Online Appendix to

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:

$$\boldsymbol{\beta}_i | A_i, C_j^e, \overline{\boldsymbol{\beta}}, \boldsymbol{\alpha}$$

$$\overline{\boldsymbol{\beta}} | \boldsymbol{\beta}_i, \Sigma$$

$$\Sigma | \boldsymbol{\beta}_i, \overline{\boldsymbol{\beta}}$$

$$\boldsymbol{\alpha} | A, I, C_j^e, \boldsymbol{\beta}_i$$

$$C_{jt}^e | A, C_{jt-1}^e, C_{jt+1}^e, \boldsymbol{\alpha}, \boldsymbol{\beta}_i$$

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

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

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 accept=1, $\boldsymbol{\beta}_i^{r+1} = \boldsymbol{\beta}_i^{Prop}$; otherwise, $\boldsymbol{\beta}_i^{r+1} = \boldsymbol{\beta}_i^{r}$.)

$$Pr(accept) \propto \frac{f\left(\boldsymbol{\beta}_i^{Prop} \middle| A_i, C_j^e, \overline{\boldsymbol{\beta}}\right)}{f\left(\boldsymbol{\beta}_i^r \middle| A_i, C_j^e, \overline{\boldsymbol{\beta}}\right)}$$

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

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

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

where

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

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

$$Z = vector\ of\ \left(1's\right) with\ length = N$$

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

The priors are specified as:

$$u_0 = vector\ of\ \left(0's\right)\ with\ length = 6$$

$$W_0 = 100 I_6$$

Step 3: Generate Σ

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

where the prior hyper-parameter $f_0$is set to 11, and G_0^{-1} is set to I_6.

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

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

$$f(\boldsymbol{\alpha}|A, I, C_j^e, \boldsymbol{\beta}) \propto |\Sigma_{\boldsymbol{\alpha}}|^{-\frac{1}{2}} exp[-1/2(\boldsymbol{\alpha} - \boldsymbol{\alpha_0})' \Sigma_{\boldsymbol{\alpha}}^{-1} (\boldsymbol{\alpha} - \boldsymbol{\alpha_0})] L(A|C_j^e, \boldsymbol{\beta}, \boldsymbol{\alpha}) L(I|\boldsymbol{\alpha}) L(C_j^e)$$

where

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

Similar to what we have done for $\boldsymbol{\beta}_i$, we use the Metropolis-Hasting methods to make draws for $\boldsymbol{\alpha}$. The probability of acceptance is

$$\Pr(accept) = \frac{f(\boldsymbol{\alpha}^{\boldsymbol{prop}}|A, I, C_j^e, \boldsymbol{\beta})}{f(\boldsymbol{\alpha}^{\boldsymbol{r}}|A, I, C_j^e, \boldsymbol{\beta})}$$

$$= \frac{|\Sigma_{\boldsymbol{\alpha}}|^{-\frac{1}{2}} exp[-1/2(\boldsymbol{\alpha}^{\boldsymbol{prop}} - \boldsymbol{\alpha}_{\boldsymbol{0}})' \Sigma_{\alpha_0}^{-1}(\boldsymbol{\alpha}^{\boldsymbol{prop}} - \boldsymbol{\alpha}_{\boldsymbol{0}})] L(A|C_j^e, \boldsymbol{\beta}, \boldsymbol{\alpha}^{\boldsymbol{prop}}) L(I|\boldsymbol{\alpha}^{\boldsymbol{prop}}) L(C_j^e)}{|\Sigma_{\boldsymbol{\alpha}}|^{-\frac{1}{2}} exp[-1/2(\boldsymbol{\alpha}^{\boldsymbol{r}} - \boldsymbol{\alpha}_{\boldsymbol{0}})' \Sigma_{\alpha_0}^{-1}(\boldsymbol{\alpha}^{\boldsymbol{r}} - \boldsymbol{\alpha}_{\boldsymbol{0}})] L(A|C_j^e, \boldsymbol{\beta}, \boldsymbol{\alpha}^{\boldsymbol{r}}) L(I|\boldsymbol{\alpha}^{\boldsymbol{r}}) L(C_j^e)}$$

