Review of Economics and Statistics Just Accepted MS.
doi:10.1162/REST_a_00631
© by the President and Fellows of Harvard College and the Massachusetts Institute of Technology

A Field Experiment in Motivating Employee Ideas

Michael Gibbs, *University of Chicago & IZA*

Susanne Neckermann, *University of Chicago & ZEW*

Christoph Siemroth, *University of Essex*

Abstract: We study a field experiment at a large technology company. Employees were encouraged to submit ideas on process and product improvements. The company randomly assigned 19 teams into treatment and control groups. Treatment team employees received rewards if their ideas were approved. Nothing changed for control team employees. Our main finding is that rewards substantially increased the quality of ideas. Rewards increased participation in the suggestion system but decreased ideas per participating employee, with zero net effect on the quantity of ideas. Broader participation persisted after the reward was discontinued, suggesting habituation. We find no evidence for motivational crowding out.

Keywords: creativity, incentives, intrinsic motivation, innovation, idea suggestion systems.

JEL classifications: C93, J24, M52, O31, O32

Acknowledgements: We are grateful to executives and employees of HCL Technologies who provided the data and spent many hours answering questions. We thank participants at numerous seminars, as well as Jed DeVaro, Christoph Engel, Kathryn Ierulli, Sacha Kapoor, Steve Levitt, Mario Macis, Henry Sauermann, Mattie Toma, Jeannine van Reeken, and the editor and referees for comments. Funding was provided by the University of Chicago Booth School of Business (Gibbs), the SEEK program of Baden-Württemberg (Neckermann), and the DFG via SFB 884 (Siemroth).

Review of Economics and Statistics Just Accepted MS.
doi:10.1162/REST_a_00631
© by the President and Fellows of Harvard College and the Massachusetts Institute of Technology

1. Introduction

In today's economy, innovation is an important determinant of competitiveness and economic growth. "Ideation" – generation of novel business-enhancing ideas – is increasingly important. One source of innovation traditionally studied by academics is R&D labs, where activities are explicitly directed towards the development of knowledge. A less frequently studied source of ideas is employees in non-R&D organizational units, who may have ideas about ways to improve processes, customer service, or product design. Outside R&D, job definitions typically do not specifically include innovation, and firms may fail to tap this potential source of innovation.

One method to access employee ideas is a formal suggestion system, which encourages employees to submit ideas for process and product improvements. Such a system can make ideas transparent to management, who can then decide about implementation, and share the knowledge within the company. Ohly et al. (2013) find that approximately 1 in 3 workplaces in the US and United Kingdom use a suggestion system. To date, little is known about how such systems work and whether rewards tied to submitting ideas increase participation and ideation.

The aim of this paper is twofold. First, we provide a detailed picture of the ideation process in a company, complementing existing literature on suggestion systems. Second, we analyze outcomes of a randomized field experiment designed to test rewards for ideation, contributing to the debate about whether rewards improve or undermine creativity. We also study the relationship between employee characteristics and the quantity and quality of their ideas.

The setting is HCL Technologies, a large, multinational business process outsourcing company based in India. It has a formal process by which employees submit new ideas, which are tracked through review and implementation. In an attempt to further improve the suggestion system, the company ran an experiment on rewards for acceptance of new ideas. The company chose 19 client

Review of Economics and Statistics Just Accepted MS.
doi:10.1162/REST_a_00631
© by the President and Fellows of Harvard College and the Massachusetts Institute of Technology

teams, comprised of more than 11,000 employees, and randomly assigned them to treatment and control groups. For 13 months, employees in treatment teams received points for ideas approved by clients, and additional points for favorable client feedback. Employees accumulated points to redeem for consumer goods. Apart from the reward, policies were identical between treatment and control groups. Our data includes the treatment period as well as 13 months before the experiment. This allows a difference-in-differences approach to estimate the effects of rewards on ideation, accounting for possible pre-treatment differences between control and treatment groups.

The experimental reward had mixed effects on ideation. Our main finding is that the quality of ideas (percent of ideas accepted for implementation, and percent pitched to the client) was substantially increased by the reward. Rewards also increased the fraction of employees participating in the suggestion system. This is in contrast to the notion that creativity is primarily a function of personality traits rather than effort-based (George & Zhou 2001). The number of ideas submitted per potential contributor fell, suggesting that individuals focused on more promising ideas, which would explain the rise in idea quality. Combined, there was no net effect on the total quantity of ideas. We find no evidence that the reward crowded out intrinsic motivation (Deci 1971; Frey & Oberholzer-Gee 1997; Frey & Jegen 2001). Finally, we analyze a 13-month window after the reward ended, and find some evidence of habituation effects.

There are extensive literatures on innovation and creativity in economics, psychology, social psychology, and management. Some research focuses on how personal characteristics affect creativity (Guilford 1950; Amabile 1983; Sauermann & Cohen 2010). Some analyzes organizational support, culture, or leadership (Amabile 1996; Robinson & Stern 1997; Tierney, Farmer & Graen 1999). Many study the effect of rewards on creativity, though findings are mixed (Amabile 1982, 1996; Eisenberger & Armeli 1997; Deci et al. 1999; Joussemet & Koestner 1999;

Review of Economics and Statistics Just Accepted MS.
doi:10.1162/REST_a_00631
© by the President and Fellows of Harvard College and the Massachusetts Institute of Technology

Eisenberger & Rhoades 2001; Eckartz et al. 2012). Some analyze how reward structure (e.g., magnitude; short or long term) affects creativity (Ariely et al. 2009; Azoulay et al. 2011; Ederer & Manso 2013) or how rewards affect different types of creativity (Charness & Grieco 2013). These studies typically do not consider the submission process. This matters as ideas need to be formulated and communicated to be effective. A field experiment is important for studying ideation. Both the setting (experienced subjects working on a familiar task, where payments may be expected) and tasks (creativity, planning, and execution) differ from typical lab experiments. This study complements existing studies on creativity in these respects.

Our study is closely related to a small number of papers on employee suggestion programs. Most use survey data to elicit employee motivations and organizational antecedents for submitting ideas. Leach et al. (2006) employ data from 182 UK organizations and show that idea quantity correlates with scheme characteristics such as centralization, publicity, and use of rewards. van Dijk and van den Ende (2002) argue that reward criteria affect number of submissions. Ohly et al. (2013) provide a conceptual discussion and literature overview. Key indicators of scheme success are participation (percent of employees), adoption (percent of ideas implemented), and savings realized – indicators we also use. Toubia (2006) models the effect of incentives on ideation, focusing on free-riding in groups. He finds that incentives may improve idea generation, in a lab experiment. We focus on similar questions, using data from an actual employee suggestion system.

To our knowledge, this is the first field study to offer causal evidence on the effect of rewards on ideation. Furthermore, we are among the first to describe and analyze a system that is becoming common as firms attempt to formalize innovation processes. In addition, we document a company's attempt to learn and improve organizational design via field experimentation. Finally, we provide field evidence on whether or not rewards crowd out intrinsic motivation and creativity.

Review of Economics and Statistics Just Accepted MS.
doi:10.1162/REST_a_00631
© by the President and Fellows of Harvard College and the Massachusetts Institute of Technology

2. INSTITUTIONAL SETTING

Company background

HCL provides a variety of services, including outsourcing, R&D, and software or hardware solutions. Most clients are large global companies. The business process outsourcing industry has become increasingly competitive and commoditized. HCL seeks to differentiate by offering greater innovation, in the hope of increasing client retention, growth, and profit margins. For that reason, all employees are encouraged to suggest ideas for process or product improvements that are reviewed for potential implementation. If an idea might have a direct effect on a client, it is shared with and approved by the client before implementation. The company encourages employees to focus on client-facing ideas, with the strategic intent of becoming the client's partner in innovation. As part of this effort, HCL set up a formal process (the "Idea Portal") for idea submission and evaluation, described below. The company also ran a field experiment to test the effect of rewards on idea generation. Rewards were given for ideas shared with and approved by clients. It is this experimental reward scheme, and the Idea Portal, that we study in this paper.

Most employees (all on teams studied here) reside in and are citizens of India. Employees are assigned to a specific client team, and all of their work is for that client. An employee is usually assigned to the same team for many years, so almost all workplace interactions take place within teams. This ensures confidentiality of sensitive client data. For our purposes this segmentation is useful because client teams work independently of each other.

The "Idea Portal"

The firm set up an intranet-based *Idea Portal* to formalize the suggestion process; it is used to collect, evaluate, and track ideas. All employees, regardless of position or level, are encouraged to devise and submit new ideas. The process is depicted in Figure 1:

Review of Economics and Statistics Just Accepted MS.
doi:10.1162/REST_a_00631
© by the President and Fellows of Harvard College and the Massachusetts Institute of Technology

[FIGURE 1 ABOUT HERE]

1. *Ideation*: one or more employees devise a new idea they deem worthy of suggesting. Employees may form ideation groups, and ideas can be submitted by 1-3 employees. An employee can be part of multiple ideation groups. Valid ideas might benefit a client directly, or indirectly by improving HCL's internal processes and services. Examples include new products or services, process improvements, or new software tools. The company's explicit goal is to increase value creation and innovation for clients.
2. *Submission:* employees submit the idea on the Portal, including a brief description, how it might be implemented, and estimates of implementation costs and projected revenue.
3. *Supervisor Input / Idea Refinement:* the supervisor is notified when an idea is submitted, and expected to review it within 3 days. The supervisor helps the employee clarify the idea, estimate costs and benefits, etc. as needed. This stage is refinement, not review.
4. *Executive Review Panel:* first formal review of the idea. Each business unit has a review panel consisting of senior executives with good understanding of the client's business. The panel meets regularly to discuss the client, innovation, and ideas from the Portal. The executives are able to place the idea in a broader business context, and judge whether it should be shared with the client (and possibly other parts of the organization). They typically do not know employees who authored the idea, or the supervisor. Once a supervisor vets the idea, it is essentially randomly assigned to the panel member with the fewest pending ideas. That executive oversees the accept / reject decision and often consults with fellow panel members. This stage results in one of four outcomes: the idea is sent back for further refinement; rejected; accepted (internal ideas); or approved internally and submitted to the client for approval (customer-facing ideas).

