Voluntary Disclosure in Bilateral Transactions

Vincent Glode Christian Opp Xingtan Zhang*

November 18, 2016

Abstract

We solve for the ex post verifiable signals that a privately informed agent would optimally send to a counterparty endowed with market power prior to a bilateral transaction. By sharing his information, the agent may reduce his own information rents, but he may also make it less likely that his counterparty screens him and jeopardizes gains to trade. We show that the agent always designs a partial disclosure plan that implements socially efficient trade in equilibrium. Our results have important implications for understanding the conditions under which asymmetric information impedes trade and for regulating information disclosure in general.

*We thank Jules van Binsbergen, Doron Levit, George Mailath, Luke Taylor, and Pavel Zryumov for helpful comments and discussions. All authors are at The Wharton School, University of Pennsylvania, 3620 Locust Walk, Philadelphia, PA 19104 and can be reached by e-mail at vglode@wharton.upenn.edu, opp@wharton.upenn.edu, and xingtan@wharton.upenn.edu, respectively.

1 Introduction

Asymmetric information can harm agents by disrupting efficient trade (e.g., Akerlof 1970, Myerson and Satterthwaite 1983, Glosten and Milgrom 1985). In this paper, we study the incentives of a privately informed agent to share some of his information with a counterparty endowed with market power prior to a bilateral transaction. Recent work by Bergemann, Brooks, and Morris (2015) highlights that information available for price discrimination plays a crucial role in determining the total surplus and its allocation in the classic monopoly pricing problem with private values, raising the important question "what forms of price discrimination will endogenously arise, and for whose benefit." We consider an environment with both private- and common-value uncertainty and analyze one natural channel determining the information a monopolist seller can use for price discrimination — the privately informed buyer can make ex-post verifiable disclosures, as in Grossman (1981), Milgrom (1981), and Shin (2003). This verifiability restriction on disclosure is an important benchmark, as it is not subject to commitment and incentive problems, even in one-shot interactions.[1] If erroneous disclosures can be verified with probability one and are penalized — perhaps by a regulator or courts — the sender optimally designs signals that are always truthful. Verifiability is further relevant in many important economic contexts with hard information, such as supply chains, mergers and acquisitions, and the trading of financial securities.

In this environment, we obtain a surprisingly strong result: the informed agent always designs a partial disclosure plan that yields socially efficient trade in equilibrium. Moreover, this disclosure plan improves both his surplus and that of the counterparty with market power. Whereas possessing superior information allows the informed agent to extract in-

[1]The early literature analyzing these types of "persuasion games" is surveyed by Milgrom (2008). Since these games focus on ex-post verifiable disclosures, they significantly differ from "cheap talk games" popularized by Crawford and Sobel (1982).

formation rents, by disclosing this information to his counterparty reduces the extent to which the agent is being inefficiently screened. We show that the agent is always able and willing to design ex post verifiable signals such that the information rents he forgoes are dominated by the additional gains to trade he extracts by preventing his counterparty from inefficiently screening him. We also characterize the optimal disclosure plan, which generally pools together multiple regions of realized valuations, sometimes allowing for gaps between them, in order to maximize the informed agent's rents and limit the counterparty's incentives to screen. Although disclosing a subset of what he knows benefits the informed agent, disclosing all his information completely eliminates the rents he can extract and is thus privately suboptimal. While we initially consider an environment where the disclosure plan is designed before any uncertainty is realized — as is common in models of Bayesian persuasion (e.g., Rayo and Segal 2010, Kamenica and Gentzkow 2011, Goldstein and Leitner 2015, Ely 2016) — we also consider the case where the disclosure plan is designed after the agent obtains private information. We show that, in line with our results above, partial disclosure leading to socially efficient trade also characterizes all informed-agent-preferred equilibria of this interim disclosure game.

