



Working Papers

www.cesifo.org/wp

Quality vs. Quantity in Information Transmission: Theory and Experimental Evidence

Jonathan Lafky
Alistair J. Wilson

CESIFO WORKING PAPER NO. 5426
CATEGORY 13: BEHAVIOURAL ECONOMICS
JUNE 2015

An electronic version of the paper may be downloaded

- *from the SSRN website:* www.SSRN.com
- *from the RePEc website:* www.RePEc.org
- *from the CESifo website:* www.CESifo-group.org/wp

ISSN 2364-1428

Quality vs. Quantity in Information Transmission: Theory and Experimental Evidence

Abstract

Information sharing has become increasingly important in helping consumers make better, more informed choices over competing products. Our project uses a novel theoretical framework and laboratory experiments to analyze three simple, commonly used incentive schemes against an unincentivized baseline. Each incentive scheme has qualitatively different theoretical predictions for behavior and efficiency, while our laboratory experiments examine the degree to which these differences manifest themselves, and the best-cast theory's robustness to human behavior. Our findings indicate the possibility for substantial efficiency gains by introducing incentives that reward information sharing, even where those incentives drive a wedge between those sending and those receiving information. In particular, our results point to a misaligned incentive commonly found in the field, sales commissions, as being a robust institution to stimulate the exchange of information.

JEL-Code: C920, D020, D830.

Keywords: rating incentives, institutional stability, social information.

Jonathan Lafky
Department of Economics
Lafayette College
110 Simon Center
USA - Easton, PA 18042
lafkyj@lafayette.edu

Alistair J. Wilson
Department of Economics
University of Pittsburgh
4907 Wesley W. Posvar Hall
230 South Bouquet Street
USA - 15260 Pittsburgh PA
alistair@pitt.edu

June 2015

Our thanks to the following: John Duffy, Emanuel Vespa, Lise Vesterlund, Stephanie Wang, and audiences at the ESA, SEA, George Mason, and Lafayette. Any mistakes within are obviously not attributable to anyone but ourselves.

1. INTRODUCTION

Information is profitably shared by friends, colleagues, even strangers. A friend may detail their positive experiences with a new car. A colleague gives their critique on a movie. A stranger on the street might point us to a coffee shop. Placing faith in the opinions of others allows us to use their experiences to make better choices. This can lead us to pay more for a recommended car than a competing model, choose a different movie, or walk slightly farther for a better espresso. Moreover, many institutions help us look beyond our immediate social networks to seek out advice from further afield. Consumers have relied on newspapers and periodicals to review new products for over a century: from books and plays, to consumer electronics and movies. More introspectively, academic journal editors have long relied upon disinterested peer reviews to advise them on publication decisions.

Our paper examines institutions for the exchange of information between peers. By varying the structure of a conditional payment to the sender of information (a consumer with product experience) we examine tradeoffs between quality and quantity for this information transfer. For quantity we consider the likelihood that the informed party provides any advice at all, while for quality we examine the content of advice and the extent to which those receiving it (consumers without product experience) make use of it.¹ We first construct a *Baseline* environment where the experienced party faces a cost of providing information, but otherwise has interests that are fully aligned with the uninformed consumers. Characterizing the best-case equilibrium outcome, we show the main failure in this environment is under-provision. Though high-quality advice is sent when the cost of communicating is small, low rates of provision lead to an inefficient outcome. To this baseline we compare three distinct alternatives, each with their own advantages and disadvantages.

The first environment, that we call *Receiver*, is the most obvious to economists: a monetary transfer from the consumer making use of the information to the party providing it. Maintaining the alignment of interest, this institution provides a transfer conditional on both the provision and acquisition of information. Examples of this institution abound: from printed reviews in newspapers, to paying referees, through to paid subscriptions for access to rating aggregates (for example, Zagat's or Angie's List). Though these payments can help increase provision while maintaining quality, there are some potential issues. The first is distributional. As the payment becomes larger, more of the gains from information transfer are shifted to the party providing information. In fact, for large enough payments, final consumers can fare worse than the baseline. The second issue is selection. Cheap-talk like environments have inherent multiplicity, and babbling equilibria always exists where no information is shared. As the size of the incentive payment increases, coordinating on efficient equilibria becomes more risky for those acquiring information.²

Our next two institutions look to parties other than the final consumer to subsidize information transfer. The second institution, labeled *Vendor*, implements payments from the marketplaces where various competing products are sold. Vendors want to increase their total sales volume,

¹In some sense quantity and quality are interchangeable, as we could simply measure the total amount of information transmitted. Practically speaking, however, there are important distinctions between the two measures, especially from the perspective of designers of ratings systems (where accuracy will typically be unobservable).

²For an example of such selection in the field, take Google Answers, a website where those seeking information directly compensated those providing it. The website was eventually closed after the community using it became too small.

though they are not concerned with exactly which of their products is sold. We show that marketplaces can benefit by offering incentives to their previous customers to encourage others to buy any product. This incentive provides a payment for each sale generated by the informed party, trading off some alignment of interest between those sharing information in exchange for greater provision. Sales-contingent incentives of this form are also prevalent. For example: salespeople at department stores providing advice over competing products or brokers offering advice over different investments (explicit sales commissions); servers providing advice on the dessert menu (implicit through a larger tip); and product-review websites that link to the “best” products in a category.³ This form of incentive distorts the preferences of those sharing information relative to those receiving information, creating a misalignment over willingness to buy (WTB). When there is full alignment, information providers and the final consumer agree not only over which of the competing products are better, but also over whether a purchase should be made at all. By contrast, an informed party with a sales-based incentive benefits from each sale, so their advice on whether or not to purchase can be biased. They can, however, still be honest in describing which of the available products is best. For example, a server’s advice on whether or not to order dessert might be disregarded, but their comparative advice between the desserts on offer can be informative.

Though sales incentives of this form drive a partial wedge between adviser and advisee, there are distinct benefits. Providing sales incentives can lead to more frequent provision of information. For online referrals, this can show up as reviews being written for a broader range of product categories, while salespeople on commission may be more inclined to take the time to offer their opinion on the products they are selling.⁴ Those utilizing advice can also be made better off. The conditional payment driving increased provision is no longer paid by final consumers (directly at least). Depending on their WTB, trading off some quality for increased quantity can improve expected outcomes, relative to both the *Baseline* and *Receiver* settings. Importantly, marketplaces can benefit from offering sales incentive in equilibrium: greater information provision will increase total sales, while the misalignment in preferences skews provision towards positively framed advice. For example, rather than reporting “Product A is bad,” which only influences the consumer’s product selection, salespeople might report “Product B is better,” which affects both the product selection and the WTB.

Finally, we examine an incentive offered by particular product manufacturers (labeled *Producer*), conditioned on product-specific sales. Such incentives create full misalignment between the parties exchanging information, so that neither information on when to buy, nor information over which particular product is better can be reliably exchanged. Whenever the size of the incentive payment dominates any altruistic motives, the equilibrium selected is pinned down, with only babbling outcomes possible. Though the long-run predictions with such incentives are unambiguous and highly inefficient, short-run gains may lead to their selection. Indeed, this form of incentive has long been used in the form of referral bonuses, such as insurance agents rewarding current customers for each friend brought in as a new customer. These incentives have become even more common with the rise in social marketing, with companies asking consumers to ‘like’ their products on social media to obtain product discounts.

³Amazon.com, the largest online retailer in the United States with yearly revenues of \$75 billion in 2013, provides payments to affiliate websites generating *any* sales through its *Amazon Associates* program. Referral payments (typically 4 percent of the sale, though this varies by category) are issued when consumers buy products after following links from an affiliate’s website to Amazon.

⁴An alternative channel would be through greater information acquisition by those providing advice. However, endogenizing the sender’s information will not be explicitly addressed here.

In terms of theoretical contributions, our paper constructs a novel framework in which to examine the tradeoff between quality and quantity. Allowing for a rich heterogeneity over the state and signal spaces, as well as final consumer WTB, we introduce our four environments—the baseline and the three conditional transfers. We characterize the best- and worst-case equilibria in each, which have a simple, intuitive structure. We then take the environments to the laboratory under a suitable parametrization, to examine equilibrium selection and the extent to which behavior adheres with theory. Using the laboratory we address questions that are not readily answered by theory, while providing a degree of observation, structure and control not present in field data.

Our questions are as follows: Can we increase the provision of information with incentives? Do incentives distort the information provided as predicted or are honesty norms too strong? How do those receiving information respond to the presence of different incentives? What are the impacts on sales and efficiency? Our findings suggest that incentives do have a significant impact on both the quantity and quality of information exchanged. However, where the theory indicates potential for both the *Receiver* and *Vendor* incentive schemes, our results point to the *Vendor* environment as superior.

More information is provided under each of our incentivized environments than in the unincentivized baseline. While ratings are useful to consumers within each of the incentive schemes, the quality of the provided information varies considerably depending on how the incentive is structured. Full misalignment of preferences, as in *Producer* where payments are contingent upon specific product sales, leads to considerable dishonesty among senders and the lowest efficiency of the four environments.

The greatest amount of information transfer occurs in an environment with full alignment in interest, where the receivers of information make a payment to those providing it. Though we find substantial coordination on informative outcomes, this setting provides those receiving information with the worst final outcomes of our four institutions, due to the payments receivers make to acquire information. We instead find that implementing a sales-driven incentive that trades off some alignment-of-interest with the informed party in exchange for greater provision leads to superior outcomes for recipients.

Moreover, where theory predicts greater changes in the quality of information provided under misaligned incentives, we instead see larger responses in quantity within the experiment. Though we do find that the dishonesty of the advice increases as we drive more of a wedge between the provider and recipient, we also find another behavioral effect. The informed party's rate of provision depends upon an interaction with the incentive and the type of information held, where information is exchanged more frequently following a good product experience. That is, subjects react more to the incentive when its conditions match their information, and they reduce their provision in those cases where the incentive would reward dishonesty.

After analyzing our experimental results, we extend our analysis through two counterfactuals. The first extrapolates our findings across a broader range of parameters by varying the distribution of consumers' WTB. The second examines the extent to which marketplaces and producers might profit in the short-run by secretly introducing misaligned incentives. In the long-run exercise, we show that when the incentives on offer are overt and commonly understood, sales-based incentives to those providing advice produce higher final-consumer welfare and increased total sales across a broad range of values for the expected outside options. In particular the sales-based incentive is still superior when consumers have low WTB, and highly trustworthy advice is necessary to generate a sale. This suggests final consumers and marketplaces do well in the long run across a range

of parametrizations. More than this though, we also show that specific producers interested in the sales of their own products do better in the long run. While highly misaligned *Producer*-type incentives are behaviorally effective at translating previous purchases into future product sales, they do so by capturing market share from other manufacturers. In contrast, the *Vendor* incentive leads to a larger increase in the producer's sales by drawing new buyers into the market as a whole, increasing the total market size. A separate analysis of the covert introduction of each type of incentive points to larger instabilities for the aligned-interest incentives: substantial short-run profits are available to manufacturers who surreptitiously offer those providing advice an incentive to gin-up sales. In contrast, these short-run gains are much smaller when introduced into the partially misaligned *Vendor* environment.

Taken together, the two counterfactuals suggest that sales-contingent incentives at the market-place level are the more-efficient stable institution. In comparison to the other environments there are large long-run gains to final consumers, to marketplaces and to producers with existing customer bases. Firms contemplating introducing *Producer*-style marketing schemes into a *Vendor* environment suffer long-run losses from reduced credulity in those providing product guidance, while the short-run gains are much smaller than they would be if interests were fully aligned.

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 WTB.

Our environment has costs for both sending and acquiring messages.⁵ However, these costs are not message specific, and the cost is incurred if any message is sent or received. Because the message cost does not vary with the sender's type, our paper is more related to the cheap talk literature than to the costly signaling literature.

Several papers have examined incentives for communication from the perspective of online ratings. Chen, Harper, Konstan, and Li (2010) use social comparison between users of a movie rating website to increase rating provision, while Wang (2010) argues that increased social identity increases provision. There is also evidence that ratings given in the absence of explicit incentives may be subject to systematic biases as in Hu, Zhang, and Pavlou (2009) and Lafky (2014), which demonstrate the tendency for raters to over-report positive or negative experiences, relative to moderate outcomes.⁶ There is also some evidence for biases due to self-selection, 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 ratings.

⁵A related 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, de Martí, and Prat (forthcoming) for models of endogenous communication with costs.

⁶Also see Bolton, Greiner, and Ockenfels (2013) with respect to managing the distortive effects of reciprocity when ratings are two-way.