Review of Economics and Statistics Just Accepted MS.
doi:10.1162/REST_a_00631
© by the President and Fellows of Harvard College and the Massachusetts Institute of Technology

5. *Client Review:* customer-facing ideas are evaluated by the client for a final decision. Estimated costs and revenue may be updated based on client input.
6. *Implementation:* accepted ideas are implemented, with results tracked on the Portal. Implementation may occur in stages; e.g., prototype development, test trials, etc. Implementation could be halted at any step, in which case the idea is listed as rejected in the data. Clients rate relevant ideas on a scale of 1-4 after implementation.

Portal ideas are accessible by all employees, to spur ideation and spread ideas across the organization. However, to ensure client confidentiality, detailed idea information is visible only to those on the client team and higher-level executives. The Portal was a mature, robust system prior to the reward experiment. It had already processed several thousand ideas.

Experimental reward program

Employees were motivated to suggest ideas prior to the experiment, due to intrinsic motivation, career concerns, and recognition from clients, teammates, supervisors, reviewing managers, and the rest of the organization. Until the experiment, however, there was no formal reward for ideas.

The company decided to run an experiment to see if rewards could further improve ideation. It wanted to reinforce its goal for the Portal, encouraging employees to become a source of innovation for clients. HCL hoped the reward would stimulate bottom-up cultural change, in which employees would ask for more and better help from their supervisor. They felt it would be more effective at generating valuable ideas than top-down management pressure. HCL also wanted to encourage employees to submit ideas on the Portal rather than implement them on their own. This would increase transparency, allowing management to observe improvements and communicate them to clients as appropriate. Moreover, the company believed that ideas entered on the Portal would spur knowledge sharing and spillovers across the company.

Review of Economics and Statistics Just Accepted MS.
doi:10.1162/REST_a_00631
© by the President and Fellows of Harvard College and the Massachusetts Institute of Technology

The company has a reward program in which employees earn points for accomplishments; e.g., work on a specific initiative. Employees accumulate points and exchange them for consumer goods (e.g., a smartphone) or vouchers for an online store.[1] Such programs are not unusual in Asian companies. It applies to all employees and had been in place several years prior to the experiment. For the experiment, the company offered new reward points for ideation on the Portal.

The experiment was run with 19 teams. Senior executives participated in selection, ensuring that well-established teams (more than $2 million revenue, at least 50 employees, and average client satisfaction of at least 25 on a scale of –100 to 100), with active use of the Portal (at least 5 ideas submitted per month, with at least 1 shared with the client), were used for the experiment. Thus, participating teams were pre-selected from the entire workforce, and our analysis is conditional on this selection. However, selection for teams that use the Portal regularly is advantageous for our purposes, since it makes study teams more homogeneous. If anything, it is likely to yield conservative estimates of the effect of rewards on ideation. Teams were randomly assigned to treatment or control (roughly 6,000 and 5,400 employees respectively).

Employees in control teams continued to receive no rewards for contributing ideas. Employees in treatment teams received an experimental reward, designed to motivate quality, not quantity. Submission of an idea alone did not merit a reward. Employees could not know ex ante if approved ideas would be implemented internally, or submitted to the client, so they had an incentive to consider the client when developing ideas. If the idea was accepted by the client, each member of the ideating team earned 2,000 points. 2,000 points was worth approximately 2.2% of monthly after-tax salary for lower level employees. After implementation, authors could earn additional

[1] Gift vouchers and goods from the category "Kids" were the most common redemptions.

Review of Economics and Statistics Just Accepted MS.
doi:10.1162/REST_a_00631
© by the President and Fellows of Harvard College and the Massachusetts Institute of Technology

points depending on client rating of the idea. If the idea received the highest rating of 4, the employee earned additional points worth approximately 40% of monthly after-tax salary of lower-level employees.[2] Therefore, rewards were significant, especially for ideas rated highly by clients.

Apart from the reward, there were no policy differences between treatment and control groups. Review processes, and supervisor and manager incentives, did not change in either group. Communications and training about innovation were corporate-wide and not affected by the experiment. Employees in treatment teams were not told the reward was an experiment. Control team employees, review panels, and clients did not know about the experiment.

A concern in such settings may be spillovers between control and treatment teams, or from HR to control teams. Such concerns are likely limited in this case. Employees generally stay in their team for many years.[3] Most or all workplace interactions are with team members and the client. Many teams are physically segregated. The experiment was not discussed broadly within the company, in part out of concern that other teams would request the reward.

The company designed and conducted the experiment on its own. The executive in charge of the Portal and experiment was a former student of one author, so we had high-level access to clarify

[2] Management believes rewards were substantial and motivated employees. One employee told us he earned an LCD TV because of his ideas, and argued this was a huge benefit given its price.

[3] Measured between the two points in time for which we have rosters, 0.9% of employees switch teams; 0.1% between control and treatment teams. According to management, no employee switched between treatment and control teams during the treatment, so the 0.1% changed prior to the experiment. A related concern might be sorting; e.g., creative employees shifting to teams with the reward. We ran regressions with employee fixed effects, and results are very similar.

Review of Economics and Statistics Just Accepted MS.
doi:10.1162/REST_a_00631
© by the President and Fellows of Harvard College and the Massachusetts Institute of Technology

questions. We visited headquarters twice to learn about the company's organization, strategy, culture, and innovation processes. We interviewed employees in some control and treatment teams about their work and innovation, without disclosing the experiment or purpose of our study. We followed up with numerous telephone discussions with management.

3. Predictions

Employees entered ideas into the Portal even without the experimental reward. We study incremental effects on innovation from introducing explicit rewards. There are three relevant dimensions to ideation, given our data: participation (attempt to develop and submit ideas); and quality and quantity of ideas conditional on participation.

Consider first *participation*. the paper, "participation" means an employee is willing and able to submit ideas. However, he or she may not submit an idea, if they do not develop a suitable idea, or do not have time to submit the idea in the Portal. Explanation of how participation is estimated is in the Methods section. Willingness to participate depends on rewards and costs from doing so. In HCL employees are encouraged to enter ideas, but many do not participate. Presumably this is because the cost of thinking about ideas or using the system, on top of normal work, outweighs the perceived benefit. The experiment provides additional benefit, which should increase the likelihood that an employee participates in ideation.

The reward might also affect idea *quality*. We use two measures, also employed by the firm in analyzing these programs: the probability an idea is *Implemented*, and the probability it is *Shared* with the client. These are natural measures, as only worthwhile ideas are implemented and HCL's goal was to generate more ideas to share with the client. They are imperfect due to the subjective nature of idea evaluation. This is a common issue in the innovation literature, which uses measures such as number of patents filed or granted.

Review of Economics and Statistics Just Accepted MS.
doi:10.1162/REST_a_00631
© by the President and Fellows of Harvard College and the Massachusetts Institute of Technology

Psychologists argue that there may be a cognitive tradeoff between quality and quantity in ideation.[4] Employees might face a choice between "exploration" – pursuing new directions – and "exploitation" – pursuing familiar directions for incremental innovation (Robbins 1952; Gittins 1979). The former is less likely to generate new ideas, but such ideas may have higher expected value or quality. The latter may generate more ideas, but of lower quality. Idea quality likely depends on time and thought spent elaborating and describing the idea, instead of on other ideas. Moreover, since the experiment rewarded quality but not quantity, we expect average quality of ideas to increase, particularly among *Prior Ideators* (employees who submitted ideas pre-treatment). A possible countervailing effect might come from employees who were not Prior Ideators. The marginal employee not ideating pre-treatment may be less creative than those who did participate. If she decides to participate due to the reward, it might lower average idea quality.

The effect of the reward on the *quantity* of ideas is similarly unclear. If participation increases, that will tend to increase quantity. However, if there is a quantity-quality tradeoff, the average quantity of ideas submitted by a given employee (conditional on participation) might fall. Thus, theory does not provide an unambiguous prediction for the direction of the treatment effect on average idea quality and quantity. One goal of the paper is to study this empirically.

It is often argued that pay for performance "crowds out" intrinsic motivation. There are several rationales. One is overjustification (Deci et al. 1999). In this view, extrinsic rewards are more salient than intrinsic motivation, and motivation shifts when extrinsic rewards are instituted. An alternative is signaling (Benabou & Tirole 2006). The idea is that employees who ideate without

---

[4] A cognitive tradeoff between idea quantity and quality suggests multitask incentives (Holmstrom & Milgrom 1991). The experiment rewarded high quality ideas rather than merely quantity.

Review of Economics and Statistics Just Accepted MS.
doi:10.1162/REST_a_00631
© by the President and Fellows of Harvard College and the Massachusetts Institute of Technology

rewards signal to themselves and / or others that they are intrinsically motivated. The theory assumes they do not have perfect knowledge about their motivations, and make inferences about their own character from their actions. Rewards make ideation attractive to extrinsically motivated types, thereby decreasing its signaling value for intrinsic motivation.

The crowding out view suggests that incentive effects of rewards may be mitigated by a reduction in intrinsic motivation, with unclear net effect. However, if there is crowding out, the net effect should vary with the extent to which employees are intrinsically motivated. The data allow us to crudely group individuals by degree of intrinsic motivation and test this. Consider employees in the pre-treatment period, without rewards. On average, intrinsic motivation among employees who ideate should be higher than among those who do not. Therefore, crowding out effects in the treatment period should be more pronounced for the first group.

Crowding out theories do not distinguish between participation in ideation, and idea quantity and quality, so we study all three. Participation is the concept that seems most closely aligned with the literature. If there is motivational crowding out, the experimental reward should have a smaller (possibly negative) effect on participation for employees who previously participated, compared to those who did not. It might have similar effects on idea quantity or quality.