where $\boldsymbol{\alpha}_{\boldsymbol{0}}$= (0,0,…0) and $\Sigma_{\alpha_0}^{-1} = 100 I_8$ 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

$$f\left(C_{jt}^e \middle| A_{jt}, C_{jt-1}^e, \boldsymbol{\alpha}, \boldsymbol{\beta}\right) \propto$$

$$\left|\upsilon_{jt}^2\right|^{-\frac{1}{2}} exp[-1/2\left(C_{jt}^e - \bar{C}_{jt}^e\right)' \left(\upsilon_{jt}^2\right)^{-1} \left(C_{jt}^e - \bar{C}_{jt}^e\right)] L(A_{jt}|C_{jt}^e, \boldsymbol{\alpha}, \boldsymbol{\beta}) L(C_{jt+1}^e|C_{jt}^e, \boldsymbol{\alpha})$$

where A_{jt} denotes the decisions all individuals make on Category *j* idea in period t. $\bar{C}_{jt}^e$ and υ_{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

$$\Pr(accept) = \frac{f(C_{jt}^{e\,Prop}|A, C_{jt-1}^e, \boldsymbol{\alpha})}{f(C_{jt}^{e\,r}|A, C_{jt-1}^e, \boldsymbol{\alpha})}$$

$$= \frac{\left|\upsilon_{jt}^2\right|^{-\frac{1}{2}} exp[-1/2\left(C_{jt}^{e\,Prop} - \bar{C}_{jt}^e\right)' \left(\upsilon_{jt}^2\right)^{-1} \left(C_{jt}^{e\,Prop} - \bar{C}_{jt}^e\right)] L(A_{jt}|C_{jt}^{e\,Prop}, \boldsymbol{\alpha}, \boldsymbol{\beta}) L(C_{jt+1}^e|C_{jt}^{e\,Prop}, \boldsymbol{\alpha})}{\left|\upsilon_{jt}^2\right|^{-\frac{1}{2}} exp[-1/2\left(C_{jt}^{e\,r} - \bar{C}_{jt}^e\right)' \left(\upsilon_{jt}^2\right)^{-1} \left(C_{jt}^{e\,r} - \bar{C}_{jt}^e\right)] L(A_{jt}|C_{jt}^{e\,\boldsymbol{r}}, \boldsymbol{\alpha}, \boldsymbol{\beta}) L(C_{jt+1}^e|C_{jt}^{e\,r}, \boldsymbol{\alpha})}$$

Appendix 2: Model Identification

We now briefly discuss some intuition as to how the parameters in our model are identified. In our model the consumers make posting decisions based on their (perceived) utility. With this assumption, we can infer individual's utility derived from posting different categories of ideas from their posting decisions. The basic logic behind the identification strategy that the "true" parameters in the utility function, as well as the "true" learning parameters, will lead to a utility function that can best predict the data we observe in the reality.

In the estimation, we fix the mean cost of one category (product ideas) and the variance of individuals' initial belief about the cost distribution and potential distribution. We have to fix the mean cost of one category because if we add a constant to all Q_is and then add the same constant to all C_js, we will obtain exactly the same utility value. When we fix C_1, we will be able to identify Q_is and C_2. As a result, the estimated values of C_2 and Q_i should be interpreted as relative to C_1. We set the initial variance of individuals' initial belief about the cost distribution and potential distribution to a large value to reflect the fact that individuals' prior believe is non-informative.

