Models*. Princeton: Princeton University Press.
- Miller, N. E., & Dollard, J. (1941). *Social Learning and Imitation*. New Haven: Yale University Press.
- Mitchell, M. (1996). *An Introduction to Genetic Algorithms*. Cambridge: MIT Press.

-
- Mowery, D. C., & Sampat, B. N. (2001). Patenting and licensing university inventions: Lessons from the history of the research corporation. *Industrial and Corporate Change*, 10(2), 317-355.
- Murray, H. A. (1938). *Explorations in Personality: A Clinical and Experimental Study of Fifty Men of College Age*. New York: Oxford University Press.
- Nemeth, C. J. (1985). Dissent, group process, and creativity. In E. J. Lawler (Eds.), *Advances in Group Processes* (Vol. 2, pp. 57–75). Greenwich: JAI Press.
- Nohria, N., & Gulati, R. (1996). Is slack good or bad for innovation?. *Academy of Management Journal*, 39(5), 1245-1264.
- Nowak, M. A., Tarnita, C. E., & Wilson, E. O. (2010). The evolution of eusociality. *Nature*, 466(7310), 1057-1062.
- Nüsslein-Volhard, C., & Wieschaus, E. (1980). Mutations affecting segment number and polarity in *Drosophila*. *Nature*, 287(5785), 795-801.
- Nutt, P. C. (1998). How decision makers evaluate alternatives and the influence of complexity. *Management Science*, 44(8), 1148-1166.
- Offerman, T., Potters, J., & Verbon, H. A. (2001). Cooperation in an overlapping generations experiment. *Games and Economic Behavior*, 36(2), 264-275.
- Okhuysen, G. A., Galinsky, A. D., & Uptigrove, T. A. (2003). Saving the worst for last: The effect of time horizon on the efficiency of negotiating benefits and burdens. *Organizational Behavior and Human Decision Processes*, 91(2), 269-279.
- Okhuysen, G. A., & Waller, M. J. (2002). Focusing on midpoint transitions: An analysis of boundary conditions. *Academy of Management Journal*, 45(5), 1056-1065.
- Ostrom, E. (1991). *Governing the commons: The evolution of institutions for collective action*. Cambridge: Cambridge University Press.
- Ounjian, M. L., & Carne, E. B. (1987). A study of the factors which affect technology transfer in a multilocation multibusiness unit corporation. *IEEE Transactions on Engineering Management*, 34(3), 194-201.

- Padgett, J. F. (2016). Evolvability of organizations and institutions. In D. S. Wilson & A. Kirman (Eds.), *Complexity and Evolution: Toward a New Synthesis for Economics* (pp. 185-199). Cambridge: MIT Press.
- Page, S. E. (1996). Two measures of difficulty. *Economic Theory*, 8(2), 321-346.
- Page, S. E. (2007). *Prologue to The Difference: How the Power of Diversity Creates Better Groups, Firms, Schools, and Societies*. Princeton: Princeton University Press.
- Page, S. E. (2016). Not half bad: A modest criterion for inclusion. In D. S. Wilson & A. Kirman (Eds.), *Complexity and Evolution: Toward a New Synthesis for Economics* (pp. 319-326). Cambridge: MIT Press.
- Palfrey, T. R., & Rosenthal, H. (1994). Repeated play, cooperation and coordination: An experimental study. *The Review of Economic Studies*, 61(3), 545-565.
- Pareto, V. (1971). *Manual of Political Economy*. New York: A. M. Kelley. (Original work published 1906)
- Pettigrew, A. (1985). *The Awakening Giant: Continuity and change in Imperial Chemical Industries*. New York: Blackwell.
- Pitariu, A. H., & Ployhart, R. E. (2010). Explaining change: Theorizing and testing dynamic mediated longitudinal relationships. *Journal of Management*, 36(2), 405-429.
- Plambeck, E. L., & Wang, Q. (2013). Implications of hyperbolic discounting for optimal pricing and scheduling of unpleasant services that generate future benefits. *Management Science*, 59(8), 1927-1946.
- Ployhart, R. E., & Vandenberg, R. J. (2010). Longitudinal research: The theory, design, and analysis of change. *Journal of Management*, 36(1), 94-120.
- Porter, M. E. (1991). Towards a dynamic theory of strategy. *Strategic Management Journal*, 12(S2), 95-117.
- Porter, M. E. (1992). Capital disadvantage: America's failing capital investment system. *Harvard Business Review*, 70(5), 65-82.