4. Data & Methods

Data were collected from the Portal. Each record contains IDs of idea authors and client team, dates of submission and latest update, verbal description, and estimates of costs of implementation and projected revenue. Estimates are provided initially by employees in consultation with their supervisor, but may be updated upon implementation with client input. We estimate idea profit as revenue minus cost. Current idea status (e.g., under review, rejected, under implementation) is recorded. Finally, the Portal records client ratings on a four-point scale. However, for most ideas

Review of Economics and Statistics Just Accepted MS.
doi:10.1162/REST_a_00631
© by the President and Fellows of Harvard College and the Massachusetts Institute of Technology

this is missing, as not all suggestions were shared, and few clients used this option. We have records for all ideas suggested during 13 months prior to treatment (*Period 1*: May 2009-May 2010), the 13 months of treatment (*Period 2*: June 2010-June 2011), and 13 months after treatment (*Period 3*: July 2011-July 2012). Data on idea acceptance / rejection was updated in January 2013.

We added information about employees who did not submit ideas, in order to study the rate of participation in ideation. Rosters were provided for April 2010 and July 2011. We do not know exactly when employees entered or exited, so cannot construct exact rosters each month. However, these provide good approximations for pre-experimental and experimental periods.[5] Finally, we collected employee gender, age, company tenure, and 9 salary groups corresponding to hierarchical levels (Level 0 is entry; Level 8 executives). The online appendix contains variable descriptions.

Summary statistics

Table 1 presents explanatory variable descriptive statistics by treatment group each period. The groups are similar in age, tenure, gender, salary, and share of prior ideators. The only statistically significant difference is that employees in treatment teams were slightly younger pre-treatment. We therefore control for individual factors in all regressions; results are not sensitive to inclusion of these variables. As is common in technology firms, employees are relatively young and tenure relatively short. Almost 80% are male. The number of employees increased substantially from

---

[5] For quantity regressions we added inactive employees from the April 2010 roster (end of pre-treatment period) to both periods, and from the July 2011 roster (end of treatment period) to the treatment period. Allocating only April 2010 employees to the pre-treatment period resulted in very similar estimates. This confirms difference-in-differences addresses the concern that we do not have exact rosters, because incomplete information affects both groups the same way.

Review of Economics and Statistics Just Accepted MS.
doi:10.1162/REST_a_00631
© by the President and Fellows of Harvard College and the Massachusetts Institute of Technology

periods 1 to 2, reflecting company growth, as well as that Period 1 data is less complete. We have only Roster 1 in the first period, whereas we can use both rosters in the second period.

[TABLE 1 ABOUT HERE]

Table 2 displays descriptive statistics on outcome variables by period and group. Despite randomization, control and treatment teams are somewhat different. Employees in control teams were more likely to suggest at least one idea, in both periods. The share of ideating employees declined in both groups in the treatment period, with a stronger drop in the control group. These differences highlight the need to control for pre-treatment differences and time trends in our statistical analyses.[6]

[TABLE 2 ABOUT HERE]

The last columns present descriptive statistics on three measures of idea quality. The first two indicate decisions made about each idea: accepted for implementation (*Imp*), and shared with the client (*Shared*). Both are the result of a review process and only reported for ideas with finished review (*Fin*). These measures reflect idea quality, since better ideas have a higher chance of acceptance, and it was the company's explicit goal to have more ideas to share with clients. The percent of ideas accepted for implementation increased in both the treatment and control groups

[6] Pre-treatment differences are not unexpected due to the small number of teams. As long as pre-treatment differences are uncorrelated with treatment response, they will not affect treatment effect estimates because we control for initial differences. We conducted robustness checks by including average team performance in Period 1 as an explanatory variable for Period 2 performance. Direction and magnitude of estimated treatment effects are similar to those below.

Review of Economics and Statistics Just Accepted MS.
doi:10.1162/REST_a_00631
© by the President and Fellows of Harvard College and the Massachusetts Institute of Technology

over time, with a stronger increase in the treatment group. The probability of sharing an idea with the client increased in the treatment group, but decreased in the control group.

We also report estimated profits (net value), the difference between estimated revenue and cost of an idea. This measure varied substantially across ideas. Some ideas had very small projected financial impact, while one idea was estimated to improve revenue by $22 million.

Unfortunately, the last potential measure, client idea rating on a 4-point scale, was given for only 17% of ideas. Moreover, there appear to be serious selection issues. 85% of ratings are 3 or 4. Only 1 of 306 ratings received the lowest score of 1. This suggests that clients only reported ratings when they were happy with the idea, and used other channels to report dissatisfaction. We therefore do not use client ratings in our analysis.

Methods

Measuring and explaining the quantity of ideas

One of our main dependent variables is *Number of Ideas*, the total number of ideas submitted by an employee in either the pre-treatment or treatment period. The unit of observation for this analysis is therefore an employee-period. A large fraction of employees (91.4%) did not submit ideas, resulting in a large number of zeros for this variable. The prevalence of zeros disqualifies simple count data models such as Poisson or negative binomial (NB) regression.[7]

---

[7] Poisson assumes equality of conditional mean μ_i and conditional variance, which results in biased estimates for overdispersed data (conditional variance larger than mean). Negative binomial regression (type II) generalizes the Poisson model, assuming $\mathrm{Var}(y_i) = \mu_i(1 + \alpha\mu_i)$, where α is an additional parameter. The NB model reduces to the Poisson for $\alpha = 0$. Estimating the NB model

Review of Economics and Statistics Just Accepted MS.
doi:10.1162/REST_a_00631
© by the President and Fellows of Harvard College and the Massachusetts Institute of Technology

We therefore analyze the data using a zero-inflated negative binomial model (ZINB), a generalization of NB used in particular for overdispersed data. ZINB is a mixture model, explaining the dependent count data variable jointly with a negative binomial process and a logit process. The intuition is that two are simultaneously at work in the data generating process. (1) The logit models participation: the employee's basic willingness or ability to submit an idea. (2) If (1) is fulfilled, the NB models the generation and submission of ideas, where the number of ideas is drawn from the negative binomial distribution. The logit accounts for "excess zeros" not explained by the negative binomial distribution. A zero observed in the data could be caused either by an employee not participating in ideation (condition (1) not fulfilled), or wanting to ideate but not having an idea that period (condition (1) fulfilled, but the draw in (2) is zero).

Formally, let f_1 = Pr(*Participation*) be parameterized as a logistic function, and $f_2(y)$ be the probability mass of count y in the negative binomial distribution. The mass of count y in the zero inflated negative binomial model is:

$$g(y) = \begin{cases} (1 - f_1) + f_1 f_2(0), & y = 0 \\ f_1 f_2(y), & y > 0 \end{cases}$$

The NB model is a special case where $f_1 = 1$.[8] The ZINB model is well suited for our data, as the two processes allow us to effectively address overdispersion, and have meaningful

---

with our data, we find significant overdispersion, indicating Poisson regression is unsuitable. We also report quasi-maximum likelihood (QML) Poisson regression estimates in Table 3, which do not impose a specific functional form for variance and can therefore account for overdispersion.

[8] Because we cluster standard errors, the objective functions of the models are pseudo-likelihoods, so common tests to discriminate between models (likelihood ratio test for NB v. ZINB; Vuong test

Review of Economics and Statistics Just Accepted MS.
doi:10.1162/REST_a_00631
© by the President and Fellows of Harvard College and the Massachusetts Institute of Technology

interpretations in our context. They allow us to model the submission decision hierarchically. The logit provides information on whether or not an employee participated in ideation. If the answer is no, the observed number of ideas is zero. These zeros indicate that an employee is not participating: they might be unmotivated to look for ideas, might shy away from using the Portal, or might work on a project with little or no room for ideation. If the answer is "yes," we may still observe zero ideas from this employee, if they did not come up with an idea, even though actively looking for one and willing to use the system. Positive counts are only observed when both processes are positive; i.e., the employee was motivated to ideate and also had an idea.

The latent class character of the zero inflated model allows us to disentangle these two causes of non-ideation, whereas standard count data models or OLS do not. The logit and negative binomial processes of the ZINB are specified as follows:

(1) Logit: $\Pr(Participation_{it}|\boldsymbol{X}_{it}) = \exp(\lambda_{it}^{L})/(\exp(\lambda_{it}^{L}) + 1)$.

(2) NB: Negative binomial model with conditional mean $\exp(\lambda_{it}^{NB})$.

Estimating treatment effects

In all regressions, we use linear or nonlinear difference-in-differences estimations to infer the causal effect of the treatment on ideation. This addresses several potential problems; e.g., pre-treatment performance differences, possible time trends (decrease over time because low-hanging fruit have been picked), and approximation of number of non-ideating employees per period. To estimate the treatment effect, we specify λ_{it}^{j} that appears in both equations (1) and (2) as follows:

---

for ZINB v. zero inflated Poisson) are not applicable. Similar tests for clustered data are not yet common. Taking pseudo-likelihoods as true likelihoods, Akaike's information criterion, likelihood ratio, and Vuong tests all favor ZINB significantly over Poisson, NB, and zero inflated Poisson.

Review of Economics and Statistics Just Accepted MS.
doi:10.1162/REST_a_00631
© by the President and Fellows of Harvard College and the Massachusetts Institute of Technology

$$\lambda_{it}^{j} = \beta_1^j Treatment_i \cdot Period\ 2_t + \beta_2^j\ Period\ 2_t + \boldsymbol{X}_{it}'\gamma^j + C_i^j \text{ for } j \in \{L, NB\},$$

where β_1 is the difference-in-differences treatment effect estimate. *Treatment* is the treatment group dummy, *Period 2* is the treatment period dummy, $\boldsymbol{X}_{it}$ is a vector of employee characteristics, and C_i is the client team fixed effect that accounts for unobservable time-invariant heterogeneity between teams. The time invariant dummy *Treatment* is not estimated separately, because it is absorbed by client team fixed effects (every team is either in the treatment group or not). Index *i* denotes the author and *t* denotes pre- or treatment period. To control for time trends or time-specific shocks, we use time fixed effects in the form of a treatment period dummy. Note that the coefficients of the logit (L) and NB process are not constrained to be equal, which allows us to determine whether treatment affects the two processes differently.

