Incentivizing Truthful Feedback on Crowdsourcing Platforms

Sponsored

RESULTS

Sponsored

NUU A9L | Unlocked 4G LTE Smartphone | 6.3" Display | 32GB + 3GB RAM | 3000 mAh Battery | Android 11 Edition | Compatible with T-Mobile AT&T| 2022 | Q3 Indigo Blue…

49

$94.90 ~~$129.90~~

Save $5.00 with coupon

prime

FREE delivery **Sun, Nov 20**

Small Business

+1 colors/patterns

Display Size	Memory	Color	Brand
6.3 inches	**32 GB**	**Blue**	**NUU**

Samsung Galaxy S20+ 5G 128GB Fully Unlocked Smartphone (Renewed)

1,761

$244.99 ~~$699.95~~

FREE delivery **Nov 17 - 21**

Climate Pledge Friendly

More Buying Choices

$211.00 (69 used & new offers)

Display Size	Memory	Color	Brand
6.7 inches	**128.0 GB**	**Cosmic Gray**	**SAMSUNG**

Moto G Stylus 5G | 2021 | 2-Day battery | Unlocked | Made for US by Motorola | 6/256GB | 48MP Camera | Cosmic Emerald

1,453

$199.99 ~~$399.99~~

prime

FREE delivery **Sun, Nov 20**

More Buying Choices

$191.99 (15 used & new offers)

Display Size	Memory	Color	Brand
6.8 inches	**256 GB**	**Cosmic Emerald**	**Motorola**

Moto G stylus | 2021 | 2-Day battery | Unlocked | Made for US by Motorola | 4/128GB | 48MP Camera | Black

Introduction

- Rating systems: Users provide feedback about their experiences with a product.
  - Goal is to rate the quality of a product.

Introduction

- Rating systems: Users provide feedback about their experiences with a product.
  - Goal is to rate the quality of a product.
- Challenge is that users who report their reviews may not be truthful.

Introduction

- Rating systems: Users provide feedback about their experiences with a product.
  - Goal is to rate the quality of a product.
- Challenge is that users who report their reviews may not be truthful.
- This paper proposes a rewarding mechanism to the users so that they are incentivized to provide truthful feedback.

A Simple Example

- Consider two types of plumbers X={Good, Bad}.
  - Distribution over types of plumbers is $P_X(x)$.

A Simple Example

- Consider two types of plumbers X={Good, Bad}.
  - Distribution over types of plumbers is $P_X(x)$.
- A user observes a plumber arriving within 1) 5 minutes or 2) 5 to 10 minutes.
  - Observation set Y = {5minutes, 5-10 minutes}

A Simple Example

- Consider two types of plumbers X={Good, Bad}.
  - Distribution over types of plumbers is $P_X(x)$.
- A user observes a plumber arriving within 1) 5 minutes or 2) 5 to 10 minutes.
  - Observation set Y = {5minutes, 5-10 minutes}
- For each type x of plumber, the distribution of observations is $p_y(x)$.

A Simple Example

- Consider two types of plumbers X={Good, Bad}.
  - Distribution over types of plumbers is $P_X(x)$.
- A user observes a plumber arriving within 1) 5 minutes or 2) 5 to 10 minutes.
  - Observation set Y = {5minutes, 5-10 minutes}
- For each type x of plumber, the distribution of observations is $p_y(x)$.
- Each agent observes the arrival time of one or more plumbers, and they report the plumbers accordingly.

A Simple Example

- Consider two types of plumbers X={Good, Bad}.
  - Distribution over types of plumbers is $P_X(x)$.
- A user observes a plumber arriving within 1) 5 minutes or 2) 5 to 10 minutes.
  - Observation set Y = {5minutes, 5-10 minutes}
- For each type x of plumber, the distribution of observations is $p_y(x)$.
- Each agent the arrival time of one or more plumbers, and they review the plumbers accordingly.
- We want to estimate $p_y(x)$ for every plumber type from the users' feedback.
  - Therefore, we want to incentivize workers to be truthful.

Notations

- X denotes the possible types of entities (in our case types of plumbers) and $P_X(x)$ denotes the distribution over these entities.

Notations

- X denotes the possible types of entities (in our case types of plumbers) and $P_X(x)$ denotes the distribution over these entities.
- Y denotes the set of observations for any task. $p_y(x)$ is the probability of an observation y given that the entity is x.

