



SCHOOL OF BUSINESS AND SOCIAL SCIENCES
AARHUS UNIVERSITY

SEARCHING FAR AWAY FROM THE LAMP- POST

Essays on Problem Solving Strategies

PhD dissertation

Oana Vuculescu

Aarhus BSS
Aarhus University
Department of Management
2016

ACKNOWLEDGEMENTS

It seems customary to start this section with a bit of self-reflection about the process and the ups and downs of the PhD student life. I'll pay tribute to the custom and I'll just say this: too often I read the PhD Comics and thought: "ah, exactly!" Overall though (that is to say "on average" but it stays true also if you look at the median): this was fun! In the following lines I tried to remember all the people who made sure it stays that way. You might notice that this section is quite lengthy. It couldn't be helped. During the past three years, I have met, worked with, or alongside a number of people from which I have learned a great deal! I am genuinely grateful for that and this is my clumsy attempt to express that.

First of all, I want to thank my advisors Lars Frederiksen and Carsten Bergenholtz. I always told my colleagues that I can't complain about my PhD life, since I have the best possible combination of advisors and that has remained true throughout these years. Carsten, we often joke about the random moments that have shaped my academic life, but if it hadn't been for the very non-random support, both in the last steps of my master thesis and throughout the years that followed, I simply wouldn't be here. Thank you! Thank you for being my advisor, my mentor, my co-author and in a few unnamed (Bamboozling) occasions, my accomplice. Lars, thank you for taking a chance on me as a PhD student. Thank you for believing in me, when I didn't, and for doubting me.. when I didn't! Most of all, thank you for giving me the opportunity to try and the freedom to fail, while at the same time always pushing me to be better.

Next, aside from my "local team", I want to thank Luigi Marengo for having me as a visiting PhD student in Sant'Anna. Thank you for the amazing dinner conversations, for sharing your knowledge on modelling with a super-ignorant PhD student. As I wrote different parts of my thesis I would often stop and ponder how my understanding was shaped by our talks and just how much that experience has formed me. Thank you for your support and always positive feedback, despite my obvious shortcomings!

I further want to thank Borge Obel and Stephan Billinger for agreeing to be part of my PhD committee, for their tremendous patience and constructive comments.

I would also like to thank the good people from the Cognition and Behavior Lab. Despite my profound inability to read guidelines and follow procedures, my time there was always fun and productive! A special "thank you" to the "Interacting Minds Center" people and Andreas Roepstorff, in particular, for welcoming me into an exceptionally warm and constructive environment: I always left with a smile and good ideas!

Finally, I want to thank Jacob Sherson and Mads Kock from the Astronomy and Physics Department and two of the people behind Science at Home. Your help with getting a good game-design that attracts (rather than repels) people is invaluable, but what I am most grateful for is the time and effort you made

to work and exchange ideas with someone outside your field who got what you were saying half of the time and googled the other. I have learned a lot in our collaboration and I do hope we can continue on that path.

Turning to my colleagues at the Department of Management, I'm not sure what the APVs results are this year, but I found a really great working environment and I am grateful to you all for contributing to that!

Thank you, Peter for being my first and last office mate. Thank you for always being available to talk either research, or politics or both and for walking me through the (sometimes utterly impenetrable) Danish culture and customs. Thank you for being an outstanding colleague and a great friend! I look forward to writing our third paper together ;). Thank you Sylvie for the cheese cake! You do awesome things and you are a great colleague, but I genuinely think all is distilled in that amazing cheese cake that always seems to pop about when I need it the most! To my third office-mate, Giacomo: thank you for some very tangible and intangible ways you've made my life better: Freeletics and work marathons, for late night sushi and Domino's pizza at the office and, of course, mangas! I look forward to the moment when we share a glass of wine after you are also done and we agree "it wasn't much!!!"

A special thanks to my Romanian colleagues: Alex, Violeta and Livia for putting up with me, for coming along in my smoke-breaks and for being there for me when the English language was not enough to express my javascript-induced frustration!

Now, I'm not done yet. Since this is getting embarrassingly long, I will just collectively thank all my amazing colleagues from the former basement: Billy, Markus, Tymen, Annmaria, Siri, Michela, Emma, Thomas, Sarah and, well, everyone else I forgot. Thank you for all the past and future fun times!

I would also like to thank my colleagues from the Innovation Management group for always being available to provide feedback on my ill-written and improperly presented drafts!

Lastly, because they are the only ones I expect to read this far: I want to thank my family: for coming along on my half-planned trips and never complaining about not seeing me for months. For always blowing up their own plans when I was in town (makes one feel quite special) and for their constant support and encouragement. Catalin, you get your own very special thanks: "Mc.".

This thesis is dedicated to you, diligent reader, whoever you are.

/oana

Executive summary

Recent years have seen the rise and expansion of new internet-based forms of organizing the search for solutions to organizational problems (Boudreau and Lakhani 2013). By relying on external contributors for new ideas and solutions to internal problems, scholars have argued that firms have now at their disposal a new mechanism to address the exploration-exploitation trade-off (Afuah and Tucci 2012). However, despite the initial attention (e.g. Leimeister et al. 2009; Poetz and Schreier 2012), evidence as to the efficacy of these platforms is missing (Dahlander and Piezunka 2014) as is an understanding of how to best design the external search for solutions (Franke et al. 2013). The external search for solutions can take many shapes (e.g. Adamczyk et al. 2012; Estellés-Arolas and González-Ladrón-de-Guevara 2012)), but recent evidence (Lakhani et al. 2013; Wooten and Ulrich 2015) bring empirical support for the idea that repeated-entries (i.e. feedback-based) tournaments might hold the key to crowdsourcing problem-solving. A subset of organizational theory on search has been relying on a computational approach to capture the interplay between the nature of the problem to be solved and the long-term dynamics of how the search for solutions takes place. However, this literature mainly addresses organizational search and has so far paid less attention to how individual search processes unfold. In this literature, while there has been a substantial stream of research looking into experiential search (i.e. search driven by reinforcement learning), there is still limited knowledge with respect to cognitive search processes (i.e. search that is driven by a solver's internal representation of the problem) (Csaszar and Levinthal 2015; Lopez-Vega et al. 2016). This is a serious limitation for individual level search since research on individual problem-solving suggests that human problem-solvers rely on a combination of cognitively-informed and experiential choices (Doll, Simon et al. 2012).

In this dissertation, I contribute to the current innovation management research by taking the discussion on search to the individual level to further our understanding of repeated-entries innovation contests. In particular, in [paper 1](#) I lay the groundwork and show how previous modelling approaches rely on a number of assumptions about stylized agents and behaviour that can be problematic and lead to results which are not supported by empirical evidence (e.g. Billinger et al. 2013; Mason and Watts 2012). [Paper 2](#) focuses on how distant-search mechanisms influence the propensity for finding the optimal solution as well as how design mechanisms might shape the search process, according to the type of problem to be solved. I find support for the fact that a lower level of “persistence” (the readiness to abandon unsuccessful paths earlier) translates into higher accuracy in terms of problem representation and improved performance, which is consistent with recent experimental work showing that performance pressure indeed can be a “double edged sword” (Gardner 2012) and that relying on multi-shot

submissions is preferable to the classical broadcast search since it allows solvers to form an accurate problem representation. In [paper 3](#), I focus on the question: if depending on the environment, certain search behaviours are preferable, what determines or rather what are the sources of heterogeneity in human search behaviour? I show that the choice between local and non-local search is not determined solely by structural considerations (e.g. the nature of the problem to be solved or incentives) but also by a substantial variation in an individual's preferred way of processing information. In addition, even with similar cognitive styles, I find that solvers have different starting points and ways of integrating evidence: some relying more on priors, others were more adept at incorporating feedback. The findings in the last paper suggest that when broadcasting a problem, a firm should be concerned not only with the type of the problem and the type of knowledge required (Felin and Zenger 2014) but also whether she can attract the contributors likely to engage in the relevant kind of search.

References

- Adamczyk, S., A. C. Bullinger, et al. (2012). "Innovation contests: A review, classification and outlook." [Creativity and Innovation Management](#) 21(4): 335-360.
- Afuah, A. and C. L. Tucci (2012). "Crowdsourcing as a solution to distant search." [Academy of Management Review](#) 37(3): 355-375.
- Billinger, S., N. Stieglitz, et al. (2013). "Search on rugged landscapes: An experimental study." [Organization Science](#)
- Boudreau, K. J. and K. R. Lakhani (2013). "Using the crowd as an innovation partner." [Harvard Business Review April](#).
- Csaszar, F. A. and D. A. Levinthal (2015). "Mental representation and the discovery of new strategies." [Strategic Management Journal](#) Forthcoming.
- Dahlander, L. and H. Piezunka (2014). "Open to suggestions: How organizations elicit suggestions to engage externals in open innovation." [Research Policy](#) Forthcoming.
- Doll, B. B., D. A. Simon, et al. (2012). "The ubiquity of model-based reinforcement learning." [Current Opinion in Neurobiology](#) 22(6): 1075-1081.
- Estellés-Arolas, E. and F. González-Ladrón-de-Guevara (2012). "Towards an integrated crowdsourcing definition." [Journal of Information Science](#) 38(2): 189-200.
- Felin, T. and T. R. Zenger (2014). "Closed or open innovation? Problem solving and the governance choice." [Research Policy](#) 43(5): 914-925.

- Franke, N., C. Lettl, et al. (2013). 'Does god play dice?' Randomness vs. deterministic explanations of idea originality in crowdsourcing. 35th Druid Celebration Conference 2013.
Barcelona, Spain.
- Lakhani, K. R., K. J. Boudreau, et al. (2013). "Prize-based contests can provide solutions to computational biology problems." Nature Biotechnology 31(2): 108-111.
- Leimeister, J. M., M. Huber, et al. (2009). "Leveraging crowdsourcing: activation-supporting components for IT-based ideas competition." Journal of Management Information Systems 26(1): 197-224.
- Lopez-Vega, H., F. Tell, et al. (2016). "Where and how to search? Search paths in open innovation." Research Policy 45(1): 125-136.
- Mason, W. and D. J. Watts (2012). "Collaborative learning in networks." PNAS 109(3): 764-769.
- Poetz, M. K. and M. Schreier (2012). "The value of crowdsourcing: Can users really compete with professionals in generating new product ideas?" Journal of Product Innovation Management 29(2): 245-256.
- Wooten, J. O. and K. T. Ulrich (2015). "Idea generation and the role of feedback: Evidence from field experiments with innovation tournaments." Available at SSRN 1838733.

Resume

I de senere år er der kommet flere nye web-baserede muligheder, hvor virksomheder kan søge hjælp til løsninger på dens problemer (Boudreau og Lakhani 2013). Forskningen har vist, at ved at satse på ekstern bistand i forbindelse med at skaffe nye ideer og løsninger på problemer, har firmaer nu nye redskaber til rådighed, når de skal afveje søgning og nyttiggørelse (Afuah og Tucci 2012). På trods af den bevågenhed dette nød i starten (fx Leimester et al. 2009; Poetz og Schreier 2012), mangler der dokumentation for effektiviteten af disse platforme (Dahlander og Piezunka 2014) såvel som en forståelse af, hvordan man bedst sætter den eksterne søgning efter løsninger op (Franke et al. 2013). Søgning eksternt efter problemløsninger kan antage mange former (fx Adamczyk et al. 2013; Estellés-Arolas og González-Ladrón-de-Guevara 2012), men på det seneste (Lakhani et al. 2013; Wooten og Ulrich 2015) er der kommet empirisk belæg for konceptet om, at *repeated-entries* turneringer (dvs. med feedback) måske er nøglen til problemløsning gennem *crowdsourcing*. Den strenghed af organisationsteorien, der omhandler søgning, har henholdt sig til en datatilgang for at fastholde samspillet mellem arten af det problem, der ønskes løst, og hvordan dynamikken i søgningen efter løsninger på lang sigt udvikler sig. Litteraturen beskæftiger sig imidlertid i hovedsagen med virksomheders søgning og har indtil videre kun beskæftiget sig ganske lidt med, hvordan individers søgeresponsers udvikler sig. Mens der har været en betydelig strøm af studier, der har set på erfaringsbaseret søgning (dvs. søgning der er drevet af forstærkende læring), ved vi stadig kun lidt om kognitive søgeresponsers (dvs. søgning der er drevet af løserens interne fremstilling af problemet) (Csaszar og Levinthal 2015; Lopez-Vega et al. 2016). Dette begrænser søgning på individniveau i alvorlig grad, da forskning om individuel problemløsning viser, at de, der skal løse et problem, er afhængige af en kombination af erkendelse og erfaring i deres valg (Doll et al. 2012).

Denne afhandling bidrager til forskningen i innovationsledelse ved at diskutere dette på individniveau med det formål at fremme forståelsen af *repeated-entries* innovationskonkurrencer. Især i artikel 1 lægger jeg fundamentet og viser, hvordan tidligere modelleringsstilgange er afhængige af nogle antagelser om *stylized agents* og tvivlsom adfærd, der kan føre til resultater, som ikke understøttes empirisk (Billinger et al. 2013; Mason og Watts 2012). Artikel 2 fokuserer på, hvordan *distant-search* påvirker, hvordan tendensen til at finde den optimale løsning såvel som opsætningen kan forme søgeresponsen, alt efter hvilken type af problem, der ønskes løst. Jeg har fundet belæg for, at mindre ihærdighed (vilje til at afvige fra blindgyder hurtigere) giver større træfsikkerhed, hvad angår problemfremstilling og bedre performance; dette er i tråd med nyere studier, der viser, at præstationspres i høj grad kan være et ”tveægget sværd” (Gardner 2012), samt at gentagne bidrag er at foretrække frem for den

klassiske spredenhaglssøgning, da den giver problemløserne et mere præcist billede af det problem, der skal løses. I artikel 3 fokuserer jeg på spørgsmålet: ”hvis en særlig søgeadfærd foretrækkes, alt efter hvilket område det drejer sig om, hvad afgør så forskellene i menneskers søgeadfærd?” Jeg viser, at valget mellem lokal og ikke-lokal søgning ikke alene afgøres af strukturelle hensyn (fx problemets karakter eller incitamenter), men også af nævneværdig variation i individets foretrukne måde at behandle information på. Endvidere har jeg fundet, at selv hvis løserne har forholdsvis mange kognitive træk til fælles, har de forskelligt udgangspunkt og forskellige måder, hvorpå de bruger dokumentation: nogle satser mere på, hvad de ved i forvejen, andre er dygtigere til at bruge feedback. Resultaterne fra den sidste artikel antyder, at når et firma fremlægger et problem, skal firmaet ikke kun bekymre sig om, hvilken type problem der er tale om, og hvilken type viden der er nødvendig (Felin og Zenger 2014), men også om det kan tiltrække bidragssydere, som vil engagere sig i relevant søgning.

Contents

Introduction	13
Innovation as search	13
1.1.Organizational-level search	14
1.2.Innovation contests	15
2. Empirical context and motivation.....	17
2.1. Overview of literature on innovation contests.....	18
3. Research approach.....	27
3.1. Problem-solving as search	28
3.2. Model-based search.....	29
3.3. Research methods	30
3.3.1. Simulation framework	31
3.3.2. Experiments	33
3.3.3. Survey and qualitative interviews.....	35
Section 4: Overview of the thesis	35
Paper 1: Fitness landscapes in organizational theory: research challenges and future directions.....	35
Paper 2: Searching far away from the lamp post, an agent-based model	36
Paper 3: Micro-foundations of problem-solving: what determines how individuals search?	38
4.1. Overview	39
4.2. Methodological contributions and limitations	41
4.3. Managerial implications.....	42
4.4. General discussion and ways forward	43
Fitness landscapes in organizational theory: Research challenges and future directions	55
Searching far away from the lamp post: an agent-based model.....	93
Micro-foundations of problem solving: What determines how individuals search?	127

Introduction

Innovation as search

Modern definitions of innovation are intrinsically and explicitly linked to search. Drucker defines systematic innovation as “the purposeful and organized search for changes, and in the systematic analysis of the opportunities such changes might offer for economic or social innovation” (Drucker 2014: p.49). Search is the first phase in Tidd and Bessant et al.’s (1997) process model of innovation, and Dosi (1988) argues that “In an essential sense, innovation concerns the search for, and the discovery, experimentation, development, imitation and adoption of new products, new production processes and new organizational set-ups” (Dosi 1988: p.222). These definitions all acknowledge that innovation is necessarily a path dependent process under conditions of uncertainty. During the last decade, *individual* search has received increased attention in part due to the rise of distributed innovation phenomena such as crowdsourcing (Howe 2008), broadcast search initiatives such as Innocentive (Lakhani 2008) or innovation contests (Boudreau et al. 2011). By focusing on the latter, this project aims to shed light on this fundamental building block of innovation: individual search behaviour.

In doing so, I shift the level of analysis from organizational to individual. This is in line with the recent call for microfoundations for organizational strategy (e.g. Felin et al. 2012). It has been long acknowledged that in order to understand how “administrative practices and organization structure” determine the generation and development of “innovative ideas” (Abernathy and Clark 1985, p. 3), management scholars have to zoom in on organizations and attempt to understand what governs the internal processes that lead to innovation. Focus is thus shifted from how organizations manage their *knowledge base*, i.e. create and transfer knowledge (Kogut and Zander 1992; Nonaka 2007) to how organizations can provide the necessary “incentives and direction” such that *individual* knowledge creation, acquisition and integration can occur (Grant 1996). As Felin et al. argue, “as fields progress, evidence suggests that assumptions about micro-level uniformity prove unsustainable and inaccurate” (Felin et al. 2012: p.1354). Indeed, although ‘search’ is a key theme in the organizational literature, and despite the fact that organizational scholars argue that the principles of organizational and individual search behaviours are “remarkably similar” (Billinger et al. 2013: p.12), our understanding regarding how these search processes take place remains limited. Empirical work (e.g. Mason and Watts 2012; Billinger et al. 2013) cautions that conceptual models so far are not sophisticated enough to capture individual problem-solving behaviours. In particular, while there has been a substantial stream of research looking into experiential search (i.e. search driven by

reinforcement learning¹), we are still in the dark when it comes to cognitive search (i.e. search that is driven by a solver's internal representation of the problem) (Gavetti and Levinthal 2000), despite the fact that research on individual problem-solving suggests that human problem-solvers rely on a combination of cognitively-informed and experiential choices (Doll et al. 2012).

The difference between the two modes of search can be illustrated by the old joke (and seemingly even older folklore story²) about a drunk looking for his keys near a lamp post. A policeman passing by stops and helps him for a few minutes without success. He then finally asks whether the man is certain he had dropped them near the lamp post. The drunk replies negatively and continues to explain that he's searching near the lamp post because "the light is so much better".

In this project I thus investigate the role, the mechanisms and the determinants of searching "far away from the lamp post", that is, search behaviours that go beyond the immediate and easy accessible solutions (i.e. searching just around the lamp post).

1.1. Organization-level search

Organizational literature on search owes a great deal to the evolutionary economics literature with its strong emphasis on variation, adaptation and selection as the main dynamics that describe innovative processes (Dosi and Nelson 1994). It is equally indebted to the same literature with respect to a formal framework (i.e. evolutionary models) that takes into account non-linearity, dynamics and self-organization (Arthur 1988).

This literature regards organizational learning as an "experiential learning" system (Levinthal and March 1981: p.308), where the focal organization responds to the environment by either "refinement" or "innovation". Search behaviours are thought to be adjusted through trial and error: "behaviour that is associated with success is repeated; behaviour which is associated with failure tends not to be repeated" (Levinthal and March 1981: p.308). March (1991) further argues that an organization's limited resources induce a choice between exploratory behaviours (i.e. distant search), which entail variation and risk, and exploitative behaviours which entail focusing on the current environment to generate predictable and short-term benefits (i.e. local search). The inherent organizational and cognitive limitations are thought to determine a higher propensity for firms' engagement in local search to the detriment of more variance inducing exploratory search (March 1991; Levinthal and March 1993). Local search is not necessarily disadvantageous (Laursen 2012), but it does imply reduced informational diversity which means that cognitively

¹ Learning occurs as a result of positive (respectively negative) feedback which encourages (or deter) the focal agent from pursuing the same course of action in the future.

² <http://quoteinvestigator.com/2013/04/11/better-light/> retrieved 31 December 2015.

or organizationally distant solutions are either simply not found or not implemented due to higher transaction costs.

Modelling organizational search as an evolutionary process has allowed scholars to address a number of relevant managerial issues, e.g., the interplay between a competitive environment and the emergence of new organizational forms (Levinthal 1997), new product development (Mihm et al. 2003) or how organizations should structure their search (Baumann and Siggelkow 2013). In addition, a number of empirical studies have looked into the structural determinants (Raisch and Birkinshaw 2008) that shape the choice between exploration and exploitation, e.g., the competitive environment (Auh and Menguc 2005) or the organizational structure (Fang et al. 2010). However, there is still considerable disagreement about what these strategies actually consist of (Katila and Ahuja 2002; Lavie et al. 2010). More recently, scholars have attempted to go beyond the traditional dichotomy (which is argued to be just one dimension of search – the ‘where’) and to focus also on ‘the how’ (Lopez et al. 2016), a different dimension which should capture the particular search strategies firms apply.

This PhD dissertation contributes to this current debate on *how* distant search can be conducted by focusing on micro-level search behaviours: individual-level problem-solving. Specifically, by drawing on recent cognitive science results that provide evidence for goal-directed behaviours in exploratory search behaviours (Daw et al. 2006), I focus on distant search that is “structured” (as opposed to random) (Cohen et al. 2007) to further our understanding of how problem-solving processes unfold.

1.2.Innovation contests

A more direct contribution of my dissertation regards the design of repeated-entries innovation contests. There are several ways an organization can overcome the downside of having a narrow search focus (Laursen 2012), while at the same time benefitting from the efficiencies of local search: resorting to internal solutions, e.g. skunkworks (Fosfuri and Rønde 2009), markets (e.g. outsourcing) or relying on solutions that are generally clustered under the umbrella term open innovation (Chesbrough 2003), i.e. distributed innovation systems. In this latter category, benefitting from the advent of the internet and ever decreasing communication costs, I find a (relatively) new way to overcome the high costs of distant search by relying on idea contests as a mechanism for generating innovation (Terwiesch and Xu 2008). In a recent paper, Lakhani et al. (2013) showcase the efficiency of such systems in a Top Coder contest. Out of over 600 submissions received, 30 were above the benchmark performance. Most remarkably, the best submissions were 1000 times better than the benchmark (Lakhani et al. 2013). Poetz and Schreier (2012) also investigate whether problem-solving via innovation contests can perform at

a comparable level with traditional problem-solving ways (e.g. experts) by setting up a real-life experiment and find empirical evidence to support this claim. The leap in performance can be further increased if the system enables knowledge sharing and interaction between solvers, i.e. knowledge brokering (Lakhani et al. 2013; Villarroel et al. 2013). However, while outsourcing the search for a solution is considered to be cheaper than in-house R&D (Howe 2008), it is by no means free, and there is no comprehensive study on how these platforms should be set up although it is becoming generally acknowledged that there's more to crowdsourcing than "turning on a website and putting up a reward" (Jouret 2009: p44). As Dahlander and Piezunka (2014) and Franke et al. (2014) show, more often than not, these set-ups do not deliver on the promise and most organizations fail to accumulate suggestions over time. This evidence, however, is addressing static, single-submission contests, while recent research in the field of innovation platforms deals with dynamic, multi-shot set-ups (Wooten and Ulrich 2015). In these tournaments, solvers have the option to submit multiple entries and receive feedback and even sometimes have access to other people's submissions (Wooten and Ulrich 2015). This and similar work (Lakhani et al. 2013) brings empirical support for the idea that repeated-entries (i.e. feedback-based) tournaments might hold the key for crowdsourcing problem-solving. The iterative process allows solvers to learn over time, which in turn means that better (or at the very least similar) solutions can be reached with a significantly lower number of solvers (Vuculescu and Bergenholz 2014).

I argue that, in a similar vein to how managers need to choose between alternative organizational forms according to the type of problem to be solved (Nickerson and Zenger 2004) so can broadcasting organizations choose between different configurations of governance mechanisms according to the type of problem to be solved and implicitly the search strategies. Thus, by disentangling the types of search behaviours that solvers engage in, I aim to provide managers with a better tool for understanding how various governance factors, such as e.g. prize- incentive should be set up, to improve efficiency.

This thesis extends previous work on problem-solving by concentrating explicitly on the role of distant search, when search takes place in the absence of firm hierarchies at an individual level. The focus is furthermore on those settings where solvers receive direct feedback. The goal is to contribute to the research stream of innovation as search by expanding our understanding of how individual problem-solving processes can be described. In the remaining sections of this introduction I first present a brief overview of the literature on innovation contests, which serves as a main motivation and empirical grounding (Section 2); I then explain the research approach as well as briefly outline the research methods that were employed in each of the three papers (Section 3) and finally I conclude by presenting an overview of the three papers and conclusions (Section 4).

2. Empirical context and motivation

Innovation contests are online competitions where a seeker organization outsources a particular problem to a (generally large and undefined) population of potential solvers. Repeated-entries innovation contests have the particularity that the competition is structured in phases (e.g. the Netflix prize, Villarroel et al. 2013) and that the seeker organization provides feedback throughout the competition, allowing solvers to improve/change previous submissions. In the Netflix prize, the challenge was to design an algorithm that would improve the company's own recommender system (Cinematch). Netflix released a database of over 100 million ratings from a sample of 480,000 subscribers. Participants were required to submit an algorithm that would make predictions over the 3 million ratings generated by subscribers from the same dataset. Throughout the contest participants were not limited in the total number of attempts they could submit. In terms of frequency, the limit was set initially as one per week and reduced subsequently to one per day. There were two sources of feedback: offline feedback – using a sub-set of the training set for which ratings were provided – and online feedback – following a submission, they would receive an immediate and accurate measure of their root mean square error in predicting the ratings (Bennett and Lanning 2007). This set-up allowed participants to gradually improve their submissions and indeed Bennet and Lanning (2007) report that, in 2007, 2000 teams made more than 13,000 submissions.

Repeated-entries innovation contests are thus a closer match to traditional problem-solving processes that might occur inside the firm where solvers hardly ever find themselves in the position of creating solutions in the absence of any feedback. Still, they are free from a number of extraneous factors that might hinder or confound effects in traditional problem-solving settings, e.g. organizational hierarchies, communication barriers across functional divisions, etc. One might say that online problem-solving platforms are to in-house innovation what stock exchanges are to markets: a less noisy environment with fewer transaction costs and better information.

Unfortunately, although the use of repeated-entries innovation contests is on the rise (Wooten and Ulrich 2011), there is not a lot of literature dedicated specifically to this type of innovation contests. A search on Web of Science for “repeated entries AND innovation AND contests”¹ actually yields no results, although a number of papers have investigated specifically innovation contests that provide feedback to solvers (e.g. Terwiesch and Xu 2008; Murray, Stern et al. 2012; Lakhani, Boudreau et al. 2013; Villarroel, Taylor et al. 2013). The remainder of this section will thus deal with the overall category of “innovation contests”.

¹ Similar searches (e.g. replacing “contests” with “tournaments” or “repeated entries” with “multi-shot”) were equally unsuccessful.

There is a large heterogeneity when it comes to the actual shape that “crowdsourcing problem-solving” takes. As such, there are authors looking into the broader phenomenon of crowdsourcing (Brabham 2008; Afuah and Tucci 2012; Estelles-Arolas and Gonzalez-Ladron-de-Guevara 2012), but there is also a parallel trend that zeros in a particular sub-type – e.g. idea competitions (Ebner, Leimeister et al. 2009; Hutter, Hautz et al. 2011; Lampel, Jha et al. 2012) and even scholars distinguishing within this latter category based on the value of the prize – Grand Innovation Prizes versus smaller scale competitions and challenges (Murray et al. 2012). Felin and Zenger (2014) advance a taxonomy of these forms of open innovation based on the characteristics of the problem to be solved and thus distinguish between innovation contests and community-directed innovation. The first approach traditionally relies on prize-based challenges to attract a large number of contributors who submit (to either an intermediary or directly to the seeker organization) solutions to a given problem. The prototypical example is Innocentive that has compiled a database of over 350,000 registered solvers and has a reach of over 13 million people. This type of crowdsourcing problem-solving, Felin and Zenger (2014) argue, is ill-suited for complex, ill-defined problems. In contrast, in user (or solver) communities contributors freely share their solutions (Jeppesen and Frederiksen 2006) and could potentially tackle more complex problems. Repeated-entries innovation contests are to be found somewhere in between these two forms, as they 1) allow for more than “broadcasting” a problem in terms of communication between the seeker organization and the solvers, and 2) they often also allow for communication between solvers (Wooten and Ulrich 2015). This, in conjunction with the fact that a number of success cases actually feature repeated-entries innovation contests (Wooten and Ulrich 2015; Murray et al. 2012; Lakhani, et al. 2013; Villarroel et al. 2013), suggests that it is timely to extend Felin and Zenger’s (2014) argument to the specific instances where feedback is available to participants.

2.1. Overview of literature on innovation contests

This literature review serves a dual purpose: first, it is evidence of the management field’s growing interest in the topic of external search for internal problems, and second, I also find strong support for an increasing awareness from scholars *outside* the management literature. This suggests that other fields (i.e. computer science, medicine, geography, physics etc.) increasingly rely on such tools to solve scientific problems. Finally, I conclude with a short outline of the current research dealing with this topic.

Getting an overview of the literature turned out to be a difficult task because the above-mentioned heterogeneity in shape is mirrored by the relatively numerous (often overlapping) variations encountered in how scholars refer to the re-discovered use of “distant search”. A

number of authors borrow, albeit some reluctantly, the term “crowdsourcing” introduced by Jeff Howe (2008). Others talk about “innovation contests” (Bullinger, Neyer et al. 2010; Boudreau, Lacetera et al. 2011), “commons based peer production” (Benkler 2002), “open collaborative innovation” (Baldwin and Hippel 2011), etc.

Since crowdsourcing is the popular term associated with this phenomenon, I did a first search using “crowdsourcing” as a “topic” on Web of Science Core Collection. Out of 1732 papers, the highest ratio in this selection falls outside the categories of “Operations Research Management Science” (2.2%) or “Business Economics” (9.3%), but to “Computer Science” with more than 50 per cent of papers being listed as pertaining to this category. The second most preeminent research area is “Engineering” with 21 per cent (373) of the total number.

After narrowing down these fields, more than 94 per cent of the papers are eliminated and within this restricted list the number of “Management” (cf. Web of Science classification) papers on the topic of crowdsourcing accounts for about 80 per cent (90 papers out of 113). Figure 1 illustrates the increasing trend in publications listing “crowdsourcing” as a “topic” in the years following Jeff Howe’s introduction of the term (2008).

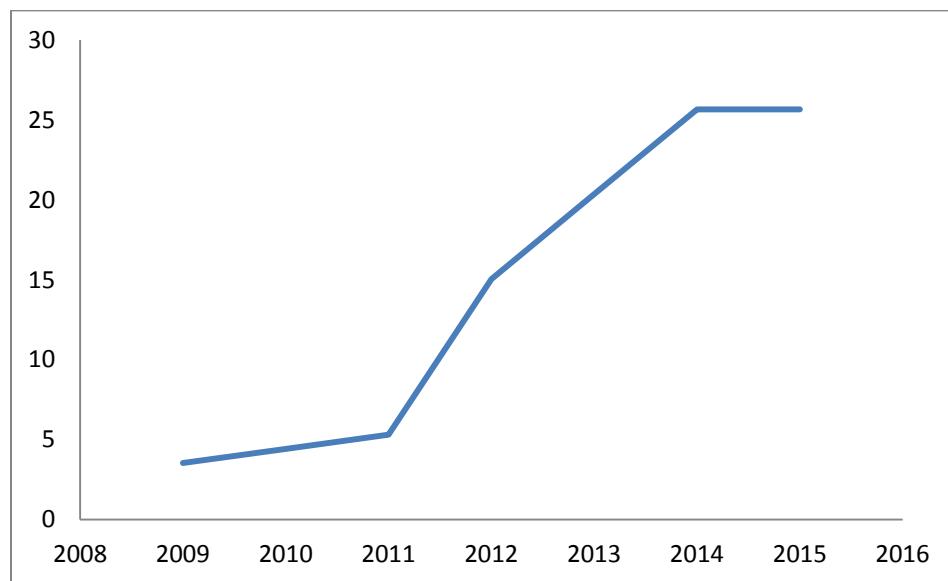


Figure 1. Percentage of total number of management publications listing over time. N=113

An outline of the methods employed in the ten most cited papers is revealing as to the young age of the topic: five are conceptual papers (Bonabeau 2009; Afuah and Tucci 2012; Bogers and West 2012; Marjanovic et al. 2012; Boudreau and Lakhani 2013), one uses a mixed-method approach combining content analysis and social networks analysis (Hutter et al. 2011) and the

remaining four are case based, e.g. an ideas competition for an ERP Software company (Ebner et al. 2009), the Atizo innovation platform (Frey et al. 2011), the Dell Idea Storm competition (Bayus 2013) and three crowd-science platforms (Franzoni and Sauermann 2014).

Research Areas	records	% of 1732
COMPUTER SCIENCE	874	50.46
ENGINEERING	373	21.53
BUSINESS ECONOMICS	162	9.35
TELECOMMUNICATIONS	140	8.08
INFORMATION SCIENCE LIBRARY SCIENCE	91	5.25
LINGUISTICS	72	4.15
SCIENCE TECHNOLOGY OTHER TOPICS	49	2.82
PSYCHOLOGY	39	2.25
OPERATIONS RESEARCH MANAGEMENT		
SCIENCE	39	2.25
REMOTE SENSING	37	2.13
ENVIRONMENTAL SCIENCES ECOLOGY	35	2.02
IMAGING SCIENCE PHOTOGRAPHIC		
TECHNOLOGY	34	1.96
HEALTH CARE SCIENCES SERVICES	32	1.84
COMMUNICATION	29	1.67
EDUCATION EDUCATIONAL RESEARCH	28	1.61
GEOGRAPHY	27	1.55
ROBOTICS	25	1.44
PUBLIC ADMINISTRATION	25	1.44
GEOLOGY	25	1.44

Table 1. Distribution of publications listing “crowdsourcing” as Topic in Web Of Science Core Collection

In a systematic literature review of the topic, Adamczyk et al. (2012) find that literature on the design and management of these platforms (as opposed to literature describing their *use* in e.g. education or sustainability) is to be found predominantly under “innovation contests”. Indeed, a search for “innovation AND contests” as “topic” in the Web of Science Core Collection yields 60 papers, with over 63 per cent published under the category of “Business Economics” (cf. Web of Science) and 18 per cent under “Operations Research Management Science”. Similar to “crowdsourcing”, I find a matching ascending trend of the use of “innovation contests” (Table 2) indicative of the recent focus on understanding how large groups of (often) anonymous individuals solve problems.

Publication Years	records	% of 60
2015	15	25
2014	11	18.33
2013	11	18.33
2012	8	13.33
2011	6	10
2008	3	5
2007	2	3.33

Table 2. Distribution of publications listing “Innovation contests” as Topic in Web Of Science Core Collection

In contrast to Adamczyk et al.’s (2012) findings, I find that the most cited papers in this selection are not primarily based on a qualitative design but rely on a quantitative set-up (e.g. (Bullinger, Neyer et al. 2010; Boudreau, Lacetera et al. 2011; Zheng, Li et al. 2011; Bayus 2013; Lakhani, Boudreau et al. 2013) and modelling (Schöttner 2008; Terwiesch and Xu 2008). Only one paper among the ten most cited papers uses a comparative case-study approach (Franzoni and Sauermann 2014). These results corroborated with the first search suggest that management scholars prefer the use of “innovation contests”, while outside the management/business field, the use of “crowdsourcing” is more common.

Given previous considerations and to acquire a reasonable knowledge of the literature I followed the keyword-based structured literature review with a quasi-structured backward and forward citation literature review starting from several of the highly cited papers identified previously. Table 3 presents a summary of the most important papers.

Author/year	Focus	Methods/data	Main findings
Afuah and Tucci 2012	Crowds	Conceptual	Theoretical framework for crowdsourcing as distant search.
Boudreau and Lakhani 2013	Crowds	Illustrative case study	The type of problem to be solved should determine the type of crowdsourcing to be used.
Boudreau 2012	Multi-sided platform	Longitudinal data on eight market leading platforms	Increasing the number of contributors leads to an increase in diversity. Increased numbers of contributors crowds out motivation. This is countered by the network effect, however the crowding out is the more powerful effect.

Author/year	Focus	Methods/data	Main findings
Boudreau, Lacetera et al. 2011	One shot contest	Quantitative research based on longitudinal data collected from TopCoder (reproducing to some extent an experimental design)	When increasing the number of participants, there are two opposing forces: 1) Decrease in the incentive to participate for and 2) Increased likelihood that a solution will be found. Main results: for ill-structured problems the second effect overcomes the first and the reverse is true for simple problems.
Bayus 2013	Problem solving community	Longitudinal quantitative study of an online community	Successful contributors (from previous "games") propose less creative ideas (due to "cognitive fixation" effects). Effects are moderated by the intensity of commenting on other people's ideas.
Boudreau, Helfat et al. 2012	Contest	TopCoder: Longitudinal mixed-methods study	Increased competition leads to worse performance outcomes, on average. This effect is amplified by the added presence of a "super-star". Negative effects are more visible in highly skilled contributors.
Bullinger, Neyer et al. 2010	Contest	Quantitative	Cooperation fosters boundary spanning behavior between competing communities. Innovativeness has an inverted U relationship with cooperation – i.e. very high and very low co-operators are very innovative.
Dahlander and Piezunka 2014	User community	Quantitative analysis of longitudinal data across multiple communities	Most organizations fail to accumulate suggestions over time. Having employees themselves participate in the community and accepting suggestions from members of the community, especially from newcomers, has a significant (positive) effect on the likelihood of that community surviving.
Ethiraj and Levinthal 2004		Simulation model (NK)	Modularization is preferred when an efficient parallel search is possible i.e. 1) genetic search is prevalent, 2) information about module performance is available and or 3) selection mechanisms are in place, 4) diversity is high.

Author/year	Focus	Methods/data	Main findings
Erat and Krishnan 2012	Contest	Analytical model	Investigates the relationship between problem specification and search scope. A narrow specification confines the search in a particular space in the landscape. On the other hand, a narrow specification increases likelihood of high numbers of contributors. A moderate level of specification is desirable.
Fang et al 2010		Simulation	Learning performance is optimal in a network which is clustered, but has a small fraction of random, cross-group link. Increasing the size of the clusters decreases performance in any setting.
Fullerton and McAfee 1999	Tournament	Formal model	For a large class of contests the optimal number of competitors is two. This insight makes designing the tournament easier and highlights the importance of selecting highly qualified contestants. This study proposes a game theory logics auction where players are incentivized to bid to win in stage one to take part in stage 2.
Füller, Hutter et al. 2011	Design competition	Mixed methods	Main factors in attracting and retaining contributors are the perceived sense of community and their overall experience. Having a sense of community and experience determines the co-creation experience which in turns influences the quality, quantity and frequency of contributions
Hong and Page 2004	Community	Simulation	Main results: diversity does trump ability when the randomly selected group is large enough that it ensures enough diversity. For small groups, best performers are preferable, but with larger groups, the increase in performance for best performers is marginal, while for random groups they are significant and ultimately outperform “experts”.

Author/year	Focus	Methods/data	Main findings
Hutter, Hautz et al. 2011	Contest	Exploratory case study	Hutter et al. identify four roles a contributor can play in a problem-solving context and their effect on performance. Communitors (people highly engaged both in terms of comments and ideas) seem to have “a large proportion” of the best solutions.
Jeppesen and Lakhani 2010	One shot contest	Quantitative analysis of 166 challenges (Innocentive)	Investigates the profile of the successful solver: successful solvers are “distant” from the knowledge domain of the problem, socially marginal and interested in fewer knowledge domains.
Lakhani, Boudreau et al. 2013	Contests	Mixed methods	Main ideas and results: contributions from the TopCoder platform outperform in-house experts by several orders of magnitude.
Felin and Zenger 2014		Conceptual	Paper advances a framework for how, according to the degree of hidden knowledge and the type of problem, focal firms should structure their search. In this framework, innovation contests are more adept to solving simple problems that require a high degree of hidden knowledge.
Knudsen and Levinthal 2007	NK	NK simulation	Expert evaluation leads to lower performance.
Mason and Watts 2012	Contest	Experimental data (wildcat wells game)	Communication between solvers improves performance. Centralized individuals are the ones who are more likely to identify the best performing solution. Efficient networks are outperforming inefficient networks for complex problems.
Terwiesch and Xu 2008	Contests	Simulation model	Via a simulation model this paper investigates the tension between the desire to have high numbers of solvers and the corresponding decrease in individual effort is mitigated by changing the award structure.
Villarroel et al. 2013	Multi-stage competing communities	Qualitative and quantitative data (survey) + longitudinal data from Netflix prize	In an innovation contest, knowledge sharing (i.e. free revealing and knowledge brokering) improves average performance, but for best performance the knowledge brokering needs to be

Author/year	Focus	Methods/data	Main findings
			moderate
Von Krogh et al. 2012	One shot contest	Multiple case study	Framework for key processes linked to formulating sharable problems: “problem generation”, “problem separation” and “problem publication”. There are four essential organizational roles: developer, advisor, coordinator and sponsor.
Winter et al. 2007	Individual search	Simulation	The most efficient searches occur when there is a moderate level of “obsession”, in that the agents take into account both local and distant information.

Table 3. Overview of selected papers dealing with innovation contests

The diverse terminology found in the literature is oftentimes misleading. As innovation contests increased in popularity, as shown earlier, scholars focused on specific instances (i.e. types) of innovation contests but were never explicit about whether or how their conclusions should apply to other cases. Table 4 presents an outline of the innovation contests found in the literature classified according to governance features.

Governance feature	Author/paper
<i>Problem</i>	Benkler 2002; Ethiraj and Levinthal 2004; Lazer and Friedman 2007; Terwiesch and Xu 2008; Boudreau, Lacetera et al. 2011; Afuah and Tucci 2012; Erat and Krishnan 2012; Felin and Zenger 2012; Bayus 2013; Boudreau and Lakhani 2013
<i>Number of contributors</i>	Page 2007; Terwiesch and Xu 2008; Boudreau, Lacetera et al. 2011; Füller, Hutter et al. 2011; Boudreau 2012; Erat and Krishnan 2012; Felin and Zenger 2012
<i>Type of contest</i>	
a. competition	Terwiesch and Xu 2008; Bullinger, Neyer et al. 2010; Füller, Hutter et al. 2011; Boudreau and Lakhani 2013; Villarroel forthcoming
b. collaboration	Bullinger, Neyer et al. 2010; Lee 2010; Frey, Lüthje et al. 2011
<i>Structure</i>	
a. single shot	Jeppesen and Lakhani 2010; Boudreau, Lacetera et al. 2011; von Krogh, Haefliger et al. 2012
b. staged	Terwiesch and Xu 2008; Murray, Stern et al. 2012

Governance feature	Author/paper
<i>Communication</i>	
a. pro	Hong and Page 2004; Rost 2011; Mason and Watts 2012
b. against	Lazer and Friedman 2007
<i>Incentives</i>	
i. intrinsic	Shah 2006; Brabham 2008; Frey, Lüthje et al. 2011; Füller, Hutter et al. 2011
ii.recognition	Jeppesen and Frederiksen 2006; Shah 2006; Fang and Neufeld 2009; Huberman, Romero et al. 2009; Dellarocas 2010; Frey, Lüthje et al. 2011
iii.awards	
a. one	Erat and Krishnan 2012
b. multiple	Terwiesch and Xu 2008; Erat and Krishnan 2012; Murray, Stern et al. 2012
<i>Selection</i>	
a. self-selection	Ethiraj and Levinthal 2004; Singh and Fleming 2010; Felin and Zenger 2012
b. by panel	Piezunka and Dahlander 2014
<i>Type of call</i>	
a. open	Terwiesch and Xu 2008
b. closed	Murray, Stern et al. 2012

Table 4. Types of innovation contests identified in the literature

Overall, although the literature on innovation contests has moved beyond the initial descriptive works described in Adamczyk et al. (2012), it seems still to be struggling to identify a comprehensive explanation of the underlying mechanisms of problem-solving or how to best design these platforms. Empirical research either focuses on a micro level – abilities, knowledge, cognition (Hong and Page 2004), the relationships between contributors – i.e. networks (Lazer and Friedman 2007), social influence (Mavrodiev, Tessone et al. 2013), etc., or on a macro level – the number of contributors (Boudreau 2012), incentive systems (Shah 2006), problem type (Lorenz et al. 2011) and the overall outcome (e.g. are these platforms successful in identifying a solution or not (Poetz and Schreier 2012)).

