

**Carleton College Department of Economics
Working Paper Series**



*Quantity Versus Quality: Experimenting with
the Margins for Social Information*

Jonathan Lafky and Alistair Wilson

No. 2018-02

Department of Economics
Carleton College
One North College Street
Northfield, MN 55057
Telephone: (507) 222-4109

September 2018

This paper represents the views of the author and does not necessarily reflect the opinion of Carleton College.

QUANTITY VERSUS QUALITY: EXPERIMENTING WITH THE MARGINS FOR SOCIAL INFORMATION

JONATHAN LAFKY AND ALISTAIR J. WILSON

ABSTRACT. People share their experiences of goods and services online, through reviews, ratings and endorsements on social networks, potentially generating welfare-improving information that can help subsequent consumers make better, more informed decisions. While the economics literature has focused on questions of alignment and the intensive quality of provided information, another tension is extensive: in the absence of an incentive, many might choose not to provide information at all. We study three different incentives that encourage information transmission on the extensive margin, examining the tradeoffs between quality and quantity of information. Our findings indicate substantial efficiency gains can be made relative to no incentives, even when the incentives damage the preference alignment between those sending and receiving information. In particular, our results point to a partially aligned incentive (similar to a referral or sales commission) as robustly encouraging the provision of information while not producing substantial reductions in quality.

The sharing of information via reviews, ratings and social networks has become an important component of online commerce, helping to reduce uncertainty about unseen products, services, and workers. Although consumers benefit from accurate information on product quality, even small nuisance costs can deter others from providing their opinions. An obvious solution to this problem is to reward those providing information. However, dependent on the parties shaping these incentives, there is the potential to influence not only the likelihood of information being provided, but also the content of that information. In this paper, we theoretically and experimentally examine the effectiveness of different incentives structures in the presence of this quality versus quantity tradeoff.

Focusing on the broad strategic tensions, we examine three incentive schemes that represent qualitatively different ways of paying for information provision. The three environments vary over the conditions (and motivating source) for a payment, changing the alignment of preference between those sending and receiving information. The end effect is an incentive that varies between common interests, only partial alignment of interests, and no alignment at all. Examining how each incentive influences both the quantity and quality of information, first in theory and then in practice

Date: August, 2018.

Lafky: Carleton College, 321 Willis Hall, Northfield, MN; jlafky@carleton.edu. *Wilson:* University of Pittsburgh, Department of Economics, 230 Bouquet Street, Pittsburgh, PA; alistair@pitt.edu. Our thanks to the following: John Asker, John Duffy, Matthew Embrey, Emanuel Vespa, Lise Vesterlund, Stephanie Wang, and audiences at the ESA, SEA, George Mason, Lafayette and the Federal Trade Commission. Any mistakes within are obviously not attributable to anyone but ourselves.

in the laboratory, we explain how the intensive and extensive margins of information transmission might be efficiently traded off.

Our attention to both the intensive and extensive margins for information-sharing is in contrast to the literature in economics, which generally focuses on the intensive margin. The literature typically measures the amount of information shared as the interests of the informed and uninformed parties grow further apart, where the provision of some form of information is taken for granted. While the comparative static here is intuitive—that preference misalignment decreases information sharing—the theoretical literature suggests that even a small biases on the sender’s part can be quite destabilizing. In contrast, the experimental literature has documented a much smaller level effect, with greater honesty than predicted.

Our contribution is to include an extensive margin, and to assess the effects of incentives for information transmission that vary on both margins. Concurring with previous work, our experimental results replicate the finding that preference misalignment is not as pernicious to information quality as predicted by theory. However, building upon this result, we show that this should make us more willing to trade off some alignment of interests in order to have the information paid for by a third party. In contrast, our incentive that maintains common interests over the provided information does lead to more total information transfer, but as the receivers have to pay for it, their net outcomes are actually made worse than the absence of any incentive. Finally, for an incentive that produces full misalignment of interests, we find the worst information transfer rates across our experiments.

While very distorted incentives to provide information do seem to fail in the long run when all parties are commonly aware of their presence, it could be that third parties still choose to covertly offer them. For example, with celebrity endorsements of products where the paid-for nature of the post is not disclosed. If receivers are unaware of the incentives supporting information (and falsely believe senders have common cause with them) third parties might be able to produce large short-run gains, while simultaneously poisoning the well in the long-run. This insight prompts the final step in our analysis, where we demonstrate that some commonly known degree of misalignment may in fact be desirable, as it can serve as an inoculation against the introduction of more-distortive incentives.

Our results highlight a previously unrecognized advantage to commission-like incentives in multi-product environments, and may help explain their prevalence. Commission-style incentives are commonplace, including brokers or realtors providing advice between two competing stocks or houses (an explicit commission on eventual purchases), servers providing recommendations over which dessert to order from a menu (an increased tip), or product review websites (referral commissions via hyperlinks). In each case, the informed party can have partially aligned interests with the decision maker over which of the competing options is better (a particular stock, house, flavor of ice cream, or consumer product), but is misaligned over the decisions to buy anything at at

all. While there certainly are cases of unbiased information provision paid for directly before final choices are made (subscription-based advice websites, fixed-fee consultants or inspectors, etc.), many sources of information that we commonly make use of are only partially aligned, and are not paid for directly before advice is proffered. Our paper posits the idea that this partial misalignment could be in our best interest.

As with most theoretical and experimental work, our paper's environment elides features present in many practical settings in order to isolate our main question, in this case the strategic tradeoffs between the quantity and quality of information. As such we focus on an abstract uninformed receiver who must make a decision between two ex ante symmetric products and an outside option (not buying either product). To aid in this decision our environment has an informed sender with a powerful, though imperfect, signal. Extensive-margin frictions over information sharing arise due to costs of providing and acquiring information, while intensive-margin frictions are generated by a conditional payment to the sender based upon the receiver's choice. Our experiments are tailored toward eliciting controlled comparative statics on the following questions: (i) How often do senders provide information?; (ii) What is the informational content of messages when provided?; (iii) Do receivers seek out the provided information?; and (iv) When viewed, how does the provided information shape the receiver's subsequent choice? We address these questions in an unincentivized baseline environment, and in three incentive structures that reward senders for providing information.

Our first incentive (labeled *Receiver*) has those making use of the information pay for it. While maintaining common interests between sender and receiver over the receiver's final choice, the treatment creates a distributional tension, where receivers must transfer some of the gains from the information exchange to the sender. Here the incentive mirrors situations where a final decision maker is choosing whether to pay for disinterested advice.

Our next two environments trade off alignment of interests by turning to a third party to pay for the provision. Our second incentive, which we call, *Marketplace*, as it aligns senders with the sales interests of the platform where competing products are sold, rewards senders contingent upon sales of any available product. While creating some misalignment between senders and receivers over whether to buy anything at all, the incentive does maintain the alignment of interests over which of the competing products is better. For example, while a broker on commission would always want you to make a purchase, they would ceteris paribus prefer you purchase the superior offering.

Finally, our third environment, which we call *Producer*, creates a full misalignment in preferences by directly incentivizing the sender to endorse a particular product, eliminating any alignment between the different sides of the information exchange, and instead aligning the sender with the interests of a particular producer.

Across these three incentive structures, our results suggest that *Marketplace* incentives do the best job at balancing the quantity/quality tradeoff for information transmission. Relative to our

baseline, many more senders choose to provide information, and while there *is* a decrease in the informational quality, it is more than made up for with increased quantity.

Below we discuss the related literature, before setting out our environment and design in Section 2. In Section 3 we outline our results, and then in Section 4 we extend them in a discussion of two counterfactuals examining third-party interests. Finally, we conclude in Section 5.

1.1. Related Literature. The theoretical starting point for the information transmission literature is Crawford and Sobel (1982), which describes the impossibility of full revelation with cheap talk when senders of information have misaligned preferences with receivers. Where the senders have state-dependent preferences, the upper bound for information revelation is shown to be decreasing in the size of the misalignment. Our aligned-incentive institutions have state-dependent preferences for senders with zero bias term, and with a larger message space full revelation would be possible. However, our misaligned-incentives bear more resemblance to the state-independent preferences with multiple dimensions in Chakraborty and Harbaugh (2010), where the sender is misaligned in one or many of these dimensions. In our setting these two dimensions are the particular product, and the receiver's willingness to pay (WTP). The main innovation in our paper's model is the joint examination of both the quality of provided information and whether it is provided at all, as our environment incorporates communication frictions.¹

A body of work has experimentally examined tensions between agents in the Crawford and Sobel environment.² The main experimental finding here is that subjects over-communicate relative to theory: senders tell the truth more often than predicted, and receivers infer honesty too much, though with a large heterogeneity explained via level- k thinking. In a setting closer to ours, Wilson (2014) examines the behavior of subjects in aligned-interest groups with similar two-sided costs. He finds subjects under-respond to the costs of sending messages, and overpay to acquire information, relative to the gains obtained. In contrast, our own paper examines tradeoffs between alignment and incentives to rate at all. Taking away the extensive margin Chung and Harbaugh (2016) examine the extent to which observed play matches the equilibrium prediction in a persuasion setting. They find that messages are persuasive, even in those environments where theory predicts they should not be, matching a result we find in our misaligned treatments.³

One natural application of our environment is to online rating systems, and several papers have examined incentives specifically in the context of online ratings. There is evidence that ratings given in the *absence* of explicit incentives exhibit bias due to differential rating frequency, as in Hu

¹A related theory paper with costs to both senders and receivers is Dewatripont and Tirole (2005), though there costs vary over the precision of articulation or interpretation of the message. See also Dessein and Santos (2006); Calvó-Armengol et al. (forthcoming) for models of endogenous communication with costs.

²See Dickhaut et al. (1995), Cai and Wang (2006), Wang et al. (2010). For extensions to multiple senders or receivers see Lai et al. (2011); Vespa and Wilson (2014) and Battaglini and Makarov (2014).

³See also Charness and Garoupa (2000) which examines the extent to which reputation affects revelation, where senders have a state-independent preference to induce sales.

et al. (2009) and Lafky (2014), which demonstrate the tendency for raters to over-report positive or negative experiences, relative to moderate outcomes.⁴ There is also some evidence for rating biases due to self-selection into the market, as in Li and Hitt (2008), where consumers who are predisposed towards a product are more likely to rate early in the product’s lifespan, leading to artificially positive information.

There is a clear incentive for firms to attempt to influence the information provided about their products, either directly by creating fraudulent ratings (Mayzlin et al., 2014) or indirectly by compensating raters. Providing explicit incentives for ratings can create obvious conflict-of-interest problems, in which compensated raters may give biased evaluations of products. The potential for such bias has led the Federal Trade Commission to enact rules requiring the disclosure of any relationship between firms and those endorsing their products via the internet.⁵ There is a small literature studying interventions designed to increase rating provision, including rebates (Li and Xiao, 2014), social comparisons (Chen et al., 2010) social identity (Wang, 2010), but to our knowledge, no existing work examines the relative performance of different incentive structures for ratings.

2. FRAMEWORK AND EXPERIMENTAL DESIGN

We will examine a representative final decision maker R (the receiver, he) making a choice between one of two (ex ante identical) risky choices, A or B , or a known-value outside option R . Each of the risky options is associated with a lottery ϕ^A and ϕ^B , represented in the experiment by two urns, *Urn A* and *Urn B*.

To make his decision, the receiver is provided with a message m from a peer (the sender, she) who has prior experience with one of the products $X \in \{A, B\}$, from which she drew a realization x . The main thrust of our experiments will be to assess how much information is transferred from the sender to the receiver, and the extent to which three alternative institutions might improve that transfer.

The timing of the game in our experiments is as follows:

- (1) Nature draws the two product lotteries ϕ^A and ϕ^B , where each is an iid draw from a set of possible lotteries Φ , with $p_r = \Pr\{\phi_r\}$.
- (2) The sender obtains a single realization x from one of the two lotteries ϕ^X , where $X \in \{A, B\}$ is selected with equal probability.
- (3) The sender decides whether or not to provide a meaningful message/rating $m \in \mathcal{M}$ to the receiver—where in our experiments \mathcal{M} is a one-to-five scale—or to send an empty

⁴Also see Bolton et al. (2013) with respect to managing the distortive effects of reciprocity when ratings are two-way.

⁵FTC press release on social media endorsements: https://www.ftc.gov/system/files/attachments/press-releases/ftc-staff-reminds-influencers-brands-clearly-disclose-relationship/influencer_template.pdf

message m_\emptyset . If the sender chooses any $m \in \mathcal{M}$ she incurs a cost c^S , where choosing $m = m_\emptyset$ is free.

- (4) The receiver observes whether there was a provided message and chooses whether or not to view if there is. If he chooses to view the message he incurs a fixed cost $c^R > 0$.
- (5) The receiver draws an outside-option value ω and makes a choice $Z \in \{A, B, R\}$.
- (6) The receiver obtains the outcome $w_R(Z, \omega)$ equal to the reservation value ω if he chooses R , and equal to an independent realization of the lottery ϕ_Z if he chooses $Z \in \{A, B\}$.

The payoff to the receiver is given by

$$u_R(Z, \rho, m) = w_R(Z; \omega) - c_R \cdot \mathbf{1}\{m \neq m_\emptyset, \text{View}\};$$

while the sender's payoff is

$$u_S(m, \rho, Z) = x + \alpha \cdot w_R(Z; \omega) - c_S \cdot \mathbf{1}\{m \neq m_\emptyset\}.$$

The sender is assumed to derive utility from her own lottery outcome x , but also experiences a positive spillover from the receiver's choice with weight $\alpha > 0$, reflecting a gain from helping the receiver.⁶

The main tension in the above model stems from the messaging costs. The sender's messaging costs (c_S) are privately incurred, but when α is small, the weight placed on the public benefits is also small. In many parameterizations (in particular our experimental ones) the resulting effect will be underprovision of information.

As an example, suppose that the sender's experienced outcome x is a realization from lottery A , denoting this information as \mathcal{I}_x^A . If the sender's information were perfectly revealed to the receiver, he would be able to make a better decision, with a benefit of $\Delta w_R(\mathcal{I}_x^A) := w_R(Z(\mathcal{I}_x^A)) - w_R(Z(\mathcal{I}_\emptyset))$. Moreover, if our receiver is a stand-in for K representative final decision makers the social benefit generated by the information transfer is therefore $(K + \alpha)\Delta w_R(\mathcal{I}_x^A)$ while the social cost from communication frictions is $(K \cdot c_R + c_S)$.

While the precise benefits of information transfer will vary based on the sender's experience x , and how informative that experience is about the lottery ϕ_X , information transfer is expected to be socially desirable even for moderate communication frictions. However, as we will discuss in further detail below, the equilibrium predictions are for inefficient communication, where senders are predicted to underprovide information, sending the empty message m_\emptyset even for low send costs.

Though the baseline environment above is predicted to be inefficient, there are potential incentive devices that might increase the transfer of socially valuable information. We examine one obvious candidate to economists (a transfer from receivers to senders) as well as two alternative incentives schemes motivated by those in common use. Below we provide more specific details

⁶Note that for the model α can be made arbitrarily small. The modeling choice here is to start out with a baseline model reflecting an informational public good, where all parties strictly benefit from the transfer of high-quality information.

on our experimental environment and parameterization, as well as the three alternative incentive schemes, where we will come back to a discussion of the equilibrium predictions in section 2.2.

2.1. Experimental Design. We use a between-subject experimental design across the four treatment environments, where a session consists of 30 rounds. The first treatment is a *Baseline* environment with no transfers. The other three incorporate conditional transfers to the sender, and affect the alignment between the interests of the sender and receiver.⁷

- A *Receiver*-funded transfer conditioned on a message being viewed. Conditional on provision, the sender and receiver share common interests, over both the superior product, and the *cardinal* willingness to pay.
- A *Marketplace*-funded transfer conditioned on a viewed message and the receiver choosing either of the products. Conditional on provision, the sender and receiver share partial interests over which product is *ordinally* better, but not over the willingness to pay for that product where they are misaligned.
- A *Producer*-funded transfer from X , conditioned on the receiver choosing the same product X . Conditional on provision, the sender and receiver have no common interests.

Determining the state. The primary uncertainties in the model are the lotteries ϕ^A and ϕ^B , which in the experiment are called *Urn A* and *Urn B*. Each urn is filled with two balls labeled with an integer between 1 and 100, so the possible realization set is $\Theta = \{1, \dots, 100\}$. In every round t and for every sender i , the two urns are independently filled via the following procedure: (i) An initial ball θ_1 is placed in the urn by a uniform draw over Θ ; ii) with probability $1/2$, the second ball θ_2 is another independent uniform draw from Θ , and with probability $1/2$ the second ball is an exact copy of the first ball, so $\theta_2 = \theta_1$. Urn X is therefore completely determined by the two balls in it, θ_1^X and θ_2^X , and a draw from the urn is a lottery $\frac{1}{2} \cdot \theta_1^X \oplus \frac{1}{2} \cdot \theta_2^X$.

