Clicks or Comments?
The quality-quantity trade-off of review systems

Matthew J. H. Murphy
University of Pennsylvania

November 16, 2025

Designing review systems: Quality vs. Frequency

- Platforms use reviews to make decisions. How should they design their review systems?

The trade-offs are natural in this context:

- The more detailed the review system, the more information is being extracted from each individual reviewer.
- It is less costly for reviewers to leave a binary review than a full review.
  - This leads to more simple reviews being reported than detailed reviews.

Is it possible to quantify the trade-off between the informativeness of reviews and how frequently they are submitted?

Different Review Systems

5-Star Reviews:

- Uber, Airbnb, Online Retailers.
- Often exhibit a “J”-shaped distribution.

Different Review Systems

5-Star Reviews:

- Uber, Airbnb, Online Retailers.
- Often exhibit a “J”-shaped distribution.

3-Star Reviews:

- Facebook Messenger, Discord.

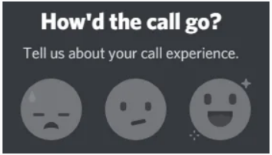

Different Review Systems

5-Star Reviews:

- Uber, Airbnb, Online Retailers.
- Often exhibit a “J”-shaped distribution.

3-Star Reviews:

- Facebook Messenger, Discord.

Binary Reviews:

- Rotten Tomatoes, ebay, Discogs.

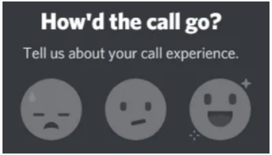

Different Review Systems

5-Star Reviews:

- Uber, Airbnb, Online Retailers.
- Often exhibit a “J”-shaped distribution.

3-Star Reviews:

- Facebook Messenger, Discord.

Binary Reviews:

- Rotten Tomatoes, ebay, Discogs.

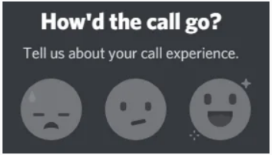

Model

Model: Reviewers' information

There are two possible values of product quality $\theta \in \Theta = \{L, H\}$ with prior $q = \mathbb{P}(L)$.

- There are N "reviewers". Reviewer i has a (conditionally independent) private signal $S_i \in \mathbb{R}$ about θ (realization s_i).

$$s_i = u_\theta + \sigma\epsilon_i, \quad \text{where} \quad u_\theta = \begin{cases} -\mu & \text{if} \quad \theta = L \\ \mu & \text{if} \quad \theta = H \end{cases}$$

- Interpret s_i as an average utility u_θ with some taste shock ϵ_i for reviewer i.
- The noise σ controls the precision of reviews.

Assumption

The ϵ_i are i.i.d., with continuous, full-support, log-concave density given by f (with CDF F).

Model: The platform's problem & review systems

- The platform has a state-dependent utility function $u : \mathcal{A} \times \Theta \to \mathbb{R}$, with action set $\mathcal{A}$ finite.
- It chooses a review system $r : \mathbb{R} \to \mathbb{R}$ that maps signals to reviews.
    - E.g., the full review system: $r_f(s) \coloneqq s$; or
    - A binary threshold system: $r_\tau(s) \coloneqq \mathbb{1}\{s > \tau\}$, for some $\tau \in \mathbb{R}$.
- Let $\mathcal{R}_r = \text{supp}(r(S_i))$ denote the set of possible reviews under r.
- Each review system r has an associated *probability of review* p_r, that is decreasing in $|\mathcal{R}_r|$.
    - More complicated reviews are less likely to be reported.
    - For today, independent of reviewers' signals.

Model: Timing

- The platform chooses a review system.
- Reviewers observe private signals s_i and decide whether to submit.
- The platform observes submitted reviews $\mathcal{S}$, updates beliefs to posterior $\pi(r, \mathcal{S})$ that $\theta = L$, and takes an action to maximize its expected utility:

$$u^*(r, \mathcal{S}) := \max_{a \in \mathcal{A}} \left\{ \pi(r, \mathcal{S}) u(a, L) + (1 - \pi(r, \mathcal{S})) u(a, H) \right\}.$$

- The platform aims to maximize $u^*(r, N) := \mathbb{E}\left[u^*(r, \mathcal{S})\right]$ by choosing r:

$$\max_r u^*(r, N).$$

- I focus on the case when N is large, to emulate online reviews.