Most of the empirical papers look at the relationship between various factors (as identified by respective literatures) to determine what accounts for the success of these platforms; but even when a correlation between an identified factor and the solution is found to be significant, the explanatory power (as indicated by R square) is reduced: e.g. in Bullinger et al. 2010 the degree

of cooperative orientation and the degree of innovativeness have been found to have a quadratic relation with an R square of 0.083, and Frey et al.'s (2011) model of knowledge diversity yields interaction effects ranging from $f^2 = 0.07$ to $f^2 = 0.12$ – small to moderate effects (Bullinger et al. 2010; Frey et al. 2011). These results are complemented by a major empirical study conducted by Franke et al. (2014) which concurrently looks at factors such as contest design elements (i.e. incentives, interaction, task framing), contributors' attributes (expertise, skills, traits and roles) and finally situational factors. They conclude "god plays dice" – that is randomness (variation introduced by increased numbers of participants) explains far more in the dependent variable variation (submission quality) than any of the factors considered. However, these results do not address and cannot be extrapolated to repeated-entries innovation contests, and empirical (Wooten and Ulrich 2015; Wooten and Ulrich 2015) and conceptual work (Vuculescu and Bergenholz 2014) suggests these particular settings have the potential to be more effective in generating high quality solutions.

3. Research approach

In line with previous work on organizational problem-solving and learning (March 1991), this project takes its point of departure in the complexity science paradigm. Unlike the traditional "Newtonian" paradigm, complexity science approaches to social systems posit that social phenomena cannot be reduced to individual components whose interactions can be captured by linear equations, but rather that such structures are highly non-linear with small initial differences often being propagated throughout the entire system (Mitchell 2009). Thus, phenomena treated from this perspective are considered to be non-reductive as interactions between components are as important as the unit's features. It is this intrinsically path dependent nature (i.e. where solutions in time t are determined by solutions in time t-1 as well as the feedback received) that makes problem-solving set-ups suitable to be treated as complex systems. It is also why they are often opaque to traditional approaches, i.e. making predictions about the overall performance of the system by looking exclusively at "the input". However, if one were to identify the attributes of the low level components and the rules governing their behaviour, via modelling and simulations, a better understanding of the emergent macro-phenomenon can be achieved and reliable predictions are possible. Certainly, as mentioned earlier, problem-solving occurs at this higher level where individual behaviours and attributes aggregate into a behaviour that is emergent in the sense that the crowd is able to outperform (in terms of speed and solution quality) its most proficient members (Lakhani et al. 2013). Individual characteristics and their rules of composition are of great importance as is the context in which they occur; this may influence negatively or positively their performance by activating

(or not) the agents' innate abilities (Felin 2012). This project focuses on individual attributes and how search behaviours aggregate to better understand how problem-solving processes can be described and also how these processes can be enhanced by creating proper governance settings under which such a system can perform optimally. In this approach I follow Felin and Zenger (2014) and use the problem and its attributes as the starting point of my analysis – i.e. each type of problem is assumed to have a corresponding optimal configuration of manageable design factors. However, I make no a priori assumptions as to which configuration is more suitable but attempt to answer that question by comparing results from running different virtual experiments of a computer model of problem-solving processes.

3.1. Problem-solving as search

By relying on the search metaphor for problem-solving, the project draws on Simon's work and the Carnegie School's subsequent efforts on bounded rationality alongside complementary insights from recent literature in cognitive science. In the following I provide a more detailed account of this approach as well as an overview of previous work.

Two ideas are recurrent in the literature when discussing problem-solving: "search" and "trial and error". They tend to be used together (i.e. "trial and error search") and sometimes used interchangeably and are so ubiquitous in organizational literature (but not limited to this realm) that it is useful to take a step back.

A recent paper traces the origins of "trial and error" in the mathematical domain (Cowles 2015). Cowles states that "trial and error" emerged as a synonym for "the rule of false positions" as a tool for students of mathematics¹ in the late eighteen century textbooks. "The rule of false positions" involves iterative numerical computations in solving an equation that might be otherwise cumbersome to solve algebraically. One can use several methods in the iteration process such as sampling, bracketing, etc. According to Cowles, by mid-1800 the idea had crossed disciplines and was to be found in psychology and biology as a way to describe both "mental life and life on earth" (Cowles 2015: p.637).

The origins of the "search" metaphor are attributed to William James (1890) who described an individual's memory search as a process similar to when "we rummage our house for a lost object" (James 1890: p.70).

The two ideas were brought together by Simon and Simon (1962) as they attempted to provide a formalized account of problem-solving in chess². In their 1962 paper, they show that even a simple computer program can reproduce the performance of expert chess players, without the

¹ The development of the method itself is most often attributed to ancient Egyptians.

² The prototypical environment for studying human problem-solving (cf. Simon and Chase 1973).

need for exceptional memory or intelligence, by resorting to a form of trial and error search “that incorporates powerful selective heuristics” (Simon and Simon 1962: p.429). This last statement is the computational equivalent to William James’ description of how human memory search “visits the probable neighborhood of that which we miss” (James 1890: p.70).

Thus, “search”, “trial and error” and subsequently “trial and error search” originally described situations in which solvers with limited information processing abilities used *heuristics* (be it the bracketing, the “program” the chess player constructs to restrict the search space from 10^{24} combinations to a more manageable number, or pre-existing memory associations that construct “the probable neighbourhood”). Trial and error search requires specifying the set of actions to be carried out as trials, whether all such actions will be conducted before feedback is integrated or the specification of an error: how should the agent interpret the feedback (Brenner 2006).

3.2. Model-based search

Despite being anchored in the literature highlighted above, for a long time the organizational literature modelling problem-solving as search was mainly influenced by the behaviourist perspective on learning. Scholars conceptualized and modelled search fundamentally on a stimulus-response basis, what computer science calls “reinforcement learning”. In this perspective solvers are assumed to search without having, developing or relying on internal representations of the problem they are solving, but rather proceed in a quasi-automatic manner following simple rules (Gavetti and Levinthal 2000, Winter et al. 2007). Thus, agents never resort to search heuristics but rely on experiential search or “ant-like” behaviour (Winter et al. 2007: p.405). Much like the behaviourist perspective (Anderson 2015), it was not that organizational scholars denied the existence of higher cognitive functions but rather the lack of empirical evidence and definite understanding of how these cognitive functions influence the problem-solving process that led them to restrict their models to simpler assumptions and mechanisms.

Indeed, even if the “cognitive revolution” has brought a focus on mental representations inspired by logic, computer science and linguistics (Barsalou 2008), their adoption was slowed down by the scarcity of empirical evidence supporting them as well as failing to make the link between cognition on the one hand and perception and action on the other (Barsalou 2008) as well as the actual biological structures that are responsible for these encodings and their manipulation (Schultz et al. 1997; Gläscher et al. 2010; Solway and Botvinick 2015). This was further fuelled by the fact that the existing mathematical apparatus was developed in physics or biology and agents in these disciplines (i.e. organisms and particles) are neither heterogeneous nor self-reflexive (Castellano et al. 2009).

However, in recent years, cognitive science and neuroscience developments have made substantial progress in advancing our understanding of the brain (Anderson 2015). Brain processes are no longer a “black-box”, and models can be expanded to include more cognitively realistic assumptions regarding human cognition since reducing human behaviour to basic *biological* evolutionary mechanisms is not without consequence. Thus, current research acknowledges that “long-term goals and expectations play a leading role” in shaping the search process (Solée et al. 2013: p.1): introducing model-based “abilities” to a searching agent drastically changes search dynamics and performance. The internal representations that an agent constructs and rely on are extremely important as they can help the agent to 1) select which information (from the environment) he is going to allocate attention to, but also 2) since most likely the current input is not enough to decide what to do next, they provide a framework where the agent can store relevant aspects from the past that could help in the future (Dayan 2012), and 3) the models allow for making predictions (Daw et al. 2006; Solway and Botvinick 2012) and thus what actions are to be taken next, a form of “offline search” (Winter et al. 2007). As Schooler et al. (2012) argue “this type of representational change seems to be the main reason why cognitive search can be more efficient than machine search” (Schooler et al. 2012: p.320). This project takes a few modest steps in uncovering the role, the mechanisms and determinants of model-based search in problem-solving by drawing from three quite different literature streams: organizational science, computer science and cognitive science.

3.3. Research methods

Overall, this project follows the approach recommended by Afuah and Tucci (2012) who propose the use of simulation modelling as a research method for investigating the phenomenon of crowdsourcing. The use of the method is justified by the low availability of longitudinal data (Davis et al. 2007; Harrison et al. 2007) but also by the particular theoretical focus which is longitudinal and nonlinear, dealing with processes rather than outcomes (Davis et al. 2007). Using an agent-based model offers both the advantages of conducting a qualitative study (i.e. allows for a greater level of detail and answers to the “HOW” question), but also that it allows for a high degree of generalization (i.e. moving away from a particular empirical setting). Agent-based simulations can provide both precision and a way to study directly how individual elements interact to create emergent phenomena (Kozlowski and Chao 2012).

In effect, via computational modelling virtual experiments can be conducted, hypothesis tested and simulations can be used as a powerful theory development tool (Davis et al. 2007; Harrison et al. 2007). Due to the complexity of the phenomenon studied, computer simulations often lack transparency – i.e. given that there are numerous non-linear interactions between different factors, the results of a simulation are not normally intuitively understandable. Therefore it is

crucial for the quality of the theory that the assumptions built into the model are realistic. The link between the model and the empirical phenomenon studied (i.e. avoiding the “toy model” fallacy) is created by a) choosing the right simulation approach in such a way that it embeds the initial framework and the fundamental assumptions – section 3.3.1 Simulation framework – as well as by b) validating simulation *assumptions* against empirical data (i.e. model calibration) – sections 3.3.2. Experiments and 3.3.3. Survey and qualitative interviews. Thus, in addition to modelling and simulation, I conducted three experimental studies, which informed the second and third essays, as well as 40 qualitative interviews (third paper). In line with Fioretti (2013), I have also attempted to validate simulation results by showing they are consistent with corresponding empirical observations.

3.3.1. Simulation framework

According to Newell and Simon “a theory of thinking and problem solving cannot predict behaviour unless it encompasses both an analysis of the structure of task environments and an analysis of the limits of the rational adaptation to task requirements” (1972: p.55).

One particular conceptualization of a problem space has gained traction in the organizational literature: fitness landscapes. Sewall Wright (1932) first introduced the idea of a “fitness landscape” as way of conceptualizing how populations of varying genetic mark-ups would “by trial and error” (Wright 1932: p.359) (i.e. in this context, natural selection) evolve into different stable combinations of high adaptiveness. Thus a fitness landscape is essentially a multi-dimensional mapping between genotype and phenotype, assuming similar genotypes are to be found closer together.

Kauffman (1993) took Wright’s idea further and advanced a mathematical model (i.e. the NK model) that allows inferences to be made about a fitness landscape based on the size of the organism’s genotype and the level of epistasis between its genes.

In the NK model, any solution to a problem (i.e. an organism) can be described by a number of attributes (i.e. the genotype), specific to the task at hand, with a number of values¹ for each attribute. The problem space is given by all possible combinations of attributes, and the successful search for a solution to a problem is equal to finding the “best combination of values”. The NK model is thus defined by N, the number of attributes a problem has and K, which is the number of elements that affect the functioning of each attribute. As such, the NK model can account for the effect that specific combinations of attribute values may have on the quality of the solution, for instance, it can be used to model conflicting designs or positive complementarities. The interplay between N and K gives rise to “smooth” or “rugged landscapes”. As such, problems within a low dimensional space (low N), with low values of K,

¹ Traditionally two.

are considered to be simple to moderately complex and easier to solve than problems within larger dimensional spaces with high values of K (with the limit case K=N-1).

The NK model is by far the most common computational model in the innovation management literature and following the introduction by Levinthal (1997), several variants have been built to allow for more complex hypothesis to be tested. For instance, Lazer and Friedman (2007) study the types of networks that are more efficient in problem-solving using such a model that includes a time dimension. They establish that the relationship between network efficiency (i.e. communication) and performance depends on the time scale of the analysis, but ultimately, in the long run high connectivity drives diversity out. In a different attempt to add more realistic assumptions to the NK model, Hebron et al. (2008) take point of departure in observed network structures (in genetics) and explore the search patterns that emerge. Their simulations showcase that imposing a scale free structure on an NK landscape leads to longer paths towards the solution and a higher degree of clustering of high fitness solutions in the landscape.

The main challenge with papers built on the NK model is that they cannot overcome the initial assumptions of the canonical model: in most papers agents can only learn via (different types of) experiential learning. As Winter et al. (2007) argue, if employed, cognitive search is relatively distant from the empirical phenomenon. For example, in Gavetti (2005) agents are simply given a set of possible problem representations from which they choose – “a cognitive memory” – while in Winter et al. (2007) the cognitive influence is “independent of experience” (Winter et al. 2007: p.404).

H-XOR

Following these considerations, the underlying function in my experimental set-up is an 8-bit “hierarchical exclusive OR” (H-XOR) (Watson and Pollack 1999). I chose this function because unlike the traditional functions used to generate NK landscapes (Kauffman 1993), the solution space is structured in several interdependent (hierarchical) modules. This is in line with the view that most problems are nearly decomposable (Simon 1962). In particular, innovation problems are closer to a hierarchical structure (Marengo 2014) in which the optimization of any given module needs to take into account all other modules, but this also “ripples down” the hierarchy.

Specifically, the H-XOR function is given by applying recursively an ‘exclusive or’ transformation onto the solution string where adjacent positions are considered starting with the leftmost. For instance, a {1010 0010} string becomes first {11 -1} and then {- -}. Once the transformation is completed, the payoff function rewards each non-null position in the hierarchy. Thus, a solution which contains an alternating pattern {1010 1010} would give a better score than a {1111 1111} since it will generate payoffs at lower levels of the hierarchy as well. The

second level transformation for the first solution is {11 11} while for the second it is {---}. The maximum score is given by {1001 0110} or symmetrically by {0110 1001}.

H-XOR's symmetrical structure also makes it more amendable for the experimental study of organizational problems, as opposed to the NK framework where the values of the fitness function are drawn randomly from a uniform distribution. In contrast, in the H-XOR setting there is an underlying pattern that they can meaningfully extrapolate from to other areas of the landscape that they have not navigated yet. This is in line with Simon (1996) who argues that, in learning, “meaningfulness is a variable of great importance. Nonsense syllables of high association value and unrelated one-syllable words are learned in about one-third of the time required for nonsense syllables of low association value”(Simon 1996: p.65).

Search algorithm

The model-based search algorithm relies on an evolutionary search heuristic developed by Iclanzan and Dumitrescu (2007). Technical details aside, the algorithm relies on a simple associative learning rule and is similar to Hayes and Simon's (1974) description of problem-solving where a solver first begins by creating a problem representation and generates a solution which he tries to improve. When a solver's memory is full, he uses the information he has collected in the search process thus far to generate a different problem representation. The agents start without any problem representation so they use their first moves to acquire a basic “understanding” of the problem. After a number of random moves, agents engage either in local search or in model-based distant search. In their model-based search they attempt solutions that incorporate previously gained knowledge about optimal configurations of the components they have identified as interdependent. In their local and model-based distant search, the agents are satisficing using their own problem representation to calculate their score. That is, they attempt to find a solution that outperforms their best score so far according to their own problem representation. If the resulting score is an improvement, they add the respective solution to their memory. Agents receive real feedback after submitting a solution and this feedback is used to compute their best score so far.

3.3.2. Experiments

Much of the organizational literature regarding the NK model is conceptual with few notable exceptions: (Fleming and Sorenson 2003; Mason and Watts 2012; Billinger et al. 2013). In fact, in an extensive overview of the literature relying on statistical physics to model social phenomena, Castellano et al. (2009) note that “there is a striking unbalance between empirical evidence and theoretical modelization in favour of the latter” (Castellano et al. 2009: p.3). Their work does not explicitly review NK models (although NK models can be thought of as a special

case of the Ising model), but the conclusion can be thought to hold more generally for such models in social science irrespective of their origin.

Taking organizational learning more broadly into account, the exploration-exploitation paradigm has received more attention, especially in recent years. In particular, recent work has looked into cognitive processes involved in decision-making under uncertainty within a specific experimental set-up – the armed bandit game (Daw et al. 2006). In a multi-armed bandit scenario, participants are asked to solve the task of maximizing the sum of rewards in a sequence of trials. At each time step they can choose one of several (two or more) “slot machines” which give a reward chosen randomly from a fixed distribution unknown to the solver. This allows for a straightforward operationalization of March’s exploration-exploitation dilemma (1991). Thus, in each trial the participant has to make the choice whether he wants to continue playing the same slot machine or switch to a different one to acquire more information. In this paradigm, solvers are assumed to develop an internal representation regarding the different underlying reward distribution. Several cognitive models have been developed based on classic decision theory – e.g. expectancy valence learning model and the prospect valence learning model (Dai et al. 2015). Still, this paradigm seems more adept at capturing the stylized exploration-exploitation behaviour (March 1991), where solvers are required to decide between a limited number of actions rather than actual problem-solving behaviour which also involves option generation (Smaldino and Richerson 2012).

Instead, the experimental set-up that informs this project allows solvers an opportunity to generate or shape new solutions. At the same time the search space is clearly defined (in terms of the number of options and the efficient way of navigating towards the optimal solution).

Aside from the fitness function (i.e. H-XOR described earlier), the set-up is similar in nature to the one developed by Billinger et al. (2013). It involves asking participants to solve a combinatorial problem presented in the form of a game with eight different tiles. The tiles can be toggled on and off (i.e. they have two possible values) leading to a problem space of 256 combinations. The participants were asked to find the combination of tiles that yields the highest score within a maximum of 25 attempts. They were informed in advance of the maximum score to reduce variation in search strategies due to uneven performance expectations. They were further instructed that they would be rewarded based on their performance (i.e. maximum achieved score and speed – in terms of number of submitted solutions).

Three separate studies were conducted: the first one involved 200 participants in the Cognition and Behaviour Lab (COBE lab in the following) at Aarhus University and represents the dataset for the second paper. The second study was conducted using Mturk, an online crowdsourcing

platform. Amazon's Mechanical Turk (Mturk) is a web-based outsourcing platform where 'requesters' set up various tasks and 'workers' select which tasks to complete for a fee corresponding on average to 6\$/hr. Its use in behavioural research is on the rise due to the easy access to a relatively large group of people that resembles the general (US) population (Berinsky et al. 2012) and the relatively low costs in set-up (Mason and Suri 2012). 270 participants completed the task in this setting. Finally, the third study was conducted in the COBE lab with a smaller sample (97 participants). The second and third studies were used in the third paper.

3.3.3. Survey and qualitative interviews

In addition to the game-play data, the third paper relies on a mixed-methods approach. Participants were first asked to fill in a 12-item survey aimed at capturing their cognitive style according to the Adaptors-Innovators theory (Kirton 1976). The 12-item scale is a reduced version of the original scale developed and validated in Miron, Erez et al. (2004), which measures the three interrelated dimensions of cognitive styles (creativity, attention to detail and conformity) independently. To complement (cf. Greene et al. 1989) 40 semi-structured interviews were collected from the COBE lab participants. In this, I follow Jack and Roepstorff (2002) who argue that, although introspection has received a lot of criticism as a method of scientific enquiry with respect to the validity of results, it is a reliable tool in the investigation of, among other, problem-solving processes which are still inscrutable by other methods (e.g. fMRI). Furthermore, by interviewing participants after completing the game without using a "think aloud" protocol, however, our approach is consistent with Ericsson and Simon (1998) who suggest that asking participants to verbalize problem-solving strategies can often (artificially) influence their behaviour. The interviews complement (Greene et al. 1989) and validate our quantitative search constructs and generate further qualitative insight into the behavioural search process.

Section 4: Overview of the thesis

Paper 1: Fitness landscapes in organizational theory: research challenges and future directions

The relatively recent and limited use of models in management has prompted few discussions with respect to their validity. Indeed, unlike other social science disciplines (e.g. economics), the use of modelling and simulations in management or organizational literature is still a niche practice. One recent paper discussing the use of modelling and simulations argues for the importance of testing model assumptions but continues by saying that indeed this is also a matter of "social acceptance" (Fioretti 2013). Understandably, when we are talking about aggregate

behaviours at the level of organizations, testing assumptions is cumbersome (if indeed possible), but at least when it comes to new organizational forms (e.g. innovation contests), which deal mainly with individual-level behaviour, by explicitly considering the plausibility of the assumptions embedded in the model, modelling and simulations can be more than shelf models and can make a meaningful and falsifiable contribution (Pfleiderer 2014).

In this paper we attempt to provide a formal background for the importance of assumption testing with respect to human search behaviour and a deeper focus on the fitness landscape features as well as a brief discussion on potential avenues for future models. In particular, we focus on a number of classical assumptions used in models in organizational theory and, via modelling and simulation, explore potential avenues for expanding the traditional models to explicitly address individual problem-solving behaviour.

We start by providing a formal definition for a fitness landscape and show how the landscape features captured by the NK model are 1) not exogenous to the search behaviour and 2) not intrinsically connected to organizational environments. We then turn to search behaviours and show how simple and uncontroversial extensions of current assumptions regarding problem-solving behaviour drastically change problem-solving dynamics and performance. In particular we argue that search behaviour is endogenous to the solution space and any model about search behaviour should start by qualifying assumptions about both the task environment and search behaviour. Furthermore, we show how qualitatively different results can be achieved by modifying assumptions about agent memory or the ability to detect and exploit patterns about the structure of the problem.

Paper 2: Searching far away from the lamp post, an agent-based model

This paper presents insights from a laboratory experiment on human problem-solving in a combinatorial task. Using the H-XOR landscape, I explore whether and how human problem-solvers are able to detect and exploit patterns in their search for an optimal solution. Empirical findings suggest that solvers do not engage only in local and random distant search, but as they accumulate information about the problem structure, they make “model-based” moves, a type of cognitive search. I then calibrate an agent-based model of search to analyse and interpret the findings from the experimental set-up and discuss implications for organizational search.

The data collection for this second paper was conducted in the Cognition and Behaviour Lab at Aarhus University. 200 participants were recruited using the lab’s pool of participants (mainly students at Aarhus University). They were asked to play two games in succession. One of them was based on a hierarchical function and the other was based on an NK function; the order of the

two games was randomized. This allowed the validation of the experimental set-up against the one used in Billinger et al. (2013).

One of the main strengths of this paper is that it brings forward a novel approach to problem-solving research that takes into account the dynamic and complex nature of the phenomenon, while at the same time being informed by empirical research, as previous research on search and problem-solving has been either purely conceptual or has relied on either qualitative research or quantitative designs.

The overall approach is bottom-up: first, the experiments allowed the study of human problem-solving behaviour, and second, the simulations attempted to qualify how different search strategies affect the propensity for finding the optimal solution. Even if the model is agnostic with respect to intrinsic or extrinsic motivation, I find that focusing on immediate results (by not engaging in more risky, distant moves) has the potential to diminish learning and the likelihood to solve complex problems.

Consistent with previous work on the balance between exploration and exploitation (March 1991), I find that solvers are more likely to engage in risk-seeking behaviour if they consistently perform below their aspiration level. However, I also distinguish a third behaviour that solvers engage in: model-based distant search, i.e. that they attempt to identify and exploit patterns as they navigate the rugged problem landscape. This different type of distant search does not seem to be influenced directly by feedback.

The findings from the experimental study were used in calibrating an agent-based model of adaptive search: First, I explored whether indeed model-based distant search can be a reliable tool in the search for the optimal solution. Results show that local search can be an efficient tool for simple problems and can lead to minor quick improvements for more complex problems. Still, in the case of (even mildly) rugged landscapes, relying on an exclusively local search strategy is inefficient. Additionally, a naïve local search and random algorithm was also outperformed by the model-based search. Secondly, I looked into what determines the success or failure of model-based distant search. The results suggest that a lower level of “persistence” (the willingness to abandon earlier on unsuccessful paths) translates into higher accuracy in terms of problem representation and improved performance. This last result is reminiscent of the “fail early, fail often” mantra of one of the most successful design companies (IDEO), an idea also captured in the information systems practice (and literature) by agile software development (Martin 2003), or scrum (Conboy 2009).

Paper 3: Micro-foundations of problem-solving: what determines how individuals search?

This last study contributes to the discussion on innovation as search by focusing on how people search and in particular on what individual antecedents determine search behaviour. Relying on a mixed-methods approach that combines an experimental set-up with surveys and interviews, we identify heterogeneous search behaviours and find evidence that individual cognitive styles and how players form and adapt internal representations can explain part of this heterogeneity.

The data collection for the third paper involved two different samples: 270 participants were recruited on Mturk (an online crowdsourcing platform) and 97 participants were recruited using the COBE lab pool of participants¹. From the second sample, 40 participants were selected upon completing the game in the COBE lab and asked to participate in the interviews. 20 participants were selected randomly during regular experimental sessions, and 20 participants were interviewed in dedicated experimental sessions where all participants were interviewed. All experimental sessions were conducted by research assistants unaffiliated with the research and who were not informed about the specific research questions and hypothesis to be investigated. The qualitative interviews were conducted by the two authors of the paper.

The first contribution in this study is explaining how individual differences with respect to information processing, i.e. cognitive styles, predict the propensity with which solvers engage in particular search behaviours. We find that participants who score high on the creativity dimension are more likely to engage in more random search and less likely to engage in local search. Players who score low on attention to detail are more likely to engage in more random search. We attribute these results to the fact that the Adaptor-Innovator constructs aim to capture a preference for structure in problem-solving (Jablokow et al. 2015).

A second contribution of this study is the focus on individual level model-based search (Doll et al. 2012). In particular we explore how internal representations are formed and used. Our analysis of the game-play indicates that players rely on top-bottom processing, but the interviews also reveal bottom-up processing (Anderson 2015), i.e. internal representations are shaped based on feedback as well as cultural expectations regarding such puzzles. Having a top-bottom internal representation is important because a) they suggest an alternative source of heterogeneity; even with similar cognitive styles, solvers have different starting points thus different paths (Solway and Botvinick 2012), and b) another point of departure from how literature has conceptualized search so far since most simulations start at a random or tabula rasa.

¹ Participants who had taken part in the first study were automatically excluded from the second.

Our findings have implications for conceptual work relying on modelling and simulation since we show that a choice between local and non-local search is not determined solely by structural considerations (e.g. the nature of the problem to be solved) but also by an individual's preferred way of processing information – i.e. the behavioural process (cf. Csaszar and Levinthal 2015). Relating this individual level to analysis of firm-based search (Lopez et al. 2016) could shed further light on how and why firms make strategic search decisions. Additionally, we show how traditional assumptions about how simulated agents either do exploration (random jumps) or exploitation does not capture how players in our game reflect about their search strategies and engage in meaningful, distant search.

4.1. Overview

My aim with the dissertation was to explore how individual search behaviours take place in the particular context of repeated-entries contests. This overall research question was based on the recent attention to distributed innovation (e.g. Afuah and Tucci 2012; Lakhani et al. 2013) as well as a rising focus on individual search behaviours (Billinger et al. 2013). In line with recent research focusing on the organizational level (Lopez et al. 2016), I suggest that the lack of focus on disentangling the various ways in which individuals engage in distant search is detrimental since it obscures an important mechanism by which individuals attempt to find solutions to a given problem.

It is widely accepted that innovation is a path-dependent process (Dosi 1988; Tidd and Bessant 1997; Drucker 2014). This dissertation is grounded in the innovation as search literature, in particular the modelling and simulation literature. The evolutionary models of innovation allow for the time dependent nature of problem-solving to be captured and explored and have yielded a number of results significant for organisational theory (e.g. Levinthal 1997; Rivkin 2000; Rivkin and Siggelkow 2003). However, in this literature, assumptions about stylized agents and behaviour can be problematic and lead to results which are not supported by empirical evidence (Billinger et al. 2013; Mason and Watts 2012). I have explored a number of such assumptions in my first paper. I maintain that the NK model has several limitations that make it difficult for organizational theorists to model meaningful problem-solving processes (be it at the organizational or individual level). “Essentially, all models are wrong, but some are useful” (Box and Draper 1987) is a commonly used quote in the modelling literature. However, models are only useful as long as they represent meaningful simplifications of the phenomenon to be modelled. I argue that there are still too many legacy elements in the way we model problem-solving and not enough accurate assumptions about what happens “between the ears” of a solver to be able to make the leap from ‘shelf models’ to models that can provide meaningful and falsifiable contributions to the literature on problem-solving (Pfleiderer 2014). The theoretical

analysis supported by simulation results suggest that the NK model is limited in terms of providing a meaningful (for organizational problems) problem structure as well as the possibility to implement more cognitively plausible search behaviours. These limitations impact our ability to capture relevant dynamics of individual problem-solving behaviour.

Indeed, one of the main limitations of previous research is its lack of focus on cognitive search (Csaszar and Levinthal 2015; Lopez et al. 2016). This is, in part, due to the limitations of the NK model, in part because previous research has mainly been focused on the organizational level. In this project, I put forward a novel research design that attempts to address some of the limitations of previous work. In particular, by relying on a different class of problems (H-XOR), where fitness values are not random but form a well-defined (albeit not obvious) hierarchical pattern, I attempt to capture in an experimental setting how human players form and develop internal representations of the problem and how these representations, in turn, affect their searching behaviour. By combining the data relative to the type and propensity for search behaviours with a simulation framework, I investigate implications for the innovation as search literature. More specifically, in the data collected in the first experiment I find evidence that human solvers engage in model-based search alongside local and distant-random search. By focusing on traditional determinants of search behaviours (i.e. time and feedback), I extract a number of stylized rules about agent behaviour, which, in conjunction with an analysis of performance and distribution of search behaviours, form the basis for the calibration of the simulation model. By running virtual experiments of the simulation model, I then investigate different propositions regarding model-based search and its efficacy in problem-solving set-ups. First, I show that local search can be an efficient tool for simple problems and can lead to minor but quick improvements for more complex problems. Still, in the case of (even mildly) rugged landscapes, relying on an exclusively local search strategy is inefficient. Second, simulations endorse empirical studies suggesting that for complex problems concentrating on immediate performance (and focusing on either local or model-based moves) increases the likelihood that solvers focus too little on learning (Manso 2011; Gardner 2012) and consequently have an increased likelihood of making spurious correlations which trap them into suboptimal peaks.

These findings motivated the question: if depending on the environment, certain search behaviours are preferable, what determines or rather what are the sources of heterogeneity in human search behaviour? Although there are scholars who note that models in organizational search literature are “remarkably non-organizational” (Knudsen and Levinthal 2007: p.39), studies do focus primarily on organizational search (Audia and Goncalo 2007) and explanations of heterogeneity in search behaviours are mainly structural (Raisch et al. 2008). Arguably, heterogeneity in terms of search behaviours has the potential to greatly influence the outcome of

the search process, as shown by Mason and Watts (2012). The last paper of this dissertation makes preliminary steps in identifying some potential sources of heterogeneity in human problem-solving behaviours. In particular, I focus on cognitive styles as they are linked to individual differences in how they perceive and process information (Sternberg and Grigorenko 1997; Miron-Spektor et al. 2011) and on the role of internal representations as an additional source of variance. The findings in this last study have implications for conceptual work relying on modelling and simulation since I show that a choice between local and non-local search is not determined solely by structural considerations (e.g. the nature of the problem to be solved or incentives) but also by a substantial variation in individuals' preferred way of processing information. In addition, even with similar cognitive styles, I find that solvers have different starting points and ways of integrating evidence: some relying stronger on priors, others were more adept at incorporating feedback. This is yet another departure from how the modelling literature has conceptualized search so far and another important cause of variability for search behaviours (Solway and Botvinick 2012).

4.2. Methodological contributions and limitations

Paradoxically, the biggest methodological strength of this project is also its biggest weakness: the experimental set-up. This dissertation aimed to bring forward a relatively novel approach to problem-solving research that takes into account the dynamic and complex nature of the phenomenon while at the same time being informed by empirical research. With a couple of notable exceptions (i.e. Billinger et al. 2013; Mason and Watts 2012), previous research on search and problem-solving has been either purely conceptual or relied on either qualitative research or quantitative designs. My use of method was thus motivated by the fact that using an agent-based model calibrated with results from an experimental study offers the advantages of getting a more detailed understanding of how people search as well as the opportunity to a) compare models of search behaviour with experimental results, and b) conduct virtual experiments with different parametrizations of the model.

Still, the empirics in this project are exclusively collected in an experimental set-up and this has obvious shortcomings: it is hard to argue that the theoretical (via simulation) and empirical findings should be extended beyond the particular environment where they were captured. Nevertheless, the fact that playing a browser game is not as different from the experience of taking part in an online innovation challenge and the fact that similar results were obtained in both the online (Mturk) and the laboratory study are grounds for optimism.

Thus, this dissertation opens up a number of avenues for future research. Fortunately, unlike more traditional innovation contexts, innovation contests provide a rich and easy (or at the very least “easier”) opportunity to access longitudinal data sources since relevant behaviours are

automatically recorded (e.g. submissions, interactions) as is the exact time and sequence. In particular, programming contests such as the ones featured by TopCoder or data-science platforms such as Kaggle could provide a fruitful environment since problems and solutions in these settings are more easily mapped in terms of changes from one iteration to the following and allow for more objective measures in terms of performance. This could enable several hypotheses derived from this dissertation to be tested, using traditional research methods. First, do contributors to such repeated-entries innovation contests indeed engage in a variety of search behaviours adapted to the feedback they receive? Second, can we find evidence of solvers detecting and exploiting problem structures outside the confinements of the stylized task provided in the lab? And finally, can we extrapolate findings regarding particular cognitive styles determining particular search behaviours in such settings?

A second limitation of this study is the operationalization of search behaviours. Although individual search behaviours are not as difficult to capture as their organizational counterparts and although I find both quantitative (paper 2) and qualitative evidence (paper 3) that solvers do detect the underlying problem structure and attempt to tailor their search behaviours accordingly, there are a number of important variables still left unexplained. My own qualitative interviews as well as recent research in cognitive science (Huys et al. 2012) suggest that human beings rely on more sophisticated evidence integration mechanisms and do not pay exclusive attention to their best solution so far nor their most recent attempt. An important step in being able to unequivocally identify individual search behaviours should then be to disentangle what players pay attention to *before* submitting a new attempt. Extending the experimental set-up to incorporate eye-tracking should allow this investigation to be carried out.

4.3. Managerial implications

Despite its limitations this dissertation holds several implications. First, this project finds support for the fact that in the context of problem-solving, performance pressure can indeed be a “double edged sword” (Gardner 2012). This complements existing work regarding incentive systems for innovation that suggests that failure should be tolerated in highly innovative settings (Holmstrom 1989). “Fail early, fail often” is a phrase that has gathered traction in recent years as the mantra of one of the most successful design companies in the world (IDEO). This idea is captured in the information systems practice (and literature) by agile software development (Martin 2003), scrum (Conboy 2009) etc. Indeed, my own results accentuate that a lower level of “persistence” (the willingness to abandon earlier on unsuccessful paths) translates into higher accuracy in terms of problem representation and improved performance.

Second, for solving complex problems, there seems to be evidence that relying on multi-shot submissions (a set-up where solvers receive feedback from the broadcasting organization and can subsequently modify their solutions) is preferable to the classical broadcast search (Vuculescu and Bergenholz 2014), and certainly there is little controversy regarding the usefulness of feedback in learning environments (Bangert-Drowns et al. 1991). The current study indirectly makes the same argument: without having the possibility to receive and incorporate feedback, solvers cannot form an accurate problem representation and their performance depends on their prior knowledge (which is what the traditional Innocentive challenge is trying to exploit).

Finally, based on a relatively short survey (the A-I scale), owners of a crowdsourcing platform could map the cognitive styles of its current users and relate this mapping to the current problem to be solved. Following Felin and Zenger's (2014) conceptualization of the problem-solving organization, the problem type and the degree of hidden (to the firm) knowledge required to solve it, determines the search methods that should be employed. For example, they argue that user and user communities might be particularly adept at solving high complexity problems, with a high degree of hidden knowledge since these problems require open, theory-guided (i.e. non experiential) search. However, I suggest that this might need further qualification: does one have the users likely to engage in the relevant kind of search? Alternatively, one could attempt to frame the problem differently in order to better match the relevant search behaviours and the problem to be solved.

Although these results do not contribute directly to the ongoing discussion about prizes on crowdsourcing platforms, e.g. Archak and Sundararajan (2009), they do offer managers a different lens in making that decision.

4.4. General discussion and ways forward

The role of the manager is to design forms of governance that can integrate the required knowledge to solve a given problem (Nickerson and Zenger 2004). Nickerson and Zenger (2004) argue that the process of knowledge development is “problem-solving”: “if a firm is to develop unique knowledge [...] through any manner other than luck, it must identify a valuable problem and conduct an efficient solution search” (Nickerson and Zenger 2004: pp.618-619). Their argument follows Simon (1991) who proposes that organizational learning literature should draw on research and terminology related to human learning and problem-solving, given its intrinsically individual level focus – “all learning takes place inside individual human heads” (Simon 1991: p.125). Of the three paradigms for addressing problem-solving (i.e. reasoning, constraint satisfaction and search) Simon (1983) reasons that search is particularly appealing since it allows for problem-solving to be conceptualized while simultaneously accounting for

bounded rationality (i.e. solvers do not have an a priori understanding of the problem to be solved nor of the full set of possible solutions). By drawing on recent developments from cognitive science and psychology and using a research design more commonly used in those fields, this dissertation follows the Simonian path and tries to answer a deceptively easy question: how do people search?

Despite the early and urgent call for focus on micro-foundations (e.g. Grant 1996), somewhat surprisingly we do not quite know *how* people solve problems. I argue that answering the small question lays the foundation for the broader (and arguably more relevant for organizational studies) question “how do organizations search”. Undisputedly how micro-level behaviours aggregate into macro-level patterns is not trivial, i.e. even if our understanding of how individuals go about solving problems was perfect, it would still leave the broader question unanswered. How people share and integrate knowledge (acquired by themselves or shared by others), whether they might be biased in their problem-solving paths by others, by hierarchies or other organizational governance elements, these are important and only partially answered questions. This dissertation stays clear from the larger question and makes modest progress with the smaller one: “how do people solve problems”.

In truth, my project raises more questions than it answers, but I believe it opens up a way to address significant questions that have the potential to further our understanding of how the search for solutions takes place. My overall approach is not new, as evidenced by the numerous references to Herbert Simon and work that has been carried out more than 30 years ago. It is, however, the recent developments in the ability to conduct experiments online (Mason and Suri 2012) and recent increase in internet-based communication and interactions (Watts 2007) as well as the unprecedented (relative) ease for a social scientist to conduct complex virtual experiments that greatly facilitate such an approach.

Finally, although my project focuses on individual problem-solving, both the model and the experimental set-up can easily be extended to incorporate collective problem-solving. As such, one can address questions such as: How do solvers shape their search paths when information from other players is revealed? Do they incorporate merely the information or do they pay attention to the strategy that generated the observed solution? Given the importance of reputation systems on online platforms (Dellarocas 2003), an important question is what is the role social cues play in collective problem solving?

In “Sciences of the Artificial” (1996), Simon recounts the story of the Dutch mathematician Simon Stevin who upon discovering the mechanics of the inclined plane “was so pleased with his construction, he incorporated it into a vignette, inscribing above it: ‘Wonder, en is gheen wonder’, that is to say ‘wonderful, but not incomprehensible’” (Simon 1996: p.4). It is the task of the “artificial” sciences “to show that the wonderful is not incomprehensible” (Simon 1996: p.4).

Ten years later Duncan Watts echoes this optimism: “three hundred years after Alexander Pope argued that the proper study of mankind should lie not in the heavens, but in ourselves, we have finally found our telescope. Let the revolution begin....” (Watts 2011: p.300)

References

- Abernathy, W. J. and K. B. Clark (1985). "Innovation: Mapping the winds of creative destruction." Research Policy 14(1): 3-22.
- Adamczyk, S., A. C. Bullinger, et al. (2012). "Innovation contests: A review, classification and outlook." Creativity and Innovation Management 21(4): 335-360.
- Afuah, A. and C. L. Tucci (2012). "Crowdsourcing as a solution to distant search." Academy of Management Review 37(3): 355-375.
- Archak, N., & Sundararajan, A. (2009). “Optimal design of crowdsourcing contests”. ICIS 2009 Proceedings, 200.
- Arthur, W. B. (1988). Self-reinforcing mechanisms in economics. The economy as an evolving complex system, 5, 9-31.
- Audia, P. G., & Goncalo, J. A. (2007). Past success and creativity over time: A study of inventors in the hard disk drive industry. Management Science, 53(1): 1-15.
- Auh, S., & Menguc, B. (2005). Balancing exploration and exploitation: The moderating role of competitive intensity. Journal of Business Research, 58(12): 1652-1661.
- Baldwin, C. and V. E. Hippel (2011). "Modeling a paradigm shift: From producer innovation to user and open collaborative innovation." Organization Science 22(6): 1399-1417.
- Bangert-Drowns, R. L., Kulik, C. L. C., Kulik, J. A., & Morgan, M. (1991). “The instructional effect of feedback in test-like events”. Review of educational research, 61(2), 213-238.
- Barsalou, L. W. (2008). "Grounded cognition." Annu. Rev. Psychol. 59: 617-645.
- Baumann, O., & Siggelkow, N. (2013). “Dealing with complexity: Integrated vs. chunky search processes.” Organization Science, 24(1), 116-132.
- Bayus, B. (2013). "Crowdsourcing new product ideas over time: An analysis of the Dell Ideastorm community." Management Science 59(1): 226-244.
- Benkler, Y. (2002). "Coase's Penguin, Or, Linux and the nature of the firm " The Yale Law Journal, 369-446.
- Bennett, J. and S. Lanning (2007). The netflix prize. Proceedings of KDD cup and workshop.
- 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., N. Stieglitz, et al. (2013). "Search on rugged landscapes: An experimental study." Organization Science.

