

Crowdsourcing New Product Ideas under Consumer Learning

Yan Huang

Heinz College
Carnegie Mellon University
yanhuang@andrew.cmu.edu

Param Vir Singh

Tepper School of Business
Carnegie Mellon University
psidhu@cmu.edu

Kannan Srinivasan

Tepper School of Business
Carnegie Mellon University
kannans@andrew.cmu.edu

Yan Huang is a PhD student at Heinz College, Param Vir Singh is Assistant Professor of Information Systems, and Kannan Srinivasan is Rohet Tolani Distinguished Professor of International Business, H.J. Heinz II Professor of Management, Marketing and Information Systems. All at Carnegie Mellon University. Authors thank seminar participants at Human Computer Institute and Tepper School of Business at Carnegie Mellon University, Smith School of Business at University of Maryland, International Conference of Information Systems, Workshop on Information Systems and Economics, Conference of Information Systems and Technology for helpful suggestions on earlier versions of this paper.

Crowdsourcing New Product Ideas under Consumer Learning

Abstract

Crowdsourcing initiatives are becoming a popular tool for new idea generation for firms. Although such initiatives are widely adopted in many different industries, the number of ideas generated often decline over time, and the implementation rates (percentage of posted ideas that are implemented by the firm) are quite low. Critics of crowdsourcing attribute these observations to three key factors: 1. individuals' limited view about firms' products, leading to the contributions of mainly niche ideas; 2. consumers' limited knowledge about firms' cost structure, leading to the proposals of mostly infeasible ideas; and 3. firms' lack of response to customers' ideas, leading to customer dissatisfaction. To investigate these criticisms in detail and to devise policies for firms to alleviate these concerns, we propose a structural model to capture individual idea contribution dynamics. We estimate the model using a rich dataset obtained from IdeaStorm.com, which is a crowdsourcing website affiliated with Dell. On this website, individuals can contribute ideas and vote on other's ideas. The firm then decides which ideas to implement.

Using the peer voting score, we are able to infer the true potential of ideas, whereas the cost to the firm for implementing the idea is indirectly imputed from the idea implementation data. We find that individuals tend to significantly underestimate the costs to the firm for implementing their ideas but overestimate the potential of their ideas in the initial stages of the crowdsourcing process. Therefore, the "idea market" is initially overcrowded with ideas that are less likely to be implemented. However, individuals learn about both their abilities to come up with high potential ideas as well as the cost structure of the firm from peer voting on their ideas and the firm response to contributed ideas. We find that the individuals learn rather quickly about their abilities to come up with high potential ideas, but the learning regarding the firm's cost structure is quite slow. We also find that an individual's discontent adversely affects the individual's continuous participation in idea contributions. As a result of the learning process, the crowdsourcing market becomes more efficient. Contributors of low potential ideas eventually drop out, while the high potential idea contributors remain active. Over time, the average potential of generated ideas increases, while the number of ideas contributed decreases. Hence, the firm can reduce the cost of screening ideas without losing high potential ideas. In our policy simulation, we show that providing more precise cost signals to individuals can accelerate the filtering process. Increasing the total number of ideas to respond to and improving the response speed will lead to more idea contributions. However, failure to distinguish between high and low potential ideas and between high and low ability idea generators lead to the overall potential of the ideas generated to drop significantly.

Keywords: Crowdsourcing, Structural Modeling, Dynamic Learning, Heterogeneity, Econometric analyses, Utility

1. Introduction

Product innovation has been an important area of business academic research. Recent advances in information technology have allowed firms to enhance their direct communication with customers, and the interaction has become an interesting source of new product ideas. Leveraging the opportunity, firms now create online idea markets where consumers can post new product ideas that are evaluated for their market potential by their peers. Jeff Howe (2006) named this new approach crowdsourcing, and he defined crowd as “the new pool of cheap labor: everyday people using their spare cycles to create content, solve problems, even do corporate R & D.” Crowdsourcing initiatives provide individuals with a platform to express their ideas, which are typically generated from their experience with actual product usage or observing others using the product. The ideas that come from the customer crowds can reveal rich information about customers’ preferences. Typical crowdsourcing platforms allow other customers to promote or demote ideas of their peers, thus providing an important early assessment of the potential of the proposed ideas. Firms can potentially obtain a large number of novel and profitable ideas at relatively low costs from such initiatives. Early adopters of this approach include some of the highly regarded business firms, such as Dell, Threadless, Starbucks, Adidas, Salesforce, BBC, BMW, Ducati, and Muji.

Although crowdsourcing initiatives have become rapidly popular in a variety of industries, the usefulness of this new approach is still under debate. Critics of such initiatives raise three important concerns. First, they argue that the individuals might be too accustomed to current consumption conditions and their own specific needs and hence, are more likely to suggest niche ideas with little market potential. Second, unlike the internal R&D teams, customers of the firm are unaware of the internal cost structure of the firm and hence, are quite likely to suggest ideas that are not viable (Schulze and Hoegl 2008). As a result, the firm typically has to invest significant effort to screen ideas, most of which have low potential and are generally infeasible. Third, individuals are often disconcerted¹ by the firm’s slow or no response to their ideas and eventually stop contributing ideas. The low implementation rate of ideas and the decline in the number of ideas posted as observed in practice seem to be consistent with the arguments against crowdsourcing. However, there is no systematic research that has investigated these issues in depth. In this study, we present an empirical framework to investigate the customers' abilities to contribute viable high potential ideas. We also highlight ways in which a firm can help its customers contribute such ideas.

¹ Individuals complain that the firm ignores their ideas; thus, they are disappointed and feel that it is a waste of time to post an idea. One individual wrote in a comment, “You’re also right, Tukulito (another individual’s ID), that Dell has NOT responded in so MANY areas. It’s been extremely frustrating.” Another individual said, “Many individuals have lost interest in IdeaStorm lately because IdeaStorm, the way it stands now is, frankly, stagnant... I’m sure many individuals have lost interest in IdeaStorm in part because they’re led to believe that their ideas are disregarded/ignored now... And it’s not just like Dell doesn’t implement any ideas now. I don’t think Dell has even commented or updated many ideas lately, even the most popular or most requested ideas...”

Enthusiastic consumers may propose new product ideas, but they have no initial idea as to how good their ideas are and may simply overestimate the potential of their ideas. Peer evaluations provide a valuable and important source of real-time feedback. A strong negative vote will let the consumer know that the idea may not be that useful after all. More importantly, when a string of new product ideas are turned down by peer consumers, the individual may conclude that, contrary to his/her initial belief, he/she is not a sophisticated generator of new product ideas. Thus, through experience and learning, those customers who are ‘bad’ at coming up with new ideas (marginal idea contributors) recognize their inabilities and may reduce the number of ideas they propose over time and eventually stop generating new ideas. In contrast, ‘good’ new product idea generators (good idea contributors)² will be encouraged to continue to provide new product ideas.

Such a learning model is entirely consistent with an overall decline in new product ideas over a period of time. Thus, a decreasing number of ideas may well reflect an efficient idea market and its resulting success rather than the ineffectiveness of the idea market. We propose and show that such learning ideas find strong empirical support.

Another important impediment to the implementation of new product idea is the cost of implementing the idea. Unfortunately, consumers have little or no understanding of this critical factor. However, consumers can learn or can infer the cost to implement ideas. Consumers cannot infer a firm’s cost structure from unimplemented ideas because firms have not made decisions on those ideas. However, when an idea is implemented, firms usually publicize their implementation decision and provide details about how they implement it. The information is broadcasted to all individuals in the community and by combining that information with the idea’s voting score, consumers can learn how costly it is for the firm to implement similar kinds of ideas. Such sophisticated learning by consumers eventually results in the generation of ideas where cost will not be an impediment for eventual implementation. Our model estimates consumers’ inferences of cost structures.

Our results show that initially contributors tend to underestimate the costs for implementing their ideas and to overestimate the potential of their ideas. Therefore, marginal idea contributors tend initially to post many low potential, unviable ideas. However, as individuals learn (update their beliefs) about the firm’s cost structure and the potential of their ideas, marginal idea contributors gradually drop out. The remaining active contributors are, for the most part, good idea contributors. Consequently, although the number of ideas generated decreases over time, the average potential of ideas posted significantly increases over time.

² In our model, an individual’s type is determined by the average potential of ideas generated by this person. The individual’s type is continuous because average potential is a continuous variable. When we say an individual is a “good idea contributor,” it means that the average potential of ideas generated by the individual falls in a higher region of the distribution. When we say an individual is a “marginal idea-contributor,” it means that the average potential of ideas generated by the individual falls in the lower region of the distribution.

These findings show that, over time, marginal idea contributors are filtered out and that the idea market becomes more efficient.

The estimation results also show that individuals learn about their own types faster than they learn about the cost structure of the firm because the cost signals the firm provides are quite imprecise. We also find that individuals feel disincentivized to contribute ideas if the firm does not reply to their submissions or takes an extended period of time to reply.

Our policy simulations indicate that Dell can accelerate the filtering out of marginal idea contributors by providing more precise cost signals. Under another policy experiment, we find that actively responding to all unimplemented ideas will adversely affect the filtering process because marginal idea-contributors who would drop out under the current policy will stay longer under the new policy. As a result, the firm would end up with more low potential ideas. In other words, the firm is better off when it selectively responds to ideas. Providing feedback on ideas with higher votes can improve the average idea potential in the later periods; however, the improvement is insignificant. The best policy is to identify good idea contributors and respond quickly to their ideas. By doing so, good idea contributors will be less disincentivized and will be encouraged to contribute more high potential ideas. Our last set of policy simulations show that if the firm wants to provide additional incentive for consumers to contribute ideas, it should reward individuals only when their ideas are implemented, rather than reward individuals when they post ideas. By doing so, the firm can achieve the same improvement on the overall potential of ideas at a lower cost.

2. Relevant Literature

Our paper is related to the emerging literature on crowdsourcing. Although crowdsourcing has attracted enormous business and media attention, there are very few academic studies on crowdsourcing. Initiatives by established firms to encourage customers for participation in the design of new products represents the most popular form of crowdsourcing being currently used and studied (Terwiesch and Xu 2008). Such crowdsourcing initiatives soliciting new product design ideas can be classified into three types. In the first type the creation of a vaguely specified product depends wholly on customer input. Threadless.com is an example of such an initiative where customers develop t-shirt designs on their own and submit the finished designs to Threadless. The second type of crowdsourcing is related to the first type in that the final product depends wholly on the customer input but differs from the first type in that the customers have to solve a specifically defined task or problem (Boudreau et al 2011, Jeppesen et al 2010). Crowdsourcing efforts at Topcoder or Innocentive correspond to this type. The first two types are also similar to each other in that in both of them the contributors typically compete with each other for a fixed monetary reward. Hence they are also classified as crowdsourcing contests. The third type of crowdsourcing corresponds to a permanent open call for contribution that is not directed towards any particular task or problem (Bayus 2010, Di Gangi et al. 2010).

Dell Ideastorm corresponds to this type. In this type of crowdsourcing consumers typically only contribute and evaluate variety of ideas and it is up to the firm to develop and implement those ideas.

Most of the studies on crowdsourcing have analyzed crowdsourcing contests where contributors compete with each other to win a prize (Archak and Sundararajan 2009, DiPalantino and Vojnovic 2009, Mo et al. 2011, Terwiesch and Xu 2008). In contrast to crowdsourcing contests, in permanent open call crowdsourcing ideation initiatives such as Ideastorm contributors do not compete with each other but help evaluate each other's contributed ideas. In the context of crowdsourcing ideation, using a reduced form approach Bayus (2010) finds that individual creativity is positively related to current effort but negatively related to past success. Di Gangi et al. (2010) find that the decision to adopt a user contributed idea is affected by the ability of the firm to understand the technical requirements and respond to community concerns regarding the idea. Lu et al. (2011) find important complementarities in crowdsourcing ideation and customer support initiatives. They find that customer support platforms provide opportunities for customers to learn about the problems other customers are facing and that helps them in suggesting better ideas for firm to implement. To our knowledge, we are the first to structurally examine the new product idea and development process based on actual crowdsourcing data.

Our paper is also related to the literature on consumer Bayesian learning. Bayesian learning models are widely applied to analyze consumers' choices under uncertainty. Erdem and Keane (1996) and Erdem et al. (2008) investigate customer learning of brand qualities from multiple resources, such as past experience, advertisement, and price. While Mehta et al. (2003) study the formation of consideration sets, Crawford and Shum (2005) and Narayanan and Manchanda (2009) examine the physicians' learning of drug prescription. Zhang (2010) develops a dynamic model of observational learning and analyzes the kidney adoption in the U.S. kidney market. In our paper, we apply the Bayesian learning model to the individual's learning of their own type and learning of the firms' cost structure to better understand the dynamics of idea posting behavior.

3. Research Context

Our data are from a crowdsourcing website, IdeaStorm.com, which is operated by Dell. Dell launched this website in February 2007 as a way to communicate directly with its customers. IdeaStorm.com was created to give a direct voice to Dell's customers and an avenue to have online brainstorming sessions, allowing the customers to share ideas and collaborate with one another and with Dell. The goal of this initiative was to hear what new products or services Dell's customers would like to see Dell develop.

