



Management Science

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

Crowding-Out or Crowding-In? The Effects of Idea Evaluation on Evaluators' Idea Generation

Johanna Schnier, Christina Raasch, Tim Schweisfurth

To cite this article:

Johanna Schnier, Christina Raasch, Tim Schweisfurth (2026) Crowding-Out or Crowding-In? The Effects of Idea Evaluation on Evaluators' Idea Generation. Management Science

Published online in Articles in Advance 24 Apr 2026

. <https://doi.org/10.1287/mnsc.2023.03335>

This work is licensed under a Creative Commons Attribution 4.0 International License. You are free to copy, distribute, transmit and adapt this work, but you must attribute this work as “Management Science. Copyright © 2026 The Author(s). <https://doi.org/10.1287/mnsc.2023.03335>, used under a Creative Commons Attribution License: <https://creativecommons.org/licenses/by/4.0/>.”

Copyright © 2026 The Author(s)

Please scroll down for article—it is on subsequent pages



With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes. For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

Crowding-Out or Crowding-In? The Effects of Idea Evaluation on Evaluators' Idea Generation

Johanna Schnier,^{a,b} Christina Raasch,^{c,d,*} Tim Schweisfurth^e

^aDepartment of Strategy and Innovation, Copenhagen Business School, 2000 Frederiksberg, Denmark; ^bDepartment of Management and Technology and ICRIOS, Bocconi University, 20136 Milan, Italy; ^cDepartment of Operations and Management, Kühne Logistics University, 20457 Hamburg, Germany; ^dKiel Institute for the World Economy, 24105 Kiel, Germany; ^eInstitute for Organizational Design and Collaboration Engineering, Hamburg University of Technology, 21073 Hamburg, Germany

*Corresponding author

Contact: jos.si@cbs.dk,  <https://orcid.org/0000-0003-4909-7520> (JS); christina.raasch@klu.org,  <https://orcid.org/0000-0003-0398-3973> (CR); tim.schweisfurth@tuhh.de,  <https://orcid.org/0000-0002-6774-3912> (TS)

Received: October 12, 2023

Revised: November 22, 2024; August 1, 2025

Accepted: October 4, 2025

Published Online in Articles in Advance:
April 24, 2026

<https://doi.org/10.1287/mnsc.2023.03335>

Copyright: © 2026 The Author(s)

Abstract. Organizations rely on their employees to produce many high-quality ideas for the improvement of products, processes, and strategies. Evaluators such as managers and technical experts need to not only assess these ideas but also face expectations to contribute ideas of their own. We investigate this task duality by asking how idea evaluation activities affect evaluators' ideation performance. Although much of the creativity and innovation literature points to a crowding-out effect, some theoretical arguments support idea crowding-in. Using data from a multinational firm's idea management system and a difference-in-differences framework, we find robust evidence of substantial idea crowding-in: Evaluators' likelihood of idea generation increases by 205% on the day that they evaluate ideas and remains elevated for the following two weeks. Extensive analyses support knowledge recombination as the key mechanism: By exposing evaluators to unfamiliar knowledge, idea evaluation creates new opportunities for recombination and idea generation. The ideas that evaluators generate postevaluation are 19% more valuable (but not more novel) than those they generate outside the postevaluation window. We contribute to innovation scholarship by identifying a hitherto overlooked effect of idea evaluation: triggering idea generation among evaluators.

History: Accepted by Karan Girotra, operations management.



Open Access Statement: This work is licensed under a Creative Commons Attribution 4.0 International License. You are free to copy, distribute, transmit and adapt this work, but you must attribute this work as "Management Science. Copyright © 2026 The Author(s). <https://doi.org/10.1287/mnsc.2023.03335>, used under a Creative Commons Attribution License: <https://creativecommons.org/licenses/by/4.0/>."

Funding: J. Schnier and C. Raasch acknowledge funding from the Deutsche Forschungsgemeinschaft [Grant RA 1798/4-1]. T. Schweisfurth acknowledges funding from the Tempowerk Technology Center Hamburg.

Supplemental Material: The online appendix and data files are available at <https://doi.org/10.1287/mnsc.2023.03335>.

Keywords: innovation • idea evaluation • idea generation • knowledge recombination

1. Introduction

Organizations cannot survive without new ideas on how to improve products, processes, and strategies (Amabile 1988, Girotra et al. 2010, Bayus 2013). Many organizations actively encourage employees at all levels to contribute ideas by using tools such as individual target agreements (Ahmadi et al. 2022), financial incentives (Ederer and Manso 2013), innovation programs (Deichmann and Jensen 2018, Cornelius et al. 2021), idea management systems (Fuchs et al. 2019, Schweisfurth and Greul 2024), internal innovation communities (Haas et al. 2015), and contests (Terwiesch and Xu 2008, Terwiesch and Ulrich 2009, Wooten and Ulrich 2017, Schweisfurth et al. 2023).

Every idea gathered by an organization needs evaluators to determine which ones to pursue and invest in

(Boudreau et al. 2016, Criscuolo et al. 2017, Lane et al. 2021). At the same time, anecdotal evidence suggests that idea evaluation can be a heavy burden for evaluators, who need to read, process, and compare a large number of ideas. For example, IBM once tasked 50 managers to take several weeks to evaluate 50,000 employee ideas from an Innovation Jam (Bjelland and Wood 2008). Such an evaluation burden may put conflicting demands on evaluators, who—much like academic reviewers—are typically required to not only evaluate others' ideas, but also to contribute new ideas themselves (Edmondson 2012, Rahmani et al. 2018). This task duality raises the question whether idea evaluation and idea generation complement or substitute for one another: Does evaluating others' ideas promote or hinder an evaluator's own ideation performance?

Much of the innovation and creativity literature suggests that idea evaluation may hinder or *crowd out* idea generation. Given that both tasks compete for employees' time and cognitive effort, the evaluation workload may deplete the time and mental resources necessary for idea generation (March and Simon 1958, Cyert and March 1963, Ocasio 1997, Agrawal et al. 2018). More importantly, however, these tasks also involve distinct thinking styles that are difficult for employees to switch between (Guilford 1967, Cropley 2006). While idea evaluation requires *convergent thinking* (applying rules to optimize and select the best option), idea generation requires *divergent thinking* (generating numerous possibilities and exploring unconventional combinations) (Berg 2016). Idea evaluation can make it difficult for employees to shift back to the *divergent thinking* necessary for successful ideation. Similarly, creativity research suggests that exposure to others' ideas may lead individuals to conform to the evaluated ideas, resulting in ideas that replicate rather than innovate (Janis 1982, Smith et al. 1993, Girotra et al. 2010).

Despite these well-established theoretical arguments, idea crowding-out remains contestable. Idea evaluation may provide epistemic, attentional, or social stimuli that increase evaluators' ability or desire to ideate, potentially *crowding in* ideas. For example, innovation research posits that individuals generate new ideas by combining existing knowledge in novel ways (Nelson and Winter 1982, Weitzman 1998, Fleming 2001). Thus, idea evaluation could give evaluators access to new knowledge that may, in turn, create more opportunities for recombination and ideation.

This paper examines whether idea evaluation either crowds out or crowds in idea generation. We use panel data from an idea management system (IMS) in a large manufacturing firm with more than 90,000 employees. The data contain the records of 185,681 ideas submitted by 16,606 and evaluated by 12,069 distinct employees, allowing us to trace every employee's daily idea generation and evaluation activities. The firm encourages employees from all functions and hierarchical levels to submit product- and process-related ideas, which are then assessed by evaluators—the employees responsible for the product or process targeted by the ideas.

We find strong evidence of *idea crowding-in*. Descriptive analyses reveal that evaluators generate 80% more ideas on the day of evaluation and 23%, 14%, and 3% more ideas on days 1, 2, and 3, respectively, after idea evaluation relative to the previous day. Event study estimates support this finding. Evaluators are 205% more likely to generate ideas on the same day of idea evaluation and are 39%, 30%, and 29% more likely to generate ideas one, two, and three days after idea evaluation than on the day preceding idea evaluation. This effect persists for roughly two weeks, after which the evaluators' idea generation returns to pre-evaluation levels. Our

findings continue to hold when we account for treatment effect heterogeneity, implement an instrumental variable strategy, use alternative samples, consider the number of business days instead of weekdays between idea evaluation and generation, run a placebo test, and apply coarsened exact matching (CEM). We also find that the extent of idea crowding-in depends on evaluators' workload, with high evaluation workloads reducing idea crowding-in but not inverting it to crowding-out.

Following an abductive approach (Sætre and Van de Ven 2021), we consider various mechanisms that could, individually or jointly, account for idea crowding-in. We find strong evidence that idea crowding-in is primarily driven by idea evaluation enhancing opportunities for knowledge recombination: evaluating ideas exposes evaluators to new knowledge, which they combine with their existing knowledge base to generate new ideas (Nelson and Winter 1982, Weitzman 1998, Fleming 2001). Consistent with knowledge recombination, we find the ideas that evaluators generate following evaluation to be similar to the evaluated ideas. We also find crowding-in to be especially strong when evaluators assess ideas on unfamiliar topics, supporting the notion that idea evaluation expands possibilities for recombination. This recombination is predominantly problem based, meaning evaluators propose new solutions to the problems addressed by the evaluated ideas.

We consider several alternative mechanisms. We find only weak evidence for an attention-shifting explanation, whereby evaluation redirects evaluators' focus from routine noncreative tasks toward idea generation. We rule out that evaluators generate ideas due to social pressure, plagiarize evaluated ideas, or bundle evaluation and generation tasks for efficiency.

Finally, we examine the quality of the crowded-in ideas by asking whether ideas generated after evaluation differ in quality from those generated without prior evaluation. Our findings reveal that ideas generated within one week of evaluation are 19% more valuable than those generated outside this one-week postevaluation period and significantly more likely to score among the top 1%, 5%, and 10% most valuable ideas submitted within a given month or year. Again, this effect persists for only two weeks following evaluation.

This paper makes several contributions to innovation scholarship. First, we extend the literature by identifying a novel function of idea evaluation: triggering idea generation among evaluators themselves. Previous research has been largely limited to two functions of evaluation—allocating scarce resources across innovation opportunities (Boudreau et al. 2016, Criscuolo et al. 2017) and providing feedback to ideators (Bayus 2013, Deichmann and van den Ende 2014, Wooten and Ulrich 2017, Piezunka and Dahlander 2019). Second, we advance the notion of *idea mobility*—the flow of ideas within an

organization—as a key driver of knowledge production. Although previous research has extensively explored how employee mobility, both within and across organizational boundaries, creates opportunities for knowledge recombination and idea generation (Singh and Agrawal 2011, Karim and Kaul 2015, Catalini 2018), we show that *idea mobility* can complement, and in some cases even replace, this process. As ideas journey through organizations on their path from inception to implementation, they ignite new knowledge recombination. Third, recent advances in natural language processing (NLP) enable us to empirically distinguish between problem-based and solution-based recombination (von Hippel and von Krogh 2016) and thereby to unpack microlevel knowledge recombination processes. Our results indicate that, rather than reusing others' solution information to solve their own problems, employees tend to build on others' problem information to propose their own solution.