Notations

- X denotes the possible types of entities (in our case types of plumbers) and $P_X(x)$ denotes the distribution over these entities.
- Y denotes the set of observations for any task. $p_y(x)$ is the probability of an observation y given that the entity is x.
- T tasks and W workers.

Notations

- X denotes the possible types of entities (in our case types of plumbers) and $P_X(x)$ denotes the distribution over these entities.
- Y denotes the set of observations for any task. $p_y(x)$ is the probability of an observation y given that the entity is x.
- T tasks and W workers.
- Worker W responds to tasks T_w, task t is evaluated by Workers W_t.
    - Assume $W_t > 1$.

Notations

- X denotes the possible types of entities (in our case types of plumbers) and $P_X(x)$ denotes the distribution over these entities.
- Y denotes the set of observations for any task. $p_y(x)$ is the probability of an observation y given that the entity is x.
- T tasks and W workers.
- Worker W responds to tasks T_w, task t is evaluated by Workers W_t.
  - Assume $W_t > 1$.
- Y_{wt} denotes the response provided by the w^{th} worker for the t^{th} task

Notations

- X denotes the possible types of entities (in our case types of plumbers) and $P_X(x)$ denotes the distribution over these entities.
- Y denotes the set of observations for any task. $p_y(x)$ is the probability of an observation y given that the entity is x.
- T tasks and W workers.
- Worker W responds to tasks T_w, task t is evaluated by Workers W_t.
  - Assume $W_t > 1$.
- Y_{wt} denotes the response provided by the w^{th} worker for the t^{th} task
- v_w (y) is rewarded to worker w for a response y.

Problem Statement

- The goal is to design a rewarding mechanism $v\colon Y \rightarrow [0, 1]$ that incentivizes truthful reviews from workers.

Problem Statement

- The goal is to design a rewarding mechanism $v: Y \rightarrow [0, 1]$ that incentivizes truthful reviews from workers.
- One way to ensure that agents are truthful is to induce a Bayes-Nash equilibrium.

Problem Statement

- The goal is to design a rewarding mechanism $v\colon Y \to [0, 1]$ that incentivizes truthful reviews from workers.
- One way to ensure that agents are truthful is to induce a Bayes-Nash equilibrium.
- A rewarding mechanism induces a Bayes-Nash equilibrium if the strategy profile (response of workers) $\{ Y_{wt} : w \in W, t \in T_w \}$ satisfies the inequality:

$$E[\, v_w (\{ Y_{wt} : w \in W, t \in T_w \})] \geq E[\, v_w (\{ Y'_{wt} : w \in W, t \in T_w\} \cup \{ Y_{w't} : w' \in W, w' \neq w, t \in T_{w'} \})]$$

for each alternative strategy $\{Y'_{wt} : t \in T_w\} \neq \{Y_{wt} : t \in T_w \}$ and all workers w.

Intuitive Explanation

- The Bayes-Nash Equilibrium says that if all the other workers adhere to the strategy profile $\{ Y_{wt} : w \in W, t \in T_w \}$ then it is beneficial to worker w to also adhere to $\{ Y_{wt} : w \in W, t \in T_w \}$.
- In our case, we would want the truthful strategy to be a Bayes-Nash Equilibrium.
- Further, if the inequality is strict, then it is called a strict Bayes-Nash Equilibrium.

Informational Requirements and Assumptions

- Assumptions and Public Knowledge
  - Generating model (P_X, p) is separated: For every $y \neq y'$,
    $$E_{X\sim P_X}\left[p_y(X)p_{y'}(X)\right] < E_{X\sim P_X}\left[p_y^2(X)\right]E_{X\sim P_X}\left[p_{y'}^2(X)\right].$$

Informational Requirements and Assumptions

- Assumptions and Public Knowledge
  - Generating model (P_X, p) is separated: For every $y \neq y'$,
    $$E_{X\sim P_X}\left[p_y(X)p_{y'}(X)\right] < E_{X\sim P_X}\left[p_y^2(X)\right]E_{X\sim P_X}\left[p_{y'}^2(X)\right].$$
  - α-separation quantifies the gap:
    $$\alpha + E_{X\sim P_X}\left[p_y(X)p_{y'}(X)\right] < E_{X\sim P_X}\left[p_y^2(X)\right]E_{X\sim P_X}\left[p_{y'}^2(X)\right].$$