-
- Porter, M. E. (1996). What is strategy?. *Harvard Business Review*, 74(6), 61-78.
- Powell, T. C., Lovallo, D., & Fox, C. R. (2011). Behavioral strategy. *Strategic Management Journal*, 32(13), 1369-1386.
- Prendergast, C. (1999). The provision of incentives in firms. *Journal of Economic Literature*, 37(1), 7-63.
- Prigogine, I. (1962). *Non-equilibrium Statistical Mechanics*. New York: Wiley.
- Quintens, L., & Matthyssens, P. (2010). Involving the process dimensions of time in case-based research. *Industrial Marketing Management*, 39(1), 91-99.
- Quirin, G. D., & Wiginton, J. C. (1981). *Analyzing Capital Expenditures: Private and Public Perspectives*. Washington: RD Irwin.
- Rabin, M. (1993). Incorporating fairness into game theory and economics. *The American Economic Review*, 83(5), 1281-1302.
- Rabin, M. (1998). Psychology and economics. *Journal of Economic Literature*, 36(1), 11-46.
- Rajagopalan, N., & Spreitzer, G. M. (1997). Toward a theory of strategic change: A multi-lens perspective and integrative framework. *Academy of Management Review*, 22(1), 48-79.
- Reilly, G., Souder, D., & Ranucci, R. (2016). Time horizon of investments in the resource allocation process: Review and framework for next steps. *Journal of Management*, 42(5), 1169-1194.
- Rendell, L., Boyd, R., Cownden, D., Enquist, M., Eriksson, K., Feldman, M. W., Fogarty, L., Ghirlanda, S., Lillicrap, T., & Laland, K. N. (2010). Why copy others? Insights from the social learning strategies tournament. *Science*, 328(5975), 208-213.
- Richard, S. F. (1975). Optimal consumption, portfolio and life insurance rules for an uncertain lived individual in a continuous time model. *Journal of Financial Economics*, 2(2), 187-203.

- Richerson P. J., & Boyd R. (2004). *Not by Genes Alone*. Chicago: University of Chicago Press.
- Riggs A. D., Martienssen R. A., & Russo V. E. A. (1996). Introduction. In V. E. A. Russo, R. A. Martienssen, & A. D. Riggs (Eds.). *Epigenetic Mechanisms of Gene Regulation* (pp. 1-4). Cold Spring Harbor: Cold Spring Harbor Laboratory Press.
- Rivkin, J. W. (2000). Imitation of complex strategies. *Management Science*, 46(6), 824-844.
- Rivkin, J. W., & Siggelkow, N. (2003). Balancing search and stability: Interdependencies among elements of organizational design. *Management Science*, 49(3), 290-311.
- Roe, R. A. (2008). Time in applied psychology: The study of “what happens” rather than “what is”. *European Psychologist*, 13(1), 37-52.
- Rogers, A. R. (1988). Does biology constrain culture?. *American Anthropologist*, 90(4), 819-831.
- Rogers, E. M. (2003). *Diffusion of Innovations*. New York: Free Press.
- Roth, A. E., & Murnighan, J. K. (1978). Equilibrium behavior and repeated play of the prisoner’s dilemma. *Journal of Mathematical Psychology*, 17(2), 189-198.
- Rousseau, D. M., & Fried, Y. (2001). Location, location, location: Contextualizing organizational research. *Journal of Organizational Behavior: The International Journal of Industrial, Occupational and Organizational Psychology and Behavior*, 22(1), 1-13.
- Rubinstein, A. (2003). Economics and psychology? The case of hyperbolic discounting. *International Economic Review*, 44 (4), 1207-1216.
- Rumelhart, D.E., & Norman D.A. (1978). Accretion, tuning, and restructuring: three modes of learning. In R.L. Cotton (Ed.), *Semantic Factors in Cognition* (pp. 37-53). Hillsdale: Erlbaum.
- Rumelhart, D.E. (1980). Schemata: the building blocks of cognition. In B.B.R. Spiro & W. Brewer (Eds.), *Theoretical Issues in Reading Comprehension* (pp. 33-58). Hillsdale: Erlbaum.

-
- Samuelson, L. (1997). *Evolutionary Games and Equilibrium Selection*. Cambridge: MIT Press.
- Samuelson, P. A. (1947). *Economic Foundations of Economic Analysis*. Harvard Economic Studies (Vol. 80), Cambridge: Harvard University Press.
- Sargent, T. J. (1993). *Bounded Rationality in Macroeconomics*. Oxford: Oxford University Press.
- Savage, L. (1954). *The Foundations of Statistics*. New York: Wiley.
- Schilling, M. A., & Phelps, C. C. (2007). Interfirm collaboration networks: The impact of large-scale network structure on firm innovation. *Management science*, 53(7), 1113-1126.
- Schreuder, E. J., & Mioch, T. (2011). The effect of time pressure and task completion on the occurrence of cognitive lockup. In Proceedings of the International Workshop on Human Centered Processes 2011, 10-11.
- Schriber, J. B., & Gutek, B. A. (1987). Some time dimensions of work: Measurement of an underlying aspect of organization culture. *Journal of Applied Psychology*, 72(4), 642-650.
- Schumpeter, J. (1934). *The Theory of Economic Development*. Cambridge: Harvard University Press.
- Seale, D. A., & Rapoport, A. (2000). Optimal stopping behavior with relative ranks: The secretary problem with unknown population size. *Journal of Behavioral Decision Making*, 13(4), 391-411.
- Seifert, C. M., Mayer, D. E., Davidson, N. Patalano, A. L., & Yaniv I. (1995). Demystification of cognitive insight: Opportunistic assimilation and the prepared mind perspective. In R.J. Sternberg & J.E. Davidson (Eds.), *The Nature of Insight* (pp. 65-124), Cambridge: The MIT Press.
- Selten, R. (1995). Autobiography. Nobel Prize in Economics documents 1994-6, Nobel Prize Committee.

