The Effects of Problem Structure and Team Diversity on Brainstorming Effectiveness *

Stylianos Kavadias† Svenja C. Sommer‡

August 10, 2009

Abstract

Since Osborn's *Applied Imagination* book in 1953, the effectiveness of *brainstorming* has been widely debated. While some researchers and practitioners consider it the standard idea generation and problem-solving method in organizations, part of the social science literature has argued in favor of nominal groups, i.e., the same number of individuals generating solutions in isolation. In this paper, we revisit this debate, and we explore the implications that the underlying problem structure and the team diversity have on the quality of the best solution as obtained by the different group configurations. We build on the normative search literature of new product development (NPD), and we show that no group configuration dominates. Therefore, nominal groups perform better in specialized problems, even when the factors that affect the solution quality exhibit complex interactions (problem complexity). In cross-functional problems, the brainstorming group exploits the competence diversity of its participants to attain better solutions. However, their advantage vanishes for extremely complex problems.

(KEYWORDS: Brainstorming process, New product development team, Ideation, Complexity)

(First version September 2007; Revised versions: October 2008, July 2009)

*We would like to thank the AE and two anonymous referees for their helpful comments, which made this manuscript a lot stronger.

†College of Management, Georgia Institute of Technology, stelios@gatech.edu

‡HEC Paris, sommers@hec.fr

1. Introduction

Since Osborn's *Applied Imagination* (1953), brainstorming has been widely used in organizational contexts as a problem-solving and concept generation technique. There are ample practitioner literature and popular business press articles that advocate brainstorming as a useful tool for problem-solving and offer suggestions for making it more effective (e.g., Fisher and Fisher 1998, Gundry and La Mantia 2001, Robbins and Judge 2007).

Despite wide adoption by practitioners, social psychology researchers have questioned, as early as 1958, whether brainstorming sessions are the most-effective way to generate problem solutions (Taylor *et al.* 1958). An abundance of experimental evidence demonstrates that brainstorming groups are outperformed by the so-called "nominal" groups (the same number of individuals generating solutions in isolation) given the same problem statement and time (Stroebe and Diehl 1994, Paulus *et al.* 1996). Based on a meta analysis of the vast amount of this research, Mullen *et al.* (1991, pg. 18) conclude: "*It appears to be particularly difficult to justify brainstorming techniques in terms of any performance outcomes, and the long-lived popularity of brainstorming techniques is unequivocally and substantively misguided.*"

The persistent character of the debate between the two viewpoints raises an important question: Has the practitioner literature been entrapped in a long-standing 'fad', or are there beneficial aspects of brainstorming that are overlooked by social psychologists? Sutton and Hargadon's (1996) study and its sequels (e.g., Hargadon 2003) suggest that the practitioners' persistence may be explained by fundamental contextual differences between laboratory experiments and real brainstorming sessions within an organization (see also Sawyer 2007).

We revisit the original research question (i.e., *are brainstorming sessions effective means for problem-solving*?) in an effort to assess the different viewpoints and to present them cohesively through an integrated theoretical framework. We build upon the normative models in the new product development (NPD) literature and we conceptualize brainstorming and nominal group sessions as multi-agent searches for a solution to a problem. We focus on problems in which the group cannot describe the performance function and all potential solutions in advance. It is this lack of algorithmic solutions that gives rise to brainstorming techniques in the first place. Drawing on the empirical social psychology literature, we impose certain behavioral rules on the search strategies. Our results show that no group configuration dominates. In fact, we show the moderating roles of problem structure and team diversity on the relative advantage of brainstorming groups.

The rest of the paper is organized as follows: Section 2 presents a brief discussion of the literature. In Section 3 and 4, we outline the key building blocks of the normative model and their specification in the simulation. In Section 5, we proceed with the simulation experiments and some analytical results. Section 6 concludes with a discussion of the hypotheses derived from the model.

2. Problem-solving Alone or in a Group? Past Research

Group brainstorming is an approach that aims to tackle a recurrent organizational challenge: problem-solving. Different scientific disciplines have tried to understand effective problem-solving, especially when it takes place at the "fuzzy front-end" of the product development processes (Wheelwright and Clark 1992). In this brief literature review, we present two parallel approaches to the topic. The first approach is based on Osborn's early claim (Osborn 1953) that groups are more effective at creating new solutions. This stream does not focus on the exact search process, but tries to determine which organizational structure generates better solutions as a result of psychological or sociological factors. The second approach builds upon the economics and operations research paradigm. Since Simon (1969), researchers have represented the problem-solving process as a search among different alternative solutions. The two research paths did not cross mainly due to their methodological distance: The first stream relies heavily on empirical and experimental work, whereas the second one applies normative methods. Although our work is normative, it draws heavily on repeatedly confirmed observations of the first group.

2.1 Brainstorming Groups: The Organizational Perspective

This stream of literature focuses mainly on the effects of team structures on creativity and problem-solving. Early on, Osborn (1953) identifies four rules for effective brainstorming: create as many ideas as possible; the wilder the ideas the better; combine and improve upon mentioned ideas; and do not criticize ideas. He proposes that building on other people's ideas and stimulating ideas based on others' arguments makes the ideation during problem-solving particularly effective.

Osborn's work stimulated a long-standing debate concerning the true effectiveness of brainstorming. Numerous experiments in psychology compare the performance of brainstorming and nominal groups. Taylor *et al.* (1958) are the first to refute Osborn's idea, followed by a large number of studies (for an overview, see Stroebe and Diehl 1994, Paulus *et al.* 1996, or Paulus 2000). The literature suggests various reasons why group brainstorming is less effective. The most-cited reasons for poor performance (see Diehl and Stroebe 1987 or Gallupe *et al.* 1991) are: *production blocking* (only one person can speak at a given time, so ideas might be forgotten while listening to and understanding the speaker); *evaluation apprehension* (members may not speak up due to the presence of peer evaluation); and *free riding* (members have less incentive to contribute due to the inability to observe effort). The results of multiple experiments performed by Diehl and Stroebe (1987) suggest that production blocking is by far the most-important factor, and it is not easily alleviated by organizational means (Diehl and Stroebe 1991). The study by Offner *et al.* (1996) provides evidence that trained facilitators may ease the effects of production blocking, but does not find evidence for overcoming the advantage of nominal groups.

If brainstorming groups perform so poorly compared to individuals, how can firms afford to use them in today's fiercely competitive business environment? Sutton and Hargadon (1996) suggest that laboratory research ignores the organizational context in which brainstorming takes place.[1] Most laboratory experiments use undergraduate or high school students (Diehl and Stroebe

[1] "*Nearly all other brainstorming research was done with participants who (1) had no past or future task inter-*

1987, 1991) as subjects, without experience in brainstorming. Further, they look at problems that do not require any prior specialized knowledge, e.g. benefits or difficulties from an extra thumb on each hand (Bouchard and Hare 1970, Gallupe *et al.* 1991). Instead, companies use group brainstorming sessions to solve more complex problems, which involve different knowledge domains and specializations and have clearer objectives (Sawyer 2007). Would such contextual differences promote brainstorming groups? Or is the evidence presented by Sutton and Hargadon (1996) just a reflection of a positive perception bias (participants believing their performance in group brainstorming is higher when, in fact, it is not) as a few studies suggest (Naquin and Tynan 2003, Paulus *et al.* 1995)?

Our study attempts to answer the question of whether brainstorming groups are more or less effective than nominal groups given specific organizational contexts. We focus our analysis on the impact of two contextual features that can be assessed by management and have been identified as potentially important factors in the prior literature: first, team diversity; and, second, the underlying structure of the problem at hand. Furthermore, we measure effectiveness or performance of each group setting by the quality of the best proposed solution (rather than by the average quality or by the number of ideas, two measures typically used in the organizational literature). This measure has also been used in other recent work (Fleming and Singh 2009, Girotra *et al* 2009).

2.2 Problem-solving as a Search Process

Simon (1969) conceptualizes problem-solving as a search through multiple solutions. Several authors follow that paradigm to determine the optimal search rule: Weitzman (1979) derives a simple reservation price rule for searching among solution alternatives with different values and search costs, and Loch *et al.* (2001) extend his findings to account for the potential of learning through imperfect experiments. Field work and empirical studies show that designers test different solutions

dependence; (2) had no past or future social relationships; (3) didn't use the ideas generated; (4) lacked pertinent technical expertise; (5) lacked skills that complement other participants; (6) lacked expertise in doing brainstorming; and (7) lacked expertise in leading brainstorming sessions" (Sutton and Hargadon 1996, pg. 688).

to spot the "best" one (Clark and Fujimoto 1989, Thomke 1998).

Drawing on Simon (1969) and his observation of the limited ability of decision-makers to know the exact solution space and to determine the optimal search, Kauffmann (1993) employs an alternative conceptualization: Agents know little about the solution space structure and they employ heuristic rules to search in "rugged" solution landscapes. Numerous other researchers analyze the effectiveness of different heuristic rules (e.g., Kauffmann *et al.* 2000, Gavetti and Levinthal 2000, Fleming 2001, Fleming and Sorenson 2004).