Employees within a client team regularly interact, so the assumption of independence between observations (employees) is likely violated. We therefore cluster standard errors in all regressions at the team level to allow for arbitrary correlation in error terms within a team.[9]

---

[9] The low number (19) of clusters is a concern because our cluster and heteroscedasticity robust standard errors are only valid asymptotically. However, we have reason to expect the standard errors are valid. Bertrand et al. (2004, Table VIII) show (for OLS) that clustered standard errors exhibit only small bias for 20 clusters. If there is no effect in artificially generated data, the null hypothesis rejection rate based on these standard errors is 5.8%, when it should be 5%. Block bootstrap, the primary alternative, rejects the null in 13% of cases (their Table V). We also estimated several linear models with the wild bootstrap procedure. Cameron et al. (2008) demonstrate in simulations that the procedure does not over-reject. The t statistics from the procedure are virtually identical to those with clustered standard errors in our linear models.

Review of Economics and Statistics Just Accepted MS.
doi:10.1162/REST_a_00631
© by the President and Fellows of Harvard College and the Massachusetts Institute of Technology

5. RESULTS

Who ideates?

Before we evaluate the experiment, we provide a brief overview of the correlates of ideation with employee characteristics. Even though this is not a causal analysis, the question is of intrinsic interest, as it provides insights into the ideation process and lays a foundation for understanding the reward effects analyzed below. The results are displayed in Table 3. Our measures of ideation are the probability of submitting at least one idea within the 13-month pre-treatment period and the number of ideas submitted. We report average marginal effects (AME).

[TABLE 3 ABOUT HERE]

We control separately for age and company tenure (marginal effects of both are based on linear and quadratic terms in all models). Age is negatively related to both outcome measures. On average, an additional year of age decreases the probability that an employee submits at least one idea by about 0.7 percent, and decreases the number of ideas per employee by about 0.017 (according to the more conservative ZINB estimate), or by 0.08 (according to the quasi-maximum likelihood Poisson estimate). These estimates suggest that older employees are less likely to contribute to the Idea Portal, and submit fewer ideas on average. An explanation may be that younger employees are more creative. Alternatively, the Portal might be more attractive to them, as they are more accustomed to using new technology. Finally, they may be more aware of recent technological developments, such as new software programming techniques.

The effect of tenure is positive and statistically significant for both measures. One might expect that employees with longer tenure already suggested ideas that came to mind and have fewer new ones, but our findings suggest the opposite. This may be because working at the company longer develops deeper understanding of the business and client, leading to more ideas.

Review of Economics and Statistics Just Accepted MS.
doi:10.1162/REST_a_00631
© by the President and Fellows of Harvard College and the Massachusetts Institute of Technology

Finally, consider the effects of hierarchical level. Senior management (salary groups 4 and above) is the reference category. Low-level groups 0 and 1 are pooled, because group 0 is too small for meaningful analysis on its own. The only statistically significant effect is that ideating low-level employees contributed about 0.5 fewer ideas in 13 months than active high-level employees, according to the more conservative ZINB estimate. The mechanism behind this might be similar to the one for tenure. Both correlate with firm-specific human capital and job match. Employees higher up in the hierarchy tend to have greater responsibility, control more resources, manage more subordinates, and have a higher-level understanding of the firm's and client's business and needs.

Treatment effects

Quantity of ideas

We now analyze effects of the experiment, starting with idea quantity. Table 4 displays OLS estimates and average marginal effects (AME) for the zero inflated negative binomial model (columns 1-2), modeling an employee's number of ideas per period.[10] In the ZINB model, quantity is modeled hierarchically with a logit capturing participation, and a negative binomial describing number of ideas, given participation. Column 2 displays marginal effects for both processes jointly. Columns 3-4 display marginal effects for logit and negative binomial processes separately. ZINB matches best with the idea generation and submission process, so we focus on those estimates.

[TABLE 4 ABOUT HERE]

Looking at partial effects first, the treatment effect is significantly positive in the logit process, suggesting rewards increased the share of employees motivated to participate – to think about ideas

[10] Table I in the online appendix displays estimates of the ZINB model rather than the average marginal effects displayed in Table 4, and nuisance parameter estimates not displayed in Table 4.

Review of Economics and Statistics Just Accepted MS.
doi:10.1162/REST_a_00631
© by the President and Fellows of Harvard College and the Massachusetts Institute of Technology

and submit them on the Portal. According to the average marginal effect, the rewards are estimated to increase the share of potential contributors by almost 18 percentage points, which is also economically significant. This accords with our prediction described above.

Interestingly, this positive effect is countered by a decrease in the number of ideas per author, as indicated by the negative binomial process. This result suggests that potential contributors submit about 0.26 fewer ideas during the 13-month treatment as a result of the reward program. The overall effect of rewards on the quantity of ideas is economically and statistically zero (see OLS and ZINB estimates in columns 1-2); the positive and negative partial effects cancel out. Consequently, ideation is spread over more employees, who concentrate on fewer ideas. This is consistent with a potential tradeoff between the quantity and quality of ideas.

The zero overall effect might be explained by crowding out, if intrinsically-motivated employees were negatively affected by the reward and consequently submitted fewer ideas. We use pre-treatment data to categorize intrinsic motivation. Intrinsic motivation should, on average, be stronger for those who ideated pre-treatment – without rewards – than among those who did not. We therefore proxy intrinsic motivation with the dummy *Prior Ideator*, which equals 1 if an employee submitted at least one idea in the pre-treatment period, and is still in the sample in the treatment period. The second condition ensures these employees have the opportunity to contribute again. About 8% of employees fall into this category in period 1. Columns 5-6 present the results of a model that allows the treatment effect to vary by whether or not the employee is a prior ideator. The models we estimate are identical to the main analyses presented above, except that:

$$\lambda_{it} = \beta_1 Treatment_i \cdot Period\ 2_t + \beta_2 Treatment_i \cdot Prior\ Ideator_i \cdot Period\ 2_t$$
$$+\beta_3\ Period\ 2_t\ + \beta_4 Prior\ Ideator_i \cdot Period\ 2_t + \beta_5 Treatment_i \cdot Prior\ Ideator_i$$
$$+\beta_6 Prior\ Ideator_i + \boldsymbol{X}'_{it}\gamma + C_i,$$

Review of Economics and Statistics Just Accepted MS.
doi:10.1162/REST_a_00631
© by the President and Fellows of Harvard College and the Massachusetts Institute of Technology

for both the logit and NB process. The average treatment effect for prior ideators evaluated at sample values of all other covariates is:

$$ATE = \frac{1}{2N}\sum_{i=1}^{N}\sum_{t=1}^{2}[F(\beta_1 + \beta_2 + \beta_3\ Period\ 2_t + \beta_4 Prior\ Ideator_i \cdot Period\ 2_t + \beta_5 Treatment_i \cdot Prior\ Ideator_i + \beta_6 Prior\ Ideator_i + \boldsymbol{X}_i'\gamma) - F(\beta_3\ Period\ 2_t + \beta_4 Prior\ Ideator_i \cdot Period\ 2_t + \beta_5 Treatment_i \cdot Prior\ Ideator_i + \beta_6 Prior\ Ideator_i + \boldsymbol{X}_i'\gamma)],$$

where *F* is the nonlinear function of the model (e.g., logistic) and *it* denotes the observation. This average treatment effect (ATE) is displayed in columns 5-6 of Table 4. Informally, the treatment effect in the logit (column 5) is the change in the fraction of potential contributors among prior ideators in the treatment group, relative to that change among prior ideators in the control group.

The results in columns 5-6 of Table 4 suggest that prior ideators in the treatment group are not more likely to participate (logit process). This is not surprising, as the logit can be interpreted as describing whether the employee actively looks for ideas and is able to use the Portal. The latter is necessarily true for those who previously submitted ideas. However, rewards appear to have induced those who did not previously suggest ideas to participate, increasing the fraction of potential contributors by about 6.5 percentage points.[11] In the negative binomial process, treatment reduced the average number of suggested ideas by 0.76 for all other participating employees, while the effect on prior ideators is not significantly different from zero. The overall effect on number of ideas submitted (not displayed) is –0.021 for prior ideators, and –0.052 for employees who had not previously suggested ideas. The evidence does not support the crowding out hypothesis that rewards reduce intrinsic motivation. Finally, the findings are consistent with a cognitive tradeoff

[11] One might expect treatment effect in column 3 to be between those of subgroups in column 5. This is a nonlinear model with separate trends for each subgroup so this need not be the case.

Review of Economics and Statistics Just Accepted MS.
doi:10.1162/REST_a_00631
© by the President and Fellows of Harvard College and the Massachusetts Institute of Technology

between quantity and quality in ideation, and a rebalancing of multitask incentives towards quality. To shed more light on these questions, we now investigate the effect of the reward on idea quality.

Quality of ideas

For quality analyses we focus on ideas rather than authors. Idea j may have several authors N_j, so we split each idea into N_j observations, each with characteristics $\boldsymbol{X}_i$ of author i, where $i = 1, \ldots, N_j$. Hence, the unit of observation is the author-idea in quality analyses. To ensure that single- and multi-author ideas get the same weight, we weight each idea j by $1/N_j$. This allows us to explain idea outcomes with both author- and idea-specific variables. We use three quality measures: whether the idea was shared with the client, whether it was implemented, and its net value.