Trading off Quality and Frequency

Quality of review systems

Each review system r induces a distribution over reviews in state L and state H. Define the measure γ_L^r over $\mathcal{R}_r$ as, for any measurable $B \subseteq \mathcal{R}_r$,

$$\gamma_L^r(B) \coloneqq \mathbb{P}_L(r^{-1}(B)) = \int_{s:r(s)\in B} \frac{1}{\sigma} f\left(\frac{s+\mu}{\sigma}\right) ds. \quad \text{Similarly define } \gamma_H^r.$$

Quality of review systems

Each review system r induces a distribution over reviews in state L and state H. Define the measure γ_L^r over $\mathcal{R}_r$ as, for any measurable $B \subseteq \mathcal{R}_r$,

$$\gamma_L^r(B) \coloneqq \mathbb{P}_L(r^{-1}(B)) = \int_{s:r(s)\in B} \frac{1}{\sigma} f\left(\frac{s+\mu}{\sigma}\right) ds. \quad \text{Similarly define } \gamma_H^r.$$

Definition

Given a review system r, define its *learning efficiency* as

$$\nu(r) \coloneqq 1 - \min_{\lambda\in[0,1]} \int_{\mathcal{R}_r} \left(\frac{d\gamma_L^r}{d\gamma_H^r}\right)^{\lambda} d\gamma_H^r.$$

Trading off between frequency and quality

- The probability of review p_r also impacts platform's learning.

Trading off between frequency and quality

- The probability of review p_r also impacts platform's learning.

Theorem

Fix a decision problem and two review systems (r, p_r) and $(r', p_{r'})$. If

$$p_r \nu(r) > p_{r'} \nu(r')$$

there exists an $\overline{N}$ such that for all $N > \overline{N}$, $u^*(r, N) > u^*(r', N)$.

- Implies that platform separably trades off between quality and frequency of review systems.

The relative information of review systems

- Can re-frame the result as

$$p_r \nu(r) > p_{r'} \nu(r') \iff \frac{\nu(r)}{\nu(r')} > \frac{p_{r'}}{p_r}.$$

- Since platforms can compute p_r for different review systems, this suggests a focus on the ratio $\nu(r)/\nu(r')$.
- Specifically, can focus on the information contained in a review system relative to the full review system $r_f(s) = s$.

The relative information of review systems

- Can re-frame the result as

$$p_r \nu(r) > p_{r'} \nu(r') \iff \frac{\nu(r)}{\nu(r')} > \frac{p_{r'}}{p_r}.$$

- Since platforms can compute p_r for different review systems, this suggests a focus on the ratio $\nu(r)/\nu(r')$.
- Specifically, can focus on the information contained in a review system relative to the full review system $r_f(s) = s$.

Definition

Given a review system r, define its *relative information* as

$$\kappa(r) := \frac{\nu(r)}{\nu(r_f)} \in [0, 1].$$

Imprecise Signals and the Optimal Binary Review System

Imprecise signals

A key feature of online reviews is that a single review is basically uninformative. This corresponds to the case when the noise in reviewers' signals σ is large.

Imprecise signals

A key feature of online reviews is that a single review is basically uninformative. This corresponds to the case when the noise in reviewers' signals σ is large.

- Normalize signals $s \mapsto \frac{s}{\sigma}$ to simplify notation.

Goal: characterize the optimal binary review system when signals are imprecise, and determine when it outperforms the full review system.

Characterizing the optimal binary review system

Call r a *threshold system* if $r(s) \coloneqq \mathbb{1}\{s > \tau\}$ for some $\tau \in \mathbb{R}$. Denote by r_τ.

Theorem

The optimal binary review system when signals are imprecise is given by the threshold review system $r_{\tau^*}(s) \coloneqq \mathbb{1}\{s > \tau^*\}$

Characterizing the optimal binary review system

Call r a *threshold system* if $r(s) := \mathbb{1}\{s > \tau\}$ for some $\tau \in \mathbb{R}$. Denote by r_τ.

Theorem

The optimal binary review system when signals are imprecise is given by the threshold review system $r_{\tau^*}(s) := \mathbb{1}\{s > \tau^*\}$, where threshold τ^* solves

$$\tau^* = \operatorname*{argmax}_{\tau} \left\{ \frac{f(\tau)^2}{F(\tau)} + \frac{f(\tau)^2}{1 - F(\tau)} \right\}.$$

Characterizing the optimal binary review system

Call r a *threshold system* if $r(s) := \mathbb{1}\{s > \tau\}$ for some $\tau \in \mathbb{R}$. Denote by r_τ.

