EXPERIMENTING WITH INCENTIVES FOR INFORMATION TRANSMISSION: QUANTITY VERSUS QUALITY

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 the quality and quantity of information. Our findings indicate substantial efficiency gains are possible relative to no incentives, even when the incentives damage the preference alignment between those sending and receiving information. In particular, our results suggest that a partially aligned incentive (similar to a referral or sales commission) can encourage 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 sharing their experiences. 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. 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 the different parties that might pay for information provision. The three environments vary over the conditions for a payment, changing the between those sending and receiving information. The end effect is an incentive that varies between full alignment, to only partial alignment of interests, to no alignment at all. Examining how each incentive influences both the quantity and quality of information, first in theory and then in practice in the laboratory, we show how the intensive and extensive margins of information transmission might be efficiently traded off.

Date: October, 2019.
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. This work was supported by research funds from the University of Pittsburgh and Lafayette College. 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. Declarations of interest: none.

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 bias on the sender's part can be quite destabilizing. In contrast, the experimental literature has documented a much smaller 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. Building upon this result, we demonstrate significant potential benefits from trading off some alignment of interest to have provision of information subsidized by a third party. In contrast, while our incentive that maintains common interests over provided information does lead to more total information transfer, we show that the net benefit to receivers can be small or even negative once the costs of making the subsidy are accounted for. Finally, for an incentive that produces full misalignment of interests, we find the worst information transfer rates across our experiments.

While incentives to provide very distorted information do seem to fail in the long run when all parties have common knowledge, it could be that third parties still choose to covertly offer such incentives. For example, where celebrities endorse products without disclosing that they have been compensated by manufacturers. If receivers are unaware of the incentives supporting information (and falsely believe senders have common cause) we show that third parties are able to produce large short-run gains, but simultaneously poison the informational well for the long-run. This insight prompts the final step in our analysis, where we demonstrate that some commonly known conflicts of interest may in fact be desirable, serving both to inoculate consumers against the introduction of more-distortive incentives, and to motivate those running marketplaces to discourage such incentives from being offered.

Our results highlight a previously unrecognized advantage to commission-like incentives in multi-product environments, and may help explain their prevalence. Commissions provide the incentives for stockbrokers and realtors to provide advice over two competing stocks or houses (an explicit commission on eventual purchases), a restaurant server providing recommendations over which dessert to order (an increased tip), or product review websites (referral commissions via hyperlink). In each case, the informed party can have an aligned interest with the decision maker over which of the competing options is better (a particular stock, house, flavor of ice cream, or consumer product), but be misaligned through the commission over the decision to buy anything 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 directly and explicitly paid for providing advice. Our paper posits the idea that this partial misalignment could in fact 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 on receivers' final 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? (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 conditional transfers that reward senders providing information.

Our first transfer (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 transfers trade off alignment of interests by turning to a third party to pay for the subsidy. Our second incentive, which we call *Marketplace,* as it aligns senders with the sales interests of a platform where competing products are sold, rewards senders for sales of any available product. While creating some misalignment between senders and receivers over whether to buy anything at all, the incentive maintains an ordinal alignment of interest over which of the competing products is better. For example, while a broker on commission would always want you to make a purchase to trigger their bonus, they would, ceteris paribus, prefer you purchase the superior option.

Finally, our third conditional transfer, which we call *Producer*, creates full misalignment in preference. It does this by directly incentivizing the sender to endorse a particular final choice, making the transfer only if a specific option or product is chosen. This causes a complete conflict of interest between the different sides of the information transfer, instead aligning the sender with the interests of a particular producer wishing to increase their own sales.

Across these three transfers, our results provide clear evidence of the potential benefit from an informational quantity/quality tradeoff. In particular, we show that the *Marketplace* incentive results in many more senders choosing 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 model and experimental design in Section 2. In Section 3 we outline our results, and then in Section 4 we extend them through a counterfactual exercise examining third-party interests across a key parameter of our model: the likely range of outside options. Finally, in Section 5 we conclude and point to areas of future research.

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 in 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]

[1]A 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.

[2]See 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); for an extention to a dynamic setting see Vespa and Wilson (2019).

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

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 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 examine a representative risk-neutral decision maker R (the receiver, he) making a choice $Z \in \{A, B, R\}$ between one of two products (ex ante identical) with associated lotteries over the delivered quality (ϕ^A and ϕ^B) or an outside option (the choice R, providing a known private-value payoff $\omega \in \Omega \subseteq \mathbb{R}$, with initial distribution H). Given their choice Z, the receiver yields a gross payoff of $w_R(Z, \omega)$, equal to the reservation ω if $Z = R$ or to a realization $z \in \mathbb{R}$ from ϕ^Z if they choose one of the two products.

As a potential aid to his decision, the receiver is matched to a peer S (the sender, she) with prior experience with one of the two products $X \in \{A, B\}$, having obtained a realization x from ϕ^X. A costly communication channel is provided between S and R: with the sender sinking a privately known cost $c_S \in \mathbb{R}$ (distribution G) when sending any non-empty message $m \in \mathcal{M}$, and the receiver incurring a fixed cost $c_R > 0$ should he choose to view a provided message ($\rho \in \{\text{View}, \text{Not}\}$). The main thrust of the theoretical model (and our experimental environment that is built upon it) is to examine the possibilities for information transfer in equilibrium between

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

the sender and receiver, and the extent to which a conditional transfer $\left(\psi^S(m,Z,\rho),\psi^R(m,Z,\rho)\right)$ to the sender or from the receiver might improve matters.

The receiver's payoff in the model is

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

while the sender's payoff is

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

The sender is assumed to derive utility from both her own lottery outcome x and sending choices, but also from the receiver's outcome. A positive spillover with weight α (with $1>\alpha>0$) reflecting an innate desire to help the receiver make a better choice.[6]

The timing of the game is as follows:

(i) Nature draws the lottery pair $\left(\phi^A,\phi^B\right)$, where each is an iid draw from $\Phi=\left\{\phi_1,\phi_2,\ldots,\phi_N\right\}$ with $p_r=\Pr\left\{\phi_r\right\}$.

(ii) The initial product $X\in\{A,B\}$ is selected with equal probability, and the sender obtains a realization x from ϕ^X and a send-cost c_S.

(iii) The sender decides whether or not to provide a meaningful message $m\in\mathcal{M}$ (and incur the send cost) or to send the costless empty message $m_\emptyset$.

(iv) The receiver is informed whether a non-empty message was sent, and if so makes a viewing decision ρ. This leads to a receiver information set $\mathcal{I}\in I$

(v) The receiver draws their outside-option value ω and makes a final choice $Z\in\{A,B,R\}$.

