



City Research Online

City St George's, University of London

Citation: Banerjee, A., Lampel, J. & Bhalla, A. (2019). Two Cheers For Diversity: An Experimental Study Of Micro-level Heterogeneity In Problemistic Search. *Strategic Organization*, 17(4), pp. 450-469. doi: 10.1177/1476127018815467

This is the accepted version of the paper.

This version of the publication may differ from the final published version. To cite this item please consult the publisher's version.

Permanent repository link: <https://openaccess.city.ac.uk/id/eprint/21012/>

Link to published version: <https://doi.org/10.1177/1476127018815467>

Copyright and Reuse: Copyright and Moral Rights remain with the author(s) and/or copyright holders. Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge, unless otherwise indicated, provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way. For full details of reuse please refer to [City Research Online policy](#).

TWO CHEERS FOR DIVERSITY: AN EXPERIMENTAL STUDY OF MICRO- LEVEL HETEROGENEITY IN PROBLEMISTIC SEARCH

ABSTRACT

In this paper, we argue for an expanded view of problemistic search. Recent behavioral theory research suggests that individual search preferences influence problemistic search. We draw on this to challenge the view of problemistic search as a centrally directed organizational process that proceeds sequentially from local to distant search. We argue that search activities in organizations are heterogeneous – some individuals will first engage in local search while others may move directly to distant search. We propose that problemistic search at the macro-organizational level is therefore the result of a mix of local and distant search activities at the micro-level that shifts towards distant search in response to negative performance evaluation. We test this idea in a laboratory experiment using a repetitive task and performance feedback.

Key words: Performance reference points, problemistic search, lab experiment, behavioral theory

INTRODUCTION

The search for solutions when organizations are performing poorly is central to our understanding of strategic change. Current views of strategic search behavior of organizations are strongly influenced by the behavioral theory of the firm as originally formulated by Cyert and March (1963). Their view of search, subsequently incorporated into performance feedback theory (Greve, 2003), argues that strategic change is often triggered when decision makers judge organizational performance to be inadequate relative to competitors. When this occurs, decision makers engage in so called ‘problemistic search’ during which they seek solutions that will close the performance gap. This perspective sees problemistic search as a reorientation in the organization’s allocation of attention (Greve, 2013; Ocasio, 1997; 2011). Thus, as long as performance is satisfactory the organization allocates most of its attention to routine operations. When decision makers conclude that organizational performance relative to competitors is inadequate, attention shifts from routine operations to problem solving mode: first to local search, and then to distant search (Tyler and Caner, 2015).

For Kacperczyk, Beckman and Moliterno (2015) this received view of search is questionable because it presents a monolithic picture of organizations where individuals allocate attention in line with instructions from top managers who decide whether problemistic search is needed, and how it should be conducted. In this paper we argue that while this received view is intuitively appealing and theoretically parsimonious, it overlooks organizational research that suggests that individuals are far more varied in their search behavior – they may engage in search because they are intrigued by anomalies, see opportunities for new products, or seek to improve work processes (Knudsen and Levinthal, 2007; Marengo, 2015). Furthermore, whereas at the macro-level it may be useful to model organizations as following satisficing logic of first engaging in local search before moving to

distant search, the same may not be expected of all individuals that make up the organization. At the micro-level where individual behavior matters, we can expect search heterogeneity – some individuals may indeed follow the satisficing logic predicted by the behavioral theory of the firm and economize cognitive resources by searching locally first, while others will move directly to distant search because they are more creative, or more willing to take risks (March, 1981).

In this paper we therefore wish to argue for models of problemistic search that are based on search heterogeneity that incorporates simultaneous local and distant search (Schunk, 2009). Furthermore, we argue that the composition of search heterogeneity will change towards greater frequency of distant search when the individuals that make up the organization, and identify themselves as responsible for organizational performance, are exposed to a stimulus that suggests a greater performance shortfall. Put differently, we would expect organizational search behavior when individuals are focused on their task performance to be a mix of local and distant search; but when individuals are told that their organization's performance is inadequate relative to external peers, the mix of local and distant search will shift towards distant search. In a sense, we are arguing that substantive research that shows at the macro-level firms engage in local search first before moving to distant search in response to performance shortfall may be correct when researchers look only at top management decisions, e.g. laying off workers, or acquiring another company, but it does not accurately describe change that is the result of micro-macro interactions that involve bottom-up innovations (Felin, Foss and Ployhart, 2015).

To study search behavior that involves micro-macro interactions in organizations is empirically challenging. Kim, Finkelstein and Halebian (2015) suggest that laboratory-based work is a promising avenue for research that permits more precise testing of how teams of individuals respond to performance feedback. Their suggestion is in line with previous studies

that use experiments such as those by Lant (1992). This is the approach that we adopt in this paper. Our choice for the production task in this experiment is the so called “number reduction task” (NRT) first pioneered by Thurstone and Thurstone (1941), and later adapted by cognitive psychologists (Novick and Sherman, 2003). The number reduction task requires participants to find the digits of a string of numbers using two simple ‘rules’. This way of completing an NRT is essentially a simple repetitive exercise in which the participant must apply the rules and get an answer. However, what the participants do not know is that there is a hidden rule that would allow them to increase their productivity by several orders of magnitude. In the experimental design, we divide participants into two groups: control and experimental, where both groups can set an aspiration i.e. the number of NRTs they aim to complete and monitor their production performance. But whereas the control group only has their own performance information over multiple rounds, the experimental group is informed that their performance is “poor” relative to other teams, regardless of their production performance. Our results show two important findings: First, 28 percent of the teams in the control group find the hidden rule. This indicates that even without any external performance reference points, some individuals in 28 percent of the teams engaged in distant search and found the hidden rule. Second, 87 percent of the teams in the experimental group find the hidden rule. This supports our main argument that organizational search is a heterogeneous combination of local and distant search activities, even without any external triggering of problemistic search; and that there is a shift towards more distant search activities when external stimulus in the form of performance shortfall information is provided.

Our paper makes two contributions. First, we seek to expand the current model of problemistic search by arguing that the deterministic two stage model of problemistic search that is the standard assumption in behavioral theory should be supplanted with a problemistic search model that sees the transition between local and distant search as a matter of degree.

Second, we illustrate the need for more multi-level organizational research by showing how heterogeneity in micro-level actions translate into macro-level organizational behavior that cannot be fully explained with only macro-level analysis (Kacperczyk et al., 2015; Klein and Kozlowski, 2000a; Moliterno, Beck and Beckman, 2014).

THEORETICAL OVERVIEW

Problemistic search: Performance feedback theory distinguishes between two types of problemistic search. The first is local search; described by Kacperczyk et al. (2015, p. 233) as a process that gives rise to “relatively simple, short-run, and incremental reactions, with search often taking place in the neighborhood of the problem”. The second is distant search which, unlike local search, calls for “discovery of new solutions” (Gavetti et al., 2012, p. 11), or as Tyler and Caner (2015) and Baum and Dahlin (2007) put it, *look for new ways of doing things*. This theory sees local and distant searches as operating in a sequence. When managers decide that inadequate performance poses a threat to the organization they first initiate local search. If local search does not result in a satisfactory solution to the problem the organization moves to distant search. Barnett and McKendrick (2004, p. 540) sum up this sequence as follows: “... search is assumed to remain “local,” restricted to solutions that are only incrementally different from current practice, and only moving to more distant possible solutions when no satisfactory local solutions are found”.