Theorem

The optimal binary review system when signals are imprecise is given by the threshold review system $r_{\tau^*}(s) := \mathbb{1}\{s > \tau^*\}$, where threshold τ^* solves

$$\tau^* = \operatorname*{argmax}_{\tau} \left\{ \frac{f(\tau)^2}{F(\tau)} + \frac{f(\tau)^2}{1 - F(\tau)} \right\}.$$

This review system is preferred to the full review system if

$$\underbrace{\left(\frac{f(\tau^*)^2}{F(\tau^*)} + \frac{f(\tau^*)^2}{1 - F(\tau^*)} \right) \cdot \frac{1}{I_f}}_{\text{relative information}} > \underbrace{\frac{p_{r_f}}{p_{r_{\tau^*}}}}_{\text{ratio of review frequencies}},$$

where $I_f := \mathbb{E}\left[\left(\frac{f'(s)}{f(s)}\right)^2\right]$ denotes the Fisher Information of f.

Intuition

$$\left(\frac{f(\tau)^2}{F(\tau)} + \frac{f(\tau)^2}{1-F(\tau)}\right) \cdot \frac{1}{I_f}$$

Intuition

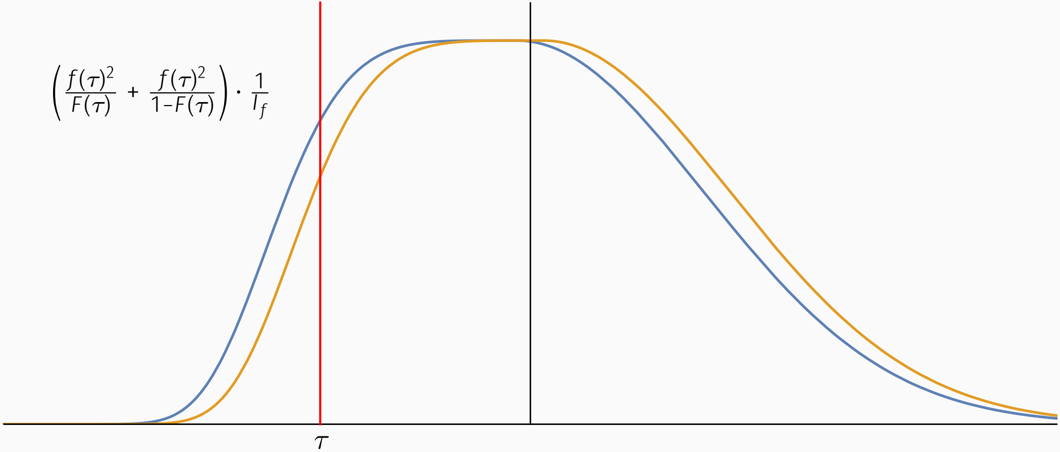
$\left(\frac{f(\tau)^2}{F(\tau)} + \frac{f(\tau)^2}{1-F(\tau)}\right) \cdot \frac{1}{I_f}$
τ

Taste Heterogeneity

Taste heterogeneity

I next study the properties of the optimal binary review system as a function of reviewer taste heterogeneity.

- Heterogeneous tastes $\Rightarrow$ extreme tastes are common (e.g., movies).
- Homogeneous tastes $\Rightarrow$ extreme tastes are uncommon (e.g., toasters).

Taste heterogeneity

I next study the properties of the optimal binary review system as a function of reviewer taste heterogeneity.

- Heterogeneous tastes $\Rightarrow$ extreme tastes are common (e.g., movies).
- Homogeneous tastes $\Rightarrow$ extreme tastes are uncommon (e.g., toasters).

When reviewers are heterogeneous, their idiosyncratic tastes have larger tails.

