



Innovative search and imitation heuristics: an agent-based simulation study

Vittorio Guida¹ · Luigi Mittone¹  · Azzurra Morreale^{2,3}

Received: 13 December 2022 / Accepted: 2 February 2024
© The Author(s) 2024

Abstract

Prominent research in strategic imitation, exploration, exploitation, and organizational learning identifies imitation as a less costly alternative to experimentation. Yet, its role in the exploration–exploitation dilemma remains underexplored in the literature. This study employs an agent-based model to examine how two distinct agent types—those who imitate and those who experiment—interact and influence each other. The model incorporates the concept of “satisficing” derived from the behavioral theory of the firm, along with insights from research on imitative heuristics. The findings reveal that overcrowding affects both agent types negatively. Imitators suffer from diminished performance due to intensified competition, which increases as more imitators join the system. Meanwhile, explorers are hindered in their attempts at radical innovation due to the presence of other explorers and clusters of imitators. This paper contributes to the field as the first to model individual agents as ‘satisficers’ within a competitive exploration–exploitation framework. By incorporating imitation, it provides novel insights into the dynamics of organizational learning and strategic decision-making.

Keywords Agent based model · Imitation, exploration, exploitation · Satisficers

✉ Luigi Mittone
luigi.mittone@unitn.it

Vittorio Guida
vittorio.guida@unitn.it

Azzurra Morreale
azzurra.morreale@ju.se

¹ Department of Economics and Management, University of Trento, via Inama 5, 38122 Trento, Italy

² Jönköping International Business School, Jönköping University, Gjuterigatan 5, 553 18 Jönköping, Sweden

³ Business School, LUT University, Lappeenranta, Finland

1 Introduction

In the exploration–exploitation trade-off, decision-makers face the challenge of balancing the pursuit of stable returns from existing knowledge (exploitation) against the need to explore new alternatives in response to environmental shocks that render prior knowledge less valuable (exploration). While the role of imitation in this trade-off is a significant aspect that warrants investigation, existing research on imitation and its relationship with exploration and exploitation, especially within organizational learning and imitation studies, remains underexplored.

This article seeks to address this research gap by proposing an agent-based simulation that explicitly models agents as ‘satisficers,’ in line with Simon’s (1955) framework. This approach sets our study apart from previous modeling efforts, such as those by Fagiolo and Dosi (2003), Boari et al. (2017), and Fagiolo et al. (2020). By incorporating aspiration levels and comparing them with realized performance (Cyert and March 1963), we introduce problemistic search attempts when an agent’s aspirations exceed their performance. This explicit incorporation of problemistic search into our agent modeling sheds light on decision-making processes where aspirations surpass performance. Problemistic search, as conceptualized by Cyert and March (1963), involves actively searching for new options or alternatives to bridge the gap between aspirations and actual performance. This aspect enriches our simulation, capturing the adaptive behavior of decision-makers who seek solutions when faced with unsatisfactory outcomes. Thus, our model provides a more nuanced understanding of how agents navigate the exploration–exploitation trade-off.

Modeling agents as “satisficers” reflects a realistic decision-making approach that aligns with how individuals and organizations set goals or aspirations for their performance. By incorporating this aspect into the simulation, we acknowledge that decision-makers have limited cognitive resources and are not able to always pursue optimal outcomes. Furthermore, this modeling approach enables us to effectively capture the essence of satisficing behavior. Satisficing is a valuable strategy that enhances efficiency, reduces decision-making time, and alleviates cognitive overload. Incorporating this behavior into our simulation acknowledges its importance in achieving satisfactory outcomes, while also mitigating the drawbacks associated with exhaustive exploration or the pursuit of perfection.

Moreover, subsequent research following March’s seminal work in 1991 has recognized the importance of individual-level inquiry (e.g., Gupta et al. 2006). This has led to several significant contributions that examine the relationships between the exploration–exploitation trade-off and other key variables, organizational outcomes, and psychological factors. These include knowledge inflows within organizations (Mom et al. 2007), the length of decision-making time horizons (Wilson et al. 2014), and working memory—which pertains to the amount of information individuals can process during decision-making (Laureiro-Martinez et al. 2019). Additionally, recent research explores the influence of performance feedback on decision-making, contingent upon the initial task complexity (Mittone et al. 2023). This body of work enriches the frameworks of organizational learning research (Levitt and March 1988; March 1991) and the Behavioral

Theory of the Firm (Cyert and March 1963). A key focus here is on how decision-makers adaptively learn and strive to 'satisfice' their aspirations, a concept initially proposed by Simon in 1955. However, a notable gap emerges in this area of study: while the organizational learning literature acknowledges the role of imitation in knowledge acquisition (Levitt and March 1988), research on exploration and exploitation at an individual level appears to overlook the dynamic interplay between imitation and exploration. This oversight is evident in the implicit assumption that decision-makers do not consider social comparisons in their processes (Lant 1992).

On the other hand, while studies investigating imitation in managerial decisions have considered the interaction between decision-makers, they have not fully integrated the concept of 'satisficing' into their modeling framework. For example, Boari et al. (2017) modeled explorers and imitators interacting in a competitive environment, but they did not include aspirations in their model. Instead, they employed decisional heuristics (Gigerenzer 1997) to represent bounded rationality.

Building on Lieberman and Asaba's (2006) emphasis on studying the interactions between experiential learning and imitation, this paper proposes an examination of the aggregate behavior of virtual agents using an Agent-Based Model (ABM). The choice of an ABM is driven by its proven ability to critically assess prior theories (Miller 2015) and to offer deep insights through the modeling exercise itself (Fioretti 2013). The general structure of the model adheres to guidelines derived from applying ABMs in organization studies (Chang and Harrington 2006; Wall 2016) and adaptive bounded rationality models (Puranam et al. 2015). The environment is modeled as a multi-armed bandit task, where the initial mean values associated with each discrete alternative remain fixed throughout the simulations, ensuring stability except for the interactions between stereotyped agents, namely explorers and imitators, who pursue innovation in different ways.

The agents in our model are defined by their search behavior, which is aligned with the behavioral theory of the firm (Cyert and March 1963). Agents engage in search when the current option fails to satisfy their aspirations (Bromiley 1991). Explorers and imitators exhibit distinct search patterns: explorers experiment with alternatives not currently adopted by other agents, aiming to innovate by "doing something new," while imitators adopt heuristics ("imitate-the-majority" and "imitate-the-best") to select alternatives chosen by other agents. Exploration is modeled as a random draw from the alternatives not currently deployed by any other agent. Our modeling of imitation is informed by insights from Nikolaeva (2014), particularly regarding imitative heuristics that stem from cognitive frames on strategic issues. This modeling choice is driven by the challenge of locating a behavioral model exclusively focused on individuals within the considered literature streams.

Our simulation results show that imitators benefit from options discovered by explorers. Compared to explorers, they are more affected by competition but exhibit a lower exit rate. Conversely, explorers face challenges in crowded environments, particularly when there are more explorers. This leads to fewer options and lower performance overall, resulting in a higher exit rate and reduced average performance for explorers. These 'overcrowding effects' are a consistent feature in our simulation results.

This article contributes to the literature by providing an initial exploration of imitation's role within the exploration–exploitation trade-off. By modeling agents as “satisficers” (Simon 1955; Cyert and March 1963), we aim to enhance this research stream. Modeling agents as “satisficers” brings significant advantages and value to understanding the exploration–exploitation trade-off and the study of imitation. This modeling choice reflects realistic decision-making approaches, balances aspirations and performance, emphasizes problemistic search, and captures the essence of satisficing behavior. By incorporating these elements into our agent-based simulation, we gain valuable insights into the dynamics of decision-making processes and their implications for exploration and exploitation.

This paper is structured as follows. The following section covers the theoretical background where exploration, exploitation, and imitation are defined, and the most relevant theoretical insights in the respective streams of literature are reported. The third section introduces our ABM in more detail. The fourth section presents simulation strategies and the results of an exploratory analysis conducted on the simulation data. Finally, in the last section, the results are critically discussed, and new directions for future research based on the limitations of the present study are proposed.

2 Theoretical background

This article draws from three key areas of previous research to develop an ABM focused on exploration and imitation. Firstly, it utilizes the conceptualizations from the BToF by Cyert and March (1963), which emphasize adaptive learning and “satisficing” (2.1). Secondly, it examines the existing literature on imitation, including its conceptualization and modeling (2.2). Lastly, the article considers examples of ABMs in management, highlighting the advantages of adopting such methodology (2.3).

2.1 Search and satisficing

Drawing from the Behavioral Theory of the Firm (BToF) by Cyert and March (1963), scholars studying individual exploration and exploitation have concentrated on fundamental aspects of search behavior. These include the concepts of satisficing and bounded rationality (Simon 1955, 2013), performance feedback (Greve 1998, 2003b; Greve and Gaba 2017), and aspirations (for instance, Bromiley and Harris 2014). Together, these elements converge to address a central issue: the nature of search. In BToF terms, search is triggered by the need to solve a problem. This “problemistic search” aims to restore performance levels that have declined recently. The outcome of such a search can lead to either local refinement, which is exploitation, or to change and experimentation, known as exploration.

Evaluating performance and aspirations is relevant in experiential learning as it can lead to knowledge creation (Levitt and March 1988; March 1991; Argote et al. 2020). Aspirations in organizations originate from two focal points or reference sources: historical (self) performance and social performance, as outlined by

Bromiley and Harris (2014). The formation of these aspirations is closely linked to their measurement methods. These methods can include weighted-average models (e.g., Cyert and March 1963; Greve 2003a), separate measures (Greve 2003b; Baum et al. 2005; Harris and Bromiley 2007), or two hierarchically ordered reference points¹ (Bromiley 1991). Satisficing occurs when performance meets or exceeds aspirations. However, in scenarios like those depicted in Bromiley's (1991) model, an agent may raise her aspirations once the satisficing condition is met. It is important to remark that the assumption of individual agents who engage in problemistic search is grounded in recent experimental research (e.g., Billinger et al. 2014, 2021)

2.2 Exploration, exploitation, and imitation

Research on organizational learning and imitation has evolved along parallel paths (Lieberman and Asaba 2006). Imitation is recognized as a key driver of knowledge acquisition (Levitt and March 1988) and as a less costly alternative to experimentation (Lieberman and Asaba 2006).

Consistently with these prior definitions, we define imitation as an alternative search strategy to exploration.

Much of the work on imitative behavior in strategy and business research is connected to institutionalism and social ecology (DiMaggio and Powell 1983) and social learning theory (Bandura and Walters 1977). Recent contributions have incorporated essential elements of institutional theory (e.g., mimetic isomorphism) to define and model imitation. For example, Duysters et al. (2020) identify imitation and legitimacy as drivers of vicarious learning, particularly for firms that adapt their learning strategies to align with those of partners and competitors.

Imitation plays a pivotal role in fields like evolutionary economics (Nelson and Winter 1982) and industrial economics (Rivkin 2000), where it is often assumed to be perfect in both target identification (deciding whom to imitate) and execution (e.g., firms as “copycats”). However, exponents of “neo-evolutionary economics” (such as Posen et al. 2013) contend that both search and imitation processes are influenced by the cognitive limitations inherent in boundedly rational agents. As a result, imperfect imitation emerges, characterized by firms' challenges in quickly identifying the right “whom” to imitate in their industry and in flawlessly replicating their strategies (“what” to imitate).

Research in this field exploits NK models (e.g., Levinthal 1997) to investigate these dynamics. The “NK” in NK models stands for the two central parameters that govern their functioning: “N” represents the number of options available for combining into a policy or strategy, and “K” denotes the number of interdependencies among these N attributes. This “K” factor essentially reflects the complexity or “ruggedness” of the system. Typically, each of the N attributes can adopt one of two

¹ In the theoretical model of Bromiley (1991), the reference points are determined sequentially and hierarchically: the first is the average performance of others (i.e., average ‘social’ performance); the second is the ‘historical’ performance, consisting of the performance of the firm itself increased by 5%. Firms aim at historical performance only if they are doing better than social performance.

values, leading to 2^N possible combinations. Each attribute value is linked to a specific payoff, resulting in an overall payoff for each combination. As the value of K increases, indicating more interdependencies among the N attributes, even a single alteration in the combination can lead to more variable overall outcomes.

Importantly, NK models have been adopted to study exploration–exploitation decisions, both in the context of individual decision-making (e.g., Billinger et al. 2014, 2021) and within team settings (Giannoccaro et al. 2020).

Lieberman and Asaba (2006) categorized mimetic behavior into two types: information-based and rivalry-based imitation, offering a framework to reconcile the extensive literature on the subject. Building on their work, Nikolaeva (2014) developed a theoretical model of imitation heuristics grounded in managerial cognitive framing. In this model, a strategic issue faced by an agent can be perceived either as an opportunity or a threat. For instance, a situation demanding change might be viewed as potentially advantageous (like entering a new market) or as a risk to the organization (such as investing in an unknown innovation).

Nikolaeva (2014) argues that firms and managers may frame the decision to imitate as either a potential opportunity or a threat. In this model, managers may select an imitation heuristic, either 'imitate-the-successful' or 'imitate-the-majority' (Gigerenzer and Selten 2002), based on their cognitive framing of the issue and the decision to imitate.

In this article, to model imitation parsimoniously, we focus on situations where both the problem and the decision to imitate are cognitively framed in a consistent manner, either as opportunities or as threats. Consequently, a manager may choose to imitate the majority when facing a strategic threat that could worsen by not imitating. Conversely, a manager might imitate the best performers when pursuing new opportunities and when imitation itself presents an opportunity.

These are direct applications of Lieberman and Asaba's (2006) categorization of imitation, which Nikolaeva (2014) summarizes as follows: "Under the information-based category, organizations copy from their peers because of the accumulated information from predecessors. On the other hand, rivalry-based imitation is competition driven, and organizations copy their rivals so they are not left behind." (p. 1761).

2.3 Modeling and ABMs in organization studies

Modeling, as a traditional research method in organization studies, has been widely recognized (Lave and March 1975). However, Miller (2015) argues that despite its seminal contributions, the exploration–exploitation field needs to devote more attention to modeling than verbal theorizing and hypothesis testing. In contrast, over the past two decades, agent-based models (ABMs) have proliferated in the organization and management literature (Chang and Harrington 2006; Wall 2016), leading to significant advancements in understanding critical issues and theories (Gomez-Cruz et al. 2017). Drawing upon the acknowledgment of the considerable learning opportunities inherent in the construction of ABM models (Fioretti 2013; Epstein 1999; Gross and Strand 2000), the primary objective

of this article is to adopt the ABM approach and critically analyze prior studies, shedding light on imitation, exploration, and exploitation in organizational learning.

Turning to previous studies focusing on imitation, Posen and colleagues (Posen et al. 2013, 2020) have made significant contributions by modeling imitation as a bounded rational process at both the individual and aggregated levels of interaction. A parallel line of research has emerged, utilizing ABMs to replicate or expand upon March's exploration and exploitation model (March 1991). Noteworthy studies in this area include those conducted by Bray and Prietula (2007), Fang et al. (2010), Kim and Rhee (2009), Miller et al. (2006), Ponsiglione et al. (2021), and Rodan (2005).

Specifically, Bray and Prietula (2007) extended March's exploration and exploitation model to investigate the impact of environmental turbulence on organizational knowledge in hierarchies. Their empirical analyses and weighted least-squares regressions revealed that a "top-down" strategy characterized by high exploitation and low exploration reduced individual knowledge accuracy, particularly in multi-tier hierarchical organizations. In contrast, flat organizations experienced a smaller decline, highlighting the resilience of a "bottom-up" approach in dynamic environments and offering insights for optimizing knowledge management.