The assumption that distant search only occurs if local search fails to deliver a satisfactory solution is consistent with conceptualizing “the organization as a monolithic entity” (Kacperczyk et al. 2015, p. 252). This has two important advantages. First, it allows researchers to speak of the organization as a unified behavioral actor that is centrally governed without having to deal with variations that may arise from internal diversity. Second, adopting a model of the organization as a unified behavioral actor means that

researchers can employ satisficing logic to predict how organizations respond to problems. A satisficing logic perspective predicts that organizations would prefer local search first because it is less resource intensive and less risky, and only move to distant search when they are obliged by lack of effective solutions to expend more resources and take more risks (Laursen, 2012).

Influence of performance reference points: Using a unified actor model to explain organizational response to performance has proven to be empirically fruitful, but this view is increasingly being questioned by researchers. In particular, the recent literature on performance feedback theory highlights a number of issues that argue against conceptualizing the organizational search as a product of centralized decision making. A crucial issue is the process that organizations use to arrive at the targets against which they later evaluate their performance. Until recently, the accepted view in performance feedback theory was that top decision makers develop targets by combining internal and external reference points into single targets. Specifically, top decision makers combine an internal reference point based on the organization's own past performance with an external reference point based on the performance of peer organizations, usually competitors (Bromiley, 1991; Greve, 2003; Lampel and Giachetti, 2010; Wiseman and Bromiley, 1996).

Some scholars have questioned this assumption, arguing that empirical research points to organizations retaining separate targets based on internal and external reference points (Baum et al., 2005; Blettner et al., 2015). A review of the performance feedback literature by Joseph and Gaba (2014) that compares studies that assume that organizations combine internal and external reference points into single targets, with studies that maintain that organizations keep reference points separate when evaluating performance, concludes that researchers often adopt different assumptions about how managers interpret performance. Washburn and Bromiley (2012) agree; they point out that contrasting assumptions about whether internal and

external reference points should be combined or operationalized separately reflect different assumptions about managerial cognition. Kim et al. (2015, p. 1364) also argue that the use of internal vs. external reference points are “associated with different underlying cognitive and organizational processes”. But they go one step further and maintain that because reference points derive from “distinct sources of performance feedback and are filtered through different cognitive and organizational processes, they may engender different interpretations, which, in turn, may induce different organizational responses” (Kim et al., 2015, p. 1364).

Search heterogeneity and individual preferences: If cognitive processes influence how organizations engage in problemistic search, this would suggest that individual cognitive processes, specifically the way they allocate attention, is also influenced by factors that are often personal in nature (see Ocasio, 2011). Recent studies seem to point in this direction. In their study of the US mutual funds industry Kacperczyk et al. (2015) find that fund managers are likely to react differently to shortfall relative to internal reference points based on comparison with the performance of other fund managers in the same organization, than to external reference points based on stock market performance of other funds. The authors argue that in the first instance the fund managers’ behavior is consistent with prospect theory; specifically, they exhibit loss aversion, whereas in the second they behave in a manner that is consistent with the behavioral theory of the firm i.e. they would aim to satisfice rather than maximize their outcomes (Gavetti, Greve, Levinthal & Ocasio, 2012). Their attention to performance feedback will therefore vary according to whether the information poses risks to their career prospects as opposed to the risks it poses to the organization as a whole.

Joseph and Gaba (2014) provide additional evidence for the proposition that managers respond differently to reference points. They attribute the difference more generally to limited attention resources that force managers to disproportionately allocate more attention to one reference point and less to another. But this allocation may change when circumstances

change, leading managers to reallocate attention. Joseph and Gaba (2014) note that under normal circumstances managers pay attention to historical performance feedback that tracks the organization's past performance. When confronted with a challenge from the environment, however, managers will dramatically increase the amount of attention allocated to social reference points that tracks how rivals are performing. Their conclusion concurs with research by Washburn and Bromiley (2012) and Bromiley and Harris (2014), who echo Ocasio's (1997; 2011) observation that limited attentional capacity implies selective attention to reference points.

To recapitulate, the dominant view of performance feedback theory is that search begins when top managers instruct subordinate organizations to direct their attention to problem solving, and proceeds from local to distant search according to a pattern dictated by economizing of resources, in particular attention resources (Laursen, 2012). However, as Li et al. (2013, p. 894) point out that, "although organizational systems, incentives, and processes can be designed to encourage managers to search, it is a manager and not an organization that is capable of searching." To accept that search often comes down to individual effort is also to accept that managers will vary in their search styles and willingness to explore: the majority of managers may respond to top management instructions by undertaking less risky and less resource costly local search, but there are others that will engage in distant search because they are motivated to go further depending on their own motivation and capabilities.

Prior research into characteristics that influence individuals' search behavior shows that indeed there is significant heterogeneity in individuals' propensity to engage in search activities. The differences in search may be due to demographic factors such as age, educational, and functional knowledge (Bantel and Jackson, 1989; Hambrick and Mason, 1984; Khan and Manopichetwattana, 1989), psychometric factors such as cognitive ability (Steyvers et al. 2009), neurological variances (Badre et al. 2012; Healey and Hodgkinson,

2014), as well as differences based on individual tendencies to seek or avoid novelty (Payzan-LeNestour, 2012). This suggests that while majority of individuals are likely to follow macro-level directives to engage in local or distant search, their search behavior is also influenced by idiosyncratic characteristics and therefore they may engage in distant search even when it is not required or mandated by the organization.

Influence of individuals' behavior on macro-level outcomes: Acknowledging that individuals can vary greatly in capacities and motivations leads us to consider how behavior heterogeneity at micro-level can be reconciled with macro-level theorization of organizational search (Felin and Foss, 2005). As Klein and Kozlowski (2000b) and Ocasio (2011) point out, while a top down approach that explicitly looks at the influence of macro-levels constructs such as firm performance on individual actions has been extensively studied, the bottom-up approach is relatively neglected. Our study takes a bottom-up approach to the macro-level construct of organizational search by aggregating the search activities of micro-level actors. Such a bottom-up approach to theorizing search allows us to understand organizational level search as a combination of both local and distant search, and to argue that as more individuals at the micro-level are incentivized by top management to engage in distant search, the mix of local and distant problemistic search at the macro-level gets more skewed towards distant problemistic search. Figure 1 presents a conceptual summary of how a bottom-up aggregation of micro-level problemistic search behavior allows us to model organizational search at the macro-level as a combination of simultaneous local and distant problemistic search.