- I focus on a specific functional form for f:

$$f_\alpha(x) \propto \exp(-|x|^\alpha), \quad \text{for} \quad \alpha \geq 1.$$

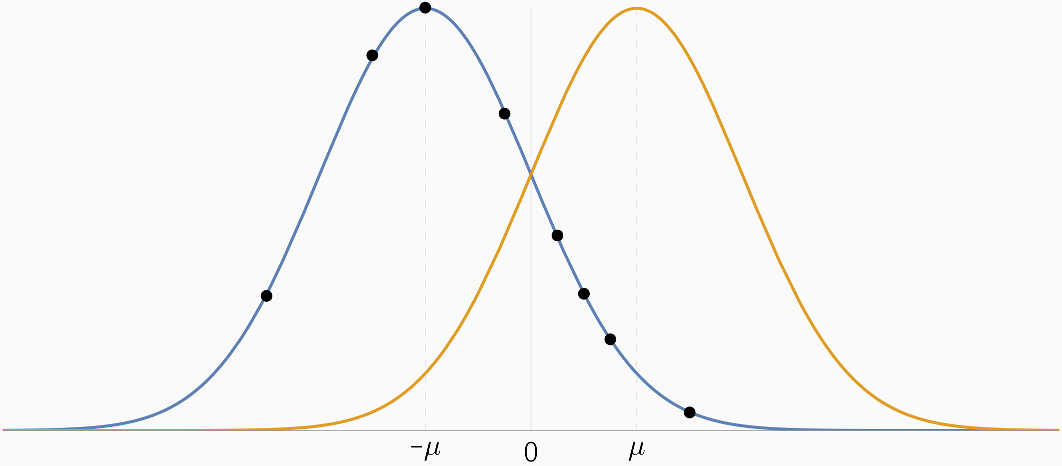

Figure 3: Illustration of possible released signals.

Illustration of different review systems

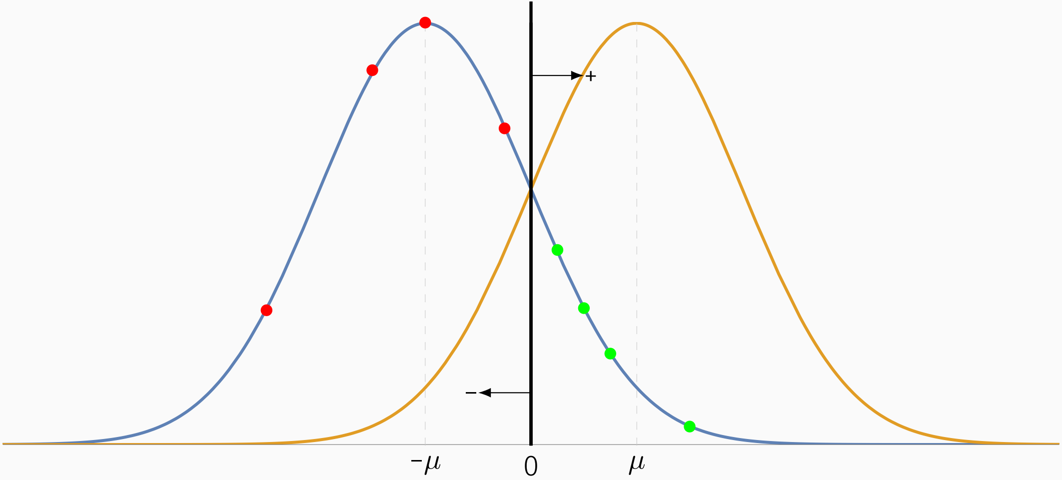

Figure 3: Illustration of the "good"/"bad" review system r_0.

Illustration of different review systems

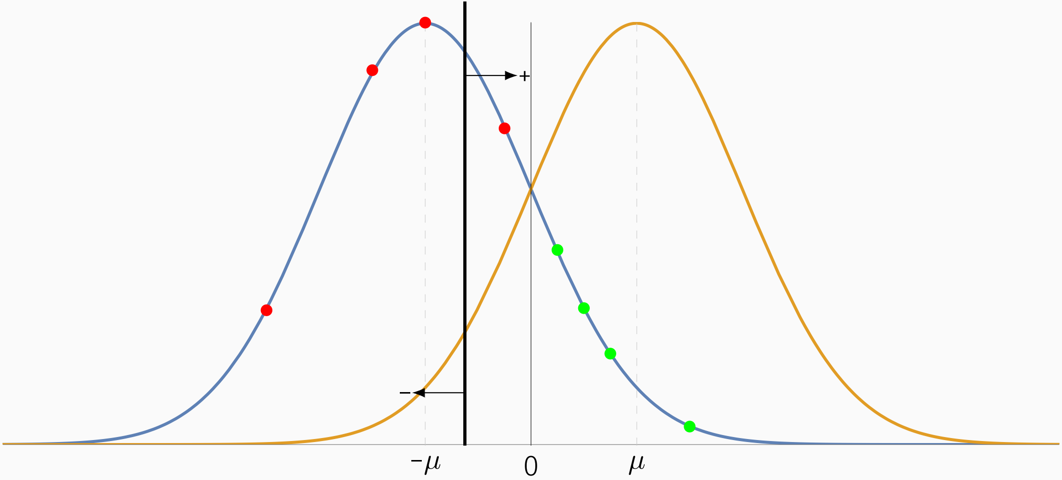

Figure 3: Illustration of the “horrible” review system r_τ, $\tau < 0$.