The set of all possible urn lotteries in our experiment is $\Phi = \{\frac{1}{2} \cdot \theta \oplus \frac{1}{2} \cdot \theta' \mid \exists \theta, \theta' \in \Theta\}$. Given the filling procedure outlined above the prior probability p_r of the lottery $\phi_r = \frac{1}{2} \cdot \theta \oplus \frac{1}{2} \cdot \theta'$ is given by

$$p_r = \begin{cases} 1/10,000 & \text{if } \theta \neq \theta', \\ 101/20,000 & \text{if } \theta = \theta'. \end{cases}$$

Without any knowledge on the specific realized lottery ϕ_r , the expected outcome is given by $\bar{\theta} = 50.5$. However, a single draw from the urn (with replacement) provides a great deal of information. The expected outcome of a new draw from the urn given knowledge of a previous

⁷In terms of defining alignment of interest over the information transferred, one easy definition is the extent to which a sender would reveal their signal in equilibrium. In our common-interest environments, the sender would fully reveal her cardinal signal if it was free to do so. In contrast, with partially aligned interests, the sender would only reveal that the signal was in one of two subsets. With no common interests, the sender would reveal no information. An alternative definition for alignment is the number of distinct posteriors possible in equilibrium.

draw x is $\nu_X(x) = 3/4 \cdot x + 1/4 \cdot \bar{\theta}$.⁸ The experimental parameterization is designed to ensure that the sender's information is very valuable to the receiver while not being perfectly informative.

Aside from the two lottery draws, our experimental environment has two other random components: a sending cost c^S and a reservation ω . For the sending cost c^S our experiment makes an iid draw from $\mathcal{G} = \{-\$0.49, \dots, \$2.50\}$, where the probability of each cost level is linearly decreasing to zero at $c^S = \$2.50$.⁹ The reservation value ω_j is a draw from $\mathcal{H} = \Theta = \{1, \dots, 100\}$, where the probability of each reservation value is decreasing linearly to zero for $\omega = 100$.¹⁰

In each experimental round, each subject performs the roles of both a sender and a receiver.¹¹ In each round t the game timing is as follows: The subject first makes a choice of urn $X_{it} \in \{A, B\}$ to draw a ball from, obtaining the realization x_{it} from the lottery ϕ_{it}^X , where the subject earns $\$0.05 \times x_{it}$ from the first stage.¹² The sender's information is therefore $\mathcal{I}_{it}^S = (X_{it}, x_{it})$, which is pertinent information to a receiver subject j who will make a choice over $\{\phi_{it}^A, \phi_{it}^B, R\}$ in the second stage.¹³

After observing their drawn ball's value x_{it} , each subject is asked to provide as a message an integer rating from one to five for either urn, so the set of available non-empty messages is: $\mathcal{M} = \{1_A, 2_A, 3_A, 4_A, 5_A\} \cup \{1_B, 2_B, 3_B, 4_B, 5_B\}$.¹⁴ After selecting a provisional rating $\hat{m}_{it} \in \mathcal{M}$, the subject was then informed of the realized cost of sending, c_{it}^S , and chose whether or not to send the selected rating, $m_{it} \in \{\hat{m}_{it}, m_\emptyset\}$, incurring the (possibly negative) cost c_{it}^S only if they chose to send. Our experiment guarantees a minimum level of alignment between the sender and the receiver by setting an explicit altruism term $\alpha = 0.1$, providing i with $\$0.10$ for every dollar earned by the matched receiver j in their second-stage choice Z_{jt} .

In addition to these baseline costs and payments for the sender, in our three transfer environments we provide an opportunity to earn an additional $T = \$2.00$, where the condition for payment

⁸Half the time the new draw is the same exact ball drawn previously (value x), a quarter of the time the new draw is a copy of the previously drawn ball (value x again), and a quarter of the time the new draw is an unrelated ball (with expectation $\bar{\theta} = 50.5$).

⁹The send-cost support includes negative values, as we wanted to have an interior level of provision in the baseline. The interpretation of a negative send cost is a sender who derives explicit benefit from sending a message.

¹⁰In the experiment we draw two send costs uniformly from \mathcal{G} , and give the subject the *lower* cost of the two, so that the ex ante probability of send cost c_i^S is approximately $\frac{2}{300} \frac{2.5-c}{3}$. For the reservation we draw two uniform draws from Θ , where the reservation draw ω_j is the lower value. The probability of a reservation draw ω is therefore approximated by $\frac{2}{100} \frac{100-\omega}{100}$.

¹¹Our choice to make all subjects both senders and receivers was made so the environment was symmetric ex ante in the sense that the entire population can both benefit from other's information provision and provide it themselves.

¹²Though x_{it} is simply a signal draw from the point of information provision, we wanted experimental senders to experience a payoff from their draw.

¹³In each round the experimental match procedure randomizes subjects to an identity $i \in \{1, \dots, n\}$ on a circle. The subject assigned to identity i acts as a sender to the clockwise subject with identity $j = \text{mod}(i, n) + 1$ and the receiver to the counter-clockwise subject with identity $k = \text{mod}(i - 2, n) + 1$ in the second stage, making their *second-stage* choice from $\{\phi_{kt}^A, \phi_{kt}^B, R\}$.

¹⁴Again, we had many choices over the message space, where our experimental choice was to use a simple one-to-five rating that is likely familiar to subjects.

$\psi^S(m, \rho, Z)$ varies by treatment. Round payoffs from the sending decision are therefore:

$$u_S(m_{it}, Z_{jt}, \rho_{jt}; c_{it}^S) = \$0.05 \times x_{it} + \$0.01 \cdot w_R(Z_{jt}) - c_{it}^S \cdot \mathbf{1}\{m_{it} \neq m_\emptyset\} + \$2.00 \times \psi^S(m_{it}, Z_{jt}, \rho_{jt}).$$

In the second stage of the game subject j receives the urns A and B from subject i .¹⁵ If there is a provided message $m_{it} \in \mathcal{M}$, the receiver is informed a rating was sent and must decide whether or not to view it, $\rho_{jt} \in \{\text{View}, \text{Not}\}$. If the message is viewed the receiver incurs a fixed cost $c^R = \$0.05$ in all four institutions, where in the *Receiver* treatment an additional amount $t = \$0.50$ brings the total view cost to $\$0.55$.¹⁶

The receiver's information set at this point is given by \mathcal{I}_{jt} —either observing the specific rating sent if they view, $\mathcal{I}_{jt} = \{m_{it}\}$; that a rating was provided but not viewed, $\mathcal{I}_{jt} = \{m_{it} \in \mathcal{M}\}$; or that no rating was provided, $\mathcal{I}_{jt} = \{m_{it} = m_\emptyset\}$. To gather more information we ask the receiver to first make a provisional decision between the two urns $\hat{Z}_{jt} \in \{A, B\}$, after which they are informed of the reservation draw ω_{jt} . They are then asked to make a final choice $Z_{jt} \in \{\hat{Z}_{jt}, R\}$: either take a realization z_{jt} from the provisionally chosen urn \hat{Z}_{jt} , or stick with the known reservation value ω_{jt} . The receiver's consumption payoff $w_R(Z_{jt})$ is then given by $\$0.10 \times z_{jt}$ or $\$0.10 \times \omega_{jt}$. Receivers therefore obtain $\$0.10$ multiplied by the value of the final ball selected. The total payoff from the receiver phase is subsequently:

$$u_R(m_{it}, \rho_{jt}, Z_{jt}) = \$0.10 \times w_R(Z_{jt}) - \$0.05 \times \mathbf{1}\{\rho_{jt} = \text{View}, m_{it} \neq m_\emptyset\} - \$0.50 \times \psi^R(m_{it}, \rho_{jt}, Z_{jt}).$$

While the first fifteen rounds exactly follow the above timeline, in the second half of each session we introduce a strategy method to collect richer strategic choices.¹⁷ Instead of showing subjects the send cost c_{it}^S and choosing whether or not they wish to send the provisional message \hat{m}_{it} at that cost, in the last fifteen rounds we ask them to specify a cutoff cost $\bar{C}_{it} \in \mathcal{G}$ below which they would be willing to send. Similarly, receivers are not informed of their specific reservation draw ω_{jt} , and are instead asked to specify a reservation cutoff $\bar{\Omega}_{jt} \in \mathcal{H}$ (a certainty equivalent) below which they would choose a draw from \hat{Z}_{jt} (and above which they would choose the realized reservation ω_{jt} .)

Data from interactions in rounds 1–15 are therefore given by $\langle (\phi_{it}^A, \phi_{it}^B, c_{it}^S, \omega_{jt}), (X_{it}, \hat{m}_{it}, m_{it}), (\rho_{jt}, \mathcal{I}_{jt}, \hat{Z}_{jt}, Z_{jt}) \rangle$; while data from rounds 16–30 are given by $\langle (\phi_{it}^A, \phi_{it}^B, c_{it}^S, \omega_{jt}), (X_{it}, \hat{m}_{it}, \bar{C}_{it}), (\rho_{jt}, \mathcal{I}_{jt}, \hat{Z}_{jt}, \bar{\Omega}_{jt}) \rangle$.

¹⁵In the experiment *Urn A* and *B* are relabeled as *Urn C* or *Urn D* for the receiver subjects. This relabeling is done with equal probability so receiver j cannot know which urn the sender chose from. For tractability though we will continue to refer to them as urns A and B rather than urns C and D .

¹⁶Here our parameterization effectively makes the implicit assumption that each receiver is representative of $K = 4$ final consumers, so that $4 \cdot t = T$.

¹⁷For a survey of papers examining (and for the most part supporting) the strategy method relative to direct-response see Brandts and Charness (2011).

2.2. Equilibrium Predictions. Our environment is a form of “cheap-talk.” Though there are costs for sending and receiving messages, there are no differential costs across the different informative messages. As such, in all of our environments there will be a “babbling” equilibrium where meaningless messages are selected and sent only for negative costs, and where receivers do not view the resulting messages. The babbling outcome reflects the worst case for information transfer, a coordination on zero information being conveyed. However, in three of our treatments there is the theoretical possibility of coordinating on more-informative outcomes. Given the parameterizations outlined above, Table 1 provides predictions for the main economic variables in the most-informative PBE (miPBE, where we enforce a symmetry restriction over the urns) in each treatment.¹⁸

Before diving into the predictions, we should note here our interpretation for the theory. While the equilibrium calculations are undoubtedly complex, the choices in our game are relatively simple. For senders: given an easy-to-interpret signal, what information would you send, and at what cost to yourself would you actually send it? For receivers: will you view a message at a fixed cost; given your information which urn do you prefer, and what is your valuation for it? As with many models we do not expect subjects’ behavior to stem from introspective calculations of the best responses given stated probability distributions and conjectures on the distributions of others play. Instead we see the miPBE here as an “as if” for the upper-bound outcomes in our four treatments given enough time to learn responses with favorable initial play.

The first-order difference in behavior across the *Baseline* and *Receiver* miPBEs is the rate at which senders provide information, $\Pr \{ \mu^* (X, x, c^S) \in \mathcal{M} \}$, given in the *Rating Provision* row in Table 1. In the *Baseline*, information is provided approximately a third of the time (up to costs of approximately \$0.10) while in the *Receiver* case provision is almost complete (at costs in excess of \$2). The extent to which decision-relevant information is exchanged is measured by Υ which is given in the *Info. Efficiency* row.¹⁹ This efficiency measure examines the expected receiver outcome $w_R(Z)$ in the equilibrium measured relative to the expected outcome under no communication \underline{W} , which is then normalized by $\overline{W} - \underline{W}$, the difference between the expected outcomes under perfect and no information exchange.

Given selection of miPBE, communication generates 36.3 percent efficiency in the *Baseline*, which increases to 98.1 percent in the *Receiver* treatment. However, gains in information efficiency do not come for free in the *Receiver* treatment. While the average sender is much better off in the

¹⁸A full specification of an equilibrium here is a tuple $\langle \mu^*, (\rho^*, \zeta^*) \rangle$: a message strategy $\mu^* : \{A, B\} \times \Theta \times \mathcal{G} \rightarrow \mathcal{M} \cup \{m_\emptyset\}$ for the sender (where $\mu^*(x, X, c_S)$ take into account the signal draw and cost of sending; and for the receiver a viewing strategy $\rho^* \in \{\text{View}, \text{Not}\}$ and a final choice rule $\zeta^* : I \times \mathcal{H} \rightarrow \{A, B, R\}$ (where $\zeta^*(\mathcal{I}, \omega)$ takes into account the precise information set $\mathcal{I} \in I$ and the reservation). More detailed derivations are relegated to the appendix, as well as a comparable prediction table formulated under a risk-averse preference.

¹⁹For both the *Baseline* and *Receiver* treatments, the most-informative equilibrium involves sending a single rating 1_X for all sender types (X, x) with $x \leq 50$ and sending the ratings 2_X through 5_X for the signals x in the ranges 51–61, 62–73, 74–86, and 87–100, respectively.

TABLE 1. Most-Informative Risk-Neutral Equilibrium Predictions

	Baseline (B)	Receiver (R)	Marketplace (M)	Producer (P)	Comp. Static
Alignment	Cardinal	Cardinal	Ordinal	None	
Distinct Ratings	10	10	2	0/Babbling	
Rating Provision	0.346	0.980	0.934	0.304	R > M > B > P
Info. Efficiency, Υ	36.3%	98.1%	87.6%	0.0%	R > M > B > P
Rec. Welfare, Υ_R	34.1%	30.6%	81.7%	0.0%	M > B > R > P
Conditional on rating					
View Rate,	1.0	1.0	1.0	0.0	B ~ R ~ M > P
Info. Efficiency Υ	106.6% [†]	100.3% [†]	93.9%	0.0%	B > R > M > P
Rec. Welfare Υ_R	100.4% [†]	31.4%	87.6%	-6.3%	B > M > R > P

Note: †-There are selection effects over quality through the sender’s decision to provide a rating or not, where senders are more likely to send given a high signal. Because of this the efficiency and receiver welfare *conditional* on no rating can be negative, while they can exceed 100 percent conditional on provision.

most-informative *Receiver* equilibrium than in the *Baseline*, the provision subsidy has to be paid for by receivers. Even though the *Receiver* treatment increases total efficiency, once we account for the payment they make to the sender, receivers actually fare worse in this setting than the *Baseline*. This is illustrated by the *Rec. Welfare* row in Table 1, which modifies the efficiency measure Υ_R to account for the incurred costs of communication to the receiver. Our receiver welfare measure uses the receiver’s expected payoff, incorporating the expected costs of the transfer, where, as above, we normalize the outcome relative to those with perfect costless messaging and zero-information transfer.

The raw cost to view a message in our experiment is $c^R = \$0.05$, which represents a 6.3 percent efficiency penalty if incurred with certainty. In the *Baseline* prediction, where ratings are provided approximately a third of the time, the presence of a cost results in receiver welfare of 34.1 percent. However, there is a much larger drop in the *Receiver* treatment, with receiver welfare falling to just 30.6 percent. There are two reasons for this larger drop: First, ratings are predicted to be provided at much higher rates so acquisition costs are incurred more often. Second, and the quantitatively larger force, much of the receiver’s gain from increased information is transferred to the sender through the incentive payment t . In fact, despite a large increase in total efficiency, receivers’ final outcomes are inferior to the most-informative *Baseline* prediction. Taken in context with the presence of multiple equilibria in the *Baseline* and *Receiver* institutions, one effect of the

reduced receiver welfare may be to make coordination on the informative equilibria riskier when the receivers are the ones footing the bill.²⁰

In contrast with the above two environments where the sender would fully reveal the cardinal level of her signal if it were free to do so, the *Marketplace* setting's conditional transfer removes some of the alignment between the sender and receiver over the available information. Instead, the two agents share a conditional alignment only over which of the two products is *ordinally* better, and have no shared interests over WTP. The *Marketplace* miPBE uses just two effective messages from the ten available. In equilibrium, sent messages are equivalent to making the statement “*Given my information, I expect A (B) to be better choice for you than B (A),*”²¹ which can be achieved by sending a five-rating on the sampled urn whenever the signal is better than expected ($x > 50$) and a five-rating on the unsampled urn otherwise. Though less useful than the ratings in the settings with fully aligned interest, this binary statement is still informative, and is predicted to increase willingness to pay. A receiver viewing a rating of 5_Y in the *Marketplace* miPBE will choose the product Y for all reservation values $\omega < 60$, and not to buy whenever $\omega \geq 60$.²²

Recall that the sender's incentive payment in the *Marketplace* treatment is conditioned on receivers choosing either one of the two products. The expected subsidy to a sender who provides a rating J is $\Pr\{z \neq R\} \cdot T$, which is smaller than the *Receiver* miPBE where the full payment is received with probability one.²³ This drop in the expected incentive reduces provision relative to the *Receiver* treatment. The 87.6 percent information efficiency in *Marketplace* is lower than the *Receiver* case—through reduced provision and coarser revealed information—but still represents a substantial improvement on the *Baseline*. For receivers though, because they are not paying for the incentive payment, they share more of the gains from increased efficiency. The 81.7 percent receiver welfare is the highest within our four treatments. Indeed, both senders and receivers do better than the *Baseline*, despite the misalignment introduced over WTP.