A strategy profile in this environment is therefore a tuple $<\mu,(\rho,\zeta)>$ with the following elements: (i) A message strategy $\mu:\Theta\times\{A,B\}\times\mathbb{R}\rightarrow\mathcal{M}\cup\{m_\emptyset\}$ for the sender, where $\mu(x,X,c_S)$ takes into account the precise signal draw (x,X) and the cost of sending (c_S). (ii) A contingent viewing decision for the receiver $\rho\in\{\text{View},\text{Not}\}$ in case a non-empty message is sent. (iii) A final choice for the receiver $\zeta:I\times\mathbb{R}\rightarrow\{A,B,R\}$, where $\zeta\left(\mathcal{I},\omega\right)$ takes into account the receiver's information set $\mathcal{I}$ and the reservation value ω. For any messaging behavior, the sequentially rational choice for the receiver (breaking ties in favor of the reservation and symmetry) is always to choose the best

[6]Note that for the model α can be made arbitrarily small. The modeling choice here is to start out with a baseline model without conditional transfers where senders are fully aligned over receiver final choices. Modulo the communication costs all parties strictly benefit from the transfer of decision-relevant information, but with the costs the environment reflects an informational public good.

of his three options:

$$\zeta^{\star}(\mathcal{I},\omega)=\begin{cases}R & \text{if } \omega \geq \max\left\{\mathbb{E}\left(\phi^{A}\,|\mathcal{I}\right),\mathbb{E}\left(\phi^{B}\,|\mathcal{I}\right)\right\},\\ A & \text{if } \mathbb{E}\left(\phi^{A}\,|\mathcal{I}\right)\text{>}\max\left\{\omega,\mathbb{E}\left(\phi^{B}\,|\mathcal{I}\right)\right\},\\ B & \text{if } \mathbb{E}\left(\phi^{B}\,|\mathcal{I}\right)\text{>}\max\left\{\omega,\mathbb{E}\left(\phi^{A}\,|\mathcal{I}\right)\right\},\\ \frac{1}{2}A\oplus\frac{1}{2}B & \text{otherwise.}\end{cases}$$

In the (B)*aseline* model without transfers (so $\psi_B^S(m,Z,\rho)=\psi_B^R(m,Z,\rho)=0$) the sender and receiver have *cardinal* alignment over any information provided. By this we mean that were there no frictions to communication and the message space rich enough, the sender would choose to fully reveal their signal to the receiver, $\mathcal{I}_x=(X,x)$. The main strategic tension in the absence of transfers is therefore whether *any* information is provided. The sender's messaging cost c_S is privately incurred, and if the weight α placed on the receiver's outcome is small, the sender will not want to incur the cost c_S. In many parameterizations (in particular our experimental ones) the resulting effect will be an underprovision of informative messages.

To formalize the underprovision, consider the information $\mathcal{I}_x$ possessed by the sender. If the sender was able to perfectly reveal this information the receiver's expected outcome would improve relative to no information by $\Delta w_R(\mathcal{I}_x):=\mathbb{E}_\omega\left[w_R\left(\zeta^\star(\mathcal{I}_x,\omega)\right)-w_R\left(\zeta^\star(\emptyset,\omega)\right)\right]$. A sender deciding whether or not to fully reveal would choose to incur the cost whenever $c_S\leq\alpha\cdot\Delta w_R(\mathcal{I}_x)$. In contrast, consider a social planner making the provision choice where the receiver is representative for $K>1$ decision makers. The social benefit generated by information revelation is $(K+\alpha)\Delta w_R(\mathcal{I}_x)$ while the social cost is $(K\cdot c_R+c_S)$. The social planner would therefore choose to reveal $\mathcal{I}_x$ whenever $c_S\leq(K+\alpha)\cdot\Delta w_R(\mathcal{I}_x^A)-K\cdot c_R$. The qualitative best-case equilibrium prediction in the *Baseline* environment without transfers is therefore that the message space $\mathcal{M}$ can be fully utilized to send high-quality information;[7] however, where α and c_R are small, information will be underprovided, relative to the first best.

2.1. **Three Alternative Institutions.** Though the environment with no transfers is predicted to be inefficient, there are simple incentive schemes that might increase the exchange of information. We examine three distinct conditional transfers of an amount T to the sender: A (R)*eceiver* transfer provided by the receiver player, conditioned on a viewed message, $\psi_R^S(m,Z,\rho)=T\cdot\mathbf{1}\left\{m\neq m_\emptyset,\rho=\text{View}\right\}$. A (M)*arketplace* transfer from the platform where A and B are sold, contingent on any product sale, $\psi_M^S(m,Z,\rho)=T\cdot\mathbf{1}\left\{Z\neq R\right\}$. And finally, a (P)*roducer* transfer

[7]With a one-to-five rating scale, one first-best usage would be to use a one signal as an indicator for "below average," with the signals two to five used to signal increasing better-than-average realizations.

from the specific producer X with which the sender has experience, contingent on a generated sale, $\psi_P^S(m, Z, \rho) = T \cdot \mathbf{1}\{Z = X\}$.[8]

The first of these transfers is likely the most intuitive to economists, where the second two are motivated by marketing incentives in common use. Below we provide the broad intuition for the equilibrium effects in the three transfer schemes. After this, we provide more details on how our experimental environment matches and stress tests the model through our chosen parameterization. We then provide a quantitative discussion of the equilibrium predictions under the parameterization.

Under the *Receiver* transfer, by making an additional transfer of T to the sender for providing a message, in any informative equilibrium the sender will provide information for costs satisfying $c_S \leq \alpha \cdot \Delta w_R(\mathcal{I}) + T$. Cardinal alignment over information is maintained, so that were the message space rich enough the sender would still choose to fully reveal. The conditional transfer simply serves to increase the quantity of information provided while maintaining common interests over the final choice Z. However, there is a potential cost through equilibrium selection. This environment is one with a generic multiplicity of equilibria, always admitting babbling equilibria where messages are meaningless. One potentially negative effect of the *Receiver* transfer is that it makes coordination on informative equilibria riskier for receivers, as they must now pay an increased viewing cost of $c_R + {}^{T}\!/_{K}$. While the transfer allows for information transfer close to the first-best in equilibrium, it also runs the risk of coordinating outcomes on the worst-case equilibrium.

In the *Marketplace* transfer, the sender is provided with an incentive to generate a sale (a final choice of either A or B) rather than an outside option choice. This mirrors common referral schemes used by large marketplaces like Amazon that provide contingent payments to affiliates for generating *any* product sale. The effect is to reduce the alignment of interest between the sender and receiver. In particular, whenever T is large relative to $\alpha \cdot \Delta w_R(\mathcal{I}_x)$, equilibrium senders can only credibly provide *ordinal* information, equivalent to making one of the two statements *"Given my information, I expect A (B) to be better choice for you than B (A)."*[9] While this reduces the information quality, the transfer means that messages are provided more frequently, and without the receiver having to risk an explicit transfer.[10]

Finally, our *Producer* transfer is motivated by referral schemes that offer discounts to customers who refer friends and colleagues. While such incentives are relatively easy to implement, for large

[8]To make the transfer net zero, the receiver transfer in (R) is $\psi_R^R(m, Z, \rho) = -\frac{1}{K} T \cdot \mathbf{1}\{m \neq m_\emptyset, \rho = \text{View}\}$, reflecting a shared payment for the message by the K receivers viewing it. For (M) and (P) we set the receiver's component to $\psi_R(m, Z, \rho) = 0$, as the transfer to the sender is facilitated by other parties.

[9]This message behavior can be characterized by a strategy that only sends 5 ratings for whichever product is believed to be better. So $\mu\left(x, X, c^S\right) = 5_Y$ for $Y \neq X$ whenever $\mathbb{E}\left(\phi^X \,|x\right)$ is below average, and $\mu\left(x, X, c^S\right) = 5_Y$ for $Y = X$, whenever $\mathbb{E}\left(\phi^X \,|x\right)$ is above average. Note that senders in our environment are able to provide a message in support (or against) both A or B, regardless of the actual product they choose.