- Bogers, M. and J. West (2012). "Managing distributed innovation: Strategic utilization of open and user innovation." *Creativity and Innovation Management* 21(1): 61-75.
- Bonabeau, E. (2009). "Decisions 2.0: The power of collective intelligence." *MIT Sloan Management Review* 50(2): 45-52.
- Boudreau, K. J. (2012). "Let a thousand flowers bloom? An early look at large numbers of software app developers and patterns of innovation." *Organization Science* 23(5): 1409-1427.
- Boudreau, K. J., N. Lacetera, et al. (2011). "Incentives and problem uncertainty in innovation contests: An empirical analysis." *Management Science* 57(5): 843-863.
- Boudreau, K., Helfat, C. E., Lakhani, K. R., & Menietti, M. E. (2012). "Field evidence on individual behavior & performance in rank-order tournaments". *Harvard Business School Working Paper*, No. 13-016, August 2012
- Boudreau, K. J. and K. R. Lakhani (2013). "Using the crowd as an innovation partner." *Harvard Business Review* 91(4): 60-69.
- Box, G. E. P., and Draper, N. R., (1987). Empirical model building and response surfaces, *John Wiley & Sons*, New York, NY.
- Brabham, D. C. (2008). "Crowdsourcing as a model for problem solving." *Convergence: The International Journal of Research into New Media Technologies* 14(1): 75–90.
- Brenner, T. (2006). "Agent learning representation: advice on modelling economic learning." *Handbook of computational economics* 2: 895-947.
- Bullinger, A. C., A.-K. Neyer, et al. (2010). "Community-based innovation contests: Where competition meets cooperation." *Creativity and Innovation Management* 19(3): 290-303.
- Castellano, C., S. Fortunato, et al. (2009). "Statistical physics of social dynamics." *Reviews of Modern Physics* 81(2): 591.
- Chesbrough, H. (2003). *Open Innovation: the new imperative for creating and profiting from technology*. Boston, Harvard Business School.
- Cohen, J. D., McClure, S. M., & Angela, J. Y. (2007). "Should I stay or should I go? How the human brain manages the trade-off between exploitation and exploration". *Philosophical Transactions of the Royal Society B: Biological Sciences*, 362(1481): 933-942.
- Conboy, K. (2009). "Agility from first principles: reconstructing the concept of agility in information systems development." *Information Systems Research* 20(3): 329-354.
- Cowles, H. M. (2015). "Hypothesis bound: Trial and error in the nineteenth century." *Isis* 106(3): 635-645.
- Dai, J., Kerestes, R., Upton, D. J., Busemeyer, J. R., & Stout, J. C. (2015). "An improved cognitive model of the Iowa and Soochow Gambling Tasks with regard to model fitting performance and tests of parameter consistency". *Frontiers in psychology*, 6.

- Davis, J. P., K. M. Eisenhardt, et al. (2007). "Developing theory through simulation methods." *Academy of Management Review* 32(2): 480-499.
- Dahlander, L., & Piezunka, H. (2014). Open to suggestions: How organizations elicit suggestions through proactive and reactive attention. *Research Policy*, 43(5): 812-827.
- Dahlander, L., & Frederiksen, L. (2012). The core and cosmopolitans: "A relational view of innovation in user communities". *Organization Science*, 23(4), 988-1007.
- Daw, N. D., O'Doherty, J. P., Dayan, P., Seymour, B., & Dolan, R. J. (2006). Cortical substrates for exploratory decisions in humans. *Nature*, 441(7095), 876-879.
- Dellarocas, C. (2003). The digitization of word of mouth: Promise and challenges of online feedback mechanisms. *Management Science*, 49(10): 1407-1424.
- Doll, B. B., D. A. Simon, et al. (2012). "The ubiquity of model-based reinforcement learning." *Current Opinion in Neurobiology* 22(6): 1075-1081.
- Dosi, G. (1988). The nature of the innovative process. *Technical Change and Economic Theory*, 2, 590-607.
- Dosi, G., & Nelson, R. R. (1994). "An introduction to evolutionary theories in economics." *Journal of evolutionary economics*, 4(3), 153-172.
- Drucker, P. (2014). *Innovation and entrepreneurship*, Routledge, London
- Ebner, W., J. M. Leimeister, et al. (2009). "Community engineering for innovations: the ideas competition as a method to nurture a virtual community for innovations." *R & D Management* 39(4): 342-356.
- Erat, S., & Krishnan, V. (2012). "Managing delegated search over design spaces". *Management Science*, 58(3), 606-623.
- Ericsson, K. A. and H. A. Simon (1998). "How to study thinking in everyday life: Contrasting think-aloud protocols with descriptions and explanations of thinking." *Mind, Culture, and Activity* 5(3): 178-186.
- Estelles-Arolas, E. and F. Gonzalez-Ladron-de-Guevara (2012). "Towards an integrated crowdsourcing definition." *Journal of Information Science* 38(2): 189-200.
- Fang, C., Lee, J., & Schilling, M. A. (2010). Balancing exploration and exploitation through structural design: The isolation of subgroups and organizational learning. *Organization Science*, 21(3): 625-642.
- Felin, T. (2012). "Cosmologies of capability, markets and wisdom of crowds: Introduction and comparative agenda." *Managerial and Decision Economics* 33(5-6): 283-294.
- Felin, T., et al. (2012). "Microfoundations of routines and capabilities: Individuals, processes, and structure." *Journal of Management Studies* 49. (8): 1351-1374.
- Felin, T. and T. R. Zenger (2014). "Closed or open innovation? Problem solving and the governance choice." *Research Policy* 43(5): 914-925.

- Fioretti, G. (2013). "Agent-based simulation models in organization science." Organizational Research Methods 16(2): 227-242.
- Fleming, L. and O. Sorenson (2003). "Navigating the technology landscape of innovation." MIT Sloan Management Review Winter 2003.
- Fosfuri, A. and T. Rønde (2009). "Leveraging resistance to change and the skunk works model of innovation." Journal of Economic Behavior & Organization 72(1): 274-289.
- Franzoni, C. and H. Sauermann (2014). "Crowd science: The organization of scientific research in open collaborative projects." Research Policy 43(1): 1-20.
- Franke, N., Lettl, C., Roiser, S., & Tuertscher, P. (2014). "Does god play dice?" - Randomness vs. deterministic explanations of crowdsourcing success. Academy of Management Proceedings 2014(1): 15164. Academy of Management.
- Frey, K., C. Lüthje, et al. (2011). "Whom should firms attract to open innovation platforms? The role of knowledge diversity and motivation." Long Range Planning 44(5): 397-420.
- Füller, J., Hutter, K., & Faullant, R. (2011). "Why co-creation experience matters? Creative experience and its impact on the quantity and quality of creative contributions". R&D Management, 41(3), 259-273.
- Fullerton, R. L., & McAfee, R. P. (1999). "Auctioning entry into tournaments". Journal of Political Economy, 107(3), 573-605.
- Gavetti, G. and D. Levinthal (2000). "Looking forward and looking backward: Cognitive and experiential search." Administrative Science Quarterly 45(1): 113-137.
- Gavetti, G. (2005). "Cognition and Hierarchy: Rethinking the Microfoundations of Capabilities' Development." Organization Science 16(6): 599-617.
- Gläscher, J., Daw, N., Dayan, P., & O'Doherty, J. P. (2010). "States versus rewards: dissociable neural prediction error signals underlying model-based and model-free reinforcement learning." Neuron, 66(4), 585-595.
- Grant, R. M. (1996). "Toward a knowledge-based theory of the firm". Strategic management journal, 17(S2), 109-122.
- Harrison, J. R., L. I. N. Zhiang, et al. (2007). "Simulation modeling in organizational and management research." Academy of Management Review 32(4): 1229-1245.
- Hayes, J. R. and H. A. Simon (1975). "Psychological differences among problem isomorphs," Carnegie-Mellon University, Department of Psychology.
- Hebbron, T., S. Bullock, et al. (2008). NKalpha: Non-uniform epistatic interactions in an extended NK model. Artificial Life XI: Proceedings of the Eleventh International Conference on the Simulation and Synthesis of Living Systems S. Bullock, MIT Press, Cambridge, MA: 234-241.

- Hong, L. and S. Page (2004). "Groups of diverse problem solvers can outperform groups of high-ability problem solvers." *Proceedings of the National Academy of Sciences of the United States of America* 101(46): 16385-16389.
- Howe, J. (2008). *Crowdsourcing: How the power of the crowd is driving the future of business*, Random House.
- Huys, Q. J., Eshel, N., O'Nions, E., Sheridan, L., Dayan, P., & Roiser, J. P. (2012). Bonsai trees in your head: how the Pavlovian system sculpts goal-directed choices by pruning decision trees. *PLoS Computational Biology*, 8(3), e1002410.
- Hutter, K., J. Hautz, et al. (2011). "Communitition: The tension between competition and collaboration in community-based design contests." *Creativity and Innovation Management* 20(1): 3-21.
- Iclanzan, D. and D. Dumitrescu (2007). "Overcoming hierarchical difficulty by hill-climbing the building block structure." *Proceedings of the 9th annual conference on Genetic and evolutionary computation*, ACM.
- Jack, A. I. and A. Roepstorff (2002). "Introspection and cognitive brain mapping: from stimulus-response to script-report." *Trends in Cognitive Sciences* 6(8): 333-339.
- James, W. (1890). *The principles of psychology*. Vols. 1, 2, Holt, New York
- Jablokow, K. W., et al. (2015) "Cognitive Style and Concept Mapping Performance." *Journal of Engineering Education* 104(3): 303-325.
- Jeppesen, L. B. and L. Frederiksen (2006). "Why do users contribute to firm-hosted user communities? The case of computer-controlled music instruments." *Organization Science* 17(1): 45-63.
- Jeppesen, L. B., & Lakhani, K. R. (2010). "Marginality and problem-solving effectiveness in broadcast search". *Organization science*, 21(5), 1016-1033.
- Jouret, G. (2009). "Inside Cisco's search for the next big idea." *Harvard Business Review* 87(9): 43-+.
- Kauffman, S. A. (1993). *The origins of order: self-organization and selection in evolution*. New York, Oxford University Press.
- Katila, R., & Ahuja, G. (2002). Something old, something new: A longitudinal study of search behavior and new product introduction. *Academy of Management Journal*, 45(6), 1183-1194.
- Kirton, M. (1976). "Adaptors and innovators: A description and measure." *Journal of Applied Psychology* 61(5): 622.
- Knudsen, T., & Levinthal, D. A. (2007). "Two faces of search: Alternative generation and alternative evaluation". *Organization Science*, 18(1), 39-54.

- Kogut, B. and U. Zander (1992). "Knowledge of the firm, combinative capabilities, and the replication of technology." *Organization science* 3(3): 383-397.
- Kozlowski, S. W. J. and G. T. Chao (2012). "The dynamics of emergence: Cognition and cohesion in work teams." *Managerial and Decision Economics* 33(5-6): 335-354.
- von Krogh, G., Wallin, M., & Sieg, J. H. (2012). "A problem in becoming: How firms formulate sharable problems for innovation contests". In *Research Policy Special Issue Conference" Open Innovation: New Insights and Evidence"*, London, UK.
- Lakhani, K. R. (2008). InnoCentive. com (A). *Harvard Business School Case*, (608-170).
- Lakhani, K. R., K. J. Boudreau, et al. (2013). "Prize-based contests can provide solutions to computational biology problems." *Nature Biotechnology* 31(2): 108-111.
- Lampel, J., P. P. Jha, et al. (2012). "Test-driving the future: How design competitions are changing innovation." *Academy of Management Perspectives* 26(2): 71-85.
- Laursen, K. (2012). "Keep searching and you'll find: what do we know about variety creation through firms' search activities for innovation?" *Industrial and Corporate Change* 21(5): 1181-1220.
- Lavie, D., Stettner, U. and Tushman, M.L., 2010. Exploration and exploitation within and across organizations. *The Academy of Management Annals*, 4(1):109-155.
- Lazer, D. and A. Friedman (2007). "The network structure of exploration and exploitation." *Administrative Science Quarterly* (52): 667-694.
- Leimeister, J. M., Huber, M., Bretschneider, U., & Krcmar, H. (2009). "Leveraging crowdsourcing: activation-supporting components for IT-based ideas competition". *Journal of management information systems*, 26(1), 197-224.
- Levinthal, D., & March, J. G. (1981). A model of adaptive organizational search. *Journal of Economic Behavior & Organization*, 2(4), 307-333.
- Levinthal, D. A. and J. G. March (1993). "The Myopia of Learning." *Strategic Management Journal* 14: 95-112.
- Levinthal, D. A. (1997). "Adaptation on rugged landscapes". *Management science*, 43(7), 934-950.
- Lopez-Vega, H., F. Tell, et al. (2016). "Where and how to search? Search paths in open innovation." *Research Policy* 45(1): 125-136.
- Lorenz, J., H. Rauhut, et al. (2011). "How social influence can undermine the wisdom of crowd effect." *Proceedings of the National Academy of Science* 108(22): 9020-9025.
- March, J. G. (1991). "Exploration and exploitation in organizational learning." *Organization Science* 2(1): 71-87.
- Marengo, L. (2014). "Representation, search, and the evolution of routines in problem solving." *Industrial and Corporate Change: dtu023*.

- Marjanovic, S., C. Fry, et al. (2012). "Crowdsourcing based business models: In search of evidence for innovation 2.0." Science and Public Policy: scs009.
- Martin, R. C. (2003). Agile software development: principles, patterns, and practices. Prentice Hall PTR.
- Mason, W. and D. J. Watts (2012). "Collaborative learning in networks." Proceedings of the National Academy of Sciences of the United States of America 109(3): 764-769.
- Mason, W., & Suri, S. (2012). Conducting behavioral research on Amazon's Mechanical Turk. Behavior Research Methods, 44(1), 1-23.
- Mavrodiev, P., C. J. Tessone, et al. (2013). "Quantifying the effects of social influence." Scientific Report. 1360(3).
- Mihm, J., Loch, C. H., Wilkinson, D., & Huberman, B. A. (2010). "Hierarchical structure and search in complex organizations". Management science, 56(5), 831-848.
- Miron, E., M. Erez, et al. (2004). "Do personal characteristics and cultural values that promote innovation, quality, and efficiency compete or complement each other?" Journal of Organizational Behavior 25(2): 175-199.
- Mitchell, M. (2009). Complexity: A guided tour. Oxford University Press
- Murray, F., S. Stern, et al. (2012). "Grand innovation prizes: A theoretical, normative, and empirical evaluation." Research Policy 41(10): 1779-1792.
- Nickerson, J. A. and T. R. Zenger (2004). "A knowledge-based theory of the firm - The problem-solving perspective." Organization Science 15(6): 617-632.
- Nonaka, I. (2007). "The knowledge-creating company." Harvard Business Review 85(7-8): 162-171.
- Pavitt, K. (1984). "Sectoral patterns of technical change: towards a taxonomy and a theory." Research Policy 13(6): 343-373.
- Pfleiderer, P (2014). "Chameleons: The Misuse of Theoretical Models in Finance and Economics", Stanford University Working paper
- Piezunka, H. and L. Dahlander (2013). A study of organizations' attention to suggestions by externals over time, Stanford University.
- Poetz, M. K. and M. Schreier (2012). "The value of crowdsourcing: can users really compete with professionals in generating new product ideas?" Journal of Product Innovation Management 29(2): 245-256.
- Raisch, S., & Birkinshaw, J. (2008). Organizational ambidexterity: Antecedents, outcomes, and moderators. Journal of Management.
- Rivkin, J. W. (2000). "Imitation of complex strategies". Management science, 46(6), 824-844.
- Rivkin, J. W., & Siggelkow, N. (2003). "Balancing search and stability: Interdependencies among elements of organizational design." Management Science, 49(3), 290-311.

- Schooler, L. J., C. Burgess, et al. (2012). Search environments, representation, and encoding. *The Evolution of Cognitive Search*. T. T. Hills and R. Dukas.
- Schöttner, A. (2008). "Fixed-prize tournaments versus first-price auctions in innovation contests." *Economic Theory* 35(1): 57-71.
- Schumpeter, J. A. (1934). *The theory of economic development*. Cambridge, Harvard University Press.
- Schultz, W., P. Dayan, and P. R. Montague (1997). "A neural substrate of prediction and reward." *Science* 275(5306): 1593-1599.
- Shah, S. K. (2006). "Motivation, governance, and the viability of hybrid forms in open source software development." *Management Science* 52(7): 1000-1014.
- Simon, H. A. and P. A. Simon (1962). "Trial and error search in solving difficult problems: Evidence from the game of chess." *Behavioral Science* 7(4): 425-429.
- Simon, H. A. (1983). Search and reasoning in problem solving. *Artificial Intelligence*, 21(1), 7-29.
- Simon, H. A. (1991). "Bounded rationality and organizational learning." *Organization Science* 2(1): 125-134.
- Simon, H. A. (1996). The sciences of the artificial. *MIT Press*
- Smaldino, P. E. and P. J. Richerson (2012). "The origins of options." *Frontiers in neuroscience* 6(50):1-11.
- Solée, R. V., S. Valverde, et al. (2013). "The evolutionary ecology of technological innovations." *Complexity* 18(4): 15-27.
- Solway, A. and M. M. Botvinick (2012). "Goal-directed decision making as probabilistic inference: a computational framework and potential neural correlates." *Psychological Review* 119(1): 120-154.
- Solway, A. and M. M. Botvinick (2015). "Evidence integration in model-based tree search." *PNAS* 112(37): 11708-11713
- Sternberg, R. J., & Grigorenko, E. L. (1997). Are cognitive styles still in style?. *American psychologist*, 52(7), 700.
- Terwiesch, C. and Y. Xu (2008). "Innovation contests, open innovation, and multiagent problem solving." *Management Science* 54(9): 1529-1543.
- Tidd, J., J. Bessant, et al. (1997). *Managing innovation*. Chichester, John Wiley.
- Tripsas, M. and G. Gavetti (2000). "Capabilities, cognition, and inertia: Evidence from digital imaging." *Strategic Management Journal* 21(10-11): 1147-1161.
- Villarroel, A. J., J. E. Taylor, et al. (2013). "Innovation and learning performance implications of free revealing and knowledge brokering in competing communities: insights from the

- Netflix Prize challenge." Computational and Mathematical Organization Theory 19: 42-77.
- Vuculescu, O. and C. Bergenholz (2014). "How to solve problems with crowds: A computer-based simulation model." Creativity and Innovation Management 23(2): 121-136.
- Watts, D. J. (2007). A twenty-first century science. Nature, 445(7127), 489-489.
- Watts, D. J. (2011). Everything is obvious: Once you know the answer. Crown Business.
- Winter, S. G., G. Cattani, et al. (2007). "The value of moderate obsession: Insights from a new model of organizational search." Organization Science 18(3): 403-419.
- Wooten, J. O. and K. T. Ulrich (2015). "Idea generation and the role of feedback: Evidence from field experiments with innovation tournaments." Available at SSRN 1838733.
- Wooten, J. O. and K. T. Ulrich (2015). "The impact of visibility in innovation tournaments: Evidence from field experiments." Available at SSRN 2214952.
- Wright, S. (1932). The roles of mutation, inbreeding, crossbreeding, and selection in evolution, 1:356-366
- Zheng, H., D. Li, et al. (2011). "Task design, motivation, and participation in crowdsourcing contests." International Journal of Electronic Commerce 15(4): 57-88.

Fitness landscapes in organizational theory: Research challenges and future directions

Oana Vuculescu*, Carsten Bergenholtz*, Mads Kock, Jacob Sherson****

***Department of Management, Aarhus University, Denmark**

****Department of Physics and Astronomy, Aarhus University, Denmark**

Keywords: search behaviours, fitness landscape, hierarchical modularity, NK model

Abstract

Using insights from computer science and biology, we argue that current computational modelling research needs to address several fundamental issues in order to generate more meaningful and falsifiable contributions. Based on comparative simulations and a new type of visualization, we address two key elements that the traditional NK framework has relied on: a) how the NK captures the complexity of organizational problems and b) search behaviours where, despite evidence, local search is often used as the dominant problem solving strategy. We show that these two components are fundamentally intertwined and outline implications for how to simulate organizational problems.

Introduction

In the current literature on search two central issues are still left unresolved: 1) in the continuum from mindless particles to perfect rationality, how to model bounded rational search behaviour and 2) how to conceptualize the space of solutions (i.e. the interdependence structure) in which search takes place (Sorenson 2002, Chang and Harrington 2006, Todd, Hills et al. 2012, Baumann 2015). In line with work from computer science (Jones 1995, Pitzer and Affenzeller 2012), as well as the original fitness model formulation (Wright 1932) we show how the two are fundamentally intertwined and in order to be able to create better models of problem solving, they must be addressed together.

Organizational theory has a long tradition of studying organizations' search for solutions to 'hard' problems, i.e. problems where it is computationally impossible or merely too expensive to list and test all possible solutions (Simon 1956; Cohen et al. 1972). The prevalent way of addressing individual or organizational search behaviour and how to conceptualize the space of solutions stems from early work on population genetics, namely the fitness landscape model (Wright 1932). Within biology, by focusing on fitness interactions between genes, Wright's framework allows for a link between low-level properties of genes and the high-level patterns of the dynamics of evolution (Altenberg 1997). The model's most famous extension, the NK model

(Kauffman 1993), explicitly models adaptive evolution as a “search in protein space” (Kauffman 1993: p. 37) which tries to find a maximum point for a chosen fitness function. This approach has grown outside the boundaries of population genetics literature and inspired a series of scholars from computer science (e.g. Pitzer and Affenzeller 2012), organizational theory (e.g. Baumann and Siggelkow 2013) and physics (e.g. Sørensen et al. 2015).

How can problem-solving be addressed in this framework? Imagine trying to solve an innovation problem, for instance designing a new educational app. As any software developer would tell you, there is no need to start from zero: nowadays there are a number of pre-defined libraries you can use, which you can think of as interconnected modules. But how are these modules interconnected and how will this affect your chances of finding a good design? (i.e. what does the task environment look like). Once you have a working prototype, should you then just go through each module, one at a time, and try to make minor improvements? (i.e. what organizational theory calls ‘local search’). What if you get stuck - should you discard everything and start again from scratch? (i.e. ‘random long-jumps’).

Levinthal (1997) introduces the NK model in order to facilitate the formal modelling and simulation of how the level of interdependence in an organization’s routines affects its long-term chances of finding the optimal configuration of such routines and thus survive in a competitive environment. Typically, the organizational literature is primarily focused on how organizations search over the space of routines for combinations leading to increased performance (e.g. Levinthal 1997), but the underlying intuition is the same. By making explicit assumptions about individual or organizational behaviour and the environment in which the agent evolves, researchers could now simulate how such agents adapt over time. Based on these assumptions one can then map the complex dynamics of organizations being embedded *in* and adapting *to* the competitive environment (Levinthal 1997). This has enabled the field to go beyond static explanations and model possible future trajectories of the current competitive situation: e.g. organizations that are highly coupled (have a high interdependence between routines), Levinthal (1997) argues, have a higher likelihood of failure in the face of changing environments. Later papers have developed this approach, addressing either different interdependence structures (e.g. Rivkin and Siggelkow 2003) or how agents (be it individuals or organizations) search or adapt (Gavetti and Levinthal 2000; Gavetti and Levinthal 2001; Winter, Cattani et al. 2007; Baumann and Siggelkow 2013, Martignoni et al. 2015). Almost all of this work, however, looks at either interdependence structures (e.g. Ethiraj and Levinthal 2004) or search behaviours in isolation (e.g. Winter et al. 2007).

This paper attempts to address these issues to shed light on the challenges of a more systematic approach to modelling organizational and individual problem solving and offers a technical and theoretical comparative analysis of a number of assumptions about search strategies and fitness

landscapes. This analysis is substantiated by a novel (to the organizational literature) type of visualization that maps how different search strategies actually 'generate' different landscapes, rather than just searching in an a priori given space. We identify and explore two main limitations with modelling problem solving via the NK framework. First, recent research (He et al. 2007) has shown that unless certain complexity-theoretical assumptions are wrong¹, for hard problems, a predictive measure of problem hardness cannot exist. This result mirrors a number of red flags already raised in the organizational literature (e.g. Frenken et al. 1999) regarding whether the NK is indeed a 'tuneable complexity landscape' and calls for a more elaborated and systematic discussion regarding the kind of problems that can be modelled via the NK framework. Second, we argue that the NK framework has several limitations with respect to modelling more plausible search behaviours, in particular due to particular assumptions regarding the fitness function.

In the following we begin, for historical reasons, to review the use of fitness landscapes in section 2 and search behaviours in section 3. Based on these insights from organizational theory and biology we discuss the implications for organizational theory of the presented visualizations and analysis of simulations in section 4.

2. Fitness landscapes

There are two main elements in the fitness landscape model that need to be specified for problem solving processes to be captured: the task structure (i.e. the problem that is to be solved) and the search behaviour (i.e. how problem solving unfolds). In order to discuss the difference between the "objectively defined task" (Simon and Newell 1971: p. 148) and the fitness landscape, which is the backbone of the NK simulations, we start by providing a short formal definition of fitness landscapes. The fitness landscape is in effect what the solver subjectively perceives or 'the problem space' (Simon and Newell 1971). We start by providing an overview of the main elements that describe an NK landscape. We caution that although we attempt to discuss the two issues in sequence, a certain amount of overlap is inevitable.

2.1. NK landscapes

In an optimization problem, a solution to a given problem is represented as a size N vector of (traditionally) binary variables². The quality of a solution, in keeping with the biological inspiration of the model, is given by a 'fitness function' (i.e. an objective function).

Let X be the space of all possible solutions to a problem. In a maximizing optimization problem¹ (i.e. a problem where the goal is to find the solution for which the function f has the highest value), for a function f , a solution $x^* \in X$ is a global optimum if

¹ P=NP respectively BPP = NP. See He et al. (2007) for an elaboration.

² The model can be easily and without loss of generality extended to larger alphabets.

$$f(x^*) \geq f(x), \text{ for any } x \in X. \quad (1)$$

where all binary variables are in their optimal position.

The landscape is a mapping between solutions and fitness values that takes into account the connectivity between solutions. In order to define connectedness (or pre-defined similarity) between solutions we need to specify a distance metric. Without a definition of a metric one cannot define a fitness landscape as Wright (1932) conceptualized it. As such, for any type of problem, a fitness landscape $L(X, f, d)$ consists of all $x \in X$, an objective function f that measures the quality of each solution and a distance measure d . Following Pitzer and Affenzeller (2012) we can define *a distance metric* as:

$$d: X \times X \rightarrow \mathbb{R}, \text{ such that } d(s, t) \geq 0,$$

$$d(s, t) = 0 \Leftrightarrow s = t, \quad d(s, t) \leq d(s, u) + d(u, t), \quad \forall s, t, u \in X. \quad (2)$$

Subsequently, the structure of the search space (Rothlauf 2011; Pitzer and Affenzeller 2012) is defined via the ε -neighbourhood of x :

$$M(x) = \{n / n \in X; n \neq x; d(x; n) \leq \varepsilon\}. \quad (3)$$

The neighbourhood function thus defines the set of all solutions that are different from the focal solution and within an ε radius of it, *as measured by the chosen distance metric* (as given by Equation 2). Thus, the fitness landscape's shape (or topology) is given only in conjunction with *a given neighbourhood function* (Jones 1995): the fitness landscape (L) is not the same as the fitness function (f) and it is likely that for any function f there can be a number of landscapes L with vastly different properties (Maier et al. 2014).

To illustrate this point, Figure 1 shows an example of the same function mapped onto three different landscapes using three different expressions for the neighbourhood function.

¹ Conversely, in a minimizing optimization problem, $x^* \in X$ is a global optimum, if $f(x^*) \leq f(x)$ for any $x \in X$.

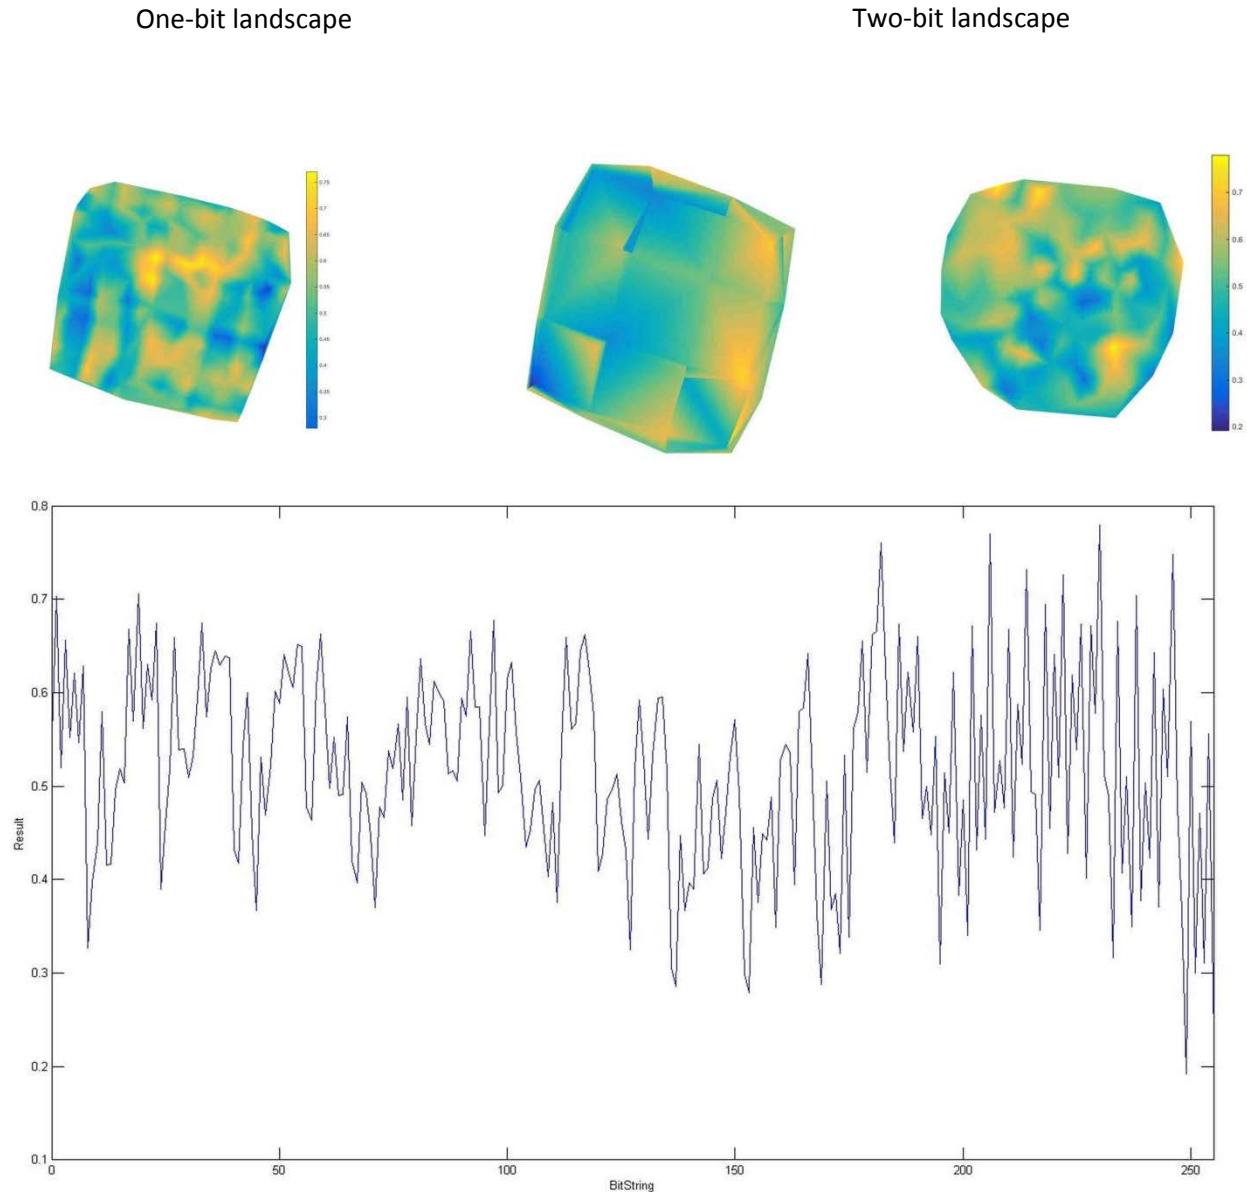


Figure 1. The same function ($N=8$, $K=3$) mapped with three different definitions of neighbourhood.

We used a dimensionality reduction method that transforms high dimensional data to low-dimensional representations while preserving pair-wise similarities (Van der Maaten and Hinton 2008) to create 3D visualizations of the multi-dimensional landscape¹. The illustration on the left

¹ Given the fact that the t-SNE (t-distributed stochastic neighbour embedding) algorithm is stochastic, as is the NK fitness function, it should be noted that this is one possible illustration of one possible NK landscape with $N=8$, $K=3$. The illustration is not a general result for all NK landscapes of $N=8$ and $K=3$.

depicts one NK landscape generated by relying on a one-bit-flip (i.e. any two solutions are considered similar if the Hamming distance¹ between them is exactly one), while the ones on the right depicts an NK landscape generated by relying on a two-bit flip neighbourhood structure (i.e. any two solutions are considered similar if the Hamming distance between them is exactly two). Both landscapes have the same underlying fitness function, generated for an NK with N=8 and K=3. Finally, the lower part of the graphic shows the same NK function mapped by converting the bit of strings to the decimal system (cf. Østman and Adami 2014) and solutions are considered to be similar if their decimal representation transformations are consecutive (i.e. 10011001 is transformed into 153 and its natural neighbours are 151 (10011000) and 154 (10011010)). The decimal representation allows for a up-front intuition regarding the distribution of high fitness peaks in the in the solution space (Østman and Adami 2014).

This visualization is telling in two ways. First, note that the two-bit-flip generates two different landscapes. Depending on the starting point, a subset of solutions is not connected in the graph. Similarly if one attempts to traverse a sequence of consecutive numbers with increments of 2, one generates two distinct and unconnected subsets: odd and even numbers. Thus, the definition of the neighbourhood function can effectively reduce (relative to the entire search space) the size of the landscape. Second, the heat-map preserves information regarding the distribution of fitness scores. One can assess (qualitatively in this case) the relative ease or difficulty of navigating towards the lighter coloured areas, in a traditional NK fashion. The three neighbourhood representations yield three different landscape topologies, i.e. smoother gradients such as the left-hand side of the two-bit flip mean that it would be easy for an agent to find the global optimum, while ‘patchier’ surfaces translate into a lower likelihood of success, such as the one-bit and the right-hand side landscape generated by the two-bit flip. Equivalently, in the 2D decimal representation, one can assess the difficulty of finding the global optimum (the highest fitness value), by looking at the shape of the generated curve. In the 2D case, since the decimal representation is arbitrary, the ‘decimal’ landscape is very “rugged”, thus a solver will likely get stuck in a suboptimal solution.

K/N ratios or epistatic interactions

While in the previous section we have shown how landscapes are defined and how the topology of the landscape changes as a result of the definition of the neighbourhood function, we now turn to how scholars have attempted to describe or more formally capture the ruggedness of a landscape, within the NK framework.

¹ The number of variables which have different values.

In the NK model, a solution to a given problem is represented as a size N vector of binary variables and the fitness function f is the average across all contributions in the genome. The details of how these contributions are computed are explained below.

$$f(x) = \frac{1}{N} \sum_{i=1}^N x_i \quad (4)$$

The parameter K gives the number of interactions between the N alleles (i.e. variables). For K=0, each contribution can take only two possible values. The single allele contributions are independent and identically distributed random variables. Each allele has one state that is preferable to the other, independently of the values of the other variables. The global optimum is the state where all alleles are in the individual optimal position. The global optimum can be reached from any initial configuration.

For K>0 the contribution of each allele depends on the position of a number of other K alleles. The choice of which alleles are interdependent is at the latitude of the modeller, but a common assumption is that neighbouring solutions influence each other. For example, for N=4, K=1, a solver cannot determine the optimal position for the first site in the solution, but has to compute all four possible combinations for the first and second site: {00, 01, 11, 10} and then choose the maximizing sequence. For K=N-1 the entire sequence appears in the argument of each single gene contribution and each step replaces the fitness with a different random number. The interdependence between alleles (or solution components) is known as *epistasis*, a term borrowed from biology where it denotes the fact that the expression of a gene is altered by the presence of another.

Early organizational studies relying on the NK model follow on the path proposed by Kauffman (1993) and study how the attributes of the search space influence the propensity of finding the optimal solution by an one-bit-flip hill climber (e.g. Levinthal 1997; Ethiraj and Levinthal 2004). Indeed, part of the NK model's popularity in organizational literature is due to the fact that it allows the investigation of different problem difficulties (Afuah and Tucci 2012), via the K/N ratio, or the level of epistatic interactions (Weise et al. 2009). Epistasis is equivalent to the non-linearity of a problem or how well a problem can be decomposed into sub-problems (Rothlauf 2011; Pitzer and Affenzeller 2012). In other words, epistasis gives a measure of signal to noise in what concerns the evaluations of the fitness function. Consider the N=4, K=1 example used before. A solver with $\{x=(0,0,0,1), f(x)=0.56\}$, might move to $\{x^*=(0,0,0,0), f(x^*)=0.58\}$, even though $\{y=(0,0,1,1), f(y)=0.72\}$. The fact that the optimal setting for the fourth allele is {1} is obscured by the epistatic interaction with its neighbour on the third position. In the light switch example, it is naturally easier to detect the optimal configurations for the light switches provided

they control distinct parts of the lighting in the room. That is, only for $K=0$, the contribution of each allele is independent and a solver can easily detect the optimal configurations and epistasis is 0.

2.2. Landscape features

Despite the fact that the notion of epistatic interactions, as outlined above, is extremely important within organizational theory, its use in quantifying the hardness of a problem has often been criticized (Mason 1995; Naudts and Verschoren 1999) in particular due to the difficulty of identifying measures of epistasis that have enough predictive power (Pitzer and Affenzeller 2012). There are several known limitations to using epistasis measures as proxies for problem complexity. First, epistatic interactions can be both positive and negative. Whether an interaction effect between two alleles is positive or negative has a significant impact on the difficulty of a problem, but epistasis measures (e.g. epistasis variance or correlation) cannot capture this distinction (Naudts and Kallel 2000). Second, empirical evidence suggests that epistatic interactions can occur at several levels (i.e. there are hierarchical interdependence structures) and this has consequences for the long-term dynamics of the system (Szendro et al. 2013). Thus, what is important is not acknowledging that a problem *has* epistatic interactions, but rather identifying the nature of those interactions.

Paralleling this trend, recent years have seen considerable development when it comes to the study of fitness landscapes (Pitzer and Affenzeller 2012; Malan and Engelbrecht 2013; McClymont 2013; Østman and Adami 2014) with a focus shift from characterizing *problem hardness* (via ruggedness measures) to characterizing *fitness landscapes* in order to determine the appropriate algorithm (McClymont 2013). The shift is due to the fact that fitness landscape analysis allows for “a deeper understanding of a whole problem class” (Pitzer and Affenzeller 2012: p. 3) rather than a specific problem instance. Current research thus aims at identifying relevant features that can describe a fitness landscape and that have known properties with respect to problem solving difficulty (Malan and Engelbrecht 2014). Malan and Engelbrecht (2014) identify three such features as potentially predictive of performance: ruggedness, neutrality and deceptiveness (Figure 2).

In keeping with the organizational theory approach, where the one-bit hill-climbing algorithm is the dominant search behaviour, in the following we describe how these features can affect the likelihood of finding the optimal solution for a classic one-bit hill-climbing algorithm, but they are not limited to this search heuristic.

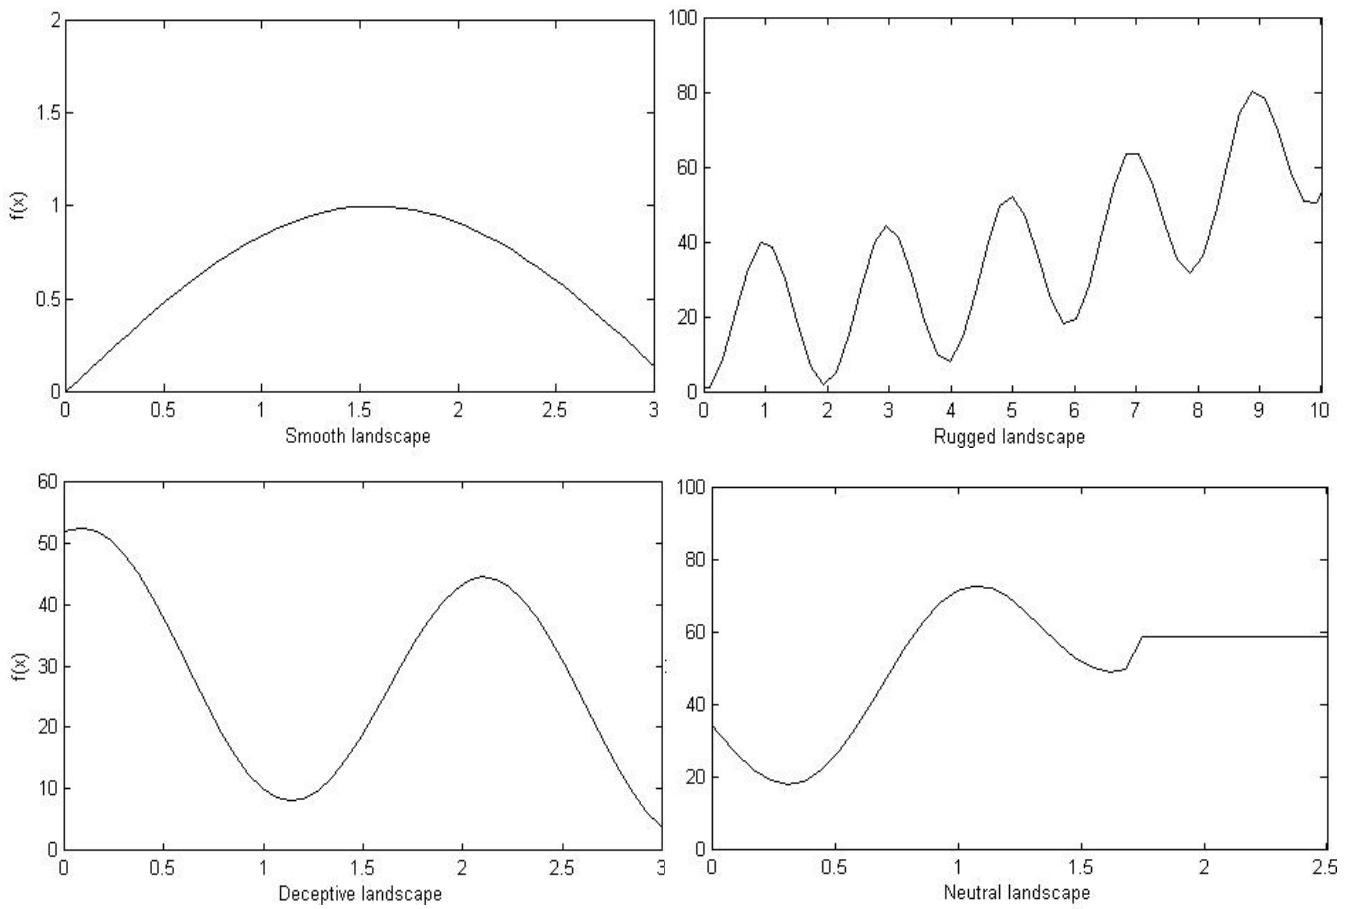


Figure 2. Landscape features adapted from Malan and Engelbrecht 2013.

2.2.1. Landscape ruggedness: modality and locality measures

In the previous section we have addressed how prior research has used K or K/N ratios as measures of landscape ruggedness or problem complexity. Indeed, Kauffman (1993) has shown that given the one-bit flip assumption, the ruggedness of an NK landscape is captured by the K parameter. However, although it is clear that highly epistatic landscapes are hard to search, it is not clear how much epistasis “is needed to make a problem difficult” (Jones 1995: p. 134). Thus, in the following, we present a number of alternative measures to capture landscape ruggedness.