Illustration of different review systems

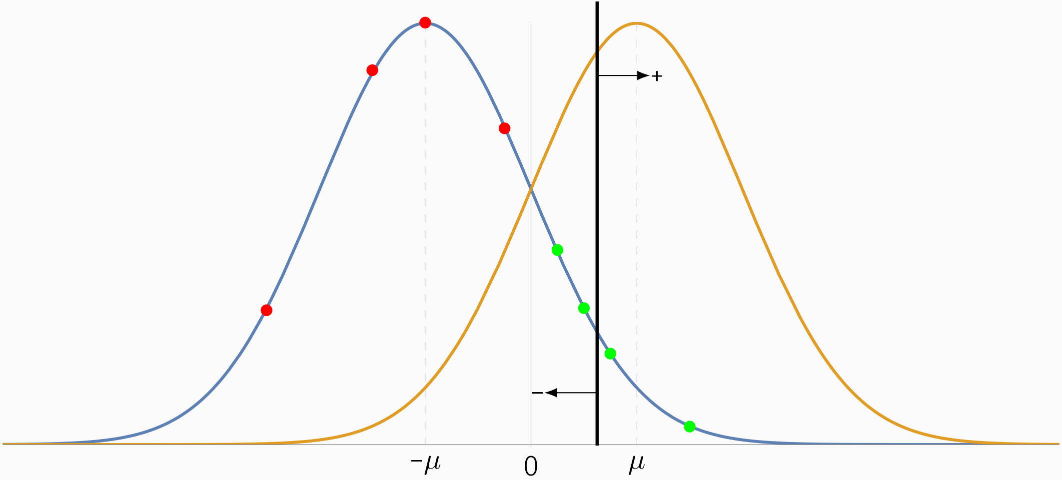

Figure 3: Illustration of the “amazing” review system r_τ, $\tau > 0$.

The optimal binary review system

Proposition

(i) The "good"/"bad" review system r_0 is optimal iff tastes are very heterogeneous ($\alpha \leq 2$).

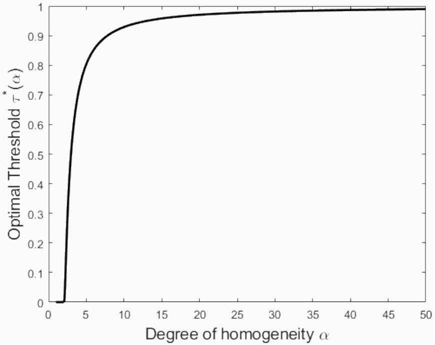

The optimal binary review system

Proposition

(i) The "good"/"bad" review system r_0 is optimal iff tastes are very heterogeneous ($\alpha \leq 2$).

(ii) As tastes become more homogeneous (α increases), the optimal review system is more "horrible" or "amazing".

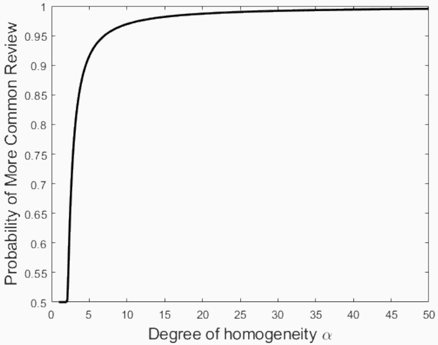

The optimal binary review system

Proposition

(i) The "good"/"bad" review system r_0 is optimal iff tastes are very heterogeneous ($\alpha \leq 2$).

(ii) As tastes become more homogeneous (α increases), the optimal review system is more "horrible" or "amazing".

(iii) The optimal binary review system is preferred to the full review system if $p_{r_\tau} > 3.25 p_{r_f}$, for all α.

