Libera Università Internazionale
degli Studi Sociali Guido Carli
PREMIO TESI D'ECCELLENZA

To Explore or to Exploit? An Experimental Study on the Effects that Autonomy and Control Exert on Individual Search Behavior in Complex Fitness Landscapes

Vittoria Di Marcantonio

2021-2022

LUISS

Libera Università Internazionale
degli Studi Sociali Guido Carli

Working Papaer n. 3/2021-2022
Publication date: January 2024
To Explore or to Exploit? An Experimental Study on the Effects that Autonomy and Control Exert on Individual Search Behavior in Complex Fitness Landscapes

ISBN 979-12-5596-087-4

Luiss Academy is an imprint of
Luiss University Press – Pola Srl
Viale Pola 12, 00198 Roma
Tel. 06 85225485
E-mail lup@luiss.it
www.luissuniversitypress.it

To Explore or to Exploit? An Experimental Study on the Effects that Autonomy and Control Exert on Individual Search Behavior in Complex Fitness Landscapes

By Vittoria Di Marcantonio

ABSTRACT

Individual search behavior represents a fundamental construct to understand the mechanisms of organizational learning and development, both in established ventures and in unfolding start-ups. In these contexts, agents do not act in complete autonomy but are guided by organizational objectives, structures and resource constrains. From these limitations, it emerges what has been defined as a dilemma between exploration efforts and exploitation activities.

This dilemma practically manifests itself in two orders of decisions. On one hand the tradeoff is reflected in the choice on *whether to search* and it is linked to the agents' settling of aspirations. On the other, it determines the decision of *where to search* in the space of possible solutions available to a firm. This research work addresses the literature gap regarding how organizations' structures, incentives and rewards, through the effect of feedback, impact on decision-makers' search behavior and on how autonomy and control structures influence individuals' ability to balance proximity and distant search. Additionally, this work accounted for the moderating effect that the introduction of a penalty has on the decisions to explore or exploit, a condition mostly unexplored in the literature.

In order to address these questions, an empirical pilot experiment, built on the basis of the NK model framework, was implemented.

The results from the experiment allow the present work to make a series of contributions. In line with the existing literature, agents stop their search process once their aspirations are met, rather than keep on searching to achieve a global optimum. Additionally, performance feedback determines where in the search space agents will look for improvements. Finally, this research enhances current understanding of individuals' search behavior with respect to the introduction of a penalty. The reported findings have also practical implications for both established firms and upcoming ventures.

INTRODUCTION

Individual search behavior represents one of the fundamental constructs to understand the mechanisms of organizational learning and development which guide the processes of innovation and change, both in established ventures and in unfolding start-ups. These structures are characterized by the presence of a series of interdependencies among their constituting elements that need to be recombined on a ceaseless basis in order to guarantee a fit between internal capabilities and external circumstances. In these contexts, in fact, individual aspirations are crucial in determining the degree of change required to adapt to ever evolving landscapes. Nonetheless, agents do not act in complete autonomy, especially when considering settled companies, but are guided by organizational objectives, structures and incentives. Additionally, these tensions between individual aspirations and settled targets take place in resource constrained settings, that allow firms to focus only on a limited set of objectives at a time. From these contrasts, it emerges what originally March (1991) defined as a dilemma between *exploration efforts* – connected to novelty, experimentation and innovation – necessary to identify future avenues and to ensure a firm's viability, and *exploitation activities* – linked to refinement and efficiency – required to leverage on current strengths.

This dilemma practically manifests itself in two orders of decisions. On one hand the tradeoff is reflected in the choice on *whether to search* (Billinger, Srikanth, Stieglitz and Schumacher, 2021). This question finds an answer in the stream of research developed around the behavioral theory of the firm (Cyert and March, 1963) and problemistic search theory (Posen, Keil, Kim and Meissner, 2018; Denrell and March, 2001; Levinthal and March, 1981). According to the behavioral theory, organizations define and adjust their objectives in accordance with a set of reference points, which can either be targets or aspirational levels. Search mechanisms, in turn, depend on targets and aspirational levels against which a firm evaluates its actual results. In line with problemistic search theory, then, a firm learns from the feedback received on its previous performance. If the target is above actual performance this will trigger search for alternative courses of action, whereas performance above the target restricts search (Posen et al., 2018; Billinger,Stieglitz and Schumacher, 2014; Denrell & March, 2001; Cyert and March, 1963; Simon, 1957, 1959).

On the other side, the decision to explore or to exploit is equal to choose *where to search* in the space of possible alternatives available to a firm, so to opt for narrow or rather distant search (Billinger et al., 2021). This line of research developed around the conceptual framework provided by the NK model (Marengo et al., 2022; Baumann et al., 2019; Billinger et al., 2014, 2021; Rivkin and Siggelkow, 2003,2007; Gavetti, 2005; Levinthal, 1997; Kauffman, 1993). As developed by Levinthal (1997) in its application in economics, the model defines a fitness landscape through two parameters N and K. An organization is defined by N attributes and each attribute can assume two possible values. The variable K determines the degree to which the fitness of the organization depends on the interrelatedness between the attributes, and therefore the complexity of the task. A general result that emerges in the literature is that as the level of interactions among organizational elements increases, the number of local optima escalates and engaging in exploration efforts becomes a successful strategy in order to escape from those optima (Kauffman, 1993; Levinthal, 1997; Rivkin and Siggelkow, 2003).

An essential contribution, providing useful insights to integrate these set of decisions, comes from the work of Billinger et al. (2021). According to the authors, it is possible to unify these views by considering these choices not as independent of each other but rather as interrelated. Specifically, the decision on whether to search comes from the aspirations-performance gap. If this gap actually exists, the subsequent decision will involve considerations on where to search in the space of alternatives available to a firm.

Additionally, in recent years there has been a growing interest on the effect that individual predispositions and characteristics have on the individuals' ability to explore and exploit. Within an organization, individuals have less autonomy on how to allocate their activities between exploration and exploitation. Nonetheless, a directive approach could be useful to orient attention, to redirect strategy and support individuals in balancing search efforts (Bidmon and Boe -Lillegraven, 2020; Tempelaar and Rosenkranz, 2019; Blettner, He,Hu and Bettis, 2015; Laureiro-Martìnez, Brusoni, Canessa and Zollo, 2015).

In line with these new developments, this research work addresses the literature gap identified by Billinger et al. (2021) regarding how organizations' structures, incentives and rewards, through the effect of feedback, impact on decision-makers' search behavior and the research path suggested by Bidmon and Boe-Lillegraven (2020) on how autonomy and control structures influence individuals' ability to balance proximity and distant search. Therefore, in the present work the following research question was addressed:

"To what extent can autonomy and control, through their effect on feedback and task complexity, influence individual decision-makers search behavior?"

Additionally, this work accounted for the moderating effect that the introduction of a penalty has on the decisions to explore or exploit, a condition mostly unexplored in the literature. Previous experiments, in fact enacted a problem of "pure search", as in the case of Billinger et al. (2014) in which engaging in additional search efforts was not associated with a downside risk. Subsequently the following research question was introduced:

"What is the effect that the introduction of a penalty has on the relationship between feedback and search breadth?"

In order to address these questions, an empirical pilot experiment was implemented. This choice addresses the lack of experimental studies considering how decision-makers search across a complex problem landscape (Baumann et al., 2019). As also evidenced by Billinger et al. (2021), in fact, experimental studies investigating how individuals maintain an equilibrium between local and more remote search strategies are limited. Additionally, as already evidenced by Gupta et al. (2006), experiments investigating on the micro-foundations of exploration and exploitation are relatively scarce.

The experiment was built on the basis of the NK Model framework. Participants had to develop a business model with the objective of reaching their aspirational level of performance in what was defined the autonomy setting, whereas they needed to update the current business model of a fictional company in order to reach a previously established target in the control setting. Additionally, throughout the rounds they were faced with three different levels of complexity delineating a smooth (K=0), complex (K=2) and maximally rugged (K=5) performance landscape. Finally they were informed that a penalty of the 10% would have been applied if, by chance, they exchanged a performative attribute with a non-performative one.

The results from the experiment allow the present work to make a series of contributions to the existing literature. First of all, it finds support for one of the main assumptions of the organizational learning literature, according to which agents stop their search process once their aspirations are met, rather than keep on searching to achieve a global optimum (Posen et al., 2018; Cyert & March, 1963; March & Simon, 1958; Simon, 1955). As performance approaches individual aspirational levels, agents will tend to satisfice and decrease their search breadth, relying onto exploitation. This relation is valid for both the autonomy and control settings.

Second, it contributes to the strategy literature through finding a confirm that performance feedback from previous rounds determines where in the search space agents will look for performance improvements. Individuals tend to concentrate search in the neighborhood of current solutions, but in highly complex task environments enlarging search breadth gives more chance to improve performance (Baumann et al., 2019; Billinger et al., 2014). As the level of complexity in the landscape increases, so it does search breadth. These tendencies are even more marked in a controlled setting since the control imposed by organizational structures has an impact on performance by directing agents' search behavior on the landscape they confront, and making at the same time the result of search more effective through a clearer direction provided by the target (Baumann et al., 2019; Rivkin and Siggelkow, 2003).

Finally, this research enhances current understanding of individuals search behavior with respect to the introduction of a penalty. According to Billinger et al. (2014) human agents are inclined towards over-exploration, interrupting local search too early and sacrificing profits from local improvements. Nonetheless, scholars agree that in a setting in which search has a cost agents will tend to stop their research once satisfying combinations are found (Billinger et al., 2021; Baillon et al., 2020; Goldstein et al., 2020; Hey et al., 2017). In the autonomy setting, it appears that for the same level of average aspirations feedback, agents sensitive to the presence of a penalty focused their research in the neighborhood of known solutions, whereas those that were not affected by the penalty looked for alternative combinations in a wider area of the search landscape. In the controlled setting, the introduction of a penalty reduces the average search breadth for the same level of average performance feedback. As also evidenced by Greve (2010) with a target to be reached, the introduction of a penalty can be used to boost exploration up to the level necessary to achieve the performance target and to simultaneously inhibit search from reaching hazardous levels.

The reported findings have also practical implications for both established firms and upcoming ventures. Within established firms, providing top-down directions, especially in complex environments and for innovation-focused organizations, has a strong effect on making the research process more effective, both reducing the level of efforts needed and in terms of reaching the desired results. Regarding upcoming startups, the present work highlights that for an entrepreneur it is crucial to calibrate his/her aspiration with a level of performance attainable in response to the environmental contingencies faced. Entrepreneurs, in order to be successful, should rely on agile and modular solutions to develop business models able to adapt in accordance with different landscapes.

THEORETICAL BACKGROUND

1.1. Exploration versus Exploitation Tradeoff

The discussion around the concepts of exploration and exploitation is grounded in March's seminal work. *Exploration* is defined by the notions of "search, variation, risk taking, experimentation, play, flexibility, discovery, innovation". Conversely, *exploitation* is captured by "refinement, choice, production, efficiency, selection, implementation, execution" (March, 1991, p.71). If there has been among scholars a general consensus on the notion of exploration, the concept of exploitation is much more blurred. There is, in fact, a lack of clarity on whether exploitation refers only to the reliance upon past knowledge or if it also involves the development of some new knowledge, even though of a different kind than in exploration (Gupta et al., 2006). One group of scholars (Baum et al., 2000; Benner and Tushman, 2002; He and Wong, 2004) recognizes learning and the development of new knowledge at the core of both exploration and exploitation. On the contrary, other studies (Rosenkpof and Nerkar, 2001; Vassolo et al., 2004) establish a connection between learning and innovation exclusively with exploration, while relating the concept of exploitation to processes relying on past knowledge (Gupta et al., 2006).

Given the contrasting but nonetheless complementary nature of the notions, organizations need to maintain an adequate balance between the two, since these are competing for scarce resources (March, 1991). The basic challenge faced by an organization, in fact, is to engage in sufficient exploitation to ensure its current viability and to devote enough efforts to exploration, in order to guarantee its future survival (Levinthal and March, 1993). Hence, the conceptualization of the relationship among the two views as a "tradeoff" or "dilemma". The scarcer the resources needed to pursue both exploration and exploitation, the greater the extent to which the two will be mutually exclusive (Gupta et al., 2006).

Recently, Billinger et al. (2021) observed that the tradeoff between exploration and exploitation applies differently to two separated but interconnected decisions: on one hand the dilemma can be seen as the decision to search (exploring) versus not searching (exploiting), which can be formulated as the decision on *whether to search*. On the other hand, the tradeoff can be theorized as the choice of undertaking radical change (exploring) or rather incremental change (exploiting). The decision in this case will be focused on *where to search*, in the neighborhood of current activities or in more remote spaces (Billinger et al., 2021). The major contribution given by Billinger et al. (2021) is reconciling these two perspectives as interrelated decisions in which the decision on *whether to search* precedes the decision of *where to search*.