Insert figure 1 here

Testing search heterogeneity using performance reference points: Ocasio (2011) points out the contrast between research that examines micro-level attention processes and macro-level

theories that make excessively simplifying assumptions about how micro-level behavior is shaped by macro-level instructions. In this paper we argue that contrary to the simplifying assumptions of top-down search model that ignores heterogeneity – both local and distant search normally take place at the micro level. The observed search behavior at the macro-level is therefore the result of macro-level instructions influencing the balance in the ecology of search further down the organization.

To test our argument that outcomes at the macro organizational level are the result of a shift the mix between local and distant search, we build on the observation that organizations respond differently to internal and external performance reference points (Kacperczyk et al., 2015). When organizations rely on internal reference points to judge performance we are likely to see efforts to keep up with the same level of past performance, or at most, efforts to improve performance incrementally. In terms of the search ecology at the micro level this means that local search will predominate, but crucially for our argument we would also expect some micro-level actors to engage in distant search for reasons that are peculiar to these individuals. However, when external reference information is introduced that suggests a significant performance gap with competitors we would expect some micro-level actors that up to this point searched locally to switch to distant search, thereby changing the balance at the macro-level from local to distant search.

EXPERIMENTAL DESIGN

In order to test our argument, we have adopted a theory-driven experimental design. Such a design is appropriate because our aim is to isolate the impact of internal and external performance reference points on performance improvement due to local and distant search.

This calls for strict control of the input: performance reference points, and measurement of the output: performance improvement due to local and distant search. In the real world, it is difficult if not impossible to manipulate the inputs and compare outputs for various reasons –

most notably because tasks vary, and measures of performance are often complicated. In an experimental setup we can manipulate the performance reference points, by controlling available information; and furthermore, we have the ability of choosing the same repeated activity for everyone that crucially allows us to clearly categorize performance outcome due to local and distant search. Thus, unlike real world setting where performance may be due to multiple factors, an experimental setting simplifies the causal processes. As Jackson and Cox (2013, p. 38) put it “the artificial environment of the laboratory is embraced in order to fully isolate causal processes of interest and thus to test elements of scientific theories”.

Theoretically, first we have proposed that while macro-level units display search heterogeneity, a reliance on only internal reference points leads to predominantly incremental improvements through local search. Crucially, however, we also maintain that there will still be performance improvement due to distant search. In the context of the experimental design this leads us to hypothesize that:

Hypothesis 1: Over time, there will be an improvement in collective performance due to distant search outcomes, even when individuals rely on only internal performance reference points.

While one can point to instances of organizations that are only guided by internal reference points – for example public organizations that set their own benchmarks; realistically most organizations operate in environments where performance is also benchmarked to competing external organizations. To take account of this context, we therefore introduce an intervention that brings to the team’s attention their performance relative to external reference point. In organizational terms, this intervention usually takes place when individuals become aware of a marked disparity between their collective

performance and that of competitors. This in turn directs search for solutions that will close the performance gap relative to external benchmarks. This gives us the following hypothesis:

Hypothesis 2: An improvement in collective performance due to distant search outcomes is more likely when individuals are informed of significant performance short fall relative to external reference points.

To test these hypotheses, our experimental design had to address several constraints. First, we had to design an experiment where problemistic search arose out of performance evaluation rather than presenting participants with a ‘problem’. We tackled this constraint by creating a simulated environment in which participants that work in teams are were asked to perform a simple and repetitive production task under time constraints, with a clear performance outcome. The key point here is that the task itself does not present a problem, but setting goals that require performance evaluation would, under certain circumstances, trigger problemistic search. Second, because cognition is difficult to observe we rely on the output of problemistic search to capture local vs. distant search. To increase measurement reliability, it was important to assign a task whose outputs from local and distant activities falls into sharply contrasting and mutually exclusive categories: one that points to local search and the other to distant search. Our choice for the production task in this experiment is the so called “number reduction task (NRT)” first pioneered by Thurstone and Thurstone (1941), and later adapted by cognitive psychologists to study ‘insight’ problem solving, or ‘pop up’ solutions (Novick and Sherman,, 2003). Novick and Sherman, (2003) point out that pop up solutions differ from conscious search solutions. Conscious search solutions are constructed incrementally, whereas pop up solutions “typically pop into mind suddenly and fully intact, accompanied by an irresistible feeling of aha!” (Novick and Sherman, 2003, p. 352). They

further note that conscious search effort is motivated by a goal, and thus involves active monitoring of progress. While pop up solutions, may solve the problem created when individuals attempt to meet a goal, but they are not produced by conscious effort. Notably, they strike the individual that comes up with the solution as novel and surprising, and thus their emergence can be seen as a manifestation of distant search in areas that are not considered initially (Novick and Bassok, 2005).

We designed the experiment with the team as our unit of analysis: we introduce treatments and observe outcomes of teams. However, while the participants are organized into teams, they performed the assigned task individually, with each team consisting of six individuals who are free to communicate with each other throughout the exercise but have no communication with the outside world except the information that is given to them. Using teams as our experimental unit of analysis reflects real world organizations where most tasks are performed by teams. Furthermore, by asking teams to report their cumulative performance we increased their motivation to set aspiration collectively and share the discovery of any insight resulting from individual idiosyncrasies. This also reflects real world organizations where teams are evaluated collectively, and information that improves team task performance is therefore shared. Our approach is in line with lab experiments involving organizational routines in repeated tasks that are often designed around teams of decision makers working towards the same performance objective (Bapuji et al. 2012; Cohen and Bacdayan, 1994; Wollersheim and Heimeriks, 2016). Furthermore, the use of teams in experiments have also been found to reduce errors from individual motivation issues in laboratory tasks (Mas and Moretti, 2009).

Task description: We use a Number Reduction Task (NRT) that is often used by researchers to study determinants of creative behavior in repeated tasks (For example see (Wagner *et al.*, 2004)). The basic idea behind the NRT is to find the ‘last digit’ of a sequence made up of the

digits the digits ‘1’, ‘4’, and ‘9’ as quickly as possible, using two simple rules. In each step of the NRT, the participant compares two digits and applies either the ‘same rule’ or the ‘different rule’. The same rule states the result of identical digits is the same digit (e.g. the result of comparing 1 and 1 gives 1, result of comparing 4 and 4 gives 4, and the result of comparing 9 and 9 gives 9), while the different rule states that the result of non-identical digits is the third digit (e.g. the result of comparing 1 and 4 gives 9, the result of comparing 1 and 9 gives 4, and so on). As the participant sequentially processes the NRT from left to right the seventh response provides the last digit. The instructions to the participants state that they are to indicate only the last digit on their response sheet. Task performance is measured by the number of NRTs correctly solved in a given time period.

What is not disclosed to the participants is that the NRTs were generated in such a way that the last three responses of every sequence always mirror the previous three responses. This mirroring means that in every NRT, the second response coincides with the last response. Figure 2 illustrates the two possible ways of solving the NRT (i) Using the two “official” rules that are communicated during the initial briefing and (ii) Using the two rules plus a third new rule based on the insight that second and last number of the new sequence are always the same. Using the third rule, participants can solve each NRT in three steps as opposed to the usual seven steps, thereby improving their productivity by at least two-folds.