- Selten, R. (1998). Features of experimentally observed bounded rationality. *European Economic Review*, 42(3-5), 413-436.
- Shannon, C. E. (1948). A mathematical theory of communication. *Bell System Technical Journal*, 27(3), 379-423.
- Shipp, A. J., & Cole, M. S. (2015). Time in ondividual-level organizational studies: What is it, how is it used, and Why isn't it exploited more often?. *Annual Review of Organizational Psychology and Organizational Behavior*, 2(1), 237-260.
- Shipp, A. J., & Jansen, K. J. (2011). Reinterpreting time in fit theory: Crafting and recrafting narratives of fit in medias res. *Academy of Management Review*, 36(1), 76-101.
- Shleifer, A., & Vishny, R. (1990). Equilibrium short horizons of investors and firms. Center for Research in Security Prices, Graduate School of Business, University of Chicago.
- Siggelkow, N., & Levinthal, D.A. (2003). Temporarily divide to conquer: Centralized, decentralized, and reintegrated organizational approaches to exploration and adaptation. *Organization Science*, 14(6), 650-669.
- Siggelkow, N., & Rivkin, J. W. (2006). When exploration backfires: Unintended consequences of multilevel organizational search. *Academy of Management Journal*, 49(4), 779-795.
- Simon, H. A. (1956). Rational choice and the structure of the environment. *Psychological Review*, 63(2), 129.
- Simon, H. A. (1957). *Models of Man, Social and Rational*. New York: Wiley.
- Simon, H. A. (1962). The architecture of complexity, *Proceedings of the American Philosophical Society*, 106(6), 467-482.
- Simon, H. A. (1969). *The Sciences of the Artificial*. Cambridge: MIT Press.
- Simon, H. A. (1978). Rationality as process and as product of thought. *The American Economic Review*, 68(2), 1-16.
- Simon, H. A. (1983). *Reason in Human Affairs*. Oxford: Blackwell.

-
- Simon, H. A. (1985). Human nature in politics: The dialogue of psychology with political science. *American Political Science Review*, 79(2), 293-304.
- Simon, H. A. (1989). The scientist as problem solver. In D. Klahr & K. Kotovsky (Eds.), *Complex information processing: The impact of Herbert A. Simon* (pp. 375-398). Hillsdale: Erlbaum.
- Simon, H. A. (1990). Invariants of human behavior. *Annual Review of Psychology*, 41(1), 1-20.
- Simon, H. A. (1991). Cognitive architectures and rational analysis. Comment. In K. VanLehn (Ed.), *Architectures for Intelligence* (pp. 25-39). Hillsdale: Erlbaum.
- Smith, A. (1759). *The Theory of Moral Sentiments*. Edinburgh: A. Kincaid and J. Bell.
- Smith, S.M., & Blankenship, S.E. (1991). Incubation and the persistence of fixation in problem solving. *The American Journal of Psychology*, 104(1), 61-87.
- Sommer, S. C., Loch, C. H., & Dong, J. (2009). Managing complexity and unforeseeable uncertainty in startup companies: An empirical study. *Organization Science*, 20(1), 118-133.
- Sonnentag, S. (2012). Time in organizational research: Catching up on a long neglected topic in order to improve theory. *Organizational Psychology Review*, 2(4), 361-368.
- Sorenson, O. (2003). Interdependence and adaptability: organizational learning and the long-term effect of integration. *Management Science*, 49(4), 446-463.
- Souder, D., & Bromiley, P. (2012). Explaining temporal orientation: Evidence from the durability of firms' capital investments. *Strategic Management Journal*, 33(5), 550-569.
- Souder, D., & Shaver, J. M. (2010). Constraints and incentives for making long horizon corporate investments. *Strategic Management Journal*, 31(12), 1316-1336.
- Sperber, D. (1994). The modularity of thought and the epidemiology of representations. In L. A. Hirschfeld & S. A. Gelman (Eds.), *Mapping the Mind: Domain specificity in Cognition and Culture* (pp. 39-67). Cambridge: Cambridge University Press.

