Cognitive microfoundations of search

Submitted to Organization Science. Please do not cite or distribute without permission.

Carsten Bergenholtz[1], Oana Vuculescu[1], Ali Amidi[2]

[1]Department of Management, BSS, Aarhus University. [2]Department of Psychology, BSS, Aarhus University

Abstract

This paper investigates the cognitive antecedents of individual search behavior in a combinatorial, complex landscape. We present results from 3 studies where 375 individuals solve a gamified problem-solving task. We complement these with measurements of cognitive styles and established neuropsychological tests of the players' cognitive abilities. The task environment allows us to distinguish between local and global search, and also to identify directed global search that takes departure in an understanding of the underlying problem structure. We document systematic heterogeneity of search, showing that individuals with certain cognitive styles and cognitive abilities engage in more local, systematic search and less undirected, global search. The archetypical assumptions of an NK model thus relies upon a particular conception of individual cognition. Integrating insights from both cognitive psychology and management, we contribute to research on the microfoundations of search, highlighting that cognitive antecedents can be as important a factor for explaining various kinds of search, as the external performance feedback that is the core of the adaptive search mechanism. These insights into the role of cognition have implications for what assumptions simulation models should rely upon and how managers can influence the search behavior of individuals.

1. Introduction

Organizations and individuals constantly face the decision of either doing the same, or doing something new in the pursuit of improving their performance. For example, an organization might consider if it should refine and market (i.e. exploit) a current technology or invest further into R&D to explore new options (He & Wong 2004), while an individual may ponder if she should choose a known or new restaurant (Schulz et al. 2019).

Recent research emphasizes adaptive search as a key mechanism via which individuals and organizations alternate between exploitation and exploration. According to this research, positive performance feedback leads actors to focus on doing the same, while negative performance leads them to try new options (Greve 2003, Billinger et al. 2014, Vuculescu 2017). In alignment with this line of thought, simulation studies have assumed homogeneous individual search behavior solely contingent on performance feedback, and heterogeneity is therefore set to stem from external factors only (Baumann et al. 2019; Smith & Rand 2018). Individuals are conceptualized as mainly carrying out local search (exploitation), while relatively rare jumps constitute undirected global search (exploration). Importantly these jumps are assumed to be in a random direction, unrelated to any understanding of the problem at hand. The above outlined research perspective on how organizations and individuals balance the trade-off, is prevalent in management (Greve 2003), biology and behavioral economics (Nowak & Sigmund 1993) as well as cognitive psychology (Cohen et al. 2007). Yet, while we acknowledge the usefulness of this type of explanation, we see two core challenges.

First, a simple binary distinction between either local or global search does not capture that global search can be of fundamentally different types. To illustrate, simulation studies have had to assume that global search is due to an undirected jump landscape (e.g. Levinthal 1997, although see Gavetti & Levinthal 2000). In contrast, studies of how individuals actually search in controlled lab-settings or in the real world, document that individuals do not merely adapt mindlessly to performance feedback, but can rely on their cognitive abilities to make sense of the environment and engage in a directed global search move (Doll et al. 2016; Gary & Wood 2011; Wilson et al. 2014). When having a poor restaurant experience, we do not merely pick out a randomly selected new restaurant the next time, but one that fits our priors and is different in type to the one that disappointed us (Schulz et al. 2019).

Second, an emerging stream of research on the microfoundations of search behavior (Felin et al. 2015) has re-established March’s (1991) original focus on the heterogeneity of individuals as an important source of organizational heterogeneity. In March’s simulation, it was the turnover of heterogeneous individuals that led to various degrees of exploration and exploitation. After more than two decades, this idea has been brought forward again, this time bringing quantitative evidence. A range of empirical studies document individual heterogeneity, questioning the assumption about homogeneous, adaptive search; risk preferences (Hills & Hertwig 2010), emotional valence (Døjbak et al. 2015) and cognitive factors (Laureiro-Martínez & Brusoni 2018; Levine et al. 2017) shape individuals' tendency to engage in local or global search. While cognitive psychology has clearly established how cognition is strongly related to a wide range of important real-world behaviors and outcomes such as job performance and happiness (Ritchie 2015), surprisingly little of this research is related to human propensities to explore or exploit (O'Doherty et al. 2017). This leaves open ample opportunities to investigate such relations in order to benefit both cognitive psychology and management research on search. We claim that identifying particular cognitive antecedents of individuals has the potential to shed further light on when and why individuals break free from local, adaptive search, enabling us to further theorize how a homogeneous explanation based solely on performance feedback is going to be inadequate. We therefore want to examine if individual variations in the core ability to process information (cognitive ability) and differences in how individuals process this information (cognitive style) are important factors shaping the propensity to engage in search. In contrast to previous work, which looks at measures of analytical and reasoning ability (such as the critical reflection test (Levine et al. 2017)), the Raven's test (Steyvers et al. 2009), we focus on a more fine-grained measure of cognition (Lezak et al. 2012), allowing us to distinguish between particular cognitive building blocks such as memory, ability to sustain attention, executive functions etc.

We present data from three studies involving 375 individuals trying to solve a novel, game-based problem-solving task (Vuculescu 2017). Like the NK, the setup is a rugged fitness landscape. Yet, in the NK any global (exploration) search is necessarily undirected, since the search space is not based on any meaningful, underlying structure (Levinthal 1997). Therefore, and in contrast to the NK framework, participants can in the present task learn the underlying problem structure, which opens up the opportunity to track: i) local search, ii) undirected global search (i.e. random), and iii) directed global search (i.e. model-based) that is aligned with the identified problem structure (Doll et al. 2016; Vuculescu 2017). This allows meaningful global search that is not just a random, long jump.

We present three interrelated findings on how cognitive characteristics constitute antecedents of search in a rugged fitness landscape. First, we document a systematic heterogeneity of search, in contrast to the homogeneous assumptions established in the literature (Billinger et al. 2014; Smith & Rand 2018). Search is partially driven by individuals' information processing approach, both in terms of their cognitive ability (Lezak et al. 2012) as well as their cognitive style (Kirton 1976). Therefore, if managers want to shape what search behavior individuals are to engage in, cognitive factors appear as important as the main mechanism of adaptive search; external, performance-based feedback. On a more general level, insight into cognitive antecedents can also help one in assessing if one has the right match between the crowd at hand and the type of problem to be solved (Felin & Zenger 2014). Second, and more particularly, the stronger one's ability to sustain attention, executive functions and aggregated overall cognitive ability, the more individuals prioritize local search over global undirected search, i.e. engage in search behaviors more akin to the assumptions in NK simulations (Levinthal 1997; Baumann et al. 2019). Such individuals carry out the more persistent search, that Billinger et al. (2014) argued individuals, usually, should prioritize. These findings have substantial, theoretical implications for how to implement behavioral assumptions in simulation models, since the more empirically realistic assumptions are likely to shape the simulated outcomes (Puranam et al. 2015; Smith & Rand 2018). Furthermore, we expand cognitive psychology theories on how cognitive antecedents shape behavior (Chan et al. 2019), by specifying what cognitive antecedents influence search for novelty (Helfat & Peteraf 2015). Third, we provide proof of concept that an inherent ability to learn a problem structure can shape future search behavior, which can be considered a microfoundational basis for Gavetti & Levinthal's (2000) simulations on the role of cognitive representations. Overall, we move beyond the more generic insight that cognition matters for search, to a more fine grained theory on how particular cognitive building blocks shape search behavior, as called for by O'Doherty et al. (2017) and Helfat and Peteraf (2015).

In the following section we present how cognitive psychology and management has approached the challenge of studying search, as well as their attempts to identify relevant antecedents of variation in search strategies. In the methods section we present the experimental task as well as the variables relied upon. We then present the results of how cognitive abilities and styles relate to search behavior and finally discuss the theoretical implications these findings.

2. Organizational search

Theorizing on organizational search is shaped by March's (1991) seminal paper on the challenge of balancing exploration and exploitation. A stream of empirical research has focused on identifying how organizational search strategies depend on industrial contexts (Yamakawa et al. 2011), performance feedback and aspiration levels (Greve 2003) or the nature of the problem to be solved (Felin & Zenger 2014; Lakhani et al. 2013). In addition to this macro-level perspective, a more recent stream of literature focuses on individuals, the ones who are carrying out the search within organizations (Li et al. 2013), and describes e.g. how their organizational roles (Nigam et al. 2016), their networks (Fleming 2002) or the search environment (MacAulay et al. 2017) influence search trajectories. The results from these studies challenge traditional simulation assumptions, where agents' behavior is assumed to be homogenous (Knudsen et al. 2019; Smith & Rand 2018) and largely engaged in local search (e.g. Levinthal 1997). We take departure in this tension and in the following sections present management's and cognitive psychology's perspective on studying search.

2.1. Perspectives on search: Management and cognitive psychology

Experimental and simulation-based studies of exploration vs. exploitation in management have generally relied upon one of two search models; the armed bandit that allows for the study of the trade-off between few uncertain options, or the fitness landscape which allows for the study of how agents adapt to feedback in a large, complex landscape (Puranam et al. 2015; Knudsen et al. 2019). An exploitation move is local; either staying at the same arm in the armed bandit (Laureiro-Martínez et al. 2015; Steyvers et al. 2009), or searching myopically in the vast search space (Billinger et al. 2014; Reypens & Levine 2018). Explorative search is, thus, categorized as all search moves that are non-local (undirected global), since these search environments generally have not allowed the participant in the experiment to create a meaningful model of how to engage in directed global search.

Cognitive psychology and life science studies in general also have a strong interest in the exploration vs. exploitation trade-off, since it constitutes the basis for how organisms are to adapt to their dynamic environments, a problem that allows no final solution (Cohen et al. 2007). However, compared to a typical approach in management, cognitive psychology employs – in some sense - a more sophisticated conceptualization of the search options that agents face. Computational models of how individuals have meaningful mental representations of the world have recently been developed, allowing for the discovery of how underlying patterns can shape future search decisions and thereby engage in more efficient search (Doll et al. 2016; Doll et al. 2012). Cognitive psychology has, thus, developed a framework where global search can both be undirected (i.e. random in simulation terminology (Kauffman 1993; Levinthal 1997)) or directed (i.e. model-based).