While most of this literature relies on simulations, Dahan and Mendelson (2001) and Erat and Kavadias (2008) offer analytical results; Dahan and Mendelson (2001) apply the results of extreme value theory (Galambos 1978) to analyze the effects of the landscape distribution's upper tail on the optimal number of sampling trials. Erat and Kavadias (2008) show how Weitzman's original search policy changes when accounting for the complex structure of solution landscapes and potentially correlated performances of alternative solutions.

Those two branches of research fit different contextual settings. The optimization models assume that some prior knowledge and valuation of the potential solutions exists. Therefore, decision-makers can assess and compare possible solutions through complex dynamic programming models. Such settings proxy the design problem-solving processes during experimentation and testing in new product development (see Thomke 2003). On the other end, models that assume search with limited cognition (bounded rationality) and imperfect value assessment (partial landscape knowledge) are better suited for the "fuzzy front-end" stage of the NPD process where ideation takes place. A dominant feature across all these models is that they do not account for the different organizational structures that could perform problem-solving, i.e. individuals searching in isolation (nominal groups) versus a team of people brainstorming to solve a problem (brainstorming group).

2.3 Recent Perspectives

Four recent papers lie close to our work. Hong and Page (2001) are the first to recognize two important dimensions in the normative search literature: the diversity in perspectives (due to reasons such as diverse degrees of understanding or different backgrounds) and the limited cognition that results in the use of simple heuristics. They show that due to the existence of both factors at once, diverse teams may reach better solutions (a claim that parallels ours). However, they do not offer any formal comparison arguments as we do.

Fleming and Singh (2009) analyze patent citations to identify whether groups of collaborators or "lone inventors" are more likely to develop breakthrough ideas. They find that collaborating teams benefit from reducing the "poor outcomes", but they do not account for "equivalence" in search capabilities (i.e., equal group size). Girotra *et al.* (2009) conduct experiments to compare the brainstorming group performance with the "nominal group technique".[2] They find that the nominal group technique outperforms brainstorming, and that organizational structures differ in their capability to reliably assess the solution quality. Unlike their study, we focus on the circumstances that make each organizational setting better at *generating* valuable ideas, but we do not provide any insights as to which setting is better at *selecting* the best one among these ideas.

Terwiesch and Xu (2008) recognize that innovation could be the outcome of decentralized search processes, such as innovation tournaments. While the focus of their paper is on single agents competing in isolation (rather than a comparison across different organizational structures), their study exhibits similar modeling assumptions to ours. We represent explicitly two different sources of capability differences among the search agents, one of which (diversity) is modeled in a similar fashion to Terwiesch and Xu (2008, see details below).

3. General Setup

In this section, we outline our model setup. We conceptualize the effects of two major enablers,

[2] A hybrid method where individuals brainstorm individually upfront and then work together as a group (Robbins and Judge 2007).

the problem structure and team diversity, on the effectiveness of problem-solving under different organizational structures (nominal versus brainstorming groups).

3.1 Problem Structure and Solution Performance

We first contemplate the problem structure. Every problem may admit a large number of solutions that constitute the solution space. We assume that the group (nominal or brainstorming) cannot describe all potential solutions in advance, thereby making it impossible to determine up front the global optimum. In fact, their inability to fully capture the performance function gives rise to brainstorming techniques in the first place.

We explicitly model two aspects of the problem structure that reflect properties of the solution space: the problem *complexity*, which summarizes the smoothness of the performance function and represents whether the solution space exhibits more or less local optima, and the degree of *functional competence diversity* required by the problem to attain certain solutions within the solution space.

Let all feasible solutions be described by a set of factors $\{w_1, w_2, w_3, ..., w_N\}$ that affect the solution performance.[3] N represents the size of the solution space. The specific value of each factor, w_j ($w_j \in \mathcal{L} = \{1, 2, 3, ..., L\}$), differentiates solutions among each other. Solutions that just differ in one of the w_i are expected to exhibit similar performances, while significantly different solutions will differ across several w_is. In summary, each feasible solution is a *point* in the solution space $\mathcal{S} = \{\vec{w} \in \mathcal{L}^N : \vec{w} = (w_1, w_2, w_3, ..., w_N)\}$.

The factors w_j determine the exact solution performance $V(\vec{w})$ as follows: $V(\vec{w}) = \sum_{i=1}^{N}(W_i)$, that is, $V(\vec{w})$ is the sum of individual performance contributions W_i from each of the N factors. The potential interactions between the different factors affect the individual performance contributions. If W_i, the performance contribution of factor w_i, depends on the values of K other factors, W_i is

[3] To give a tangible example of such factors, contemplate the project to redesign the supermarket shopping cart undertaken by IDEO ("Nightline," ABC News, July 13, 1999). During the ideation and concept generation phase, various factors are addressed, such as user behavioral traits, design for manufacturing concerns, and anti-theft challenges.

a function of a (K+1)-dimensional vector, $(w_i, w_{i1}, w_{i2}, ..., w_{iK})$. One can also interpret W_i as a subsystem performance influenced by the various factors (w_i).

This representation allows us to rigorously characterize the *complexity* of the performance function $V(\vec{w})$ through the number of interactions K. A larger K increases the complexity of the performance function. For $K = 0$ each factor contributes independently to $V(\vec{w})$; therefore, the performance of neighboring points (i.e., points that differ by only one w_i) differs by only one W_i value and the points are highly correlated. If $K = N - 1$, the small difference of a single w_i between two neighboring points alters all W_j $(j = 1, ..., N)$ values and makes the neighboring solutions highly uncorrelated (see also Gavetti and Levinthal 2000).

Our choice to model the interactions via the individual performance contributions aligns with prior literature: it maps directly on the performance function of the NK simulation method (e.g., Kauffman 1993), and it resonates with Mihm *et al.* (2003), who also model such interactions within a new product development problem via individual performance contributions.

Our representation offers a nice characterization of the degree of *cross-functionality* of the problem. While some specialized problems only require the competencies from one functional area, most organizational problems require diverse functional perspectives. For example, the commercialization of any new product requires the synthesis of consumer need knowledge (marketing competence), the definition and realization of technical specifications (engineering competence), and the scaling of lab prototypes to large production runs (manufacturing competence).

The degree of cross-functionality defines how the factors w_i stem from functional disciplines. For simplicity, we consider two basic settings: a problem that requires a single functional competence, and an alternative setting where two functional competencies are necessary. In the latter case, we assume that factors w_1 to w_d relate to the first functional expertise, whereas w_{d+1} to w_N stem from the other functional expertise.

We consider that the organizational objective is to obtain the highest $V(\vec{w})$ possible for a problem at hand. However, the task is difficult because the inherent complexity makes optimization impossible.

3.2 Team Diversity

As in a typical organizational setting, a group of M members is assigned the task to generate solution ideas for a product design problem.[4] Different members might report to different functional disciplines and, therefore, exhibit *functional background diversity* (Pelled *et al.* 1999), which implies thorough understanding of some factors that affect the solution performance (see problem structure) and minimal understanding of the rest. However, even members of the same functional discipline might adopt diverse perspectives when solving a problem, due to different past organizational and work experience; they may, therefore, exhibit *cognitive diversity* (Taylor and Greve 2006). We capture cognitive diversity through the variation in the initial solutions (original starting points) and through the systematic variation during idea generation (details below). For simplicity, hereafter the term "diversity" refers to cognitive styles unless otherwise stated.

We should note here that there is no consensus regarding the link between diversity and higher performance. For example, Dougherty (1992) shows that the different interpretive schemes ("thought worlds") that come with diversity inhibit creativity due to intra-group conflicts and excessive coordination (see also Milliken and Martins 1996). Others identify a weak positive relationship between team diversity in terms of their *ex-ante* ideas and the quality of the solution produced (e.g., Yetton and Bottger 1982, Wanous and Youtz 1986, and Libby *et al.* 1987). Sawyer (2007) suggests that diversity helps only if "group flow factors" (e.g., open communication, well-defined goals, equal participation) keep the conflict from spiraling out of control. We abstract away from intergroup conflicts, assuming reliable "group flow" in order to isolate the diversity effects.

[4] We focus on a new product development setting, but our insights apply to any organizational challenge that necessitates cross-functional expertise.

3.3 The Search Process

In this section, we conceptualize the search process during idea generation and problem-solving. We ground our assumptions in the experimental facts recorded in the social psychology literature.

Similar to Terwiesch and Xu (2008), we assume that all individuals have some (randomly determined) initial solution. In our notation, this is $\vec{w}_i$, a point in the solution landscape. In their effort to generate ideas and solve the problem (challenge), team members' new ideas are triggered by prior solutions (either the initial reference point or subsequently generated ideas). That is, we posit that idea generation cannot happen without some prior starting knowledge or prior ideas as a stimulus; in other words, individuals cannot magically come up with solutions, but instead they need to *build* upon one of the earlier proposed solutions. How do they build? There are two distinct components of the idea generation process: first, the starting point based on which the idea will be built; second, the actual transformation that happens and brings the new idea to life.