In computer science, a frequently used measure of landscape ruggedness is the number of local maxima, or the modality of a landscape. The *modality* of a given landscape is often computed relative to the size of the fitness landscape: the higher the density of such deceiving optima, the more complex the problem, i.e. the higher the likelihood that a solver will be stuck and unable to find the optimal solution. Note that the definition of a distance metric (and implicitly the neighbourhood function) affects the number of local optima, since, by definition, for a problem

(X, f) and a *neighbourhood function* M , a solution x^* is called locally optimal *with respect to M*, if

$$f(x) \leq f(x^*) \text{ for all } x \in M(x). \quad (5)$$

The *locality* of a landscape is given by how closely together (with respect to the distance d) solutions with similar fitness values are located (Rothlauf 2011). In general, the lower the distance, the higher the locality and the easier it is to find a global optimum, since better solutions are located closer together (Pitzer and Affenzeller 2012).

Another measure of locality was proposed by Jones and Forrest in 1995 where they proposed a *fitness distance correlation coefficient*.

$$\rho_{FDC} = \frac{C_{fd}}{\sigma(f)\sigma(d_{opt})} \quad (6)$$

Where

$$C_{fd} = \frac{1}{m} \sum_{i=1}^m (f_i - \bar{f})(d_{i,opt} - \bar{d}_{opt}) \quad (7)$$

with $\sigma(f)$ and $\sigma(d_{opt})$ as the standard deviations for the fitness values, respectively the distances to the optimal solution, f is the mean value for the fitness function, d_{opt} is the mean value for the distance to the optimal solution, f_i the fitness value for solution i and finally $d_{i,opt}$ is the distance of solution i , to the optimal solution x^* .

The fitness-distance correlation coefficient, allows Jones and Forrest (1995) to distinguish between three classes of landscapes:

- a. Straightforward, for $\rho_{FDC} \leq -0.15$. This is the ideal case where the closer a solver gets to the global optimum, the higher the fitness and are roughly correspondent to “smooth” landscapes. NK problems where $K \leq 3$, fall in this category.
- b. Difficult $-0.15 < \rho_{FDC} < 0.15$. There is limited correlation between the fitness difference and the distance to the optimal solution. This makes such optimization problems very hard to solve and renders the search heuristics to random search. According to Jones and Forrest (1995) as K increases over 3, NK landscapes quickly become uncorrelated and ρ_{FDC} approaches 0. These are “rugged” landscapes, with limited or uncorrelated ruggedness.
- c. Misleading $\rho_{FDC} \geq 0.15$. There is an inverse correlation between the fitness difference and the distance to the optimal solution. Thus, the solver is “drawn” away from the global optimum.

According to Malan and Engelbrecht's (2014) classifications, these would be "deceptive landscapes".

2.2.2. Deceptiveness

Recent advances in biology point to the existence of higher-order epistatic interactions which generate multidimensional landscapes (Segre et al. 2005; Kondrashov et al. 2015). These interactions seem to be organized hierarchically in functional modules that interact with each other (Segre et al. 2005, Jaimovich et al. 2010). This type of interaction structure is reminiscent of the hierarchical structure which has been argued to be an essential feature of organizational problems, at least when it comes to innovation problems (Pelikan et al. 2000; Gavetti 2005). In this context, hierarchy, is seen as the composition of systems out of subsystems with each subsystem in turn having its own hierarchy (Yu et al. 2009), until a certain level of fine grained modularity is achieved. This is a qualitatively different kind of 'problem complexity' (as compared to landscape 'ruggedness') and the one most likely to be encountered in real-life design problems (Pelikan et al. 2000; Martin 2001; Yu et al. 2009). Note however that hierarchical decomposition and hierarchical interdependence are different from the one-level interdependence, which is captured by NK-like landscapes - see also Marengo et al. (2000) for a more detailed account. The latter assumes that the task of solving a problem can be reduced to several low order modules that have intertwined contributions to the overall fitness. In hierarchical problems the interdependence (or interactions) *between* levels is also present and this obstructs single-level decomposition (Pelikan 2005). This description is also in line with Simon's description of complexity (Simon 1962; 1996).

Such problems are likely to generate *deceptive* landscapes, according to Malan and Engelbrecht's (2014) classification, since they generate so-called hierarchical traps (Watson and Pollack 1999; Martin 2001; de Jong et al. 2005). The interactions between building blocks make hierarchical problems deceptive (i.e. misleading according to Jones and Forrest 1995) in Hamming space (at lower hierarchical levels), but fully non-deceptive at higher hierarchical levels (Iclanzan and Dumitrescu 2007) – i.e. at higher hierarchical levels (better problem representations), solvers are able to attain better solutions by making incremental changes. In biological terms: the lowest hierarchical level describes "how a mutation in a given gene affects the phenotypic consequence of another mutation and the highest level describes how altered functionality of a given module of genes affects the phenotypic consequence of altered functionality of another module." (Segre et al. 2005: p. 81)

One example of such function is illustrated in Figure 3.

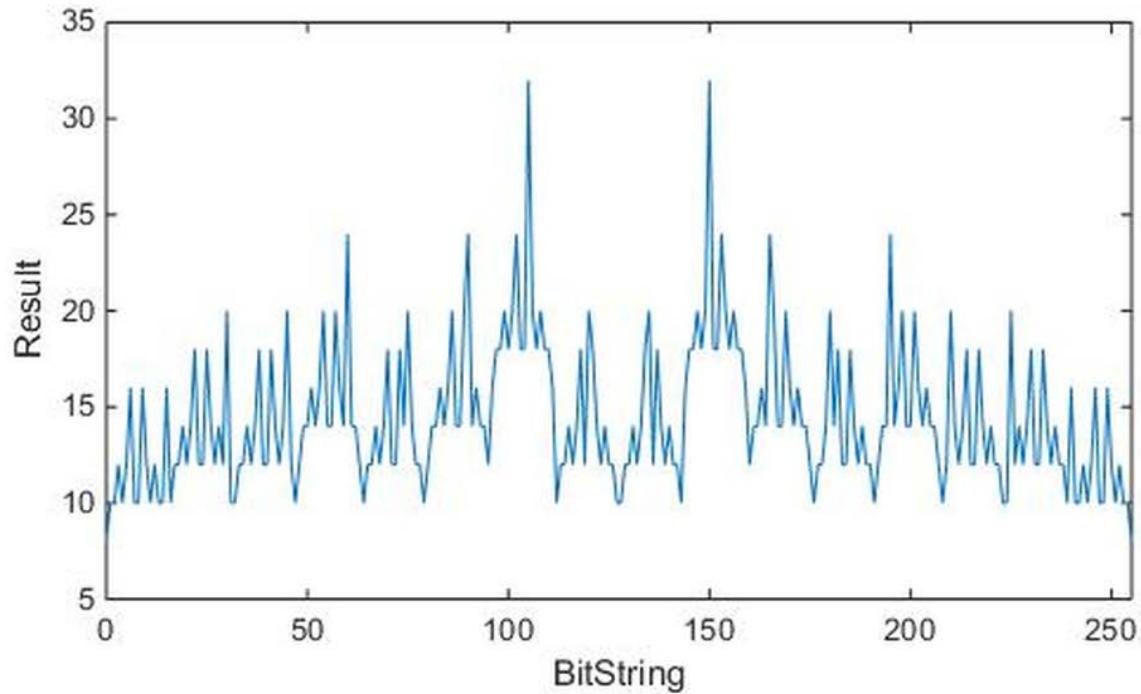
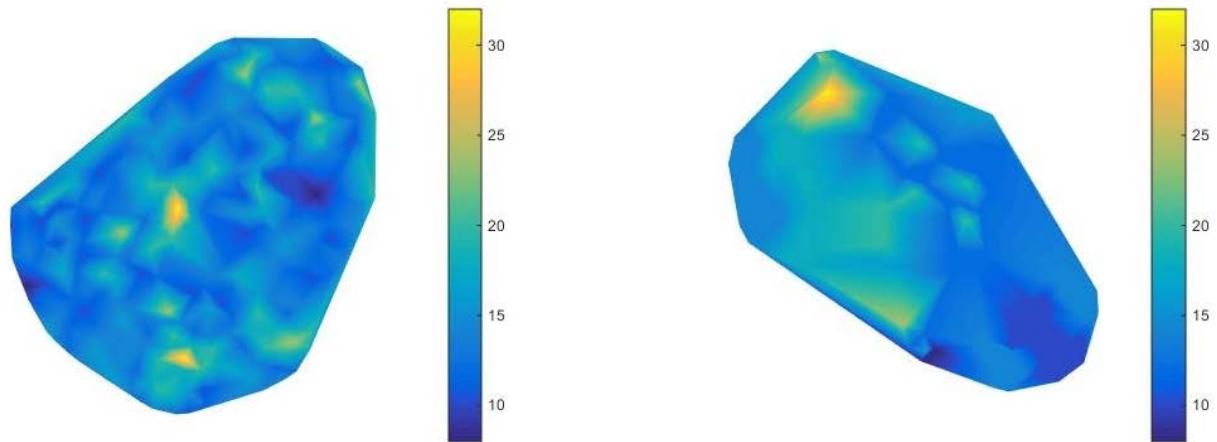


Figure 3. Same function (HXOR N=8) mapped with three different definitions of neighbourhood.

Figure 3 shows the visualization of a hierarchical problem using a one-bit flip hill-climber (left) and a “chunking” algorithm that was tailored specifically for this problem (see Appendix 1 for a description). Notice that the chunking landscape is ‘smoother’ than the one-bit-flip landscape.

Since the H-XOR¹ function has $2^{N/2}$ local optima for the one-bit flip hill-climber, the probability that a given point in the one-bit landscape is connected with a path to the global optimum is significantly lower (Figure 4) as compared to the chunking algorithm showing that this problem is ‘deceptive’ for a one-bit-hill-climber but not for an algorithm that can exploit the problem structure.

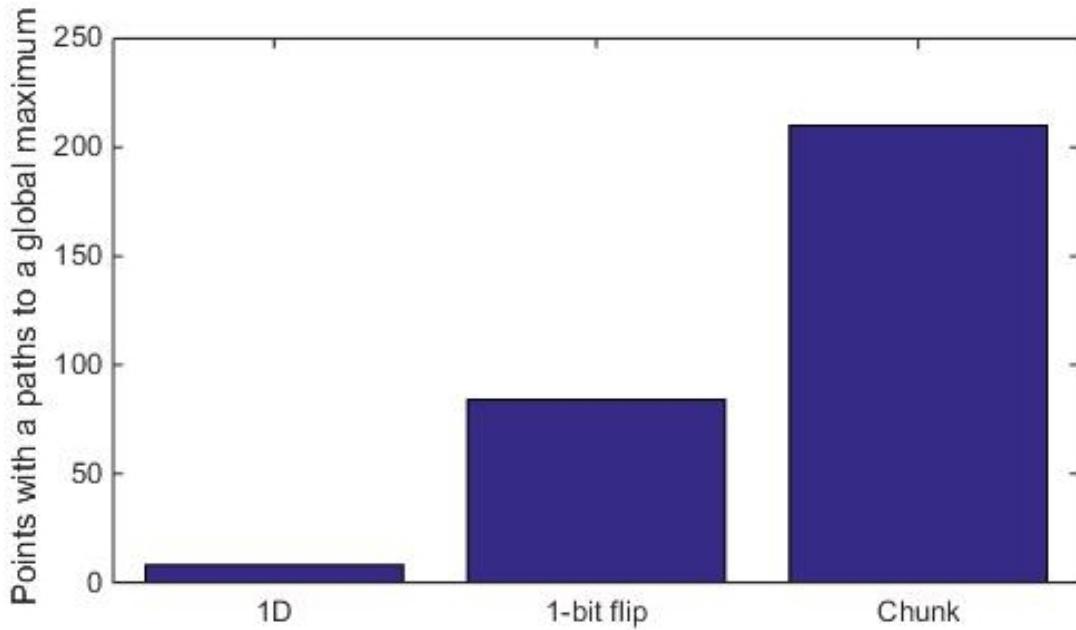


Figure 4. Comparative performance of 1D, 1-bit flip and chunking.

2.2.3. Neutrality

So far we have only looked at the “smooth vs rugged” distinction and different means of capturing *ruggedness*. A different intuition about how evolutionary dynamics might be influenced by the underlying fitness function comes from models that consider the possibility that some solutions have equal fitness. This was fuelled by developments in molecular biology which have questioned the “rugged landscape” metaphor, in particular its explanation of speciation (Barnett 1997; Gavrilets 1999). This work was largely driven by the neutral theory of molecular evolution and in particular the observation that the majority of mutations at a molecular level do not affect the phenotype (Galván-López et al. 2011). The previous framework assumed that once a population became stuck in a suboptimal peak it could only escape it if the fitness function was changed (e.g. shifting balance theory) or via a long jump. The neutral theory of molecular evolution relies on the conjunction that there must be a series of fitness neutral

¹ See Appendix for a detailed explanation of HXOR,

mutations that would allow even organisms that were currently located in a suboptimal peak to “escape” and undergo further evolution.

In an NKq (Newman and Engelhardt 1998) landscape, for a landscape L we define the *neutral neighbours of x* :

$$M_n(x) = \{x^* \in N(x) / f(x) = f(x^*)\}. \quad (8)$$

where fitness contributions are integers drawn from $[0, q]$. The total fitness in this case is given by:

$$f(x) = \frac{1}{N(q-1)} \sum_{i=1}^N x_i \quad (9)$$

Figure 5 shows a 3D reduction for a 1 bit-flip hill-climber as well as the decimal representation used earlier to depict two neutral functions with $N=8$, $K=1$, $q=2$ and $N=8$, $K=1$, $q=3$, respectively.

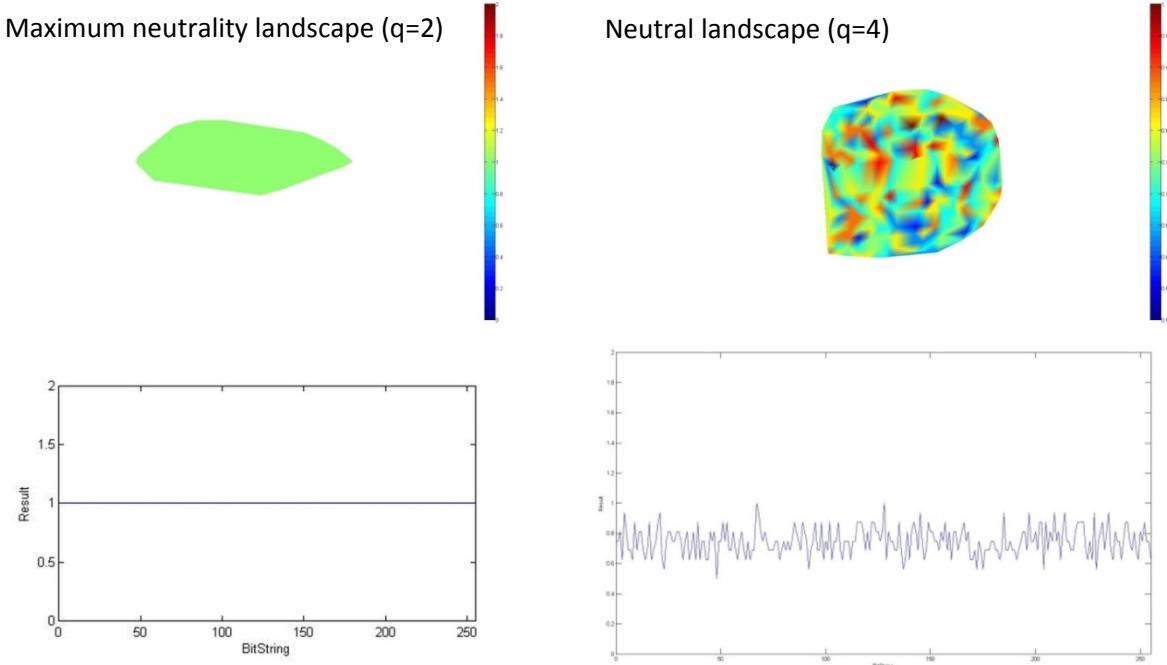


Figure 5. The same function ($N=8$, $K=3$, $q=4$) mapped with two different definitions of neighbourhood.

The left-hand side of the picture corresponds to the maximally neutral landscape and depicts a very simple, flat fitness landscape, without any local maxima. The right-hand side picture, corresponding to the $q=3$ neutral landscape, has a number of ridges of high fitness states, as well

as valleys of low fitness states. Thus, neutral landscapes are not necessarily beneficial for adaptation, since, even for low values of K, a solver is likely to find himself trapped in these “stretches of lethal states” (Franke et al. 2011: p.4).

A number of authors have introduced neutral extensions of the NK landscape and investigate how the new topology might influence the evolutionary processes (e.g. Barnett 1997; Newman and Engelhardt 1998; Lobo et al. 2004). The implementations vary in both details and conclusions regarding the influence of neutrality on the features of the landscape (Geard et al. 2002), but they do conclusively show that neutrality is an important feature that influences search performance and is not captured by traditional measures of ruggedness (Pitzer and Affenzeller 2012), commonly used in NK studies.

Figure 6 provides an illustration of how introducing neutrality can change the dynamics of adaptation. We compare the performance of a one-bit hill climber on an NK landscape ($N=8$, $K=3$) and a NKq landscape ($N=8$, $K=3$, $q=4^1$). Simulations show a higher success ratio on the NKq landscape. The success ratio is defined by the ratio between the number of paths and the number of successful paths, and is thus a measure of the likelihood of finding the optimal solution. Simulations were conducted on 1000 different NK (and correspondingly 1000 different NKq landscapes) and the difference was found to be significant ($p=0.04$) with an effect size $r=0.1$.

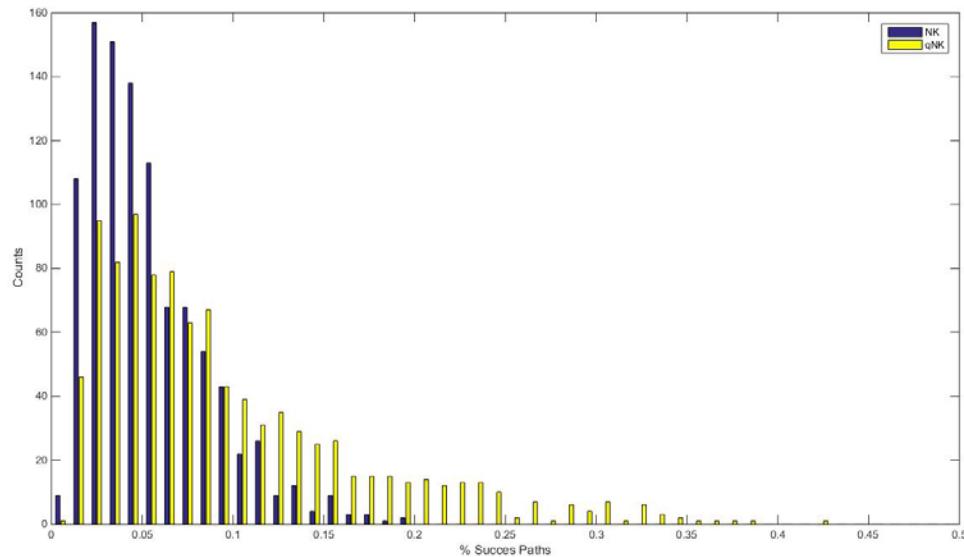


Figure 6. Comparative performance of a 1-bit hill climber on two landscapes: $N=8$, $K=3$ and its neutral counterpart $N=8$, $K=3$, $q=4$.

¹ We purposely chose a value for q higher than the lowest possible ($Q=2$) which yields maximum neutrality.

Thus, if neutrality is a feature that characterizes social science problems, caution should be used when characterizing the fitness landscape by relying on fitness distance correlations or K/N ratios (Galván-Lopez and Poli 2006). As Huyen et al. (1996) argue, a small value for the fitness distance correlation (i.e. $-0.15 < \rho_{FDC} < 0.15$) that would normally be connected with a very rugged landscape, is not informative as to the ease/difficulty of finding the global optimum since local optima, when connected, are no longer local (Huynen et al. 1996). This is further explored by Lobo et al. (2004) who conclude that there is an interplay between the ruggedness and neutrality of the landscape. Their simulations suggest that the desirability of neutrality is contingent on the former. For instance, for rugged landscapes, neutrality is beneficial, but for smooth landscapes neutrality just makes adaptation slower.

In consequence, the measures detailed in the previous section do not necessarily capture the relative ease or difficulty an adaptive solver would have on a landscape that does have neutral ridges.

So far we have ignored issues pertaining to the search behaviour, or rather, following the NK literature, we have taken the one-bit flip as a reference. This assumption however isn't as innocuous as it may seem. As such, the features described above (i.e. either metrics such as the number of local optima that a given problem has, or neutrality or deceptiveness) can only be defined with respect to a neighbourhood function M (Pitzer and Affenzeller 2012). We address these concerns in the following section.

3. Search behaviours

Simon (1956) describes agents of increasing intelligence: from the “simple-minded” organism that is driven by a basic stimulus response rule to a more complex, cognitively endowed, actor. Newell and Simon (1976) further introduce the hypothesis that in order for intelligent search to be better than random search, the space of solution has to “exhibit[s] at least some degree of order and pattern” (Newell and Simon 1976: p. 121). Since it is in the interplay between the structure of the problem and the search heuristic that is the focus in these models, it is not just the fitness function that is important, but also the particular search behaviours with which agents are endowed. Furthermore, since search behaviours are not as well defined in the NK model for organizational search as they are in the biological equivalents (i.e. selection, genetic drift, mutation and recombination cf. Huxley 2010) scholars need to build a different empirical foundation in what concerns search behaviours. Still, the most prevalent models in organizational theory seem closer to the ‘mindless particle’ end of the spectrum (cf. Winter et al. 2007; Csaszar and Levinthal 2015), in what concerns an agent’s “ability to store and manipulate symbols” (Newell and Simon 1976: p.115).

Several authors (Jones 1995; Frenken et al. 1999) have cautioned that the NK ruggedness is in fact a property of the landscape and not a property of the task environment, as defined above: the ruggedness is a property of L and not f or its domain. That is, ruggedness, as defined in Kauffman's original model, is assumed to be given by a one bit mutation of the candidate solution (1993). However, as Frenken et al. (1999) point out, the assumption of one-bit flip is of limited relevance in the context of human search behaviours, since such an one-bit conception does not fit human behaviour: human problem solvers are less likely to engage in small, incremental changes. Billinger et al. (2013) e.g. find that the average search distance is above two. As such, the ruggedness of the NK landscape does not allow for intuitions to be formed about problem hardness *in general* (i.e. the likelihood that a solver can find the optimal solution efficiently, provided that local search is not the only or the dominant search heuristic). For example, with a basic hill climbing technique (without random long jumps), a rugged NK landscape ($K>0$) is reduced to one single peak (likely a local optima) where the solver gets stuck. Thus, even simple heuristics as "hill climbing with long jumps" vastly improve the search process as they are able to cover the entire rugged landscape. In a problem-solving context the heuristic chosen is of outmost importance and the task environment's statistical features (Kauffman 1993) do not preclude the existence of a powerful search heuristic that can in fact resolve to a flat, single-peaked landscape.

3.1. Search in organizational theory

Following the early studies, recent NK model extensions take a more nuanced view on what organizations do in their attempt to find solutions and thus focus on different search behaviours and their performance on landscapes of different complexity (as judged by the {N, K} pair).

In the canonical NK model, the search heuristic is inspired by a simple evolutionary mechanism: adaptive mutation. This very simple search algorithm is actually very efficient. Hill-climbing is one of the most powerful domain-general search algorithms (Russel and Norvig 2010). It also provides earlier models with a straight-forward way of implementing bounded rationality assumptions. A solver endowed with such a simple heuristic is clearly light years away from the all-knowing *homo economicus*, but this simple solver can still solve complex problems, such as identifying the optimal configuration of organizational routines (Levinthal 1997).

Another type of search is where the solver is capable of evaluating all solutions in the one-bit mutation set and chooses the one that maximizes his performance (i.e. 'offline search' (Gavetti 2005)). Other scholars have taken into account the fact that the assumption of bounded rationality is not violated if solvers are endowed with more intelligent heuristics (Winter et al. 2007). For example, Gavetti and Levinthal (2000) allow agents to be directed in their search by

representations of the search space that are attributed to them a priori, while Winter et al. (2007) assume that agents have exogenously attributed “preferred direction”.

3.2. Alternative search behaviours

Wright (1932) argues that the fundamental mechanism behind speciation must be a non-adaptive one, i.e. it cannot be that hill climbing alone can account for the tremendous variety in species (Gavrilets 1999).

Natural computing was quick to adopt biological mechanisms and adapt: stochastic hill-climbing, first-choice climbing and random-restart hill-climbing were the first natural successors that already showed a marked improvement over the performance of the canonical hill-climber. Although still extremely simple, these algorithms capture fundamental dynamics of adaptation. For example, Figure 7 illustrates the exploration/exploitation trade-off via the variance in performance for a random restart one-bit hill-climber on an NK landscape ($N=8$, $K=3$). Since such a landscape has a number of local optima, a hill-climber with zero probability of restarting would quickly climb up the nearest peak and the search would stop. The probability of identifying the optimal solution is strictly dependent on the number and size of basins of attraction¹ for these local optima. For large numbers of local optima and large basins of attraction, the likelihood that the agent finds himself in the vicinity of the global optimum decreases and so does his probability of success. As the probability of restart increases, the solver also increases his chances of “landing” in the right part of the landscape. Evidently, a high probability of restart (in this case $p>0.2$) decreases performance since the agent engages excessively in exploration (sampling the landscape) and not enough in exploitation (hill-climbing).

¹ The areas around local optima that lead a hill-climbing algorithm directly to the local peak.

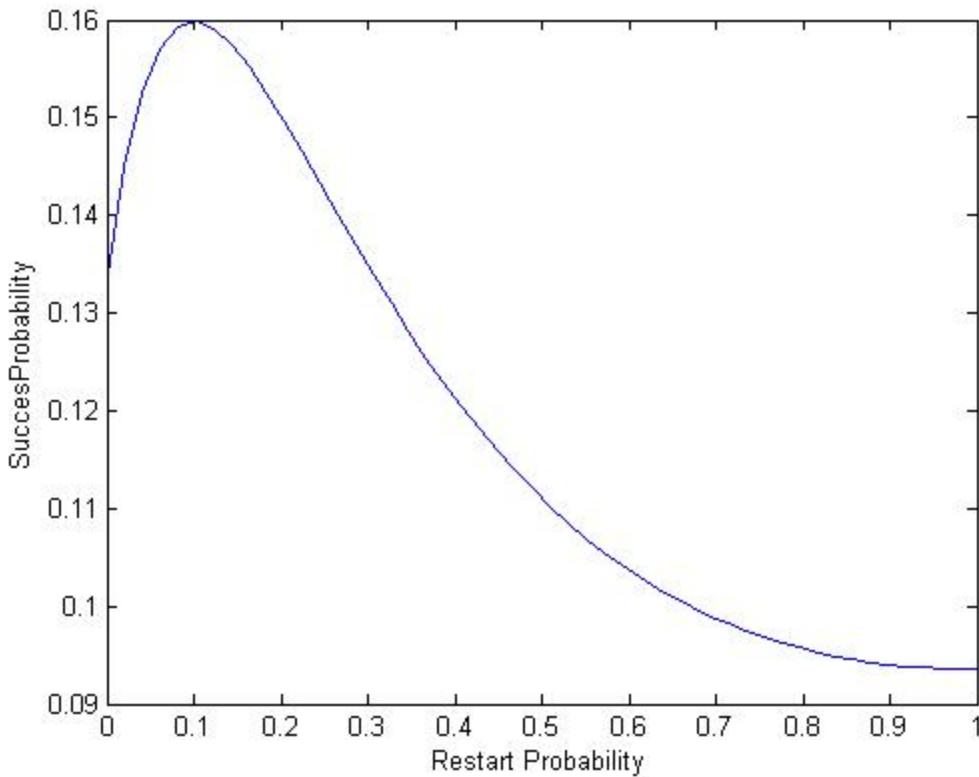


Figure 7. Performance varying with restart probability for a stochastic hill-climber.

Genetic algorithms (Holland 1992) are another class of algorithms that is inspired by evolutionary mechanisms. The major difference between the former and the latter was the use of genetic operators: in addition to selection and mutation, genetic algorithms rely on gene recombination as well (Affenzeller et al. 2009). Genetic algorithms are widely used in practice (e.g. Matthey et al. 2007) as are a number of various other computing tools ranging from fuzzy logic and belief calculus to machine learning like inductive logic programming (Konar 1999).

However, as shown by Wolpert and Macready (1999), an algorithm's average performance is determined by how much knowledge regarding the optimization function is incorporated into the search heuristic (the “No Free Lunch Theorem” for optimization). Thus, computer science has moved away from general purpose evolutionary algorithms with their limited knowledge of the problem space to algorithms that are designed specifically for the problem at hand. To illustrate this point, we compare the performance of three different algorithms on a hierarchical landscape (Figure 8): one-bit flip, and the same “chunking” algorithm that relies on 26 operations that are derived by taking into account the particularities of the hierarchical problem (H-XOR, N=8) and random search.

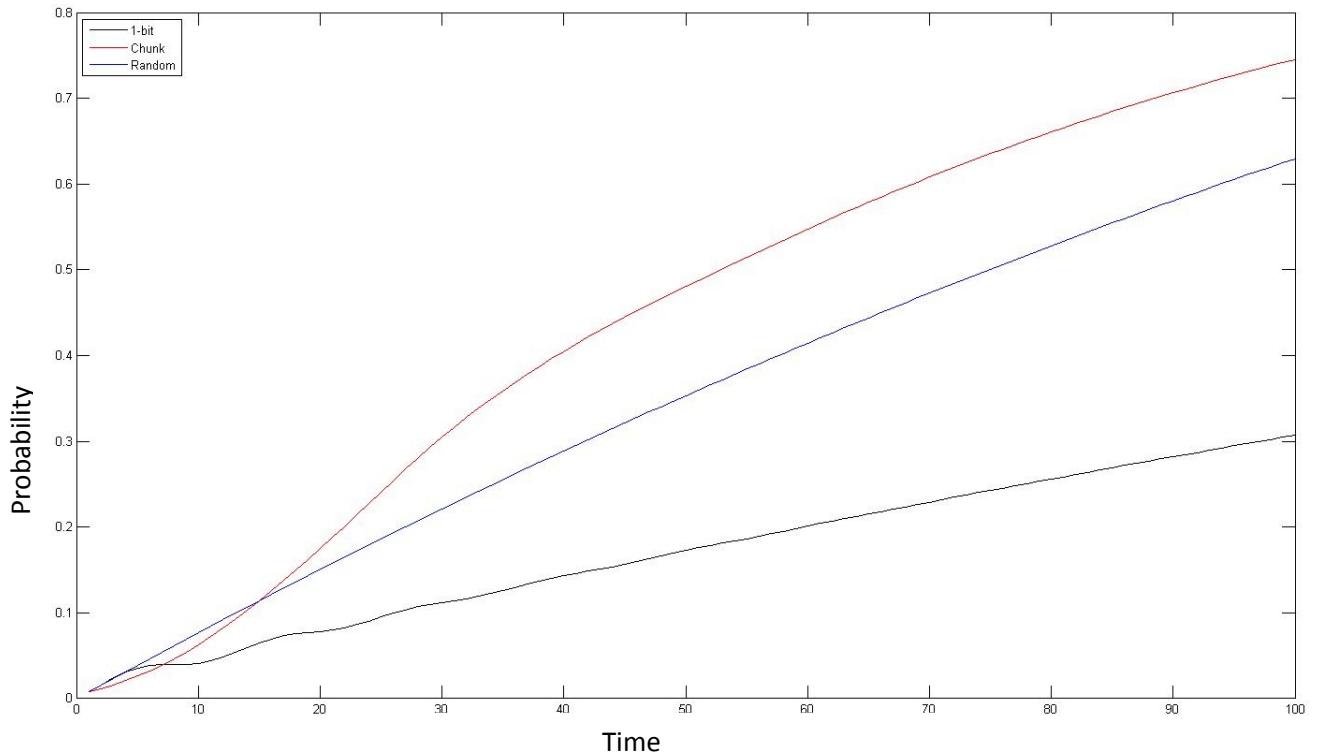


Figure 8. Performance of three different search algorithms on an HXOR landscape ($N=8$): 1-bit flip, chunking and random.

We show how the chunking algorithm significantly outperforms both the hill-climber and a random search. Computational experiments also endorse this: algorithms that embed these principles out-perform traditional recombination (genetic algorithms) or local search, since a solver that uses an inappropriate problem decomposition effectively generates a rugged landscape (Cioffi-Revilla et al. 2012).

Memory in computational agents

Finally, most computational approaches discussed so far rely on agents which do not have memory¹. The process of problem solving they describe is path dependent and solvers attempt to improve their current best performance, but, for example, in an eight-bit NK problem, at each time-step, a solver still has 255 ($2^8 - 1$) possible combinations to choose from. Naturally, the search heuristic implemented restricts the set of available moves, but the assumption is still strong and unjustified. The equivalent claim: ‘draws from an urn with or without replacement yield identical probabilities of success’ highlights the problem with disregarding memory.

¹ Note this is different from Gavetti (2005) who allows agents to have a ‘cognitive memory’, that is a set of different partial representations of the landscape.

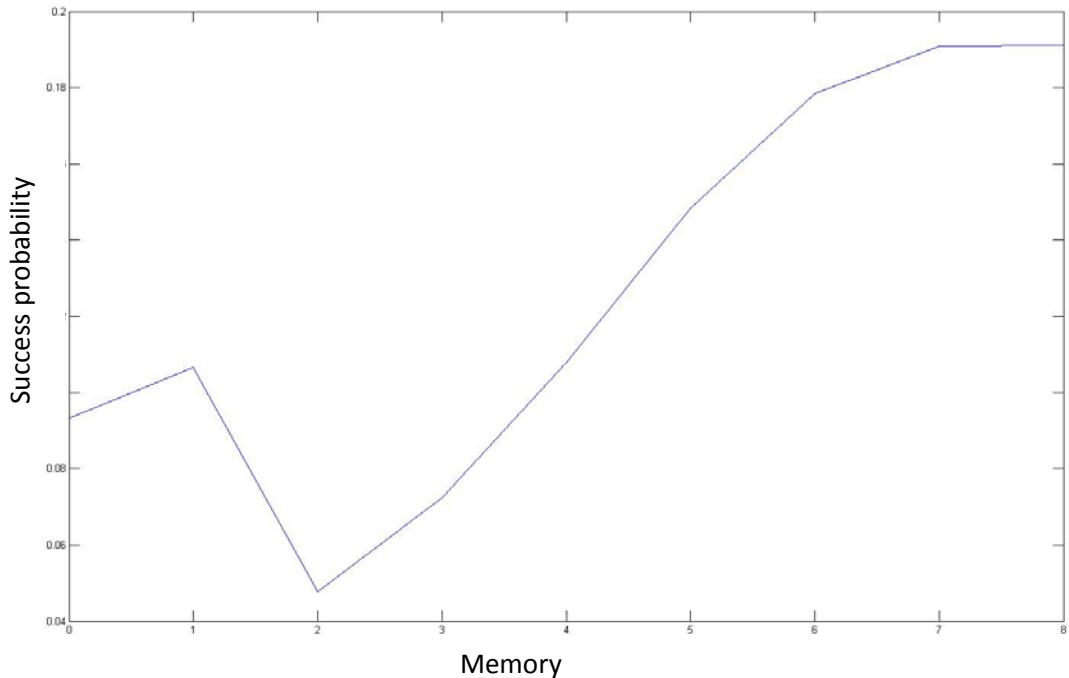


Figure 9. Performance of a 1-bit flip hill climber as a function of the number of remembered previous solutions.

Figure 9 illustrates how the performance of a hill-climber changes as a function of the number of solutions he remembers. Note that memory is not ‘universal a recipe for success’, since remembering past success restricts the following moves that the hill-climber can make, which in turn means that the agent is more easily stuck in a suboptimal solution.

It is beyond the scope of this paper to give a comprehensive overview of the history of natural computing. It should be noted that there is a great variety of increasingly sophisticated algorithms for solving optimization problems: for instance, in addition to algorithms inspired by evolutionary theories, computer science has also developed methods inspired by the human nervous system (artificial neural networks), the collective behaviour of groups of organisms (e.g. particle swarm optimization) or quantum physics (Rozenberg et al. 2011).

Some of these algorithms are directly related to the previous discussion on features of fitness landscapes and rely on different measures of search performance such as e.g. search dispersion to adapt the search as it progresses (Maier et al. 2014). Others rely on more classical mappings of the search space, such as decision trees (Huys et al. 2012) or Bayesian models (Pelikan 2005). These mappings are continuously adapted during the search process, using the information the agent gathers by interacting with the environment.

To sum up, fuelled by the development of Artificial Intelligence, in computer science literature there has been a shift in focus from agents that can be thought of as ‘mindless particles’, to ‘smarter’ agents, which, without being perfectly rational, are capable of observing or partially observing their environment and constructing “beliefs” that allow them to generate subsequent moves (Russel and Norvig 2010).

4. Discussion

Based on the review of organizational theory, computer science and biology perspectives, in the following we try to outline future potential avenues for expanding the models as well as discuss potential implications. Although we acknowledge that models are to be thought of as useful simplifications of reality, in this section we highlight how previously made simplifications and NK legacy elements actually restrict our ability to rely on these models to further our understanding about human/organizational problem-solving dynamics.

Modelling landscapes and complexity

Irrespective of the NK model’s biological origins, there are no clear empirical specifications or constraints about how fitness landscapes should be conceptualized. In fact, Kauffman acknowledges that the contributions of individual alleles’ are drawn randomly from a uniform distribution, since the exact contribution of a gene to an organism’s fitness is not known (Kauffman 1993). The use of the NK canonical function has been justified in a similar manner by the fact that modellers are interested in how the evolutionary process “typically” unfolds (Rivkin 2000: p.828). However, as Jones and Forrest (1995) and Frenken et al. (1999) show, for $K>3$ NK landscapes become quickly uncorrelated. Given that organizational problems and innovation problems are rarely the real-life correspondent of either completely smooth ($K=0$) or completely uncorrelated landscapes, this questions whether NK functions are meaningful in modelling the dynamics of problem solving systems.

While a mapping between the various components of a given solution and its performance (e.g. the modules of the educational app described in the introduction and its computational performance) is not trivial to make in an organizational setting, it can be even harder to see how it can be argued that such mapping is random and more importantly how the NK captures the interdependence structure of a *typical organizational problem*.

Indeed a number of scholars have attempted more meaningful extensions of the NK model (Rivkin and Siggelkow 2003; Ethiraj and Levinthal 2004; Ethiraj et al. 2008). For example, acknowledging the problematic nature of the random interdependence structure in NK functions, Ethiraj and Levinthal (2004) impose a block near-modular design on top of the NK matrix and then allow solvers to optimize inside the modules as well as resort to recombinative practices.

They show that erring on the side of too much decomposition is detrimental to search efficacy. However, Watson and Pollack (2005) caution that according to Simon's definition of modularity, it is only on the *short-term* that modules are quasi-independent and on the contrary, long-term dynamics should assume stronger inter-module interdependence. This kind of interdependence, they argue, is not captured by structural (as opposed to functional) models of modularity.

One further limitation of these extensions is that the properties of these pseudo-NK landscapes are not as established as results for the canonical NK, so it is not obvious whether assumptions about the structure of the problem or assumptions about the search behaviours are driving the simulation results. For instance, recent work shows that imposing a block structure on the NK interactions qualitatively changes the structure of the landscape by actually *diminishing* the number of evolutionary paths towards the global maximum, under the SSWM¹ condition (Schmiegel and Krug 2014), while Hebron et al. (2008) show that imposing a scale free structure on an NK landscape leads to longer adaptive walks and more clustering of optima in the landscape. In addition, all the extensions referenced in this study (e.g. Gavetti and Levinthal 2000; Rivkin and Siggelkow 2003; Ethiraj and Levinthal 2004; Gavetti 2005; Ethiraj et al. 2008) have as a basis the canonical form for the NK fitness function, which averages across individual fitness contributions. This, as McKelvey et al. (2013) show, inevitably generates the same result: with the increase in N and K, the value for the fitness function converges towards the mean of the uniform distribution (0.5) and this skews the interpretation of the simulation findings. The particular way the fitness function is generated in the NK model is also what Szendro et al. (2013) argue makes such models less amendable to being able to capture different levels of epistatic effects.

In biology, the empirical evidence towards the existence of multi-modal landscapes with numerous epistatic interactions continues to increase (Østman and Adami 2014) with scholars inquiring if it is reasonable to assume that adaption is taking place on a highly uncorrelated landscapes and if it is meaningful to assume there are no 'neutral ridges' (Gavrilets 1999) or hierarchical interdependence (Segre et al. 2005). With few exceptions (e.g. Fleming and Sorenson 2001; 2004), a similar, empirically grounded, discussion about how can we create meaningful landscapes for organizational problems seems to be missing.

Paradoxically, we argue that one way forward, as suggested by current developments in computer science, is to revert to what Wright (1932) and Kauffman (1993) originally proposed: relying on a fitness landscape to first acquire a "rough" image of a problem class, instead of investigating specific instances (Pitzer and Affenzeller 2012, Malan and Engelbrecht 2014). Investigations into the topology of the fitness landscape would allow for a better understanding

¹ Strong selection, weak mutation.

of the dynamic processes of adaptation in a similar vein to previous considerations: e.g. rugged landscapes are likely to trap solvers in suboptimal peaks, deceptive landscapes are likely to attract solvers towards suboptimal optima etc. Alongside the few methods introduced in this paper, a number of methods for landscape analysis have been comprehensively developed in this literature.

Modelling search behaviours

While previous paragraphs discuss issues related to how we conceptualize landscapes, we now turn to search behaviours. We have already discussed how several studies have questioned and attempted to expand the human problem solving models of search beyond its biological origins. Their approach is largely driven by theoretical concerns, but recent research is attempting to do the same driven by empirical results. This empirically driven approach allows for a better specification of search behaviours which in turn results in better models. The hill-climber most often used in management science is stochastic-restart hill climbing, rather than the canonical hill-climber. Before Billinger et al. (2013) few if any scholars spend time explaining some of the subsequent (seemingly innocuous) modelling decisions. However, as shown earlier (Figure 7), the performance of a hill-climber with random restarts differs significantly from the performance of a hill-climber without random restarts.

These results mirror a previous study conducted by Mason and Watts (2012). By comparing the performance of actual solvers and computational agents, Mason and Watts (2012) show that heterogeneity in terms of search behaviours has the potential to greatly influence the outcome of the search process. Whether relying on constructs such as attention control (Laureiro-Martínez et al. 2015), intelligence (Steyvers et al. 2009) or cognitive styles (Kirton 1976), there is empirical evidence that there is a great heterogeneity when it comes to human search behaviours, but such heterogeneity is rarely taken into account in modelling approaches (Miller and Page 2009). This is also endorsed by empirical results which suggest that humans are capable of solving hard computational problems (Carruthers and Stege 2013), evidence to the fact that humans have far more sophisticated search strategies.

We have shown how assumptions about memory influence performance (Figure 8). Memory is not only important when it comes to restricting the search space (e.g. Gavetti 2005), but also as one way of advancing more plausible assumptions into modelling human problem-solving. Human problem solvers are not guided in their search merely by immediate feedback, but also by representations of the problem they form over time, via accumulated experience (Doll et al. 2012). The modelling community has struggled to capture this via ‘preferred direction’ (Winter et al. 2007) or partial representations of the solution (Gavetti and Levinthal 2000, Gavetti 2005),

but the attempt is complicated by the fact that one assumes an agent has a problem representation *before* solving a problem. Where could this representation come from?