- Staw, B. M., Sandelands, L. E., & Dutton, J. E. (1981). Threat rigidity effects in organizational behavior: A multilevel analysis. *Administrative Science Quarterly*, 26(4), 501-524.
- Stein, J. C. (1988). Takeover threats and managerial myopia. *Journal of Political Economy*, 96(1), 61-80.
- Stein, J. C. (1989). Efficient capital markets, inefficient firms: A model of myopic corporate behavior. *The Quarterly Journal of Economics*, 104(4), 655-669.
- Sterman, J. D. (1987). Testing behavioral simulation models by direct experiment. *Management Science*, 33(12), 1572-1592.
- Strang, D., & Macy, M. W. (2001). In search of excellence: Fads, success stories, and adaptive emulation. *American Journal of Sociology*, 107(1), 147-182.
- Sunstein, C. (2003). *Why Societies Need Dissent*. Cambridge: Harvard University Press.
- Sutton, R. S., & Barto, A. G. (1981). Toward a modern theory of adaptive networks: expectation and prediction. *Psychological Review*, 88(2), 135-170.
- Sutton, R. S., & Barto, A. G. (1998). *Introduction to Reinforcement Learning*. Cambridge: MIT press.
- Svenson, O., & Benson, L. (1993a). Framing and time pressure in decision making. In O. Svenson & A.J. Maule (Eds.), *Time Pressure and Stress in Human Judgment and Decision Making* (pp. 133-144). Boston: Springer.
- Svenson, O., & Benson, L. (1993b). On experimental instructions and the inducement of time pressure behavior. In O. Svenson & A.J. Maule (Eds.), *Time Pressure and Stress in Human Judgment and Decision making* (pp. 157-165). Boston: Springer.
- Szathmáry, E., & Smith, J. M. (1995). The major evolutionary transitions. *Nature*, 374(6519), 227-232.
- Tarde, G. D. (1903). *The Laws of Imitation*. New York: H. Holt.

-
- Thakor, A.V. (1990). Investment myopia and the internal organization of capital allocation decisions. *Journal of Law, Economics & Organization*, 6(1), 129-154.
- Thaler, R. H.(1981). Some empirical evidence on dynamic inconsistency. *Economics Letters*, 8(3), 201-207.
- Thaler, R. H. (1991). *Quasi Rational Economics*. New York: Russell Sage Foundation.
- Thaler, R. H., Tversky, A., Kahneman, D., & Schwartz, A. (1997). The effect of myopia and loss aversion on risk taking: An experimental test. *The Quarterly Journal of Economics*, 112(2), 647-661.
- Todd, P. M., Gigerenzer, G. E., & ABC Group (2012). *Ecological Rationality: Intelligence in the World*. New York: Oxford University Press.
- Tomasello, M. (1999). The human adaptation for culture. *Annual Review of Anthropology*, 28(1), 509-529.
- Tullock, G. (1967). The welfare costs of tariffs, monopolies, and theft. *Western Economic Journal*, 5(3), 224-232.
- Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, 185(4157), 1124-1131.
- Tversky, A., & Kahneman, D. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2), 263-291.
- Waddington, C. H. (1942), The epigenotype, *Endeavour*, 41(1), 18-20.
- Wallas G. (1926). *The Art of Thought*. New York: Harcourt Brace Jovanovich.
- Wallsten, T. S. (1993). Time pressure and payoff effects on multidimensional probabilistic inference. In O. Svenson & A.J. Maule (Eds.), *Time Pressure and Stress in Human Judgment and Decision Making* (pp. 167-179). Boston: Springer.
- Walras, L. (1954). *Elements of Pure Economics; or, the Theory of Social Work*. American Economic Association Translation Series. Homewood: Richard D. Irwin (Original work published 1874)

- Wanchisen, B. A., Tatham, T. A., & Hineline, P. N. (1992). Human choice in “counterintuitive” situations: Fixed-versus progressive-ratio schedules. *Journal of the Experimental Analysis of Behavior*, 58(1), 67-85.
- Watts, D. J. (1999). Networks, dynamics, and the small-world phenomenon. *American Journal of Sociology*, 105(2), 493-527.
- Watts, D. J. (2002). A simple model of global cascades on random networks. *Proceedings of the National Academy of Sciences*, 99(9), 5766-5771.
- Watts, D. J., & Strogatz, S. H. (1998). Collective dynamics of ‘small-world’ networks. *Nature*, 393(6684), 440-442.
- Wejnert, B. (2002). Integrating models of diffusion of innovations: A conceptual framework. *Annual Review of Sociology*, 28(1), 297-326.
- White, B.Y., & Frederiksen, J.R. (1986). Progressions of Qualitative Models as a Foundation for Intelligent Learning Environments. Report No. 6277.
- Wilson, D. S. (2015). *Does Altruism Exist? Culture, Genes, and the Welfare of Others*. New Haven: Yale University Press.
- Wilson, D. S. (2016). Two meanings of complex adaptive systems. In D. S. Wilson & A. Kirman (Eds.), *Complexity and Evolution: Toward a New Synthesis for Economics* (pp. 31-46). Cambridge: MIT Press.
- Wilson, D. S., & Sober, E. (1994). Reintroducing group selection to the human behavioral sciences. *Behavioral and Brain Sciences*, 17(4), 585-608.
- Wilson, D. S., & Wilson, E. O. (2007). Rethinking the theoretical foundation of sociobiology. *The Quarterly Review of Biology*, 82(4), 327-348.
- Wilson, D. S., Kirman, A., Lupp, J. (2016). Introduction. In D. S. Wilson & A. Kirman (Eds.), *Complexity and Evolution: Toward a New Synthesis for Economics* (pp. 1-11). Cambridge: MIT Press.

