

Improving Online Idea Generation Platforms and Customizing the Task Structure on the Basis of Consumers' Domain-Specific Knowledge

The authors explore how firms can enhance consumer performance in online idea generation platforms. Most, if not all, online idea generation platforms offer all consumers identical tasks in which (1) participants are granted access to ideas from other participants and (2) ideas are classified into categories, but consumers can navigate freely across idea categories. The former is linked to stimulus ideas, and the latter may be viewed as a first step toward problem decomposition. The authors propose that the effects of both stimulus ideas and problem decomposition are moderated by consumers' domain-specific knowledge. In particular, concrete cues such as stimulus ideas are more beneficial to low-knowledge consumers, and high-knowledge consumers are better served with abstract cues such as the ones offered by problem decomposition. The authors' hypotheses are supported by an extensive empirical investigation involving more than 6,000 participants. The findings suggest that online idea generation platforms should use problem decomposition more explicitly and that firms should not immediately show other participants' ideas to high-knowledge consumers when they access the platform. In other words, online idea generation platforms should customize the task structure on the basis of each participant's domain-specific knowledge.

Keywords: idea generation, online idea generation platforms, consumer knowledge, stimulus ideas, problem decomposition

Online Supplement: <http://dx.doi.org/10.1509/jm.13.0212>

Idea generation is used routinely and extensively across a wide range of organizations and industries. With the emerging trend of new product development outsourcing in recent years (e.g., Carson 2007; Raassens, Wuyts, and Geyskens 2012), idea generation has also taken a new direction with an increased level of consumer involvement in innovation projects (e.g., Bayus 2012; Fuchs and

Schreier 2011; Hoyer et al. 2010; Magnusson 2009; Slot, Srinivasan, and Wuyts 2013; Soukhoroukova, Spann, and Skiera 2012; Terwiesch and Ulrich 2009; Toubia 2006; Toubia and Florès 2007). For example, companies such as Dell and Starbucks have launched online open innovation platforms that allow consumers to propose ideas to the company (see, e.g., www.ideastorm.com, mystarbucksidea.force.com). Other platforms, such as InnoCentive (www.innocentive.com), give companies access to online communities of “problem solvers.” With the availability of user-friendly, free hosting services such as IdeaScale (www.ideascale.com), open idea generation platforms are now accessible to any organization or company.

Most, if not all, online idea generation platforms offer identical tasks to all consumers. In particular, they (1) give participants access to other participants' ideas and (2) classify ideas into categories (e.g., “Products,” “Experience,” “Involvement”) but let consumers navigate freely across categories. The former is linked to the literature on stimulus ideas, and the latter may be viewed as a first step toward a manipulation called “problem decomposition,” in which

Lan Luo is Associate Professor of Marketing, Marshall School of Business, University of Southern California (e-mail: lluo@marshall.usc.edu). Olivier Toubia is Glaubinger Professor of Business, Columbia Business School, Columbia University (e-mail: ot2107@columbia.edu). Address correspondence to Olivier Toubia. The authors thank Dan Bartels; Ellie Kyung; Dilip Soman; Gülden Ülkümen; and seminar participants at the INFORMS Marketing Science Conference, the Product Development and Management Association Annual Research Forum, and the Columbia Marketing Brown Bag Series for their insightful comments on previous drafts of this article. The authors also thank the USC Marshall Center for Global Innovation for the financial support it has provided for this research. Wayne Hoyer served as guest editor and Sanjay Sood served as area editor for this article.

participants are instructed to work on subcategories of the problem separately.

Our research is among the first to explore how online idea generation platforms can enhance consumer performance by deviating from this standard structure (i.e., by not showing stimulus ideas and/or by explicitly decomposing the problem) and customizing their structure on the basis of each consumer's characteristics. As Hoyer et al. (2010) discuss, although consumers nowadays are willing and able to share their ideas to firms benevolently, research on how firms may effectively enhance consumer performance at idea generation has been scarce.

Theoretically, our research extends the idea generation literature by showing that domain knowledge moderates the impacts of stimulus ideas and problem decomposition. Extensive research has examined the impact of stimulus ideas and the conditions under which they are beneficial versus detrimental. However, very little is known about whether stimulus ideas have a different impact on different types of idea generation participants. Similarly, although research has documented a positive main effect of problem decomposition, little is known about the moderators of this effect.

To derive our hypotheses, we adopt a cognitive view of idea generation, according to which both stimulus ideas and problem decomposition provide search cues in the memory retrieval step of the idea generation process. Drawing on the psychology literature, we propose that low-level, concrete cues such as stimulus ideas are more beneficial to low-knowledge consumers. In contrast, consumers with more abundant domain knowledge are more prone to the negative cognitive fixation induced by such ideas. We further suggest that problem decomposition provides high-level, abstract cues that are more beneficial to high-knowledge consumers compared with low-knowledge consumers. Our hypotheses are supported by an extensive empirical investigation involving more than 6,000 participants.

From a managerial perspective, our results have clear implications related to the design of online idea generation platforms. By specifically comparing a typical online idea generation platform with some variants in which stimulus ideas are not presented and/or the task is explicitly decomposed, we propose the following changes to improve extant online generation platforms. First, we show that online idea generation platforms should use problem decomposition more explicitly (rather than merely classifying ideas). Second, they should either not show stimulus ideas to high-knowledge consumers, or at least not show stimulus ideas to these consumers until they have generated a few ideas on their own. In other words, online idea generation platforms should customize the task structure based on each participant's domain-specific knowledge. Indeed, it is now easy to collect preliminary data on each consumer and customize the idea generation task on the fly based on the consumer's domain-specific knowledge. In a subsequent section, we illustrate that such a customized system has the potential to considerably enhance the value of online idea generation platforms.

Hypothesis Development

Cognitive View of Idea Generation

Our hypotheses are based on a cognitive view of idea generation, which has become increasingly popular in the literature. Cognitive research in idea generation has been concerned primarily with modeling and understanding the cognitive processes involved in idea generation and investigating ways in which the format and the structure of the task influence the output (Nagasundaram and Dennis 1993). For example, the Search for Ideas in Associate Memory model (Nijstad, Diehl, and Stroebe 2003; Nijstad and Stroebe 2006; Nijstad, Stroebe, and Lodewijkx 2002, 2003, 2006), based on Raaijmakers and Shiffrin's (1981) Search of Associate Memory model, views idea generation as a two-stage process. In the first stage, the knowledge-activation stage, a search cue is assembled in short-term memory and used to probe long-term memory, resulting in the activation of an "image" in long-term memory.¹ Search cues typically consist of examples of previously generated ideas and/or elements of the problem definition (Nijstad and Stroebe 2006; Nijstad, Stroebe, and Lodewijkx 2002). In the second stage, the idea production stage, the features of the image activated in the first stage are used to generate new ideas by combining knowledge, forming new associations, or applying knowledge to a new domain. Indeed, these operations are widely believed to be critical to the creative process (Goldenberg and Mazursky 2002; Mednick 1962; Simonton 2003). We note that the Search for Ideas in Associate Memory model is closely related to the Geneplore model proposed by Finke, Ward, and Smith (1992; for applications of the Geneplore model in marketing, see, e.g., Moreau and Dahl 2005; Sellier and Dahl 2011).

This cognitive view of idea generation provides a useful framework for studying interventions aimed at improving consumer performance in idea generation. The interventions we consider herein influence the knowledge-activation stage of the process. More specifically, they provide cues to consumers that influence the knowledge retrieved from long-term memory during their search for ideas. As mentioned previously, the two most common types of search cues are stimulus ideas and elements of the problem definition (Nijstad and Stroebe 2006; Nijstad, Stroebe, and Lodewijkx 2002). These two types of cues may be mapped onto the two manipulations that we consider in this article, stimulus ideas and problem decomposition. A key difference between these two types of cues is that whereas stimulus ideas provide low-level, concrete search cues in the form of specific solutions to the problem, problem decomposition provides high-level, abstract cues that are elements of the problem definition. In this article, we propose that the

¹The term "image" does not imply a visual or spatial representation of information. Images are knowledge structures that consist of a central concept and several features and associations related to that concept (e.g., features of the image "hotel" include "has rooms" and "has lobby," and that image may be associated with the image "restaurant").

effects of both manipulations are moderated by domain-specific knowledge, but in opposite directions.

Stimulus Ideas

A sizable literature stream has explored the conditions under which exposure to stimulus ideas may be beneficial versus detrimental. Research has shown that stimulus ideas benefit participants by providing a set of cues that activate knowledge that otherwise might not be accessed (Brown et al. 1998). Nijstad and Stroebe (2006) suggest that stimulus ideas can lead to productivity gains in idea generation because they make related knowledge more accessible. Other studies have also shown that exposure to others' ideas leads to cognitive stimulation, which in turn may lead to additional ideas as long as the inefficiencies inherent to face-to-face groups are reduced and participants generate ideas individually while being exposed to stimulus ideas (Dugosh et al. 2000; Nijstad, Stroebe, and Lodewijkx 2002; Paulus and Yang 2000). This is achieved in practice by using electronic idea generation mechanisms (Gallupe et al. 1992; Nunamaker, Applegate, and Konsynski 1987).