For the first step, we assume that each individual builds upon the current best idea with some likelihood (say α), or upon some of their previous ideas with $1 - \alpha$ probability. Our assumptions allow us to capture two important structural features: the likelihood α represents the capability of individuals to assess the *solution quality* of the "floating" ideas during the ideation session. α could range from 0 to 1, indicating that there are contexts where individuals may have some understanding of the running $max\{V(\vec{w})\}$, but, in some other cases, the ideation process could resemble a random walk process. The second feature is the memory effect that plays a role. Individuals hold recent solutions more vividly in memory and, as such, they are more likely to use recent solutions as starting points from where they build on. The influence of the short-term memory is well-recorded in the psychology and decision-making literature (Kahneman *et al.* 1993, Kahneman 2000).

For the second step, we assume that individuals change a part of their initial point and keep other factors unchanged.[5] We, therefore, model idea generation during the problem-solving process

[5] As an example, Thomke *et al.* (1999) describe the crash test process in BMW, where building better performance

as a *random* change of a subset of the factors the particular individual is knowledgable about. Our assumption about *randomly* changing some factors captures the fact that mapping between the solution vectors and their performance ($V(\vec{w})$) cannot be *ex-ante* described and optimized over. Therefore, the change of a particular factor cannot be decided based on some underlying optimization process.

3.4 Comparison of Brainstorming and Nominal Groups

Aligned with the extant literature in social psychology, we consider two basic organizational configurations for the problem-solving process. The **brainstorming group** configuration considers that all members are in the same room and work together on generating problem solutions; the **nominal group** configuration keeps the team members isolated during the solution generation process and eventually adopts and implements the best solution among all resulting ones. Since time is usually a limiting factor, we constrain both to search for the same fixed amount of time T.

Individuals who brainstorm in isolation (nominal groups) start by building on their own initial solution and continue expanding their own generated ideas over time. We capture the team diversity by the number of different starting points, as well as the number of distinct ideas generated in each period by the individuals in nominal groups.

In brainstorming groups, the members are influenced by and can build on ideas mentioned by others in the discussion. In other words, their new ideas may be triggered by other members' ideas. The notion of building upon the ideas of others is often encouraged (suggested already by Osborn 1953; IDEO's deep dive sessions have it as a rule). The first person speaking up, therefore, influences the path everybody takes, capturing the observation that groups may fixate quickly on previously mentioned ideas and focus longer on the same topic (e.g., Larey and Paulus 1999). Sawyer (2007) calls this "topic fixation" (Sawyer, 2007, pg. 65).[6]

implied keeping the rest of the structural elements of a car body fixed, and changing the dimensions of a pillar.

[6] Team members could definitely depart from the group string of thoughts at any given point in time. However, by

Finally, in accordance with the psychology literature, we incorporate the effects of *production blocking*, *evaluation apprehension* and *free riding*. As the experimental evidence has shown, all three effects imply that fewer members speak up during the group brainstorming process. For the effect of production blocking and free riding, we assume that they reduce the number of ideas generated, independently of the quality of ideas. For the evaluation apprehension phenomenon, we assume that it has consequences for the quality of ideas being mentioned. Building on Cottrell's (1972) evaluation apprehension model, we assume that the group presence will induce members with a 'better than the current best' idea to speak-up and contribute, and will induce members with 'worse than the current best' idea to hold off. Since we assume that individuals are capable of assessing the solution quality (with likelihood α), such phenomena are allowed to happen in our model.

Based on the above discussion, in brainstorming groups, only a subset M' of the M group members communicate their ideas in each time period. The remaining members $M - M'$ may generate ideas but they do not influence the outcome since these ideas stay in the shadow. Without evaluation apprehension, M' corresponds to any subset of potential ideas being mentioned. However, the presence of evaluation apprehension implies that M' includes at least one idea that is better than the current best, thereby skewing the distribution of the mentioned ideas. Depending on the training of the facilitator and the rules s/he enforces, production blocking may be reduced (but not eliminated, see Diehl and Stroebe 1991). Such rules may include pre-specified speaking times, notepads, or even electronic brainstorming.

Our overall objective is to compare the performance of brainstorming groups in generating valuable ideas relative to that of *equivalent* nominal groups.[7] Therefore, we assume an equal number of individuals (M) and an equal amount of time spent in the two setups (T idea generation

assuming that team members are more likely to build on a previously mentioned idea compared to isolated individuals, we are able to capture the essential difference with the nominal group: the cross influence of co-located individuals.

[7] Note that we focus on the question of which setting is better at generating valuable ideas, but cannot provide any insights into the question of which setting is better at selecting ideas.

phases). More importantly, unlike the majority of the experiments performed in the psychology literature, we focus on the expected value of the *best* performance of the generated solutions, as opposed to the number or average value of the solutions. While our problem setup cannot be solved analytically in its general format, we complement our numerical experiments with specific analytical results.

4 Specification of the Numerical Experiments

In this section, we specify the parameters we chose for the simulation experiments presented in Section 5. We start by describing the creation of problem instances (the solution landscape), followed by a description of the search processes of the nominal and brainstorming groups.

4.1 Creating Problem Instances

All numerical experiments (unless otherwise stated) follow the same set of specifications: a team of $M = 8$ individuals generates ideas for $T = 15$ time periods for a problem with a solution space of size $N = 10$, where the different factors w_i take one of two distinct values ($\mathcal{L} = \{0, 1\}$) and contribute W_i to the solution performance, i.e., $V(\vec{w}) = \sum_{i=1}^{N}(W_i)$. Each performance contribution W_i is drawn from the same distribution function $F(W)$. As a robustness check, we vary the dispersion of the performances W_i. We consider both a uniform distribution ($U[0, 1]$, a standard formulation of the NK model), and an exponential distribution with the same mean ($\lambda = 2$). The exponential distribution fits better our ideation context as it does not bound from above the quality of the proposed solutions. Since the results are qualitatively the same, we present the latter results.

Regarding the main constructs introduced in our general setup, we make the following assumptions: For the experiments of the cross-functional problems, we consider that the functional groups are of equal size (four individuals each) and that each group is knowledgable about the same number of factors, i.e., group 1 about $w_1, ..., w_5$ and group 2 about $w_6, ..., w_{10}$.[8]

[8] We could naturally imagine a setting where expertise sets may overlap. However, to identify the difference in our

For the complexity K, first recall that if W_i depends on K other factors, a change in w_i or any of the other K factors will result in a new draw of W_i from $F(W)$. We explore two settings: one in which the interactions are equally likely across functional boundaries (integrated problem structure), and one in which the complexity lies mostly within the functional disciplines (modular problem structure). In the former case, the source of the K interactions is randomly chosen; that is, the performance contribution W_i depends on K randomly determined factors w_j $(j \neq i)$. In the latter case, we control the locus of interactions, and K_1 is the number of interactions (randomly chosen) within the functional discipline (say W_i depends on K_1 other factors from the $w_1, ..., w_5$ set), whereas K_2 is the number of interactions across the two functional disciplines (again W_i depends on K_2 other factors from the $w_6, ..., w_{10}$ set).

4.2 The Search Process

In the beginning of the problem-solving process, we 'assign' each individual an initial solution or point in the landscape. As discussed above, the degree of difference among the starting points captures the amount of cognitive diversity. On one extreme, all members have the same random starting point (no diversity); on the other extreme, all members have different, randomly generated starting points (total diversity).

In the **nominal group**, an individual builds in each idea generation step on his prior ideas by changing one randomly chosen factor at a time.[9] In the simulation, each individual builds with a large probability ($\alpha = 70\%$) on his best idea; otherwise, he builds on any of the last three ideas he generated. Individuals stop if they run out of time or if they run out of ideas, that is, if they cannot generate new ideas from any of the ideas in their memory. Note that individuals starting from the

results between specialized and cross-functional settings, we limit ourselves to mutually exclusive expertise sets, in order to present the sharpest contrast. Note that we would obtain results similar to a (lower dimensional) specialized problem, if all 8 individuals belonged to the same functional group. Since it is obviously better to include members of all functional areas concerned, we do not discuss this case further.

[9] In reality, individuals may change more than one aspect of a solution at a time. However, varying a larger subset of factors together with a larger solution landscape (larger N or larger $\mathcal{L}$) would be computationally more cumbersome, without altering the basic insights.

same or a similar point also generate many identical ideas, since they search the same area in the solution landscape; in contrast, those starting from very different points tend to generate few or no identical ideas. Thus, cognitive diversity affects both the number of different ideas individuals start out with, and the number of different ideas generated in the search process.