Our results shed light on the conditions under which we should (and should not) observe inefficient trading among asymmetrically informed agents. These results thus have important implications for regulating information disclosure in bilateral transactions with imperfect competition and asymmetric information problems, such as corporate takeovers, real estate transactions, and over-the-counter trading.[2] Once information disclosure is restricted

[2]For empirical evidence that these types of bilateral transactions often feature imperfect competition, see Ambrose, Highfield, and Linneman (2005), Glaeser, Gyourko, and Saks (2005), Boone and Mulherin (2007), King, Osler, and Rime (2012), Atkeson, Eisfeldt, and Weill (2013), Li and Schürhoff (2014), Begenau, Piazzesi, and Schneider (2015), Hendershott et al. (2015), Di Maggio, Kermani, and Song (2016), Li, Taylor, and Wang (2016), and Siriwardane (2016). For empirical evidence that these types of bilateral transactions often involve heterogeneously informed traders, see Eckbo, Giammarino, and Heinkel (1990), Garmaise and Moskowitz (2004), Green, Hollifield, and Schürhoff (2007), Hollifield, Neklyudov, and Spatt (2014), Jiang and Sun (2015), Menkhoff et al. (2016), and Stroebel (2016).

to ex post verifiable signals the regulator does not need to mandate what information traders should disclose nor does he need to produce additional information for uninformed market participants. All the regulator needs to do is to enforce the truthfulness of disclosure by disciplining traders who send signals that ex post (i.e., given the information is revealed after the transaction occurs) prove to violate their own disclosure standards.

Related Literature. Our paper contributes to the theoretical literature that studies optimal information sharing among traders. An important result in that literature goes back to Grossman (1981) and Milgrom (1981) who show that, when disclosures are restricted to be ex post verifiable, an agent may find it optimal to fully reveal his private information to his counterparties. However, unlike in our model the agent making the disclosure decision is not being screened by counterparties with market power — either all traders take the price as given (as in Milgrom 1981) or it is the informed agent who sets the price (as in Grossman 1981). In those environments, an informed seller always benefits from improving his customers' perception of product quality (which is not necessarily the case once counterparties have market power). He then finds it optimal to fully disclose his private information, since any information he withholds is interpreted to be unfavorable (see also Grossman and Hart 1980, Milgrom and Roberts 1986). Shin (2003), Acharya, DeMarzo, and Kremer (2011), and Guttman, Kremer, and Skrzypacz (2014), however, show that in these types of persuasion games full disclosure becomes suboptimal once we allow for uncertainty in the existence of private information. In Admati and Pfleiderer (2000) where disclosure is assumed to be costly, a firm picks a socially optimal disclosure plan when that firm is a monopolist that captures all gains to trade. Unlike in these settings, the information designer in our model is the responder to an ultimatum offer and his private information is thus his only source of profits. As a result, his optimal disclosure plan is partial — despite the existence of his private information being common knowledge — yet we show that it always yields socially efficient trade in equilibrium. Our framework therefore speaks to how voluntary information sharing can solve the classic monopoly pricing problem.

In that sense, our paper is closely related to Bergemann, Brooks, and Morris (2015) who analyze how signals yielding monopolists additional information for price discrimination affect total surplus and its allocation. Bergemann, Brooks, and Morris (2015) show in a setting with private value uncertainty that general information structures (including randomization) exist such that total surplus can be increased to any level less than or equal to the one from efficient trade, and any allocation of the incremental surplus is attainable. Information available for price discrimination thus critically determines efficiency and the allocation of surplus, raising the question of which information structure should arise endogenously. The objective of our paper greatly differs from theirs. We analyze a setting with both private- and common-value uncertainty and show that if information disclosure by the informed agent must be (a) voluntary and (b) ex post verifiable (with randomization not being possible), precise predictions for both total surplus and its allocation obtain: (1) total surplus is unique and equal to the surplus generated by efficient trade (whether the disclosure plan is designed at an ex ante or interim stage), and (2) the surplus allocation is unique and both agents benefit from the optimal disclosure plan (in the case of ex ante disclosure design).

More broadly, our focus on market power also relates our paper to Gal-Or (1985) who models oligopolistic firms that can commit ex ante to sharing noisy signals of their private information about the uncertain demand for their products. Since sharing information increases the correlation of firms' output decisions, thereby lowering their expected profits, the unique symmetric pure-strategy equilibrium is characterized by no information sharing among firms. Lewis and Sappington (1994) model a setting similar to ours and investigate whether an uninformed seller with market power would like to help his prospective buyer(s) acquire private information about the value of the asset (see also Eső and Szentes 2007, who assume that trading occurs through an auction). Under general conditions, the seller in Lewis and Sappington (1994) either wants his prospective buyer(s) to be fully informed or completely uninformed about how much they value the asset. Finally, Roesler and Szentes (2016) solve for the buyer's optimal information acquisition in a setting similar to ours and show that the buyer always finds it optimal to limit his information acquisition and avoid the seller inefficiently screening him (see also Glode, Green, and Lowery 2012).