While relying on insights from cognitive psychology is useful, we note that the field often relies on relatively simple problems characterized by uncertainty, such as armed bandits (Doll et al. 2012; Steyvers et al. 2009) or, literally, selecting dots on a screen (Bahrami et al. 2010). The problems are, thus, rarely complex involving interdependence across search options. Since the field's aim is to uncover basic cognitive processes and structures, such problems are useful since they allow researchers to isolate the mechanisms of interest. Yet, in response to this perspective, Smaldino & Richerson (2012) point out that this is not entirely uncontroversial, because even when solving a relatively mundane problem such as how to navigate in a restaurant, individuals often have to generate options rather than merely selecting among pre-established options. Most situations require some form of option generation (Smaldino & Richerson 2012). Furthermore, the viable options to be generated will rarely be arbitrarily global, but meaningful new options that reflect the focal actor's understanding of the environment; as Smaldino & Richerson (2012) point out, hardly any of us consider 'punching the waiter' as a viable solution among the set of options, while we might try to engage in non-local search and explore a completely different menu. The importance of the difference between selecting between pre-defined options vs. shaping meaningful new options is further supported by recent developments in neuroscience: "choice behavior […}differ[s] between self-generated and externally provided options" (Kaiser et al. 2013: p. 815). When studying individual search behavior, we can't necessarily expect to transfer insights form one search paradigm to another.

Rather than targeting the nature of the internal cognitive structures and processes of individuals, the aim of management is to be able to account for more complex scenarios. The NK fitness landscape has been widely used to understand how individuals and organizations should search, since it presents a context-independent, large combinatorial search space and allows for complexity in terms of interdependence between choices (Baumann et al. 2019). Similar to findings at the organizational level (Greve 2003), empirical studies reveal that individuals engage in performance dependent adaptive search. Individuals make small (local) adjustments when doing well and bigger (global) jumps when performing poorly (Billinger et al. 2014; Vuculescu 2017). In this framework, however, the adaptive search concept is model-free (undirected global), in the sense that there is no underlying problem structure, which prohibits any form of model-based (directed global) search that relies on a meaningful mental representation of the problem. Nevertheless, one can find support for implementing rationalized assumptions about search behavior, where agents primarily engage in local search and occasionally explore (Levinthal 1997).

The importance of studying directed global search is also recognized in management research, since scholars have discussed the importance of acknowledging that individuals and organizations do not merely do something (anything) different, but are able to follow a mental map (cf. Gary & Wood 2011; Helfat & Peteraf 2015). The creator of the NK model, Stuart Kauffman, makes a similar point when arguing that individuals do not merely make a random choice among a range of options, but choose what is considered the best option (Gabora & Kauffman 2016). Furthermore, Gavetti & Levinthal (2000) presented simulations where cognition is conceptualized as forward-looking search, i.e. agents were encoded with a priori representations of the search space which enabled directed global search behavior.

Cognition has been speculated to enable this different and more efficient search. A long tradition of research in psychology has uncovered how a range of basic cognitive abilities[i] shape an individual's ability to engage in higher level reasoning when solving complex problems (Dams-O'Connor & Gordon 2013; Luria 1976; Murray et al. 2017; Ritchie 2015). In other words, stronger cognitive abilities should lead to an improved ability to generate useful mental models, further facilitating more efficient search (Chan et al. 2019; Helfat & Peteraf 2015). However, a clear link between cognitive abilities and model-based search behavior has been difficult to establish, since the applied empirical frameworks (be it experimental or qualitative) have not allowed for a quantitative categorization of global directed search.

The thorny problem is that even when we have categorized a participant's search behavior, it is difficult to reverse engineer from this individual's behavior to their internally hidden mental model. Therefore, cognitive scientists have favored artificial and simple tasks that not only allow for an exhaustive mapping of all available actions, but also for a mathematical approximation of how available information can lead to different models of the world, which in turn are directly linked to certain actions. Moving beyond the artificial laboratory setup, a recent study enabled physicists and citizen scientists to engage with a highly complex, real world physics landscape (Heck, Vuculescu et al. 2018). In this study, participants were asked to cool down atoms via a gamified and remote interface. In line with previous findings, it was documented that adaptive search is a key driver of search behavior, even in this – both to the physicists and citizen scientists - unknown landscape. Now, the citizen scientists did not have relevant theoretical maps, but the physicists could engage in model-based search based on their theoretical insight into the behavior of ultra-cold atoms. Interestingly, citizen scientists engaged in far more explorative search than the model-based algorithms used by experts (Heck, Vuculescu et al. 2018). This setup illustrates the challenge that management research faces; to quantitatively study how agents search in an environment where meaningful mental models of the underlying problem-structure can be generated (Helfat & Peteraf 2015).

Vuculescu (2017) presented a further step towards this goal, since the experimental task used in her study allows for meaningful search. Based on a rugged fitness landscape, one can trace search strategies employed and distinguish, not merely between local and global search, but also between undirected (random) global search and directed (model-based) global search. While microfoundational antecedents of engaging in either local or global undirected search have been established (Steyvers et al. 2009; Laureiro-Martinez et al. 2015; Døjbak et al. 2015; Levine et al. 2017; Reypens & Levine 2018), we are not aware of any empirical studies that have investigated heterogeneous individual propensities to engage in these three search strategies. In other words, what are the cognitive microfoundations of search behavior in a combinatorial, rugged fitness landscape that also allows directed global search? Being able to identify which cognitive microfoundations (e.g. cognitive abilities) are related to a certain type of search behavior, may help a manager optimize the organization of search. For example, individuals with poorer executive functions may be less prone to engage in local, systematic search and less able to shift their search, requiring managerial intervention in order to enable this (see Chan et al. (2019) for a literature review on the role of executive functions in management.

2.2. Individual antecedents of search behavior

As pointed out by Laureiro-Martinez et al. (2010), the original March (1991) study took departure in individual heterogeneity in order to simulate variations in organizational outcomes. An emerging stream of research has investigated the microfoundations of this heterogeneity of search (Reypens & Levine 2018), the idea being that one can understand higher level outcomes by understanding micro-level processes and antecedents (Felin et al. 2015). Empirical research on individual antecedents of managerial search behavior has shown how variation in age, education and experience influences performance (Finkelstein et al. 2009; Hambrick & Mason 1984). Helfat & Peteraf (2015) further ask what the cognitive underpinnings of managerial capabilities are, and call for empirical research to explore and validate such links. The importance of individual antecedents is well established by now, since neurological activity (Laureiro-Martinez et al. 2015), risk sensitivity (Hills & Hertwig 2010), the emotional valence (Døjbak et al. 2015), cognitive flexibility (Laureiro-Martínez & Brusoni 2018) and analytical ability (Levine et al. 2017) shape search behavior.

We again draw upon cognitive psychology to advance the paradigmatic approach of studying how variation in cognition might influence organizational search behavior. Cognitive- and neuropsychology has a long history of mapping cognitive abilities. Whereas the aforementioned studies have mapped brain activity (Laureiro-Martinez et al. 2015), proxies for analytical skills such as the critical reflection test (Levine et al. 2017) and cognitive flexibility (Laurerio-Martinez & Brusoni 2018), less focus has been on identifying which lower or higher level cognitive building blocks that constitute the foundation for complex problem solving reasoning processes that we are interested in (Lezak et al. 2012). Even though cognitive psychology has carried out extensive work on search behavior, curiously little is known about how cognitive antecedents are related to exploration vs. exploitation: "*Almost nothing is known about the role played by the brain's varied control systems with respect to exploration… Issues pertinent to the brain's engagement with exploratory decision making are ripe for both theoretical and experimental research*." (O'Doherty et al. 2017: p. 91). We therefore intend to draw upon the methodological insights from cognitive psychology, in order to examine individual antecedents of search behavior. Such antecedents can refer to variation in cognitive ability but also variation in cognitive style, both of which will be covered in the following two sections.

2.3. Cognitive styles

The concept of cognitive styles has been developed to capture how individuals differ in how they perceive and process information (Miron-Spektor et al. 2011; Sternberg & Grigorenko 1997) and integrate this information in their "mental models" (Hayes & Allinson 1998: p. 850), rather than capturing creativity in a quantitative sense, e.g. in terms of number of ideas created. The concept thus has explicit links to our conceptualization of how individuals search.

A number of measures of cognitive styles have been developed, e.g. the adaption-innovation (Kirton 1976), analytic-intuitive (Hayes & Allinson 1998) and field dependence-independence (Witkin & Goodenough 1977). All show that cognitive styles are similar in kind to personality traits, in that they are stable over time and involve no one optimal style. The Adaptors-Innovators theory (Kirton 1976) is developed in the context of problem-solving and allows for a distinction between adaptors and innovators as given by an individual's preferred strategy for problem-solving: 'highly adaptive' people tend to rely on established solutions ('do things better') while 'highly innovative' people tend to do the reverse ('do things differently') (Kirton 1976). Innovators are more likely to reframe the given problem, while adaptors are more likely to accept and be preoccupied with how the problem is represented at the moment (Kirton 2003). Importantly, cognitive styles appear unrelated to cognitive ability (Kirton 2003). The cognitive style measure has become widely used in organizational theory, since it has been repeatedly shown to influence behavior and performance in organizations (Carnabuci & Dioszegi 2015; Miron-Spektor et al. 2011).

In this study, we aim to investigate whether individual differences in information processing styles are related to how individuals tend to navigate (i.e. search) a solution landscape, rather than assessing creativity or the final ability to solve the problem as such. More specifically, we theorize that the 'wider search' associated with innovators not only leads to a more divergent output (Kirton 2003) but translates into a continuous lower propensity to engage in local search during an ongoing search process (or higher propensity to engage in global search). In contrast, individuals that focus on 'doing things' better have a higher propensity to engage in local search (or less propensity to engage in global search).