-
- Wilson, R. C., Geana, A., White, J.M., Ludvig, E.A., & Cohen, J.D. (2014). Humans use directed and random exploration to solve the explore-exploit dilemma. *Journal of Experimental Psychology, General*, 143(6), 2074.
- Winter, S. G. (1982). Binary choice and the supply of memory. *Journal of Economic Behavior & Organization*, 3(4), 277-321.
- Winter, S. G. (2000). The satisficing principle in capability learning. *Strategic Management Journal*, 21(10-11), 981-996.
- Wisdom, T. N., Song, X., & Goldstone, R. L. (2013). Social learning strategies in networked groups. *Cognitive Science*, 37(8), 1383-1425.
- Woolley, A. W. (1998). Effects of intervention content and timing on group task performance. *The Journal of Applied Behavioral Science*, 34(1), 30-46.
- Wright, S. (1932). The roles of mutation, inbreeding, crossbreeding and selection in evolution. *Proceedings of the 6th International Congress of Genetics*, 1, 356-366.
- Wu, C. M., Schulz, E., Gerbaulet, K., Pleskac, T. J., & Speekenbrink, M. Under pressure: The influence of time limits on human exploration. *Feedback*, 62(620), 620.
- Yaari, M. E. (1965). Uncertain lifetime, life insurance, and the theory of the consumer. *The Review of Economic Studies*, 32(2), 137-150.
- Yaniv, I., & Meyer, D.E. (1987). Activation and metacognition of inaccessible stored information: potential bases for incubation effects in problem solving. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 13(2), 187-205.
- Young, D.L., Goodie, A.S., Hall, D.B., & Wu, E. (2012). Decision making under time pressure, modeled in a prospect theory framework. *Organizational Behavior and Human Decision Processes*, 118(2), 179-188.
- Zaheer, S., Albert, S., & Zaheer, A. (1999). Time scales and organizational theory. *Academy of Management Review*, 24(4), 725-741.

Appendix A

Experimental Instructions

The following are the instructions for each condition of the laboratory experiment on time horizon effects. Since the relevant sessions were held at the Cognitive and Experimental Economics Laboratory at the University of Trento (CEEL), the original instructions were in Italian. Here below is the English translation.

Short-Term Condition

Preliminary Information

Welcome, thank you for agreeing to participate in this experiment! You are taking part in a study on decision making. During the experiment you can earn a sum of money. To this sum will be added €5 for participation. Your income will depend on your decisions and the decisions of the other participants in the experiment.

The choices you make and the answers you provide will remain absolutely anonymous. Experimenters will not be able to associate your choices and answers to your name.

For the duration of the experiment, please do not communicate with the other participants (under penalty of exclusion from the experiment) and pay close attention to the instructions that follow, which will be read aloud by one of the researchers. If you have questions, raise your hand: one of the researchers will answer your question.

At the end of the experiment we will ask you to fill in a short questionnaire. We will then proceed with the payment, which will be made privately and in cash.

General Instructions

A space traveler is an intelligent creature from another planet. He likes to buy images made up of combinations of abstract symbols and pays different sums of money for different images.

Your task is to create images by selecting and/or deselecting the symbols in the row marked by the orange background (see the screenshot below). There are 10 different symbols and you don't know which ones the space traveler likes. However, in each round you can produce the combination you wish by selecting and/or deselecting all the symbols you want.

To sell an image you need to click on the "Sell" button (see the screenshot below). Once done, the price paid by the space traveler for the image you have produced is indicated under the heading "Payment". You don't know what the price is until you have sold the image concerned. In subsequent rounds, the space traveler will be willing to buy the same image for the same price. However, he will not negotiate with you or sell you any pictures. So when you see the price the image was sold.

You will not receive any information regarding the images produced by the other participants or their payment. You will encounter two space travelers from two different planets, each of whom will purchase 11 images from you. The two do not know each other and, coming from different planets, they have completely different tastes and preferences. As a result, they are willing to pay different prices for the same image. Both pay with the currency of the space, the value of which you do not know.

You will only encounter the second space traveler after completing all sales with the first. The two do not communicate with each other, so neither of them know anything about your exchanges with each other. When you encounter a space traveler for the first time, the starting combination is already available on the computer screen and is automatically sold. Starting from the second round, you can change it as you prefer and then sell it by clicking on "Sell".

For each space traveler, your goal is to maximize the "Cumulative Payment" obtained by selling the images you produced during the 11 rounds. Both the current round and the "Cumulative Payment" are indicated in the blue bar located at the top (see the screenshot above).

The three participants with the highest "Total Earnings", that is the highest sum of cumulative payments obtained in the two different matches, will receive a prize. The first prize is €15, the second €10 and the third €5 (Note: if more participants obtain the same "Total Earnings", the winner of the prize in question will be determined by means of a lottery).