A body of work has experimentally examined tensions between agents in the Crawford and Sobel environment.⁷ The main experimental finding 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 our own, 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 institutional changes that influence the cost of sending and receiving messages, and the extent to which senders are aligned with receivers. Taking away the costs of communication, and independent of our own work, Chung and Harbaugh (2012) examine the extent to which observed play matches the equilibrium prediction in a similarly structured sender-receiver environment. Receivers face a choice between two products and an outside option. 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.⁸

2. THEORETICAL FRAMEWORK

In this section we introduce our formal framework. We then introduce the four environments we will analyze, and provide an informal discussion of the most-informative equilibria, where readers interested in more formal constructions are referred to the appendix. After the informal discussion of equilibria for general parameters, we move on to the specifics in the next section, where we introduce the experimental parametrization, and indicate the most-informative equilibrium outcomes.

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

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

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

⁷See Dickhaut, McCabe, and Mukherji (1995), Cai and Wang (2006), Wang, Spezio, and Camerer (2010). For extensions to multiple senders or receivers see Lai, Lim, and Wang (2011); Vespa and Wilson (2014) and Battaglini and Makarov (2014).

⁸See also Charness and Garoupa (2000) which examines the extent to which reputation affects revelation, where senders have a state-independent preference to induce sales.

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

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

The focus of our paper is social-information transmission, the degree to which R 's peers, other consumers with relevant product experience, will provide information to the marketplace that improves R 's decision making. To this end, we introduce a consumer with product experience S (the sender, she). The sender's experience with the product will be fairly natural, she has previously bought one of the two competing products, $X \in \{A, B\}$, and received a product of quality level $x \in \Theta$.⁹ The experienced quality level x provides information on product X 's quality distribution, ϕ^X , so S 's signal (X, x) can be helpful to R 's choice.

Given full information sharing, R updates his prior p_l on product X 's process being ϕ_l to the posterior $q_l(x)$. The three different choices in $\{A, B, R\}$ will now be associated with three expected qualities: the product X with expected quality $\nu_X(x) := \sum_{l=1}^L q_l(x) \cdot \bar{\theta}_l$; the unsampled product $X' \in \{A, B\} \setminus \{X\}$ at $\nu_{X'}(\emptyset) := \bar{\theta}$; and the outside option at ω .¹⁰ A fully informed consumer can simply choose the best option of the three, where the information exchange has two components:

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

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

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

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

A second-best outcome emerges in equilibrium given sender's provision costs, one where consumers contemplating a purchase, the marketplaces where products are sold, and the manufacturers of specific products would all conceivably benefit from greater provision. Our analysis will consider three plausible institutions for increasing provision, and examine equilibrium predictions and experimental behavior in each regime. Under a particular institution ψ , we will examine the expected gross outcome for receivers, $W(\psi) = \mathbb{E}w_R(Z; \psi)$, normalized to indicate the efficiency of

⁹The environment can be extended to either multiple senders $\{S_1, S_2, S_3, \dots\}$ who provide simultaneously, or a sequence of players (R_0, R_1, R_2, \dots) , where receivers $\{R_j\}_{j=0}^i$ act as sender for R_{i+1} .

¹⁰Given an information set \mathcal{I} , we will use the notation $\nu_Z(\mathcal{I}) := \mathbb{E}(z | \mathcal{I})$, where expectations are taken over the process ϕ^Z and the realization $z \sim \phi^Z$.

information transfer as:

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

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

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

- (i) Nature draws a state of the world $(\phi^A, \phi^B, c^S, \omega)$ where ϕ^A and ϕ^B are iid draws from a set of lotteries $\Phi \subset \Delta\Theta$; $c^S \in \mathbb{R}$ is a cost of provision drawn independently from a CDF G ; and $\omega \in \mathcal{H} \subset \mathbb{R}$ is a reservation, drawn independently from CDF H .
- (ii) S observes her message cost c^S , and obtains a single draw x from ϕ^X , where the initial product $X \in \{A, B\}$ is selected with equal probability. Given the signal (X, x) , her choice is a rating/message $m \in \mathcal{M} \cup \{m_\emptyset\}$, where \mathcal{M} is a set of meaningful ratings, and m_\emptyset is the empty message, choosing not to provide. Choosing any $m \in \mathcal{M}$ incurs the cost c^S , while choosing not to provide, $m = m_\emptyset$, is free.
- (iii) R observes whether there was a provided rating, learning either that $\{m \in \mathcal{M}\}$ or that $\{m = m_\emptyset\}$ has occurred. If $\{m \in \mathcal{M}\}$ he first chooses whether or not to acquire the information, $\rho \in \{\text{View}, \text{Not}\}$, incurring a fixed cost $c^R > 0$ only when he views.
- (iv) R observes his private outside option ω , and the precise rating $m \in \mathcal{M}$ if viewing. He then makes a choice $Z \in \{A, B, R\}$, with a realization z from ϕ^Z if a product is chosen, and ω if he selects the reservation.¹¹

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

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

where $\alpha > 0$ is a preference parameter reflecting a prosocial incentive to help receivers make better product choices.¹² The institution ψ is reflected by a conditional transfer $T > 0$ to ($t > 0$ from) the sender (receiver), conditioned on the specific events indicated through $\psi_S(m, \rho, Z) \in \{0, 1\}$ (and $\psi_R(m, \rho, Z) \in \{0, 1\}$). The transfer conditions can therefore respond to any of the action choices made in the game, where our paper will focus on four simple variations.

A pure strategy for the sender is a rating choice $\mu : \{A, B\} \times \Theta \times \mathbb{R} \rightarrow \mathcal{M} \cup \{m_\emptyset\}$, a decision on the rating to send given the signal (X, x) and provision cost c^S . The strategy for the receiver is the tuple $(\rho, \{\zeta_m\}_{m \in \mathcal{M}}, \zeta_{\mathcal{M}}, \zeta_\emptyset)$: a listening decision $\rho \in \{\text{View}, \text{Not}\}$, and a product choice $\zeta_{\mathcal{I}} : \mathcal{H} \rightarrow \Delta\{A, B\} \cup \{R\}$ for every possible information set, $\mathcal{I} \in \bigcup_{k \in \mathcal{M}} \{m = k\} \cup \{m \in \mathcal{M}\} \cup \{m = m_\emptyset\}$.¹³ The relevant choices for R are: i) viewing the provided rating m , and responding with ζ_m ; ii) knowing that a rating was provided, $m \in \mathcal{M}$, but not viewing it, with a response $\zeta_{\mathcal{M}}$; and iii) knowing that no rating was provided, and responding with ζ_\emptyset . Beliefs at every information

¹¹We assume that the realization of ω happens after rating viewing for tractability. The first-order effect from alternating the order is to reduce viewing behavior for those R -types with high reservations.

¹²The sender having a lexicographic preference over the receiver outcome, subordinate to her own outcome, would also suffice for our needs. The required assumption is simply that conditional on sending a non-empty message, the sender and listening receiver are strongly aligned in interest over the choice Z .

¹³We allow for receivers to randomize over the two products A and B to maintain symmetry in the case where no information is revealed. However the focus of the paper is on how ratings provided by S allow us to break symmetry.

set \mathcal{I} for S and R are given by $\lambda_S(\mathcal{I})$ and $\lambda_R(\mathcal{I})$, which are conditional distributions over the entire state $(\phi^A, \phi^B, c^S, \omega)$.

Our solution concept will be Perfect Bayesian equilibrium (PBE), $\langle \mu^*, \rho^*, \zeta^*; \lambda_S^*, \lambda_R^* \rangle$, where we focus on illustrating the most-informative and least-informative PBEs under two symmetry restrictions. The first restriction is that we restrict attention to symmetric messaging strategies over the two products, so that there exist complementary messages, and similar quality experiences (the specific draw x) lead to similar ratings, regardless of the precise sampled product identity ($X = A$ or $X = B$). The second restriction is that when agents are indifferent over $Z \in \{A, B\}$, any resulting ties are broken by an equal randomization between them, where all other strategy components are pure. The reasoning behind these restrictions is to focus on symmetry between the two products being broken by ratings ex post, and not through an ex ante coordination on a particular product.¹⁴

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

Baseline. No conditional transfers, so $\psi^S(m, \rho, Z) = \psi^R(m, \rho, Z) = 0$ for all possible action choices. So long as c^R is not too large and there are senders with low enough provisions costs, information transfer is possible in equilibrium.¹⁵ However, because the provision cost c^S is privately incurred by senders, there may be draws for which senders do not provide a rating, and select m_\emptyset instead. The main tension in the *Baseline* environment is on provision, where rational senders will under-provide information relative to the benefit receivers derive.¹⁶ In an environment with transferable utility where α is small, a Pareto improvement will be possible at costs $c^S > 0$, where receivers pay for rating provision, which leads us to the next institution.

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

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

In any equilibrium where receivers view, the effect is to reduce the cost of provision. When the drawn cost is c^S , the effective cost to the sender is $c^S - T$, with the effective result that the cost distribution G is shifted to the left by T . So the institution can increase provision, while maintaining alignment in the sender's incentive. However, there are potential downsides to this incentive. Transfers of this form do not preclude coordination on low-information outcomes—in particular

¹⁴Moreover, in our experimental environment, we will explicitly prohibit any such coordination by randomly relabeling A and B between the sender and receiver.

¹⁵A sufficient condition for the second part is that $G(0) > 0$, so a positive fraction of senders have no cost or enjoy sending. In the experiment subjects do receive negative cost draws for sending, with the interpretation being that these senders (net any costs) enjoy sending information to others.

¹⁶In informative equilibria, the decision to provide or not can be characterized by a signal-specific cutoff $c^*(\nu_X(x); \psi)$ where the sender provides a rating in $m \in \mathcal{M}$ for all $c^S \leq c^*(\nu_X(x); \psi)$. The theoretical appendix provides additional details.

babbling equilibria—where the selected ratings have low information content and receivers choose not to view given the acquisition cost c^R . In fact, the equilibrium coordination problem becomes riskier for receivers than the baseline case. If R mis-coordinates on the informative equilibrium he pays the additional cost t whenever S is coordinating on babbling.¹⁷ Even with perfect coordination on informative outcomes, depending on the size of the transfer t , the *Receiver* institution alters the distribution of the surplus generated by information transfer, giving more of the gains to senders. Receivers have an option to view or not, which ensures that they cannot be forced below the lower-bound expected outcome \underline{W} . However, if the institution is fixed and the most-informative equilibrium always selected, receivers can be made worse off in the *Receiver* setting than they were in the *Baseline*. By requiring receivers to pay t to view ratings, receiver’s expected outcomes can be reduced to the lower bound \underline{W} . Though this institution can generate greater efficiency than the *Baseline*, it can also fully redistribute the full surplus, $W(\text{Rec.}) - \underline{W}$, not just the additional gain $W(\text{Rec.}) - W(\text{Base.})$.

Vendor Transfer. Our third environment considers an intermediary vendor player C , who sells the products A and B , producing a fixed profit when either is sold, but making no profits when consumers stick with their outside option. The vendor therefore has an interest in increasing WTB, but is indifferent over which specific product is purchased. To drive sales volume, the vendor will be interested in incentivizing senders to generate sales. We assume that the form of the transfer is conditioned on the sale of either product and a viewed rating.¹⁸ For each sale generated by the sender, the vendor player transfers t , for a total of $T = \eta \cdot t$. The condition for the transfer to the sender is

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

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

Similar to the receiver transfer, this environment has the effect of reducing sender’s provision costs in informative equilibria, and therefore increasing the quantity of information. However, if the conditional transfer T is large in relation to the altruism term α , the transfer has a negative effect on alignment of interest between S and R : they are now only partially aligned. Senders derive a benefit from receivers buying a product, regardless of their particular reservation ω , though they are still aligned with the receiver over which product A or B should be chosen. For α small, the most-informative equilibria restrict senders to providing just two effective ratings, with the ordinal interpretations “ $A \succ B$ ” and “ $B \succ A$.” The statement that $A \succ B$ tells the receiver that the sender’s signal is in $\{(X, x) \mid \mathbb{E}[w_R(A) \mid x] > \mathbb{E}[w_R(B) \mid x]\}$.¹⁹

¹⁷In a babbling equilibrium, ratings in \mathcal{M} are still provided whenever $c^S < 0$. However if the same rating m_X is provided at all signals (X, x) , then provided ratings have zero information content, and it is a best response not to view.

¹⁸All of this information is commonly available to marketplaces. Consider an offer where the marketplace provides previous customers with a unique hyperlink to give to friends when they recommend a product. When purchases are made using the link, the marketplace observes the sale, and that the link was transmitted between the two.

¹⁹If α is very large and T small, more-informative equilibria exist. In cases where T is not large enough for senders to provide at every c^s (and α is not negligibly small) changes in the probability of providing a rating given the signal x must also be incorporated. This effect is very small under our experimental parametrization at the equilibrium, and so we will not focus on it in our discussion of the theory.

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

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

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

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

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

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

More developed and detailed constructions for the theory, and a formal description of the equilibrium concept used to generate the theoretic predictions are provided in the appendix. We now illustrate the specifics through our experimental parametrization.

