Strategic Reviews

James Siderius joint with **Mohamed Mostagir**

(Preliminary Work)

INFORMS 2019
Seattle, WA, October 22, 2019

Reviews and Influence

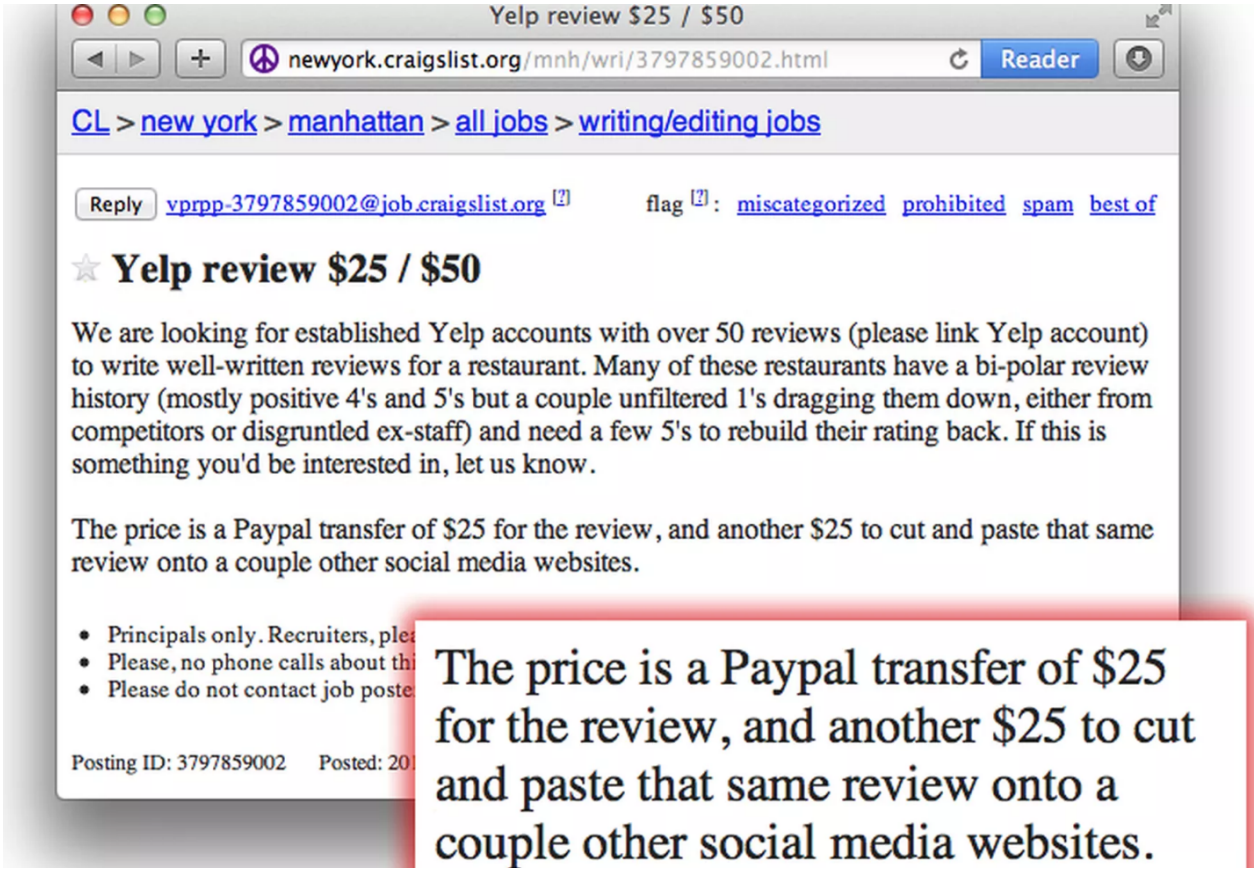

What We Do

- Consumers want to buy good products. Firms want consumers to buy *their* product. What role do reviewers play?
  - do reviewers have an incentive to bias reviews?
  - can consumers be influenced by these signals, even with these biases?
  - are firms willing to pay for this persuasion?

- Try to understand strategic incentives between reviewers and firms, and how influence arises endogenously through reputation.

- How can platform limit incentives to accept "bribes," lie about reviews, and lose influence.