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 average learning speed (how much adjustment individuals make to their perceptions) of the potential of the ideas is affected by both $\sigma_{\delta_i}^2$ and the slope parameter φ. In addition, from Equation (10), we know the relationship between $\sigma_{\xi_i}^2$, the variance of the voting scores individual i's ideas receive, which can be directly estimated from the voting score data, and the variance of potential of the individuals i's ideas, $\sigma_{\delta_i}^2$, is $\sigma_{\xi_i}^2 = \varphi^2 \sigma_{\delta_i}^2$. Therefore, individuals' learning speed (how much their behavior change after receiving a potential signal) observed in the data can help us identify φ. Once φ is identified, $\sigma_{\delta_i}^2$ is also identified. Note that φ is a population level parameter. It is possible that there still remain variations in individuals' learning speed of the potential of the ideas, after controlling for $\sigma_{\delta_i}^2$. Thses remaining variations will be captured by θ_{ij}, which we will explain in detail later.

The overall frequency that an individual i posts Category j ideas is jointly determined by θ_{i0} and θ_{ij}. However, we are able to separately identify θ_{i0} and θ_{ij} because they enter the utility function in different ways. If we observe an individual who posts frequently, it could be because 1) he/she incurs low cost to post an idea; or 2) he/she receives higher payoffs when his/her Category j ideas are implemented. θ_{i0} is the constant term in the utility function, which does not change as individuals receive signals over time; while θ_{ij} is multiplied by the perceived probability of individuals' ideas being implemented. For example, when the firm implements a Category j idea and so all individuals' perceive costs of implementing Category j idea are updated. Individuals whose θ_{ij} is larger will be affected more significantly. In addition, the magnitude of θ_{ij} is also reflected in the changes in individuals posting behavior after they receive a potential signal. For example, consider two hypothetical individuals, namely A and B. From the voting score data, we find the mean and variance of their ideas' voting score are very similar. This implies that A and B updates their perception of the potential of their ideas in a similar way. However, individual A's probability of posting a Category j idea changes dramatically after she receives a new potential signal, while individual B's probability of posting Category j idea does not change a lot. The only cause of this difference is different θ_{ij}. Therefore, such variation help identify θ_{ij}. Similar logic can be applied for the identification of θ_{ij} for the same individual. Assume that after receiving a potential signal, individual A's probability of posting a Category 1 idea changes significantly, while her probability of posting a Category 2 idea only changes slightly, we can conclude that $\theta_{A1} > \theta_{A2}$. Once θ_{ij} is controlled, θ_{i0} can be identified from the overall frequency that individual i posts ideas (after controlling for θ_{ij}). 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 $D_{it} = 1$ identifies d_i. The binary construction of D_{it} can help disentangle the effects of learning and dissatisfaction.

The identification of Q_i and σ_γ^2 relies on two sets of observations. 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, is an important source of the identification of Q_i and σ_γ^2,. Note that we observe the voting score an idea receives is a linear function of the idea's potential, or $V_i = cons + \varphi Q_i$. V_i can be easily estimated by averaging all individual i's ideas' voting scores; and the identification of φ has been discussed previously. Given V_i and φ, identifying Q_i is equivalent to identifying $cons$. Consider a hypothetical "well-informed" individual's probability of posting a Category 1 idea is 0.1, i.e. $\exp(\tilde{U}_{ijt})/(1+\exp(\tilde{U}_{ijt})) = 0.1$. Solving for $\tilde{U}_{ijt}$, we get $\tilde{U}_{ijt}$=-2.303. Given θ_{i0}, θ_{i1}, d_i and C_1, as well as the variance parameter $\sigma_{\delta_i}^2$, Equation (19) is an equation of two unknown parameters is Q_i and σ_γ^2, or equivalently $cons$ and σ_γ^2. Another source of identifying Q_i and σ_γ^2 is the likelihood of observed implementation decisions on all Category 1 ideas. From our dataset, we observe the decisions the firm makes on each idea, given its voting score. In Equation 18, Q_{mjt} can be calculated by $Q_{mjt} = (V_{mjt} - cons)/\varphi$. Assume that a Category 1 idea with log-voting score equally 2 has 0.01 chance to be implemented, then $[(V_{mjt} - cons)/\varphi + C_1]/\sigma_\gamma^2$ = -2.326. Given V_{mjt} and φ, it is also an equation with two unknowns parameters $cons$ and σ_γ^2. Combining these two constraints, Q_i (or equivalently $cons$) and σ_γ^2 can be identified. Once Q_i is identified, C_2 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. C_0 can be identified through the probability of posting in the first seven weeks as no idea was implemented before the seventh week. 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 an 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.