3. EXPERIMENTAL DESIGN

3.1. Experimental Procedures. We utilize a *between-subject* experimental design over four treatment environments, paralleling a *Baseline* environment with no transfers, and the three conditional transfers to senders described above: a *Receiver* transfer conditioned on a viewed rating; a *Vendor* transfer conditioned on a viewed rating and a sale; and a *Producer* transfer from producer X , conditioned on a sale of product X . In each treatment, subjects participate in 30 rounds of the fixed environment, with random anonymous matching. We now detail the chosen parametrization for the experiment.

Determining the state. The primary uncertainties in the model are the quality distributions ϕ^A and ϕ^B . The experiment’s proxy for these distributions are two urns, *Urn A* and *Urn B*. Each urn is filled with two balls labeled with an integer between 1 and 100, so the possible set of realizations for each ball is $\Theta = \{1, \dots, 100\}$. In every round t and for every sender i , the two urns are independently filled via the following procedure: i) An initial ball θ_1 is placed in the urn, where each integer value from 1 to 100 is equally likely; ii) with probability $1/2$, the second ball θ_2 is another independent uniform draw from 1 to 100, and with probability $1/2$ the second ball is a copy of the first ball, so $\theta_2 = \theta_1$. An urn X is completely determined by the two balls in the urn, θ_1^X and θ_2^X , and choosing urn A or B is equivalent to choosing the lotteries $\frac{1}{2} \cdot \theta_1^A \oplus \frac{1}{2} \cdot \theta_2^A$ or $\frac{1}{2} \cdot \theta_1^B \oplus \frac{1}{2} \cdot \theta_2^B$, respectively. The set of all possible urns, $\Phi = \{\phi_r \in \Delta\Theta \mid \exists \theta, \theta' \in \Theta \text{ s.t. } \phi_r = \frac{1}{2} \cdot \theta \oplus \frac{1}{2} \cdot \theta'\}$, has 5,050 possible lotteries, where the prior probability p_r of the urn being the lottery $\phi_r = \frac{1}{2} \cdot \theta \oplus \frac{1}{2} \cdot \theta'$ is given by

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

The particular balls in an urn are correlated—an urn with two identical balls is as likely as an urn with two independent balls—and the expected subsequent draw given a previous draw x is $\nu_X(x) = 3/4 \cdot x + 1/4 \cdot \bar{\theta}$. 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 expectation $\bar{\theta}$).

The other components of the state are the sending cost c^S and the reservation ω , which are determined as follows: the cost c^S is an iid draw from $\mathcal{G} = \{-\$0.49, \dots, \$2.50\}$, where the probability of each cost level is linearly decreasing to zero at $c^S = \$2.50$; the reservation value ω_j is a draw from $\mathcal{H} = \Theta = \{1, \dots, 100\}$, where the probability of each reservation value is decreasing linearly to zero for $\omega = 100$.²⁰

Senders. In the first stage of each round t , each of the n subjects in the session is randomly matched to an identity $i \in \{1, \dots, n\}$ and chooses an urn $X_{it} \in \{A, B\}$ to draw a ball from, observing a ball with value x_{it} , drawn from the distribution ϕ_{it}^X , so the sender’s signal is (X_{it}, x_{it}) . The receiver for the subject with identity i is the subject with identity $j = \text{mod}(i, n) + 1$ (where i acts as the receiver for $k = \text{mod}(i - 2, n) + 1$ in the same round). In this way, the n subjects are randomly

²⁰In the experiment we draw two sending costs uniformly from \mathcal{G} , and give the subject the lower cost of the two. The probability of getting the cost $c_i^S \in \mathcal{G}$ is approximately $\frac{2}{300} \frac{2.5-c}{3}$, which can be thought of as an affine transformation of a Beta(1,2) distribution to the support $[-0.5, 2.50]$. For the reservation we draw two uniform draws from Θ , and the reservation draw ω_j is the smaller value. The probability of a reservation draw ω is therefore approximated by $\frac{2}{100} \frac{100-\omega}{100}$, a Beta(1,2) over $[0, 100]$.

formed into a circle, and urns passed clockwise around it. Each subject in each round acts first as a sender and then as a receiver, but for two different randomly matched participants.

As a sender the subject receives $\$0.05 \times x_{it}$ for their first decision X_{it} . Though from the perspective of the sender x_{it} is simply a signal draw, we wanted senders to experience a payoff from their signal. After observing the drawn ball's value, the subject is asked to provide an integer rating from one to five for either urn, so the set of available non-empty ratings is: $\mathcal{M} = \{1_A, 2_A, 3_A, 4_A, 5_A\} \cup \{1_B, 2_B, 3_B, 4_B, 5_B\}$. After selecting their provisional rating $\hat{m}_{it} \in \mathcal{M}$, the subject is informed of the drawn provision cost c_{it}^S , and must choose whether or not to send the selected rating, $m_{it} \in \{\hat{m}_{it}, m_\emptyset\}$, where they incur the (possibly negative) cost c_{it}^S only if they send. The experimentally induced altruism term α is set to 0.1 by providing i with $\$0.10$ for every dollar received by the matched receiver j for the receiver choice Z_{jt} .²¹ In all institutions where provision has an incentive, the conditional payment to the sender is fixed at $T = \$2.00$, while the conditions for the payment $\psi^S(m, \rho, Z)$ vary by treatment. The impact on round payoffs from the sending decision is therefore

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

Receivers. In the second stage of the game the subject in role j receives the urns A and B from the subject with identity $i = \text{mod}(j - 2, n) + 1$.²² If there is a provided rating $m_{it} \in \mathcal{M}$, then j is informed a rating exists and j 's first decision as a receiver is ρ_{jt} , whether or not to view. If they view they incur the fixed cost $c^R = \$0.05$ in all four institutions, while in the *Receiver* environment an additional amount $t = \$0.50$ brings the total view cost to $\$0.55$.

After realizing their information set \mathcal{I}_{jt} —either observing the specific rating if they view, $\mathcal{I}_{jt} = \{m_{it}\}$, that a rating was provided if they do not view, $\mathcal{I}_{jt} = \{m_{it} \in \mathcal{M}\}$, or that no ratings was provided, $\mathcal{I}_{jt} = \{m_{it} = m_\emptyset\}$ —the receiver then makes a provisional decision between the two urns $\hat{Z}_{jt} \in \{A, B\}$. Finally, they are informed of their reservation draw ω_{jt} , and asked to make a final choice $Z_{jt} \in \{\hat{Z}_{jt}, R\}$ between the ball of unknown value z_{jt} drawn from urn \hat{Z}_{jt} , or the reservation ball of certain value ω_{jt} . Their consumption payoff $w_R(Z_{jt})$ is either $\$0.10 \times z_{jt}$ or $\$0.10 \times \omega_{jt}$, respectively, so receivers get $\$0.10$ multiplied by the value of the final ball selected. The total payoff from the receiver phase is subsequently

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

Last 15 rounds. The second half of each session has a similar structure to the above. However, after the first 15 rounds, subject decisions are collected through a strategy method to collect a richer set of information.²³ Instead of showing subjects the message cost c_{it}^S and choosing whether or not they wish to send the provisional rating \hat{m}_{it} , they now specify a cutoff cost $\bar{C}_{it} \in \mathcal{G}$ below which they would be willing to send. Similarly, receivers are not informed of their specific reservation draw ω_{jt} , and they instead specify a reservation cutoff $\bar{\Omega}_{jt} \in \mathcal{H}$ (the certainty equivalent) below

²¹The theoretical characterizations for the *Vendor* and *Producer* scenarios might have qualitative changes if subjects coming into the lab had large enough innate altruism towards other subjects. Moreover, all informative equilibria would have quantitative shifts in the provision levels were the true α different from the induced one.

²²Urn A from subject i is relabeled as either *Urn C* or *Urn D* for subject j (with equal probability), so receiver j cannot coordinate on knowing the urn X_{it} that i chose from. For tractability though we will continue to refer to them as urns A and B rather than urns C and D .

²³For a survey of papers examining (and for the most part supporting) the strategy method relative to direct-response see Brandts and Charness (2011).

TABLE 1. Most-Informative Risk-Neutral Equilibrium Predictions

	Baseline (B)	Receiver (R)	Vendor (V)	Producer (P)	Comp. Static
Distinct Ratings	10	10	2	0/Babbling	
Rating Provision	0.346	0.980	0.934	0.304	R>V>B>P
Info. Efficiency, Υ	36.3%	98.1%	87.6%	0.0%	R>V>B>P
Rec. Welfare, Υ_R	34.1%	30.6%	81.7%	0.0%	V>B>R>P
Sales (Any Product)	0.776	0.819	0.827	0.750	V>R>B>P
Sales (Same Product)	0.403	0.445	0.414	0.375	R>V>B>P
Conditional on no provided rating					
Info. Efficiency	-1.0% [†]	-5.1% [†]	-2.2% [†]	0.0%	P>B>V>R
Conditional on provided rating					
View Rate,	1.0	1.0	1.0	0.0	B~R~V>P
Info. Efficiency Υ	106.6% [†]	100.3% [†]	93.9%	0.0%	B>R>V>P
Rec. Welfare Υ_R	100.4% [†]	31.4%	87.6%	-6.3%	B>V>R>P
Sales (Any Product)	0.824	0.821	0.832	0.750	V>B>R>P
Conditional on provided rating and sale					
Sales (Same Product)	0.552	0.543	0.501	0.500	B>R>V>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.

which they would choose the risky draw z_{jt} , and above which they would choose the realized reservation ω_{jt} .

Data from interactions in rounds 1–15 are therefore given by $\langle (\phi_{it}^A, \phi_{it}^B, c_{it}^S, \omega_{jt}), (X_{it}, \hat{m}_{it}, m_{it}), (\rho_{jt}, \mathcal{I}_{jt}, \hat{Z}_{jt}, Z_{jt}) \rangle$; while data from rounds 16–30 are given by $\langle (\phi_{it}^A, \phi_{it}^B, c_{it}^S, \omega_{jt}), (X_{it}, \hat{m}_{it}, \bar{C}_{it}), (\rho_{jt}, \mathcal{I}_{jt}, \hat{Z}_{jt}, \bar{\Omega}_{jt}) \rangle$.²⁴

3.2. Equilibrium Predictions. Given the parameters chosen for the model above, Table 1 provides predictions for the main economic variables in the most informative PBE (under the symmetry restrictions) in each treatment. These predictions assume risk neutrality with respect to the final monetary gambles, where a comparable table in the appendix presents predictions under a risk-averse formulation.

For both the *Baseline* and *Receiver* treatments, the most-informative equilibrium involves sending a single rating 1_X for all senders types (X, x) with $x \leq 50$ and with the ratings 2_X-5_X selected after the signals x in 51–61, 62–73, 74–86, and 87–100, respectively. As such, both treatments utilize all five ratings for both products, and all ten ratings are sent with positive probability.²⁵ The

²⁴The corresponding realized decisions are given by: $m_{it} = \hat{m}_{it}$ if $c_{it}^S \leq \bar{C}_{it}$ and otherwise $m_{it} = m_0$; and $Z_{jt} = \hat{Z}_{jt}$ if $\omega_{jt} < \bar{\Omega}_{jt}$ and $Z_{jt} = R$ otherwise,

²⁵Less efficient equilibria exist where, for instance, the messages 3, 4 and 5 imply products that are better than average. With just one sender and homogenous risk preferences, messages implying a below-average product are

first-order difference in behavior across the *Baseline* and *Receiver* treatments is the rate at which senders provide information, $\Pr \{ \mu^* (X, x, c^S) \in \mathcal{M} \}$, given in the *Rating Provision* row in Table 1. In the *Baseline*, information is provided approximately a third of the time, while in the *Receiver* case, provision is almost complete.²⁶ The extent to which decision-relevant information is exchanged is measured by Υ which is given in the *Info. Efficiency* row. This measure examines the expected receiver outcome $w_R(Z)$ relative to no communication, set against an upper bound derived from perfect information exchange.

Given selection of the best-case PBE, communication generates 36.3 percent information efficiency in the *Baseline*, which increases to 98.1 percent in the *Receiver* treatment.²⁷ However, the gains in information efficiency do not come for free. While the average sender is much better off in the most-informative *Receiver* equilibrium than in the *Baseline*, the provision subsidy has to be paid for. Even though the *Receiver* treatment increases total efficiency, once we account for the payment they make to the sender, receivers actually fare worse in this parametrization than the *Baseline*. This is illustrated by the *Rec. Welfare* row in Table 1, which modifies the efficiency measure to account for the incurred costs of communication to the receiver,

$$\Upsilon_R = \frac{\mathbb{E}_{\theta_X^S, c^S, \omega} \left[u_R \left(\zeta_{\mathcal{I}(\mu^*)}^*(\omega), \rho^*, \mu^*(\theta_X^S) \right) \right] - \underline{W}}{\overline{W} - \underline{W}} = \Upsilon - \frac{(c^R + t_\psi) \cdot \Pr \{ \rho = \text{View}, m \in \mathcal{M} \}}{\overline{W} - \underline{W}}.$$

That is, Υ_R modifies the information efficiency Υ to include the receiver's viewing costs where the t_ψ component is the transferred amount, and is directly proportional to a receiver's expected utility.