- Today:
  - three-tier model of reviews
  - characterize how reviewers and consumers make decisions
  - investigate how firms can benefit with bribes

What We Do

- Consumers want to buy good products. Firms want consumers to buy *their* product. What role do reviewers play?
  - do reviewers have an incentive to bias reviews?
  - can consumers be influenced by these signals, even with these biases?
  - are firms willing to pay for this persuasion?

- Try to understand strategic incentives between reviewers and firms, and how influence arises endogenously through reputation.

- How can platform limit incentives to accept "bribes," lie about reviews, and lose influence.

- Today:
  - three-tier model of reviews
  - characterize how reviewers and consumers make decisions
  - investigate how firms can benefit with bribes

What We Do

- Consumers want to buy good products. Firms want consumers to buy *their* product. What role do reviewers play?
  - do reviewers have an incentive to bias reviews?
  - can consumers be influenced by these signals, even with these biases?
  - are firms willing to pay for this persuasion?

- Try to understand strategic incentives between reviewers and firms, and how influence arises endogenously through reputation.

- How can platform limit incentives to accept "bribes," lie about reviews, and lose influence.

- Today:
  - three-tier model of reviews
  - characterize how reviewers and consumers make decisions
  - investigate how firms can benefit with bribes

What We Do

- Consumers want to buy good products. Firms want consumers to buy *their* product. What role do reviewers play?
  - do reviewers have an incentive to bias reviews?
  - can consumers be influenced by these signals, even with these biases?
  - are firms willing to pay for this persuasion?

- Try to understand strategic incentives between reviewers and firms, and how influence arises endogenously through reputation.

- How can platform limit incentives to accept "bribes," lie about reviews, and lose influence.

- Today:
  - three-tier model of reviews
  - characterize how reviewers and consumers make decisions
  - investigate how firms can benefit with bribes