Appendix 3: Derivation of the Updating Rules

We begin with the Bayes rule. The Bayes rule is

$P(A|B) = \frac{P(B|A)P(A)}{P(B)} \propto P(B|A)P(A)$. Now let us explain how we use the Bayes rule in coming up with our updating functions. In the Bayesian updating process, A represents people's belief about a certain parameter, B represents signal. Let us begin with the learning process of the implementation cost. Assume that individuals' prior belief about the mean of implementation cost in period *t* follows a normal distribution N $(C_{jt-1}^e, \sigma_{C_{jt-1}}^2)$ and the cost signal individuals receive in period *t* is $C_{kjt} \sim N(C_j, \sigma_\mu^2)$. The updated (posterior) distribution of the cost distribution is N $(C_{jt}^e, \sigma_{C_{jt}}^2)$.

The prior of the mean implementation cost in period t follows a normal distribution $N\ (C^{e}_{jt-1}, \sigma^2_{C_{jt-1}})$, and this is similar to the term $P(A)$ in the equation above. As we are dealing with a continuous distribution, instead of probability mass P, we use the probability density function of the normal distribution $N\ (C^{e}_{jt-1}, \sigma^2_{C_{jt-1}})$ below:

$$p\left(C_j \middle| C^{e}_{jt-1}, \sigma^2_{C_{jt-1}}\right) = \left(2\pi\sigma^2_{C_{jt-1}}\right)^{-\frac{1}{2}} exp\left[-\frac{1}{2\sigma^2_{C_{jt-1}}}\left(C_j - C^{e}_{jt-1}\right)^2\right].$$

As we assume that the cost signal C_{kjt} follows a normal distribution $C_{kjt} \sim N(C_j, \sigma^2_{\mu})$, the probability density of observing a cost signal of a value C_{kjt} is:

$$p\left(C_{kjt} \middle| C_j\right) = \left(2\pi\sigma^2_{\mu}\right)^{-\frac{1}{2}} exp\left[-\frac{1}{2\sigma^2_{\mu}}\left(C_{kjt} - C_j\right)^2\right]$$

Let $D = (C_{1jt}, \ldots, C_{k_{Cjt}jt})$ be the cost signals individuals receive in period t, with k_{Cjt} indicating the number of such signals. The likelihood of observing D, given C_j and σ^2_{μ} is simply the product of the $p\left(C_{kjt} \middle| C_j\right)$ over k=1 to k_{Cjt}.

$$p\left(D \middle| C_j\right) = \prod_{k=1}^{k_{Cjt}} p\left(C_{kjt} \middle| C_j\right) = \left(2\pi\sigma^2_{\mu}\right)^{-\frac{k_{Cjt}}{2}} exp\left[-\frac{1}{2\sigma^2_{\mu}}\sum_{k=1}^{k_{Cjt}}\left(C_{kjt} - C_j\right)^2\right]$$

$$\propto exp\left[-\frac{1}{2\sigma^2_{\mu}}\sum_{k=1}^{k_{Cjt}}\left(C_{kjt} - C_j\right)^2\right]$$