Finally, the *Producer* treatment only has a babbling equilibrium outcome as the sender and receiver share no ex ante interests over information transmission. Though ratings are still provided (for negative costs c^S) receivers choose not to view them in equilibrium and the ratings convey no useful information. Because the sender incentive is large and conditioned on the sales of a

²⁰Assuming a benign form of the babbling equilibrium (where less harmful on miscoordination 1_X -ratings are sent for $c^S < 0$), we consider a coordination problem for receivers between the babbling and the miPBE outcomes. Receivers in *Baseline* will choose to view whenever they place more than a 0.16 probability on senders coordination on the miPBE. In contrast, in the *Receiver* treatment, the additional cost of viewing means receivers must place more than 0.69 probability on the miPBE, and so receiver-coordination on not-viewing/babbling is risk dominant.

²¹The equilibrium message behavior can be characterized by a rating strategy such as $\mu(X, x, c^S) = 5_Y$, where $Y = X$ for $x > 50$ and $Y \neq X$ when $x \leq 50$, sent whenever the provision cost c^S is not too large.

²²The exact choice stems from the expected quality of urn Y given the message 5_Y being $5/8\bar{\theta} + 3/8\mathbb{E}(x|x > 50) = 59.875$. Note that unlike the previous cases, the rating here is also informative about the other urn $Y' \neq Y$, which has an expected quality of 41.125.

²³In section 4 we come back to thinking through why a Marketplace player might provided such incentives.

particular product, sender preferences are effectively state independent. Without any common interests with the sender the receivers cannot view messages as credible in equilibrium. As such, the *Producer* incentives are predicted to be counterproductive in the long run, and it is clear that neither producer *A* nor *B* would want to pay for information incentives in equilibrium.

While the equilibrium analysis of the *Producer* setting indicates the incentive is counterproductive in the long-run, in the short-run such incentives may be quite powerful. If producers deployed such incentives to senders in secret within the *Baseline* or *Receiver* environments it would be possible to substantially increase the same-product sales figures until receivers learned that five-ratings were largely dishonest. A similar short-run benefit might induce marketplaces to offer secret incentives for marketplace-level incentives, as a five-rating increases the WTP (and therefore sales) for a credulous receiver. We come back to an experimental comparison of the short- and long-run incentives in section 4.

3. RESULTS

Below we present results from sessions conducted at the Pittsburgh Experimental Economics Laboratory. We used the z-Tree (Fischbacher, 2007) experimental software to collect data from 176 unique subjects across 12 sessions, with 3 sessions per treatment. Subjects were paid for two randomly selected rounds out of the 30 they played, with average payments per subject of \$18.09. We first outline treatment averages, and test the comparative statics derived from our equilibrium predictions. We then analyze subject behavior in more detail, and outline the “why” of the results.

3.1. **Efficiency.** Our main outcome results can be summarized as follows:

Result 1 (Efficiency).

- (1) *Total information transfer is substantially lower than predicted in all treatments, though observed behavior matches the best-case equilibrium comparative statics.*
- (2) *Receivers are better off under the Marketplace incentive than the Receiver incentive.*

Evidence for the first result is contained in Table 2 which summarizes the experimental data, where each row provides sample analogs to the predictions in Table 1. While our theory predicts receivers either always or never view, we additionally include averages for the subsample with viewed ratings. The first five rows in Table 2 outline unconditional means: i) the rate of provision by senders; ii) the overall efficiency of the information transfer;²⁴ and iii) the receiver welfare, which accounts for the expected cost of viewing. The remaining rows report conditional outcomes, first viewing frequency conditional on a sent rating, and second the efficiency and receiver welfare conditional on a rating being sent and viewed.

²⁴All efficiency measures are calculated by a recombinant procedure, matching all receivers against the entire sender population consistent with the information set. In this way, we integrate out much of the exogenous noise (variation in $x_{it} | m_{it}$ and c_{it}) while retaining the observed strategic variation in sender and receiver behavior.

TABLE 2. Experimental Outcomes

	Baseline ($S = 46$)		Receiver ($S = 40$)		Market. ($S = 47$)		Producer ($S = 44$)		Comparative Static					
	Avg.	Std. Err	N	Avg.	Std. Err	N	Avg.	Std. Err	N	Theory (miPBE)	Data			
Provision	0.369	0.013	1380	0.605	0.014	1200	0.634	0.013	1410	0.488	0.014	1320	R \succ M \succ B \succ P	M \succ R \succ P \succ B
Info. Efficiency	9.0%	3.9%		28.1%	4.1%		17.7%	3.8%		1.0%	3.9%		R \succ M \succ B \succ P	R \succ M \succ B \succ P
Rec. Welfare	7.3%	3.8%		-2.9%	4.1%		14.1%	3.8%		-1.1%	3.9%		M \succ B \succ R \succ P	M \sim B \succ P \sim R
Conditional on provided rating														
Viewing	0.735	0.020	509	0.742	0.016	726	0.899	0.010	894	0.688	0.018	644	B \sim R \sim M \succ P	M \succ R \sim B \succ P
Conditional on viewed rating														
Info. Efficiency	58.4%	7.6%	374	82.4%	6.3%	539	57.6%	5.2%	804	45.1%	7.0%	443	B \succ R \sim M \succ P	R \succ B \sim M \succ P
Rec. Welfare	52.1%	7.6%		13.4%	7.9%		51.3%	6.7%		38.8%	7.9%		B \succ M \succ R \succ P	B \sim M \succ P \succ R

Note: Tests are derived from Wald tests after regression (probit) results for Efficiency/Welfare (Provision/Viewing/Sales decisions) on treatment dummies. The relationship $A \succ B$ (\sim , \succ) represent rejection of mean equality for treatments A and B with the one sided alternative $A \succ B$ at 99 percent confidence (95, 90 percent confidence, respectively). The relationship $A \sim B$ reflects failure to reject at 90 percent confidence. We list binary pairs of tests according to the data averages; comparative static relationships are ordered by the size of the effects, and are transitive except for Inf. Efficiency given a viewed rating where $B \sim V$.

Efficiency of information transfer. Ordinarily, unconditional efficiency matches the comparative statics predicted by theory. At one end, the *Receiver* incentive produces the most information transfer, with significantly greater efficiency than any other treatment. The *Producer* treatment meanwhile has the lowest information transfer, with an average efficiency just below zero. Between these two extremes are the *Marketplace* and *Baseline* treatments, in line with their theoretical rankings. However, once we compare the cardinal levels in Table 1 and 2, we see quantitatively large differences. The *Receiver* and *Marketplace* treatments theoretically allow for efficiency levels close to the frictionless upper bound, but in realization we end up with just a quarter of the total efficiency. One reason for the drop is the lower than predicted rates of provision and viewing in *Receiver* and *Marketplace*. To control for this difference we instead look at the information efficiency conditional on a provided and viewed message. Examining this subsample in Table 2 we do see a much larger efficiency increase. *Receiver* and *Producer* are again the best and worst treatments, respectively. For *Receiver* we observe an average efficiency of 82 percent of the upper bound, while even the *Producer* treatment is conveying useful information, with 45 percent efficiency.

Conditional on a viewed rating, *Baseline* fares significantly worse than *Receiver*, where the best-case theory predicts comparable levels. Instead, we observe a large and significant efficiency gap between *Receiver* and *Baseline*, and no significant difference between *Marketplace* and *Baseline*. Instead of the 94 percent efficiency possible after an equilibrium rating is provided, we observe just 58 percent in *Marketplace*.

Receiver outcomes. For the three treatments where receivers do not pay any additional amount to the sender, the immediate effect of viewing a rating is the same: the receiver pays a 6.3 percent efficiency penalty due to the cost of viewing. However, in the *Receiver* treatment, the transfer to senders reduces receivers' share of the information surplus substantially— with a 69 percent penalty to efficiency on viewing. Despite this large cost, receiver welfare conditional on viewing is still positive. Net of costs, receivers who view a rating improve their outcomes relative to their expected outcome with no sender present in the game.

The most important result with regard to receiver outcomes is that their maximum payoffs are achieved in the *Marketplace* treatment. Conditional on a viewed rating, the receiver welfare row of Table 2 shows that receivers do just as well in the *Baseline* and *Marketplace* treatments, so the effect of misalignment over WTP in the *Marketplace* incentive is small. In other words, trading off some quality in the information provided in order to subsidize ratings leads to offsetting increases in quantity. The unconditional receiver welfare figures therefore attain a maximum in *Marketplace*, however, the overall difference in receiver welfare compared to the *Baseline* is only marginally significant.

TABLE 3. Decision to send a rating.

	Baseline		Receiver		Marketplace		Producer	
	Provision	Cutoff	Provision	Cutoff	Provision	Cutoff	Provision	Cutoff
Signals:								
x , Low	0.344 *** (0.019)	\$0.069 ** (0.031)	0.563 *** (0.041)	\$0.646 *** (0.084)	0.572 *** (0.036)	\$0.684 *** (0.090)	0.396 *** (0.027)	\$0.237 *** (6.4)
x , Average	0.337 *** (0.019)	\$0.078 ** (0.031)	0.556 *** (0.041)	\$0.634 *** (0.084)	0.627 *** (0.036)	\$0.750 *** (0.090)	0.481 *** (0.027)	\$0.315 *** (6.4)
x , High	0.422 *** (0.018)	\$0.144 *** (0.030)	0.691 *** (0.039)	\$0.828 *** (0.084)	0.716 *** (0.033)	\$0.867 *** (0.090)	0.563 *** (0.030)	\$0.598 *** (6.4)
Costs:								
$c_{it} \leq 0$	0.977 *** (0.008)		0.999 (0.001)		0.993 (0.004)		0.982 *** (0.007)	
Cost ($\frac{\partial}{\partial c}$)	-0.469 *** (0.093)		-0.333 *** (0.095)		-0.327 *** (0.056)		-0.567 *** (0.043)	

Note: Results for *Sent* specifications are recovered using a random-effects probit model on the binary dependent variable $1 \{m_{it} \neq m_0\}_{it}$, whether a rating was sent; results for *Cutoff* specifications from a random-effect least squares model with dependent variable \bar{C}_{it} , the specified cost cutoff to send. All variables are predicted levels for dummy conditions, and the marginal effect of a one-dollar increase in rating cost from zero. Standard errors on the predicted levels/marginal effects are in parentheses where stars represent significance relative to a null at the babbling provision level (0.306/\$0.00 with full provision for negative costs) at: ***-99 percent confidence; **-95 percent confidence; and *-90 percent confidence.

3.2. Sender Behavior. We now move to an analysis of sender and receiver behavior. Mirroring the sequential nature of the experiment, we first examine senders, analyzing their behavior in the same order in which they made choices. Sender behavior is first summarized in our second result, then discussed in detail below.

Result 2 (Sender Behavior).

- (1) *Senders' choices respond on the extensive and intensive margins as predicted by theory, though with a smaller response than predicted on the intensive margin.*
- (2) *Sender's with very positive information are more likely to provide a rating in all four settings.*
- (3) *Total information provided is largest under the Receiver and Marketplace incentives, and smallest under the Producer incentive.*

Incentives increase provision. The differing incentive structures heavily influence the extensive margin, substantially increasing the number of ratings sent. The *Provision* row of Table 2 indicates significant treatment variation over the frequency with which ratings are sent, ranging from just

over one-third of the time in the *Baseline*, just under half the time in *Producer*, and almost two-thirds of the time in the *Receiver* and *Marketplace* treatments. Offering incentives to rate products increases the quantity of ratings, with all three incentive schemes generating significantly more ratings than the *Baseline*. Table 3 provides estimation results indicating the factors that influence the decision to provide a rating. For each treatment we provide two regressions. The first is a random-effects probit on the decision to send, where the right-hand-side variables are dummies for low, average and high signal draws ($x_{it} \leq 33$, $33 < x_{it} < 67$ and $67 \leq x_{it}$, respectively), a dummy for a non-positive rating cost ($c_{it}^S \leq 0$), and the actual cost of rating when it is positive (c_{it}^S). The second, using the cutoff-strategy data from the last 15 rounds, provides results from a random-effects panel estimate, regressing the specified cost cutoff on the same three signal-quality dummies. Taking both specifications together, the table provides the estimated rating frequencies and cost cutoffs for signals drawn in each quality region, alongside the predicted rating frequencies given a negative rating cost, and the marginal effect on rating from increasing the sender's cost from zero.²⁵

Senders in all treatments are sensitive to their signal, with a rating more likely to be provided following a good draw, but the size of the effect is strongest in the incentivized environments. In *Baseline*, going from a low draw to a high draw increases the likelihood of rating by 6.7 percentage points, and the send cutoff by 7.5 cents. This almost doubles for the *Receiver* treatment, with a 12.8 percentage point increase for a high draw, where the specified cutoffs differ by 18 cents. In both the *Baseline* and *Receiver* the effect of increased provision is only significant when the draw is high. However, in the *Marketplace* and *Baseline* treatments there are significant increases in rating frequency from low draws to average, and from average draws to high. The final two rows of Table 3 show that subjects are sensitive to the cost of rating, with subjects overwhelmingly choosing to rate following a negative rating cost, and exhibiting a large and significant decrease in rating frequency as costs rise.

Incentives distort information provided. Having identified the effect of incentives on the extensive margin, we now focus on the intensive margin, specifically the prevalence of dishonest ratings. We say that a rating is dishonest if it satisfies any of three conditions: i) *Overstatements*, a positive rating (a four or five) on the drawn urn X for a low-quality signal ($x \leq 50$); ii) *Understatements*, a negative rating (a one or two) on urn X for a high-quality signal ($x > 50$); or iii) *Other Urn*, any rating where the rated urn Y is not the urn X sampled by the sender. We then measure the extent to which the three incentivized treatments have different rates of dishonesty when compared to the *Baseline*. Table 4 summarizes the dishonesty over both the initially selected ratings and those that were ultimately sent. We focus on the last 15 rounds where subjects have already had experience

²⁵In the three incentivized treatments the majority of the variation in provision is between subject through the rating costs. In the appendix we illustrate variation in the subject-level averages for the rate of provision, and the specified send-cost cutoff (Figure 4).

TABLE 4. Frequency of each type of dishonesty (last 15 rounds)

		Baseline	Receiver	Marketplace	Producer
Chosen Ratings	Overstatement	0.023	0.035 *	0.079 ***	0.147 ***
	Understatement	0.039	0.003 ***	0.011 ***	0.027
	Other Urn	0.012	0.000 ***	0.043 ***	0.064 ***
	Any lie	0.074	0.038 ***	0.133 ***	0.238 ***
Sent Ratings	Overstatement	0.024	0.036	0.055 ***	0.158 ***
	Understatement	0.028	0.003 ***	0.013 *	0.029
	Other Urn	0.016	0.000 ***	0.035 ***	0.065 ***
	Any lie	0.069	0.039 **	0.103 ***	0.252 ***

Note: Confidence levels from two-sided binomial test of each treatment mean against null of same dishonesty as Baseline: ***-99 percent; **-95 percent; *-90 percent.

with the rating environment, so that we can be more confident that observed behavior is dishonesty rather than simply initial miscoordination on the meaning of each rating (though results are similar for the entire sample).

While incentives to exaggerate an urn’s quality exist in both the *Marketplace* and *Producer* treatments, overstatements are much more common in the *Producer* treatment. Dishonesty with respect to which urn is rated is also substantial in both of the misaligned environments. In the *Marketplace* treatment, ratings sent for the unobserved urn are most often favorable fours and fives, while in the *Producer* treatment they are overwhelmingly negative, most often ones. Rather than attempt to influence the receiver’s WTP with an overstatement on the incentivized product—which we will later show is only marginally effective—some senders in *Producer* use their rating to influence the receiver’s choice between products through a negative rating on the unsampled urn.²⁶

Though the dishonesty results do indicate significant changes across treatments, the levels of dishonesty in the misaligned treatments are low in comparison to theory—which is a common finding in the experimental literature.²⁷ Across all four treatments, the majority of provided information is honest—greater than three quarters of ratings in all treatments.

Efficiency given sender behavior. Finally, we ask how much information in total is being transmitted in each treatment. What is the change in informational content due to dishonesty? Given the observed accuracy and frequency of ratings, what is the attainable efficiency? To answer these questions we construct the empirical best response for receivers, given observed sender outcomes.

²⁶As an example of this type of rating dishonesty, Italian antitrust authorities fined the website TripAdvisor for a series of false negative reviews for hotels (“TripAdvisor fined \$600,000 by Italian Anti-Trust”, Associated Press, Dec. 22, 2014) while promoting its reviews as “authentic and genuine.”

²⁷See Gneezy (2005) for a simple experimental identification of this, and Kartik (2009) for a theoretical treatment.