The structure of Ideastrom.com is quite simple, yet effective. Any individual (not necessarily a customer) can register on the website to participate in the initiative. Once registered, an individual can then

post any relevant idea. Dell assigns 500 Dell points to the contributor for each idea³. Once an idea is posted, all the other individuals can vote on the idea. They can either promote the idea, which yields an additional ten points for the idea contributor, or demote the idea, which results in a ten point deduction. Individuals are also allowed to comment on ideas and express their opinions in greater detail. Dell uses the peer voting and comments to gauge the potential of contributed ideas. Dell assigns web managers to maintain the website, and their job is to pass the ideas generated by the individuals on to the corresponding groups within the company for review. The web managers communicate with the individuals through 1) direct comments about the ideas and 2) changes in the status of the idea. Typically, the evolution of an idea’s status is as follows.

Most of the posted ideas posted “Acknowledged” within 48 hours. If the web managers find an idea is already part of their existing product or services, they will change the status to “Already offered”. Among the remaining ideas, the web managers selectively pass ideas to related departments for review, and the status is changed to “Under Review”. After carefully evaluating these ideas, Dell makes one of three decisions: “Implemented”, “Partially Implemented” or “Not Planned”. Once an idea is “Implemented”, it is closed for votes and comments. Dell also provides details regarding the decision through comments or blogs. “Partially Implemented” and “Not Planned” ideas are not closed, which means that individuals can still vote and comment on these ideas, and it is possible that at some point, Dell will re-evaluate the ideas. Ideas that do not receive any comments within a year are “Archived” and thus no longer available for individuals to view (IdeaStorm.com). All individuals can see how many peer votes an idea received and which ideas have been implemented by the firm. In this way, our modeling framework allows the individuals to learn from these two observations.

Dell broadly categorizes all the ideas into three categories: Product Ideas, Dell Ideas, and Topic Ideas (see Table 1). Each idea could be related to up to three sub-categories. When an individual posts an idea on IdeaStorm, he/she selects the Category as well as the sub-categories to which the idea belongs.

Table 1. Idea Classification

Categories	Sub Categories
Product Idea	Accessories (Keyboards, etc.); Adamo; Alienware; Broadband and Mobility; Desktops; Dimension; Inspiron; Laptop Power; Laptops; Latitude; Linux; Mobile Devices; Monitors and Displays; Desktops and Laptops; Netbooks; New Product Ideas; Operating Systems; OptiPlex; Precision Workstations; Printers and Ink; Servers and Storage; Software; Studio; Vostro; XPS
Dell Idea	Advertising and Marketing; Dell; Dell Community; Dell Website; IdeaStorm; Retail; Sales Strategies; Service and Support
Topic Idea	Digital Nomads; Education; Enterprise; Environment; Gaming; Healthcare and Life Sciences; PartnerStorm; Small Business; Storm Session Topics; Women's Interest

³ This policy was changed in December 2008.

Data

Our data have expanded from the initiation of IdeaStorm.com in early 2007 to the end of 2010. By the end of 2010, more than 12,000 ideas had been contributed and more than 400 had been implemented. However, we only use the data from the initiation of IdeaStorm.com to September 2008. During October 2008, a large number of material changes were made to the initiative, and therefore, we restrict our attention to data prior to these changes. We also exclude data from the first two weeks because the number of ideas contributed during this period was extremely small (≤ 5), perhaps due to the public's lack of awareness of the website. Furthermore, most of the initial ideas during this period were suggested by Dell's employees. After the elimination of the initial period, we have 84 weeks of data (Week 3 to Week 86). In our dataset, most of the ideas fall into the first two categories, though there are a few ideas that belong to Category 3 (less than 10% of the number of ideas in Categories 1 and 2, see Table 2). Therefore, our analysis focuses on the first two categories of ideas.

Table 2. Summary Statistics by Category

Category	1	2	3
Category Name	Product idea	Dell idea	Topic idea
# Posted	5337 (1419)*	4243(1565)	392(108)
# Implemented	100(41)	110(54)	10(3)
% Implemented	1.87 (2.89)	2.59(3.45)	2.55(2.78)
Average log (votes)	4.626(5.286)	4.580(5.600)	4.352(4.556)
SD of log (votes)	2.160 (1.875)	2.147(1.696)	2.720(2.742)

*Numbers outside the parentheses are full sample statistics; numbers inside the parentheses are statistics of the sample of 490 selected individuals

Table 3. Summary Statistics for Individuals

Variables	Mean	Std. dev	Min	Max
Mean Log(votes)	4.819	1.513	-4.000	7.667
Number of Ideas Contributed	7.269	19.411	2	164
First Time Post (Week)	22.41	21.76	3	83

Figure 1 shows the distribution of the number of ideas generated by each individual (the total number of individuals in the full sample is 11,811). It follows the power law distribution, which is common in the user generated content settings. A majority of individuals on the website only vote but never post any new product idea. In addition, among those who posted an idea, most posted only one idea during these 84 weeks. The notion of learning is meaningful only when a respondent posts at least two ideas. The 490 individuals who posted two or more ideas constitute fewer than 5% of the consumers on the site but account for nearly 40% of all new product ideas. Table 3 shows the important statistics of these individuals who posted two or more ideas. We observe that there is significant variation among individuals in terms of mean log(votes), number of ideas generated, and first time posting.

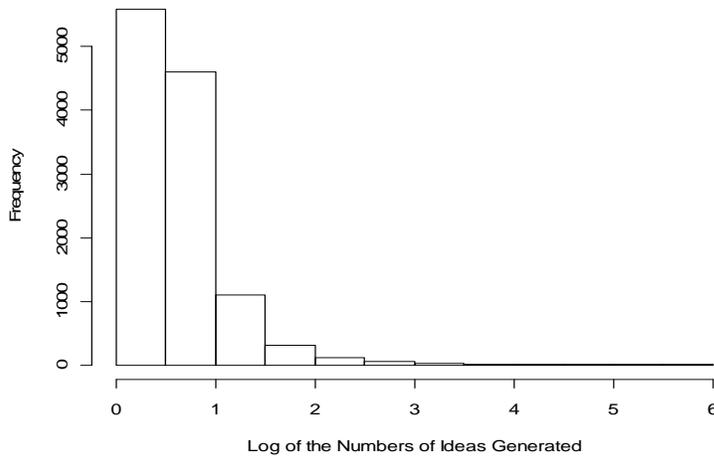


Figure 1. Histogram of the Number of Ideas Generated by an Individual

Dynamics in Idea Generation

The dynamics of individual participation on the crowdsourcing website are shown in Figures 2-4. From Figure 2, it is evident that the number of the ideas posted early on was very high; however, the number declined quickly over time and then stabilized. By comparing the weekly idea generation activities of the full sample and the 490 individual subsamples, we can conclude that the behavior of the individuals in the subsample is also similar.

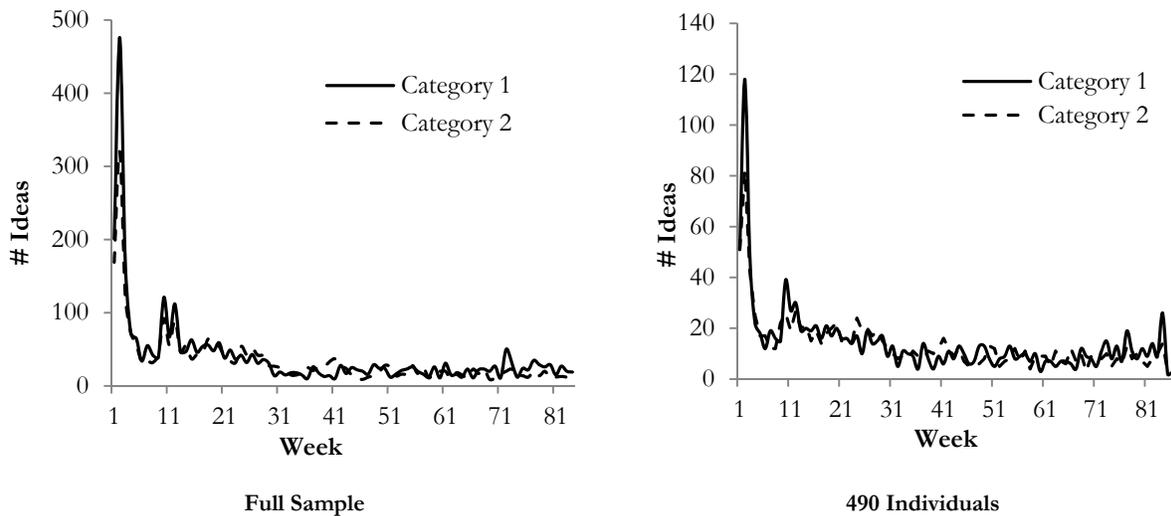


Figure 2. Numbers of Ideas Generated in Each Week

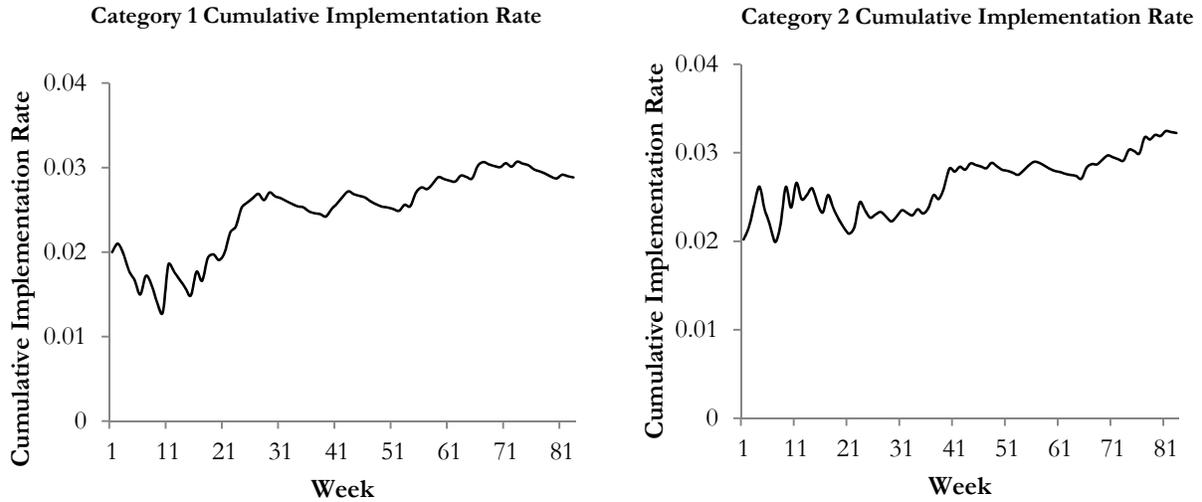


Figure 3. Cumulative Implementation Rate

In addition, if we look at the implementation rates of different categories of ideas (Figure 3), we note that the implementation rates of both Category 1 and Category 2 ideas increase over time. In Figure 4, we note that despite some random disturbance, the weekly average log votes tend to increase over time. The data patterns shown in Figures 3 and 4 suggest that although the number of ideas generated decreases over time, the quality/potential of the ideas seems to increase. These preliminary data analyses indicate that the decline in the number of ideas posted over time is only one side of the story and that the criticisms of ideation websites similar to IdeaStorm.com must be examined carefully.

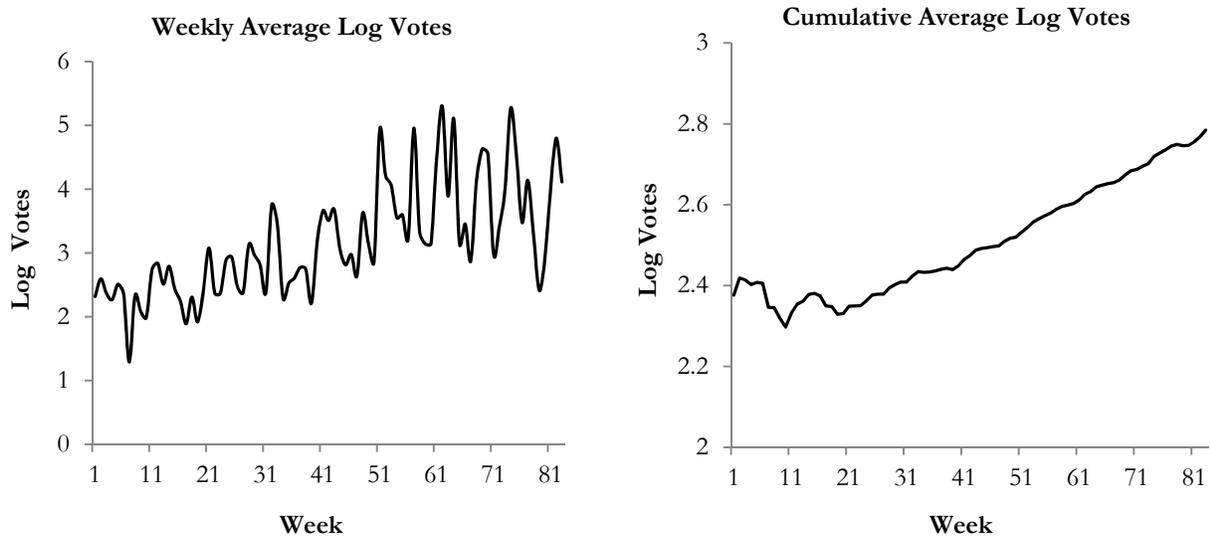


Figure 4. Weekly Average Log Votes and Cumulative Average Log Votes

4. Model

In this section, we develop a structural model to understand the dynamics of the participation behavior of individuals. The model incorporates learning about a firm's cost structure and the potential of their ideas as well as the dissatisfaction that was mentioned quite frequently by the participants.