The next section presents a classic problem of a monopolist that screens a privately informed agent. In Section 3, we study the agent's incentives to commit to sharing some of his private information with the monopolist and how the resulting disclosure plan affects the efficiency of trade. Section 4 shows that our main insights survive when the agent designs his disclosure plan after obtaining private information rather than before. The last section concludes. All proofs are collected in the Appendix.

2 The Inefficiency of Trade

The monopolist seller of an asset (or good) chooses the price he will quote to a prospective buyer (or customer) in a take-it-or-leave-it offer.[3] The seller is uncertain about how much the buyer is willing to pay for the asset and only knows that the buyer's valuation of the asset, which we denote by v, has a cumulative distribution function (CDF) denoted by $F(v)$. This CDF is continuous and differentiable and the probability density function (PDF), denoted by $f(v)$, takes strictly positive values everywhere on the support $[v_L, v_H]$. The function $c(\cdot)$ is assumed to be weakly increasing and continuous. Both agents are risk neutral and the functions $F(\cdot)$ and $c(\cdot)$ are common knowledge.[4]

Whenever the buyer's valuation is greater than the seller's, trade creates a surplus and

[3] The buyer/seller roles could be reversed without affecting our results.

[4] See Hirshleifer (1971), Diamond (1985), and Kurlat and Veldkamp (2015), among many others, for analyses of the costs and benefits of disclosure when traders are risk-averse.

is efficient. However, the seller may find it privately optimal to use his market power and inefficiently screen the informed buyer, thereby jeopardizing the gains to trade $[v - c(v)]$. We assume that whenever indifferent between two strategies, an agent picks the one that maximizes the social surplus in the resulting subgame-perfect Nash equilibrium. The buyer only accepts to pay a price p in exchange for the asset if $v \geq p$; otherwise, the seller must retain the asset, which is worth $c(v) \geq 0$ to him. The seller's expected payoff from quoting a price p is thus given by:

$$\Pi(p) = [1 - F(p)]p + F(p)\mathbb{E}[c(v)|v < p]. \tag{1}$$

When picking a price, the seller considers the trade-off between the probability that a sale occurs and the profit he gets if a sale occurs. The seller's marginal profit of increasing the price p is:

$$\Pi'(p) = [1 - F(p)] - f(p)p + f(p)\mathbb{E}[c(v)|v < p] + F(p)\frac{\partial}{\partial p}\mathbb{E}[c(v)|v < p], \tag{2}$$

which simplifies to:

$$\Pi'(p) = [1 - F(p)] - f(p)[p - c(p)]. \tag{3}$$

The first term on the right-hand side of equation (3) is the seller's expected benefit from collecting a higher price when trade occurs. The second term is the expected cost from reducing the probability of trade and destroying the gains to trade. We impose the following condition on the surplus from trade $[v - c(v)]$:

Assumption 1. *The surplus from trade $[v - c(v)]$ crosses zero at most in one point (from below).*

This condition is implied by any of the following assumptions common in the literature: (i) the seller's valuation for the asset is a constant $c < v_H$; (ii) the surplus from trade

$[v - c(v)]$ is a constant $\Delta > 0$; (iii) the ratio of the above-mentioned cost and benefit of marginally increasing the price, i.e., $H(v) = \frac{f(v)}{1-F(v)}[v - c(v)]$, is strictly increasing in v.[5] Hence, any one of these fairly standard assumptions is sufficient to derive our results.

Assumption 1 implies that there exists a unique cutoff $\hat{v} \in [v_L, v_H]$ for which trade is socially efficient as long as it occurs if and only if $v \geq \hat{v}$. Thus, for trade to be efficient, the seller must find it optimal to quote a price that is accepted by the buyer if and only if $v \geq \hat{v}$. Since $f(v)$ is strictly positive everywhere on the support $[v_L, v_H]$, the maximum price that can maintain efficient trade is $p = \hat{v}$. As a result, trade can be efficient only if:

$$\Pi'(\hat{v}) \leq 0. \tag{4}$$

This necessary condition for efficient trade can be interpreted as follows. Efficient trade requires that $\hat{v} - c(\hat{v}) \geq \frac{1-F(\hat{v})}{f(\hat{v})}$, which means that either the gains to trade are large or that the seller's beliefs about v are concentrated (i.e., the density $f(v)$ is high enough) when the surplus from trade becomes positive. If instead $\Pi'(\hat{v}) > 0$, the seller inefficiently screens the buyer and jeopardizes gains to trade.