How to Proceed

1. Answer the comprehension questions below. Only when all the answers are correct it will be possible to proceed with the experiment. If you have further questions, please raise your hand: one of the researchers will come to you right away.
2. When you have correctly answered the comprehension questions, the encounter with the first space traveler will begin. It is a good idea to spend a minute learning how to select and/or deselect symbols. When you are ready to sell your first image click on "Sell."
3. After completing the encounter with the first space traveler, click on "Proceed" to start doing business with the second.
4. Finally, after having completed the encounter with the second space traveler, click on "Proceed" to go to the next page and answer the questionnaire.

Long-Term Condition

Preliminary Information

Welcome, thank you for agreeing to participate in this experiment! You are taking part in a study on decision making. During the experiment you can earn a sum of money. To this sum will be added €5 for participation. Your income will depend on your decisions and the decisions of the other participants in the experiment.

The choices you make and the answers you provide will remain absolutely anonymous. Experimenters will not be able to associate your choices and answers to your name.

For the duration of the experiment, please do not communicate with the other participants (under penalty of exclusion from the experiment) and pay close attention to the instructions that follow, which will be read aloud by one of the researchers. If you have questions, raise your hand: one of the researchers will answer your question.

At the end of the experiment we will ask you to fill in a short questionnaire. We will then proceed with the payment, which will be made privately and in cash.

General Instructions

A space traveler is an intelligent creature from another planet. He likes to buy images made up of combinations of abstract symbols and pays different sums of money for different images.

Your task is to create images by selecting and/or deselecting the symbols in the row marked by the orange background (see the screenshot below). There are 10 different symbols and you don't know which ones the space traveler likes. However, in each round you can produce the combination you wish by selecting and/or deselecting all the symbols you want.

To sell an image you need to click on the "Sell" button (see the screenshot below). Once done, the price paid by the space traveler for the image you have produced is indicated under the heading "Payment". You don't know what the price is until you have sold the image concerned. In subsequent rounds, the space traveler will be willing to buy the same image for the same price. However, he will not negotiate with you or sell you any pictures. So when you see the price the image was sold.

You will not receive any information regarding the images produced by the other participants or their payment.

You will encounter two space travelers from two different planets, each of whom will purchase 22 images from you. The two do not know each other and, coming from different planets, they have completely different tastes and preferences. As a result, they are willing to pay different prices for the same image. Both pay with the currency of the space, the value of which you do not know.

You will only encounter the second space traveler after completing all sales with the first. The two do not communicate with each other, so neither of them know anything about your exchanges with each other. When you encounter a space traveler for the first time, the starting combination is already available on the computer screen and is automatically sold. Starting from the second round, you can change it as you prefer and then sell it by clicking on "Sell".

For each space traveler, your goal is to maximize the "Cumulative Payment" obtained by selling the images you produced during the 22 rounds. Both the current round and the "Cumulative Payment" are indicated in the blue bar located at the top (see the screenshot above).

The three participants with the highest "Total Earnings", that is the highest sum of cumulative payments obtained in the two different matches, will receive a prize. The first prize is €15, the second €10 and the third €5 (Note: if more participants obtain the same "Total Earnings", the winner of the prize in question will be determined by means of a lottery).

How to Proceed

1. Answer the comprehension questions below. Only when all the answers are correct it will be possible to proceed with the experiment. If you have further questions, please raise your hand: one of the researchers will be come to you right away.
2. When you have correctly answered the comprehension questions, the encounter with the first space traveler will begin. It is a good idea to spend a minute learning how to select and/or deselect symbols. When you are ready to sell your first image click on "Sell."
3. After completing the encounter with the first space traveler, click on "Proceed" to start doing business with the second.
4. Finally, after having completed the encounter with the second space traveler, click on "Proceed" to go to the next page and answer the questionnaire.

Appendix B

Experimental Instructions

The following are the instructions for each condition of the laboratory experiment on time horizon effects. Since the relevant sessions were held at the Cognitive and Experimental Economics Laboratory at the University of Trento (CEEL), the original instructions were in Italian. Here below is the English translation.

Independent Problem-Solving Condition

Preliminary Information

Welcome, thank you for agreeing to participate in this experiment! You are taking part in a study on decision making. During the experiment you can earn a sum of money. To this sum will be added €5 for participation. Your income will depend on your decisions and the decisions of the other participants in the experiment.

The choices you make and the answers you provide will remain absolutely anonymous. Experimenters will not be able to associate your choices and answers to your name.

For the duration of the experiment, please do not communicate with the other participants (under penalty of exclusion from the experiment) and pay close attention to the instructions that follow, which will be read aloud by one of the researchers. If you have questions, raise your hand: one of the researchers will answer your question.

At the end of the experiment we will ask you to fill in a short questionnaire. We will then proceed with the payment, which will be made privately and in cash.

General Instructions

A space traveler is an intelligent creature from another planet. He likes to buy images made up of combinations of abstract symbols and pays different sums of money for different images.

Your task is to create images by selecting and/or deselecting the symbols in the row marked by the orange background (see the screenshot below). There are 10 different symbols and you don't know which ones the space traveler likes. However, in each round encounter you can produce the combination you wish by selecting and/or deselecting all the symbols you want.