Consider first the dummy variable *Shared*, which equals 1 if an idea was shared with the client, and zero otherwise. Every idea is first reviewed internally, and the panel decides whether or not to share it with the client. *Shared* is a meaningful measure of quality for two reasons. First, it would be unwise of the company to bother clients with bad or trivial ideas. Second, the firm's goal was to increase its perceived value to the customer. Value-added is more salient for ideas that are explicitly communicated to the client than for ideas implemented without communication to clients. In Table 5 we report OLS coefficients and logit marginal effects of the following models:

OLS: $Shared_{ijm} = \mu_{ijm} + \varepsilon_{ijm}$ if $Finished_j = 1$,

Logit: $\text{Prob}\left(Shared_{ijm} = 1 \middle| \boldsymbol{X}_{im}, \boldsymbol{I}_j, Finished_j = 1\right) = \frac{\exp(\mu_{ijm})}{\exp(\mu_{ijm})+1}$,

where $\mu_{ijm} = \beta_1\, Treatment_i \cdot Period\ 2_m + \boldsymbol{X}'_{im}\gamma + \boldsymbol{I}'_j\delta + C_i + T_m$. *Treatment* is the treatment group dummy, $\boldsymbol{X}_{im}$ is a vector of employee characteristics at time of submission, $\boldsymbol{I}_j$ is a vector of idea-specific variables including project type, C_i is the client team fixed effect of employee i, and

Review of Economics and Statistics Just Accepted MS.
doi:10.1162/REST_a_00631
© by the President and Fellows of Harvard College and the Massachusetts Institute of Technology

T_m is the monthly time fixed effect.[12] Index *i* denotes the author, *j* the idea, and *m* the month of submission. Results are robust to estimating trend by period or month fixed effects. Analyses only include ideas with a final decision as *Shared* and *Implemented* are only available after review.

[TABLE 5 ABOUT HERE]

Columns 1-2 of Table 5 report results for the probability an idea is shared with the client. We find a substantial and statistically significant positive treatment effect. Depending on the model, the treatment increased the probability of sharing the idea by 19-21 percentage points. This suggests the treatment worked as intended. This could be due to employees shifting focus towards customer-related ideas, or increasing idea quality without a shift in focus. Both would increase the likelihood an idea is shared with the customer.

A second quality measure is *Implemented*: 1 if the idea is accepted for implementation, 0 if not. This is a measure of quality as the company has incentives to only approve worthwhile ideas. *Implemented* and *Shared* differ because an idea may be implemented without being shared with the client, or approved internally but rejected by the client (shared but not implemented). Results are in columns 3-4 of Table 5. Models are the same as above for *Shared*, with *Implemented* as the dependent variable. Once more we find a large positive and statistically significant treatment effect. Rewards improve the likelihood of implementation by 15-18 percentage points.

---

[12] A priori, one might think treatment affects the number of idea authors, which in turn might influence the probability an idea is shared. Including number of authors as a control would then bias treatment effect estimates. We found no statistically or economically significant treatment effect on number of authors, so include it in all quality regressions to improve fit. Treatment effect estimates are even larger when it is excluded. See Table II of the online appendix.

Review of Economics and Statistics Just Accepted MS.
doi:10.1162/REST_a_00631
© by the President and Fellows of Harvard College and the Massachusetts Institute of Technology

A third potential measure of idea quality is estimated *Profit*. This is part of the idea submission on the Portal and often uses supervisor input. Employees and supervisors receive training on how to assess costs and benefits. Estimates may be modified during executive review, and updated after implementation (our data are the most recent estimates). The company compared realized with predicted values. They told us they found estimates to be "very accurate," reflecting the training.

One might think the company would accept all ideas with positive profit, and vice versa. However, this is not the case. Strategic considerations, such as development of new competencies, client relationships, or synergies with other ideas induce the company to accept some ideas with negative net value and reject some with positive value. In general, the company places no emphasis on high dollar-value ideas, but rather wants employees to think about all ideas that may have any client benefit. This was true in treatment and control groups. Therefore, we have no reason to believe net profits would be affected by the treatment but include the regression for completeness.

Table 5 column 5 shows OLS results for projected profit, measured as the log of revenue minus cost. Note that ideas with zero or negative net value drop out with a log-transformation. We therefore also ran a linear model; it yielded the same qualitative results with inferior fit. Point estimates of the treatment effect are substantial and positive. An estimate of 0.308 in the log-linear specification implies a treatment effect of about +36%. However, the effect is not significantly different from zero, so we cannot reject the hypothesis that treatment had no effect on profit.

We tested for motivational crowding out, using quality measures and expected profit, by comparing treatment effects on prior ideators to those on all other employees. None of the regressions (available upon request) showed a significant difference in the treatment effect. Hence, there is no evidence of crowding out with respect to idea quality.

Review of Economics and Statistics Just Accepted MS.
doi:10.1162/REST_a_00631
© by the President and Fellows of Harvard College and the Massachusetts Institute of Technology

Regressions in Table 5 include author characteristics. Two points are worth mentioning. First, tenure is positively correlated with the probability an idea is shared, and with profit. Since these regressions control for salary group, this is in line with our finding of a positive correlation between tenure and quantity of ideas. Both results suggest firm- and / or client-specific human capital may improve ideation. Second, ideas with more authors have higher average quality for all measures. This is consistent with research finding that collaboration, particularly between those with different perspectives and skills, improves the quality of innovation (Ford & Sullivan 2004; Shin & Zhou 2007). An alternative explanation is that authors pool ideas and submit the best among them. Suppose N authors randomly form a group, and each author i has one idea of quality Y_i. Group quality, the N^{th} order statistic $Y_{(N)} = \max\{Y_1, \ldots, Y_N\}$, increases in N even without collaboration.

As a robustness check we ran the regressions as in Table 5 without employee controls (Table II, online appendix). Results are very similar, which shows that the difference-in-differences specification and controlling for unobserved time-invariant heterogeneity between client teams are important, but employee or idea level controls are not.

In sum, we find large positive treatment effects on idea quality, measured by the probability an idea is implemented, and probability it is shared with the client. We cannot reject the null hypothesis that estimated profits are unaffected by the experiment. In line with our findings on quantity, there is no evidence of motivational crowding out. The increase in idea quality might explain the drop in quantity per author detected above, reflecting a cognitive tradeoff between exploration and exploitation, and a response to multitask incentives rebalanced towards quality.

Post-treatment effects

The company discontinued the reward scheme after the 13-month experiment. We collected data for the 13 months that followed. This allows us to examine whether the positive treatment

Review of Economics and Statistics Just Accepted MS.
doi:10.1162/REST_a_00631
© by the President and Fellows of Harvard College and the Massachusetts Institute of Technology

effects on participation and idea quality persisted, which might suggest habituation, or disappeared. Alternatively, one might expect a reduction in idea quantity and quality – either because the incentive was removed, or due to perceived unfairness associated with removing the reward. Table 6 displays results on quantity and quality, now including a post-treatment effect (DID Post-treatment) for the 13 months following the experiment.

The post-treatment effect on the share of potential contributors (logit process of the ZINB model) remains positive and significant. Employees who became potential contributors due to rewards tended to continue participation after the reward was discontinued. This result might reflect habituation. Another potential explanation is that rewards changed the working culture of treatment teams, fostering greater attention to, or peer pressure for, ideation. Alternatively, the reward might have better communicated the importance to the firm of employee ideas, thereby changing implicit incentives. Unfortunately, the data do not allow us to disentangle these alternative interpretations. The effect on the number of ideas per potential contributor (negative binomial process, column 2) is insignificant in the post-treatment period. The overall effect on the quantity of ideas remains statistically and economically insignificant.

The table also reports results for one of our quality measures, *Shared* (effects are similar when using *Implemented* as the dependent variable). Interestingly, the treatment effect on quality vanished after elimination of the reward scheme, even though the treatment effect on participation remains. Apparently, employees continue to think about and submit ideas, but they no longer focus their effort on ideas with a high probability of client acceptance. Instead, they return to baseline with respect to idea quality. This could also explain why the effect on the number of ideas per participating employee becomes insignificant (column 2). Taken together, these findings are consistent with changes in multitask incentives during and after the treatment. In this view, the

Review of Economics and Statistics Just Accepted MS.
doi:10.1162/REST_a_00631
© by the President and Fellows of Harvard College and the Massachusetts Institute of Technology

rewards induced employees to spend additional effort fine-tuning and improving ideas to increase their odds of receiving a reward. Once the rewards were eliminated, relative incentives to emphasize quality returned to their original level. This is interesting for both theory and application, as it suggests that the quality of ideas can be manipulated with rewards.

Discussion of results

Organizational policies sometimes create unintended side effects that might confound our findings. Here we discuss three issues: distorted multitask incentives (Holmstrom & Milgrom 1991), influence costs (Milgrom & Roberts 1988), and favoritism (Prendergast & Topel 1996).

First consider distorted incentives. The Portal or experimental reward might have motivated employees to submit ideas that barely pass the implementation threshold, or that differ from ideas in non-treatment accounts in other ways (e.g., implementation risk). In addition, employees might have been induced to distort their description of an idea, how it would be implemented, or its costs and benefits, in order to increase odds of acceptance. Several features of the Portal mitigate such concerns. Supervisor consultation is a check that the description is not overstated, so the panel and client have the proper information for decisions. Career concerns give supervisors implicit incentives to avoid pushing marginal ideas. Panel members have significant human capital and high-level perspectives, so are in an excellent position to evaluate an idea's costs and tangible and intangible benefits. Being responsible for business unit success, they also have strong motivation to approve only valuable ideas. Further, the executive assigned to oversee the idea often consults other panelists, ensuring multiple perspectives and reducing biases. The most significant ideas are then vetted by the client, which has every incentive to reject poor ideas. Finally, ideas with substantial costs and complexity might be cancelled during implementation, reducing risk.

Review of Economics and Statistics Just Accepted MS.
doi:10.1162/REST_a_00631
© by the President and Fellows of Harvard College and the Massachusetts Institute of Technology

The experimental reward was designed to avoid distorting incentives. It did not reward mere suggestion of an idea, which would likely motivate employees to submit more ideas with little concern for cost, risk, or effect on the client. It only rewarded ideas accepted by clients. Moreover, the largest reward was for high client rating after implementation. These factors motivated employees to submit ideas of high quality from the client's perspective.

An additional concern is influence costs: employees might lobby the supervisor, review panel, or client to accept their ideas. A closely related concern is favoritism, in which supervisors, the review panel, or clients might be biased into accepting ideas suggested by employees whom they know and work with personally. Once more, the evaluation process makes such concerns unlikely to be significant, with or without the experimental reward. Control and treatment teams may have several thousand employees. Most employees who suggest ideas are several hierarchical levels below members of the review panel and do not personally know panel members. Ideas are randomly assigned to a review panel executive, and reviewed within 20 days, so employees cannot anticipate who will handle their idea, and have little time to maneuver once the idea is assigned. If their idea is shared, an employee's client contacts are usually several levels below the client decision maker, giving little room for lobbying. Further, review panel members and client decision makers have strong incentives to promote business goals, and little incentive to favor ideas from one employee over another, even if they personally knew an ideator. Thus, there is virtually no scope for influence costs or favoritism between the employee and either the review panel or client.