2.4. Cognitive ability

Cognitive abilities are considered to be hierarchical in nature as lower-level cognitive building blocks such as processing speed, attention, and working memory have been shown to facilitate the ability to engage in more complex, higher level reasoning (Lezak et al. 2012; Luria 1976). A wide range of studies have shown that strong cognitive abilities lead to better individual outcomes such as higher job performance (F. L. Schmidt & Hunter 1998), higher grade point averages, higher salary and even longer life expectancy (Ritchie 2015).

In terms of search more specifically, a recent stream of studies has shown that individuals with better cognitive abilities are able to identify and cope with strategic uncertainty more efficiently (Benito-Ostolaza et al. 2016), e.g. when engaged in an armed bandit scenario (Steyvers et al. 2009) or making sequential investment decisions in a market (Levine et al. 2017). However, these studies have relied on the Cognitive Reflection test (Frederick 2005), which is a measure of an individual's tendency to override one's intuitive response, or general intelligence tests such as the Raven's test, which is a non-verbal assessment of fluid intelligence (Raven et al. 2003). While the Raven's test has been shown to be a relatively strong predictor of intelligence (capturing 50% of the variance (Gignac 2015)), it is a measure of higher order reasoning ability. These kinds of tests capture a general ability to problem solve and reasoning, and not the underlying microfoundations and building blocks of higher-level cognitive abilities. Chan et al. (2019) emphasize that one should not conflate general intelligence with more fine-grained cognitive abilities and call for more nuanced insights. Granted, cognitive abilities are generally correlated, but first of all, systematic variation across abilities do exist, and second, cognitive efforts can be improved by framing or training (Chan et al. 2019), which opens up the potential for managerial interventions.

Overall, research has clearly documented that cognitive abilities impact behavior and can improve search performance. However, to the best of our knowledge, no studies have linked the widely used NK framework with any measure of cognitive abilities such as the Raven's test, not to mention more fine-grained neurocognitive tests. It is therefore an important open question if the prevalent assumption in NK models of individuals individual' adaptive search behavior is dependent on cognitive ability. This also implies that no studies have linked the underlying building blocks of higher-level cognitive abilities with any rugged fitness landscape.

One can theorize that higher-level cognitive abilities such as reasoning (a general high intelligence score and executive functions) should facilitate the generation of a more appropriate model of the problem-structure and, thus, reduce inefficient search. That is, these individuals will engage in less undirected global search, searching more akin to the relatively greedy, local search portrayed in the NK model (Levinthal 1997). Furthermore, people with a strong ability to learn should not merely identify that undirected global search is inefficient but have a bigger likelihood to uncover the underlying problem-structure, and thus engage in directed global search.

3. Experimental setting, methods and data

We utilize four different types of data from three different studies in order to quantitatively analyze how 375 individuals navigate the search landscape in the 'Alien Game' (see Vuculescu 2017). One type of data is based on a quantitative coding of the search behaviour of individuals trying to solve the experimental task. A second type of data is based on a quantitative survey on cognitive styles that 336 individuals completed just before playing the game. In order to also capture cognitive ability we added another round of data collection where we included a third type of data, which captures the cognitive ability of 39 individuals, relying on established neuropsychological tests. Finally, we have also interviewed 40 individuals just after they finished the game, in order to complement (Greene et al. 1989) and validate our quantitative analysis as well as provide qualitative insight into the search process. Insights from interviews are anecdotally reported in the paper, as well as in more detail in appendix 2.

3.2. Participants in the studies

In the first study we relied on Amazon's Mechanical Turk (Mturk) platform to collect data on 270 participants' cognitive style as well as how they navigated the Alien Game. Mturk is a web-based outsourcing platform which is widely used in behavioural research, due to the ease of access to a relatively large group of people that has a closer resemblance to the general (US) population than a student sample (Berinsky et al. 2012; Hauser et al. 2019) and relatively low costs in setup. Before the statistical analysis was carried out, 28 participants were dropped from the Mturk sample (leading to an N of 244), since an inspection of their answers and game play revealed they had not actually engaged in the game (e.g. submitted the same solution many times in a row) or they provided incoherent replies to the survey, e.g. clicking the same options throughout the survey. The Mturk participants (42% women, average age 34, 9.89 SD) report that 98% have completed at least their secondary education, while 48% have at least a bachelor degree, indicative of a relative high level of education, which is consistent with previous work on Mturk demographics (Berinsky et al. 2012). In the second study we replicated study 1 in the lab and collected data from 97 participants, who were students at a large European University. The study was conducted by research assistants, not co-authoring our paper nor familiar with any research expectations. Participants are recruited using an internal recruitment system. Three lab participants were dropped based on the same criteria as in the Mturk study (leading to an N of 94). The lab-sample were all students at the university and consists of 49% women, with an average age of 24 years (5.84 SD), which matches the overall pool of participants at the lab. On average participants spent 20 min. on the game and survey. In the third study we collected data from 39 participants at the same university. Participants first played the game in one lab-session, and at least five days later came back to the lab to complete six batteries of neuropsychological tests of cognitive ability, which took approximately 45 min. to complete and were administered by two trained research assistants under the supervision of one of the co-authors, an expert in neuropsychology. Individuals were paid a flat-rate for the cognitive ability test.

3.2. Experimental task

Before playing the Alien Game, participants watched a video with game instructions and played a tutorial level to become familiar with the game environment. Participants have 25 attempts to toggle 8 tiles in any of two positions. They could decide to change one tile at a time, or all 8, which in total allows 256 possible combinations (see Figure 1 and https://youtu.be/b8BCkq93ovc for an illustration of how the game is played). After each attempt players get feedback in the form of a points score, which ranges from 18 to 42 points.

The underlying function of the experimental task is an 8-bit ‘hierarchical exclusive OR’ (H-XOR) (Watson & Pollack 1999). This function implies an underlying structure to the landscape, in contrast to setups that rely on the NK model, where solutions depend on random draws from a uniform distribution (Billinger et al. 2014). The solver is not directly informed about the H-XOR function, but merely sees the sequence of eight tiles and can attempt to learn from the variation in points obtained for the solutions submitted. In other words, participants could extrapolate from their past attempts. Players were paid based on performance when playing the game; i.e. the higher the score and the quicker the high score is achieved, the better the reward. In order to eliminate variance due to different performance expectations, players are informed in advance about the maximum score as well as the monetary performance reward. Depending on performance participants received between 8 and 20 Euros.

Qualitative studies have studied how individuals rely on mental models to navigate a complex environment (Gary & Wood 2011), but our game design allows an opportunity to study how mental models might influence subsequent search behaviours in a quantitative framework. The design positions us somewhere in between the two extremes of previous quantitative work: Participants don’t merely select between very few options (Laureiro-Martinez et al. 2015), but in contrast to Heck, Vuculescu et al.’s (2018) problem of how to cool down atoms, the space of solution is mapped out and there is an optimal solution and an efficient way of navigating towards it. Therefore, our task involves options that are not all obvious from the start but have to be shaped in a sequential, feedback-based process. This also makes our task computationally demanding, since it involves (exponentially) many combinations of future actions. While the experimental task relied upon is artificial, limiting external validity, we note that our ambition is to study the search strategies employed while navigating the search space, rather than who ends up at a certain solution. Unless solvers due to sheer luck find the optimal solution in their first attempt, they have to try out a number of actions to understand a) what the solution space looks like and b) how to search this solution space.

3.3. Variables: Search strategies

In order to study how individual antecedents explain variation in human search behavior, we first need to distinguish between different types of search strategies. Building on recent cognitive psychology literature (Doll et al. 2012) and previous research on this type of game (Vuculescu 2017) we rely on the following three strategies: local, directed global (model-based) and (undirected) global search. We thus differentiate between two fundamentally different search mechanisms: *model-free* mechanisms (local, undirected global) which operate without an internal representation of the problem space and *model-based* (directed global) mechanisms which rely on the agent having (acquired) an approximate representation of the problem space (Doll et al. 2012). We code local search moves as moves involving exactly one-bit flip from their reference point (be it their own best score so far or their most recent solution), model-based (or directed global) moves as moves involving exactly two-bit flips from their reference point without violating the underlying problem structure[ii], and global undirected search as everything else. We attempt to capture the fact that solvers can form a model of the problem they are attempting to solve and let their subsequent moves attempt to exploit that. One limitation in this coding is that directed global search is not normative. Participants can engage in directed global search moves that violate the underlying problem structure and thus will not be captured by the coding. However, this scheme is preferred since it is the most conservative.

We thus have three binary variables (one for each search strategy) which constitute our dependent variables. Note that the dependent variable does not refer to performance in the game as such, but how individuals navigate the solution space. We only analyse submissions 3-25, where all three search strategies are available[iii].

3.4. Variables: AI cognitive styles

We use Miron et al.'s (2004) 12-item scale to capture the adaptors-innovators constructs. Their questionnaire relies on a 7-point Likert-type scale that captures the three factors that Kirton (1976) also identifies: 1) creativity 2) conformity and 3) efficiency which Miron et al. (2004) label 'attention-to-detail'. Since we have used the instrument exactly as it is presented in their work we also borrow their terminology. Table 1 lists the questions.

Following a first analysis of the three factors, one item has been dropped, as it results in a relatively low Cronbach's alpha (0.64)[iv] for the respective factor as well as contributing to an overall poorer fit for the model (RMSEA = 0.0848 and GFI=0.923). The deleted item is related to the 'conformity' factor; "I avoid cutting corners" (see table 1 for item loadings). We attribute the poor fit to social desirability bias. The resulting Cronbach's alpha for the revised factor is 0.69 which is satisfactory given our sample size and the exploratory nature of this work (Flynn et al. 1994). The Cronbach's alpha score for the other two factors indicate reliable measures: 0.859 for 'attention to detail' and 0.90 for 'creativity', respectively[v].