1.2. Problemistic Search Theory and the Role of Feedback

Problemistic search theory defines a behavioral process through which a firm learns from the feedback received on its previous performance. The fundamental idea is that the process of decision-making within organizations cannot be represented by the selection of an optimal course of action among a set of known alternatives, but rather as a process of sequential sampling to identify alternative actions (Denrell and March, 2001; Posen et al., 2018; Billinger et al., 2021). As explained by Simon in his work on bounded rationality (1957) the set of alternatives considered is not given but is developed through searching processes. In bounded rationality search models, an organization responds to success or failure through varying the intensity of search, the level of organizational slack and the aspiration level for performance (Cyert and March, 1963). Success lowers search and stimulates slack and targets, whereas failure triggers search and lowers slack and targets in order to restore the aspiration/performance equilibrium (Levinthal and March, 1993). Individuals then stop their search process when they meet their aspirations rather than keeping on searching to achieve a global optimum (Posen et al., 2018; Cyert & March, 1963; March & Simon, 1958; Simon, 1955). This performance assessment is realized in relation to an aspiration level, which in turn is influenced by past performance (Cyert and March, 1963; Lant, 1992). A key point is represented by understanding how a decision maker establishes expectations about what outcome can be classified as satisfactory. In the absence of previous knowledge or social comparison, an agent forms its aspirations based on the feedback received on its own actions (Lant, 1992; Billinger et al., 2021). Search is sparkled when a firm recognizes performance to be below its aspiration levels and it ends when a satisfactory solution is found, bringing back performance to the aspired values. Organizations initially concentrate their efforts in proximity of current practices and possibilities. Only when this process has proven unfruitful, they start looking for solutions in more distant domains (Cyert and March, 1963; Posen et al. 2018).

According to the behavioral theory of the firm organizations determine and adapt their aspirations in accordance with a set of reference points (Cyert and March, 1963). According to March's model (1988), a decision maker moves among two reference points – a lower point that ensures survival and a success point which depends on aspirational levels. The following steps in the search landscape are represented by efforts to close the gap between aspirations and performance. Organizational changes are, in fact, evaluated on their ability to restore performance levels (Simon 1955; Greve, 2003; Posen et al., 2018). The most robust description of aspiration formation is based on an elementary decision rule of adjustment to performance feedback (Lant, 1992). Feedback has a central role in Billinger et al. (2021) model, since immediate and historical assessment on performance has an influence both on the decision to stop searching, once an agent is pleased with his/her performance, and if not, on the decision of search breadth, through recombining attributes to test alternatives, enlarging the search domain (Billinger et al., 2021).

1.3. Individual versus Organizational Ambidexterity

According to Simon (1959) at the individual level, within organizations when performance falls below aspirations, this triggers search for alternative courses of action. Within an organization, due to standardized procedures and behavioral expectations, individuals have less autonomy on how to allocate their activities between exploration and exploitation (Bidmon & Boe-Lillegraven, 2020). Nonetheless, a directive approach could be useful to determine when an individual should explore or exploit within teams and business units. Individuals, in fact, have naturally different inclinations towards ambidexterity, with someone requiring more support to balance search efforts (Bidmon & Boe-Lillegraven, 2020; Tempelaar and Rosenkranz, 2019; Laureiro-Martìnez et al., 2015). Individuals manifest a strong tendency towards adaptive search, meaning that failures activate exploration whereas successes trigger exploitation. Moreover successes curb search for new alternatives in the neighborhood of existing ones, while failure prompts more distant and exploratory search. Individuals, in order to respond to feedback, tend to interrupt neighborhood search too early, overlooking the possibility to achieve local improvements (Billinger et al., 2014). Additionally, within organizations individual behavior is influenced by the use of incentives. When performance-based incentives are reduced, individuals, especially high-performing ones, engage in more exploratory activities. Furthermore, lowering performance-based incentives leads to a higher exploration performance obtained through experiential learning (Lee and Meyer-Doyle, 2017).

The ability to balance exploration and exploitation within organizations partially depends on the afore mentioned attention shifts in adaptive aspirations. Managers can deviate reference points to purposely orient attention throughout companies to adjust or redirect strategy (Blettner et al., 2015). Especially in more complex environments, with extensive interdependencies among organizational elements, organizations will need to lean more on organizational features promoting a more extensive search (Rivkin and Siggelkow, 2003). Within organizations, performance feedback can be used to boost employees' efforts and to activate search for improvements in work tasks. Additionally, establishing goals at a central level allows to align goals with firm strategy and for aspiration levels to be concrete and high enough to trigger efforts to augment performance (Greve, 2010).

HYPOTHESES

In order to investigate the research questions, a series of hypotheses has been developed, addressing the dimensions of feedback, search breadth and task complexity. The hypotheses will account for possible differences in these mechanisms with respect to the autonomy of a decision maker, establishing his/her own reference points, or control imposed to an agent, represented by the settling of a target serving as the conditions of a survival point and of an aspirational level.

Feedback

Autonomy Setting – It is necessary to start from feedback since search behavior critically depends on it. Search mechanisms at the individual level, in fact, progressively readjust to performance feedback (Billinger et al., 2014). Performance assessment, in relation to an aspiration level, in the absence of previous knowledge is, in fact, based on the feedback received on an agent's own actions (Billinger et al., 2021; Lant, 1992; Cyert and March, 1963).

In line with previous literature, the main idea behind this research project is that individuals perceive performance feedback as a success or failure on the basis of a reference point (Billinger et al., 2014; Bromiley, 1991; March, 1988; Markowitz, 1952). In an autonomous setting, without a previous benchmark to hang on or directions regarding targets provided, feedback received in the early stages of the search process has a strong influence in setting expectations in the absence of prior assumptions on possible performances (Billinger et al., 2021). The aspiration formation process is, then, based on a rule of adjustment to performance feedback. Aspiration will adjust upwards in response to positive feedback, whereas these will settle downwards in response to negative feedback (Lant, 1992). Nonetheless, it needs to be taken into account that responsiveness to feedback may depend on whether performance is evaluated with reference to own previous performance or to peer performance. Performance feedback at the individual level is, therefore, subject to multiple interpretations (Joseph and Gaba, 2015). Subsequently, aspiration levels act as a guide to encode performance. Through feedback, aspirations respond to past performance and consequently adjust behavior, which becomes less sensitive to performance outcomes. Direct experience is, in fact, the main driver of aspirations change that is realized as a consequence of successful outcomes (March, 1988). In particular, agents receiving a positive feedback at the end of the first trial may gain confidence and become less likely to stop search early (Billinger et al., 2021). On the contrary, a general conclusion in the literature prescribes that individuals will stop their search process once their aspirations are met, rather than keep on searching to achieve a global optimum (Posen et al., 2018; Cyert & March, 1963; March & Simon, 1958; Simon, 1955). In line with Billinger et al. (2014), as a reference point it is possible to consider the highest-performing combination, the one that obtained the highest payoff, found by an individual in prior trials. Individuals' reference points are reciprocally influenced through dynamic tradeoffs based on how easily individual aspiration levels can be reached and on how much benefit they bring. The search landscape, in fact, depends on agents' attributes, determining their aspiration and survival levels and to the state of the population, since through feedback individuals compare their results in relation to the performance of all the other agents (Marengo et al., 2022). Connected to the work of Billinger et al. (2021), the aspiration level is linked to the decision on whether to search. In fact, once a satisfying choice is identified in relation to this point, an individual will cease looking for performance improvements. Decision-makers, in fact, rather than pursuing a global optimum, are mainly responsive to whether they encounter a reference point during their search activity (Marengo et al., 2022). The most difficult challenge for an individual is, then, to stabilize around the good choices, while at the same time keeping on searching for further improvements (Rivkin and Siggelkow, 2003; Baumann et al., 2019). Individuals search behavior, in fact, is not uniform but rather mixed, alternating elements of local and more distant search (Billinger et al., 2014). A positive performance feedback decreases search breadth and focuses search in proximity of the areas in which an agent has experienced a performance increase (Billinger et al., 2014). Positive local feedback generates a strong path dependence and it may tie agents to suboptimal equilibria (March, 1991). Additionally, in a sequential search process what is learned at a particular point in time affects what can be learned at a later point in the research. (Rivkin and Siggelkow, 2003; Baumann et al., 2019). From previous studies, in fact, it emerges that performance substantially increases in the number of search trials (Billinger et al., 2014).

Performance feedback received at the end of the first trials shapes agents' aspirations. Agents receiving a positive feedback will, then, gain confidence and indulge in subsequent exploration, enlarging their search space. Once agents' aspirations are fulfilled, they will cease looking for improvements and will stick to the combinations found. It is then, possible to hypothesize that:

H1a: "Positive feedback, relative to an agent's aspirations, will lead to an increase of search breadth in the initial trials."
H1b: "Positive feedback aligned to an agent's aspirations will result in a reduction of search breadth."

Nonetheless, decision-makers evaluate their performance on the basis of two reference points: an aspiration level and a survival point (Marengo et al., 2022). As explained, the degree of search depends on the individual perception of how well an agent is doing. On this basis individuals categorize results into success and failure. As a consequence, decision makers move among two reference points, a lower point ensuring survival and a success point that depends on the previously analyzed aspiration levels (March, 1988). The following steps in the search landscape are represented by efforts to close the gap between aspirations and performance. The survival point determines the evaluation of discovered alternatives. If an alternative is found above this reference, this will be adopted despite potentially losing fitness in the current period. As a consequence, individuals are not affected by additional improvements in-between these subjective references. Agents, in fact, do not keep on searching for continuous improvements but rather rely on reference points as heuristics of search (Marengo et al., 2022). When performance feedback is below a level that is considered as acceptable, this will trigger explorative search in order to reach a level comprised between the aspirational and survival levels. Therefore it follows that:

H1c: "Negative feedback in relation to an agent's survival point will result in an increase of search breadth."

Control Setting – On the other hand, in line with the findings by Marengo et al. (2022), in the control setting reference points are set by firms, rather than being adaptively defined by agents, in order to match with their environment. A firm, actually, tries to gravitate around what Rivkin and Siggelkow (2003) define as a "sticking point" – "a configuration of choices from which it will not change" (p.292). In a controlled setting, managers can establish reference points to effectively direct attention throughout a company to adjust or redirect strategy (Blettner et al., 2015). Additionally, establishing goals through a central authority allows to align goals with firm strategy and for aspiration levels to be concrete and high enough to trigger efforts to increase performance (Greve, 2010). Within organizations, reference points filter both whether and how agents search. Establishing a target equals for individuals to search on a subjective landscape that is reduced and much smoother than the underlying performance landscape. This reduced landscape is made of peaks, satisfying aspirations, connected by ridges to survival points and separated by performance holes (Marengo et al., 2022). As explained by March (1988), establishing a target has a strong influence in directing attention, defining the subjective reference points for success in search behavior. The subsequent level of search depends on the individual perception of how well the agent is doing. Additionally, through this perception a decision maker classifies results into the categories of success and failure (March, 1988).

Consequently, individuals will search on a ridge between a preferable and undesirable performance, combined with performance holes - combinations leading to performance below the survival point - and peaks - combinations above aspirations (Marengo et al., 2022). In line with the findings of Simon (1959) individuals will look for alternative courses of action when performance falls below aspirations within organizations. As already mentioned, problemistic search is affected by a heuristic rule of searching in the neighborhood of current activities. Establishing a target to be reached could made the search process more successful through the identification of where the problem resides within the organization (Greve, 2010).

In a controlled setting, the aspirational level and survival point are not individually defined by agents but externally provided. Therefore feedback will not shape individuals' aspirations but it will reflect if an agent performance is in line with the established target. Therefore it follows that: l reflect if an agent performance is in line with the established target. Therefore it follows that:

H2a: "Positive feedback in relation to the established target will result in a reduction of the search breadth."
H2b: "Negative feedback in relation to the established target will result in an increase of search breadth."

Complexity

Autonomy Setting – According to Billinger et al. (2014) complexity of the search landscape does not directly influence search behavior, but rather indirectly through performance feedback. Individuals behavior adapts to task complexity, since as task difficulty increases so it does search breadth (Billinger et al., 2014). Local search can be seen as a sequence of trials that involve changing only one attribute at a time and learning from the resulting performance feedback. In rugged landscapes, as long as experimentation is local and fails to consider interdependencies, it will only lead to a low local peak (Baumann et al., 2019). Performance feedback from previous trials has a crucial role in determining where in the search space individuals will look for higher-performing configurations (Billinger et al., 2014). As represented by the NK model, the fact that superior solutions are often far from the starting point, since they require a change in more than one choice, poses a challenge to the sequential search process. Boundedly rational individuals, in fact, cannot identify on their own higher-performing solutions that are deeply different from the solutions they know. On the contrary, individuals tend to search in the neighborhood of current solutions, by changing one dimension at time (Baumann et al., 2019). In highly complex task environments local search can easily get stuck on local optima and undertaking more distant search gives more chances to improve performance (Billinger et al., 2014). One effective way to search is to simultaneously change several choices, which has been identified as executing "long jumps" (Baumann et al., 2019; Levinthal, 1997). A more trustworthy alternative could be to detect superior solutions by collecting available insights, as learning from the observation of others' solutions (Baumann et al., 2019; Rivkin, 2000). Reference points, in fact, are subjective, depending on the individual performance in previous rounds, and therefore can change over time and across individuals in the same setting (Billinger et al., 2014). High complexity in the landscape and changes in the strategy rewards tend to penalize agents with high aspiration levels (Marengo et al., 2022). Problem representations can change throughout the search process. A shift in representations can be considered as an higher-level experimentation, broadening search but without relying on any superior insight (Baumann et al., 2019). Agents' search behavior responds to task complexity, even though not in a straightforward manner. Complexity, in fact, induces agents to mix local and distant search, not to opt exclusively for one of the two. Performance increases with the number of trials, whereas task complexity negatively affects the recognition of improvements, negatively influencing performance in turn (Billinger et al., 2014). If the new combination found delivers a superior performance it will be implemented, otherwise it will be discarded (Baumann et al., 2019). Complexity of the search landscape will be reflected in the performance feedback received that will impact on the aspirations of decision-makers. Agents will, then, need to engage in a more explorative research on the performance landscape, as interdependencies among attributes increase, in order to reach their aspirational level and satisfice. It follows that:

H3a: "As complexity - represented by the interdependencies among attributes – increases, agents will engage in a more explorative search."