Utility Function

The benefits an individual derives from posting an idea include the expected benefit received when the idea is implemented or needs are satisfied (Franke and von Hippel, 2003, Kuan 2001, Lakhani and von Hippel 2003). The reputation gain is indicated by the 500 Dell points. These points are shown in the individual's profile, but they cannot be cashed in or used for discounts and thus have no monetary value. However, online reputations may translate into a number of benefits, including job offers by established companies (Kumar et al. 2011, Huang et al 2010). Specifically, an individual's utility function (U_{ijt}) is given by the following equation

$$U_{ijt} = \begin{cases} c_i + r_i + d_i D_{it} + \theta_{ij} + \varepsilon_{ijt} & \text{if the idea is implemented} \\ c_i + r_i + d_i D_{it} + \varepsilon_{ijt} & \text{if the idea is not implemented} \end{cases}$$

where j represents the idea category. We adopt the classification on the website and set idea categories as Product ideas (Category 1) and Dell ideas (Category 2). The parameter, c_i , represents the cost incurred by individual i when he/she posts an idea, and r_i is the reputation gain the individual receives from the 500 IdeaStorm points. The parameter, θ_{ij} , measures individual i 's additional utility from the implementation of his/her Category j idea. The error term, ε_{ijt} , captures the individual choice specific random shock in period t . D_{it} is a binary variable with 1 meaning individual i is disincentivized in period t , and 0 otherwise. d_i denotes the extent to which such disincentive adds to individual i 's cost to post an idea or how it harms the utility the individual receives from posting an idea.

It is obvious that we cannot identify c_i and r_i simultaneously because they enter linearly in the utility function. Therefore, we combine these two terms and define $\theta_{i0} = c_i + r_i$. Thus, the individual's utility function reduces to

$$U_{ijt} = \begin{cases} \theta_{i0} + d_i D_{it} + \theta_{ij} + \varepsilon_{ijt} & \text{if the idea is implemented} \\ \theta_{i0} + d_i D_{it} + \varepsilon_{ijt} & \text{if the idea is not implemented} \end{cases} \quad (1)$$

where θ_{i0} is individual specific. In each period, individuals make decisions on whether or not to post ideas in a category. Before they post their ideas, they do not know whether the idea will be implemented. However, they form an expectation on the probability of their ideas being implemented. Let $E(U_{ijt} | Info(t))$ denote

the expected utility individual i can obtain from posting Category j idea in period t , conditional on the information individual i has up to period t . $E(U_{ijt}|Info(t))$ then can be expressed as

$$E(U_{ijt}|Info(t)) = \tilde{U}_{ijt} + \varepsilon_{ijt} = \theta_0 + d_i D_{it} + \theta_{ij} P_{ijt}|Info(t) + \varepsilon_{ijt} \quad (2)$$

where $P_{ijt}|Info(t)$ represents the conditional probability of implementation. We further elaborate $d_i D_{it}$ and $\theta_{ij} P_{ijt}|Info(t)$ in the following subsections.

Individual's Disincentive

One of the major criticisms about crowdsourcing websites is that the firms usually respond rather slowly to ideas of contributors. Currently, among all ideas posted, the status of 93.4% is only “Acknowledged”, while 2.2% ideas are “Implemented”.

To capture this characteristic, we introduce a new individual-level variable “disincentive” (D_{it}), which is a binary variable that equals 1 as long as there is one idea posted by an individual that has not moved to any status other than “Acknowledged” 12 weeks after it was originally posted. We chose twelve weeks as the criterion because the vast majority of the ideas that eventually moved to the next level in our dataset received the first status change within twelve weeks. If one sees that his/her ideas remain at “Acknowledged” status for more than twelve weeks, he/she may assume that this idea has little chance to be seriously reviewed, and the firm will likely not provide any feedback on the idea.

The selection of the cutoff point is subjective. We also use other time points as cutoff points, but the nature of the estimation results remains unchanged with different cutoff points, and only the magnitude of the estimates slightly changes. We defer these details to the section where we discuss the robustness of our findings. The effect of disincentive is denoted as d_i . We allow d_i to be different across individuals, i.e., some individuals could feel extremely disincentivized under situations where $D_{it} = 1$, while others may not share that feeling.

Individual's Learning Process

Idea contribution decisions of individuals are based on their beliefs of the probability of implementation; thus, $P_{ijt}|I(t)$ is updated from two types of learning. The first type of learning is learning by individuals about the firm’s cost structure, and the second type of learning is learning the potential of their own ideas, referred to herein as “their type”. Individuals can learn about the two components of their utility function from their experiences and observations in a Bayesian manner (Erdem and Keane, 1996).

Learning about the Firm's Cost Structure

Suppose that, at the moment when the website is launched, an individual's prior belief of the firm's cost of implementing a Category j idea is

$$C_{j0} \sim N(C_0, \sigma_{C_0}^2) \quad (3)$$

In Equation (3), C_0 is the prior mean of the cost of implementing an idea in Category j , and $\sigma_{C_0}^2$ measures prior belief about the variation of the cost associated with the implementation of different ideas within Category j . Individuals learn the firm's cost structure by observing the implementation of contributed ideas, including their own ideas and their peers' ideas. Whenever one idea is implemented, all individuals receive a common signal about the cost the firm incurs. Note that this learning process is common across individuals. C_{kjt} in Equation (4) denotes the cost signal all individuals receive when one Category j idea is implemented in period t . The difference between each specific cost signal and the mean implementation cost of ideas in the same Category is captured through the parameter μ_{kjt} , which is a zero mean normally distributed random variable, and its variance, σ_μ^2 , measures the variance of the implementation cost signals of ideas within the same category.

$$C_{kjt} = C_j + \mu_{kjt} \quad (4)$$

$$\mu_{kjt} \sim N(0, \sigma_\mu^2)$$

If there are k_{Cjt} Category j ideas implemented in period t , then the cumulative signal that individuals receive is C_{sjt} . C_{sjt} is simply the average of the k_{Cjt} signals $(C_{1jt}, \dots, C_{k_{Cjt}jt})$, and it has the following distribution

$$C_{sjt} \sim N\left(C_j, \frac{\sigma_\mu^2}{k_{Cjt}}\right) \quad (5)$$

Let C_{jt}^e denote individual's prior mean of Category j idea's implementation cost in the beginning of period t . By definition, conditional on the cumulative information he/she has received by the beginning of period t , individuals update C_{jt}^e using the following Bayesian rule (DeGroot, 1970)

$$C_{jt}^e = C_{jt-1}^e + (C_{sjt} - C_{jt-1}^e) \frac{\sigma_{C_{jt-1}}^2}{\sigma_{C_{jt-1}}^2 + \frac{\sigma_\mu^2}{k_{Cjt}}} \quad (6)$$

$$\sigma_{C_{jt}}^2 = \frac{1}{\frac{1}{\sigma_{C_{jt-1}}^2} + \frac{k_{Cjt}}{\sigma_\mu^2}} \quad (7)$$

The prior in period $t=0$ is $C_{j0}^e = C_0$, $\sigma_{C_{j0}}^2 = \sigma_{C_0}^2$.

Learning about the Potential of One's Own Ideas (Individual Type)

IdeaStorm.com allows individuals to vote on their peers' ideas, and the voting score is used as a measure of the potential of the ideas. A high voting score means that many customers would like to see this idea implemented, while a low voting score means the idea is probably a niche or limited idea that is favored by few. We model the average potential of ideas to be different across individuals but common across categories. Thus, individuals are heterogeneous with respect to their ability to generate ideas of high potential (good or marginal idea contributors). However, we model individual type to be idiosyncratic across different categories of ideas and invariant over time. These assumptions are required for the model identification. We validate the first assumption by regressing the logarithm of voting scores for ideas on categories with individual fixed effects. The test proves that the difference in the votes received by the two categories of ideas posted by the same individual is not significant (p-value = 0.926). We test the second assumption by regressing the logarithm of voting scores for ideas on week, on cumulative number of ideas posted, and on both, all with individual-level fixed effects. We find that none of the three models are statistically significant (F-test p-values for test 1, 2 and 3 are 0.365, 0.970 and 0.491, respectively). This test rules out the possibility that the individual's type evolves over time.

Let Q_i denote the mean potential of ideas posted by individual i ; then Q_{sit} , the potential of an idea posted by individual i in period t , is

$$Q_{kit} = Q_i + \delta_{kit} \quad (8)$$

$$\delta_{kit} \sim N(0, \sigma_{\delta_i}^2)$$

where δ_{kit} is the deviation of the potential of a specific idea posted by individual i in period t from the average potential of his/her idea. The variance of δ_{kit} is individual specific, which means that individuals learn about the potential of their ideas at different rates over time. Note that individuals learn their potential by observing the voting scores their ideas receive. We assume that the natural logarithm of the votes (V_{sit}) that an idea receives is linearly correlated with the potential of the idea

$$V_{kit} = cons + \varphi Q_{kit} \quad (9)$$

Thus, by plugging (8) into (9) we obtain

$$V_{kit} = V_i + \xi_{kit} \quad (10)$$

$$V_i = cons + \varphi Q_i$$

$$\xi_{kit} = \varphi \delta_{kit}$$

$$\xi_{kit} \sim N(0, \sigma_{\xi_i}^2)$$

$$\sigma_{\xi_i}^2 = \varphi^2 \sigma_{\delta_i}^2$$

where V_i is the mean value of the logarithm of votes that individual i 's idea receives and ξ_{kit} is its random shock. Again, if individual i posts k_{Qit} ideas in period t , then the cumulative signal that he/she receives is Q_{sit} . Q_{sit} is the average of the k_{Qit} signals ($Q_{1it}, \dots, Q_{k_{Qit}it}$), and it has the following distribution

$$Q_{sit} \sim N(Q_i, \frac{\sigma_{\delta_i}^2}{k_{Qit}}) \quad (11)$$

When the website was launched, individuals' prior beliefs of the potential of their ideas and the log voting scores their ideas receive were

$$Q_{i0} \sim N(Q_0, \sigma_{Q_0}^2) \quad (12)$$

$$V_{i0} \sim N(\varphi Q_0, \varphi^2 \sigma_{Q_0}^2) \quad (13)$$

Similar to the learning process of the firm's cost structure, individuals update their beliefs about V_{it}^e and Q_{it}^e together when they post an idea and observe the voting scores their ideas receive. The updating rules for V_{it}^e and Q_{it}^e are (Erdem, Keane and Sun, 2008)

$$V_{it}^e = V_{it-1}^e + (V_{sit} - V_{it-1}^e) \frac{\sigma_{V_{it-1}}^2}{\sigma_{V_{it-1}}^2 + \frac{\sigma_{\xi_i}^2}{k_{Qit}}} \quad (14)$$

$$Q_{it}^e = Q_{it-1}^e + (V_{sit} - V_{it-1}^e) \frac{\varphi \sigma_{Q_{it-1}}^2}{\varphi^2 \sigma_{Q_{it-1}}^2 + \frac{\sigma_{\xi_i}^2}{k_{Qit}}} \quad (15)$$

where

$$\sigma_{V_{it}}^2 = \frac{1}{\frac{1}{\sigma_{V_{it-1}}^2} + \frac{k_{Qit}}{\sigma_{\xi_i}^2}} \quad (16)$$

$$\sigma_{Q_{it}}^2 = \frac{1}{\frac{1}{\sigma_{Q_{it-1}}^2} + \frac{k_{Q_{it}}\varphi^2}{\sigma_{\xi_i}^2}} \quad (17)$$

In addition, we denote the priors for potential and for log-votes at the moment that the website was launched to be $Q_{i0}^e = Q_0$, $\sigma_{Q_{i0}}^2 = \sigma_{Q_0}^2$, and $V_{i0}^e = \varphi Q_0$, $\sigma_{V_{i0}}^2 = \varphi^2 \sigma_{Q_0}^2$.