TABLE 5. Possible Efficiency Given Sender Behavior

Variable	Baseline Receiver Marketplace Producer			
Info. efficiency	32.2%	54.1%	48.5%	21.3%
Info. efficiency rating sent	91.0%	96.1%	88.2%	62.2%
Info. efficiency no rating sent	-2.2%	-10.1%	-20.2%	-17.8%
Rec. welfare rating sent, not viewed	9.2%	5.8%	9.0%	13.3%
Rec. welfare rating sent and viewed	84.7%	27.1%	82.0%	56.0%

Note: Information efficiency possible, is calculated where receivers use the risk-neutral empirical best response with accurate beliefs on the empirical expectation of quality for each urn, given their information set.

The first row in Table 5 indicates the maximum efficiency possible given sender behavior.²⁸ The second and third rows break out the sample into the efficiency possible with and without a rating. Finally, the last two rows hold constant the provision of a rating, but examine the receiver welfare (inclusive of viewing costs and payments to senders) from viewing the rating or choosing not to view. A comparison of the bottom two rows therefore indicates whether a best-responding receiver should view the rating or not.

When a rating is provided, the attainable efficiency levels are very high—over 90 percent of that attainable if receivers perfectly observed the sender’s draw from the urn in *Baseline* and *Receiver*, and just under that in *Marketplace*. Though there is a substantial drop off in the *Producer* treatment, the total information content is still substantial.²⁹

3.3. Receiver Behavior. We now turn to the other side of the market, and ask how receivers respond to sender behavior. Our main results on receiver behavior are summarized as:

Result 3 (Receiver Behavior).

- (1) *Receivers are less likely to acquire information when the sender’s preference is completely misaligned in Producer, and most likely to acquire under the partial alignment of Marketplace.*
- (2) *Viewed ratings have substantial effects on the choice between products in all four treatments.*
- (3) *In all treatments except Producer a positive rating increases the likelihood of a sale.*

²⁸The expected efficiency given sender’s behavior for a best-responding receiver is calculated as $\hat{\Upsilon}_S = \frac{1}{\bar{W} - W} \left[\left(\sum_{m \in \mathcal{M} \cup \{m_\emptyset\}} \hat{\Pr}\{m\} \cdot \int \max\{\hat{\nu}_A(m), \hat{\nu}_B(m), \omega\} dH(\omega) \right) - W \right]$, where the best-response action is simply to choose the option with the highest expected outcome given the information available m and the reservation ω . Similar calculations lead to estimates for the attainable efficiency (and receiver welfare) conditioned on receiving a rating or not.

²⁹The size of the provision selection effect in each treatment can be calculated by comparing the row in which no rating was sent to that in which a rating was sent but not viewed.

Viewing Ratings. Conditional on a sender providing a rating, the receiver’s first decision is whether or not to view that rating. In three of the treatments viewing incurs a cost \$0.05, while in the *Receiver* treatment the effective viewing cost is \$0.55. Conditional on a provided rating, the fraction of rounds where receivers choose to view and incur the acquisition cost is given in the *View Rate* row of Table 2. Sent ratings are equally likely to be viewed in the *Baseline* and *Receiver* treatments. Relative to the *Baseline*, ratings are more likely to be viewed in the *Marketplace* treatment, and less likely to be viewed in the *Producer* treatment. Broadly speaking, we see that ratings are likely to be viewed in all treatments, however there are two behaviors worth emphasizing. First we see similar viewing rates in *Baseline* and *Receiver*, despite the much higher cost of viewing in *Receiver*. Second, the viewing rate in *Producer*, while lower than any of the other treatments, is much higher than predicted by theory, with well-over half of receivers choosing to view. This discrepancy relative to the theory can be explained by inspecting Table 5, where we showed substantial information provision in *Producer*. In fact, given the difference in receiver outcomes between viewing and not viewing a message, the more remarkable feature of the behavior here is that more receivers do not choose to view in *Producer*. Relative to the *Receiver* treatment, there is more to gain from viewing a message relative to the costs. As such, the data is indicative of increased wariness by receivers when the sender is known to have fully misaligned incentives.

Product Choice. After deciding whether to view, the receiver next selects one of the two urns, representing which of the two products they would choose if they do not select the outside option. Receivers at this point differ in their information sets, either knowing the specific rating sent for one of the two urns; knowing that a rating was provided, but they chose not to view it; or knowing that no rating was sent. Accounting for symmetry across the urns leads to seven distinct information sets in each treatment—the five rating values, not viewing the rating, and no rating sent.

Looking only at those receivers who are provided with a rating, we use random-effect probits to estimate the likelihood that the receiver chooses the rated urn rather than the unrated urn. Table 6(A) shows that receivers overwhelmingly choose urns rated four or five, and avoid those rated one or two. The intermediate rating of three leads to just over a third of subjects choosing the rated urn. For those who choose not to view a sent rating, the fraction choosing the rated urn is not significantly different from $\frac{1}{2}$.³⁰

While Table 6(A) indicates which of the urns is chosen after rating, Table 6(B) completes the picture by indicating the WTP for the selected urn. Here we look at the final choice between the selected urn and the outside option. Similar to our analysis of the sender’s provision decision, we break our estimation into two parts, the first over the entire data, and the second over just the last fifteen rounds of the experiment using the elicited WTP cutoffs.

³⁰However, this is mechanical, as subjects who do not view the rating do not know which urn was rated, and the interface randomly locates the two urns on their screen, so there is no possibility of coordination on urn location

TABLE 6. Receiver behavior.

(A) Probability of selecting rated product

	Baseline	Receiver	Marketplace	Producer
$m= 1$ or 2	0.059 *** (0.019)	0.014 *** (0.008)	0.083 *** (0.018)	0.079 *** (0.025)
$m= 3$	0.413 *** (0.057)	0.362 *** (0.049)	0.378 *** (0.041)	0.338 *** (0.059)
$m= 4$ or 5	0.973 *** (0.014)	0.981 *** (0.009)	0.968 *** (0.009)	0.981 *** (0.08)
Not viewed	0.504 (0.016)	0.492 (0.023)	0.491 (0.022)	0.493 (0.017)

(B) Choice to buy selected product

	Baseline		Receiver		Marketplace		Producer	
	Buy	WTP	Buy	WTP	Buy	WTP	Buy	WTP
$m= 1$ or 2	0.756 (0.036)	49.68 (1.96)	0.745 (0.033)	49.86 (1.93)	0.764 (0.029)	53.73 (2.18)	0.801 (0.039)	52.98 (2.54)
$m= 3$	0.691 (0.055)	51.82 (2.34)	0.808 (0.040)	52.38 (2.24)	0.838 *** (0.032)	50.51 (2.47)	0.724 (0.057)	46.83 (3.05)
$m= 4$ or 5	0.856 *** (0.030)	59.24 *** (1.99)	0.895 *** (0.021)	61.31 *** (1.85)	0.805 ** (0.022)	58.26 *** (2.01)	0.802 ** (0.026)	52.04 (2.28)
Not viewed	0.646 ** (0.043)	49.62 (2.12)	0.762 (0.034)	48.93 (2.13)	0.656 * (0.055)	50.97 (3.46)	0.696 (0.036)	47.56 (2.49)
No rating	0.762 (0.016)	51.20 (1.47)	0.721 (0.022)	49.06 (1.65)	0.773 (0.021)	53.75 * (1.97)	0.765 (0.019)	49.47 (2.11)

Note: For Table 6(A) the probability of product choice \hat{Z} being the rated product Y is derived from a random-effect probit estimate. For Table 6(B): Buy probability is the estimate of $\Pr\{Z \neq R | \mathcal{I}\}$, derived from a random-effect probit estimate for each information set \mathcal{I} . WTP figures are estimate for $\mathbb{E}(\Omega | \mathcal{I})$, which are derived from a random-effects panel estimate. Standard errors on all probabilities/cutoffs are in parentheses; stars represent significant differences from babbling levels ($\frac{1}{2}$ for product choices, 0.75 for the sales rate, and 50.5 for the reservation cutoffs) at: ***-99 percent confidence; **-95 percent confidence; and *-90 percent confidence.

The *Buy* columns reflect the rate at which the selected urn is chosen over the outside option, which we assess with a random-effects probit. In contrast, the cutoff estimates reflect the predicted reservation cutoff, obtained using a random-effects panel estimate. Predicted levels are indicated for the following receiver types: i) viewed a negative rating of one or two; ii) viewed the intermediate rating of three; iii) viewed a positive rating of four or five; iv) was provided a rating but did not view; and v) was not provided with a rating.

The WTP regressions, which use data from the last half of the experiment only, show a reservation cutoff close to the unconditional expected quality of $\bar{\theta} = 50.5$ following a positive rating in *Producer*. In other words, receivers in the *Producer* treatment seem to learn not to react to positive ratings with increased WTP. Though receivers eventually ignore favorable *Producer*-based ratings when deciding whether to participate in the market (i.e., whether to buy *any* product), ratings are still somewhat persuasive, as they continue to have a strong effect on the choice *between* products.

4. DISCUSSION: THIRD PARTIES AND THE STABILITY OF RATING INCENTIVES

Our analysis so far has focused on how outcomes relate to the consumers in our experimental environment, the sender and receiver of information. However, in motivating our incentives there are two other interested parties: the marketplace, that aims to maximize total sales across all products; and the producer whose product was sampled by the sender, and aims to maximize sales of their particular product. In this section we show that, taking the goals of these third parties into account, that *Marketplace* is the most-attractive long-run incentive. We summarize our analysis here, and provide supporting evidence below:

Result 4.

- (1) *Across a large range of outside-option distributions, the Marketplace incentive maximizes outcomes both for the marketplace and producer third parties.*
- (2) *Relative to the Baseline and Receiver setting, Marketplace is more stable to the covert introduction of a Producer incentive.*

We start by showing that across a range of parameterizations the *Marketplace* incentive is preferred both by the marketplaces that sell competing products, but also by the particular product manufacturer the sender has information on (represented by the urn X in our experiment). The results outlined in the previous section were tied to the specific parameterizations used in our experiments. However, as we collected willingness to pay information in the last half of our experiments, we can use this data to extrapolate our results. Of particular importance for the third parties we consider is the reservation distribution $H(\omega)$, which affects the number of receivers who are likely to buy one of the two products even without information, and those who are unlikely to buy even with a very positive rating. We now show that the receiver WTP data suggests the *Marketplace* incentives are beneficial to both third parties for a large range of reservation distributions.

We parameterize a family of variance-preserving beta distributions $\{H_\sigma(\omega)\}$ that vary the expected reservation σ , and which intersects with our experimental parameterization at $\sigma = \frac{100}{3}$. Using the WTP data and behavioral elasticities estimated across treatments, we calculate the total




FIGURE 1. Percentage Point Sales Difference from *Baseline* under reservation $H_\sigma(\omega)$

sales and same-product sales under the counterfactual distribution $H_\sigma(\omega)$, which we graph as the difference from the *Baseline* amounts in Figures 1(A) and 1(B).³¹

Figure 1(A) indicates that across a broad range of possible reservation distributions, the *Marketplace* treatment generates more total sales than any of the other incentives. Quantitatively, the marketplace benefits more when the expected reservation is higher. This behavior is in contrast with the *Receiver* environment, which has little effect at increasing total sales over the *Baseline* across the range of σ , and the *Producer* environment, which has a mostly negative effect for the marketplace.

Figure 1(B) shows *Marketplace* is also superior in the long-run for generating same-product sales. For the specific producer the sender has information on, the *Producer* incentive only outperforms *Marketplace* when expected reservations are very low. The reasoning for this can be seen by re-examining receiver reactions in Table 6. The *Producer* incentive is very effective at changing the consumer’s preferred product from the two, but bad at increasing their willingness to pay. As such, the incentive is most effective when most consumers will buy one of the two products regardless of the information provided. In contrast, the *Producer* incentive will be ineffective if the receiver needs to be convinced of a product’s absolute quality.

In the long run the *Marketplace* incentive is therefore superior for both third parties across a wide variety of outside-option distributions when receivers are fully aware of the incentives provided. However, third parties may also benefit substantially in the short run by *covertly* introducing incentives. Our next counterfactual examines the extent to which third parties benefit in the short run by offering incentive *B* to senders where receivers believe they are offered incentive *A*.³² As an

³¹Full details of the counterfactual model are provided in Appendix B. The main features of the model are that hold constant the observed message distribution (conditional on the signal x) and the receiver’s willingness to pay (conditional on each observed message), and use experimental variation in the environments to estimate subjects’ provision and viewing elasticities with respect to changes in the expected benefits of sending and viewing, respectively.

³²The Federal Trade Commission has recently cracked down on precisely this behavior (“FTC Staff Reminds Influencers and Brands to Clearly Disclose Relationship”, Federal Trade Commission, April 19, 2017). Producers can also

TABLE 7. Short-run effects

		(A) Total Sales Difference				(B) Same-Product Sales Difference				
		Sender Incentive				Sender Incentive				
		Base.	Rec.	Mkt.	Prod.	Base.	Rec.	Mkt.	Prod.	
Receiver Incentive	Base.	–	0.1%	0.3%	0.5%	Base.	–	-0.7%	1.7%	4.3%
	Rec.	-0.8%	–	0.4%	0.1%	Rec.	0.0%	–	3.0%	5.6%
	Mkt.	1.2%	0.2%	–	0.8%	Mkt.	-2.0%	-2.4%	–	2.4%
	Prod.	-0.1%	0.6%	0.4%	–	Prod.	-4.8%	-4.2%	-1.8%	–

Note: Values show percentage point changes in the frequency of sales when receivers under the row treatment face senders from the column treatment relative to the frequency when both sender and receiver are in the row treatment.

approximation for the short-run benefits, we combine the provision and ratings from one treatment with the viewing and final choice behavior from another. That is, fixing the receiver response in treatment A , we assess the third-party sale rates $\Pi(A, B)$ when matched to the sender behavior in treatment B .³³ The short-run sales effect is then calculated as $\Delta\Pi(A, B) = \Pi(A, B) - \Pi(A, A)$.

Given the receiver’s perceived environment A , Table 7 indicates the change $\Delta\Pi(A, B)$ for both total and same-product sales when the sender’s environment is covertly changed to B . Using data from the last ten rounds of the experiment we assess the rating distributions, view rates, and the relevant sales rates conditional on a viewed rating in all four treatments. The *total sales* effects are small across each comparison, and so we conclude that short-run considerations by the marketplace (at the experimental parameterization) are unlikely to trump the long-run considerations. However, we do see much stronger effects when examining same-product sales. Secretly offering senders a *Producer* incentive when receivers believe they are in the *Baseline* or *Receiver* environment leads to substantial sales increases (a 12 and 16 percent relative sales increase, respectively).

In contrast, the introduction of *Producer*-type incentives into the *Marketplace* environment generates a much smaller benefit, as receivers are already more suspicious of positive ratings, though the overall effect is still a 6 percent relative increase in sales volume. Combined with the results from our counterfactual over the long-run effects, we can conclude that *Marketplace* is the most-stable of the informative incentive structures.³⁴ Long-run profits to specific manufacturers are reduced if they introduce product-specific incentives to a *Marketplace* environment, which may serve to offset the short-run gains. In contrast, for the *Baseline* and *Receiver* environments

go one step further, avoid paying any incentives and simply fraudulently rate their own products, as in Mayzlin et al. (2014).

³³One subtlety in this counterfactual is that senders believe receivers are aware of the incentives they are being offered. Our approximation therefore assumes that if senders were aware of the asymmetries they would not change their messaging behavior.

³⁴*Producer* is the most stable of the four, but also the most inefficient.

the long-run same-product sales are increased through the introduction of *Producer* incentives, and there are much larger short-run gains from their covert introduction.

5. CONCLUSION

The effectiveness of information exchange is determined not only by the quality or accuracy of information provided, but also the quantity. A persistent problem in many information environments is one of underprovision, where parties with socially useful information fail to provide it. To solve this problem, those with relevant information can be incentivized to provide it to others. However the party offering incentives can heavily influence the total benefits from information, and how they are distributed. By examining both the extensive (quantity) and intensive (quality) margins of information transfer under four qualitatively different incentive structures, we show that all parties may actually benefit from incentives that increase the quantity at the cost of a slight decrease in quality.

Our *Receiver* environment maintains full alignment-of-interest between senders and receivers, with consumers who provide information being compensated by those using the information. This leads to a large increase in the quantity of information. However, this environment is not the optimal arrangement for the consumers making use of information, as the burden of paying for it can erode much of the benefits. The best outcome for receivers is instead found in our *Marketplace* incentive, which maintains only partial alignment-of-interests between the senders and receivers, offering a reward to senders whenever they generate a sale. This type of sales-contingent incentive has a corrosive effect on the quality of provided information (both theoretically and in our data), creating a motivation for informed parties to provide only favorable information. However, ratings under this incentive are still useful to consumers, as they truthfully reveal an ordinal feature: which of the competing products is better. Consumers, marketplaces, and producers are all potentially better off thanks to a trade-off between quality and quantity.