In the **brainstorming group**, due to colocation, individuals build on each other's ideas. For the purpose of our model, the first person proposing a solution determines the group's initial point. This point is randomly chosen among the starting points assigned to the individuals at the outset.[10] Like in the nominal groups, during the brainstorming session, individuals generate ideas by varying one randomly determined factor at a time. They again build either with 70% probability on the best idea found along the current search path, or on one of the ideas mentioned *by any group member* in the last period. To incorporate production blocking, we restrict the number of ideas mentioned in each period to a subgroup $M' = 4$ of the ideas generated by the individuals. Further, we capture evaluation apprehension (individuals with better ideas are more likely to speak up) by ensuring that at least one of the M' individuals mentions an idea which improves on the current best idea, as long as someone finds an improvement. If the group runs out of ideas along a given search path, they start from a different starting point of another randomly chosen individual (if such a different point exists). Therefore, groups stop if they either run out of time or have explored all starting points of the individuals.[11]

The outcome of the nominal group is the best solution found across all M individual search efforts. The outcome of the brainstorming group is the best solution found by the group during its search effort. We finally compare the best performance found by the brainstorming group to the

[10] In reality, given the logic of evaluation apprehension, individuals with better starting points might be more likely to speak up first. Since it is not clear which rule (best one?, k^{th} order statistic?) would adequately capture this in the numerical experiments, we show only the results based on randomly determined starting points. Our sensitivity analysis showed that having group members mention their ideas in decreasing order of solution quality would improve the results of the brainstorming groups; in that light, the simulation results presented here might somewhat overstate the relative performance of the nominal groups, but the qualitative insights remains unchanged.

[11] Given our limit on the number of periods, the latter happens typically only if there is only one starting point.

best performance found by any of the individuals in the nominal search.

5. Brainstorming or Nominal? Effective Problem-Solving Strategies

In this section, we analyze the relative effectiveness of the two different organizational configurations. We first present the results for a specialized problem, where all individuals can vary all N factors given that they come from the same functional competence background. We then turn to the results for a cross-functional problem. The numerical results presented below are an average over 100 different landscape realizations and 100 runs within each landscape.

5.1 Results for Specialized Problems

We consider a problem residing within a specialized knowledge area: all individuals can modify all solution factors w_i. Such a setting could represent a problem within an organizational function.

The graph in Figure 1 shows the performance advantage of nominal groups over brainstorming groups. If the problem is very simple ($K = 0$), the global optimal solution can be found by simply varying one factor at a time. Both the brainstorming and nominal groups perform equally well; that is, for the given time horizon and number of possible solutions, both organizational forms tend to identify the best solution. Overall, in this special case, given enough time, even one individual alone can come up with the best idea, while saving the cost of employing several individuals.

As the number of interactions K among the factors increases, so does the problem complexity. As evident in many realistic contexts, a complex problem may admit a number of "good" solutions (i.e., local optima) in the solution space. The brainstorming group generates, on average, fewer ideas than the equivalent nominal group. However, more complexity in the landscape drives an exponentially higher number of "good" solutions (local peaks). Therefore, the total number of ideas generated becomes the main driver for the quality of the best solution found. Overall, the advantage of nominal groups over brainstorming groups increases with the problem complexity.

The advantage tends to grow larger if the individuals have diverse starting points. This diversity allows the nominal group to search an even larger and less correlated space of the solution landscape, while production blocking prevents the brainstorming group from equally benefiting from this diversity.[12]

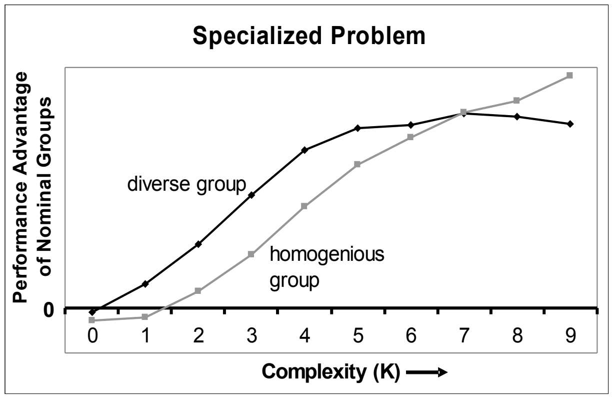

Figure 1: Relative performance of nominal over brainstorming groups in specialized problems. (All figures use an average over 100 runs in a landscape and 100 landscapes for N=10, M=8, M'=4, T=15, prob. to build on the best solution = 0.7.)

The sensitivity of our results to the rest of the parameters of the problem (not explicitly shown in the figure) follows intuitively: As the production blocking effect diminishes (M' increases), the advantage of the nominal groups decreases. However, their advantage is never overcome (see Section 5.2 for a more detailed discussion). In the absence of evaluation apprehension, the advantage of nominal groups further increases. Evaluation apprehension de facto reduces the production blocking effect for the brainstorming groups by enabling a subset of solutions (M') with better-than-average solution qualities. Similarly, the advantage of nominal groups increases if individuals are likely to build only on the best solution rather than on other recent ideas (α increases). The reason is that brainstorming groups rely solely on such departures from the main path of advancement to broaden their search path, while nominal groups already search a broader set of solutions due to

[12]Note that for very high K, the advantage of homogenous nominal groups is larger than that of diverse groups, not because homogenous nominal groups perform better, but rather because homogenous brainstorming groups explore only one path in which they get quickly stuck, and, therefore, they perform even poorer. Recall that individuals in nominal groups explore different paths even when they have the same starting point. See discussion after the analytical results.

their ability to search different paths in parallel.[13]

The experimental results can be analytically shown for two extreme cases: a smooth solution space with no interactions ($K = 0$), and an extremely complex problem with the maximum level of interactions ($K = 9$). For $K = 9$, a change in one factor w_i changes all performance contributions W_j; therefore, the overall solution performance is a new random draw from the convolution of the individual performance multivariate distributions. Although individuals build on prior solutions, and, therefore, they may generate ideas in a particular order, when $K = 9$, the order is irrelevant for determining the expected value of the highest quality solution. Eventually, the expected value will depend only on the number of draws.

Let $G(V)$ denote the distribution function of $V(\cdot)$ (i.e., the aforementioned convolution). Similar to the distributions assumed in the simulation, we assume that $G(V)$ has a positive support. To keep the notation simple, let I and I' denote the total number of different ideas explored by the nominal and the brainstorming group, respectively. Note that, the larger the degree of production blocking, the less ideas are generated by the brainstorming groups, and thus the smaller I'. Further, the cognitive diversity of the participating individuals affects the number of starting points. It also affects the number of different ideas generated during the search process, since search areas overlap to a smaller or larger degree. Thus, diversity results in an increase in I, but may or may not increase I', depending on whether production blocking limits the group from exploring further ideas.

Proposition 1 *For the specialized problem, the expected advantage $\Delta = E[V^{NG}] - E[V^{BG}]$ of the nominal group over the brainstorming group for $K = 9$ is positive and :*

(a) increases as the degree of production blocking increases (I' decreases);

(b) increases as the degree of diversity of the individuals increases, if production blocking prevents

[13] Note that for $\alpha = 0$, the idea generation process would represent a random walk. In this case, nominal groups will again outperform brainstorming group, since only the total number of generated ideas matters.

the group from benefiting from it (only I increases), and decreases if both I and I′ increase equally.

Proof. For detailed proofs, see the Appendix.

It can be shown that the advantage of nominal groups is given by $\Delta = \int_0^\infty (G^{I'}(V) - G^I(V))dV$. Only if $I'=I$ will nominal and brainstorming groups perform equally well; with production blocking ($I' < I$), $\Delta \geq 0$, and thus nominal groups outperform brainstorming groups. Further, it is easy to show that Δ decreases in I' (and thus increases in the degree of production blocking) and increases in I (and thus increases in the diversity of the individuals, if I' is unaffected). Instead, if both I and I' increase equally, that is, if both nominal and brainstorming groups benefit from the diversity equally, then the advantage Δ decreases. The last observation offers an explanation for our numerical results for $K = 9$, where homogenous brainstorming groups run out of potential paths to pursue before the end of the T periods; hence, the lack of diversity, rather than production blocking, limits their performance.