Here,$p\left(D \middle| C_j\right)$ is similar to $P(B|A)$ in the first equation. Finally, $p\left(C_j \middle| D\right)$ corresponds to $P(A|B)$. Following the Bayes rule $P(A|B) = \propto P(B|A)P(A)$, the posterior of mean cost of implementation is

$$p\left(C_j \middle| D\right) \propto p\left(D \middle| C_j\right) p\left(C_j \middle| C^{e}_{jt-1}, \sigma^2_{C_{jt-1}}\right)$$

$$\propto \exp\left[-\frac{1}{2\sigma^2_{\mu}}\sum_{k=1}^{k_{Cjt}}\left(C_{kjt} - C_j\right)^2\right] \exp\left[-\frac{1}{2\sigma^2_{C_{jt-1}}}\left(C_j - C^{e}_{jt-1}\right)^2\right]$$

$$= \exp\left[-\frac{1}{2\sigma^2_{\mu}}\sum_{k=1}^{k_{Cjt}}\left({C_{kjt}}^2 + {C_j}^2 - 2C_{kjt}C_j\right)^2 - \frac{1}{2\sigma^2_{C_{jt-1}}}\left({C_j}^2 + {C^{e}_{jt-1}}^2 - 2C_jC^{e}_{jt-1}\right)^2\right]$$

$$\propto exp\left[-\frac{{C_j}^2}{2}\left(\frac{1}{\sigma^2_{C_{jt-1}}} + \frac{k_{Cjt}}{\sigma^2_{\mu}}\right) + C_j\left(\frac{C^{e}_{jt-1}}{\sigma^2_{C_{jt-1}}} + \frac{\sum_{k=1}^{k_{Cjt}} C_{kjt}}{\sigma^2_{\mu}}\right) - \left(\frac{{C^{e}_{jt-1}}^2}{2\sigma^2_{C_{jt-1}}} + \frac{\sum_{k=1}^{k_{Cjt}} {C_{kjt}}^2}{2\sigma^2_{\mu}}\right)\right] \quad (*)$$

So far, we have derived the posterior distribution of C_j in terms of the prior and the signals. The posterior is also normally distributed and parameterized by two parameters—mean and variance. Therefore, if we can recover the mean and the variance of the posterior distribution, we can fully describe the posterior distribution. We do this in the following steps. Given $C_{jt}^{e}$and $\sigma_{C_{jt}}^{2}$

$$p(C_j|D) = \left(2\pi\sigma_{C_{jt}}^{2}\right)^{-\frac{1}{2}} exp\left[-\frac{1}{2\sigma_{C_{jt}}^{2}}\left(C_j - C_{jt}^{e}\right)^2\right] \propto exp\left[-\frac{1}{2\sigma_{C_{jt}}^{2}}\left({C_j}^2 + {C_{jt}^{e}}^2 - 2C_jC_j^{e}\right)^2\right]. \text{(**)}$$

This is just the definition of the posterior distribution. That is, Equation (*)≡Equation (**). To connect prior distribution and signals with posterior variance $\sigma_{C_{jt}}^{2}$, we match coefficients of ${C_j}^2$ in Equation (*) and Equation (**). We then have

$$\frac{-{C_j}^2}{2\sigma_{C_{jt}}^{2}} = \frac{-{C_j}^2}{2}\left(\frac{1}{\sigma_{C_{jt-1}}^{2}} + \frac{k_{Cjt}}{\sigma_{\mu}^{2}}\right)$$

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

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

which is Equation (7) in the paper. Similarly, we match coefficients of C_j in equation (*) and equation (**), and then get

$$\frac{2C_jC_{jt}^{e}}{2\sigma_{C_{jt}}^{2}} = C_j\left(\frac{C_{jt-1}^{e}}{\sigma_{C_{jt-1}}^{2}} + \frac{\sum_{k=1}^{k_{Cjt}} C_{kjt}}{\sigma_{\mu}^{2}}\right)$$

$$\frac{C_{jt}^{e}}{\sigma_{C_{jt}}^{2}} = \left(\frac{C_{jt-1}^{e}}{\sigma_{C_{jt-1}}^{2}} + \frac{\sum_{k=1}^{k_{Cjt}} C_{kjt}}{\sigma_{\mu}^{2}}\right)$$

$$C_{jt}^{e} = \left(\frac{C_{jt-1}^{e}}{\sigma_{C_{jt-1}}^{2}} + \frac{\sum_{k=1}^{k_{Cjt}} C_{kjt}}{\sigma_{\mu}^{2}}\right)\sigma_{C_{jt}}^{2}$$