Meanwhile, research has shown one of the main drawbacks of exposure to stimulus ideas to be cognitive fixation: participants may fixate on the knowledge captured by stimulus ideas, limiting their ability to identify original ideas (e.g., Bayus 2012; Cardoso and Badke-Schaub 2011; Jansson and Smith 1991; Marsh, Landau, and Hicks 1996; Smith 2003; Smith, Ward, and Schumacher 1993). Smith and Blankenship (1991) demonstrate that the distracting stimuli work as memory retrieval blocks. Indeed, the retroactive inhibition literature has shown that when retrieving items from memory, the activation of cues related to a subset of items inhibits retrieval of the remaining items (e.g., Nickerson 1984; Watkins and Allender 1987).

In summary, previous research has shown both positive and negative effects of stimulus ideas. On the positive side, stimulus ideas provide cues activating knowledge that may not be readily activated otherwise. On the negative side, stimulus ideas may induce cognitive fixation. Notably, the literature has focused primarily on how stimulus ideas are presented to participants but has largely ignored the issue of to whom the ideas are presented. In particular, very little is known about the type of participants who benefit more from the positive effects of stimulus ideas and/or suffer less from the ideas' negative effects.

We propose that because stimulus ideas provide concrete examples of solutions to the problem, their positive effects should be more pronounced for low-knowledge consumers than for high-knowledge consumers. Following Ratchford (2001), we conceptualize domain knowledge as a person's embodied skill or expertise that is acquired through past experiences, formal or informal training, or education. Note that by "low-knowledge consumers," we refer to participants who have relatively limited knowledge about a specific domain but still have enough basic knowledge to be able to generate ideas on the topic. Indeed, it is well accepted that ideas cannot be generated without reference to prior knowledge (e.g., Goldenberg and Mazursky 2002; Mednick 1962; Simonton 2003).

Research has shown that novices (resp., experts) have a more concrete (resp., abstract) way of reasoning and tend to approach problems in a bottom-up (resp., top-down) manner (Chase and Simon 1973; Chi, Feltovich, and Glaser 1981; Schmidt and Boshuizen 1993; Wiley 1998). Indeed, experts classify knowledge using a more abstract, principle-based organization (Adelson 1981, 1984; Chase and Simon 1973; Chi, Glaser, and Rees 1982; Schmidt and Boshuizen 1993). Moreover, it has been shown that abstract representations become more accessible in people's memory as they accumulate knowledge (Ericsson and Kintsch 1995; Kyung, Menon, and Trope 2013; Schmidt and Boshuizen 1993). In contrast, domain knowledge of less knowledgeable people tends to be less structured (Chi, Glaser, and Rees 1982) and more episodic (Mitchell and Dacin 1996). These studies suggest that concrete (resp., abstract) cues such as stimulus ideas fit better with the cognitive structures that characterize low-knowledge (resp., high-knowledge) consumers. Therefore, we expect that low-knowledge people should benefit more from the concrete examples provided by stimulus ideas.

We have argued thus far that the positive effects of stimulus ideas should be more pronounced for low- versus high-knowledge consumers. Next, we propose that high-knowledge consumers, compared with low-knowledge consumers, should also be more prone to cognitive fixation induced by stimulus ideas (i.e., the negative effects of stimulus ideas are more pronounced for high- vs. low-knowledge consumers).

Indeed, Wiley (1998) demonstrates that domain knowledge acts as a mental set that promotes fixation in problem solving (i.e., consumers who have greater knowledge in the relevant domain are actually more prone to fixation). This finding is consistent with another stream of literature in industrial design, which has shown that fixation is more likely to happen when the provided stimuli or examples are more familiar to the participants. Perttula and Sipilä (2007) and Purcell and Gero (1992) find that fixation is more severe when participants are presented with more common examples of solutions. Purcell and Gero (1996) report that mechanical engineers become fixated when shown examples of solutions that are characteristic of the knowledge base of their discipline.

It is important to note that research has shown fixation to be mostly unconscious and unintentional. Smith, Ward, and Schumacher (1993) find that fixation remains even when participants are explicitly instructed to diverge from the examples, and Marsh, Ward, and Landau (1999) report a similar result. Wiley (1998) finds that warning participants of fixation is not enough to remove it. As Smith (2003, pp. 27–28) notes, "the implicitly retrieved examples cannot be voluntarily rejected to make way for more appropriate responses." Therefore, it seems unlikely that consumers—and, in particular, high-knowledge consumers, who are more likely to be prone to fixation—would be able to voluntarily ignore stimulus ideas to avoid the negative effects of fixation.

This discussion suggests that the positive effects of stimulus ideas should be more pronounced for low- versus

high-knowledge consumers, and their negative effects should be more pronounced for high- versus low-knowledge consumers. Therefore, the net effect of stimulus ideas should be more favorable to low- versus high-knowledge consumers, as our first hypothesis states:

H₁: There is a negative interaction between consumer knowledge and stimulus ideas. Namely, the presence of stimulus ideas reduces the performance of high-knowledge consumers relative to low-knowledge consumers.

Problem Decomposition

Problem decomposition consists of simply decomposing the problem into subproblems and instructing participants to consider each subproblem separately. The creator of the brainstorming technique himself, Alex F. Osborn, argues that breaking down (i.e., decomposing) a problem into subcategories and instructing participants to work on the components separately improves the output of the session (Osborn 1963, p. 174). This prediction has been confirmed by Dennis et al. (1996), among others, in the context of electronic brainstorming. In one of their experiments, for example, they either asked participants, “What can elected officials, business leaders, and the general public do to encourage a higher level of leadership in the community?” (intact problem) or asked them first to focus on “What can elected officials do to encourage...?” followed by “What can business leaders do...?” and “What can the general public do...?” (decomposed problem). The authors show that participants who worked on the decomposed problem produced more ideas and more good ideas compared with participants who worked on the intact problem. Coskun et al. (2000) and Pitz, Sachs, and Heerboth (1980), among others, replicated this finding.

As we have indicated, most online idea generation platforms classify ideas into categories, but they do not explicitly instruct participants to consider these categories separately. Given the finding that problem decomposition tends to be beneficial, it seems relevant to consider whether online idea generation platforms should use this manipulation more explicitly and whether it will affect different types of consumers differently.

Unlike stimulus ideas, for which both positive and negative effects have been documented in the literature, the literature on problem decomposition has not identified any particular negative effect of this manipulation. Therefore, we focus on the positive effects, which we argue should be greater for high- versus low-knowledge consumers. From a cognitive perspective, problem decomposition provides another set of search cues to participants. Instead of being concrete examples of ideas, the cues are now elements of the problem definition (Nijstad and Stroebe 2006; Nijstad, Stroebe and Lodewijkx 2002). One key difference between these two types of cues is that whereas stimulus ideas provide low-level, concrete cues, problem decomposition provides high-level, abstract cues that do not reflect any low-level, specific solution to the problem.

As we have discussed, high-level, abstract cues fit better with the cognitive structures and problem-solving strate-

gies that characterize high-knowledge consumers. Moreover, previous research has suggested that abstract cues should help high-knowledge consumers activate a broader set of images in their search for new ideas. Indeed, according to theories of spreading activation, higher-level concepts tend to activate lower-level ones (Collins and Loftus 1975). Therefore, higher-level cues are more beneficial (i.e., lead to the activation of a broader set of relevant lower-level concepts) for high-knowledge consumers, whose knowledge is more hierarchical and structured (Chase and Ericsson 1981; Cowley and Mitchell 2003; Kyung, Menon, and Trope 2013). This leads to our second hypothesis:

H₂: There is a positive interaction between consumer knowledge and problem decomposition. Namely, problem decomposition increases the performance of high-knowledge consumers relative to low-knowledge consumers.

Experimental Setup

Testing our hypotheses requires the following elements: (1) soliciting ideas from consumers, (2) varying the task structure by providing stimulus ideas and/or decomposing the problem into subproblems, (3) obtaining an indicator of the quality of each generated idea, (4) measuring the performance of each consumer, and (5) assessing the knowledge level of each consumer in our idea generation task. We discuss each of these elements in the following subsections.

Idea Generation Task

Study 1 is related to idea generation for possible applications of “EasyCode,” a technology that allows camera cell phones to scan two-dimensional bar codes. In particular, the participants were told, “We are interested in new applications for the EasyCode technology. Please enter any idea you may think of.” This technology is well known today as “QR codes,” but at the time of our studies (Fall 2009 to Fall 2010) it was still emerging in the United States (some have argued that the tipping point for this technology happened in the fourth quarter of 2010; see, e.g., Cohen 2011; Mobio 2011). A detailed description of the technology presented to the respondents appears in Web Appendix A1.

In Study 2 (conducted in Fall 2014), we recruited respondents to submit ideas related to enhancing consumer experience in each of the following four domains: fast-food restaurants, personal banking, movie theaters, and social media platforms. In this study, each respondent was prompted to complete four idea generation tasks (one per domain, presented in random order).

Varying Task Structure

We examined the following four conditions of idea generation task structure: stimulus ideas, not decomposed (Condition 1); no stimulus ideas, not decomposed (Condition 2); stimulus ideas, decomposed (Condition 3); and no stimulus ideas, decomposed (Condition 4). Given that Condition 1 is the most similar to the task structure employed by most extant online idea generation platforms (which give participants access to others’ ideas and classify ideas into categories

but do not instruct consumers to work on each category separately), we deem the first condition as the control and the other three as treatment conditions.

Using Study 1 as an example, in the stimulus ideas conditions, participants were exposed to idea examples during idea generation (e.g., “a user could scan an EasyCode in a magazine and see a movie’s trailer & show times,” “a user could scan an EasyCode on a TV screen and download and play a game from a TV show”). When the task was not decomposed, we instructed participants to generate ideas for the EasyCode technology that could be “printed on any type of paper, carton, or electronic screens.” In the decomposed conditions, respondents were prompted to generate ideas for the three different types of support (“paper,” “carton,” and “electronic screens”) in sequence.