2. Literature

2.1. Drivers of Idea Generation

Valuable ideas are the source of organizational renewal and innovation (Amabile 1988, Kornish and Ulrich 2011, Bayus 2013, Kornish and Hutchison-Krupat 2017). What ultimately matters for organizations is that they produce highly valuable “breakthrough” or “extreme-value” ideas that significantly boost their performance (Girotra et al. 2010). Generating such ideas requires both quantity and quality: Large numbers of ideas, high average idea quality, or both expand the right tail of the value distribution and increase the likelihood of extreme-value outcomes (Dahan and Mendelson 2001, Singh and Fleming 2010, Boudreau et al. 2011).

Both research and practice have long sought to understand how to support prolific and high-quality idea generation. Creativity scholars have examined the effectiveness of tools and methods such as brainwriting and brainstorming (Osborn 1953, Nijstad et al. 2002, Rietzschel et al. 2007) and, more recently, artificial intelligence (Bell et al. 2024), whereas innovation scholars have explored the roles of contests (Terwiesch and Xu 2008, Boudreau et al. 2011), communities (Jeppesen and Frederiksen 2006, Riedl et al. 2024), incentives (Ederer and Manso 2013, Forderer and Burtch 2025), and reward structures (Gallus 2017, Burtch et al. 2022) in fostering idea generation.

More pertinent to this paper, a very influential body of research has investigated how feedback on previously submitted ideas impacts the quantity and quality of future ideas (Dahlander and Piezunka 2014, Wooten and Ulrich 2017, Mihm and Schlapp 2019, Burtch et al. 2022). Idea feedback can take various forms, from the simple selection or rejection of ideas (Deichmann and van den Ende 2014) to detailed evaluations of particular aspects of such ideas (Camacho et al. 2019). The impact of feedback on the quantity and quality of future ideas

varies strongly, depending on the level of detail (Piezunka and Dahlander 2019), valence (Chan et al. 2021), and timing (Mihm and Schlapp 2019). Nonetheless, research shows that ideators who receive feedback generate more (Wooten and Ulrich 2017, Camacho et al. 2019, Piezunka and Dahlander 2019), although not necessarily better (Wooten and Ulrich 2017), ideas than those who do not.

Although this literature primarily focuses on how ideators respond to receiving feedback, there has been little research into how evaluators respond to providing feedback, particularly how the activity of evaluating the ideas of others impacts their own idea generation. Given the widespread organizational practice of having members engage in both idea generation and evaluation, this lack of research is surprising. For example, employees are expected to develop their own ideas while also evaluating those of their colleagues (Rahmani et al. 2018); scientists are tasked to write papers and to review others' work (Boudreau et al. 2016, Li 2017, Lane et al. 2021); and online community members contribute their own solutions to problems while also assessing the solutions of their peers (Riedl et al. 2024). We therefore ask the following question: How does idea evaluation affect evaluators' own ideation productivity in terms of idea quantity and quality?

2.2. Idea Crowding-Out vs. Crowding-In

Much of the literature has indicated that, because idea evaluation is time-consuming and cognitively demanding, its likely consequence is idea crowding-out. Idea evaluation involves reading, processing, and comparing numerous ideas based on diverse criteria, including potential value (Girotra et al. 2010, Fuchs et al. 2019) and novelty (Boudreau et al. 2016, Lane et al. 2024). Evaluators must often gather additional information and must discuss the strengths and weaknesses of ideas with other evaluators (Criscuolo et al. 2017, Lane et al. 2021). After these deliberations, and often constrained by budget limitations, evaluators decide whether to accept or reject ideas. This time spent on evaluation takes up valuable slack time—time free from tight schedules and pressures—which is known to be essential for successful ideation. For example, Agrawal et al. (2018) show that, during break weeks, students produce more ideas, which are also more likely to be highly valuable. Similarly, Schweisfurth and Greul (2024) find that unexpected interruptions in production provide employees with time away from core tasks, leading to the generation of more and better ideas. For these reasons, idea evaluation activities may be expected to compromise evaluators' ability to generate ideas of their own.

Even in the remaining time, evaluators may find it more difficult to develop new ideas because idea generation and idea evaluation involve distinct thinking styles that are hard to combine and switch between. Idea

generation requires *divergent thinking*, which involves making unexpected associations and combinations (Guilford 1967, Cropley 2006, Berg 2016). In contrast, idea evaluation requires *convergent thinking*, which emphasizes economic concerns, logic, accuracy, and speed (Cropley 2006). As employees increasingly engage in idea evaluation, they may become too anchored in a *convergent thinking* style to freely explore novel associations and combinations, making them less productive ideators.

Contrary to this crowding-out perspective, some innovation research could be taken to indicate that idea evaluation may in fact stimulate or crowd in idea generation. A well-established view in innovation is that individuals generate ideas by recombining diverse knowledge components in novel ways (Nelson and Winter 1982, Weitzman 1998, Fleming 2001). To create opportunities for recombination, individuals must continually expand the pool of knowledge available to them. Organizations can support this process by facilitating interactions among employees with disparate knowledge. Key strategies include rotating employees across different parts of the organization (Cornelius et al. 2021), hiring individuals with nonredundant knowledge from outside the organization (Singh and Agrawal 2011), and reconfiguring organizations to connect employees who have not previously collaborated (Karim and Kaul 2015). Idea evaluation may be yet another way to increase opportunities for recombination: As evaluators assess others' ideas, they may be exposed to unfamiliar knowledge, increasing combinatorial opportunities and idea generation.

3. Empirical Context

3.1. Firm's IMS

Our data come from the Idea Management System (IMS) of a large manufacturing firm with more than 90,000 employees operating on around 50 sites across the globe. Its IMS is an online tool designed to elicit, evaluate, and reward innovative ideas of employees. It tracks all activities relating to employee ideas, ensuring process reliability and transparency.

The firm implemented the IMS in 2004. Employees have since submitted 400,000 ideas worth more than €800 million. Roughly every second employee submits at least one idea every year. The firm's IMS resembles those of other firms in the manufacturing industry, as described for example by Cornelius et al. (2021) and Fuchs et al. (2019). Figure 1 illustrates the IMS process, which comprises idea submission, evaluation, selection, and implementation. We will now provide greater detail

on each of these stages and will explain the roles of evaluating employees therein.

3.2. Idea Submission

The firm encourages all its employees—regardless of their experience and level in the hierarchy—to submit ideas they consider valuable. These ideas mostly target existing products and processes, suggesting how to make them more efficient by using fewer steps, fewer material resources, fewer human resources, less effort, and/or different machines. Employees may also submit ideas for entirely new products and processes that replace or complement existing ones. Although nondisclosure agreements prevent us from revealing the content of the submitted ideas, we provide two fictitious yet highly representative examples to illustrate the nature of the ideas:

1. To assemble its products, the firm requires metal sheets. These need to be sliced to fit their particular use in the production process. The slicing creates waste—parts of metal sheets that can no longer be used. To reduce this waste, employees devise a smarter way of slicing metal sheets, reducing production costs by 5%.

2. The frontline workers use a wide variety of tools along the assembly line, making it hard to quickly locate smaller tools during production. An employee proposes using color-coded bins at every workstation for better organization. This increases frontline workers' productivity by 1%.


To submit ideas, employees are required to describe (1) the problem they seek to address (What is the current state of the art and why is it a problem?), (2) the proposed solution (How does the idea address the problem?), and (3) the expected benefits and savings (How does the idea create value?). Employees must also provide a short title of their ideas and (since 2016) must indicate which of seven topic categories their idea pertains to.

Idea submission is entirely voluntary and may result in a bonus payment. Employees are free to submit ideas at any time, either individually or in teams. They are not constrained to submit ideas during particular ideation contest periods. Although the firm does occasionally organize such contests, our data do not cover these.

3.3. Idea Evaluation

Once submitted, ideas are evaluated by the employees primarily responsible for the product or process targeted by the ideas (lead evaluators). Because employees

Figure 1. Idea Management Process



typically submit ideas relating to their own work area and tasks, the lead evaluator is often their direct supervisor. Lead evaluators may involve other employees (coevaluators) whose knowledge they deem relevant in the evaluation of ideas. Coevaluators are typically technical experts from within the lead evaluator's unit.

The evaluators must ensure that ideas satisfy three key criteria. The first is novelty (Boudreau et al. 2016, Criscuolo et al. 2017): The idea must be new and must not duplicate previously submitted or implemented ideas. If an idea is identical to one that has been rejected, it is re-evaluated in light of any changes in circumstances. The second is technical feasibility (Lane et al. 2022, 2024; Franzoni and Stephan 2023): The company must have the tools, staff, and financial resources necessary to implement the idea, and it must meet safety standards. The third is value (Girotra et al. 2010, Fuchs et al. 2019, Ghosh and Wu 2023): The idea must provide tangible benefits over the current approach, with its expected economic value surpassing its expected implementation costs. At this stage, the value criterion does not require a precise calculation of expected economic value, which comes only after the idea has been implemented (see below). If any of these criteria are not met, an idea is rejected. To document how well ideas meet these criteria, evaluators write reports, which are accessible to the submitting employees and all other evaluators involved in evaluating the idea through the IMS. The final decision on ideas is taken by the lead evaluator based on these evaluation reports.

For our analysis, we define an evaluator as any employee involved in the evaluation of ideas. This includes lead evaluators and coevaluators, irrespective of their organizational role or hierarchical position. We define idea evaluation as the reading, processing, and assessing of ideas that result in the production of evaluation reports. Evaluators may produce multiple reports per idea, for instance when revising an initial evaluation based on others' reports or in response to new information.

The evaluators have discretion over when to evaluate ideas. They receive e-mails whenever they are assigned a new evaluation task, and they are expected to submit their evaluation reports within three weeks. Unless they complete the evaluation within this timeframe, they receive a first reminder, followed by additional reminders every four weeks. There are no formal penalties for missing these deadlines.

3.4. Idea Implementation and Bonus Payment

Lead evaluators implement the ideas they have selected using their own budgets, which gives them a vested interest in ensuring that ideas are executed efficiently and promptly. After implementation, the value of an idea is calculated as the expected cost savings or additional revenue generated in the first year minus the

implementation costs. Standard values provided by the controlling department are used for this calculation, including hourly rates, energy costs, and material costs. Some selected ideas are assigned a value of zero, especially when their value is too small to justify the effort associated with calculation.

Idea submitters may receive bonus payments for implemented ideas, contingent on idea value and alignment with the employees' core responsibilities. Although bonus payments increase with idea value, they decrease when ideas closely align with the employees' core responsibilities, because the employees are expected to develop such ideas as part of their roles.

4. Data

4.1. Description of the Data

Our data comprise the idea generation and evaluation activity of 23,408 distinct employees who were employed in one of the firm's major plants between January 2004 and March 2018. Together, they generated and evaluated a total of 185,681 ideas. Among these employees, 5,267 (22.5%) both evaluated and generated ideas; 6,802 (29.1%) only evaluated ideas; and 11,339 (48.8%) only generated ideas (see Table A.1 in the Online Appendix).