[10]This particular institution is capable of creating greater sales for the marketplace in equilibrium than the *Baseline* or *Receiver* incentives. However, this result critically depends on the precise distribution of the outside option, $H(\omega)$.

transfers the effect is to remove all alignment between the sender and receiver over the receiver's final choice. Best responding sender's will choose any message that increases the chance of a sale for product X. As such, the equilibrium prediction is that the value of information dips below the cost of viewing messages. The equilibrium prediction is the babbling outcome. While the *Producer* equilibrium is clearly inferior, such incentives can have very large short-run effects in generating sales if introduced covertly. While our theory model focuses on the long-run equilibrium effects, in our discussion section we come back to consider the potential short-run gains if conflicts of interest are not commonly known.

2.2. **Experimental Parameterization.**

Parameterization. The primary uncertainties in the model are the product-specific lotteries ϕ^A and ϕ^B, which in the experiment are called *Urn* A and *Urn* B. Each urn is composed of two balls, where each ball has an attached value θ to the receiver in $\Theta := \{\$0.10, \$0.20, \ldots, \$9.90, \$10.00\}$.[11] The set of product lotteries in our experiment is all possible equal lotteries over two values from Θ, so the lottery set is $\Phi = \left\{\frac{1}{2} \cdot \theta \oplus \frac{1}{2} \cdot \theta' \,|\exists \theta, \theta' \in \Theta\right\}$. The prior probability of each distinct lottery $\phi_r = \frac{1}{2} \cdot \theta \oplus \frac{1}{2} \cdot \theta'$ is given by:

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

The selection process is therefore biased toward the selection of homogeneous urns with two identical balls.[12]

Without any knowledge on the selected lottery ϕ_r for either product, the expected outcome from a draw is given by $\bar{\theta} = \$5.05$. However, knowledge of a single realization $x \in \Theta$ from the lottery provides a great deal of information. Knowing x, the conditional expectation for a new realization is $\nu(x) = {}^{3}\!/_{4} \cdot x + {}^{1}\!/_{4} \cdot \bar{\theta}$.[13]

Aside from the two product lottery draws, our experimental environment has two other random components: a send cost c^S and a reservation value ω. The send cost c^S is drawn over $\{-\$0.49, -\$0.48, \ldots, \$2.49, \$2.50\}$,[14] while the reservation value ω is a drawn from the set of

[11] In the experiment the balls are labeled with the integers 1 to 100, with values multiplied by \$0.10 when drawn by the receiver .

[12] We biased the process in this way to ensure that the sender's information would be valuable to the receiver while not being perfectly informative.The urns filling process is explained to subjects as follows: (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 $\theta_2 = \theta_1$.

[13] Half 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 expected value $\bar{\theta}$.

[14] We allow the send cost to have potentially negative values—senders who derive pleasure from communicating—as we wanted to have an interior level of provision in the *Baseline*.

possible ball values Θ. Both private information draws follow a linearly decreasing probability, so that lower values are more likely.[15]

The remaining scalar parameters are: a receiver cost of viewing a message of $c_R = \$0.05$; a sender-receiver alignment parameter of $\alpha = 1/10$ (senders derive one-tenth the value of the receiver's final choice as an explicit payment); a conditional transfer to the sender of $T = \$2.00$; and each receiver is modeled as being representative for $K = 4$ receivers. As a final experimental choice we structure the set of available messages to be coarse, mirroring a one-to-five rating scheme. The set of available non-empty messages in every round is therefore given by $\mathcal{M} = \{1_A, 2_A, 3_A, 4_A, 5_A\} \cup \{1_B, 2_B, 3_B, 4_B, 5_B\}$, a one-to-five rating for either urn (not just the one they sampled from).

Experimental Design. We use a between-subject experimental design across the four conditional transfer environments, where each session consists of 30 rounds. In each experimental round, every subject performs the roles of both a sender and a receiver.[16,17] In every round t the game timing is as follows: subject i as a sender 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.[18] As a sender the subject therefore has information $\mathcal{I}_{it}^S = (X_{it}, x_{it})$, which is pertinent information to a matched receiver subject j whose final outcome is defined by $\left\{\phi_{it}^A, \phi_{it}^B, \omega_{jt}\right\}$ in the second stage.[19]

Finally, our design augments the strategic choice environment to collect additional information to facilitate our counterfactual analyses. For each sender i, after realizing their signal x_{it} we ask them to select a provisional costly message $\hat{m}_{it} \in \mathcal{M}$. After selecting a message the subject is then prompted on whether or not they want to incur the send cost c_S. For the first 15 rounds this is simply a binary choice between the empty message or $\hat{m}_{it}$ after the drawn send cost c_{it}^S is revealed. For rounds 16–30 of the experiment we instead elicit the maximum send-cost $\overline{c}_{it}$ for which they would incur the cost.

[15] We do this via an order statistic. Two costs/reservations are drawn, where the lower of the two is given to subjects. The ex ante probability of send cost c_i^S is then approximately $\frac{2}{300}\frac{2.5-c}{3}$ (median of \$0.38, mean of \$0.51), while the reservation draw ω can be approximated by $\frac{2}{100}\frac{100-\omega}{100}$ (median of \$3.00, mean of \$3.38).

[16] Our 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.

[17] In the experiment *Urns 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 *C* and *D*.

[18] Though x_{it} is simply a signal draw from the point of information provision, we wanted experimental senders to experience a payoff from their draw, however we pay them only half the receivers value.

[19] In terms of matching subjects are given a random identity $i \in \{1, \ldots, n\}$ each round. Subject i acts as the sender to the subject $j = \mod(i, n) + 1$ and as the receiver to subject $k = \mod(i-2, n) + 1$ in the second stage—making their receiver choice over a different set of product lotteries $\left\{\phi_{kt}^A, \phi_{kt}^B, R\right\}$ than the ones they drew from.

For each receiver j, the first choice is always whether or not to view a message if one was sent, $\rho_{jt} \in \{\text{View}, \text{Not}\}$. After the viewing choice any information is revealed to the receiver,[20] and we ask them to make a provisional choice between the two urns $\hat{Z}_{jt} \in \{A, B\}$. For the first 15 rounds we then reveal the reservation ω_{jt} and the subject makes a final choice $Z_{jt} \in \left\{\hat{Z}_{jt}, R\right\}$. In rounds 16–30 we switch to a strategy method and elicit a reservation cutoff $\overline{\omega}_{jt} \in \Theta$ below which they would choose a draw from $\hat{Z}_{jt}$ and above which they would choose the realized reservation.

Data from interactions in rounds 1–15 are therefore given by $\left\langle \left(\phi^A_{it}, \phi^B_{it}, c^S_{it}, \omega_{jt}\right), \left(X_{it}, \hat{m}_{it}, m_{it}\right), \left(\rho_{jt}, \mathcal{I}_{jt}, \hat{Z}_{jt}, Z_{jt}\right)\right\rangle$; while data from rounds 16–30 are given by $\left\langle \left(\phi^A_{it}, \phi^B_{it}, c^S_{it}, \omega_{jt}\right), \left(X_{it}, \hat{m}_{it}, \overline{c}_{it}\right), \left(\rho_{jt}, \mathcal{I}_{jt}, \hat{Z}_{jt}, \overline{\omega}_{jt}\right)\right\rangle$.