All studies were conducted on the web, and respondents completed the task independently from one another. The web interface was such that respondents could enter as many or as few ideas as desired, and no time limit was imposed on the task. This setup was in line with the norm in existing online idea generation platforms.

In Study 1, we employed a between-subjects design by randomly assigning respondents to one of the idea generation conditions. In Study 2, we adopted a mixed design: the four domains were randomly assigned to the four conditions at the respondent level (between subjects), and each respondent completed the four conditions in random order (within subject). Table 1 provides an overview of the tasks used across studies, how each task was decomposed, and the source of the stimulus ideas. We provide more details on our research design when we discuss these studies in subse-

quent sections. Web Appendices A1–A2 provide screenshots of the web interfaces. Tables 2–3 show summary statistics from our two studies.

Idea Evaluation

Because an accurate assessment of idea quality is essential to our empirical investigation, we adopted several measures from the literature (Girotra, Terwiesch, and Ulrich 2010; Kornish and Ulrich 2011; Toubia and Flores 2007) to evaluate idea quality in our two studies. In Study 1, we evaluated consumers’ adoption intent associated with each idea using respondents recruited from Amazon.com’s Mechanical Turk (MTurk) panel. After being introduced to the technology, each respondent was asked to rate 20 ideas in terms of how likely they “would be to use it if it were available” on a ten-point scale adopted from Morrison (1979). The 20 ideas presented to each consumer were selected randomly from the set of ideas generated in that study, among the ideas that had received the fewest evaluations up to that point. This selection mechanism ensured that, by the end of the study, all ideas had received approximately the same number of evaluations.

In Study 2, we collected both consumer and business value evaluations of all ideas generated. Specifically, we obtained consumer evaluations from respondents recruited from MTurk. The online paradigm was similar to the one used to collect adoption intent ratings. After a brief introduction of the task context, the respondents were asked to rate 20 ideas on five-point scales based on whether the idea was novel, insightful, valuable for consumers, and well articulated. We collected consumer evaluation rather than

TABLE 1
Overview of Studies

| Study | Topic | Decomposition | Stimulus Ideas |
|------------------------|---|---|--|
| Study 1 | New applications for the Easy-Code technology | Code printed on <ul style="list-style-type: none"> •Paper •Carton •Electronic screens | Six ideas (two for each support) |
| Study 2 | | | |
| Fast-food restaurants | Enhance user experience at fast-food restaurants | Ideas related to <ul style="list-style-type: none"> •Service •Products •Dine-in experience | Other participants’ ideas from pretest |
| Personal banking | Enhance user experience with personal banking | Ideas related to <ul style="list-style-type: none"> •Service •Products •Education | Other participants’ ideas from pretest |
| Movie theaters | Enhance user experience in movie theaters | Ideas related to <ul style="list-style-type: none"> •Seats, screens, and sound •Experience management •Food and snacks •Ticket processing •Special offers and programs | Other participants’ ideas from pretest |
| Social media platforms | Enhance user experience with social media platforms | Ideas related to <ul style="list-style-type: none"> •Features •Management •Integration •New platforms or technologies | Other participants’ ideas from pretest |

adoption intent scores because the latter was not directly applicable to the idea generation domains considered in this study. Business value evaluations were assessed using a panel of senior business major undergraduate students who participated in the idea evaluation task as part of a class assignment. As senior business majors, these students had received formal training in evaluating the business value of new product ideas through a series of business classes. We further refreshed their memory with a lecture on this particular topic shortly before distributing this assignment. Following Girotra, Terwiesch, and Ulrich (2010), we instructed the students that an idea's technical feasibility, novelty, specificity, and potential market demand should be accounted for when the idea is being evaluated for its business value.

TABLE 2
Study 1 Summary Statistics

| | M | SD |
|--------------------------------------|--------|-------|
| Idea Generation | | |
| Participant type | MTurk | |
| Number of participants | 407 | |
| Consumer knowledge score (1–5) | 2.812 | .781 |
| Number of ideas | 1.799 | 1.990 |
| Idea quality score metric (1–10) | 6.635 | 1.766 |
| Consumer performance metric | 11.933 | 8.697 |
| Idea Evaluation | | |
| Idea quality metric: adoption intent | MTurk | |
| Participant type | MTurk | |
| Number of participants | 1,216 | |
| Number of evaluations per idea | 32.425 | 1.423 |

TABLE 3
Study 2 Summary Statistics

| | Fast-Food Restaurants | | Personal Banking | | Movie Theaters | | Social Media Platforms | |
|-----------------------------------|------------------------|-------|------------------------|-------|------------------------|--------|------------------------|-------|
| | M | SD | M | SD | M | SD | M | SD |
| Idea Generation | | | | | | | | |
| Participant type | MTurk | | MTurk | | MTurk | | MTurk | |
| Number of participants | 301 | | 301 | | 301 | | 301 | |
| Consumer knowledge score (1–5) | 3.560 | .834 | 3.351 | .837 | 3.567 | .843 | 3.621 | .874 |
| Number of ideas | 2.903 | 2.378 | 2.568 | 1.961 | 3.950 | 3.176 | 2.425 | 1.909 |
| Idea quality score metric 1 (1–5) | 3.474 | .398 | 3.615 | .417 | 3.561 | .408 | 3.540 | .423 |
| Consumer performance metric 1 | 10.085 | 8.217 | 9.284 | 7.084 | 14.068 | 11.273 | 8.584 | 6.894 |
| Idea quality score metric 2 (1–5) | 3.808 | .579 | 3.653 | .449 | 3.402 | .624 | 3.379 | .582 |
| Consumer performance metric 2 | 11.054 | 8.945 | 9.380 | 7.011 | 13.438 | 10.681 | 8.194 | 6.578 |
| Idea Evaluation | | | | | | | | |
| Idea Quality Metric 1: | | | | | | | | |
| Consumer Evaluation | | | | | | | | |
| Participant type | MTurk | | MTurk | | MTurk | | MTurk | |
| Number of participants | 666 | | 622 | | 918 | | 632 | |
| Number of evaluations per idea | 14.667 | .926 | 15.369 | .488 | 14.773 | .468 | 16.552 | .602 |
| Idea Quality Metric 2: | | | | | | | | |
| Business Value Evaluation | | | | | | | | |
| Participant type | Senior business majors | | Senior business majors | | Senior business majors | | Senior business majors | |
| Number of participants | 91 | | 81 | | 115 | | 71 | |
| Number of evaluations per idea | 15.038 | 1.047 | 12.569 | 1.844 | 14.698 | 1.709 | 16.530 | 1.460 |

Measuring Consumer Performance

Our empirical investigation requires examining how the interplay of consumer knowledge and task structure affects consumer performance in idea generation. Following the extant idea generation literature (e.g., Dennis and Valacich 1994; Dennis et al. 1996; Diehl and Stroebe 1987; Gallupe et al. 1992; Lamm and Trommsdorff 1973; Valacich, Dennis, and Nunamaker 1992), we define a participant's performance as the sum of the average quality ratings of his or her ideas. Specifically, let s_i^{mj} be the average quality score on metric m (e.g., adoption intent) received by idea i submitted by respondent j , and let N_j be the number of ideas submitted by respondent j . The performance of consumer j on metric m is measured as $Y_j^m = \sum_{i=1}^{N_j} s_i^{mj}$. Given that we collected more than one idea quality metric in Study 2, we tested our hypotheses in this study using standardized consumer performance scores calculated from (1) consumer evaluations, (2) business value evaluations, and (3) the average of these two, respectively.

Measuring Consumer Knowledge

We adapted the measurement scales developed by Mitchell and Dacin (1996) to gauge domain-specific consumer knowledge (for more details on the measurement scales used in our studies, see the Appendix). In Study 1, we based our scale items on knowledge about technological products and used the average of each respondent's replies to these items to compute a knowledge score. We decided to measure knowledge about technological products because many ideas related to the EasyCode technology also relate to technological products in general. We acknowledge that whether this is the best way to measure consumer knowl-

edge in this study is still an open question. Nevertheless, our findings from Study 1 provide a lower bound on the magnitude of the effects predicted by our hypotheses and on the benefits from customizing the task based on the degree of consumer knowledge.

In Study 2, we addressed this issue by measuring domain-specific knowledge scale items related to the specific domains of fast-food restaurants, personal banking, movie theaters, and social media platforms. Namely, we computed four domain-specific knowledge scores for each respondent, corresponding to each of the four domains.

In our empirical investigation, we further control for differences in the raw knowledge scores across domains by mean-centering the knowledge scores using their corresponding domain-specific means. This is particularly necessary for Study 2 because the average knowledge scores may exhibit noticeable differences across domains (e.g., as Table 3 shows, the average knowledge scores are much higher for social media platforms than for personal banking).²

To empirically verify that our results were related to consumer knowledge, as suggested by our theoretical arguments, and not to alternative consumer characteristics, we also included the domain-specific lead user scale (adapted from Hoffman, Kopalle, and Novak 2010), the emergent consumer scale (adapted from Hoffman, Kopalle, and Novak 2010), and the domain-specific consumer innovativeness scale (adapted from Goldsmith and Hofacker 1991) in Study 1 (for more details about the scale items, reliability, and discriminant validity assessments, see the Appendix). We found that our hypotheses did not hold when consumers were characterized on the basis of any of these alternative characteristics, suggesting that consumer knowledge is indeed the relevant construct for our analysis.