Table A.2 in the Online Appendix compares the group of employees who evaluated ideas (evaluators) to the group who did not (nonevaluators). Evaluators tend to be further along in their careers. They are more likely to hold permanent positions, have moved between organizational units more often, and hold roles higher up in the organization than nonevaluators. Besides being more advanced in their organizational careers, evaluators exhibit a higher ideation productivity than nonevaluators. They generate more than twice as many ideas per year, and the ideas they generate are more valuable on average.¹

On average, evaluators assess 3.0 ideas and write 3.9 evaluation reports per year (Table A.3 in the Online Appendix). The evaluation workload is distributed unevenly among evaluators. For example, 50% of the evaluators assess only one idea and write only one evaluation report per year; 10% of the evaluators are responsible for more than 70% of the evaluated ideas (Figure A.1 in the Online Appendix). The lead evaluators involve coevaluators for 62.0% of the ideas (Figure A.2 in the Online Appendix). The entire evaluation process, from submission to final decision, takes on average 127.0 days. The evaluators complete their reports within five weeks on average, with more than half of these reports submitted within the three-week evaluation deadline.

4.2. Construction of the Sample

4.2.1. Sampling of Employees. A key question is which employees to include in our sample. Although one approach is to sample all the employees, a problem with

this approach is that the assignment of evaluation tasks is not random, as seen above, but depends on employees' career stage and demonstrated ideation productivity, among other factors. Extant research warns that, if the treatment is endogenously assigned, sampling both treated and never-treated units may bias estimates (Kim 2022, Miller 2023). Specifically, including both evaluators and nonevaluators in our sample would have made it hard for us to assess whether estimates are driven by a causal effect of idea evaluation or by evaluators and nonevaluators differing in the first place. Even if we matched evaluators with nonevaluators on the basis of observed characteristics (e.g., their past ideation productivity), they likely still fundamentally differ in unobserved ways. Given these concerns, we only include evaluators in our sample, omitting nonevaluators.

4.2.2. Choice-Based Sampling. We construct a panel data set at the evaluator-day level. For every evaluator, we sample the days between their first and last IMS interaction, with a one-month lead and lag. This gives us a panel data set with more than 26,425,003 evaluator-day observations. Because there are many more days on which evaluators do not generate ideas than days on which they do, any outcome variables related to evaluators' idea generation take the value zero in the vast majority of cases. Regression analysis, including logistic regression, is known to produce biased estimates if applied to such data (King and Zeng 2001a, b). Following past innovation research (Sorenson and Stuart 2001, Singh 2005, Singh and Marx 2013), we apply choice-based sampling: We sample all ones (i.e., days on which evaluators generate at least one idea) and a subset of zeroes (i.e., days on which evaluators do not generate any ideas). Specifically, for every day on which an evaluator generates at least one idea, we randomly sample two other days on which the same evaluator does not generate any ideas. This gives us a choice-based sample consisting of twice as many days on which evaluators do not generate ideas than days on which they do. Summary statistics and pairwise correlations of this choice-based sample are provided in Tables A.4 and A.5 in the Online Appendix. When running regressions using the choice-based sample, we account for the oversampling of ones by including evaluator-specific weights. For more details, see Section A.1 in the Online Appendix.

5. Main Analysis

In this section, we analyze how idea evaluation affects the number of ideas that evaluators generate themselves. We outline our empirical strategy in Section 5.1 and present evidence of idea crowding-in in Section 5.2. Section 5.3 explores the boundary condition of crowding-in.

5.1. Empirical Strategy

Ideally, to identify the effect of idea evaluation on evaluators' own idea generation, we would randomly assign evaluation tasks to some employees (*treated*) and not others (*control*) and compare the changes in idea generation by treated employees before and after idea evaluation with concurrent changes in idea generation by control employees. However, as seen above, the assignment of evaluation tasks to employees is endogenous.

We address endogeneity using a two-way fixed effects (TWFE) event study framework that exploits variation in the timing of evaluation activities within and across the evaluators. That is, we compare the idea generation by the evaluators who evaluate ideas in t with that of the evaluators who do not evaluate ideas in t while also accounting for different baseline idea generation levels across the evaluators and idea generation trends common to all evaluators (Sandler and Sandler 2019, Schmidheiny and Sieglöck 2023, Borusyak et al. 2024). Thus, the same evaluator switches from being a treated unit in periods when they evaluate ideas to a control unit in periods when they do not evaluate ideas.

Our main specification reads as follows:

$$y_{it} = f \left(\sum_{\tau=-7, \tau \neq -1}^{\tau=7} \alpha^{\tau} E_{it}^{\tau} + \beta X_{it} + \gamma_i + \delta_t + \epsilon_{it} \right), \quad (1)$$

where i denotes the evaluator and t the day; y_{it} represents evaluator i 's idea generation on day t . The lead and lag variables $E_{it}^{-7}, \dots, E_{it}^7$ are our key independent variables, and E_{it}^{τ} takes the value of one if evaluator i evaluated an idea τ weekdays ago and zero otherwise. We omit E_{it}^{-1} as a reference period such that α^{τ} indicates changes in evaluator i 's idea generation relative to the day before idea evaluation.

To account for the effect of multiple evaluation activities on evaluators' idea generation, Specification (1) allows several leads and lags to be "turned on" simultaneously. We choose an event window that runs from six days prior to an evaluation event to six days thereafter. As is common practice, we bin the event dummies at the endpoints of the event window in E_{it}^{-7} and E_{it}^7 (Schmidheiny and Sieglöck 2023). To ensure that the estimates for α^{-7} and α^7 are not driven by the different number of leads and lags observed for every evaluator, we trim events that were more than 20 days away from the focal day. Thus, E_{it}^{-7} sums up evaluation events that are more than 6 but fewer than 20 days into the future, and E_{it}^7 sums up evaluation events that are more than 6 but fewer than 20 days in the past.

Variables γ_i are evaluator fixed effects that adjust for differences in evaluators' baseline levels of idea generation; δ_t are day fixed effects that control for idea generation trends common to all the evaluators. We also include a vector of time-variant evaluator characteristics, X_{it} . To account for lifecycle effects on the evaluators' idea

generation (Levin and Stephan 1991, Jones 2010), we control for the number of ideas previously generated by the evaluators (*No. of Prev. Ideas*) and the cumulative value of these ideas (*Cum. Idea Value*). As early-career evaluators may have incentives to submit more ideas to advance their careers than their late-career counterparts, we control for whether the evaluators are permanently employed by the firm (*Permanent*), the number of times they have moved between organizational units (*Prev. Mobility*), and their level in the firm's hierarchy (*Hierarchy*). An overview of the variables used in our analysis appears in Table A.6 in the Online Appendix.

Our event study framework requires that there are no pretrends in evaluators' idea generation prior to the evaluation event. Any anticipation effects would raise doubts about the causal effect of idea evaluation on idea generation. We will demonstrate the absence of pretrends in the next section by examining the coefficient estimates for the lead indicator variables.

The event study specification is useful to capture the dynamic effect of idea evaluation on idea generation. However, to simplify our results, we will often report findings from the following single difference-in-differences specification:

$$y_{it} = f(\alpha * Post\ Evaluation_{it} + \beta X_{it} + \gamma_i + \delta_t + \epsilon_{it}). \quad (2)$$


This specification is identical to the event study specification as given in Equation 1 but collapses the lag indicators to the dummy variable *Post Evaluation_{it}* that takes the value of one if day *t* is within one week after evaluator *i* evaluated ideas and zero otherwise. Thus, α indicates changes in the evaluators' likelihood of idea generation if day *t* is within one week after idea evaluation.

5.2. Idea Evaluation and Idea Quantity

5.2.1. Main Results. We begin by descriptively exploring how the number of ideas generated by the evaluators on any given day changes after they evaluated others' ideas. Figure 2 plots the percentage change in ideas generated against the days after idea evaluation. We observe strong crowding-in effects: Evaluators generate 79.8% more ideas on the same day of evaluation and 22.7%, 13.8%, 3.2%, and 4.3% more ideas on week-days 1, 2, 3, and 4, respectively, after evaluation compared with the day before idea evaluation.

Next, we provide more systematic evidence of idea crowding-in by estimating our event study specification. We report the results from logistic regression with *Idea Generation* as the dependent variable as well as from Poisson regression with *No. of Ideas* as the dependent variable in Table A.7 in the Online Appendix. We find that evaluators are 2.62 times more likely to generate ideas on the same day of evaluating others' ideas and 1.30, 1.24, and 1.19 times more likely to generate ideas one, two, and three days, respectively, after idea evaluation relative to the day before idea evaluation (column 1 of Table A.7 in

Figure 2. Effect of Evaluation on Idea Generation: Descriptive Evidence




Notes. This figure illustrates postevaluation changes in the number of generated ideas based on the full sample of generated and evaluated ideas. Changes are relative to the day before idea evaluation.

the Online Appendix). We plot coefficient estimates from column 1 of Table A.7 in Figure 3, which demonstrates the absence of pretrends. Poisson regressions produce very similar results (columns 3 and 4 of Table A.7 in the Online Appendix). We also estimate the event study specification using ordinary least squares (OLS) regression. Column 1 of Table A.8 in the Online Appendix indicates that idea evaluation increases evaluators' likelihood of idea generation by 205.0% on the same day of idea evaluation and by 39.9%, 30.4%, and 29.1% at one, two, and three days, respectively, after idea evaluation relative to the day before idea evaluation. We plot coefficient estimates of column 1 of Table A.8 in Figure A.3 in the Online Appendix, which again confirms the absence of pretrends.

Further, we estimate Specification (2), which collapses the event study lags to the binary treatment variable

Figure 3. (Color online) Effect of Idea Evaluation on Idea Generation: Event Study—Logistic Regression



Notes. This figure displays changes in the odds of idea generation as reported in column 1 of Table A.7 in the Online Appendix. Error bars indicate 95% confidence intervals.

Post Evaluation. Logistic regression suggests that the evaluators are 1.67 times more likely to generate ideas within one week after idea evaluation than outside this postevaluation window (column 1 of Table 1). Similarly, OLS regression predicts a 66.3% increase in the likelihood of idea generation within one week after evaluation (column 1 of Table A.9 in the Online Appendix).

5.2.2. Long-Term Effects. We also examine whether idea evaluation has any long-term effects on idea generation. To do so, we estimate the event study specification at the evaluator-week (rather than evaluator-day) level. Figure A.4 in the Online Appendix demonstrates that coefficients are positive and statistically significant until week 2 following idea evaluation, suggesting that idea evaluation raises the evaluators' idea generation for a sustained period but not permanently.

5.2.3. Heterogeneous Treatment Effects. A growing body of literature highlights that TWFE estimates may be inconsistent if treatment effects vary across time and adoption cohorts (de Chaisemartin and D'Haultfoeuille 2020, Callaway and Sant'Anna 2021, Goodman-Bacon 2021, Sun and Abraham 2021, Borusyak et al. 2024). In our context, the effect of idea evaluation on idea generation may vary over time as employees become more familiar with the IMS submission process. Likewise, evaluators' response to idea evaluation may vary across cohorts owing to varying evaluation workloads.