Also note that, since we assume that whenever indifferent an agent picks the strategy that maximizes the social surplus from trade, we can rule out any equilibrium where the seller inefficiently mixes between quoting multiple prices $p_n \in [v_L, v_H]$. If he were to mix over many prices, the seller would have to be indifferent between quoting all these prices and quoting only one of these prices (taking into account the buyer's best response to each price). The tie-breaking rule implies that the seller should instead play the pure strategy of quoting the price that socially dominates all other prices. Similarly, we can rule out equilibria where the buyer inefficiently mixes between accepting or not a price quote. A

[5] See, e.g., Glode and Opp (2016) and Glode, Opp, and Zhang (2016) who specifically impose this condition, Fuchs and Skrzypacz (2015) who define a "strictly regular environment" in a similar way and Myerson (1981) who similarly assumes that bidders' virtual valuation functions are strictly increasing.

tie-breaking rule based on social optimality thus ensures that we can restrict our attention to pure-strategy subgame-perfect Nash equilibria in our model.

We further illustrate the seller's incentives to price the asset inefficiently through a simple parameterized example that we will revisit later.

Example 1. *Suppose the buyer values the asset at* $v \sim U[1,2]$ *and the seller values it at a constant* $c \leq 1$*. The surplus from trade is always positive (i.e.,* $\hat{v} = 1$*) and trade is thus efficient if and only if it occurs with probability* 1*. The seller's optimization problem when picking a price can be written as:*

$$\max_{p \in [1,2]} \Pi(p) = Pr(v \geq p)p + Pr(v < p)c = (2-p)p + (p-1)c. \tag{5}$$

When $\Pi'(1) = c \leq 0$*, the seller quotes a price* $p = 1$ *that is always accepted by the buyer and trade is therefore efficient. However, when* $c \in (0,1]$ *the seller finds it optimal to screen the buyer by quoting a price* $p = 1 + \frac{c}{2}$*, which inefficiently destroys the surplus from trade with probability* $\frac{c}{2}$*.*

The example above shows a simple case where $\hat{v} = v_L$, that is, the surplus from trade is positive for any realization of v. In cases like that, efficient trade requires that $v_L - c(v_L) \geq \frac{1}{f(v_L)}$. For cases where $\hat{v} \in (v_L, v_H)$ however, efficient trade can never be sustained in equilibrium since by definition $\hat{v} - c(\hat{v}) = 0 < \frac{1-F(\hat{v})}{f(\hat{v})}$, and the seller always finds it optimal to quote a price that is at least marginally higher than the efficient price $p = \hat{v}$. This situation arises, for example, whenever the seller values the asset at a constant $c \in (v_L, v_H)$.

3 Information Disclosure prior to Trading

In this section, we analyze the buyer's decision to share a subset of his information with the seller before trade occurs. Clearly, if trade is already socially efficient without disclosure, the buyer is only paying $\hat{v}$ for the asset and information disclosure is suboptimal — with additional information the seller would only consider raising the price he quotes the buyer. Thus, for the remainder of the paper we focus on situations where trade would be socially inefficient if the buyer did not disclose any of his private information. Sharing information might hurt the buyer since possessing private information yields informational rents, but it might also reduce the seller's incentives to charge inefficient mark-ups that reduce the expected gains from trade.

For now, we assume that the agent must design his disclosure plan prior to acquiring private information and that he can commit to not manipulating the signal later, as is common in models of Bayesian persuasion. Assuming that the buyer is uninformed at the time of the information design facilitates our analysis as it eliminates the existence of signaling concerns. However, we relax this assumption in Section 4.

We also restrict our attention to ex-post verifiable disclosures or signals, as in Grossman (1981), Milgrom (1981), and Shin (2003). In our model, these types of signals are characterized as follows.