Insert figure 2 here

Procedure: The experiment was run with participants at three executive education workshops. We followed the recommendations of designing extra-laboratory experiments to collect data from classroom participants (Charness et al. 2013; Loyd et al. 2005). The authors were the instructors for these workshops, and therefore organized and supervised the experimental

setup. Over the three workshop cohorts, there were 33 teams with six participants in each team i.e. 198 individuals participated in the exercise. Four teams had less than six participants; their responses have not been counted in this study. Participants' average age was 35 years, with an average work experience of 11 years. 76 percent of the participants were men and 24 percent women. All the three workshops were conducted in close succession within 6 weeks, at the same location, and followed exactly the same procedure with the same people involved in the procedure. We compared the results from the three workshops and found no significant difference.

The participants were randomly pre-assigned to 33 teams. Participants were informed that data collected during the exercise would be used in research, that anonymity of individual performance is guaranteed, and that they can decline to participate in the experiment if they so wish. The participants were given an initial briefing that included briefing for informed consent, an introduction to the NRT, and a practice round with five NRTs to ensure that everyone was familiar with the two rules that are used to solve the NRT. The participants were given a copy of these written instructions. Each team was informed at the outset that they would be given 70 new NRTs in each round. These would be provided on seven sheets of paper with 10 NRTs on each sheet. The teams were also instructed that they had to agree on a team target and an overall approach they will use to meet this target, and that their team's performance will be measured by the number of NRT sequences correctly solved in 60 seconds. The participants were also informed that there would be three rounds with the same conditions and that there will be a ten-minute break between rounds during which their performance would be calculated, that they had to remain in the room and use the time to reflect and discuss their task strategy. At the end of every round, after counting the number of correct responses, the NRT response sheets were given back to the teams. Participants were also instructed that they were not allowed to leave the room or communicate with the 'outside

world' by phone or internet. The teams were assigned separate rooms with a moderator present in each room.

Control and experimental treatments: To test the two hypotheses, we take advantage of a mixed design: within group and between group comparisons. Since our participants engage in this activity over three rounds, we have the opportunity to test hypothesis 1 by employing a within group comparison, using the teams from the control group that only have their own internal performance reference points. And to test hypothesis 2, we employ a between group comparison that enables us to compare performance of teams in the control group that only have information about their own performance, with teams in the experimental groups that have information about their own performance plus information about their performance relative to other teams.

This reflects assumptions about how naturally occurring organizations operate. Thus, organizations may have information about their own performance, and may or may not have information about how other organizations operate. But clearly it is rare, if not unrealistic, to expect organizations to have information about external organization but no information about their own performance. Translated to our own experimental design this means that information about own performance allows teams in both control and experimental groups to form internal reference points for evaluating performance, while information about performance relative to other teams allows the teams in the experimental group to form both internal and external reference points. The control group is not given any other information and set their internal reference point based on their own performance, but the experimental group is told after each round that their performance is “poor” relative to the performance of other teams as “most other teams had solved all 70 sequences”. This comparative information can be used by teams in the experimental group to form external reference point. The experimental group therefore has two pieces of information that it can use to set reference

points: their own performance, and the performance of peers. Figure 3 summarizes the overall design of the experiment.

Insert figure 3 here

FINDINGS

Observations to support hypothesis 1: To test the first hypothesis, we investigate the improvement in performance due to distant search within the control group. To begin with, figure 4 shows the rate at which teams in the two groups improved their performance. Specifically, for the teams in the control group, we find that performance in terms of task productivity rates improved over the three rounds, in line with ‘learning effects’ prediction (Epple et al. 1991) – albeit the average performance is far less than the performance of the experimental group. However, table 1 indicates that by the end of the third round, 5 of the 13 teams i.e. 28 percent of the control group show evidence of performance improvement due to distant search, even when they were not given any external performance reference points. As we had randomly assigned individuals to teams, and the treatment was also assigned randomly, the improvement in performance due to distant search in these teams of the control group is attributable to micro-level heterogeneity in propensity to engage in distant search. This shows that while the predominant state in the teams of the control group is to improve performance by local search, over time there is evidence of performance improvement from distant search, indicating that both local and distant search takes place – even when the teams rely on only internal performance reference points. This observation is also clear from table 2 and figure 5, which show the difference-in-difference effect between the control and experimental groups from round 1 and 3. We find that by the third round, the probability of a team in the control group improving performance by engaging in distant search is 28%.

Observations to support hypothesis 2: To test the second hypothesis, we observe the differences in performance due to distant search between the experimental groups and control groups. Table 1 shows the actual count of teams in the control and experimental groups that show evidence of improving performance by distant search. The disparity in the rates at which teams belonging to each group discover the hidden rule points towards the relationship between reference information and improvement in productivity due to distant search. We test whether the proportions for experimental and control group are statistically different using a chi-square test of independence. Furthermore, from the results of a difference-in-difference model (table 2 and figure 5), we observe that the probability of a team improving performance due to distant search increases by 28% even when there is no external reference point, while as a result of the introduction of the external reference point, the probability of a team improving performance due to distant search increases by an additional 59%.

Insert tables 1 & 2 and figures 4 & 5

DISCUSSION

Relevance for research on problemistic search: In this paper we call for an expanded view of problemistic search, as originally formulated by Cyert and March (1963), by challenging two assumptions about this construct that have come to dominate discussions of strategic search. First, we challenge the assumption that problemistic search proceeds in two distinct stages: first local and then distant. Second, we challenge the view of problemistic search as a macro construct that treats search as an organization wide activity under the direction of top managers.

Our alternative assumption is that problemistic search at the macro-level is a composition of search activities at the micro-level (Kozlowski and Klein, 2000). More specifically, we argue that the macro-view of problemistic search as originally formulated by Cyert and March's (1963) behavioral theory is based on a monolithic view of organizations, and therefore we should adopt a multi-level view that takes account of the heterogeneity of searchers with diverse capabilities and preferences at the micro-level. Individuals are likely to differ in their search strategies, i.e. some will stick to local search, while others will move directly to distant search, indicating a mix of local and distant search strategies at the macro-level.

Results from the experiment illustrate this argument and shows how the mix of local and distant search strategies skews towards distant search with negative external reference information. To begin with, 87 percent of the teams in the experimental groups found the hidden rule as opposed to only 28 percent in the control groups. This demonstrates that the mix of search strategies in teams that were repeatedly told that their performance was inferior compared to other teams. These teams are far more likely to shift to distant search than teams that only had information about their own performance. But it is also important to note that 28 percent of the teams in control groups also came up with the hidden rule, providing evidence for the presence of heterogeneous search strategies even when negative external reference information is not provided. This points to predominance of local search, as predicted by the search models originally postulated by behavioral theory, but also to the presence of distant search. A mix of search strategies is present in both experimental and control groups, but the mix may shift dramatically if environmental conditions change.