Our last incentive shows that trading off too much alignment-of-interest leads to sub-optimal results. Theory predicts that when senders are provided with a *Producer*-driven incentive to generate product-specific sales, the misalignment of interests between sender and receiver is complete and no information is transmitted. While our experimental data suggests a less extreme outcome, it is the least efficient of our four environments. Though inefficient, producers do have a clear short-run motivation to introduce this type of misaligned incentive for referral sales, as the deceptively favorable recommendations it generates lead to increased sales.

Finally, we use our experimental data to point out another potential benefit for our *Marketplace* institution: a little bit of commonly known misalignment can help inoculate the environment from the introduction of fully misaligned incentives. The *Marketplace* incentive creates less-credulous consumers who are not as easily exploited in the short run by secretly introduced *Producer* payments. Further research will hopefully examine the effects from other important channels excluded

from our analysis—in particular reputation, aggregation on multiple sources, and repeated interaction. Nonetheless, our results demonstrate that in a controlled one-shot environment there are potential benefits from offering incentives that decrease the accuracy of individual sources of information in exchange for an increasing in the total number of available sources.

REFERENCES

- Battaglini, Marco and Uliana Makarov**, “Cheap talk with multiple audiences: An experimental analysis,” *Games and Economic Behavior*, 2014, 83, 147–164.
- Bolton, Gary, Ben Greiner, and Axel Ockenfels**, “Engineering trust: reciprocity in the production of reputation information,” *Management Science*, 2013, 59 (2), 265–285.
- Brandts, Jordi and Gary Charness**, “The strategy versus the direct-response method: a first survey of experimental comparisons,” *Experimental Economics*, 2011, 14 (3), 375–398.
- Cai, Hongbin and Joseph Tao-Yi Wang**, “Overcommunication in strategic information transmission games,” *Games and Economic Behavior*, 2006, 56 (1), 7–36.
- Calvó-Armengol, Antoni, Joan de Martí, and Andrea Prat**, “Communication and influence,” *Theoretical Economics*, forthcoming.
- Chakraborty, Archishman and Rick Harbaugh**, “Persuasion by Cheap Talk,” *American Economic Review*, 2010, 100 (5), pp. 2361–2382.
- Charness, Gary and Nuno Garoupa**, “Reputation, honesty, and efficiency with insider information: An experiment,” *Journal of Economics & Management Strategy*, 2000, 9 (3), 425–451.
- Chen, Yan, Max Harper, Joseph Konstan, and Sherry Xin Li**, “Social Comparisons and Contributions to Online Communities: A Field Experiment on MovieLens,” *American Economic Review*, 2010, 100 (4), pp. 1358–1398.
- Chung, Wonsuk and Rick Harbaugh**, “Biased Recommendations,” August 2016. mimeo, Indiana University.
- Crawford, Vincent P. and Joel Sobel**, “Strategic Information Transmission,” *Econometrica*, 1982, 50 (6), 1431–1451.
- Dessein, Wouter and Tano Santos**, “Adaptive organizations,” *Journal of Political Economy*, 2006, 114 (5), 956–995.
- Dewatripont, Mathias and Jean Tirole**, “Modes of Communication,” *Journal of Political Economy*, 2005, 113 (6), 1217–1238.
- Dickhaut, John W., Kevin A. McCabe, and Arijit Mukherji**, “An experimental study of strategic information transmission,” *Economic Theory*, 1995, 6, 389–403. 10.1007/BF01211783.
- Fischbacher, Urs**, “z-Tree: Zurich toolbox for ready-made economic experiments,” *Experimental economics*, 2007, 10 (2), 171–178.
- Gneezy, Uri**, “Deception: The Role of Consequences,” *American Economic Review*, 2005, 95 (1), 384–394.

- Holt, Charles A and Susan K Laury**, “Risk aversion and incentive effects,” *American Economic Review*, 2002, 92 (5), 1644–1655.
- Hu, Nan, Jie Zhang, and Paul A Pavlou**, “Overcoming the J-shaped distribution of product reviews,” *Communications of the ACM*, 2009, 52 (10), 144–147.
- Kartik, Navin**, “Strategic communication with lying costs,” *The Review of Economic Studies*, 2009, 76 (4), 1359–1395.
- Lafky, Jonathan**, “Why do people rate? Theory and evidence on online ratings,” *Games and Economic Behavior*, 2014.
- Lai, Ernest K., Wooyoung Lim, and Joseph Tao-Yi Wang**, “Experimental Implementations and Robustness of Fully Revealing Equilibria in Multidimensional Cheap Talk,” November 2011. mimeo, National Taiwan University.
- Li, Lingfang and Erte Xiao**, “Money Talks: Rebate Mechanisms in Reputation System Design,” *Management Science*, 2014.
- Li, Xinxin and Lorin M Hitt**, “Self-selection and information role of online product reviews,” *Information Systems Research*, 2008, 19 (4), 456–474.
- Mayzlin, Dina, Yaniv Dover, and Judith Chevalier**, “Promotional reviews: An empirical investigation of online review manipulation,” *The American Economic Review*, 2014, 104 (8), 2421–2455.
- Vespa, Emanuel and Alistair J. Wilson**, “Communication with Multiple Senders: An Experiment,” February 2014. mimeo, University of Pittsburgh.
- Wang, Joseph Tao-Yi, Michael Spezio, and Colin F. Camerer**, “Pinocchio’s Pupil: Using Eye-tracking and Pupil Dilation to Understand Truth Telling and Deception in Sender-Receiver Games,” *American Economic Review*, 2010, 100 (3), 984–1007.
- Wang, Zhongmin**, “Anonymity, social image, and the competition for volunteers: a case study of the online market for reviews,” *BE Journal of Economic Analysis & Policy*, 2010, 10 (1).
- Wilson, Alistair J.**, “Costly Communication in Groups: Theory and an Experiment,” 2014. University of Pittsburgh working paper.

APPENDIX A. FOR ONLINE PUBLICATION: THEORETICAL FRAMEWORK

In this section we introduce our formal framework. We then introduce the four environments we will analyze, and provide a discussion of the most-informative equilibria.

To begin constructing our framework we first describe the uninformed consumer's problem: A representative consumer R (the receiver, he) faces a choice between two initially symmetric, non-divisible products, product A or product B . The consumer has a unit demand for either product, but can also choose to purchase neither, and consume an outside option with a privately known value $\omega \in \mathbb{R}$, drawn according to a CDF H . If R chooses a product $Z \in \{A, B\}$, he forgoes his outside option, and receives a product with some random quality/utility level $z \in \Theta = \{\theta_1, \dots, \theta_N\} \subset \mathbb{R}$. We denote the outside-option choice as R (as in they choose themselves), so the overall choice set for the receiver is $Z \in \{A, B, R\}$, with payoffs given by

$$w_R(Z; \omega) = \begin{cases} z & \text{if } Z \in \{A, B\} \\ \omega & \text{otherwise.} \end{cases}$$

The realized product quality z for product Z is governed by a quality control process $\phi^Z \in \Delta\Theta$, a lottery over quality levels. The consumer R 's problem is that he has no specific information on the quality control processes ϕ^A and ϕ^B for the two products. Though he has a symmetric expectation on the delivered qualities, this comes from knowledge that each producer's quality process is an iid draw from a constellation of processes, $\Phi = \{\phi_1, \dots, \phi_L\} \subset \Delta\Theta$. A producer with the process ϕ_l occurs with prior probability p_l and has an expected delivered quality $\bar{\theta}_l := \mathbb{E}(z | z \sim \phi_l)$. So without further information, the expected quality of either product is $\bar{\theta} = \sum_{l=1}^L p_l \cdot \bar{\theta}_l$.

Without additional information a risk-neutral R will choose to purchase either of the two products when $\bar{\theta} > \omega$, and abstain from a purchase when $\omega > \bar{\theta}$. Ex ante, the expected outcome for R in the absence of any other information is

$$\underline{W} = \mathbb{E}_{\omega, \phi^A, \phi^B} [\max\{\omega, \bar{\theta}\}] = H(\bar{\theta}) \cdot \bar{\theta} + \int_{\bar{\theta}}^{\infty} \omega \cdot dH(\omega).$$

The focus of our paper is social-information transmission, the degree to which R 's peers, other consumers with relevant product experience, will provide information to the marketplace that improves R 's decision making. To this end, we introduce a consumer with product experience S (the sender, she). The sender's experience with the product will be fairly natural: she has previously bought one of the two competing products, $X \in \{A, B\}$, and received a product of quality level $x \in \Theta$.³⁵ The experienced quality level x provides information on product X 's quality distribution, ϕ^X , so S 's signal (X, x) can be helpful to R 's choice.

³⁵The environment can be extended to either multiple senders $\{S_1, S_2, S_3, \dots\}$ who provide simultaneously, or a sequence of players (R_0, R_1, R_2, \dots) , where receivers $\{R_j\}_{j=0}^i$ act as sender for R_{i+1} .

Given full information sharing, R updates his prior p_l on product X 's process being ϕ_l to the posterior $q_l(x)$. The three different choices in $\{A, B, R\}$ will now be associated with three expected qualities: the product X with expected quality $\nu_X(x) := \sum_{l=1}^L q_l(x) \cdot \bar{\theta}_l$; the unsampled product $X' \in \{A, B\} \setminus \{X\}$ at $\nu_{X'}(\emptyset) := \bar{\theta}$; and the outside option at ω .³⁶ A fully informed consumer can simply choose the best option of the three, where the information exchange has two components:

- (i) Helping R discern between the products A and B . Where $\bar{\theta} > \nu_X(x)$ the information focuses R 's choice on the unsampled product X' , and for $\nu_X(x) > \bar{\theta}$ the focal product is X .
- (ii) Discerning when to buy. In those cases where $\nu_X(x) > \bar{\theta}$ the provided information increases R 's WTP, and the set of buyers would be all those with reservations $\omega < \nu_X(x)$. When $\bar{\theta} > \nu_X(x)$, the WTP is unaltered, and the set of buyers is those with reservations $\omega < \bar{\theta}$.

Full exchange of the signal (X, x) provides the upper-bound for information transfer in this environment, where the expected payoff is given by

$$\bar{W} := \mathbb{E}_{\omega, \phi^A, \phi^B, x} [\max \{ \omega, \nu_X(x), \bar{\theta} \}].$$

Against these two extremes (full frictionless exchange and no information transfer) we will consider equilibrium predictions and experimental behavior where information transmission between the consumers has frictions: where the provision of information incurs private costs to the sender; where the language available to communicate information is limited; and where acquisition of provided information to the receiver is costly.

Given sender provision costs, a second-best outcome emerges in equilibrium in which consumers, merchants, and the manufacturers of specific products benefit from greater provision. Our analysis will consider three plausible institutions for increasing provision, and examine equilibrium predictions and experimental behavior in each regime. Under a particular institution ψ , we will examine the expected gross outcome for receivers, $W(\psi) = \mathbb{E}w_R(Z; \psi)$, normalized to indicate the efficiency of information transfer as:

$$\Upsilon(\psi) = \frac{W(\psi) - \underline{W}}{\bar{W} - \underline{W}}.$$

That is, we measure efficiency as the gain in $W(\psi)$ relative to the individually rational, no-communication lower bound \underline{W} , as a fraction of the the maximal information exchange possible in a frictionless setting, $(\bar{W} - \underline{W})$.

The full sender-receiver game's timing is as follows:

³⁶Given an information set \mathcal{I} , we will use the notation $\nu_Z(\mathcal{I}) := \mathbb{E}(z | \mathcal{I})$, where expectations are taken over the process ϕ^Z and the realization $z \sim \phi^Z$.

- (i) Nature draws a state of the world $(\phi^A, \phi^B, c^S, \omega)$ where ϕ^A and ϕ^B are iid draws from a set of lotteries $\Phi \subset \Delta\Theta$; $c^S \in \mathbb{R}$ is a cost of provision drawn independently from a CDF G ; and $\omega \in \mathcal{H} \subset \mathbb{R}$ is a reservation, drawn independently from CDF H .
- (ii) S observes her send cost c^S , and obtains a single draw x from ϕ^X , where the initial product $X \in \{A, B\}$ is selected with equal probability. Given the signal (X, x) , her choice is a rating/message $m \in \mathcal{M} \cup \{m_\emptyset\}$, where \mathcal{M} is a set of meaningful ratings, and m_\emptyset is the empty message, choosing not to provide. Choosing any $m \in \mathcal{M}$ incurs the cost c^S , while choosing not to provide, $m = m_\emptyset$, is free.
- (iii) R observes whether there was a provided rating, learning either that $\{m \in \mathcal{M}\}$ or that $\{m = m_\emptyset\}$ has occurred. If $\{m \in \mathcal{M}\}$ he first chooses whether or not to acquire the information, $\rho \in \{\text{View}, \text{Not}\}$, incurring a fixed cost $c^R > 0$ only when he views.³⁷
- (iv) R observes his private outside option ω , and the precise rating $m \in \mathcal{M}$ if viewing. He then makes a choice $Z \in \{A, B, R\}$, with a realization z from ϕ^Z if a product is chosen, and ω if he selects the reservation.³⁸

Completing the specification of the game over the action choices (m, ρ, Z) , S and R 's preferences are modeled through the net utility functions

$$\begin{aligned} u_S(m, \rho, Z) &= x + \alpha \cdot w_R(Z; \omega) - c^S \cdot \mathbf{1}\{m \neq m_\emptyset\} + \psi_S(m, \rho, Z) \cdot T, \\ u_R(Z, \rho, m) &= w_R(Z; \omega) - c^R \cdot \mathbf{1}\{\rho = \text{View}, m \neq m_\emptyset\} - \psi_R(m, Z, \rho) \cdot t, \end{aligned}$$

where $\alpha > 0$ is a preference parameter reflecting a prosocial incentive to help receivers make better product choices.³⁹ The institution ψ is reflected by a conditional transfer $T > 0$ to ($t > 0$ from) the sender (receiver), conditioned on the specific events indicated through $\psi_S(m, \rho, Z) \in \{0, 1\}$ (and $\psi_R(m, \rho, Z) \in \{0, 1\}$). The transfer conditions can therefore respond to any of the action choices made in the game, where our paper will focus on four simple variations.

A pure strategy for the sender is a rating choice $\mu : \{A, B\} \times \Theta \times \mathbb{R} \rightarrow \mathcal{M} \cup \{m_\emptyset\}$, a decision on the rating to send given the signal (X, x) and provision cost c^S . The strategy for the receiver is the tuple $(\rho, \{\zeta_m\}_{m \in \mathcal{M}}, \zeta_{\mathcal{M}}, \zeta_\emptyset)$: a listening decision $\rho \in \{\text{View}, \text{Not}\}$, and a product choice $\zeta_{\mathcal{I}} : \mathcal{H} \rightarrow \Delta\{A, B\} \cup \{R\}$ for every possible information set, $\mathcal{I} \in \bigcup_{k \in \mathcal{M}} \{m = k\} \cup \{m \in \mathcal{M}\} \cup \{m = m_\emptyset\}$.⁴⁰ The relevant choices for R are: i) viewing the provided rating m , and responding with ζ_m ; ii) knowing that a rating was provided, $m \in \mathcal{M}$, but not viewing it, with a response $\zeta_{\mathcal{M}}$;

³⁷The cost c^R can be thought of as a small nuisance cost of viewing a rating. Intuitively, this is the small amount of time and effort it takes to read and comprehend a message.

³⁸We assume that the realization of ω happens after rating viewing for tractability. The first-order effect from alternating the order is to reduce viewing behavior for those R -types with high reservations.

³⁹The sender having a lexicographic preference over the receiver outcome, subordinate to her own outcome, would also suffice for our needs. The required assumption is simply that conditional on sending a non-empty message, the sender and listening receiver are strongly aligned in interest over the choice Z .

⁴⁰We allow for receivers to randomize over the two products A and B to maintain symmetry in the case where no information is revealed. However the focus of the paper is on how ratings provided by S allow us to break symmetry.

and iii) knowing that no rating was provided, and responding with ζ_\emptyset . Beliefs at every information set \mathcal{I} for S and R are given by $\lambda_S(\mathcal{I})$ and $\lambda_R(\mathcal{I})$, which are conditional distributions over the entire state $(\phi^A, \phi^B, c^S, \omega)$.

Our solution concept will be Perfect Bayesian equilibrium (PBE), $\langle \mu^*, \rho^*, \zeta^*; \lambda_S^*, \lambda_R^* \rangle$, where we focus on illustrating the most-informative and least-informative PBEs under two symmetry restrictions. The first restriction is that we restrict attention to symmetric messaging strategies over the two products, so that there exist complementary messages, and similar quality experiences (the specific draw x) lead to similar ratings, regardless of the precise sampled product identity ($X = A$ or $X = B$). The second restriction is that when agents are indifferent over $Z \in \{A, B\}$, any resulting ties are broken by an equal randomization between them, where all other strategy components are pure. The reasoning behind these restrictions is to focus on symmetry between the two products being broken by ratings ex post, and not through an ex ante coordination on a particular product.⁴¹

Given the description of the environment, we now describe the four environments we will compare. We first introduce the specific transfer rules ψ , and then provide a qualitative description of the most-informative equilibrium. The environments we will study are:

Baseline. No conditional transfers, so $\psi^S(m, \rho, Z) = \psi^R(m, \rho, Z) = 0$ for all possible action choices. So long as c^R is not too large and there are senders with low enough provisions costs, information transfer is possible in equilibrium.⁴² However, because the provision cost c^S is privately incurred by senders, there may be draws for which senders do not provide a rating, and select m_\emptyset instead. The main tension in the *Baseline* environment is on provision, where rational senders will under-provide information relative to the benefit receivers derive.⁴³ In an environment with transferable utility where α is small, a Pareto improvement will be possible at costs $c^S > 0$, where receivers pay for rating provision, which leads us to the next institution.