Model

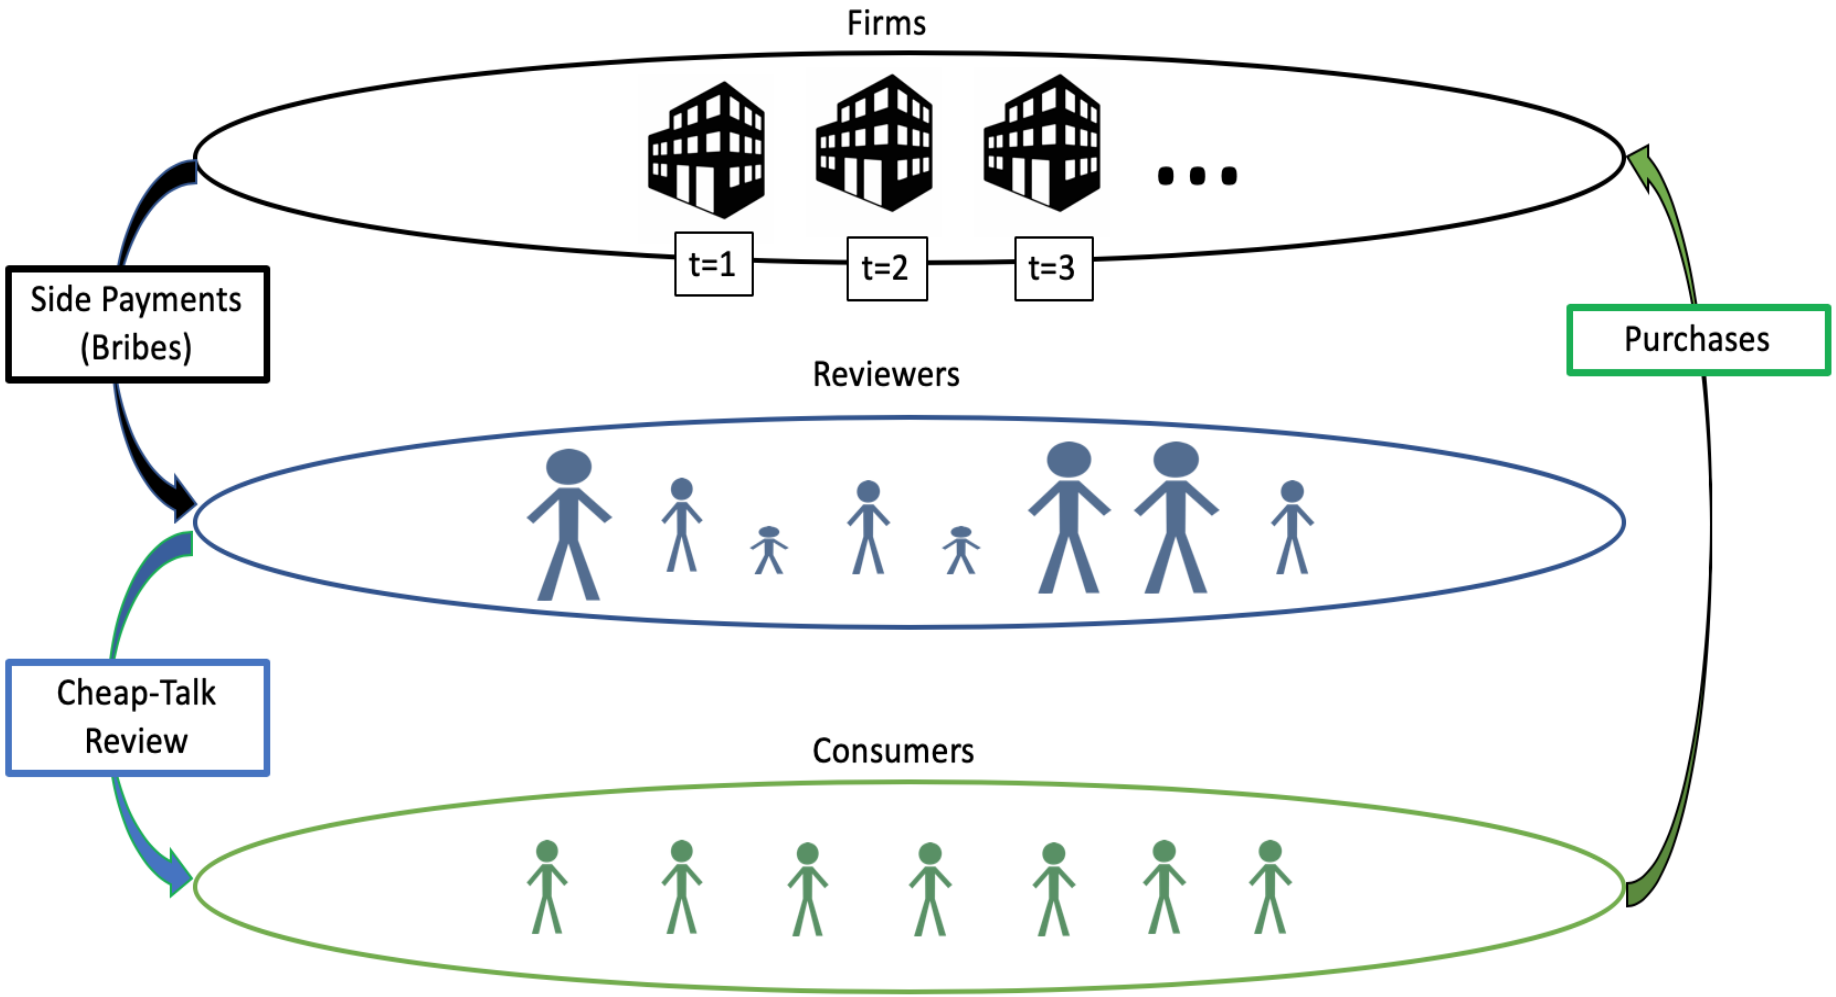
Firms
t=1
t=2
t=3
Side Payments
(Bribes)
Purchases
Reviewers
Cheap-Talk
Review
Consumers

Model: Players

(a) **Firms**: Arrive sequentially at each time $t = 1, 2, \ldots$ and live for only one period.

- Each new firm has a quality q_t of its product which is not known (with certainty) to anyone.

(b) Reviewers: The same set of agents over time who consume and review each of the products at all times t.

- Each reviewer j has a type ω_j which is either high-skill (H) or low-skill (L), where she is high-skill with probability p.
- High-skill types receive more precise signals of the product quality than do the low-skill types.
- For simplicity, firms and reviewers know all reviewers' skill types.

(c) Consumers: There is a continuum of consumers who have heterogenous preferences for quality. Formally, each consumer has an outside option ϕ_i which it can obtain instead of purchasing the product, where:

- ϕ_i is increasing in i with $\lim_{i \to 0} \phi_i = -\infty$ and $\lim_{i \to 1} \phi_i = \infty$.

Model: Players

(a) **Firms**: Arrive sequentially at each time $t = 1, 2, \ldots$ and live for only one period.