The post-exercise debriefing with all the participants reveal some patterns of individual search behavior within the teams that is useful in developing an understanding of the differences in search strategies. Teams that were given the treatment i.e. a high external

performance reference point, indicated that individuals in these teams were all engaged in “looking for a radically different way to do it”. They reported a sense of collective realization that “it’s simply not possible to do it [solve all 70 NRT sequences] ... if someone [another team] is doing it then they are either cheating or there is a pattern in the numbers”. This realization focused their effort towards distant search that led to the discovery of the hidden rule used to solve all the sequences in the next round. As a participant observed “we cracked the code in five minutes [once they were looking for the pattern] ... but we were not sure if this was going to work in the next round”.

In contrast, teams in the control group that found the insight (28 percent), indicate a different approach. In these instances, distant search effort was more likely to have been led by individuals as opposed to a collective team effort. As a participant reports “I was sure [that] there is a trick ... I am generally good at these things”. A striking feature reported by some individuals in teams from the control group that did not find the hidden rule, was how they were unable to pursue distant search due to the team climate. One participant reports “I even told [name] ... there has to be a pattern ... but we decided to help each other do it faster”. Another participant reported how the team’s approach was designed around individuals and not the task “[name] was the fastest, she was doing it in her head, so our approach was to let her do it and [name] can cross check ... you know because she was doing it so fast ... there could be mistakes”.

What is of particular significance is how micro-level actors respond to mandates that are often results of an external stimulus such as competitive performance benchmarks. On the one hand while a strong external stimulus as the one used in our experimental design triggered a higher frequency of distant search – thereby shifting the mean search behavior towards distant search in order to solve this specific problem, it also appears to have reduced the variance in search activities. This implies that freedom from specific organizational mandates may indeed

lead to a wider variety of problem formulation and search compared to searching within mandates handed down from top management. However, search activities within mandates are likely to be more productive for the organization if the mandates are chosen correctly. Not surprisingly, firms like Google that once prided themselves on employee freedom and resource abundance as central to generating breakthrough innovation are shifting to more focused R&D mandates in order to improve productivity (Carney and Getz, 2013).

Relevance for ‘problem-framing’ in multi-level organizational research: At the heart of our discussion lies the difference between the act of ‘problem framing’ by micro-level actors and the organization’s top management interpretation of the problem leading to search mandates for the rest of the organization. To view the actions of micro-level actors as purely governed by the instructions of top managers who have an organization-wide view, is not only at odds with realism, but also leads to a model of organizational search that has difficulty explaining phenomena such as intrapreneurial innovation (Menzela et al. 2007) or organizational ingenuity (Lampel et al. 2014). In contrast, by considering that micro-level actors are much more varied in their process of sensemaking we obtain variance in problem framing and search strategies that can explain the emergence of innovation even when there is no specific search mandate.

Research on microfoundations of problem formulation also suggests that an increase the comprehensiveness of alternative problem formulations positively influences the scope of search (Baer et al. 2013). However, the same body of research also points out that problem formulation is bounded by cognitive constraints on the decision maker (Eisenhardt and Zbaracki, 1992; Levinthal, 2011; Schwenk, 1988). The sources of these constraints can be traced to cognitive capabilities (Helfat and Peteraf, 2015), heuristics (Busenitz and Barney, 1997), group dynamics (Schweiger et al. 1986), and organizational processes (Hambrick and Mason, 1984) that all point towards developing a better understanding of how decision

makers respond to the demands of cognitive effort in problem formulation and improve the diversity of solutions (Schwenk, 1984).

The post-exercise debriefing of the results indicate two interesting aspects of how cognitive effort is expended at the micro-level. First, we observe that micro-level knowledge stored in routines can be a cognitive impediment to problem (re)formulation. As a participant from the control group points out “... we didn’t know we had to find a third rule” – suggesting that they framed the problem of increasing productivity as an ‘optimization issue’ and went about activities to fine-tune their production techniques over multiple rounds. In contrast, in the experiment group the announcement of high performance shortfall led the participants to formulate the increasing production problem as a problem with the production process itself that required search for novel techniques. In organizational terms we could argue that when confronted with the information that their performance was far below that of their peers – teams in the experimental group came to realize that their task capabilities were inadequate to meet the challenge, whereas teams in the control group that were not pressed externally to increase production were satisfied with their existing task capabilities. This suggests that in the absence of strong external stimuli decision makers may prefer to align problem formulation with existing capabilities.

Our results also show that over a quarter of the teams in the control group (28 percent) invested in increasing cognitive effort without any external performance benchmarks. Interestingly, Wagner et al. (2005) who use the same task but under different conditions, also found more than 20 percent of their control sample were able to find the hidden rule i.e. indication of cognitive effort without any external stimulus. This further indicates that while the dominant state of micro-level actors is to reduce cognitive effort in the absence of any stimulus, individuals still make idiosyncratic investments in cognitive effort to reframe problems. Indeed, as Katila and Thatchenkery (2015) point out “distant search may arise

from human curiosity to explore the unknown. So, rather than portray distant search as a difficult goal for most organizations, given appropriate permission we may be more likely to explore than is commonly thought.”

LIMITATIONS AND DIRECTIONS FOR FUTURE RESEARCH

In our design we use performance reference points as opposed to financial incentives to influence search behavior. As in real organizations, research using experimental designs shows that decision makers are likely to change behavior when they are presented with financial incentives (Agarwal, Croson, and Mahoney, 2010; Hossain and List, 2012) as well as performance reference points (Kuhnen and Tymula, 2012). However, these interventions work differently. While incentives appeal to the decision makers’ desire for a larger financial reward for performance, reference points appeal to the agents’ self-esteem and the desire of not falling behind a benchmark. In lab experiments, even though financial incentives have been shown to motivate behavior (Tafkov, 2013), scholars have also presented research that shows small financial rewards for performance may in fact have a negative influence on expected outcomes (Gneezy and Rey-Biel, 2014), as participants may perceive accepting financial reward to change behavior dilutes their social image amongst peers (Ariely et al. 2009). Therefore, while we do not use any financial incentives but rely solely on historical and social performance benchmarks, there is scope for further research on how individual’s desire for larger rewards influence search.

A major limitation of designing an experiment that tests problemistic search is avoiding the bias that may result from explicitly confronting subjects with a problem. We designed our experiment around a repetitive task, and instructed the subjects in how to perform this task without any indication that a problem may occur. The teams discussed and set the targets, and performance evaluation was based on team output. Beyond collectively deciding on targets,

the team did not coordinate task execution. Team performance was therefore the sum of individual task performance, but this may not be the case for most teams in a natural setting.