Firm's Decision Rule to Implement Ideas

The firm selectively implements ideas generated by individuals. In general, the firm will consider the potential (market demand) of the ideas as well as the costs of implementing the ideas. Assume that a firm only implements ideas that provide a positive net profit. The net profit the firm generated from implementing the m^{th} Category j idea posted in period t can be expressed as

$$\pi_{mjt} = Q_{mjt} + C_{mjt}$$

where Q_{mjt} represents the potential of the idea and C_{mjt} represents the firm's cost associated with implementing the idea. Then, the probability that an idea will be implemented is

$$P_{mjt} = Pr(\pi_{mjt} > 0)$$

At the point that the firm makes implementation decisions, C_{mjt} is observed only by the firm, and not by consumers or researchers. However, the exact Q_{mjt} is observed by all three parties, given cons and φ . Therefore, for researchers, the likelihood that an idea with observed potential Q_{mjt} is eventually implemented is

$$P_{mjt} = Pr(Q_{mjt} + C_{mjt} > 0 | Q_{mjt}) = 1 - \Phi\left(\frac{Q_{mjt} + C_j}{\sigma_\gamma}\right) \quad (18)$$

where σ_γ represents the true standard deviation of the cost for the firm to implement ideas in the same category. Let I_{mjt} denote the decision the firm makes on the m^{th} Category j idea posted in period t , with value 1 indicating that the idea is implemented and 0 otherwise. The likelihood that we observe I_{mjt} given Q_{mjt} , C_j and σ_γ is

$$L(I_{mjt}) = \Phi\left(\frac{Q_{mjt} + C_j}{\sigma_\gamma}\right)^{I_{mjt-1}} (1 - \Phi\left(\frac{Q_{mjt} + C_j}{\sigma_\gamma}\right))^{1-I_{mjt}}$$

Individual's Decision Making Problem

As previously mentioned, individuals make decisions on whether to post an idea in a Category based on their expectation of the utility they can possibly derive from each choice. We assume individuals' decisions on idea posting are independent of each other across categories and that the individuals are aware that the firm makes implementation decisions by comparing the potential of the ideas and the implementation costs to the firm. Then, the \tilde{U}_{ijt} in Equation (2) can be expressed as

$$\tilde{U}_{ijt} = \theta_0 + d_i D_{it} + \theta_{ij} Pr(\pi_{ijt} | Info(t) > 0) \quad (19)$$

where $\pi_{ijt} | Info(t) | I(t) \sim N(Q_{it}^e + C_{jt}^e, \sigma_{Q_{it}}^2 + \sigma_{C_{jt}}^2 + \sigma_{\delta_i}^2 + \sigma_{\gamma}^2)^4$. We assume further that ε_{ijt} follows a Type 1 extreme value distribution and that the probability that individual i will post a Category j idea in period t takes the standard logit form. In this case, the likelihood of observing action A_{ijt} can be expressed as

$$L(A_{ijt}) = \left(\frac{\exp(\tilde{U}_{ijt})}{1 + \exp(\tilde{U}_{ijt})} \right)^{A_{ijt}} \left(\frac{1}{1 + \exp(\tilde{U}_{ijt})} \right)^{(1 - A_{ijt})} \quad (20)$$

5. Estimation

In the literature, most of the Bayesian learning models are estimated by (simulated) maximum likelihood estimation methods. However, in our case, due to the individual-specific Q_i , $\sigma_{\delta_i}^2$ and θ_{ij} , the frequentist estimation methods are inconvenient. Following Netzer et al. (2008) and Narayanan and Manchanda (2009), we apply Markov Chain Monte Carlo (MCMC) methods to estimate the individual-specific parameters. We use the Gibbs sampler to recursively make draws from the following conditional distribution of the model parameters. For complete details, see the Appendix.

Model Hierarchy

The parameters in our model are summarized in Table 4. For identification purposes, θ_{i0} , C_1 , $\sigma_{C_0}^2$, and $\sigma_{Q_0}^2$ are fixed (model identification will be briefly discussed in later sections and elaborated in Appendix 2). Among the remaining parameters, parameter vector $\beta = [C_0, C_2, \sigma_{\gamma}^2, \sigma_{\mu}^2, Q_0, cons, \varphi]$ is common across individuals, while parameter vector $\beta_i = [Q_i, \log(\sigma_{\delta_i}^2), d_i, \theta_{i1}, \theta_{i2}]$ is heterogeneous across individuals. We further assume that β_i follows the following distribution

⁴ One concern may be that individuals may not know the actual variation in firm's costs to implement ideas in one Category (σ_{γ}^2). We try other specifications and find that dropping σ_{γ}^2 or dropping both σ_{γ}^2 and $\sigma_{\delta_i}^2$ has little effect on the estimation results.

$$\beta_i = \begin{pmatrix} Q_i \\ \log(\sigma_{\delta_i}^2) \\ d_i \\ \theta_{i1} \\ \theta_{i2} \end{pmatrix} \sim MVN(\bar{\beta}, \Sigma)$$

where $\bar{\beta}$ denotes the mean of β and Σ denotes the variance and covariance matrix of β .

Table 4. Summary of the Parameters in the Model

Notation	Explanation
θ_{i0}	Cost for individual i to post an idea. (Fixed to -5).
θ_{ij}	Payoffs individual i receives when his/her Category j ideas are implemented
d_i	Level of disincentive individual i receives when the status of one or more of i 's ideas stays as "Acknowledged" for more than 12 weeks
D_{it}	Binary variable that takes a value of 1 when there is at least one idea posted by individual i that has not moved to any status other than "Acknowledged" for more than 12 weeks after it is originally posted in period t .
C_0	Individual's initial prior mean costs for implementing each Category of ideas
$\sigma_{C_0}^2$	Individual's initial prior variance of the costs for implementing each Category of ideas (set to 50, assume the prior is uninformative)
C_j	The firm's mean cost for implementing Category j ideas (the mean cost for Category 1 is fixed at -6)
σ_{γ}^2	The variance of true distribution of the costs for the firm to implement ideas in the same Category
σ_{μ}^2	Variance of cost signal
Q_0	Individuals' initial prior mean of the potential of their ideas
$\sigma_{Q_0}^2$	Individuals' initial prior variance of the potential their ideas (set to 50, assume prior is uninformative)
Q_i	Mean potential of ideas generated by individual i
$\sigma_{\delta_i}^2$	Variability of potential of ideas generated by individual i
$cons$	Intercept of linear function between log votes and the potential
φ	Slope coefficient between log votes and potential

Conditional on $cons$, φ and $\sigma_{\delta_i}^2$, the updating process of the potentials of individuals' ideas is deterministic because we explicitly observe the potential signal (votes). The updating process of the variance of cost belief is also deterministic, given σ_{μ}^2 . Only the updating process of C_{jt}^e is stochastic. Following Narayanan and Manchanda (2009), the distribution of C_{jt+1}^e , conditional on C_{jt}^e , can be expressed as

$$C_{jt+1}^e | C_{jt}^e \sim N(\bar{C}_{jt+1}^e, v_{jt+1}^2)$$

where

$$\bar{C}_{jt+1}^e = \frac{\sigma_{C_{jt+1}}^2}{\sigma_{C_{jt}}^2} \bar{C}_{jt}^e + k_{C_{jt}} \frac{\sigma_{C_{jt+1}}^2}{\sigma_{\mu}^2} C_j \quad (21)$$

$$v_{jt+1}^2 = k_{Cjt} \frac{\sigma_{C_{jt+1}}^4}{\sigma_{\mu}^2} \quad (22)$$

Therefore, the unobserved cost belief can be drawn from the following natural hierarchy

$$\begin{aligned} C_{jt}^e | C_{jt-1}^e &\sim N(\bar{C}_{jt}^e, v_{jt}^2) \\ C_{jt-1}^e | C_{jt-2}^e &\sim N(\bar{C}_{jt-1}^e, v_{jt-1}^2) \\ &\dots \\ C_{j1}^e | C_{j0}^e &\sim N(\bar{C}_{j1}^e, v_{j1}^2) \end{aligned}$$

The full hierarchical model can be specified as

$$\begin{aligned} A_{ijt} | C_{jt}^e, \sigma_{C_{jt}^e}^2, \sigma_{\gamma}^2, \theta_{i1}, \theta_{i2}, V_{si}, cons, \varphi, Q_0, \sigma_{\delta_i}^2, d_i, D_{it}; I_{mjt} | C_j, \sigma_{\gamma}^2, V_{mjt}, cons, \varphi, \\ C_{jt}^e | C_{jt-1}^e, C_j, \sigma_{\mu}^2, k_{Cjt} \\ \beta_i | \bar{\beta}, \Sigma \end{aligned}$$

where the additional notation V_{si} denotes a vector of the log voting scores that all ideas generated by i receive in all periods.

Identification

We now briefly discuss some intuition as to how the parameters in our model are identified (see Appendix 2 for full discussion). We cannot identify all Q_i and C_j at the same time because if we add a constant to all Q_i 's and then subtract the same constant from all C_j 's, we will obtain exactly the same likelihood value. For identification purposes, we fixed $C_1 = -6$. As a result, the estimated values of C_2 and Q_i should be interpreted as relative to C_1 . We further fix θ_0 to -5 to achieve the identification of other parameters of interest. The parameter θ_{ij} should then be interpreted as the value individual i derives from the implementation of his Category j idea relative to the cost he/she incurs when posting the idea. d_i can be easily identified because D_{it} is observed for every i in every period. As the potential for the ideas posted by the same person is the same, the systematic difference in the number of ideas in each Category generated before the first implementation of ideas yields the difference between θ_{i1} and θ_{i2} , at least for the individuals who post before Week 7.

Ideally, after controlling for θ_{i1} and θ_{i2} , Q_i can be identified from the behavior of “well-informed” individuals, whose perception about the firm’s cost structure and their potential of ideas is very close to the true value. Because C_1 is fixed, Q_i should be easily identified. However, in our case, we cannot guarantee that at the end of the observation period, the perception of the contributor about the potential of his/her ideas and the implementation cost for the firm are closest to the true values. To solve this problem, we use the implementation of the ideas given the votes as another source of identification of Q_i and C_2 . Combining these two constraints, Q_i and C_2 can be identified. Once C_1 , C_2 and Q_i are identified, σ_r^2 can be easily identified from the likelihood of one idea being implemented. C_0 can be identified through the probability of posting before the first cost signal is broadcast to the consumers. Given C_j and ρ , ρ can be identified through the probability of posting for the latecomers throughout the whole observation period. Before one individual posts any ideas on the website for the first time, his/her belief about the potential of his/her ideas is always ρ , while his/her beliefs about the implementation cost is updated. Given the different ρ for different t s, ρ can then be identified. ρ and φ are identified from the relationship between ρ and φ . As ρ and φ are pooled parameters, given all individuals’ ρ and observed φ , ρ and φ determine a linear curve that best approximates the relationship between ρ and φ .

Finally, the variance parameters σ_r^2 and $\sigma_{\delta_i}^2$ are both identified from the dynamics of the posting behavior of individuals over time. We are able to identify σ_r^2 and $\sigma_{\delta_i}^2$ simultaneously because the signals of the implementation costs and potentials are generated from different events. σ_r^2 is identified through the dynamics of the choice probabilities at the population level. Similarly, the learning speed of the potential of the ideas is affected by both $\sigma_{\delta_i}^2$ and the slope parameter φ . φ affects every individual in the same way, while $\sigma_{\delta_i}^2$ affects individuals differently. Therefore, the different learning speeds of different individuals observed in the data can help us identify $\sigma_{\delta_i}^2$.

6. Estimation Results

The estimates of the pooled parameters are presented in Table 5. Comparing the estimate of C_2 with C_1 (fixed to -6), we see that C_2 is smaller in terms of absolute value. Thus, the cost that the firm incurs when implementing Category 2 ideas is lower than the cost of implementing Category 1 ideas, which is consistent with the higher implementation rate of Category 2 ideas as compared to Category 1 ideas. The estimate for C_0 is significantly higher than both C_1 and C_2 , indicating that individuals initially tend to underestimate the implementation costs incurred by the firm. It is consistent with the observation that the numbers of ideas created decrease over time.

Table 5. Pooled Parameter Estimates

Notation	Parameter Estimates	Standard Deviation
C_0	-1.178	0.159
$\sigma_{C_0}^2$	50	-- (Fixed)
C_1	-6	-- (Fixed)
C_2	-5.813	0.081
$\log(\sigma_\mu^2)$	6.337	0.139
$\log(\sigma_r^2)$	1.086	0.069
Q_0	3.359	0.265
$\sigma_{Q_0}^2$	50	-- (Fixed)
<i>cons</i>	-0.427	0.073
φ	2.170	0.033

The estimate of $\log(\sigma_\mu^2)$ is 6.337, which is equivalent to saying that σ_μ^2 is $\exp(6.337)=565$. This variance is quite large compared to the absolute values of C_1 and C_2 , indicating that the implementation cost signals the firm provides to individuals are imprecise and consequently, that individuals cannot learn quickly about the implementation costs to the firm. Remember that $\exp(\sigma_\mu^2)$ is the variance of one signal and that there are cases where several ideas are implemented within a week. In those weeks, the variance of the cumulative signal individuals receive will be $\exp(\sigma_\mu^2)$ divided by the number of ideas implemented in each week; thus, the learning regarding the implementation could still be significant.

Relative to the variance of the signal, the estimate of $\log(\sigma_r^2)$ is much smaller (1.086); that is, $\sigma_r^2 = 2.962$. Q_0 is also higher than the average level of Q_i , indicating that most of the individuals overestimated the potential of their ideas before their ideas were voted on. *cons* and φ determine the linear relationship between log votes and potential. The slope coefficient is 2.170, meaning that when the potential of the idea increases by 1, the log of one idea's vote increases by 2.170.

The estimation results of the mean and standard deviation of individual-level parameters are summarized in. Additionally, histograms of the distribution of the 5 individual-level parameters are shown in. We can see that the population average of the potential of the ideas is significantly lower than the cost of implementing both categories of ideas. This is consistent with the low implementation rate we observe in the data. We also observe significant differences among individuals with respect to the ability to generate ideas with high potential (individual type).

The population average of variance of the potentials of ideas by one individual is very low, $\exp(\overline{\log(\sigma_\delta^2)})= 0.264$. This result suggests that the potentials for the ideas posted by the same person are

relatively consistent. High potential idea contributors consistently generate ideas with high potential, while marginal idea contributors rarely come up with high potential ideas. This variance also implies the learning speeds of individuals with respect to the potential of their ideas. The small average variance indicates that, on average, individuals learn quickly about the potential of their ideas. When the website was launched, many individuals, i.e., idea providers, entered the market. As they learn about their own potential of their ideas and the cost for the firm to implement their ideas, marginal idea contributors dropped out, and the “idea market” became more efficient in a short time. In other words, the crowdsourcing mechanism is quite effective in filtering idea providers, and the “idea market” reaches efficiency quickly. However, the standard deviation of $\sigma_{\delta_i}^2$ is relatively large (1.698), indicating that some individuals have better consistency in terms of the potential of their ideas, while others have a lower consistency.