- Each new firm has a quality q_t of its product which is not known (with certainty) to anyone.

(b) Reviewers: The same set of agents over time who consume and review each of the products at all times t.

- Each reviewer j has a type ω_j which is either high-skill (H) or low-skill (L), where she is high-skill with probability p.
- High-skill types receive more precise signals of the product quality than do the low-skill types.
- For simplicity, firms and reviewers know all reviewers' skill types.

(c) Consumers: There is a continuum of consumers who have heterogenous preferences for quality. Formally, each consumer has an outside option ϕ_i which it can obtain instead of purchasing the product, where:

- ϕ_i is increasing in i with $\lim_{i \to 0} \phi_i = -\infty$ and $\lim_{i \to 1} \phi_i = \infty$.

Model: Players

(a) **Firms**: Arrive sequentially at each time $t = 1, 2, \ldots$ and live for only one period.

- Each new firm has a quality q_t of its product which is not known (with certainty) to anyone.

(b) Reviewers: The same set of agents over time who consume and review each of the products at all times t.

- Each reviewer j has a type ω_j which is either high-skill (H) or low-skill (L), where she is high-skill with probability p.
- High-skill types receive more precise signals of the product quality than do the low-skill types.
- For simplicity, firms and reviewers know all reviewers' skill types.

(c) Consumers: There is a continuum of consumers who have heterogenous preferences for quality. Formally, each consumer has an outside option ϕ_i which it can obtain instead of purchasing the product, where:

- ϕ_i is increasing in i with $\lim_{i \to 0} \phi_i = -\infty$ and $\lim_{i \to 1} \phi_i = \infty$.

Model: Information

- At each time t, the firm's quality and the signals $s_{j,t}$ of each reviewer t are drawn according to the following process:
  - Quality is drawn from the standard normal, $q_t \sim \mathcal{N}(0,1)$.
  - Each reviewer's signal is an unbiased, noisy signal of the quality q_t. High-skill reviewers have less noise around the truth than low-skill reviewers.

- Conditional on $s_{j,t}$, all reviewers j (simultaneously) send reviews $r_{j,t} \in \mathbb{R}$, which are publicly observable to all players.

- Each consumer chooses to either purchase the product ($x_{i,t} = 1$) or not ($x_{i,t} = 0$) at unit price. Consumers receive independent experiences $e_{i,t} = q_t + \eta_{i,t}$, for some noise term $\eta_{i,t}$, where $\eta_{i,t}$ are iid, distributed symmetrically around 0, and have finite variance.

Model: Information

- At each time t, the firm's quality and the signals $s_{j,t}$ of each reviewer t are drawn according to the following process:
  - Quality is drawn from the standard normal, $q_t \sim \mathcal{N}(0,1)$.
  - Each reviewer's signal is an unbiased, noisy signal of the quality q_t. High-skill reviewers have less noise around the truth than low-skill reviewers.

- Conditional on $s_{j,t}$, all reviewers j (simultaneously) send reviews $r_{j,t} \in \mathbb{R}$, which are publicly observable to all players.

- Each consumer chooses to either purchase the product ($x_{i,t} = 1$) or not ($x_{i,t} = 0$) at unit price. Consumers receive independent experiences $e_{i,t} = q_t + \eta_{i,t}$, for some noise term $\eta_{i,t}$, where $\eta_{i,t}$ are iid, distributed symmetrically around 0, and have finite variance.

Model: Information

- At each time t, the firm's quality and the signals $s_{j,t}$ of each reviewer t are drawn according to the following process:
  - Quality is drawn from the standard normal, $q_t \sim \mathcal{N}(0,1)$.
  - Each reviewer's signal is an unbiased, noisy signal of the quality q_t. High-skill reviewers have less noise around the truth than low-skill reviewers.

- Conditional on $s_{j,t}$, all reviewers j (simultaneously) send reviews $r_{j,t} \in \mathbb{R}$, which are publicly observable to all players.

- Each consumer chooses to either purchase the product ($x_{i,t} = 1$) or not ($x_{i,t} = 0$) at unit price. Consumers receive independent experiences $e_{i,t} = q_t + \eta_{i,t}$, for some noise term $\eta_{i,t}$, where $\eta_{i,t}$ are iid, distributed symmetrically around 0, and have finite variance.