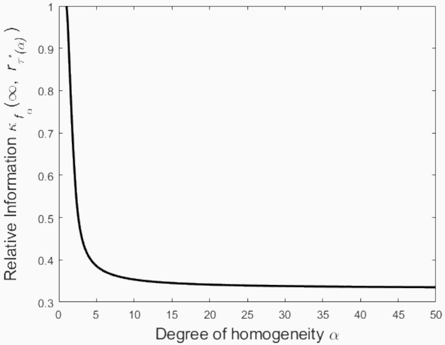

The optimal binary review system

Proposition

(i) The "good"/"bad" review system r_0 is optimal iff tastes are very heterogeneous ($\alpha \leq 2$).

(ii) As tastes become more homogeneous (α increases), the optimal review system is more "horrible" or "amazing".

(iii) The optimal binary review system is preferred to the full review system if $p_{r_\tau} > 3.25 p_{r_f}$, for all α.

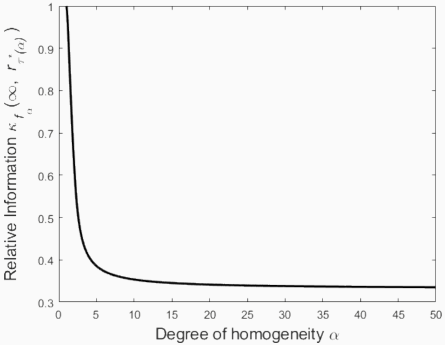

Both the "horrible" and "amazing" asymmetric review systems are optimal.

When to use the “horrible” review system over the “amazing”

If one side of the distribution is more homogeneous than the other, the platform should always set the threshold to isolate the homogeneous side.

Proposition

If negative signals are more homogeneous than positive signals ($\alpha_L > \alpha_H$), the optimal binary review system is “horrible”.

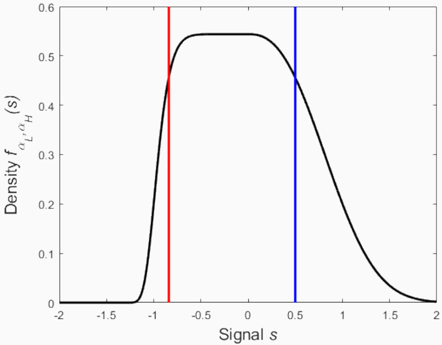

Conclusion

- I study how a platform designs optimal review systems, trading off between the informativeness of reviews and how frequently they are submitted.
- I study how reviewer heterogeneity affects the performance of simple review systems.
- My results suggest an explanation for why positively-skewed review systems are common in certain settings, and not in others.

In the paper:

- Beyond binary review systems.
- Dependence of reporting rates on reviewers’ signals.

Extensions

3 Bin Comparative Statics

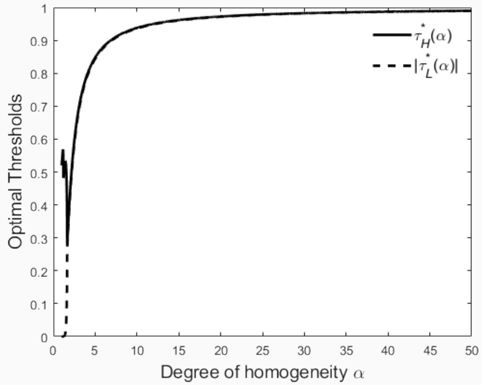

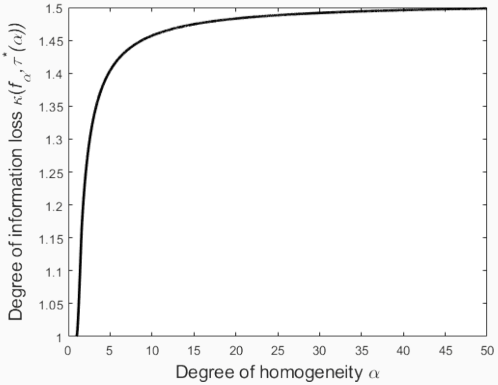

Non-Uniform Reporting Rates: Notation

Suppose now that the probability that a reviewer decides to leave a review is given by

$$p(r,s) = p_r q(s).$$

Now, define the distribution of reviews in state θ as

$$\gamma_L^{(p,r)}(B) := p_r \int_{s:r(s)\in B} q(s) f_\theta(s) ds$$

Then the generalization of $p_r \nu(r)$ is given by

$$\rho(\mu; r, p) := 1 - \min_{\lambda\in[0,1]} \left[\int_{R_r} \left(d\gamma_L^{(p,r)}\right)^\lambda \left(d\gamma_H^{(p,r)}\right)^{1-\lambda} + \left(\gamma_L^{(1-p,r)}(\mathbb{R})\right)^\lambda \left(\gamma_H^{(1-p,r)}(\mathbb{R})\right)^{1-\lambda} \right].$$