At the same time, a negative performance feedback may result in a downwards update of individuals' aspirations and it may lead agents to satisfice on a lower point and subsequently stop search earlier. So, it is possible to additionally hypothesize that:

H3b: "As complexity – represented by the interdependencies among attributes – increases, agents will satisfice on a lower payoff and reduce their search breadth."

Control Setting – In a controlled environment, cognitive representations of the problem space can, indeed, improve the effectiveness of search by providing intuitions into potentially superior solutions and by suggesting an understanding of the structural characteristics of the problem (Baumann et al., 2019). Representations can be defined as coarse, since these are approximations of the real problem structure (Baumann et al., 2019; Gavetti and Levinthal, 2000). Search breadth is not, in fact, influenced exclusively by feedback but it also depends on the features of the environment that an organization faces, since these influence their reference points. Establishing aspiration levels influences the extent of feasible options and opportunities for development available for an organization. Reference points, indeed, influence individuals' perceptions and evaluation of performance and guide the search process. As recalled, agents, in a controlled setting, do not face the entire performance landscape but what they see are peaks – points above the aspiration level –, valleys – combinations below the survival point – and ridges, connecting aspirations and survival levels (Marengo et al., 2022). Organizational design determines the number of these points and their associated payoffs. Additionally, it also affects the probability that a firm will actually attain such equilibrium. Organizations, especially those facing complex environments, need to achieve a balance between elements that support search and elements supporting stability. In the presence of extensive interdependencies among organizational attributes, organizations will need to rely more on features supporting a more extensive search (Rivkin and Siggelkow, 2003).

The risks connected with a disequilibrium between aspiration levels and survival points is not symmetrical. In fact, decision makers with high aspiration points and low survival point may search in a too wide area of the landscape and consequently not reach their desired performance. On the contrary, agents with high survival points but low aspirations experience only a small portion of the performance landscape and may settle on mediocre solutions, threatening their long term survival (Marengo et al., 2022). The control imposed by organizational structures impacts a firm's performance by directing agents' search behavior on the landscape they confront. Design features affect the degree to which a firm searches in its environment to find successful combinations of coordinated choices and if the organization is able to stabilize around those combinations once identified (Rivkin and Siggelkow, 2003).

Representations are effective since they allow for a cognitive or "offline" evaluation of possible solutions, so that superior combinations can be found without testing them with experimentation (Baumann et al., 2019; Gavetti and Levinthal, 2000). Coarse representations restrict the search space to choices that have higher expected performance, since these establish a higher starting point for subsequent experiential search efforts (Baumann et al., 2019; Gavetti and Levinthal, 2000). Additionally, coarse insights are particularly useful when a complex problem cannot be divided in smaller modules, so when there are high interdependencies within the problem attributes (Baumann et al., 2019; Gavetti et al., 2005). Moreover, representations of the underlying problem structure facilitate problem decomposition (Baumann et al., 2019).

Therefore, when an organization is facing a complex problem, the presence of a central coordinator or "strategist" providing insights into the problem structure, allows for a more effective search (Baumann et al., 2019).

As interdependencies among organizational attributes increase, the subsequent complexity will be reflected in performance feedback. A positive feedback, resulting in a payoff belonging to the target, will lead agents to reduce search efforts. On the contrary, a negative feedback in the presence of an externally imposed target, rather than updating downwards agents aspirations, might have the effect of stimulating search in order to reach the target itself. It follows that:

H4a: "As complexity – represented by the interdependencies among attributes – increases, agents will reduce search breadth in response to a positive performance feedback."
H4b: "As complexity – represented by the interdependencies among attributes – increases, agents will increase search breadth in response to a negative performance feedback."

Introduction of a penalty

Autonomy Setting – In both the autonomy and control setting, this research work will try to account for the possible influence that the introduction of a penalty exerts on the decision to stop search.

Individuals are inclined toward over-exploration, as evidenced by Billinger et al. (2014) experiment, since they tend to cease neighborhood search too early. Moreover, in the initial rounds agents tend to engage in distant search, meaning that they will change multiple attributes altogether (Billinger et al., 2014). Local search allows immediate and incremental gains in proximity of existing alternatives, bearing the risk of localizing on a local optimum. On the contrary, distant search is riskier and gives agents the chance to discover better alternatives in the search landscape. In simple tasks the better option for agents would be to engage more in local search. Nonetheless, human decision makers interrupt local search in favor of more distant search too early in simple tasks, sacrificing potential gains from local improvements (Billinger et al., 2014). In a setting in which additional search has a cost, in order to reduce regret agents will tend to stop as soon as they meet reasonably high valued combinations (Billinger et al., 2021; Baillon et al., 2020; Goldstein et al., 2020; Hey et al., 2017).

As previously mentioned, feedback is classified as positive or negative on the basis of subjective reference points (Marengo et al., 2022; Billinger et al., 2014; March, 1988). When taking decisions under risk the security level – the maximum of the minimal outcomes for a possible choice – represents one of the most common reference points (Baillon et al., 2020).

As also evidenced by Labianca et al. (2009), in competitive comparison, when an agent confronts its performance with its competitors and its relative outcome is not satisficing, he/she will engage in exploration and radical changes. At the same time, highly performative agents will engage in explorative and riskier changes in order to reach the combinations they strive to (Labianca et al. 2009). It is, then, necessary to consider that aspirations change on the basis of individual experiences, with an impact on risk taking attitude and the subsequent decisions on search breadth. Success in relation to aspiration levels induces a preference for smaller risks whereas failure induces agents to take greater risks (March, 1988).

The tendency of agents to excessively rely on exploration in response to aspirational levels and negative feedback is well documented in the literature. The introduction of a penalty should reduce this tendency, however this effect is unlikely to completely vanish. Therefore it is possible to hypothesize that:

H5a: "The introduction of a penalty moderates the relationship between aspirations and search breadth."
H5b: "The introduction of a penalty moderates the relationship between performance feedback and search breadth."

Control Setting – Regarding the introduction of a penalty in a controlled setting, it is relevant to recall that within an organization managers focus on organizational goals sequentially and different aspiration levels are attached to each goal. Firms, therefore, will tend to expand more when they find themselves below aspiration levels (Greve, 2008). Therefore, even if riskier, when an agent achieves a performance below the established target, he/she will engage in a more exploratory search, despite the increased risk in undertaking it. This is in line with previous findings by Greve (2010), according to which decision makers are willing to bear more risks when their performance results below their aspiration level.

Within a controlled setting, in which an organization sets a target to be reached, risk taking attitude can be influenced through the use of a penalty up to a level necessary to implement strategic changes in order to achieve the performance target and at the same time to inhibit search from reaching hazardous levels (Greve, 2010). As highlighted by the work of Lee and Meyer-Doyle (2017), incentives can be used within organizations to influence individual behavior. In particular, when performance-based incentives are reduced, individuals will engage in more exploratory search. On this regard, March (1991) proposed that incentives could represent a factor shaping individuals' decision to explore or exploit, but arranging them to promote individual's exploration is particularly difficult due to the uncertainty and remote gains associated with these activities. In line with this argument, actually lowering performance-based incentives has a beneficial effect on exploration performance obtained through experiential learning. This effect is particularly powerful for agents operating in complex task environments, since these settings require higher levels of exploration to attain a satisficing performance (Lee and Meyer-Doyle, 2017). As recalled, according to Billinger et al. (2014) individuals tend to break off neighborhood search too early, wasting the possibilities offered by local improvements. The introduction of a penalty, within a controlled setting, could be used to discourage the excessive relying on exploration, in response to a negative feedback. Therefore it is possible to hypothesize that:

H6: "The introduction of a penalty in a controlled setting, moderates the relationship between performance feedback and search breadth."

METHOD

The Model

In order to address the research question and to test the hypotheses reported above an empirical experiment has been conducted. In line with previous literature (Gavetti and Levinthal, 2000; Gavetti,2005; Billinger et al., 2014, 2021; Marengo et al., 2022) an implementation of the NK model (Kauffman, 1993) has been used.

As developed by Levinthal (1997) in its application in economics, the model defines a fitness landscape through two parameters N and K. A *fitness landscape* is a multidimensional space in which each attribute of a system is represented by a dimension of the space and a final dimension that implies the fitness level of the system (Levinthal, 1997).

In a NK Model an organization is defined by N attributes and each attribute can assume two possible values. Therefore, the fitness space is constituted by the 2N possible combinations of attributes. The values of the N decision variables are determined as random draws from a uniform distribution and the overall payoff of a combination is given by the average of the values assigned to each of the N variables (Levinthal,1997).

The variable K determines the degree to which the fitness of the organization depends on the interrelatedness between the attributes, and therefore the complexity of the task. In fact, the contribution of a single attribute to the overall fitness depends on the other K attributes. If K=0, the contribution of each attribute is independent from all the other elements, whereas when K assumes the highest value of N-1, then the contribution of each attribute to the fitness of the organization depends on all the other attributes. The value of the parameter K determines the smoothness or ruggedness of the landscape. If K=0 each attribute contributes independently to the overall fitness and the resulting landscape is smooth, since a change in one attribute does not affect the fitness contribution of the other N-1 attributes. When K rises up to the maximum value of N-1, the landscape becomes more rugged and in this case a change in one attribute affects the value of the K other attributes. In particular, when K =N-1 a change in just one attribute affects the fitness contribution of all the other attributes. Moreover the value of K affects the number of peaks in the fitness landscape. If K is equal to 0 the fitness space will be a single-peaked one. This means that, since attributes are independent of each other, the environmental fitness can always be improved by shifting a single attribute. On the contrary when K is higher than 0 the fitness space will be characterized by multiple peaks. Given the high interdependence between the attributes, a change in a single attribute may actually lower the overall fitness but a change in multiple attributes may result in an increase of fitness (Levinthal, 1997).

The process of adaptation allows organizations to modify their structure in order to increase their fitness. The initial configuration of an organization will have an enduring effect on its future structure when the environment has multiple peaks, that is extremely complex, since the specific peak that an organization can reach is, for the majority, decided by the starting place in the space of alternative organizational forms. These effects endure as a result of the path dependence of the search process (Levinthal, 1997).

This landscape increases the risk that boundedly rational individuals may be dragged towards low-performing peaks. Managers need therefore to construct a search process that facilitates reaching higher peaks while escaping lower ones (Baumann, 2019; Baumann, 2015; Siggelkow, 2002). As explained by Baumann (2019) "a peak represents a choice combination in which performance cannot be improved by changing only one choice" (p. 289). The greater the degree of interrelatedness among decisions, the more rugged will be the landscape faced by an organization. This ruggedness will serve as a stabilizing factor and, in order to counterbalance this stasis, organizations will need to design organizational features in favor of search (Rivkin and Siggelkow, 2003).

As discussed by Friedman (1953), adaptive search heuristics, like aspiration levels and survival points, can lead to global optima when decisions to undertake are relatively simple, but these are not enough when decisions involve highly interdependent elements (Marengo et al., 2022). In a "smooth" landscape, as interdependencies are absent, whatever the "starting point" is, an "hill-climbing" local search will lead to the optimum solution, the only peak in the landscape, regardless of the order in which dimensions are changed. On the contrary in a "rugged "landscape, local search from the initial starting point will lead to a suboptimal equilibrium, a "local" peak. In order to reach the best solution in the landscape, so the "global" peak, the appropriate changes will need to be undertaken in both the dimensions of the landscape. However, when dimensions can be changed only sequentially, it could be necessary to move downhill in order to reach higher peaks and the process becomes exponentially more difficult. In general, the effectiveness of a search process can be improved by systematically enlarging search. A search process that explores a larger part of the space allows to recognize superior solutions that could have not been reached through a local "hill climbing" process (Baumann et al., 2019).

Searching for good combinations is complicated in rugged performance landscapes as managers do not know in advance which ones are better. It is, therefore, impossible to immediately direct efforts towards a high performing peak (Baumann et al., 2019). On the contrary, better combinations in the performance landscape must be found, as explained by the literature, through a sequential search process. Managers start with an initial combination of attributes and look for better solutions by modifying their current selection over time (Baumann et al., 2019; Nelson and Winter, 1982; Simon, 1955).