To address potential treatment effect heterogeneity, we re-estimate our event study and single difference-in-differences models using the robust estimator proposed

by de Chaisemartin and D'Haultfoeuille (2026). This estimator is particularly well suited for our empirical setting, because it accommodates nonabsorbing treatments—those that can be turned on and off—and allows for dynamic treatment effects. As shown in Figure A.5 and Table A.10 in the Online Appendix, the robust estimates are very close to (and in some cases even larger than) the original estimates, alleviating concerns about treatment effect heterogeneity. We provide further details in Section A.2 in the Online Appendix.

5.2.4. Robustness Checks. We take several steps to verify the robustness of our findings, which we outline in the Online Appendix. Specifically, we implement an instrumental variable strategy (Section A.3.1), replicate our main analysis using alternative samples (Section A.3.2), conduct a placebo test (Section A.3.3), consider the number of business days instead of weekdays between idea evaluation and generation (Section A.3.4), and apply CEM (Section A.3.5). Taken together, all these estimates confirm our results, indicating significant idea crowding-in from idea evaluation; that is, evaluators become more productive ideators following idea evaluation.

5.3. Boundary Condition: Is There Too Much Idea Evaluation?

Naturally, there must be limits to crowding-in effects from idea evaluation. Evaluators' days contain only so many hours; if they are all spent on idea evaluation, the number of ideas they generate must drop. Thus, we examine how idea crowding-in varies with evaluators' evaluation workload. We expect idea crowding-in to decrease

Table 1. Effect of Idea Evaluation on Idea Generation: Single Difference-in-Differences—Logistic and Poisson Regression

	Logistic regression		Poisson regression	
	<i>Idea generation</i>		<i>No. of ideas</i>	
	(1)	(2)	(3)	(4)
<i>Post Evaluation</i>	0.512*** [1.669]*** (0.029)	0.477*** [1.611]*** (0.028)	0.523*** [1.687]*** (0.041)	0.479*** [1.614]*** (0.032)
Day fixed effects	X	X	X	X
Evaluator fixed effects	X	X	X	X
Controls		X		X
Dependent variable weighted mean	0.005	0.005	0.007	0.007
No. of evaluators	5,267	5,267	5,267	5,267
No. of days	4,446	4,446	4,446	4,446
No. of observations	412,332	412,332	412,332	412,332
Pseudo R^2	0.152	0.155	0.182	0.185
Wald	303.074***	109.829***	162.928***	94.966***

Notes. This table reports results from Specification (2) using logistic and Poisson regression. The analysis is at the evaluator-day level. Coefficient estimates indicate changes in the log odds of idea generation (columns 1 and 2) and the log number of ideas generated (columns 3 and 4). Exponentiated coefficient estimates, shown in square brackets, indicate changes in the odds of idea generation (columns 1 and 2) and the number of ideas generated (columns 3 and 4). Robust standard errors, clustered at the evaluator level, are in parentheses. All regressions include weights to adjust for choice-based sampling.

* $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$.

with growing evaluation workloads because this leaves evaluators less time to generate ideas of their own.

We begin by visually inspecting how evaluators' evaluation workload affects idea crowding-in. We define evaluators' daily evaluation workload as the number of ideas they evaluate on day t and their weekly evaluation workload as the number of ideas they evaluate in the week preceding day t . We relate evaluators' likelihood of idea generation on day t to their daily and weekly evaluation workloads in panels A and B of Figure A.6 in the Online Appendix. Likelihoods are relative to no idea evaluation (zero evaluation workload). We make two observations. First, moderate evaluation workloads are ideal to trigger idea generation. Second, even high evaluation workloads—at least those that we observe—do not crowd out idea generation, suggesting that high evaluation workloads are still more effective in eliciting evaluator ideas than no evaluation workload at all.

Next, we examine more systematically how idea crowding-in varies with daily evaluation workloads by extending Specification (1). We regress *Idea Generation* on two sets of evaluation lead and lag indicators: one capturing low evaluation workload and the other higher evaluation workload. We define $E_{it}^{\tau, low}$ to take the value of one if evaluator i evaluated fewer than three (but at least one) ideas τ days ago and zero otherwise. Analogously, we define $E_{it}^{\tau, high}$ to take the value of one if evaluator i evaluated three or more ideas τ days ago and zero otherwise. We choose an evaluation workload of three ideas as the threshold based on Figure A.6 in the Online Appendix. Figure A.7 in the Online Appendix compares changes in the likelihood of *Idea Generation* predicted by a one-standard-deviation increase in high versus low daily evaluation workloads. Although both workloads crowd in ideas, low evaluation workloads do so more than high ones. For example, a one-standard-deviation increase in $E_{it}^{0, low}$ raises evaluators' likelihood of idea generation by 75.1%, whereas the same increase in $E_{it}^{0, high}$ raises it by only 38.7%. We arrive at very similar conclusions when examining weekly—rather than daily—evaluation workloads (see Section A.4 in the Online Appendix). Together, this confirms that moderate evaluation workloads are ideal to trigger idea crowding-in.

6. Mechanisms

Next, we examine potential (not mutually exclusive) mechanisms. We find that idea crowding-in is primarily driven by knowledge recombination: Idea evaluation exposes evaluators to unfamiliar knowledge, which they recombine with their extant knowledge to generate ideas. We also explore alternative mechanisms, including a shift in evaluators' attention (Section 6.2), social comparison (Section 6.3), idea plagiarism (Section 6.4), and task bundling (Section 6.5), but find weak evidence for attention shifting only.

6.1. Knowledge Recombination

Idea crowding-in may derive from enhanced opportunities for knowledge recombination. Idea generation is widely conceptualized as a recombinatorial process, in which individuals generate ideas by combining existing knowledge components in unprecedented ways (Nelson and Winter 1982, Weitzman 1998, Fleming 2001). However, because there are only so many ways in which existing knowledge components can be recombined, accessing unfamiliar knowledge components is key to sustaining idea generation (Fleming and Sorenson 2004). Innovation scholars have identified various ways of expanding individuals' knowledge pool by such unfamiliar components. Notably, teamwork (Sutton and Hargadon 1996, Wuchty et al. 2007, Singh and Fleming 2010) and employee mobility—both within organizations (e.g., across units, teams, and sites) (Karim and Kaul 2015, Cornelius et al. 2021, Kang and Eklund 2025) and between organizations (Singh and Agrawal 2011)—have been shown to be effective in pooling diverse knowledge components and facilitating knowledge recombination.

We propose that idea evaluation may offer evaluators a unique opportunity to access unfamiliar knowledge. By assessing others' ideas, they may encounter new methods and processes and may recall forgotten knowledge (Haunschild et al. 2015, Ramdas et al. 2018, Argote et al. 2021). We provide several pieces of evidence that support knowledge recombination as the key mechanism.

6.1.1. Idea Similarity. If idea crowding-in was driven by evaluators building on the ideas they evaluate, the ideas generated postevaluation should resemble the evaluated ideas. To test this resemblance hypothesis, we examine whether and how the similarity between idea k , evaluated by evaluator i , and idea l , generated by the same evaluator i , depends on whether evaluator i generated idea l within one week after evaluating idea k or outside this one-week postevaluation window. Specifically, we estimate the following specification:

$$\text{Similarity}_{iklt} = \alpha \text{Post Evaluation}_{it} + \beta \mathbf{X}_k + \gamma \mathbf{Z}_l + \epsilon_{iklt}. \quad (3)$$

We define *Post Evaluation* as usual; $\mathbf{X}_k$ is a vector of evaluated idea controls, and $\mathbf{Z}_l$ is a vector of generated idea controls. We use three measures of *Similarity*. The first is the *Cosine Similarity* between the titles of evaluated and generated ideas. The second, *Share Same Topic—NLP-Assigned*, captures the share of topics common to evaluated and generated ideas, based on topics derived from NLP models. The third, *Share Same Topic—Pre-Assigned*, also captures the share of common topics, but uses topics preassigned by the firm. All *Similarity* measures range from zero (low similarity) to one (high similarity). Further details on the construction of these measures are provided in Section A.6 in the Online Appendix.

We find strong evidence for the resemblance hypothesis. Figure A.8 in the Online Appendix compares the distribution of *Cosine Similarity* for evaluated-generated idea dyads where evaluators generated the idea within one week after evaluation to the distribution for dyads where idea generation occurred outside this time window. The distribution of ideas generated postevaluation is significantly more right-skewed, suggesting that evaluators build on the ideas they evaluate. OLS regressions of *Idea Similarity* on *Post Evaluation* further support the resemblance hypothesis (Table A.12 in the Online Appendix). For example, we find that ideas generated postevaluation are 10.6% more similar (column 2) to the evaluated ideas and have 44.7% (column 4) more topics in common with the evaluated ideas than those generated outside the postevaluation window. In the Online Appendix, we demonstrate that these findings are robust to different measures of *Share Same Topic—NLP-Assigned* (Table A.13).

6.1.2. Familiar vs. Unfamiliar Knowledge. Knowledge recombination suggests that idea crowding-in is driven by evaluation exposing evaluators to unfamiliar knowledge. Thus, crowding-in should be stronger when evaluators assess ideas containing unfamiliar rather than familiar topics. We test this prediction, defining a topic

as unfamiliar to an evaluator if none of their previous idea generation or evaluation activities are connected to it. To infer evaluators' (lack of) familiarity with topics, we rely on the topics assigned to ideas using NLP models (see Section A.5). Building on Specification (2), we regress *Idea Generation* on *Post Evaluation Unfamiliar* and *Post Evaluation Familiar*, which capture the share of unfamiliar and familiar topics, respectively, contained in the ideas that evaluator i evaluated during the preceding week. We normalize these variables to have a mean of zero and a standard deviation of one for easier comparison across coefficient estimates.