Generalization of Learning loss near Uninformative Signals

$$E(g) := \int_{-\infty}^{\infty} g(s)q(s)f(s)ds \quad \text{and} \quad E(g;\tilde{r}) := \int_{-\infty}^{\infty} g(s)\mathbb{1}\{r(s) \in \tilde{r}\}q(s)f(s)ds$$

Theorem

Suppose that $p(r,s) = p_r q(s)$ where q is continuous. Let $\ell_f := \frac{f'}{f}$ denote the linear score of f. Then learning-loss for a finite review system r is given by

$$\frac{p_r}{p_f} \cdot \lim_{\mu \to 0} \frac{\rho(\mu; p_f \cdot q)}{\rho_q(\mu; r, p_r \cdot q)} = \frac{E\left(\ell_f^2\right) + \frac{p_f E(\ell_f)^2}{1-p_f E(1)}}{\sum_{\tilde{r} \in R_r} \frac{E(\ell_f;\tilde{r})^2}{E(1;\tilde{r})} + \frac{p_r E(\ell_f)^2}{1-p_r E(1)}}.$$

Null Reviews are informative, but not as $p_r \to 0$

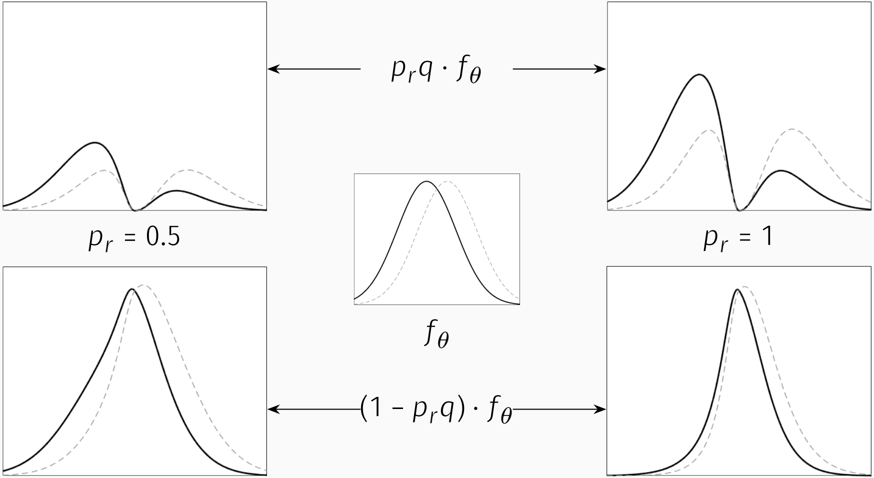

Figure 5: Example with $q(s) = 1 - \frac{1}{1+2(x^2+4\mathbb{1}\{x\leq 0\}x^2)}$.

Existing Information

We have thus far focused on the decision of the platform when eliciting information from reviewers. What if this information has already been collected?

Lemma

Fix a consumer with a decision problem, and two review systems r, r'. Then there exists an n large such that for all $n_r, n_{r'} > n(\mu)$, n_r reviews of type r is preferred to $n_{r'}$ reviews of type r' if

$$\frac{n_r}{n_{r'}} > \frac{\log(1 - \nu(\mu; r'))}{\log(1 - \nu(\mu; r))}.$$

Relationship between Production and Consumption of Information

We have the following:

Lemma

Let $\zeta(\mu) := \frac{\nu(\mu;r')}{\nu(\mu;r)}$ and let $\iota(\mu) := \frac{\log(1-\nu(\mu;r'))}{\log(1-\nu(\mu;r))}$. Then, we have that

$$\iota(\mu) > \zeta(\mu) \quad \text{if and only if} \quad \nu(\mu;r') > \nu(\mu;r).$$

Moreover, we have that

$$\lim_{\mu\to 0} \zeta(\mu) = \lim_{\mu\to 0} \iota(\mu).$$