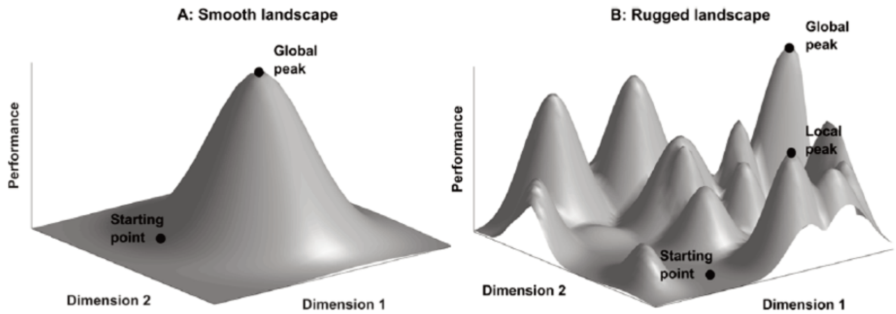

Figure 1. Representation of a Fitness Landscape
Source: Baumann et al. (2019)

Experimental Setting

Following the works of Billinger et al. (2014, 2021) a field experiment has been implemented. Adopting an experiment to study search processes in a complex combinatorial task is particularly appropriate. Experimental settings allow to control and to modify factors like task complexity or information available to decision makers (Billinger et al., 2014; Sterman, 1987). Additionally, empirical experiments allow to gain valuable insights into human behavior in a controlled environment (Billinger et al., 2014; Sterman; 1987). Decision rules in simulation models of human behavior attempt to describe decision making aptitudes as they are rather than how they should be. Direct experiments can be used to corroborate or contradict decision rules in simulated settings. In empirical experiments, the use of an interactive game allows individuals to have a role in the framework being modeled. Participants play the game in a controlled environment adapted to the model being tested, and are given the same information set, but they can take decisions as they want. Human behavior can, therefore, be immediately confronted with the expected decision-making behavior from the model (Sterman, 1987).

Only a small number of experimental studies considers how individuals or groups search throughout a complex problem system. Experiments are particularly useful since a stylized theoretical model can be translated into an empirical setting to examine the degree to which individuals or groups act as expected by the model. Similarly to experimental game-theory studies, observed search behavior in a complex problem landscape may differ from the results of simulation studies (Baumann et al., 2019; Camerer, 2003).

Based on this theoretical insights and in line with the works of previous scholars, a model has been designed to develop the performance landscape and examine diverse aspects of effective search processes (Baumann et al., 2019). Through the use of the NK model, it is possible to link a firm's choices to its payoffs and build a landscape in the space of decisions. Organizations can be modeled as trying to reach and maintain a peak on this landscape, given by combinations of interrelated elements that together grant a high payoff (Rivkin and Siggelkow, 2003). The performance landscape created through the use of the NK algorithm defines the task environments that agents will face through the empirical experiment (Billinger et al., 2014).

The experiment was structured on two different setups an Autonomy Setting and a Control Setting, in order to address the research question on whether the autonomy or control in executing a certain task had an influence on the search behavior of an agent. According to Gavetti (2005), in an Autonomy regime decision-makers can independently define a representation of the strategic landscape they face, realize a strategy based on it and implement the strategy through local search. On the contrary, in a Controlled setting corporate executives determine how to frame the search landscape, define strategies based on their perceptions and demand on their subordinates to implement it (Gavetti, 2005).

These conditions where reflected in the questionnaires provided to the experiment participants. In the autonomy setting participants were asked to develop a business model that would allow them to reach a leading position in the market – to reach the global peak. Agents were informed about what was the global peak in the landscape. Consequently they were able to autonomously settle their aspirations and decide how to search for solutions through the performance landscape.

Conversely, within the control setting participants worked in a company for which they needed to update the current business model in order to reach an established target. Decision-makers knew what was the global peak, but they were also explicitly provided a strategy to follow when searching, defined by the target imposed by the fictional company CEO.

In the experiment, the business model is described by six factors which can assume two possible dimensions. The N factors on which the experimental setting is built have been taken from Morris et al. (2005) six-component framework for characterizing a business model. The development of a business model, in fact, requires coordination among functionally specialized units and the NK model represents a valid structure to represent complexity coming from interdependency patterns among alternatives (Baumann et al., 2019; Andries et al., 2013). As defined by Morris et al. (2005) "A business model is a concise representation of how an interrelated set of decision variables [...] are addressed to create sustainable competitive advantage in defined markets" (p.727). The development of new strategies, technologies, products or business models requires to address complex problems, involving a large number of highly interdependent choices. Managers are, indeed, boundedly rational individuals that need to find a high-performing combination of increasingly interdependent choices. This equals to find a "peak" in a rugged performance landscape that managers can explore only through sequential search (Baumann et al., 2019). Additionally, it is suitable to study a business model since through experimentation with a specific configuration and the respective feedback from the environment, decision makers can actively learn from the environment. If feedback received is negative, the initial business model is reshaped and a new configuration is implemented. Enterprises will therefore change their initial configurations as they learn about and incorporate information throughout the experimental process (Andries et al., 2013; Gruber et al., 2008; Minniti and Bygrave, 2008).

In the experiment, the business model is described by six factors which can assume two possible dimensions. For each factor, participants were asked to choose among two options, accounting for a total of 6 binary choices. Each dimension could assume two possible values 0 or 1. Participants did not know the payoff of the single options. Therefore, the entire search landscape is made of 26 = 64 possible alternatives. The number of search trials, for each scenario, is limited to 6. Participants were provided with the same initial combination. In each round, in response to the feedback received from the previous one, agents could decide on whether to change none, some or all the attributes from their previous combination.

Additionally, in line with the description of the NK model, complexity was introduced through the use of the parameter K. In line with previous literature (Marengo, 2022; Billinger et al., 2021, 2014; Csaszar and Levinthal, 2016; Gavetti, 2005; Gavetti & Levinthal, 2000), three different levels of complexity are considered. In the first two trials, the level of complexity was at 0 (K=0), meaning that there were no interactions among the different attributes. In the third and fourth round a more complex landscape was developed with some degree of interrelatedness among attributes (K=2). In the last two rounds, a maximally rugged and highly complex landscape was built (K=5).

Finally, this research work will try to account for a condition mostly unexplored in the literature. As explained by Billinger et al. (2021), the previous work of Billinger et al. (2014) enacts a problem of "pure search" in which search is not associated to a downside risk. This assumption limits the extension of their results to many real-life settings, in which the exploration of different alternatives is associated to a high-level of risk, such as developing new products or viable business models. The work of Billinger et al. (2021) to account for the dimension of risk-taking proposes to adopt an opportunity cost of changing the current combination, since the final reward for participants depends on the sum of payoffs accumulated through the different rounds. In the setting developed, since the objective for participants is to reach the higher possible or established payoff in the current round, a different penalty was introduced. The penalty consisted in a payoff reduction of the 10% for each attribute in which participants accidentally changed the alternative with the higher payoff (valued at 1) with the lower performing one (valued at 0) (-0,1 as the payoff associated with said attribute).

Implementation

The experimental setting described above has been developed to test this research work's hypotheses. In order to implement it, a pilot experiment has been undertaken. The experiment involved 20 participants and three separate sessions were arranged (two sessions with 7 participants and one with 6). Each session lasted approximately two hours and took place online, through the platform of Google Meets. At the beginning of the call, a Word copy of the Instructions, Questions and Final Questionnaire was sent via mail to each participant. The copies sent had all the same initial combination with a payoff equal to 0,5. In two sessions the first file was the one based on the Autonomy Setting (see Appendix 1), whereas in one session I sent first the Control Setting one (see Appendix 2). Participants read the instructions and for each round answered to questions 1-6 and, as an exemplification of the aspirational levels, they had to write their expected payoff at the end of the round. They had to choose among the same alternatives for 6 rounds.

Regarding the issue of complexity, the first two rounds were set in a smooth landscape. Starting from round 3, complexity was introduced. Between rounds 3 and 4, agents faced the moderately complex environment (K=2) and through rounds 5 and 6 agents faced the maximally complex rugged landscape (K=5). Agents did not know that they were going to face increasingly complex landscapes. This condition was reflected in the feedback they received for their performance.

At the end of each round, participants communicated privately their combination and expected payoff. Once all the responses were collected, the average payoff was publicly announced whereas individual feedback was communicated separately through chat messages, so that participants could make their own evaluations on how to proceed in the tasks. At the end of the last round in each scenario, participants also answered to the final questionnaire.

Participants were then asked to reiterate the whole procedure in the alternative scenario. The total observations collected in both scenarios amounts to 40.

ANALYSIS AND RESULTS

Empirical Analysis

Dependent Variable – The aim of this research work is to understand to what extent exogenous factors influence individuals' search behavior, so their inclination towards exploitation – applying previously successful solutions in order to solve current tasks –, or rather exploration – relying on new mixes of choices to acquire the necessary knowledge to face present contingencies. Therefore, the principal construct that will be analyzed through this analysis is search breadth. Search breadth serves as a proxy to qualify an observed search behavior as exploitative or explorative. It is measured as the number of attributes changed between each round. This variable can assume a value between 0 and 6, as the number of attributes that participants in the experiment were allowed to change in each trial. On average, agents changed 2,10 attributes per trial in the autonomy setting (standard deviation: 0,21), whereas in the control setting the average was 1,76 (standard deviation: 1,34).

Independent Variables – In order to test this research work's hypotheses, a series of variables has been developed. In the autonomy setting to test the relationship between performance feedback and search breadth, it was first of all established which was the reference point against which agents confronted their performance. Answering Question 2) from the final questionnaire, 40% of participants declared to compare their performance to the payoff achieved in the previous rounds, whereas the remaining 60% measured their results with respect to the average performance achieved by all the other participants. Therefore, a measure to codify this tendency has been introduced named Feedback Reference. Additionally, the variable Performance Feedback was construed to encode performance as a success or failure in comparison to the feedback reference for each agent. Finally, it was interesting to compare the payoff achieved at the end of each round with the aspirational level of agents, so to their expected payoff, through the measure of Aspirations Feedback.

On the contrary, in the control setting 85% of participants affirmed to compare their performance to the payoff achieved in previous rounds, making the comparison with the payoff achieved by other participants much less relevant. Therefore, the payoff achieved by an agent at the end of each round was compared to the Target to be reached. In this setting, the variable Performance Feedback classified success when an agent's payoff fell within or above the Target, whereas failure was encoded when an agent's payoff was below the target.

With reference to the relationship between search breadth and complexity, in the autonomy setting the variable Updated Aspirations was introduced to classify whether aspirations adjusted upwards or downwards with respect to the previous round.

Control Variables – Within the experimental setting, it was possible to control for several factors that may have an impact on individuals' search behavior. First of all, as one of the central aspects of this research, it was possible to distinguish between an Autonomy Setting, characterized by the absence of previously determined reference points, and a Control Setting, in which a target to be reached by agents was clearly established. Additionally, the experimental setup allowed to control for the Complexity of the search space. Moreover, it was possible to define the number of Rounds available for each participant. Finally, through the introduction of a Penalty it was possible to introduce an opportunity cost of exploration.

The tables below summarize the variables used to conduct the analysis of experimental results.

Name	Type	Min	Max	Mean	SD	Description
Search Breadth	Count	0	6	2,10	0,21	Number of attributes changed between each round
Feedback Reference	Dummy	0	1	0,40	0,5	Feedback compared with own previous performance is coded 1; with average of other participants 0
Payoff	Scale	-0,1	1	0,56	0,23	Payoff achieved at the end of each round
Performance Feedback	Dummy	0	1	0,54	0,50	Payoff equal or above the feedback reference is coded 1;below 0
Aspirations	Scale	0	1	0,64	0,18	Expected Payoff at the beginning of each round
Aspirations Feedback	Dummy	0	1	0,45	0,50	Payoff equal or above aspirations is coded 1; below 0
Updated Aspirations	Dummy	0	1	0,71	0,45	Aspirations equal or above the previous round are coded 1; below 0
Complexity K = [0;2;5]	Categorical	0	5	-	-	Task complexity
Round	Count	1	6	-	-	Number of trials available to each participant
Penalty	Categorical	-	-	-	-	Cost of exchanging a performative attribute with a non performative one equals -0,1 for each attribute

Table 1.a – Descriptive Statistics for the Autonomy Setting

Name	Type	Min	Max	Mean	SD	Description
Search Breadth	Count	0	6	1,76	1,34	Number of attributes changed between each round
Feedback Reference	Dummy	0	1	0,85	0,36	Feedback compared with own previous performance is coded 1; with average of other participants 0
Payoff	Scale	-0,1	1	0,55	0,20	Payoff achieved at the end of each round
Target	Scale	0,6	0,8	-	-	Target payoff to be achieved by agents
Performance Feedback	Dummy	0	1	0,47	0,50	Payoff within or above target is coded with 1; below 0
Complexity K = [0;2;5]	Categorical	0	5	-	-	Task complexity
Round	Count	1	6	-	-	Number of trials available to each participant
Penalty	Categorical	-	-	-	-	Cost of exchanging a performative attribute with a non performative one equals -0,1 for each attribute

Table 1.b – Descriptive Statistics for the Control Setting

Results

Relationship between Search Breadth and Performance Feedback – Regarding the relationship between performance feedback and search breadth, it is necessary to distinguish findings related to the autonomy setting and those related to the control setting. For both scenarios, in order to test the effect that performance feedback actually had on search breadth, average performance feedback was computed between rounds 1 to 5 (excluding the last round, since the relative performance feedback could not be reflected in the number of attributes changed in subsequent rounds). For the same reason, average search breadth was based on the number of attributes changed between rounds 2 to 6 (excluding the first round, since the number of attributes changed did not depend on performance feedback from the previous round).