One answer lies in the fact that the mental models or representations that solvers use to guide their search are themselves adaptive and subject to reinforcement learning (Miller and Page 2009). It is then not only in the generation and evaluation of *solutions* that feedback loops are important (Bonabeau 2009), but the same mechanism can account for the emergence and evolution of problem representations as solvers engage in “imagining future events” (Schacter et al. 2007: p. 659). This approach insures that solvers do not have to start with an ex-ante map of the landscape, but gradually formulate it, or as Simon and Newell describe it (1972): they incorporate knowledge into their search heuristic. The computational solutions involve the implementation of machine learning techniques (Rand 2006) which can range from Bayesian algorithms (Pelikan et al. 2003), decision trees (Huys et al. 2012) and more recently deep learning (Mnih et al. 2015). Irrespective of the details, these methods have the potential to bring forward a middle way in modelling problem solving that is in keeping with the bounded rationality assumption, but at the same time allows more than “ant-like” behaviour (Winter et al. 2007) consistent with some of the theories regarding human cognition (Le 2013). The underlying idea is that while engaging in problem solving, as more information is available human solvers are able to detect and abstract essential features, in the same way visual pattern recognition works (Roland and Gulyás 1995). As a result of that, in structured environments subsequent variations or mutations are not random, but closer to what biologists call “facilitated variation”: generated new solutions are potentially *useful* (Parter et al. 2008: p.2). This however, is not possible in NK landscapes where fitness contributions are drawn randomly from an underlying distribution since, as Watson et al. (2011) argue the environment has to display a certain degree of regularity that the agents can exploit.

Conclusion

While the simulation approach has gained attention in high-status outlets within organizational theory, we acknowledge and address the sceptical concerns still being raised about theoretical assumptions (Fioretti 2013) and the weak empirical grounding of these assumptions (Chang and Harrington 2006; Mason and Watts 2012). Much like in the original biological setup, the organizational literature has had simplistic assumptions about agent behaviour, embedded in a relatively undefined fitness landscape (Ganco and Hoetker 2008, McKelvey et al. 2013). However, unlike micro-biology, where evolutionary forces are well known (Huxley 2010), defining human search behaviours in this conceptual framework turns out to be elusive: how do

we model and define what constitutes intelligent boundedly rational behaviour; for instance, do agents have memory and how good are they at interpreting the landscape?

Given our limited understanding about the genotype-phenotype mapping in a technological setting (Solée et al. 2013), we suggest that the focus should not be on the statistical features of the landscape to be searched under the one-bit flip condition, but on how the interplay of search behaviours and the different natures of interdependence structures translates into problem solving performance. Only then we can focus on how the search can be best organized in such a way that solvers effortlessly find themselves in the vicinity of the optimal solution (cf. Felin and Zenger 2014).

Additionally, we argue that any model of organizational learning should allow for more plausible (and if possible, empirically validated) assumptions regarding learning and expertise. For uninformed solvers, the fitness landscape will be extremely large and rugged, as they have to deal with a seemingly unconstrained search space (Zhang and Norman 1994). However, a problem representation works by effectively constraining the search space, generating a different set of possible solutions. Early investigations in the use of problem representations (Kotovsky and Simon 1990) show that knowledge about the landscape changes the structure of the landscape: “*<the easy problem> problem*” (Winter 2004)¹. These ideas are not foreign to organizational literature which has a long tradition of looking at managerial decisions through the lenses of cognitive frames (or schemas) which seem to be the primary source of difficulty for organizations in turbulent environments (Kaplan 2008; Bingham and Kahl 2013). Still, most modelling approaches do not take into account this perspective about how cognitive frames change the search behavior *and implicitly* the landscape.

We argue that moving away from “armchair speculations” Simon (1982) regarding human search behaviour and the nature of the problem is essential in these settings as seemingly innocuous assumptions can drastically change the problem solving performance. We further identify two potential avenues for future research: focusing on different landscape features and creating “smarter” agents by relying on the recent developments in computer science, both of which, we argue, should be endorsed by empirical calibration and validation.

References

- Affenzeller, M., S. Wagner, et al. (2009). Genetic algorithms and genetic programming: modern concepts and practical applications, Crc Press.

¹ Learning and practice as well as context influence whether a problem is perceived as “easy” by a given solver.

- Afuah, A., & Tucci, C. L. (2012). "Crowdsourcing as a solution to distant search" Academy of Management Review, 37(3), 355-375.
- Altenberg, L. (1997). "NK fitness landscapes." Handbook of evolutionary computation 7: 5-B2.
- Barnett, L. (1997). "TANGLED WEBS--Evolutionary Dynamics on Fitness Landscapes with Neutrality."
- Barnett L. (1998). "Ruggedness and neutrality-the NKp family of fitness landscapes". In Artificial Life VI: Proceedings of the sixth international conference on Artificial life Aug 28 (pp. 18-27).
- Billinger, S., N. Stieglitz, et al. (2013). "Search on rugged landscapes: An experimental study." Organization Science 25(1): 93-108.
- Bingham, C. B. and S. J. Kahl (2013). "The process of schema emergence: Assimilation, deconstruction, unitization and the plurality of analogies." Academy of Management Journal 56(1): 14-34.
- Bonabeau, E. (2009). "Decisions 2.0: The power of collective intelligence." MIT Sloan Management Review 50(2): 45-52.
- Baumann, O. and N. Siggelkow (2013). "Dealing with Complexity: Integrated vs. Chunky Search Processes." Organization Science 24(1): 116-132.
- Baumann, Oliver. "Models of complex adaptive systems in strategy and organization research." Mind & Society (2015): 1-15.
- Burton, R. M. and B. Obel (1980). "A Computer Simulation of the M-Form Hypothesis." Administrative Science Quarterly 25(3): 457-466.
- Carruthers, S. and U. Stege (2013). "On evaluating human problem solving of computationally hard problems." The Journal of Problem Solving 5(2): 4.
- Chang, M. H., & Harrington, J. E. (2006). Agent-based models of organizations. Handbook of computational economics, 2, 1273-1337.
- Cioffi-Revilla, C., K. De Jong, et al. (2012). "Evolutionary computation and agent-based modeling: biologically-inspired approaches for understanding complex social systems." Computational and Mathematical Organization Theory 18(3): 356-373.
- Cohen, M. D., J. G. March, et al. (1972). "A garbage can model of organizational choice." Administrative science quarterly: 1-25.
- Csaszar, F. A. and D. A. Levinthal (2015). "Mental representation and the discovery of new strategies." Strategic Management Journal Forthcoming.
- de Jong, E. D., R. A. Watson, et al. (2005). A generator for hierarchical problems. Proceedings of the 2005 workshops on Genetic and evolutionary computation, ACM.
- Doll, B. B., D. A. Simon, et al. (2012). "The ubiquity of model-based reinforcement learning." Current opinion in neurobiology 22(6): 1075-1081.

- Ethiraj, S. K. and D. Levinthal (2004). "Modularity and innovation in complex systems." *Management Science* 50(2): 159-173.
- Ethiraj, S. K., Levinthal, D., & Roy, R. R. (2008). "The dual role of modularity: innovation and imitation". *Management Science*, 54(5), 939-955.
- Felin, T. and T. Zenger (2014). "Closed or open innovation? Problem solving and the governance choice". *Research Policy* 43(5): 914-925.
- Fioretti, G. (2013). "Agent-based simulation models in organization science." *Organizational Research Methods* 16(2): 227-242.
- Fleming, L., & Sorenson, O. (2001). "Technology as a complex adaptive system: evidence from patent data". *Research Policy*, 30(7), 1019-1039.
- Fleming, L. and O. Sorenson (2004). "Science as a map in technological search." *Strategic Management Journal* 25(8-9): 909-928.
- Franke, J., Klözer, A., de Visser, J. A. G., & Krug, J. (2011). "Evolutionary accessibility of mutational pathways". *PLoS Comput Biol*, 7(8), e1002134.
- Frenken, K., L. Marengo, et al. (1999). Interdependencies, nearly-decomposability and adaptation. *Computational techniques for modelling learning in economics*, Springer: 145-165.
- Galván-López, E., R. Poli, et al. (2011). "Neutrality in evolutionary algorithms... What do we know?" *Evolving Systems* 2(3): 145-163.
- Galván-López E, Poli R. (2006) An empirical investigation of how and why neutrality affects evolutionary search. In *Proceedings of the 8th annual conference on Genetic and evolutionary computation 2006 Jul 8*: 1149-1156. ACM.
- Ganco, M. and G. Hoetker (2009). "NK modeling methodology in the strategy literature: bounded search on a rugged landscape." *Research methodology in strategy and management* 5(2009): 237-268.
- Gavetti, G. and D. Levinthal (2000). "Looking forward and looking backward: Cognitive and experiential search." *Administrative Science Quarterly* 45(1): 113-137.
- Gavetti, G. and D. Levinthal (2001). "Bringing Cognition Back In and Moving Forward." *Journal of Management and Governance* 5(3): 213-216.
- Gavetti, G. (2005). Cognition and hierarchy: Rethinking the microfoundations of capabilities' development. *Organization Science*, 16(6), 599-617.
- Gavrilets, S. (1999). "Evolution and speciation in a hyperspace: the roles of neutrality, selection, mutation and random drift." *Evolution*: 1.
- Geard, N., J. Wiles, et al. (2002). *A comparison of neutral landscapes-NK, NKp and NKq*. Evolutionary Computation, 2002. CEC'02. Proceedings of the 2002 Congress on, IEEE.

- Golberg, David E. (1989) Genetic algorithms in search, optimization, and machine learning.
Addion Wesley, Reading (MA)
- Goldberg, D. (2002). The Design of Innovation (Genetic Algorithms and Evolutionary Computation), Springer New York.
- Hargadon, A. and R. I. Sutton (1997). "Technology Brokering and Innovation in a Product Development Firm." Administrative Science Quarterly 42(4): 716-749.
- He, J., C. Reeves, et al. (2007). "A note on problem difficulty measures in black-box optimization: Classification, realizations and predictability." Evolutionary Computation 15(4): 435-443.
- Hebbron, T., S. Bullock, et al. (2008). NKalpha: Non-uniform epistatic interactions in an extended NK model. Artificial Life XI: Proceedings of the Eleventh International Conference on the Simulation and Synthesis of Living Systems, MIT Press, Cambridge, MA: 234-241
- Hedström, P. and G. Manzo (2015). "Recent Trends in Agent-based Computational Research: A Brief Introduction." Sociological Methods & Research 44(2): 179-185.
- Holland, J. H. (1992). Adaptation in natural and artificial systems: an introductory analysis with applications to biology, control, and artificial intelligence, MIT press.
- Huxley, J.S.: Evolution (2010): The modern synthesis, MIT Press
- Huynen, Martijn A., Peter F. Stadler, and Walter Fontana. "Smoothness within ruggedness: the role of neutrality in adaptation." Proceedings of the National Academy of Sciences 93.1 (1996): 397-401.
- Huys, Q., N. Eshel, et al. (2012). "Bonsai trees in your head: how the Pavlovian system sculpts goal-directed choices by pruning decision trees." PLoS Comput. Biol. 8(3): e1002410.
- Iclanzan, D. and D. Dumitrescu (2007). Overcoming hierarchical difficulty by hill-climbing the building block structure. Proceedings of the 9th annual conference on Genetic and evolutionary computation, ACM.
- Jaimovich, A., Rinott, R., Schuldiner, M., Margalit, H., & Friedman, N. (2010). Modularity and directionality in genetic interaction maps. Bioinformatics, 26(12), i228-i236.
- Jones, T. (1995). Evolutionary algorithms, fitness landscapes and search, Citeseer.
- Jones, T. and S. Forrest (1995). Fitness Distance Correlation as a Measure of Problem Difficulty for Genetic Algorithms. ICGA.
- Kaplan, S. (2008). "Framing Contests: Strategy Making Under Uncertainty." Organization Science 19(5): 729-752.
- Karaboga, D. (2005). An idea based on honey bee swarm for numerical optimization, Technical report-tr06, Erciyes university, engineering faculty, computer engineering department.

- Katada, Y. and K. Ohkura (2006). Estimating the Degree of Neutrality in Fitness Landscapes by the Nei's Standard Genetic Distance-An Application to Evolutionary Robotics. Evolutionary Computation, 2006. CEC 2006. IEEE Congress on, IEEE.
- Kauffman, S. A. (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.
- Knudsen, T., and D. A. Levinthal. "Two faces of search: Alternative generation and alternative evaluation." Organization Science 18.1 (2007): 39-54.
- Konar, A. (1999). Artificial intelligence and soft computing: behavioral and cognitive modeling of the human brain, CRC press.
- Kondrashov, D. A., & Kondrashov, F. A. (2015). Topological features of rugged fitness landscapes in sequence space. Trends in Genetics, 31(1), 24-33.
- Kotovsky, K. and H. A. Simon (1990). "What makes some problems really hard: Explorations in the problem space of difficulty." Cognitive Psychology 22(2): 143-183.
- 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.
- Le, Q. V. (2013). Building high-level features using large scale unsupervised learning. Acoustics, Speech and Signal Processing (ICASSP), 2013 IEEE International Conference on, IEEE.
- Levinthal, D. A. (1997). "Adaptation on rugged landscapes." Management Science 43(7): 934-950.
- Lobo, J., J. H. Miller, et al. (2004). "Neutrality in technological landscapes." Sante Fe Institute Working Paper.
- Maier, H., Z. Kaplan, et al. (2014). "Evolutionary algorithms and other metaheuristics in water resources: current status, research challenges and future directions." Environmental Modelling & Software 62: 271-299.
- Malan, K. M. and A. P. Engelbrecht (2013). "A survey of techniques for characterising fitness landscapes and some possible ways forward." Information Sciences 241: 148-163.
- Malan, K. M. and A. P. Engelbrecht (2014). Fitness landscape analysis for metaheuristic performance prediction. Recent advances in the theory and application of fitness landscapes, Springer: 103-132.
- Marengo, L. and G. Dosi (2005). "Division of labor, organizational coordination and market mechanisms in collective problem-solving." Journal of Economic Behavior & Organization 58(2): 303-326.

- Marengo, L., G. Dosi, et al. (2000). "The structure of problem-solving knowledge and the structure of organizations." *Industrial and Corporate Change* 9(4): 757-788.
- Martignoni, D., Menon, A., & Siggelkow, N. (2015). "Consequences of Misspecified Mental Models: Contrasting Effects and the Role of Cognitive Fit". *Strategic Management Journal*.
- Martin, W. (2001). *Foundations of Genetic Algorithms 2001 (Foga 6)*, Morgan Kaufmann.
- Mason, A. (1995). Non-linearity measure of a problem's crossover suitability. Proceedings of the IEEE Conference on Evolutionary Computation, IEEE.
- Mason, W. and D. J. Watts (2012). "Collective problem solving in networks." *Proceedings of the National Academy of Sciences*.
- Mata, R. and B. von Helversen (2015). "Search and the Aging Mind: The Promise and Limits of the Cognitive Control Hypothesis of Age Differences in Search." *Topics in Cognitive Science* 7(3): 416-427.
- Matthey, D., J. Wang, et al. (2007). "Enhanced bonding of gold nanoparticles on oxidized TiO₂ (110)." *Science* 315(5819): 1692-1696.
- McClymont, K. (2013). Recent advances in problem understanding: Changes in the landscape a year on. Proceedings of the 15th annual conference companion on Genetic and evolutionary computation, ACM.
- McKelvey, B., Li, M., Xu, H., & Vidgen, R. (2013). Re-thinking Kauffman's NK fitness landscape: From artifact and groupthink to weak-tie effects. *Human Systems Management*, 32(1), 17-42.
- Miller, J. H. and S. E. Page (2009). *Complex adaptive systems: an introduction to computational models of social life: an introduction to computational models of social life*, Princeton university press.
- Miner, A. S., P. Bassof, et al. (2001). "Organizational improvisation and learning: A field study." *Administrative science quarterly* 46(2): 304-337.
- Mnih, V., K. Kavukcuoglu, et al. (2015). "Human-level control through deep reinforcement learning." *Nature* 518(7540): 529-533.
- Naudts, B. and A. Verschoren (1999). "Epistasis and deceptivity." *Simon Stevin-Bulletin of the Belgian Mathematical Society* 6: 147-154.
- Naudts, B., & Kallel, L. (2000). "Comparison of summary statistics of fitness landscapes." *IEEE Trans. Evol. Comp*, 4, 1-15.
- Newell, A. and H. A. Simon (1972). *Human problem solving*. Englewood Cliffs, N.J., Prentice-Hall.
- Newell, A. and H. A. Simon (1976). "Computer science as empirical inquiry: Symbols and search." *Communications of the ACM* 19(3): 113-126.

- Newman, M. and R. Engelhardt (1998). "Effects of selective neutrality on the evolution of molecular species." *Proceedings of the Royal Society of London B: Biological Sciences* 265(1403): 1333-1338.
- Nickerson, J. A. and T. R. Zenger (2004). "A knowledge-based theory of the firm—The problem-solving perspective." *Organization science* 15(6): 617-632.
- Østman, B. and C. Adami (2014). Predicting evolution and visualizing high-dimensional fitness landscapes. *Recent Advances in the Theory and Application of Fitness Landscapes*, Springer: 509-526.
- Østman, B., A. Hintze, et al. (2010). "Critical properties of complex fitness landscapes." *arXiv preprint arXiv:1006.2908*.
- Parter, M., Kashtan, N., & Alon, U. (2008). Facilitated variation: how evolution learns from past environments to generalize to new environments. *PLoS Comput Biol*, 4(11), e1000206.
- Pelikan, M. (2005). Bayesian Optimization Algorithm. *Hierarchical Bayesian Optimization Algorithm*, Springer: 31-48.
- Pelikan, M., D. E. Goldberg, et al. (2000). *Hierarchical Problem Solving and the Bayesian Optimization Algorithm*. GECCO.
- Pelikan, M., D. E. Goldberg, et al. (2003). *Hierarchical Bayesian optimization algorithm: Toward a new generation of evolutionary algorithms*. New York, Ieee.
- Pitzer, E. and M. Affenzeller (2012). A comprehensive survey on fitness landscape analysis. *Recent Advances in Intelligent Engineering Systems*, Springer: 161-191.
- Rand, W. (2006). *Machine learning meets agent-based modeling: When not to go to a bar*. the Proceedings of Agent.
- Reidys, C. M. and P. F. Stadler (2002). "Combinatorial landscapes." *SIAM review* 44(1): 3-54.
- Rivkin, J. W. (2000). "Imitation of complex strategies." *Management Science* 46(6): 824-844.
- Rivkin, J. W. and N. Siggelkow (2003). "Balancing search and stability: Interdependencies among elements of organizational design." *Management Science* 49(3): 290-311.
- Rivkin, J. W., & Siggelkow, N. (2007). "Patterned interactions in complex systems: Implications for exploration". *Management Science*, 53(7), 1068-1085.
- Roland, P. and B. Gulyás (1995). "Visual memory, visual imagery, and visual recognition of large field patterns by the human brain: functional anatomy by positron emission tomography." *Cerebral Cortex* 5(1): 79-93.
- Romero, Philip A., and Frances H. Arnold. "Exploring protein fitness landscapes by directed evolution." *Nature Reviews Molecular Cell Biology* 10.12 (2009): 866-876.
- Rothlauf, F. (2011). *Design of modern heuristics: principles and application*, Springer Science & Business Media.

- Rozenberg, G., T. Bäck, et al. (2011). Handbook of natural computing, Springer Publishing Company, Incorporated.
- Russel, S. and P. Norvig (2010). "Artificial Intelligence: A Modern Approach (3d edition)." EUA: Prentice Hall.
- Schacter, D. L., D. R. Addis, et al. (2007). "Remembering the past to imagine the future: the prospective brain." Nature Reviews Neuroscience 8(9): 657-661.
- Schmiegelt, B. and J. Krug (2014). "Evolutionary accessibility of modular fitness landscapes." Journal of Statistical Physics 154(1-2): 334-355.
- Segre, D., DeLuna, A., Church, G.M. and Kishony, R., 2005. Modular epistasis in yeast metabolism. Nature genetics, 37(1), pp.77-83.
- Siggelkow, N. and D. A. Levinthal (2003). "Temporarily Divide to Conquer: Centralized, Decentralized, and Reintegrated Organizational Approaches to Exploration and Adaptation." Organization Science 14(6): 650-669.
- Shoham, Y., R. Powers, et al. (2007). "If multi-agent learning is the answer, what is the question?" Artificial Intelligence 171(7): 365-377.
- Simon, H. A. (1956). "Rational choice and the structure of the environment." Psychological review 63(2): 129.
- Simon, H. A. (1962). "The architecture of complexity." Proceedings of the American philosophical society 106(6): 467-482.
- Simon, H. A. (1974). "The structure of ill structured problems." Artificial intelligence 4(3): 181-201.
- Simon, H. A. (1982). Models of bounded rationality: Empirically grounded economic reason (Vol. 3). MIT press.
- Simon, H. A. (1996). The sciences of the artificial. Cambridge, Mass., MIT Press.
- Simon, H. A. and A. Newell (1971). "Human problem solving: The state of the theory in 1970." American Psychologist 26(2): 145.
- Simon, H. A., M. Augier, and J. G. March. Models of a man: Essays in memory of Herbert A. Simon. MIT Press, 2004.
- Smith, T., P. Husbands and M. O'Shea. (2002). "Fitness landscapes and evolvability." Evolutionary computation 10.1 : 1-34.
- Solée, R. V., S. Valverde, et al. (2013). "The evolutionary ecology of technological innovations." Complexity 18(4): 15-27.
- Sorenson, O. (2002). Interorganizational complexity and computation. Companion to organizations, 664-685.
- Sternberg, R. J., & Grigorenko, E. L. (1997). Are cognitive styles still in style? American psychologist, 52(7), 700.

- Sørensen, J. J. W., M. K. Pedersen, et al. (2015). "Exploring the Quantum Speed Limit with Computer Games." [arXiv preprint arXiv:1506.09091](#).
- Szendo, I. G., Schenk, M. F., Franke, J., Krug, J., & De Visser, J. A. G. (2013). Quantitative analyses of empirical fitness landscapes. [Journal of Statistical Mechanics: Theory and Experiment](#), 2013(01), P01005.
- Steyvers, M., Lee, M. D., & Wagenmakers, E. J. (2009). A Bayesian analysis of human decision-making on bandit problems. [Journal of Mathematical Psychology](#), 53(3), 168-179.
- Todd, P. M., T. T. Hills, et al. (2012). [Cognitive search: Evolution, algorithms, and the brain](#), MIT press.
- Thrun, S. (1995). "Learning to play the game of chess." [Advances in neural information processing systems](#) 7.
- Valente, M. (2008). Pseudo-NK: An enhanced model of complexity, LEM Working Paper Series.
- Van der Maaten, L. and G. Hinton (2008). "Visualizing data using t-SNE." [Journal of Machine Learning Research](#) 9(2579-2605): 85.
- Vuculescu, O. (2015). Searching far away from the lamp-post: an agent based model, presented at [Organizing Crowds and Innovation Conference](#), Oxford 2015
- Watson, R. A. and J. B. Pollack (1999). [Hierarchically consistent test problems for genetic algorithms](#). Evolutionary Computation, 1999. CEC 99. Proceedings of the 1999 Congress on, IEEE.
- Watson, R. A., & Pollack, J. B. (2005). Modular interdependency in complex dynamical systems. [Artificial Life](#), 11(4), 445-457.
- Watson, R. A., Buckley, C. L., & Mills, R. (2011). Optimization in “self-modeling” complex adaptive systems. [Complexity](#), 16(5), 17-26.
- Weinberger, E. D. (1996). Np completeness of kauffman's nk model, a tuneable rugged fitness landscape.
- Weise, T., M. Zapf, et al. (2009). Why is optimization difficult? [Nature-Inspired Algorithms for Optimisation](#), Springer: 1-50.
- Winter, S, G. "The easy problem problem" in Simon et al 2004 *Models of a man: Essays in memory of Herbert A. Simon*. MITPress, 2004
- Winter, S. G., G. Cattani, et al. (2007). "The value of moderate obsession: Insights from a new model of organizational search." [Organization Science](#) 18(3): 403-419.
- Wolpert, David H., and William G. Macready. "No free lunch theorems for optimization." [Evolutionary Computation, IEEE Transactions on](#) 1.1 (1997): 67-82.
- Wood, S., J. Busemeyer, et al. (2005). "Older adults as adaptive decision makers: evidence from the Iowa Gambling Task." [Psychology and aging](#) 20(2): 220.

Wright, S. (1932). The roles of mutation, inbreeding, crossbreeding, and selection in evolution, na.

Yu, T. L., D. E. Goldberg, et al. (2009). "Dependency Structure Matrix, Genetic Algorithms, and Effective Recombination." Evolutionary Computation 17(4): 595-U593.

Zhang, J. and D. A. Norman (1994). "Representations in distributed cognitive tasks." Cognitive science 18(1): 87-122.

Zott, C. (2002). "When Adaptation Fails An Agent-based Explanation of Inefficient Bargaining Under Private Information." Journal of Conflict Resolution 46(6): 727-753.

Appendix 1

H-XOR function

The H-XOR function (Watson and Pollack 1999) is given by applying recursively an ‘exclusive or’ transformation onto the solution string where adjacent positions are considered starting with the leftmost. For instance, a {1010 0010} string becomes first {11 -1} and then {- -}. Once the transformation is completed, the payoff function rewards each non-null position in the hierarchy. Thus, a solution which contains an alternating pattern {1010 1010} would give a better score than a {1111 1111} since it will generate payoffs at lower levels of the hierarchy as well. The second level transformation for the first solution is {11 11} while for the second it is {-- --}. The maximum score is given by {1001 0110} or symmetrically by {0110 1001} (see a more extensive description in paper 1).

Operations for the ‘chunking algorithm’

Chunks 8

Inverse all	e.g. 01111111->10000000
Mirror all	e.g. 01111111->11111110

Chunks 4 4

Inverse the 1 st chunk	e.g. 0111 1111->1000 1111
Inverse the 2 nd chunk	e.g. 0111 1111->0111 0000
Mirror the 1 st chunk	e.g. 0111 1111->1110 1111
Mirror the 2 nd chunk	e.g. 0111 1110->0111 0111
Permute the 1 st and 2 nd chunk	e.g. 0111 1111->1111 0111

Chunks 3 2 3

Inverse the 1 st chunk	e.g. 011 11 111->100 11 111
Inverse the 2 nd chunk	e.g. 011 11 111->011 00 111
Inverse the 3 rd chunk	e.g. 011 11 111->011 11 000
Mirror the 1 st chunk	e.g. 011 11 111->110 11 000
Mirror the 2 nd chunk	e.g. 011 10 111->011 01 000

Mirror the 3rd chunk e.g. 011 11 011->011 11 110

Permute the 1st and 3rd chunk e.g. 011 11 111->111 11 011

Chunks 2 2 2

Inverse the 1st chunk e.g. 01 11 11 11->10 11 11 11

Inverse the 2nd chunk e.g. 01 11 11 11->01 00 11 11

Inverse the 3rd chunk e.g. 01 11 11 11->01 11 00 11

Inverse the 4th chunk e.g. 01 11 11 11->01 11 11 00

Mirror the 1st chunk e.g. 01 01 01 01->10 01 01 01

Mirror the 2nd chunk e.g. 01 01 01 01->01 10 01 01

Mirror the 3rd chunk e.g. 01 01 01 01->01 01 10 01

Mirror the 4th chunk e.g. 01 01 01 01->01 01 01 10

Permute the 1st and 2nd chunk e.g. 01 11 11 11->11 01 11 11

Permute the 1st and 4th chunk e.g. 01 11 11 11->11 11 11 01

Permute the 2nd and 3rd chunk e.g. 11 01 11 11->11 11 01 11

Permute the 3rd and 4th chunk e.g. 11 11 01 11->11 11 11 01

Watson, R. A. and J. B. Pollack (1999). Hierarchically consistent test problems for genetic algorithms. Evolutionary Computation, 1999. CEC 99. Proceedings of the 1999 Congress on, IEEE.

Searching far away from the lamp post: an agent-based model¹

Oana Vuculescu

Department of Management, Aarhus University, Denmark

Abstract

This paper presents insights from a laboratory experiment on human problem solving in a combinatorial task. I rely on a hierarchical rugged landscape to explore whether and how human problem solvers are able to detect and exploit patterns in their search for an optimal solution. Empirical findings suggest that solvers do not engage only in local and random distant search, but as they accumulate information about the problem structure, solvers make “model-based” moves, a type of cognitive search. I then calibrate an agent based model of search to analyse and interpret the findings from the experimental setup and discuss implications for organizational search.

1. Introduction

A central idea in the problem solving literature is that the process of discovering a solution can be conceptualized as an adaptive search through a space of alternatives (Simon 1962). Research so far has emphasized the clear distinction between local and distant search, focusing primarily on “the where” (local/distant) to the detriment of “the how” (search heuristics) (Lopez-Vega et al. Forthcoming). That is, despite the fact that scholars have acknowledged that “even simple models of the world have a tremendous potential to guide search processes” (Gavetti and Levinthal 2000: p.137), less literature deals with search strategies that are non-local and cognition driven (Gavetti and Levinthal 2000). There is less attention on further identifying individual search behaviours, which are the basic elements of organizational search, and this limits our understanding about how these processes aggregate into macro-level behaviour (e.g. a collective problem solving approach such as crowdsourcing).

Herbert Simon’s seminal (1956) work on bounded rationality and the metaphor of problem solving as a search have inspired a large number of organizational scholars and works from the influential March (1991), to organizational learning (Levitt and March 1988) and finally

¹ The current version of this paper was presented at Organizing Crowds and Innovation Conference, Oxford 2015 and received a conditional acceptance for a special issue of “Strategic Organization”: *Organizing Crowds and Innovation*, edited by Teppo Felin (University of Oxford), Karim Lakhani (Harvard University) and Michael Tushman (Harvard University)

evolutionary economics (Nelson and Winter 2002; Beinhocker 2006). One of the most prominent streams of research derived from this tradition relies on the metaphor of search on a rugged landscape i.e. the NK model (Kauffman 1993; Levinthal 1997) as a way to conceptualize how problems of varying complexity pose qualitatively different challenges to a solver.

One difficulty with work built on the NK model is that it can seldom overcome the initial assumptions derived from the canonical model. With notable exceptions (Winter et al. 2007), in most NK variants, experiential search- local search and long jumps (randomly selecting a completely different solution) - are the only types of search that agents can engage in. Research on individual (as opposed to organizational) problem-solving suggests that human problem-solvers have far more sophisticated search strategies (Doll et al. 2012) and rely on a combination of cognitive (reflective) and experiential (reflexive) moves (Gavetti and Levinthal 2000).

However, if employed, the problem representation in the NK models is relatively distant from the empirical phenomenon (Winter et al. 2007). For example, in Gavetti (2005) agents are simply given a set of possible problem representations from which they choose -“a cognitive memory”-, while in Winter et al. (2007) the cognitive influence is “independent of experience” (Winter et al. 2007: p.404).

In this paper, I put forward a different model of adaptive search that allows solvers to make moves relying on problem representations that they themselves have shaped using their gradually accumulated expertise: “model-based” search. This approach is consistent with the idea that problem formulation is not a top-down one-shot process, but rather an adaptive, iterative process where different “models” of the problem are formulated as more information is gathered about the problem structure (Newell and Simon 1972).

To address this research agenda, first, I conducted a series of lab experiments where solvers were asked to play successively two games with payoffs described by the NK and a hierarchical function (Watson and Pollack 1999). Similar to Billinger et al. (2013), the task was to solve a puzzle by generating combinations of two different tokens. Each participant had a limited number of attempts and received immediate and accurate feedback (i.e. their score) after each attempt. The experiment provided a realistic benchmark for the agent-based model parametrization as well as insight into real (as opposed to simulated) problem solving behaviour. Additionally, the experimental setup was constructed to be a one-to-one match with the agent-based model, facilitating the comparison between the two data-sets. Using a family of the same hierarchical functions as a problem space, I use simulations to explore the efficiency of problem solving processes in different settings (i.e. problem sizes and complexities).

The main contribution of this paper is twofold: the first contribution is the identification of a distinct type of search behaviour that solvers rely on as they attempt to navigate the problem

space. I find evidence that model-based distant search is an important strategy that changes the adaptive search process. The agent-based model derived from the empirical study reproduces well-known results in the innovation management literature and informs the second contribution. Main results from the simulation suggest that “fail often, fail early” is contingent on problem complexity, in line with the problem solving standpoint, (Nickerson and Zenger 2004; Nickerson et al. 2012; Felin and Zenger 2014). Furthermore, for simple and small problems relying on model-based distant search considerably lengthens the problem solving process, while for problems that are more complex model-based distant search is actually optimal.

The remainder of this paper is organized into four sections: following an overview of previous work in Section 2, in Section 3.1 I describe the experiment and the results from the lab, while Section 3.2. is centred around the model and a number of virtual experiments. Finally, in Section 4, I discuss these results focusing on theoretical and managerial implications.

2. Theory

A relatively large body of theoretical work on problem solving in strategic organization research e.g. Levinthal (1997), Siggelkow and Levinthal (2003), was built around the powerful search metaphor put forward by Newell and Simon (1972), in particular around the NK model (Kauffman 1993) and March’s exploration-exploitation model (March 1991). Extensions and variants of these models have been used subsequently to study the dynamics of problem solving.

Thus there appears to be two potential avenues for further development in what concerns the applicability of the model for problem solving: 1) the landscape (i.e. the problem space) and 2) the way the search is conducted (Rivkin and Siggelkow 2007).

One of the limitations of the NK model is that it only allows for single-level interdependence, while most innovation problems can be thought of as hierarchical problems (Yu et al. 2009) with interdependencies *between* levels, as well. A problem space with interdependencies between levels is still rugged, but has discernible patterns (i.e. the ruggedness is not random). The random ruggedness of the NK landscape (McKelvey et al. 2013) restricts the use of the model further in what concerns implementing more cognitively plausible assumptions regarding search. As mentioned earlier, the problem-solving literature largely focuses on experiential search (cf. Gavetti 2005; Gavetti et al. 2007; Billinger et al. 2013). For example, in Levinthal’s seminal paper (1997), search is “non-greedy local” (imitate the first nearby alternative that is superior) or “non-greedy random long-jump” (imitate the first randomly selected distant alternative that is superior).

Still, trial and error learning is but one of two fundamental cognitive mechanisms that humans engage in. Individuals rely exclusively on trial and error only in very uncertain environments (Doll et al. 2012) and, as they explore, they actually do form “world models” and their subsequent actions are driven by model-based search (Graybiel 2008). Model-based search is thus a mode of search in which agents employ previously acquired information to form so called cognitive maps of the search space and rely on these maps to determine and evaluate future actions. Consequently, model-based search is both *influenced by* and *influences* the way agents perceive their environment.

Imagine the fictitious and stylized task of designing a smart phone. A very simple problem decomposition would identify two modules: screen and battery. These two modules are naturally interdependent: i.e. screen size affects battery life, hence the bigger the screen, the larger the electric output (amps) from the battery should be. This would ostensibly come with a cost in terms of battery size. For a non-trivial problem, an engineer cannot really optimize both components in isolation and *then* match them up. The best possible battery is, we assume, also the smallest possible one and this is not enough to power the best possible screen which should be of high-fidelity and of at least medium size. The NK model allows for this intuition to be formed. Furthermore, it tells us that with more components (N) and more interdependent components (K) it becomes harder for an engineer to figure out the right combination. However, the NK does not tell us the whole story. Going one level lower, batteries are also made up of several interconnected components and there are in fact multiple ways of designing a battery, all optimal, by some criterion. Similarly, there are multiple ways of designing a screen, all perfectly valid and (as the wide variety of smartphones on the market clearly illustrates) all appealing in some way to a different consumer need. Again, following the NK intuition, one engineer would start randomly with some combination of components and then tweak them one by one, until no more improvements can be made. For large Ns and Ks the task becomes exponentially long. Fortunately, there is an alternative way of solving such problems. Like a chess-master who is no longer concerned with individual positions on the chess board, but rather with configurations of pieces, after a number of failed attempts, our engineer *learns* how viable batteries and screens can be designed. He now has a number of configurations he can choose from trying to simply match the different viable configurations for screens and hardware. There is no need to exhaustively evaluate all possible combinations of screws and bolts that make up a phone, but rather focus his attention on what he has filtered through prior experience as ‘good possible combinations’. This search heuristic bears resemblance to descriptions of new product development teams in Takeuchi and Nonaka (1986). It is an adaptive advantage for solvers to assume some underlying structure about the task environment and attribute a set of

interdependencies between the building blocks. As the problem becomes more complex (there is a higher degree of interdependence), these heuristics often fall short and re-decomposition is required (Luo and Knoblich 2007).

Searching far away from the lamp-post

Despite Levitt and March (1988: p.516) initial formulation of organizational learning as “encoding inferences from history into routines that guide behaviour”, a great deal of the early literature focuses on local uninformed search, as highlighted in a recent review by Mckelvey et al. (2013). Relying on the NK model, previous research often assimilated exploration with random distant search. Distant search was therefore seen as exclusively random (e.g. “long random jumps” in Levinthal (1997)). It held the somewhat uncertain promise of introducing enough variance into the genetic mark-up of the organization, such that it may lead the search away from a suboptimal peak and towards the global maximum. In this perspective, the larger the distance, the more “exploratory” the agent.

However, later developments, in particular with respect to the cognitive micro foundations of organizational change have brought back the spot-light on model-based distant search (Gavetti and Levinthal 2000; Tripsas and Gavetti 2000; Gavetti 2005). This work is informed by cognitive science research which distinguishes between habit and goal directed learning (Graybiel 2008). Psychology and related disciplines have a long tradition of postulating two qualitatively different ways of processing information and decision making: e.g. Kahneman’s system 1 and system 2 (2011). The theories and concepts do not necessarily overlap, but there is consensus over a substantial difference to be noticed between reflective and reflexive decision making (Dolan and Dayan 2013). The former takes into account possible prospective courses of actions and consequent actions (i.e. model-based search) and the latter describes a type of search where the assessment of future actions is based exclusively on immediate experience (Doll et al. 2012; Dolan and Dayan 2013). Management scholars have approached model-based distant search mainly at an organizational level (Gavetti and Rivkin 2007; Chen 2008) focusing on the development or emergence (Cohen and Bacdayan 1994) of capabilities as a result of upper-management strategic decision-making. For an extensive review see Eggers and Kaplan (2013). Notably, a number of authors have looked into strategies that could mitigate the pitfalls of relying exclusively on local search and introduced more sophisticated behavioural assumptions. For instance, in an influential paper, Gavetti and Levinthal (2000) endow their agents with partial representations of the landscape that allow a different interpretation of decision-making. As such, while still relying on local search (only looking to nearby neighbours), under his own problem representation, the agent is capable of cognitive search as he searches the landscape guided by

imperfect cognitive maps. Furthermore and more significantly for the present argument, the agent is capable of changing its problem representation, but only when his performance is decreasing and only by learning from peers or random change (Gavetti and Levinthal 2000).

These results are complemented by Winter et al.(2007) and Baumann (2010) who suggest that there is a desirable middle level of persistence in the face of failure. Using different assumptions and implementations, these models have in common an attempt to describe the search behaviour of agents who are not dominated in their actions by immediate feedback. This allows their models to distance themselves from the “ant-like” (cf. Winter et al. 2007) model of search. However, while search behaviour is sometimes informed by cognitive problem representations, these representations are exogenous to the problem solving process.

Modularity and problem representation

Model-based distant search assumes that agents are informed in their moves by a representation of the landscape they are navigating through. This problem representation is thought to determine what and how information is perceived and what further structures are identified (Zhang 1997). A stream of research in innovation management focuses on the structure a problem should have to make problem solving efficient. As such, the literature largely agrees on the fact that a modular problem is desirable e.g. (Benkler 2002; Lakhani et al. 2013) especially in “collaborative innovation”(Baldwin and von Hippel 2011), with modularity being acknowledged as a tool to manage complexity. Further inquiries take a more dynamic perspective and suggest that *modularization* is indeed beneficial, but caution specifically against excessive modularization which seems to have more pitfalls (Ethiraj and Levinthal 2004); subsequent research, however, qualifies the relationship between the degree of modularization and the given problem (Brusoni, Marengo et al. 2007).

However, the relationship between problem representation and the knowledge an agent acquires is not unidirectional. Perfect decomposability of a problem can only occur insofar as the solver has a coherent representation of the underlying structure (Greeno 1977), which cannot happen in early stages of solving so called ill-structured problems. This knowledge is formulated by the solver by gaining familiarity with a problem (Hayes and Simon 1975) and thus problem solving becomes a mix of general search heuristics and task-specific information (problem structure refinement) acquired through trial and error. Problem representations do not just come *ex nihilo*, but rather by gradually accruing experience and increased ability to organize information into larger chunks (Ericsson et al. 2006).

Therefore, it would seem that at least when it comes to ill-defined problems formulating a representation does not *preclude* problem solving, but rather it is part of an iterative, interactive process (Hayes and Simon 1974). The circular relationship between the two is emphasized by Klein (1998) and more recently, in the collective problem solving context by Sieg et al.(2010) who note that the initial problem representation can in fact be refined or even profoundly changed.

Refining is the process by which a solver changes his problem representation. Therefore, in an information processing perspective, so-called “A-HA” moments are simply the discovery of a *more effective* problem representation (Kaplan and Simon 1990; Gigerenzer and Edwards 2003; Hodgkinson et al. 2008). Returning to the landscape metaphor, by resorting to these mappings, landscapes are reduced, making the search more effective. As illustrated by Kaplan and Simon (1990) nobody searches completely at random for their car keys but resorts to a heuristic to constrain the search space, for example: “think of the last time you had them” (Kaplan and Simon 1990). This heuristic can later on be refined, using information gathered in the search process. This is in line with recent evidence on how human problem solving actually occurs – that is, a mix between experiential and model-based trials (Dolan and Dayan 2013).

Acquiring extensive expertise can be both beneficial (Ward 1994) in situations where the focal agent is able to make use of maps and find a solution faster and detrimental (Wiley 1998), since the same expertise can promote fixation and trap the solver in a local optimum.

By focusing on a third type of search behaviour: model-based distant search, I set out in this paper to understand how changing problem representations restructure the search space and how this in turn affects the problem solving process.

3. Research methods

Recent work dealing with search on rugged landscapes cautions against using modelling and simulation in isolation from empirical evidence. In particular, Mason and Watts (2012) and Billinger et al. (2013) show that traditional assumptions regarding solver behaviour are not sufficient to capture the dynamics involved in problem solving. For example, Billinger et al. (2013) find that solvers are prone to over-exploration and break off the local search (exploitation) too early, while Mason and Watts (2012) claim that current models lack heterogeneity when it comes to search strategies and behaviours. To mitigate the risk of embedding unsupported assumptions into the agent based model, a laboratory experiment was conducted prior to conducting the simulations.

3.1. Experimental data

3.1.1. Experimental design

The lab experiment is similar to the set-up used in Billinger et al. (2013). The data-collection was designed to be a one-to-one match with the simulation model so that results are easily interpreted in the modelling framework. 200 people were recruited from a pool of participants and asked to come to the lab for one-hour sessions (approx. 40 min of actual playing time). Although the participants were primarily students at various faculties at Aarhus University, the particular research question as well as recent work on using such samples (Exadaktylos et al. 2013) suggest that this does not raise serious concerns with respect to the out-of-sample validity of the experiment.

Participants played two combinatorial games in succession (a hierarchical and an NK¹ problem, corresponding to the two different functions). They were instructed to solve a puzzle by repeatedly submitting different combinations of two different tokens (Figure 1). A participant could toggle on or off any of eight different slots to construct a given solution. At the outset of the game, participants were told that they had 25 attempts to identify the combination of tokens that would yield the best score. There was no upper or lower limit to how many changes could be made from one submission to the next. In order to avoid additional variation in the search behaviours due to different expectation levels, participants were informed before starting the hierarchical task of the highest possible score. To insure that the order of the games did not have any significant effects on the search strategies, the order of the two games was randomized. Finally, the participants were informed that their reward would depend on their performance².