Table 2 attests to the importance of unfamiliar knowledge for idea crowding-in. Column 4 indicates a 4.0% increase in the evaluators' likelihood of idea generation in response to a one-standard-deviation increase in *Post Evaluation Familiar* but a 11.3% increase in this likelihood in response to a one-standard-deviation increase in *Post Evaluation Unfamiliar*. That is, evaluating ideas containing unfamiliar knowledge is almost three times as effective in eliciting evaluator ideas than evaluating ideas containing familiar knowledge. By controlling for the numbers of ideas that evaluators previously evaluated and generated, we can rule out that our finding is driven by evaluators naturally encountering more unfamiliar topics early on in their organizational careers, which

Table 2. Mechanism: Recombination—Exposure to Unfamiliar Knowledge

	OLS regression: Dependent variable = <i>Idea Generation</i>					
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Post Evaluation Familiar</i>	0.0003*** [6.1159]*** (0.0001)		0.0002** [4.0789]** (0.0001)	0.0002*** [4.0030]*** (0.0001)	0.0002** [3.5285]** (0.0001)	0.0002** [3.7179]** (0.0001)
<i>Post Evaluation Unfamiliar</i>		0.0007*** [12.8379]*** (0.0001)	0.0006*** [11.9639]*** (0.0001)	0.0006*** [11.3432]*** (0.0001)	0.0013*** [24.2794]*** (0.0002)	0.0030*** [56.8652]*** (0.0004)
<i>Post Evaluation Unfamiliar</i> × <i>Idea Team Size</i>					−0.0002* [−4.4547]* (0.0001)	
<i>Post Evaluation Unfamiliar</i> × <i>Cross-Unit Ideation</i>						−0.0018*** [−34.6048]*** (0.0003)
Evaluator fixed effects	X	X	X	X	X	X
Day fixed effects	X	X	X	X	X	X
Controls				X	X	X
Dependent variable weighted mean	0.0052	0.0052	0.0052	0.0052	0.0052	0.0052
No. of evaluators	12,069	12,069	12,069	12,069	12,069	12,069
No. of days	4,446	4,446	4,446	4,446	4,446	4,446
No. of observations	419,134	419,134	419,134	419,134	419,134	419,134
Adjusted R^2	0.0215	0.0215	0.0216	0.0217	0.0232	0.0232

Notes. This table reports results from Specification (2) using OLS regression. The analysis is at the evaluator-day level. We define *Post Evaluation Unfamiliar* and *Post Evaluation Familiar* as the share of unfamiliar and familiar topics, respectively, contained in the ideas that an evaluator evaluated during the preceding week. We normalize these variables to have a mean of zero and a standard deviation of one such that coefficient estimates indicate percentage point changes in the likelihood of idea generation in response to a one-standard-deviation increase in the independent variables. We define *Idea Team Size* as the average number of employees involved in idea generation in the evaluator's unit, and *Cross-Unit Ideation* as the average number of organizational units involved in the evaluator's unit. Percentage changes, shown in square brackets, are calculated by dividing percentage point changes by the weighted mean of the dependent variable. Robust standard errors, clustered at the evaluator level, are reported in parentheses. All regressions include weights to adjust for choice-based sampling.

* $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$.

may also be when they are most eager to generate ideas of their own. This finding is robust to alternative measures of topic (un)familiarity, including whether topics were covered within the past three years or one year, as well as to variations in the number of topics assigned to ideas (Tables A.14 and A.15 in the Online Appendix).

As additional support for the knowledge recombination mechanism, we find the evaluators with otherwise limited access to unfamiliar knowledge to benefit most from evaluating unfamiliar topics. Previous research suggests that individuals can gain access to unfamiliar knowledge through teamwork (Girotra et al. 2010, Singh and Fleming 2010) and cross-unit collaboration (Singh 2008, Lahiri 2010, Cornelius et al. 2021). If access to these alternative sources is limited, the effect of evaluating unfamiliar ideas should be especially strong. Conversely, evaluators with greater access to such sources should benefit less from evaluating unfamiliar ideas. To test this argument, we interact *Post Evaluation Unfamiliar* with *Idea Team Size*, which captures the average number of employees per generated idea (column 5 of Table 2), and *Cross-Unit Ideation*, which captures the average number of organizational units per generated idea (column 6 of Table 2). We measure both variables at the level of the evaluator's unit—that is, they reflect the average collaboration patterns in idea generation within that unit. We find that the benefits of evaluating unfamiliar ideas decline by 4.5% and 34.6% for every additional employee and organizational unit contributing to an idea, respectively—consistent with the notion that evaluation substitutes for other sources of unfamiliar knowledge.

Finally, to further validate the importance of unfamiliar knowledge for crowding-in, we adopt a more implicit measure of topic (un)familiarity, namely whether the evaluation of an idea involves coevaluators. Recall that the lead evaluators may call in coevaluators whenever they feel that they lack the knowledge needed to evaluate ideas. Thus, the involvement of coevaluators implies that an idea spans disparate knowledge domains, with some of them unfamiliar to the lead evaluator and possibly also to coevaluators. We find coevaluation to be almost twice as effective in eliciting evaluator ideas as solo evaluation (columns 1–3 of Table A.16 in the Online Appendix). Coevaluation is especially effective when evaluators join the evaluation process at a later stage and have access to previous evaluation reports (columns 4–6 of Table A.16 in the Online Appendix), suggesting that previous reports help evaluators to understand and contextualize unfamiliar topics.

Taken together, these results provide strong evidence that access to unfamiliar knowledge plays a central role in driving idea crowding-in.

6.1.3. Problem-Based vs. Solution-Based Recombination. Finally, we explore the nature of knowledge recombination, distinguishing between problem-based

recombination (when evaluators build on the problems addressed in others' ideas and propose alternative solutions) and solution-based recombination (when they apply the solutions proposed in others' ideas to problems they face elsewhere) (von Hippel and von Krogh 2016). Problem- and solution-based recombination have distinct implications for organizational performance. Problem-based recombination helps organizations to find superior solutions to well-defined existing problems but does not expand the number of problems being addressed. In contrast, solution-based recombination helps solve previously unaddressed problems by transferring solutions across domains but does not improve existing solutions.

Empirically, to assess the relative importance of problem- versus solution-based recombination, we examine whether evaluated and generated ideas become more similar along their problem versus solution components following evaluation. In Section A.7 in the Online Appendix, we explain how we measure the problem and solution similarities of idea dyads. We find that evaluators draw more heavily on the problems than on the solutions of the ideas they evaluated. Figure A.9 in the Online Appendix shows that problem similarity becomes a stronger predictor of overall idea similarity after evaluation (left), whereas the predictive power of solution similarity remains unchanged before and after evaluation (right). OLS regressions further support the dominance of problem-based recombination. Although both problem and solution similarity explain the total similarity between idea dyads, the explanatory power of problem similarity increases by 31.1% within the week following evaluation, whereas that of solution similarity remains unchanged (column 5 of Table A.17 in the Online Appendix). This suggests that idea evaluation is particularly effective in transferring problem-related knowledge from evaluatees to evaluators, that is, in surfacing problems, for which evaluators then go on to find better solutions.

We observe that problem-based recombination is especially likely when the problems are relatively unexploited—that is, when they have received little attention in the recent past—compared with problems that have been heavily addressed (columns 1 and 2 of Table A.18 in the Online Appendix). This finding aligns with research emphasizing the value of directing search toward underexplored areas to sustain organizational learning and adaptation (March 1991, Levinthal and March 1993). In addition, problem-based recombination is most likely when the problem addressed in the evaluated idea is relatively distant from those the evaluator previously explored in their own idea generation (columns 3 and 4 of Table A.18 in the Online Appendix). In such cases, evaluation appears to expand the set of problems that evaluators are able to address.

Together, our findings indicate that idea crowding-in operates by granting evaluators access to unfamiliar knowledge, which facilitates (problem-based) recombination.

6.2. Shift in Attention

Idea crowding-in may also be driven by evaluation shifting employees' attention to idea generation. Much research has documented employees' tendency to allocate their attention to easy routine tasks at the expense of more complex nonroutine ones such as idea generation (Repenning 2001, Sullivan 2010). March and Simon (1958, p. 206) refer to this as Gresham's law: "When an individual is faced both with highly programmed and highly unprogrammed tasks, the former tend to take precedence over the latter." How, then, can employees' attention be shifted toward nonroutine tasks such as idea generation? One way is to expose them to stimuli that increase the salience of nonroutine tasks (Ocasio 2011). Idea evaluation may provide such stimuli, increasing the salience of idea generation and thus employees' likelihood to generate ideas of their own.

This shift-in-attention explanation suggests stronger idea crowding-in when evaluators have not interacted with the IMS for a long time. In contrast, if their most recent interaction was just yesterday, their attention will already be focused on idea generation. Table A.19 in the Online Appendix provides evidence that a shift in attention accounts, at least to some extent, for idea crowding-in. In columns 1 and 2, we interact *Post Evaluation* with *Years Since IMS Interaction*, which is the number of years since evaluators last interacted with the IMS, either by submitting or evaluating ideas. We observe stronger crowding-in with increasing temporal distance to evaluators' last IMS interaction. The effect of *Post Evaluation* on *Idea Generation* increases by 27.1% with every year of no IMS interaction (column 2). A split-sample analysis confirms this finding (columns 3–6). We split the sample based on the median value of *Years Since IMS Interaction*, which is 0.10 years or 37 days. Although we observe idea crowding-in in both subsamples, it is much stronger in the sample where the last IMS interaction was longer ago.

Although this provides evidence that attention-shifting contributes to idea crowding-in, we find its role to be small compared with recombination, as can be seen from the coefficient estimates across Table 2 and Table A.19 in the Online Appendix. Because the evaluation treatments are standardized in both analyses, with a mean of zero and standard deviation of one, direct comparison is feasible. According to column 4 of Table 2, a one-standard-deviation increase in *Post Evaluation Unfamiliar* raises the likelihood of idea generation by 11.3%, whereas an equivalent increase in *Post Evaluation Familiar* raises it by only 4.0%—a difference of 7.3 percentage points. Thus, to match the effect of access to unfamiliar knowledge would require evaluators to remain inactive for an additional 0.27 years ($= 7.3/27.1$), or just over three months—extending the average inactivity period from 1.5 to 4.5 months, a duration of inactivity observed in fewer than 15% of cases.

6.3. Social Comparison

Idea crowding-in may derive from social comparison processes (Festinger 1954). Employees who observe their coworkers and subordinates generate ideas may feel pressure to contribute their fair share to ideation (Mas and Moretti 2009, Bandiera et al. 2010). Or, to the extent that they perceive their coworkers to be competitors, they may want to outperform them (Chan et al. 2014, Baumann et al. 2019). In both cases, we would expect idea crowding-in to increase with the social proximity between evaluators and evaluatees.

To test this prediction, we compare the differential effects of evaluating ideas by socially proximate versus distant evaluatees. We use various measures of social proximity, including evaluators and evaluatees working in the same unit, sharing the same gender, or having the same record of intraorganizational mobility. Across various specifications in Table A.20 in the Online Appendix, we find coefficient estimates for the evaluation of ideas by socially proximate and distant evaluatees to be highly similar, suggesting that social comparison is unlikely to explain idea crowding-in.

6.4. Idea Plagiarism

Our results may be driven by evaluators plagiarizing the ideas they evaluate. If this was the case, idea crowding-in would occur primarily after evaluators reject ideas. After all, evaluators can only receive credit for ideas that are not already selected for implementation. To test for plagiarism, we re-estimate Specification (1) but distinguish between two sets of lead and lag indicators: one capturing evaluators selecting ideas and the other capturing evaluators rejecting ideas. We compare coefficient estimates for the selection versus rejection of ideas in Figure A.10 in the Online Appendix. The coefficient estimates are virtually identical, suggesting that plagiarism is an unlikely explanation for idea crowding-in.

6.5. Efficient Bundling

The observed patterns may also be explained by evaluators scheduling idea evaluation and idea submission activities to achieve efficiency gains. First, they may handle the evaluation and submission of ideas in a single IMS session to minimize logon costs. In this case, evaluators may have previously generated ideas but postponed their submission until evaluation tasks accumulated, allowing them to complete both activities in one session. However, such *system-logon bundling* can account only for the increase in submissions on the day of evaluation; it cannot explain the continued increase on the following days, each of which requires a separate IMS log-on.

Second, the patterns may reflect *task bundling*, where evaluators complete one set of tasks first before turning to the other to reduce switching costs. However, *task bundling* does not imply a fixed task order. It cannot

explain why evaluators consistently evaluate ideas before submitting their own but not the reverse. Thus, *task bundling* fails to account for our results.