Autonomy Setting – With reference to the autonomy setting, the first step, as before explained in the variables section, consisted in identifying on which reference point agents anchored their aspirational levels. Once these were defined, the payoff received in each round was confronted with the payoff obtained in the previous round (with the expected payoff and actual payoff in the first round) for agents focused on their previous performance, or with the average payoff at the end of the round for agents that were interested in their performance with respect to the other participants. For each participant it was, then, computed the average performance feedback between rounds 1-5 and the average search breadth in rounds 2-6. What emerges it is the relationship depicted in Figure 2.a. There exists a negative relation between average performance feedback and average search breadth. This result is in line with an accepted finding in the literature, according to which agents stop their search process once their aspirations are met, rather than keep on searching to achieve a global optimum (Posen et al., 2018; Cyert & March, 1963; March & Simon, 1958; Simon, 1955). As performance approaches individual aspirational levels, agents will tend to satisfice and decrease their search breadth, relying onto exploitation. On the contrary, as performance feedback decreases, agents will concentrate their efforts on exploration in order to meet their aspirational level.

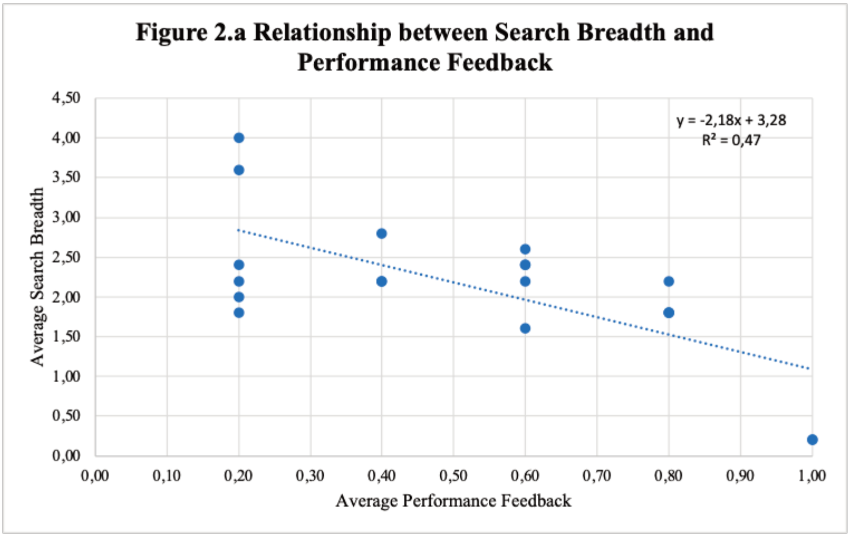

Regression Statistics	
Multiple R	0,687
R Square	0,472
Standard Error	0,658
Significance F	0,001
Observations	20

Table 2.a – Regression Statistics for the Autonomy Setting

As shown by the value of R Square this linear correlation is true for the 47% of observations collected, with a confidence interval of 95%. The multiple R value suggests that this relation is strong. This result is also statistically significant, since the value of Significance F is less than 0.05 and the p-value for the average performance feedback is 0,0008 (<0,05). Nonetheless, the value of the standard error is quite high, possibly due to the small sample size. With this limitation, it is possible to accept H1b and H1c. A positive feedback with respect to an agent's aspirations results in a reduction of search breadth, whereas a negative feedback leads to an enlargement of search breadth.

However, it is also noteworthy to conduct an analysis focused on the relationship between search behavior and aspirations with reference to the agents receiving a positive performance feedback. Part of the literature, in fact, in contrast with the previous findings, suggests that a positive feedback may adjust aspirations upwards. Agents would then become greedy and unlikely to stop search, especially in the initial trials (Billinger, 2021; Lant, 1992). In order to test this assumption, for each agent the payoff obtained from round 1 to 5 was compared with the expected payoff through the variable Aspirations Feedback. Then, for each round, the average number of changes made by all participants (the Average Search Breadth) receiving a positive Aspiration Feedback (Positive Asp. Fb.) and a negative Aspiration Feedback (Negative Asp. Fb.) were computed. The resulting relation is shown in Figure 2.b.

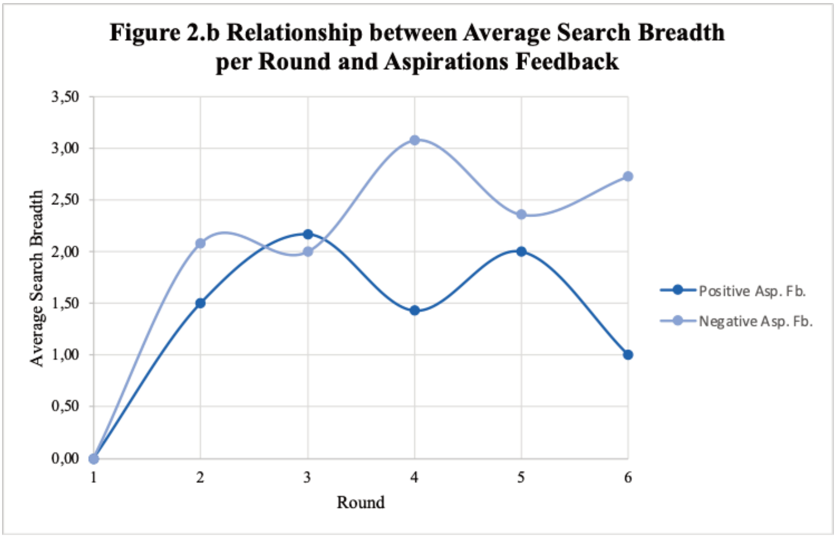

*For the 1st Round, since the number of changes was not influenced by Aspirations Feedback, Average Search Breadth is normalized at 0.

As expected, and in line with the findings for H1b and H1c, agents obtaining a payoff below their aspirations, on average, registered a higher average search breadth than those meeting their aspirations. Focusing on agents receiving a positive aspirations feedback, it is possible to see that that after the first two trials average search breadth increases in line with an increase in expectations, with the average expected payoff passing from 0,59 to 0,7. But, after the third round, as also complexity increases, average search breadth first decreases, then it increases in round 5, to decrease again in the last round, where successful agents decrease their average search breadth and satisfice. If it is true that after a successful performance in the initial rounds agents increase their expectations, and subsequently enlarge their search breadth, it is not clear why after satisficing and receiving a positive aspirations feedback agents increase again their search breadth. It needs to be considered that, after round 4, agents that achieved a payoff equal or above their expectations on average changed 2 attributes with a standard deviation of 1,88, suggesting a quite varied response to aspirations feedback. Based on the observed data, there is not enough evidence to support H1a, which is consequently rejected.

Control Setting – In order to test the relationship between performance feedback and search breadth in the control setting, the payoff achieved by an agent at the end of each round was compared with the established target. For each agent was then computed the average performance feedback for rounds 1 to 5 and the number of attributes changed between rounds 2 to 6. What emerges is that there exist a negative relation between Average Performance Feedback and Average Search Breadth, as shown in Figure 2.c.

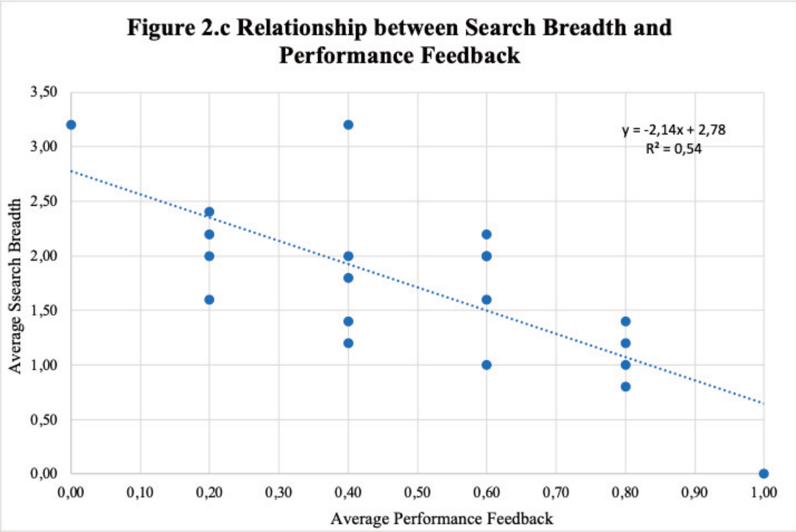

Regression Statistics	
Multiple R	0,732
R Square	0,536
Standard Error	0,538
Significance F	0,0002
Observations	20

Table 2.b – Regression Statistics for the Control Setting

This relation is even stronger than in the autonomy scenario, as demonstrated by the higher value of the R square indicator. In line with previous findings, an established target strongly influences attention, defining the reference points for success in search behavior (March, 1988). Individuals, in organizations, will therefore look for alternative courses of action when performance falls below this reference (Simon, 1959). Through this relationship is, in fact, possible to explain the 54% of observations collected, with a confidence interval of 95%. The value of the multiple R indicator points that the relationship among the two variables is strong. Additionally, this result is statistically significant as the value of the indicator of significance F is less than the critical value (<0,05), as the p-value for the average performance feedback (0,00024). However, also in this setting, the value for the standard error is high, as a possible effect due to the small sample size. With this limitation, it is then possible to accept H2a and H2b. A high average performance feedback in relation to the established target will result in a decrease of search breadth, favoring exploitation, whereas a negative performance feedback will increase search breadth, leading to exploration.

Relationship between Search Breadth and Complexity

Autonomy Setting – In order to test the relationship between complexity and search breadth, it is necessary to start from performance feedback. As evidenced by Billinger et al. (2014), in fact, complexity of the search landscape indirectly influence search behavior through performance feedback. Performance feedback from previous rounds determines where in the search space agents will look for performance improvements. Therefore, it was first necessary to distinguish for each level of complexity agents achieving positive and negative performance feedback. Subsequently, average search breadth for each level of complexity was computed for both clusters. The relationship derived between the level of complexity – represented by the interrelationships among attributes – and search breadth is shown in Figure 3.a.

Figure 3.a Relationship between Complexity and Search Breadth

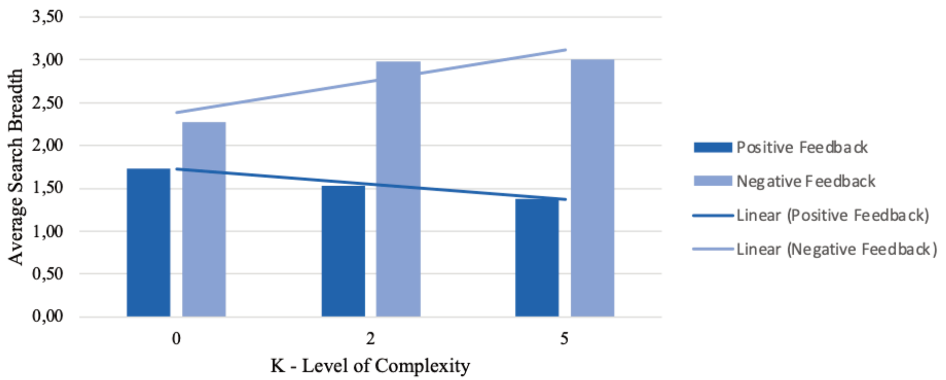

As it is possible to observe, complexity of the landscape, results in a negative performance feedback, but decision-makers will nonetheless still strive to reach a higher aspirational level through a more explorative research on the performance landscape. Performance feedback from previous rounds determines where in the search space agents will look for performance improvements. Individuals tend to concentrate search in the neighborhood of current solutions, but in highly complex task environments enlarging search breadth gives more chance to improve performance (Baumann et al., 2019; Billinger et al., 2014). It is therefore possible to accept H3a, since as the observations suggest as complexity increases, so it does search breadth.

The impact that negative performance feedback exerts on aspirations is measured through the variable Updated Aspirations. For each round, the average number of attributes changed was computed for agents receiving a negative feedback and that at the same time updated downwards their aspirations. Rather than stopping search early, participants, on average, tried to change a greater number of attributes, increasing their search breadth rather than reducing it and satisfice in line with their new expectations, as shown in Figure 3.b. Participants tended to change several attributes altogether, executing what has been defined as “long jumps” (Baumann et al., 2019; Levinthal, 1997). Based on the observations collected H3b, suggesting that an increase in complexity is reflected on the decision to satisfy on a lower payoff and reduce search breadth, is rejected.