The average level of the disincentive effect is -1.940, meaning that when individuals’ ideas are not responded to in a timely manner, individuals will be less likely to post ideas, and the average level of this effect is equivalent to increasing the cost of posting an idea by 2/5. Given the low overall probability of posting, the impact of such discouragement is quite large. The mean payoff individuals receive when their Category 1 ideas are implemented is slightly higher than when their Category 2 ideas are implemented. This is consistent with the numbers of ideas posted in these two categories during the first few weeks. This finding is also intuitive because ideas in Category 1 are more about product improvement, while ideas in Category 2 are related to customer services and marketing strategies. It is not surprising that individuals receive greater payoffs when the firm improves the product design according to an individual’s suggestion than when the firm improves services and communications with their customers, as suggested.

Table 6. Individual-level Parameter Estimates

Notation	Mean Among Individuals*	Standard Deviation Among Individuals *
Q_i	2.531	0.277
$\log(\sigma_{\delta_i}^2)$	-1.333	1.697
d_i	-1.940	1.732
θ_{i1}	3.186	1.244
θ_{i2}	3.093	1.434

* For each individual, the posterior distribution of each parameter has a mean and standard deviation. The mean and standard deviation reported here are the mean and standard deviation of the individual-level parameter means.

To explore the correlation between individual mean potential (Q_i) and other individual-level parameters, we draw the scatter plots of all other individual-level parameters against individual mean potentials, respectively (Figure 6). Interestingly, the correlation of Q_i and $\sigma_{\delta_i}^2$ is negative, indicating that the potential of ideas generated by good idea contributors are more consistent, and thus, these individuals tend to

learn more quickly about their ability. Another interesting finding is the correlation between Q_i and d_i . In other words, good idea contributors would not be as disappointed as marginal idea contributors. We explore the policy implications of this finding later.

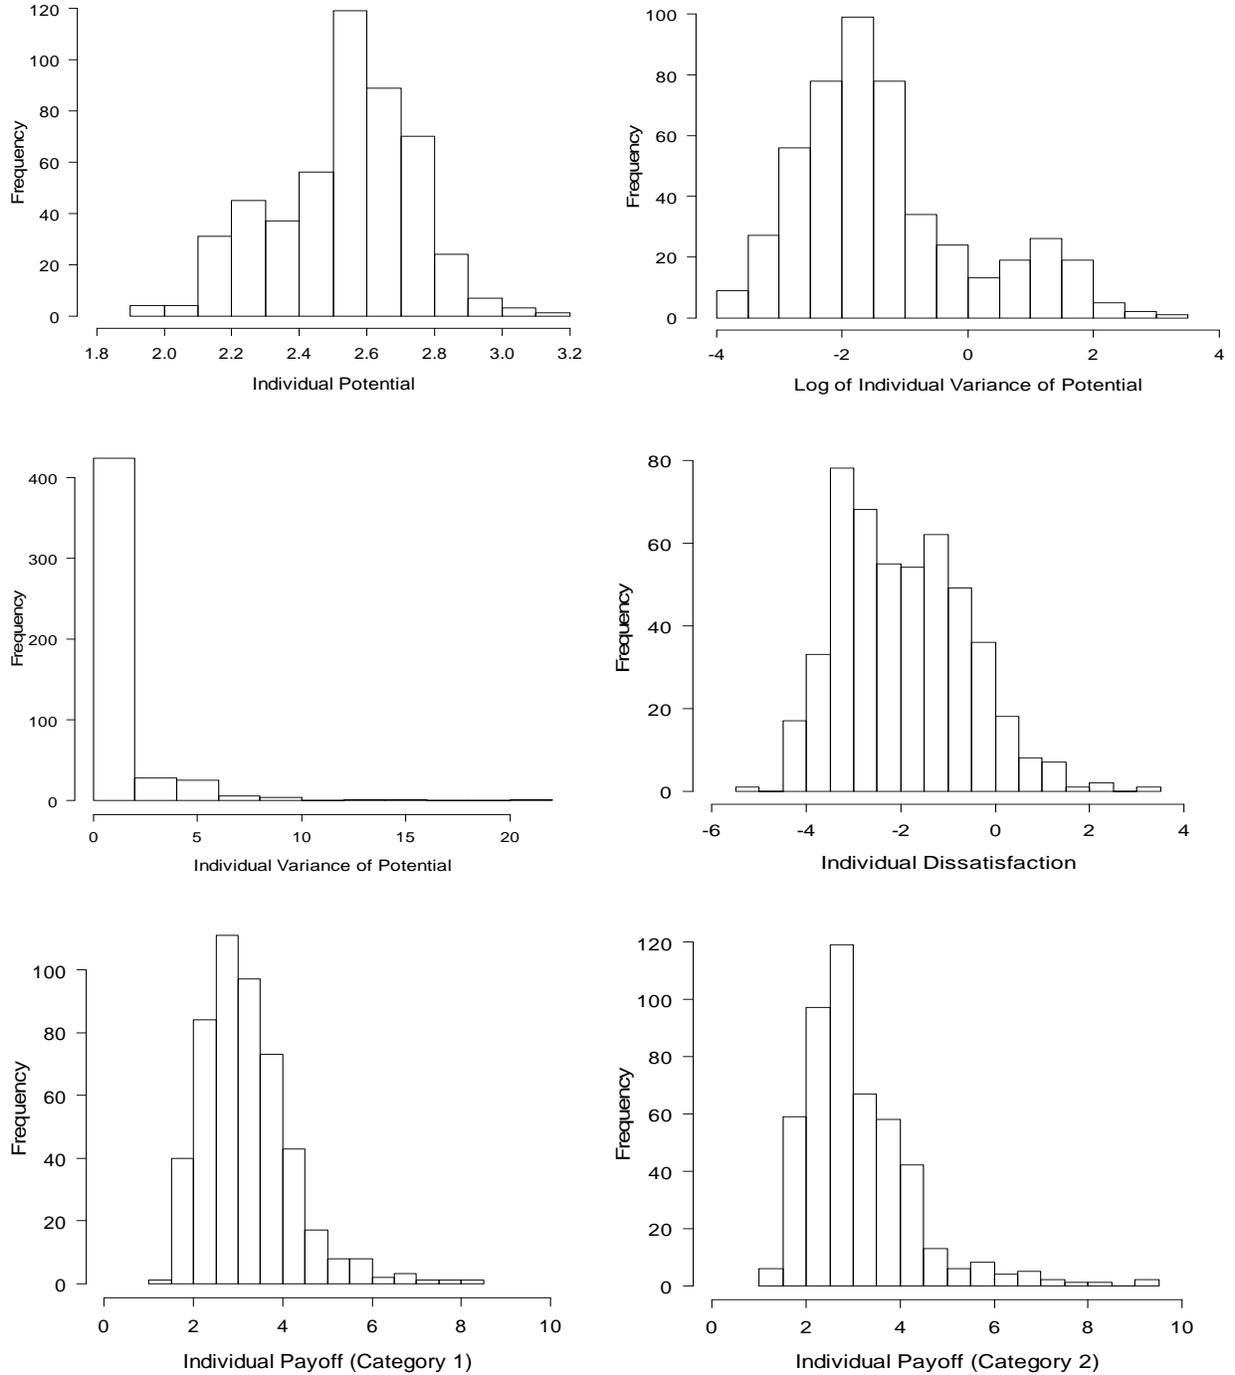


Figure 5. Distributions of Individual-level Parameters

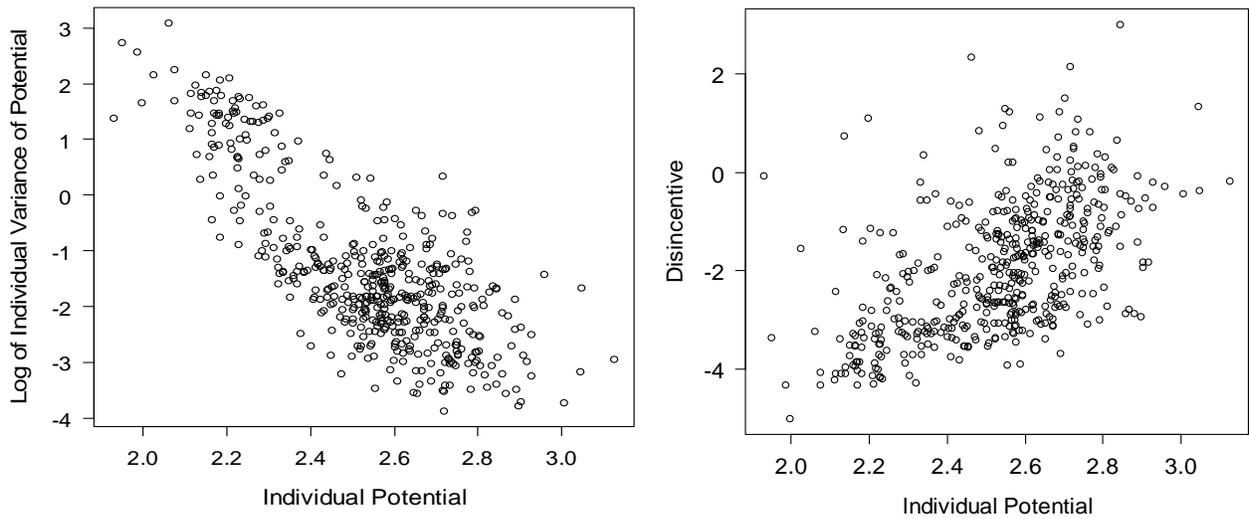


Figure 6. Scatter Plot of Individual-level Parameters

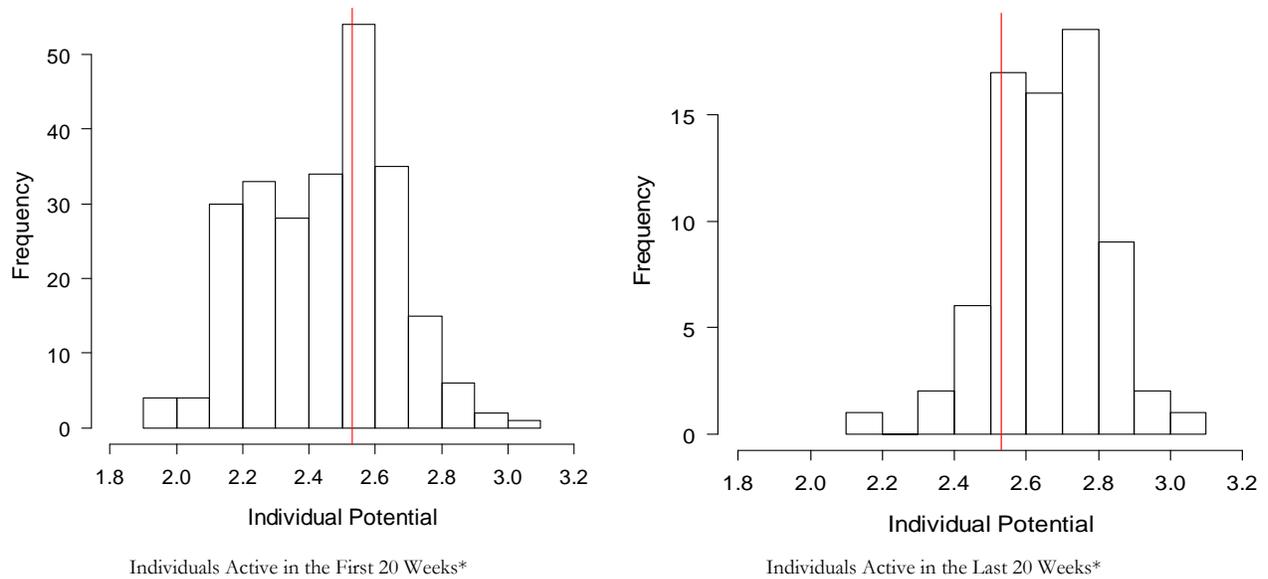


Figure 7. Comparison of the Mean Potential of Individuals Active in the First and Last 20 Weeks

*“Active individual” is defined as one who posts 2 or more ideas in the periods under discussion. Vertical lines in both plots are the mean potentials across individuals.

Our estimation process produces the posterior mean of an individual’s ability to generate ideas with good potential. This allows us to explicitly examine the filtering process of idea providers in the market. Figure 7 visualizes the comparison between the mean potential of individuals active in the first 20 weeks (“Active individual” is defined as those who post ≥ 2 ideas in the periods under discussion) and that of individuals active in the last 20 weeks. The vertical line in both plots is the average mean potential of the 490

individuals in our sample. From the two plots, it is evident that the distribution shifts toward the right. The majority of the individuals who are active in the last 20 periods are those who have been identified as good idea contributors.

Model Fit and Model Comparison

We compare our model with three null models, which include the random coefficient model (no learning), the cost learning only model, and the potential learning model. From Table 7, we note that our full model outperforms all other alternative models with respect to marginal likelihood and deviance information criterion (DIC). We also find that only the cost learning model slightly improves marginal likelihood when compared to the no learning model. However, the DIC is even larger than the no learning model. It appears that only including cost learning does not significantly improve model fit. This is because, on the one hand, cost learning is relatively slow and therefore has limited contribution to model fit, and on the other hand, in the cost learning only model, the learning dynamics of individuals are forced to be homogenous, which deviates from reality. Not surprisingly, we find that when we include learning about their own type, both marginal likelihood and DIC improve significantly. This suggests that learning about their own type explains a good amount of the individuals posting dynamics. By comparing the full model and the potential learning only model, we find that after we control for individual learning about personal type, adding cost learning will improve the performance of the model. This suggests that although the firm only provides imprecise cost signals and individuals learn slowly about the firm’s cost structure, the effect of cost learning still explains a significant degree of the remaining dynamics.