For $K = 0$, there are no interactions between the influence factors and $V(\vec{w}) = \sum_{i=1}^{N} W_i(w_i)$. Thus, for $K = 0$, a change in one factor w_i only results in a change in the performance contribution W_i. Due to this independence, the expected value of the brainstorming group's performance depends only on the number of ideas generated along each dimension, and not on the order in which these ideas are being generated. In case of independence, it is optimal for the nominal groups to determine the best idea for each influence factor separately (pick the best w_i based on the performance contribution W_i) and to combine them into a new solution. If nominal groups recombine their solutions in this way, their expected performance also depends only on total number of ideas generated along each dimension. Let I_i and I'_i denote the total number of ideas generated along factor w_i by nominal groups and brainstorming groups, respectively (with I and I' representing the total number of ideas generated). Then, the best solution of the nominal groups is given by: $Max(V) = Max_{I_1}(W_1) + Max_{I_2}(W_2) + ... + Max_{I_N}(W_N)$ (and similarly for the brainstorming groups). If, as in the simulation, a draw along a particular performance dimension is equally likely in the nominal and the brainstorming group, the following holds:

Proposition 2 *For the specialized problem, the expected advantage* $\Delta = E[V^{NG}] - E[V^{BG}]$ *of the nominal group over the brainstorming group for* $K = 0$ *is positive and:*

(a) increases as the degree of production blocking increases (I' *decreases);*

(b) increases as the degree of diversity of the individuals increases, if production blocking prevents the group from benefiting from it (only I *increases), and decreases if both* I *and* I' *increase equally.*

The proof is based on a sample path argument (see the Appendix for details). Note the difference to the simulation results: In the simulation, brainstorming groups perform equally well or even slightly better when $K = 0$. The reason is that, in the simulation, we assume that the nominal and brainstorming groups ignore the fact that the dimensions are independent; that is, they ignore that none of the parameter changes resulted in a value change for other parameters. This results in a large number of possible parameter changes before an individual finds the best solution. The brainstorming group, on the other hand, builds on the best solution available in every step; in addition, it generates several ideas to pick from in each period, thereby moving up the solution landscape quicker. Therefore, it is more likely to find the peak in fewer steps; thus, it might outperform the nominal group whenever the latter does not find the peak in the given time limit. However, there is a second reason for the discrepancy. Even if nominal groups recombine their best parameter choices, nominal groups do not outperform brainstorming groups in the simulation (results available from the authors). The reason is that, in the simulation, the set of solutions ($N = 10$ or $\mathcal{L} = \{0, 1\}$) is limited relative to the given time; hence the brainstorming group tends to find the best possible solution. With a sufficiently large set of solutions or very limited time, production blocking and diversity would become again relevant, and nominal groups that recombine their solutions would outperform brainstorming groups for $K = 0$.

5.2 Results for Cross-Functional Problems

Next, we turn to the analysis of a cross-functional problem, which requires input of two functional areas. This implies that experts are able to recognize, assess and manipulate the impact of factors that non-experts do not know about or understand. As discussed above, we consider two competence groups with four members each, and assume that each group can manipulate five factors. In addition, we analyze two forms of cognitive diversity: completely different starting points (total diversity), versus starting points that differ only along the parameters that individuals are knowledgeable about (partial diversity). The former represents settings where the participants from one functional field have diverse beliefs about the parameters from the other function, whereas the latter implies a common belief about those parameters.

The graph in Figure 2 shows that the performance of the nominal groups, relative to that of brainstorming groups, again increases with the underlying complexity of the problem.[14] However, the performance difference does not preserve the same sign for all levels of complexity; for low-to-medium complexity, brainstorming groups outperform nominal ones. Their ability to build on solutions proposed by individuals with different functional backgrounds enables the co-located team members to acquire a holistic perspective of the problem and to find superior performance solutions. However, their superiority does not hold across the entire range of complexity levels. As in the specialized problem, increased complexity pronounces the production blocking effect, and, eventually, nominal groups outperform brainstorming groups. In other words, in very complex problems, it is more important to generate many ideas than to build on each other's ideas. For high K, neighboring ideas exhibit a very low correlation among their performances; therefore, building on each other's ideas is almost identical (for $K = 9$ identical) to new random draws from the solution landscape. Hence, only the number of ideas generated affects who performs better.

[14]The lower difference at $K = 0$ over $K = 1$ is again due to the fact that the set of solutions is limited; thus, for a very simple landscape ($K = 0$), both tend to find decent solutions (though the individuals in nominal groups only within the set of factors they are knowledgable about).

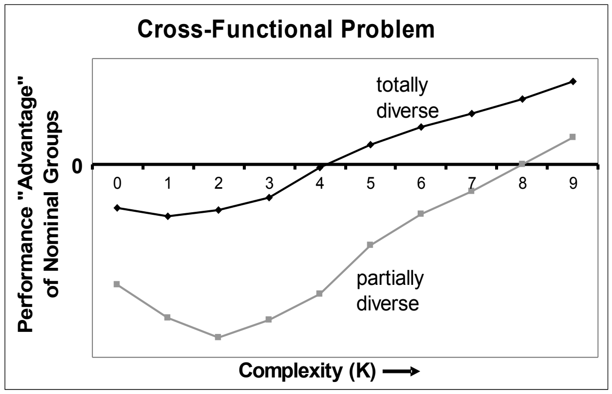

Figure 2: Relative performance of nominal over brainstorming groups in cross-functional problems

The nominal groups reduce the performance gap faster when the team members exhibit total diversity, since the latter enables the individuals to explore radically different solution concepts. The managerial insight from our observation is important. It points to the value of breaking down the stereotype impressions that one area may have about the other area, potentially through different project assignments.

The effects of both production blocking and evaluation apprehension present patterns similar to the ones observed for specialized problems, as shown in Figure 3. The figure illustrates that production blocking and evaluation apprehension may have different effects on the quality of the best idea. (It is well-established in the social psychology literature that both phenomena limit the total number of ideas generated). The underlying reason rests with the fact that evaluation apprehension affects the quality of the subset of mentioned ideas. At the same time though, another important insight emerges concerning the effect of "topic fixation", which does not allow the performance difference to collapse to zero, even without production blocking present. In other words, the fact that groups follow a common train of thought has implications for their performance when combined with underlying problem complexity.

So far, we have assumed that the interactions between the solution performance drivers exhibit a random pattern. However, the association of subproblems with functional areas is typically not random. For example, one could posit that a consumer behavioral trait (e.g., exclusivity-

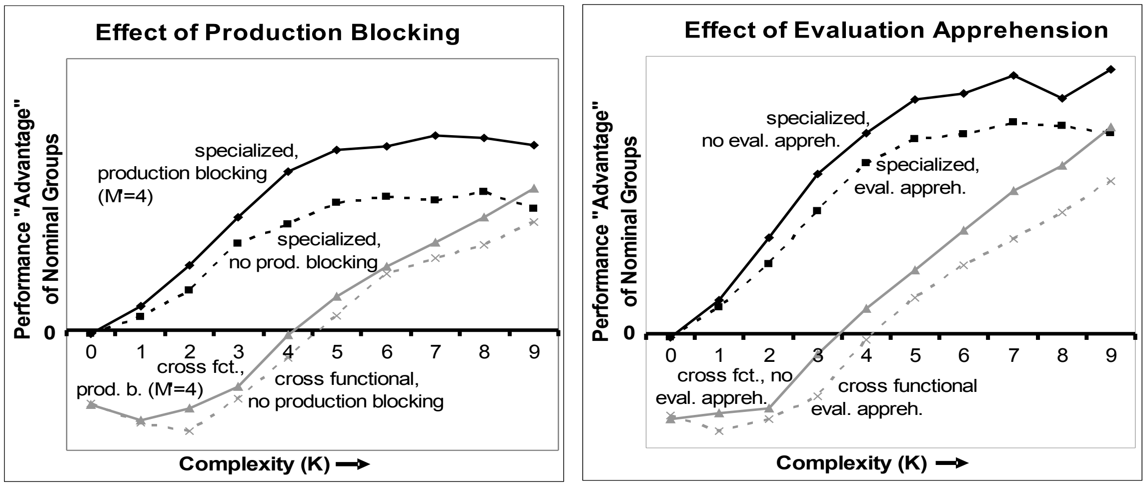

Figure 3: Sensitivity analysis of relative performance of nominal over brainstorming groups in cross-functional problems

seeking) that influences the market acceptance of a proposed design interacts more with other marketing factors (say an advertising campaign plan that promotes status association) than with production process attributes (e.g., delivery lead-times).[15] For that reason, we explore how the distribution of the interactions across the two functional areas affects the robustness of our results. In Figure 4, $K1$ denotes the number of interactions among the performance factors within the functional competence, and $K2$ the number of interactions across the two sets of factors w_i. Figure 4 again plots the relative performance of nominal over brainstorming groups. Below the graph, we represent examples of the interactions in the form of a design structure matrix, frequently used in NPD to illustrate the equivalent product architecture.

As shown in Figure 4, to determine the overall solution for the nominal groups, first pick the best solution in each subgroup (each functional area) as is (ignoring the effects of any cross-functional interactions), and then combine the solution dimensions of the first subgroup with those of the second subgroup.[16] We present this variant of the nominal group technique for two reasons: we can identify conditions under which it is beneficial, and we can verify the simulation results

[15] A similar association between problem structure and team structure has been made earlier, e.g., Sosa et al. 2004.

[16] We show the results for partial diversity, since the choice of the best solution within a functional area is better defined (all use the same assumptions about the other functional area). The sensitivity of the results to the specifications is qualitatively the same as before.