Figure 3.b Focus on the Relationship between Search Breadth and Complexity for Agents receiving a Negative Performance Feedback

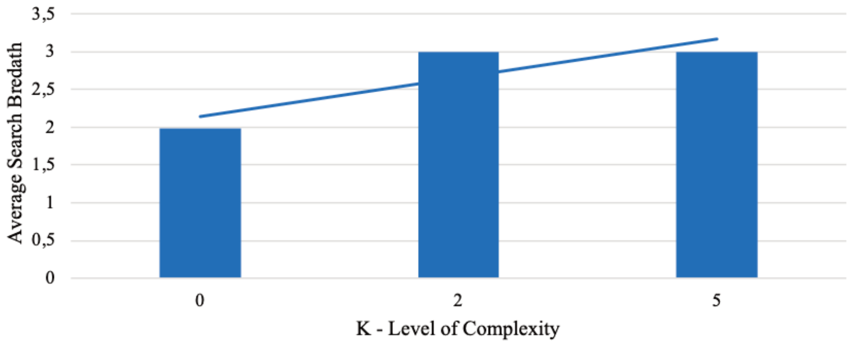

Control Setting – Performance feedback acted as a guide also to test the relationship between complexity and search breadth in the Control scenario. Positive and negative performance feedback, with reference to the established target, were assessed for decision-makers in each round. Then, average search breadth was computed for each level of complexity, distinguishing between agents receiving positive and negative performance feedback. As in the autonomy scenario, as complexity increases and this condition is reflected on performance feedback, a positive feedback ,resulting in a payoff belonging to the target, will lead agents to reduce search breadth. Negative performance feedback in relation to an established target, on the contrary will spur search efforts in order to reach the same target. Organizations, in the presence of extensive interdependencies among their attributes need to rely on features supporting a more extensive search and establishing a target, indeed, influences individuals’ understanding and evaluation of feedback and guides the search process (Marengo et al., 2022; Rivkin and Siggelkow, 2003).

Based on the trend observed through the data collected, it is possible to accept H4a and H4b, according to which as complexity increases, agents will reduce search breadth in response to a positive performance feedback and will enlarge their search space in response to a negative feedback.

It is possible to notice that these tendencies are even more marked in a controlled setting. In line with extant literature, the control imposed by organizational structures has an impact on performance by directing agents’ search behavior on the landscape they confront and making at the same time the result of search more effective through a clearer direction provided by the target (Baumann et al., 2019; Rivkin and Siggelkow, 2003).

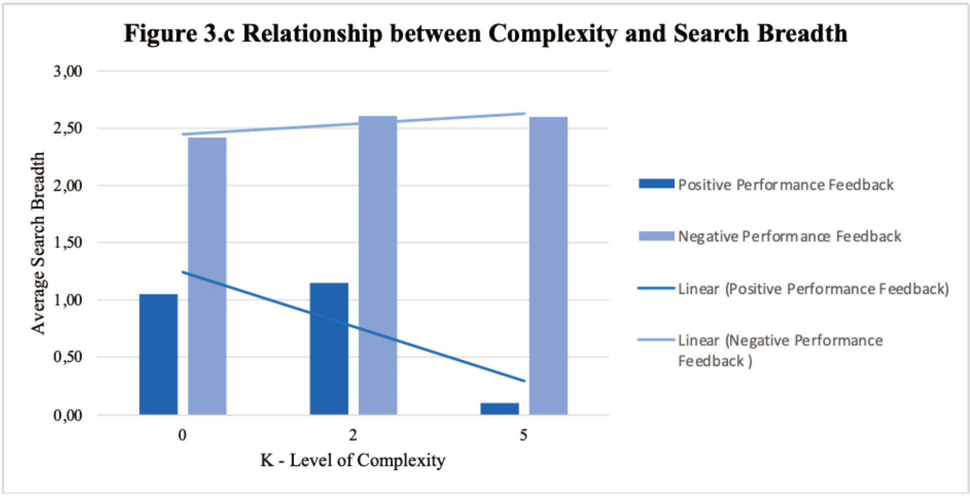

Effect of Introducing a Penalty

In order to account for the effect that a penalty has on the relationship between feedback and search in both the autonomy and control scenarios, it is first necessary to distinguish between agents that were or not affected by the presence of a penalty. From the answers collected in response to Question 5) in the final questionnaire, it emerges that the introduction of a penalty inhibited 65% of participants in the autonomy setting and 70% in the control setting from changing a greater number of attributes in between rounds.

Autonomy Setting – One of the main findings from the work of Billinger et al. (2014) is that human agents are inclined towards over-exploration, interrupting local search too early and sacrificing profits from local progresses. Nonetheless, according to the literature, in a setting in which search has a cost agents will tend to stop their research for better combinations once satisfying combinations are found (Billinger et al., 2021; Baillon et al., 2020; Goldstein et al., 2020; Hey et al., 2017). This research work tried to verify if the introduction of a penalty reduced the tendency of decision makers towards relying on over-exploration with reference to aspirational levels and performance feedback. To study the impact on aspirations, after distinguishing between agents affected or not by the penalty, we observed for the two clusters of agents what was the average level of search breadth for the same level of aspirations feedback. What emerges is the relationship presented in Figure 4.a.

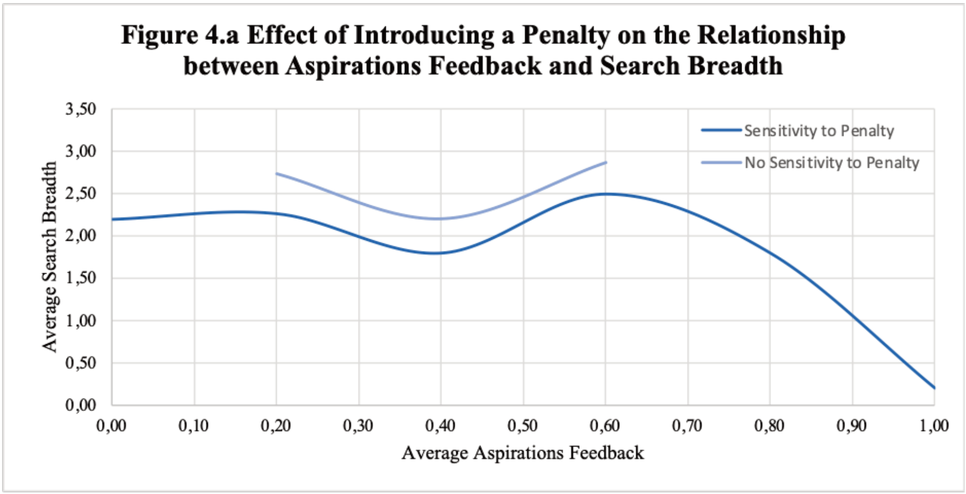

The different length of the lines is due to the fact that a larger portion of participants were affected by the penalty, and therefore their aspirations feedback fluctuates among a larger range. Nonetheless, it appears clearly that for the same level of average aspirations feedback, agents sensitive to the presence of a penalty on average focused their research in the neighborhood of solutions known, whereas those that were not affected by the penalty looked for alternative solutions on a wider area of the search landscape. It is therefore possible to accept H5a.

Regarding the moderating effect of a penalty on the relationship between performance feedback and search breath, the effect is not straightforward, as represented in Figure 4.b.

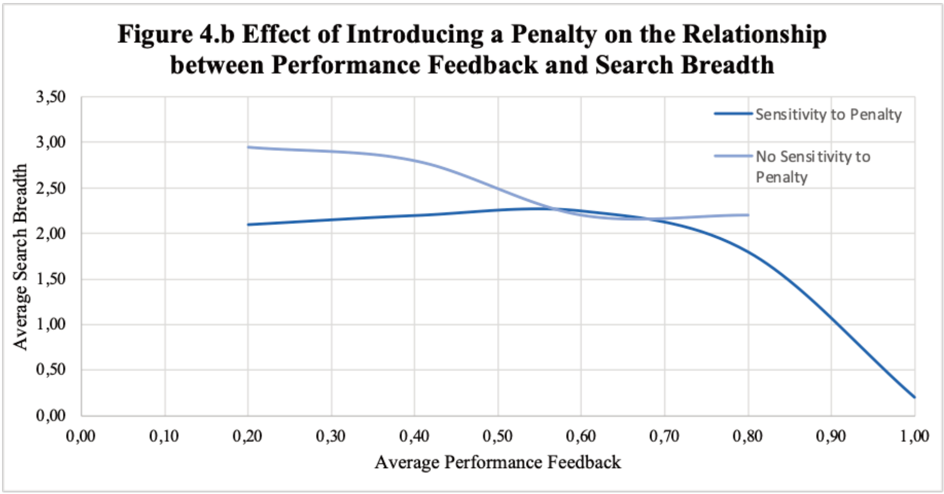

Also in this case, average performance feedback stretches on a larger set of values for agents sensitive to penalty because of their larger proportion on total participants. What is interesting to notice is that agents that received a medium-high average performance feedback, despite being sensitive to the introduction of a penalty, had an average search breadth slightly higher than agents not sensitive to the penalty. This may be due to the fact that receiving on average a positive performance feedback made human agents greedy, overcoming the moderating effect exerted by the introduction of the penalty. Therefore, on the basis of the observations collected, it is not possible to accept H5b, according to which the introduction of a penalty moderates the relationship between performance feedback and search breadth.

Control Setting – In order to test the effect that the introduction of a penalty had on the relationship between performance feedback and search breadth, the average number of attributes changed for a determined level of average performance feedback was observed. Since 14 participants out of 20 admitted that the presence of a penalty had an inhibiting effect on their decision to change the number of attributes in between rounds, the higher number of observations extended average performance feedback on a larger set of values. Nonetheless, considering the interval of average performance feedback between 0,2 and 0,6 (agents achieving on average a payoff above the target from the 20% to 60% of the rounds) it is possible to observe the moderating effect of the introduction of a penalty as shown in Figure 4.c.

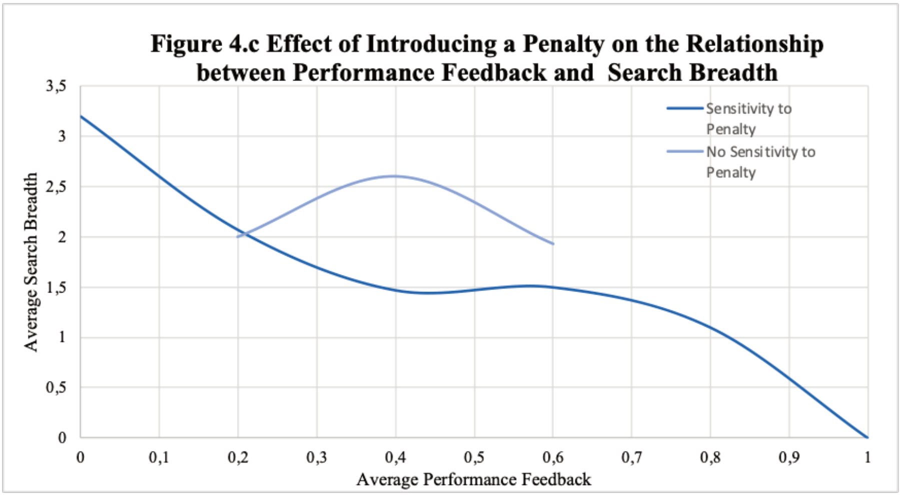

In between this interval, agents that showed no sensitivity to the introduction of a penalty reached a consistently higher value of the average search breadth relative to the average performance feedback than agents sensitive to the introduction of a penalty. As also evidenced by Greve (2010), in a controlled setting, with a target to be reached, the introduction of a penalty can be used to boost exploration up to a level necessary to achieve the performance target and to simultaneously inhibit search from reaching hazardous levels. As shown in the graph in Figure 4.5, it was possible to effectively reach the same level of average performance feedback with a lower level of average search breadth. It is, therefore, possible to accept H6, since the introduction of a penalty in a controlled setting reduces the average search breadth for the same level of average performance feedback.

DISCUSSION & LIMITATIONS

The aim of this research work was to understand how autonomy and control influence human decision-makers' search behavior. In particular, it was observed how performance and aspirations feedback and different levels of complexity impacted on the average search breadth of agents. Subsequently, it was of interest to examine how the introduction of a penalty affected the relationship between feedback and search breadth.

This research work draws insights from two main streams of literature. On one hand, it adds to the stream of literature of the behavioral theory of the firm (Cyert and March, 1963), according to which an organization determines and subsequently adapts its aspirations on the basis of a reference point, and to the connected problemistic search theory that models the behavior of a firm as learning from the feedback received on its previous performance in order to achieve a fit between its capabilities and the environment (Denrell and March, 2001). On the other hand, it also builds on the branch of strategy literature based on the NK model (Levinthal, 1997; Kauffman, 1993), that provides a framework to study agents' search behavior, in terms of the choice between narrow versus distant search, in complex landscapes. These two streams converge in and originate from March (1991) fundamental definitions of exploration and exploitation and the subsequent implications of what these concepts entail and why their difficult balancing generates what has been defined as a tradeoff or dilemma.

In order to present the results of this study, the dependent variable search breadth was introduced to condense the two step decision process described by Billinger et al. (2021). Regarding the decision of whether to search, a value of search breadth equal to 0 implied that the agent decided to not make any changes to the status quo since he/she did not recognize any discrepancy in between his/her aspirations and the performance feedback received. Each value of the search breadth dimensions from 1 onwards identifies some degree of mismatch between aspirations and feedback, which is reflected in the decision to engage in narrow or distant search, defined by the number of attributes changed in-between rounds.