Kozlowski and Klein (2000) argue that lower level phenomena combine to create higher-level property along two extremes that run between composition and compilation. Composition is the coalescing of identical lower-level properties to yield higher-level phenomena that are essentially the same in terms of characteristics. Compilation describes combination of elements at lower-level that gives rise to phenomena at higher level that is different than its constituents. They illustrate this distinction using the examples of a rowing crew who exert efforts individually, with the combined performance representing the sum of these efforts. In contrast, compilation is represented by a baseball team, where team members interact with each other, and their combined performance is more a product of their ability to coordinate performances in an interactive fashion rather than the sum of their individual efforts. Our experiment is closer to the compositional end of the spectrum outlined by Kozlowski and Klein (2000). By confining ourselves to compositional design we increase the validity of our findings by avoiding the problem of multiple interaction patterns, which in turn can give rise to multiple emergent phenomena at a higher level. But this also limits the generalizability of our findings. Interaction in most teams, and arguably all organizations, is closer to the compilation end of the spectrum. Future research that adopts a multi-level view of problemistic search should identify different patterns of interaction at the micro-level in order to see if the same pattern of local and distant search holds at the macro-level.

Finally, in this paper, while we have used extant research to argue that individuals differ in their propensity to search – we did not design the experiment to determine individual characteristics that influence their propensity to engage in distant search. Future scholars of strategic management exploring the determinants of individual search behavior would benefit by taking into account a large body of literature in related fields that is gaining momentum in

the study of managerial decision making (Hodgkinson & Sadler-Smith, 2018). For instance, cognitive psychologists provide a strong theoretical and empirical basis to understand conscious and nonconscious cognitive processes that influence individual and collective search behavior (See for instance Evans & Stanovich, 2013; Kahneman & Frederick, 2002; Lieberman, 2007).

CONCLUSION

Our paper calls for an expanded view of problemistic search by scrutinizing and reconsidering two assumptions that play a central role in how performance feedback theory sees problemistic search. The first assumption is that problemistic search consists of two stages that operate sequentially. The first stage is local search, involving search for solutions near the organization's current stock of known solutions to problems; if local search fails the organization launches distant search for solutions that extend beyond current stock of known solutions. The two-stage model of problemistic search is linked to a second assumption that is not always explicitly stated, namely that top managers alone are charged with the task of triggering local search, and subsequently it is top managers who decide, based on further evaluation of performance, whether it is necessary to move to distant search.

Both assumptions tend to dominate current views of strategic search in part because they provide a parsimonious explanation of a wide variety of organizational behavior. Calling for an expanded view involves departure from parsimony by incorporating a more realistic view of how organizations operate. To begin with, it is clear that the monolithic view of organizational search may have to be modified in light of recent empirical research of how organizations evaluate feedback from multiple performance benchmarks. Such modification does not require an abandonment of the proposition that top managers play a role in organizational search. Rather, it pushes us towards a more complex multi-level view of

problemistic search. In this view, individuals in organizations who are engaged in problemistic search may use their own initiative to move from local to distant search, or even go directly to distant search. In the aggregate, this means that local and distant search occur simultaneously, but without a sense of urgency caused by failure to address performance shortfall the distribution of local to distant search is skewed towards the former. As our study suggests, a realistic view of problemistic search is therefore one in which local search predominates first, as performance feedback theory predicts. The transition from predominantly local to predominantly distant search occurs when top managers who monitor overall performance alert the organization to the inadequacy of local search. Expanding problemistic search therefore not only challenges the sequential view of local to distant search, it also suggests, as Greve (2011, p.88) notes, that performance feedback theory needs an expanded view of top managers' role in problemistic search.

REFERENCES

- Agarwal R, Croson R and Mahoney JT (2010) The role of incentives and communication in strategic alliances: An experimental investigation. *Strategic Management Journal*, 31(4): 413-437.
- Ariely D, Bracha A and Meier S (2009). Doing good or doing well? Image motivation and monetary incentives in behaving prosocially. *The American Economic Review*, 99(1): 544-555.
- Badre D, Doll BB, Long NM and Frank MJ (2012). Rostrolateral prefrontal cortex and individual differences in uncertainty-driven exploration. *Neuron*, 73(3): 595-607.
- Baer M, Dirks KT and Nickerson JA (2013). Microfoundations of strategic problem formulation. *Strategic Management Journal*, 34(2): 197-214.
- Bantel KA and Jackson SE (1989) Top management and innovations in banking: Does the composition of the top team make a difference? *Strategic Management Journal*, 10(S1): 107-124.
- Bapuji H, Hora M and Saeed AM (2012) Intentions, intermediaries, and interaction: Examining the emergence of routines. *Journal of Management Studies*, 49(8): 1586-1607.
- Barnett WP and McKendrick DG (2004) Why are some organizations more competitive than others? evidence from a changing global market. *Administrative Science Quarterly*, 49(4): 535-571.

- Baum JA and Dahlin KB (2007) Aspiration performance and railroads' patterns of learning from train wrecks and crashes. *Organization Science*, 18(3): 368-385.
- Baum JA, Rowley TJ, Shipilov AV and Chuang Y (2005) Dancing with strangers: Aspiration performance and the search for underwriting syndicate partners. *Administrative Science Quarterly*, 50(4): 536-575.
- Blettner DP, He Z, Hu S and Bettis RA (2015) Adaptive aspirations and performance heterogeneity: Attention allocation among multiple reference points. *Strategic Management Journal*, 36(7): 987-1005.
- Bromiley P (1991) Testing a causal model of corporate risk taking and performance. *Academy of Management Journal*, 34(1): 37-59.
- Bromiley P and Harris JD (2014) A comparison of alternative measures of organizational aspirations. *Strategic Management Journal*, 35(3): 338-357.
- Busenitz LW and Barney JB (1997) Differences between entrepreneurs and managers in large organizations: Biases and heuristics in strategic decision-making. *Journal of Business Venturing*, 12(1): 9-30.
- Carney B and Getz I (2013) *Google's 20% mistake*. Wall Street Journal.
<http://www.wsj.com/articles/SB10001424127887323906804579038451227863412>
Accessed online 21.08.2016
- Charness G, Gneezy U and Kuhn MA (2013) Experimental methods: Extra-laboratory experiments-extending the reach of experimental economics. *Journal of Economic Behavior & Organization*, 91: 93-100.

- Cohen MD and Bacdayan P (1994) Organizational routines are stored as procedural memory: Evidence from a laboratory study. *Organization Science*, 5(4): 554-568.
- Cyert RM and March JG (1963). A behavioral theory of the firm. *Englewood Cliffs, NJ*, 2.
- Eisenhardt KM and Zbaracki MJ (1992) Strategic decision making. *Strategic Management Journal*, 13(S2): 17-37.
- Epple D, Argote L and Devadas R (1991) Organizational learning curves: A method for investigating intra-plant transfer of knowledge acquired through learning by doing. *Organization Science*, 2(1): 58-70.
- Evans JSB and Stanovich KE (2013) Dual-process theories of higher cognition: Advancing the debate. *Perspectives on psychological science*, 8(3): 223-241.
- Felin T and Foss NJ (2005) Strategic organization: A field in search of micro-foundations. *Strategic Organization*, 3(4): 441.
- Felin T, Foss NJ and Ployhart RE (2015) The microfoundations movement in strategy and organization theory. *The Academy of Management Annals*, 9(1): 575-632.
- Gavetti G, Greve HR, Levinthal DA and Ocasio W (2012) The behavioral theory of the firm: Assessment and prospects. *The Academy of Management Annals*, 6(1): 1-40.
- Giachetti C and Lampel J (2010) Keeping Both Eyes on the Competition: Strategic Adjustment to Multiple Targets in The UK Mobile Phone Industry *Strategic Organization*, 8(4): 347-376.