The raw cost to view a message in our experiment, $c^R = \$0.05$, represents a 6.3 percent efficiency penalty if incurred with certainty. In the *Baseline* prediction, where ratings are provided approximately a third of the time, the cost leads to a receiver welfare of 34.1 percent. However, there is a much larger drop in the *Receiver* treatment, with receiver welfare falling to just 30.6 percent. There are two reasons for this larger drop: First, ratings are predicted to be provided at much higher rates, in almost every round, so acquisition costs are incurred more often. Second, and the quantitatively larger force, much of the receiver's gain from increased efficiency is transferred to the sender through the incentive payment t . In fact, despite a large increase in total efficiency, receivers' final outcomes are inferior to the most-informative *Baseline* prediction. As there are multiple equilibria in the *Baseline* and *Receiver* institutions, another effect of reduced receiver welfare is to make coordination on the most-informative equilibrium riskier.²⁸

strategically identical, so though ratings of one and two can be sent for 1–25 and 26–50, respectively, there is no incentive for this split as each rating carries a similar interpretation of “do not buy.” However, with more than one sender, ratings distinguishing between levels of below-average quality may help to aggregate information on the urn's composition.

²⁶In informative equilibria, fixing the signal (X, x) , the same rating $\mu^*(X, x, c^S) \in \mathcal{M}$ is sent whenever the cost of provision c^S is less than or equal to a cutoff $c^*(\nu_X(x))$, and for all send costs greater than this the empty rating m_\emptyset is chosen. So long as $\alpha > 0$, the cutoff $c^*(\nu_X(x))$ is increasing in $|\nu_X(\theta_X^S) - \bar{\theta}|$. Given the parametrization, and selection of the most-informative symmetric PBE, this leads to a U-shaped effect on provision conditional on the drawn signal x . The effect is strongest in the *Baseline*, and completely absent in the *Producer* treatment.

²⁷Constraints on the message space, that $|\mathcal{M}| \ll |\{A, B\} \times \Theta|$, does not substantially affect information transfer here. The four ratings with inferences $\nu_X(m = k_X) > \bar{\theta}$ are used to distinguish between choosing the product X and the reservations in $\omega \in [51, 87]$.

²⁸Assume that babbling equilibrium ratings are the (less harmful on mis-coordination) 1_X -type messages that warn receivers away from an urn, and are sent whenever $c^S < 0$. Receivers in the *Baseline* treatment will choose to view whenever they place more than 0.158 probability on ratings $m \in \mathcal{M}$ being informative. In contrast, in the *Receiver*

In contrast with the above two environments, the *Vendor* setting’s most-informative equilibrium uses just two effective ratings from the ten available. The equilibrium can be characterized by the rating strategy $\mu(X, x, c^S) = 5_Y$, where $Y = X$ for $x > 50$ and $Y \neq X$ when $x \leq 50$, so long as the provision cost c^S is not too large. These ratings are equivalent to the statement “Given my information, I expect A to be better than B.” Though less useful than ratings that convey more refined information on the sender’s signal, this statement is still informative. An equilibrium receiver viewing the rating 5_Y in *Vendor* will choose the product Y for all reservation values $\omega < 60$, and not to buy whenever $\omega \geq 60$.²⁹

Because the incentive payment in the *Vendor* treatment is conditioned on receivers choosing one of the two products, the expected subsidy to a sender providing a rating J is $\Pr\{\zeta_J(\omega) \neq R\} \cdot T$, which is smaller than the *Receiver* case where the full payment is received. This drop in the expected incentive reduces provision relative to the *Receiver* treatment. The 87.6 percent information efficiency in *Vendor* is lower than the *Receiver* case—through reduced provision and coarser revealed information—but is still a substantial improvement on the *Baseline*. For receivers though, because they not paying for the incentive payment, they share more of the gains from increased efficiency, and 81.7 percent receiver welfare, the highest within our four treatments. Indeed, both senders and receivers do better than the *Baseline*, despite the misalignment introduced over WTB.

The gains in *Vendor* are facilitated by a transfer from by an unmodeled *Vendor* player. To understand whether vendors would want to subsidize provision in equilibrium, Table 1 also provides the fraction of receivers making a purchase in each environment, both unconditionally and conditional on a provided rating, in the *Sales (Any Product)* rows.³⁰ Quantitatively, the predicted equilibrium effects are not large, which is an artifact of the particular distribution H chosen for the experiment.³¹ However, relative to the other environments, the *Vendor* treatment does produce the highest equilibrium sales rate (both overall, and conditional on provision). The reasons behind this are greater provision and the fact that every viewed rating increases receivers’ WTB. In contrast, for the *Receiver* setting, despite a higher rate of provision, approximately half of the provided ratings will convey negative information, which do not affect WTB. Given increased sales relative to both *Baseline* and *Receiver*, a large enough profit per sale (or equivalently a large enough quantity of receivers, η) would allow *Vendor* subsidies to pay for themselves in equilibrium.³²

Finally, the *Producer* treatment only has a babbling equilibrium outcome. Though ratings are still provided (for non-positive costs c^S), receivers choose not to view them, as $c^R > 0$ and the rating conveys no useful information—where this same babbling equilibrium is present in all four treatments. Payments conditioned on sales of a particular product are predicted to be counterproductive in the long run, and it is clear that the producers A and B would not want to

treatment, the additional cost of viewing means receivers must place more than 0.694 probability on ratings being informative to view, and so babbling and not viewing is risk dominant.

²⁹The exact choice stems from the expected quality of product Y given the message 5_Y being $\frac{5}{8}\bar{\theta} + \frac{3}{8}\mathbb{E}(x|x > 50) = 59.875$. Note that unlike the previous cases, the rating here is also informative about the product $Y' \neq Y$, which has an expected quality of 41.125.

³⁰The expected product sales per consumer without the possibility of a rating is $H(\bar{\theta}) = 0.75$.

³¹Whether or not a vendor player is willing to make this transfer in equilibrium depends on: i) new sales created when paying the subsidy, here $H(60) - H(\bar{\theta})$; ii) the profitability of each sale relative to the subsidy cost; iii) whether or not the ratings would have been sent anyway without the subsidy.

³²Vendor players have an incentive to secretly provide sale-conditioned incentives in the *Receiver* and *Baseline* environments, as sending glowing reviews is more effective in increasing sales if receivers are completely credulous, and a five rating increases the WTB beyond 60.

pay for incentives in equilibrium.³³ To get a sense for how each institution affects the producer X , the *Sales (Same Product)* rows of Table 1 indicate the fraction of product sales where the receiver chooses the same product as the preceding sender (where $Z = X$)

The same-product recurrence rates are somewhat counter-intuitive: in the *Baseline* and *Receiver* treatments, senders are completely aligned with receivers, yet these are the only institutions where viewing a rating leads to significantly greater same-product sales than random choice. There are two reasons for this. The first (and much smaller effect) is that ratings are more likely to be provided when the sender has had a positive experience. The second (and quantitatively larger) component is through the influence of differing ratings on receiver’s WTB, where this effect only exists in the *Baseline* and *Receiver* equilibria. When ratings are provided, just under half are negative, but negative ratings still result in an $H(\bar{\theta})$ fraction of the receivers making a purchase (just for the alternative, unrated product). Positive ratings on the other hand, which are sent just over half the time, lead to $H(\nu(m))$ sales for the rated product, as the positive information increases the receiver’s WTB. Theoretically, all of the increases in total sales from communication in the *Baseline* and *Receiver* equilibria accrue to the same product sampled by the sender. In contrast, for *Vendor* and *Producer*, the different ratings provided in equilibrium have no effect on receiver’s WTB, and so there is no effect on same-product sales. Within the table, there is an almost negligible same-product sales effect in *Vendor*, where this stems from slightly higher provision rates when senders observe better-than-average product draws.

Our paper’s main hypotheses are the comparative-static orderings between treatments for each outcome variable in Table 1 where the orderings are not affected by moderate levels of risk aversion. Table 10 in the appendix provides similar comparisons for risk averse consumers.

4. RESULTS

Below we present summary results and analysis of data from sessions conducted at the Pittsburgh Experimental Economics Laboratory. We collected 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 the treatment averages for the economic variables of interest, and test the comparative statics derived from the most-informative equilibrium predictions. We subsequently go on to detail the behavior of subjects, and outline the “why” of the economic results by examining the component behavior. Finally, in the next section we extend the results with a counterfactual exercise examining long- and short-run effects from changes to our experimental environments.

4.1. Broad Economic Outcomes. Table 2 provides a summary of the main outcome variables in the experiment. Each of the rows provides the sample analog to a prediction in Table 1. However, because the most-informative equilibria have receivers either always or never viewing, we additionally include averages for the subsample with viewed ratings ($m_{it} \in \mathcal{M}$ and $\rho_{jt} = \text{View}$).

The first five rows in Table 2 outline the unconditional effects for the main economic variables: i) the rate of provision by senders; ii) the overall efficiency of the information transfer, Υ ;³⁴ iii)

³³However, such incentives deployed secretly in any of the other treatment environments would increase same-product sales in the short run, holding constant receiver’s behavior. We discuss such shifts in short-run incentives in section 5.

³⁴All 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. We could further

TABLE 2. Experimental Outcomes

	Baseline ($S = 46$)			Receiver ($S = 40$)			Vendor ($S = 47$)			Producer ($S = 44$)			Comparative Static	
	Avg.	Std. Err	N	Avg.	Std. Err	N	Avg.	Std. Err	N	Avg.	Std. Err	N	Theory	Data
Provision	0.369	0.013	1380	0.605	0.014	1200	0.634	0.013	1410	0.488	0.014	1320	$R \succ V \succ B \succ P$	$V \succ^* R \succ^{***} P \succ^{***} B$
Info. Efficiency	9.0%	3.9%		28.1%	4.1%		17.7%	3.8%		1.0%	3.9%		$R \succ V \succ B \succ P$	$R \succ^{**} V \succ^* B \succ^* P$
Rec. Welfare	7.3%	3.8%		-2.9%	4.1%		14.1%	3.8%		-1.1%	3.9%		$V \succ B \succ R \succ P$	$V \succ B \succ^* P \sim R$
Sales (Any Product)	0.754	0.006		0.771	0.012		0.774	0.011		0.759	0.012		$V \succ R \succ B \succ P$	$V \sim R \sim P \sim B$
Sales (Same Product)	0.386	0.013		0.383	0.014		0.417	0.013		0.425	0.014		$R \succ V \succ B \succ P$	$P \sim V \succ^{**} B \sim R$
Conditional on no provided rating														
Info. Efficiency	-10.8%	4.5%	871	-11.8%	6.1%	474	-39.2%	5.8%	516	-30.5%	5.4%	676	$P \succ B \succ V \succ R$	$B \sim R \succ^{***} P \succ^{\circ} R$
Conditional on provided rating														
Viewing	0.735	0.020	509	0.742	0.016	726	0.899	0.010	894	0.688	0.018	644	$B \sim R \sim V \succ P$	$V \succ^{**} R \sim B \succ^{**} P$
Conditional on viewed rating														
Info. Efficiency	58.4%	7.6%	374	82.4%	6.3%	539	57.6%	5.2%	804	45.1%	7.0%	443	$B \succ R \succ V \succ P$	$R \succ^{**} B \sim V \succ^* P$
Rec. Welfare	52.1%	7.6%		13.4%	7.9%		51.3%	6.7%		38.8%	7.9%		$B \succ V \succ R \succ P$	$B \sim V \succ^* P \succ^{**} R$
Sales (Any Product)	0.781	0.021		0.820	0.017		0.792	0.014		0.788	0.020		$V \succ B \succ R \succ P$	$R \succ^{\circ} V \sim P \sim B$
Conditional on viewed rating and sale														
Sales (Same Product)	0.510	0.029	292	0.534	0.024	442	0.576	0.020	637	0.688	0.025	349	$B \succ R \succ V \succ P$	$P \succ^{**} V \succ^* R \sim B$

Note: Tests are derived from Wald tests after regression (probit) results for Efficiency/Welfare (Provision/Viewing/Sales decisions) on treatment dummies. The relationship $A \succ^{***} B$ (\succ^{**} , \succ^* and \succ°) represent rejection of mean equality for treatments A and B with the one sided alternative $A \succ B$ at 99 percent confidence (95, 90 and 80 percent confidence, respectively). The relationship $A \sim B$ reflects failure to reject at 80 percent confidence. We list binary pairs of tests according to the data averages; all statistical relationships are transitive at the relevant level except for Inf. Efficiency given a viewed rating where $B \succ^{\circ} V$. Where $A \sim B \sim C$ in many cases $A \succ C$, etc.

the receiver welfare, which accounts for the expected cost of viewing, Υ_R ; iv) the proportion of receivers opting to make a purchase of any product, $\hat{\Pr}\{Z_{jt} \neq R\}$; and v) the proportion of receivers opting to purchase the same product as the sender, $\hat{\Pr}\{Z_{jt} = X_{it}\}$.³⁵

After the unconditional outcomes we provide conditional averages assessed over the relevant subsamples. The first conditional measure examines outcomes without a provided rating, where we indicate the information efficiency $\Upsilon|m_\emptyset$. The second conditional measure is restricted to rounds with a rating, $m_{it} \in \mathcal{M}$, where we indicate the proportion of receivers who view, $\hat{\Pr}\{\rho_{jt} = \text{View}|\mathcal{M}\}$. We next consider only those rounds with exchanged information, and examine the subsample with viewed ratings. Here we parallel the unconditional results, indicating the information efficiency, receiver welfare and both overall and same-product sales levels. However, to make clear the degree to which ratings persuade receivers to purchase the same product, we also indicate same-product sales conditional on both a viewed rating and any sale, so $\hat{\Pr}\{Z_{jt} = X_{it} | Z_{jt} \neq R, \mathcal{M}, \text{View}\}$ —where random choice would select the same product $1/2$ the time.