This research contributes to the extant literature in the following ways. First, with regards to the behavioral theory of the firm and problemistic search theory, in both the autonomy and control scenarios, average performance feedback and average search breadth are negatively correlated. Agents that throughout the experiment on average achieved a positive performance feedback - with respect to their own previous performance, to their peers performance or to the established target - registered lower levels of average search breadth. On the contrary, a negative performance feedback is related to a greater level of average search breadth. This result is in line with the findings from other scholars, according to which agents stop their search process once their aspirations are met, rather than keep on searching to achieve a global optimum (Posen et al., 2018; Cyert & March, 1963; March & Simon, 1958; Simon, 1955). Search behavior depends on performance feedback, which assesses individuals' reference points, demarking successes and failures. Individuals, then, manifest a strong tendency towards adaptive search, since success restrains search of new alternatives in proximity of existing ones, therefore supporting exploitation (Billinger et al., 2014).

Second, this research contributes to the literature on the NK model with its findings on the relationship between search breadth and the level of complexity. Following the directions provided by Billinger et al. (2014), it is necessary to consider that the complexity of tasks faced does not impact on search behavior, but rather on feedback received from searching for new alternatives. Performance feedback from previous rounds determines where in the search space agents will look for performance improvements. In the autonomy setting, complexity of the landscape results in a negative performance feedback which will trigger explorative research on the performance landscape Therefore, as complexity increases, so it does search breadth. This result is confirmed, and it is even more clear, in the control scenario. Organizations, in the presence of pervasive interdependencies among their attributes need to rely on features supporting a more extensive search and establishing a target, indeed, influences individuals' understanding and evaluation of feedback and guides the search process (Marengo et al., 2022; Rivkin and Siggelkow, 2003). The control imposed by organizational structures has an impact on performance by directing agents' search behavior on the landscape they confront and making at the same time the result of search more effective through a clearer direction provided by the target (Baumann et al., 2019; Rivkin and Siggelkow, 2003). According to Bidmon & Boe-Lillegraven (2020) top-down guidance and specified behavioral expectations may improve ambidextrous behavior in employees. Especially in settings in which individuals are inhibited from deciding autonomously on how to balance proximity and distance search, it can be useful for agents to decrease autonomy even further in order to meet organizational expectations. In settings in which individuals are constrained from following their natural aptitudes, a closer control may help them conducting ambidextrous tasks (Bidmon & Boe-Lillegraven, 2020).

Finally, through the experimental setup, it was possible to account for a condition not widely explored in the literature. Introducing a penalty, in fact, associates an opportunity cost to the decision of exploring. As Billinger et al. (2014) experiment evidenced, individuals are inclined toward over-exploration, as they tend to cease neighborhood search too early. In their setting, in fact, a problem of "pure search" was enacted in which search was not associated with a downside risk.

In a setting in which additional search has a cost, in the autonomy setting, it appears clearly that for the same level of average aspirations feedback, agents sensitive to the presence of a penalty on average focused their research in the neighborhood of solutions known, whereas those that were not affected by the penalty looked for alternative solutions on a wider area of the search landscape. At the same time, in the control setting agents that showed no sensitivity to the introduction of a penalty reached a consistently higher value of the average search breadth relative to the average performance feedback than agents sensitive to the introduction of a penalty. As also evidenced by Greve (2010), the introduction of a penalty, with a target to be reached stimulates exploration up to a level necessary to achieve the performance target and simultaneously inhibits search from reaching hazardous levels. Therefore, the introduction of a penalty in a controlled setting reduces the average search breadth for the same level of average performance feedback.

The findings of the present work have also practical implications for established firms and emerging start-ups alike. Decision-makers like entrepreneurs, managers or even employees must deal with two difficult challenges. They need to understand what level of performance can be reached and what actions and plans need to be implemented in order to reach it. Receiving feedback helps in shaping expectations, mitigating overly optimistic or pessimistic options. Additionally, it helps decision-makers in deciding how a current competitive position or business models needs to be adapted (Billinger et al., 2021). In an established firm, alongside the fundamental function performed by feedback, defining a target to be reached by agents allows to direct innovation processes in a more effective way. As shown, in the controlled setting, it was possible to effectively reach the same level of average performance feedback with a lower level of average search breadth. This may help firms, especially those focused on innovative technologies and operating in complex environments to reach the same results with reduced efforts.

Nonetheless, the process of aspirations' formation in relation to the feedback received, affected, in turn, by the conditions of the environment faced may be a useful guide for entrepreneurs launching future ventures. For example, starting from individual aspirations, the findings of this study suggest that unfolding start-ups, should first understand what level of performance can realistically be expected, as resulting from feedback and the complexity of the landscape faced, and then on this basis develop a plan to achieve the desired results, rather than investing resources in testing solutions that may later result unfeasible. As an example, a successful approach when developing a business model may be the one based on the lean start-up methodology (Blank, 2013) according to which emerging businesses should test their hypotheses, collect frequently customers' feedback and on this basis developing "minimum viable products".

As with all research work, this study suffers from a series of limitations. First of all, as recalled, the findings are based on a pilot experiment. It would be interesting to replicate and adapt the same experiment to a larger sample in order to find a stronger evidence to support its main findings. In particular, with respect to the relationship between average search breadth and average performance feedback it would be interesting to see if with a larger sample the standard error would decrease in order to have a more precise analysis and eventually generalize its findings. Moreover, regarding the effect of introducing a penalty, due to sample restrictions it was not possible to replicate the experimental setting with and without the penalty. A bigger sample would allow to better account for the moderating effect of the penalty by distinguishing between clusters in which the penalty was or not introduced. Additionally for the control setting, it would be of the utmost interest to understand how to effectively set the reference target, by which internal and external considerations management is moved in establishing an objective rather than another and on what basis firms operating in the same landscape may decide to settle on different levels of performance.

CONCLUSIONS

This research work had the objective to investigate on the effects that autonomy – a setup in which agents are able to independently settle and reshape their aspirations in accordance with the performance feedback received – and control – a setting in which agents need to reach an externally imposed target – exert on individual search behavior. Additionally, this study tried to depict some of the effects that the introduction of a cost of exploration – a penalty – had on agents, under both conditions.

The search concept indicates the degree of change with respect to the initial status quo undertaken by an agent when confronting a complex performance landscape, constituted by a series of attributes and their respective intensity of interrelatedness. The construct represents a proxy to define if an observed behavior can be qualified as exploitative or explorative.

An experiment has been implemented, in order to observe the effect that the afore mentioned factors had on individuals. A crucial role in the empirical setting was played by the feedback that agents received in between the different phases, reflecting the conditions of autonomy and control and the complexity of the landscape faced.

The findings from this work contribute to the activities of scholars and practitioners alike. The results on the relationship between search breadth and performance feedback add to the literature on the behavioral theory of the firm and problemistic search. Additionally, the present research enriches the literature on the NK model through the findings related to the relationship between search breadth and the level of complexity. Finally, the present work addresses a dimension previously neglected by scholars and observes how the introduction of a penalty moderates the previous relationships. The present work opens future research paths for authors interested in testing how theoretical assumptions are actually reflected in agents' behaviors, in particular it would be interesting to test the moderating effect of a penalty on a larger sample to get valuable insights starting from this work contributions.

Finally, managers, especially those operating in innovative and complex contexts, could draw on the results of this study to implement organizational structures and objectives supporting a guided innovation process to reach their targets with a reduced deployment of resources, whereas entrepreneurs could rely on the presented findings and their underlying theoretical framework to structure a successful process of business model development.

REFERENCES AND BIBLIOGRAPHY

Andries, P., Debackere, K., & Van Looy, B. (2013). Simultaneous experimentation as a learning strategy: Business model development under uncertainty. *Strategic entrepreneurship journal*, 7(4), 288-310.

Audia, P.G. & Greve, H.R. (2006). Less likely to fail: Low performance, firm size, and factory expansion in the shipbuilding industry. *Management Science*, 52(1), 83-94.

Baillon, A., Bleichrodt, H., & Spinu, V. (2020). Searching for the reference point. *Management Science*, 66(1), 93-112.

Baum, J.A., Li, S.X., & Usher, J.M. (2000). Making the next move: How expeperiential and vicarious learning shape the locations of chains' acquisitions. *Administrative Science Quarterly*, 45(4), 766-801.

Baumann, O. (2015). Models of complex adaptive systems in strategy and organization research. *Mind & Society*, 14(2), 169-183.

Baumann, O., Schmidt, J., & Stieglitz, N. (2019). Effective search in rugged performance: A review and outlook. *Journal of Management*, 45(1), 285-318.

Benner, M.J., & Tushman, M.L. (2002). Process management and technological innovation: A longitudinal study of the photography and paint industries. *Administrative science quarterly*, 47(4), 676-707.

Bidmon, C. M., & Boe-Lillegraven, S. (2020). Now, switch! Individuals' responses to imposed switches between exploration and exploitation. *Long Range Planning*, 53(6), 101928

Billinger, S., Srikanth, K., Stieglitz, N., & Schumacher, T. R. (2021). Exploration and exploitation in complex search tasks: How feedback influences whether and where human agents search. *Strategic Management Journal*, 42(2), 361-385.

Billinger, S., Stieglitz, N., & Schumacher, T. R. (2014). Search on rugged landscapes: An experimental study. *Organization Science*, 25(1), 93-108.

Blank, S. (2013). Why the lean start-up changes everything. *Harvard business review*, 91(5), 63-72.

Blettner, D.P., He, Z.L., Hu, S. & Bettis, R.A. (2015). Adaptive aspirations and performance heterogeneity: Attention allocation among multiple reference points. *Strategic Management Journal*, 36 (7), 987-1005

Bromiley, P. (1991). Testing a causal model of corporate risk taking and performance. *Academy of Management journal*, 34(1), 37-59.

Camerer, C. (2003). *Behavioral Game Theory: Experiments in strategic interaction*. Princeton, NJ: Princeton University Press.

Chen, W.R. & Miller, K.D. (2007). Situational and institutional determinants of firms' R&D search intensity. *Strategic Management Journal*, 28(4), 369-381.

Csaszar, F. A., & Levinthal, D. A. (2016). Mental representation and the discovery of new strategies. *Strategic Management Journal*, 37(10), 2031-2049.

Cyert, R.M. & March, J.G. (1963). *A Behavioral Theory of the Firm*. Engelwood Cliffs, NJ: Prentice Hall.

Denrell, J., & March, J. G. (2001). Adaptation as information restriction: The hot stove effect. *Organization science*, 12(5), 523-538.

Duncan, R.B. (1976). *The ambidextrous organization: Designing dual structures for innovation*. In R.H. Kilmann, L.R. Pondy, & D. Slevin (Eds.), *The management of organization design: Strategies and implementation* (pp. 167-188). New York: North Holland

Friedman, M. (1953). *Essays in Positive Economics*. Chicago: Chicago University Press.

Gaba, V., & Greve, H. R. (2019). Safe or profitable? The pursuit of conflicting goals. *Organization Science*, 30, 647-667.

Gavetti, G., & Levinthal, D. (2000). Looking forward and looking backward: Cognitive and experiential search. *Administrative science quarterly*, 45(1), 113-137.

Gavetti, G. (2005). Cognition and hierarchy: Rethinking the microfoundations of capabilities' development. *Organization Science*, 16(6), 599-617.

Gibson, C.B., Birkinshaw, J. (2004). The antecedents, consequences, and mediating role of organizational ambidexterity. *Academy of management Journal*, 47(2), 209-226.

Goldstein, D. G., McAfee, R. P., Suri, S., & Wright, J. R. (2020). Learning when to stop searching. *Management Science*, 66(3), 1375-1394.

Greve, H.R. (2003). *Organizational learning from performance feedback: A behavioral perspective on innovation and change*. Cambridge, MA: Cambridge University Press.

Greve, H.R. (2003). A behavioral theory of R&D expenditures and innovations: Evidence from shipbuilding. *Academy of Management Journal*, 46(6), 685-702

Greve, H.R. (2008). A behavioral theory of firm growth: Sequential attention to size and performance goals. *Academy of Management Journal*, 51(3), 476-494.

Greve, H.R. (2010). Designing performance feedback systems to guide learning and manage risk. *Organizational Dynamics*, 39(2), 104-114.

Gruber, M., MacMillan, I. C., & Thompson, J. D. (2008). Look before you leap: Market opportunity identification in emerging technology firms. *Management science*, 54(9), 1652-1665.

Gupta, A. K., Smith, K. G., & Shalley, C. E. (2006). The interplay between exploration and exploitation. *Academy of Management Journal*, 49(4), 693-706.

He, Z.-L., & Wong, P.-K. (2004). Exploration vs. exploitation: An empirical test of the ambidexterity hypothesis. *Organization Science*, 15(4), 481-494.

Herriott, S. R., Levinthal, D., & March, J. G. (1985). Learning from experience in organizations. *The American Economic Review*, 75(2), 298-302.

Hey, J. D., Permana, Y., & Rochanahastin, N. (2017). When and how to satisfice: an experimental investigation. *Theory and Decision*, 83(3), 337-353.

Joseph, J. & Gaba, V. (2015). The fog of feedback: Ambiguity and firm responses to multiple aspiration levels. *Strategic Management Journal*, 36(13), 1960-1978.

Kauffman, S. A. (1993). *The origins of order: Self-organization and selection in evolution*. Oxford University Press, USA.

Labianca, G., Fairbank, J. F., Andrevski, G., & Parzen, M. (2009). Striving toward the future: aspiration – performance discrepancies and planned organizational change. *Strategic Organization*, 7(4), 433-466.