Definition 1. *A signal whose realization belongs to a countable set S is called "ex-post verifiable" if it can be represented by a measurable function $g : ([v_L, v_H], \mathcal{B}([v_L, v_H])) \to (S, 2^S)$, where $\mathcal{B}([v_L, v_H])$ is the Borel algebra on $[v_L, v_H]$.*

What this definition implies is that any signal $s \in S$, $g^{-1}(s) \equiv \{v : g(v) = s\}$ is a Borel set in $[v_L, v_H]$. Since a Borel set of $[v_L, v_H]$ must be characterized by unions of intervals, designing a disclosure plan implies combining partitions to inform the seller about the true v. Ex-post verifiability thus rules out the disclosure of noisy versions of the buyer's information. If the buyer sends him a signal, the seller must be able to confirm once all uncertainty is resolved that the true realization of v was indeed possible given the signal sent. This restriction strikes us as a natural one to impose if we are going to assume, as is common in the literature on persuasion games, that the "sender" of the information does not manipulate his signal. If the receiver only trusts ex post verifiable disclosures and erroneous disclosures are penalized, the sender has incentives to design truthful signals even if manipulation is allowed. Hence, assuming ex post verifiable disclosures can substitute for the assumption that lying is heavily penalized ex post.

Before going further, we summarize the timeline in this baseline model. First, the buyer designs a disclosure plan to send ex-post verifiable signals to the seller. Then the buyer learns his private valuation and the seller receives a signal consistent with the chosen disclosure plan. Finally, the seller quotes a price and the buyer decides whether to accept or not. We can now state our main result.

Proposition 1. *If the buyer can commit to any disclosure plan that sends ex-post verifiable signals to the seller, he designs a partial disclosure plan that results in efficient trade.*

Proposition 1 states that if private information can only be shared in a verifiably truthful manner, the privately informed agent finds it optimal to choose a disclosure plan that results in efficient trade. In other words, the incentives of the privately informed buyer, who is being screened by the seller, are aligned with social welfare. By sharing a subset of his information with the seller, the buyer is making sure that he will be quoted more efficient prices, which leads to a larger social surplus. Our paper shows that the improvement in the realized surplus that can be split between the two agents always swamps any potential loss in information rents associated with disclosure. The proposition also states that it is never optimal for the buyer to share all his information with the seller, as such a disclosure plan would drive the buyer's rents to zero. Unlike in Grossman (1981) where full disclosure is optimal, the informed trader in our model does not have market power and can only collect a surplus if he conceals some information from his counterparty. We now return to our earlier parameterized example to illustrate this result.

Example 2. *As earlier, we assume the buyer values the asset at $v \sim U[1,2]$ and the seller values it at a constant $c \leq 1$. We have already shown that whenever $c \in (0,1]$ the seller quotes a price $p = 1 + \frac{c}{2}$, which destroys the gains to trade with probability $\frac{c}{2}$. The buyer acquires the asset whenever $v \geq p$ and he collects an expected profit of:*

$$Pr\left(v \geq 1 + \frac{c}{2}\right)\left[E\left(v|v \geq 1 + \frac{c}{2}\right) - \left(1 + \frac{c}{2}\right)\right] = \frac{(2-c)^2}{8}. \tag{6}$$

Now, consider what happens if the buyer promises to share some of his information with the seller, for example, by disclosing whether $v \in \left[1, 1 + \frac{c}{2}\right)$ or $v \in \left[1 + \frac{c}{2}, 2\right]$. The seller's optimization problem when quoting a price to the buyer then adjusts to the signal being sent. If the seller learns that $v \geq 1 + \frac{c}{2}$, his optimization problem becomes:

$$\max_{p \in \left[1+\frac{c}{2}, 2\right]} \Pr\left(v \geq p | v \geq 1 + \frac{c}{2}\right) p + \Pr\left(v < p | v \geq 1 + \frac{c}{2}\right) c = \left(\frac{2-p}{1-\frac{c}{2}}\right) p + \left(\frac{p - (1+\frac{c}{2})}{1-\frac{c}{2}}\right) c, \tag{7}$$

and if he learns that $v < 1 + \frac{c}{2}$, it becomes:

$$\max_{p \in \left[1, 1+\frac{c}{2}\right)} \Pr\left(v \geq p | v < 1 + \frac{c}{2}\right) p + \Pr\left(v < p | v < 1 + \frac{c}{2}\right) c = \left(\frac{1 + \frac{c}{2} - p}{\frac{c}{2}}\right) p + \left(\frac{p-1}{\frac{c}{2}}\right) c. \tag{8}$$