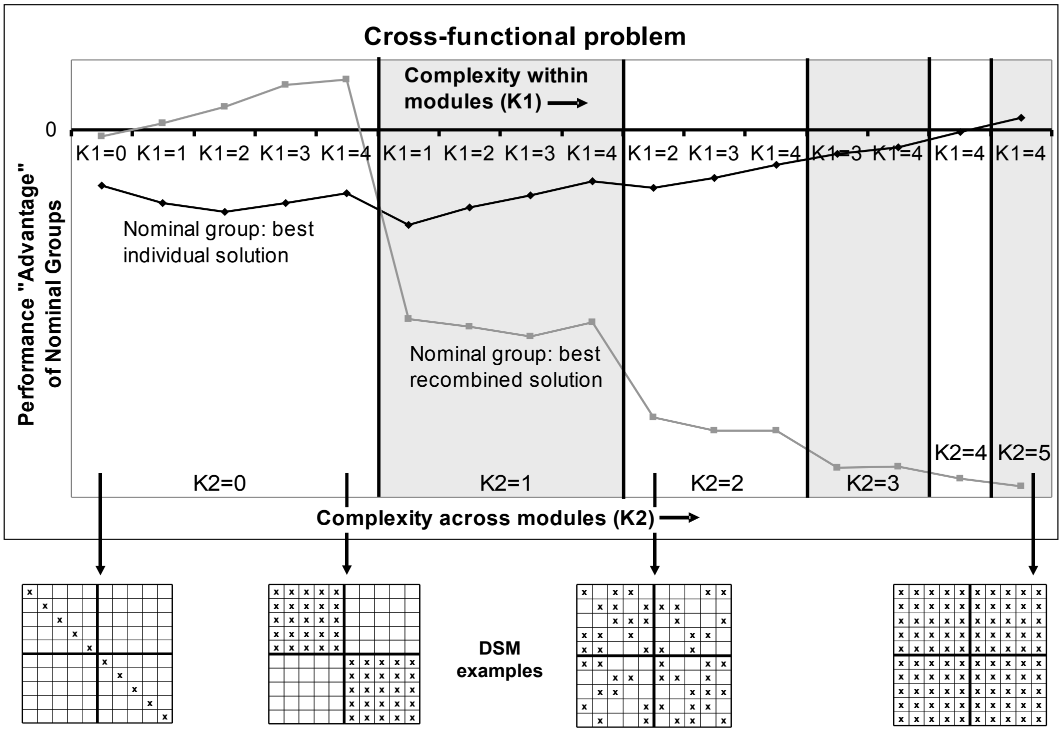

Figure 4: Relative performance of nominal over brainstorming groups with modular interactions (K1 interactions within and K2 interactions across functional areas) and partial diversity

analytically for a scenario without cross-functional interactions ($K1 = 4$ and $K2 = 0$).

The simulation results demonstrate the importance of the locus of complexity on the relative performance of the two organizational settings. For the same level of complexity (same $K = K1 + K2$), the benefits of group brainstorming are larger if most interactions are within a competence domain (larger $K1$). Hence, brainstorming groups seem to perform better in more modular problems. One explanation is that a more modular problem structure reduces the difficulty of building on each other, since interactions are more confined. For example, it is easier to solve a problem if decision w_1 influences the performance of w_2 and w_2 influences the performance of w_1 than if w_1 influences the performance of w_2 and w_2 influences the performance of say w_7.

Further, we find that in modularized problem structures, nominal groups can outperform brainstorming groups in cross-functional problems in two scenarios: The first scenario is that of extremely high complexity ($K = K1 + K2 > \overline{K}$), so that building on others' ideas is equivalent to random ideas and has no value. Here, our result for $K1 = 4$, $K2 = 5$ mirrors our earlier result for $K = 9$. The second scenario involves very modular problems with low, or rather no interactions across functional competence areas ($K2 = 0$). Here, nominal groups can obtain an advantage by choosing the best solutions from each functional area and combining them. Note that this holds independent of the level of complexity within each functional area (for all $K1$). This latter scenario differs from our earlier results for random interactions, where brainstorming groups always had an advantage in cross-functional problems of low and medium complexity. Note, however, that brainstorming groups remain the dominant organizational structure for most realistic settings: As long as there is the slightest level of interaction across functional knowledge domains, and as long as the problem is not extremely complex, brainstorming groups are at an advantage.

As before, we complement the experimental results with analytical results. Note first that for $K = 9$ ($K1 = 4$ and $K2 = 5$), the results of Proposition 1 also apply (details not provided). Further, we can provide results for one specific case, $K1 = 4$ and $K2 = 0$, under the premise that nominal groups recombine their solutions. Since $K1 = 4$, a change in one factor w_i changes all performance contributions W_j belonging to the same functional area. Therefore, the total contribution (across the five factors) is again a random draw from the convolution of the five individual performance distributions. Let $G_i(V_i)$ denote the distribution function of the aforementioned convolution of functional group i, ($G_i(V_i)$ has a positive support). Since $K2 = 0$, there are no interactions between the two knowledge domains, i.e., $V(\vec{w}) = V_1 + V_2$.[17] Given the outlined structure, Corollary 1 repeats the logic of Proposition 1 for the functional areas.

Corollary 1 *For the cross-functional problem with a modular problem structure and* $K1 = 4$ *and* $K2 = 0$*, the results of Theorem 1 (*$\Delta = E[V^{NG}] - E[V^{BG}] \geq 0$ *and the sensitivity results) apply.*

These results confirm our findings from the simulations; namely, nominal groups can also

[17] The main qualitative results of the analytic part hold for more general functions, as long as the performance contributions are independent, e.g., for multiplicative performance function $V = W_1^{\alpha} W_2^{\beta}$. Details available from authors.

outperform brainstorming groups for cross-functional problems, if there are either no (or few) interactions across functional competence areas ($K2 = 0$) and ideas are recombined, or if the problem is extremely complex ($K = 9$). The following corollary shows that not only does the degree of production blocking ($I - I'$) matter, but it also matters how the production blocking is distributed across the functional areas.

Corollary 2 *Assume a given level of production blocking (given $I' = I'_1 + I'_2$) and two equally important problem dimensions (same improvement potential or $G_i(V)$). The performance advantage of the nominal groups grows larger, as the production blocking effect becomes more heterogenous (that is as I'_1 and I'_2 deviate from $\frac{I'}{2}$).*

Our results suggest that the simulation results would be worse for the brainstorming group under a more one-sided production blocking effect. Thus, it calls for managerial attention, as it implies that a 'balanced' participation from each functional area can improve the brainstorming group's performance (e.g., equal number of individuals and a trained facilitator, to ensure that members with different knowledge competencies participate approximately equally).

6. Discussion and Conclusion

In this study, we develop a formal model to analyze the effectiveness of different organizational structures for problem-solving during the ideation phase of new product development (NPD). Since the late 1950s, the brainstorming group method has found both strong support and fierce opposition as to its effectiveness in problem-solving. Numerous academic studies have questioned whether a brainstorming group is more effective compared to the same number of individuals attempting to find a solution in isolation, a.k.a. *nominal* group. Literature in social psychology identifies several contextual factors, primarily *production blocking*, that impede the performance of co-located brainstorming teams (e.g., Stroebe and Diehl 1994, Paulus *et al.* 1996). On the opposite end of the spectrum, the organizational theory literature has recorded brainstorming as very effective in 'real' contexts (e.g., Sutton and Hargadon 1996, Hargadon 2003).

We model idea generation as a search process and we account explicitly for two contextual features: First, the problem structure, as manifested by the complexity and the degree of cross-functionality of the problem; second, the diversity among the team members, as attested by their ability to explore different parts of the solution landscape. We compare the expected best solution obtained by the same number of individuals working either separately (nominal group) or together (brainstorming group) to solve the same problem.

Our results show that there is no dominant organizational structure. In that regard, we provide a possible explanation for the earlier contradictory findings. Due to feasibility challenges, experimental work considers either very simple problems with input from only one functional expertise (e.g., "benefits or difficulties from an extra thumb on each hand"), or very complex problems where many factors interact in non-linear ways and create a multitude of uncorrelated solutions (e.g., "solve the campus parking problem"). Our results agree with their findings that nominal groups perform best in these settings. The organizational theory literature, on the other hand, looks at case studies of moderately complex settings which require input from various functional areas; our results support their claim that brainstorming groups perform best in these settings.

Our analysis reveals a few important managerial insights. First, nominal groups always obtain better solutions to specialized problems. However, once we shift to the organizational reality of cross-functional problems, brainstorming groups can exploit the competence diversity of their members and their ability to "build on ideas" of others to generate better solutions. Still, this diversity-driven advantage washes out once the factors that determine the quality of the resulting solution exhibit complex interactions and relationships. In addition, the locus of complexity plays a role: complexity within a functional competence domain provides brainstorming groups with a greater advantage compared to complexity across competence domains.