Table 7. Model Comparison

Model	Random Coefficient	Cost Learning Only	Potential Learning Only	Full Model
Log Marginal Likelihood	-8131.6	-8121.9	-7928.5	-7909.5
Difference in DIC (wrt. Full Model)*	305.2	309.4	25.0	0.0

* Difference in DIC=DIC of the model-DIC of the full model. Smaller DIC is preferred.

Robustness Checks

We perform several tests to check the robustness of our results. The comparisons among the different alternative models are shown in Tables 8 and 9. As the cutoff points we use to define an individual’s dissatisfaction are arbitrary, we employ different cutoff points to investigate how the estimation results will change. The first three columns in both Table 8 and Table 9 show that the parameter estimates, both the pooled estimates and the mean of the individual-level parameter estimates, are quite stable. Thus, the estimation results are not sensitive to the selection of the cutoff points.

Table 8. Pooled Parameter Estimates*

Parameter	Main Model (Cutoff=12)	Cutoff=8 weeks	Cutoff=16 weeks	Various Enter Time
C_0	-1.178 (0.159)	-0.919 (0.212)	-2.451 (0.147)	-3.057 (0.211)
C_2	-5.813 (0.081)	-5.810 (0.151)	-5.889 (0.086)	-5.759 (0.116)
$\log(\sigma_\mu^2)$	6.337 (0.139)	6.395 (0.164)	6.504 (0.128)	6.027 (0.217)
$\log(\sigma_r^2)$	1.086 (0.069)	1.235(0.083)	1.114 (0.068)	1.836 (0.072)
Q_0	3.359 (0.265)	2.916 (0.305)	3.608 (0.161)	12.450 (0.271)
<i>cons</i>	-0.427(0.073)	-0.288 (0.094)	-0.721 (0.175)	2.017(0.265)
φ	2.170 (0.033)	2.318 (0.043)	2.261 (0.080)	2.591(0.072)

*Standard deviations are in parentheses.

Table 9. Individual-level Parameter Estimates*

Parameter	Main Model (Cutoff=12 Weeks)	Cutoff=8 Weeks	Cutoff=16 Weeks	Various Enter Time
Q_i	2.531 (0.277)	2.306 (0.286)	2.564 (0.255)	1.080 (0.266)
$\log(\sigma_{\delta_i}^2)$	-1.333 (1.697)	-1.475 (1.754)	-1.455 (1.793)	-1.579 (1.670)
d_i	-1.940 (1.732)	-1.909 (1.657)	-2.221 (1.736)	-2.407(2.108)
θ_{i1}	3.186 (1.244)	3.427 (1.362)	3.807 (1.221)	5.561 (2.087)
θ_{i2}	3.093 (1.434)	3.324 (1.629)	3.626 (1.423)	5.256 (2.209)

*The mean and standard deviation (in parentheses) reported here are the mean and standard deviation of the individual-level parameter means.

Another potential concern about the model estimation is that, as we do not have information on when each individual joined the community, in the estimation, we assume that everyone joined in the first period. One may argue that, in reality, individuals may join at different time points. To address this concern, we employ another extreme case. From our dataset, we cannot determine the voting activities of individuals, but we do have information on when an individual commented on others' posts. We combine the comment data and idea posting data and find the first record, either idea posting or commenting, for each individual. Therefore, we set the time that the first record occurred to be the individual's registration time. Among the 490 individuals, only 74 commented before their first idea posting. We estimate our model again under this assumption, and the results are shown in the last column of Tables 7 and 8. The main difference between this set of results and the main model results is that the estimate of Q_0 is much larger. We also note that the estimates of payoff coefficients θ_1 and θ_2 significantly increase. These differences are not surprising because most individuals posted in the first period when they joined the community. Although the values of the parameter estimates are different, our main claims hold; that is, early on, individuals overestimate the potential of their ideas and underestimate the cost to the firm associated with the implementation of their ideas. Individuals continue to learn more quickly about the potential of their ideas than about the firm's cost structure. As before, the firm's failure to respond or its untimely response to idea postings negatively affects the continuing participation of contributors.

7. Policy Simulations

We conduct three sets of policy simulations to determine how a firm can improve the overall performance of its crowdsourcing website by accounting for the heterogeneous learning by individuals about the firm's cost structure and their ability to generate ideas with good potential. The simulation results are the average across 1000 simulation iterations. The average potential and the number of Category 1 ideas generated in each period are reported. The number of Category 2 ideas generated in each period has a similar pattern as those in Category 1.

Should the Firm Provide More Precise Information on their Cost Structure?

We simulate the evolution of the average potential of individual's posts over time and the numbers of the two categories of ideas contributed each week under different standard deviations of cost signal. As shown in Figure 8, if the firm can provide the cost signals with smaller standard deviations, the average potential of ideas will be significantly improved after 30 weeks. We also observe a significant decrease in the numbers of ideas in each Category posted each week, which can further help reduce the screening costs that the firm incurs. When the firm provides individuals with more precise cost signals, individuals learn more quickly about the implementation costs. As their initial belief is much lower than the absolute value of the true cost, a faster learning speed means that in each period, an individual's perceived costs will be greater in terms of absolute value. Therefore, their *ex ante* utility will decrease, and thus, individuals with lower individual mean potential will drop out sooner. In other words, by providing more detailed feedback about their implementation costs, the firm can improve the efficiency of the idea market. We visually show this impact in the graphs for the reduction in variance of the cost of implementation. It is worth noting that in this analysis we ignore the firm's incentive to be imprecise in signaling its implementation costs due to competitive reasons.

Should the Firm Respond to Individuals' Ideas More Quickly to Reduce Disincentive?

Our estimation results show that the firm's lack of or untimely response to idea posts negatively affects an individual's participation in idea contribution. To deal with this type of disincentive, the firm may attempt to improve the number of ideas to which it responds and to reduce the time between the posting of the idea and the firm's valuable feedback to the contributor. Although we do not know how Dell selects the ideas to which it replies, the extremely low response rate makes the effect of their selection strategy immaterial. Almost every individual is disincentivized in the latter part of the observation period. In this set of simulations, we examine the various policies that aim to reduce an individual's feeling of dissatisfaction and disincentiveness.

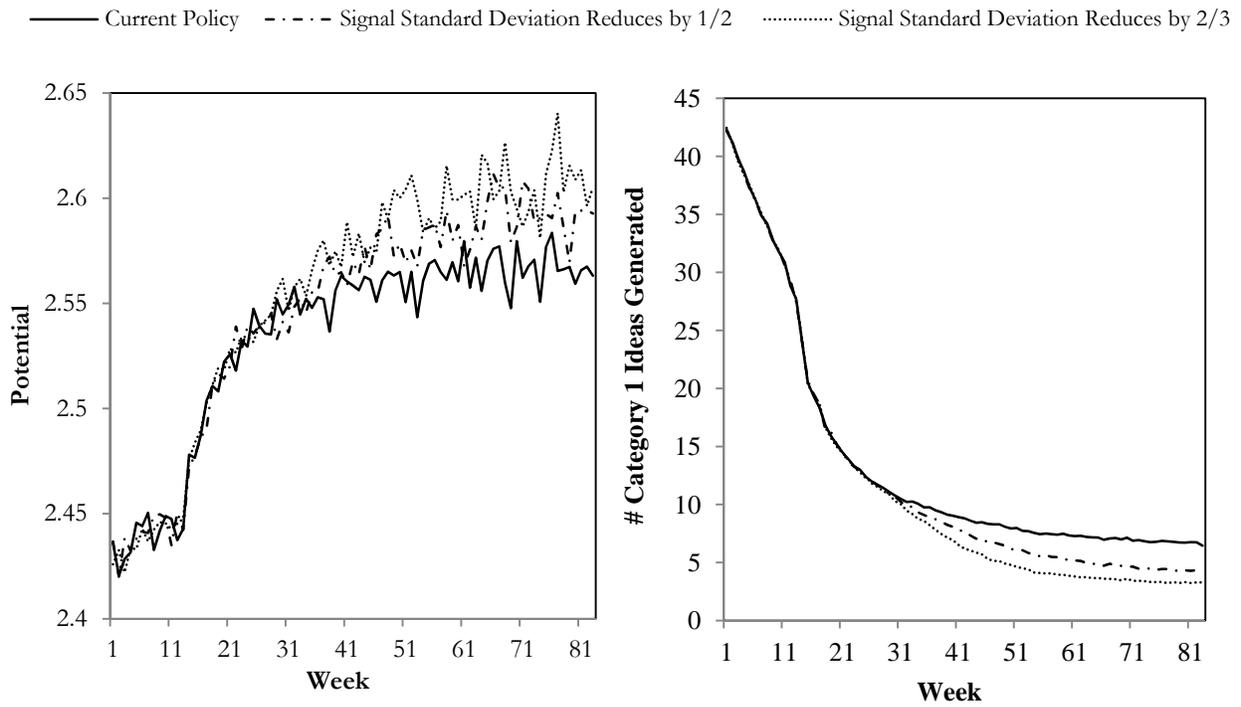


Figure 8. Simulation Results When the Firm Provides More Precise Signals

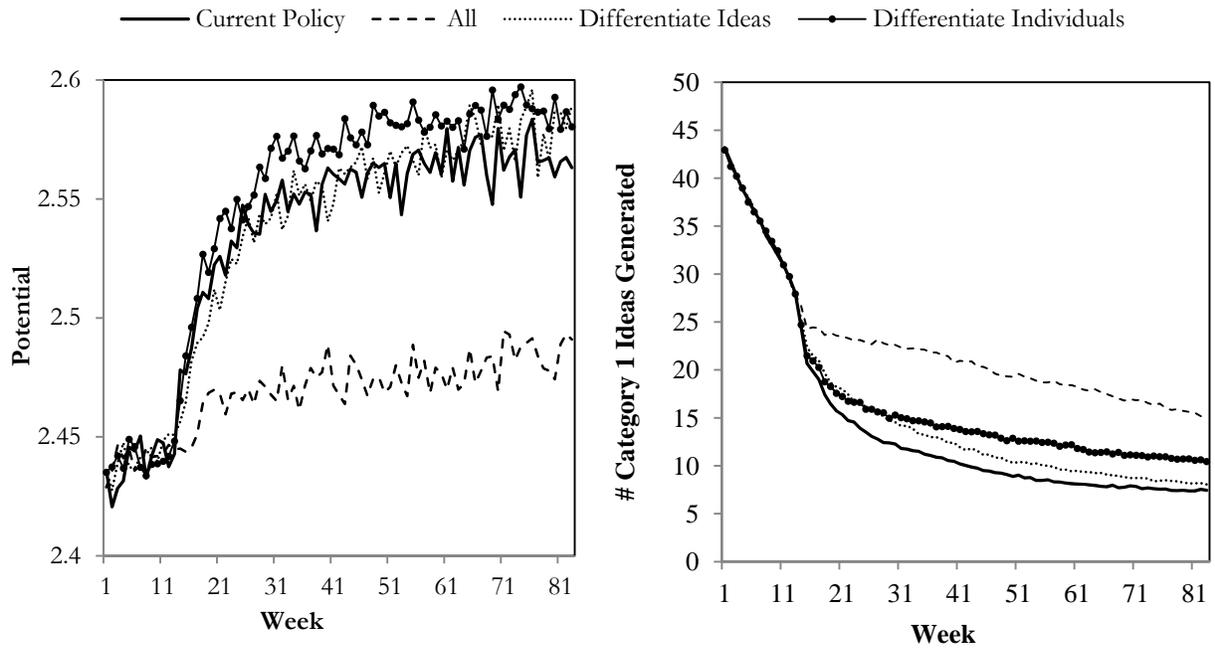


Figure 9. Simulation Results When the Firm Replies to More Ideas in a More Timely Manner

We experimented with three new policies. In Figure 9, “All” represents the policy under which the firm responds to all non-implemented ideas within 11 weeks. “Differentiate Ideas” represents the policy under which the firm only replies to ideas that have a potential above the average. “Differentiate Individuals” represents the policy under which the firm identifies good idea contributors and only replies to ideas generated by them. Not surprisingly, we find that under the three policies, the number of ideas generated increases. This is intuitive because all the policies can reduce the individual’s disincentive and thus encourage them to post ideas. On the contrary, the effects that the three policies have on the average potential of ideas posted over time are very different. Interestingly, we find that in the left plot, the curve labeled “All” is below the curve representing current policy everywhere, indicating that if the firm improves the response time and response rate, it completely removes the disincentive, and the firm is worse off because it receives more ideas with significantly lower potential. Therefore, the firm should strategically select the ideas to which it responds. It is easier to implement the “Differentiate Ideas” strategy because all the firm needs to do is to look at the votes and respond to the ideas for which the log of votes is above average. Furthermore, this strategy leads to the submission of only slightly more ideas, and it outperforms the “All” strategy in terms of the potential of the ideas.