Study 1

Method

We recruited 407 participants from MTurk to complete the online idea generation task. The flow of the study was as follows. Each respondent first answered the knowledge measurement scale items and the open-ended questions discussed previously (see the “Measuring Consumer Knowledge” subsection and the Appendix). Next, a description of the technology was given. Finally, the respondent was randomly assigned to one of the four experimental conditions, with each asking him or her to generate some ideas for possible applications of the EasyCode technology.

In the control condition (stimulus ideas, not decomposed), we provided all respondents with an identical set of six stimulus ideas and instructed them to generate some

²To further verify the reliability of the respondents’ self-reported responses to the scale items, we also included some open-ended questions in Study 1 to assess the respondents’ general knowledge and usage of technology products. We found that respondents with higher knowledge scores described significantly more situations. Due to the length of Study 2 and the results from Study 1, we did not include such open-ended questions in the second study.

ideas for possible applications of the EasyCode technology that could be “printed on any type of paper, carton, or electronic screens.” To show the respondents possible applications related to the technology, stimulus ideas associated with all three types of support were illustrated (two ideas for paper, two for carton, and two for electronic screens).

In the first treatment condition (no stimulus ideas, not decomposed), we removed stimulus ideas from the survey interface, while keeping other aspects constant. In the second treatment condition (stimulus ideas, decomposed), we prompted the respondents to submit ideas for the three types of support (i.e., paper, carton, and electronic screens) in sequence, with the two stimulus ideas associated with each support type being the same as the ones used in the control condition. In the last treatment condition (no stimulus ideas, decomposed), we removed stimulus ideas and asked respondents to submit ideas for each type of support sequentially.

A total of 750 ideas were generated across the four experimental conditions. We obtained quality score ratings of these ideas by collecting adoption intent ratings from a different set of respondents from MTurk (N = 1,216), as described in the “Idea Evaluation” subsection. We pre-screened these respondents to ensure that none of them was involved in the idea generation task. The quality ratings were used to compute a consumer performance metric score for each respondent in our idea generation task (as described in the “Measuring Consumer Performance” subsection). In summary, 1,623 respondents participated in either the idea generation or idea evaluation task in this study.

Results

We tested our hypotheses by estimating the following regression:

$$(1) \quad Y_j^a = \beta_0 + \beta_K \times K_j + \beta_{ST} \times ST_j + \beta_{DE} \times DE_j + \beta_{K \times ST} \times K_j \times ST_j + \beta_{K \times DE} \times K_j \times DE_j + \epsilon_j,$$

with Y_j^a being respondent j ’s performance score on adoption intent, K_j being the respondent’s domain-specific knowledge score (as described in the “Measuring Consumer Knowledge” subsection), ST_j being an indicator variable denoting whether respondent j was exposed to stimulus ideas, and DE_j being an indicator variable denoting whether this respondent completed a decomposed task.³ Table 4 presents the results.

In line with H_1 , the regression analysis showed a significant negative interaction between consumer knowledge and stimulus ideas. In addition, we found a significant positive interaction between consumer knowledge and problem decomposition, confirming H_2 . Our analysis also revealed a positive main effect of problem decomposition, which is consistent with findings from prior research (e.g., Coskun et al. 2000; Dennis et al. 1996).

³We did not include the interaction term between ST_j and DE_j in Equations 1 or 2 because we did not hypothesize an interaction between these two manipulations. Nevertheless, all conclusions held if we included this additional term in our analyses.

TABLE 4
Study 1 Estimation Results

| Parameter | Consumer Performance Metric Based on Adoption Intent | |
|----------------------------|--|------|
| | Coefficient | SE |
| Intercept | | |
| Intercept | -.344* | .073 |
| Main Effects | | |
| Knowledge | .141 | .089 |
| Stimulus | -.019 | .089 |
| Decomposition | .797* | .089 |
| Interaction Effects | | |
| Knowledge × Stimulus | -.339* | .115 |
| Knowledge × Decomposition | .294* | .116 |

*Significant at .05.

To explore the interaction effects further, we examined the effects of stimulus ideas and problem decomposition at one standard deviation below and above the mean knowledge score using a spotlight analysis (Fitzsimons 2008; Irwin and McClelland 2003). Specifically, we centered knowledge score at one standard deviation below (resp., above) the mean and ran a similar regression as in Equation 1 to examine the effects of stimulus ideas and problem decomposition on low-knowledge (resp., high-knowledge) consumers.

As we expected, low-knowledge consumers performed significantly better when stimulus ideas were present than when they were absent ($\beta_{ST}^{LowKnow} = .245, t = 1.96, p = .05$). We also find that high-knowledge consumers performed significantly worse when they were exposed to stimulus ideas ($\beta_{ST}^{HighKnow} = -.284, t = 2.27, p < .05$). Our spotlight analyses reveal that problem decomposition enhanced the performance of both high- and low-knowledge consumers, but significantly more so for high-knowledge consumers ($\beta_{DE}^{HighKnow} = 1.026, t = 8.13, p < .01$; $\beta_{DE}^{LowKnow} = .568, t = 4.45, p < .01$).⁴ In summary, our spotlight analyses confirm that (1) stimulus ideas enhance the performance of low-knowledge consumers and are detrimental to the performance of more knowledgeable consumers and (2) problem decomposition enhances the performance of both high- and low-knowledge consumers but is more beneficial for high-knowledge consumers.

Robustness Checks

Given that our idea generation task did not impose any constraint on the time spent on the task or the number of ideas submitted (similar to most existing idea generation platforms), we conducted the following two robustness checks. In the first robustness check, we omitted respondents with task completion time less than one standard deviation below the average completion time. The rationale is that if a

⁴We tested whether the two beta regression coefficients across high- and low-knowledge consumers were statistically different in all spotlight analyses across studies. All reported differences were statistically significant at $p < .05$.

respondent completes the task unusually quickly, (s)he may not have paid adequate attention in the study. For similar reasons, in the second robustness check, only respondents who submitted at least one idea were retained in the analysis. Web Appendix A3 presents the results of both robustness checks. Our conclusions hold under both robustness checks.

In addition, we carried out a third robustness check in which we replicated our tests of H_1 and H_2 with a different sample of respondents and using consumer performance metrics based on alternative idea quality ratings (adoption intent, overall attractiveness, and business value) from three panels of evaluators (MTurk consumers, freshman and sophomore undergraduate students, and senior business majors). We also accounted for the existence of identical or nearly identical ideas in this robustness check to control for potential noise in our performance measures. A total of 1,651 respondents participated in this robustness check. In the interest of space, we provide the detailed description of this effort in Web Appendix A4.

Discussion

Compared with the extant literature, this study provides the first empirical evidence that stimulus ideas and problem decomposition have differential impacts on low- and high-knowledge consumers. Consistent with our hypothesis, consumers who are less knowledgeable about the focal problem benefit more from the low-level, concrete search cues provided by stimulus ideas. Because this type of search cue is at odds with the knowledge structures of high-knowledge consumers, and because these consumers are more prone to cognitive fixation, exposure to stimulus ideas had an overall detrimental effect on high-knowledge consumers. In contrast, because problem decomposition provides high-level, abstract search cues that fit better with the cognitive structures that characterize high-knowledge consumers, decomposing the idea generation problem into subproblems is considerably more beneficial for consumers with greater domain-specific knowledge. An important managerial implication from this study is that firms can greatly foster consumer performance in idea generation by customizing the task for low- and high-knowledge consumers. We explore this insight further in Study 2 and the “Customizing Idea Generation Tasks: Empirical Assessment” section.

Study 2

Method

In Study 2, we specifically examine the impact of varying the task structure from a typical online idea generation platform. Diverging from Study 1, we employed a mixed design in which each respondent was invited to submit ideas in four different idea generation domains in four distinct experimental conditions, with the domains randomly matched to experimental conditions. Consequently, each respondent completed four idea generation tasks, with each domain and each condition shown exactly once in a random order.

In particular, each respondent was invited to submit ideas on how to enhance consumer experience in the fol-

lowing domains: fast-food restaurants, personal banking, movie theaters, and social media platforms. We chose these four domains because (1) experiences in these domains are relevant to most consumers and (2) domain-specific knowledge varies considerably across them (both within and across individuals), thus providing a suitable setting to examine whether findings from our Study 1 are generalizable in highly distinctive domains.

To test how online idea generation platforms may be enhanced on the basis of our findings, we developed our own idea generation platforms using PHP (a programming language mainly used for web development) and MySQL databases (an open-source database used in web applications). We mimicked the structure and the look and feel of commercial idea generation platforms while allowing for specific manipulations to be tested in a controlled environment.

Commercial idea generation platforms use ideas submitted by previous participants as stimulus ideas. However, this creates potential confounds because participants are not all exposed to the same set of stimulus ideas. To ensure that all participants were exposed to the same set of stimulus ideas while staying as close as possible to the user interface of a commercial online idea generation community, we first ran a pretest in which we collected an initial set of ideas and their popularity scores. In particular, we collected ideas in the four domains from 401 consumers recruited from MTurk, receiving 181, 200, 213, and 210 ideas for fast-food restaurants, personal banking, movie theaters, and social media platforms, respectively.

After removing ideas that were clearly off topic, we classified stimulus ideas into subcategories as in most idea generation communities. For example, in the idea generation task related to fast-food restaurants, stimulus ideas were classified into subcategories of “service,” “products,” and “dine-in experience.” These three subcategories emerged as a natural way to classify the ideas generated in the pretest, and we obtained them by manually going through all the ideas. We used a similar approach to determine subcategories in all four domains.