Lant, T. K. (1992). Aspiration level adaptation: An empirical exploration. *Management science*, 38(5), 623-644.

Laureiro-Martínez, D., Brusoni, S., Canessa, N., & Zollo, M. (2015). Understanding the exploration–exploitation dilemma: An fMRI study of attention control and decision-making performance. *Strategic management journal*, 36(3), 319-338

Lee, S., & Meyer-Doyle, P. (2017). How performance incentives shape individual exploration and exploitation: Evidence from microdata. *Organization Science*, 28(1), 19-38.

Levinthal, D. A. (1997). Adaptation on Rugged Landscapes. *Management Science*, 43 (7), 934-950

Levinthal, D. A., & March, J. G. (1993). The myopia of learning. *Strategic management journal*, 14(S2), 95-112.

Levinthal, D., & March, J. G. (1981). A model of adaptive organizational search. *Journal of economic behavior & organization*, 2(4), 307-333.

Levitt, B., & March, J. G. (1988). Organizational learning. *Annual review of sociology*, 14(1), 319-338.

Lounamaa, P. & March, J. G. (1987). Adaptive coordination of a learning team. *Management Science*, 33(1), 107-123.

March, J. G. (1991). Exploration and exploitation in organizational learning. *Organization science*, 2(1), 71-87.

March, J. G. (1988). Variable risk preferences and adaptive aspirations. *Journal of Economic Behavior & Organization*, 9(1), 5-24.

March, J. G., & Shapira, Z. (1992). Variable risk preferences and the focus of attention. *Psychological review*, 99(1), 172.

March, J. G. & Simon, H. A. (1958). *Organizations*. New York: Wiley.

Marengo, L., Romagnoli, M., Zeijen A. (2022). *Aspiration Levels and Survival Points: Search on Holey Performance Landscapes*. Unpublished manuscript.

Markowitz, H. (1952). The utility of wealth. *Journal of political Economy*, 60(2), 151-158.

Minniti, M., & Bygrave, W. (2001). A dynamic model of entrepreneurial learning. *Entrepreneurship theory and practice*, 25(3), 5-16.

Morris, M., Schindehutte, M., & Allen, J. (2005). The entrepreneur's business model: toward a unified perspective. *Journal of business research*, 58(6), 726-735.

Nelson, R.R. & Winter, S. G. (1982). *An evolutionary theory of economic change*. Cambridge, MA: Belknap Press of Harvard University Press.

O'Reilly, C.A & Tushman, M.L. (2004). The ambidextrous organization. *Harvard Business Review*, 82(4), 74-81

Ossenbrink, J., Hoppmann, J., & Hoffmann, V. H. (2019). Hybrid ambidexterity: How the environment shapes incumbents' use of structural and contextual approaches. *Organization Science*, 30(6), 1319-1348.

Posen, H. E., Keil, T., Kim, S., & Meissner, F. D. (2018). Renewing research on problemistic search – A review and research agenda. *Academy of Management Annals*, 12(1), 208-251.

Posen, H. E., & Levinthal, D. A. (2012). Chasing a moving target: Exploitation and exploration in dynamic environments. *Management science*, 58(3), 587-601.

Ref, O. & Shapira, Z. (2017). Entering new markets: The effect of performance feedback near aspiration and well below and above it. *Strategic Management Journal*, 38(7), 1416-1434.

Rivkin, J. W. (2000). Imitation of complex strategies. *Management science*, 46(6), 824-844.

Rivkin, J. W., & Siggelkow, N. (2003). Balancing search and stability: Interdependencies among elements of organizational design. *Management science*, 49(3), 290-311.

Rivkin, J. W., & Siggelkow, N. (2007). Patterned interactions in complex systems: Implications for exploration. *Management science*, 53(7), 1068-1085.

Rosenkopf, L. & Nerkar, A. (2001). Beyond local search: Boundary-spanning, exploration, and impact in the optical disk industry. *Strategic Management Journal*, 22 (4), 687-699.

Siggelkow, N. (2002). Evolution toward fit. *Administrative science quarterly*, 47(1), 125-159.

Simon, H. A. (1955). A behavioral model of rational choice. *The quarterly journal of economics*, 69(1), 99-118.

Simon, H.A. (1957). *Models of Man: Social and Rational*. New York: John Wiley & Sons.

Simon, H.A. (1959). Theories of decision-making in economics and behavioral science. *American Economic Review*, 49(3), 253-283.

Sterman, J. D. (1987). Testing behavioral simulation models by direct experiment. *Management Science*, 33(12), 1572-1592.

Tempelaar, M. P., & Rosenkranz, N. A. (2019). Switching hats: The effect of role transition on individual ambidexterity. *Journal of Management*, 45(4), 1517-1539.

Tushman, M. L., & O'Reilly III, C. A. (1996). Ambidextrous organizations: Managing evolutionary and revolutionary change. *California management review*, 38(4), 8-29.

Tushman, M. L., & Romanelli, E. (1985). *Organizational evolution: A metamorphosis model of convergence and reorientation*. Research in organizational behavior.

Vassolo, R. S., Anand, J., & Folta, T. B. (2004). Non-additivity in portfolios of exploration activities: A real options-based analysis of equity alliances in biotechnology. *Strategic Management Journal*, 25(11), 1045-1061.

Vissa, B., Greve, H.R., & Chen, W.R. (2010). Business group affiliation and firm search behavior in India: Responsiveness and focus of attention. *Organization Science*, 21(3), 696-712

APPENDICES

APPENDIX 1 – AUTONOMY SETTING EXPERIMENT NARRATIVE
TO EXPLORE OR TO EXPLOIT? AN EXPERIMENTAL STUDY ON THE EFFECTS THAT AUTONOMY AND CONTROL EXERT ON INDIVIDUAL SEARCH BEHAVIOR IN COMPLEX FITNESS LANDSCAPES

Instructions

Thank you for your time in participating in today's experiment. Please do not talk with other participants and do not communicate with other means.

The experiment

You are launching GREEN, a small company operating in the organic cosmetic industry producing sustainable face & body cleansers, creams and lotions. The GREEN products are targeted to high and medium earning female consumers sensitive to environmental problems interested in buying effective but responsibly sourced and produced products. Your objective is to develop a viable and successful business model that will allow you to reach a leading position in your market. In order to define your business model you will need to combine the different attributes provided. Please note that attributes are not important per se but it is how you combine the attributes that will determine if you will succeed. Their combination is what will define your final payoff.

Task

Consider that the highest payoff offered by the market in each trial is equal to 1. Given the competitive nature of the market in which you are entering, your objective is to maximize your payoff in each round. Finally, keep in mind that changing attributes has a cost, therefore a penalty of the 10% will be applied if by making these adjustments you will accidentally substitute the higher performative alternative with the less performing one.

How does it work?

In order to execute your task you will need to answer the following six questions regarding the business model of the company. You will be provided with an initial combination that will be the same for all the experiment participants. Through a process of trial-and-error, no economic or previous knowledge required, you will need to select one of the alternatives proposed in each question. During each round you can choose to change none, some or all the attributes with respect to the initial combination or previous round. At the end of each round, please provide an answer to the point of what you expect your payoff to be. After all of the participants will submit their questions, you will receive feedback on your performance. The same questions will be repeated for 6 rounds. At the end of the last round, you will have to answer a short questionnaire on the decision-making process followed throughout the experiment.

Questions

1) *How does the company creates value?*

   a. Focusing on the R&D efforts for its innovative products (0)
   b. Focusing on a high customization of its products for its targeted customers (1)

2) *Who does the company create value for?*

   c. Enlarge the potential market by extending distribution abroad (1)
   d. Keep focusing on the initial niche market (0)

3) *What is the company source of competence?*

   e. Investing in marketing efforts (0)
   f. Improving the supply chain management of sustainable feedstock (1)

4) *How does the company competitively position itself?*

   g. Stressing on the intrinsic quality of its products (1)
   h. Developing tight customer relationships (0)

5) *How does the company make money?*

   a. Focusing on competitive pricing and volumes (0)
   l. Relying on high retail margins (1)

6) *What are the company's ambitions?*

   m. Growth Model: focusing on long-term strategy to generate a capital gain for investors (1)
   n. Income Model: focusing on a medium-term strategy to invests up to the point that the business is able to generate a stable income stream (0)

Please indicate what you believe your payoff to be at the end of the round:

Final Questionnaire

1) How many attributes did you change on average during each round ?
   a. 0-1
   b. 2-4
   c. 5-6

2) When receiving feedback on the previous rounds, did you put more weight on your payoff in comparison to your own past performance or in comparison to the average payoff relative to your own?
   a. I put more weight on my payoff in comparison to my previous performance
   b. I put more weight on my payoff in comparison to the average performance

3) How many attributes did you change when your performance was below your expectations?
   a. 0-1
   b. 2-4
   c. 5-6

4) How many attributes did you change when your performance was above your expectations ?
   a. 0-1
   b. 2-4
   c. 5-6

5) Knowing that changing a performative attribute could heavily affect your final score, has this feature inhibited you from changing a greater number of attributes?
   a. Yes
   b. No

APPENDIX 2 – CONTROL SETTING EXPERIMENT NARRATIVE TO EXPLORE OR TO EXPLOIT? AN EXPERIMENTAL STUDY ON THE EFFECTS THAT AUTONOMY AND CONTROL EXERT ON INDIVIDUAL SEARCH BEHAVIOR IN COMPLEX FITNESS LANDSCAPES

Instructions

Thank you for your time in participating in today's experiment. Please do not talk with other participants and do not communicate with other means.

The experiment

You are working for GREEN, a small company operating in the organic cosmetic industry producing sustainable face & body cleansers, creams and lotions. . The GREEN products are targeted to high and medium earning female consumers sensitive to environmental problems interested in buying effective but responsibly sourced and produced products. Due to a loss of market share and the subsequent financial distress in which the enterprise finds itself, the founder and CEO Bill is asking you, his employees and collaborators, suggestions to update the current business model and increase its profitability. Your objective is to update the company business model to reach an established target.
In order to adjust your business model you will need to combine the different attributes provided. Please note that attributes are not important per se but it is how you combine the attributes that will determine if you will succeed. Their combination is what will define your final payoff.

Task

Consider that the highest payoff offered by the market in each trial is equal to 1. Given the niche market in which GREEN operates, the CEO wants to maximize the profits of the company in relation to its direct competitors. Therefore, your objective it is not to obtain the highest payoff possible, but to identify a combination of attributes that guarantees in each trial a payoff between 0.6 and 0.8. Given the innovative nature of the company, Bill believes in the importance of fostering intrapreneurship within his organization and he is asking you to propose the necessary adjustments to improve its performance. Finally, keep in mind that changing attributes has a cost, therefore a penalty of the 10% will be applied if by making these adjustments you will accidentally substitute the higher performative alternative with the less performing one.

How does it work?

In order to execute your task you will need to answer the following six questions regarding the business model of the company. You will be provided with an initial combination that will be the same for all the experiment participants. Through a process of trial-and-error, no economic or previous knowledge required, you will need to select one of the alternatives proposed in each question. During each round you can choose to change none, some or all the attributes with respect to the initial combination or previous round. At the end of each round, please provide an answer to the point of what you expect your payoff to be. After all of the participants will submit their questions, you will receive feedback on your performance. The same questions will be repeated for 6 rounds. At the end of the last round, you will have to answer a short questionnaire on the decision-making process followed throughout the experiment.

Questions

1) *How does the company creates value?*

   b. Focusing on the R&D efforts for its innovative products (0)
   c. Focusing on a high customization of its products for its targeted customers (1)

3) *Who does the company create value for?*

   d. Enlarge the potential market by extending distribution abroad (1)
   e. Keep focusing on the initial niche market (0)

4) *What is the company source of competence?*

   f. Investing in marketing efforts (0)
   g. Improving the supply chain management of sustainable feedstock (1)

5) *How does the company competitively position itself?*

   h. Stressing on the intrinsic quality of its products (1)
   i. Developing tight customer relationships (0)

6) *How does the company make money?*

   i. Focusing on competitive pricing and volumes (0)
   l. Relying on high retail margins (1)

7) *What are the company's ambitions?*

   m. Growth Model: focusing on long-term strategy to generate a capital gain for investors (1)
   n. Income Model: focusing on a medium-term strategy to invests up to the point that the business is able to generate a stable income stream (0)

Please indicate what you believe your payoff to be at the end of the round:

Final Questionnaire

1) How many attributes did you change on average during each round?
   a. 0-1
   b. 2-4
   c. 5-6

2) When receiving feedback on the previous rounds, did you put more weight on your payoff in comparison to your own past performance or in comparison to the average payoff relative to your own?
   a. I put more weight on my payoff in comparison to my previous performance
   b. I put more weight on my payoff in comparison to the average performance

3) How many attributes did you change, in the following round, when your performance feedback was below your expectations?
   a. 0-1
   b. 2-4
   c. 5-6

4) How many attributes did you change, in the following round, when your performance was above your expectations?
   a. 0-1
   b. 2-4
   c. 5-6

5) Knowing that changing a performative attribute could heavily affect your final score, has this feature inhibited you from changing a greater number of attributes?
   a. Yes
   b. No