Receiver Transfer. One simple policy to ameliorate the provision failure in the *Baseline* is a transfer from the receiver to the sender. Should a non-empty rating be provided, each receiver viewing a rating instead pays the cost $c^R + t$, where the additional t is a transfer to the sender. Given η representative receivers for every sender and no institutional costs, the transfer received by the sender is $T = \eta \cdot t$. Transfers are conditioned on a receiver viewing the provided rating, so the institution's

⁴¹Moreover, in our experimental environment, we will explicitly prohibit any such coordination by randomly relabeling A and B between the sender and receiver.

⁴²A sufficient condition for the second part is that $G(0) > 0$, so a positive fraction of senders have no cost or enjoy sending. In the experiment subjects do receive negative cost draws for sending, with the interpretation being that these senders (net any costs) enjoy sending information to others.

⁴³In informative equilibria, the decision to provide or not can be characterized by a signal-specific cutoff $c^*(\nu_X(x); \psi)$ where the sender provides a rating in $m \in \mathcal{M}$ for all $c^S \leq c^*(\nu_X(x); \psi)$. The theoretical appendix provides additional details.

conditions are

$$\psi^S(m, \rho, Z) = \psi^R(m, \rho, Z) = \begin{cases} 1 & \text{if } m \neq m_\emptyset \text{ and } \rho = \text{View,} \\ 0 & \text{otherwise.} \end{cases}$$

In any equilibrium where receivers view, the effect is to reduce the cost of provision. When the drawn cost is c^S , the effective cost to the sender is $c^S - T$, with the effective result that the cost distribution G is shifted to the left by T . So the institution can increase provision, while maintaining alignment in the sender’s incentive. However, there are potential downsides to this incentive. Transfers of this form do not preclude coordination on low-information outcomes—in particular babbling equilibria—where the selected ratings have low information content and receivers choose not to view given the acquisition cost c^R . In fact, the equilibrium coordination problem becomes riskier for receivers than the baseline case. If R miscoordinates on the informative equilibrium he pays the additional cost t whenever S is coordinating on babbling.⁴⁴ Even with perfect coordination on informative outcomes, depending on the size of the transfer t , the *Receiver* institution alters the distribution of the surplus generated by information transfer, giving more of the gains to senders. Receivers have an option to view or not, which ensures that they cannot be forced below the lower-bound expected outcome \underline{W} . However, if the institution is fixed and the most-informative equilibrium always selected, receivers can be made worse off in the *Receiver* setting than they were in the *Baseline*. By requiring receivers to pay t to view ratings, receiver’s expected outcomes can be reduced to the lower bound \underline{W} . Though this institution can generate greater efficiency than the *Baseline*, it can also fully redistribute the full surplus, $W(\text{Rec.}) - \underline{W}$, not just the additional gain $W(\text{Rec.}) - W(\text{Base.})$.

Marketplace Transfer. Our third environment considers an intermediary marketplace player C , who sells the products A and B , producing a fixed profit when either is sold, but making no profits when consumers stick with their outside option. The market therefore has an interest in increasing WTP, but is indifferent over which specific product is purchased. To drive sales volume, the marketplace will be interested in incentivizing senders to generate sales. We assume that the form of the transfer is conditioned on the sale of either product and a viewed rating.⁴⁵ For each sale generated by the sender, the marketplace transfers t , for a total of $T = \eta \cdot t$. The condition for the transfer to the sender is

⁴⁴In a babbling equilibrium, ratings in \mathcal{M} are still provided whenever $c^S < 0$. However if the same rating m_X is provided at all signals (X, x) , then provided ratings have zero information content, and it is a best response not to view.

⁴⁵All of this information is commonly available to marketplaces. Consider an offer where the marketplace provides previous customers with a unique hyperlink to give to friends when they recommend a product. When purchases are made using the link, the marketplace observes the sale, and that the link was transmitted between the two.

$$\psi^S(m, \rho, Z) = \begin{cases} 1 & \text{if } m \neq m_\emptyset, \rho = \text{View and } Z \neq R, \\ 0 & \text{otherwise,} \end{cases}$$

while $\psi^R = 0$ at all decisions.

Similar to the receiver transfer, this environment has the effect of reducing sender’s provision costs in informative equilibria, and therefore increasing the quantity of information. However, if the conditional transfer T is large in relation to the altruism term α , the transfer has a negative effect on alignment of interest between S and R : they are now only partially aligned. Senders derive a benefit from receivers buying a product, regardless of their particular reservation ω , though they are still aligned with the receiver over which product A or B should be chosen. For α small, the most-informative equilibria restrict senders to providing just two effective ratings, with the ordinal interpretations “ $A \succ B$ ” and “ $B \succ A$.” The statement that $A \succ B$ tells the receiver that the sender’s signal is in $\{(X, x) \mid \mathbb{E}[w_R(A) \mid x] > \mathbb{E}[w_R(B) \mid x]\}$.⁴⁶

Whether receivers are better off here than in the most-informative equilibria in the *Baseline* and *Receiver* environments depends on the specifics. On the one hand, the provision of informative ratings increases, and receivers are not paying for the provision subsidy. On the other hand, the amount of information conveyed with each rating is more restricted. Similarly, providing such a transfer may or may not help increase product sales for the marketplace. From the market platform’s perspective, this conditional transfer can be sensible both as a short-run marketing response to the *Baseline* setting, and in the long run. If a large mass of final consumers have reservations just above $\bar{\theta}$, such a policy will be very effective. Every provided rating is predicted to have a positive influence on sales. A rating with the content “ $A \succ B$ ” leads to expected qualities satisfying $\nu_A(A \succ B) > \bar{\theta} > \nu_B(A \succ B)$. As such, every provided and viewed rating is predicted to increase WTP to $\nu_A(A \succ B)$ from $\bar{\theta}$. In comparison, in the aligned-interest setting, many provided ratings will simply say that a particular product is below average, having no effect on total sales.

The sales benefits from this type of incentive have a limit though, and in environments where the large mass of receivers have reservations far above $\nu_A(A \succ B)$ (or for that matter below $\bar{\theta}$), a sales-incentive will be ineffective at generating increased market revenue. In the *Baseline* and *Receiver* cases there are equilibria where senders with very good product experiences can influence high-reservation receivers. In *Marketplace* though, the lack of alignment between S and R over WTP limits the extent to which a high-reservation receivers can be persuaded to purchase—no receiver with $\omega > \nu_A(A \succ B)$ will ever make a purchase in equilibrium. Similarly, for any receivers

⁴⁶If α is very large and T small, more-informative equilibria exist. In cases where T is not large enough for senders to provide at every c^s (and α is not negligibly small) changes in the probability of providing a rating given the signal x must also be incorporated. This effect is very small under our experimental parameterization at the equilibrium, and so we will not focus on it in our discussion of the theory.

with $\omega < \bar{\theta}$, sales-conditioned incentives are ineffective marketing tools for marketplaces, as these receivers would have purchased a product even in the absence of any provided information.

Producer Transfer. Our final institution considers a transfer to senders provided by the specific producer $X \in \{A, B\}$ with which the sender has product experience (through a unique coupon in product X 's packaging, say). Here we consider the transfer T to the sender only if the receiver purchases the specific product X ,

$$\psi^S(m, \rho, Z) = \begin{cases} 1 & \text{if } Z = X, \\ 0 & \text{otherwise,} \end{cases}$$

where we again set $\psi^R = 0$ in all situations, as the transfer is paid for by the producer X . We could strengthen the condition and require that the rating be viewed as well, however the equilibrium outcome would be essentially the same.

Whenever the transfer T is large relative to α , the only equilibria are babbling and involve zero-information transfer (where the appendix provides a sufficient condition). If receivers view and alter their purchasing decisions based on ratings (either substituting between products, or increasing their WTP) then senders have a signal-independent incentive to choose the rating that increases the likelihood of an X sale. As such, in any equilibrium, ratings convey no meaningful information about X , and receivers choose not to view as $c^R > 0$. Though such incentives may generate a significant sales boost to producers in the short-run if introduced without receiver's knowledge, the effect is to poison the well in the long run as receivers learn of their presence. Any transfers made from the producer X to the sender that are linked to sales are therefore predicted to be ineffective in equilibrium.

More developed and detailed constructions for the theory, and a formal description of the equilibrium concept used to generate the theoretic predictions are provided in the appendix.

We will examine equilibria of this game (and modifications where we provide conditional transfers of utility to and from S and R) within a simple message space $\mathcal{M} = \{1_A, 2_A, \dots, K_A\} \cup \{1_B, 2_B, \dots, K_B\}$. That is, we allow a rating from 1 to K to be sent for either product. We will examine equilibria of the game that satisfy the following symmetry requirements.

Definition. A *symmetric rating equilibrium* (SRE) is a Perfect Bayesian equilibrium with actions $(\mu^*, (\rho, \zeta_m^*, \zeta_M^*, \zeta_\emptyset^*))$ and a corresponding belief system λ_S, λ_R with the property that:

- (i) If $\mu^*(X, x, c^S) = M_Y$ for $X, Y \in \{A, B\}$ then for $X' \neq X$ and $Y' \neq Y$, $\mu^*(X', x, c^S) = M_{Y'}$.
- (ii) If for $X \in \{A, B\}$ the provision strategy is $\mu^*(X, x, c^S) = m_\emptyset$ then it must be that for $X' \neq X$ that the strategy satisfies $\mu^*(X', x, c^S) = m_\emptyset$.

- (iii) For any receiver information set \mathcal{I} , whenever $\nu_A(\mathcal{I}) = \nu_B(\mathcal{I}) > \omega$, R 's final choice is a $(\frac{1}{2}, \frac{1}{2})$ mixture over $\{A, B\}$. When $\omega = \max\{\nu_A(\mathcal{I}), \nu_B(\mathcal{I})\}$, the receiver chooses R .

This definition is a refinement over the set of equilibria, and asks both players S and R to treat producers A and B symmetrically. For S the restriction asks that her selected ratings responds symmetrically to similar information on products A and B . For example, given a send cost c^S , if she gets a signal (A, x) and sends the rating m_A , then the refinement requires that given the signal (B, x) she sends the rating m_B . Note that this equilibrium refinement does not restrict the specific rating chosen. For instance, the strategy could be $\mu^*(A, x, c^S) = m_B$ (sending a B rating given a signal from A), with the symmetry requirement that $\mu^*(B, x, c^S) = m_A$. The refinement similarly enforces symmetric decisions on whether or not to *provide* a message given comparable quality draws from A and B —though this would endogenously arise given parts (1) and (3). Finally, the last part breaks indifference symmetrically, indifference between a product and the outside option is broken in favor of the outside option; indifference over the products results in an equal randomization over a product purchase

The refinement eliminates asymmetric equilibria. Such equilibria are potentially interesting, and can be more efficient than the symmetric case, as the cheap empty message becomes more informative. However, such asymmetric equilibria require fairly significant coordination, and our laboratory data indicate such outcomes are not focal. Our focus is on cases where communication from S to R breaks the initial product symmetry, rather than examining equilibrium coordination.

Proposition. *An SRE always exists for the rating game*

A babbling PBE, $(\dot{\mu}, \text{Not View}, \dot{\zeta})$ can always be constructed that meets the symmetry restrictions. For example, senders with costs $c^S \leq 0$ send the rating 1_X for all signals (X, x) , and so satisfying parts (1) and (2). Because the ratings are uninformative about the sender's observed quality level x , receivers do not listen, as $c_R > 0$. The receiver's subsequent choice is simply to maximize expected utility given consistent beliefs, and he responds at all histories by randomizing equally over the two products when $\omega < \bar{\theta}$, and choosing the outside option whenever $\omega \geq \bar{\theta}$, and so meeting part (3). However, more informative SRE can exist, and our focus now shifts to characterizing the best-case.

Given the assumption that $\alpha > 0$, senders and receivers fully agree on the choice $Z \in \{A, B, R\}$ at all possible information sets. Conditional on providing a rating $m \neq m_\emptyset$, senders and receivers have a common interest in revealing information. In fact, full-revelation of all decision-relevant information can be communicated even when the number of available ratings is lower than the number of possible signals ($K < N$). This is because whenever S 's draw (X, x) implies that product X' is strictly better than product X , then all that needs to be conveyed is the ordinal information “do not buy X .” Any space \mathcal{M} that is rich enough to distinguish between each of the

signals that imply $\nu_X(x) > \bar{\theta}$, plus a single additional rating implying $\nu_X(x) \leq \bar{\theta}$, conveys all the necessary information.⁴⁷

More specifically, given K ratings for each product, and c^R small enough for viewing to be rational, the most-informative SRE, has the following intuitive form: Given an (X, x, c^S) -type sender the rating strategy is characterized by a continuous provision cutoff $c^* : \nu_X(\Theta) \rightarrow \mathbb{R}$ and an interval partition $[\nu_0^*, \nu_1^*), [\nu_1, \nu_2^*), \dots, [\nu_{K-1}^*, \nu_K^*]$ such that

$$\mu^*(X, x, c^S) = \begin{cases} J_X & \text{if } \nu_X(x) \in [\nu_{J-1}, \nu_J) \text{ and } c_S \leq c^*(\nu_X(x); J), \\ m_\emptyset & \text{otherwise.} \end{cases}$$

This strategy for S transforms the observed signal into an expected quality level for R if he chooses product X , $\nu_X(x)$. Breaking the expected quality into K -intervals through $K - 1$ strategy parameters $\{\nu_1^*, \dots, \nu_{K-1}^*\}$ (where $\nu_0 = \min \Theta$ and $\nu_K = \max \Theta$). The sender provides the relevant rating X_k whenever the effective cost c^S/α is not too high. Defining the expected quality given a rating J_X as $\bar{\nu}_J := \mathbb{E}\{z_X \mid \mu(X, x, c^S) = J_X\}$, the receiver's decision is to always view so long as c_R is less than the expected benefit of viewing. All beliefs are pinned down by Bayes' rule, and the receiver's sequentially rational response to the message X_k is

$$\zeta_{J_X}^*(\omega) = \begin{cases} R & \text{if } \omega \geq \max\{\bar{\nu}_J, \bar{\theta}\} \\ X & \text{if } \bar{\nu}_J > \max\{\omega, \bar{\theta}\} \\ X' \neq X & \text{otherwise.} \end{cases}$$

That is, the receiver chooses: i) the reservation if it is better than both alternatives; ii) the rated product X if the ratings implies it is better than average and the drawn reservation; and iii) the unrated product if m indicates X is below average, and the reservation draw is smaller than $\bar{\theta}$. Completing the specification of the game, sequential rationality implies there are two reservation cutoffs $\nu_\emptyset = \frac{1}{2} \cdot \mathbb{E}\{\nu_X(x) \mid \mu(X, x, c^S) = m_\emptyset\} + \frac{1}{2} \cdot \bar{\theta}$ and $\nu_{\mathcal{M}} = \frac{1}{2} \cdot \mathbb{E}\{\nu_X(x) \mid \mu(X, x, c^S) \neq m_\emptyset\} + \frac{1}{2} \bar{\theta}$, where the strategies for receivers who do not receive a message, or who do not listen are

$$\zeta_\emptyset^*(\omega) = \begin{cases} R & \text{if } \omega \geq \nu_\emptyset \\ \frac{1}{2} \cdot A \oplus \frac{1}{2} \cdot B & \text{otherwise} \end{cases} \quad \text{and} \quad \zeta_{\mathcal{M}}^*(\omega) = \begin{cases} R & \text{if } \omega \geq \nu_{\mathcal{M}} \\ \frac{1}{2} A \oplus \frac{1}{2} B & \text{otherwise,} \end{cases}$$

respectively. When α is small, these two decisions will be close to the symmetric decisions in a babbling equilibrium, where $\nu_\emptyset \simeq \nu_{\mathcal{M}} \simeq \bar{\theta}$, though they must adapt somewhat to selection effects in provision as the cutoff c^* is signal-dependent.⁴⁸

⁴⁷Where there are multiple senders, ratings which delineate between differing below average levels will aid convergence.

⁴⁸Numerically solving for an equilibrium involves searching for the fixed point vector $(\bar{\nu}_1, \dots, \bar{\nu}_K, \nu_\emptyset)$ that satisfies the sender ICC and cutoff calculations below, where $\nu_{\mathcal{M}}$ can be calculated subsequently. In the limit, as $\alpha \rightarrow 0$, the provision decision becomes independent of the signal, and a fixed-point calculation can be avoided.