We also followed common practice in commercial online idea generation platforms by assigning a popularity score to each stimulus idea. We obtained the popularity scores of the stimulus ideas on the basis of binary votes (“thumbs up/thumbs down”) on pretest ideas from 907 consumers from MTurk who each evaluated 20 ideas. The ideas shown to each participant were selected such that the distribution of the number of votes per idea would resemble the heavily skewed distributions found in commercial online idea generation platforms. In particular, ideas with a high score (as defined by [number of positive votes – number of negative votes] × 10, as in many commercial platforms) were more likely to be shown to new participants. At the conclusion of this pretest, we had a set of stimulus ideas classified into subcategories, with a popularity score for each idea.

Three hundred one consumers recruited from MTurk participated in the main study. Each respondent was presented with four idea generation tasks in a random order, and each task randomly matched a domain and a condition.

Within each task, respondents first answered the domain-specific consumer knowledge scale items (see the “Measuring Consumer Knowledge” subsection and the Appendix), after which they were asked to submit ideas related to enhancing consumer experience in the corresponding domain, based on the corresponding experimental condition.

We designed the control condition (stimulus ideas, not decomposed) in this study to resemble a typical online generation platform. In particular, respondents were told that they had entered an idea generation community where they could submit their own ideas and browse ideas from other participants. These stimulus ideas were grouped according to the subcategories identified in the pretest. Consistent with the setup in extant idea generation platforms, the subcategory tabs were ordered by the number of stimulus ideas in the subcategory, and ideas in each tab were sorted in decreasing order of their popularity scores obtained in the pretest. Five ideas were shown per page, with the option of seeing additional pages of stimulus ideas in each tab. The popularity score of each idea was displayed next to it, along with “thumbs up” and “thumbs down” icons that allowed participants to vote on the idea. To make the voting feature as realistic as possible, the popularity score shown to the respondent increased or decreased by ten points when (s)he clicked on the thumbs-up or thumbs-down button. However, to ensure that such votes would not alter the popularity scores seen by other respondents, the score changed only for that participant; future participants would see the initial score based on the pretest. In other words, the survey interface was independent across respondents.

During the idea generation task, respondents were free to submit their own ideas and browse/vote on others’ ideas at any time. When clicking on the “Enter Idea Here” button, respondents were taken to a pop-up window where they were prompted to submit one idea at a time. Respondents were also required to categorize their idea into a subcategory when submitting it. We employed such a classification requirement to resemble the common practice in most online idea generation platforms. Consistent with the norm in existing platforms, respondents could submit as many or as few ideas as they desired, and there was no time limit on the task. For screenshots from this study, see Web Appendix A2.

We included three treatment conditions to examine the impact of varying the task structure from a typical online idea generation platform. The first treatment condition (no stimulus ideas, not decomposed) differed from the control condition only in that respondents were not shown any stimulus idea (all other aspects were identical to the control condition). The second treatment condition (stimulus ideas, decomposed) varied from the control condition by explicitly instructing respondents to submit ideas related to each subcategory in sequence (e.g., “How could we improve service at fast-food restaurants?” followed by “How could we improve products offered at fast-food restaurants?” and “How could we improve dine-in experience at fast-food restaurants?”). Within each decomposed task, respondents were exposed to stimulus ideas from the corresponding subcategory (i.e., ideas related to “service,” “products,” or “dine-in experience,” as in the previous example), using the

same set of ideas as in the control condition. All other aspects were identical to the control condition. The third treatment condition (no stimulus ideas, decomposed) differed from the control in that no stimulus ideas were presented and the problem was decomposed. Again, all other aspects were similar to the control condition.

A total of 3,566 ideas were generated in this study (fast-food restaurants: 874; personal banking: 773; movie theaters: 1,189; and social media platforms: 730 ideas). Given that the results of our hypothesis testing in Study 1 did not change after controlling for identical or nearly identical ideas, for simplicity we did not identify such ideas in Study 2. Nevertheless, ideas that were apparently unrelated to the topic were removed manually.

We collected two sets of idea quality scores (see the “Idea Evaluation” subsection). We obtained the consumer evaluation scores from a different panel of 2,838 consumers recruited from MTurk. In addition, 358 senior business major students were recruited to evaluate these ideas’ business value. Altogether, 4,805 respondents participated in the pretest, the main study, or the idea evaluation task in this study.

Results

We employed a mixed-effects model to test our two hypotheses while controlling for heterogeneity in both respondents and domains. In particular, we used a random-effect coefficient to capture heterogeneity across respondents. We employed fixed effects to capture heterogeneity across domains. We estimate the following model in this study:

$$(2) Y_{jk}^m = \beta_{0j} + \beta_{\text{bank}} \times \text{Bank}_{jk} + \beta_{\text{movie}} \times \text{Movie}_{jk} + \beta_{\text{social}} \times \text{Social}_{jk} + \beta_K \times K_{jk} + \beta_{ST} \times ST_{jk} + \beta_{DE} \times DE_{jk} + \beta_{K \times ST} \times K_{jk} \times ST_{jk} + \beta_{K \times DE} \times K_{jk} \times DE_{jk} + \epsilon_{jk},$$

where Y_{jk}^m represents respondent j ’s performance score in task k on metric m ; Bank_{jk} , Movie_{jk} , and Social_{jk} are indicator variables denoting whether the k th idea generation task completed by respondent j was related to personal banking, movie theaters, or social media platforms, respectively; K_{jk} is respondent j ’s domain-specific knowledge score in task k ; ST_{jk} is an indicator variable denoting whether respondent j was exposed to stimulus ideas in task k ; and DE_{jk} indicates whether problem decomposition was used for task k of respondent j .

As we discuss in the “Measuring Consumer Knowledge” subsection, because the knowledge score measures are domain specific, we mean-centered the K_{jk} measure in Equation 2 for each specific domain separately. We used standardized consumer performance scores in our analysis as before. Given that Y_{jk}^m was measured on the basis of both consumer evaluations and business value evaluations of ideas generated, we ran three separate mixed-effects models with the dependent variable calculated from (1) consumer evaluation scores, (2) business value evaluation scores, and (3) the average of these two.

Table 5 presents the results. Overall, the three consumer performance metrics gave rise to identical conclusions. The random-effects coefficient reveals that there was considerable heterogeneity among respondents in their performance at the idea generation tasks. The fixed-effects coefficients indicate that, in general, respondents performed the best in generating ideas related to movie theaters. In contrast, enhancing consumer experience in personal banking seemed to be the most challenging idea generation task.

Consistent with Study 1, we found a significant positive main effect of explicitly decomposing the idea generation task into subtasks. In addition, in line with our previous study, we found no main effect of stimulus ideas. Notably, whereas Study 1 did not reveal a significant main effect of domain-specific consumer knowledge, it played a signifi-

TABLE 5
Study 2 Estimation Results

| | Consumer Performance Metric Based on: | | | | | |
|-----------------------------|---------------------------------------|------|---------------------------|------|------------------------|------|
| | Consumer Evaluation | | Business Value Evaluation | | Average of Two Metrics | |
| | Coefficient | SE | Coefficient | SE | Coefficient | SE |
| Intercept | | | | | | |
| Intercept | -.575* | .053 | -.462* | .054 | -.519* | .053 |
| SD of intercept | .490* | .029 | .497* | .029 | .493* | .029 |
| Idea Generation Task | | | | | | |
| Personal banking | -.164* | .052 | -.263* | .052 | -.213* | .052 |
| Movie theaters | .490* | .052 | .311* | .052 | .400* | .052 |
| Social media platforms | -.112* | .052 | -.261* | .052 | -.187* | .052 |
| Main Effects | | | | | | |
| Knowledge | .189* | .042 | .186* | .042 | .182* | .042 |
| Stimulus | -.009 | .037 | -.017 | .037 | -.013 | .037 |
| Decomposition | 1.056* | .037 | 1.052* | .037 | 1.054* | .037 |
| Interaction Effects | | | | | | |
| Knowledge × Stimulus | -.193* | .046 | -.202* | .047 | -.197* | .046 |
| Knowledge × Decomposition | .122* | .053 | .107* | .046 | .114* | .046 |

*Significant at .05.

cant positive role in consumer performance in Study 2. This might be because, compared with our previous study, which measures consumer knowledge about technological products in general (rather than about the specific technology of EasyCode), Study 2 measures domain knowledge more specifically.

The bottom panel of Table 5 illustrates results related to our hypothesis testing. In line with findings from our previous study, we found a significant negative interaction between domain-specific consumer knowledge and exposure to stimulus ideas. Furthermore, as predicted in H₂, we found a significant positive interaction between domain knowledge and problem decomposition. We further carried out spotlight analysis to explore these interaction effects. Because we obtained identical conclusions when the three consumer performance metrics were used in the mixed-effects model, we employed spotlight analysis using the average of the two standardized consumer performance metrics. Specifically, we ran two additional mixed-effects models with knowledge scores centered at one standard deviation below and above the domain-specific means. Our analysis revealed that stimulus ideas were highly beneficial for low-knowledge consumers ($\beta_{ST}^{LowKnow} = .154, t = 2.87, p < .01$). In contrast, exposure to such ideas impeded the performance of high-knowledge consumers ($\beta_{ST}^{HighKnow} = -.180, t = 3.35, p < .01$). In addition, explicitly decomposing the idea generation task into subtasks had a significantly stronger positive impact on high-knowledge consumers than on their low-knowledge counterparts ($\beta_{DE}^{HighKnow} = 1.151, t = 21.41, p < .01$; $\beta_{DE}^{LowKnow} = .957, t = 17.87, p < .01$).