- Gneezy U and Rey - Biel P (2014) On the relative efficiency of performance pay and noncontingent incentives. *Journal of the European Economic Association*, 12(1): 62-72.
- Greve HR (2003) *Organizational learning from performance feedback: A behavioral perspective on innovation and change* Cambridge University Press.
- Greve HR (2013) Microfoundations of management: Behavioral strategies and levels of rationality in organizational action. *The Academy of Management Perspectives*, 27(2): 103-119.
- Hambrick DC and Mason PA (1984) Upper echelons: The organization as a reflection of its top managers. *Academy of Management Review*, 9(2): 193-206.
- Healey MP and Hodgkinson GP (2014) Rethinking the philosophical and theoretical foundations of organizational neuroscience: A critical realist alternative. *Human Relations*: 0018726714530014.
- Helfat CE and Peteraf MA (2015) Managerial cognitive capabilities and the microfoundations of dynamic capabilities. *Strategic Management Journal*, 36(6): 831-850.
- Hodgkinson G and Sadler-Smith E (2018) The dynamics of intuition and analysis in managerial and organizational decision making. *The Academy of Management Perspectives*, Article in press.
- Hossain T and List JA (2012) The behavioralist visits the factory: Increasing productivity using simple framing manipulations. *Management Science*, 58(12): 2151-2167.

- Jackson M and Cox DR (2013) The principles of experimental design and their application in sociology. *Annual Review of Sociology*, 39: 27-49.
- Joseph J and Gaba V (2015) The fog of feedback: Ambiguity and firm responses to multiple aspiration levels. *Strategic Management Journal*, 36(13): 1960-1978.
- Kahneman D and Frederick S (2002) Representativeness revisited: Attribute substitution in intuitive judgment. *Heuristics and biases: The psychology of intuitive judgment*, 49: 81.
- Kacperczyk A, Beckman CM and Moliterno TP (2015) Disentangling risk and change internal and external social comparison in the mutual fund industry. *Administrative Science Quarterly*, 60(2): 228-262.
- Katila R and Thatchenkery S (2015) Local search. In D. Teece and M. Augier (Ed.), *The palgrave encyclopedia of strategic management* Palgrave Macmillan: Basingstoke.
- Khan AM and Manopichetwattana V (1989). Innovative and noninnovative small firms: Types and characteristics. *Management Science*, 35(5): 597-606.
- Kim JJ, Finkelstein S and Haleblian JJ (2015). All aspirations are not created equal: The differential effects of historical and social aspirations on acquisition behavior. *Academy of Management Journal*, 58(5): 1361-1388.
- Klein KJ and Kozlowski SWJ (Eds.) (2000a). *Multilevel theory, research, and methods in organizations: Foundations, extensions, and new directions*. San Francisco: Jossey-Bass.

- Klein KJ and Kozlowski SW (2000b) From micro to meso: Critical steps in conceptualizing and conducting multilevel research. *Organizational Research Methods*, 3(3): 211-236.
- Knudsen T and Levinthal DA (2007). Two faces of search: Alternative generation and alternative evaluation. *Organization Science*, 18(1): 39-54.
- Kozlowski SW and Klein KJ (2000) A multilevel approach to theory and research in organizations: Contextual, temporal, and emergent processes. In K. Klein and S. Kozlowski (Ed.), *Multilevel theory, research and methods in organizations: Foundations, extensions, and new directions*. Jossey-Bass.
- Kuhnen CM and Tymula A (2012) Feedback, self-esteem, and performance in organizations. *Management Science*, 58(1): 94-113.
- Lampel J, Honig B and Drori I (2014) Organizational ingenuity: Concept, processes and strategies. *Organization Studies*, 35(4): 465-482.
- Lant TK (1992) Aspiration level adaptation: An empirical exploration. *Management Science*, 38(5): 623-644.
- 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.
- Levinthal DA (2011) A behavioral approach to strategy—what's the alternative? *Strategic Management Journal*, 32(13): 1517-1523.

- Li Q, Maggitti PG, Smith KG, Tesluk PE and Katila R (2013) Top management attention to innovation: The role of search selection and intensity in new product introductions. *Academy of Management Journal*, 56(3): 893-916.
- Lieberman MD (2007) Social cognitive neuroscience: a review of core processes. *Annual Review of Psychology* 58: 259-289.
- Loyd DL, Kern, MC and Thompson L (2005) Classroom research: Bridging the ivory divide. *Academy of Management Learning & Education*, 4(1): 8-21.
- March JG (1981) Footnotes to organizational change. *Administrative Science Quarterly*: 563-577.
- Marengo L (2014) Representation, search, and the evolution of routines in problem solving. *Industrial and Corporate Change*: dtu023.
- Mas A and Moretti E (2009) Peers at work. *The American Economic Review*, 99(1): 112-145.
- Menzel HC, Aaltio I and Ulijn JM (2007) On the way to creativity: Engineers as intrapreneurs in organizations. *Technovation*, 27(12): 732-743.
- Molitero TP, Beck N, Beckman CM and Meyer M (2014) Knowing your place: Social performance feedback in good times and bad times. *Organization Science*, 25(6): 1684-1702.
- Novick LR and Bassok M (2005) Problem solving. In K. Holyoak and B. Morrison (Ed.), *The Cambridge handbook of thinking and reasoning*. Cambridge Univ Press.

- Novick LR and Sherman SJ (2003) On the nature of insight solutions: Evidence from skill differences in anagram solution. *The Quarterly Journal of Experimental Psychology: Section A*, 56(2): 351-382.
- Ocasio W (1997) Towards an attention-based view of the firm. *Strategic Management Journal*: 18: 187-206.
- Ocasio, W (2011) Attention to Attention. *Organization Science*: 22 (5): 1286-1296.
- Payzan-LeNestour E (2012) Learning to choose the right investment in an unstable world: Experimental evidence based on the bandit problem. *Swiss Finance Institute Research Paper* (10-28).
- Schunk D (2009) Behavioral heterogeneity in dynamic search situations: Theory and experimental evidence. *Journal of Economic Dynamics and Control*, 33(9): 1719-1738.
- Schweiger DM, Sandberg WR and Ragan JW (1986) Group approaches for improving strategic decision making: A comparative analysis of dialectical inquiry, devil's advocacy, and consensus. *Academy of Management Journal*, 29(1): 51-71.
- Schwenk CR (1984) Cognitive simplification processes in strategic decision - making. *Strategic Management Journal*, 5(2): 111-128.
- Schwenk CR (1988) The cognitive perspective on strategic decision making. *Journal of Management Studies*, 25(1): 41-55.
- Steyvers M, Lee MD and Wagenmakers E (2009) A Bayesian analysis of human decision-making on bandit problems. *Journal of mathematical psychology*, 53(3): 168-179.