Model: Payoffs and Bribes

- Consumers are myopic and maximize their current-period utility given posted reviews $\mathbf{r}_t$:

$$x^*_{i,t}(\mathbf{r}_t) = \arg \max_{x_{i,t} \in \{0,1\}} \mathbb{E}[(e_{i,t} - \phi_i) x_{i,t} | \mathbf{r}_t]$$

  which has the cutoff strategy $x^*_{i,t}(\mathbf{r}_t) = 1$ iff $\mathbb{E}[q_t | \mathbf{r}_t] \geq \phi_i$.

  - Let $X^*_t(\mathbf{r}_t)$ be the total amount of the product purchased, conditional on $\mathbf{r}_t$.
  - We can define the influence index $I_{j,t}$ of reviewer j at time t as:

$$I_{j,t} = \frac{\partial X^*_t(\mathbf{r}_t)}{\partial r_{j,t}}$$

  - Influence of reviewer j is the sensitivity of a consumer's decision from j's review.

- Assume **firm** t may offer a bribe schedule, $b_{j,t}(r_{j,t}) \geq 0$, for each reviewer j.

- Let $\mathbf{r}^*_t$ denote the reviews posted at time t. Then the payoff of **firm** t is given by:

$$U_t = X^*_t(\mathbf{r}_t) - \sum_{j=1}^{n} b_{j,t}(\mathbf{r}^*_t)$$

Model: Payoffs and Bribes

- Consumers are myopic and maximize their current-period utility given posted reviews $\mathbf{r}_t$:

$$x^*_{i,t}(\mathbf{r}_t) = \arg \max_{x_{i,t} \in \{0,1\}} \mathbb{E}[(e_{i,t} - \phi_i) x_{i,t} | \mathbf{r}_t]$$

  which has the cutoff strategy $x^*_{i,t}(\mathbf{r}_t) = 1$ iff $\mathbb{E}[q_t | \mathbf{r}_t] \geq \phi_i$.

  - Let $X^*_t(\mathbf{r}_t)$ be the total amount of the product purchased, conditional on $\mathbf{r}_t$.
  - We can define the influence index $I_{j,t}$ of reviewer j at time t as:

$$I_{j,t} = \frac{\partial X^*_t(\mathbf{r}_t)}{\partial r_{j,t}}$$

  - Influence of reviewer j is the sensitivity of a consumer's decision from j's review.

- Assume **firm** t may offer a bribe schedule, $b_{j,t}(r_{j,t}) \geq 0$, for each reviewer j.

- Let $\mathbf{r}^*_t$ denote the reviews posted at time t. Then the payoff of **firm** t is given by:

$$U_t = X^*_t(\mathbf{r}_t) - \sum_{j=1}^{n} b_{j,t}(\mathbf{r}^*_t)$$

Model: Payoffs and Bribes

- Consumers are myopic and maximize their current-period utility given posted reviews $\mathbf{r}_t$:

$$x^*_{i,t}(\mathbf{r}_t) = \arg \max_{x_{i,t} \in \{0,1\}} \mathbb{E}[(e_{i,t} - \phi_i) x_{i,t} | \mathbf{r}_t]$$

  which has the cutoff strategy $x^*_{i,t}(\mathbf{r}_t) = 1$ iff $\mathbb{E}[q_t | \mathbf{r}_t] \geq \phi_i$.

  - Let $X^*_t(\mathbf{r}_t)$ be the total amount of the product purchased, conditional on $\mathbf{r}_t$.
  - We can define the influence index $I_{j,t}$ of reviewer j at time t as:

$$I_{j,t} = \frac{\partial X^*_t(\mathbf{r}_t)}{\partial r_{j,t}}$$

  - Influence of reviewer j is the sensitivity of a consumer's decision from j's review.

- Assume **firm** t may offer a bribe schedule, $b_{j,t}(r_{j,t}) \geq 0$, for each reviewer j.

- Let $\mathbf{r}^*_t$ denote the reviews posted at time t. Then the payoff of **firm** t is given by:

$$U_t = X^*_t(\mathbf{r}_t) - \sum_{j=1}^{n} b_{j,t}(\mathbf{r}^*_t)$$

, cont.

- Some reviewers are behavioral and always truthfully report (with small probability $\epsilon > 0$) (i.e., $r_{j,t} = s_{j,t}$). Otherwise, the reviewer is strategic.

- Assume reviewers also care about their influence over consumers' decisions according to propensity $\beta > 0$ (e.g., status or preferential treatment).