2.3. **Predictions.** Our environment is a form of "cheap-talk," as although 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 whenever $c_S < 0$, and where receivers choose not view. The babbling outcome reflects the worst-case equilibrium for information transfer. However, three of our treatments allow for the possibility of more-informative equilibria. Table 1 provides quantitative predictions for the main economic variables under the most-informative PBE (miPBE, where we enforce a symmetric response to the urns in the absence of information) in each experimental treatment.[21]

Before outlining the qualitative predictions, we should note here the interpretation for the miPBE. While the equilibrium calculations are complex, the choices in our game are simple. For senders: given an easy-to-interpret signal, what information would you want to send, and at what cost to yourself would you actually send it? For receivers: will you view a message at a small fixed cost; given your information which of the two products do you prefer, and what is your valuation for it? While we would not expect subjects' behavior to start out at the miPBE through an introspective expected best-response calculation, we do see this equilibrium as one that can be reached in the long run. That is, we view the miPBE as a plausible upper bound on outcomes, given favorable initial play and enough time for subjects to learn.

For the predictions in Table 1, rather than focus on the precise miPBE constructions, we provide the measures that speak to the economics of the issue. The *Alignment* between the sender and receiver over what information would be revealed about the signal were the message free. *Rating Provision*, the fraction of rounds where a non-empty message is predicted to be sent. *Informational*

[20]The receiver's information set at this point is therefore given one of the following three options: the specific rating sent if viewed, $\mathcal{I}_{jt} = \{m_{it}\}$; that a rating was provided if not viewed, $\mathcal{I}_{jt} = \{m_{it} \in \mathcal{M}\}$; or that no rating was provided, $\mathcal{I}_{jt} = \{m_{it} = m_{\emptyset}\}$

[21]A comparable prediction table formulated under a risk-averse preference is provided in the appendix. The comparative statics across all variables in Table 1 continue to hold under moderate risk aversion.

TABLE 1. Most-Informative Risk-Neutral Equilibrium Predictions

	(B)aseline	(R)eceiver	(M)arketplace	(P)roducer	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, $\Upsilon^{\star}$	36.3%	98.1%	87.6%	0.0%	R≻M≻B≻P
Rec. Welfare, $\Upsilon^{\star}_R$	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 $\Upsilon^{\star}$	106.6%†	100.3%†	93.9%	0.0%	B≻R≻M≻P
Rec. Welfare $\Upsilon^{\star}_R$	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.

Efficiency, a measure of how much information is transferred via the receiver's gross expected payoff, and *Receiver Welfare*, a net measure for receivers that additionally accounts for their viewing costs. Finally, because one of our main channels for inefficiency is through the underprovision of information, we also provide the informational efficiency and receiver welfare measures conditional on a message being sent, alongside the receiver's predicted *View Rate* in the best-case equilibrium.

The rating provision row indicates the rate at which senders provide informative messages in the miPBE, $\Pr\left\{\mu^{\star}\left(X,x,c^S\right)\in\mathcal{M}\right\}$. In the *Baseline*, information is predicted to be sent approximately a third of the time (up to costs of approximately \$0.10 to the sender for extreme signals), where the *Receiver* and *Marketplace* transfers increase this to more than 90 percent (incurring send costs in excess of \$2). In contrast, because the *Producer* treatment has no viewing in the babbling outcome, the rate drops to 30.4 percent (babbling messages are sent only when the cost is negative). In terms of the quantity of information provided, our qualitative comparative static prediction for the four treatments is therefore: R≻M≻B≻P. However, given the quantitative levels, it is clear that the main effect here is the separation between *Receiver* and *Marketplace* on one side with high provision, and *Baseline* and *Producer* on the other.

The extent to which decision-relevant information is exchanged is measured by the informational efficiency, $\Upsilon^{\star}$.[22] This efficiency measure examines the expected receiver outcome $\mathbb{E}_{\omega,\mathcal{I}}\left[w_R\left(\zeta^{\star}\left(\mathcal{I},\omega\right)\right)\right]$

[22]For 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.

under the miPBE, relative to the expected outcome under no communication $\underline{W}$ (the babbling payoff), and rescaled to an efficiency through $\overline{W}-\underline{W}$, the upper bound difference between the expected choice outcome under full revelation and babbling.[23] The qualitative prediction for the best-case equilibria is therefore the same as the provision. However, the quantitative gap between treatments is now more pronounced. Where *Receiver* leads to possible informational efficiency of 98 percent, this figure falls to 88 percent for *Marketplace*, 34 percent for the *Baseline*, and zero for *Producer*.[24]

However, gains in informational efficiency do not come for free. While for the *Marketplace* and *Producer* treatments the transfer comes from a party outside the model—where we examine the gains to these parties in our discussion section—in the *Receiver* treatment the transfer is facilitated by the receiver player. Our receiver welfare measure $\Upsilon^{\star}_{R}$ modifies the informational efficiency to account for the receiver's viewing costs (again rescaling via $\overline{W} - \underline{W}$). 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 just over a third of the time, the presence of the cost results in a 2.2 percentage point efficiency reduction, leading to a receiver welfare of 34.1 percent. While the viewing cost is the same in *Marketplace*, informative messages are viewed over nine-tenths of the time, so the net effect is a larger decrease to 81.7 percent receiver welfare.

The largest reduction relative to the gross informational efficiency is predicted for the *Receiver* treatment. There are two reasons for this larger drop: First, ratings are predicted to be provided at high rates so viewing costs are incurred more often. Second, and the quantitatively larger force, much of the gain from increased information transfer is distributed to the sender through the transfer. In fact, by parameterizing the receiver's viewing cost at \$0.55, we purposefully chose to stress test the *Receiver* environment. Despite the large increase in total efficiency, receivers' final outcomes are predicted to be inferior to those in the most-informative *Baseline* prediction. As mentioned earlier, given the presence of multiple equilibria in the *Baseline* and *Receiver* institutions, one effect of the reduced receiver welfare could be to make coordination on the informative equilibria riskier, as receivers are the ones footing the bill.[25] In contrast to our efficiency predictions, the receivers are predicted to do best under the *Marketplace* incentive. While quantitatively

[23]So the efficiency measure is defined as $\Upsilon^{\star} := \Delta w_R\left(\mu^{\star}\left(X, x, c_S\right)\right)/\Delta w_R\left((X,x)\right)$, the difference in receiver's outcomes under the equilibrium, relative to the difference under full revelation .

[24]For both the *Baseline* and *Receiver* treatments, the intuitive 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. For Marketplace, a simple equilibrium strategy is to send the message 5_X when $x > 50$ and the message $5_{X'}$ for $X' \neq X$ (sending a five rating for the unsampled urn) when $x \leq 50$.