Influence costs and favoritism are most likely play a role between the employee and supervisor. However, as described above, that stage is idea refinement, not evaluation. The company hopes that employees "lobby" their supervisor for help improving the idea. Conceivably, the supervisor could help the employee write the idea description to understate concerns and overstate benefits,

Review of Economics and Statistics Just Accepted MS.
doi:10.1162/REST_a_00631
© by the President and Fellows of Harvard College and the Massachusetts Institute of Technology

but doing so would risk damaging the supervisor's reputation, evaluation, and prospects. Such a discovery is highly likely given the considerations described above. There may be some tendency towards favoritism, but it seems likely that supervisors will focus more on helping employees improve ideas rather than distort how they are presented.

The reward design also mitigates influence costs and favoritism. Clients were not aware of the experiment and, hence, did not know that rewards were at stake for ideators. Incentives for supervisors and review panel executives were not changed. Employees were rewarded for idea quality, not quantity. Therefore, it seems likely that the experimental reward motivated employees to ask supervisors for help in improving ideas, if anything.

Summing up, the Portal and reward were designed thoughtfully, and seem likely to avoid serious issues with distorted incentives, influence costs, or favoritism. However, we conducted several tests (available on request) seeking evidence for these problems. If lobbying drove results, the treatment effect should be larger for those with good contacts and expertise for successful lobbying. We investigated whether treatment effects varied by company tenure or hierarchical level, and found no differential effects for these groups. We checked whether other idea attributes, such as risk, changed with introduction of the reward. The data have no direct measure of idea risk, but we investigated estimated cost and log cost. We find no evidence that ideas become more or less costly due to the reward. Indeed, regressions on cost and log cost yield different signs for treatment effects, suggesting no clear directional effect. Further, as shown above, idea estimated net value is not affected by the reward, reinforcing the conclusion that idea quality did not suffer.

6. CONCLUSIONS

This paper provides a unique statistical analysis of employee innovation. Data from a company's employee suggestion system allow analysis of effects of a reward for fostering ideation,

Review of Economics and Statistics Just Accepted MS.
doi:10.1162/REST_a_00631
© by the President and Fellows of Harvard College and the Massachusetts Institute of Technology

implemented as a randomized field experiment. Our findings are of interest for the debate on the effects of incentives on creativity and intrinsic motivation. Most prior literature uses survey data, case studies, or laboratory experiments to study creativity. This study analyses experienced individuals performing their regular work. Moreover, random assignment to control and treatment groups, and use of difference-in-differences estimation, facilitate causal inferences.

Our findings suggest that rewards are a suitable tool to induce employees to think about process and product improvements, and to use a formal ideation system. The reward scheme substantially increased the likelihood that employees participate in the Idea Portal. Broadening participation was, in fact, one of the company's objectives. In fact, once familiar with the system, employees continued using it even in the absence of the reward scheme, suggesting habituation.

A main finding is that the treatment succeeded in increasing idea quality, measured by percent of ideas shared with the client, and percent of ideas implemented. At the same time the quantity of submitted ideas was not affected. Interestingly, employees returned to baseline with respect to the number and quality of ideas they suggest when the rewards were discontinued. Treatment and post-treatment responses on quality suggest that employees can fine-tune ideation and are responsive to incentives when doing so. The results are consistent with the notion of a quantity-quality tradeoff that may be addressed with multitask incentives. Proper choice of a performance measure is likely important. If the company had rewarded each submission rather than idea acceptance, we would likely have observed an increase in idea quantity at the expense of quality.

Employees suggested ideas even without formal rewards, indicating intrinsic motivation or implicit incentives. However, our results do not support the view that rewards crowd out intrinsic motivation. The reward motivated more employees to participate in ideation and focus on idea quality; those ideating prior to the experiment did not reduce innovation when offered the reward.

Review of Economics and Statistics Just Accepted MS.
doi:10.1162/REST_a_00631
© by the President and Fellows of Harvard College and the Massachusetts Institute of Technology

The findings illustrate the utility from experimentation in organizational policies. The firm was able to test a new incentive before rolling it out to the entire organization. In addition to providing useful information about whether the idea was sound, the experiment provided insights into the specific design of the policy. Conceivably the firm could have gone further. For example, they might have tried different rewards, perhaps to increase the total number of ideas or encourage collaboration in ideation. There may be substantial costs to implementing poorly-designed policies, and bureaucratic inertia might make it costly to change or eliminate a policy after it is implemented. Moreover, changes to an existing policy might reduce management's credibility with employees. Experimentation can help firms reduce these risks.

Our findings suggest a tradeoff between quantity and quality in ideation. How strong is this tradeoff? What are the implications for job design, team structure, and incentive plan design? Data on employee tasks and time management might allow for deeper understanding of this issue. Nevertheless, our findings suggest that incentives can motivate creativity if designed appropriately.

The company discontinued the experiment after 13 months of treatment. Its own analysis – a comparison of simple averages between treatment and control groups – suggested the incentive was ineffective, and it decided not to implement the reward scheme. There are three reasons why they reached a different conclusion. First, our approach controls for initial differences between the two groups. Second, we look at the subset of ideas that finished review, while the company included ideas still under review. That might have introduced a confounding factor if review speeds varied (teams with faster reviews might have misleadingly-higher acceptance rates). Third, we control for team composition, such as tenure, level, age and gender. Differences between groups in those characteristics may affect idea acceptance, but are not easily disentangled from the treatment effect using non-regression methods. After seeing our analysis, the company

Review of Economics and Statistics Just Accepted MS.
doi:10.1162/REST_a_00631
© by the President and Fellows of Harvard College and the Massachusetts Institute of Technology

implemented a similar reward scheme for the entire company. The only difference between the experimental reward and the one now used is that HCL added additional points for submitting ideas, to motivate more participation and a higher quantity of ideas in addition to higher quality.

REFERENCES

Amabile, Teresa (1982). Children's Artistic Creativity Detrimental Effects of Competition in a Field Setting. *Personality & Social Psychology Bulletin*, 8(3):573-578.

Amabile, Teresa (1983). *The Social Psychology of Creativity*. New York: Springer-Verlag.

Amabile, Teresa (1996). *Creativity in Context*. Westview Press, New York.

Ariely, Dan, Gneezy, Uri, Loewenstein, George & Mazar, Nina (2009). Large Stakes and Big mistakes. *Review of Economic Studies*, 76(2):451-469.

Azoulay, Pierre, Graff Zivin, Joshua & Manso, Gustavo (2011). Incentives and Creativity: Evidence from the Academic Life Sciences. *RAND Journal of Economics*, 42(3):527-554.

Benabou, Roland & Tirole, Jean. 2006. Incentives and Prosocial Behavior. *American Economic Review*, 96(5):1652-1678.

Bertrand, Marianne, Duflo, Esther & Mullainathan, Sendhil (2004). How Much Should We Trust Differences-in-differences Estimates? *Quarterly Journal of Economics*, 119(1):249-275.

Cameron, A. Colin, Gelbach, Jonah & Miller, Douglas (2008). Bootstrap-Based Improvements for Inference with Clustered Errors. *Review of Economics & Statistics*, 90(3):414-427.

Charness, G. & Grieco, D. (2013). Individual Creativity, Ex-ante Goals and Financial Incentives. Working paper, University of California, Santa Barbara.

Deci, Edward (1971). Effects of Externally Mediated Rewards on Intrinsic Motivation. *Journal of Personality & Social Psychology*, 18(1):105-115.

Review of Economics and Statistics Just Accepted MS.
doi:10.1162/REST_a_00631
© by the President and Fellows of Harvard College and the Massachusetts Institute of Technology

Deci, Edward, Koestner, Richard & Ryan, Richard (1999). A Meta-analytic Review of Experiments Examining Effects of Extrinsic Rewards on Intrinsic Motivation. *Psychological Bulletin*, 125(6):627–668.

Eckartz, Katharina, Kirchkamp, Oliver & Schunk, Daniel (2012). How do Incentives Affect Creativity? CESifo working paper 4049.

Ederer, Florian & Manso, Gustavo (2013). Is Pay for Performance Detrimental to Innovation? *Management Science*, 59(7):1496-1513.

Eisenberger, Robert & Armeli, Stephen (1997). Can Salient Reward Increase Creative Performance Without Reducing Intrinsic Creative Interest? *Journal of Personality & Social Psychology*, 72(3):652-663.

Eisenberger, Robert & Rhoades, Linda (2001). Incremental effects of reward on creativity. *Journal of Personality & Social Psychology*, 81(4):728-741.

Frey, Bruno & Jegen, Reto (2001). Motivation Crowding Theory. *J. Econ. Surveys* 15(5):589–611.

Frey, Bruno & Oberholzer-Gee, Felix (1997). The Cost of Price Incentives: An Empirical Analysis of Motivation Crowding-out. *American Economic Review,* 87(4):746–755.

Ford, Cameron & Sullivan, Diane (2004). A Time for Everything. *J. Org. Behavior*, 25(2):279-92.

George, Jennifer & Zhou, Jing (2001). When Openness to Experience and Conscientiousness Are Related to Creative Behavior. *Journal of Applied Psychology*, 86(3):513-524.

Gittins, John (1979). Bandit Processes and Dynamic Allocation Indices. *Journal of the Royal Statistical Society, Series B (Methodological)*, 41(2):148–177.

Guilford, Joy Paul (1950). Creativity. *American Psychologist,* 5(9):444–454.

Holmstrom, Bengt & Milgrom, Paul (1991). Multitask Principal-agent Analyses: Incentive Contracts, Asset ownership, and Job Design. *J. of Law, Economics & Organization*, 7:24-52.