¹ I decided to use the NK problem as well as an additional validity check for the experimental design

² Ranging from 6\$-24\$

Game 1

[See instructions again](#)

Solutions

The screenshot shows a user interface for a game. At the top is a search bar with the placeholder "refresh submitted solutions". Below it is a table with three columns: "Solution", "Score", and "Number". A single row is shown with the solution "□△△△△△□□", a score of 10, and the number 1. Below the table is a navigation bar with buttons "<<", "<", "1", ">", and ">>". Underneath the navigation bar is a row of eight triangular icons, each containing an alien head. A "Create" button is located at the bottom left of this row.

Solution	Score	Number
□△△△△△□□	10	1

<< < 1 > >>

Create

Figure 1. Screen-shot of the game

The solvers filled in a short survey (testing their understanding of the game instructions) and were afterwards redirected to the browser-games. They each received immediate and accurate scores after each submitted attempt. Players had access to their own previous attempts. Note that both puzzles had each 256 possible combinations with one (NK function), respectively two (hierarchical function) optimal solutions. The data from the experiments is both qualitative and quantitative in its nature (Boero and Squazzoni 2005), in that the first allows for calibrating the rules of behaviour for the agents, while the second is used for direct parametrization of the variables: such as the number of agents, the size of the problem and the distributions of strategies.

3.1.2. Landscape

Most problems are nearly decomposable (Simon 1962). In particular, innovation problems are closer to a hierarchical structure (Marengo 2014). In the smartphone example, the optimization of any given module needs to take into account all other modules, but this also “ripples down” the hierarchy. A choice for a particular type of screen is influenced by the choice of battery and vice versa, but it also influences several other interconnected components which make up the screen (e.g. digitizer and LCD panel) and so on. Accordingly, I use in this paper a different

implementation of the rugged landscape metaphor, relying on a class of problems known as hierarchical problems.

In the traditional NK landscape, the fitness score is determined by averaging over a number of random draws from a uniform distribution. Unlike the NK function, the HXOR function¹ is defined by a transformation function (1) and a fitness function (2). The transformation function insures the “validity” of the strings of symbols, while the fitness function assigns the corresponding score. The transformation function is applied recursively to transform the chosen string into one single symbol hereby reducing the effective dimensionality of the problem. The function is given by:

$$h(a, b) = \begin{cases} 1, & \text{if } h(a) = h(\bar{b}) \text{ and } a = 1 \\ 0, & \text{if } h(a) = h(\bar{b}) \text{ and } a = 0 \\ \text{null}, & \text{otherwise} \end{cases} \quad (1)$$

where \bar{X} is the bitwise negation of a string X (e.g. $\overline{01} = 10$).

Based on (1) the fitness function is defined as:

$$f(a) = \begin{cases} 1, & \text{if } a = 1, \text{ or } a = 0 \\ 0, & \text{otherwise} \end{cases} \quad (2)$$

Every *correctly* identified pair of attributes is transformed and passed on to the next level and its fitness is added to the overall fitness. Thus a pair 01 contributes to the level of fitness in the lower level with $1+1=2$ and is also transformed into a “higher order” 0 (according to the transformation function h) and passed into the next level where it has a contribution of 2 (the sum of its components).

For example, the string 00010010 will yield a fitness of 12, according to the following calculations:

Transformation	Fitness score
0 0 0 1 0 0 1 0 - 0 - 1 - - - - - - - -	$= 8+$ $+2+2+$ $+Null +$ $+Null = 12$

¹ The same function that was used in the laboratory experiment and the simulations

With the basic formula, HXOR is a bimodal landscape, with the two optimal solutions either 0110 1001 or 1001 0110. Forming chunks allows for large distances to be covered quickly – e.g. going from 0110 0110 (fitness 24) to 0110 1001 (fitness 32) would be at the lowest hierarchical level (e.g. without chunking) very hard, but working with chunks of size 4 is quicker – hold one fixed and go through the possible combinations of the other.

The more general formula (3 and 4) for a *biased HXOR* allows the model to introduce landscapes of varying complexity:

$$h(a, b) = \begin{cases} 1, & \text{if } h(a) = h(\bar{b}) \text{ and } a = 1 \\ 0, & \text{if } h(a) = h(\bar{b}) \text{ and } a = 0 \\ \text{null, otherwise} \end{cases} \quad (3)$$

Based on (3) the fitness function for a variable a is defined as:

$$f(a) = \begin{cases} 1, & \text{if } a = 1, \\ \text{bias, if } a = 0 \\ 0, & \text{otherwise} \end{cases} \quad (4)$$

As such, for $bias = 0$, the landscape becomes unimodal and for $bias=1$, the landscape is the one already detailed in the earlier paragraphs and the one played in the experimental sessions. For bias values between 0 and 1, the complexity of the problem is increased since payoffs for 1s are larger than payoffs for 0s, but the optimal solution requires a combination of both.

3.1.3. Experimental results

It is a daunting task to make inferences about what happens “between the ears” of the participants and evidently the data collected does not allow for a one to one mapping between recorded behaviour and the variables of interest (i.e. whether the players were engaging in “local search” or “distant search” had to be inferred from the sequence of submitted solutions). To ensure additional validity for the coding of the search behaviours, I also collected anecdotal evidence consisting of 8 semi-structured interviews. These were collected after players finished their sessions and focused specifically on the strategies the players used in their attempt to find the optimal solution (i.e. one of the questions they answered was “what kind of strategies did you engage in during the game”), but also on their reference point (i.e. “how did you construct a new solution”).

I initially constructed a simple operationalization of local search as moves of exactly 1 bit from the reference point, and distant search as everything else. E.g. a succession from {0, 0, 0, 0, 0, 0,

$\{0, 0\}$ to $\{0, 0, 0, 0, 0, 0, 0, 1\}$ or to any configuration involving exactly one attribute set to “1” and all other attributes remaining “0” is considered a “local search” move, while all others are considered “distant moves”. The reference point was the individual’s best performance so far (Billinger et al. 2013) due to the human inherent predisposition to compare future states with the status quo (Kahneman and Tversky 1979). The interviews revealed that, indeed, solvers were paying attention to the solution which gave them the highest score so far.

However, further data analysis, as well as the interviews, reveal that this coding is not as clear-cut as it might seem at a first glance. A number of respondents indicated that although they did judge their *performance* based their best score so far (i.e. “I always benchmarked my score against the solution where I did best, to see if I was improving”), they generated new solutions with their latest submission in mind (“I paid attention to what I did last, since that gave me information about how my moves would affect my score”). Therefore, the coding was adjusted to allow for the reference point to be either the solution giving the player’s best score so far or the latest solution. As expected, the number of distant search moves decreased (See Table 1)¹. The distribution of strategies is similar to the one found by Billinger et al. (2013), where they report values ranging from 33 to 40 percent for local search, contingent on the problem complexity.

Hierarchical			NK		
Reference	Reference1	Reference2		Reference1	Reference2
Local Search	25.10%	35.20%	Local Search	23.50%	35.50%
Distant search	74.90%	64.80%	Distant search	76.50%	64.50%

Reference1 refers to the coding which involves the best solution so far, while *Reference2* refers to the coding which involves both the best solution so far and the latest solution

Table 1. Distributions of search behaviour

Turning to what determines the two search strategies: the solvers engage in local search progressively as they advance in the game — a kind of exploitation. The probability of engaging in local search is increased if solvers receive positive feedback with respect to the reference point or with respect to their previous move. That is, as long as they are performing above their reference level, the further away they are in the game, solvers will be risk averse and engage only in minor tweaks of their solutions. There is also robust evidence pointing at a curvilinear relationship with time: i.e. the solvers were less likely to engage in local search in the last couple of moves. Additionally, the more successive unsuccessful attempts a player made, the more likely he was to break off the local search.

¹ The issue of reference points will be addressed in a future study, but this coding was preferred since it is the most conservative with respect to distant search moves, thus allowing for the smallest number of false positives.

Determinants of local search						
Independent variables	Model 1	Model 2	Model 3	Model 4	Model 5	Model 6
Feedback(ref=0)	-0.3988*** (0.08291)	-0.6594*** (0.08687)	-0.3254** (0.0967)	-0.2793*** (0.09766)	-0.3008** (0.09856)	
Immediate_feedback (ref=0)		-0.8691*** (0.06461)	-0.8341*** (0.06510)	-0.8795*** (0.06610)	-0.8844*** (0.0638)	-0.8384* (0.06444)
Number of successive unsuccessful trials			-0.06224*** (0.008234)	-0.1141*** (0.009909)	-0.1110*** (0.01001)	-0.1232* (0.09865)
Trial				0.05572*** (0.005618)	0.1758*** (0.02124)	0.1740*** (0.02114)
Trial*Trial					- 0.00454*** (0.000771)	0.00445** (0.00076)
Constant	-0.2033** (0.07593)	0.5154*** (0.09388)	0.4863*** (0.09401)	-0.02327 (0.1071)	-0.6024*** (0.1462)	0.8249*** (0.1266)
-2 Log Likelihood	6091.56	5906.76	5845.98	5745	5710	20487.86
N	4646	4646	4646	4646	4646	4646

*p<0.05, **p<0.005, ***p<0.001

Models 1-5 are logistic regressions. The last model fitted is marginal logistic regression model:

$$\log \left\{ \frac{\Pr(LS_{ij}=1)}{\Pr(LS_{ij}=0)} \right\} = \beta_1 + \beta_2 immediate_feedback_{ij} + \beta_3 trial_{ij} + \beta_4 unssucessful_{ij} + \beta_5 trial_{ij} * trial_{ij} + e_{ij}$$
, where $\beta_1 - \beta_5$ are the fixed-coefficients (intercept β_1) and e_{ij} is the error for observation j of subject i . 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. To account for potential autocorrelation in the observations, we included a multiplicative over-dispersion parameter ($\phi \neq 1$). Since -2 Res Log Pseudo-Likelihood is not comparable even among nested models different modelling assumptions with respect to within subjects variance were tested and reported results are robust. *Immediate_feedback* and *feedback* are binary variables which capture whether a participant has received negative (coded as 0) or positive (1) feedback with respect to either his latest (immediate feedback) or his best (feedback) submission. The number of successive unsuccessful trials is a continuous variable that captures the number of trials since the latest improve in performance, with respect to the best achieved score so far.

Table 2. Determinants of local search

Mirroring the results for local search determinants, players seem more likely to engage in distant search, early on in the game. Positive feedback results in a lower probability of solver engaging in distant search, and with a long streak of unsuccessful moves, solvers will be more likely to

break off the search and engage in distant search. This seems to be the “exploratory” behaviour expected.

3.1.4. Model-based distant search

But are all distant-search moves the same and are they random? While the experimental setup cannot decisively answer this question, the qualitative interviews as well as the coding show that solvers engage in a more systematic exploration of the search space. Specifically in later stages, solvers do seem to make (on average) less distant moves that violate the underlying problem structure, suggestive of the fact that in later stages they have a better representation of the problem interdependence structure (See Table 3)¹. Such moves, without being “wrong” as such, yield little useful information to a player, since the interdependence structure confounds the interpretation of feedback. Note that this approximate understanding does involve “chunking”, but the patterns in the data suggest that this is not a result of players attempting to manage the complexity of a problem, but rather a result of them being able to form an approximate understanding of the underlying function. Successful players seem to be better able to use the feedback they receive in earlier stages as confirmed by a Cox proportional hazards regression (see Figure 2, which shows that violating the underlying problem structure decreases the probability of finding the optimal solution during the 25 trials).

Evolution of interdependence violation over time		
Independent variables	Model 1	Model 2
Trial	-0.00884** (0.003147)	-0.00884** (0.002163)
Constant	-0.5210*** (0.04625)	-0.5210*** (0.3179)
-2 Log Likelihood	7669.04	
-2 Res Log Pseudo-Likelihood		12122.38
N	4646	4646

Model 1 is a logistic regression. The last model fitted is a hierarchical linear model which allowed for respondent specific variance, but the -2 Res Log Pseudo-Likelihood is not comparable even among nested models. However, different modelling assumptions with respect to within subjects variance were tested and reported results are robust

Table 3. Evolution of interdependence violation over time

¹ For example, a move from 1001 0000 to 1011 0001 would be considered such a violation since the correct problem decomposition involves “chunking” the first and last four bits. Similarly, going one level lower in the problem decomposition a transition from 1001 1001 to 1001 1100 is a violation of the underlying structure since it involves crossing the boundaries of lower level chunks

The probability of survival (not finding the optimal solution) for interdependence violation

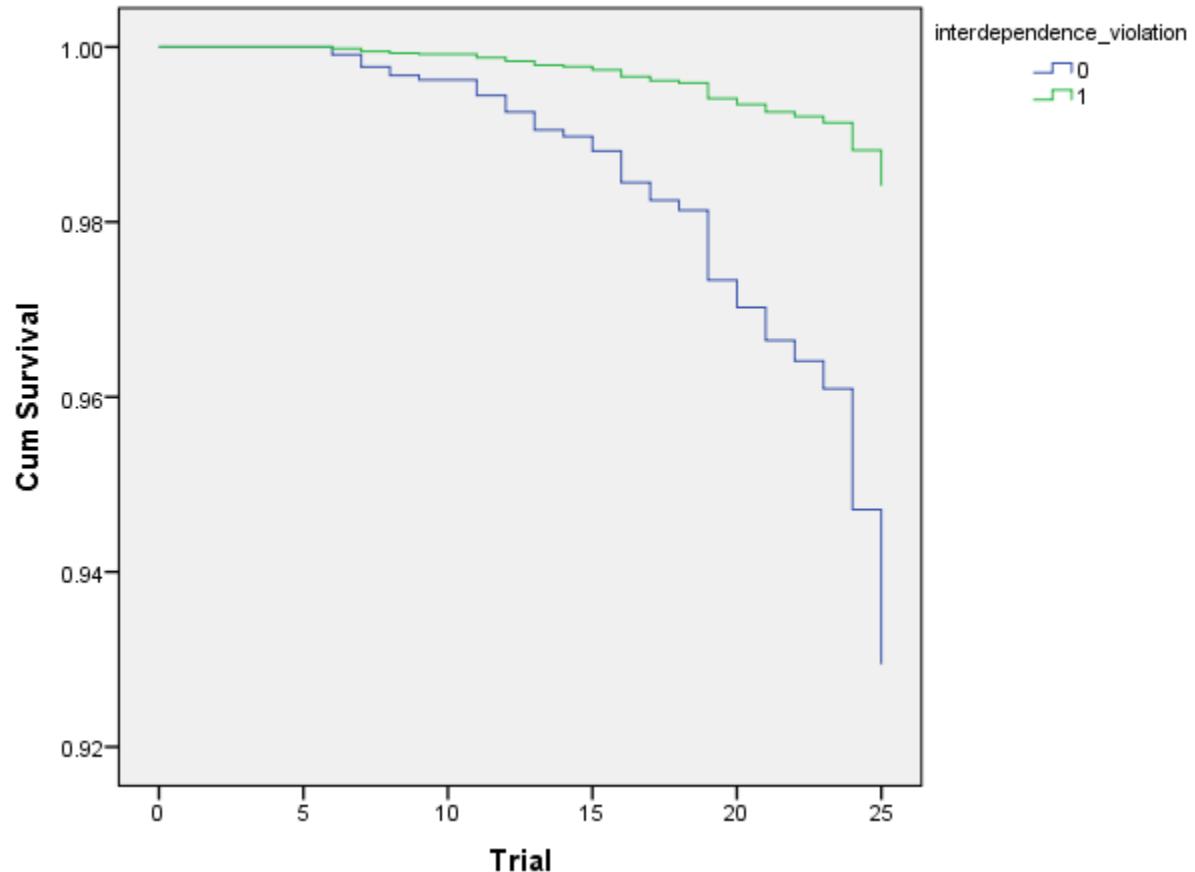


Figure 2. Hazard plot for Cox's regression comparing the ratio of reaching the endpoint for solvers who make more moves consistent with the underlying structure (*interdependenceViolation*=0) vs solvers who make less such moves (*interdependenceViolation*=1). The cut-off point was 11.

Indeed, in the interviews, while some respondents showed an obvious preference for local search (“I kept everything but one symbol constant because that was the only way to understand what impacted what”) others were able to detect and exploit the symmetries of the hierarchical landscape (“I started by playing some different combinations [...] and then [...] by looking how the score changed, I tried to understand the pattern. I was [making moves] trying to see if it was symmetry or adjacency that was [important] ”).

Consequently, I continue by making a further distinction in what concerns distant search moves: model-based search and random search. Consistent with the evidence presented earlier, moves were coded as “model-based search” if they involved distances of exactly 2 in the Hamming

space¹ and if they didn't involve simultaneous changes in the first order problem decomposition. This coding means that at any given time step, only twelve out of 255 possible new combinations can be considered "model-based search". Thus, if players engaged exclusively in local and random search, the odds that a move coded as "model-based search" would be randomly chosen by a respondent are 12/255, yielding a ratio of 4.7% such "model-based" moves. In effect, the ratio is considerably higher. According to this coding the distributions of three search strategies are²: model-based search: 27.6%, local search: 37.6% and random search: 34.7%.

This differentiation allows for a further investigation of the two different types of distant search. Random distant search does not seem to be influenced by feedback alone (whether immediate or with respect to the reference point), but primarily by the number of successive unsuccessful trials: the longer the unsuccessful streak, the more likely will be that a player will engage in random search (Table 4). The number of unsuccessful trials also captures feedback, but the fact that no significant relationship exists between the likelihood of engaging in random search and these two variables suggests that players will primarily engage in random search only if more such failed attempts are made.

¹ The Hamming distance is a measure of the number of positions at which the values are different. E.g. a move from (0000) to (0011) corresponds to a Hamming distance of 2.

² Since, according to this coding scheme, a move can only be considered "model-based" after time=2, all reported results refer to the remaining 22 attempts.

Determinants random distant serch (RS)				
Independent variables	Model 1	Model 2	Model 3	Model 4
Feedback	0.3547*** (0.08320)	0.6048*** (0.08711)	0.1753 (0.09817)	
Immediate_feedback		0.8301*** (0.06572)	0.7804*** (0.06646)	0.7517 (0.06433)
Number of succesive unsuccesful trials			0.07665*** (0.008559)	0.08367** (0.007655)
Constant	0.2033*** (0.08320)	-0.4830*** (0.09430)	-0.4417*** (0.09453)	-0.3103*** (0.05928)
-2 Log Likelihood	5863.97	5701.28	5614.53	
-2 Res Log Pseudo-Likelihood				-19350.54
N	4646	4646	4646	4646

*p<0.05, **p<0.005, ***p<0.001

Models 1-3 logistic regressions. Model 4 is a marginal logistic regression model:

$\log \left\{ \frac{\Pr(RS_{ij}=1)}{\Pr(RS_{ij}=0)} \right\} = \beta_1 + \beta_2 immediate_feedback_{ij} + \beta_3 unssuccesful_{ij} + e_{ij}$, where where $\beta_1-\beta_3$ are the fixed-coefficients (intercept β_1) and e_{ij} is the error for observation j of subject i . 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. To account for potential autocorrelation in the observations, we included a multiplicative over-dispersion parameter ($\phi \neq 1$). Since -2 Res Log Pseudo-Likelihood is not comparable even among nested models different modelling assumptions with respect to within subjects variance were tested and reported results are robust. Immediate_feedback and feedback are binary variables which capture whether a participant has received negative (coded as 0) or positive (1) feedback with respect to either his latest (immediate feedback) or his best (feedback) submission. The number of successive unsuccessful trials is a continuous variable that captures the number of trials since the latest improve in performance, with respect to the best achieved score so far.

Table 4. Determinants of random search.

The same cannot be said of model-based search. The data does not show significant relations between model-based search and any of the variables of interest: feedback, immediate feedback or the number of successive unsuccessful trials. There is, however, similar to the case of local search, to be noted a significant ($p<0.001$) relationship between the type of search behaviour the player was using in his previous attempt. I thus find evidence of “strategic persistence” (Lant et al. 1992), pointing to the fact that, controlling for other factors, a player will be more likely to show some consistency in his search behaviour.

3.2. An agent based model of search

The agent-based model examines how search behaviours are influenced by embedding additional knowledge in the search heuristic and how that in turn affects the likelihood of finding an optimal or near optimal solution. The aim is to explore the interplay between different problem solving strategies, problem formulation and refinement and performance. Following the analysis of the experiments, I calibrate the agent-based model, using the relationships and distributions from the experimental data. Therefore, for the parametrization, I used both the descriptive results from the lab experiments as well as the determinants of each strategy to generate rules for agent behaviour.

By running simulations on different model parameterizations, several virtual experiments are conducted in the attempt to identify significant dynamics that can describe the problem solving processes. Specifically, a twofold dependent variable is measured: average performance and the time required to reach it in several distinct scenarios corresponding to variants of the hierarchical problem.

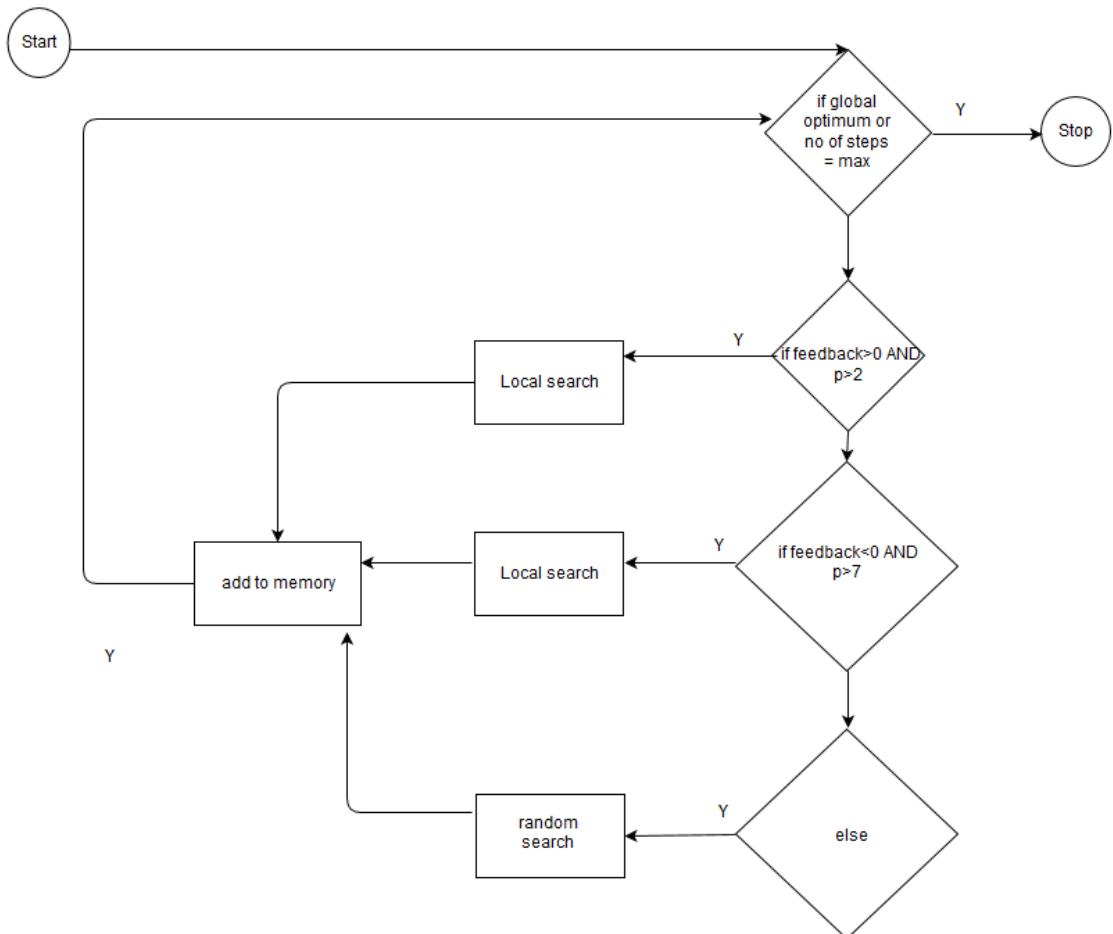
3.2.1 Problem representation and problem restructuring

One of the reasons behind this particular choice of rugged landscape is the fact that it allows for a simplified, albeit non-trivial implementation of problem representations as well as how they might be formed. As the upper limit of their memory is reached, agents attempt to detect regularities about the landscape, relying on their best attempts so far. As such, if an agent's best solution so far has "01" for the first two positions and all the other solutions he "remembers" have the values of either "01" or "10" for these loci, the two positions are "chunked". For subsequent model-based distant search moves the agent will only use these values as possible values.

If an agent starts with chunks of size 1, for an 8 bit HXOR of bias=1, all possible strings will yield the same payoff =8 and are as such completely un-informative with respect to feedback. Finding the right solution is rendered to sampling all 256 possible combinations. For chunks of size 2, however, the landscape has only 2^4 local optima, since the agent can optimize at this level of representation all strings of length 2 i.e. "10" and "01" yield better payoffs than "11" and "00". The task of finding the right solution is easier, albeit there are still 16 possible combinations that, at this level of representation are equally good. Finally, for chunks of size 4, the search is further constrained and the agent needs to navigate through only 2^2 local optima (Watson and Pollack 1999).

This type of problem representation accounts, for example, for competency traps (Siggelkow and Levinthal 2005) in organizational learning as well as a closer mapping between organizational

problems and model (Goldberg 2000). A hierarchical structure of the problem also allows for a more cyclical problem solving process to be modelled. As such, the computational agents start “tabula rasa”, but as they explore the landscape, they encode the feedback they receive and form cognitive maps of the space. These imperfect representations are used in subsequent searches and are refined.



Legend:



Figure 3: Flowchart for the model-based search model

3.2.2. Model-based search

Random distant search occurs in early stages (time<2) (before any patterns are detectable). *Persistence* is a variable that measures the number of consecutive moves that lead to scores lower than the current best performance an agent has made so far (i.e. if at time $t-1$ an agent makes a move that leads to decreasing performance, the variable gets incremented by 1. However, if at time t the same agent makes a move that marks an improvement of his best performance the persistence counter is set again to 0)¹. This parameter captures the variable “successive unsuccessful trials” and was initially set to 3 (the median value obtained from the laboratory experiments).

In later stages (time>2) only if the persistence threshold has been reached, the agent engages in random search. More specifically, if a previous move has generated positive feedback, the agent will be more likely to engage in local search. With a lower probability, if a previous move has generated negative feedback the agent will still engage in local search. Otherwise he engages in model-based distant search, the details of which are explained below.

In this model, the implementation of model-based distant relies on an evolutionary search heuristic developed by Iclanzan and Dumitrescu (2007). Technical details aside, the algorithm relies on a simple associative learning rule and is similar to Hayes and Simon’s (1974) description of problem solving where a solver first begins by creating a problem representation and generates a solution which he tries to improve. When a solver’s memory is full, he uses the information he has collected in the search process thus far to generate a different problem representation.

The agents start without any problem representation (i.e. looking at all possible combinations of bolts and screws, in this case 1s and 0s) so they use their first moves to acquire a basic “understanding” of the problem. After a number of random moves, agents engage either in local search or in model-based distant search. In their model-based search they attempt solutions that incorporate previously gained knowledge about optimal configurations of the components they have identified as interdependent². In their local and model-based distant search, the agents are satisficing using their own problem representation to calculate their score. That is, they attempt to find a solution that outperforms their best score so far, according to their own problem

¹ “Persistence” is similar to the “moderate obsession” described in Winter, S. G., G. Cattani, et al. (2007). “The value of moderate obsession: Insights from a new model of organizational search.” *Organization Science* **18**(3): 403-419. and it is calibrated based on the empirically derived “number of successive unsuccessful trials”

² If an agent has identified a correlation between the first and the second loci, those and only those configurations which have been deemed successful in the past are possible model-based moves.

representation¹. If the resulting score is an improvement, they add the respective solution to their memory. Agents receive real feedback after submitting a solution and this feedback is used to compute their best score so far.

As agents engage in different learning actions, they remember the solutions they encounter. The memory an agent has is a fixed ratio of the size of the current problem representation reflecting human capacity limitations within attention and working memory²(Collins and Koechlin 2012). Once an agent's memory is full, he will attempt to develop his own map of the search space: i.e. which components are interdependent with which.

3.2.3 Simulation results

The following experiments were conducted on various parameter configurations and the search strategies performance was compared both in terms of speed and performance. If not explicitly stated, the eight-bit version of the problem was used (i.e. the task, for the computational agents, similar to that of the lab participants was to find the optimal configuration of 1s and 0s for a problem with 8 binary attributes). For each reported result (unless otherwise specified) 2000 virtual experiments were conducted to account for the stochastic nature of the model and ensure that results do indeed reflect underlying dynamics of modelled processes. Analyses of variance suggest that 2000 runs were sufficient.

Drawbacks and virtues of local search

The first round of experiments attempted to benchmark the adaptive search described in earlier paragraphs against local search. For simplicity, I refer to this algorithm as *model-based distant search*, but I want to caution the reader that, as explained in the previous section, it describes a complex search algorithm that relies on model-based distant search moves, *alongside* local and random moves.

For simple problems (bias=0), local search outperforms in terms of efficiency – that is, for *simple* problems, local search operates faster, and reaches the global optimum, even for problems of increasing size³ (e.g. for N=16: Speed average local search = 56.52; std. dev= 7.12; Speed average model-based = 260; std. dev= 0; sig at p<0.001).

¹ I used this assumption to approximate the fact that, presumably, solvers in the lab did not submit solutions they knew would be poorer than their best performance so far.

² For different parametrizations for memory size, except the limit case of size of memory=0, the results detailed below are robust.

³ In order to make a meaningful comparison between model-based and local search, speed here denotes the number of "steps" an agent makes, not the number of function evaluations.

However, as the ruggedness of the landscape increases, much like in a traditional NK scenario, relying on local search traps the solver in sub-optimal solutions. For increasing size and complexity, an exclusively local strategy (local search) quickly gets stuck in local optimum.

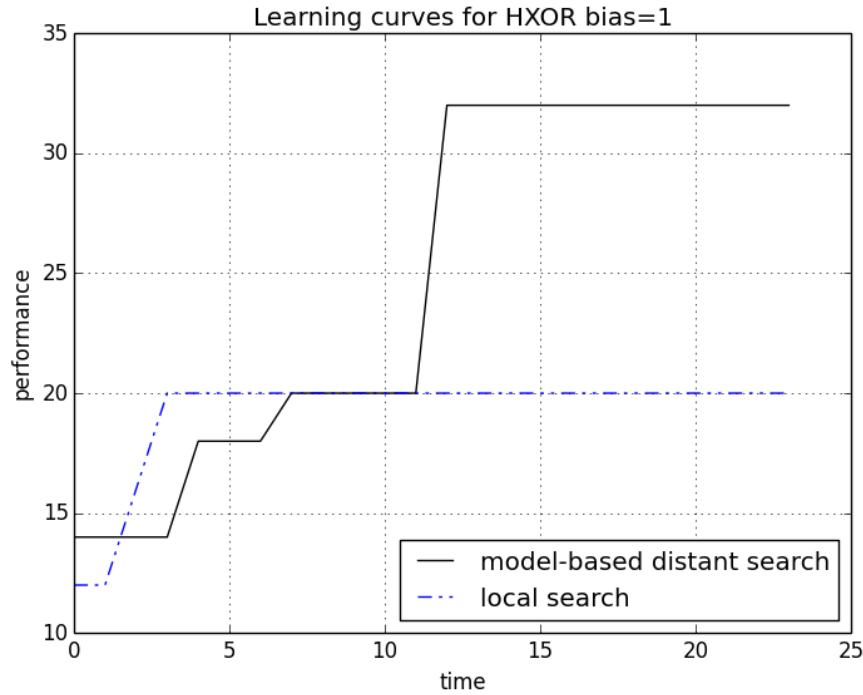
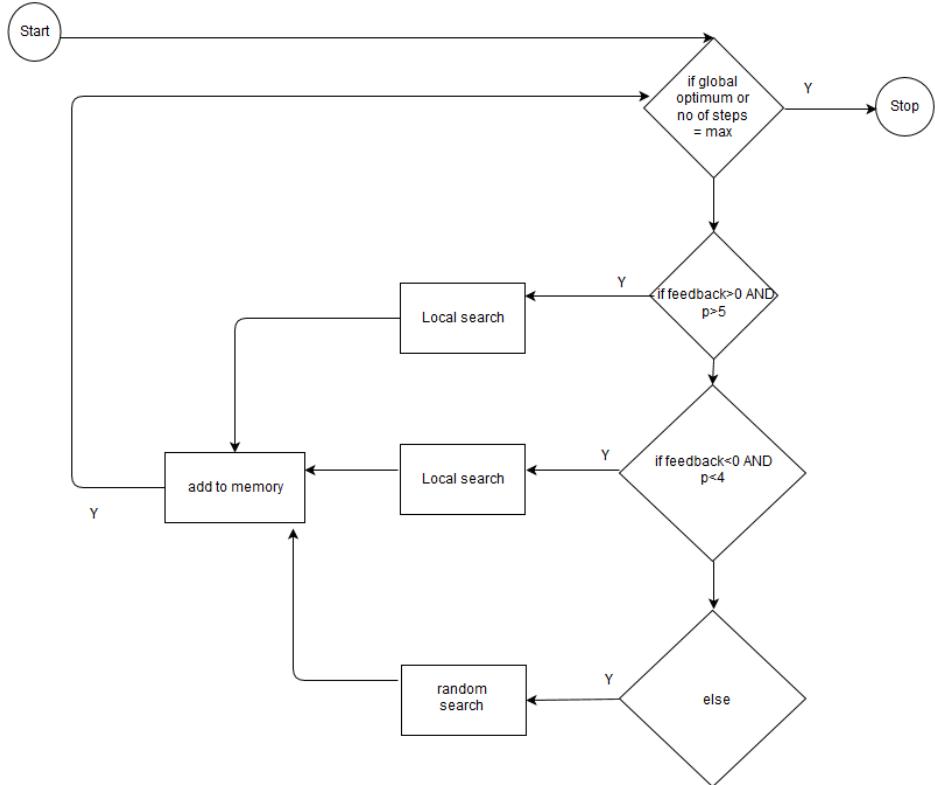


Figure 4. Illustrative run (N=8, HXOR, bias=1): Local search can outperform model-based distant search on short time intervals (vertical axis= performance, horizontal axis = time steps)

As illustrated in Figure 4, in the short run, even for more complex problems ($\text{bias}>0$) local search quickly improves performance and can outperform model-based distant search. This is because it is relatively easy in hierarchical problems to identify local optima, but without a proper problem representation, the agent is unable to escape those local optima.

But wouldn't a mix of local and random search be enough to account for the results seen in the lab sessions? Virtual experiments suggest the answer is “no”. I have parametrized “a naïve” agent based model (Figure 5) such that local search would occur with increased likelihood if positive feedback had been received, while random jumps would occur if local search steps had not been made.



Legend:

process	decision
---------	----------

Figure 5: Flowchart for the local and random search model

A parameter sweep allowed for the identification of the parameters such that the distribution of search behaviours matches the empirical data. Still, the model's average performance is significantly lower than the one observed in the lab (Average lab performance = 26.86; std. dev.=4.59; Average simulation performance = 24.24; std. dev.=5.79, sig. at $p<0.001$). This suggests that human players are more efficient when it comes to their distant search moves. Indeed, running simulations with the full model yields an average performance that, while lower, is not significantly different from the empirical results (Simulation performance = 25.96; std. dev=5.318, $p>0.05$).

Furthermore, model-based distant search out-performs the naïve local search and random algorithm across different levels of problem complexity (Table 5).

Bias	0		0.1		0.3		0.7		0.8	
	MB	LS/R	MB	LS/R	MB	R	MB	R	MB	LS/R
Average standard deviation N=1000	32** *	28.26 6	31.9702** * 6.506	28.028 9	28.4* * 6.6002		22.1** 27.5		23.2** 21.1	22.4 6
	0	4	0.66615	6	6.35	6.87	5.46	5.14	5.2	4.99

Performance comparison across different levels of problem complexity between model-based search (MB) and the naïve algorithm: local search with long jumps (LS/R)

Table 5. Comparison between the two parametrized models of search

As such, on a given problem of any level of complexity (i.e. a non-trivial hierarchical problem, bias>0), applying a simple heuristic, that ignores the structure of the problem yields a search landscape which has many local optima. Allowing for random distant search does improve the performance, but the path towards the solution becomes very long, because, without adequate decomposition, the search involves evaluating numerous large lower order chunks. Additionally, the additional variability in the search behaviour hinders the performance of the naïve search even for unimodal landscapes (Table 5, bias=0).

Persistence and problem representation

To explore previous results further, a batch of simulations was conducted on different levels of “persistence” the agents might have. The answer seems to be quite counter-intuitive: agents who are *less* persistent perform better and faster (in terms of the number of evaluations).

There is no significant difference between the groups with respect to the degree of chunking. When an agent gets stuck on a suboptimal peak it is not necessarily as a consequence of operating with high dimensional chunks because (similar to Bauman and Siggelkow (2013)’s chunky search) the agents do not start with a coarse problem representation, but gradually learn it. Rather, prior successful experience is trapping agents in suboptimal solutions. A computational agent might spuriously correlate the first and the seventh loci of the problem. This in turn will restrict the space of possible model-based moves. Since this correlation is not consistent with the accurate problem decomposition, the feedback is not informative (leading to more potential spurious correlations) and the likelihood of being stuck in a sub-optimal solution increases. Computational agents who are less persistent have an increased likelihood to build a “database” with sufficiently diverse information that allows them to find the accurate problem representation and the global optimum.

These simulations highlight a connection between two previously reported but uncorrelated results: problem representation and persistence (e.g. “moderate obsession” (Winter, Cattani et al. 2007)). For complex problems ($\text{bias} > 0$), less persistent agents (i.e. agents who switch to random search after a lower number of moves that do not outperform their best result thus far), acquire more diverse information and are consequently less likely to form incorrect high order representations of the problem. It is then not that chunking in itself is a good/bad predictor of performance, but rather that performance is driven by developing an accurate problem decomposition.

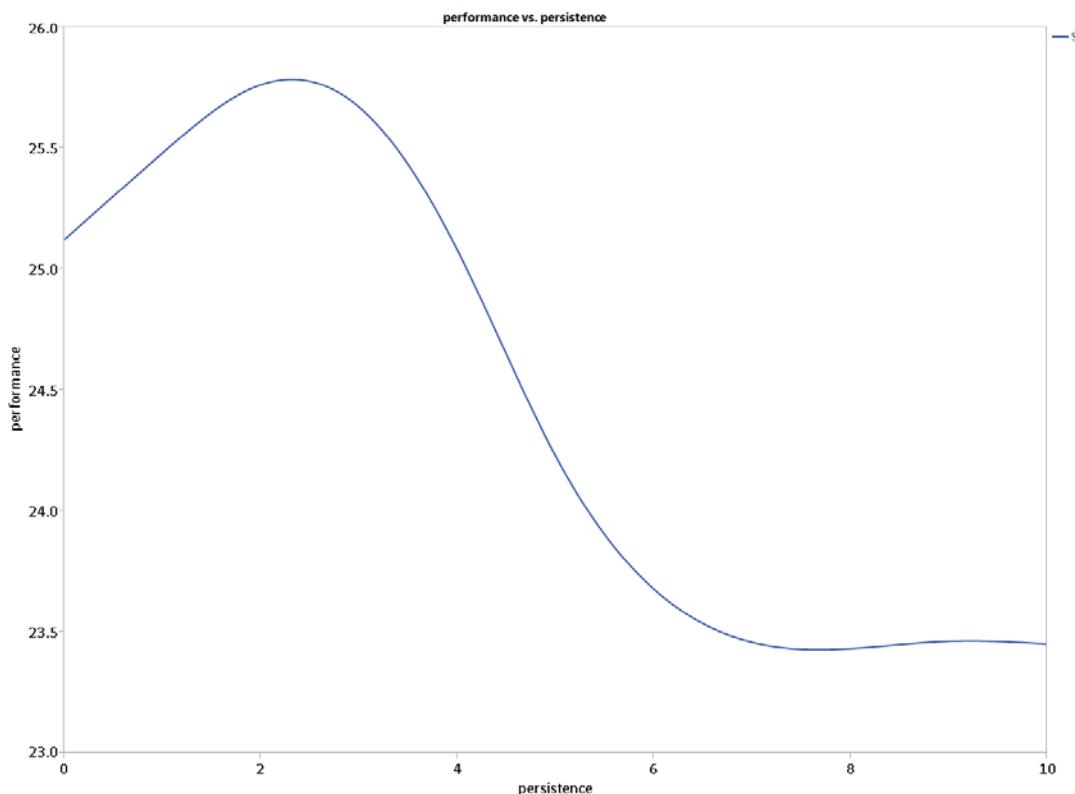


Figure 6: Average performance at different levels of persistence

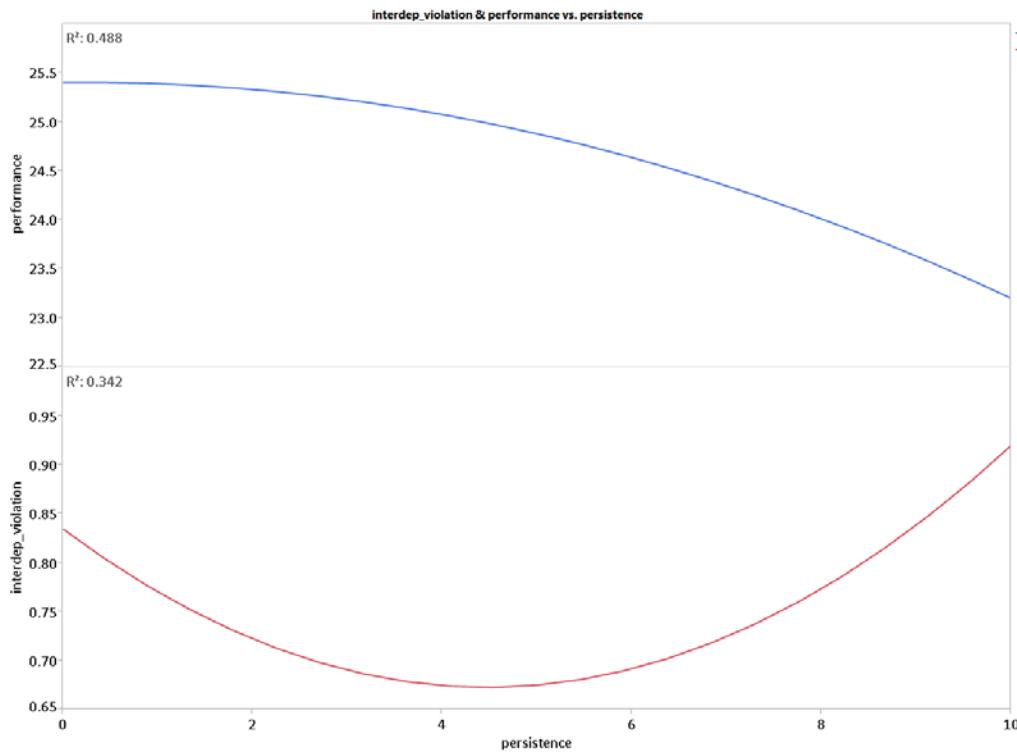


Figure 7. Average performance and average number of interdependence violations across different levels of persistence

This study aimed to advance our understanding of how human problem solving processes aggregate and whether they can be reliably designed to successfully solve problems of varying complexity. The overall approach was bottom-up: first, the experiments allowed the study of human problem solving behaviour and second, the simulations attempted to qualify how different search strategies affect the propensity of finding the optimal solution. Even if the model is agnostic with respect to intrinsic or extrinsic motivation, I find that focusing on immediate results (by not engaging in more risky, distant moves) has the potential to diminish learning and the likelihood to solve complex problems.