This resulted in a model with acceptable fit (X^2 = 110.67, d.f.=41, goodness-of-fit index (GFI)=0.94, root-mean-square error of approximation (RMSEA) = 0.0682). While a RMSEA of 0.05 or less would indicate a close fit, our values are still below 0.08 and this indicates a "reasonable error in approximation" (Browne & Cudeck 1992: p. 239). Item loadings are all highly significant ($p<0.001$, cf. table 1).

[Insert table 1 around here]

3.5. Variables: Cognitive abilities

Cognitive abilities (or functions) are assessed with a battery of standardized neurocognitive tests. Neurocognitive testing with standardized administrations and available normative data remains the gold standard for the examination of cognitive functions in both clinical and healthy populations (Casaletto & Heaton 2017). The test battery consisted of six independent tests assessing different cognitive domains. The tests were chosen based on a hierarchical understanding of cognition with lower level cognitive abilities being foundational for higher-level and more complex cognitive processes (Lezak et al., 2012). For example, processing speed and attention are lower-level abilities, which higher-level cognitive processes such as executive functions rely upon. The cognitive domains assessed are processing speed (i.e. the ability to rapidly process information), sustained attention (the ability to sustain cognitive focus), working memory (i.e. the ability to manipulate temporarily stored information), learning and memory (i.e. the ability to learn and retain novel information), long-term memory (i.e. the ability to retain information for a longer period of time), and executive functioning (i.e. the ability to plan, organize, regulate, and monitor goal-directed behaviour (Denckla 1994)). Processing speed was assessed with the Trail Making Test (TMT) – Part A (Reitan 1958) and the Coding subtest of Wechsler's Adult Intelligent Scale - fourth edition (WAIS-IV) (Wechsler 2008). Attention was assessed with the Paced Auditory Serial Addition Test (PASAT) (Wiens et al. 1997); working memory was assessed with the Digit Span subtest of WAIS-IV (Wechsler, 2008); Learning & memory (trial 1 to 5) and long-term memory (30 min delayed trial) was assessed with the RAVLT test (Schmidt 1996); executive functioning was assessed with TMT – part B (Reitan, 1958) and a computerized version of the Wisconsin Card Sorting Test (Heaton 1993). Both tests are widely considered valid measures of executive functions. In addition to these cognitive domains we also calculated a global cognition score based on performance scores from each of the individual cognitive tests.

4. Results

4.1. Cognitive styles and search behavior

As a simple descriptive measure, we report that the average search (Hamming) distance is 2.68 (1.98 SD), i.e. on average players toggle more than two (out of the eight) tiles at a given time. This result is remarkably similar to Billinger et al. (2014), who report an average search distance of 2.65 (1.99 SD) in a game with 10 tiles and 1024 options. Overall out of 7223 categorized search moves (excluding move 1), 52.7% were local, 18.4% directed global, while 28.7% were undirected global.

To explore whether cognitive style (as expressed by the three A-I factors) has an impact on each of the three search strategies[vi], we first analysed the Mturk dataset. Tables 2, 3 and 4 report results from the analysis, for each of these three dependent variables. The first model we analyse is a standard generalized model for independent binomial counts. In the second model, we assume that one possible source of correlations among observations is time and we model time as a random effect. Since the variance from the random effect is rather small (e.g. for undirected global search the estimate is = 0.01069 and the standard error = 0.01384), we estimate a third model, a marginal logistic regression model. To account for potential autocorrelation in the observations, we included a multiplicative over-dispersion parameter in our model.

The same model was fitted for directed and undirected global search. The within-subject association among the vector of responses is modelled by specifying time as an R-side effect for each participant, with a standard compound symmetric structure.

[Insert tables 2, 3 and 4 around here]

The same marginal logistic model is subsequently applied on data collected in the university lab and we are thus able to re-test the results from the Mturk sample.

[Insert table 5 around here]

Table 5 illustrates the lab results which replicate results from the Mturk study. Participants who score high on 'creativity' are less likely to engage in local search and more likely to engage in undirected global search. Note that undirected global search is not necessarily deleterious to problem solving processes, since it can potentially help a solver escape from local optimal solutions, i.e. solutions where no minor changes can lead to an improvement. Additionally, we find that individuals who score low on attention to detail are more likely to do undirected global search moves, while individuals who score high on attention to detail are more likely to do local search. These results seem consistent with the description of 'innovators' (i.e. individuals who score high on creativity) as having a less structured approach to problem solving, while 'adaptors' (i.e. individuals who score low on creativity) prefer a more systematic approach. Indeed, in interviews some respondents equate local search moves with 'systematic search'. Although the first (Mturk) dataset does not reveal the same pattern with respect to the second result, the lab results are supported by the multinomial modelling of the combined dataset (Appendix 1). Results from this analysis show that players who score high on attention to detail are, relative to undirected global search, less likely to do either local search or directed global search. We attribute this discrepancy to the particular nature of our first sample: Mturkers receive ratings according to their performance in a given task, which shapes their chances of being approved for future tasks. For this particular study, we chose a sub-sample of the Mturker population with a high acceptance rate (>99%). We conjecture that either successful Mturkers are simply higher on this dimension or they are more likely to self-report higher levels of attention. Indeed the scores on this factor differ significantly across samples (lab average = 5.1, Mturk average = 5.47, p=0.0013, t-test), a discrepancy which we do not find for our other predictor (creativity).

4.2. Cognitive ability and search behavior

With respect to the relationship between cognitive abilities and search strategies, we first report here, similarly as with cognitive styles, a simple generalized model for each of the three search strategies and search distance (using as before a standard binomial model). Table 6 reports all p values obtained by investigating the potential effect of every cognitive ability on the likelihood of choosing each of the search strategies.

[Insert table 6 around here]

The main results in table 6 show that individuals with strong executive functions and individuals that have a strong global cognition score (an aggregation of all cognitive ability scores), engage in more local search behavior, and less undirected global search behavior. Undirected global search can help an individual escape from a local optimum, but is generally an inefficient form of search. Yet, not only higher level cognitive abilities shape behavior, since sustained attention is strongly related to undirected global search. The less one is able to continuously pay attention to the current task, the less systematic and the more undirected search becomes. The fact that working memory does not play a role, is likely due to the fact that the game interface clearly shows the players' last attempts (see figure 1), which means that an ability to memorize the past few attempts is less useful. In contrast, learning and memory turns out to be important for engaging in search behavior that is aligned with the underlying problem structure (i.e. global directed search). Learning & memory is based on a test that requires the participant to learn and memorize 15 words across 5 trials, while long-term memory requires the participant to retain as many words as possible after a period of 30 minutes. Since the game play lasted about 10 min., it was expected that learning & memory outcomes would be the most meaningful.

We engaged in more in-depth analysis of the relationship between sustained attention, learning and memory, executive functions and global cognition score by subjecting the above outlined relationships to further robustness checks. We employ the same approach as when analysing cognitive styles, taking time and the individual into account. Specifically, we first report a generalized linear model with the cognitive abilities as independent variables (Model 1 in tables 7-9) and subsequently we fit a marginal logistic model (Models 2 and 3 in Table 7 and 2-5 in Tables 8, 9).

[Insert table 7-9 around here]

In general, the relationships are robust and reveal similar significance levels, however, they also paint a more nuanced picture. Thus, while a high global cognition score *is* a good predictor for a diminished likelihood of engaging in local search, it seems that the executive functions component is mainly driving these results. Table 9 (model 5) outlines that feedback is a far stronger predictor for undirected global search than any of the higher or lower level cognitive factors (cf. the estimates). In other words, if one is not doing well, the propensity to engage in global undirected search overwhelms cognitive tendencies, in line with the adaptive search hypothesis. In contrast, Table 8 (model 5) shows that cognitive factors (executive functions) are more important factors for explaining local search, than the performance feedback one receives. The relative impact on behavior of feedback vs. cognitive factors is thus dependent on the kind of search behavior one sets out to predict.

5. Discussion

A central question in strategy research is how to switch from local, adaptive search to global search that can facilitate innovation. Based on a NK environment, mindless feedback-based search algorithms can be effective. In order to be able to capture the role of cognitive heterogeneity, we rely on a search framework where such an adaptive search mechanism is relatively ineffective (Vuculescu 2017). We integrate insights from cognitive psychology to identify both lower and higher level building blocks of complex problem solving abilities (Lezak et al. 2012), as well as measurements of cognitive style (Kirton 1976). To our knowledge, we are the first to investigate the individual cognitive antecedents of search behavior in a rugged, combinatorial landscape that also provides a quantitative, empirical basis for the cognitive representations Gavetti & Levinthal (2000) argue facilitate efficient search. This enables us to study the cognitive microfoundations of search, providing insight into the relative importance of individual antecedents and a more fine-grained understanding of how to shape individuals' search processes.

We present three interrelated findings and their associated theoretical contributions. First, we document substantial variation in search behavior dependent on the cognitive style of how one processes information. To illustrate, participants who score high on creativity (vs. those that score low) are about 40% more likely to engage in undirected global search than local (45%) or directed global search (38%), cf. table 2 in appendix 1. Second, to our knowledge we are the first to show a link between specific building blocks of cognitive ability and propensities to search in a rugged landscape. Individuals with low abilities to sustain attention, strong executive functions and overall high cognition scores carry out less undirected global search, engaging in more 'systematic' local search. These two findings have four important theoretical implications.

First, we contribute to research on the role of microfoundations in explaining the adaptive nature of search. In a controlled, combinatorial landscape we identify that it is not merely external performance feedback that shapes search behavior, since cognitive factors appear as important a variable as the feedback received from the environment. In other words, an unexplained residual in a statistical model based on performance feedback is not just noise, but can be partially attributed to the heterogeneity of cognitive antecedents, providing further empirical support for the microfoundational insistence on understanding how individuals shape macro-level behavior (Felin et al. 2015).