Informational Requirements and Assumptions

- Assumptions and Public Knowledge
  - Generating model (P_X, p) is separated: For every $y \neq y'$,
    $$E_{X\sim P_X}\left[p_y(X)p_{y'}(X)\right] < E_{X\sim P_X}\left[p_y^2(X)\right]E_{X\sim P_X}\left[p_{y'}^2(X)\right].$$
  - α-separation quantifies the gap:
    $$\alpha + E_{X\sim P_X}\left[p_y(X)p_{y'}(X)\right] < E_{X\sim P_X}\left[p_y^2(X)\right]E_{X\sim P_X}\left[p_{y'}^2(X)\right].$$
  - Workers and host do not know the generating model (P_X, p)
    - Know that the model is separated.
  - Workers and host assume that everyone is statistically identical.

Square Root Agreement Rule

- W_t = set of workers that respond to task t.
- T_w = set of tasks that worker w submitted a response to.
  - T = All tasks; $|T_w| \leq T \ \forall w$.
- Y_{wt} = Response of worker w to task t.

W_t = set of workers that evaluated task t.
T_w = set of tasks that worker w submitted a response to.
T = All tasks; $|T_w| \leq T \; \forall w$.
Y_{wt} = Response of worker w to task t.

Square Root Agreement Rule

- Repeat the following for every worker $w \in [n]$:
  - Let T_w be the set of tasks answered by worker w.
  - For every task $t \notin T_w$ ***not*** answered by this worker, select two different workers i, j who replied to task t.
  - Compute their #agreements on every possible observation y:

$$A(y) = \sum_{t \notin T_w} 1\{Y_{it} = Y_{jt} = y\}$$

  - Scale appropriately to define a popularity index $I_w(y)$ for every observation y:

$$I_w(y) = \frac{1}{T - |T_w|}\left(1 + \sum_{t \notin T_w} 1\{Y_{it} = Y_{jt} = y\}\right).$$

W_t = set of workers that evaluated task t.
T_w = set of tasks that worker w submitted a response to.
T = All tasks; $|T_w| \leq T \ \forall w$.
Y_{wt} = Response of worker w to task t.

Square Root Agreement Rule

- Repeat the following for every worker $w \in [n]$:
  - Let T_w be the set of tasks answered by worker w.

W_t = set of workers that evaluated task t.
T_w = set of tasks that worker w submitted a response to.
T = All tasks; $|T_w| \leq T \; \forall w$.
Y_{wt} = Response of worker w to task t.

Square Root Agreement Rule

- Repeat the following for every worker $w \in [n]$:
  - Let T_w be the set of tasks evaluated by worker w.
  - For every task $t \notin T_w$ ***not*** evaluated by this worker, select two different workers i, j who replied to task t.

W_t = set of workers that evaluated task t.
T_w = set of tasks that worker w submitted a response to.
T = All tasks; $|T_w| \leq T \ \forall w$.
Y_{wt} = Response of worker w to task t.

Square Root Agreement Rule

- Repeat the following for every worker $w \in [n]$:
  - Let T_w be the set of tasks evaluated by worker w.
  - For every task $t \notin T_w$ ***not*** evaluated by this worker, select two different workers i, j who replied to task t.
  - Compute their #agreements on every possible observation y:

$$A(y) = \sum_{t \notin T_w} 1\{Y_{it} = Y_{jt} = y\}$$

W_t = set of workers that evaluate task t.
T_w = set of tasks that worker w submitted a response to.
T = All tasks; $|T_w| \leq T \ \forall w$.
Y_{wt} = Response of worker w to task t.

Square Root Agreement Rule

- Repeat the following for every worker $w \in [n]$:
  - Let T_w be the set of tasks evaluated by worker w.
  - For every task $t \notin T_w$ ***not*** evaluated by this worker, select two different workers i, j who replied to task t.
  - Compute their #agreements on every possible observation y:

$$A(y) = \sum_{t \notin T_w} 1\{Y_{it} = Y_{jt} = y\}$$

  - Scale appropriately to define a popularity index $I_w(y)$ for every observation y:

$$I_w(y) = \frac{1}{T - |T_w|}\big(1 + A(y)\big).$$

W_t = set of workers that evaluated task t.
T_w = set of tasks that worker w submitted a response to.
T = All tasks; $|T_w| \leq T \; \forall w$.
Y_{wt} = Response of worker w to task t.