To sell an image you need to click on the "Sell" button (see the screenshot below). Once done, the price paid by the space traveler for the image you have produced is indicated under the heading "Payment". You don't know what the price is until you have sold the image concerned. In subsequent rounds, the space traveler will be willing to buy the same image for the same price. However, he will not negotiate with you or sell you any pictures. So when you see the price the image was sold.

You will not receive any information regarding the images produced by the other participants or their payment.

You will encounter three space travelers from three different planets, each of whom will purchase 25 images from you. The three do not know each other and, coming from different planets, have completely different tastes and preferences. As a result, they are willing to pay different prices for the same image. All three pay with the currency of the space, the value of which you do not know.

You will only encounter the second space traveler after completing all sales with the first and, similarly, the third only after completing all sales with the second. The three do not communicate with each other, so none of them know anything about your exchanges with others. When you encounter a space traveler for the first time, the starting combination is already available on the computer screen and is automatically sold. Starting from the second round, you can change it as you prefer and then sell it by clicking on "Sell".

For each space traveler, your goal is to maximize the "Cumulative Payment" obtained by selling the images you produced during the 25 rounds. Both the current round and the "Cumulative Payment" are indicated in the blue bar located at the top (see the screenshot above).

The three participants with the highest "Total Earnings", ie with the highest sum of cumulative payments obtained in the three different matches, will receive a prize. The first prize is €15, the second €10 and the third €5 (Note: if more participants obtain the same "Total Earnings", the winner of the prize in question will be determined by means of a lottery).

How to proceed

1. Answer the comprehension questions below. Only when all the answers are correct it will be possible to proceed with the experiment. If you have further questions, please raise your hand: one of the researchers will come to you right away.
2. When you have correctly answered the comprehension questions, the encounter with the first space traveler will begin. It is a good idea to spend a minute learning how to select and/or deselect symbols. When you are ready to sell your first image click on "Sell".
3. After completing round 25 with the first space traveler, click on "Go to next" to start doing business with the second. Similarly, after completing round 25 with the second, click on "Go to the next" to start doing business with the third.
4. Finally, after completing round 25 with the third space traveler, click on "Continue" to go to the next page and answer these

Parallel Problem-Solving w/ Perfect Imitation Condition

Preliminary Information

Welcome, thank you for agreeing to participate in this experiment! You are taking part in a study on decision making. During the experiment you can earn a sum of money. To this sum will be added €5 for participation. Your income will depend on your decisions and the decisions of the other participants in the experiment.

The choices you make and the answers you provide will remain absolutely anonymous. Experimenters will not be able to associate your choices and answers to your name.

For the duration of the experiment, please do not communicate with the other participants (under penalty of exclusion from the experiment) and pay close attention to the instructions that follow, which will be read aloud by one of the researchers. If you have questions, raise your hand: one of the researchers will answer your question.

At the end of the experiment we will ask you to fill in a short questionnaire. We will then proceed with the payment, which will be made privately and in cash.

General Instructions

A space traveler is an intelligent creature from another planet. He likes to buy images made up of combinations of abstract symbols and pays different sums of money for different images.

Your task is to create images by selecting and/or deselecting the symbols in the row marked by the orange background (see the screenshot below). There are 10 different symbols and you don't know which ones the space traveler likes. However, in each round encounter you can produce the combination you wish by selecting and/or deselecting all the symbols you want.

To sell an image you need to click on the "Sell" button (see the screenshot below). Once done, the price paid by the space traveler for the image you have produced is indicated under the heading "Payment". You don't know what the price is until you have sold the image concerned. In subsequent rounds, the space traveler will be willing to buy the same image for the same price. However, he will not negotiate with you or sell you any pictures. So when you see the price the image was sold.

In each round, you will be able to view the payment obtained in the previous round by 2 other randomly drawn participants, without however being able to view the specific image associated with the payment concerned. If you deem it appropriate, you can decide to imitate one of the two participants in question by clicking on the "Imitate the image produced by this participant" button. Once done, the image associated with this payment is automatically sold and you will have access to the next round, where you can view the imitated image as the last image you produced.

You will encounter three space travelers from three different planets, each of whom will purchase 25 images from you. The three do not know each other and, coming from different planets, have completely different tastes and preferences. As a result, they are willing to pay different prices for the same image. All three pay with the currency of the space, the value of which you do not know.

You will only encounter the second space traveler after completing all sales with the first and, similarly, the third only after completing all sales with the second. The three do not communicate with each other, so none of them know anything about your exchanges with

others. When you encounter a space traveler for the first time, the starting combination is already available on the computer screen and is automatically sold. Starting from the second round, you can change it as you prefer and then sell it by clicking on "Sell".

For each space traveler, your goal is to maximize the "Cumulative Payment" obtained by selling the images you produced during the 25 rounds. Both the current round and the "Cumulative Payment" are indicated in the blue bar located at the top (see the screenshot above).

The three participants with the highest "Total Earnings", ie with the highest sum of cumulative payments obtained in the three different matches, will receive a prize. The first prize is €15, the second €10 and the third €5 (Note: if more participants obtain the same "Total Earnings", the winner of the prize in question will be determined by means of a lottery).