Second, in addition to shedding light on the relative importance of microfoundations, the above outlined empirical findings on the heterogeneity of search contribute to a discussion about the assumptions that agent-based models in management rely on when investigating how to optimize search (March 1991; Levinthal 1997; Puranam et al. 2015). While individuals in general might tend to stop local search prematurely (Billinger et al. 2014), we are the first to identify that individuals with higher cognitive abilities or certain cognitive styles constitute more persistent (local) searchers, more akin to the assumptions in the seminal NK model (Levinthal 1997). For example, those with strong executive functions and ability to maintain sustained attention are less sensitive to immediate negative feedback, and thus seem to engage in a more long-term search strategy. Put differently, these individuals are less likely to do undirected global search, a kind of search that might be inefficient depending on search costs and the size of the landscape (Heck, Vuculescu et al. 2018) or the complexity of the search tasks (Billinger et al. 2014). Granted, assumptions about individual behavior in simulation models have to simplify reality to some degree (Davis et al. 2007; Knudsen et al. 2019; Puranam et al. 2015) and Billinger et al. (2014) provided support for the generic adaptive search assumption. However, we offer further, valuable insight on relevant, cognitive antecedents of the heterogeneity of search that would likely impact simulation results on optimization of search. For example, a collective of cognitive diverse solvers are likely to be more efficient than the homogeneous searchers simulations usually rely on, since more diverse search behaviors would be represented. We consider this an opportunity to compare simulations that rely on different assumptions, by contrasting the rationalized, adaptive searcher with more empirically informed parameters (cf. Puranam et al. 2015; Smith & Rand 2018).

Third, we contribute to research on how managers can optimize the organizational context of search. Incentives can of course influence search behavior (Ederer & Manso 2013), yet managerial interventions can also shape what type of search individuals engage in. If systematic, local search is the goal, then factors that impair the ability to engage in sustained attention and executive functions should be considered. For one, *some* cognitive abilities can be improved by training, e.g. ability to shift or maintain attention, Chan et al. (2019). Furthermore, how the search process is organized can also alleviate interruptions, shaping not only the performance but potentially also the kind of search behavior being conducted. Finally, recent studies indicate that sustained attention and executive functions can be shaped by contextual factors (Chan et al. 2019): Sleep (Lowe et al. 2017) and physical activity (Radel et al. 2018) can benefit an individual's ability to sustain attention and engage their executive functions, further influencing their likelihood to generate new, good ideas (Gish et al. 2019). Insight into cognition thus opens up opportunities for organizations to not merely manage search directly (e.g. via incentives), but also indirectly by realizing the importance of certain cognitive factors. Individual search behavior thus turns out be shaped by a complex interplay between external performance feedback, cognitive antecedents and external factors shaping how individual cognition unfolds. Overall, we speculate that this more persistent search behavior could constitute one path through which the stronger cognitive ability leads to superior individual performance outcomes, as macro-level data clearly documents (Ritchie 2015; Hunter & Schmidt 1998).

Fourth, our study also contributes to an emerging stream of research on matching the governance of search with the type of problem one is facing (Afuah & Tucci 2012; Felin & Zenger 2014; Lakhani et al. 2013). Organizations like InnoCentive or TopCoder that rely on crowds to generate relevant perspectives on how to solve problems (Boudreau & Lakhani 2013) could, based on a relatively short survey (the A-I scale) or established cognitive ability tests, map how employees or crowds are likely to navigate a given search space. For example, individuals categorized as creative would engage in more undirected global search, while those with strong executive functions would engage in the aforementioned, persistent local search. Based on this insight, one can either match the crowd to the given problem, or decompose and formulate the problem differently (von Krogh et al. 2013) in order to create a better match with the kind of search one can expect to obtain.

The third main finding is that individuals who score high on the variable ‘learning and memory’ seem to grasp the underlying problem structure, a result that an unstructured problem task does not allow. While we remind the reader about the nature of the exploratory analysis, we still consider it a proof-of-concept of the experimental task. We thus provide an empirical foothold for studying how individuals search a combinatorial, rugged, non-random landscape (Csaszar & Levinthal 2016; Gavetti & Levinthal 2000). This constitutes a methodological contribution to studying directed global search and allows more realistic simulations, which fit how people rely on models to navigate the world. Computer science grapples with a related challenge, when developing machine-learning based algorithms able to search – for the algorithm – unknown landscapes. Recent successful solutions have relied on a model-free approach, being able to beat humans in relatively complex games such as Chess, Go and Dota (Guez et al. 2019; Silver et al. 2018). However, a different perspective argues that the algorithms should be inspired by the human ability to create relevant causal, mental models and extrapolate from small samples (Marcus 2018). In any case, this endeavor might benefit from the field of management’s insight into how to identify directed search behavior and organize search in unknown solution spaces.

6. Conclusion and future work

We elaborate on how cognition shapes not just overall search performance or a generic exploration vs. exploitation tendency, but offer a more fine-grained theory on how particular cognitive antecedents shape individual search in rugged landscapes (cf. O’Doherty et al. 2017; Helfat & Peteraf 2015). We propose that the insights provided in this study can help pave the way for developing further experimental tasks and a simulation framework that allows directed global search for innovation (Gavetti & Levinthal 2000). The setup also points to the opportunity to move beyond homogeneous, rationalized assumptions about search behavior and integrate empirically informed insights into simulation models (Smith & Rand 2018). In order to not solely rely on this particular search framework, future studies could vary the complexity of search frameworks (Levinthal 1997), consider how benign the search environment is (MacAulay et al. 2017) as well as draw upon less artificial problems (cf.

Heck, Vuculescu et al. 2018). Furthermore, in order to isolate the mechanism of interest we excluded social context. Yet, organizational problems are often solved in a social context and future studies could develop further insight into how individuals engage in model-based social learning. Finally, while a well-defined task such as the Alien Game or physics challenges (Heck, Vuculescu et al. 2018) imply the benefit of being able to track search behavior, one could also develop how to identify search behaviors in real-world organizational datasets, such as the online coding platform TopCoder (Boudreau & Lakhani 2013).

Figure 1. Screenshot of the game

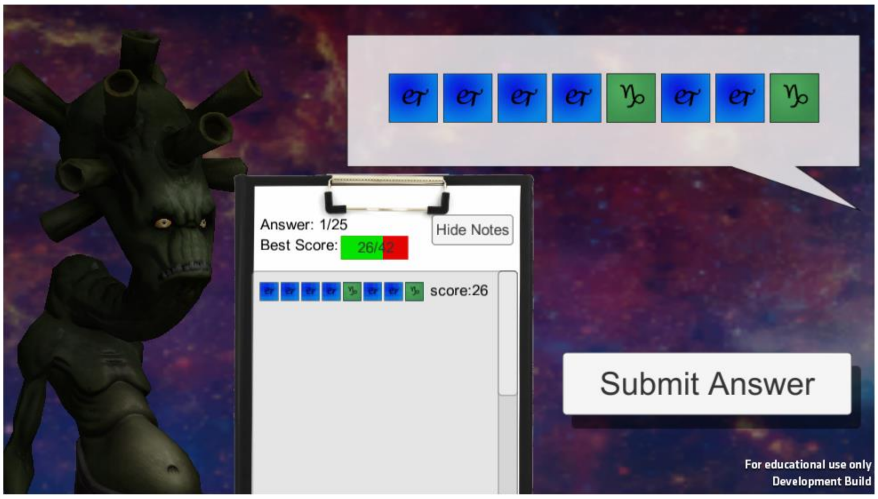

Table 1. Item loadings for the KAI 3-Factor model

Factor Loading Matrix:	Estimate (StdErr) conformity	detail	creativity
I try not to oppose team members	0.658*** (0.86)		
I adapt myself to the system	0.669*** (0.064)		
I adhere to accepted rules in my area of work	0.876*** (0.064)		
Thorough when solving problems		0.790*** (0.05)	
Addresses small details needed to perform the task		0.866*** (0.052)	
Performs the task precisely over a long time		1.004*** (0.054)	
Good in tasks that require dealing with details		0.941*** (0.058)	
I have a lot of creative ideas			1.268*** (0.064)
I prefer tasks that enable me to think creatively			1.370*** (0.063)
Innovative			1.249*** (0.06)
I like to do things in an original way			1.019*** (0.061)

N=336, *** p<0.001

Table 2. Effect of the three A-I factors on the likelihood of doing undirected global search (Mturk study)

Response variable = undirected global search (binomial)

		Estimate (StdErr)	p value	-2LL	Cov param (StdErr)
Model 1	creativity(ref=low)	-0.289 (0.095)	0.0024	4332.03	
	detail(ref=low)	0.100 (0.111)	0.3683	4331.21	
	conformity(ref=low)	0.152 (0.10)	0.1582	4328.04	
Model 2	creativity(ref=low)	-0.289 (0.095)	0.0024		0.01 (0.002)
	conformity(ref=low)	0.153 (0.11)	0.1557		0.01 (0.014)
	detail(ref=low)	-0.062 (0.097)	0.5207		0.01 (0.014)
Model 3	creativity(ref=low	-0.289 (0.095)	0.0024		1.00 (0.02)
	conformity(ref=low)	0.152 (0.11)	0.1080		0.99 (0.02)
	detail(ref=low)	0.059 (0.114)	0.6079		0.99 (0.02)

N=242

Table 3. Effect of the three A-I factors on the likelihood of local search (Mturk study)

Response variable = local search (binomial)

		Estimate (StdErr)	p value	-2LL	Cov param (StdErr)
Model 1	creativity(ref=low)	0.212 (0.072)	0.0032	6960.02	
	detail(ref=low)	0.105 (0.071)	0.1383	6966.45	
	conformity(ref=low)	-0.038 (0.076)	0.6198	6968.39	
Model 2	creativity(ref=low)	0.213 (0.072)	0.0033		0.0005 (0.001)
	detail(ref=low)	0.100 (0.071)	0.1568		0.0005 (0.001)
	conformity(ref=low)	-0.044 (0.076)	0.5634		0.0005 (0.001)
Model 3	creativity(ref=low)	0.213 (0.072)	0.0033		1.000 (0.02)
	detail(ref=low)	0.007 (0.081)	0.9311		1.001 (0.02)
	conformity(ref=low)	-0.08 (0.078)	0.3009		1.001 (0.02)