Review of Economics and Statistics Just Accepted MS.
doi:10.1162/REST_a_00631
© by the President and Fellows of Harvard College and the Massachusetts Institute of Technology

Joussemet, Mireille & Koestner, Richard (1999). Effect of Expected Rewards on Children's Creativity. *Creativity Research Journal*, 12(4):231-239.

Leach, Desmond, Stride, Christopher & Wood, Stephen (2006). The Effectiveness of Idea Capture Schemes. *International Journal of Innovation Management,* 10(3):1-26.

Milgrom, Paul & Roberts, John (1988). An Economic Approach to Influence Activities in Organizations. *American Journal of Sociology*, 94:S154-S179.

Ohly, Sandra, Leach, Des & Zou, Qin (2013). Suggestion-Making in Organizations: A Review and Multilevel Model. Manuscript, University of Kassel.

Prendergast, Canice & Topel, Robert (1996). Favoritism in Organizations. *Journal of Political Economy*, 104(5):958-978.

Robinson, Alan & Stern, Sam (1997). *Corporate Creativity*. San Francisco: Berrett-Koehler.

Robbins, Herbert (1952), Some Aspects of the Sequential Design of Experiments. *Bulletin of the American Mathematical Society*, 58(5):527–535.

Sauermann, Henry & Cohen, Wesley (2010). What Makes Them Tick? Employee Motives and Firm Innovation. *Management Science*, 56(12):2134-2153.

Shin, Shung Jae & Zhou, Jing (2007). When is Educational Specialization Heterogeneity Related to Creativity in Research and Development Teams? *J. Applied Psychology*, 92(6):1709-1721.

Tierney, Pamela, Farmer, Steven & Graen, George (1999). An Examination of Leadership and Employee Creativity. *Personnel Psychology*, 52(3):591-620.

Toubia, Olivier (2006). Idea Generation, Creativity & Incentives. *Marketing Science* 25(5):411-5.

van Dijk, Christiaan & van den Ende, Jan (2002). Suggestion Systems: Transferring Employee Creativity Into Practicable Ideas. *R&D Management,* 32(5):387-395.

Review of Economics and Statistics Just Accepted MS.
doi:10.1162/REST_a_00631
© by the President and Fellows of Harvard College and the Massachusetts Institute of Technology

Figure 1. Process for Evaluating New Ideas

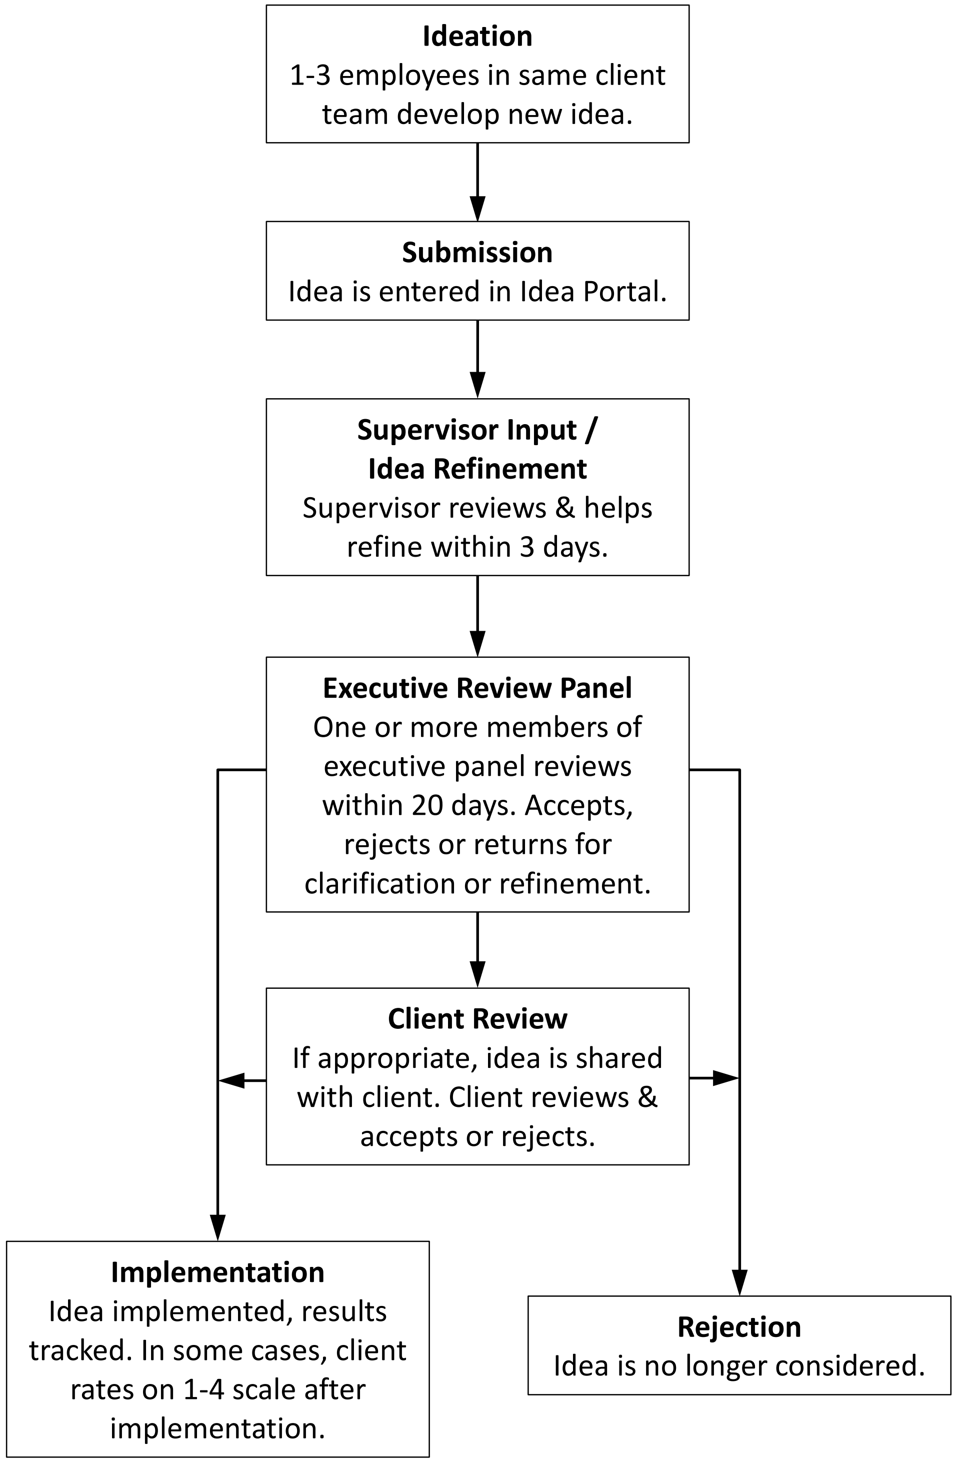

Review of Economics and Statistics Just Accepted MS.
doi:10.1162/REST_a_00631
© by the President and Fellows of Harvard College and the Massachusetts Institute of Technology

Table 1: Descriptive statistics

	Period	Treatment Group	Control Group	Combined
Number of Employees	1	3015	3185	6200
	2	5260	5881	11141
	3	4061	4511	8572
Mean Age	1	29.4**	29.9	29.7
		(4.70)	(4.87)	(4.80)
	2	29.7	29.6	29.7
		(5.38)	(5.41)	(5.40)
	3	30.5	30.3	30.4
		(5.50)	(5.54)	(5.52)
Mean Tenure	1	2.99	2.90	2.94
		(2.32)	(2.34)	(2.33)
	2	2.97	2.80	2.88
		(2.48)	(2.37)	(2.42)
	3	3.73	3.53	3.62
		(2.53)	(2.30)	(2.41)
Share of Men	1	0.78	0.80	0.79
		(0.42)	(0.40)	(0.41)
	2	0.76	0.77	0.76
		(0.43)	(0.42)	(0.43)
	3	0.76	0.77	0.76
		(0.43)	(0.42)	(0.43)

Review of Economics and Statistics Just Accepted MS.
doi:10.1162/REST_a_00631
© by the President and Fellows of Harvard College and the Massachusetts Institute of Technology

Table 1: Descriptive statistics (cont.)

	Period	Treatment Group	Control Group	Combined
Mean Salary Group	1	1.43	1.45	1.44
		(0.73)	(0.77)	(0.75)
	2	1.47	1.47	1.47
		(0.79)	(0.82)	(0.81)
	3	1.49	1.49	1.49
		(0.81)	(0.84)	(0.82)
Share of Prior Ideators	1	0.087	0.063	0.075
		(0.28)	(0.24)	(0.26)
	2	0.049	0.035	0.042
		(0.22)	(0.18)	(0.20)
	3	0.064	0.045	0.054
		(0.24)	(0.21)	(0.23)

Note: Standard deviations of the means are displayed in parentheses. In each line, the difference of group means is tested with a t-test using standard errors that are clustered at client team level. ***Significant at the 1% level; **significant at the 5% level; *significant at the 10% level. *Number of employees* in period 1 (pre-treatment) is based on employment roster 1; *Number of employees* in period 2 (treatment) is based on rosters 1 and 2; *Number of employees* in period 3 (post-treatment) is based on roster 2. *Age* and *tenure* are measured at the end of the respective period.

Review of Economics and Statistics Just Accepted MS.
doi:10.1162/REST_a_00631
© by the President and Fellows of Harvard College and the Massachusetts Institute of Technology

Table 2: Summary statistics on outcome variables by group for pre-treatment and treatment period

Group	Period	# of ideas	Ideator*	Authors	Finished	Imp|Fin*	Shared|Fin*	Log(Net Value)
Treatment	1	517	0.195	1.716	0.745	0.410	0.449	8.705
			(0.355)	(0.893)	(0.436)	(0.379)	(0.440)	(2.030)
	2	566	0.083	1.309	0.643	0.516	0.597	9.428
			(0.264)	(0.550)	(0.480)	(0.461)	(0.450)	(2.023)
Control	1	361	0.336	1.903	0.634	0.638	0.415	8.865
			(0.285)	(1.118)	(0.482)	(0.249)	(0.398)	(2.083)
	2	363	0.112	1.402	0.510	0.712	0.339	9.257
			(0.204)	(0.667)	(0.501)	(0.366)	(0.446)	(2.146)

Note: Ideator is the share of employees who submitted at least one idea in the respective period. *Authors* is the mean number of authors per idea. *Finished* is the share of ideas with finished review; those ideas are either accepted for implementation or rejected. *Imp|Fin* and *Shared|Fin* denote the shares of ideas selected for implementation or for sharing with the client, respectively, among ideas with finished review. Log(*Net Value*) is the mean of the logarithm of the projected profit in US dollar terms. Period 1 is pre-treatment, period 2 is the treatment period. * denotes client team means (pooled standard deviations).