Incentive compatibility requires that for all $k = 1, \dots, K - 1$ the sender's messaging interval parameters satisfy

$$(1) \quad v_k^* = \mathbb{E} \left\{ \omega \mid \bar{\nu}_{k+1} \geq \omega > \max \{ \bar{\nu}_k, \bar{\theta} \} \right\}.$$

This condition essentially requires that indifference between messages k_X and $(k + 1)_X$ is set so that the loss made through receivers with reservations $\omega \in (\nu^*, \bar{\nu}_{k+1})$ choosing the product is equal to the gain from convincing receivers with $\omega \in (\bar{\nu}_k, \nu^*)$ to choose the product over the reservation (where for $\bar{\nu}_1$ we replace the average implied quality with the unconditional average). The cost cutoff for providing any rating at all is given by

$$(2) \quad c^*(\nu_X(x); J) = \alpha \cdot [\Pr \{ \omega < \bar{\nu}_J \} \max \{ \theta, \nu_X(x) \} - \Pr \{ \omega < \nu_\emptyset \} (\frac{1}{2}\bar{\theta} + \frac{1}{2}\nu_X(x)) + \kappa_J^B]$$

where

$$\kappa_J^B = \Pr \{ \omega \geq \bar{\nu}_J \} \mathbb{E} (\omega \mid \omega \geq \bar{\nu}_J) - \Pr \{ \omega \geq \nu_\emptyset \} \mathbb{E} (\omega \mid \omega \geq \nu_\emptyset).$$

The cutoff here is continuous (because of equation 1) and piecewise linear. For all signals with an implied quality for X of $\nu \leq \bar{\theta}$, the cutoff is linearly decreasing in the implied quality $\nu_X(x)$ with slope $-\frac{1}{2} \cdot \Pr \{ \omega < \nu_\emptyset \}$, the likelihood the provided rating changes the receiver's decision. For $\nu > \bar{\theta}$, the cost cutoff is increasing, where the slope depends on the specific rating J_X selected, as the rating choice changes R 's WTP. Additionally, given the form of μ^* , as $\alpha \rightarrow 0$, the cost cutoff below which the sender will provide a rating tends to zero for every possible signal x .

A.1. Transfer Mechanisms. Behavior is fairly intuitive in the baseline model, and in addition, SRE are not terribly difficult to compute. The main tension present in the model is that even the most-informative equilibria of the game are inefficient. Where the equilibrium has senders provide a message when $c^S \leq \alpha \cdot c^*(\nu_X(x))$, a sender maximizing the *joint* utility of the sender and η representative receivers would provide a message when $c^S \leq (\alpha + \eta) \cdot c^*(\nu_X(x))$. The remainder of our theory section will sketch out equilibria in the three alternative transfer mechanisms.

Receiver Transfer. In any informative equilibria where receivers view messages, this policy reduces the effective send costs for S . When a message costs c^S to provide, the effective cost when the message is viewed is $c^S - \eta \cdot T$. The most-informative equilibrium here is similar to that in the baseline with the cost distribution G shifted the left by $\eta \cdot T$. That is the cost cutoffs for providing in the receiver setting are given by

$$(3) \quad c^*(\nu_X(x); J) = \eta \cdot T + \alpha \cdot [\Pr \{ \omega < \bar{\nu}_J \} \max \{ \theta, \nu_X(x) \} - \Pr \{ \omega < \nu_\emptyset \} (\frac{1}{2}\bar{\theta} + \frac{1}{2}\nu_X(x)) + \kappa_J^R],$$

where $\kappa_J^R = \kappa_J^B$.

The first-order change here is to increasing the number of meaningful messages provided. Given the shift in the cutoffs, there are second-order effects from the shape of $G(\cdot)$ that require recalculation of the expected qualities $(\bar{\nu}_1, \dots, \bar{\nu}_K, \nu_\emptyset, \nu_{\mathcal{M}})$ and rating cutoffs.⁴⁹

Marketplace Transfer. Similar to the *Receiver* transfer, in informative equilibria the *Marketplace* incentive has the effect of reducing sender costs, thereby increasing provision. However, if the transfer is large in relation to the altruism term α , the transfer has a negative effect on the ratings possible. The sender and receiver are now only partially aligned in interests. Senders benefit from receivers increasing their WTP, though they are still aligned with the receiver over which particular product A or B should be chosen. The most informative SSRE is now characterized by senders providing just two message, with the effective meanings “ $A \succ B$ ” and “ $B \succ A$.”

Calculation of these equilibria is easier than the above and requires calculating the fixed points ν_\emptyset and $\bar{\nu}$. The cost-cutoff for provision is given by

$$c^*(\nu_X(x)) = \Pr\{\omega < \bar{\nu}\} \cdot \eta T + \alpha \cdot [\Pr\{\omega < \bar{\nu}\} \max\{\bar{\theta}, \nu_X(x)\} - \Pr\{\omega < \nu_\emptyset\} (\frac{1}{2}\bar{\theta} + \frac{1}{2}\nu_X(x)) + \kappa^V]$$

where

$$\kappa^V = \Pr\{\omega \geq \bar{\nu}\} \mathbb{E}(\omega | \omega \geq \bar{\nu}) - \Pr\{\omega \geq \nu_\emptyset\} \mathbb{E}(\omega | \omega \geq \nu_\emptyset).$$

The WTP given a provided message is calculated as

$$\bar{\nu} = \Pr\{\nu_X(x) \leq \bar{\theta} | m \in \mathcal{M}\} \cdot \bar{\theta} + \Pr\{\nu_X(X) > \bar{\theta} | m \in \mathcal{M}\} \cdot \mathbb{E}\{\nu_X(x) | m \in \mathcal{M}, \nu_X(X) > \bar{\theta}\},$$

where in the limit as $\alpha \rightarrow 0$ we can remove the conditioning on provision.

Producer Transfer. For the producer transfer, if no message breaks the belief’s symmetry, then there is no value in viewing, and the equilibrium is uninformative. For any rating m that is viewed in an SRE that does break the receiver’s symmetry of belief, this must produce a gain in sales of at least $\frac{1}{2}H(\bar{\theta})$ —one of the two products must have an expectation of at least $\bar{\theta}$, and sequential rationality implies this product is bought $H(\nu(m))$ of the time. For our experimental parameterization, a sufficient (though not necessary) condition for babbling being the only outcome is met by $\frac{1}{2} \cdot T \cdot H(\bar{\theta}) > \alpha \cdot \max_{\theta \in \Theta} |\theta - \bar{\theta}|$. That is the minimum gain through the incentive from making the owned product focal exceeds the maximal altruistic difference in expected quality.

⁴⁹ In our experimental setting, these effects are only non-negligible for ν_\emptyset , due to a very small likelihood of no provision.

APPENDIX B. RESERVATION DISTRIBUTION SHIFTS

Different choices for the reservation distribution $H(\omega)$ will directly affect the rate of sales, and therefore the desirability of the different institutions to each party in the market. Changing this distribution affects both babbling and informative equilibria. Efficiency, receiver welfare and the same-product sales rate are all directly related to the distribution H .

As a counterfactual exercise, we examine outcomes in the best-case equilibria across a family of reservation distributions. Where the experimental distribution H is (approximately) parameterized by a beta distribution (rescaled to a $[0, 100]$ support) with a mean reservation of $100/3$ and standard deviation of $50\sqrt{2}/3$, we now look at the effects across a family of beta distributions $\{\hat{H}_\sigma(\omega)\}_\sigma$ indexed by a scalar parameter σ . Setting the beta distribution's two shape parameters we examine variance-preserving shifts of the mean, where the expected reservation is set to the parameter σ , where the standard deviation is fixed to that in the experimental distribution H .

Figures 2(A,C,E,G) indicate theoretical differences (in the most-informative equilibrium) relative to the *Baseline* environment over efficiency, receiver welfare, total sales and same-product sales, respectively. For instance, in Figure 2(A) we indicate the efficiency difference $\Upsilon(\sigma; \psi) - \Upsilon(\sigma; \text{Base.})$, and in Figure 2(C) the receiver welfare difference $\Upsilon_R(\sigma; \psi) - \Upsilon_R(\sigma; \text{Base.})$.⁵⁰ In each figure, the dashed vertical line at $\sigma = 100/3$ indicates the experimental parameterization.

The theoretical comparisons here assume the most-informative equilibrium, however we know that subjects are more honest than predicted, and also deviate from theory in their rating, viewing and WTP behavior. We now extrapolate from the observed experimental behavior across the family of WTP distributions using the elicited reservation cutoffs in the last half of each experimental session. To do this, we make some assumptions on how behavior changes as we alter H . First, we assume changes in H do not affect the ratings selected \hat{m}_{it} after each signal x_{it} , or the reservation cutoff $\bar{\Omega}_{jt}$ selected at each information set \mathcal{I}_{jt} .⁵¹ Second, we incorporate changes to both the provision and viewing rates in each environment through a logit response to the expected benefits/costs at σ . Using the experimental variation in the environment, we estimate subjects' provision and viewing elasticities with respect to changes in the expected monetary incentives and realized monetary benefits of viewing, respectively. Table 8 contains details of the estimation procedure and

⁵⁰The scale and normalization for efficiency are also functions of σ . That is, the expected upper bound given full information, \bar{W}_σ , and the expected outcome under babbling, \underline{W}_σ , are both affected by the choice of $H(\cdot)$, where $\bar{W}_\sigma - \underline{W}_\sigma$ is decreasing in σ . The discontinuity in the *Receiver* comparison stems from a critical σ after which it is no longer rational to view ratings, so that the babbling equilibrium is the most-informative.

⁵¹Given the specified reservations we calculate the expected total sales rate as $\text{Sales}(\sigma) = \frac{1}{|\mathcal{S}|} \sum_{j \in \mathcal{S}} \frac{1}{10} \sum_{t=21}^{30} \hat{H}_\sigma(\bar{\Omega}_{jt})$, while the rate of same product sales is calculated as $\text{Product}(\sigma) = \frac{1}{|\mathcal{S}|} \sum_{j \in \mathcal{S}} \frac{1}{10} \sum_{t=21}^{30} \hat{H}_\sigma(\bar{\Omega}_{jt}) \cdot \mathbf{1}\{Z_{jt} = X_{it}\}$. We here focus on the last ten rounds, allowing rounds 16 to 20 for subjects to gain experience with the WTP elicitation.

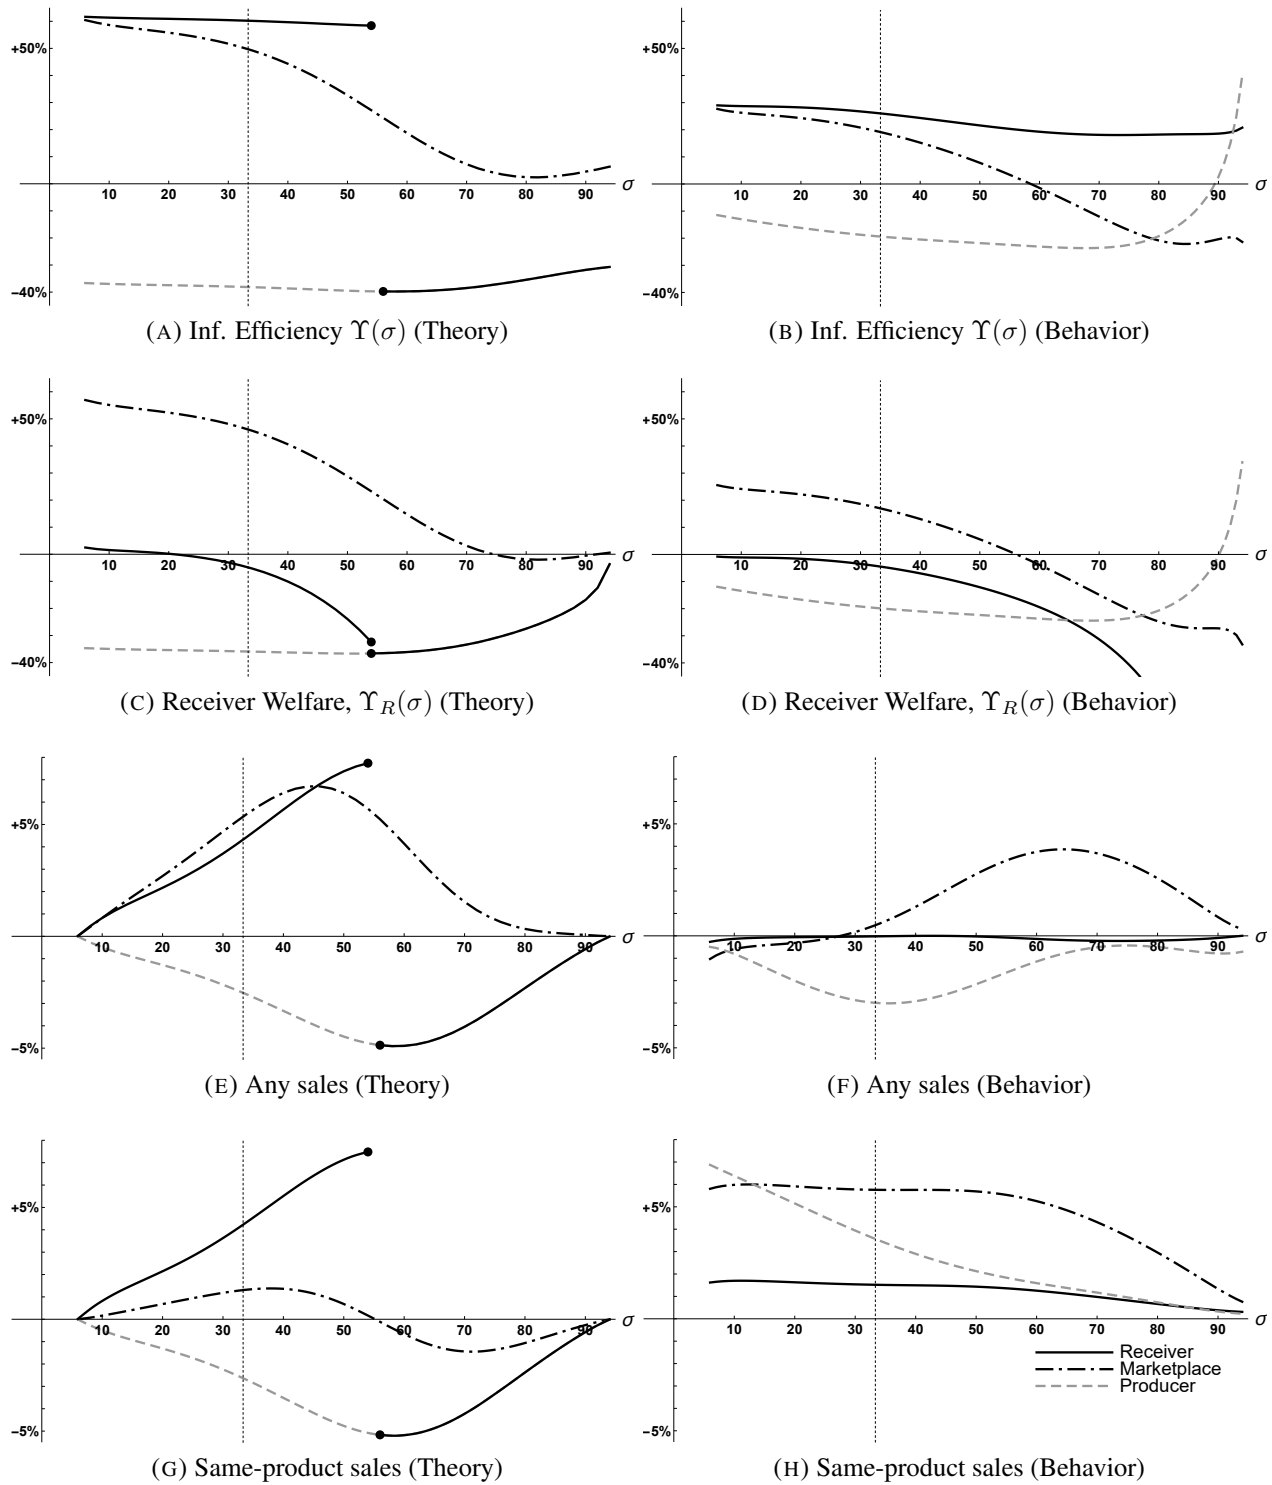


FIGURE 2. Changes relative to *Baseline* with a reservation distribution \hat{H}_σ

assessed parameters. We use this estimated logit response to extrapolate how the provision and viewing rates vary as σ changes.⁵²

TABLE 8. Counterfactual Provision/View Response Model

	Provision		Viewing	
Constant, C	-0.405 (0.072)	***	0.356 (0.166)	***
Incentive Difference, $\Delta U(\psi, \lambda)$	0.00674 (0.0071)	***		
Payoff Difference, $\Delta V(\psi, \lambda)$			0.0211 (0.0041)	***
Reciprocity, $\delta_V(\psi)$			$\left\{ \begin{array}{ll} 0.677 & \text{if R or M treatment} \\ 0 & \text{otherwise} \end{array} \right.$	***
Probability Model	$\frac{\exp \{C_p(\psi) + 0.00674 \cdot \Delta U(\psi, \lambda)\}}{\exp \{C_p(\psi) + 0.00674 \cdot \Delta U(\psi, \lambda)\} + 1}$		$\frac{\exp \{C_V(\psi) + 0.0211 \cdot \Delta V(\psi, \lambda)\}}{\exp \{C_V(\psi) + 0.0211 \Delta V(\psi, \lambda)\} + 1}$	

Note: Results are estimated from a logit regression in the last 10 rounds where $\Delta U(\psi, 33.3)$ (the difference in the expected incentive payment from providing/not) and $\Delta V(\psi, 33.3)$ (the expected payoff difference from viewing/not) are calculated from the empirical response in rounds 21–30. The counterfactual model uses changes in the monetary incentive to provide/view to model the change in each. N.B. The ensuing $\Pr \{\text{Provision} | \lambda, \psi\}$ incorporates the viewing probability $\Pr \{\text{View} | \lambda, \psi\}$ through the $\Delta U(\psi, \lambda)$ term, as the incentives can be a functions of viewing rates.