[25]Assuming 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 have a belief that senders coordinate on the miPBE in excess of 0.16. In contrast, for the *Receiver* treatment, the additional cost of viewing means receivers must have a belief in excess of 0.69. As such the *Receiver* transfer causes the miPBE to become risk dominated by the babbling outcome.

similar for receivers, *Baseline* is superior to *Receiver*, while the Producer treatment is predicted to have the worst outcomes for receiver players.

The last three rows indicate the predictions conditional on information provision, where the figures show the similar best-case information transfer in *Baseline* and *Receiver*, modulo the sender's provision choice. While each non-empty message in the *Marketplace* treatment is predicted to have less information content, the quantitative predictions are not too different. Despite only ordinal alignment over provided information, *Marketplace* messages convey information with 94 percent efficiency. The reason for this is partly a function of our parameterization, and partly structural. Equilibrium *Marketplace* messages are predicted to be binary, taking the form of a 5-rating for one of the two urns. Though less useful than information about the precise signal x, messages under *Marketplace* do change the receiver's expected quality for the rated urn. Where uninformed receivers would choose the outside option for all $\omega \geq \bar{\theta}$, those in possession of an equilibrium *Marketplace* rating choose the outside option for $\omega \geq \overline{\omega}^{\star}_{M} = \$5.99$.[26]

Though binary in nature, the *Marketplace* information fully reveals sender's information about which of the two products is relatively better. For receivers with reservations below 50.5, the message fully reveals all decision-relevant information to the receiver. Inefficiencies for the information revealed under the *Marketplace* scheme only come about for receivers with reservations above 50.5, where the decision is between the revealed-better product and the reservation. As such, the informational efficiency of *Marketplace* depends critically on the cumulative distribution for reservation values, $H(\omega)$. In particular, were $H(\bar{\theta}) = 1$ (reservations always less than the unconditional product quality) the Marketplace transfer would be informationally fully efficient. In contrast, were $1 - H(\overline{\omega}^{\star}_{M}) = 1$ (reservations always greater than $\overline{\omega}^{\star}_{M}$, the implied quality of the rated urn) the informational efficiency of a *Marketplace* message would be zero. Because our parameterization sets $H(\bar{\theta}) = 0.75$ (and $1 - H(\overline{\omega}^{\star}_{M}) = 0.17$), the quality/quantity tradeoffs inherent in the Marketplace incentive are skewed towards a good outcome. In our discussion, we come back to this point, where we use the strategy methods in our data collection to examine a family of counterfactual reservation distributions.

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

[26]The exact choice stems from the expected quality of urn Y given the message 5_Y being ${}^{5}/{}_{8}\bar{\theta} + {}^{3}/{}_{8}\mathbb{E}\left(x \,|x > \$5.05\right) = \$5.99$. Note that unlike the other treatments, the 5_Y rating is also informative about the other the unrated urn $Y' \neq Y$, which has an expected quality of \$4.11.

predictions. We then analyze subject behavior in more detail, and examine the subject decisions that underpin our results.

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

Result 1 (Efficiency)**.**

(i) *Total information transfer is substantially lower than predicted in all treatments, though observed behavior matches the best-case equilibrium comparative statics.*

(ii) *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;[27] 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.

Efficiency of information transfer. Ordinally, 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

[27] All 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	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≻M≻B≻P	$M\overset{\star}{\succ}R\overset{\star\star\star}{\succ}P\overset{\star\star\star}{\succ}B$
Info. Efficiency	9.0%	3.9%		28.1%	4.1%		17.7%	3.8%		1.0%	3.9%		R≻M≻B≻P	$R\overset{\star\star}{\succ}M\overset{\star}{\succ}B\overset{\star}{\succ}P$
Rec. Welfare	7.3%	3.8%		-2.9%	4.1%		14.1%	3.8%		-1.1%	3.9%		M≻B≻R≻P	$M\sim B\overset{\star}{\succ}P\sim R$
	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∼R∼M≻P	$M\overset{\star\star\star}{\succ}R\sim B\overset{\star\star}{\succ}P$
	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≻R≻M≻P	$R\overset{\star\star\star}{\succ}B\sim M\overset{\star}{\succ}P$
Rec. Welfare	52.1%	7.6%		13.4%	7.9%		51.3%	6.7%		38.8%	7.9%		B≻M≻R≻P	$B\sim M\overset{\star}{\succ}P\overset{\star\star\star}{\succ}R$

Note: Tests are derived from Wald tests after regression (probit) results for Efficiency/Welfare (Provision/Viewing/Sales decisions) on treatment dummies. The relationship $A\overset{\star\star\star}{\succ}B$ ($\overset{\star\star}{\succ}$, $\overset{\star}{\succ}$) represent rejection of mean equality for treatments A and B with the one sided alternative A>B at 99 percent confidence (95, 90 percent confidence, respectively). The relationship A∼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∼V.

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.

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)**.**

(i) *Senders' choices respond on the extensive and intensive margins as predicted by theory, though with a smaller response than predicted on the intensive margin.*

(ii) *Senders with very positive information are more likely to provide a rating in all four settings.*

(iii) *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*, to 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

TABLE 3. Decision to send a rating.

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

Note: Results for *Sent* specifications are recovered using a random-effects probit model on the binary dependent variable $1\left\{m_{it} \neq m_{\emptyset}\right\}_{it}$,whether a rating was sent; results for *Cutoff* specifications from a random-effect least squares model with dependent variable $\overline{c}_{it}$, 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.

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 ($x_{it} \leq \$3.30$, $\$3.30 < x_{it} < \$6.70$ and $\$6.70 \leq x_{it}$, respectively), a dummy for a non-positive rating cost ($c^S_{it} \leq 0$), and the actual cost of rating when it is positive (c^S_{it}). 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.[28]

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

[28]In 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 5).

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.

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 ($x \leq \$5.00$); ii) *Understatements*, a negative rating (a one or two) on urn X for a high-quality signal ($x > \$5.00$); 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 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.[29]

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.[30] 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. The first row in Table 5 indicates the maximum efficiency possible given sender behavior.[31] 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.[32]

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)**.**

[29]As 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."

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

[31]The expected efficiency given sender's behavior for a best-responding receiver is calculated as $\hat{\Upsilon}_S = \frac{1}{\overline{W}-\underline{W}}\left[\left(\sum_{m\in\mathcal{M}\cup\{m_\emptyset\}}\hat{\Pr}\{m\}\cdot\int\max\left\{\hat{\nu}_A(m),\hat{\nu}_B(m),\omega\right\}dH(\omega)\right)-\underline{W}\right]$, 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.

[32]The 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.

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.

(i) *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*.*

(ii) *Viewed ratings have substantial effects on the choice between products in all four treatments.*

(iii) *In all treatments except* Producer *a positive rating increases the likelihood of a sale.*

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 of \$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 will 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

TABLE 6. Receiver behavior.