In the first case, it is easy to verify that the seller finds it optimal to quote $p_h = 1 + \frac{c}{2}$, just as he did without disclosure. However, in the second case, the seller finds it optimal to quote $p_l = \max\{\frac{1}{2} + \frac{3}{4}c, 1\}$. Under this disclosure plan, the buyer collects an expected profit of:

$$\begin{aligned} & Pr\left(v \geq 1 + \frac{c}{2}\right)\left[E\left(v|v \geq 1 + \frac{c}{2}\right) - \left(1 + \frac{c}{2}\right)\right] \\ + \quad & Pr\left(p_l \leq v < 1 + \frac{c}{2}\right)\left[E\left(v|p_l \leq v < 1 + \frac{c}{2}\right) - p_l\right]. \end{aligned} \tag{9}$$

The first term is equal to the expected profit the buyer would collect without disclosure. The second term is the profit the buyer collects when $v < 1 + \frac{c}{2}$ and is strictly positive whenever $c > 0$. Thus, the buyer is strictly better off under this disclosure plan than without any *disclosure. Moreover, if $c \leq \frac{2}{3}$ the seller quotes $p_l = 1$ when $v < 1 + \frac{c}{2}$, which implies that trade is efficient regardless of the signal realization.*

If, however, $c > \frac{2}{3}$, the seller quotes $p_l = \frac{1}{2} + \frac{3}{4}c$ when $v < 1 + \frac{c}{2}$, which leads to trade that is more efficient than without disclosure, but that remains inefficient. What the proof of Proposition 1 shows is that in such a case we can apply the same reasoning again and construct an alternative disclosure plan that splits the region of inefficient trade $\left[1, 1 + \frac{c}{2}\right)$ into $\left[1, \frac{1}{2} + \frac{3}{4}c\right)$ and $\left[\frac{1}{2} + \frac{3}{4}c, 1 + \frac{c}{2}\right)$, such that the buyer is strictly better off and trade is more efficient than under the first disclosure plan proposed. Hence, we can always construct a disclosure plan that strictly dominates, from the buyer's perspective, any proposed disclosure plan that does not yield efficient trade.

Before solving for the optimal disclosure plan in our parameterized example, we analyze the tradeoff the buyer faces when designing the signals he will share with the seller under a general CDF $F(\cdot)$.

Suppose the disclosure plan $g(v)$ leads to a signal s_a when $v \in A \equiv [a_L, a_H)$ and to a signal s_b when $v \in B \equiv [b_L, b_H)$, where $b_L \geq a_H$ and $a_L \geq \hat{v}$. Since we know from Proposition 1 that the optimal disclosure plan always leads to efficient trade, we focus on a situation where the seller quotes an efficient price after receiving either of these signals. When designing his disclosure plan, the buyer must decide whether to keep these two regions separated so that $s_a \neq s_b$ or to pool them so that $s_a = s_b$. If $s_a \neq s_b$, the signal-dependent price quote, which we denote as $x(s)$, must satisfy these two conditions to allow for efficient trade: $x(s_a) = a_L$ and $x(s_b) = b_L$. The buyer then expects to collect a surplus of:

$$\int_{a_L}^{a_H} (v - a_L) dF(v) + \int_{b_L}^{b_H} (v - b_L) dF(v). \tag{10}$$

If $s_a = s_b$, the buyer's expected surplus depends on the seller's response to the disclosure, that is, the price he quotes after receiving a signal that $v \in A \cup B$. We can rule out situations where this price is weakly greater than $x(s_b) = b_L$ since if it is the case the buyer is strictly better off keeping the two regions separate. In situations where the price quote $p \in [a_L, b_L)$ however, the buyer can expect to collect a surplus of:

$$\int_p^{a_H} (v - p)dF(v) + \int_{b_L}^{b_H} (v - p)dF(v). \tag{11}$$

The net benefit of pooling these regions and have $s_a = s_b$ can thus be written as:

$$\begin{aligned} & \left[\int_p^{a_H} (v-p)dF(v) + \int_{b_L}^{b_H} (v-p)dF(v)\right] - \left[\int_{a_L}^{a_H} (v-a_L)dF(v) + \int_{b_L}^{b_H} (v-b_L)dF(v)\right] \\ = \; & (b_L - p)[F(b_H) - F(b_L)] - (p - a_L)[F(a_H) - F(p)] - \int_{a_L}^{p} (v - a_L)dF(v). \end{aligned} \tag{12}$$