Fang et al. (2010) conducted a comprehensive study using simulations to examine the influence of interpersonal network structure on organizational learning. Their findings emphasized the relationship between network structure and performance levels in organizations. Moderate cross-group linking led to higher equilibrium performance, indicating the positive impact of interconnectedness between subgroups on organizational performance. The study also highlighted the importance of preserving subgroup heterogeneity for facilitating broader exploration of ideas and beliefs, underscoring the role of network structure in optimizing learning processes.

Kim and Rhee (2009) focused on the choice between exploration and exploitation and its implications for organizational performance, considering environmental dynamism and internal variety. Their simulation models demonstrated the relationship between organizational practices, internal variety, and knowledge adaptation. The findings underscored the importance of managing internal variety through strong complementary practices to achieve a balanced approach between exploration and exploitation, providing valuable insights for organizational adaptation to dynamic environments.

Miller et al. (2006) expanded March's (1991) agent-based model by incorporating direct interpersonal learning, spatial context, tacit knowledge, and personnel turnover. Their study provided a comprehensive understanding of organizational learning dynamics, highlighting the significance of informal knowledge exchange and the impact of geographical factors. The insights gained from their research offer valuable guidance for optimizing learning processes in organizations.

Ponsiglione et al. (2021) utilized an agent-based computational laboratory, the Computational Laboratory of Organizational Design (CLOD), to conduct simulative experiments. Their study investigated the advantages of employing natural language as an informal coordination mechanism for enhancing organizational performance, particularly in turbulent environments. The findings emphasized the benefits

of natural language-based coordination, providing practical insights for practitioners and researchers.

Rodan (2005) extended March's (1991) model by examining individual- and organizational-level processes influencing belief variation, selection/retention, and experimentation. The author's analysis sheds light on factors impacting organizational learning, such as experimentation propensities, restraints, and turnover in membership. The research enriched the understanding of organizational learning processes and offered practical implications for organizations.

Nevertheless, it is worth noting that the studies above predominantly concentrate on the isolated aspects of imitation or exploration–exploitation individually. In contrast, our particular interest lies in exploring the intricate interplay between these factors (Lieberman and Asaba 2006). In this regard, Fagiolo and Dosi (2003) proposed an agent-based model where firms explore a lattice representing a Schumpeterian technological space. This model incorporated exploitation, exploration, and imitation elements, offering valuable insights into economic growth. Similarly, Fagiolo et al. (2020) expanded upon this model by introducing a financial system investigating economic growth with and without financial institutions and resources.

While these models are well-described and finely implemented, their treatment of bounded rationality relies primarily on the stochasticity of the models, needing a solid theoretical basis in the literature, such as satisficing. In contrast, Boari et al. (2017) approached bounded rationality from a more theoretical perspective in their model, providing an interesting definition of exploration and exploitation based on vicarious and experiential learning. Their model did not directly implement or program the proposed framework, suggesting that interesting findings could arise from the modeling exercise itself. This theoretical depth, however, was grounded in a "downstream" approach to modeling bounded rationality, drawing on Gigerenzer's (1997) work and avoiding the modeling of aspirations and satisfaction.

From the literature above, two interconnected lessons can be drawn. Firstly, while starting from different assumptions, research questions, and theoretical approaches, previous studies have consistently excluded "satisficing" from their models. Secondly, modeling inherently requires simplification of reality (Lave and March 1975), and the studies discussed above exemplify balanced and well-designed models.

Building upon these observations, we acknowledge a gap in the literature resulting from legitimate modeling choices, specifically excluding agents as "satisficers." Consequently, our modeling choices are based on an opposing logic to previous studies, explicitly modeling agents as "satisficers" while maintaining simplicity in other model aspects.

3 The model

Wall (2016) provides an extensive literature review on the adoption of ABMs in the management field. It shows that these models consist of three building blocks: the environment, the agents, and the possible interactions between the agents themselves or the agents and the environment. We will introduce the model by referring to each of these building blocks.

3.1 Environment

The environment is designed as a multi-armed bandit task with one hundred options $o \in O = (1, 2, \dots, 100)$. Each option o is one of the discrete arms of the machine, and its performance distribution is only known if the agent has already experimented with it. Conceptually, the environment can be considered the space of viable solutions to a generic management problem. According to this definition, the agents must explore this random space without prior knowledge.

The performance of each option is drawn from a normal distribution with mean μ and variance ν .

$$P_o \sim N(\mu, \nu)$$

where the mean μ is drawn from a discrete uniform distribution within the interval [55, 95]; the variance ν is set to 1 for each alternative to avoid generating excessive “noise” in the data.

The means of the distributions associated with each option are either kept stable or re-drawn at each period, depending on the parameter modeling the presence of *turbulence*. As a result, our setting comprises two types of systems: one with perfectly stable alternatives for agents to discover and adopt, and the other marked by constantly changing alternatives—an extremely turbulent one.

Three parameters define the *initial conditions* of the model, namely, *density* (d), the initial proportion of *explorers* (e), and the percentage loss due to competition with another agent (c).

Density ($0 \leq d \leq 1$) indicates the number of agents in the environment proportional to the number of options (i.e., 100). For example, if $d=0.2$, there will be 20 agents in the system. The density is kept constant during each simulation: an agent is replaced after it exits the simulation.

The proportion of explorers ($0 \leq e \leq 1$) represents the initial fraction of explorers in the simulation. For example, if $d=0.2$ and $e=0.5$, half of the twenty agents in the system (i.e., 10) will be explorers at the simulation’s outset.

The parameter e represents the probability of a new agent being an explorer, which helps maintain the initial proportions in the system. However, it does not guarantee that the proportion of explorers will remain unchanged from the start of the simulation.

Losses due to competition (c) occur when agents select the same alternative. The parameter c represents the percentage loss an agent incurs for each other agent choosing the same alternative. More details about competition and losses are provided in Sect. 3.3 below, which introduces interactions between agents.

Finally, agents with an average performance of 60 or lower exit the model. This exit threshold is a parameter that can be adjusted in future simulations. While the exit threshold is not the primary focus of this study, setting it too high could introduce a confounding factor in assessing the role of competition in the simulations. Consequently, we have set it to a value approximating the 10th percentile of the possible mean values of the options, which, as previously explained, are uniformly distributed between 55 and 95.

3.2 Agents

There are a finite number of agents $I = \{1, 2, \dots, N\}$ in the model, where $N \cong d \cdot 100$. Each agent $i \in I$ is defined by their search type s_i , qualitative aspiration $a_{i,t}$, and quantitative aspiration $A_{i,t}$ on performance.

Agents search when they experience a performance that is not “satisficing” (Simon 1955; Cyert and March 1963); otherwise, they stop searching and exploit the current option (Cyert and March 1963; Posen et al. 2018).

Fagiolo and Dosi (2003) modeled exploration using a probability parameter representing the propensity to explore, while imitation was represented as an event that occurs according to a conditional probability of receiving a signal from other miners (who, in their model, are engaged in exploitation. This model leverages the stochastic properties of the probability distributions involved and serves as an excellent example of defining and using theoretical parameters.

In our model, an agent i can be generated as either an explorer ($s_i = \text{“explorer”}$) or an imitator ($s_i = \text{“imitator”}$), where e is the probability of being an explorer and $(1 - e)$ is the probability of being an imitator. This defines the only type of search agent i will use if it is not “satisficed” with the performance generated by the option at hand.²

Before delving into the mechanisms of imitation and exploration employed by our agents, it’s essential to establish clear definitions of aspirations and performance. In our model, all agents participate in a “learning cycle” L_c , where they sample the current option over ten periods. At each period t within the learning cycle $l \in L_c = \{1, 2, \dots, L, \dots, 10\}$, an agent i samples the current option and computes an average performance $\bar{P}_{i,t,l}$.

$$\bar{P}_{i,t,l} = \frac{1}{L} \cdot \sum_{l=1}^L P_{i,t,l}$$

Conceptually, in our model, “performance” encompasses any metric that reflects the financial returns associated with an organization’s chosen strategy. While there is a rich history of using various performance measures in the study of organizational aspirations, as extensively discussed in the literature (Bromiley and Harris 2014), this paper primarily focuses on theoretical aspects. Therefore, we define “performance” in more abstract terms as a metric that captures the environmental response to the chosen option, aligning with the requirements of our theoretical framework.

In line with BToF (Cyert and March 1963), at the beginning of each L_c (i.e., $l=1$), agent i compares $\bar{P}_{i,t,l}$ with its aspirations. Aspirations are modeled according to Bromiley’s (1991) switching model. Let $\bar{P}_{S,t-1}$ be the average social performance of all agents in the system at time $t-1$, which is the information available to each agent at time t .

² An exciting development would be to calibrate a set of “propensity” parameters with experimental data to study how and when agents search and in what manner (e.g., propensity to imitate or propensity to explore). As we discuss in the concluding remarks of this paper, this may be a promising direction for future research.

$$\bar{P}_{S,t-1} = \frac{1}{100} \cdot \sum_{i=1}^{100} P_{i,t}$$

If $\bar{P}_{i,t,l} < \bar{P}_{S,t-1}$, the agent's target will be $\bar{P}_{S,t-1}$, hence, $a_{i,t} = \text{social}$, and $A_{i,t} = \bar{P}_{S,t-1}$. If, on the contrary, $\bar{P}_{i,t,l} \geq \bar{P}_{S,t-1}$, $a_{i,t} = \text{historical}$, and $A_{i,t} = 1.05 \cdot \bar{P}_{i,t,l}$, with a 0.5 probability; $\bar{P}_{i,t,l}$, otherwise.

According to Bromiley's (1991) model, the new aspiration level will be equal to the historical performance increased of its 5% (i.e., $1.05 \cdot \bar{P}_{i,t,l}$). However, as we do not model risk attitudes in this study, we randomize the possibility of having such an increase in aspirations by making the increased historical performance and the historical performance equiprobable.

Satisficing (Simon 1955) occurs when the average performance of the current option is at least equal to the aspirations.

$$\bar{P}_{i,t,l} \geq A_{i,t}$$

In such a case, agents will continue to exploit the current option; otherwise, they will search by exploring or imitating.

Exploration is assumed to be innovative for the system. Thus, it is modeled as randomly choosing a "free" alternative. In contrast, imitation involves copying alternatives currently adopted by others. The choice of imitation heuristic depends on the agents' aspirations. Imitators will adopt an "imitate-the-majority" heuristic if they aim for the average social performance or an "imitate-the-best" heuristic if they seek opportunities to improve their historical performance (Nikolaeva 2014). In the model, the majority corresponds to the "mode" option, which is chosen by most agents. If there are multiple majorities, the imitator randomly selects one of them.

In the simulations performed in this paper, the top performers correspond to the top 5% of agents ranked by their descending performance. This top 5% includes agents whose performance ranks above the 95th percentile of the performance distribution. This parameter is adjustable within the model's program and can be readily modified for future testing. The imitator randomly selects one of these top-performing agents to imitate in the next learning cycle.

3.3 Interactions

As mentioned above, agents may interact with others in such an environment due to competition. Specifically, a loss is introduced when multiple agents adopt the same strategy/behavior. Losses due to competition are defined as follows.

$$losses_{i,o,t} = (c \times others_{o,t}),$$

where c is a parameter of the proportional loss from competition, and "others" is the number of other agents adopting the option o at time t . The parameter l is set for each simulation to 1%. Therefore, the performance of the i -th agent at time t is

Table 1 Fixed and Variable Parameters in the simulated models

Parameters	Value(s)
<i>Fixed</i>	
Number of Alternatives	100
Top Performance Quantile	0.95
Exit Average Performance	60
New Agents	Yes
<i>Variable</i>	
Density (d)	{0.2, 0.5, 0.8}
Explorers Proportion (e)	{0, 0.2, 0.5, 0.8, 1}
Losses to Competition (c)	{0, 0.01}
Turbulent Environment (t)	{0, 1}

$$P_{i,o,t} = P_{i,o,t} \cdot (1 - losses_{o,t}).$$

By selecting an option, agent i receives a final performance value determined by a random draw from the relative distribution, adjusted for a proportional decrease due to competition.

These competitive interactions can affect the agent's perception of the option (Puranam et al. 2015). The greater the competition for an alternative, the lower the agent's performance, and the less likely it is that satisficing will occur. In essence, competition reduces the reward the agent associates with a strategy, making it more likely to be abandoned.

4 Simulations, data, and results

4.1 Simulation strategy

The data for this study were generated by simulating the ABM with varying initial conditions through adjustments to four parameters. A summary of both fixed and variable parameters, along with their respective values, is provided in Table 1. The choice of fixed parameters was guided by logical reasoning, emphasizing their role as background variables in this study. Detailed explanations of these choices can be found in "Appendix 1".

All the simulated parameters and their possible levels are included in this table. By combining the fixed and the variable parameters, 60 model specifications emerge.

The combinations of the parameters listed in Table 1 result in 60 different model specifications. However, since the focus of this study is on the aggregate interaction between explorers and imitators, we restrict the data analysis to the nine configurations in which this interaction is at play (see Table 2 below). In these configurations, competition is always present, and the proportion of explorers varies. In addition, we exclude the extreme initial conditions where either explorers or imitators are absent.

Table 2 Parameter configurations of the models

Model	d	e	c	Turbulence
4	0.2	0.2	0.01	No
6	0.2	0.5	0.01	No
8	0.2	0.8	0.01	No
14	0.5	0.2	0.01	No
16	0.5	0.5	0.01	No
18	0.5	0.8	0.01	No
24	0.8	0.2	0.01	No
26	0.8	0.5	0.01	No
28	0.8	0.8	0.01	No
34	0.2	0.2	0.01	Yes
36	0.2	0.5	0.01	Yes
38	0.2	0.8	0.01	Yes
44	0.5	0.2	0.01	Yes
46	0.5	0.5	0.01	Yes
48	0.5	0.8	0.01	Yes
54	0.8	0.2	0.01	Yes
56	0.8	0.5	0.01	Yes
58	0.8	0.8	0.01	Yes

We ran one thousand simulations for each configuration, with each simulation running for five hundred periods. This led to a final dataset comprising thirty million rows, resulting from five hundred observations per simulation, conducted one thousand times for each model configuration. Consistently with our focus on the models listed in Table 2, the analysis was also restricted to a subset of the original dataset, which counted nine million rows.

During each of the five hundred periods in a single simulation, we collected aggregated measures, including mean values of agents' performance, aspirations, losses, the the performance-aspiration gap the exits from the simulations by agent type, and agent counts by type.

In line with the logic of conducting an exploratory analysis of the simulation data, the results are reported in the following sequence.

First, this section presents an overview of the model's key outcomes, focusing on three variables: agents' choices of exploitation and search, the number of agents that exit or remain in the simulations, and agents' concentration among alternatives. Additionally, we calculate the Herfindahl–Hirschman Index (HHI) concentration index and the average sum of the most numerous five clusters formed during the simulations to provide a further understanding of imitative dynamics.