(A) Probability of selecting rated product

	Baseline	Receiver	Marketplace	Producer
m= 1 or 2	0.059***	0.014***	0.083***	0.079***
	(0.019)	(0.008)	(0.018)	(0.025)
m= 3	0.413	0.362***	0.378***	0.338***
	(0.057)	(0.049)	(0.041)	(0.059)
m= 4 or 5	0.973***	0.981***	0.968***	0.981***
	(0.014)	(0.009)	(0.009)	(0.08)
Not viewed	0.504	0.492	0.491	0.493
	(0.016)	(0.023)	(0.022)	(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	\$4.97	0.745	\$4.99	0.764	\$5.37	0.801	\$5.30
	(0.036)	(0.20)	(0.033)	(0.19)	(0.029)	(0.20)	(0.039)	(0.25)
m= 3	0.691	\$5.18	0.808	\$5.24	0.838***	\$5.05	0.724	\$4.68
	(0.055)	(0.23)	(0.040)	(0.22)	(0.032)	(0.25)	(0.057)	(0.31)
m= 4 or 5	0.856***	\$5.92***	0.895***	\$6.13***	0.805**	\$5.83***	0.802**	\$5.20
	(0.030)	(1.99)	(0.021)	(0.19)	(0.022)	(0.20)	(0.026)	(0.23)
Not viewed	0.646**	\$4.96	0.762	\$4.89	0.656*	\$5.10	0.696	\$4.76
	(0.043)	(0.21)	(0.034)	(0.21)	(0.055)	(0.35)	(0.036)	(0.25)
No rating	0.762	\$5.12	0.721	\$4.91	0.773	\$5.38*	0.765	\$4.95
	(0.016)	(0.15)	(0.022)	(0.17)	(0.021)	(0.20)	(0.019)	(0.21)

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 $\mathbb{E}(\bar{\omega}\,|\mathcal{I})$, 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 \$5.05 for the reservation cutoff) at: ***-99 percent confidence; **-95 percent confidence; and *-90 percent confidence.

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 $\frac{1}{2}$.[33]

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.

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 $\bar{\theta} = \$5.05$ 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. Moreover, the reservation in response to a positive rating in *Producer* is not significantly different from the response to a negative rating in the other three treatments (though the large majority of receivers there choose the unrated urn).[34] It is as if a positive rating in Producer is entirely disregarded.

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, *Marketplace* is the most-attractive long-run incentive. We summarize our analysis here, and provide supporting evidence below:

[33]However, 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

[34]That the reservations for unrated urns when selected are so close to \$5.05 indicates that subjects are close to risk neutral.

Result 4. *Leveraging our experimental design, we examine the effects of an important parameter in our experiment: the distribution of outside options $H(\omega)$. Our analysis indicates:*

(i) *The* Marketplace *incentive maximizes sales for both the marketplace and producer third parties for a large range of high-expectation distributions.*

(ii) *Relative to the* Baseline *and* Receiver *setting, the* Marketplace *incentive is more stable to the covert introduction of* Producer *incentives.*

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 strategy choices (cutoffs) in the last half of our experiments, we can use this data to extrapolate our results to counterfactual parameterizations.

As mentioned in the design section, one of the key parametric choices in our experiment is the distribution of outside options, $H(\omega)$. The realization of the reservation value has a clear effect on which margin information is most important on: trying to choose which of the two products is better ($\omega < \bar{\theta}$), or trying to decide whether to give up the reservation for the better product ($\omega \geq \overline{\theta}$). Theoretically, because of the reduced sender alignment, the *Marketplace* incentive is only capable of convincing receivers to forgo the reservation up to a limit $\overline{\omega}^{\star}_{M}$, where all receivers with reservations below $\overline{\theta}$ will purchase a product regardless of the sender's message. The predicted gains from the *Marketplace* incentive are therefore particularly sensitive to $H\left(\overline{\omega}^{\star}_{M}\right) - H\left(\overline{\theta}\right)$, the mass of receivers capable of changing their minds. We now show that the receiver WTP data suggests the *Marketplace* incentives are beneficial to both the marketplace and producer third parties for a parametric family of reservation distributions.

We parameterize a family of variance-preserving beta distributions $\{H_{\sigma}(\omega)\}$ that vary the expected reservation σ, intersecting with our experimental parameterization at $\sigma = \$3.33$. Using the WTP data and estimating behavioral elasticities for the sending and viewing decisions, we calculate the informational efficiency, total sales and same-product sales under the counterfactual distribution $H_{\sigma}(\omega)$, which we graph as the difference from the *Baseline* amounts in the three plots for Figure 1.[35]

Figure 1(A) makes clear that, regardless of the reservation distribution, the *Receiver* incentive provides the highest quality information for the receivers across our three transfer environments. However, because the costs of consuming information in our *Receiver* parameterization are large

[35]Full details of the counterfactual model are provided in Appendix B. The main features of the model are that we hold constant the observed message distribution (conditional on the signal x) and the receiver's willingness to pay (conditional on the observed message). We then make use of experimental variation in the environments to estimate subjects' provision and viewing elasticities in a logit model, modeling how the sending and viewing decisions changes under the expected payoff for the counterfactual $H_{\sigma}(\omega)$.

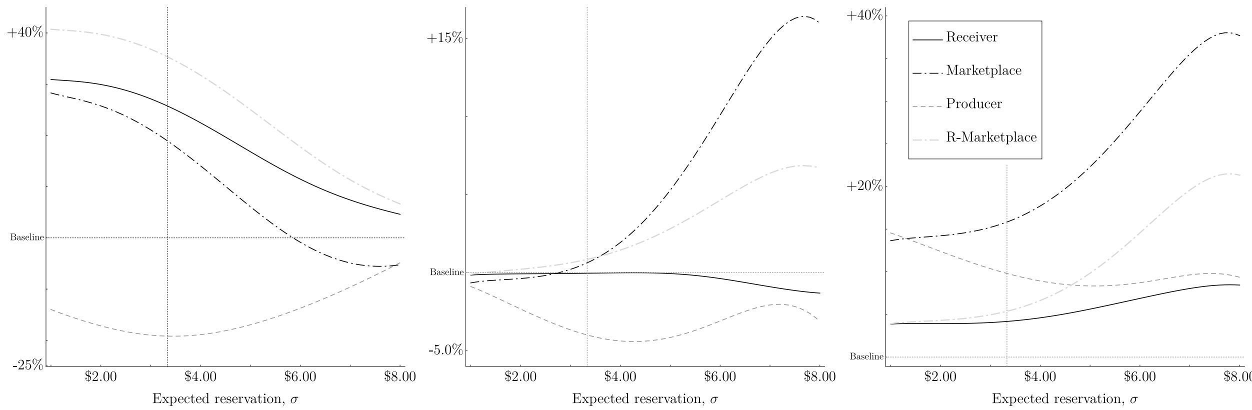

(A) Inf. Efficiency, Υ (B) Total sales (C) Same-product sales

FIGURE 1. Change from *Baseline* under counterfactual distribution $H_{\sigma}(\omega)$

relative to the benefits, many receivers choose not to view messages. As a counterfactual transfer we also consider the effects if the marketplace player were to provide the Receiver transfer to the sender. This transfer, which we call *R-Marketplace*, leads to higher viewing behavior, which further increases the behavioral informational efficiency over *Receiver*.

In Figures 1(B) and 1(C) we illustrate, respectively, how the three transfers and the *R-Marketplace* counterfactual affect the interests of the marketplace and producer third parties. Figure 1(B), indicates the proportional change in total product sales over the *Baseline* at each value of σ. Inspecting the change in sales, the figure shows that under reservation distributions with moderate to high expected draws, the *Marketplace* transfer generates more sales than the other options. This shows up as a proportional change in total sales of up to 17 percent. In contrast, the *Receiver* and *Producer* transfers fail to grow *Baseline* sales for any value of σ, where *Producer* reduces sales by close to 5 percent for middling values of σ. The *R-Marketplace* alternative where the third-party pays for an aligned *Receiver*-like transfer is the dominant choice for the marketplace third party when σ is low. However, the proportional gain in sales over *Baseline* never exceeds 1 percent in this region.