The buyer pays a lower price when $v \in B$, but he might have to pay a higher price when $v \in A$. Using derivations from Section 2, we know that trade can be efficient once the seller receives a signal s_a (i.e., $p = a_L$) only if:

$$a_L - c(a_L) \geq \frac{1}{f(a_L|s_a)}. \tag{13}$$

When $s_a = s_b$, this condition becomes:

$$a_L - c(a_L) \geq \frac{Pr(v \in A \cup B)}{f(a_L)}, \tag{14}$$

which is strictly more restrictive than the analog condition for when $s_a \neq s_b$:

$$a_L - c(a_L) \geq \frac{Pr(v \in A)}{f(a_L)}. \tag{15}$$

This means that the seller is more likely to deviate to a price $p > a_L$ when the regions A and B are pooled under the same signal than when they are not. But as long as pooling the two regions still allows for efficient trade, the buyer is strictly better off setting $s_a = s_b$ as the benefit of pooling from equation (12) simplifies to:

$$(b_L - a_L)[F(b_H) - F(b_L)] > 0. \tag{16}$$

For similar reasons, the buyer may benefit from pooling intervals that are far from each other. When regions A and B are far apart, the information rents associated with obtaining an asset worth at least b_L at a price of a_L are large. Thus the optimal disclosure plan might pool regions of v that include gaps between them. But on the other hand, doing so tends to increase the seller's incentives to quote a high, inefficient price.

Overall, when designing a disclosure plan, the buyer wants to pool together multiple regions of v, potentially with gaps between them, in order to minimize the price he pays even when the asset is highly valuable to him, but he must do so while limiting the seller's incentives to screen him. In order to fully characterize the buyer's optimal disclosure plan, we return to our parameterized example for which the $F(\cdot)$ distribution is specified.

Example 3. *In earlier analyses of the parameterized example with $v \sim U[1, 2]$, we analytically showed that the buyer is better off disclosing a subset of his private information whenever trade is otherwise inefficient. However, solving for the buyer's unique optimal disclosure plan is much more involved and requires the use of numerical methods. To do so, we discretize the interval $[1, 2]$ equally into $n = 51$ points v_i. Given each possible disclosure plan, the optimal price quote $x(s)$ must live on one of those n points. We can without loss of generality assume that the realization of the signal is the price the seller would quote upon receiving the signal, by the information revelation argument in the Bayesian persuasion literature (e.g., Kamenica and Gentzkow 2011). The signal can be represented by the condition density function $(q_{ij})_{1\leq i,j\leq n}$ where r_{ij} represents the probability of the buyer sending signal $p = v_j$ conditional on his true valuation being v_i. Since we focus on ex post verifiable disclosures, $r_{ij} \in \{0, 1\}$, that is, every possible realization v_i can be* *associated with only one signal $p = v_j$. Thus, $q_{ij} = \frac{1}{n} r_{ij}$ denotes the joint density for the buyer choosing signal $p = v_j$ and the true valuation being v_i. The buyer's optimization problem boils down to picking r_{ij} to minimize the expected transaction price, conditional on trade being efficient. Since the choice variables are integers and the system is linear, the problem is an integer linear programming, which we can solve numerically. Figure 1 shows the optimal disclosure plan when $c = 0.5$.*