Second, we analyze the distinctions between imitators and explorers. These results are presented as general descriptions, accompanied by the corresponding graphical representations of the identified differences included in "Appendix 1". Differences were tested on a value-by-value basis between the two groups of agents in each simulated model, ensuring precise data analysis. Analyzing differences

1. Environment Setup (executed only once, at the beginning of the simulation)	2. Agents Activation	3. Computations	4. Data Collection	5. Substituting Agents
<ul style="list-style-type: none"> 100 options are generated as Gaussian distributions <ul style="list-style-type: none"> Means are drawn from a discrete uniform distribution $U(55,95)$ Variance is set to 1 $N \approx d$ - 100 agents are generated of which: <ul style="list-style-type: none"> Explorers $(1 - e)$ Imitators Competition e is set to be either: <ul style="list-style-type: none"> at play (1% loss in performance for any other agent using the same option) or not (no losses) Turbulent Environment (t): <i>turbulence</i> \rightarrow the mean of each option changes at each period - redrawn from $U(55,95)$ environment is <i>stable</i> otherwise 	<ul style="list-style-type: none"> Each agent is activated in random order Agents "behave" and "decide" based on information from the previous period Each agent performance is computed <p>For more details on the agent algorithm, see Appendix 1.</p>	<p style="text-align: center;">sequence repeated at each time step/period</p> <p>Statistics about agents' aggregate behavior are computed. For example:</p> <ul style="list-style-type: none"> average performance in the period average aspirations Total number of agents Etc. <p>Lists of reference groups are updated:</p> <ul style="list-style-type: none"> Best performers (parameterized as "Top Performance") 	<ul style="list-style-type: none"> Data is collected after agents have made their decisions and statistics have been computed At each period, data is collected at the aggregated level After data collection, some counters and accessory variables are reset to 0 Turbulence - if present - is triggered after the agents' activation and the data collection 	<ul style="list-style-type: none"> Agents who obtain an average performance below 60 exit the model. Agents that exited the model are substituted with new ones who will operate in the next round
Parameters involved				
<p>Density (d): number of agents with respect to the number of options</p> <p>Explorers Proportion (e): proportion of explorers of the total number of agents</p> <p>Turbulent Environment (t)</p>	<p>Losses to Competition (c): either 0 or 0,01 depending on the initial setup of the simulation. It is multiplied by the number of other agents choosing the same option and corresponds to the share of performance that the agent loses due to competition.</p>	<p>Top Performance Percentile: a fixed parameter indicating what percentile of agents sorted by their performance level can be considered as top performers, hence, the best competitors in a given step. It is set to 0,95.</p>	<p>Exit Average Performance: the average performance that an agent obtains over time for a given option. It is set at 60, being it approximately the 10th percentile of the possible draws from the $U(55,95)$ used to generate the options.</p>	

Fig. 1 Simulation behavior. The simulation process is represented starting with the setup phase of the environment. Options and all the other initial conditions are generated only once when there is no turbulence in the simulations. In contrast, if turbulence is included, the mean values of the options will be re-drawn at each period. Steps from 2 to 5 are repeated at each period

between explorers and imitators is critical to interpreting aggregate variables at the aggregate level.

Third, we evaluate the aggregate variables that characterize agents as “satisficers” in a competitive environment. This assessment includes performance, losses, aspirations, and the gap between performance and aspirations. The various models underwent sensitivity testing regarding changes in their structural parameters, d and e . The outcomes of these tests are detailed in “Appendix 2”.

Finally, we compare the outcomes outlined so far with those obtained from additional simulations that introduced turbulence to the environment.

The simulation procedure is illustrated in Fig. 1, depicting the role of each parameter at each step of the simulation. Additionally, the agent behavior/algorithm is displayed in Fig. 14 in “Appendix 1”.

4.2 Overview of the model's outcomes: number of agents and their decisions

Figure 2 shows the mean number of agents per type of decision, specifically the mean counts of decisions to exploit, explore, or imitate. Interestingly, the number of agents who chose exploitation stabilizes early in each simulation. Imitation tends to increase throughout the simulations and even surpasses exploration in models where imitators were initially set to be a minority.

This occurs because, as shown in Fig. 3, imitators tend to replace explorers. Although the total number of agents in the simulations is kept constant due to the substituting mechanism, explorers and imitators exit the simulations at rates that may be higher or lower than their generation probabilities. Regrettably, these

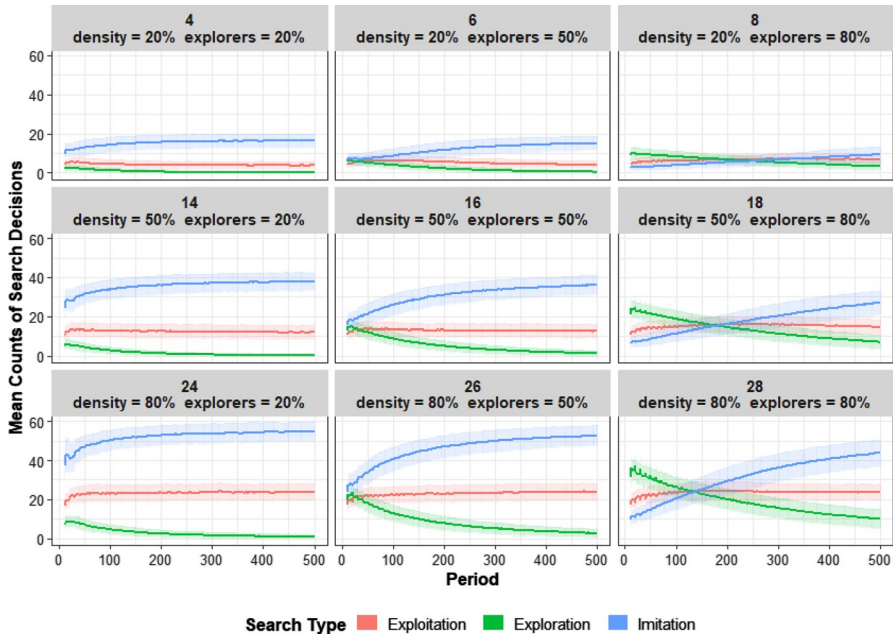


Fig. 2 Mean counts of search decisions. Mean values computed for each search behavior at each period (i.e., exploitation, exploration, or imitation) across one thousand simulations of each model specification

dynamics may not yield interesting levels of exploration and imitation for comparative analysis.

In contrast, exploitation concerns those agents who reach satisficing options and choose to cease searching, irrespective of their type. The stable pattern of exploitation confirms that agents are modeled in a manner consistent with our slightly modified version of Bromiley’s (1991) model: exploitation occurs when agents aiming to improve their historical performance do not automatically raise their aspirations. On the other hand, agents systematically search for alternatives when they are either below the social average performance or have triggered an increase in their own historical performance.

Figure 4 shows that the cumulative sum of exits in the simulated models increases as the two parameters increase. A joint reading of Figs. 3 and 4 indicates that explorers tend to “survive” less in the simulations as their number increases both in relative (*e*) and absolute terms (*d*).

In summary, these data show that in the simulated models, imitation prevails over exploration due to the replacement of imitators with explorers, favoring the former category. On the other hand, we observe that the number of exploiting agents remains stable, with one exception. In models where the initial number of explorers is the majority (models 8, 18, and 28), exploitation seems to increase when there is a sufficiently high number of explorers in the system. However, when the number of imitators consistently exceeds the number of explorers, exploitation stabilizes again.

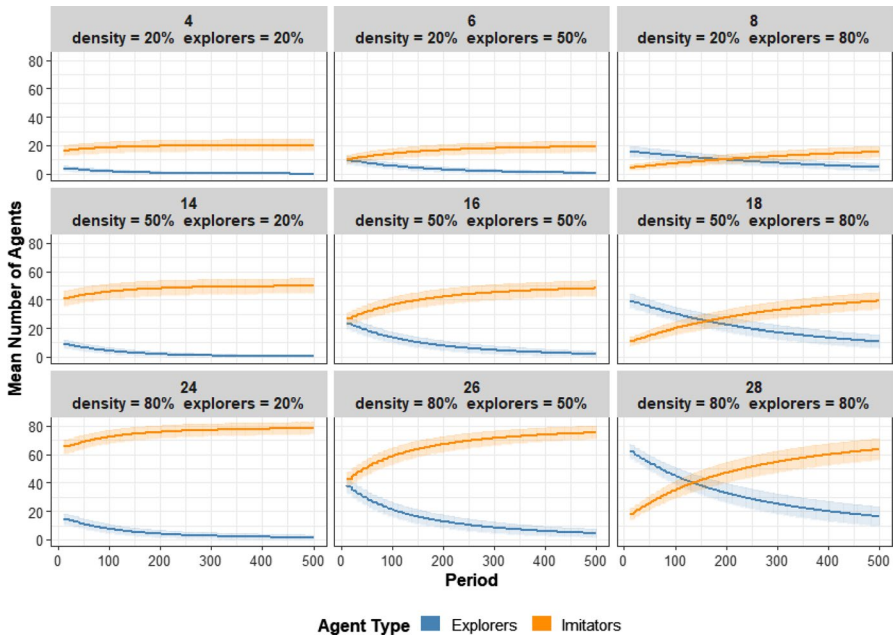


Fig. 3 Mean counts of agents per type. Mean counts of explorers and imitators at each step across the one thousand simulations run for each model specification (i.e., at varying density levels and share of explorers)

An additional piece of information that helps readers to understand better what happened in the simulated models is the concentration of agents around the options.

Figure 5 illustrates the mean HHI computed across all model simulations. The agents exhibit a lower concentration on the same options as the parameters d and e increase. Although the concentration across models decreases with these increasing parameters, the mean HHI increases as the simulations progress. This trend may result from a shift in the composition of agents from the initial conditions to a setting mostly populated by imitators, who search by aggregating around an option.

Similarly, Fig. 6 shows the cumulative relative number of agents in the most significant clusters of agents adopting the same alternative. This measure was obtained by sorting the counts of agents associated with each alternative in descending order. The dynamics mirror those observed for the HHI. The cumulative sum across clusters allows us to observe the numerosity of individual groups through their vertical differences and assess the overall numerosity of agents in the top five chosen options relative to the entire population. In addition, Fig. 6 shows that the higher the values of d and e , the less numerous the most populated clusters become.

We observed that imitators proliferate and replace explorers. This explains the increase in the cumulative sums of the agents in the clusters over time, especially in the last stages of the simulations.

It is interesting to note, however, that the agents in the early stages of the simulations achieve a remarkably high level of concentration on the first five options.

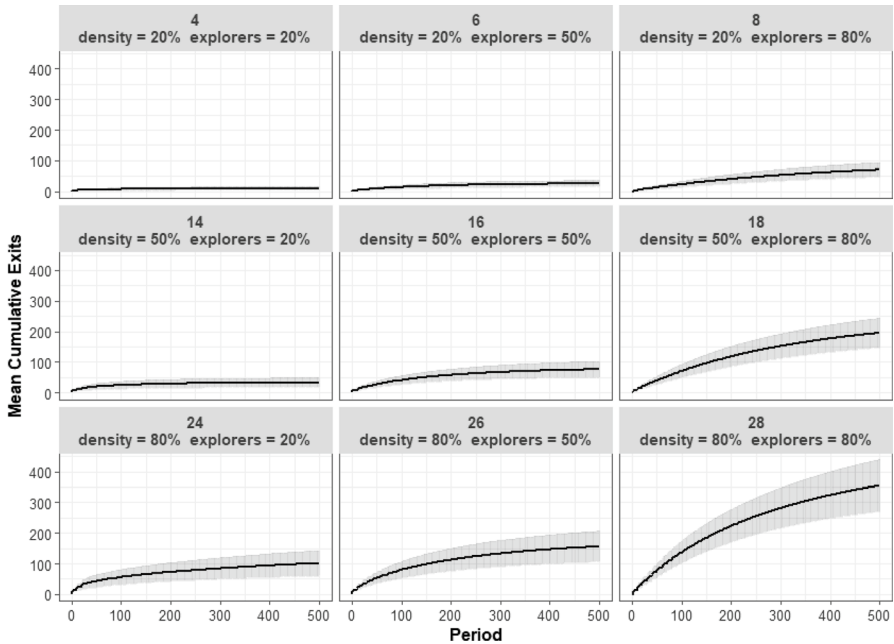


Fig. 4 Mean cumulative exits. Agents’ exits from the model due to them reaching the performance threshold of 60 have been cumulated over each of the one thousand simulations per model, and the mean value has been computed at each step

This reflects the speed of convergence between the two imitation heuristics. Indeed, an "imitate the majority" heuristic might lead agents targeting average social performance to imitate the larger clusters created by imitators targeting top performers. Thus, the initial concentration of agents in a few clusters at the beginning of the simulations is unsurprising.. Interestingly, as d increases, imitators engage with a broader range of options provided by explorers over time, resulting in slower clustering.

4.3 Explorers and imitators

The differences between explorers and imitators are graphically shown in the figures in “Appendix 1” of this paper, from Fig. 15 to 26. For the sake of brevity, we offer a series of comments on these figures as follows. Compared to imitators, explorers exhibit lower and less variable performance. Nevertheless, the mean performance of imitators tends to decline over time.

Losses are experienced almost exclusively by imitators, and their magnitude increases with higher values of d and e . Explorers experience losses when they are in the minority (i.e., at low levels of e), and these losses tend to be higher as d increases. These losses experienced by explorers illustrate the scenario where an explorer becomes the target of imitation.




Fig. 5 Mean Herfindahl–Hirschman Index. The index has been computed at each step for each of the one-thousand simulations by summing the square of the number of agents per option over the total number of agents. It has been then divided by 10,000 (theoretical maximum value) to make the value relative to 1

Mean aspirations reflect mean performance. However, the differences between explorers and imitators are less pronounced in this regard than they are for performance. Moreover, the differences in mean aspirations between explorers and imitators are more significant as e increases, with imitators having, on average, higher aspirations than explorers.

Regarding the gap between performance and aspirations (i.e., satisficing), we find that imitators are, on average, less satisfied than explorers. However, imitators benefit from the presence of more explorers, i.e., their gap between performance and aspirations is, on average, less negative as e increases. Conversely, explorers suffer from very crowded environments (when d is high), especially when they are in the majority (when e is high), leading to increasing levels of the explorers' mean performance-aspirations gap.

This sensitivity of agents to crowded environments is a crucial feature of the whole model. Both explorers and imitators perform worse when d is high, regardless of the e level. In other words, both categories are negatively affected by being in large populations, and this effect is exacerbated when they are in the majority.

Finally, it can be noticed that imitators very rarely leave the simulations, while explorers do so more frequently as both parameters increase, as we understood from Figs. 3 and 4.

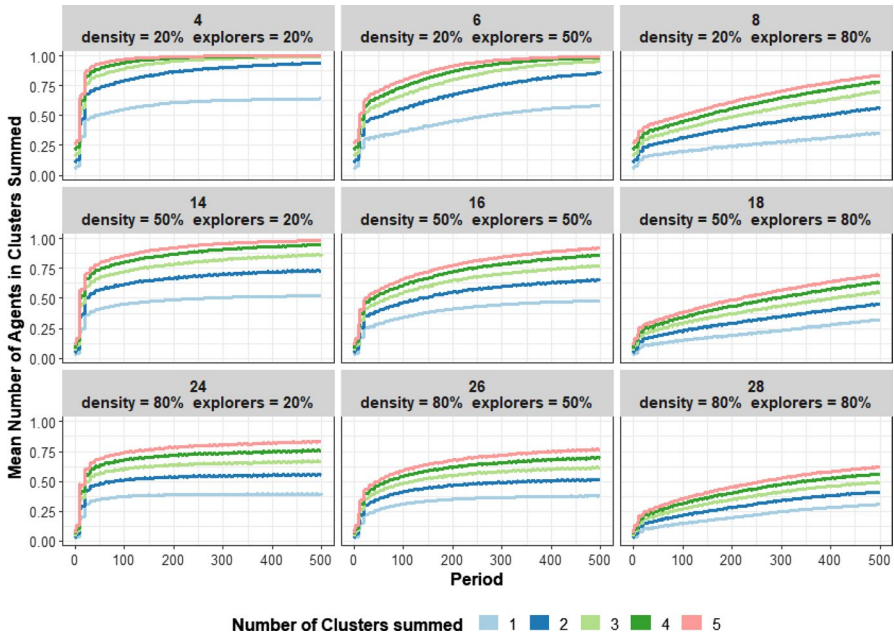


Fig. 6 The mean number of agents in the five largest clusters combined. It is calculated by summing cluster sizes in descending order. Each number in the graph represents the cumulative sum of the clusters, where 1 is the largest cluster, 2 is the sum of the largest and second-largest clusters, and so on. The vertical gaps between the lines indicate the respective sizes of the clusters

4.4 Satisficing: performance, losses, and aspirations.

This part of the analysis focuses on the pooled aggregated performance measures, losses, aspirations, and the performance-aspirations gap.