To begin with, the experimental findings are consistent with previous research on problem solving under uncertainty (e.g. Billinger, Stieglitz et al. (2013)). Consistent with previous work on the balance between exploration and exploitation (March 1988), I find that as solvers are more likely to engage in risk seeking behaviour if they consistently perform below their aspiration level. However, I also distinguish a third behaviour that solvers engage in: model-based distant search. In line with psychology and cognitive science literature (Gläscher, Daw et al. 2010; Doll, Simon et al. 2012), but also organizational literature (Gavetti and Levinthal 2000), I find that solvers do not engage merely in local or random-distant search, but attempt to identify and exploit patterns as they navigate the rugged problem landscape. This different type

of distant search is not influenced directly by feedback. Works such as Gavetti and Levinthal (2000) or Winter et al. (2007) have already brought cognition into the spotlight, but to the best of my knowledge, empirical work has not yet focused on the mechanisms or the propensity with which, when solving complex problems, solvers are likely to engage in model-based distant search. Additionally, while this finding is in itself interesting, being able to investigate what determines different types of searching behaviour is an important first step towards being able to aggregate these behaviours and to identify macro-level patterns.

Consequently, the findings from the experimental study were used in calibrating an agent-based model of adaptive search. The simulations carried allowed for different propositions to be investigated.

First, I explored whether indeed model-based distant search, can be a reliable tool in the search for the optimal solution. Virtual experiments pit local search against the agent-based model search. Results show that local search can be an efficient tool for simple problems and can lead to minor quick improvements for more complex problems. Still, in the case of (even mildly) rugged landscapes, relying on an exclusively local search strategy is inefficient. Additionally, a naïve local search and random algorithm was also outperformed by the model-based search. This result is interesting as it is telling of the importance of model-based distant search. By relying on (often imperfect) maps of the landscape, agents are able to greatly shorten the paths towards the global optimum as well as to increase their likelihood of finding it.

Secondly, I looked into what determines the success or failure of model-based distant search. “Fail early, fail often” is a phrase that has gathered quite a bit of traction in recent years as the mantra of one of the most successful design companies in the world (IDEO). This idea is captured in the information systems practice (and literature) by *agile software development* (Martin 2003), scrum (Conboy 2009) etc. The results suggest that a lower level of “persistence” (the willingness to abandon earlier on unsuccessful paths) translates into higher accuracy in terms of problem representation and improved performance. These results add to previous research regarding solving complex problems. On the one hand, the model does reproduce previous work and shows that problem decomposition is desirable (Ethiraj and Levinthal 2004) and that relying on gradually larger modules correlates with higher performance (Baumann and Siggelkow 2013). On the other hand, the specific simulation framework allows for a more detailed explanation of the search process. As such, results show, low persistence becomes important insofar as computational agents fail to detect an accurate problem representation. Simulations endorse empirical studies suggesting that for complex problems concentrating on immediate performance (and focusing on either local or model-based moves) increases the likelihood that solvers focus too little on learning (Manso 2011; Gardner 2012), and

consequently have an increased likelihood of making spurious correlations which trap them into suboptimal peaks.

Although these results do not contribute directly to the ongoing discussion about prizes on crowdsourcing platforms e.g. Archak and Sundararajan (2009), they do offer managers a different lens in making that decision. For example, they confirm that in the context of problem solving performance pressure indeed can act like a “double edged sword” (Gardner 2012) and complement work regarding incentive systems for innovation that suggests failure should be tolerated in highly innovative settings (Holmstrom 1989).

Additionally, another key contribution of this paper is the agent based model which is the backbone of this study. Overall, this paper contributes to the strategic organization literature and to the innovation-as-search literature in particular by proposing a model that allows for more plausible (and empirically validated) assumptions regarding solver behaviour to be taken into account. This is important because, especially in the context of individual problem solving, I have shown that model-based distant search is a prevalent search strategy and an important determinant for success.

A limitation of this study is its focus on *adaptive* search. For solving complex problems, there seems to be evidence that relying on multi-shot submissions (a set-up where solvers receive feedback from the broadcasting organization and can subsequently modify their solutions) is preferable to the classical broadcast search (Vuculescu and Bergenholtz 2014) and certainly there is little controversy regarding the usefulness of feedback in learning environments (Bangert-Drowns, Kulik et al. 1991). The current study is indirectly making the same argument: without having the possibility to receive and incorporate feedback, solvers cannot form an accurate problem representation and their performance depends on their prior knowledge (which is what the traditional Innocentive challenge is trying to exploit). While this is certainly not the norm, there are several instances where online platforms resorted to multi-shot submissions (e.g. the Netflix prize) and, recently, even Innocentive has resorted to setting up two or multiple phases for its challenges.

Finally, this study aims to bring forward a novel approach to problem solving research that takes into account the dynamic and complex nature of the phenomenon, while at the same time being informed by empirical research. Previous research on search and problem-solving has been either purely conceptual or relied on either qualitative research or quantitative designs. Using an agent based model offers both the advantages of conducting research that allows for a greater level of detail and answers the “how” question and for a higher degree of generalization (moving away from a particular empirical setting). Furthermore, the model is easily adaptable to other avenues of research. So far I have focused on external search and tailored the model accordingly, but the

model could be adjusted to replicate and advance knowledge regarding internal search, organizational learning and even collective cognition.

References

- Archak, N. and A. Sundararajan (2009). "Optimal Design of Crowdsourcing Contests." *ICIS 2009 Proceedings*: 200:229
- Baldwin, C. and E. von Hippel (2011). "Modeling a Paradigm Shift: From Producer Innovation to User and Open Collaborative Innovation." *Organization Science* 22(6): 1399-1417.
- Bangert-Drowns, R. L., Kulik, C. L. C., Kulik, J. A., & Morgan, M. (1991). "The instructional effect of feedback in test-like events". *Review of educational research*, 61(2), 213-238.
- Baumann, O. (2010). "Search, Failure, and the Value of Moderate Patience." *Schmalenbach Business Review* 62: 234-259.
- Baumann, O. and N. Siggelkow (2013). "Dealing with complexity: integrated vs. chunky search processes." *Organization Science* 24(1): 116-132.
- Beinhocker, E. D. (2006). *The origin of wealth : evolution, complexity, and the radical remaking of economics*. Boston, Harvard Business School Press.
- Benkler, Y. (2002). "Coase's penguin, or, Linux and The Nature of the Firm." *Yale Law Journal* 112(3): 369-+.
- Billinger, S., N. Stieglitz, et al. (2013). "Search on rugged landscapes: An experimental study." *Organization Science* 25(1): 93-108.
- Boero, R. and F. Squazzoni (2005). "Does Empirical Embeddedness Matter? Methodological Issues on Agent-Based Models for Analytical Social Science." *Journal of Artificial Societies and Social Simulation* 8(4): 6.
- Boudreau, K. J., N. Lacetera, et al. (2011). "Incentives and problem uncertainty in innovation contests: An empirical analysis." *Management Science* 57(5): 843-863.
- Brusoni, S., L. Marengo, et al. (2007). "The value and costs of modularity: a problem-solving perspective." *European Management Review* 4(2): 121-132.
- Chen, W.-R. (2008). "Determinants of firms' backward-and forward-looking R&D search behavior." *Organization Science* 19(4): 609-622.
- Cohen, M. D. and P. Bacdayan (1994). "Organizational routines are stored as procedural memory: Evidence from a laboratory study." *Organization Science* 5(4): 554-568.
- Collins, A. and E. Koechlin (2012). "Reasoning, learning, and creativity: frontal lobe function and human decision-making." *PLoS Biol* 10(3): e1001293.
- Conboy, K. (2009). "Agility from first principles: reconstructing the concept of agility in information systems development." *Information Systems Research* 20(3): 329-354.
- Dolan, R. J. and P. Dayan (2013). "Goals and habits in the brain." *Neuron* 80(2): 312-325.

- Doll, B. B., D. A. Simon, et al. (2012). "The ubiquity of model-based reinforcement learning." *Current Opinion in Neurobiology* 22(6): 1075-1081.
- Eggers, J. P. and S. Kaplan (2013). "Cognition and Capabilities: A Multi-Level Perspective." *Academy of Management Annals* 7(1): 295-340.
- Ericsson, K. A., N. Charness, et al. (2006). *The Cambridge handbook of expertise and expert performance*, Cambridge University Press.
- Ethiraj, S. K. and D. Levinthal (2004). "Modularity and innovation in complex systems." *Management Science* 50(2): 159-173.
- Exadaktylos, F., A. M. Espín, et al. (2013). "Experimental subjects are not different." *Scientific Reports* 3: 1213.
- Felin, T. and T. R. Zenger (2014). "Closed or open innovation? Problem solving and the governance choice." *Research Policy* 43(5): 914-925.
- Gardner, H. K. (2012). "Performance pressure as a double-edged sword enhancing team motivation but undermining the use of team knowledge." *Administrative Science Quarterly* 57(1): 1-46.
- Gavetti, G. (2005). "Cognition and Hierarchy: Rethinking the Microfoundations of Capabilities' Development." *Organization Science* 16(6): 599-617.
- Gavetti, G. and D. Levinthal (2000). "Looking forward and looking backward: Cognitive and experiential search." *Administrative Science Quarterly* 45(1): 113-137.
- Gavetti, G., D. Levinthal, et al. (2007). "Neo-carnegie: The carnegie school 's past, present, and reconstructing for the future." *Organization Science* 18(3): 523-536.
- Gavetti, G. and J. W. Rivkin (2007). "On the origin of strategy: Action and cognition over time." *Organization Science* 18(3): 420-439.
- Gigerenzer, G. and A. Edwards (2003). "Simple tools for understanding risks: from innumeracy to insight." *BMJ: British Medical Journal* 327(7417): 741.
- Gläscher, J., N. Daw, et al. (2010). "States versus rewards: dissociable neural prediction error signals underlying model-based and model-free reinforcement learning." *Neuron* 66(4): 585-595.
- Goldberg, D. E. (2000). "The design of innovation: Lessons from genetic algorithms, lessons for the real world." *Technological Forecasting and Social Change* 64(1): 7-12.
- Graybiel, A. M. (2008). "Habits, rituals, and the evaluative brain." *Annual Review of Neuroscience*. 31: 359-387.
- Greeno, J. G. (1977). "Process of understanding in problem solving." *Cognitive Theory* 2: 43-83.
- Hayes, J. R. and H. A. Simon (1975). *Psychological differences among problem isomorphs*, Carnegie-Mellon University, Department of Psychology.

- Hodgkinson, G. P., J. Langan-Fox, et al. (2008). "Intuition: A fundamental bridging construct in the behavioural sciences." *British Journal of Psychology* 99(1): 1-27.
- Holmstrom, B. (1989). "Agency costs and innovation." *Journal of Economic Behavior & Organization* 12(3): 305-327.
- Iclanzan, D. and D. Dumitrescu (2007). *Overcoming hierarchical difficulty by hill-climbing the building block structure*. Proceedings of the 9th annual conference on Genetic and evolutionary computation, ACM.
- Kahneman, D. (2011). *Thinking, fast and slow*, Macmillan.
- Kahneman, D. and A. Tversky (1979). "Prospect theory: An analysis of decision under risk." *Econometrica: Journal of the Econometric Society*: 263-291.
- Kaplan, C. A. and H. A. Simon (1990). "In search of insight." *Cognitive Psychology* 22(3): 374-419.
- Kauffman, S. A. (1993). *The origins of order: Self-organization and selection in evolution*, Oxford university press.
- Klein, G. A. (1998). *Sources of power: How people make decisions*, MIT press.
- Lakhani, K. R., H. Lifshitz-Assaf, et al. (2013). "Open innovation and organizational boundaries: Task decomposition, knowledge distribution, and the locus of innovation." *Handbook of economic organization: Integrating economic and organizational theory*: 355-382.
- Lant, T. K., F. J. Milliken, et al. (1992). "The role of managerial learning and interpretation in strategic persistence and reorientation: An empirical exploration." *Strategic Management Journal* 13(8): 585-608.
- Levinthal, D. A. (1997). "Adaptation on rugged landscapes." *Management Science* 43(7): 934-950.
- Levitt, B. and J. G. March (1988). "Organizational Learning." *Annual Review of Sociology* 14: 319-340.
- Lopez-Vega, H., F. Tell, et al. (Forthcoming). "Where and how to search? Search paths in open innovation." *Research policy*.
- Luo, J. and G. Knoblich (2007). "Studying insight problem solving with neuroscientific methods." *Methods* 42(1): 77-86.
- Manso, G. (2011). "Motivating innovation." *The Journal of Finance* 66(5): 1823-1860.
- March, J. G. (1991). "Exploration and exploitation in organizational learning." *Organization Science* 2(1): 71-87.
- Marengo, L. (2014). "Representation, search, and the evolution of routines in problem solving." *Industrial and Corporate Change*: dtu023.
- Martin, R. C. (2003). *Agile software development: principles, patterns, and practices*, Prentice Hall PTR.

- Mason, W. and D. J. Watts (2012). "Collective problem solving in networks." Proceedings of the National Academy of Sciences.
- McKelvey, B., L. Meng, et al. (2013). "Re-thinking Kauffman's NK fitness landscape: From artifact & groupthink to weak-tie effects." Human Systems Management 32: 17-42.
- Nelson, R. R. and S. G. Winter (2002). "Evolutionary Theorizing in Economics." Journal of Economic Perspectives 16(2): 23.
- Nickerson, J., C. J. Yen, et al. (2012). "Exploring the problem-finding and problem-solving approach for designing organizations." The Academy of Management Perspectives 26(1): 52-72.
- Nickerson, J. A. and T. R. Zenger (2004). "A knowledge-based theory of the firm—The problem-solving perspective." Organization Science 15(6): 617-632.
- Rivkin, J. W. and N. Siggelkow (2007). "Patterned interactions in complex systems: Implications for exploration." Management Science 53(7): 1068-1085.
- Sieg, J. H., M. W. Wallin, et al. (2010). "Managerial challenges in open innovation: a study of innovation intermediation in the chemical industry." R&D Management 40(3): 281-291.
- Siggelkow, N. and D. A. Levinthal (2003). "Temporarily Divide to Conquer: Centralized, Decentralized, and Reintegrated Organizational Approaches to Exploration and Adaptation." Organization Science 14(6): 650-669.
- Siggelkow, N. and D. A. Levinthal (2005). "Escaping real (non-benign) competency traps: Linking the dynamics of organizational structure to the dynamics of search." Strategic Organization 3(1): 85-115.
- Simon, H. A. (1956). "Rational choice and the structure of the environment." Psychological Review 63(2): 129-138.
- Simon, H. A. (1962). "The architecture of complexity." Proceedings of the American philosophical society 106(6): 467-482.
- Takeuchi, H. and I. Nonaka (1986). "The new new product development game." Harvard Business Review 64(1): 137-146.
- Terwiesch, C. and Y. Xu (2008). "Innovation contests, open innovation, and multiagent problem solving." Management Science 54(9): 1529-1543.
- Tripsas, M. and G. Gavetti (2000). "Capabilities, cognition, and inertia: Evidence from digital imaging." Strategic Management Journal 21(10-11): 1147-1161.
- Vuculescu, O. and C. Bergenholz (2014). "How to Solve Problems with Crowds: A Computer-Based Simulation Model." Creativity and Innovation Management: 23.2 (2014): 121-136.
- Ward, T. B. (1994). "Structured imagination: The role of category structure in exemplar generation." Cognitive Psychology 27(1): 1-40.

- Watson, R. A. and J. B. Pollack (1999). Hierarchically consistent test problems for genetic algorithms. Evolutionary Computation, 1999. CEC 99. Proceedings of the 1999 Congress on, IEEE.
- Wiley, J. (1998). "Expertise as mental set: The effects of domain knowledge in creative problem solving." Memory & Cognition 26(4): 716-730.
- Winter, S. G., G. Cattani, et al. (2007). "The value of moderate obsession: Insights from a new model of organizational search." Organization Science 18(3): 403-419.
- Zhang, J. (1997). "The nature of external representations in problem solving." Cognitive Science 21(2): 179-217.

Micro-foundations of problem solving: What determines how individuals search?
Oana Vuculescu*, Carsten Bergenholz*

***Department of Management, Aarhus University, Denmark**

Keywords: individual search, exploration and exploitation, antecedents of search behaviour, heterogeneity of search behaviour, mental maps

Abstract

This study contributes to the theory on problem solving as search by focusing on how individuals search and what individual antecedents determine search behaviour. So far research has mainly addressed the strategic distinction between exploration and exploitation at an organizational level, and has done less to disentangle the heterogeneity of individual level search behaviour processes. Combining data collected from individuals solving an experimental task with a quantitative survey we identify heterogeneous search behaviours as well as evidence that individual cognitive styles explain part of this heterogeneity. Relying on modelling and simulations, we also investigate how different search strategy preferences yield different performances for different types of problems.

Introduction

How do we explain the heterogeneity of search behaviour? March (1991) has framed the discussion on organizational search, by establishing exploration and exploitation as fundamentally distinct search strategies. Even though in March's (1991) seminal work, the focus is on the influence of *individuals*, in the succeeding discussion heterogeneity of search is primarily associated with the *organizational level*: e.g. organizational history and aspiration levels (Greve 2003), situational and institutional factors (Chen and Miller 2007) or the type of problem the organization is facing (Felin & Zenger 2014, Lakhani et al. 2013). Acknowledging some recent exceptions, individual level variables have not been at the center of attention (Audia and Goncalo 2007, Laureiro-Martínez et al. 2015). Still, search at an individual level has received increased attention in part due to the rise of crowdsourcing phenomena (Lakhani et al. 2013), in part due to recent calls for micro-foundations of organizational strategy (Felin and Foss 2005). Modelling and simulation have attempted to study search from a different angle, as this perspective aims to capture how search unfolds over time as well as non-linear dynamics and path-dependencies (Levinthal 1997; Csaszar and Levinthal 2015; Martignoni et al. 2015). However, just as the field in general, research relying on modelling and simulation assumes

homogeneous behaviour and that heterogeneity stems not from differences between agents, but changes in the competitive environments that organizational agents search in (Levinthal 1997, Rivkin 2000, Ethiraj and Levinthal 2004).

In line with emerging calls for uncovering micro-foundations of the behavioural process of individual search (Mason and Watts 2012; Billinger et al. 2013; Csaszar and Levinthal 2015) and in particular the individual antecedents of heterogeneous search strategies (Laureiro-Martínez et al. 2015, Puranam et al. 2015) a number of recent studies investigate the potential impact of heterogeneity of individuals on search behaviour (Audia and Goncalo 2007, Gary and Wood 2011, Laureiro-Martínez et al. 2015, Døjbak et al. 2015). In this paper we aim to develop this alternative approach. We focus on the information processing level of individuals and in order to go beyond more theoretical assumptions (see e.g. Martignoni et al. 2015) we empirically study the individual antecedents and the actual behavioural processes of individual search behaviour.

We collect quantitative data on individual traits and search behaviour from 367 participants as well as qualitative interviews (40) from individuals that have played an experimental game, similar in style to Billinger et al. (2013). We show how the cognitive style of an individual, a stable trait (Kirton 1976, Miron et al. 2004), can explain substantial variation of the heterogeneity of search behaviour, adding to explanations focusing on the environment, such as the problem (Felin & Zenger 2014) or the competitive dynamics (Levinthal 1997). Results show that more creative people are more likely to engage in more random search and less likely to engage in local search, while participants who score low on attention to detail are more likely to engage in random search. Thus, we argue that Felin and Zenger's (2014) framework might require further qualification, since the organization might not have access to the appropriate individuals to carry out the type of search the given problem calls for. Furthermore, we connect cognitive science literature on mental models with theoretical conceptions of managerial mental models within organizational theory (Cornelissen and Werner 2014). In conjunction with these theoretical insights we use qualitative interviews to illustrate the actual behavioural process and indicate how search behaviour unfolds over time.

Overall, this theoretical and empirical setup enables us to contribute to theory on problem solving by going beyond a mere strategic focus on either exploring or exploiting, in order to consider how antecedents influence search behaviours and how the mental models of individuals lead to meaningful exploration.

In section 2 we present an overview of how organizational theory conceptualizes mental maps and individual search behaviour as well as potential individual level antecedents of heterogeneous search behaviour. In section 3 we describe the data and methods, while in section 4 we present the results from our quantitative and qualitative analysis. In section 5 we conclude with a discussion on heterogeneity of individual search, and the implications for future research.

2. Theoretical background of search and mental maps

Problem solving as search has its roots in March and Simon's (1958) bounded rationality perspective on how organizations process information. This has given rise to two main traditions in organizational theory (Lant and Shapira 2001): the interpretivist approach and the computational perspective. The former focuses on how “‘meaning’ is created around information in a social context” (Lant and Shapira 2001: p 6), while the latter laid the foundation for the information processing perspective of problem solving and decision making. Both types of research, the first one mainly qualitative (e.g. Kaplan and Tripsas 2008, Gavetti and Rivkin 2007) and the second mainly simulation-based (Gavetti 2005, Martignon et al. 2015, Csazsar and Levinthal 2015) have argued for the importance of mental models (or cognitive maps) for explaining the strategic behaviour of managers and there is increasing evidence that mental models have an impact on e.g. strategic choices (Gavetti and Levinthal 2000; Gavetti 2005; Gavetti and Rivkin 2007; Kaplan and Tripsas 2008).

The large and heterogeneous (both in terms of scope and level of analysis) body of recent work based on the interpretivist approach is primarily concerned with the influence of ‘cognitive structures’(cf. Narayanan et al. 2011). Less attention is given to the information processing limitations of the focal agents (cf. Hodgkinson and Healey 2008). This work revolves around cognitive frames (i.e. mental models) as knowledge structures that influence how agents perceive the environment and thus their subsequent actions. Thus, the focus is on study of how actors engage in what current cognitive science research denotes as “model-based search” - actions which rely on “*expectancy processes in the deliberative system*” (Redish 2012: p.92). Generally these studies have struggled to take into account the inherent dynamic nature of mental models, i.e. how individuals continuously update their internal representation and adjust their behaviour accordingly (cf. Hodgkinson and Healey 2008; Cornelissen and Werner 2014, but see also Rivkin and Gavetti 2007; Tripsas and Kaplan 2008).

Despite sharing its roots with cognitive science (Lant and Shapira 2001), the second stream of literature, the computational perspective, has been largely dominated by the behaviourist perspective. This literature mainly relies on the ‘exploration vs exploitation trade-off categorization’: “Exploration includes things captured by terms such as search, variation, risk taking, experimentation, play, flexibility, discovery, innovation. Exploitation includes such things as refinement, choice, production, efficiency, selection, implementation, execution.” (March 1991: p.71). The trade-off is captured in a local vs. and distant search distinction, generally operationalized as local vs. random. This second stream of literature is thus not without limitations, as it has struggled with going beyond a very limited view on bounded rationality that entails that most search behaviours are quasi-automatic and that agents cannot overcome

cognitive, information processing limitations which means their actions are largely determined by the environment in which they act (cf. Lant and Shapira 2001, Csazar and Levinthal 2015). In cognitive science, this type of automatic behaviour is known as “model-free” search (Doll et al. 2012), since it implies the learner never acquires an understanding (model) of the ‘world’. Representations of the environment or knowledge structures are not as central in this literature. This is probably largely due to the fact that even cognitive science has struggled to get a clear grasp of how mental models are formed (e.g. how individuals generate and evaluate new solutions relying on internal representations), as well as what is the underlying biological basis for these maps (Barsalou 2008). This is then in contrast with the previous perspective which assumes a more intelligent agent that makes deliberative moves that involve a degree of prediction of the outcomes of a given action (Doll et al. 2012).

More recently these two fundamentally different accounts (i.e. model-based and model free) of how humans search and make decisions have been reconciled under a common framework (i.e. a two-systems approach) and the general agreement is that they coexist as complementary mechanisms (Solway and Botvinick 2012). Although little research focuses on how mental models are formed (cf. Gläscher et al. 2010; Smaldino and Richerson 2012), there is an emerging consensus that the same fundamental learning mechanisms are responsible for how the models are subsequently adapted. Thus, mental models are not only informed by search behaviours, search behaviours are also determined by mental models.

However, organizational literature so far has not focused on both types of search behaviours and the interplay between them. Acknowledging a few exceptions (e.g. Gary and Wood 2011) studies have not captured the path dependent process, i.e. the feedback loop between a focal agent’s mental model of the environment and its subsequent search behaviours (Narayanan et al. 2011).

2.1. Individual level search behaviours

In the computational perspective, if one searches locally in well-known surroundings, his actions are classified as exploitation, while distant, random searches are considered exploration (Levinthal 1997, Csaszar and Levinthal 2015). Much like “exploration vs exploitation”, this distinction is subject to interpretation (Lavie et al. 2010) since it is not immediately clear how ‘local’ is defined (Rosenkopf 2008). Nonetheless, this operationalization allows for very simple encodings of agent behaviour: e.g. when receiving positive rewards from the environment an agent will be less prone to switch strategies (since success leads to exploitation) and when receiving negative feedback the agent will engage in distant search which would allow him to explore different parts of the landscape.

In addition to the simulation literature, a key stream of research has tried to explain cognitive search processes that underlie the strategic choice between different search behaviours, by experimentally investigating the decision to either explore or exploit. More specifically, focus

has been on a very limited set of a priori options (e.g. four in Laureiro-Martínez et al. 2015) and how people either explore the unknown or exploit the known options (Solway and Botvinick 2012). The multi-armed bandit (i.e. gambling/slot machine) is a typical example. Game-players can successively choose between the four options (levers) with unknown payoff distributions, in order to generate as many points as possible. Laureiro-Martínez et al. (2015) argue that if one switches between the four levers, an explorative move has been made. But, such an experimental setup is quite different from a situation where a solver is faced with a complex problem and has to create a representation of the problem and shape the options to be considered, out of many possible (potentially infinite) options available (Smaldino and Richerson 2012). This difference between selecting between pre-defined options (of exploration or exploitation) vs. shaping meaningful new options seems relevant, given recent neurological studies: “choice behaviour [...] differ[s] between self-generated and externally provided options” (Kaiser, Simon et al. 2013: p.815).

Lopez et al. (2016) offer a more nuanced categorization of organizational search in order to go beyond a simple two-fold distinction and capture firm search processes. In particular, by simultaneously focusing on the ‘where’ (local/distant) and the ‘how’ (cognitive/experiential) they identify four distinct search paths that firms can engage in: situated, analogical, scientific and sophisticated paths. Similarly, on the individual level recent work has tried to study if our stylized conceptualization of how search occurs fits empirical data (Billinger et al. 2013). In Billinger et al.’s (2013) study players were trying to solve a puzzle with 1024 options. They show that, compared to the stylized conceptualization of search, human problem solvers tend to engage in less local search moves and break out of exploitation sooner than expected in simulation models (Billinger et al. 2013). Mason and Watts (2012) expose human solvers and computer-based agents to the same problem of balancing explorative and exploitative search in a landscape. They conclude that search behaviour should not merely be captured by local vs. distant search, since human beings display more heterogeneous and sophisticated search behaviours than our current models of simulated agents assume. Finally, in a similar attempt to uncover heterogeneity of search behaviours, Laureiro-Martínez et al. (2014) identify differences in how managers and entrepreneurs balance the exploitation - exploitation trade-off, when playing the armed bandit.

2.2. Antecedents of search behaviour

If we accept heterogeneity in search behaviours this raises the question whether this heterogeneity is generated by individual level variables such as traits (Felin and Foss 2005, Foss & Pedersen 2014). Indeed, Hills and Hertwig (2010) speculate that the differences they identified in terms of different search behaviours can be attributed to e.g. differences in terms of risk sensitivity (Hills and Hertwig 2010). In line with this perspective, a recent stream of research

tries to capture individual antecedents that can explain variation in adoption of search strategies search. These studies focus on, for instance, neurological activity (Laureiro and Martinez et al. 2015) or the emotional basis (Døjbak et al. 2015) for balancing exploration and exploitation. But, taking the cognitivist approach to search, i.e. assuming that the focal agent is an information processing entity which takes in ‘inputs’ from the environment and generates responses, various individual propensities for processing information should be a potentially important factor that influences search behaviour. Studies on such individual differences have been developed by psychology, either relying on a measure of how creative an individual is or how individuals differ in their cognitive style. The former measure, e.g. the Creative Personality Scale (Gough 1979), captures creativity operationalized as the number of new ideas one develops (Shalley et al. 2004).

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

A number of such measures of cognitive styles have been developed, e.g., adaption–innovation (Kirton 1976), analytic–intuitive (Allinson and Hayes 1996) and field dependence–independence (Witkin and Goodenough 1977). All show that cognitive styles are similar in kind to personality traits, i.e. they are stable over time, involve no optimal or best configuration and they constitute a continuum rather than a binary scale. 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 “liable to indulge in wider solution search” (Kirton 2003: p. 49), i.e. 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. In other words, adaptive people are less likely to change their idea about how to solve a problem, while innovators tend to see the “boundaries between different ideas as being more porous and flexible” (Jablokow et al. 2015: p. 308). Importantly, the A-I theory maintains that innovators are not per se “better” at being creative nor do adaptors make “just reactive” changes (Kirton 2003: p. 125), and that, furthermore, cognitive styles are unrelated to cognitive ability (Kirton 2003).

The cognitive style measure has become more widely used in organizational theory, since it has been repeatedly shown to influence behaviour and performance in organizations (Carnabuci and Dioszegi 2015). Carnabuci and Dioszegi (2015) e.g. show that an adaptive individual in a social

network characterized by structural holes receives better performance ratings, than the innovative individual in a similar network position. They thus argue that a complementary fit exists between the context of an individuals' position in a social network, and their cognitive style. Miron-Spektor et al. (2011) find support for how the A-I theory relates to different phases of the innovation process and show that teams with a higher number of members who score high on creativity and conformity are associated with radical innovation.

In this study, we investigate how cognitive styles relate to different types of search strategies in the ongoing process of search, rather than creativity or the final ability to solve the problem as such. We aim to investigate whether individual differences in information processing can be mapped out in how individuals tend to navigate (i.e. search) a solution landscape. More specifically, we show 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.

3. Setting, data and methods

We rely on a mixed methods approach (Greene et al. 1989), utilizing three different types of data, collected at the same time in order to quantitatively and qualitatively analyse how 367 individuals try to solve the ‘Alien Game’ (see <http://scienceathome.org/games/aliengame/>). 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 all individuals completed just before playing the game. Finally, we have also interviewed 40 individuals just after they finished playing the game, in order to complement (Greene et al. 1989) and validate our quantitative search constructs and generate further qualitative insight into the behavioural search process.

The experimental task

Relying on an experimental design we position ourselves somewhere in between the two extremes of previous work: in our experiment, we have allowed participants more freedom with respect to generating solutions, but the space of solutions is clearly defined as 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 (Huys et al. 2012). Such a setup allows us to focus on mental models and their influence on subsequent search behaviours - 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 in this solution space.

We use ‘The Alien Game’ (see <http://scienceathome.org/games/aliengame/>), a task where a solution to the game is given by a sequence of eight tiles. The tiles can be toggled in any of two positions, yielding 256 possible combinations. Participants have 25 attempts to search in an environment where they do not merely have very few options available, but have to create and test their ideas about how to solve the problem.

Thus, at any given attempt, a participant toggles any number of tiles and submits the constructed solution. The average search distance is 2.68 (meaning on average players toggle more than 2 tiles at a given time) with a standard deviation of 1.98, a result which is remarkably similar to the one in Billinger et al. (2013), where they report values ranging from 2.58 to 2.72 in a game with 1024 options. The feedback is a point score (ranging from 18 to 42) which relies on a hierarchical function described below and in more detail in Vuculescu (2015). In order to eliminate variance due to different performance expectations (Kahneman and Tversky 1979), players are informed in advance about the maximum score, as well as that, on top of a baseline pay, they will receive rewards based on their performance; i.e. the higher the score and the quicker the high score is achieved, the better the reward. The rewards based on performance range from 50 to 250 per cent of the baseline pay. The data collected from the experimental task consists of the sequence of submitted solutions and time stamps.

Task structure

The underlying function in the experimental set-up is an 8-bit ‘hierarchical exclusive OR’ (H-XOR) (Watson and Pollack 1999). We choose this function because unlike the traditional functions used to generate NK landscapes (Kauffman 1993), the solution space is structured thus allowing players to more easily detect and exploit the non-random patterns. In other words, in the NK setting players attempt to solve a randomly generated problem, where extrapolating from past experience to generate useful new solutions is not really possible¹ (Parter et al. 2008). In contrast, in the H-XOR setting players can detect an underlying pattern that they can meaningfully extrapolate from to other areas of the landscape that they have not navigated yet. Specifically, the H-XOR function is given by applying recursively an ‘exclusive or’ transformation onto the solution string where adjacent positions are considered starting with the left-most². The data collection was carried out in two phases with two distinct studies: one conducted via Mturk (270) and one conducted during several experimental sessions in the lab

¹ Given that the contribution of interdependent components simply assumes new draws from the uniform distribution.

² For instance, a {1010 0010} string becomes first {11 -1} and then {--}. Once the transformation is completed, the payoff function rewards each non-null position in the hierarchy. Thus, a solution which contains an alternating pattern {1010 1010} would give a better score than a {1111 1111} since it will generate payoffs at lower levels of the hierarchy as well. The second level transformation for the first solution is {11 11} while for the second it is {-- --}. The maximum score is given by {1001 0110} or symmetrically by {0110 1001}.

(97), enabling us to test findings from the first sample. Amazon's Mechanical Turk (Mturk) is a web-based outsourcing platform where 'requesters' set up various tasks and 'workers' select which tasks to complete for a pay corresponding on average to 6\$/hr. Its use in behavioural research is on the rise due to ease of access to a relatively large group of people that resembles the general (US) population (Berinsky et al. 2012) and relatively low costs in setup (Mason and Suri 2012).

Before the statistical analysis was carried out, seven participants were dropped from the Mturk sample after an analysis of their irregular answers and game play. The sample consists of 42 per cent women and 58 per cent men, with an average age of 34 years (standard deviation = 9.89). In terms of the highest level of education, 98 per cent of the Mturk participants reported having completed at least their secondary education and over 48 per cent reported having 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). On average they spent 18 minutes on completing the tasks. The laboratory sessions were conducted in a social science lab at a large University and conducted by a research assistant, not co-authoring our paper nor familiar with any research expectations. Participants are recruited using an internal recruitment system. The sample consists of 49 per cent women (51 per cent men), with an average age of 24 years (standard deviation = 5.84). Lab participants are students at the university, being either bachelor or master students in social sciences, humanities or natural sciences. The laboratory sessions are accompanied by further quantitative and qualitative data collection. The procedure for the quantitative data collection is identical for both samples: participants are first asked to fill in a short survey containing the A-I items and then redirected to the browser game¹. Then they watch the game instructions and play a tutorial level to get familiarized with the game environment and thereafter move on to the actual game. Figure 1 shows a screenshot of the game design.

¹ Available at: <http://scienceathome.org/play/aliengame>.

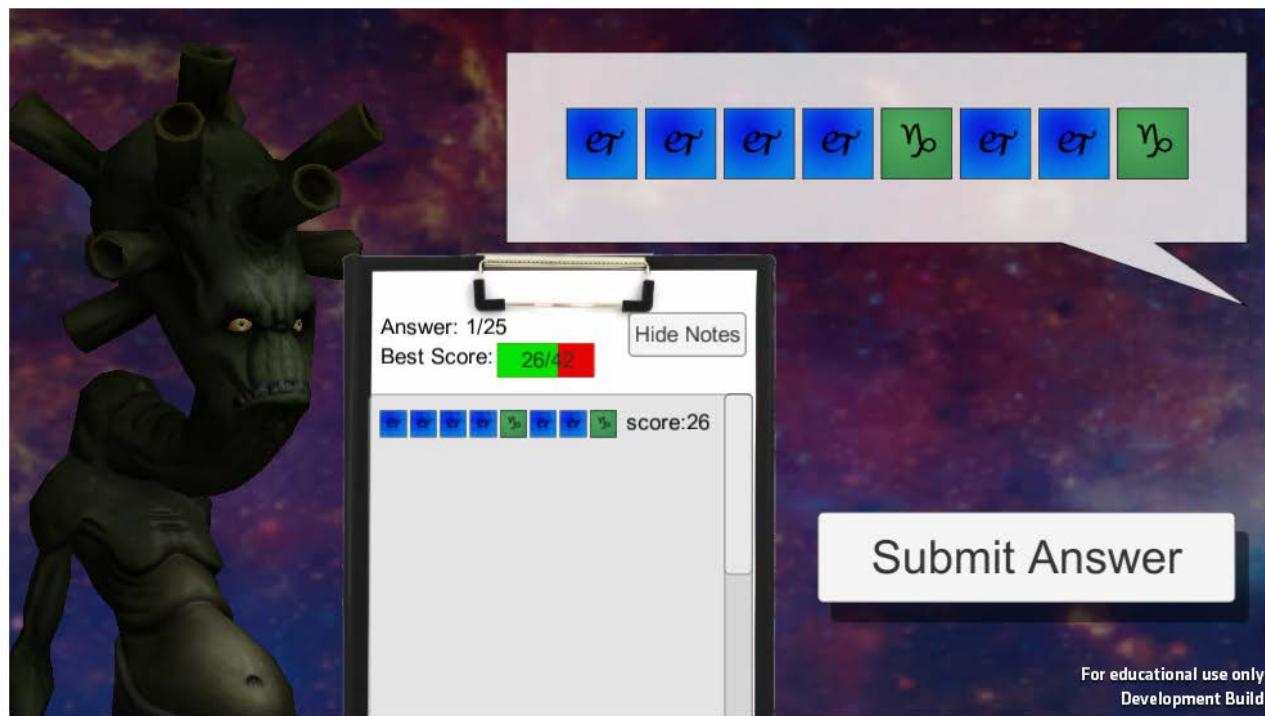


Figure 1. Screenshot of the game

3.2. Variables

Search strategies

In order to study how individual antecedents explain variation in human search behaviour, we need to distinguish between different types of search strategies. Traditionally organizational research operates with a twofold distinction: local and distant search. Building on recent cognitive science literature (Doll et al. 2012) and previous research on a similar game (Vuculescu 2015) we distinguish between three search strategies: local, model-based and random search. We thus differentiate between two fundamentally different search mechanisms: *model-free* mechanisms which operate without an internal representation of the problem space and *model-based* 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 moves as moves involving exactly two-bit flips from their reference point without violating the underlying problem structure¹ and random search as

¹ 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.

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 model-based search is not normative. Presumably our solvers make a number of model-based moves that *do* violate the underlying problem structure (since their representation of the problem space is erroneous) and thus are not captured by this coding. However, this scheme is preferred since it is the most conservative. It is further validated by comparing the solvers' own qualitative accounts of their moves and the coding scheme.

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 have tried to search through the solution space. To avoid biasing the results we only analyse submissions 3-25¹.

A-I

We use Miron et al.'s (2004) 12-item scale for capturing 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 all the questions.

Following a first analysis of the three factors, one item has been dropped, as it results in a relatively low Cronbach alpha (0.64)² for the respective factor as well as contributing to an overall poorer fit for the model ($\text{RMSEA} = 0.0848$ and $\text{GFI}=0.923$). The deleted item is related to the 'conformity' factor; "I avoid cutting corners". We attribute the poor fit to a potential effect of social desirability bias. The resulting Cronbach alpha for the revised factor is 0.69 which is satisfactory (Flynn et al. 1994) given our sample size and the exploratory nature of this work. For the other two factors Cronbach alpha indicates reliable measures: 0.859 for 'attention to detail' and 0.90 for 'creativity', respectively. To further investigate the robustness of the scale, the factor analysis is performed separately for each of the two samples (Mturk and lab) and a similar pattern emerges: low reliability for the 'conformity' factor that increases once the item is removed, suggestive of the fact that the instrument (without the specific item) is reliable³.

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 an

¹ Since the coding of search moves involved search distances, the first submission would serve as reference point, while in the second attempt the only possible search strategy that could be identified would be local search.

² 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.

³ As a robustness check, subsequent analyses were conducted both including and excluding the item and resulting estimates do not change.

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 and Cudeck 1992: p. 239). Item loadings are all highly significant ($p<0.001$) cf. Table 1.

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	0.6581*** (0.0857)		
I adapt myself to the system	0.6687*** (0.0635)		
I adhere to accepted rules in my area of work	0.8757*** (0.064)		
Thorough when solving problems		0.7896*** (0.0485)	
Addresses small details needed to perform the task		0.8662*** (0.052)	
performs the task precisely over a long time		1.004*** (0.0542)	
Good in tasks that require dealing with details		0.9411*** (0.0576)	
I have a lot of creative ideas			1.2682*** (0.0643)
I prefer tasks that enable me to think creatively			1.3695*** (0.0628)
Innovative			1.2486*** (0.0603)
I like to do things in an original way			1.0189*** (0.0607)

N=366, *** $p<0.001$

As a further validation of the instrument we have conducted an analysis of the factor covariation which yields results similar to those reported by Miron et al. (2004). Table 2 illustrates that creativity and conformity are not significantly correlated (0.3), but there is a significant correlation between conformity and attention to detail (0.44) and between creativity and attention to detail (0.37). Finally, the three factors (creativity, attention to detail and conformity) were recoded as binary variables, following Miron et al. (2011), by using the top 20% values as a cut-

off point. Thus, respondents who scored above the 20% threshold were coded as being “high” on that particular dimension.

Table 2. Factor correlation matrix for the KAI 3-Factor model

Factor Correlation Matrix: Estimate (StdErr)		conformity	attention to detail	creativity
conformity	1	0.4461*** (0.0556)	-0.059 (0.0621)	
attention to detail	0.4461*** (0.0556)	1	0.3726*** (0.0513)	
creativity	-0.059 (0.0621)	0.3726*** (0.0513)	1	

N=366, ***p<0.001

Interviews

Following lab sessions, we randomly selected 40 participants to take part in short (on average 7.5 minutes) semi-structured interviews. Following 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).

4. Results

We first present the regression analysis as well as a descriptive analysis of the most typical submissions. The section concludes with an analysis of the qualitative interviews that shed further light on the behavioural search process.

4.1. The antecedents of search behaviour

To explore whether cognitive style (as expressed by the three A-I factors) has an impact on each of the three search strategies, we first analysed the Mturk dataset. Tables 3, 4 and 5 reports results from the analysis, for each of the three dependent variables. The first model we analysed 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 random search the estimate is = 0.01069 and the standard error = 0.01384), we estimate a third model, a marginal logistic regression model, expressed in Equations 1-3¹:

$$Var(Y_{ij}) = \varphi v(\mu_{ij}) \quad (1)$$

To account for potential autocorrelation in the observations, we included a multiplicative over-dispersion parameter ($\varphi \neq 1$) in our model and the variance function is:

$$v(\mu_{ij}) = \mu_{ij}(1 - \mu_{ij}) \quad (2)$$

And the link function is the logit function:

$$\log(\mu_{ij}/(1 - \mu_{ij})) = X_{ij}\beta, \quad (3)$$

Since we model the data as a marginal model, the conditional mean of the j^{th} response, given X_{i1}, \dots, X_{in_i} , depends solely on X_{ij} . Given that our covariates are not time-dependent, this assumption poses no difficulties.

Thus, for example, for an individual i to choose local search at time j the conditional expectation is given by:

$$\log \left\{ \frac{\Pr(LS_{ij}=1)}{\Pr(LS_{ij}=0)} \right\} = \beta_1 + \beta_2 creativity_i + \beta_3 conformity_i + \beta_4 detail_i + e_{ij} \quad (4)$$

The same model was fitted for model-based (Equation 5) and random search (Equation 6).

$$\log \left\{ \frac{\Pr(MB_{ij}=1)}{\Pr(MB_{ij}=0)} \right\} = \beta_1 + \beta_2 creativity_i + \beta_3 conformity_i + \beta_4 detail_i + e_{ij} \quad (5)$$

$$\log \left\{ \frac{\Pr(R_{ij}=1)}{\Pr(R_{ij}=0)} \right\} = \beta_1 + \beta_2 creativity_i + \beta_3 conformity_i + \beta_4 detail_i + e_{ij} \quad (6)$$

¹ In Appendix 1 we include an alternative way of modelling the same data, by using a multinomial regression model.