Figure 1(C) provides a proxy for the preferences of the producer X, who has access to a group of prior customers as the sender. Here the figure's vertical axis indicates the proportionate change in same-product sales over the *Baseline*. Similar to the any-product sales figure, the plot illustrates that the *Marketplace* transfer is superior for generating same-product sales. Though the *Producer* and *Marketplace* transfers have similar effects when expected reservations are very low, as σ increases the *Marketplace* transfer clearly dominates. The reasoning for this follows the receiver reactions in Table 6. While the *Producer* incentive is highly effective at changing the consumer's provisionally preferred product, it is bad at increasing willingness to pay. As such, the incentive is most effective when consumers will buy one of the two products regardless of the provided information, and least effective when the receiver needs to be convinced of the product's expected quality. In contrast, the *Marketplace* incentive does a good job of balancing the two forces.

In the long run—when receivers are fully aware of the incentives provided to senders—our counterfactual analysis suggests that the *Marketplace* incentive is the superior choice for both third parties across a wide range of reservation distributions. In particular, individual producers create more same-product sales under the *Marketplace* incentive than they do under a *Producer* incentive targeted. In Figures 2(A) and 2(B) we illustrate the long-run proportional change in total and same-product sales, respectively, in a shift to the *Producer* incentive from each of the three alternatives. The (A) subfigure shows that in the long-run, the marketplace loses in the shift to *Producer*, where the loss is particularly acute when moving from *Marketplace* when receivers have high expected reservations. In contrast, Figure 2(B) indicates the long-run interests for the producer player, where a move from both aligned environments to Producer generates sales gains for *X* across the entire range of reservation distributions. However, the move from *Marketplace* creates a long-run sales loss (decreasing in σ) for the producer player. The implication then is that sufficiently farsighted producers would not want to move from *Marketplace* to *Producer* due to long-run sales loss. However, producer players with more short-run motives might still create benefits for themselves by *covertly* introducing *Producer*-like incentives to senders.

Our final counterfactual examines the extent to which a producer player can benefit in the short run by offering the producer incentive to the sender while receivers make choices believing the senders to have more-aligned incentives.[36] To make this calculation, we can combine the provision and ratings from the *Producer* treatment with the viewing and final choice behavior from each of the other three environments.[37] The short-run proportional sales effect is then calculated relative to sales generated in the prior environment with common knowledge.

Figures 2(C) and 2(D) illustrates the short-run total and same-product sales gains from a shift to *Producer*, again looking across the reservation distribution $H_\sigma(\omega)$. While the marketplace player does achieve a small short-term gain from a move from *Baseline* to *Producer* (through a combination of increased provision and more favorable messages), the short term effects from *Receiver* and *Marketplace* are flat or negative. In contrast, the covert introduction of a *Producer* incentive to senders leads to substantial sales increases for the producer third party when receivers believe they

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

[37] One 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 this asymmetry they would not change their messaging behavior.

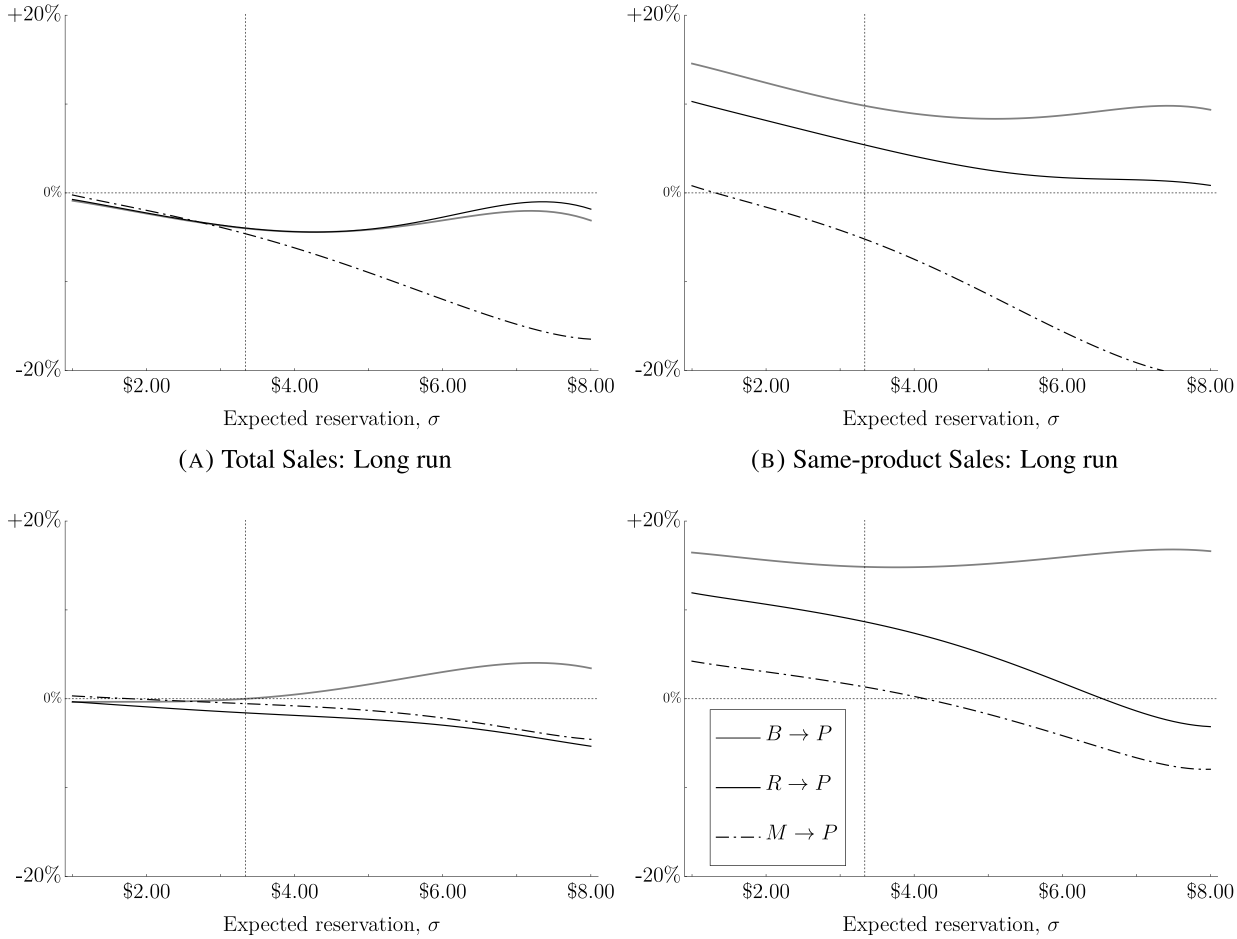

(A) Total Sales: Long run

(B) Same-product Sales: Long run

(C) Total Sales: Short Run

(D) Same-product Sales: Short Run

FIGURE 2. Sales changes from a shift to *Producer* incentives

are in the *Baseline,* increasing sales by up to 17 percent. For all but the highest expected reservations the move from *Receiver* to *Producer* also produces large short-run increases in same-product sales, with a 10 percent increase at the experimental parameterization.