- Strategic reviewers are infinitely patient and maximize their average payoff:

$$V_j = \sum_{t=0}^{\infty} \delta^t (\beta_j I_{j,t} + b_{j,t})$$

  where we have suppressed the RHS dependence on the history of reviews.

- Classify "pure-strategy" perfect Bayesian equilibria as $\delta \to 1$ (where reviewers are infinitely patient).

, cont.

- Some reviewers are behavioral and always truthfully report (with small probability $\epsilon > 0$) (i.e., $r_{j,t} = s_{j,t}$). Otherwise, the reviewer is strategic.

- Assume reviewers also care about their influence over consumers' decisions according to propensity $\beta > 0$ (e.g., status or preferential treatment).

- Strategic reviewers are infinitely patient and maximize their average payoff:

$$V_j = \sum_{t=0}^{\infty} \delta^t (\beta_j I_{j,t} + b_{j,t})$$

  where we have suppressed the RHS dependence on the history of reviews.

- Classify "pure-strategy" perfect Bayesian equilibria as $\delta \to 1$ (where reviewers are infinitely patient).

, cont.

- Some reviewers are behavioral and always truthfully report (with small probability $\epsilon > 0$) (i.e., $r_{j,t} = s_{j,t}$). Otherwise, the reviewer is strategic.

- Assume reviewers also care about their influence over consumers' decisions according to propensity $\beta > 0$ (e.g., status or preferential treatment).

- Strategic reviewers are infinitely patient and maximize their average payoff:

$$V_j = \sum_{t=0}^{\infty} \delta^t (\beta_j I_{j,t} + b_{j,t})$$

  where we have suppressed the RHS dependence on the history of reviews.

- Classify "pure-strategy" perfect Bayesian equilibria as $\delta \to 1$ (where reviewers are infinitely patient).

, cont.

- Some reviewers are behavioral and always truthfully report (with small probability $\epsilon > 0$) (i.e., $r_{j,t} = s_{j,t}$). Otherwise, the reviewer is strategic.

- Assume reviewers also care about their influence over consumers' decisions according to propensity $\beta > 0$ (e.g., status or preferential treatment).

- Strategic reviewers are infinitely patient and maximize their average payoff:

$$V_j = \sum_{t=0}^{\infty} \delta^t (\beta_j I_{j,t} + b_{j,t})$$

  where we have suppressed the RHS dependence on the history of reviews.

- Classify "pure-strategy" perfect Bayesian equilibria as $\delta \to 1$ (where reviewers are infinitely patient).

Summary: Timing

For each time t:

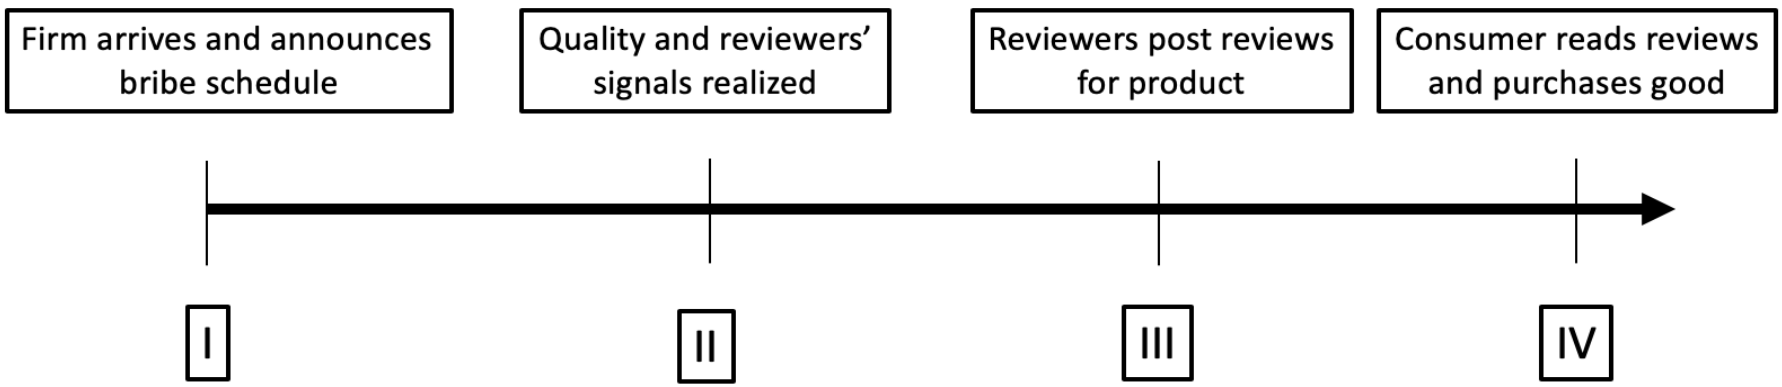

Try to understand equilibrium behavior for all $t > T$, for some large T (referred to as "eventually").