Square Root Agreement Rule

$$I_w(y) = \frac{1}{T - |T_w|}\left(1 + \sum_{t \notin T_w} 1\{Y_{it} = Y_{jt} = y\}\right).$$

- For a task t that worker w answered (say y), reward the worker only if she agrees with an arbitrarily selected co-worker w':

$$v_w(Y_w = s) = \frac{1\{Y_{w'} = s\}}{\sqrt{I_w(y)}}, s \in \mathcal{Y}$$

inversely proportional to the popularity of the agreement.

Intuition 1: Why Agreement Rule?

- Agent w is to choose between reporting y and $y' \neq y$.

Intuition 1: Why Agreement Rule?

- Agent w is to choose between reporting y and $y' \neq y$.
- Expected reward for the choice of report (strategy) $s \in \mathcal{Y}$:

$$E\left[1\{Y_{w'} = s \,|Y_w = y\}\right] = \begin{cases} P\left(Y_{w'} = y|Y_w = y\right): s = y \text{ (truthful)} \\ P\left(Y_{w'} = y'\middle|Y_w = y\right): s = y' \neq y \end{cases}.$$

Intuition 1: Why Agreement Rule?

- Agent w is to choose between reporting y and $y' \neq y$.
- Expected reward for the choice of report (strategy) $s \in \mathcal{Y}$:

$$E\left[1\{Y_{w'} = s \,|Y_w = y\}\right] = \begin{cases} P\left(Y_{w'} = y|Y_w = y\right): s = y \text{ (truthful)} \\ P\left(Y_{w'} = y'|Y_w = y\right): s = y' \neq y \end{cases}.$$

- Agents are statistically identical; If I observe y, it is likely that others observe y.

Intuition 1: Why Agreement Rule?

- Agent w is to choose between reporting y and $y' \neq y$.
- Expected reward for the choice of report (strategy) $s \in \mathcal{Y}$:

$$E\left[1\{Y_{w'} = s \,|Y_w = y\}\right] = \begin{cases} P\left(Y_{w'} = y|Y_w = y\right): s = y \text{ (truthful)} \\ P\left(Y_{w'} = y'\middle|Y_w = y\right): s = y' \neq y \end{cases}.$$

- Agents are statistically identical; If I observe y, it is likely that others observe y.
  - This incentivizes truthful behavior, based on an agent's observation.

Intuition 1: Why Agreement Rule?

- Agent w is to choose between reporting y and $y' \neq y$.
- Expected reward for the choice of report (strategy) $s \in \mathcal{Y}$:

$$E\left[1\{Y_{w'} = s \,|Y_w = y\}\right] = \begin{cases} P\left(Y_{w'} = y | Y_w = y\right) : s = y \text{ (truthful)} \\ P\left(Y_{w'} = y' \middle| Y_w = y\right) : s = y' \neq y \end{cases}.$$

- Agents are statistically identical; If I observe y, it is likely that others observe y.
  - This incentivizes truthful behavior, based on an agent's observation.
- Issue: Workers may find it more likely for others to submit a highly popular observation.

Intuition 2: Importance of Popularity Index

- Popularity index, asymptotic:

$$I_w(y) = \frac{1}{T - |T_w|}\left(1 + \sum_{t \notin T_w} 1\{Y_i = Y_j = y\}\right)$$
$$\rightarrow^{T_w, T \rightarrow \infty} P(Y_i = Y_j = y).$$

Intuition 2: Importance of Popularity Index

- Popularity index, asymptotic:

$$I_w(y) = \frac{1}{T - |T_w|}\left(1 + \sum_{t \notin T_w} 1\{Y_i = Y_j = y\}\right)$$
$$\rightarrow^{T_w, T \rightarrow \infty} P\left(Y_i = Y_j = y\right).$$

  - Repels workers from submitting an answer based on its popularity.

Result: SRA induces a Bayes-Nash Eq.

- Suppose all other agents $\neq w$ are truthful.

$$E[v_w(s)] = \begin{cases} \dfrac{P\left(Y_{w'} = y \middle| Y_w = y\right)}{I_w(y)} : s = y \text{ (truthful)} \\ \dfrac{P\left(Y_{w'} = y' \middle| Y_w = y\right)}{I_w(y')} : s = y' \neq y \end{cases}.$$

Result: SRA induces a Bayes-Nash Eq.