Robustness Checks

In line with Study 1, we further conducted robustness checks in which we omitted respondents with task completion time less than one standard deviation below the mean or respondents who submitted zero ideas in an idea generation task. All our previously discussed conclusions remain intact (Web Appendix A3).

Discussion

This study provides the first empirical investigation of the impact of varying the task structure from a typical online idea generation platform. We find that consumer performance in idea generation can be enhanced considerably by modifying the task structure from the extant setup used in most idea generation platforms. In particular, we find that although most online idea generation platforms classify ideas into different subcategories, explicit instructions to guide respondents to work on each subcategory separately lead to substantial gains in consumer performance. Furthermore, because high-knowledge consumers are better served with high-level, abstract search cues offered by problem decomposition, we find that such an improvement in consumer performance is even more pronounced with high-knowledge consumers.

Furthermore, while most extant idea generation platforms offer identical task structures to all consumers, our study suggests that customizing the task structure on the basis of the participant's domain-specific consumer knowl-

edge can be highly beneficial. Indeed, exposure to idea examples from other participants in an idea generation platform can significantly decrease the performance of participants with more abundant domain-specific knowledge. Therefore, our study suggests that online idea generation platforms should be customized by domain knowledge such that others' ideas are not shown to high-knowledge consumers immediately after they access the platform. Although our experiment compared the two extremes of showing versus not showing stimulus ideas, firms may experiment with showing stimulus ideas to high-knowledge consumers only after they have had an opportunity to submit a set of initial ideas. It is also important to note that to the extent that firms value votes submitted by consumers, high-knowledge consumers may still be given the opportunity to provide feedback on ideas submitted by other participants, but preferably only after having submitted a few ideas of their own.

Customizing Idea Generation Tasks: Empirical Assessment

After obtaining empirical support for our hypotheses, we further examined the benefits of customizing the idea generation task conditional on each consumer's domain knowledge. As per our empirical findings, we consider a customized idea generation system that would (1) assess the participant's knowledge level of the focal problem, (2) categorize low- versus high-knowledge consumers on the fly, and (3) customize the task so that (a) low-knowledge consumers are presented with a decomposed idea generation problem in which stimulus ideas are offered for each subproblem and (b) high-knowledge consumers are assigned to a decomposed idea generation task without stimulus ideas.

Although our hypothesis testing employed a continuous consumer knowledge measure, task customization inevitably requires classifying consumers into low- and high-knowledge types. To mimic a situation in which each consumer is classified on the fly, we used the median split of respondents' domain-specific knowledge scores to classify respondents as high- versus low-knowledge consumers in the corresponding domain. We used the median split approach because this method has been widely used to classify consumers into high- and low-knowledge types (e.g., Bettman and Susan 1987; Dahl and Moreau 2007; Mandel and Johnson 2002).

Both Studies 1 and 2 enable us to assess the value of such a customized idea generation system *ex post*. In particular, in each study we compare the performance of low-knowledge consumers in the stimulus ideas, decomposed condition (Condition 3) combined with high-knowledge consumers in the no stimulus ideas, decomposed condition (Condition 4) with the performance of consumers in any of the four conditions (i.e., all consumers assigned to the same task structure).

Because Study 2 includes idea generation tasks in four distinct domains, we used the domain-specific median knowledge scores to categorize each respondent at the domain level. For example, the same respondent may be classified as a high-knowledge consumer in fast-food restaurants but a low-knowledge consumer in personal bank-

ing. In addition, in our comparisons presented subsequently, consumer performance scores in Study 2 were based on the average of consumer and business value evaluations.

Table 6 provides the results of such comparisons.⁵ In both studies, we find that the average performance of low-knowledge consumers in Condition 3 combined with high-knowledge consumers in Condition 4 is significantly superior to that of consumers (both types combined) in all four conditions. Notably, Table 6 also reveals that the commonly adopted task structure in most extant online generation platforms (Condition 1: stimulus ideas, not decomposed, assigned to everyone) yielded among the worst consumer performance across the four conditions, indicating that the current idea generation task structure used by many firms or agencies is indeed quite suboptimal. Overall, our findings suggest that a customized system has the potential to considerably enhance the value firms derive from involving consumers in idea generation.

Conclusions

Despite the increasing popularity of consumer involvement in idea generation, research on how firms may foster consumer performance in such tasks has been scarce. To our knowledge, our research is among the first attempts to investigate this underresearched yet important topic. Our research uncovers clear and simple modifications that would enhance online idea generation platforms and shows that customized idea generation systems can be highly beneficial managerially.

⁵Note that consumer performance scores from Studies 1 and 2 are not directly comparable. While, in general, respondents in Study 2 submitted more ideas in each domain than those in Study 1, the idea quality scores in Study 2 were based on five-point scales rather than on ten-point scales as in Study 1. As robustness checks, we also carried out similar comparisons after omitting respondents with task completion time less than one standard deviation below the mean and after excluding respondents with zero idea submissions in an idea generation task (Web Appendix A3).

Specifically, we demonstrate that low-level, concrete search cues such as stimulus ideas are considerably more beneficial to low-knowledge consumers than to high-knowledge consumers. In addition, although explicitly decomposing the idea generation task into subtasks leads to substantial improvement in the performance of both low- and high-knowledge consumers, problem decomposition is significantly more beneficial for high-knowledge consumers. Our hypotheses were supported by two studies involving more than 6,000 respondents in total.

We further outline a modified idea generation platform in which problem decomposition is used more explicitly and other participants' ideas are not shown to high-knowledge consumers. Our empirical results suggest that firms can significantly improve consumer performance in idea generation by employing such a system. With the increased availability of web-based online idea generation platforms and the readily available web technology to classify consumers on the fly on the basis of their domain knowledge, this process may be readily implemented in a timely and cost-effective manner.

We conclude by highlighting some avenues for further research other than those already mentioned. First, future studies may identify optimal ways of classifying ideas into subcategories and decomposing the task accordingly. Although our hypotheses are robust to the range of decompositions that we tested in our studies, future work could further explore how many subcategories should be used and what principles should be used to determine these categories. Second, although we adopted well-established approaches to measuring consumer performance in idea generation, we acknowledge that the ultimate performance metric should be the market performance of the new products that are developed on the basis of these ideas. However, collecting this metric requires developing and launching such products, which was not feasible in our context. Further research may explore how to incorporate such metrics into performance evaluations.

TABLE 6
Benefits from Customizing Idea Generation Tasks

| Consumer Performance | Low Knowledge | | High Knowledge | | All Consumers | |
|---|---------------|-------|----------------|--------|---------------|--------|
| | M | SD | M | SD | M | SD |
| Study 1 | | | | | | |
| Condition 1: Stimulus ideas, not decomposed | 9.033 | 6.034 | 7.953 | 5.857 | 8.357 | 5.929 |
| Condition 2: No stimulus ideas, not decomposed | 6.308 | 4.678 | 9.289 | 6.543 | 7.811 | 5.829 |
| Condition 3: Stimulus ideas, decomposed | 15.095* | 6.929 | 16.201 | 8.805 | 15.699 | 7.929 |
| Condition 4: No stimulus ideas, decomposed | 10.565 | 5.612 | 20.701* | 10.806 | 16.073 | 8.029 |
| Customized: Low knowledge in Condition 3 + High knowledge in Condition 4 | — | — | — | — | 18.078* | 8.106 |
| Study 2 | | | | | | |
| Condition 1: Stimulus ideas, not decomposed | 6.493 | 4.832 | 5.491 | 4.015 | 5.947 | 4.678 |
| Condition 2: No stimulus ideas, not decomposed | 5.051 | 4.191 | 7.393 | 5.102 | 6.280 | 4.370 |
| Condition 3: Stimulus ideas, decomposed | 15.043* | 9.146 | 15.742 | 10.811 | 15.433 | 10.106 |
| Condition 4: No stimulus ideas, decomposed | 11.933 | 8.235 | 19.065* | 13.580 | 15.404 | 10.134 |
| Customized: Low knowledge in Condition 3 + High knowledge in Condition 4 | — | — | — | — | 17.148* | 11.129 |

*Best in column at .05.

Appendix: Measurement Scales Used to Categorize Consumer Type

Domain-Specific Consumer Knowledge Scale Used in Study 1

(Adapted from Mitchell and Dacin 1996; Cronbach's alpha = .835)

- Compared to the average person, I do not know much about technology products. (reverse-coded)
- I am very familiar with technology products.
- I am not skilled at utilizing technology products. (reverse-coded)
- I am very interested in technology products.
- I own a lot of technology products.
- My friends own a lot of technology products.
- I read articles related to technology products all the time.

Domain-Specific Consumer Knowledge Scales Used in Study 2

(Adapted from Mitchell and Dacin 1996)

Fast-Food Restaurants (Cronbach's alpha = .834)

- Compared to the average person, I do not know much about fast-food restaurants. (reverse-coded)
- I am very familiar with fast-food restaurants.
- I am not knowledgeable about fast-food restaurants. (reverse-coded)
- I am very interested in fast-food restaurants.
- I go to fast-food restaurants a lot.
- My friends go to fast-food restaurants a lot.
- I read about fast-food restaurants (e.g., reviews, blogs, inserts, ads, flyers) all the time.