We further show that the brainstorming group's effectiveness shrinks with higher degrees of production blocking and cognitive diversity; yet, it may increase in the presence of evaluation apprehension. While both evaluation apprehension and production blocking may pose similar limitations regarding the total number of ideas generated, evaluation apprehension may, in fact, help by improving the solution quality of the ideas mentioned. As such, management should reflect to what extent evaluation apprehension should be at the center of countermeasures.

Finally, our findings suggest that the difference between nominal and brainstorming groups may not be solely due to production blocking and evaluation apprehension. Instead, we show another basic trade-off between the benefit of building on each other's ideas (which allows a diverse group to attain solutions that members may not reach individually), and the topic fixation disadvantage (which prevents the group from reaching out to different parts of the solution landscape).

Figure 5 summarizes the results for the cross-functional problems.

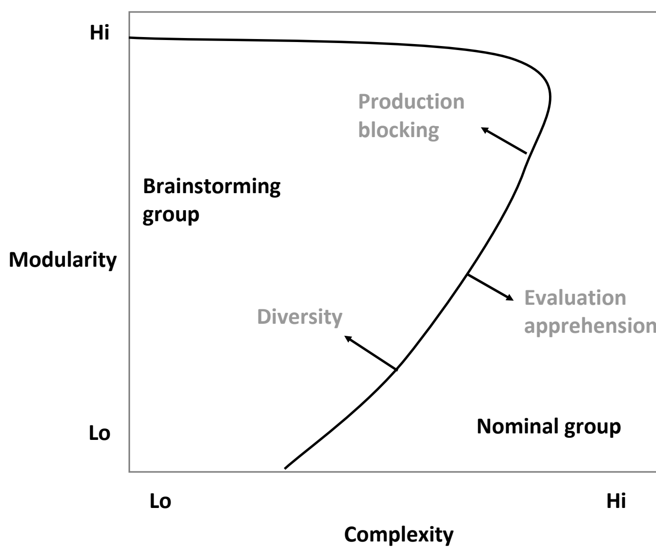

Figure 5: Comparison of brainstorming and nominal groups

The results lend themselves naturally to hypotheses which could be tested within the context of a lab experiment. We leave this aside for future research. As in any abstract modeling effort, our setup presents limitations. We chose to 'group' all dimensions of personality diversity defined in psychology and organizational behavior together in two meaningful sets: cognitive diversity (similarity of ideas generated), and functional background diversity (the capability to manipulate certain performance drivers). In addition, we do not analyze the nominal group technique (Robbins and Judge 2007). Instead, we chose to follow the social psychology literature in studying the two organizational configurations separately, in order to isolate the effect of contextual parameters on their effectiveness. We are aware that our effort is only a stepping stone to further understanding of the ideation stage of the NPD process. However, it offers a coherent theoretical framework that brings together seemingly opposing perspectives, and adds important managerial insights.

7. References

Bouchard, T.J., M. Hare. 1970. Size, performance, and potential in brainstorming groups. *Journal of Applied Psychology* 54 51-55.

Bowers, N.L., H.U. Gerber, J.C. Hickman, D.A. Jones, C.J. Nesbitt. 1986. *Actuarial Mathematics.* First Edition. Society of Actuaries: Itasca, Illinois.

Clark, K.B., T. Fujimoto. 1989. Lead time in automobile product development explaining the japanese advantage. *Journal of Engineering and Technology Management* 6(1) 25-56.

Cottrell, N.B. (1972). Social facilitation. In C. McClintock (ed.), *Experimental Social Psychology* (pp. 185-236). New York: Holt, Rinehart & Winston.

Dahan, E., H. Mendelson. 2001. An extreme-value model of concept testing. *Management Science* 47(1) 102-116.

Diehl, M., W. Stroebe. 1987. Productivity loss in brainstorming groups: Toward the solution of a riddle. *Journal of Personality and Social Psychology* 53 497-509.

Diehl, M., W. Stroebe. 1991. Productivity loss in idea-generating groups: tracking down the blocking effect. *Journal of Personality and Social Psychology* 61(3) 392-403.

Dougherty D. 1992. Interpretive barriers to successful product innovation in large firms. *Organization Science* 3(2) 179-202.

Erat, S., S. Kavadias. 2008. Sequential testing of product designs: Implications for learning. *Management Science* 54(5) 956-968.

Fisher, F., M.D. Fisher. 1998. *The Distributed Mind: Achieving High Performance Through the Collective Intelligence of Knowledge Work Teams.* New York: AMACOM.

Fleming, L. 2001. Recombinant uncertainty in technological search. *Mgmt Sci.* 47(1) 117-132.

Fleming, L., J. Singh. 2009. Lone inventors as sources of breakthroughs: Myth or reality? *Management Science*, forthcoming.

Fleming, L., O. Sorenson. 2004. Science as a map in technological search. *Strategic Management Journal* 25(8-9) 909-928.

Galambos, J. 1978. *The Asymptotic Theory of Extreme Order Statistics.* John Wiley & Sons.

Gallupe, R.B., L.M. Bastianutti, W.H. Cooper. 1991. Unblocking brainstorms. *Journal of Applied Psychology* 76 137-142.

Gavetti, V., D. Levinthal. 2000. Looking forward and looking backward: Cognitive and experiential search. *Admin. Sci. Quart.* 45 113-137.

Girotra K., C.Terwiesch, K.T. Ulrich. 2009. Idea generation and the quality of the best idea. *INSEAD working paper 2009/32/TOM.*

Gundry, L., L. La Mantia. 2001. *Breaksthrough Teams for Breakneck Times: Unlocking the Genius of Creative Collaboration.* Chicago: Dearborn Trade.

Hargadon 2003. *How Breakthroughs Happen: The Surprising Truth About How Companies Innovate*, HBS Press, Cambridge, MA.

Hong, L., S. Page. 2001. Problem solving by heterogeneous agents. *Journal of Economic Theory* 97, 123-163.

Kahneman, D., D.L. Fredrickson, C.A. Schreiber, D.A. Redelmeier. 1993. When more pain is preferred to less: Adding a better end. *Psychological Science* 4 401-405.

Kahneman, D. 2000. Evaluation by moments: Past and future. In: Kahneman, D., A. Tversky, A. (eds.) 2000. *Choices, Values and Frames.* New York: Cambridge University Press and the Russell Sage Foundation, New York, 293-308.

Kauffman, S.A. 1993. *The Origins of Order.* Oxford University Press, New York.

Kauffman, S.A., J. Lobo, W.G. Macready. 2000. Optimal search on a technology landscape. *J. Econom. Behavior Organ.* 43(2) 141-166.

Larey, T.S., P.B. Paulus. 1999. Group preference and convergent tendencies in small groups: A content analysis of group brainstorming performance. *Creativity Research Journal* 12(3) 175-184.

Libby R., K.T. Trotman, I. Zimmer. 1987. Member variation, recognition of expertise and group performance. *Journal of Applied Psychology* 72(1) 81-87.

Loch, C. H., C. Terwiesch, S. Thomke. 2001. Parallel and sequential testing of design alternatives. *Management Science* 47(5) 663-678.

Mihm, J., C.H. Loch, A. Huchzermeier. 2003. Problem-solving oscillations in complex engineering projects. *Management Science* 49(6) 733-750.

Milliken, F., L. Martins. 1996. Searching for common threads: Understanding the multiple effects of diversity in organizational groups. *Academy of Management Review* 21 402-433.

Mullen, B., C. Johnson, E. Salas. 1991. Productivity loss in brainstorming groups: A meta-analytic integration. *Basic and Applied Social Psychology* 12(1) 3-23.

Naquin, C.E., R.O. Tynan. 2003. The team halo effect: Why teams are not blamed for their failures. *Journal of Applied Psychology* 88 332-340.

Offner A.K., T.J. Kramer, J.P. Winter. 1996. The effects of facilitation, recording and pauses upon group brainstorming. *Small Group Research* 27 283-298.

Osborn, A.F. 1953. *Applied Imagination: Principles and Procedures of Creative Thinking.* New York, NY: Charles Scribner's Sons.

Paulus, P.B. 2000. Groups, teams, and creativity: The creative potential of idea-generating groups. *Applied Psychology: An International Review* 49(2) 237-262.

Paulus, P.B., V. Brown, A.H. Ortega. 1996. Group creativity. In: R.E. Purser and A. Montuori (eds.) Social Creativity in Organizations. Cresskill, NJ: Hampton Press.

Paulus, P.B., T.S. Larey, A.H. Ortega. 1995. Performance and perceptions of brainstormers in an organizational setting. *Basic and Applied Social Psychology* 17 249-265.

Pelled, L.H., K.M. Eisenhardt, K.R. Xin. 1999. Exploring the black box: An analysis of work group diversity, conflict, and performance. *Administrative Science Quarterly* 44 1-28.

Robbins S.P., Judge T.A. 2007. *Organizational Behavior.* Prentice Hall.