The mean performance (as shown in Fig. 7) decreases as the parameter d increases and is higher as e increases. This aligns with our earlier observations: more agents in the system lead to increased competition for imitators and fewer alternatives for explorers to experiment with. Furthermore, a higher value of e implies a greater number of explorers in the initial conditions of the simulation, indicating the presence of a more extensive set of discoveries made by them, which then become available to imitators.

The exit of explorers and the resulting increase in imitators explains the the observed decreasing pattern at higher levels of d and e .

Figure 8 confirms these results in terms of mean losses due to competition. The increase in losses over time, even in models with a high parameter e , reflects the shift in the majority towards imitators, thereby increasing the concentration on selected alternatives.

As with the mean performance, and in line with what was described in the "The Model" section, the mean aspirations (as shown in Fig. 9) decrease and increase as d and e increase, respectively. However, on average, the aspirations are always higher than performance at all time steps, which is by design.

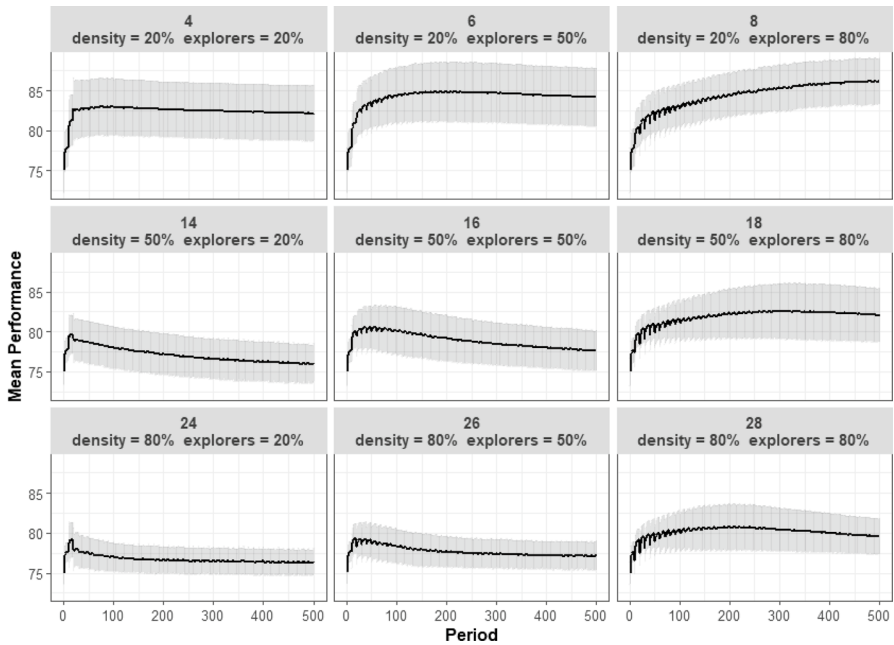


Fig. 7 Mean Performance. Mean values computed for each period across one thousand simulations per model specification (i.e., at varying density levels and share of explorers)

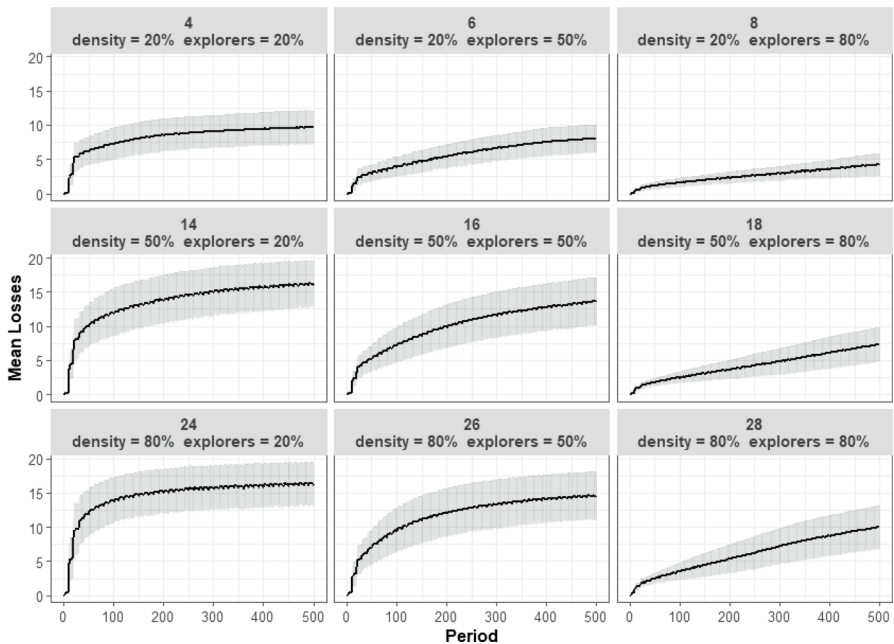


Fig. 8 Mean losses experienced due to competition. Mean values computed for each period across one thousand simulations per model specification (i.e., at varying density levels and share of explorers)

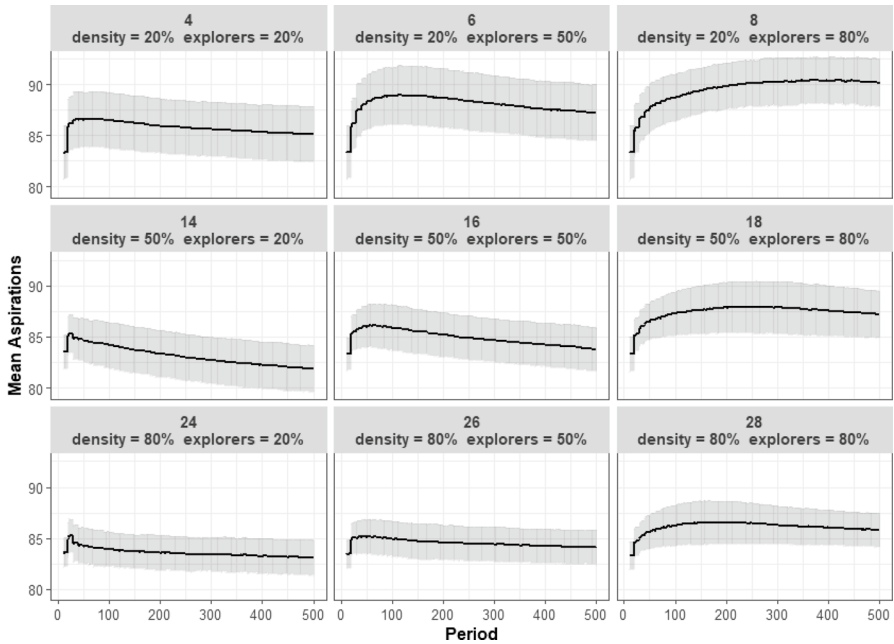


Fig. 9 Mean aspirations. Mean values computed for each period across one thousand simulations per model specification (i.e., at varying density levels and share of explorers). Since each agent can compute aspirations only after sampling an option for ten rounds, the first ten rounds of the simulations have been excluded from the figure

Notably, mean aspirations decrease over time in all models, reflecting the lower experienced performance of imitators due to increased competition.

Finally, Fig. 10 shows the mean gap between agents’ performance and aspirations. Different from the aggregate variables examined so far, the performance-aspiration gap is less intuitive to interpret for at least two reasons. First, and most obviously, the fact that the gap is always negative on average forces us to interpret its values in the opposite direction. A less negative gap indicates a higher average level of satisfaction. Second, and more importantly, this aggregate variable brings together many of the dynamics previously mentioned, such as the different ways crowding affects explorers and imitators.

For simplicity and brevity, we will assess this variable in qualitative terms, describing the agents’ experience in the simulated models. Furthermore, a more insightful approach to analyzing this variable involves contrasting its interpretation between the initial and final stages of the simulations.

In the early stages, the gap between performance and aspirations is influenced by the level of the competition experienced by imitators. For example, Model 4 shows the initial level of the gap closest to zero among the models. This model features few agents ($d=20\%$) with a higher proportion (80%) of imitators ($e=20\%$). Nevertheless, the limited alternatives discovered by the explorers, combined with those randomly assigned to the agents at the model’s initiation,

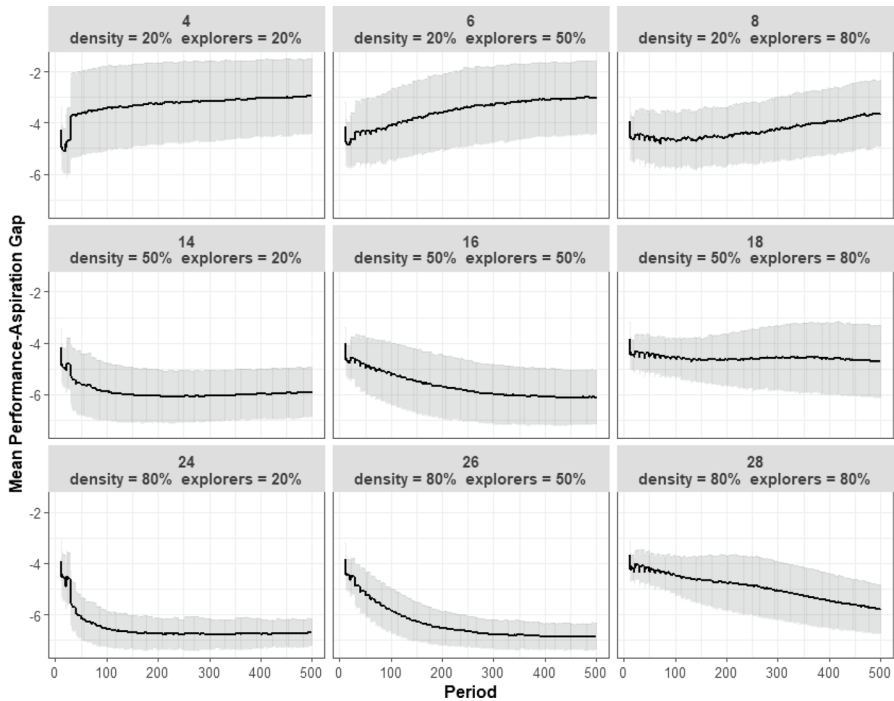


Fig. 10 Mean Performance-Aspiration Gap. Mean values computed for each period across one thousand simulations per model specification (i.e., at varying density levels and share of explorers). Since each agent can compute aspirations only after sampling an option for ten rounds, the first ten rounds of the simulations have been excluded from the figure

are sufficient to compensate for the losses due to competition. This is evidenced by the model showing the highest HHI (as depicted in Fig. 5) and the highest concentration of agents in the first five clusters (as seen in Fig. 6). As explorers find new alternatives, imitators subsequently follow.

The same logic can be applied to other models in which the initial steps show a less negative gap between performance and aspirations, followed by increasing or stable trends (e.g., models 6, 8, 16, 18, and 28). In these models, the search process proves beneficial for agents in the early periods.

In the later stages of the simulations, the replacement of explorers with new imitators, a common occurrence across all models, determines the trend of the performance-aspiration gap. This trend depends on the "legacy" left by explorers in terms of discovering more rewarding options.

The gap decreases (becomes less negative) when there are enough good options to compensate for concentration and losses; otherwise, it increases (becomes more negative). For example, Model 28, which contrasts with Model 4 in terms of parameters d and e , exhibits a complete reversal in the satisficing trend. In this model, competition is so intense that losses reduce performance, even when agents are less concentrated on the options compared to other models.

Table 3 Summary of the main results of the simulations with stable environment

Aspect	Main results
Decision Dynamics	Exploitation becomes stable early in simulations, with a consistent number of agents choosing this strategy Imitation tends to increase and surpass exploration, even when imitators are initially a minority (high e)
Agent Dynamics	Imitators proliferate and replace explorers over time Concentration on the first five options is fast in the early stages Reflecting the speed of convergence between imitation heuristics With higher d and e , imitators deal with a wider range of options from explorers, slowing down clustering
Explorers vs. Imitators	Explorers show lower performance and higher variability than imitators Losses are primarily experienced by imitators, increasing as parameters d and e increase Imitators have higher aspirations on average than explorers, with differences more pronounced at higher values of e Both explorers and imitators perform worse in crowded environments, especially when they are the majority
Satisficing (various measures)	Mean performance decreases with higher d and increases with higher e , reflecting increased competition for imitators and more initial discoveries by explorers Mean aspirations decrease over time due to increased competition and lower experienced performance of imitators The performance-aspiration gap is imitators-driven, yet imitators depend on explorers Early stages, less negative the more favourable conditions are for imitators Later stages, less negative the bigger the “legacy” left by explorers to imitators (more good alternatives discovered)

In summary, the performance-aspiration gap highlights the significance of the early knowledge provided by explorers for imitators throughout the simulations. This significance is contingent on the intensity of competition, primarily driven by the total number of agents in the simulated models.

Table 3 summarizes all the results reported so far.

4.5 Introduction of environmental turbulence

Additional simulations were conducted employing the previously detailed models with a notable modification: the environment was rendered more turbulent, signifying continuous change in all aspects. Specifically, the mean performance of each option was regenerated at every simulation period. In this dynamic setting, the stochastic nature of the initial model setup was evident in the variability of all variables.

To streamline our presentation, we concentrate on two primary metrics: the number of agents per type and the mean performance levels observed within these turbulent simulations.

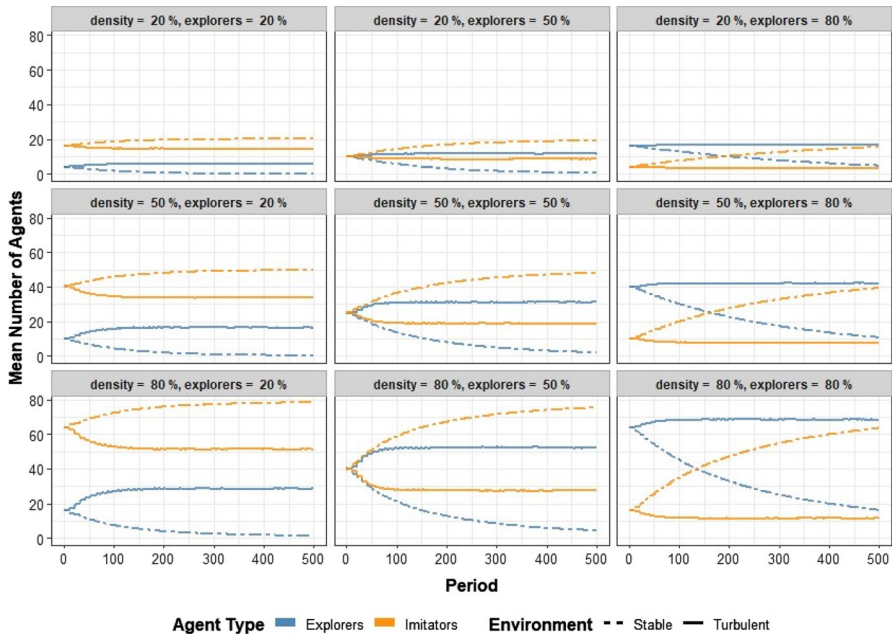


Fig. 11 Mean counts of agents per type: stable vs turbulent environments

A more comprehensive analysis and discussion of these findings will be elaborated upon in the subsequent section of this manuscript.

Our agent modeling approach places “satisficers” in a challenging scenario characterized by a dynamically fluctuating environment. Figure 11 compares the counts of the two agent types computed in the simulations with a turbulent environment juxtaposed with the previously reported counts in a stable one (Fig. 2).

The number of explorers and imitators differs slightly from that dictated by the initial conditions set by parameter e . In particular, the number of explorers tends to increase, while the number of imitators decreases. This discrepancy can be easily explained by the tendency of imitators to aggregate around random options, thus promoting competition. Similarly, the mean aggregated performance (Fig. 12) shows initial variability attributed to this initial competition before aligning with the expected value of the uniform distribution of means inherent in the data-generating process of alternative options.