7. Idea Evaluation and Idea Quality

As shown in the previous section, evaluation leads to a significant increase in idea *quantity*. The present analysis explores whether these additional ideas differ systematically regarding their *quality*. We focus on two key dimensions of idea quality: novelty (Boudreau et al. 2016, Criscuolo et al. 2017, Lane et al. 2024) and value (Girotra et al. 2010, Kaplan and Vakili 2015, Ghosh and Wu 2023).

To assess the impact of idea evaluation on quality, we compare the quality of ideas from the same evaluator, where some were generated within one week of evaluation and others outside this timeframe. Specifically, we regress the novelty and value of ideas on $Post\ Evaluation_{it}$, which—as above—is one if day t is within one week since evaluator i evaluated an idea and zero otherwise. Positive and statistically significant coefficient estimates for $Post\ Evaluation$ would suggest that idea evaluation is linked to higher idea novelty and value. We provide more details on the empirical approach in Section A.8 of the Online Appendix.

7.1. Idea Novelty

We follow a recombinatorial view of novelty (Boudreau et al. 2016, Criscuolo et al. 2017, Lane et al. 2024), measuring idea novelty as the extent to which ideas combine topics in unprecedented ways. We assign topics to ideas using NLP methods (for details, see Section A.5 in the Online Appendix). Every idea is mapped to N topics—those that best describe it—selected from a set of M possible topics. From these N topics, we derive the $[N \times (N - 1)]/2$ topic combinations that characterize the idea and assess how many of these combinations have not been made before. The novelty score reflects the share of

topic combinations that are new: A score of one indicates that all topic combinations are novel, whereas a score of zero indicates that none are novel.


Figure A.11 in the Online Appendix shows descriptively that the ideas generated within one week after evaluation hardly differ in their novelty from those generated at other times. Table A.21 in the Online Appendix confirms this finding: *Post Evaluation* has no statistically significant effect on idea novelty. This result is robust to different parameter choices in our topic assignment procedure, including different values for $M \in \{500, 1,000\}$ and $N \in \{3, 5, 7\}$. We therefore conclude that idea evaluation does not affect the novelty of generated ideas.

7.2. Idea Value

Next, we examine how idea evaluation affects the value of ideas. As noted, idea value is calculated as the costs saved or revenues generated within the first year after implementation minus the implementation costs. Figure 4 suggests that the ideas generated within one week after evaluation are more valuable than the ideas generated outside this postevaluation window. As seen in Figure 4(a), the share of zero-valued ideas is 61.2% for ideas generated postevaluation but 71.7% for ideas generated outside the postevaluation window. Figure 4(b) shows that postevaluation ideas are disproportionately represented in the right tail of the value distribution.

Regression analysis confirms this finding. Table 3 shows that the ideas generated within one week after evaluation are 19.5% ($= 100 * (e^{0.178} - 1)$) more valuable, on average, than the ideas generated outside this time window (column 1). We also examine whether idea evaluation increases the evaluators' likelihood of generating ideas that the organization finds particularly valuable—so-called extreme-value ideas (Girotra et al. 2010, Singh and Fleming 2010). *Top 1*, *Top 5*, and *Top 10* take the value of one if an idea ranks among the 1%, 5%, and 10% most valuable ideas among all ideas generated in a given year,

Figure 4. Relationship Between Evaluation and Value of Generated Ideas



Notes. This figure compares the value distribution of ideas generated within one week after idea evaluation against the value distribution of ideas generated outside this time window. It is based on the sample of selected ideas, as the value is only known for these ideas. To preserve confidentiality, we transformed $\log(\text{Idea Value} + 1)$ by normalizing it to have a mean of one and a standard deviation of one.

Table 3. Effect of Idea Evaluation on Idea Value: OLS Regressions

	OLS regression				
	<i>Log(Idea Value + 1)</i>		<i>Top 1</i>	<i>Top 5</i>	<i>Top 10</i>
	(1)	(2)	(3)	(4)	(5)
<i>Post Evaluation</i>	0.178** [19.475]** (0.062)	0.192** [21.189]** (0.062)	0.006** [58.781]** (0.002)	0.007+ [13.983]+ (0.004)	0.017*** [16.849]** (0.005)
Day fixed effects	X	X	X	X	X
Evaluator fixed effects	X	X	X	X	X
Controls		X	X	X	X
No. of evaluators	4,633	4,633	4,633	4,633	4,633
No. of days	4,284	4,284	4,284	4,284	4,284
No. of observations	108,938	108,938	108,938	108,938	108,938
Adjusted R ²	0.382	0.386	0.238	0.257	0.262

Notes. This table reports results from Specification 4 using OLS regression. The analysis is at the idea-evaluator-day level. Coefficient estimates indicate percentage point changes in the respective dependent variable. Percentage changes, shown in square brackets, are calculated by dividing percentage point changes by the weighted mean of the dependent variable. Robust standard errors are in parentheses and clustered at the evaluator level. Estimations are based on the sample of selected ideas, as the value is only known for these ideas.

+ $p < 0.1$; * $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$.

respectively. Columns 3–5 of Table 3 indicate that evaluators have a 58.8%, 14.0%, and 16.8% higher probability of generating *Top 1*, *Top 5*, and *Top 10* ideas, respectively, within one week after idea evaluation than outside this time window. These findings are very robust to alternative definitions of extreme-value ideas, for example, when comparing ideas generated in the same *month* or *week* rather than year. The positive effect of idea evaluation on idea value is similarly short lived as its effect on idea quantity, persisting for only about two weeks (see Tables A.22 and A.23 in the Online Appendix).

To examine whether these findings are driven by ideas becoming more valuable across the entire idea distribution (upward mean shift) or mainly extreme-value ideas becoming more valuable (upper quantile shift), we estimate quantile regressions (Girotra et al. 2010, Singh and Fleming 2010, Boudreau et al. 2011). Coefficient estimates from quantile regressions indicate the impact of idea evaluation on particular quantiles of the value distribution of ideas rather than on its conditional mean (Koenker and Bassett 1978). As seen in Table A.24 in the Online Appendix, all coefficient estimates are statistically significant and positive, reflecting the general positive effect of idea evaluation on idea value. We do not find larger effect sizes for higher value quantiles, suggesting that our findings derive from a mean rather than an upper quantile shift.

8. Discussion

8.1. Contributions

This paper makes several contributions. Our first contribution is to highlight a hitherto unrecognized function of idea evaluation, which is to encourage idea generation among evaluators. Extant research has emphasized the importance of idea evaluation in resource allocation

(Terwiesch and Ulrich 2009, Boudreau et al. 2016, Kornish and Hutchison-Krupat 2017), with a particular focus on how biased evaluation can lead to inefficient resource allocation (Mueller et al. 2012, Criscuolo et al. 2017, Li 2017). In a second stream, scholars have explored how idea evaluations, once communicated to ideators, shape their subsequent ideation behaviors (Deichmann and van den Ende 2014, Wooten and Ulrich 2017, Piezunka and Dahlander 2019). We establish a third key function of idea evaluation beyond resource allocation and feedback, which is to encourage idea generation among evaluators. We thereby also extend a sparse literature on interactions between the idea generation and evaluation phases (Berg 2016, Perry-Smith and Mannucci 2017).

Second, we contribute to innovation scholarship by introducing the concept of *idea mobility*—the movement of ideas within an organization. Research has long recognized the significance of employee mobility—the movement of employees across teams, units, or organizations—for knowledge production (Choudhury 2022). As employees transition between contexts, they carry sticky knowledge, making it accessible to others and facilitating knowledge recombination and generation (Argote and Ingram 2000, Singh and Agrawal 2011, Stadler et al. 2022). For example, the relocation of employees to new office sites (Catalini 2018) or organizational units (Karim and Kaul 2015), even if only for some time (Cornelius et al. 2021, Kang and Eklund 2025), has been shown to enhance innovation performance. However, as employees increasingly work remotely, the traditional benefits of employee mobility—such as spontaneous interaction, knowledge exchange, and recombination—may be diminished. We show that *idea mobility* can complement, and perhaps even substitute for, employee mobility by facilitating knowledge transfer across

organizational members. Unlike employee mobility, idea mobility operates virtually: Ideas can circulate rapidly across individuals, often within days, and enable recombination even in remote settings. Our findings suggest that idea mobility has the greatest recombinatorial potential when ideas travel long distances—passing through many evaluators who, through their evaluations, make new knowledge accessible to other evaluators. In the end, it is not only the original idea but also the knowledge codified through these evaluations that becomes a potent source for recombination and innovation.

Our third contribution is to unpack the microprocesses of knowledge recombination by utilizing advances in NLP. Recent research highlights the potentials of NLP models for developing and testing management theory (Arts et al. 2023, Guzman and Li 2023, Aceves and Evans 2024). Although the concept of knowledge recombination has been a cornerstone of the innovation literature for decades (Nelson and Winter 1982, Weitzman 1998, Fleming 2001), capturing how it unfolds at the individual level has remained challenging owing to methodological constraints. For instance, researchers have had to rely on predefined patent classifications (Fleming and Sorenson 2004) or manually coded idea topics (Girotra et al. 2010) to track recombination.

NLP models enable us to trace knowledge recombination with unprecedented precision. This allows us, first, to confirm the crucial role of access to unfamiliar knowledge for successful recombination (Fleming and Sorenson 2004). Second, we are able to empirically assess the relative importance of problem- versus solution-based recombination. Although these two forms have been conceptually distinguished before (von Hippel and von Krogh 2016), to our best knowledge, this study is the first to demonstrate that evaluation makes employees more likely to build on others' problem information to propose their own solutions rather than apply existing solutions to address problems they encounter elsewhere.

8.2. Managerial Implications

Practice and research have used various tools and incentives to engage employees in idea generation. Many follow the direct approach of inviting or requiring more ideas. Our research suggests that having employees evaluate others' ideas is an indirect yet powerful way of tapping into employees' creativity. To fully reap the benefits of idea evaluation for employee ideation, we recommend that organizations reconsider who they assign evaluation tasks to and when.

First, we recommend that organizations assign ideas to multiple evaluators who, together, cover all relevant areas but individually are unfamiliar with some aspects of the idea. Although previous research highlights the importance of expertise for effective evaluation (Li 2017), our findings point to a tradeoff: Evaluators are

more likely to generate ideas of their own when exposed to unfamiliar rather than familiar content. To balance evaluation quality with the potential for idea crowding-in, organizations can adopt a model akin to academic peer review, where reviewers are selected for complementary expertise (e.g., theory versus method). This approach ensures high-quality evaluation while exposing evaluators to novel content that can spark knowledge recombination.