Sawyer, R.K. 2007. *Group Genius: The Creative Power of Collaboration.* Chapter 4. Cambridge, MA: Basic Books, Perseus Books Group.

Simon, 1969. *The Science of the Artificial.* 2nd. ed., Boston: MIT Press.

Stroebe, W., M. Diehl. 1994. Why are group members less effective than their members: On productivity losses in idea-generating groups. *European Review of Social Psychology* 5 271-303.

Sosa, M., S. Eppinger, C. Rowles. 2004. The misalignment of product architecture and organizational structure in complex product development. *Mgmt. Sci.* 50(12) 1674-1689.

Sutton, R.I., A. Hargadon. 1996. Brainstorming groups in context: Effectiveness in a product design firm. *Administrative Science Quarterly* 41 685-718.

Taylor D.W., P.C. Berry, C.H. Block. 1958 Does group participation when using brainstorming facilitate or inhibit creative thinking? *Administrative Science Quarterly* 3 23-47.

Taylor, A., H.R. Greve. 2006. Superman or the fantastic four? Knowledge combination and experience in innovative teams. *Academy of Management Journal* 49(4) 723-740.

Terwiesch, C., Y. Xu. 2008. Economic models for open innovation systems and multi-agent problem solving. *Management Science.* 54(9) 1529-1543.

Thomke, S.H. 1998. Simulation, learning and R&D performance: Evidence from automotive development. *Research Policy.* 27 55-74.

Thomke, S.H. 2003. *Experimentation Matters: Unlocking the Potential of New Technologies for Innovation.* Harvard Business School Press.

Wanous J.P., M.A. Youtz. 1986. Solution diversity and the quality of group decisions. *Academy of Management Journal* 29(1) 149-159.

Weitzman, M.L. 1979. Optimal search for the best alternative. *Econometrica* 47 641-654.

Wheelwright, S.C., K.B. Clark. 1992. *Revolutionizing Product Development*, Chapter 2 and 3. The Free Press, NY.

Yetton P.W., P.C. Bottger. 1982. Individual versus group problem solving: An empirical test of a best member strategy. *Organizational Behavior and Human Performance* 29(3) 307-321.

Appendix: Proof of Propositions

Proof of Proposition 1 :

Since all draws are made from the same distribution, the cdf of $max(V^{NG})$ is given by $G^I(V)$ and that of $max(V^{BG})$ is given by $G^{I'}(V)$. Further, since the pdfs have a positive support, the expected values are given by $E[max(V^{NG})] = \int_0^\infty (1 - G^I(V))dV$ and $E[max(V^{BG})] = \int_0^\infty (1 - G^{I'}(V))dV$ respectively (Bowers *et al.* 1986, p.62). Therefore, $\Delta = E[max(V^{NG})] - E[max(V^{BG})] = \int_0^\infty (G^{I'}(V) - G^I(V))dV$. Since $G(V)$ is the cdf, $G(V) \leq 1$ and, hence, for all $I' \leq I, \Delta \geq 0$. Thus, nominal groups outperform brainstorming groups.

Further, it is straightforward to show that Δ decreases in I' (part a) and increases in I as long as I' remains fixed (part b1). If the diversity results in an increase of both I and I' by the same number of ideas i (part b2), then $\Delta = \int_0^\infty G^i(V)(G^{I'}(V) - G^I(V))dV \leq \int_0^\infty (G^{I'}(V) - G^I(V))dV$. Thus, if both nominal and brainstorming groups benefit equally from the increase in diversity, the advantage of nominal groups decrease.

Proof of Proposition 2 :

For $I = I'$ (no production blocking) $\Delta = 0$. Hence, we only need to show that Δ decreases in I' to show that nominal groups perform at least as well as brainstorming groups (see part a).

Proof of Part (a): Since $E[V^{NG}]$ is not affected by I', all we need to show is that $E[V^{BG}]$ increases in I'. Consider a fixed value of I' and a specific sample path with I'_i out of I' draws being taken along dimension w_i. The expected value of this sample path is given by $V_0 = \sum_i \int_0^\infty [1 - F_i^{I'_i}(W_i)]dW_i$, which, in turn, is realized with the probability of the path being taken (say q).

Now, let us consider what happens if production blocking is reduced and one more idea is mentioned, that is the number of ideas is $I' + 1$. The additional idea can be generated along any dimension w_i with probability p_i (in our simulation $p_i = \frac{1}{N}$) and achieve an expected value of $V_i = \sum_{j \neq i} \int_0^\infty [1 - F_j^{I'_j}(W_j)]dW_j + \int_0^\infty [1 - F_i^{I'_i+1}(W_i)]dW_i$. The new expected value of this sample path (with the additional idea being voiced) is given by $\sum p_i V_i$. Since all $V_i > V_0$, this new expected value is greater than V_0. Note that the probability of this sample path occurring is still q; the chance of drawing from one of the N dimensions in the $(I' + 1)^{st}$ contribution step has been accounted for in the expected value. Since this holds true for any I' and any I'_i (any sample path for the first I' draws), $E[V^{BG}]$ increases in I' and hence the advantage of nominal groups Δ is decreasing in I'.

Proof of Part b (constant I'): $E[V^{BG}]$ is not affected by the value of I; all we need to show is that $E[V^{NG}]$ increases in I. With a similar argument as in Part (a), it is easy to show that $E[V^{NG}]$ and hence Δ increases in I and hence in the degree of diversity of the individuals, if I' remains unchanged (brainstorming groups cannot take advantage of it due to production blocking).

Proof of Part b (I and I' affected equally): For expositional simplicity, we show this part of the proof for two dimensions and one additional draw only. The argument remains qualitatively the same for higher dimensions, but it is combinatorially complex to write formally. Consider again a particular sample path with I_i and I'_i draws along dimension i for the nominal and brainstorming groups respectively. Then, the expected performance difference is $\Delta_0 = \int_0^\infty (F_1^{I'_1}(W_1) - F_1^{I_1}(W_1))dW_1 + \int_0^\infty (F_2^{I'_2}(W_2) - F_2^{I_2}(W_2))dW_2$. If the diversity increases so that one additional idea is being generated, the idea is being generated along dimension i with probability p_i.

Consider first the cases, in which both the nominal and the brainstorming group generate a new idea along the same dimension i, which occurs with probability p_i^2: The performance difference for these two cases is given by $\Delta_1 = \int_0^\infty F_i(W_i)(F_i^{I'_i}(W_i) - F_i^{I_i}(W_i))dW_i + \int_0^\infty (F_j^{I'_j}(W_j) - F_j^{I_j}(W_j))dW_j$. $\Delta_1 \leq \Delta_0$ for both cases since $F_i \leq 1$, and thus the difference decreases for these scenarios.

Consider next the two cases, in which the nominal and brainstorming group generate a new idea along two different dimensions. With probability $p_1 * p_2$, the brainstorming group builds on dimension one and the nominal group on dimension two, and with the same probability the brainstorming group builds on dimension two and the nominal group on dimension one. The sum of the performance differences of these two cases is given by:

$\Delta_2 = [\int_0^\infty (F_1(W_1)F_1^{I'_1}(W_1) - F_1^{I_1}(W_1))dW_1 + \int_0^\infty (F_2^{I'_2}(W_2) - F_2(W_2)F_2^{I_2}(W_2))dW_2]$
$+ [\int_0^\infty (F_1^{I'_1}(W_1) - F_1(W_1)F_1^{I_1}(W_1))dW_1 + \int_0^\infty (F_2(W_2)F_2^{I'_2}(W_2) - F_2^{I_2}(W_2))dW_2]$
$= \Delta_0 + [\int_0^\infty F_1(W_1)(F_1^{I'_1}(W_1) - F_1^{I_1}(W_1))dW_1 + \int_0^\infty F_2(W_2)(F_2^{I'_2}(W_2) - F_2^{I_2}(W_2))dW_2]$. Since $F_i \leq 1$ the average performance difference of these two cases $\frac{\Delta_2}{2}$ is again less or equal to Δ_0, and hence

in expectation the difference between brainstorming and nominal groups decreases, if the diversity affects nominal and brainstorming groups equally.

Proof of Corollary 1 : Follows directly from Propositions 1 and 2.

Proof of Corollary 2 : $V^{BG} = \int_0^\infty (1 - G^i(V))dV + \int_0^\infty (1 - G^{I'-i}(V))dV$.
$\frac{\partial V^{BG}}{\partial i} = ln(G) \int_0^\infty (G^{I'-i}(V) - G^i(V))dV$ and $\frac{\partial^2 V^{BG}}{\partial i^2} = -[ln(G)]^2 \int_0^\infty (G^{I'-i}(V) + G^i(V))dV < 0$. Thus, V^{BG} is concave and maximized at $i^* = \frac{I'}{2}$. Since V^{NG} is not influenced by production blocking, the difference increases if the effect of production blocking becomes more heterogeneous.