Consequently, imitators exit the model faster than explorers until the cumulative losses due to competition, combined with the average performance level, fall below the exit threshold. Changing this threshold would induce a corresponding shift in the curve and a reciprocal change in the number of agents per type concerning the parameter change. In other words, a higher exit threshold would result in more imitators leaving the system and vice versa, thus influencing the initial variation in performance. The same logic applies to the losses due to the competition parameter (c).

To conclude our exploration of turbulence dynamics, it is intriguing to observe that explorers outperform imitators. Specifically, the results reveal a fascinating




Fig. 12 Mean performance in simulations with a turbulent environment

nuance: while explorers outperform imitators, they do not reach the performance levels observed in a perfectly stable environment. In such an environment, explorers gradually disappear from the system. This intriguing result can be summarized by the phrase "poor but happy": despite their relatively lower performance in the turbulent environment, explorers persist in the simulations. Figure 13 provides a visual representation of this comparative analysis of results.

A methodological note is warranted. Given that we have modeled the maximum turbulence allowable in our environment, competition becomes the sole variable contributing to the observed differences. It is essential to recognize that alternative methods of modulating turbulence in the environment might lead to different results. A more detailed discussion of these methodological nuances is provided in the following section.

5 Discussion and conclusions

This study delves into the dynamics between explorers and imitators in a competitive environment, drawing upon frameworks from organizational learning (Levitt and March 1988) and strategic imitation (Lieberman and Asaba 2006). These frameworks typically view imitation as a less resource-intensive alternative to experimentation for knowledge acquisition. Only a few studies focus on the interaction between agents engaging in innovation and those adapting through imitation (Fagiolo and Dosi 2003; Boari et al. 2017; Fagiolo et al. 2020). Although

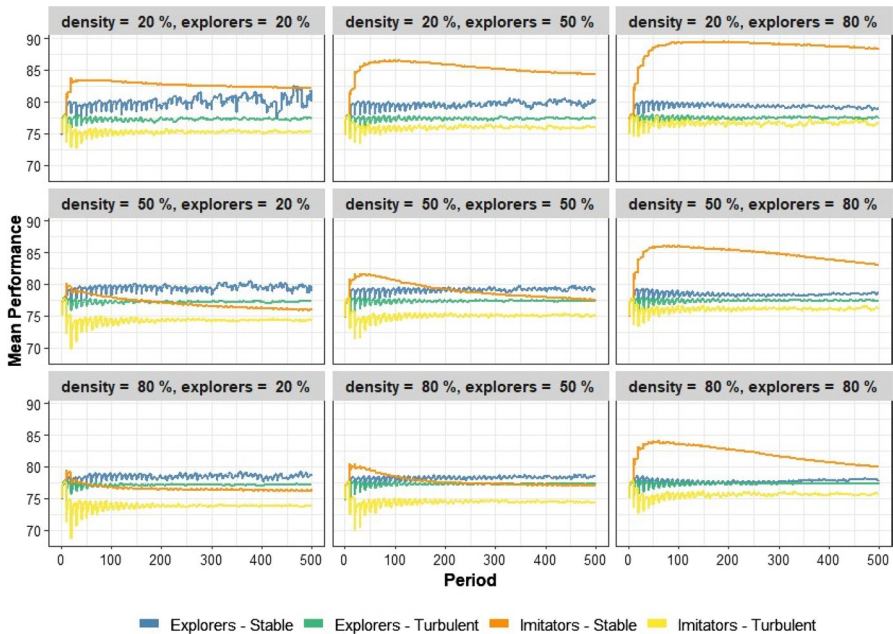


Fig. 13 Mean performance of explorers and imitators in stable and turbulent environments

these notable works have addressed the interaction between explorers and imitators, they have overlooked a crucial theoretical building block: "satisficing."

Our study utilizes agent-based modeling (ABM), an approach widely adopted in organizational studies to investigate the exploration–exploitation trade-off and imitation. Pioneering applications can be seen in the research of Posen et al. (2013, 2020), Bray and Pietrula (2007), Fang et al. (2010), Kim and Rhee (2009), Miller et al. (2006), Ponsiglione et al. (2021), and Rodan (2005). However, these studies typically examine exploration and imitation separately.

Our model synthesizes these elements and incorporates the concept of agents as "satisficers," aligning with Simon's (1955) theory. This novel approach seeks to address a gap in existing literature where the concept of "satisficing" has been largely neglected. In doing so, we reflect the realistic way individuals and organizations establish performance goals. Our approach recognizes the limitations of decision-makers' cognitive resources. The inclusion of satisficing behavior, an efficient strategy that reduces decision-making time and cognitive load, highlights its role in reaching acceptable results without the exhaustive search for perfection.

Our model delineates two distinct agent types: explorers, who undertake a random search for new alternatives, and imitators, who adopt a strategy of either majority imitation or best imitation, contingent on their comparative performance within the system. This dichotomy is rooted in the problemistic search theory of Cyert and March (1963) and influenced by the aspiration adjustment framework developed by Bromiley (1991).

This structure allows for a nuanced exploration of strategic behaviors in innovation and imitation, providing a fresh perspective on the exploration–exploitation trade-off.

The results of our simulations reveal two primary dynamics. First, the random search approach of explorers often leads to their exit from the model due to performance falling below a critical threshold. Over time, this shifts from the environment from one rich in innovation to one dominated by imitators, echoing the high failure rates observed in innovative markets, as Redmond (1995) noted. Second, the model shows that both explorers and imitators are adversely impacted by overcrowding effects, albeit in different ways. Explorers struggle due to the limited availability of rewarding alternatives, while imitators face intensified competition as the model evolves.

These findings highlight the nuanced and complex nature of agent interactions in competitive settings and demonstrate the value of incorporating the concept of satisficing into agent-based models. By doing so, our study offers a deeper understanding of the strategies and behaviors that drive innovation and imitation in competitive environments.

Moreover, we investigated agent behavior in simulations of highly turbulent environments, focusing on the adaptability and success of two types of agents: explorers and imitators. We modeled the turbulent environment to undergo continuous change, challenging the agents' ability to adapt.

Our findings revealed that due to competition, explorers and imitators deviated slightly from their initial numbers. Earlier competition led to significant performance fluctuations, eventually aligning with the expected values of the uniform distribution governing the environment's dynamism. A notable aspect of the results was the consistent superiority of explorers over imitators in performance, particularly in turbulent environments. Despite this, explorers did not achieve the performance levels observed in stable environments. This observation led to an intriguing finding: although explorers showed relatively lower success in terms of performance in the turbulent setting, they exhibited greater persistence compared to those in a stable environment. This phenomenon, termed "poor but happy," highlights the resilience of explorers (as a category) in the face of environmental uncertainty.

Our methodological approach strongly emphasized the role of competition and the search styles of agents in determining the outcomes observed in the simulations. This approach suggests that different methods of introducing turbulence could lead to varied outcomes, indicating that the findings are specific to the simulated conditions of maximum turbulence. Such insights are essential for understanding how different agent types adapt and succeed under various environmental conditions. In conclusion, this study opens to future investigations of different types of environmental turbulence and their impact on the adaptability of agents who search by either exploring or imitating.

5.1 Limitations and directions for future studies

While contributing to understanding competitive dynamics between explorers and imitators, this study presents several limitations and opens avenues for future research.

A primary limitation is the modeling of bounded rationality for imitators and explorers. They choose options randomly or heuristically, with a comprehensive knowledge of the search landscape. For instance, explorers focus solely on free options, and imitators perfectly replicate others' choices, assuming an unrealistic awareness of the overall social performance. This simplified representation contrasts with the concept of boundedly rational agents having an imperfect perception of their environment, as Puranam et al. (2015) noted. Moreover, adopting heuristics only partially fits Simon's original formulations (e.g., Simon 1955). Future research could enrich this model by introducing a more realistic, "blurred vision" for agents, limiting their awareness of others' actions. Specifically, the literature on imperfect imitation suggests incorporating sampling errors to enhance the bounded rationality of agents, offering a more authentic competitive environment. Innovations such as "first-mover advantages," in terms of intellectual property rights or knowledge benefits, could also be explored to introduce delays in the accessibility of new knowledge, echoing the work of Posen et al. (2013).

Another aspect that warrants attention is the execution of known strategies-exploitation. The current model assumes immediate maximization of returns upon choosing an option, a simplistic approach that overlooks the nuances of behavioral theory. Modeling exploitation as a refinement process could yield significant insights, aligning more closely with the principles outlined by Cyert and March (1963).

The study also reveals a limitation in the learning capabilities of agents attributed to the design of the agent-based model. Agents adapt to immediate options and general performance levels rather than engaging in comprehensive learning. Addressing this could involve implementing a more intricate individual learning model, as suggested by Lieberman and Asaba (2006) and Posen et al. (2013, 2020), where imitators sample reference groups from society, mirroring the concept of adaptive bounded rationality.

Moreover, this study's distinction between exploration and imitation differs from previous research. For instance, Fagiolo and Dosi (2003) and Fagiolo et al. (2020) treat exploration as a probability parameter, with imitation being a secondary outcome. Our model, by contrast, assigns fixed behavioral types to agents, limiting their ability to choose between exploration and imitation. This presents a theoretical gap in understanding what motivates agents to select one strategy over the other, an area ripe for exploration in future studies.

Additionally, ABMs often lack individual-level parameters for agents, such as risk attitudes, which are crucial determinants in exploration-exploitation decisions (March 1991). Incorporating such idiosyncrasies could facilitate a more nuanced understanding of exploratory and imitative behaviors at an individual level.

This study represents an initial foray into modeling imitation within a competitive exploration-exploitation environment using "satisficers." Despite its shortcomings, it merges various theoretical frameworks, providing a foundation for more comprehensive models. Future research could significantly benefit from calibrating ABMs with experimental data, fostering a collaborative approach between experimentalists and modelers to illuminate the complex interplay of exploration, imitation, and exploitation.

Appendix 1: Additional information and simulations

Fixed Parameters and their values

We decided to keep some parameters fixed for two general reasons.

First, these parameters are not the central focus of the paper, so we decided to be parsimonious in the number of varying model characteristics we included in the model.

Second, given the lower relevance of these parameters with respect to the research question of the present paper, we chose values that were appropriate in this sense. In other words, we chose values that would have kept the effects of these parameters reasonably mild with respect to those of the variable parameters.

We justify these choices below.

The “number of alternatives” was arbitrarily chosen to be 100. However, we believe that choosing different values for this parameter would produce results quite similar to those obtained by using the value 100, but on a different scale. This is because most of the parameters in this agent-based model are relative to the number of options.

Both the “Top Performance Quantile” and “Exit Average Performance” parameters were briefly introduced and motivated in the “The Model” section of the paper. Both parameters are related to competition.

On the one hand, the number of best competitors is higher when the percentile is lower (e.g., 0.90) and decreases when the selected percentile is higher (e.g., 0.99). Choosing the 95th percentile ensures that imitation does not dilute or exacerbate competition due to the number of best competitors.

On the other hand, the performance threshold that determines an agent’s exit from the simulation models the severity of the competition itself. We chose a reasonably low percentile of the possible means generated for the options. The value of 60 approximates the 10th percentile (the actual value is 59) of the theoretical distribution of means of the discrete uniform distribution, which ranges from 55 to 95.

To evaluate how initial conditions affect the number of imitators and explorers in the system, we chose to replace agents that exited the models by introducing a conservative tendency of the model to reproduce the initial conditions themselves. In other words, we chose to replace agents to observe dynamics that we could not otherwise: see, for example, the tendency of explorers to exit the model despite high values of the e parameter.

The algorithm of the single agent

Agents in our simulations executed the algorithm shown in Fig. 12. At each period, agents were activated in a random order and started by drawing a performance level from the distribution of the option at hand. In the very first period of the simulation, a random option was assigned to each agent (Fig. 14).

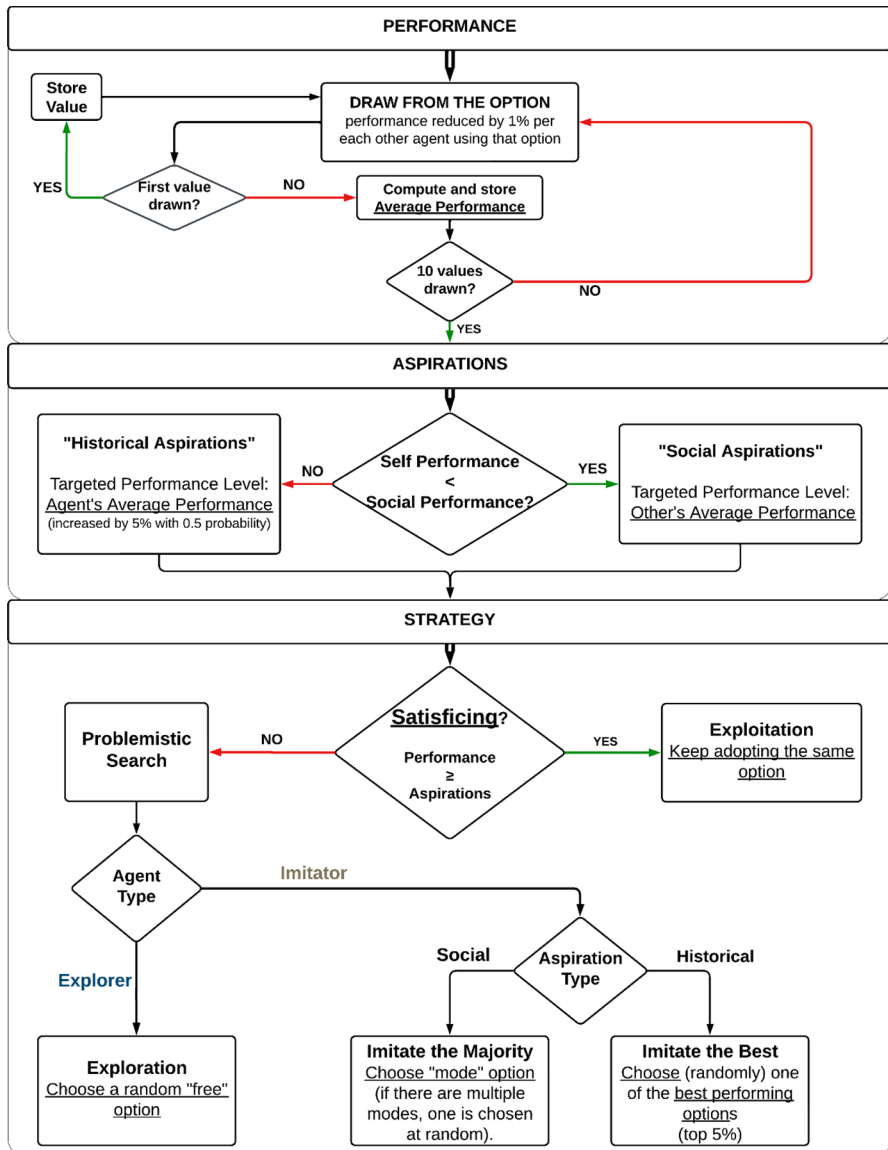


Fig. 14 Agent algorithm. Each agent in the simulation executed the algorithm illustrated in the flowchart.

Comparisons between explorers and imitators

The following are the graphical representations of the variables of interest in the model discussed in the paper computed for explorers and imitators.

For each variable, two graphical outcomes were produced: the first illustrates the mean values of that variable at each step across simulations as well as the standard

deviations; the second depicts those very same mean values that are shown conditional to being significantly different. Statistical significance has been assessed by computing paired Wilcoxon tests between the series.

Exploitation decisions

See Figs. 15 and 16.




Fig. 15 Mean Exploitation of Explorers and Imitators




Fig. 16 Results of Paired Wilcoxon Test on mean exploitation. Non-significant differences are omitted

Performance

See Figs. 17 and 18.