- Suppose all other agents $\neq w$ are truthful.

$$E[v_w(s)] = \begin{cases} \dfrac{P\left(Y_{w'} = y \middle| Y_w = y\right)}{I_w(y)} : s = y \text{ (truthful)} \\ \dfrac{P\left(Y_{w'} = y' \middle| Y_w = y\right)}{I_w(y')} : s = y' \neq y \end{cases}.$$

- Truthful reward $\geq$ False reward if

$$P\left(Y_w = Y_{w'} = y\right) P\left(Y_w = Y_{w'} = y'\right) \geq P\left(Y_{w'} = y', Y_w = y\right)^2$$

$$\Leftrightarrow \sum_x P_X(x) p_y^2(x) \sum_x P_X(x) p_{y'}^2(x) \geq \left(\sum_x P_X(x) p_y(x) p_{y'}(x)\right)^2$$

for every $y' \neq y$.

Result Cont'd: Strict Bayes Nash

$$P\left(Y_w = Y_{w'} = y\right)P\left(Y_w = Y_{w'} = y'\right) > P\left(Y_{w'} = y', Y_w = y\right)$$

- SRA is **strictly Bayes-Nash incentive compatible** *in expectation* if model is separable for some $\alpha > 0$.
  - Difference in reward for submitting true vs. false response is lower bounded by a value

$$\omega = \omega(T, \alpha, |\mathcal{Y}|, n)$$

  when T is sufficiently large.

Extensions Overview

- Ex-ante collaboration: Workers may agree to submit a strategy

$$\sigma_w(s|y) \in \Delta(\mathcal{Y})$$

before the game starts.

- Symmetric strategy: $\sigma_w = \sigma \ \forall w$.
  - Everyone submits $s = 1$ regardless of observation y.
- [Theorem] Truthful responses strictly maximize rewards over all symmetric strategies when #tasks $\to \infty$.

Collusion and Equilibrium

- Workers are not allowed to communicate throughout the game.
- It may be possible that workers communicate/collaborate beforehand to reduce effort while maintaining high reward.
- One such implication is a symmetric strategy profile, where all workers agree on a fixed modification $y \mapsto q(y)$.
  - For example, we want to get rid of cases when submitting trivial answers can achieve a high payoff. Otherwise, workers can receive high reward with low effort.

Collusion and Equilibrium

- Uninformativeness of a symmetric strategy profile q is define as

$$\Omega(q) = \frac{1}{|\mathcal{Y}|(|\mathcal{Y}|-1)} \sum_{y,y'} \sum_{y'':y''\neq y'} \sqrt{q_y(y')q_y(y'')}$$

and say that a strategy q is ω-uninformative if $\Omega(q) \geq \omega$.

  - $\Omega(q) = 0 \Leftrightarrow q$ is fully-informative: $q_{\cdot}(y)$ has disjoint supports.
  - $\Omega(q)$ is maximized if the report q is chosen independently of the true answer.

- Given enough tasks, a fully-informative strategy profile maximizes reward. In other words, SRA is strongly truthful across symmetric equilibria, asymptotically.

Strong Truthfulness Over All Equilibria

- Crowdsourcing host can choose how to assign tasks.
- Suppose $\frac{wN}{n} < M$ (fix to our notation).
- Asymmetric strategies that may arise (Sec. 5.3).
- State theorem 4.

Limitations and Our Extension

- Because workers are assumed to be statistically identical, a majority vote is the most accurate estimate of the underlying truth.
- When workers are heterogeneous, e.g. follow the Dawid-Skene model

$$p_w = P(Y_{wt} = y^*) \ \forall t,$$

then SRA fails to incentivize truthful behavior.

- Most algorithms that efficiently aggregate worker responses are designed for the Dawid-Skene model, and it is of interest how to incentivize truthful behavior under such settings.
- Mechanisms designed for heterogeneous workers require extraneous reports, i.e. is not minimal.

- We are currently designing a mechanism that incentivizes truthful response when workers are heterogeneous.

Summary

- SRA incentivizes truthful responses without requiring extraneous reports.
- Under SRA, an honest response maximizes workers’ rewards over symmetric equilibria.