Second, organizations should consider assigning idea evaluation tasks to a broader pool of employees, including those without a previous track record in ideation. Expanding participation in evaluation can help reap the benefits of idea crowding-in and can contribute to a fuller idea pipeline. It can also promote idea diversity. We have seen that evaluation tasks tend to be assigned to employees with strong ideation records, which can create a self-reinforcing cycle: Successful ideators are more likely to evaluate, which further enhances their own ideation, increasing their chances of being selected as evaluators again, and so forth. Over time, ideation and evaluation could become concentrated among a small subset of employees, potentially at the expense of idea diversity. This is problematic, because previous research shows that newcomers and peripheral actors are often key sources of outlier ideas (Jeppesen and Lakhani 2010, Sgourev 2013, Cattani et al. 2017). To sustain diversity in idea generation, organizations should therefore consider assigning evaluation tasks not only to established contributors but also to employees with little or no ideation experience. Along these lines, some firms have begun experimenting with distributed evaluation models, where employees—regardless of their ideation history and hierarchical position—evaluate one another's ideas (Schweisfurth et al. 2023).

8.3. Limitations and Future Research

Our study has limitations, which open avenues for future research. First, we have presented a *partial* analysis of employee behavior; that is, we examined the effect of idea evaluation on idea generation but disregarded its effects on employees' other activities. Specifically, because employees' time and attention are limited, idea crowding-in must come at the expense of some third, competing activity. Unfortunately, we cannot tell what these competing activities are, or whether employees would have added more value by engaging in these competing activities than they did by generating ideas. We encourage researchers to explore the costs and benefits of idea crowding-in in this light.

Second, we may have underestimated the true magnitude of idea crowding-in, because we captured only the ideas submitted and evaluated through the firm's IMS while missing out on those that bypassed the system. Employees are likely to submit many of their ideas through the IMS, not least because this is a simple way

for them to secure bonus payments. Still, some employees, especially those with ties to upper management, may prefer to propose and advance their ideas through more informal channels, for example, over lunch or in board meetings. Like the ideas submitted through the IMS, these ideas may crowd in additional ideas, which we failed to observe.

Third, we were unable to fully account for why idea evaluation has no effect on the novelty of subsequently generated ideas. At first glance, this null result may seem counterintuitive, because we identified access to unfamiliar knowledge as a key driver of idea crowding-in. One may reasonably expect such exposure to foster novel combinations of topics. A possible explanation for the lack of an effect on idea novelty is that, although evaluation may allow evaluators to recombine knowledge components that are *novel to themselves*, these combinations may not be *novel to the organization*—which is how the literature conceives of idea novelty and what our novelty measure captures (Boudreau et al. 2016, Lane et al. 2024). Further investigation of the relationship between idea evaluation and idea novelty remains a promising direction for future research.

Finally, we caution against interpreting the observed crowding-in effect as definitive evidence of a causal effect of idea evaluation on idea generation. The assignment of evaluation tasks was not random, which makes it impossible to completely rule out that unobserved factors influence both the likelihood of being selected to evaluate and the likelihood to generate ideas. We employed rigorous identification strategies to strengthen causal inference, including recent advances in two-way fixed effects models, instrumental variables, and matching. Still, we encourage future research to build on our findings through randomized field experiments that reinforce the causal effect of evaluation on ideation.

8.4. Generalizability of Idea Crowding-In

We have documented idea crowding-in within a single manufacturing firm, which calls for caution when generalizing our findings. However, many features of this firm's IMS are typical of innovation programs in multinational companies (Deichmann and van den Ende 2014, Fuchs et al. 2019, Cornelius et al. 2021). For example, idea submission is voluntary and incentivized, and evaluators are also expected to contribute their own ideas. This dual role—where employees both generate and evaluate ideas—is increasingly common as innovation becomes more distributed and is central to understanding when crowding-in is most likely to occur (Schweisfurth et al. 2023).

Individuals engage in idea evaluation and generation also outside their company context. In online communities (Smirnova et al. 2022) and crowdsourcing (Klapper et al. 2024, Riedl et al. 2024), for example, participants

submit their own ideas while assessing others' submissions; similarly, the scientific community depends on researchers contributing ideas while reviewing manuscripts and grant proposals (Boudreau et al. 2016, Li 2017, Ayoubi et al. 2025). We anticipate that idea evaluation will also foster idea generation in these contexts, given that the core mechanism behind idea crowding-in—knowledge recombination—should also operate in these settings. We encourage researchers to explore whether this holds true.

Acknowledgments

The authors thank the editor, associate editor, and three anonymous reviewers for exceptionally constructive and thoughtful feedback; Prisca Brosi, Paola Criscuolo, Thomas Gillier, Tom Grad, Carolin Haeussler, Lars Bo Jeppesen, H. C. Kongsted, Henry Sauermann, and Arne Thomas for helpful feedback; and participants at the Annual Meeting of the Academy of Management (2023), the DRUID Conference (2023), the Open and User Innovation Conference (2023), the VHB Technology, Innovation, and Entrepreneurship Conference (2023), and the Innovation and Product Development Conference (2022).

Endnote

¹ For confidentiality reasons, we cannot disclose the value of the ideas.

References

- Aceves P, Evans JA (2024) Mobilizing conceptual spaces: How word embedding models can inform measurement and theory within organization science. *Organ. Sci.* 35(3):788–814.
- Agrawal A, Catalini C, Goldfarb A, Luo H (2018) Slack time and innovation. *Organ. Sci.* 29(6):1056–1073.
- Ahmadi S, Jansen JJP, Eggers JP (2022) Using stretch goals for idea generation among employees: One size does not fit all! *Organ. Sci.* 33(2):671–687.
- Amabile T (1988) A model of creativity and innovation in organizations. *Res. Organ. Behav.* 10:123–167.
- Argote L, Ingram P (2000) Knowledge transfer: A basis for competitive advantage in firms. *Organ. Behav. Human Decision Processes* 82(1):150–169.
- Argote L, Lee S, Park J (2021) Organizational learning processes and outcomes: Major findings and future research directions. *Management Sci.* 67(9):5399–5429.
- Arts S, Cassiman B, Hou J (2023) Position and differentiation of firms in technology space. *Management Sci.* 69(12):7253–7265.
- Ayoubi C, Zume Dumlao JM, Teplitskiy M (2025) Learning by evaluating: Evidence from academic peer review. Preprint, submitted July 14, <http://dx.doi.org/10.2139/ssrn.5339282>.
- Bandiera O, Barankay I, Rasul I (2010) Social incentives in the workplace. *Rev. Econom. Stud.* 77(2):417–458.
- Baumann O, Eggers JP, Stieglitz N (2019) Colleagues and competitors: How internal social comparisons shape organizational search and adaptation. *Admin. Sci. Quart.* 64(2):275–309.
- Bayus BL (2013) Crowdsourcing new product ideas over time: An analysis of the Dell IdeaStorm community. *Management Sci.* 59(1):226–244.
- Bell JJ, Pescher C, Tellis GJ, Füller J (2024) Can AI help in ideation? A theory-based model for idea screening in crowdsourcing contests. *Marketing Sci.* 43(1):54–72.
- Berg JM (2016) Balancing on the creative highwire. *Admin. Sci. Quart.* 61(3):433–468.

- Bjelland OM, Wood RC (2008) An inside view of IBM's 'innovation jam'. *MIT Sloan Management Rev.* 50(1):32–40.
- Borusyak K, Jaravel X, Spiess J (2024) Revisiting event-study designs: Robust and efficient estimation. *Rev. Econom. Stud.* 91(6):3253–3285.
- Boudreau KJ, Lacetera N, Lakhani KR (2011) Incentives and Problem uncertainty in innovation contests: An empirical analysis. *Management Sci.* 57(5):843–863.
- Boudreau KJ, Guinan EC, Lakhani KR, Riedl C (2016) Looking across and looking beyond the knowledge frontier: Intellectual distance, novelty, and resource allocation in science. *Management Sci.* 62(10):2765–2783.
- Burtch G, He Q, Hong Y, Lee D (2022) How do peer awards motivate creative content? Experimental evidence from Reddit. *Management Sci.* 68(5):3488–3506.
- Callaway B, Sant'Anna PH (2021) Difference-in-differences with multiple time periods. *J. Econom.* 225(2):200–230.
- Camacho N, Nam H, Kannan PK, Stremersch S (2019) Tournaments to crowdsource innovation: The role of moderator feedback and participation intensity. *J. Marketing* 83(2):138–157.
- Catalini C (2018) Microgeography and the direction of inventive activity. *Management Sci.* 64(9):4348–4364.
- Cattani G, Ferriani S, Lanza A (2017) Deconstructing the outsider puzzle: The legitimization journey of novelty. *Organ. Sci.* 28(6):965–992.
- Chan TY, Li J, Pierce L (2014) Compensation and peer effects in competing sales teams. *Management Sci.* 60(8):1965–1984.
- Chan KW, Li SY, Ni J, Zhu JJ (2021) What feedback matters? The role of experience in motivating crowdsourcing innovation. *Production Oper. Management* 30(1):103–126.
- Choudhury P (2022) Geographic mobility, immobility, and geographic flexibility: A review and agenda for research on the changing geography of work. *Acad. Management Ann.* 16(1):258–296.
- Cornelius PB, Gokpinar B, Sting FJ (2021) Sparking manufacturing innovation: How temporary interplant assignments increase employee idea values. *Management Sci.* 67(4):2231–2250.
- Criscuolo P, Dahlander L, Grohsjean T, Salter A (2017) Evaluating novelty: The role of panels in the selection of R&D projects. *Acad. Management J.* 60(2):433–460.
- Cropley A (2006) In praise of convergent thinking. *Creativity Res. J.* 18(3):391–404.
- Cyert R, March JG (1963) *A Behavioral Theory of the Firm* (Prentice Hall/Pearson Education, Upper Saddle River, NJ).
- Dahan E, Mendelson H (2001) An extreme-value model of concept testing. *Management Sci.* 47(1):102–116.
- Dahlander L, Piezunka H (2014) Open to suggestions: How organizations elicit suggestions through proactive and reactive attention. *Res. Policy* 43(5):812–827.
- de Chaisemartin C, D'Haultfoeuille X (2020) Two-way fixed effects estimators with heterogeneous treatment effects. *Amer. Econom. Rev.* 110(9):2964–2996.
- de Chaisemartin C, D'Haultfoeuille X (2026) Difference-in-differences estimators of intertemporal treatment effects. *Rev. Econom. Statist.*
- Deichmann D, Jensen M (2018) I can do that alone ... or not? How idea generators juggle between the pros and cons of teamwork. *Strategic Management J.* 39(2):458–475.
- Deichmann D, van den Ende J (2014) Rising from failure and learning from success: The role of past experience in radical initiative taking. *Organ. Sci.* 25(3):670–690.
- Ederer F, Manso G (2013) Is pay for performance detrimental to innovation? *Management Sci.* 59(7):1496–1513.
- Edmondson AC (2012) *Teaming: How Organizations Learn, Innovate, and Compete in the Knowledge Economy* (Jossey-Bass, NJ).
- Festinger L (1954) A theory of social comparison processes. *Human Relations* 7(2):117–140.
- Fleming L (2001) Recombinant uncertainty in technological search. *Management Sci.* 47(1):117–132.
- Fleming L, Sorenson O (2004) Science as a map in technological search. *Strategic Management J.* 25(8–9):909–928.
- Forderer J, Burtch G (2025) Estimating career benefits from online community leadership: Evidence from Stack exchange moderators. *Management Sci.* 71(3):2443–2466.
- Franzoni C, Stephan PE (2023) Uncertainty and risk-taking in science: Meaning, measurement and management in peer review of research proposals. *Res. Policy* 52(3):104706.
- Fuchs C, Sting FJ, Schlickel M, Alexy O (2019) The ideator's bias: How identity-induced self-efficacy drives overestimation in employee-driven process innovation. *Acad. Management J.* 62(5):1498–1522.
- Gallus J (2017) Fostering public good contributions with symbolic awards: A large-scale natural field experiment at Wikipedia. *Management Sci.* 63(12):3999–4015.
- Ghosh S, Wu A (2023) Iterative coordination and innovation: Prioritizing value over novelty. *Organ. Sci.* 34(6):2182–2206.
- Girotra K, Terwiesch C, Ulrich KT (2010) Idea generation and the quality of the best idea. *Management Sci.* 56(4):591–605.
- Goodman-Bacon A (2021) Difference-in-differences with variation in treatment timing. *J. Econom.* 225(2):254–277.
- Guilford JP (1967) *The Nature of Human Intelligence* (McGraw-Hill, New York).
- Guzman J, Li A (2023) Measuring founding strategy. *Management Sci.* 69(1):101–118.
- Haas MR, Criscuolo P, George G (2015) Which problems to solve? Online knowledge sharing and attention allocation in organizations. *Acad. Management J.* 58(3):680–711.
- Haunschild PR, Polidoro F, Chandler D (2015) Organizational oscillation between learning and forgetting: The dual role of serious errors. *Organ. Sci.* 26(6):1682–1701.
- Janis IL (1982) *Groupthink: Psychological Studies of Policy Decisions and Fiascos* (Houghton Mifflin, Boston).
- Jeppesen LB, Frederiksen L (2006) Why do users contribute to firm-hosted user communities? The case of computer-controlled music instruments. *Organ. Sci.* 17(1):45–63.
- Jeppesen LB, Lakhani KR (2010) Marginality and problem-solving effectiveness in broadcast search. *Organ. Sci.* 21(5):1016–1033.
- Jones BF (2010) Age and great invention. *Rev. Econom. Statist.* 92(1):1–14.
- Kang H, Eklund JC (2025) Fast friends: The impact of short-term visits on firms' invention outcomes. *Organ. Sci.* 36(3):1154–1176.
- Kaplan S, Vakili K (2015) The double-edged sword of recombination in breakthrough innovation. *Strategic Management J.* 36(10):1435–1457.
- Karim S, Kaul A (2015) Structural recombination and innovation: Unlocking intraorganizational knowledge synergy through structural change. *Organ. Sci.* 26(2):439–455.
- Kim JD (2022) Startup acquisitions, relocation, and employee entrepreneurship. *Strategic Management J.* 43(11):2189–2216.
- King G, Zeng L (2001a) Explaining rare events in international relations. *Internat. Organ.* 55(3):693–715.
- King G, Zeng L (2001b) Logistic regression in rare events data. *Political Anal.* (Oxford) 9(2):137–163.
- Klapper H, Piezunka H, Dahlander L (2024) Peer evaluations: Evaluating and being evaluated. *Organ. Sci.* 35(4):1363–1387.
- Koenker R, Bassett G (1978) Regression quantiles. *Econometrica* 46(1):33–50.
- Kornish LJ, Hutchison-Krupat J (2017) Research on idea generation and selection: Implications for management of technology. *Production Oper. Management* 26(4):633–651.
- Kornish LJ, Ulrich KT (2011) Opportunity spaces in innovation: Empirical analysis of large samples of ideas. *Management Sci.* 57(1):107–128.
- Lahiri N (2010) Geographic distribution of R&D activity: How does it affect innovation quality? *Acad. Management J.* 53(5):1194–1209.
- Lane JN, Cai T, Menietti M, Weber G, Guinan E (2024) Greenlighting innovative projects: How evaluation format shapes the