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

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

$$= C^e_{jt-1}\left(1 - \frac{k_{Cjt}\sigma^2_{C_{jt-1}}}{k_{Cjt}\sigma^2_{C_{jt-1}} + \sigma^2_\mu}\right) + C_{sjt}\left(\frac{k_{Cjt}\sigma^2_{C_{jt-1}}}{k_{Cjt}\sigma^2_{C_{jt-1}} + \sigma^2_\mu}\right)$$

$$= C^e_{jt-1} + \left(C_{sjt} - C^e_{jt-1}\right)\frac{\sigma^2_{C_{jt-1}}}{\sigma^2_{C_{jt-1}} + \frac{\sigma^2_\mu}{k_{Cjt}}}$$

which is Equation (6) in the paper.

The proof of the updating rules of V^e_{it} and $\sigma^2_{V_{it-1}}$ is almost identical to the proof we derive above. The derivation of the updating rules of Q^e_{it} and $\sigma^2_{Q_{it-1}}$ is only slightly different. Assume for a moment that people directly observe the potential signals Q_{sit}, then the updating rules for Q_i will be

$$\sigma^2_{Q_{it}} = \frac{1}{\frac{1}{\sigma^2_{Q_{it-1}}} + \frac{k_{Qit}}{\sigma^2_{\delta\,i}}}$$

$$Q^e_{it} = Q^e_{it-1} + (Q_{sit} - Q^e_{it-1})\frac{\sigma^2_{Q_{it-1}}}{\sigma^2_{Q_{it-1}} + \frac{\sigma^2_{\delta\,i}}{k_{Qit}}}$$

However, in reality, the potential signal Q_{sit} is not directly observed. Instead, individuals observe the voting score V_{kit} and then use the linear relation between V_{kit} and Q_{kit} to recover the potential signal Q_{kit}. As in the paper, we assume the relationship between V_{kit} and Q_{kit} as

$$V_{kit} = \text{cons} + \varphi Q_{kit}$$

$$V_{kit} = V_i + \xi_{kit}$$

$$\sigma^2_{\xi\,i} = \varphi^2\sigma^2_{\delta\,i}$$

Therefore, $Q_{kit} = (V_{kit} - cons)/\varphi$, $Q^e_{it-1} = (V^e_{it-1} - cons)/\varphi$ and $\sigma^2_{\delta\,i} = \sigma^2_{\xi\,i}/\varphi^2$. Now the two updating rules discussed above can be rewritten as:

$$\sigma^2_{Q_{it}} = \frac{1}{\frac{1}{\sigma^2_{Q_{it-1}}} + \frac{k_{Qit}}{\sigma^2_{\delta\,i}}} = \frac{1}{\frac{1}{\sigma^2_{Q_{it-1}}} + \frac{\varphi^2 k_{Qit}}{\sigma^2_{\xi\,i}}}$$

$$Q^e_{it} = Q^e_{it-1} + (Q_{sit} - Q^e_{it-1})\frac{\sigma^2_{Q_{it-1}}}{\sigma^2_{Q_{it-1}} + \frac{\sigma^2_{\delta\,i}}{k_{Qit}}}$$

$$= Q^e_{it-1} + \left(\frac{V_{kit} - cons}{\varphi} - \frac{V^e_{it-1} - cons}{\varphi}\right)\frac{\sigma^2_{Q_{it-1}}}{\sigma^2_{Q_{it-1}} + \frac{\sigma^2_{\xi\,i}}{\varphi^2 k_{Qit}}}$$

$$= Q^e_{it-1} + (V_{sit} - V^e_{it-1})\ \frac{\varphi\sigma^2_{Q_{it-1}}}{\varphi^2\sigma^2_{Q_{it-1}} + \frac{\sigma^2_{\xi\ i}}{k_{Qit}}}$$

These are equation (17) and (15) in the paper.