When comparing the average potential in the “under the current policy” and the “Differentiate Ideas” strategy, we note that, in the beginning, the latter performs no better than the current policy. However, at approximately week 30, the “Different Ideas” strategy outperforms the current policy. The “Differentiate Individuals” strategy is obviously the best policy in terms of the potential of the ideas contributed by individuals, standing out immediately after 12 weeks. It also leads to more idea generation, especially in later periods. If we do not consider the screening cost, this is undoubtedly the best policy because the firm will have more high potential ideas from which to choose, and furthermore, it will profit from these good ideas. Nevertheless, if we take both the screening cost and the cost of more frequent responses into account, the effectiveness of this policy must be re-evaluated. Another caveat of this policy is that it is not easy to implement; both the firm and the individuals themselves must learn about the potential of the ideas.

Should Individuals be Rewarded for Posting or for Posting Ideas that are Implemented?

Two commonly observed reward structures used on this type of crowdsourcing website include giving a reward as long as an individual posts an idea (the 500 IdeaStorm points in our case) and giving a reward to contributors only when an idea is implemented (IdeaStorm is currently applying this reward structure). In this set of policy simulations, we aim to investigate which reward structure performs better.

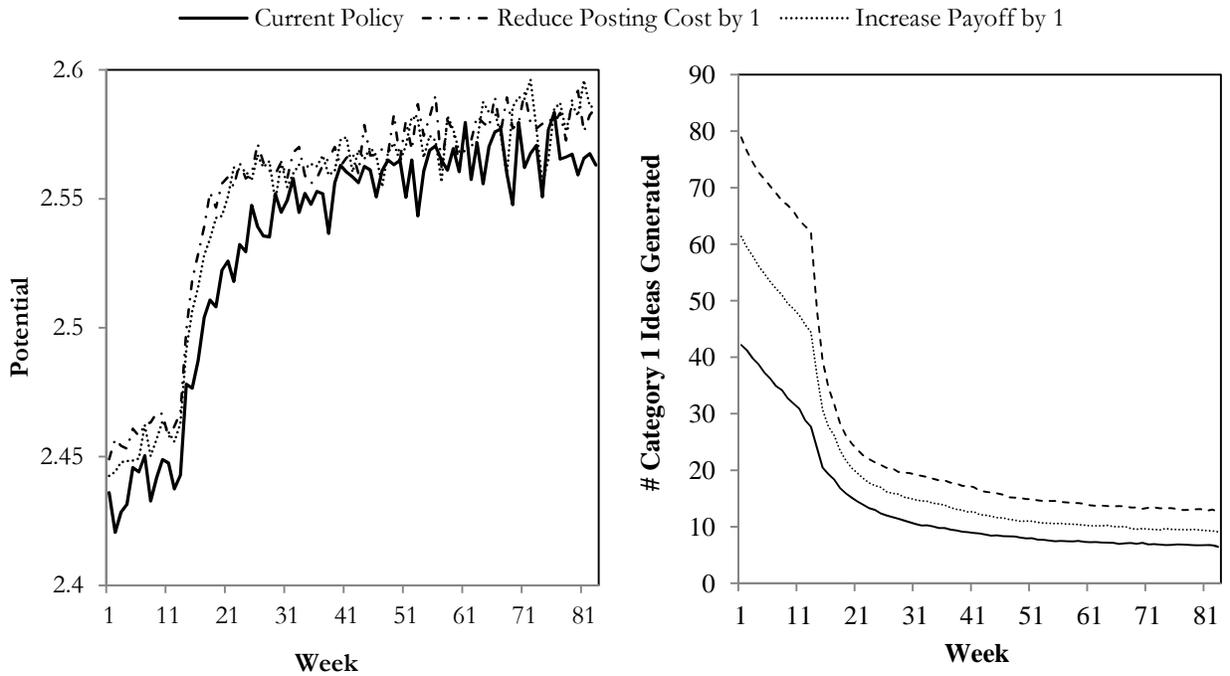


Figure 10. Simulation Results for Different Reward Structures

In Figure 10, “Reducing Posting Cost” represents the policy under which individuals are rewarded as long as they post an idea. This policy will add a positive constant to the utility function of individuals, thus reducing the cost of posting by the same constant. “Increase Payoff by 1” represents the policy under which individuals are rewarded only when their ideas are implemented. In our model, θ_{ij} is raised by 1. From the figures, it is evident that the effects of these two policies on the evolution of average potential are similar. Although both policies will increase postings, the “Reducing Posting Cost” policy will lead to a greater number of ideas. To determine which policy is better from the firm’s perspective, we consider the cost of screening and the cost of the reward. It is obvious that the “Reducing Posting Cost” policy will cost the firm much more than the “Increase payoff” policy if the firm offers a monetary award. The screening cost will also be higher under the “Reducing Posting Cost” policy.

8. Conclusion

Our analysis of crowdsourcing data yields several important insights.

Why Does the Number of Contributed Ideas Decrease over Time?

Before IdeaStorm was launched there was no channel for customers to propose ideas to Dell for improving its products and services. When IdeaStorm was launched customers got a channel to propose ideas to Dell and the initial huge number of contributed ideas could represent the backlog of accumulated ideas that the

customers had “in stock” initially due to lack of such a channel. However, this does not explain why the potential of contributed ideas increases over time.

Our results show that, initially, individuals not only overestimate the potential of their ideas, but they also underestimate the cost of implementing their ideas. Hence, individuals tend to overestimate the probability that their idea will be implemented, and therefore, they initially post many ideas. As individuals learn about the true cost structure of the firm and the potential of their own ideas the expected utility of idea posting for marginal idea contributors decreases. These learning processes cause the low potential idea contributors to stop posting ideas. Hence, the two learning processes perform a filtering function.

As we explained earlier, an individual's ability to come up with high potential idea does not vary over time. The average potential of contributed ideas increases over time because over time the fraction of high potential idea contributors increases as the low potential idea provides stop contributing. Our results indicate that it is a combination of this initial response to emergence of the new channel along with the two learning processes that explain the decrease in ideas over time.

Why Does the Fraction of Ideas that are Implemented Increase over Time?

Individuals overestimate the potential of their ideas and underestimate the cost the firm incurs to implement their ideas. Once the website is launched, many individuals enter the “idea market”, and thus, the market is crowded by both high potential ideas and low potential ideas. As individuals learn the potential of their ideas from their experiences, marginal idea contributors tend to post fewer ideas. Consequently, at the aggregate level, the overall potential of ideas generated improves over time. From the firm’s point of view, the cost associated with implementing ideas is not changed. However, the implementation rate should increase over time.

Facilitated by technology, crowdsourcing has become an intriguing platform for direct idea generation and implementation. The attraction of the business model lies in the real-time assessment of ideas by peers (consumers). As the business headlines on this potentially powerful new product idea source shift from hype, a sobering reality has set in as a declining number of ideas are posted and few ideas are implemented. The observed empirical trend is seen as undermining the potential of crowdsourcing. On the contrary, our analysis suggests that this can be fully justified as a natural outcome of improving the efficiency of these markets. The findings bode well for these emerging new product idea generation methods. Based on these understandings, we propose several policies that may potentially improve the performance of these crowdsourcing initiatives and simulate the overall potential of the ideas and the number of ideas generated under these proposed policies. Our policy simulations indicate that providing more precise cost signals and rewards can help a firm procure higher potential ideas. Furthermore, associating rewards with implementation

is more cost-effective than offering rewards just for posting an idea. Another interesting finding in our policy simulation is that purely increasing the numbers of ideas to respond to and shortening the time to respond without differentiating the ideas negatively affects the overall potential of ideas. In fact, a better strategy is to respond only to the ideas of good idea contributors.

Our paper also has certain limitations. First, our dataset is limited. For example, one important variable, the registration time of individuals is not available in our data. In our analysis, we employ two extremes and argue that the reality should fall between the two extremes. If we had registration data, we would be able to obtain a more precise estimation for parameters in the learning dynamics. From the data, we know little about how the voting score of a particular idea is obtained, as we only observe the final score each idea receives at the end of 2010. We have no information on how many people promote an idea and how many people demote the idea, which is information that may allow us to interpret the voting scores more precisely. Second, in our current model, the firm is myopic and isolated. It implements ideas with positive net potential. However, in reality, the firm can be forward-looking. The firm may strategically implement ideas and send signals by taking consumers' reactions into consideration. The firm is also competing with other firms, and thus, it may also account for its competitors' reactions to its implementation decisions. Finally, our analysis is based on a dataset from a single crowdsourcing site, IdeaStorm.com. To test the external validity of this study, our model should be tested on datasets from other crowdsourcing ideation sites. Despite the limitations, our paper is the first to provide a dynamic structural framework that analyzes consumers' participation in the crowdsourcing ideation websites, helping both practitioners and researchers to understand this popular web application. We hope that our work can pave the way for future research in this important area.

References

- Archak, N., Sundararajan, A. 2009. "Optimal Design of Crowdsourcing Contests," *ICIS 2009 Proceedings*.
- Bayus, B. L. 2010. "Crowdsourcing and Individual Creativity Over Time: The Detrimental Effects of Past Success," *Working Paper*.
- Boudreau, K., Nicola L., Lakhani, K. 2011. "Incentives and Problem Uncertainty in Innovation Contests: An Empirical Analysis," *Management Science* (57:5) 843-863.
- Crawford, G., Shum, M. 2005. "Uncertainty and Learning in Pharmaceutical Demand," *Econometrica* (73:4) 1137-1173.
- DeGroot, M. H. 1970. *Optimal Statistical Decisions*, New York, McGraw-Hill.
- Di Gangi, P. M., Wasko, M., Hooker, R. 2010. "Getting customers' ideas to work for you: Learning from Dell how to succeed with online user innovation communities," *MIS Quarterly Executive* (9) 213-228.

- DiPalantino, D., Vojnovic, M. 2009. "Crowdsourcing and All-Pay Auctions," *Proceedings of the 10th ACM International Conference on Electronic Commerce* 119-128.
- Erdem, T., Keane, M. P. 1996. "Decision-Making under Uncertainty: Capturing Dynamic Choice Processes in Turbulent Consumer Goods Markets," *Marketing Science* (15:1) 1-20.
- Erdem, T., Keane, M. P., Sun, B. 2008. "A Dynamic Model of Brand Choice When Price and Advertising Signal Product Quality," *Marketing Science* (27:6) 1111-1125.
- Franke, N., von Hippel, E. 2003. "Satisfying Heterogeneous User Needs via Innovation Toolkits: the Case of Apache Security Software," *Research Policy* (32:7) 1199-1215.
- Howe, J. 2006. "The Rise of Crowdsourcing." *Wired Magazine* 14, 1-5.
- Huang, Y., Singh, P., Ghose, A., 2010 "Show Me the Incentives: A Dynamic Structural Model of Employee Blogging Behavior" *ICIS 2010 Proceedings*.
- Jeppesen, L., Lakhani, K. 2010. "Marginality and Problem-Solving Effectiveness in Broadcast Search," *Organization Science* (21) 1016-1033.
- Kuan, J. 2001. "Open Source Software as Consumer Integration into Production," *Working Paper*.
- Kumar, V., Gordon, B., Srinivasan, K. 2011 "Competitive Strategy for Open Source Software," forthcoming at *Marketing Science*.
- Lakhani, K. R., von Hippel, E. 2003. "How Open Source Software Works: "Free" User-to-user Assistance," *Research Policy* (32:6) 923-943.
- Mehta, N., Rajiv, S., Srinivasan, K. 2003. "Price Uncertainty and Consumer Search: A Structural Model of Consideration Set Formation," *Marketing Science* (22:1) 58-84.
- Mo, J., Zheng, Z, Geng, X. 2011. "Winning Crowdsourcing Contests: A Micro-Structural Analysis of Multi-Relational Social Network," *Workshop on information Management (CSWIM)*, Harbin, China.
- Narayanan, S, Manchanda. P. 2009. "Heterogeneous Learning and the Targeting of Marketing Communication for New Products Heterogeneous Learning and the Targeting of Marketing Communication for New Products," *Marketing Science* (28:3) 424-441.
- Netzer, O., Lattin, J. M., Srinivasan, V. 2008. "A Hidden Markov Model of Customer Relationship Dynamics," *Marketing Science* (27: 2) 185-204.
- Schulze, A., Hoegl M. 2008. "Organizational Knowledge Creation and the Generation of New Product Ideas: A Behavioral Approach, " *Research Policy* (37:10) 1742-1750.
- Terwiesch, C., Xu, Y. 2008. "Innovation Contests, Open Innovation, and Multiagent Problem Solving," *Management Science* (54:9) 1529-1543.
- Zhang, J. 2010. "The Sound of Silence: Observational Learning in the U.S. Kidney Market," *Marketing Science* (29:2) 315-335.

Crowdsourcing New Product Ideas under Consumer Learning

Appendix 1: Hierarchical Bayesian Estimation

As mentioned in the estimation strategy section, we use MCMC methods to estimate parameters in our model. To be more specific, the Gibbs sampler is applied to recursively make draws from the following conditional distribution of the model parameters:

$$\begin{aligned} & \boldsymbol{\beta}_i | A_i, C_j^e, \bar{\boldsymbol{\beta}}, \boldsymbol{\alpha} \\ & \bar{\boldsymbol{\beta}} | \boldsymbol{\beta}_i, \Sigma \\ & \Sigma | \boldsymbol{\beta}_i, \bar{\boldsymbol{\beta}} \\ & \boldsymbol{\alpha} | A, I, C_j^e, \boldsymbol{\beta}_i \\ & C_{jt}^e | A, C_{jt-1}^e, C_{jt+1}^e, \boldsymbol{\alpha}, \boldsymbol{\beta}_i \end{aligned}$$

The additional notation A_i denotes the vector of actions individual i takes in all periods, A denotes the decisions all individuals make in all periods, I denotes the decision the firm makes on all ideas posted within the observation period, $\boldsymbol{\beta}$ denotes $\boldsymbol{\beta}_i$ for all individuals, and C_j^e denotes the vector of the mean implementation cost beliefs in all periods. Further, the posterior distributions of $\boldsymbol{\beta}_i$, $\boldsymbol{\alpha}$ and C_j^e do not belong to any conjugate family, and therefore, we use the Metropolis-Hasting method to generate new draws. Each iteration involves five steps.