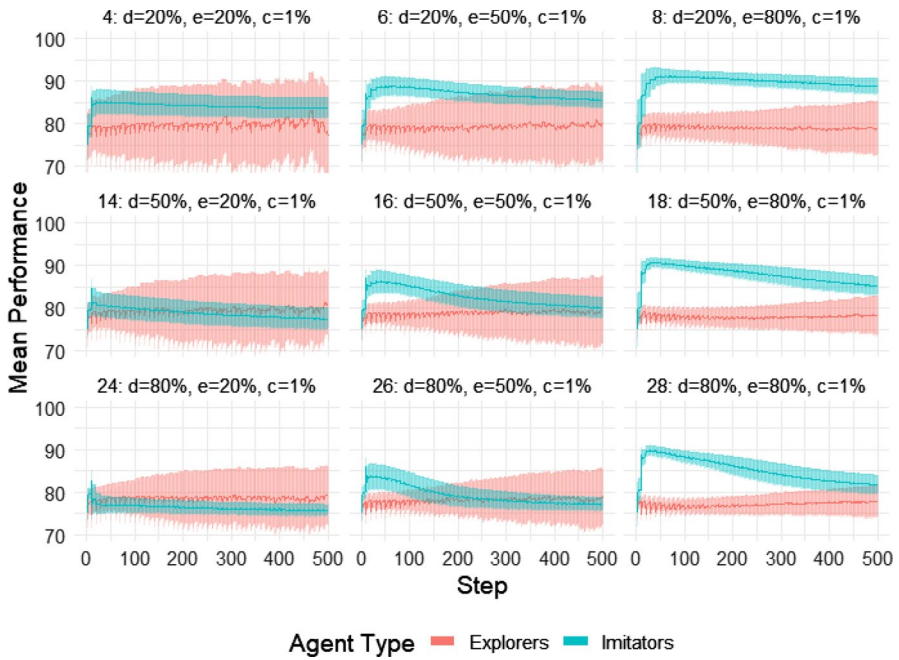


Fig. 17 Mean performance of explorers and imitators. Mean values computed for each period across one thousand simulations per model specification (i.e., at varying levels of density and share of explorers)

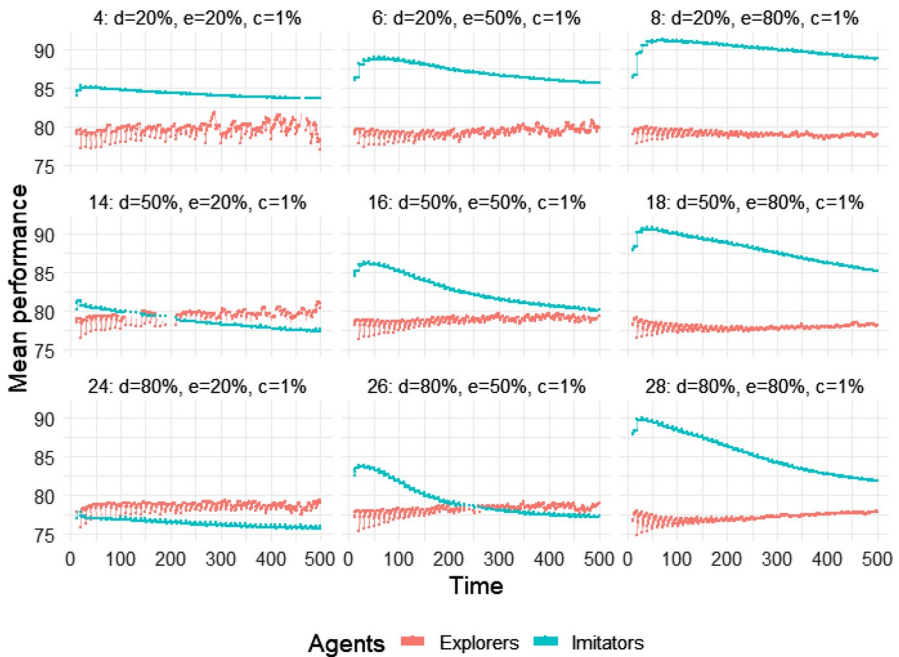


Fig. 18 Results of Paired Wilcoxon Test on mean performance. Non-significant differences are omitted

Aspirations

See Figs. 19 and 20.

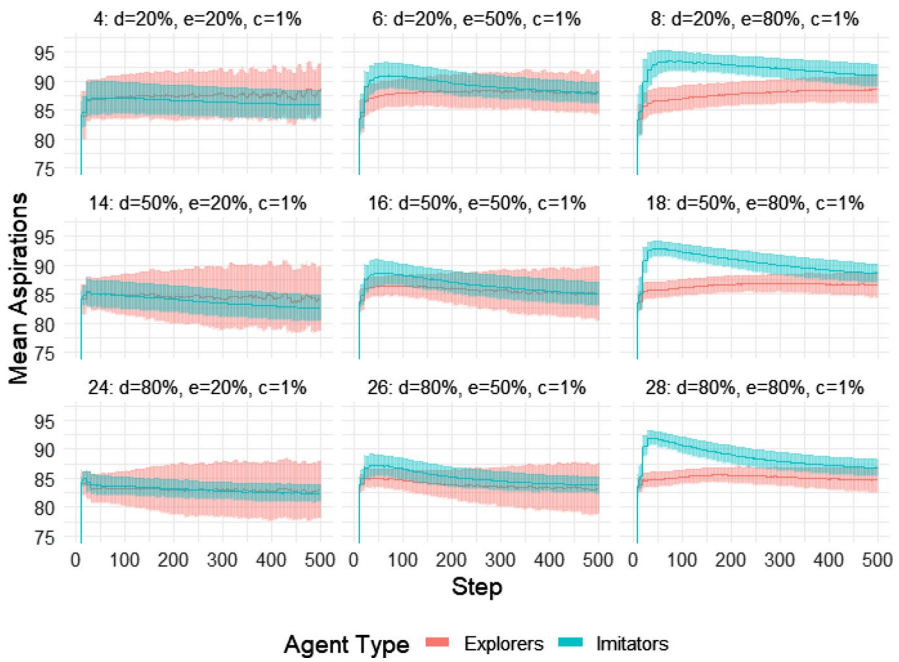


Fig. 19 Mean Aspirations of Explorers and Imitators

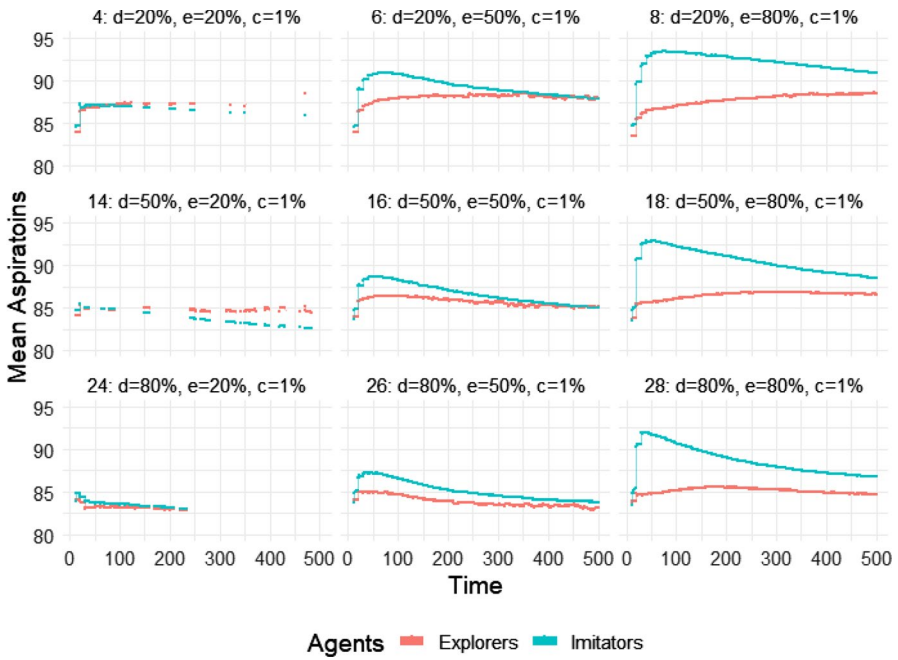


Fig. 20 Results of Paired Wilcoxon Test on mean aspirations. Non-significant differences are omitted

Performance-Aspirations gap

See Figs. 21 and 22.




Fig. 21 Mean Performance-Aspirations Gap of Explorers and Imitators




Fig. 22 Results of Paired Wilcoxon Test on mean performance-aspirations gap. Non-significant differences are omitted

Losses

See Figs. 23 and 24.

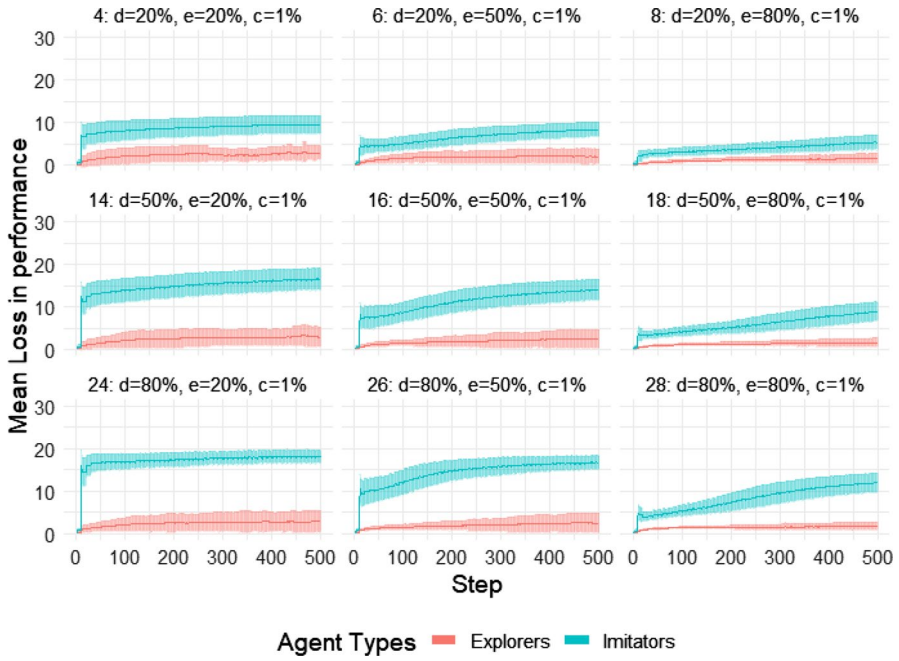


Fig. 23 Mean Losses of Explorers and Imitators

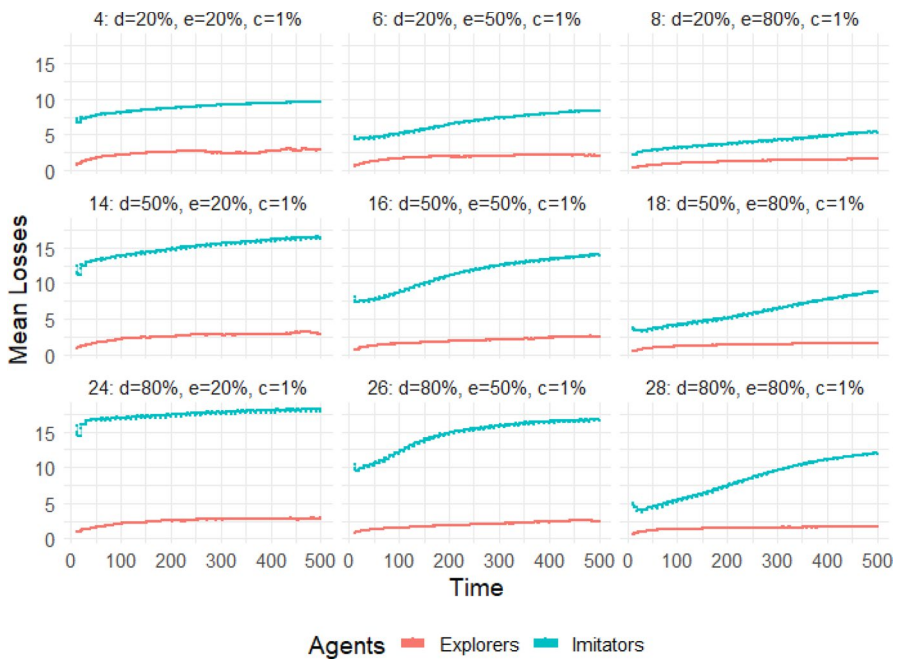


Fig. 24 Results of Paired Wilcoxon Test on mean losses. Non-significant differences are omitted

Cumulative exits

See Figs. 25 and 26.




Fig. 25 Mean Cumulative Exits of Explorers and Imitators

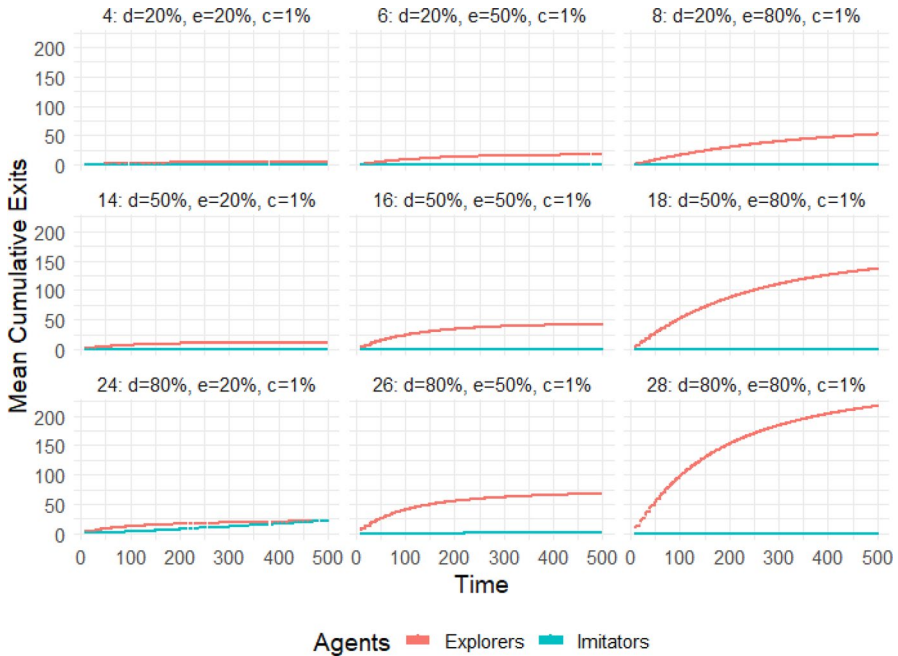


Fig. 26 Results of Paired Wilcoxon Test on mean cumulative exits. Non-significant differences are omitted

Appendix 2: Sensitivity checks

The following is a set of statistical comparisons between those simulated models that have been reported in the present paper. Specifically, these sensitivity checks have been conducted for the nine models with static environments described in the Results section. We have analyzed the statistical differences between each simulated model as the parameters d and e varied across the different specifications.

For each of the analyzed variables, we ran two types of statistical tests and provided a graphical representation of their results. We performed a Kruskal–Wallis Test and a Paired Wilcoxon Test between “triplets” of model specifications. These triplets consist of sets of three models differing for one parameter while the other is kept constant.

For example, triplet A consists of models with a constant value of $d=2$ and, hence, all the considered values of e , namely, 0.2, 0.5, and 0.8.

In contrast, triplet D includes three models with a constant parameter $e=0.2$, while d varies across the three models by taking the values of 0.2, 0.5, and 0.8, respectively.

Finally, a figure representing the compared models in terms of the considered variable is provided. In such a figure, solid line points indicate that the three models have emerged as statistically different in terms of that variable at that step. On the contrary, discontinuities in the lines indicate that at least two out of the three models are not statistically different when that variable is considered at that time step.