How to proceed

1. Answer the comprehension questions below. Only when all the answers are correct it will be possible to proceed with the experiment. If you have further questions, please raise your hand: one of the researchers will come to you right away.
2. When everyone has correctly answered the comprehension questions, the encounter with the first space traveler will begin. It is a good idea to spend a minute learning how to select and/or deselect symbols. When you are ready to sell your first image, click on "Sell". Alternatively, to imitate the image produced by another participant, click on "Imitate this combination" next to the image concerned.
3. After completing round 25 with the first space traveler, click on "Go to next" to start doing business with the second. Similarly, after completing round 25 with the second, click on "Go to the next" to start doing business with the third.
4. Finally, after completing round 25 with the third space traveler, click on "Continue" to go to the next page and answer the questionnaire.

Parallel Problem-Solving w/ Imperfect Imitation Condition

Preliminary information

Welcome, thank you for agreeing to participate in this experiment! You are taking part in a study on decision making. During the experiment you can earn a sum of money. To this sum will be added €5 for participation. Your income will depend on your decisions and the decisions of the other participants in the experiment.

The choices you make and the answers you provide will remain absolutely anonymous. Experimenters will not be able to associate your choices and answers to your name.

For the duration of the experiment, please do not communicate with the other participants (under penalty of exclusion from the experiment) and pay close attention to the instructions that follow, which will be read aloud by one of the researchers. If you have questions, raise your hand: one of the researchers will answer your question.

At the end of the experiment we will ask you to fill in a short questionnaire. We will then proceed with the payment, which will be made privately and in cash.

General Instructions

A space traveler is an intelligent creature from another planet. He likes to buy images made up of combinations of abstract symbols and pays different sums of money for different images.

Your task is to create images by selecting and/or deselecting the symbols in the row marked by the orange background (see the screenshot below). There are 10 different symbols and you don't know which ones the space traveler likes. However, in each round encounter you can produce the combination you wish by selecting and/or deselecting all the symbols you want.

To sell an image you need to click on the "Sell" button (see the screenshot below). Once done, the price paid by the space traveler for the image you have produced is indicated under the heading "Payment". You don't know what the price is until you have sold the image concerned. In subsequent rounds, the space traveler will be willing to buy the same image for the same price. However, he will not negotiate with you or sell you any pictures. So when you see the price the image was sold.

In each round, you will be able to view the payment obtained in the previous round by 2 other randomly drawn participants, without however being able to view the specific image associated with the payment concerned. However, you will receive information about how many symbols of the last image you produced are different from the image associated with the payments of each of the 2 participants. If you deem it appropriate, you can decide to imitate one of the two participants in question by clicking on the "Imitate the image produced by this participant" button. Once done, the image associated with this payment is automatically sold and you will have access to the next round, where you can view the imitated image as the last image you produced.

It is important to keep in mind that the imitation process is not foolproof but provides for a margin of error. This means that if you decide to imitate the image of another participant, the relative image is not necessarily reproduced correctly but can present imperfections. In particular, the probability that the combination of another participant is imitated incorrectly is higher the greater the differences between the last image you produced and the image you intend to imitate.

You will encounter three space travelers from three different planets, each of whom will purchase 25 images from you. The three do not know each other and, coming from different planets, have completely different tastes and preferences. As a result, they are willing to pay different prices for the same image. All three pay with the currency of the space, the value of which you do not know.

You will encounter the second space traveler only after completing all sales with the first and, similarly, the third only after completing all sales with the second. The three do not communicate with each other, so none of them know anything about your exchanges with others. When you encounter a space traveler for the first time, the starting combination is already available on the computer screen and is automatically sold. Starting from the second round, you can change it as you prefer and then sell it by clicking on "Sell".

For each space traveler, your goal is to maximize the "Cumulative Payment" obtained by selling the images you produced during the 25 rounds. Both the current round and the "Cumulative Payment" are indicated in the blue bar located at the top (see the screenshot above).

The three participants with the highest "Total Earnings", ie with the highest sum of cumulative payments obtained in the three different matches, will receive a prize. The first prize is €15, the second €10 and the third €5 (Note: if more participants obtain the same

"Total Earnings", the winner of the prize in question will be determined by means of a lottery).

How to proceed

1. Answer the comprehension questions below. Only when all the answers are correct it will be possible to proceed with the experiment. If you have further questions, please raise your hand: one of the researchers will be right away.
2. When everyone has correctly answered the comprehension questions, the encounter with the first space traveler will begin. It is a good idea to spend a minute learning how to select and/or deselect symbols. When you are ready to sell your first image, click on "Sell". Alternatively, to imitate the image produced by another participant, click on "Imitate this combination" next to the image concerned.
3. After completing round 25 with the first space traveler, click on "Go to next" to start doing business with the second. Similarly, after completing round 25 with the second, click on "Go to the next" to start doing business with the third.
4. Finally, after completing round 25 with the third space traveler, click on "Continue" to go to the next page and answer the questionnaire.