Personal Banking (Cronbach's alpha = .862)

- Compared to the average person, I do not know much about personal banking. (reverse-coded)
- I am very familiar with personal banking.
- I am not skilled at utilizing various personal banking products and services. (reverse-coded)
- I am very interested in personal banking.
- I use a lot of personal banking products and services.
- My friends use various personal banking products and services.
- I read about different personal banking products and services all the time.

Movie Theaters (Cronbach's alpha = .845)

- Compared to the average person, I do not know much about movie theaters. (reverse-coded)
- I am very familiar with movie theaters.
- I am not knowledgeable about movie theaters. (reverse-coded)
- I am very interested in going to the movies.
- I go to movie theaters a lot.
- My friends go to movie theaters a lot.
- I read about movie theaters (e.g., reviews, blogs, inserts, flyers) all the time.

Social Media Platforms (Cronbach's alpha = .870)

- Compared to the average person, I do not know much about social media platforms. (reverse-coded)

- I am very familiar with various social media websites.
- I am not skilled at the various functions/features offered by different social media platforms. (reverse-coded)
- I am very interested in social media websites.
- I participate in a lot of social media websites.
- My friends use social media websites a lot.
- I read about social media websites (e.g., reviews, blogs) all the time.

Alternative Consumer Characteristic Scales Collected in Study 1

Domain-Specific Lead Users Scale (Adapted from Hoffman, Kopalle, and Novak 2010; Cronbach's alpha = .658)

- Other people consider me as "leading edge" with respect to technology products.
- I tend not to look for new and different usages of technology products. (reverse-coded)
- I have suggested to my friends and family members some new and different ways of utilizing new technology.
- I normally do not participate in store offers/promotions to try out new technology products. (reverse-coded)
- I have come up with some new and different solutions to satisfy my unmet needs by using technology products.

Emergent Consumers Scale (Adapted from Hoffman, Kopalle, and Novak 2010; Cronbach's alpha = .817)

- When I hear about a new technology product, it is easy for me to come up with ideas of how to apply this technology.
- If I don't see an immediate use for a new technology, I normally do not think about how I might use it in the future. (reverse-coded)
- When I see a new technology product, it is easy for me to visualize how it might fit into the life of an average person in the future.
- If someone gave me a new technology product with no clear application, I could "fill in the blanks" so someone else would know what to do with it.
- Even if I don't see an immediate use for a new technology, I like to imagine how people might use it in the future.
- I try to avoid experimenting with new ways of using new technology products. (reverse-coded)
- I do not like to find patterns in complexity. (reverse-coded)
- I can picture how new technology products could be applied to improve an average person's life.

Domain-Specific Consumer Innovativeness Scale (Adapted from Goldsmith and Hofacker 1991; Cronbach's alpha = .790)

- In general, I am among the last in my circle of friends to adopt new technology products. (reverse-coded)
- If I heard that a new technology product was available, I would be interested to try it out.
- Compared to my friends, I do not adopt a lot of technology products. (reverse-coded)
- In general, I am the first in my circle of friends to know about a new technology product.
- I will adopt a new technology product, even if I haven't heard of it before.
- I get to know many technology products before other people do.

Discriminant Validity Checks

In Study 1, we also conducted discriminant validity checks of the previously listed consumer characteristic scales in a pretest with 123 respondents following the guidelines proposed by Churchill (1979). For the seven-item consumer knowledge scale, we submitted the corresponding items to an exploratory factor analysis with oblique rotation. The analysis suggested one dominant factor with loadings above .570 for all the items. We conducted similar analysis for the five-item domain-specific lead user scale. The exploratory factor analysis suggested one dominant factor with loadings above .5 for all the items. With regard to the eight-item emergent consumer scale, the exploratory factor analysis revealed a single factor with factor loading above .62 for all items. Finally, for the six-item domain-specific consumer innovativeness

scale, the exploratory factor analysis suggested a single factor with factor loading above .563 for all items. A correlation analysis revealed moderate correlations among the four consumer type categorizations (ranging from .382 between lead users and emergent consumers to .429 between emergent consumers and high-knowledge consumers). We further used confirmatory factor analysis to formally test the discriminant validity of these four constructs. We found that a four-factor structural model fits the data significantly better (goodness-of-fit index = .917, root mean square error of approximation = .053) than a single-factor model (goodness-of-fit index = .6749, root mean square error of approximation = .095), suggesting that domain-specific knowledge, domain-specific leader user, emergent nature, and domain-specific innovativeness are four distinct constructs.

REFERENCES

- Adelson, Beth (1981), "Problem Solving and the Development of Abstract Categories in Programming Languages," *Memory & Cognition*, 9 (4), 422–33.
- (1984), "When Novices Surpass Experts: The Difficulty of a Task May Increase with Expertise," *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 10 (3), 483–95.
- Bayus, Berry (2012), "Crowdsourcing New Product Ideas over Time: An Analysis of the Dell IdeaStorm Community," *Management Science*, 59 (1), 226–44.
- Bettman, James and Mita Sujun (1987), "Effects of Framing on Evaluation of Comparable and Noncomparable Alternatives by Expert and Novice Consumers," *Journal of Consumer Research*, 14 (2), 141–54.
- Brown, Vincent, Michael Tumeo, Timothy S. Larey, and Paul B. Paulus (1998), "Modeling Cognitive Interactions During Group Brainstorming," *Small Group Research*, 29 (4), 495–526.
- Cardoso, Carlos and Petra Badke-Schaub (2011), "Fixation or Inspiration: Creative Problem Solving in Design," *Journal of Creative Behavior*, 45 (2), 77–82.
- Carson, Stephen J. (2007), "When to Give Up Control of Outsourced New Product Development," *Journal of Marketing*, 71 (January), 49–66.
- Chase, William G. and K. Anders Ericsson (1981), "Skilled Memory," in *Cognitive Skills and Their Acquisition*, John R. Anderson, ed. Hillsdale, NJ: Lawrence Erlbaum Associates.
- and Herbert A. Simon (1973), "Perception in Chess," *Cognitive Psychology*, 4 (1), 55–81.
- Chi, Michelene T.H., Paul J. Feltovich, and Robert Glaser (1981), "Categorization and Representation of Physics Problems by Experts and Novices," *Cognitive Science*, 5 (2), 121–52.
- , Robert Glaser, and Ernest Rees (1982), "Expertise in Problem Solving," in *Advances in the Psychology of Human Intelligence*, Vol. 1, R.J. Sternberg, ed. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Churchill, Gilbert A., Jr. (1979), "A Paradigm for Developing Better Measures of Marketing Constructs," *Journal of Marketing Research*, 16 (February), 64–73.
- Cohen, Heidi (2011), "QR Codes: 26 Must-Have Facts," HeidiCohen.com, (April 4), (accessed May 27, 2015), [available at <http://heidicohen.com/qr-code-data>].
- Collins, Allan M. and Elizabeth F. Loftus (1975), "A Spreading-Activation Theory of Semantic Processing," *Psychological Science*, 82 (6), 407–428.
- Coskun, Hamit, Paul B. Paulus, Vincent Brown, and Jeffrey J. Sherwood (2000), "Cognitive Stimulation and Problem Presentation in Idea-Generating Groups," *Group Dynamics: Theory, Research, and Practice*, 4 (4), 307–329.
- Cowley, Elizabeth and Andrew A. Mitchell (2003), "The Moderating Effect of Product Knowledge on the Learning and Organization of Product Information," *Journal of Consumer Research*, 30 (3), 443–54.
- Dahl, Darren W. and C. Page Moreau (2007), "Thinking Inside the Box: Why Consumers Enjoy Constrained Creative Experiences," *Journal of Marketing Research*, 44 (August), 347–69.
- Dennis, Alan R. and Joseph S. Valacich (1994), "Group, Sub-Group, and Nominal Group Idea Generation: New Rules for a New Media?" *Journal of Management*, 20 (4), 723–36.
- , ———, Terry Connolly, and Bayard E. Wynne (1996), "Process Structuring in Electronic Brainstorming," *Information Systems Research*, 7 (2), 268–77.
- Diehl, Michael and Wolfgang Stroebe (1987), "Productivity Loss in Brainstorming Groups: Toward the Solution of a Riddle," *Journal of Personality and Social Psychology*, 53 (3), 497–509.
- Dugosh, Karen Leggett, Paul B. Paulus, Evelyn J. Roland, and Huei-Chuan Yang (2000), "Cognitive Stimulation in Brainstorming," *Journal of Personality and Social Psychology*, 79 (5), 722–35.
- Ericsson, K. Anders and Walter Kintsch (1995), "Long-Term Working Memory," *Psychological Review*, 102 (2), 211–45.
- Finke, Ronald A., Tomas B. Ward, and Steven M. Smith (1992), *Creative Cognition*. Cambridge, MA: MIT Press.
- Fitzsimons, Gavan J. (2008), "Death to Dichotomizing," *Journal of Consumer Research*, 35 (1), 5–8.
- Fuchs, Christoph and Martin Schreier (2011), "Customer Empowerment in New Product Development," *Journal of Product Innovation Management*, 28 (1), 17–32.
- Gallupe, R. Brent, Alan R. Dennis, William H. Cooper, Joseph S. Valacich, Lana M. Bastianutti, and Jay F. Nunamaker (1992), "Electronic Brainstorming and Group Size," *Academy of Management Journal*, 35 (2), 350–69.
- Girotra, Karan, Christian Terwiesch, and Karl T. Ulrich (2010), "Idea Generation and the Quality of the Best Idea," *Management Science*, 56 (4), 591–605.
- Goldenberg, Jacob and David Mazursky (2002), *Creativity in Product Innovation*. Cambridge, UK: Cambridge University Press.
- Goldsmith, Ronald E. and Charles F. Hofacker (1991), "Measuring Consumer Innovativeness," *Journal of the Academy of Marketing Science*, 19 (3), 209–221.
- Hoffman, Donna L., Praveen K. Kopalle, and Thomas P. Novak (2010), "The 'Right' Consumers for Better Concepts: Identifying Consumers High in Emergent Nature to Develop New Product Concepts," *Journal of Marketing Research*, 47 (October), 854–65.
- Hoyer, Wayne D., Rajesh Chandy, Matilda Dorotic, Manfred Krafft, and Siddharth S. Singh (2010), "Consumer Cocreation in