N=242

Table 4. Effect of the three A-I factors on the likelihood of directed global search (Mturk study)

Response variable = directed global search (binomial)

		Estimate (StdErr)	p-value	-2LL	Cov par (StdErr)
Model 1	creativity(ref=low)	-0.079 (0.083)	0.3377	5716.15	
	detail(ref=low)	-0.086 (0.080)	0.2825	5715.92	
	conformity(ref=low)	-0.049 (0.086)	0.5641	5716.74	
Model 2	creativity(ref=low)	-0.0784 (0.083)	0.3449		0.001 (0.002)
	detail(ref=low)	-0.08 (0.081)	0.3237		0.001 (0.002)
	conformity(ref=low)	-0.041 (0.086)	0.6375		0.001 (0.002)
Model 3	creativity(ref=low)	-0.079 (0.082)	0.3377		1.0004 (0.02)
	detail(ref=low)	-0.086 (0.08)	0.2825		1.0004 (0.02)
	conformity(ref=low)	-0.049 (0.086)	0.5642		1.0004 (0.02)

N=242

Table 5. Effect of the three A-I factors on the likelihood of the three search strategies: (Lab study, Marginal logistic model)

Type of search behaviour/KAI Factor	Local search	Directed global search	Undirected global search
Creativity (Low)			
estimate	0.2444*	0.2732	-0.3660***
standard deviation	0.1069	0.1691	0.1087
p value	0.0223	0.1063	0.0008
Detail(low)			
estimate	-0.3542***	0.2648	0.3574**
standard deviation	0.09837	0.1524	0.1329
p value	0.0003	0.0825	0.0072
Conformity (low)			
estimate	0.02859	0.07902	-0.08792
standard deviation	0.1014	0.1532	0.1056
p value	0.7781	0.6061	0.4052

N=94, * p<0.05, **p<0.01, *** p<0.001

Table 6. Cognitive abilities and search strategies

Cognitive ability	Directed global	Local search	Undirected global
Processing speed	0.55	0.058	0.22
Sustained attention	0.21	0.075	0.009-
Working memory	0.64	0.616	0.713
Learning & memory	0.018+	0.87	0.282
Long-term memory	0.267	0.807	0.500
Executive functions	0.607	0.0004+	0.004-
Global cognition score	0.231	0.032+	0.024-

N=39, p values

Marginally significant values (<0.10) are highlighted in light grey, while statistically significant values (<0.05) are highlighted in dark grey.
- Negatively correlated, i.e. a significant p value indicates that a higher cognitive ability score leads to a lower likelihood of engaging in the particular search behavior.
+ Positively correlated, i.e. a significant p value indicates that a higher cognitive ability score leads to a higher likelihood of engaging in the particular search behavior.

Table 7. Effect of cognitive abilities on the likelihood of directed global search

Response variable = directed global search (binomial)

		Estimate (StdErr)	p value	-2LL	Cov param (StdErr)
Model 1	Learning memory	0.259 (0.1094)	0.0180	661.088	
Model 2	Learning memory	0.259 (0.1097)	0.0186		10.052 (0.0531)
Model 3	Learning memory	0.253 (0.1098)	0.0213		10.074 (0.0533)
	Feedback (ref = 0)	-0.473 (0.2079)	0.0233		

N=39

Table 8. Effect of cognitive abilities on the likelihood of local search

Response variable = local search (binomial)

		Estimate (StdErr)	p value	-2LL	Cov param (StdErr)
Model 1	Global cognition score	0.351 (0.163)	0.0320	910.23	
Model 2	Global cognition score	0.354 (0.163)	0.0297		0.096 (0.053)
Model 3	Executive functions	0.469 (0.134)	0.0005		0.098 (0.074)
Model 4	Global cognition score	-0.333 (0.268)	0.2150		0.097 (0.074)
	Executive functions	0.674 (0.215)	0.0018		
Model 5	Executive functions	0.470 (0.135)	0.0005		0.085 (0.701)
	Feedback (ref = 0)	-0.335 (0.180)	0.0639		

N=39

Table 9. Effect of cognitive abilities on the likelihood of undirected global search

Response variable = undirected global search (binomial)

		Estimate (StdErr)	p value	-2LL	Cov param (StdErr)
Model 1	Global cognition score	-0.344 (0.152)	0.0243	965.09	
Model 2	Global cognition score	-0.348 (0.152)	0.0224		0.106 (0.0731)
Model 3	Executive functions	-0.309 (0.111)	0.0057		0.104 (0.0727)
Model 4	Sustained attention	-0.210 (0.083)	0.0121		0.106 (0.0730)
Model 5	Global cognition score	0.324 (0.327)	0.3226		
	Executive functions	-0.383 (0.196)	0.0507		0.012 (0.048)
	Sustained attention	-0.191 (0.119)	0.1080		
	Feedback (ref = 0)	0.943 (0.186)	<0.0001		

N=39

References

Afuah, A., & Tucci, C. L. (2012). Crowdsourcing as a Solution to Distant Search. Academy of Management Review, 37(3), 355-375.
Bahrami, B., Olesen, K., Latham, P. E., Roepstorff, A., Rees, G., & Frith, C. D. (2010). Optimally Interacting Minds. Science, 329, 1081-1085.
Baumann, O., Schmidt, J., & Stieglitz, N. (2019). Effective Search in Rugged Performance Landscapes: A Review and Outlook. Journal of Management, 45(1), 285-318.
Benito-Ostolaza, J. M., Hernandez, P., & Sanchis, L. (2016). Do individuals with higher cognitive ability play more strategically? Journal of Behavioral and Experimental Economics, 64, 5-11.

Berinsky, A. J., Huber, G. A., & Lenz, G. S. (2012). Evaluating online labor markets for experimental research: Amazon. com's Mechanical Turk. Political Analysis, 20(3), 351-368.
Billinger, S., Stieglitz, N., & Schumacher, T. R. (2014). Search on Rugged Landscapes: An Experimental Study. Organization Science, 25(1), 93-108.
Boudreau, K. J., & Lakhani, K. R. (2013). Using the Crowd as an Innovation Partner. Harvard Business Review, April.
Browne, M. W., & Cudeck, R. (1992). Alternative ways of assessing model fit. Sociological Methods & Research, 21(2), 230-258.
Carnabuci, G., & Dioszegi, B. (2015). Social networks, cognitive style, and innovative performance: A contingency perspective. Academy of Management Journal, 58(3), 881-905.
Casaletto, K. B., & Heaton, R. K. (2017). Neuropsychological Assessment: Past and Future. Journal of the International Neuropsychological Society, 23(9-10), 778-790.
Chan, T., Wang, I., & Ybarra, O. (2019). Leading and managing the workplace: The role of executive functions. Academy of Management Perspectives, In press.
Cohen, J. D., McClure, S. M., & Yu, A. J. (2007). Should I stay or should I go? How the human brain manages the trade-off between exploitation and exploration. Philosophical Transactions B, 362, 933-942.
Csaszar, F. A., & Levinthal, D. A. (2016). Mental representation and the discovery of new strategies. Strategic Management Journal, 37(10), 2031-2049.
Dams-O'Connor, K., & Gordon, W. A. (2013). Integrating Interventions after Traumatic Brain Injury: A Synergistic Approach to Neurorehabilitation. Brain Impairment, 14(1), 51-62.
Davis, J. P., Eisenhardt, K. M., & Bingham, C. B. (2007). Developing Theory Through Simulation Methods. Academy of Management Review, 32(2), 480-499.
Denckla, M. B. (1994). Measurement of executive function. In G. R. Lyon (Ed.), Frames of reference for the assessment of learning disabilities: New views on measurement issues. Baltimore, MD: Paul H Brookes Publishing Co.
Doll, B., Bath, K., Daw, N. D., & Frank, M. (2016). Variability in dopamine genes dissociates model-based and model-free reinforcement learning. Journal of Neuroscience, 36(4), 1211-1222.
Doll, B., Simon, D. A., & Daw, N. D. (2012). The ubiquity of model-based reinforcement learning. Current Opinion in Neurobiology, 22, 1075-1081.
Døjbak, D. H., Eskildsen, J. K., Argote, l., Mønster, D., Burton, R. M., & Obel, B. (2015). Exploration versus exploitation: Emotions and performance as antecedents and consequences of team decisions. Strategic Management Journal, Early View.
Ederer, F., & Manso, G. (2013). Is Pay for Performance Detrimental to Innovation? Management Science, 59(7), 1496-1513.
Felin, T., Foss, N. J., & Ployhart, R. E. (2015). The Microfoundations Movement in Strategy and Organization Theory. The Academy of Management Annals, 9(1), 575-632.
Felin, T., & Zenger, T. R. (2014). Closed or open innovation? Problem solving and the governance choice. Resarch Policy, 43(5), 914-925.
Finkelstein, S., Hambrick, D. C., & Cannella, A. (2009). Strategic Leadership: Theory and Research on Executives, Top Management Teams, and Boards. New York: Oxford University Press.
Fleming, L. (2002). Finding the organizational sources of technological breakthroughs: the story of Hewlett-Packard's thermal ink-jet. Industrial and Corporate Change, 11(5), 1059-1084.
Flynn, B. B., Schroeder, R. g., & Sakakibara, S. (1994). A framework for quality management research and associated measurement instrument. Journal of Operations Management, 11(4), 339-366.
Frederick, S. (2005). Cognitive Reflection and Decision Making. Journal of Economic Perspectives, 19(4), 25-42.
Gabora, L., & Kauffman, S. (2016). Toward an Evolutionary-Predictive Foundation for Creativity. Psychonomic Bulletin & Review, Forthcoming.
Gary, M. S., & Wood, R. (2011). Mental models, decision rules, and performance heterogeneity. Strategic Management Journal, 32, 569-594.
Gavetti, G., & Levinthal, D. (2000). Looking Forward and Looking Backward: Cognitive and Experiental Search. Administrative Science Quarterly, 45(1), 113-137.
Gignac, G. E. (2015). Raven's is not a pure measure of general intelligence: Implications for g factor theory and the brief measurement of g. Intelligence, 52, 71-79.