Performance

See Figs. 27, 28 and 29.

P-values are plotted.




Fig. 27 Results of Kruskal–Wallis’s test on the performance levels across triplets. P-values are plotted

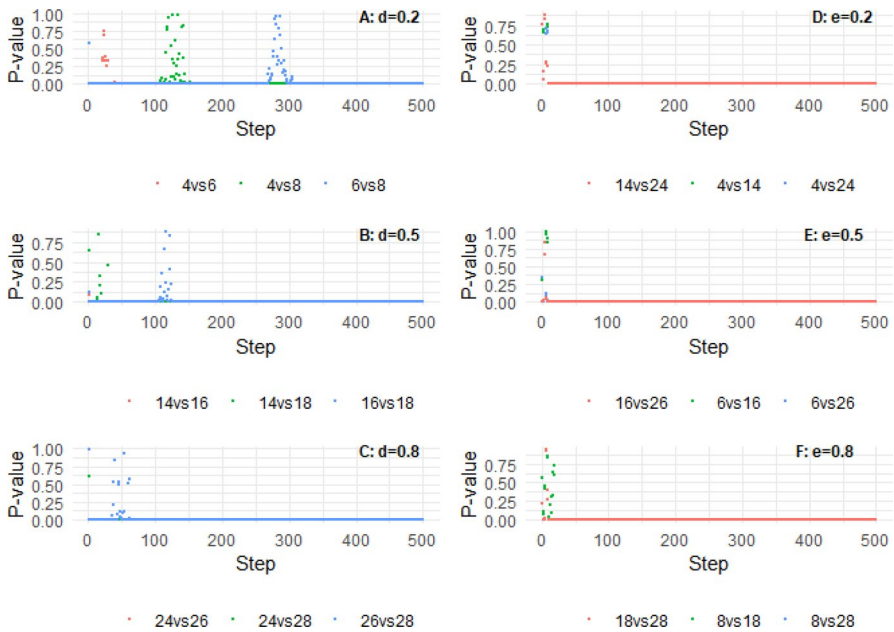


Fig. 28 Results of paired Wilcoxon test on the performance levels across triplets of models

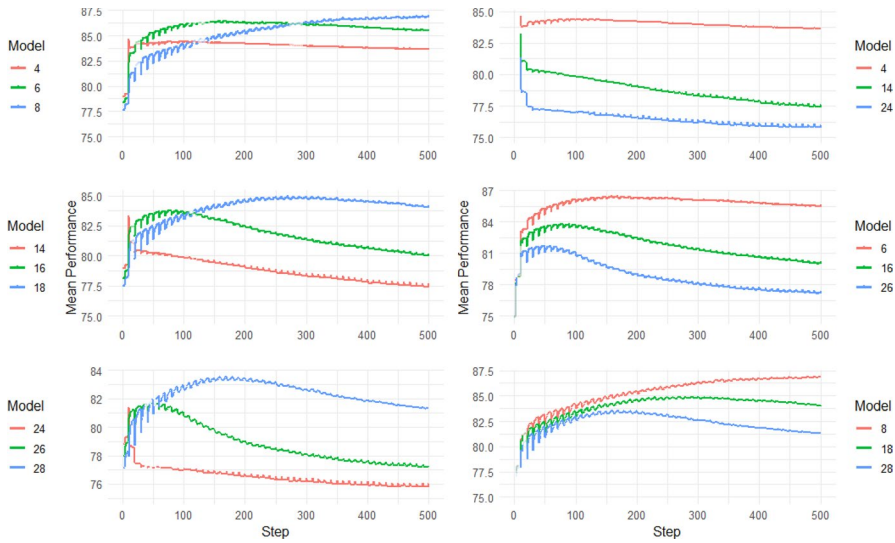


Fig. 29 Comparisons between models in terms of mean performance. Transparent dots indicate absence of statistical significance in the comparison between at least two of the three models

Aspirations

See Figs. 30, 31 and 32.

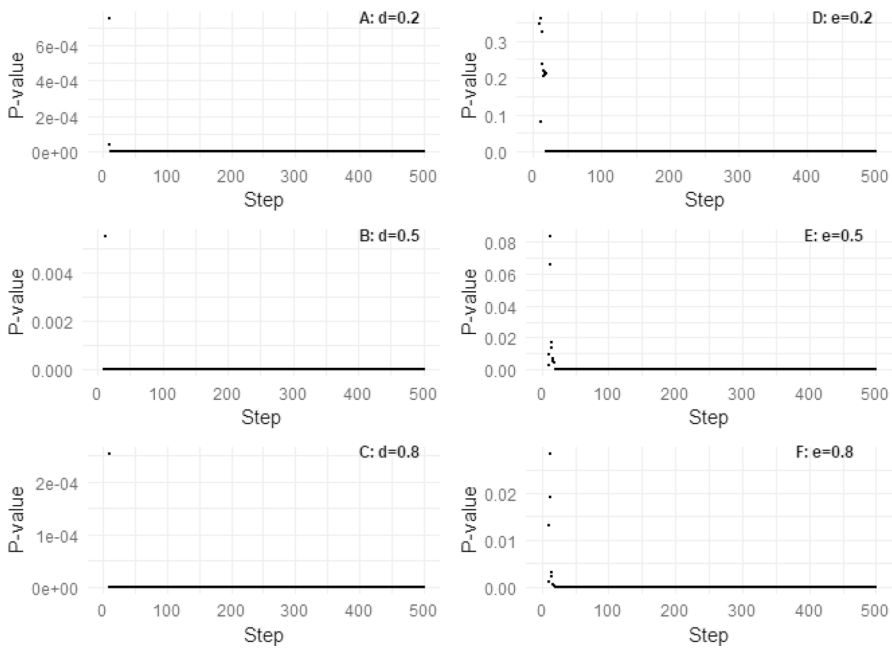


Fig. 30 Results of Kruskal–Wallis’s test on the aspirations levels across triplets of models. P-values are plotted

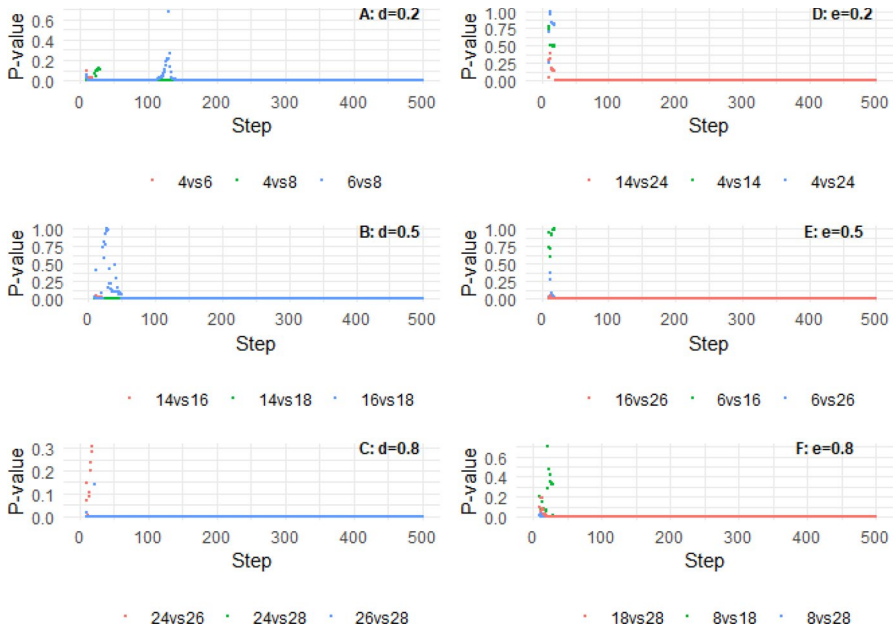


Fig. 31 Results of paired Wilcoxon test on the aspirations levels across triplets of models where parameters d and e are constant as indicated in each plot. P-values are plotted




Fig. 32 Comparisons between models in terms of mean aspirations. Transparent dots indicate absence of statistical significance in the comparison between two of the three models

Performance-Aspirations gap

See Figs. 33, 34 and 35.

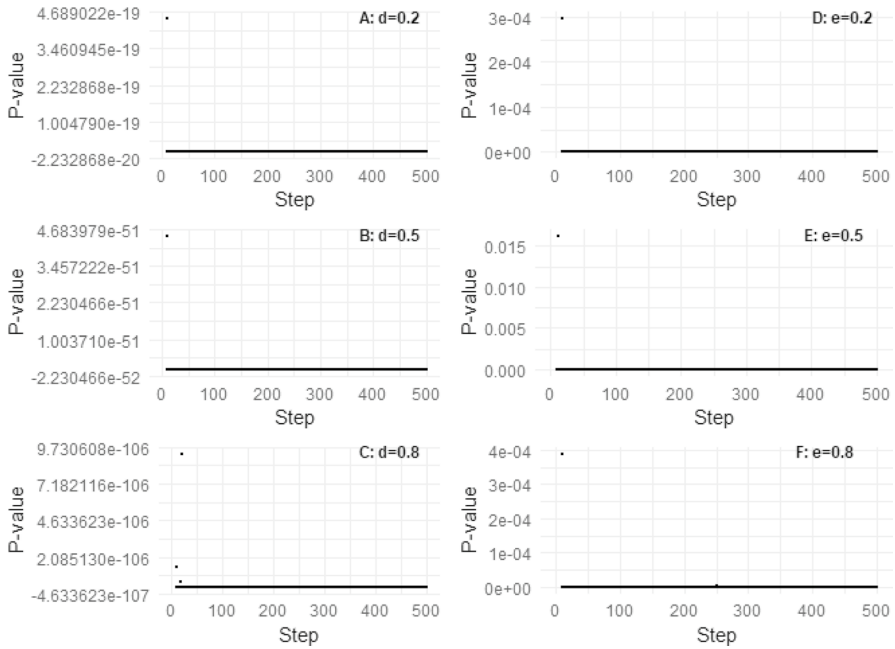


Fig. 33 Results of Kruskal–Wallis’s test on the performance-aspirations gap across triplets of models where parameters d and e are constant as indicated in each plot. P-values are plotted

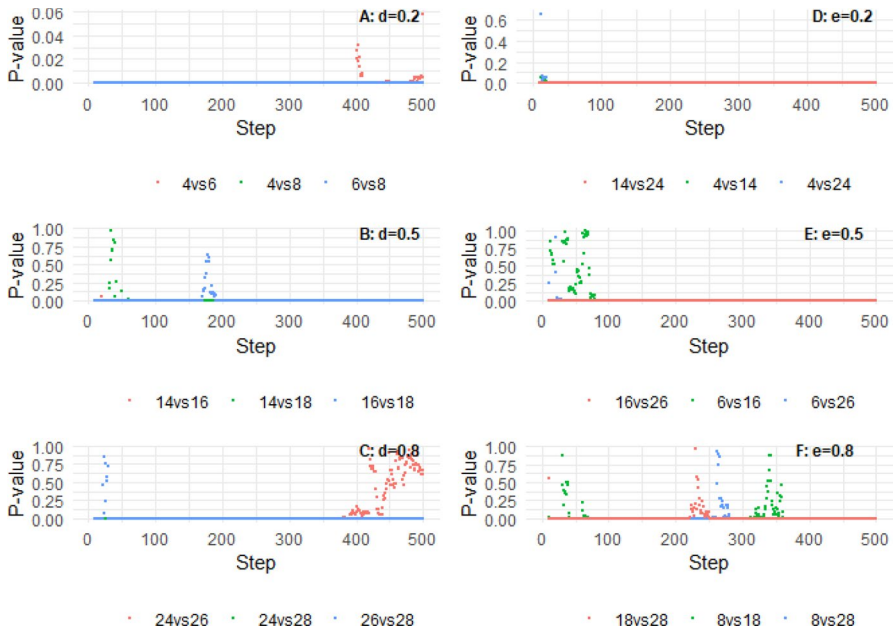


Fig. 34 Results of paired Wilcoxon test on the performance-aspirations gap across triplets of models where parameters d and e are constant as indicated in each plot. P-values are plotted

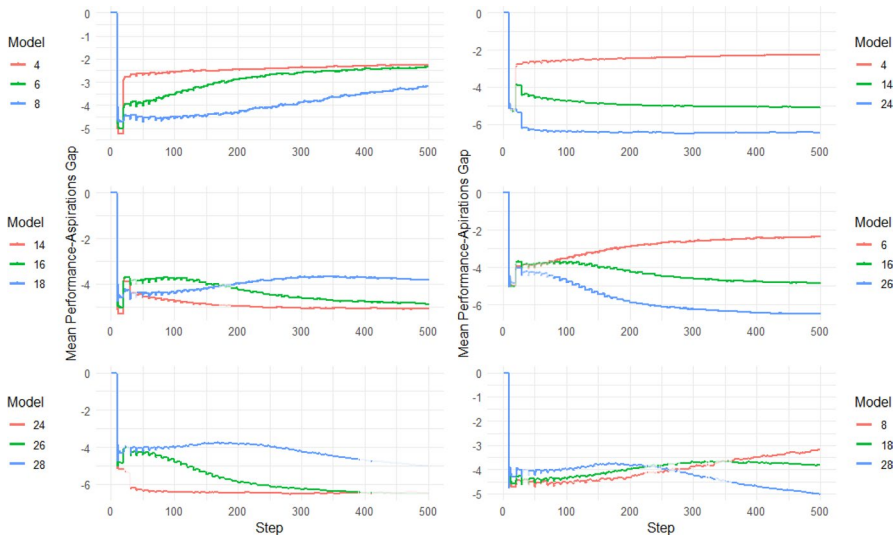


Fig. 35 Comparisons between models in terms of mean performance-aspirations gap. Transparent dots indicate absence of statistical significance in the comparison between two of the three models

Losses

See Figs. 36, 37 and 38.