- Tafkov ID (2012) Private and public relative performance information under different compensation contracts. *The Accounting Review*, 88(1): 327-350.
- Thurstone LL and Thurstone TG (1941) Factorial studies of intelligence. *Psychometric monographs*.
- Tyler BB and Caner T (2015) New product introductions below aspirations, slack and R&D alliances: A behavioral perspective. *Strategic Management Journal*, 37(5): 896–910
- Wagner U, Gais S, Haider H, Verleger R and Born J (2004) Sleep inspires insight. *Nature*, 427(6972): 352-355.
- Washburn M and Bromiley P (2012) Comparing aspiration models: The role of selective attention. *Journal of Management Studies*, 49(5): 896-917.
- Wiseman RM and Bromiley P (1996) Toward a model of risk in declining organizations: An empirical examination of risk, performance and decline. *Organization Science*, 7(5): 524-543.
- Wollersheim J and Heimeriks KH (2016) Dynamic Capabilities and Their Characteristic Qualities: Insights from a Lab Experiment. *Organization Science* 27(2):233-248

Table 1: Evidence of distant search as observed in the creation of the new, hidden rule in experimental and control groups

Condition	Total	Distant search		Success rate
		No	Yes	
Experimental groups:				
Internal and external performance reference points	15	2	13	87%
Control groups:				
Only internal performance reference point	18	13	5	28%
Total Observations	33	15	18	

$\chi^2 = 11.444^{***}$

*** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$

Table 2: Results of a difference-in-differences regression between rounds 1 & 3

VARIABLES	OLS Distant Search
DID Term (Experiment * Time)	0.59*** (0.15)
Experiment group effect	0.00 (0.10)
Time (Round 3)	0.28** (0.10)
Constant	-0.00 (0.07)
Observations	66
R-squared	0.59

Standard errors in parentheses

*** p<0.001, ** p<0.01, * p<0.05

Figure 1: Multilevel approach to theorizing problemistic search

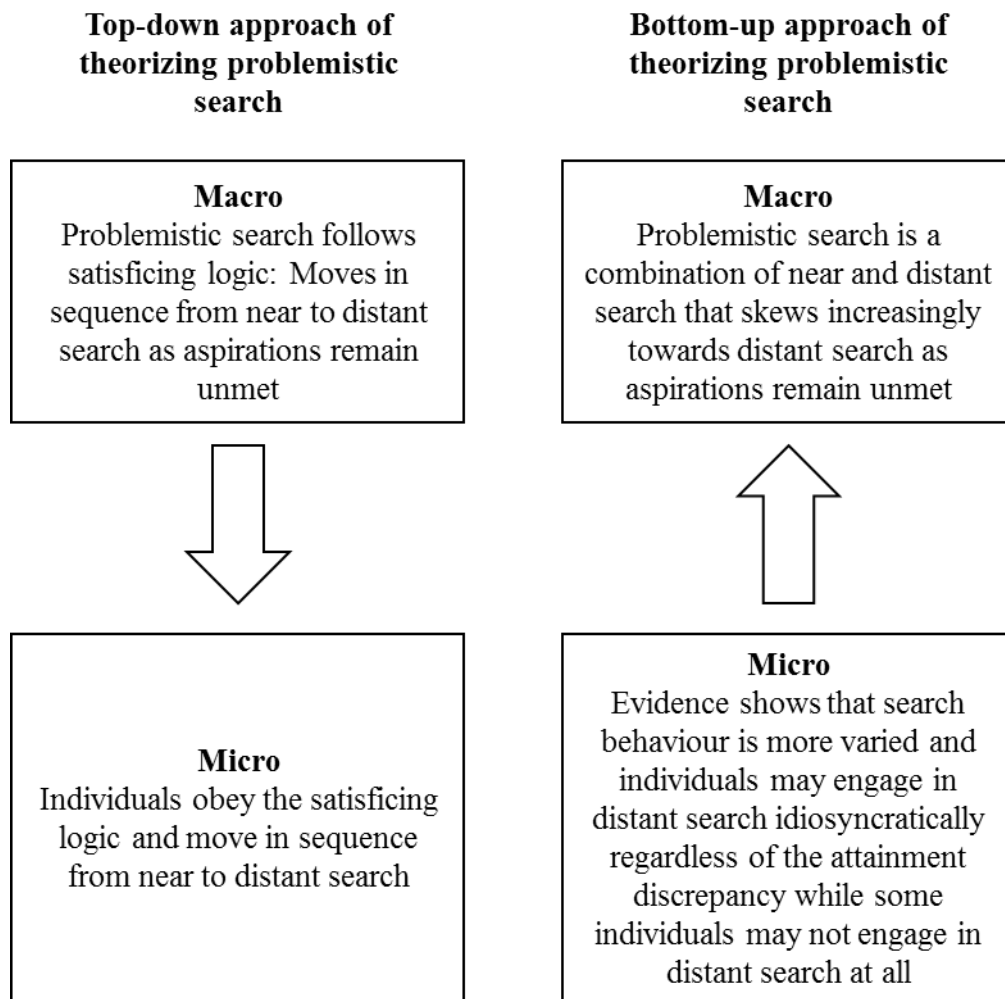


Figure 2 The two ways of solving a NRT

(i) Left-hand side shows how to solve the NRT by sequentially executing the task using the two known rules in seven steps (ii) Right-hand side shows how to solve the NRT by using the two known rules and by creating a third new rule in three steps. (For further details see (Wagner *et al.*,2004))




Figure 3: Experimental design

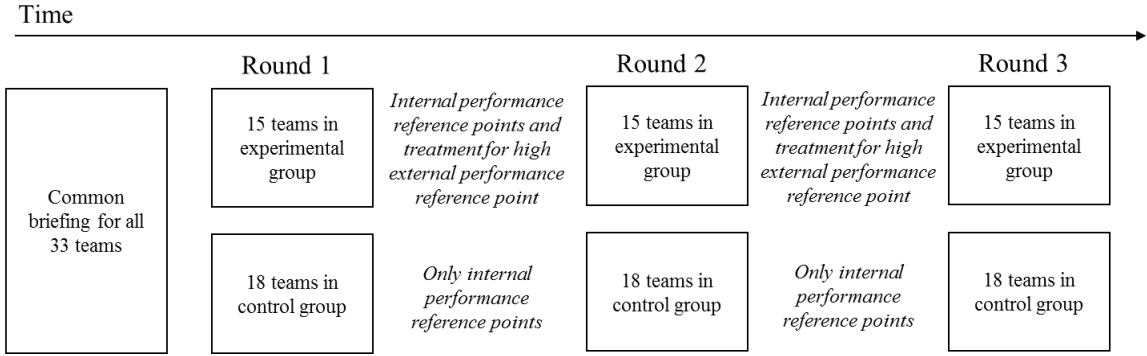


Figure 4: Performance of control and experimental groups

Bars indicate the mean performance and the stock lines represent the maximum, 75th percentile, 25th percentile, and minimum values within a group. N = 18 for control groups and N = 15 for experiment groups




Figure 5: Linear predictions of distant search