Efficiency of information transfer. Our unconditional efficiency results match the comparative-statics generated by the most-efficient equilibrium predictions. At one end, the *Receiver* institution 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 *Vendor* and *Baseline* treatments, in line with their theoretical efficiency rankings. Ordinarily, the results indicate that the most-informative equilibrium prediction does quite well. However, once we compare the cardinal levels in Table 1 and 2, we see quantitatively large differences. The *Receiver* and *Vendor* 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 possible.

One reason for the drop is the large difference in provision and viewing rates from the predicted level in *Receiver* and *Vendor*. To control for this we can instead look at the information efficiency conditional on a provided and viewed rating. 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, the *Baseline* is significantly worse than the *Receiver* at just 58 percent efficiency, where the best-case theory predicts comparable levels.³⁶ Instead, we observe a large and significant efficiency gap between *Receiver* and *Baseline*, and no significant difference between *Vendor* and *Baseline*. Instead of the 94 percent efficiency possible after an equilibrium ratings is provided, we observe just 58 percent in *Vendor*.

Result 1 (Efficiency). *A transfer from Receivers to Senders to defray provision costs leads to significant increases in information transmission. However, this gain is less than theoretically predicted due to lower provision and viewing rates.*

reduce the exogenous noise by focusing on the last 15 rounds and using the supplied cutoffs \bar{C}_{it} and $\bar{\Omega}_{jt}$ to remove variation introduced through G and H but here we provide results over all 30 rounds.

³⁵Here, and henceforth, $\hat{\Pr}\{\mathcal{T}|\mathcal{S}\}$ will denote the sample proportion $\frac{1}{|\mathcal{S}|} \sum_{s \in \mathcal{S}} \mathbf{1}\{s \in \mathcal{T}\}$.

³⁶The *Baseline* is actually predicted to do better than *Receiver*, conditional on provision, as ratings are relatively more likely to be provided following a positive draw x than a negative one, where these differences in prediction depend on the shape of cost CDF, G .

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 through the viewing cost c^R . However, in the *Receiver* treatment, the transfer t 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 the costs involved, receivers who view a rating improve their outcomes relative to the expected outcome with no sender, \underline{W} . Moreover, we will later show that viewers can expect to do better than those choosing not to view, accounting for any selection effects.

The most important result with regard to receiver outcomes is that their maximum payoffs are achieved in the *Vendor* treatment. Conditional on a viewed rating, the receiver welfare row of Table 2 shows that receivers do just as well in the *Baseline* and *Vendor* treatments, so the effect of misalignment over WTB in the *Vendor* incentive seems to be small. Trading off some quality in the information provided in order to subsidize ratings leads to large quantity increases. The unconditional receiver welfare figures therefore attain a maximum in *Vendor*. However, the overall difference in receiver welfare compared to the *Baseline* is only marginally significant. The reason for this (and for the comparable efficiencies conditional on a viewed rating) are due to negative efficiency for those receivers who are not provided with a rating. That is, those receivers without a rating do worse than the lower bound level predicted with zero information transfer. This effect is strongest in the *Vendor* treatment, where the absence of a rating leads to an efficiency level of -39 percent, though the magnitude of this effect in *Producer* is comparable. As we will later demonstrate when we examine behavior in more detail, the cause of these drops is that ratings in the incentivized environments are less likely to be provided by senders with bad signals, which depresses final outcomes for receivers without a provided rating (while increasing the outcome for those with a rating).

Result 2 (Receiver Welfare). *Decreasing the alignment-of-interest between the sender and receiver can improve receiver welfare, provided the misalignment leads to a compensating increase in rating provision.*

Product Sales and Recurrence. Consumers are obviously not the only parties of interest in our setting, as vendors and producers will be affected by changes in purchasing behavior. However, without further assumptions, we are only able to partially address the desirability of each mechanism to these other participants. The costs incurred in incentivizing ratings relative to any additional revenues through increased sales are not identified, and so we will consider the benefit to marketplaces (and producers) through the total (same product) sales rate.

Conditional on a provided rating, we find that the *Receiver* institution produces the most sales (significantly more than *Baseline* and *Producer*, though only marginally significant against *Vendor*), where theory indicates *Vendor* should drive more sales given a rating. Unconditionally, the greatest sales rate is in the *Vendor* treatment—as predicted by the theory, though the primary driver here is the greater fraction of viewed ratings—though this difference is not significant. While we do observe that the institutions with the largest sales volumes are indeed *Vendor* and *Receiver*, the size of the increase over *Baseline* and *Producer* is smaller than predicted.

One potential reason for the small sales effect is the specific reservation distribution H chosen for the experiment. Receivers are predicted to choose the outside option only a quarter of the time in uninformative settings, with this number decreasing as the setting becomes more informative. As a result, the possibility for large sales differences is somewhat limited. In section B we examine counterfactual behavior across a family of reservation distributions using the elicited WTB cutoffs

in the experiment, where we show that the *Vendor* environment is more effective at generating new sales as the average reservation increases.

Result 3 (Total Sales). *Conditional on a viewed rating, the Receiver environment generates the most sales. However, greater provision and viewing in Vendor lead to similar unconditional sales rates.*

Result 3 gives us a coarse description of the desirability of each institutions from the point of view of the vendor player—frequently the party with explicit control over the rating system’s design. The results indicates that the *Receiver* and *Vendor* are most effective at inducing new consumers to purchase a new product rather than the outside option.

Where Result 3 was about the purchase of *any* product, other parties might be concerned with the extent to which ratings induce receivers to purchase the same product as the sender. Here our experimental data indicates the polar opposite of the equilibrium predictions. Conditional on a viewed rating, the *Producer* treatment has the greatest degree of product recurrence, where we should expect to see no difference from random uniform choice across the two competing products. In contrast, the *Baseline* and *Receiver* treatments experience the lowest same-product sales rates, where theory indicates they should have the most.³⁷ The success at generating increased product sales suggests that *Producer*-like incentives can be profitable to specific manufacturers. Conditional on a viewed rating, the *Producer* environment is far and away the best at generating repeat purchases. However, unconditionally the *Vendor* and *Producer* treatments are more comparable, with each generating significantly more same-product purchases than the other institutions. Later we will show that though *Producer* incentives are very effective at capturing market share for the specific product, they are ineffective at growing the overall market size through increased consumer WTB. In comparison, the *Vendor* incentive does moderately well at both of these goals, generating comparable product recurrence overall.

Result 4 (Repeat Sales). *Product-specific incentives to provide ratings lead to a significant increase in the rate at which subsequent consumers purchase the same product as the rater, and can be profitably offered by producers.*

4.2. Sender Behavior. Taking the broad economic results above as a starting point, we now move to more-detailed analysis of behavior for each component choice. Given the sequential nature of the experiment we start with senders, analyzing and summarizing their behavior in the same order as they make their choices. We first examine the effect that each incentive has on senders’ selected ratings given their observed product-quality signals. We then address how the cost c^s and the incentive influence rating provision. Finally, we detail a composite result for senders by calculating the information transfer possible given their response.

Sender behavior is summarized in a series of histograms in Figure 1, with each column corresponding to one of the four treatments. The first five rows provide histograms for the signals drawn by the senders that selected each specific rating one to five, where we bin the draws x_{it} into five equal width intervals: $[1, 20]$, $[21, 40]$, \dots , $[81, 100]$. For each *selected* rating value J , the height of the each white bar indicates the fraction of all rounds with a signal x_{it} in the bin (so the selected rating was $\hat{m}_{it} = J_Y$). The gray bars indicate the fraction of rounds where the selected rating was

³⁷The same results hold if we look at recurrence unconditionally, where we observe the order $P \succ^{**} V \succ^{\circ} B \succ^{*} R$, ranging from 42.7 percent sales in *Producer* to 38.4 percent sales in *Receiver*.

ultimately sent (so the final selection was $m_{it} = J \neq m_0$).³⁸ The white shaded regions therefore indicate selected ratings that were ultimately not sent after observing the send cost c_{it}^S .

Observed Behavior (Rating Selection). *The most-common rating selections are characterized by the division of the signal space into five equal width regions.*³⁹

The data indicates a common agreement over the “honest” meaning of each rating one to five in the *Baseline* and *Receiver* treatments. The modal coordination over ratings’ meanings are through the division of the signal space into five equal regions. Moreover, for the *Vendor* and *Producer* treatments, where senders have misaligned incentives, the figure still indicates substantial honesty in contrast to the equilibrium predictions.

The bottom two rows in figure 1 illustrate the relative frequency and the empirical content of each provided rating. The penultimate row indicates the empirical likelihood of each rating, conditional on provision, $\hat{\Pr}\{m_{it}|m_{it} \in \mathcal{M}\}$. The height of each bar therefore indicates the rate at which a receiver viewing ratings encounters each particular rating value. In both the *Baseline* and *Receiver* treatments these distributions are close to uniform, so that each rating is equally likely. However, in the *Vendor* and *Producer* treatments we see sharp increases in the frequency of fives, where the proportion of five ratings has the empirical order $P \succ^{***} V \succ^{***} B \sim R$.

To demonstrate that this increase in positive ratings is at least partly due to distortions in honest response, and not just differential provision rates for differing signals, the last row in Figure 1 indicates the empirical content of each rating. Given a rating of J_Y , the figure indicates the expected outcome for a receiver choosing urn Y , the conditional quality $\hat{v}_Y(J)$.⁴⁰ Despite the higher-than-predicted honesty, the expected quality of urns rated as a five drops significantly with sender-receiver misalignment.⁴¹ Moreover, in the *Producer* treatment we see a flatter overall profile, where there are smaller differences in expected quality between urns rated three, four and five.

Incentives increase provision. The differing incentive structures heavily influence subjects’ decisions to provide a rating. The *Provision* row of Table 2 indicates that the rate at which ratings are sent varies significantly by treatment, with ratings provided just over one-third of the time in the *Baseline*, just under half the time in *Producer*, and almost two-thirds of the time in the *Receiver* and *Vendor* treatments. Offering incentives to rate products increases the quantity of ratings, with all three incentive schemes generating significantly more rating provision than the *Baseline*. However, in comparison to the best-case theory, the *Baseline* is the only treatment with an informative equilibrium where the rating volume is similar to the equilibrium level. Far fewer ratings are provided in the *Vendor* and *Receiver* treatments than predicted by the best-case theory.

³⁸The rated urn Y need not be the urn X the sender drew from, where we discuss this in the next sub-section.

³⁹Given a rating J , the modal region for the sender’s draw is the x is $[1, 20] + 20 * (J - 1)$, in all treatments. This is significant at the one-percent level. From 2,500 bootstrapped histograms for each treatment, none falsifies the result for the chosen rating (jointly). For sent ratings a single bootstrapped histogram in *Producer* does not satisfy the result.

⁴⁰The empirical value is calculated as $\hat{\Pr}\{\text{Rate Urn } X | m = J\} \cdot (3/4 \cdot \hat{\mathbb{E}}(x | m = J_X) + 1/4 \cdot \hat{\theta}) + \hat{\Pr}\{\text{Rate Urn } Y \neq X | m = J\} \cdot \hat{\theta}$, taking into account both the fraction of senders rating the other urn ($X \neq Y$), and the likelihood that an unrelated ball is drawn. N.B. given a rating J_Y , there are slight changes in the expected value of a draw from the unrated urn Y' from $\hat{\theta} = 50.5$. The largest difference is in the *Producer* treatment where receipt of a one-rating on urn Y implies an expected quality on the unrated urn of 48.1 (significantly lower than 50.5 at the 99 percent level).

⁴¹The expected quality of the rated urn given a five has the order $R \succ^{***} V \succ^{***} P$, with $R \overset{*}{\succ} B \overset{\circ}{\succ} V$ due to lower power in the *Baseline* given fewer provided ratings.