- New Product Development,” *Journal of Service Research*, 13 (3), 283–96.
- Irwin, Julie R. and Gary H. McClelland (2003), “Negative Consequences of Dichotomizing Continuous Predictor Variables,” *Journal of Marketing Research*, 40 (August), 366–71.
- Jansson, David G. and Steven M. Smith (1991), “Design Fixation,” *Design Studies*, 12 (1), 3–11.
- Kornish, Laura J. and Karl T. Ulrich (2011), “Opportunity Spaces in Innovation: Empirical Analysis of Large Samples of Ideas,” *Management Science*, 57 (1), 107–128.
- Kyung, Ellie J., Geeta Menon, and Yaacov Trope (2014), “Construal Level and Temporal Judgments of the Past: The Moderating Role of Knowledge,” *Psychonomic Bulletin & Review*, 21 (3), 734–39.
- Lamm, Helmut and Gisela Trommsdorff (1973), “Group Versus Individual Performance on Tasks Requiring Ideational Proficiency (Brainstorming): A Review,” *European Journal of Social Psychology*, 3 (4), 361–88.
- Magnusson, Peter R. (2009), “Exploring the Contributions of Involving Ordinary Users in Ideation of Technology-Based Services,” *Journal of Product Innovation Management*, 26 (5), 578–93.
- Mandel, Naomi and Eric J. Johnson (2002), “When Web Pages Influence Choices: Effects of Visual Primes on Experts and Novices,” *Journal of Consumer Research*, 29 (2), 235–45.
- Marsh, Richard L., Joshua D. Landau, and Jason L. Hicks (1996), “How Examples May (and May Not) Constrain Creativity,” *Memory & Cognition*, 24 (5), 669–80.
- , Thomas B. Ward, and Joshua D. Landau (1999), “The Inadvertent Use of Prior Knowledge in a Generative Cognitive Task,” *Memory & Cognition*, 27 (1), 94–105.
- Mednick, Sarnoff A. (1962), “The Associative Basis of the Creative Process,” *Psychological Review*, 69 (3), 220–32.
- Mitchell, A. and Peter A. Dacin (1996), “The Assessment of Alternative Measures of Consumer Expertise,” *Journal of Consumer Research*, 23 (3), 219–39.
- Mobio (2011), “The Naked Facts: QR Barcode Scanning in 2H-2010,” research report, (accessed May 27, 2015), [available at <http://docslide.us/documents/the-naked-facts-2h20102.html>].
- Moreau, C. Page and Darren W. Dahl (2005), “Designing the Solution: The Impact of Constraints on Consumers’ Creativity,” *Journal of Consumer Research*, 32 (1), 13–22.
- Morrison, Donald G. (1979), “Purchase Intentions and Purchase Behavior,” *Journal of Marketing*, 43 (April), 65–74.
- Nagasundaram, Murli and Alan R. Dennis (1993), “When a Group Is Not a Group: The Cognitive Foundation of Group Idea Generation,” *Small Group Research*, 24 (4), 463–89.
- Nickerson, Raymond S. (1984), “Retrieval Inhibition from Part-Set Cuing: A Persisting Enigma in Memory Research,” *Memory & Cognition*, 12 (6), 531–52.
- Nijstad, Bernard A., Michael Diehl, and Wolfgang Stroebe (2003), “Cognitive Stimulation and Interference in Idea-Generating Groups,” in *Group Creativity: Innovation Through Collaboration*, Paul B. Paulus and Bernard A. Nijstad, eds. New York: Oxford University Press.
- and Wolfgang Stroebe (2006), “How the Group Affects the Mind: A Cognitive Model of Idea Generation in Groups,” *Personality & Social Psychology Review*, 10 (3), 186–213.
- , ———, and Hein F.M. Lodewijckx (2002), “Cognitive Stimulation and Interference in Groups: Exposure Effects in an Idea Generation Task,” *Journal of Experimental Social Psychology*, 38 (6), 535–44.
- , ———, and ——— (2003), “Production Blocking and Idea Generation: Does Blocking Interfere with Cognitive Processes?” *Journal of Experimental Social Psychology*, 39 (6), 531–48.
- , ———, and ——— (2006), “The Illusion of Group Productivity: A Reduction of Failures Explanation,” *European Journal of Social Psychology*, 36 (1), 31–48.
- Nunamaker, Jay F., Jr., Lynda M. Applegate, and Benn R. Konsynski (1987), “Facilitating Group Creativity: Experience with a Group Decision Support System,” *Journal of Management Information Systems*, 3 (4), 5–19.
- Osborn, Alex F. (1963), *Applied Imagination, Revised Edition*. New York: Scribner.
- Paulus, Paul B. and Hwei-Chuan Yang (2000), “Idea Generation in Groups: A Basis for Creativity in Organizations,” *Organizational Behavior and Human Decision Processes*, 82 (1), 76–87.
- Perttula, Mati and Pekka Sipilä (2007), “The Idea Exposure Paradigm in Design Idea Generation,” *Journal of Engineering Design*, 18 (1), 93–102.
- Pitz, Gordon F., Natalie J. Sachs, and Joel Heerboth (1980), “Procedures for Eliciting Choices in the Analysis of Individual Decisions,” *Organizational Behavior and Human Performance*, 26 (3), 396–408.
- Purcell, A. Terry and John S. Gero (1992), “Effects of Examples on the Results of a Design Activity,” *Knowledge-Based Systems*, 5 (1), 82–91.
- and ——— (1996), “Design and Other Types of Fixation,” *Design Studies*, 17 (4), 363–83.
- Raaijmakers, Jeroen G.W. and Richard M. Shiffrin (1981), “Search of Associative Memory,” *Psychological Review*, 88 (2), 93–134.
- Raassens, Néomie, Stefan Wuyts, and Inge Geyskens (2012), “The Market Valuation of Outsourcing New Product Development,” *Journal of Marketing Research*, 49 (October), 682–95.
- Ratchford, Brian T. (2001), “The Economics of Consumer Knowledge,” *Journal of Consumer Research*, 27 (4), 397–411.
- Schmidt, Henk G. and Henny P.A. Boshuizen (1993), “On the Origin of Intermediate Effects in Clinical Case Recall,” *Memory & Cognition*, 21 (3), 338–51.
- Sellier, Anne-Laure and Darren W. Dahl (2011), “Focus! Creative Success Is Enjoyed Through Restricted Choice,” *Journal of Marketing Research*, 48 (December), 996–1007.
- Simonton, Dean Keith (2003), “Scientific Creativity as Constrained Stochastic Behavior: The Integration of Product, Person, and Process Perspectives,” *Psychological Bulletin*, 129 (4), 475–94.
- Slot, Johanna H., Raji Srinivasan, and Stefan Wuyts (2013), “Managing the Crowd: Prize Structure and Creativity in Online Idea Generation Contests,” working paper, Tilburg University.
- Smith, Steven M. (2003), “The Constraining Effects of Initial Ideas,” in *Group Creativity: Innovation Through Collaboration*, Paul B. Paulus and Bernard A. Nijstad, eds. New York: Oxford University Press, 15–31.
- and Steven E. Blankenship (1991), “Incubation and the Persistence of Fixation in Problem Solving,” *American Journal of Psychology*, 104 (1), 61–87.
- , Thomas B. Ward, and Jay S. Schumacher (1993), “Constraining Effects of Examples in a Creative Generation Task,” *Memory & Cognition*, 21 (6), 837–45.
- Soukoroukova, Arina, Martin Spann, and Bernd Skiera (2012), “Sourcing, Filtering, and Evaluating New Product Ideas: An Empirical Evaluation of the Performance of Idea Markets,” *Journal of Product Innovation Management*, 29 (1), 100–112.
- Terwiesch, Christian and Karl T. Ulrich (2009), *Innovation Tournaments: Creating and Selecting Exceptional Opportunities*. Boston: Harvard Business School Press.
- Toubia, Olivier (2006), “Idea Generation, Creativity, and Incentives,” *Marketing Science*, 25 (5), 411–25.
- and Laurent Florès (2007), “Adaptive Idea Screening Using Consumers,” *Marketing Science*, 26 (3), 342–60.
- Valacich, Joseph S., Alan R. Dennis, and Jay F. Nunamaker Jr. (1992), “Group Size and Anonymity Effects on Computer-Mediated Idea Generation,” *Small Group Research*, 23 (1), 49–73.
- Watkins, Michael J. and Laurel E. Allender (1987), “Inhibiting Word Generation with Word Presentations,” *Journal of Experimental Psychology: Learning, Memory and Cognition*, 13 (4), 564–68.
- Wiley, Jennifer (1998), “Expertise as Mental Set: The Effects of Domain Knowledge in Creative Problem Solving,” *Memory & Cognition*, 26 (4), 716–30.