Review of Economics and Statistics Just Accepted MS.
doi:10.1162/REST_a_00631
© by the President and Fellows of Harvard College and the Massachusetts Institute of Technology

Table 3: Who ideates? Influence of employee characteristics on ideation

	(1) Logit AME	(2) ZINB AME	(3) QML Poisson AME
Dependent variable	Ideator	Number of Ideas	Number of Ideas
Age	-0.007***	-0.017***	-0.080***
	(0.001)	(0.004)	(0.023)
Tenure	0.018***	0.030***	0.310***
	(0.003)	(0.009)	(0.059)
Male	0.004	-0.056	-0.298
	(0.009)	(0.042)	(0.221)
Salary Groups 0 & 1 pooled	-0.094**	-0.516*	-1.308***
	(0.039)	(0.291)	(0.408)
Salary Group 2	0.005	-0.119	0.009
	(0.034)	(0.100)	(0.339)
Salary Group 3	0.075	0.022	0.531
	(0.048)	(0.074)	(0.429)
Client FE	yes	yes	yes
Log Pseudo likelihood	-1416.53	-2177.45	-2954.93
Clusters	18	19	19
Observations	5887	5916	5916

Note: The regressions use data from the pre-treatment period only, where both groups have identical incentives. *Ideator* is a dummy indicating whether an employee submitted at least one idea in the given period. *Number of Ideas* is the number of ideas submitted within the 13 pre-treatment months. *Salary group* is an indicator for an employee's position in the company hierarchy. The reference category is (upper) management, that is, salary groups 4 and above. The marginal effects of *Age* and *Tenure* are based on linear and quadratic terms. Standard errors are clustered at the client team level. ***Significant at the 1% level; **significant at the 5% level; *significant at the 10% level.

Review of Economics and Statistics Just Accepted MS.
doi:10.1162/REST_a_00631
© by the President and Fellows of Harvard College and the Massachusetts Institute of Technology

Table 4: Number of suggested ideas per employee and period

	(1) Overall effect	(2) Overall effect	(3) ZINB	(4) ZINB	(5) ZINB	(6) ZINB
	OLS	ZINB AME	Logit AME	NB AME	Logit AME	NB AME
Dependent variable	NumIdeas	NumIdeas	Pr(Participation)	NumIdeas|Part.	Pr(Part.)	NumIdeas|Part.
DID Treatment	0.001	-0.004	0.178***	-0.264**	0.065***	-0.761***
	(0.063)	(0.051)	(0.049)	(0.132)	(0.019)	(0.281)
DID Treat. Prior Ideator					0.003	-0.207
					(0.056)	(0.331)
Age	-0.007	-0.005**	-0.008***	-0.004	-0.003**	0.014
	(0.008)	(0.002)	(0.002)	(0.008)	(0.001)	(0.010)
Age^2	-0.000					
	(0.000)					
Tenure	0.061***	0.027***	0.060***	0.005	0.018***	0.005
	(0.012)	(0.004)	(0.010)	(0.030)	(0.004)	(0.029)
$Tenure^2$	-0.004**					
	(0.002)					
Male	-0.016	-0.015	0.052***	-0.165*	0.013	-0.303**
	(0.014)	(0.015)	(0.013)	(0.086)	(0.009)	(0.141)

Review of Economics and Statistics Just Accepted MS.
doi:10.1162/REST_a_00631
© by the President and Fellows of Harvard College and the Massachusetts Institute of Technology

Table 4: Number of suggested ideas per employee and period (cont.)

	(1) Overall effect	(2) Overall effect	(3) ZINB	(4) ZINB	(5) ZINB	(6) ZINB
	OLS	ZINB AME	Logit AME	NB AME	Logit AME	NB AME
Dependent variable	NumIdeas	NumIdeas	Pr(Participation)	NumIdeas|Part.	Pr(Part.)	NumIdeas|Part.
Controls salary groups	yes	yes	yes	yes	yes	yes
Client FE	yes	yes	yes	yes	yes	yes
Time FE	period	period	period	period	period	period
R^2	0.088					
Log Pseudo likelihood		-5266.22	-5266.22	-5266.22	-4558.93	-4558.93
Clusters	19	19	19	19	19	19
Observations	17045	17045	17045	17045	17045	17045

Note: The table reports the results from OLS and zero inflated negative binomial models used to explain the number of ideas per author and period. The zero inflated model is a mixture model, where a logit and negative binomial process jointly explain the dependent variable. Columns (1) and (2) display the effect on the number of ideas per author and period. Columns (3) and (4) display the average marginal effects for the two processes separately. Columns (4) and (5) display the marginal effects for the logit and the negative binomial process, allowing the treatment effect to vary between *Prior ideators* and the rest. The unit of observation is the author-period, i.e., all ideas within a period of 13 months are summed up. Marginal effects of *Age* and *Tenure* are based on linear and quadratic terms. Standard errors are clustered at the client team level. ***Significant at the 1% level; **significant at the 5% level; *significant at the 10% level.

Review of Economics and Statistics Just Accepted MS.
doi:10.1162/REST_a_00631
© by the President and Fellows of Harvard College and the Massachusetts Institute of Technology

Table 5: Treatment effects on different measures of idea quality

	(1) OLS	(2) Logit AME	(3) OLS	(4) Logit AME	(5) OLS
Dependent variable	Shared	Shared	Implemented	Implemented	Log(Net Value)
DID Treatment	0.209**	0.188**	0.153*	0.177**	0.308
	(0.087)	(0.079)	(0.084)	(0.087)	(0.390)
Number of Authors	0.082***	0.092***	0.060**	0.064***	0.185**
	(0.020)	(0.025)	(0.023)	(0.018)	(0.069)
Age	-0.036	-0.016***	-0.037	-0.012*	0.341*
	(0.025)	(0.005)	(0.025)	(0.006)	(0.182)
Age^2	0.000		0.000		-0.005*
	(0.000)		(0.000)		(0.003)
Tenure	0.044**	0.032***	0.014	0.008	0.129*
	(0.019)	(0.012)	(0.022)	(0.011)	(0.061)
$Tenure^2$	-0.002		-0.001		-0.005
	(0.002)		(0.002)		(0.006)
Male	0.038	0.037	0.114***	0.111***	-0.060
	(0.055)	(0.053)	(0.040)	(0.034)	(0.273)
Controls salary groups	yes	yes	yes	yes	yes
Controls project type	yes	yes	yes	yes	yes
Client FE	yes	yes	yes	yes	yes
Time FE	month	month	month	month	month
R^2	0.687		0.753		0.962
Log Pseudo likelihood		-546.45		-457.07	
Clusters	19	15	19	17	19
Observations	1779	1697	1779	1747	1912

Review of Economics and Statistics Just Accepted MS.
doi:10.1162/REST_a_00631
© by the President and Fellows of Harvard College and the Massachusetts Institute of Technology

Table 5: Treatment effects on different measures of idea quality (cont.)

Note: The table reports estimates of OLS and logistic regressions using as outcome variables the probability that an idea is shared with the customer (columns 1 and 2), the probability that an idea is accepted for implementation (columns 3 and 4), and the logarithm of the projected net value (profit from the idea) (column 5). The treatment effect is the difference-in-differences estimator. The unit of observation is the author-idea. Each observation is weighted by 1/(*Number of Authors*), where *Number of Authors* represents the number of employees who submit the idea together. Only ideas with finished review process (either accepted or rejected) are included in the samples for columns (1) to (4). Marginal effects of *Age* and *Tenure* are based on linear and quadratic terms. Standard errors are clustered at the client team level. ***Significant at the 1% level; **significant at the 5% level; *significant at the 10% level.

Review of Economics and Statistics Just Accepted MS.
doi:10.1162/REST_a_00631
© by the President and Fellows of Harvard College and the Massachusetts Institute of Technology

Table 6: Idea quantity and quality, treatment and post treatment effects

	(1) ZINB	(2) ZINB	(3) Quantity effect	(4) Quality
	Logit AME	NB AME	ZINB AME	Logit AME
Dependent variable	Pr(Part.)	NumIdeas|Part.	NumIdeas	Shared
DID Treatment	0.187***	-0.197**	-0.014	0.180***
	(0.062)	(0.094)	(0.040)	(0.066)
DID Post Treatment	0.255**	-0.160	0.026	0.019
	(0.106)	(0.130)	(0.056)	(0.202)
Age	-0.008***	0.006	-0.002	-0.012***
	(0.002)	(0.008)	(0.002)	(0.004)
Tenure	0.069***	0.009	0.028***	0.034***
	(0.009)	(0.015)	(0.003)	(0.010)
Male	0.053***	-0.090	-0.007	0.021
	(0.019)	(0.065)	(0.015)	(0.045)
Controls salary groups	yes	yes	yes	yes
Client FE	yes	yes	yes	yes
Time FE	period	period	period	month
Log Pseudo likelihood	-7767.68	-7767.68	-7767.68	-810.98
Clusters	19	19	19	16
Observations	25152	25152	25152	2310

Note: The table reports marginal effects for a zero inflated negative binomial model explaining the number of ideas per author and period, and for a Logit model explaining the probability of sharing an idea with the client. Marginal effects of *Age* and *Tenure* are based on linear and quadratic terms. Standard errors are clustered at the client team level. ***Significant at the 1% level; **significant at the 5% level; *significant at the 10% level.