Babbling-Trigger Equilibrium

- Babbling equilibrium: Consumers ignore the reviews and instead choose X_t^* according to their prior, $X_t^* = \phi^{-1}(0)$.
  - On-path play: Reviewers abstain or send defunct signals and consumers do not listen.
  - Off-path play: Consumers assume any posted reviews are spurious and do not reflect true quality.

- When is babbling an equilibrium?
  - Because small probability $\epsilon > 0$ the reviewer will post $r_{j,t} = s_{j,t}$ always, babbling forever is not a best-response from the beginning (for a consumer).
  - Once a reviewer has been determined to not be truthtelling type (with sufficiently high probability), consumers credibly can switch to babbling.

- Babbling-trigger represents a harsh consumer who punishes reviewers who, beyond a reasonable doubt, can be identified as untruthful in their reviews.
  - This equilibrium provides a best-case scenario for the efficacy of the platform.

Reviewer-Consumer Reputation Game

- Assume firms do not offer any bribes; that is, the bribe schedules are given exactly by $b_{j,t}(r_{j,t}) \equiv 0$ for all j, t.

Theorem

In the babbling-trigger equilibrium, all reviewers are eventually honest (i.e., $r_{j,t} = s_{j,t}$) and consumers eventually infer the true types of every reviewer.

- Consumers eventually use the inverse-variance weighted average to infer expected quality:

$$\mathbb{E}[q_t|\mathbf{r}_t] = \frac{\sum_{j=1}^{n} r_{j,t}/\sigma_{\omega_j}^2}{1 + \sum_{j=1}^{n} 1/\sigma_{\omega_j}^2}$$

  and where $X_t^*(\mathbf{r}_t) = \phi^{-1}(\mathbb{E}[q_t|\mathbf{r}_t])$. Influence index is higher for high-skill types.

- Intuition: Suppose I'm a low-skill type $\sigma_L = 100$ and want to match high-skill type $\sigma_H = 1$.
  - Correlation between my s_t and q_t will be $1/101$ instead of $1/2$.
  - Can I improve my correlation by biasing my s_t?

Fixed Bribe Schedules

- Suppose we take the bribe schedules $b_{j,t}(r_{j,t})$ as given but not necessarily equal to zero. How does the equilibrium change?

Theorem

When bribe schedules are fixed, in the babbling-trigger equilibrium every low-skill type reviewer eventually reports truthfully ($r_{j,t} = s_{j,t}$). On the other hand, every high-skill reviewer eventually either: (i) reports truthfully ($r_{j,t} = s_{j,t}$) or (ii) plays a strategy where $r_{j,t} = s_{j,t} + \varepsilon'_{j,t}$ with $\mathbb{E}[\varepsilon'_{j,t}] = 0$, $\mathbb{E}[(\varepsilon'_{j,t})^2] = \sigma_L^2 - \sigma_H^2$, and $\varepsilon'_{j,t} \perp s_{j,t}$.

Key Takeaway: High-skill reviewer can choose to mimic a low-skill reviewer. The reviewer accepts the bribe, and pretends to "inject noise" into his review but instead biases her signal.

- Tradeoff between bribe payment and influence: high-skill reviewer will have the influence of a low-skill reviewer, but receive payments.
- Consumers still make purchase decisions in the *exact same way*!
- But, valuable information from reviewer to consumer destroyed.

Strategic Firms

- Bribe schedules are endogenous decisions on the part of the firm.
  - Reputation game between firm and reviewer as well.
  - Reviewer's influence has more than just intrinsic value, as influence translates into payoffs for firm.

- For example, take β to be very small (i.e., little intrinsic value derived for influence). Should the reviewer accept large bribes in exchange for very biased reviews?
  - No, in the babbling-trigger equilibrium, eventually this reviewer will lose all influence. No future firms will offer bribes.