- perceived feasibility of novel ideas. Harvard Business School Technology & Operations Mgt. Unit Working Paper 24-064. Harvard Business School, Boston.
- Lane JN, Ganguli I, Gaule P, Guinan E, Lakhani KR (2021) Engineering serendipity: When does knowledge sharing lead to knowledge production? *Strategic Management J.* 42(6):1215–1244.
- Lane JN, Teplitskiy M, Gray G, Ranu H, Menietti M, Guinan EC, Lakhani KR (2022) Conservatism gets funded? A field experiment on the role of negative information in novel project evaluation. *Management Sci.* 68(6):4478–4495.
- Levin SG, Stephan PE (1991) Research productivity over the life cycle: Evidence for academic scientists. *Amer. Econom. Rev.* 81(1):114–132.
- Levinthal DA, March JG (1993) The myopia of learning. *Strategic Management J.* 14(S2):95–112.
- Li D (2017) Expertise versus bias in evaluation: Evidence from the NIH. *Amer. Econom. J. Appl. Econom.* 9(2):60–92.
- March JG (1991) Exploration and exploitation in organizational learning. *Organ. Sci.* 2(1):71–87.
- March JG, Simon HA (1958) *Organizations* (Wiley, New York).
- Mas A, Moretti E (2009) Peers at work. *Amer. Econom. Rev.* 99(1):112–145.
- Mihm J, Schlapp J (2019) Sourcing innovation: On feedback in contests. *Management Sci.* 65(2):559–576.
- Miller DL (2023) An introductory guide to event study models. *J. Econom. Perspective* 37(2):203–230.
- Mueller J, Melwani S, Goncalo JA (2012) The bias against creativity: Why people desire but reject creative ideas. *Psych. Sci.* 23(1):13–17.
- Nelson RR, Winter SG (1982) *An Evolutionary Theory of Economic Change* (Harvard University Press, Cambridge, MA).
- Nijstad BA, Stroebe W, Lodewijkx HF (2002) Cognitive stimulation and interference in groups: Exposure effects in an idea generation task. *J. Experiment. Soc. Psych.* 38(6):535–544.
- Ocasio W (1997) Towards an attention-based view of the firm. *Strategic Management J.* 18(S1):187–206.
- Ocasio W (2011) Attention to attention. *Organ. Sci.* 22(5):1286–1296.
- Osborn AF (1953) *Applied Imagination: Principles and Procedures of Creative Problem Solving* (Charles Scribners Sons, New York).
- Perry-Smith JE, Mannucci PV (2017) From creativity to innovation: The social network drivers of the four phases of the idea journey. *Acad. Management Rev.* 42(1):53–79.
- Piezunka H, Dahlander L (2019) Idea rejected, tie formed: Organizations' feedback on crowdsourced ideas. *Acad. Management J.* 62(2):503–530.
- Rahmani M, Roels G, Karmarkar US (2018) Team leadership and performance: Combining the roles of direction and contribution. *Management Sci.* 64(11):5234–5249.
- Ramdas K, Saleh K, Stern S, Liu H (2018) Variety and experience: Learning and forgetting in the use of surgical devices. *Management Sci.* 64(6):2590–2608.
- Repenning NP (2001) Understanding fire fighting in new product development. *J. Production Innovation Management* 18(5):285–300.
- Riedl C, Grad T, Lettl C (2024) Competition and collaboration in crowdsourcing communities: What happens when peers evaluate each other? *Organ. Sci.* 35(6):1957–1976.
- Rietzschel EF, Nijstad BA, Stroebe W (2007) Relative accessibility of domain knowledge and creativity: The effects of knowledge activation on the quantity and originality of generated ideas. *J. Experiment. Soc. Psych.* 43(6):933–946.
- Sætre AS, Van de Ven A (2021) Generating theory by abduction. *Acad. Management Rev.* 46(4):684–701.
- Sandler DH, Sandler R (2019) Multiple event studies in public finance and labor economics: A simulation study with applications. *J. Econom. Soc. Measures* 39(1):31–57.
- Schmidheiny K, Siegloch S (2023) On event studies and distributed-lags in two-way fixed effects models: Identification, equivalence, and generalization. *J. Appl. Econometrics* 38(5):695–713.
- Schweisfurth TG, Greul A (2024) Unexpected interruptions, idle time, and creativity: Evidence from a natural experiment. *Organ. Sci.* 35(1):116–137.
- Schweisfurth TG, Schöttl CP, Raasch C, Zaggl MA (2023) Distributed decision-making in the shadow of hierarchy: How hierarchical similarity biases idea evaluation. *Strategic Management J.* 44(9):2255–2282.
- Sgourev SV (2013) How Paris gave rise to cubism (and Picasso): Ambiguity and fragmentation in radical innovation. *Organ. Sci.* 24(6):1601–1617.
- Singh J (2005) Collaborative networks as determinants of knowledge diffusion patterns. *Management Sci.* 51(5):756–770.
- Singh J (2008) Distributed R&D, cross-regional knowledge integration and quality of innovative output. *Res. Policy* 37(1):77–96.
- Singh J, Agrawal A (2011) Recruiting for ideas: How firms exploit the prior inventions of new hires. *Management Sci.* 57(1):129–150.
- Singh J, Fleming L (2010) Lone inventors as sources of breakthroughs: Myth or reality? *Management Sci.* 56(1):41–56.
- Singh J, Marx M (2013) Geographic constraints on knowledge spillovers: Political borders vs. spatial proximity. *Management Sci.* 59(9):2056–2078.
- Smirnova I, Reitzig M, Sorenson O (2022) Building status in an online community. *Organ. Sci.* 33(6):2519–2540.
- Smith SM, Ward TB, Schumacher JS (1993) Constraining effects of examples in a creative generation task. *Memory Cognition* 21(6):837–845.
- Sorenson O, Stuart TE (2001) Syndication networks and the spatial distribution of venture capital investments. *Amer. J. Sociol.* 106(6):1546–1588.
- Stadler C, Helfat CE, Verona G (2022) Transferring knowledge by transferring individuals: Innovative technology use and organizational performance in multiunit firms. *Organ. Sci.* 33(1):253–274.
- Sullivan BN (2010) Competition and beyond: Problems and attention allocation in the organizational rulemaking process. *Organ. Sci.* 21(2):432–450.
- Sun L, Abraham S (2021) Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *J. Econom.* 225(2):175–199.
- Sutton RI, Hargadon A (1996) Brainstorming groups in context: Effectiveness in a product design firm. *Admin. Sci. Quart.* 41(4):685–718.
- Terwiesch C, Ulrich KT (2009) *Innovation Tournaments: Creating, Selecting, and Developing Exceptional Opportunities* (Harvard Business Review Press, Boston).
- Terwiesch C, Xu Y (2008) Innovation contests, open innovation, and multiagent problem solving. *Management Sci.* 54(9):1529–1543.
- von Hippel E, von Krogh G (2016) CROSSROADS—Identifying viable “need–solution pairs”: Problem solving without problem formulation. *Organ. Sci.* 27(1):207–221.
- Weitzman ML (1998) Recombinant growth. *Quart. J. Econom.* 113(2):331–360.
- Wooten JO, Ulrich KT (2017) Idea generation and the role of feedback: Evidence from field experiments with innovation tournaments. *Production Oper. Management* 26(1):80–99.
- Wuchty S, Jones BF, Uzzi B (2007) The increasing dominance of teams in production of knowledge. *Science* (1979) 316(5827):1036–1039.