Figures 2(B,D,F,H) indicate the differences in efficiency, receiver welfare and sales relative to the *Baseline* in these behavioral extrapolations. Again, in each panel, the vertical dotted-line indicates the distribution used in the experiment, $H(\omega)$. The behavioral counterfactuals indicate fairly constant efficiency gains in the *Receiver* environment, where Figure 2(B) indicates increasing gains in *Receiver* over *Marketplace* as σ increases. This effect stems from greater WTP following positive ratings in both aligned-interest environments, which becomes more and more important as average reservations increase. However, other than total efficiency, which reflects the sender’s preference for *Receiver* across the entire range of σ , the behavioral extrapolations indicate the other participants (receivers, marketplaces, producers with existing consumers) prefer the *Marketplace* institution for the large central range of σ . Though the preference ordering changes somewhat when average reservations are especially high or especially low, we see the results as favoring the use of sales-linked incentives for information sharing, if marketplaces are willing to pay for these programs.

⁵²As reservations increase, receivers viewing behavior decreases as ratings are less useful to those unlikely to purchase a product. The combined effect of reduced viewing and reduced sales also alters the benefits to providing a rating in each of our environments, and so provision falls. The two effects are modeled in the counterfactual through a response to the expected difference in payoff from providing/viewings where $\text{Provision}(\sigma; \psi)$ is therefore a function of both $\text{Viewing}(\sigma)$ and the environment ψ . This exercise is analogous to calculating a logit quantal response equilibrium, where we fix the observed rating and WTP behavior.

Figure 2(F) indicates the total sales difference between the *Baseline* and each incentive environment, measured as the percentage of receivers who buy any product. Similarly, Figure 2(H) provides information on the change in sales for the specific product X_{it} sampled by the sender. In the experiment we see similar total sales in the *Baseline*, *Receiver* and *Marketplace* treatments in the last ten rounds, and somewhat reduced sales in *Producer*. Figure 2(H) illustrates the relative sales increases in *Marketplace* as we increase the average reservation away from the experimental parameterization (though, in absolute terms sales decrease as σ increases).

Though the sales gains in *Marketplace* are relatively small at our parameterization, as outside options increase, so too does the relative sales benefit to the marketplace. Once we raise the expected outside option to $\sigma = \bar{\theta}$, the *Marketplace* treatment's total sales fall to 0.543, where this is significantly larger than either the *Baseline* or *Receiver* treatments at 0.514 and 0.513, respectively. The relative effect in *Marketplace* grows even larger as average reservations increase. For $\sigma = 64$, though the total fraction of the population making a purchase is lower in all treatments, the relative gain in the *Marketplace* setting is largest, with a 12 percent increase in sales volume over the *Baseline*. By contrast, the *Receiver* environment produces a negative and insignificant sales effect relative to the *Baseline* for higher values of σ , primarily due to reduced viewing behavior.⁵³

Moreover, the same-product sales rate in Figure 2(H) is larger in *Marketplace* than *Producer* when average reservations are high. From the point of view of manufacturers, *Producer* is only superior to *Marketplace* when consumers are very likely to purchase anyway. Though specific producers might benefit in the short-run from covertly offering such incentives to their previous customers, there is little or no long-run benefit when the status quo is a *Marketplace* environment.

APPENDIX C. FOR ONLINE PUBLICATION: BEHAVIORAL EFFICIENCY LOSSES

Given the observed sender and receiver behavior, we can examine which elements of that behavior produce the largest losses or gains in overall information efficiency. We use the strategy-method data provided in the last 15 rounds of each session to examine how efficiency and receiver welfare change as we modify individual components of observed behavior, holding the rest of the behavior constant. We examine three shifts in behavior: i) senders using an empirical best response (provision and rating selection); ii) receivers always choosing to view provided ratings; and iii) receivers using the risk-neutral empirical best response WTP cutoffs for low ratings, high ratings and when no rating is sent. We summarize the changes in total information efficiency for each counterfactual behavior in Table 9.

⁵³Our model for provision and listening decreases with σ at a greater rate in the *Receiver* treatment. In the experiment, the difference in the rate of provision in the last ten rounds between *Receiver* and *Baseline* is 0.268, while listening rates are near identical. At $\sigma = 65$ the counterfactual model reduces the provision difference to 0.216 (reflecting a drop in the expected receiver transfer stemming from reduced viewing). Listening rates are modeled as falling by 0.104 in *Baseline*, and 0.157 in *Receiver*.

TABLE 9. Efficiency Gains/Losses from Component Best-Responses (last 15 rounds)

	Baseline	Receiver	Marketplace	Producer
Observed Info. Efficiency	7.58%	25.72%	17.36%	-7.41%
Sender Empirical BR	7.56%	22.90%	40.37%	-9.66%
All Viewing	7.98%	9.10%	3.15%	6.19%
All Viewing (Rec. Welfare)	7.38%	-1.01%	2.85%	5.35%
Receiver BR Cutoffs: 1,2	2.06%	1.87%	5.82%	2.43%
Receiver BR Cutoffs: 4,5	4.84%	9.56%	14.66%	8.62%
Receiver BR Cutoffs: m_\emptyset	6.33%	3.13%	4.98%	10.51%

Note: All counterfactuals represent changes in efficiency relative to the first row. Each counterfactual holds constant the observed sender and receiver behavior in the last 15 rounds, except for: i) the Sender Best Response counterfactual considers a change to sender behavior where the optimal cost cutoff and rating choice is used; ii) the All Viewing counterfactual assumes receivers view all sent ratings. iii) the Cutoff counterfactuals consider the gain from changing receiver cutoffs to the empirical best response cutoff.

The very first row provides the analog to the information efficiency given in Table 2, where we calculate the value using just the last 15 rounds (and where we incorporate the additional precision on intensities from the cutoffs into the calculation). This treatment efficiency will be the level against which we measure the changes in efficiency for each counterfactual.

Our first shift in behavior holds constant the receivers' response (viewing, urn and reservation decisions) but has senders use an empirical best response, one that maximizes their expected payoff given the observed receiver behavior. That is, we calculate the best response rating $\hat{\mu}(X, x, c_t)$ for all signals x and costs c^S , where the rating strategies are summarized in the appendix. Changing the sender's response in this way affects both the quantity and quality of provision. Though there are small positive effects in the *Baseline*, the largest benefits from this change are within the *Receiver* and *Marketplace* treatments, where the vast majority of the efficiency gains are from increased rating provision. The relative efficiency decrease in the *Producer* treatment stems from optimally behaving senders being dishonest more often than we observe, thereby reducing rating quality and final receiver outcomes.

In the second counterfactual, *All Viewing*, we force every receiver to view all provided rating, fixing their response to each particular rating to the average observed behavior. There are moderate gains here in the *Receiver* and *Baseline* treatments, however once we account for the cost of viewing (in the receiver welfare row), the potential efficiency gains in *Receiver* actually become a small loss.⁵⁴

⁵⁴Though it is an empirical best response to view a rating in all treatments, it is not so if receivers react sub-optimally to the provided rating.

Finally, the last three rows examine efficiency changes if receivers use the empirical best-response WTP in reaction to: low ratings (ones and twos); high ratings (fours and fives); and no rating. Here the largest absolute gains are through the response to high ratings, as subjects exhibit lower WTP after seeing a four or five than is optimal in all treatments. Here the figures reflect unconditional efficiency increases, but conditional on a positive rating being provided the gains in each treatment are more comparable: a 49 percentage point gain in *Baseline*, 48 points in both *Receiver* and *Marketplace*, and a 42 percentage point gain in *Producer*.

The last row in the table indicates the efficiency gains from changes to the receivers response following no provision, where the largest total efficiency gains are possible in the *Baseline* and *Producer*, which have the lowest probabilities of provision. Controlling for the rate at which information is not provided, the conditional efficiency gains from receivers best responding are a ten percentage point efficiency gain in *Baseline*, an eight point gain in *Receiver*, and much larger gains of fourteen and twenty percentage points in *Marketplace* and *Producer*, respectively. These results therefore point to large increases if receivers were to understand the provision selections effects under the misaligned incentives.

Result 5 (Improved Efficiency). *The largest potential increase in total efficiency from participants best responding to observed play stems from increases to rating provision in the Receiver and Marketplace environments.*

TABLE 10. Risk-Averse Equilibrium Predictions

	Baseline (B)	Receiver (R)	Marketplace (M)	Producer (P)	Comp. Static (M)
Distinct Messages	10	10	2	0/Babbling	
Message Provision	0.315	0.975	0.899	0.304	R>M>B>P
Info. Efficiency, Υ	31.6%	96.8%	68.2%	0.0%	R>M>B>P
Rec. Welfare, Υ_R	29.7%	29.6%	62.6%	0.0%	M>B>R>P
Sales	0.746	0.801	0.764	0.7191	R>M>B>P
Prod. Recurrence	0.396	0.463	0.403	0.360	R>M>B>P
Conditional on no provided rating					
Info. Efficiency, Υ	-1.3%	-4.46%	-0.7%	0.0%	P>M>B>R
Conditional on provided rating					
View Rate,	1.0	1.0	1.0	0.0	B~R~M>P
Info. Efficiency Υ	102.8%	99.4%	76.0%	0.0%	B> R>M>P
Rec. Welfare Υ_R	96.6%	30.6%	69.8%	-6.3%	B>M>R>P
Sales	0.805	0.803	0.770	0.719	M>B>R>P
Conditional on provided rating and sale					
Prod. Recurrence	0.588	0.580	0.530	0.500	B>R>M>P

Note: Figures assume expected-utility Bernoulli function $v(x) = \frac{x^{1-r}}{1-r}$ for $r = 0.3$ (see Holt and Laury 2002 for empirical evidence of typical risk-aversion levels). Babbling given this risk aversion level has an efficiency of -0.6 percent relative to the risk-neutral baselines, while full communication has an efficiency of 99.4 percent. 10-message equilibria send regions 1–47, 48–60, 61–72, 73–86, 87–100. Marketplace equilibrium is to send 5_Y for $Y \neq X_5$ and θ_S^X in 1–47, and 5_X for θ_S^X in 48–100.

TABLE 11. Risk-Neutral Sender’s Empirical Best Response

	Baseline	Receiver	Marketplace	Producer
Provision	0.334	0.904	0.882	0.568
Optimal Ratings	4, Same Urn, $x > 50.5$ 4, Other Urn, $x \leq 50.5$	5, Same Urn, $x > 55$ 4, Same Urn ,55 $\geq x > 50.5$ 4, Other Urn, $x \leq 50.5$	4, Same Urn, $x > 50.5$ 4, Other Urn, $x \leq 50.5$	5, Same urn, all x

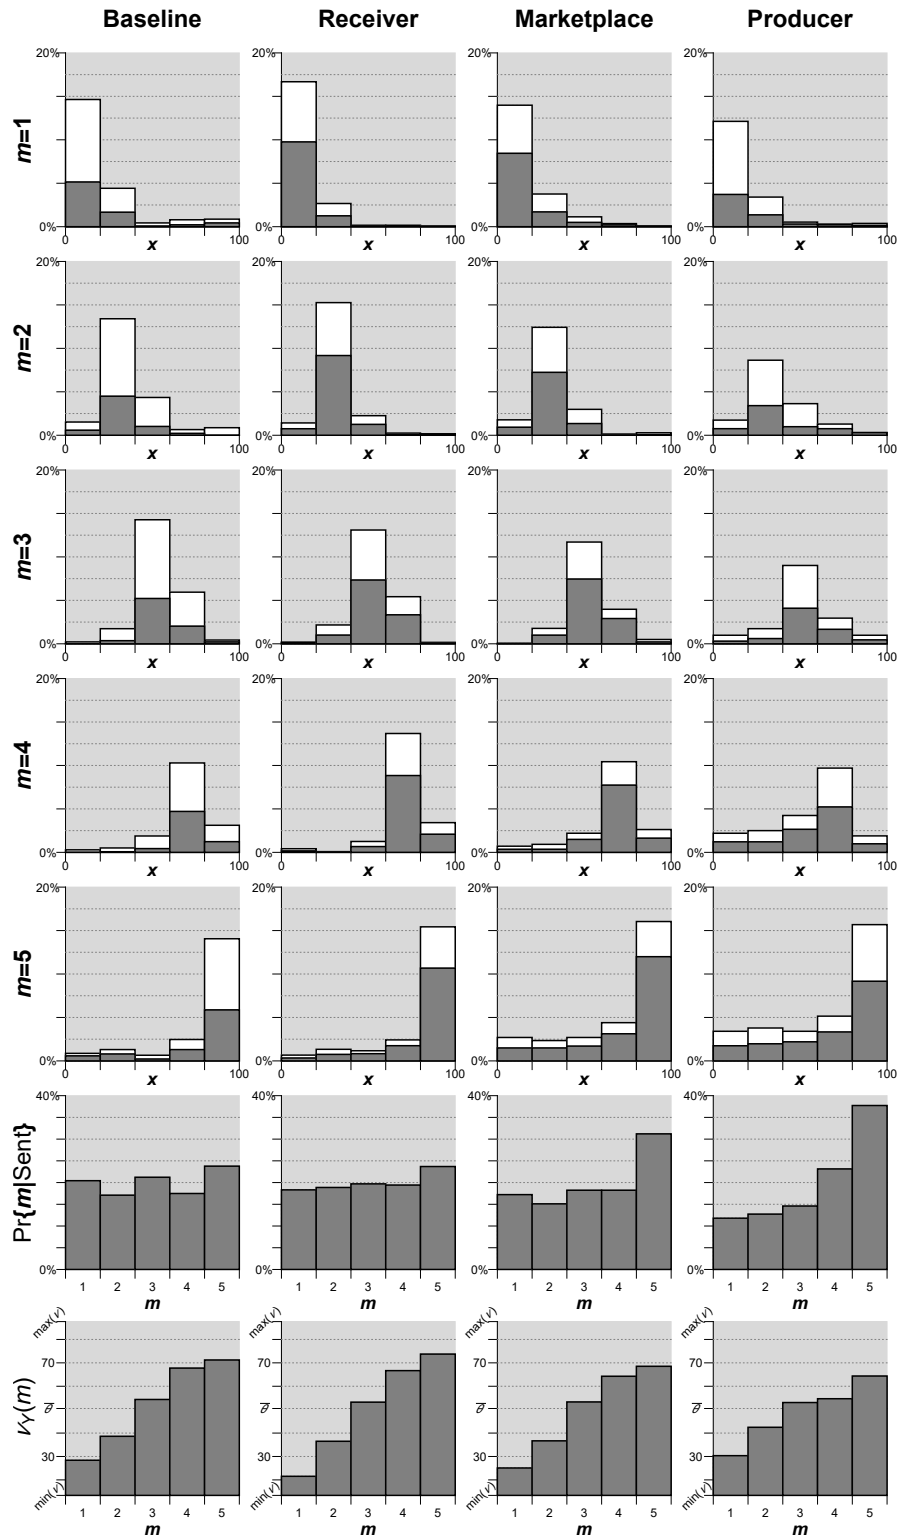


FIGURE 3. Sender behavior

Note:

Each column corresponding to a treatment. The first five rows show histograms of the signals drawn by senders who choose each of the ratings one through five. For each *selected* rating, the height of each white bar indicates the fraction of all rounds with a signal in that bin. Gay bars indicate the fraction of rounds where the selected rating was ultimately sent. White shaded regions indicate selected ratings that were ultimately not sent after the sender observed the cost of rating.




FIGURE 4. Provision as a function of rating cost .

Note: Probability of provision as a function of the send cost is recovered via a Gaussian kernel regression. Subject cutoffs in (B) are estimated from the final 5 rounds of the experiment; distributions include smoothing for clarity