- Each entering firm must choose $b_1(r_1), \cdots, b_n(r_n)$, then reviewers observe signals, and post reviews.
  - Look for a stationary equilibrium where all firms post the same bribe schedules and high-skill reviewers decide whether to report truthfully or mimic low-skill.
  - Consumers still continue to use inverse-variance weights to make purchase decisions.

Complementarity of Bribes

- Assume there are just two reviewers both with high-precision. How does bribe schedule $b_1(r_1)$ affect the decisions of reviewers 1 and 2 to mimic low-precision?
  - Clearly, increasing the slope of $b_1(r_1)$ will increase incentives for reviewer 1 to mimic imprecision.
  - But, an ambiguous effect on reviewer 2...

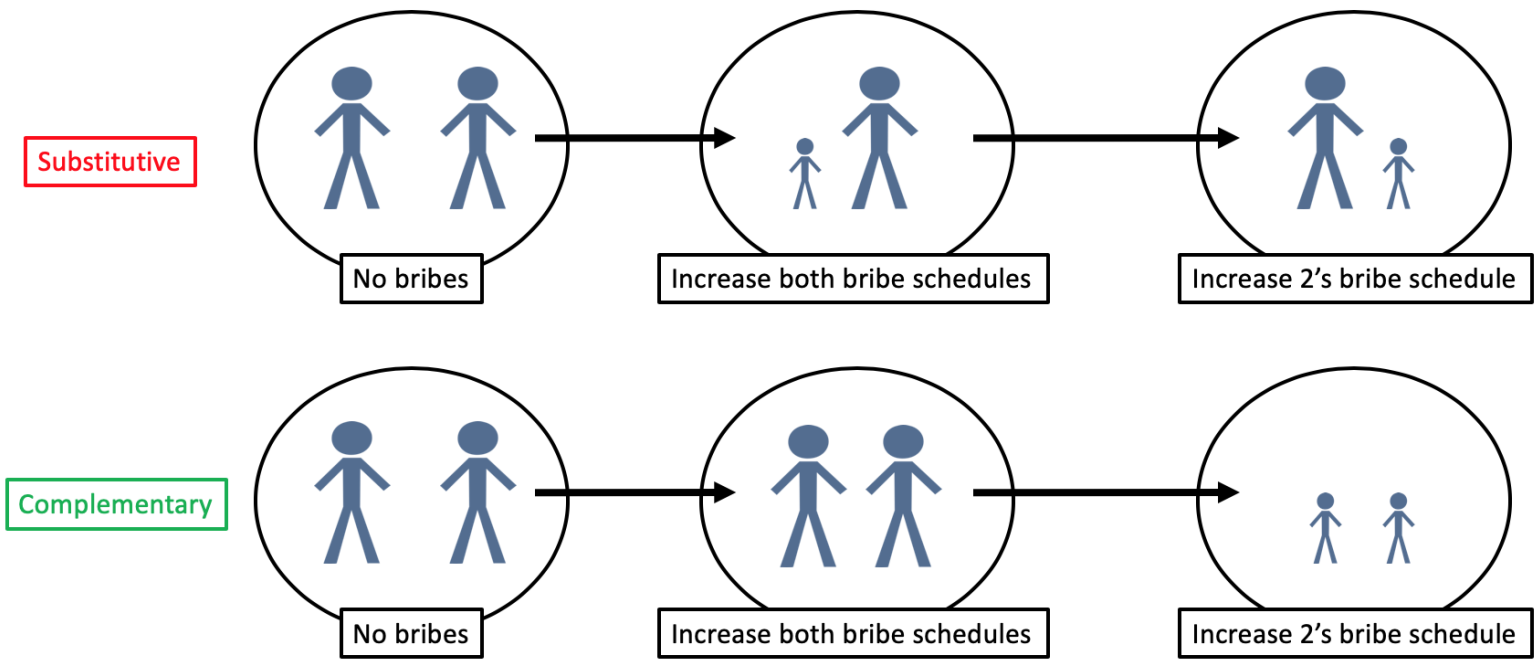

Platform Incentives and Ongoing Work

- Characterization of firm's optimal bribe schedule, given reputation game between reviewers and consumer.

- Optimal policy: Can the platform reward influence (and push up β for some reviewers)?
  - Decrease incentives to mimic low-skill and instead report truthfully.
  - Externalities: possible my decision to report truthfully can nudge others to do the same.
  - Relatively inexpensive revenue-sharing can restore substantial amounts of information on the platform.