Gish, J., Wagner, D. T., Gréggoire, D. A., & Barnes, C. M. (2019). Sleep and entrepreneurs' abilities to imagine and form initial beliefs about new venture ideas. Journal of Business Venturing, Published online.
Greene, J. C., Caracelli, V. J., & Graham, W. F. (1989). Toward a Conceptual Framework for Mixed-Method Evaluation Designs. Educational Evaluation and Policy Analysis, 11(3), 255-274.
Greve, H. R. (2003). A behavioral theory of R&D expenditures and innovations: Evidence from shipbuilding. Academy of Management Journal, 46(6), 685-702.
Guez, A., Mirza, M., Gregor, K., Kabra, R., Racaniere, S., Weber, T., . . . Lillicrap, T. (2019). An investigation of model-free planning. arXiv.
Hambrick, D. C., & Mason, P. A. (1984). Upper Echelons: The Organization as a Reflection of Its Top Managers. Academy of Management Review, 9(2), 193-206.
Hauser, D. J., Paolacci, G., & Chandler, J. J. (2019). Common concerns with MTurk as a participant pool: Evidence and solutions. In F. R. Kardes, P. P. Herr, & N. Schwarz (Eds.), Handbook of Research Methods in Consumer Psychology.
Hayes, J., & Allinson, C. W. (1998). Cognitive Style and the Theory and Practice of Individual and Collective Learning in Organizations. Human Relations, 51(7), 847-871.
He, Z.-L., & Wong, P.-K. (2004). Exploration vs. Exploitation: An Empirical Test of the Ambidexterity Hypothesis. Organization Science, 15(4), 481-494.
Heaton, R. K. (1993). Wisconsin Card Sortin Test: Computer version-2 research edition. Odessa, FL: Psychological Assessment Resources.
Heck, R., Vuculescu, O., Bergenholtz, C., Sherson, J. F., & al., E. (2018). Remote optimization of an ultra-cold atoms experiment by experts and citizen scientists. PNAS.
Helfat, C. E., & Peteraf, M. A. (2015). Managerial cognitive capabilities and the microfoundations of dynamic capabilities. Strategic Management Journal, 36, 831-850.
Hills, T. T., & Hertwig, R. (2010). Information Search in Decisions from Experience: Do our Patterns of Sampling Foreshadow our Decisions? Psychological Science, 21(12), 1787-1792.
Kaiser, S., Simon, J., Kalis, A., Schweizer, S., Tobler, P., & Mojzisch, A. (2013). The cognitive and neural basis of option generation and subsequent choice. Cogn Affect Behav Neurosci, 13(4), 814-829.
Kauffman, S. (1993). The origins of order: self-organization and selection in evolution. New York: Oxford University Press.
Kirton, M. (1976). Adaptors and innovators: A description and measure. Journal of Applied Psychology, 61(5), 622-629.
Kirton, M. (2003). Adaptation and innovation in the context of diversity and change. London: Routledge.
Knudsen, T., Levinthal, D. A., & Puranam, P. (2019). Editorial: A Model Is a Model. Strategy Science, 4(1), 1-3.
Lakhani, K. R., Lifshitz, H. A., & Tushman, M. L. (2013). Open Innovation and Organizational Boundaries: Task Decomposition, Knowledge Distribution and the Locus of Innovation. In A. Grandori (Ed.), Handbook of Economic Organization: Integrating Economic and Organization Theory (pp. 355-382). Northampton, MA: Edward Elgar Publishing.
Laureiro-Martínez, D., & Brusoni, S. (2018). Cognitive flexibility and adaptive decision-making: Evidence from a laboratory study of expert decision makers. Strategic Management Journal, 39, 1031-1058.
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.
Laureiro-Martínez, D., Brusoni, S., & Zollo, M. (2010). The Neuroscientific Foundations of the Exploration-Exploitation Dilemma. Journal of Neuroscience, Psychology and Economics, 3(2), 95-115.
Levine, S. S., Bernard, M., & Nagel, R. (2017). Strategic Intelligence: The Cognitive Capability to Anticipate Competitor Behavior. Strategic Management Journal, 38(2390-2423).
Levinthal, D. A. (1997). Adaption on Rugged Landscapes. Management Science, 43(7), 934-950.
Lezak, M. D., Howieson, D. B., Bigler, E. D., & Tranel, D. (2012). Neuropsychological assessment (5th ed.). New York, NY: Oxford University Press.

Li, Q., Maggitti, P. G., Smith, K. G., Tesluk, P. E., & Katila, R. (2013). Top management attention to innovation: The role of search selection and intensity in new product introductions. Academy of Management Journal, 56(3), 893-916.
Lowe, C., Safati, A., & Hall, P. (2017). The neurocognitive consequences of sleep restriction: A meta-analytic review. Neurosci Biobeh Rev, 80, 586-604.
Luria, A. R. (1976). The working brain: an introduction to neuropsychology. USA: Basic Books.
MacAulay, S., Steen, J., & Kastelle, T. (2017). The search environment is not (always) benign: reassessing the risks of organizational search. Industrial and Corporate Change, 1-23.
March, J. G. (1991). Exploration and Exploitation in Organizational Learning. Organization Science, 2(1), 71-87.
Marcus, G. (2018). Deep Learning: A Critical Appraisal.
Miron-Spektor, E., Erez, M., & Naveh, E. (2011). The Effect of Conformist and Attentive-To-Detail Members on Team Innovation: Reconciling the Innovation Paradox. Academy of Management Journal, 54(4), 740-760.
Miron, E., Erez, M., & Naveh, E. (2004). Do personal characteristics and cultural values that promote innovation, quality, and efficiency compete or complement each other? Journal of Organizational Behavior, 25, 175-199.
Murray, J., Jaramillo, J., & Wang, X. (2017). Working Memory and Decision-Making in a Frontoparietal Circuit Model. Journal of Neuroscience, 37(50), 12167-12186.
Nigam, A., Huising, R., & Golden, B. (2016). Explaining the Selection of Routines for Change during Organizational Search. Administrative Science Quarterly, 61(4), 551-583.
Nowak, M. A., & Sigmund, K. (1993). A strategy of win-stay, lose-shift that outperforms tit-for-tat in the Prisoner's Dilemma game. Nature, 364(6432).
O'Doherty, J. P., Cockburn, J., & Pauli, W. M. (2017). Learning, Reward, and Decision Making. Annu. Rev. Psychol., 68, 73-100.
Puranam, P., Stieglitz, N., Osman, M., & Pillutla, M. M. (2015). Modelling Bounded Rationality in Organizations: Progress and Prospects. Academy of Management Annals, 9(1), 337-392.
Radel, R., Tempest, G., & Brisswalter, J. (2018). The long and winding road: Effects of exercise intensity and type upon sustained attention. Physiol Behav, 195, 82-89.
Raven, J., Raven, J. C., & Court, J. H. (2003). Manual for Raven's Progressive Matrices and Vocabulary Scales. San Antonio, TX: Harcourt Assessment.
Reitan, R. M. (1958). Validity of the Trail Making Test as an indicator of organic brain damage. Percept Mot Skills, 8, 271-276.
Reypens, C., & Levine, S. S. (2018). Behavior in Behavioral Strategy: Capturing, Measuring, Analyzing. Behavior in Behavioral Strategy, 221-246.
Ritchie, S. (2015). Intelligence: All that matters. London: Hodder & Stoughton.
Schmidt, F. L., & Hunter, J. E. (1998). The validity and utility of selection methods in personnel psychology: Practical and theoretical implications of 85 years of research findings. Psychological Bulletin, 124(2), 262-274.
Schmidt, M. (1996). Auditory and Verbal Learning Test. A handbook. Los Angeles: Western Psychological Services.
Schulz, E., Bhui, R., Love, B. C., Brier, B., Todd, M. T., & Gershman, S. J. (2019). Structured, uncertainty-driven exploration in real-world consumer choice. PNAS, 116(28), 13903-13908.
Silver, D., Hubert, T., Schrittwieser, J., Antonoglou, I., Lai, M., Guez, A., . . . Hassabis, D. (2018). A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play. Science, 362(6419), 1140-1144. doi:10.1126/science.aar6404
Smaldino, P. E., & Richerson, P. J. (2012). The origins of options. frontiers in Neuroscience, 6(50), 1-11.
Smith, E. B., & Rand, W. (2018). Simulating Macro-Level Effects from Micro-Level Observations. Management Science, 64(11), 4967-5460.
Sternberg, R. J., & Grigorenko, E. (1997). Are cognitive styles still in style? American Psychologist, 52, 700-712.
Steyvers, M., Lee, M. D., & Wagenmakers, E.-J. (2009). A Bayesian analysis of human decision-making on bandit problems. Journal of Mathematical Psychology, 53, 168-179.

von Krogh, G., Wallin, M. W., & Sieg, J. H. (2013). A problem in becoming: How firms formulate sharable problems for innovation contests. Zurich.
Vuculescu, O. (2017). Searching far away from the lamp-post: An agent-based model. Strategic Organization, 15(2), 242-263.
Watson, R. A., & Pollack, B. (1999). Hierarchically consistent test problems for genetic algorithms. Paper presented at the Proceedings of the 1999 Congress on Evolutionary Computation (CEC), Washington DC, USA.
Wechsler, D. (2008). Wechsler Adult Intelligence Scale-Fourth Edition. San Antonio, TX: Pearson.
Wiens, A. N., Fuller, K. H., & Crossen, J. R. (1997). Paced Auditory Serial Addition Test: adult norms and moderator variables. J Clin Exp Neuropsychol, 19, 473-483.
Wilson, R. C., Geana, A., White, J. M., Ludvig, E. A., & Cohen, J. D. (2014). Humans Use Directed and Random Exploration to Solve the Explore-Exploit Dilemma. J Exp Psych Gen, 143(6), 2074-2081.
Witkin, H. A., & Goodenough, D. R. (1977). Field dependence revisited. ETS Research Bulletin Series, 2(1-53).
Yamakawa, Y., Yang, H., & Lin, Z. J. (2011). Exploration versus exploitation in alliance portfolio: Performance implications of organizational, strategic, and environmental fit. Research Policy, 40(2), 287-296.