Step 1: Generate $\boldsymbol{\beta}_i$

The conditional distribution of $\boldsymbol{\beta}_i$ is

$$f(\boldsymbol{\beta}_i | A_i, C_j^e, \bar{\boldsymbol{\beta}}, \boldsymbol{\alpha}) \propto |\Sigma|^{-1/2} \exp[-1/2(\boldsymbol{\beta}_i - \bar{\boldsymbol{\beta}})' \Sigma^{-1}(\boldsymbol{\beta}_i - \bar{\boldsymbol{\beta}})] L(A_i | C_j^e, \boldsymbol{\beta}_i, \boldsymbol{\alpha})$$

Clearly, this posterior distribution does not have a closed form; therefore, we use the Metropolis-Hasting method to generate new draws with a random walk proposal density. The increment random variable is multivariate normally distributed with its variances adapted to obtain an acceptance rate of approximately 20% (Atchade, 2006). The probability that proposed $\boldsymbol{\beta}_i$ will be accepted is calculated using the following formula (the superscript *Prop* represents the proposed new $\boldsymbol{\beta}_i$ in this current iteration, i.e., iteration r . When $\text{accept}=1$, $\boldsymbol{\beta}_i^{r+1} = \boldsymbol{\beta}_i^{\text{Prop}}$; otherwise, $\boldsymbol{\beta}_i^{r+1} = \boldsymbol{\beta}_i^r$.)

$$Pr(\text{accept}) \propto \frac{f(\boldsymbol{\beta}_i^{Prop} | A_i, C_j^e, \bar{\boldsymbol{\beta}})}{f(\boldsymbol{\beta}_i^r | A_i, C_j^e, \bar{\boldsymbol{\beta}})}$$

$$= \frac{|\Sigma|^{-1/2} \exp[-1/2(\boldsymbol{\beta}_i^{Prop} - \bar{\boldsymbol{\beta}})' \Sigma^{-1} (\boldsymbol{\beta}_i^{Prop} - \bar{\boldsymbol{\beta}})] L(A_i | C_j^e, \boldsymbol{\beta}_i^{Prop}, \boldsymbol{\alpha})}{|\Sigma|^{-1/2} \exp[-1/2(\boldsymbol{\beta}_i^r - \bar{\boldsymbol{\beta}})' \Sigma^{-1} (\boldsymbol{\beta}_i^r - \bar{\boldsymbol{\beta}})] L(A_i | C_j^e, \boldsymbol{\beta}_i^r, \boldsymbol{\alpha})}$$

Step 2: Generate $\bar{\boldsymbol{\beta}}$

$$\bar{\boldsymbol{\beta}} | \boldsymbol{\beta}_i, \Sigma \sim MVN(u, W)$$

where

$$W = (Z'Z \otimes \Sigma^{-1} + W_0^{-1})^{-1}$$

$$u = V[(Z' \otimes \Sigma^{-1}) \text{vec}(\boldsymbol{\beta}) + W_0^{-1} u_0]$$

$Z = \text{vector of } (1\text{'s}) \text{ with length} = N$

$$\text{vec}(\boldsymbol{\beta}) = (\boldsymbol{\beta}_1, \boldsymbol{\beta}_2, \dots, \boldsymbol{\beta}_N)'$$

The priors are specified as:

$u_0 = \text{vector of } (0\text{'s}) \text{ with length} = N \times 5$

$$W_0 = 100I_5$$

Step 3: Generate Σ

$$\Sigma | \boldsymbol{\beta}_i, \bar{\boldsymbol{\beta}} \sim IW(f_0 + N, G_0^{-1} + \sum_{i=1}^N (\boldsymbol{\beta}_i - \bar{\boldsymbol{\beta}})' (\boldsymbol{\beta}_i - \bar{\boldsymbol{\beta}}))$$

where the prior hyper-parameter f_0 is set to 10, and G_0^{-1} is set to I_5 .

Step 4: Generate $\boldsymbol{\alpha}$

The conditional distribution of $\boldsymbol{\alpha}$ is

$$f(\boldsymbol{\alpha} | A, I, C_j^e, \boldsymbol{\beta}) \propto |\Sigma_{\boldsymbol{\alpha}}|^{-1/2} \exp[-1/2(\boldsymbol{\alpha} - \boldsymbol{\alpha}_0)' \Sigma_{\boldsymbol{\alpha}}^{-1} (\boldsymbol{\alpha} - \boldsymbol{\alpha}_0)] L(A | C_j^e, \boldsymbol{\beta}, \boldsymbol{\alpha}) L(I | \boldsymbol{\alpha}) L(C_j^e)$$

where

$$L(C_j^e) = \prod_{t=1}^T L(C_{jt}^e | C_{jt-1}^e, \boldsymbol{\alpha})$$

Similar to what we have done for β_i , we use the Metropolis-Hasting methods to make draws for α . The probability of acceptance is

$$\begin{aligned} \Pr(\text{accept}) &= \frac{f(\alpha^{prop}|A, I, C_j^e, \beta)}{f(\alpha^r|A, I, C_j^e, \beta)} \\ &= \frac{|\Sigma_\alpha|^{-\frac{1}{2}} \exp[-1/2(\alpha^{prop} - \alpha_0)' \Sigma_{\alpha_0}^{-1}(\alpha^{prop} - \alpha_0)] L(A|C_j^e, \beta, \alpha^{prop}) L(I|\alpha^{prop}) L(C_j^e)}{|\Sigma_\alpha|^{-\frac{1}{2}} \exp[-1/2(\alpha^r - \alpha_0)' \Sigma_{\alpha_0}^{-1}(\alpha^r - \alpha_0)] L(A|C_j^e, \beta, \alpha^r) L(I|\alpha^r) L(C_j^e)} \end{aligned}$$

where $\alpha_0 = (0, 0, \dots, 0)$ and $\Sigma_{\alpha_0}^{-1} = 100I_7$ are diffused priors.

Step 5: Generate C_j^e

Finally, we sequentially draw C_{jt}^e for $t=1$ to T . The conditional distribution of C_{jt}^e is

$$\begin{aligned} f(C_{jt}^e|A_{jt}, C_{jt-1}^e, \alpha, \beta) &\propto \\ &|v_{jt}^2|^{-\frac{1}{2}} \exp[-1/2(C_{jt}^e - \bar{C}_{jt}^e)' (v_{jt}^2)^{-1} (C_{jt}^e - \bar{C}_{jt}^e)] L(A_{jt}|C_{jt}^e, \alpha, \beta) L(C_{jt+1}^e|C_{jt}^e, \alpha) \end{aligned}$$

where A_{jt} denotes the decisions all individuals make on Category j idea in period t . \bar{C}_{jt}^e and v_{jt}^2 in the equation above are calculated using Equation (21) and (22). Again, because the posterior distribution does not have a close form, we have to use the Metropolis-Hasting methods to draw new C_{jt}^e .

The probability of acceptance is

$$\begin{aligned} \Pr(\text{accept}) &= \frac{f(C_{jt}^{eProp}|A, C_{jt-1}^e, \alpha)}{f(C_{jt}^{eR}|A, C_{jt-1}^e, \alpha)} \\ &= \frac{|v_{jt}^2|^{-\frac{1}{2}} \exp[-1/2(C_{jt}^{eProp} - \bar{C}_{jt}^e)' (v_{jt}^2)^{-1} (C_{jt}^{eProp} - \bar{C}_{jt}^e)] L(A_{jt}|C_{jt}^{eProp}, \alpha, \beta) L(C_{jt+1}^e|C_{jt}^{eProp}, \alpha)}{|v_{jt}^2|^{-\frac{1}{2}} \exp[-1/2(C_{jt}^{eR} - \bar{C}_{jt}^e)' (v_{jt}^2)^{-1} (C_{jt}^{eR} - \bar{C}_{jt}^e)] L(A_{jt}|C_{jt}^{eR}, \alpha, \beta) L(C_{jt+1}^e|C_{jt}^{eR}, \alpha)} \end{aligned}$$

Appendix 2: Model Identification

We now briefly discuss some intuition as to how the parameters in our model are identified. We cannot identify all Q_i and C_j at the same time because if we add a constant to all Q_i s and then subtract the same constant from all C_j s, we will obtain exactly the same likelihood value. For identification purposes, we fixed $C_1 = -6$. As a result, the estimated values of C_2 and Q_i should be interpreted as relative to C_1 . We observe individual's actions from which we can infer individual's utility derived from posting different categories of

ideas. Once we know \tilde{U}_{ijt} , we can infer parameters in the utility function. However, we cannot identify θ_{i0} and θ_{ij} simultaneously because if we observe an individual who posts frequently, especially when he/she learns about the potential of his/her ideas, we cannot tell whether it is because he/she incurs low cost for posting an idea or because he/she receives higher payoffs when his/her ideas are implemented. Therefore, θ_0 is fixed to -5, and θ_{ij} should then be interpreted as the value individual i derives from the implementation of his Category j idea relative to the cost he/she incurs when posting an idea. d_i can be easily identified because D_{it} is observed for every i in every period. The difference in individual i 's posting behavior between cases where $D_{it} = 0$ and $D_{it} = 1$ gives us the identification of d_i . The binary construction of D_{it} can help disentangle the effects of learning and dissatisfaction.

The difference in the number of ideas generated in Categories 1 and 2 contribute to the identification of C_2 and the population mean of Q_i and θ_{ij} . In the data, we observe that individuals post more Category 1 ideas than Category 2 ideas. This difference may be caused by $C_1 < C_2$ or by $\theta_{i1} > \theta_{i2}$. However, the first implementation of Category 1 ideas occurred in Week 11, while the first implementation of Category 2 ideas occurred in Week 7. As the potential of ideas posted by the same person is the same, the systematic difference in the number of ideas of each Category generated before Week 7 gives the difference between θ_{i1} and θ_{i2} , at least for the individuals who posted before Week 7. Note that we fix θ_{i0} to achieve the identification of θ_{ij} . An alternative identification strategy is that we could assume the payoff that individuals derive from the implementation of their ideas is the same while the cost for them to post different categories of ideas is different.

Ideally, after controlling for θ_{i1} and θ_{i2} , Q_i can be identified from the behavior of “well-informed” individuals whose perception about the firm’s cost structure and potential of their ideas is very close to the true value. As C_1 is fixed to -6, the only unknown parameter in Equation (19) for $j=1$ is Q_i . However, in our case, we cannot guarantee that at the end of the observation period, the individuals’ perceptions about the potential of their ideas and the implementation cost for the firm are closest to the true values. From our dataset, we can observe votes and the decisions the firm makes on each idea, and we can use the likelihood of idea implementation given the votes as another source of identification of Q_i and C_2 . Combining these two constraints, Q_i can be identified. Once Q_i is identified, the only unknown parameter in Equation (18) for $j=2$ is C_2 . This parameter can be identified through the probabilities that individuals post Category 2 ideas and the firm’s decisions on Category 2 ideas, given the votes each idea receives. Once C_1 , C_2 and Q_i are identified, σ_r^2 can be easily identified from the likelihood for one idea to be implemented. C_0 can be identified through the probability of posting in the first seven weeks. In these seven weeks, individuals have not received any cost signals, and their beliefs about the cost structure stay at C_0 , but they receive signals about the potential of

their ideas when they post. Given Q_i , C_0 can be easily identified. Given C_j and C_0 , Q_0 can be identified through the probability of posting for the latecomers throughout the whole observation period. Before one individual posts any ideas on the website for the first time, his/her beliefs about his/her idea's potential is always Q_0 , while his/her beliefs about the implementation cost is updated. Given the different C_{jt}^e for different t 's, Q_0 can then be identified. $cons$ and φ are identified from the relationship between Q_i and V_i : As $cons$ and φ are pooled parameters, given all individuals' Q_i and observed V_{sit} , $cons$ and φ determine a linear curve that best approximates the relationship between Q_i and V_{sit} .

Finally, the variance parameters σ_μ^2 and $\sigma_{\delta_i}^2$ are both identified from the dynamics of the posting behaviors of individuals over time. We are able to identify σ_μ^2 and $\sigma_{\delta_i}^2$ simultaneously because the signals of the implementation costs and the potentials are generated from different events. σ_μ^2 is identified through the dynamics of the choice probabilities at the population level. For example, if one idea is implemented in period t , the perceived cost of implementation for all individuals will be updated. For those who do not post in this period, their perception about the potential of their ideas has not changed before or after the period, and the changes in the probability of posting ideas after they receive the cost signal help us to identify σ_μ^2 . If σ_μ^2 is very small, which means that the cost signals individuals receive are precise, then individuals can learn faster, their perceptions converge to the true value quickly, and vice versa. Similarly, the learning speed of the potential of the ideas is affected by both $\sigma_{\delta_i}^2$ and the slope parameter φ . φ affects every individual in the same way, while $\sigma_{\delta_i}^2$ affects individuals differently. Therefore, the different learning speeds of different individuals observed in the data can help us identify $\sigma_{\delta_i}^2$.