FIGURE 1. Sender behavior

TABLE 3. Decision to send a rating.

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

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

Result 5 (Provision and Incentives). *Provision is greater than the Baseline in all three treatments with a provision incentive. Rating frequency is similar in the Receiver and Vendor treatments, but significantly lower in the Producer treatment.*

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 ($1 \{x_{it} \leq 33\}$, $1 \{33 < x_{it} < 67\}$ and $1 \{x_{it} \geq 67\}$), a dummy for a non-positive rating cost ($1 \{c_{it}^S \leq 0\}$), and the actual cost of rating when it is positive ($c_{it}^S \cdot 1 \{c_{it}^S > 0\}$). 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 provision rates and cost cutoffs for signals drawn in each quality region, alongside the predicted provision rates given a negative rating cost, and the marginal effect on provision from increasing the sender's cost from zero.⁴²

Senders in all treatments are sensitive to their signal x_{it} , 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 the *Baseline* going from a low draw ($x_{it} \leq 33$) to a high draw ($x_{it} \geq 67$) increases provision

⁴²In the three sender-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 3).

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

		Baseline	Receiver	Vendor	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.

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 in provision 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 *Vendor* and *Baseline* treatments there are significant provision increases from low draws to average, and from average draws to high, leading to a stronger selection effect, which we quantify below.

In absolute terms, the cutoff regressions indicate the effects from sender incentives on provision. The average cost cutoff increases from approximately \$0.10 in the *Baseline* to just over \$0.70 in *Receiver*. However, this represents an under-reaction relative to the size of the incentive, where the average incentive payment to the sender is \$1.48. Similarly, the average subsidy from providing a rating in the *Vendor* treatment is \$1.41, while the average cutoff specified is \$0.77.

For the probit regressions, the dummy for negative provision costs indicates almost *all* subjects send a rating if their payoff increases from doing so. The last row in Table 3 indicates the response to positive costs in the probits, the marginal effect on provision from a one dollar increase in the cost of rating. Here the effect is negative and quantitatively large across all four treatments.

Observed Behavior (Provision). *Senders in the incentivized environments provide ratings more often when their signal is good, and are highly responsive to the cost c^S . Senders under-provide ratings in the incentivized treatments relative to the average incentive payment they obtain from providing a rating.*

Incentives distort information provided. We now examine the quality and overall efficacy of ratings, detailing the prevalence of dishonest ratings and quantifying the effects from sender behavior on a sequentially rational receiver.

We say that a rating is dishonest if it satisfies any of three conditions: i) *Overstatements*, a positive rating (a four or five) on the drawn urn X for a low-quality signal ($x \leq 50$); ii) *Understatements*, a negative rating (a one or two) on urn X for a high-quality signal ($x > 50$); or iii) *Other Urn*, any rating where the rated urn Y is not the urn X sampled by the sender. We then measure the extent to which the three incentivized treatments have different rates of dishonesty when compared to the *Baseline*. Table 4 summarizes the dishonesty over both the initially selected

ratings and those that were ultimately sent. We focus on the last 15 rounds where subjects have already had experience with the rating environment, so that we can be more confident that observed behavior is dishonesty rather than simply initial mis-coordination on the meaning of each rating (though results are similar on the entire sample).

While incentives to overstate an urn’s quality exist in both the *Vendor* 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 *Vendor* 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 receivers WTB with an overstatement on the incentivized product—which we will later demonstrate 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.⁴³

Observed Behavior (Rating Dishonesty). *Sender dishonesty increases significantly when incentives are misaligned. Greater misalignment leads to greater dishonesty, while having receivers pay senders for provided ratings leads to greater honesty than the Baseline. The most common form of dishonesty is an overstatement of quality.*

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.⁴⁴ Across all four treatments, the majority of provided information is honest—greater than three quarters of ratings in all treatments. In contrast the theoretical (and empirical) best response for senders in the *Vendor* treatment is to select a rating for the unsampled urn half the time, while for *Producer* the empirical best response is to always rate the urn X as a five (overstatements half the time).

Efficiency given sender behavior. Given the behavior documented, it is natural to ask how much information in total is being transmitted in each setting. What is the change in informational content due to dishonesty? Given the greater provision of ratings following good signals, how informative is it simply to know that a rating was provided? Given the provision rates and rating quality, what is the attainable efficiency? To answer these questions we construct the empirical best response for receivers, given consistent beliefs for the average product outcome $\hat{\nu}_Z(\mathcal{I})$, at all possible information sets \mathcal{I} .

The first row in Table 5 indicates the total efficiency possible given sender behavior.⁴⁵ We then break out the sample into the efficiency possible without a rating ($m = m_\emptyset$) and with a rating ($m \in \mathcal{M}$). Finally, the last two rows hold constant provision of a rating, but examine the receiver welfare from viewing the rating (subtracting the cost c^R and any transferred amount t from viewing) or choosing not to view (incurring no cost and best responding to the information

⁴³Related to this type of rating dishonesty, Italian antitrust authorities recently 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.”

⁴⁴See Gneezy (2005) for a simple experimental identification of this, and Kartik (2009) for a theoretical treatment.

⁴⁵The expected efficiency given sender’s behavior for a best-responding receiver is calculated as $\hat{\Upsilon}_S = \frac{1}{\bar{W} - W} \left[\left(\sum_{m \in \mathcal{M} \cup \{m_\emptyset\}} \hat{\Pr}\{m\} \cdot \int \max\{\hat{\nu}_A(m), \hat{\nu}_B(m), \omega\} dH(\omega) \right) - W \right]$, where the best-response action is simply to choose the option with the highest quality product 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.

TABLE 5. Possible Efficiency Given Sender Behavior

Variable	Baseline Receiver Vendor Producer			
Info. Efficiency	32.2%	54.1%	48.5%	21.3%
Info. Efficiency $ m \in \mathcal{M}$	91.0%	96.1%	88.2%	62.2%
Info. Efficiency $ m = m_0$	-2.2%	-10.1%	-20.2%	-17.8%
Info. Efficiency $ m \in \mathcal{M}, \rho = \text{Not}$	9.2%	5.8%	9.0%	13.3%
Rec. Welfare $ m \in \mathcal{M}, \rho = \text{View}$	84.7%	27.1%	82.0%	56.0%

Note: Information efficiency possible, $\hat{\Upsilon}_S$, is calculated where receivers use the risk-neutral empirical best response with accurate beliefs on the empirical expectation $\hat{\nu}_Y(\mathcal{I})$ for each urn Y , given their information set \mathcal{I} .

set $\{m_{it} \in \mathcal{M}\}$ rather than $\{m_{it} = J_Y\}$). A comparison of the bottom two rows therefore indicates whether a best-responding receiver should view the rating or not. The table supports the following result:

Result 6 (Rating Efficiency). *Sender behavior is more informative than babbling in all four treatments. In all treatments, the benefits from viewing the rating and reacting optimally exceed the acquisition costs.*

When a rating is provided, the attainable efficiency levels are very high—over 90 percent of that attainable if receivers perfectly viewed the draw x_{it} in *Baseline* and *Receiver*, and just under that in *Vendor*. Though there is a substantial drop off in the *Producer* treatment, the total information content is still substantial.⁴⁶ Conditional on provision, the efficiency levels possible in all treatments except *Producer* are quantitatively close to the best-case levels predicted by theory.

Sender Summary. Introducing incentives for ratings significantly influences both the quantity and quality of ratings. Senders respond to incentives by providing ratings more frequently, especially after observing better signals. The *Producer* treatment, and to a quantitatively lesser extent the *Vendor* treatment, exhibit significant levels of dishonesty, most often through overstatements, providing a good rating to a low quality product. However, the observed levels of dishonesty are small relative to theory. Subjects in all four treatments provide meaningful information through ratings, and in all treatments except *Producer* the primary reduction in total efficiency stems from a lower quantity of ratings than predicted in the best-case theory.

4.3. Receiver Behavior. Having documented sender behavior, we now examine receiver choices, following the sequence of their decisions in the experiment; i) the decision on whether or not to view a provided rating, ρ_{jt} ; ii) the decision on which urn to pick a ball from, \hat{Z}_{jt} ; and iii) the final decision Z_{jt} , between the selected urn \hat{Z}_{jt} and the outside option, ω_{jt} .

Viewing Ratings. In all rounds where the sender provides a rating, the receiver’s first decision is whether or not to pay the cost to view it. In three of the treatments viewing incurs a cost $c_R = \$0.05$, while in the *Receiver* treatment there is the additional cost of the incentive payment $t_\psi = \$0.50$, so

⁴⁶The 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.

the effective viewing cost is \$0.55.⁴⁷ 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 2.⁴⁸

Observed Behavior (Viewing). *Conditional on provision, 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 Vendor treatment, and less likely to be viewed in the Producer treatment.*

Though we are not able to precisely isolate this given our data, our interpretation of the viewing behavior is that reciprocity helps receivers coordinate on informative outcomes and viewing ratings. As consumers act as both senders and receivers—which is a natural feature, given our focus on information exchange among peers—the requirement that the receiver views the rating for the sender to obtain the transfer in *Receiver* and *Vendor* helps subjects coordinate on viewing. Such an explanation helps accounts for the increased viewing in *Vendor* over the *Baseline*. Though we see no substantial increase in viewing for *Receiver*, we also see no drop relative to the *Baseline*, despite the large additional cost for acquiring the rating. As such, viewing in *Receiver* can be interpreted as two effects canceling one another out: an increase in viewing due to reciprocity, and a compensating decrease in viewing from greater acquisition costs. Encouraging viewing through the incentive’s structure might therefore help increase the coordination on informative equilibria whenever participants are active on both sides of the information transfer.

Product Choice. After any viewing decision, the receiver’s next task is to select one of the two urns to draw a ball from, representing which of the two products they will choose if they do not select the outside option. Receivers at this point will differ in their information sets, either knowing the specific rating J sent for one of the two urns, $\{m_{it} = J_Y\}$; knowing that a rating was provided, but they chose not to view it, $\{m_{it} \in \mathcal{M}\}$; or knowing that no rating was sent, $\{m_{it} = m_\emptyset\}$. Accounting for symmetry leads to seven distinct information sets across each of the four treatments—the five rating values, not viewing the rating, and no provision.

Looking only at those receivers who are provided with a rating, we use random-effect probits to estimate the likelihood that the product decision \hat{Z}_{jt} is the rated urn, Y . Table 6(A) indicates the estimated probability that the receiver selects the rated urn, where receivers overwhelmingly choose urns rated four or five, and avoid urns rated one or two. The intermediate rating of three leads to just over a third of subjects choosing the rated urn. For those who do not view the rating, the fraction choosing the rated urn is not significantly different from $\frac{1}{2}$.⁴⁹

Observed Behavior (Product Choice). *Ratings are persuasive on the choice of product, with a similar effect across all four treatments: Receivers are very likely to select the rated urn after receiving a four or five, and very likely to select the unrated urn given a rating of one and two.*

While Table 6(A) indicates which product becomes the status quo given a rating, Table 6(B) completes the picture by indicating the WTB for the selected product. Here we look at the final choice Z_{jt} , between the selected product \hat{Z}_{jt} and the outside option ω_{jt} . Similar to our analysis of the sender’s provision decision, we break our estimation into two parts, one estimation over

⁴⁷There is no explicit mention of a transfer in the experimental instructions, the receivers in the *Receiver* treatment face a viewing cost of \$0.55 instead of \$0.05.

⁴⁸The vast majority of differences in aggregate viewing behavior are attributed to subject-level random effects, with many subjects choosing either to never or always view in the last 15 rounds.

⁴⁹However, this is mechanical, as subjects who do not view the rating do not know which urn was rated, and the interface randomly locates the two urns on their screen, so there is no possibility of coordination on urn location

TABLE 6. Receiver behavior.