In this learning model, we impose two assumptions. First, the implementation cost is normally distributed which is a continuous distribution. Second, we assume individuals' prior belief about the mean implementation cost is also normally distributed, which is also a continuous distribution. The normal prior assumption provides tractability benefits, as the posterior will also have a closed form representation. For this reason, in learning literature, most models use this formulation. This type of learning model has some nice features. For example, from Equation (7), $\sigma^2_{C_{jt}} = \frac{1}{\frac{1}{\sigma^2_{C_{jt-1}}} + \frac{k_{Cjt}}{\sigma^2_{\mu}}}$, we can see that $\sigma^2_{C_{jt}}$ is monotonically decreasing. This means that as individuals receive more signals, the variance of the posterior distribution keeps decreasing, and so their uncertainty is reduced. From Equation (6), $C^e_{jt} = C^e_{jt-1} + \left(C_{sjt} - C^e_{jt-1}\right)\frac{\sigma^2_{C_{jt-1}}}{\sigma^2_{C_{jt-1}} + \frac{\sigma^2_{\mu}}{k_{cjt}}}$, we can see that individuals' new belief about the mean of the cost distribution is affected by their prior belief and the new signal they receive. Individuals adjust their belief by comparing the new signals they receive and their prior belief. $\frac{\sigma^2_{C_{jt-1}}}{\sigma^2_{C_{jt-1}} + \frac{\sigma^2_{\mu}}{k_{cjt}}}$ tells us the weight individuals assume to the new signals. $\frac{\sigma^2_{C_{jt-1}}}{\sigma^2_{C_{jt-1}} + \frac{\sigma^2_{\mu}}{k_{cjt}}}$ is always between 0 and 1. σ^2_{μ} represents the variance of the signals. When σ^2_{μ} is small, which means that the signal is precise, individuals assign a larger weight to the new signals and so their beliefs get updated faster. In addition, when σ^2_{μ} is fixed, the weight assigned to the difference is bigger when the variance of the prior ($\sigma^2_{C_{jt-1}}$) is large. This indicates that individuals learn very quickly in the beginning. As $\sigma^2_{C_{jt}}$ becomes smaller individuals' learning progress slows down, their belief will tend to stabilize. These features match individuals' real-world behavior well.

Appendix 4 Convergence of the Markov Chain

In our model, we have two sets of parameters and we will show the convergence of the chains for the two sets of parameters separately. Parameter vector $\alpha = [\ C_0, C_2, \sigma^2_{\gamma 1}\ \sigma^2_{\gamma 2}, \sigma^2_{\mu}, Q_0, cons, \varphi]$ is common across individuals, while parameter vector $\beta_i = \beta_i = [Q_i, log(\sigma^2_{\delta\ i}), d_i, \theta_{i0}, \theta_{i1}, \theta_{i2}]$ is heterogeneous across individuals. We further assume that β_i follows the following distribution

$$\beta_i = \begin{pmatrix} Q_i \\ log(\sigma^2_{\delta\ i}) \\ d_i \\ \theta_{i0} \\ \theta_{i1} \\ \theta_{i2} \end{pmatrix} \sim MVN(\bar{\beta}, \Sigma)$$

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

We plot the series of draws of α and $\bar{\beta}$ separately. The Markov chain was run a total of 45,000 iterations, and plotted is every 30th draw of the chain. The figures indicate that chain converged after about 9,000 iterations. And the convergence of $\bar{\beta}$ is slightly faster.