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.

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

Response variable = random search (binomial)		Estimate /st. error	p value	-2LL	Cov param /st.error
Model 1	creativity(ref=low)	-0.2889 0.09508	0.0024	4332.03	
	detail(ref=low)	0.1001 0.1113	0.3683	4331.21	
	conformity(ref=low)	0.1524 0.10	0.1582	4328.04	
	creativity(ref=low)	-0.2889 0.0951	0.0024		0.01069 0.0024
	conformity(ref=low)	0.1533 0.11	0.1557		0.01097 0.01394
Model 2	detail(ref=low)	-0.06217 0.09678	0.5207		0.01101 0.01395
	creativity(ref=low)	-0.2889 0.0951	0.0024		1.0004 0.01964
	conformity(ref=low)	0.1524 0.11	0.1080		0.9997 0.01963
	detail(ref=low)	0.05852 0.1141	0.6079		0.9999 0.01963
Model 3	creativity(ref=low)	-0.2889 0.0951	0.0024		
	conformity(ref=low)	0.1524 0.11	0.1080		
	detail(ref=low)	0.05852 0.1141	0.6079		

Table 4. Effect of the three A-I factors on the likelihood of engaging in local search (Mturk study)

Response variable = local search (binomial)					
		estimate (st.err)	p value	-2LL	Cov param/ st. error
Model 1	creativity(ref=low)	0.2128 0.07226	0.0032	6960.02	
	detail(ref=low)	0.1047 0.07065	0.1383	6966.45	
	conformity(ref=low)	-0.03753 0.07564	0.6198	6968.39	
	creativity(ref=low)	0.2126 0.07243	0.0033		0.000463 0.000681
	detail(ref=low)	0.1003 0.07082	0.1568		0.00046 0.000677
Model 2	conformity(ref=low)	-0.04381 0.07583	0.5634		0.000465 0.000684
	creativity(ref=low)	0.2128 0.07227	0.0033		1.0004 0.01964
	detail(ref=low)	0.00697 0.08066	0.9311		1.0006 0.01965
	conformity(ref=low)	-0.07969 0.07703	0.3009		1.0006 0.01965
Model 3	creativity(ref=low)	0.2128 0.07227	0.0033		
	detail(ref=low)	0.00697 0.08066	0.9311		
	conformity(ref=low)	-0.07969 0.07703	0.3009		

Table 5. Effect of the three A-I factors on the likelihood of engaging in model-based search (Mturk study)

Response variable = model-based search (binomial)

		Estimate/ st.error	p value	-2LL	Cov param/st error
Model 1	creativity(ref=low)	-0.07924	0.3377	5716.15	
		0.08264			0.00104
	detail(ref=low)	-0.08638	0.2825	5715.92	1
		0.08036			0.00150
	conformity(ref=low)	-0.04943	0.5641	5716.74	6
		0.08571			0.00103
Model 2	creativity(ref=low)	-0.07839	0.3449		7
		0.08298			0.00150
	detail(ref=low)	-0.07964	0.3237		1
		0.0807			0.00103
	conformity(ref=low)	-0.04056	0.6375		8
		0.08607			0.00150
Model 3	creativity(ref=low)	-0.07924	0.3377		3
		0.08265			1.0004
	detail(ref=low)	-0.08638	0.2825		0.01964
		0.08037			1.0004
	conformity(ref=low)	-0.04943	0.5642		0.01964
		0.08573			1.0004

The same marginal logistic model is subsequently applied on data collected in the university lab and we are thus able to test results based on the Mturk sample (Table 6).

**Table 6. Effect of the three A-I factors on the likelihood of engaging in the three search strategies:
(Lab study)**

Type of search behaviour/KAI Factor	Local search	Model-based search	Random 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

Table 6 shows that the laboratory dataset replicates the results from the first study, i.e. participants who score high on ‘creativity’ will be less likely to engage in local search and more likely to engage in random search. Note that random 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 random search moves, and also individuals who score high on attention to detail are more likely to do local search. These results seem to be 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 approach¹. Although the first dataset does not reveal the same pattern with respect to the second result, the lab results are supported by the multinomial modelling of the pooled dataset (Appendix 1). Results from this analysis show that players who score high on attention to detail are, relative to random search, more likely to do either local search or model-based search. We attribute this discrepancy to the particular nature of our first sample: Mturkers receive ratings according to their performance in a given tasks and this impact 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 (i.e. attention to detail) or they are more likely to self-report higher levels of attention.

¹ Indeed, in interviews some respondents equate ‘local search’ moves with ‘systematic search’.

Indeed the scores on this factor differ significantly across samples (Average lab=5.1, Average Mturk=5.47, $p=0.0013$), a discrepancy which we do not find for our other predictor (creativity).

4.2. Priors in the game

After showing the antecedents of search behaviour focus is shifted to the actual game-play. According to Smaldino and Richerson (2012), alongside psycho-biological determinants, solution generation is influenced by socio-cultural factors or prior experience (Eggers and Kaplan 2013). In this section we thus investigate the extent and nature of priors, i.e. early typical solution that indicate biases in solver's prior cognitive mappings. First of all we identify the priors in the overall distribution of all submitted solutions (Figure 2 and Table 7).

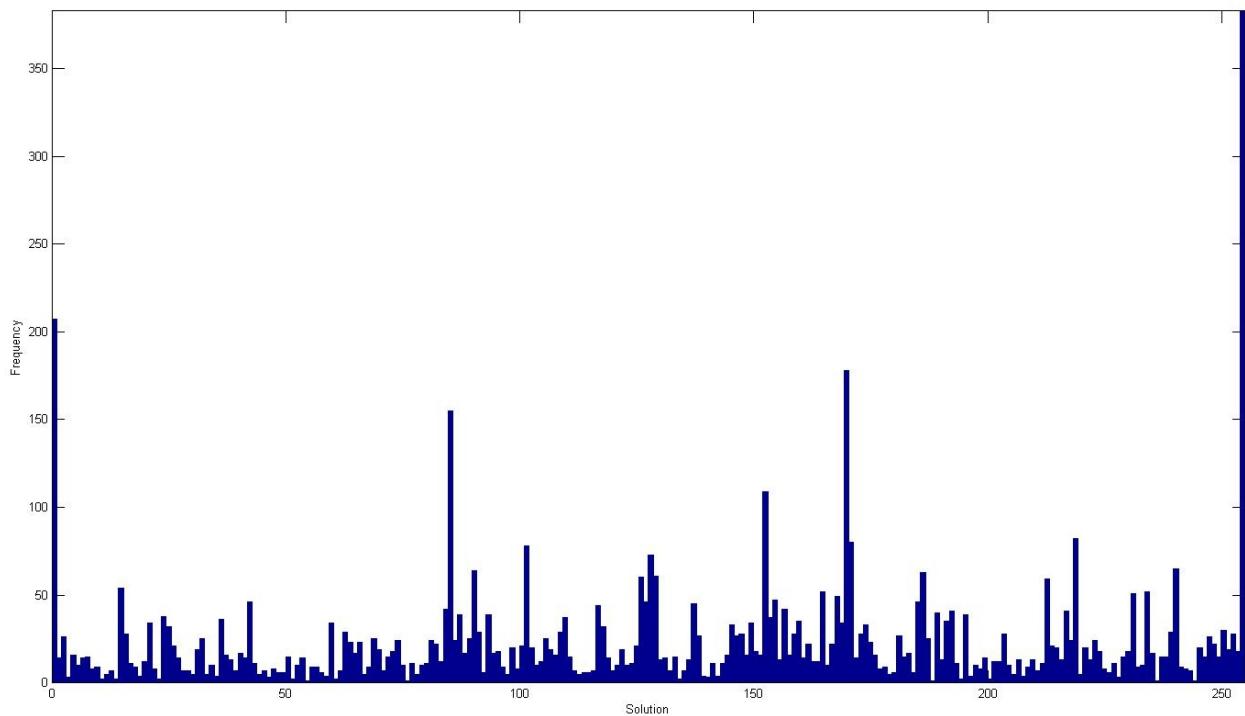


Figure 2. Solution distribution throughout the entire game. Submissions were recorded as binary strings which were converted into decimal numbers.

Table 7. Frequency for the most used solutions throughout the game

11111111	332	Frequency for most used solution:	332
10101010	178	95% of solutions are used less than:	60
00000000	172		
01010101	155		
11011011	82		
10101011	80		
10011001	79		
10000000	73		
11110000	65		
01011010	64		
10111010	63		
01100110	62		
10000001	61		
01111110'	60		

This overall picture shows that a number of solutions have a much higher frequency as compared to others and most of these solutions are not local optima in the game. This suggests that irrespective of the feedback a number of solvers are more likely to try out submissions which one participant calls ‘the usual and common ones’ (see Table 8). They are mainly those of a particular, stereotypical nature, either involving all 1s and 0s, or alternating 1s and 0s. Notice that these priors entail a score of 18 or 20, which are the two lowest possible scores in the game. The priors are thus not attractive in terms of the immediate score.

Table 8. First 6 submissions by one of the interviewed participants. He referred to his first attempts as “trying out the usual ones, the common ones”.

solution								trial number	score
0	0	0	0	0	0	0	0	1	18
1	1	1	1	1	1	1	1	2	18
0	1	0	1	0	1	0	1	3	26
1	1	1	1	0	0	0	0	4	18
1	1	0	0	1	1	0	0	5	18
1	0	0	0	0	0	0	1	6	22

Subsequently we refine this analysis by focusing on the first submission (Table 9 and Figure 3) in the game; we confirm that a number of solutions (i.e. ‘priors’) overwhelmingly outweigh other solutions in terms of frequency. The chance of one participant randomly choosing any one of the top nine priors is approximatively 3.5 per cent. If we assume our participants are making

independent draws, we should only find a total of 12 such solutions among all the first submissions. In contrast, 250 of 366 solutions are one of these priors.

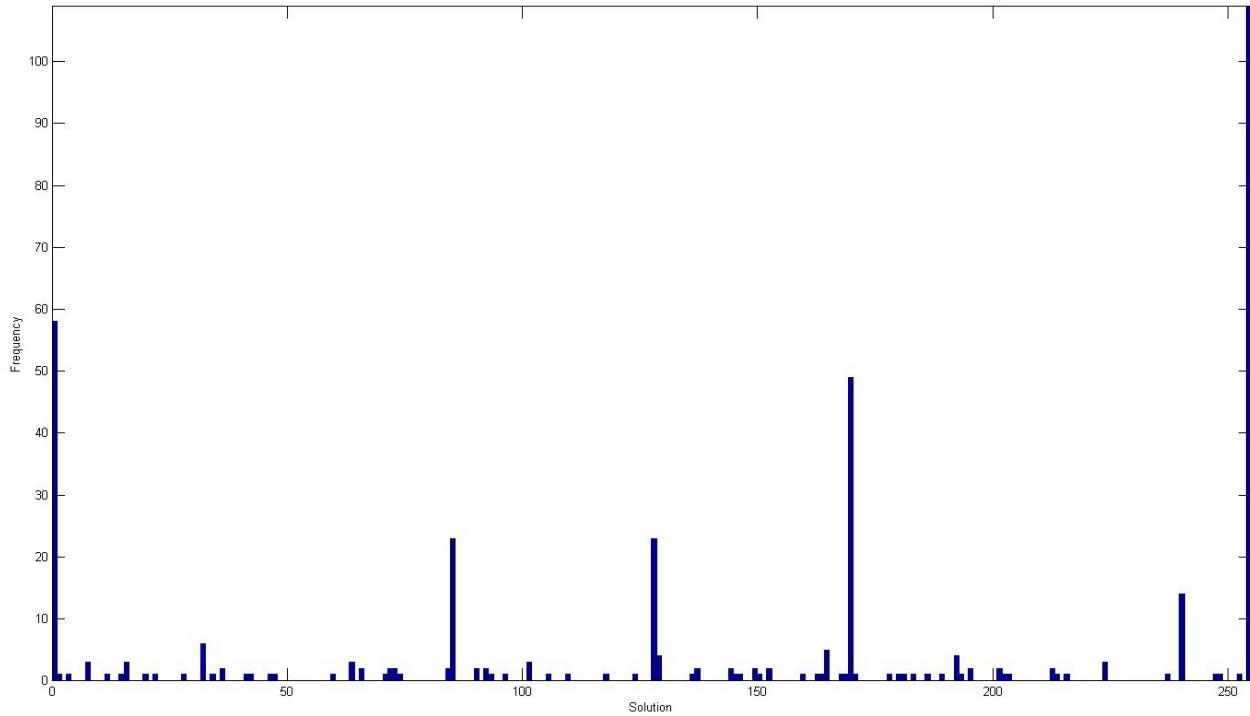


Figure 3. Solution distribution for the first step. Submissions were recorded as binary strings which were converted into decimal numbers.

Table 9. Frequency for the most used solutions in the first step of the game

	Frequency for most used solution: 90% of solutions are used less than:	107
11111111	107	
00000000	52	
10101010	49	
01010101	23	
10000000	23	
11110000	14	
00000001	6	
00100000	6	
10100101	5	

We further analyse submissions in the fifth step only to see the same pattern emerge again (Table 10 and Figure 4). As the tables and figures highlight, these priors we identified in the overall

distribution of solutions and confirmed by isolating the first step are still the most frequently used solutions. The case of the participant highlighted above (Table 7) is rather prototypical: players did not just start out with *one* prior solution (e.g. all 1s or all 0s), but rather with a set of possible solutions (mental maps) based on preconceived ideas about what solutions to such puzzles would usually look like.

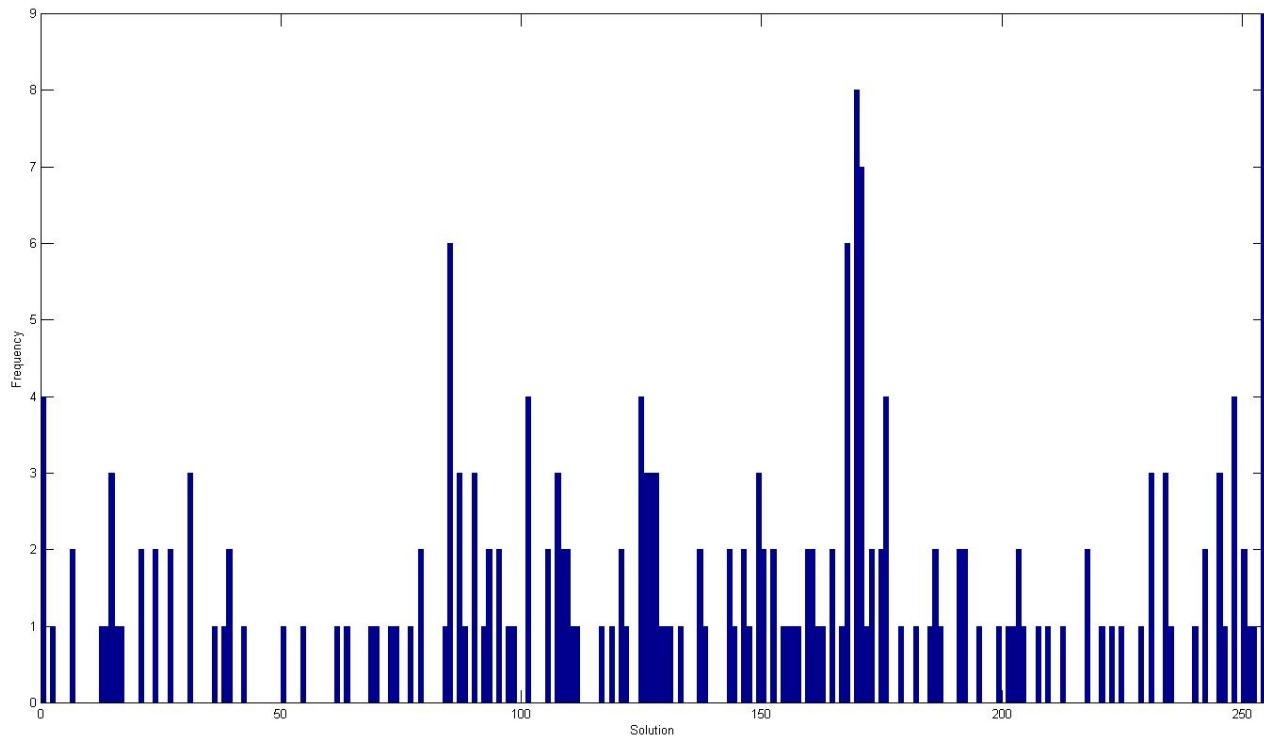


Figure 4. Solution distribution for step five. Submissions were recorded as binary strings which were converted into decimal numbers.

Table 10. Frequency for the most used solutions in the fifth step of the game

11111111	9	Frequency for most used solution: 95% of solutions are used less than:	9
10101010	8		4
10101011	7		
01010101	6		
10101000	6		
01100110	4		
01111101	4		
10110000	4		
11111000	4		

4.3. Analysis of qualitative interviews

4.3.1. Validating search strategies

The following section attempts to shed further light on the actual behavioural process (cf. Csaszar and Levinthal 2015) relying on interviews of participants.

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). Instead, they generally address systematic, random and local search. What interviewees characterize as systematic search could be local search, i.e. changing one bit at a time, but also distant search, e.g. symmetrically changing similar bits in both halves of the eight bits (i.e. mirroring). Interestingly, what respondents characterize as random search fairly seldom was random; comparing interview responses with actual game-play behaviour reveals that random search usually would either be the priors identified above: players reported reverting to these patterns when other approaches did not yield the anticipated payoffs: e.g. a participant reports he “randomly tap[ped] whatever”, but his actual game-play in that particular situation is a typical prior {0,1,0,1,0,1,0,1}. Another search behaviour they described as ‘random’ was randomly changing one bit at a time, what the literature would categorize as local search. 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...¹”.

4.3.2. Switching search strategies and mental representations

Most players talk about how they engage in different search strategies, and change strategy fairly frequently. This is contrary to the common assumption in the simulation literature where agents are generally assumed to search primarily locally (e.g. Lazer and Friedman 2007) and only with a low probability engage in long jumps. One interviewee remarks: “I tried to make all the things green [i.e. 0]...to see a pattern...then turn the first one [blue]...turn the second one [blue]...before that [strategy] I tried several random things”. The switching behaviour is mainly explained in

¹ During the game, solvers had access to their previous attempts via a virtual ‘clipboard’.

terms of their internal representations: “You test something, and then you find out if you are wrong, and then you have to adjust...”

While ‘the logic’ of the game has eluded some of the players, most of our respondents did form articulate mental models of the underlying problem. The interviews reveal qualitative differences in their nature. For instance, some players report testing the typical priors, while others have quite different mental models, e.g. “So each square is a letter, spelling: ‘I am smart’”. We also encounter more abstract representations, such as the solution is “a Fibonacci sequence” or a (correct) realization that the solution would have to be (inversely) symmetrical. 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 local-search only 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.

To sum up, in contrast to typical simulation models (Levinthal 1997; Lazer and Friedman 2007; Csaszar and Levinthal 2015) 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.

5. Discussion and implications for future research

This study contributes to theory on problem solving as search by focusing on the determinants of search behaviours at an individual level and studying how individuals search and adapt their search strategies. Based on empirical data we advance and refine a number of assumptions

regarding search behaviours, in order to go beyond the exploration-exploitation distinction (March 1991). We argue the latter better fits an armed bandit kind of scenario (Laureiro-Martínez et al. 2015) where one has to make a strategic decision to balance exploration and exploitation when selecting between a very limited number of options (Smaldino and Richerson 2012). Instead we investigate what individuals actually do, when one has decided to either ‘explore’ or ‘exploit’. We therefore examine a situation where an individual has to generate or shape new solutions for a complex problem relying on internal problem representations. In this way we provide an empirical foothold for organizational scholars studying the influence on problem solving performance of, for instance, mental representations (Gavetti and Levinthal 2000; Gavetti 2005; Csaszar and Levinthal 2015) or more particularly how problem representations should be modularized (Brusoni et al. 2007; Baumann and Siggelkow 2013). Our first major contribution is to explain how individual differences with respect to information processing, i.e. relying on the three-factor adaptors-innovators scale (Miron et al. 2004) influence the propensity with which solvers engage in different search strategies throughout the process. Our finding has implications for conceptual work relying on modelling and simulation since we show that a choice between local and non-local search is not determined solely by structural considerations (e.g. the nature of the problem to be solved or incentives) but also by a substantial variation in individuals’ preferred way of processing information. We provide an opportunity for simulation literature to base their modelling on heterogeneous agents, rather than simply assuming that heterogeneity stems from changes in the competitive environment that organizational agents search in (e.g. Levinthal 1997). Furthermore, we add to former types of individual level explanations such as the neurological basis (Laureiro-Martínez et al. 2015) or the influence of emotions (Døjbak et al. 2015) on strategic search behaviour. Relating our individual level to analysis of firm-based search (Lopez et al. 2016) could shed further light on how and why firms make strategic search decisions and supplement existing work on the influence of manager’s mental models on decision making (Kaplan and Tripsas 2008).

In addition, the above theoretical contribution also implies an opportunity for practice. Based on a relatively short survey (the A-I scale), owners of a crowdsourcing platform could map the cognitive styles of its current users and relate this mapping to the current problem to be solved. Following Felin and Zenger’s (2014) conceptualization of the problem-solving organization, the problem type and the degree of hidden (to the firm) knowledge required to solve it, can determine the search methods that should be employed. For example, they argue that user and innovation contests might be particularly adept at solving low complexity problems involving a high degree of hidden knowledge, since these problems require decentralized, trial and error (i.e. experiential) search. However, our study suggests that this might need further qualification: does one have the users likely to engage in the relevant kind of search? Since problems are not given

(Jeppesen and Lakhani 2010) and the seeker organization does have a certain degree of control on problem formulation (von Krogh et al 2012), one could formulate the problem differently, in order to decompose the problem and better match the relevant search behaviours and the problem to be solved.

A second major contribution is to provide a nuanced framework for making sense of individual search behaviours that is in line with recent cognitive science literature (e.g. Doll et al. 2012; Smaldino and Richerson 2012; Kaiser, Simon et al. 2013) and organizational literature (Rosenkopf 2008) and takes individual mental models into account. By relying on a simple task and clearly operationalized search variables we can go beyond a metaphorical conception of mental models (Porac and Thomas 2002) and we can investigate how individuals reflect on their mental models and their search strategies and engage in meaningful, distant searches.

Finally, we develop the current conceptualization of individual level model-based search (Doll et al. 2012). In particular we explore how internal representations are formed and used in conjunction with model-free moves. Our analysis of the initial submissions (i.e. the priors) indicates that players rely on top-down processing, but the interviews also reveal bottom-up processing (Anderson 2015), i.e. internal representations are shaped based cultural expectations regarding such puzzles, as proposed by Smaldino and Richerson (2012), but also based on feedback. Having a top-bottom internal representation is important because it suggests a) an alternative source of heterogeneity, since even with similar cognitive styles, solvers have different starting points and thus different paths (Solway and Botvinick 2012), and b) another departure from how simulation literature has conceptualized search so far, since most simulations start at a random position or tabula rasa. However, we also find support for players consistently engaging in model-free search (i.e. search behaviours that are not derived from their current problem representation) and complement the literature on mental maps (e.g. Kaplan and Tripsas 2008; Gavetti and Rivkin 2007), which has largely focused on model-based search. Bringing together both streams of literature, in this study, we emphasise the importance of taking into account the actual search process when studying strategy making. Our experimental set-up where the task is well defined also lays the ground for more informed evolutionary models that could, via virtual experiments, explore the complementary fit (Carnabuci and Dioszegi 2015) between the type of problem to be solved and the type of potential solvers (and subsequently their dominant search strategy) as well as their starting points.

Implications for future research

Our work identifies several research directions that can further our understanding of individual problem solving. First, one interesting finding not fully developed in this study is that solvers seem to rely on different evidence integration mechanisms (Solway and Botvinick 2015) than

previously assumed. In particular, the organizational literature assumes that solvers will pay attention to their best solution so far (e.g. Levinthal 1997; Lazer and Friedman 2007; Billinger et al. 2013), but the qualitative interviews suggest that participants pay attention to both their most recent and their best submission, but also to how the acquired information fares with respect to their internal representations. We speculate that dependent on cognitive styles different types of information will be sampled and integrated, an hypothesis that can be tested within an eye-tracking setup (Salvucci et al. 2000).

Second, this study examines cognitive styles, but there are a number of cognitive ability measures that could plausibly determine how individuals perceive and integrate information in a problem solving context. In particular, working memory (Kotovsky et al. 1985), attention (Sweller 1988) or cognitive control (Sweller 1983) have been connected with performance in various cognitive tasks and assessing their influence onto particular search behaviours might further explain the heterogeneity reported.

Overall, we argue that by relying on a task with a known and well defined interdependence structure allows us to get a step closer to actually being able to capture how individuals shape their search by relying on cognitive maps. Thus, we bring forward a research approach (i.e. game and setup) that is uniquely suited for and easily adaptable to futures studies of individual and/or collective problem solving.

References

- Allinson, C., and J. Hayes (1996). "The cognitive style index." *Journal of Management Studies* 33(1): 119-135.
- Afuah, A., & Tucci, C. L. (2012). "Crowdsourcing as a solution to distant search". *Academy of Management Review*, 37(3), 355-375.
- Ambrosini, V., C. Bowman, and N. Collier (2009). "Dynamic capabilities: an exploration of how firms renew their resource base." *British Journal of Management* 20(1): 9-24.
- Anderson, J.R. (2015). *Cognitive psychology and its implications*. New York: Worth Publishers.
- Antonietti, A. and M. A. Giolletta (1995). "Individual differences in analogical problem solving." *Personality and Individual Differences* 18(5): 611-619.
- Audia PG, Goncalo JA. (2007) "Past success and creativity over time: A study of inventors in the hard disk drive industry". *Management Science*. 2007 Jan;53(1):1-5.
- Auh, Seigyoung, and Bulent Menguc (2005). "Balancing exploration and exploitation: The moderating role of competitive intensity." *Journal of Business Research* 58.12 (2005): 1652-1661.
- Baumann, O, and N. Siggelkow (2013). "Dealing with complexity: integrated vs. chunky search processes." *Organization Science* 24(1): 116-132.

- Barsalou, L. W. (2008). "Grounded cognition." *Annu. Rev. Psychol.* 59: 617-645.
- Berinsky, A. J., G A. Huber, and G. S. Lenz (2012). "Evaluating online labor markets for experimental research: Amazon. com's Mechanical Turk." *Political Analysis* 20(3): 351-368.
- Billinger, S. and N. Stieglitz and T. R. Schumacher (2013). "Search on Rugged Landscapes: An Experimental Study" in *Organization Science* 25(1): 93-108.
- Bingham, C. B., & Kahl, S. J. (2013). "The process of schema emergence: Assimilation, deconstruction, unitization and the plurality of analogies". *Academy of Management Journal*, 56(1), 14-34.
- Browne, M. W., and R. Cudeck (1992). "Alternative ways of assessing model fit." *Sociological Methods & Research* 21(2): 230-258.
- Brusoni, S., L. Marengo, A. Prencipe and M. Valente (2007). "The value and costs of modularity: a problem-solving perspective." *European Management Review* 4(2): 121-132.
- Carnabuci, G., & Diószegi, B. (2015). "Social networks, cognitive style and innovative performance: A contingency perspective". *Academy of Management Journal*, amj-2013.
- Chen, W. R., & Miller, K. D. (2007). "Situational and institutional determinants of firms' R&D search intensity". *Strategic Management Journal*, 28(4), 369-381.
- Choi, J. N. (2004). "Individual and contextual predictors of creative performance: The mediating role of psychological processes." *Creativity Research Journal* 16(2-3): 187-199.
- Cornelissen, J. P., & Werner, M. D. (2014). "Putting framing in perspective: A review of framing and frame analysis across the management and organizational literature". *The Academy of Management Annals*, 8(1), 181-235.
- Csaszar, F. A. and D. A. Levinthal (2015), Mental representation and the discovery of new strategies. *Strategic Management Journal*. Early view: doi: 10.1002/smj.2440
- Daw, N. D. et al. (2011) "Model-based influences on humans' choices and striatal prediction errors." *Neuron* 69(6): 1204-1215.
- Dayan, P. (2012). "How to set the switches on this thing." *Current opinion in neurobiology* 22(6): 1068-1074.
- Denrell, J. (2007). "Adaptive learning and risk taking." *Psychological Review* 114(1): 177-187.
- Djamshidian, A, et al. (2011). "Novelty seeking behaviour in Parkinson's disease." *Neuropsychologia* 49(9): 2483-2488.
- Doll, B. B., A. Dylan, A. Simon, and N. D. Daw (2012) "The ubiquity of model-based reinforcement learning." *Current opinion in neurobiology* 22(6): 1075-1081.

- Døjbak, D. H., J. K. Eskildsen, et al. (2015). "Exploration versus exploitation: Emotions and performance as antecedents and consequences of team decisions." Strategic Management Journal. Early view: 10.1002/smj.2380.
- Eggers, J. P., & Kaplan, S. (2013). "Cognition and capabilities: a multi-level perspective". The Academy of Management Annals, 7(1), 295-340.
- Ethiraj, S. K., & Levinthal, D. (2004). "Modularity and innovation in complex systems". Management Science, 50(2), 159-173.
- Felin, T. and N. J. Foss (2005): "Strategic Organization: A Field in Search of Microfoundations". Strategic Organization 3(4): 441-455.
- Felin, T., & Zenger, T. R. (2014). "Closed or open innovation? Problem solving and the governance choice." Research Policy, 43(5), 914-925.
- Foss, N. J., & Pedersen, T. (2014). "Microfoundations in strategy research". Strategic Management Journal. Early view: 10.1002/smj.2362
- Flynn, B. B., Schroeder, R. G., & Sakakibara, S. (1994). "A framework for quality management research and an associated measurement instrument". Journal of Operations Management, 11(4), 339-366.
- Gavetti, G., and D.A. Levinthal (2000). "Looking forward and looking backward: Cognitive and experiential search." Administrative Science Quarterly 45(1): 113-137.
- Gavetti, G., D. A. Daniel Levinthal, and J. W. Rivkin (2005). "Strategy-making in novel and complex worlds: The power of analogy". Strategic Management Journal 26: 691-712.
- Gavetti, G., & Rivkin, J. W. (2007). "On the origin of strategy: Action and cognition over time". Organization Science, 18(3), 420-439.
- Gary, M. S., & Wood, R. E. (2011). "Mental models, decision rules, and performance heterogeneity". Strategic Management Journal, 32(6), 569-594.
- Gino, F. and D. Ariely (2012). "The dark side of creativity: original thinkers can be more dishonest." Journal of Personality and Social Psychology 102(3): 445-459.
- Gläscher, J, et al. (2010). "States versus rewards: dissociable neural prediction error signals underlying model-based and model-free reinforcement learning." Neuron 66(4): 585-595.
- Gough, Harrison G. "A creative personality scale for the Adjective Check List." Journal of personality and social psychology 37.8 (1979): 1398.
- Greene, Jennifer C., Valerie J. Caracelli, and Wendy F. Graham. 1989. "Toward a Conceptual Framework for Mixed-Method Evaluation Designs." Educational Evaluation and Policy Analysis 11 (3):255-274.
- Greve, H. R. (2003). "Organizational learning from performance feedback: A behavioral perspective on innovation and change". Cambridge University Press.

- Gupta, A. K., K. G. Smith, C. E. Shalley (2006): "The interplay between exploration and exploitation". Academy of Management Journal 49(4): 693-706.
- Hanna, R., S. Mullainathan, et al. (2014). "Learning through noticing: Theory and evidence from a field experiment." The Quarterly Journal of Economics 129(3): 1311-1353.
- 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.
- Hills, T. T., & Hertwig, R. (2010). "Information Search in Decisions From Experience Do Our Patterns of Sampling Foreshadow Our Decisions?". Psychological Science.
- Huys, Q. J., et al. (2012) "Bonsai trees in your head: how the Pavlovian system sculpts goal-directed choices by pruning decision trees." PLoS Comput. Biol 8(3): e1002410.
- Jablokow, K. W., et al. (2015) "Cognitive Style and Concept Mapping Performance." Journal of Engineering Education 104(3): 303-325.
- Johnson, A. and D. A. Crowe (2009). "Revisiting Tolman, his theories and cognitive maps." Cogn Crit 1: 43-72.
- Jones, M., and F. Canas (2010). "Integrating reinforcement learning with models of representation learning." Proceedings of the 32nd Annual Conference of the Cognitive Science Society.
- Kahneman, D. and A. Tversky (1979). "Prospect theory: An analysis of decision under risk." Econometrica: Journal of the Econometric Society 47(2): 263-292.
- Kaiser, S., J. J. Simon, et al. (2013). "The cognitive and neural basis of option generation and subsequent choice." Cognitive, Affective, & Behavioural Neuroscience 13(4): 814-829.
- Kaplan, S., & Tripsas, M. (2008). "Thinking about technology: Applying a cognitive lens to technical change". Research Policy, 37(5), 790-805.
- Stuart A. Kauffman. (1993). "The origins of order: Self organization and selection in evolution". Oxford university press.
- Kirton, M. (1976). "Adaptors and innovators: A description and measure." Journal of Applied Psychology 61(5): 622.
- Kotovsky, K., J. R. Hayes, and H.A. Simon (1985). "Why are some problems hard? Evidence from Tower of Hanoi." Cognitive Psychology 17(2): 248-294.
- Kozlowski, S. W. J. and G. T. Chao (2012). "The Dynamics of Emergence: Cognition and Cohesion in Work Teams." Managerial and Decision Economics 33(5-6): 335-354.
- von Krogh, G., Wallin, M., & Sieg, J. H. (2012). "A problem in becoming: How firms formulate sharable problems for innovation contests". In Research Policy Special Issue Conference "Open Innovation: New Insights and Evidence", London, UK.

- Lant, T. K., & Shapira, Z. (2001). "Introduction: Foundations of research on cognition in organizations". Organizational Cognition. Mahwah, NJ: Lawrence Erlbaum Associates, 1-14.
- Lakhani, K. R., Boudreau, K. J., Loh, P. R., Backstrom, L., Baldwin, C., Lonstein, E., Lydon, M., MacCormack, A., Arnaout, R.A. and Guinan, E. C. (2013). "Prize-based contests can provide solutions to computational biology problems. Nature biotechnology, 31(2), 108-111.
- Laureiro-Martínez, D., et al. (2014) "Frontopolar cortex and decision-making efficiency: comparing brain activity of experts with different professional background during an exploration-exploitation task." Frontiers in human neuroscience 7: 927.
- Laureiro-Martínez, D., S. Brusoni, et al. (2015). "Understanding the exploration-exploitation dilemma: An fMRI study of attention control and decision-making performance." Strategic Management Journal 36(3): 319-338.
- Lavie, D., Stettner, U., & Tushman, M. L. (2010). "Exploration and exploitation within and across organizations". The Academy of Management Annals, 4(1), 109-155.
- Lazer, D. and A. Friedman (2007). "The network structure of exploration and exploitation." Administrative Science Quarterly 52(4): 667-694.
- Levinthal, D. A. (1997): "Adaption on Rugged Landscapes". Management Science 43(7): 934-950.
- Lopez-Vega, H., F. Tell, et al. (2016). "Where and how to search? Search paths in open innovation." Research Policy 45(1): 125-136
- March, J. G. (1991). "Exploration and exploitation in organizational learning." Organization Science 2(1): 71-87.
- Martignoni, D., Menon, A., & Siggelkow, N. (2015). "Consequences of Misspecified Mental Models: Contrasting Effects and the Role of Cognitive Fit". Strategic Management Journal.
- Mason, W, and D. J. Watts (2012) "Collaborative learning in networks." PNAS 109(3): 764-769.
- Mason, W.. Winter, and Suri (2012). "Conducting behavioural research on Amazon's Mechanical Turk." Behaviour Research Methods 44(1): 1-23.
- Mckelvey, B., M. Li, et al. (2013). "Re-thinking Kauffman's NK fitness landscape: From artifact & groupthink to weak-tie effects." Human Systems Management 32: 17-42.
- Miles, M.B., and M.A. Huberman. (1994). "Qualitative Data Analysis". Vol. 2. ed. Thousand Oaks, CA: Sage.
- Miron-Spektor, E., M. Erez, et al. (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., M. Erez, et al. (2004). "Do personal characteristics and cultural values that promote innovation, quality, and efficiency compete or complement each other?" Journal of Organizational Behaviour 25(2): 175-199.
- Narayanan, V. K., Zane, L. J., & Kemmerer, B. (2011). The cognitive perspective in strategy: An integrative review. Journal of Management, 37(1), 305-351.
- Newell, A., & Simon, H. A. (1972). "Human problem solving" (Vol. 104, No. 9). Englewood Cliffs, NJ: Prentice-Hall.
- Nickerson, J. A., and T. R. Zenger (2004). "A knowledge-based theory of the firm—The problem-solving perspective." Organization Science 15(6): 617-632.
- Parter, M., Kashtan, N., & Alon, U. (2008). "Facilitated variation: how evolution learns from past environments to generalize to new environments". PLoS Comput Biol, 4(11), e1000206.
- Pfleiderer, P (2014). "Chameleons: The Misuse of Theoretical Models in Finance and Economics", Stanford University Working paper
- Porac, J. F., & Thomas, H. (2002). "Managing cognition and strategy: issues, trends and future directions". Handbook of Strategy and Management, 165-181.
- Puranam, P., Stieglitz, N., Osman, M., & Pillutla, M. M. (2015). "Modelling bounded rationality in organizations: Progress and prospects". The Academy of Management Annals, 9(1), 337-392.
- Raisch, Sebastian, and Julian Birkinshaw.(2008) "Organizational ambidexterity: Antecedents, outcomes, and moderators." Journal of management
- Redish, D. A. (2012). "Search processes and hippocampus." in Cognitive search: Evolution, algorithms, and the brain, p: 81-95.
- Rivkin, J. W. (2000). "Imitation of complex strategies". Management science, 46(6), 824-844.
- Rosenkopf, L. (2008), "Enriching Research on Exploration and Exploitation", in Barry and Hansen (eds.), The SAGE Handbook of New Approaches in Management and Organization. SAGE Publications: California.
- Salvucci, D. D., & Goldberg, J. H. (2000). "Identifying fixations and saccades in eye-tracking protocols". In Proceedings of the 2000 symposium on Eye tracking research & applications (pp. 71-78). ACM.
- Schacter, D. L., D.R. Addis, and R. L. Buckner (2007). "Remembering the past to imagine the future: the prospective brain." Nature Reviews Neuroscience 8(9): 657-661.
- Schwartzstein, J. (2014). "Selective attention and learning." Journal of the European Economic Association 12(6): 1423-1452.
- Schultz, W., P. Dayan, and P. R. Montague (1997). "A neural substrate of prediction and reward." Science 275(5306): 1593-1599.

- Shalley, C. E., Zhou, J., & Oldham, G. R. (2004). "The effects of personal and contextual characteristics on creativity: Where should we go from here?." *Journal of management*, 30(6), 933-958.
- Simon, H. A. (1977) "The structure of ill-structured problems." in *Models of discovery*. p: 304-325.
- Simon, H. A. (1983). "Search and reasoning in problem solving." *Artificial Intelligence* 21(1): 7-29.
- Smaldino, P. E. and P. J. Richerson (2012). "The origins of options." *Frontiers in neuroscience* 6(50):1-11.
- Solway, A. and M. M. Botvinick (2012). "Goal-directed decision making as probabilistic inference: a computational framework and potential neural correlates." *Psychological Review* 119(1): 120-154.
- Solway, A. and M. M. Botvinick (2015). "Evidence integration in model-based tree search." *PNAS* 112(37): 11708-11713.
- Sternberg, R. J., and Grigorenko, E. (1997). Are cognitive styles still in style? *American Psychologist* 52: 700–712.
- Sweller, J. (1983). "Control mechanisms in problem solving." *Memory & Cognition* 11(1): 32-40.
- Sweller, J. (1988). "Cognitive load during problem solving: Effects on learning." *Cognitive Science* 12(2): 257-285
- Taggar, S. (2002). "Individual creativity and group ability to utilize individual creative resources: A multilevel model." *Academy of Management Journal* 45(2): 315-330.
- Teece, D. J. (1996) "Firm organization, industrial structure, and technological innovation." *Journal of Economic Behaviour & Organization* 31(2): 193-224.
- Todd, P. M., T. T. Hills, and T. W. Robbins (2012). "Building a foundation for cognitive search." in *Cognitive search: Evolution, algorithms, and the brain* p: 1-7.
- Tolman, E. C (1948). "Cognitive maps in rats and men." *Psychological Review* 55(4): 189-208.
- Vuculescu, O. (2015) "Searching far away from the lamp-post: an agent based model of search", presented at *Organizing Crowds and Innovation Conference*, Oxford 2015
- Watson, R. A. and J. B. Pollack (1999). Hierarchically consistent test problems for genetic algorithms. Evolutionary Computation. CEC 99. *Proceedings of the 1999 Congress on IEEE*.
- Witkin, H. A., and D. R. Goodenough (1977). "Field dependence revisited." *ETS Research Bulletin Series* 1977(2): i-53.

Appendix 1: Multinomial modelling of dependent variable. Robustness check

In this formulation, the dependent variable has three possible outcomes corresponding to the three possible search behaviours: local, model-based and random search. We use the third strategy (random) as the reference category, because we feel the distinction between local and model-based on the one hand and random search on the other is 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 the interviews. We thus expect to find a significant difference between random search and the two other categories. The predictor variables are the same as in the marginal logistic regression model.

Ordered Value	Strategy	Total Frequency
Model-based	1	1749
Local search	2	4178
Random	3	1240

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

Table 1. Frequencies of the response variable categories. Pooled data.

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 equation 3 becomes:

$$\log \left\{ \frac{\pi_{ijr}}{\pi_{ij1}} \right\} = \beta_1 + \beta_2 creativity_i + \beta_3 conformity_i + \beta_4 detail_i + \beta_5 time + b_i + e_{ij}$$

$\pi_{ijr} = P(Y_{ij} = r)$ 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 β_i . The random effect b_i is assumed to have a univariate normal distribution with zero mean and compound symmetric covariance matrix.

	estimate	std err	p values	Odds Ratio Estimates	95% Confidence Limits
creativity (ref=LOW)					
Model-based (ref=Random)	0.3275	0.1098	0.0029	1.387	1.119 1.721
Local search (ref=Random)	0.3774	0.147	0.0103	1.458	1.093 1.945
<hr/>					
detail (ref=LOW)					
Model-based (ref=Random)	-0.3251	0.1184	0.0061	0.722	0.573 0.911
Local search (ref=Random)	-0.3197	0.1566	0.0412	0.726	0.534 0.987
<hr/>					
conformity (ref=LOW)					
Model-based (ref=Random)	-0.181	0.1076	0.0926	0.834	0.676 1.03
Local search (ref=Random)	-0.07779	0.144	0.5892	0.925	0.698 1.227

Table 2. Relative effect of cognitive style on the odds of choosing one of the three strategies.

Results from the multinomial regression. Pooled data.

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 random than local or model-based search (45% and 38%) and likewise, players who score high on 'attention to detail' will be more likely to engage in model-based or local search than random. 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.

References

Fitzmaurice, G. M., N. M. Laird, et al. (2012). Applied longitudinal analysis, John Wiley & Sons.