(A) Probability of choosing rated product

	Baseline	Receiver	Vendor	Producer
$m=1$ or 2	0.059 *** (0.019)	0.014 *** (0.008)	0.083 *** (0.018)	0.079 *** (0.025)
$m=3$	0.413 (0.057)	0.362 (0.049)	0.378 *** (0.041)	0.338 *** (0.059)
$m=4$ or 5	0.973 *** (0.014)	0.981 *** (0.009)	0.968 *** (0.009)	0.981 *** (0.08)
$m \in \mathcal{M}, \rho = \text{Not}$	0.504 (0.016)	0.492 (0.023)	0.491 (0.022)	0.493 (0.017)

(B) Choice of selected product versus outside option

	Baseline		Receiver		Vendor		Producer	
	Sales	WTB	Sales	WTB	Sales	WTB	Sales	WTB
$m=1$ or 2	0.756 (0.036)	49.68 (1.96)	0.745 (0.033)	49.86 (1.93)	0.764 (0.029)	53.73 (2.18)	0.801 (0.039)	52.98 (2.54)
$m=3$	0.691 (0.055)	51.82 (2.34)	0.808 (0.040)	52.38 (2.24)	0.838 *** (0.032)	50.51 (2.47)	0.724 (0.057)	46.83 (3.05)
$m=4$ or 5	0.856 *** (0.030)	59.24 *** (1.99)	0.895 *** (0.021)	61.31 *** (1.85)	0.805 ** (0.022)	58.26 *** (2.01)	0.802 ** (0.026)	52.04 (2.28)
$m \in \mathcal{M}, \rho = \text{Not}$	0.646 ** (0.043)	49.62 (2.12)	0.762 (0.034)	48.93 (2.13)	0.656 * (0.055)	50.97 (3.46)	0.696 (0.036)	47.56 (2.49)
$m = m_\emptyset$	0.762 (0.016)	51.20 (1.47)	0.721 (0.022)	49.06 (1.65)	0.773 (0.021)	53.75 * (1.97)	0.765 (0.019)	49.47 (2.11)

Note: For Table 6(A) the probability of product choice \hat{Z} being the rated product Y is derived from a random-effect probit estimate. For Table 6(B): Sale probability is the estimate of $\Pr\{Z \neq R | \mathcal{I}\}$, derived from a random-effect probit estimate for each information set \mathcal{I} . WTB 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 50.5 for the reservation cutoffs) at: ***-99 percent confidence; **-95 percent confidence; and *-90 percent confidence.

the entire data, and one over just the last fifteen rounds of the experiment using the elicited WTB cutoffs.

The sale probability columns reflect the rate at which the selected product is chosen over the outside option, which we assess with a random-effects probit on the dependent variable $\mathbf{1}\{Z_{jt} = \hat{Z}_{jt}\}$, where the regressors are dummies for the receiver's information type. In contrast, the cutoff estimates reflect the predicted reservation cutoff, obtained using a random-effects panel estimate with the dependent variable $\bar{\Omega}_{jt}$, where the right-hand side again consists of dummies for the receiver's

information type. 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.

Observed Behavior (WTB). *Receivers in all treatments are more likely to purchase a product after observing a positive rating. This effect is strongest in the Baseline and Receiver treatments where senders and receivers are fully aligned in their interests. Negative ratings do not influence total sales in any treatment, and only affect the receiver’s choice between products.*

Examining the sales rate and WTB for receivers without a rating or those who do not view, the data suggest receivers do not understand the provision selection effects. As discussed earlier, if a rating is provided but not viewed, the selection effect leads to an increase in both products’ expected quality. Conversely, when no rating is provided, the selection effect leads to a decrease in the expected quality of both products in comparison to the babbling level of $\bar{\theta}$. Examining Table 6b, we see no evidence that receivers are aware of these effects—if anything we observe the reverse effect. The results suggest that subjects do not take the provision selection effects into account.

Examining the WTB regressions, which use data from the last half of the experiment only, we observe that the reservation cutoff falls toward $\bar{\theta}$ following a positive rating in *Producer*. That is, receivers in the *Producer* treatment seem to learn not to react to positive ratings with increased WTB. Though receivers eventually ignore favorable *Producer*-based ratings when deciding whether to participate in the market, ratings are still persuasive as the ratings continue to have strong effects on product choice.

Receiver Summary. Provided ratings are mostly viewed in all treatments, though more so in the *Baseline*, *Receiver*, and *Vendor* treatments, and less so in *Producer*. Ratings are persuasive with regards to which product is selected in all treatments; low ratings lead to the unrated product being chosen, while high ratings push receivers toward the rated product. While low ratings are influential in the choice of product, they do not influence the probability of any product purchase. In three of the institutions, high ratings influence both the choice of product and the WTB of the receivers viewing them. In *Producer*, however, where dishonesty is a more common response and the quality of ratings lowest, we see that ratings are only effective at influencing the choice between products, and not WTB. Though the observed levels of dishonesty are low relative to theory, receivers are particularly wary of advice in this treatment.

5. INSTITUTIONAL STABILITY

For any one of our institutions to be focal, they would ideally be resistant to the introduction of other forms of incentive. Suppose the *Receiver* institution is the focal one for a particular society, and consumers commonly pay to acquire others product experiences. Producers (if there were no legal sanctions for doing so) might covertly offer *Producer*-like incentives to previous customers. The calculus for producers here is the extent to which their introducing such incentives will profit them in the short and long run. From our experiment we have shown that in the long-run—where there is common-knowledge of the precise incentives on offer—that there is some support for *Vendor*-like commission incentives. Receivers, marketplaces and specific manufacturers all obtain long-run benefits from this institution.

Some of our comparative-static predictions are tied to the particular parameter choices used in our design, and it may be that the *Receiver* institution is better in other parametrizations. However, given observations from our experiment, and our theoretical framework, it is possible to combine

TABLE 7. Short-run effects

		(A) Total Sales				(B) Same-Product Sales				
		Senders				Senders				
		B	R	V	P	B	R	V	P	
Receivers	B	–	0.001	0.003	0.005	B	–	-0.007	0.017	0.043
	R	-0.008	–	0.004	0.001	R	0.000	–	0.030	0.056
	V	0.012	0.002	–	0.008	V	-0.020	-0.024	–	0.024
	P	-0.001	0.006	0.004	–	P	-0.048	-0.042	-0.018	–

Note: Values show the change in the frequency of sales when receivers from the row type face ratings sent by senders in the column type.

the two to examine potential counterfactuals. While it is impractical to characterize behavior for all possible parametrizations, in the appendix we consider counterfactual long-run effects from varying a key parameter in our model, the reservation distribution $H(\omega)$. We show that across a very broad range of possible distributions for receiver’s outside options, that our findings extend, that the *Vendor* treatment remains superior to the other three in the long run. This finding is across receiver welfare, total sales and recurrent sales—meaning that many stakeholders will be happy with overt sales-commissions for providing information.

A separate counterfactual exercise allows us to examine short-run incentives, answering the following: To what extent might vendors and producers benefit from *covertly* introducing mis-aligned incentives to those rating products? To examine this question we leverage the rich choice data collected in our experiment, and combine the rating selection and provision from one treatment with the viewing and final choice behavior from another. That is, fixing the receiver response to ratings (their viewing, product and WTB decisions at all information sets) in environment Ψ , we assess whether there might be short-run sales gains to either marketplaces or specific producers by offering senders the incentives from environment Ψ' .

Given the receiver’s perceived environment, Table 7 indicates the change to both total and same-product sales when the sender’s environment is changed. We again use data from the last ten rounds of the experiment to assess the sent rating distributions, viewing rates, and the relevant sales rates conditional on each viewed rating. The effect on *total sales* is small across each comparison, and we therefore conclude that short-run considerations by a vendor at the experimental parametrization are unlikely to trump the long-run considerations. However, we see much stronger effects when examining repeat sales. Offering senders a *Producer* incentive while receivers believe they are in the *Baseline* or *Receiver* environment leads to substantial sales increases of 12 and 16 percent, respectively. These figures would likely be even larger if senders knew that receivers were unaware of the true incentives at play.

The introduction of *Producer*-type incentives into the *Vendor* environment generates a much smaller benefit, as receivers are already more suspicious of positive ratings, though the overall effect is still a 6 percent increase in sales volume. Combined with the results from our counterfactual in the appendix, the long-run generalization, we conclude that *Vendor* is the more-stable informative institution.⁵⁰ Long-run profits to specific manufacturers are reduced if they introduce

⁵⁰Where *Producer* is the most-stable of the four, but is also the most inefficient.

product-specific incentives to a *Vendor* environment, which may serve to offset the short-run gains. In contrast, for the *Baseline* and *Receiver*, long-run producer profits are increased through the introduction of product-specific rating incentives, and there are also much larger short-run gains from their covert introduction.

6. CONCLUSION

Effective institutions for facilitating information exchange need to focus not only on the quantity of information provided, but also the quality. The obvious problem in many information-sharing environments is one of underprovision, in which not enough information is shared. To increase the flow of information, those with relevant experience can be incentivized to provide information to the marketplace, where natural sources of funding for these incentives are: i) the *Receiver* making use of the information; ii) the *Vendors* and marketplaces in which competing products are sold; or iii) the specific *Producers* that manufacture each product. However, the interests of each funding source and the manner in which they might construct these incentives differ. Our paper examines three alternative incentives (and an unincentivized *Baseline*) within a rich and readily extensible theoretical framework.

In our *Receiver* environment we require consumers who view ratings to pay those generating them. This leads to a large quantity of high-quality ratings, and is the most-preferred scheme for senders. However, this environment is not the optimal arrangement for the consumers making use of the provided information, as the burden from paying for ratings can erode the benefits from increased information. Most preferred by receivers is our *Vendor* incentive, which shifts the costs of paying for ratings to an interested third party (the vendor) that pays senders whenever they generate a sale, a commission system. This type of sales-contingent incentive corrodes the quality of provided information (both theoretically and in our data), as it drives a wedge between those communicating. However, the misalignment in interests is not complete, and some information can still credibly be sent as to which product is better. Final consumers, the marketplaces, and producers are all potentially better off here thanks to a tradeoff between quality and quantity.

Our last institution demonstrates that too much of a tradeoff in alignment leads to sub-optimal results. When senders are provided with a *Producer* incentive to generate sales of a specific product, the wedge between the sender and receiver of information is complete. Neither information on which of the competing products is better, nor when to purchase any product at all can be credibly exchanged in theory. In contrast, our experimental data suggests a less extreme outcome, though it is the least efficient of our four environments. However, this incentive scheme is the most effective of the four at increasing product sales. Our results indicate large gains to the producers of specific products from introducing this type of misaligned incentive for ratings, particularly in the short run.

Policymakers or platforms that host and control rating aggregates should be wary of distortive incentives being covertly offered. However, our data points to another benefit for commission incentives in our *Vendor* institution: by incorporating partial misalignments in interest, this institution is the most inoculated against *Producer*-type incentives. Further research will hopefully examine other important effects outside of the present paper—in particular reputation and repeated interaction. Nonetheless, that we show both theoretical and empirical support for an incentive that drives a wedge between the two parties exchanging information in a one-shot setting shows why it might be such a common mechanism in the field.

REFERENCES

- BATTAGLINI, M., AND U. MAKAROV (2014): “Cheap talk with multiple audiences: An experimental analysis,” *Games and Economic Behavior*, 83, 147–164.
- BOLTON, G., B. GREINER, AND A. OCKENFELS (2013): “Engineering trust: reciprocity in the production of reputation information,” *Management Science*, 59(2), 265–285.
- BRANDTS, J., AND G. CHARNNESS (2011): “The strategy versus the direct-response method: a first survey of experimental comparisons,” *Experimental Economics*, 14(3), 375–398.
- CAI, H., AND J. T.-Y. WANG (2006): “Overcommunication in strategic information transmission games,” *Games and Economic Behavior*, 56(1), 7–36.
- CALVÓ-ARMENGOL, A., J. DE MARTÍ, AND A. PRAT (forthcoming): “Communication and influence,” *Theoretical Economics*.
- CHAKRABORTY, A., AND R. HARBAUGH (2010): “Persuasion by Cheap Talk,” *American Economic Review*, 100(5), pp. 2361–2382.
- CHARNNESS, G., AND N. GAROUPA (2000): “Reputation, honesty, and efficiency with insider information: An experiment,” *Journal of Economics & Management Strategy*, 9(3), 425–451.
- CHEN, Y., M. HARPER, J. KONSTAN, AND S. X. LI (2010): “Social Comparisons and Contributions to Online Communities: A Field Experiment on MovieLens,” *American Economic Review*, 100(4), pp. 1358–1398.
- CHUNG, W., AND R. HARBAUGH (2012): “Biased Recommendations,” mimeo, Indiana University.
- CRAWFORD, V. P., AND J. SOBEL (1982): “Strategic Information Transmission,” *Econometrica*, 50(6), 1431–1451.
- DESSEIN, W., AND T. SANTOS (2006): “Adaptive organizations,” *Journal of Political Economy*, 114(5), 956–995.
- DEWATRIPONT, M., AND J. TIROLE (2005): “Modes of Communication,” *Journal of Political Economy*, 113(6), 1217–1238.
- DICKHAUT, J. W., K. A. MCCABE, AND A. MUKHERJI (1995): “An experimental study of strategic information transmission,” *Economic Theory*, 6, 389–403, 10.1007/BF01211783.
- GNEEZY, U. (2005): “Deception: The Role of Consequences,” *American Economic Review*, 95(1), 384–394.
- HOLT, C. A., AND S. K. LAURY (2002): “Risk aversion and incentive effects,” *American Economic Review*, 92(5), 1644–1655.
- HU, N., J. ZHANG, AND P. A. PAVLOU (2009): “Overcoming the J-shaped distribution of product reviews,” *Communications of the ACM*, 52(10), 144–147.
- KARTIK, N. (2009): “Strategic communication with lying costs,” *The Review of Economic Studies*, 76(4), 1359–1395.
- LAFKY, J. (2014): “Why do people rate? Theory and evidence on online ratings,” *Games and Economic Behavior*.
- LAI, E. K., W. LIM, AND J. T.-Y. WANG (2011): “Experimental Implementations and Robustness of Fully Revealing Equilibria in Multidimensional Cheap Talk,” mimeo, National Taiwan University.
- LI, X., AND L. M. HITT (2008): “Self-selection and information role of online product reviews,” *Information Systems Research*, 19(4), 456–474.
- VESPA, E., AND A. J. WILSON (2014): “Communication with Multiple Senders: An Experiment,” mimeo, University of Pittsburgh.