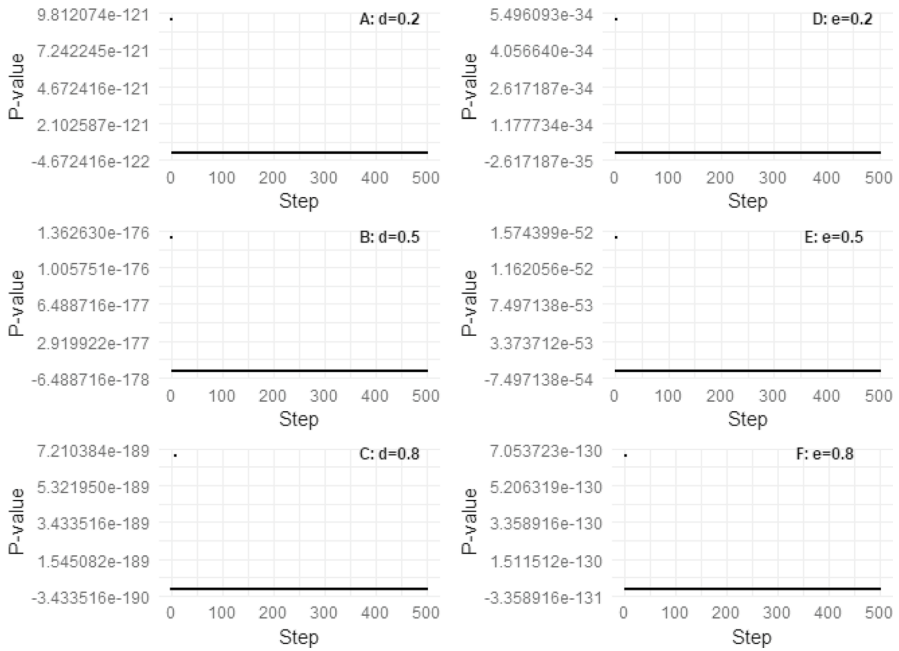


Fig. 36 Results of Kruskal–Wallis’s test on the mean losses across triplets of models where parameters d and e are constant as indicated in each plot. P-values are plotted

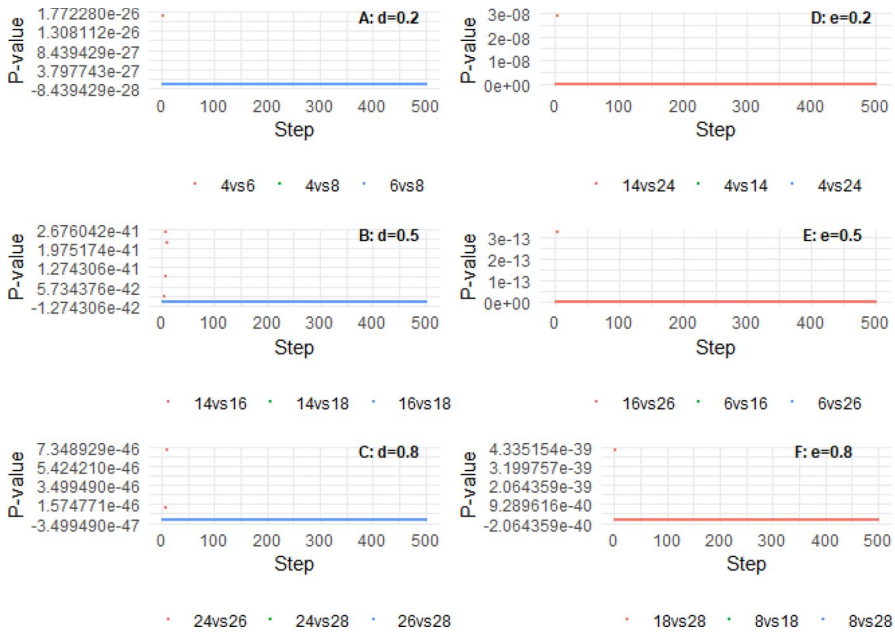


Fig. 37 Results of paired Wilcoxon test on the mean losses across triplets of models where parameters d and e are constant as indicated in each plot. P-values are plotted




Fig. 38 Comparisons between models in terms of mean losses. Transparent dots indicate absence of statistical significance in the comparison between two of the three models

Exits

See Figs. 39, 40 and 41.

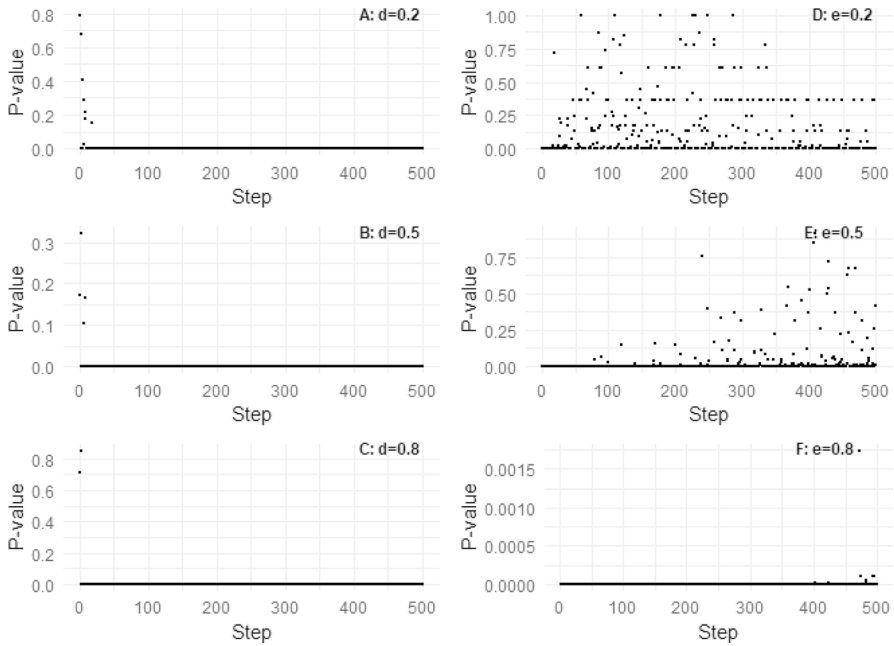


Fig. 39 Results of Kruskal–Wallis’s test on the mean exits across triplets of models where parameters d and e are constant as indicated in each plot. P-values are plotted

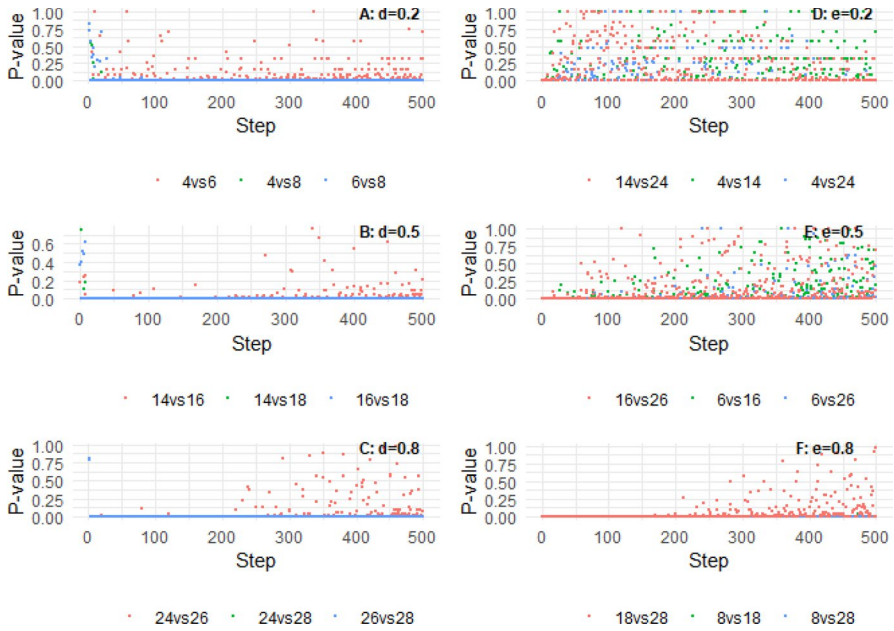


Fig. 40 Results of paired Wilcoxon test on the mean exits across triplets of models where parameters d and e are constant as indicated in each plot. P-values are plotted

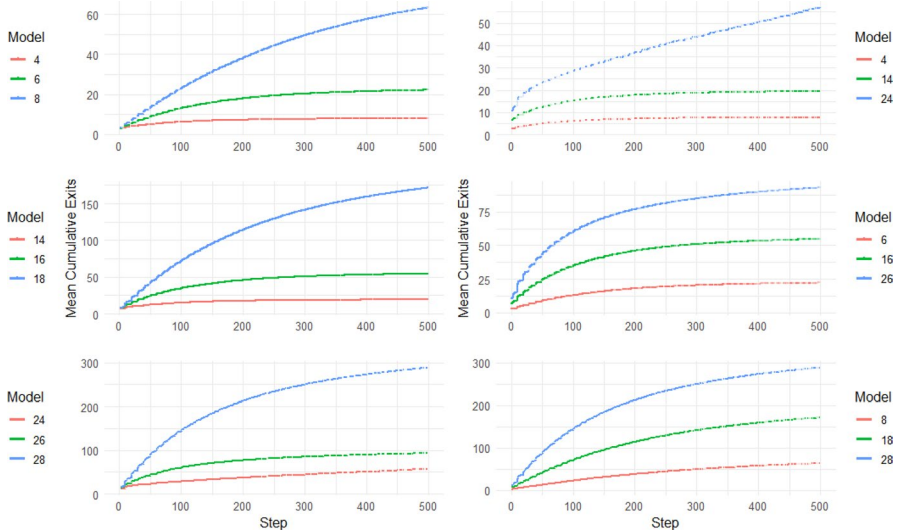


Fig. 41 Comparisons between models in terms of mean exits. Transparent dots indicate absence of statistical significance in the comparison between two of the three models

Funding Open access funding provided by Università degli Studi di Trento within the CRUI-CARE Agreement.

Declarations

Conflict of interest The authors have no conflicts of interest to declare.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Argote L, Lee S, Park J (2020) Organizational learning processes and outcomes: major findings and future research directions. *Manag Sci* 67:5399–5429
- Bandura A, Walters RH (1977) *Social learning theory*, vol 1. Prentice Hall, Englewood cliffs
- Baum JA, Rowley TJ, Shipilov AV, Chuang YT (2005) Dancing with strangers: aspiration performance and the search for underwriting syndicate partners. *Adm Sci Q* 50(4):536–575
- Billinger S, Stieglitz N, Schumacher TR (2014) Search on rugged landscapes: an experimental study. *Organ Sci* 25(1):93–108
- Billinger S, Srikanth K, Stieglitz N, Schumacher TR (2021) Exploration and exploitation in complex search tasks: how feedback influences whether and where human agents search. *Strateg Manag J* 42(2):361–385
- Boari C, Fioretti G, Odorici V (2017) A model of innovation and knowledge development among boundedly rational rival firms. *Team Perform Manag Int J* 23:82–95
- Bray D, Prietula M (2007) Extending March's exploration and exploitation: managing knowledge in turbulent environments
- Bromiley P (1991) Testing a causal model of corporate risk-taking and performance. *Acad Manag J* 34(1):37–59
- Bromiley P, Harris JD (2014) A comparison of alternative measures of organizational aspirations. *Strateg Manag J* 35(3):338–357
- Chang MH, Harrington JE Jr (2006) Agent-based models of organizations. *Handb Comput Econ* 2:1273–1337
- Cyert RM, March JG (1963) *A behavioral theory of the firm*. Englewood Cliffs, NJ, 2(4):169–187
- DiMaggio PJ, Powell WW (1983) The iron cage revisited: institutional isomorphism and collective rationality in organizational fields. *Am Sociol Rev* 48:147–160
- Duysters G, Lavie D, Sabidussi A, Stettner U (2020) What drives exploration? Convergence and divergence of exploration tendencies among alliance partners and competitors. *Acad Manag J* 63(5):1425–1454
- Epstein JM (1999) Agent-based computational models and generative social science. *Complexity* 4(5):41–60
- Fagiolo G, Dosi G (2003) Exploitation, exploration and innovation in a model of endogenous growth with locally interacting agents. *Struct Chang Econ Dyn* 14(3):237–273
- Fagiolo G, Giachini D, Roventini A (2020) Innovation, finance, and economic growth: an agent-based approach. *J Econ Interac Coord* 15(3):703–736
- Fang C, Lee J, Schilling MA (2010) Balancing exploration and exploitation through structural design: the isolation of subgroups and organizational learning. *Organ Sci* 21(3):625–642
- Fioretti G (2013) Agent-based simulation models in organization science. *Organ Res Methods* 16(2):227–242

- Giannoccaro I, Galesic M, Massari GF, Barkoczi D, Carbone G (2020) Search behavior of individuals working in teams: a behavioral study on complex landscapes. *J Bus Res* 118:507–516
- Gigerenzer G (1997) Bounded rationality: Models of fast and frugal inference. *Swiss J Econ Stat* 133(2/2):201–218
- Gigerenzer G, Selten R (eds) (2002) Bounded rationality: the adaptive toolbox. MIT Press, Cambridge
- Gómez-Cruz NA, Saa IL, Hurtado FFO (2017) Agent-based simulation in management and organizational studies: a survey. *Eur J Manag Bus Econ* 26(3):313–328
- Greve HR (1998) Performance, aspirations, and risky organizational change. *Admin. Sci. q.* 43:58–86
- Greve HR (2003a) A behavioral theory of R&D expenditures and innovations: evidence from shipbuilding. *Acad Manag J* 46(6):685–702
- Greve HR (2003b) Organizational learning from performance feedback: a behavioral perspective on innovation and change. Cambridge University Press, Cambridge
- Greve HR, Gaba V (2017) Performance feedback in organizations and groups: common themes. *The Oxford handbook of group and organizational learning*, pp 1–45
- Gross D, Strand R (2000) Can agent-based models assist decisions on large-scale practical problems: a philosophical analysis. *Complexity* 5(5):26–33
- Gupta AK, Smith KG, Shalley CE (2006) The interplay between exploration and exploitation. *Acad Manag J* 49(4):693–706
- Harris J, Bromiley P (2007) Incentives to cheat: The influence of executive compensation and firm performance on financial misrepresentation. *Organ Sci* 18(3):350–367
- Kim T, Rhee M (2009) Exploration and exploitation: Internal variety and environmental dynamism. *Strateg Organ* 7(1):11–41
- Lant TK (1992) Aspiration level adaptation: an empirical exploration. *Manage Sci* 38(5):623–644
- Laureiro-Martinez D, Brusoni S, Tata A, Zollo M (2019) The manager's notepad: Working memory, exploration, and performance. *J Manag Stud* 56(8):1655–1682
- Lave CA, March JG (1975) An introduction to models in the social sciences. Harper and Row, New York
- Levinthal DA (1997) Adaptation on rugged landscapes. *Manage Sci* 43(7):934–950
- Levitt B, March JG (1988) Organizational learning. *Ann Rev Sociol* 14(1):319–338
- Lieberman MB, Asaba S (2006) Why do firms imitate each other? *Acad Manag Rev* 31(2):366–385
- March JG (1991) Exploration and exploitation in organizational learning. *Organ Sci* 2(1):71–87
- Miller KD (2015) Agent-based modeling and organization studies: a critical realist perspective. *Organ Stud* 36(2):175–196
- Miller KD, Zhao M, Calantone RJ (2006) Adding interpersonal learning and tacit knowledge to March's exploration-exploitation model. *Acad Manag J* 49(4):709–722
- Mittone L, Morreale A, Ritala P (2023) Initial conditions and path dependence in explorative and exploitative learning: an experimental study. *Technovation* 129:102895
- Mom TJ, Van Den Bosch FA, Volberda HW (2007) Investigating managers' exploration and exploitation activities: the influence of top-down, bottom-up, and horizontal knowledge inflows. *J Manag Stud* 44(6):910–931
- Nelson RR, Winter SG (1982) The Schumpeterian tradeoff revisited. *Am Econ Rev* 72(1):114–132
- Nikolaeva R (2014) Interorganizational imitation heuristics arising from cognitive frames. *J Bus Res* 67(8):1758–1765
- Ponsiglione C, Cannavacciuolo L, Primario S, Quinto I, Zollo G (2021) The ambiguity of natural language as resource for organizational design: a computational analysis. *J Bus Res* 129:654–665
- Posen HE, Lee J, Yi S (2013) The power of imperfect imitation. *Strateg Manag J* 34(2):149–164
- Posen HE, Keil T, Kim S, Meissner FD (2018) Renewing research on problemistic search—a review and research agenda. *Acad Manag Ann* 12(1):208–251
- Posen HE, Yi S, Lee J (2020) A contingency perspective on imitation strategies: when is “benchmarking” ineffective? *Strateg Manag J* 41(2):198–221
- Puranam P, Stieglitz N, Osman M, Pillutla MM (2015) Modeling bounded rationality in organizations: progress and prospects. *Acad Manag Ann* 9(1):337–392
- Redmond WH (1995) An ecological perspective on new product failure: the effects of competitive overcrowding. *J Prod Innov Manag* 12(3):200–213
- Rivkin JW (2000) Imitation of complex strategies. *Manag Sci* 46(6):824–844
- Rodan S (2005) Exploration and exploitation revisited: extending March's model of mutual learning. *Scand J Manag* 21(4):407–428
- Simon HA (1955) A behavioral model of rational choice. *Q J Econ* 69(1):99–118
- Simon HA (2013) Administrative behavior. Simon and Schuster, New York

Wall F (2016) Agent-based modeling in managerial science: an illustrative survey and study. *RMS* 10(1):135–193

Wilson RC, Geana A, White JM, Ludvig EA, Cohen JD (2014) Humans use directed and random exploration to solve the explore–exploit dilemma. *J Exp Psychol Gen* 143(6):2074–2081

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.