Contrasted to the above, the introduction of *Producer*-type incentives into the *Marketplace* environment generates much smaller benefits to the producer player, and for a smaller range of distributions. Receiver players are already more suspicious of positive ratings, and provision is already high. Though there are small short-run gain to the producer at the experimental parameterization, it is small (a 0.5 percent gain), and becomes negative for reservation distributions with higher expectations. Combined with the results from the long-run comparison in Figure 2(C) we can conclude that *Marketplace* is the most stable informative incentive structure. Across a large range of reservation distributions, the move from *Marketplace* to *Producer* causes both a short-and a long-run loss for the producer third party. Moreover, for almost the entire range the negative long-run effect is quantitatively stronger.

In contrast to *Marketplace*, moving to *Producer* from the *Baseline* and *Receiver* environments offers larger short-run gains set against smaller long-run losses as receivers become aware. Both aligned-interest incentive environments are therefore more unstable when faced with the introduction of distorting incentives. Moreover, while advocates for consumer interests might have strong motives to police the environment to stop covert incentives from being introduced, few channels for enforcement outside regulation exist. Our results provide some optimism, however, showing that platforms over which trade happens have strong incentives to regulate and administer information markets. The marketplace player has a clear long-run interest in preventing a move to *Producer*—checking Figure 2(A), this is particularly true when the status quo is their most-preferred *Marketplace* incentive. And while the incentive is there to monitor and police ratings, so too is a strong policy lever: exclusion of those sellers manipulating ratings from the platform. [38]

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. While individuals with relevant information can be incentivized to provide it to others, the party offering those incentives can heavily influence the total benefits from information sharing, as well as how those benefits are distributed. We examine a setting where both the extensive (quantity) and intensive (quality) margins of information can vary, and assess the extent to which three qualitatively different incentive structures improve outcomes. In particular, we examine an experimental parameterization where all interested parties can theoretically benefit from an incentive that increases quantity at a slight cost to quality. Though there are some differences with the theoretical predictions in levels, the qualitative behavior in our experiments validates the idea that quality and quantity can be beneficially traded off in an information context.

Our *Receiver* environment maintains full alignment-of-interest between senders and receivers, with consumers who provide information compensated by those making use of the information.

[38] Examining the terms of service for sellers on Amazon—available from sellercentral.amazon.com/gp/help/external/G1801—the sanction threats are for “cancellation of listing, suspension or forfeiture of payments, and removal of selling priveleges.” In terms of restrictions on seller behavior, while sellers are allowed to “request feedback and reviews from your own customers in a neutral manner” they are prohibited from offering any incentive (including coupons) in exchange for reviews. Moreover, because sellers are barred from communicating with customers outside of the Amazon buyer-seller messaging tool, there is a clear channel for Amazon to monitor.

This leads to a large increase in the quantity of information without a reduction in quality. However, though information can be increased with this transfer, receivers' total welfare can be decreased—which is the case in our parameterization—as the payment can defray much of the benefits. Instead, receivers may be better off by allowing another interested party to provide the transfer that increases quantity. In our *Marketplace* incentive, the transfer reduces the alignment-of-interests between senders and receivers by rewarding senders that generate sales. This type of sales-contingent incentive has a corrosive effect on the quality of provided information (both theoretically and in our data) by creating a motive for the provision of sales-favorable information. Importantly, this incentive still maintains some alignment between informed and uninformed consumers over an ordinal feature of choice: which of the competing products is better. In fact, we show that at our experimental parameterization (and beyond) there are cases where consumers, marketplaces, and producers would all be better off by trading between quality and quantity in this manner.

The last of our incentives illustrates that trading off all alignment-of-interest can and does lead to sub-optimal results. Theory predicts that when senders are provided with a *Producer*-driven incentive to generate product-specific sales no information can be transmitted. While our experimental data echoes other work in suggesting a less extreme outcome than babbling, it is nonetheless the least efficient of our environments.

Although our experimental results favor the *Marketplace* incentive, commission-style incentives are not necessarily superior to the other incentives outside of our parameterization. In particular, were the costs of the incentive smaller—divided across a larger body of representative receivers, or subsidized by another party—maintaining full alignment under a *Receiver*-like incentive would be more effective. That is, while our theory and experiments show the *possibility* of trading off the quality of provided information for a greater chance of provision, such a trade-off will not always be productive for final consumers, where variants of *Receiver* might plausibly achieve first-best outcomes. However, though our results will obviously be sensitive to choices in our experimental design, through features of our data collection choices and the structure provided by the theory, some extrapolation to other environments is possible. In our discussion, we focus on the aspect of our design to which our results are most sensitive: the distribution of reservations. The likely range of outside options has a strong effect on the strength of information needed by decision makers. If reservations are very low, provided information can be coarse (which product is better); if reservations are high information needs to be more precise (how much better than average is this product). Extrapolating our behavioral data across this parameter shows that *Marketplace* incentives generate superior outcomes for the parties with the clearest channels to facilitate the sharing of experiences from prior customers: the platforms where products are sold, and the specific manufacturers who make the products. In the long run, where consumers are likely to need extensive information to make a purchase, we show that the *Marketplace* incentive remains the superior incentive for all parties.

As a final counterfactual, and leveraging the same extension over a family of reservation distributions, we use our experimental data to point out another less-obvious benefit for our sales-contingent *Marketplace* transfer: a little bit of known misalignment helps inoculate the environment from the introduction of fully misaligned incentives that poison the well in the long run. That is, if consumers are aware that the parties providing information are somewhat biased, the incentive for firms to secretly introduce *Producer*-like side payments is greatly reduced. By combing messages from senders under the *Producer* incentive with receiver behavior from the other incentives, we show that producers are least likely to benefit—in both the long and short run for a range of values—when the receivers already believe senders are a little biased. The implication here is that the social benefits from an acknowledge partial conflict of interest is greater than it first appears. Not only does the skepticism that comes from a partial conflict of interest help decision makers parse the information provided to them, the presence of commonly known misalignment also helps to stop the wedge from being driven further.

While our results provide an important example of what could be, more study is clearly required. In particular, it would be useful to more directly identify how acknowledging potential conflicts of interest affects decision makers. How can we most effectively generate skepticism without completely poisoning the well? In the age of "fake news," there are clear policy questions to be answered about how effective declarations of conflicts of interest are in the presence of uncertainty–both self-declarations by those providing information, and from other motivated parties. Beyond this, further research should examine the effects from variation excluded from our analysis. In particular, it would be useful to study the sensitivity of both provision and acquisition of information to costs, aggregation across multiple sources of advice, and reputation effects from repeated interaction. While there are many questions left to answer, our results function as a clear demonstration of the potential for efficiency gains from distorting senders' incentives. Substantial benefit can be generated by giving up some informational quality in exchange for an increase in the frequency with which information is provided.

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.

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.

______ **and** ______, "Information Transmission under the shadow of the future: An experiment," April 2019. University of Pittsburgh working paper.

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.