Appendix 1: Multinomial modelling of dependent variable. Robustness check

In this approach, the dependent variable has three possible outcomes corresponding to the three possible search behaviours: local, directed and undirected global search. We use the third strategy (undirected global) as the reference category, because we consider the distinction between local and directed global on the one hand and undirected global search on the other to be more meaningful given our theoretical framework. In addition, even though local search moves can be both model-free and model-informed, players often describe local search as "systematic", in interviews carried out after game play. We thus expect to find a significant difference between undirected global search and the two other categories. The predictor variables are the same as in the marginal logistic regression model.

Table I. Frequencies of the response variable categories. Pooled data (N=336)

Ordered Value	Strategy	Total frequency
Directed global	1	1336
Local search	2	3809
Undirected global	3	2076

In modelling category probabilities, strategy='3' serves as the reference category.

In the multinomial model, the estimate for the parameter can be identified compared to the baseline category. We further introduce time as a fixed effect and model subject variance as a random effect. Thus the equation is:

$$log\left\{\frac{\pi_{ijr}}{\pi_{ij1}}\right\} = ß_1 + ß_2 creativity_i + ß_3 conformity_i + ß_4 detail_i + ß_5 time + b_i + e_{ij}$$

$\pi_{ijr} = P(Y_{ij} = r)$ are the response probabilities for individual *i* to choose strategy *r* at time *j*,. The influence of the covariates is assessed, as before, through the coefficients $ß_i$. The random effect b_i is assumed to have a univariate normal distribution with zero mean and compound symmetric covariance matrix.

Table II. Relative effect of cognitive style on the odds of choosing one of the three strategies. Results from multinomial regression. Pooled data (N=336)

creativity (ref=LOW)	Estimate	StdErr	P value	Odds Ratio	95% Conf Limits	
Directed global (ref=Undir global)	0.3275	0.1098	0.0029	1.387	1.119	1.721
Local search (ref=Undir global)	0.3774	0.147	0.0103	1.458	1.093	1.945
detail (ref=LOW)						
Directed global (ref=Undir global)	-0.3251	0.1184	0.0061	0.722	0.573	0.911
Local search (ref=Undir global)	-0.3197	0.1566	0.0412	0.726	0.534	0.987
conformity (ref=LOW)						
Directed global (ref=Undir global)	-0.181	0.1076	0.0926	0.834	0.676	1.03
Local search (ref=Undir global)	-0.07779	0.144	0.5892	0.925	0.698	1.227

This alternative modelling strategy serves as a robustness check for the quantitative results presented in the main text. They show that players who score high on ’creativity’ would be more likely to engage in undirected global than local or directed global search (45% and 38%) and likewise, players who score high on ’attention to detail’ will be more likely to engage in directed global or local search than undirected global. Although we acknowledge that this model formulation a) makes unwarranted assumptions regarding the fact that players have a stable and intransitive preference structure for the three search strategies and b) that a marginal rather than a mixed model is more meaningful given our dataset and this generates relatively larger estimates (cf. Fitzmaurice, Laird et al. 2012) we think these results further support our overall findings.

Appendix 2: Insights from interviews

Method and data

Following lab sessions, we randomly selected 40 participants to take part in a short (on average 7.5 minutes) semi-structured interview. According to Greene’s et al.’s (1989) conceptualization of different mixed methods approaches, we rely on a complementary approach, aiming to “increase the interpretability, meaningfulness and validity of constructs and inquiry results…” (Greene et al. 1989: p.259). We thus explore how players create mental models of the solution space and aim to validate that solving the experimental task is meaningful for the players in order to make sure they did not perceive themselves to be stumbling around in a random solution space. Furthermore, we establish face validity of the search behaviour constructs, by matching the quantitative coding with player accounts. Half of these interviews are collected with the participants sitting in front of their game play history, which allowed us to ask specific questions regarding various submissions and mitigates recall bias. By allowing participants to refer to solutions they have tried out as well as transitions from one submission to the next, we are able to get a better grasp of what influences their search behaviours. The interviews are semi-structured and contain five main questions which address: i) overall search strategies, ii) how players switch strategies, iii) what information players sample, iv) how they try to mentally represent the problem and v) changes in such mental representation. The interviews were collected by the two authors, who both individually and collaboratively listened to and analysed the interviews, in order to extract main patterns (Miles & Huberman 1994).

Analysis of qualitative interviews

We explored if players actually reflect upon and follow any of the search strategies that we argue to have identified in the quantitative coding of game plays. We compared the database with the actual accounts of the participants and found no large discrepancies between the coding and the players' explanations, thus finding qualitative support for the coding. When verbalizing their search strategies, very few participants make a two-fold distinction between local and distant search (cf. the traditional exploration vs. exploitation division). Furthermore, what respondents characterize as undirected global search fairly seldom was random; comparing interview responses with actual game-play behaviour reveals that undirected global search usually would be 'systematic' in some sense, e.g. all green [1,1,1,1,1,1,1,1] or all blue [0, 0, 0, 0, 0, 0, 0, 0], or alternating green-blue [0,1,0,1,0,1,0,1]. A participant e.g. reports he "randomly tap[ped] whatever", but his actual game-play in that particular situation was [0,1,0,1,0,1,0,1]. In any case, what they usually describe as a random submission is a solution that is not anchored in the feedback they have received so far. One respondent did acknowledge the difficulty of truly random behaviour, and actually "looked away from the screen" to take 4-5 guesses in order to diversify his search path. The reference points that players rely on are usually the last submitted attempt or their best attempt so far: "What I did was mostly based on the immediately preceding one...but if my score goes down too much I went back to the notes [the virtual clipboard, cf. figure 1]…".

Most of our respondent formed articulate mental models of the underlying problem, without necessarily revealing 'the logic' of the game. For instance, some players report testing typical priors such as all green or all blue options. Some had very sophisticated (albeit wrong) mental models, e.g. "So each square is a letter, spelling: 'I am smart'". We also encountered more abstract representations, such as the solution is "a Fibonacci sequence". We note that at least one player correctly inferred that the solution would have to be (inversely) symmetrical and managed to solve the game in 11 moves. A couple of players realize that the game had many (256) combinations and thus the entire search space could not be covered in the limited number of attempts they had (25) so they tailored their strategies accordingly, e.g. avoiding solely engaging in local-search strategies.

Another distinction is worth emphasizing: Some of these mental representations are clearly top-down: "...such games usually involve a structured solution", but others were feedback based as the following passages from different interviewees highlight: "...at the seventh attempt I noticed", "...I [realized] you can't have too many tiles in a row green", "...there should be four of each". The degree of flexibility with respect to these representations also varies since a number of players continuously develop and adapt a model, e.g. keep the first five tiles constant and adapt the last three or focusing on that four should be green, and the rest blue. Therefore, the behaviour is not merely based on the current condition and the received feedback, but an overall idea and mental representation of what the core pattern of the game could be. This difference is difficult to capture in the complex settings of organizational strategy-making where it is not immediately obvious whether mental models reflect acquired experience or prior expectations with respect to the environment.

Overall, players followed the adaptive rational model outlined in Puranam et al. (2015); they generally created representations of the task, responded to the feedback they received based on their actions, and either maintained or changed their former representation, dependent on the feedback. However, in contrast to typical simulation models (Levinthal 1997; Lazer & Friedman 2007; Csaszar & Levinthal 2016) one-bit flips are not the only strategy and certainly not the baseline behaviour for human search. We find consistent evidence that individual search behaviours are much more heterogeneous than typically assumed in the literature and that players' submissions are primarily based on some form of mental model, rather than trying to adapt the last solution (cf. a one-bit flip 'hill-climber'). These mental models influence search behaviours and determine, for instance, how “patient” (Winter et al. 2007) a solver will be with respect to negative feedback or how distant search moves are carried out.

References

Fitzmaurice, G. M., N. M. Laird, et al. (2012). Applied longitudinal analysis, John Wiley & Sons.
Lazer, D., & Friedman, A. (2007). The network structure of exploration and exploitation. Administrative Science Quarterly, 52(4), 667-694.
Miles, M. B., & Huberman, M. A. (1994). Qualitative Data Analysis (Vol. 2. ed). Thousand Oaks, CA: Sage.
Winter, S. G., Cattani, G., & Dorsch, A. (2007). The value of moderate obsession: Insights from a new model of organizational search. Organization Science, 18, 403-419.

---

[i] Since we depend on cognitive psychology construct, we use the term cognitive ability, rather than the more management oriented term cognitive capabilities (Helfat & Peteraf 2015; Levine et al. 2017).
[ii] In the H-XOR function the first four and last four variables (the halves) have a stronger interdependence within vs. between each other. Thus, ‘first four’ and ‘last four’ is a natural problem decomposition. While the problem is very difficult, interviews show that at least one participant managed to learn the basic structure of the problem and generate an appropriate mental representation, leading the participant to solve the game in just 11 attempts.
[iii] Since the coding of search moves involved search distances, the first submission would serve as reference point, while due to the coding scheme the only possible search strategy that could be identified in the second attempt would be local search.
[iv] Such a coefficient is considered to be “acceptable” (Flynn et al. 1994), but in order to have a more conservative measure we have decided to remove the item.
[v] As a robustness check, subsequent analyses were conducted both including and excluding the item and resulting estimates do not change.
[vi] The operationalization of how search behavior was transformed into three search strategies was provided in section 3.3. Local search moves: 1 bit flip, model-based (or directed global) search: 2 bit flips that don’t violate underlying problem structure, random (undirected global) search: Everything else.