- WANG, J., M. SPEZIO, AND C. F. CAMERER (2010): “Pinocchio’s Pupil: Using Eyetracking and Pupil Dilation to Understand Truth Telling and Deception in Sender-Receiver Games,” *American Economic Review*, 100(3), 984–1007.
- WANG, Z. (2010): “Anonymity, social image, and the competition for volunteers: a case study of the online market for reviews,” *BE Journal of Economic Analysis & Policy*, 10(1).
- WILSON, A. J. (2014): “Costly Communication in Groups: Theory and an Experiment,” University of Pittsburgh working paper.

APPENDIX A. FOR ONLINE PUBLICATION: THEORY

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

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

- (i) If $\mu^*(X, x, c^S) = M_Y$ for $X, Y \in \{A, B\}$ then for $X' \neq X$ and $Y' \neq Y$, $\mu^*(X', x, c^S) = M_{Y'}$.
- (ii) If for $X \in \{A, B\}$ the provision strategy is $\mu^*(X, x, c^S) = m_\emptyset$ then it must be that for $X' \neq X$ that the strategy satisfies $\mu^*(X', x, c^S) = m_\emptyset$.
- (iii) For any receiver information set \mathcal{I} , whenever $\nu_A(\mathcal{I}) = \nu_B(\mathcal{I}) > \omega$, R 's final choice is a $(\frac{1}{2}, \frac{1}{2})$ mixture over $\{A, B\}$. When $\omega = \max\{\nu_A(\mathcal{I}), \nu_B(\mathcal{I})\}$, the receiver chooses R .

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

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

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

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

Given the assumption that $\alpha > 0$, senders and receivers fully agree on the choice $Z \in \{A, B, R\}$ at all possible information sets. Conditional on providing a rating $m \neq m_\emptyset$, senders and receivers have a common interest in revealing information. In fact, full-revelation of all decision-relevant information can be communicated even when the number of available ratings is lower

than the number of possible signals ($K < N$). This is because whenever S 's draw (X, x) implies that product X' is strictly better than product X , then all that needs to be conveyed is the ordinal information “do not buy X .” Any space \mathcal{M} that is rich enough to distinguish between each of the signals that imply $\nu_X(x) > \bar{\theta}$, plus a single additional rating implying $\nu_X(x) \leq \bar{\theta}$, conveys all the necessary information.⁵¹

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

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

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

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

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

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

respectively. When α is small, these two decisions will be close to the symmetric decisions in a babbling equilibrium, where $\nu_\emptyset \simeq \nu_{\mathcal{M}} \simeq \bar{\theta}$, though they must adapt somewhat to selection effects in provision as the cutoff c^* is signal-dependent.⁵²

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

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

⁵¹Where there are multiple senders, ratings which delineate between differing below average levels will aid convergence.

⁵²Numerically solving for an equilibrium involves searching for the fixed point vector $(\bar{\nu}_1, \dots, \bar{\nu}_K, \nu_\emptyset)$ that satisfies the sender ICC and cutoff calculations below, where $\nu_{\mathcal{M}}$ can be calculated subsequently. In the limit, as $\alpha \rightarrow 0$, the provision decision becomes independent of the signal, and a fixed-point calculation can be avoided.

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

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

where

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

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

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

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

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

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

The first-order change here is to increasing the number of meaningful messages provided. Given the shift in the cutoffs, there are second-order effects from the shape of $G(\cdot)$ that require recalculation of the expected qualities $(\bar{\nu}_1, \dots, \bar{\nu}_K, \nu_\emptyset, \nu_M)$ and rating cutoffs.⁵³

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

⁵³ In our experimental setting, these effects are only non-negligible for ν_\emptyset , due to a very small likelihood of no provision.

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

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

where

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

The WTB given a provided message is calculated as

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

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

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

APPENDIX B. RESERVATION DISTRIBUTION SHIFTS

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

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

Figures 2(A,C,E,G) indicate theoretical differences (in the most-informative equilibrium) relative to the *Baseline* environment over efficiency, receiver welfare, total sales and same-product sales, respectively. For instance, in Figure 2(A) we indicate the efficiency difference $\Upsilon(\sigma; \psi) - \Upsilon(\sigma; \text{Base.})$, and in Figure 2(C) the receiver welfare difference $\Upsilon_R(\sigma; \psi) - \Upsilon_R(\sigma; \text{Base.})$.⁵⁴ In each figure, the dashed vertical line at $\sigma = 100/3$ indicates the experimental parametrization.

The theoretical comparisons here assume the most-informative equilibrium, however we know that subjects are more honest than predicted, and also deviate from theory in their rating, viewing and WTB behavior. We now extrapolate from the observed experimental behavior across the family of WTB distributions using the elicited reservation cutoffs in the last half of each experimental session. To do this, we make some assumptions on how behavior changes as we alter H . First, we assume changes in H do not affect the ratings selected \hat{m}_{it} after each signal x_{it} , or the reservation cutoff $\bar{\Omega}_{jt}$ selected at each information set \mathcal{I}_{jt} .⁵⁵ Second, we incorporate changes to both the provision and viewing rates in each environment through a logit response to the expected benefits/costs at σ . Using the experimental variation in the environment, we estimate subjects' provision and viewing elasticities with respect to changes in the expected monetary incentives and realized monetary benefits of viewing, respectively. Table 8 contains details of the estimation procedure and assessed parameters. We use this estimated logit response to extrapolate how the provision and viewing rates vary as σ changes.⁵⁶

⁵⁴The scale and normalization for efficiency are also functions of σ . That is, the expected upper bound given full information, \bar{W}_σ , and the expected outcome under babbling, \underline{W}_σ , are both affected by the choice of $H(\cdot)$, where $\bar{W}_\sigma - \underline{W}_\sigma$ is decreasing in σ . The discontinuity in the *Receiver* comparison stems from a critical σ after which it is no longer rational to view ratings, so that the babbling equilibrium is the most-informative.

⁵⁵Given the specified reservations we calculate the expected total sales rate as $\text{Sales}(\sigma) = \frac{1}{|\mathcal{S}|} \sum_{j \in \mathcal{S}} \frac{1}{10} \sum_{t=21}^{30} \hat{H}_\sigma(\bar{\Omega}_{jt})$, while the rate of same product sales is calculated as $\text{Product}(\sigma) = \frac{1}{|\mathcal{S}|} \sum_{j \in \mathcal{S}} \frac{1}{10} \sum_{t=21}^{30} \hat{H}_\sigma(\bar{\Omega}_{jt}) \cdot \mathbf{1}\{Z_{jt} = X_{it}\}$. We here focus on the last ten rounds, allowing rounds 16 to 20 for subjects to gain experience with the WTB elicitation.

⁵⁶As reservations increase, receivers viewing behavior decreases as ratings are less useful to those unlikely to purchase a product. The combined effect of reduced viewing and reduced sales also alters the benefits to providing a rating in each of our environments, and so provision falls. The two effects are modeled in the counterfactual through a response to the expected difference in payoff from providing/viewings where $\text{Provision}(\sigma; \psi)$ is therefore a function of both $\text{Viewing}(\sigma)$ and the environment ψ . This exercise is analogous to calculating a logit quantal response equilibrium, where we fix the observed rating and WTB behavior.




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

TABLE 8. Counterfactual Provision/View Response Model

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

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

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

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

Though the sales gains in *Vendor* are relatively small at our parametrization, as outside options increase, so too does the relative sales benefit to the marketplace. Once we raise the expected outside option to $\sigma = \bar{\theta}$, the *Vendor* treatment’s total sales fall to 0.543, where this is significantly larger than either the *Baseline* or *Receiver* treatments at 0.514 and 0.513, respectively. The relative

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

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

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

effect in *Vendor* grows even larger as average reservations increase. For $\sigma = 64$, though the total fraction of the population making a purchase is lower in all treatments, the relative gain in the *Vendor* setting is largest, with a 12 percent increase in sales volume over the *Baseline*. By contrast, the *Receiver* environment produces a negative and insignificant sales effect relative to the *Baseline* for higher values of σ , primarily due to reduced viewing behavior.⁵⁷

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

APPENDIX C. FOR ONLINE PUBLICATION: BEHAVIORAL EFFICIENCY LOSSES

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

⁵⁷Our model for provision and listening decreases with σ at a greater rate in the *Receiver* treatment. In the experiment, the difference in the rate of provision in the last ten rounds between Receiver and Baseline is 0.268, while listening rates are near identical. At $\sigma = 65$ the counterfactual model reduces the provision difference to 0.216 (reflecting a drop in the expected receiver transfer stemming from reduced viewing). Listening rates are modeled as falling by 0.104 in *Baseline*, and 0.157 in *Receiver*.

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

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

In the second counterfactual, *All Viewing*, we force every receiver to view all provided rating, fixing their response to each particular rating to the average observed behavior. There are moderate gains here in the *Receiver* and *Baseline* treatments, however once we account for the cost of viewing (in the receiver welfare row), the potential efficiency gains in *Receiver* actually become a small loss.⁵⁸

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

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

Result 7 (Improved Efficiency). *The largest potential increase in total efficiency from participants best responding to observed play stems from increases to rating provision in the Receiver and Vendor environments.*

⁵⁸Though it is an empirical best response to view a rating in all treatments, it is not so if receivers react sub-optimally to the provided rating.

TABLE 10. Risk-Averse Equilibrium Predictions

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

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

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

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

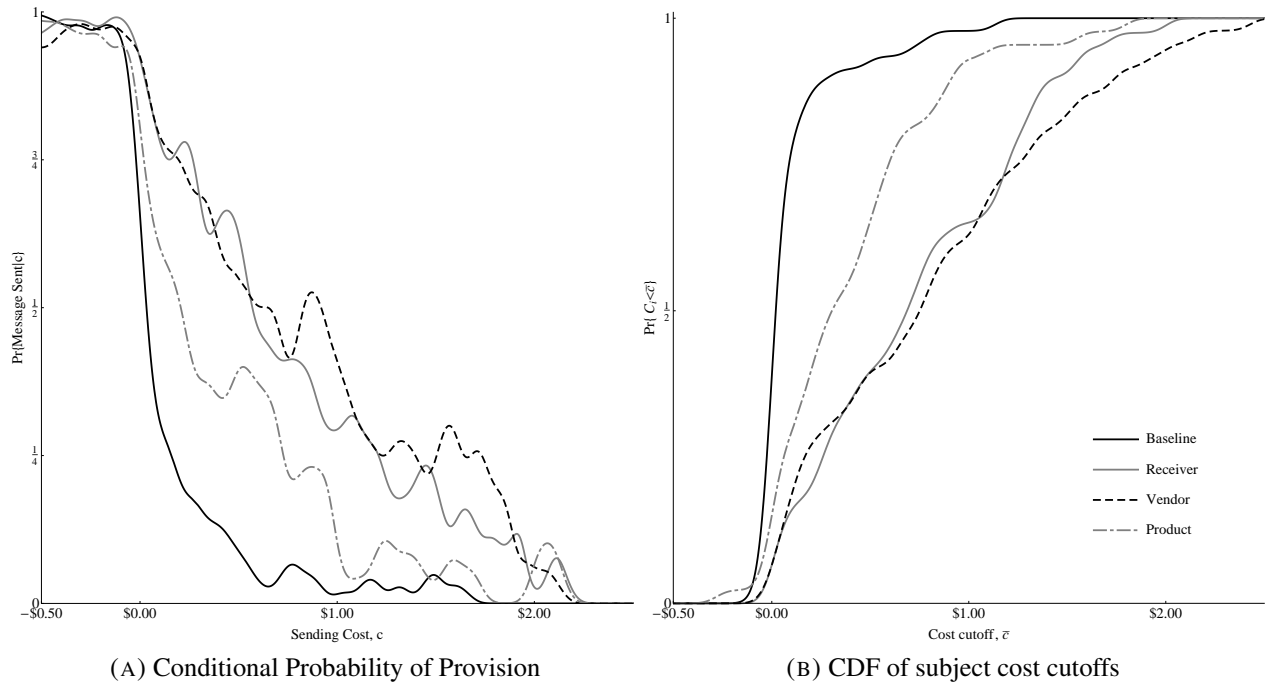


FIGURE 3. Provision as a function of rating cost .

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