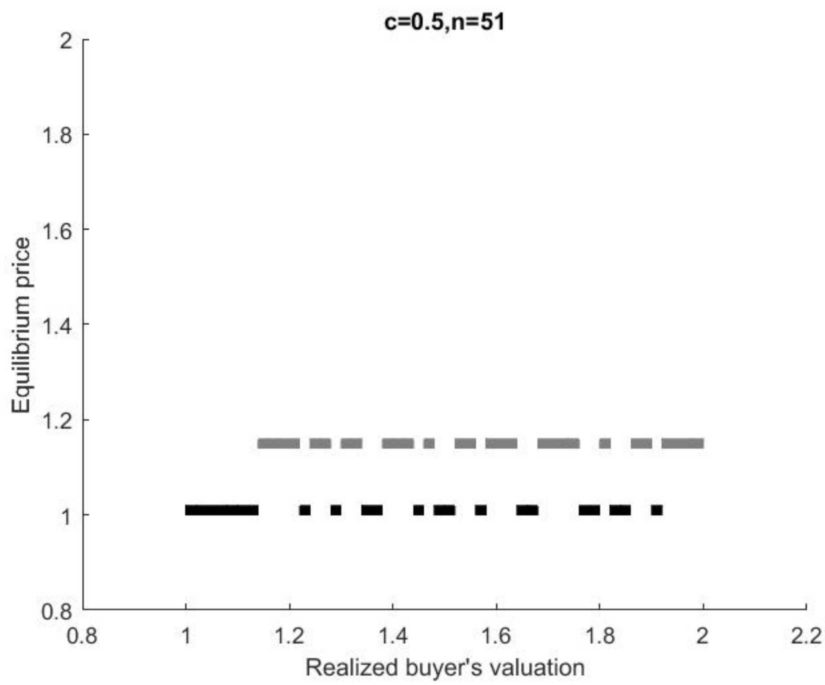

Figure 1: **Optimal disclosure plan when** $c = 0.5$ **and** $v \sim U[1, 2]$**.** In this parameterization, the buyer finds it optimal to release two signals. When the realization of v belongs to the lower (and darker) combination of sub-intervals the seller receives a signal that leads him to quote a price 1. Otherwise, the seller receives a signal that leads him to quote a price of 1.14. Importantly, trade is efficient under the buyer's optimal disclosure plan.

The buyer finds it optimal to split the interval $[1, 2]$ into two combinations of sub-intervals. When the seller receives a signal that v belongs to the lower/darker combination of sub-intervals, he responds by quoting a price $p = 1$. When the seller instead receives a signal that v belongs to the higher/paler combination of sub-intervals, he responds by *quoting a price* $p = 1.14$*. In both cases, these prices are the lowest possible realizations of* v *and the buyer always accepts to pay them in exchange for the asset. Figure 1 also shows that the optimal signal structure involves gaps between regions. These gaps allow the buyer to pay low prices and extract large information rents. In particular, the low signal arises with probability* 0.43 *and results in the buyer paying* 1 *in exchange for an asset worth, on average,* 1.41 *to him. The high signal arises with probability* 0.57 *and results in the buyer paying* 1.14 *in exchange for an asset worth, on average,* 1.57 *to him. Overall, the buyer collects an expected surplus of* 0.42 *whereas the seller collects an expected surplus of* 0.58 *under this disclosure plan. The expected social surplus is* 1*, since trade occurs with probability* 1 *and the seller values the asset at* $c = 0.5$*. Without disclosure, the seller would have quoted an inefficient price* $p = 1.25$*, which would have resulted in an expected surplus of* 0.56 *going to the seller and an expected surplus of* 0.28 *going to the buyer.*

4 Interim Disclosure

In the last section, we assumed that the buyer designs his disclosure plan prior to collecting private information, and he can then commit to not manipulating the signal later. We now discuss the robustness of our results to "interim" disclosure, that is, disclosure that is chosen after the buyer collects private information, but before the associated uncertainty is resolved for the seller. Specifically, the sequential game we now study is played as follows. First, the buyer privately observes v. Second, he designs an ex-post verifiable signal to send to the seller. Then, the seller quotes a price and the buyer decides whether to accept or reject it. We will show that in any "buyer-preferred" equilibrium, information disclosure is partial and leads to efficient trade, just as in Proposition 1 from Section 3.

We start by describing each agent's strategy profiles. Consistent with earlier notation, let $g(v)$ denote the signal that the buyer sends when his valuation for the asset is v. In the context of interim disclosure, ex-post verifiability requires that $g(v)$ is a Borel set in $[v_L, v_H]$ and that $v \in g(v)$ for any v. Since $g(v)$ is now designed by the buyer after he observes v, we can interpret $g(v)$ as the message that the buyer sends in this signaling game. Upon receiving a signal s, the seller forms a belief about the buyer's types (i.e., valuations), which we denote as $\mu(s) \in \Delta([v_L, v_H])$.[6] Then the seller quotes a price $x(s)$ to maximize his profits. Finally, the buyer decides whether to accept or reject the seller's price quote. A buyer's optimal strategy in that last stage is simply to accept the offer if and only if the quoted price is no greater than his true valuation. For ease of exposition, we do not introduce extra notation for that final stage and directly assume that the buyer follows this optimal strategy.

Overall, we now have a signaling game where the buyer sends a message and the seller chooses an action based on the message. We dub this signaling game as the *interim disclosure game*. We can now state the definition of an equilibrium in this setting.

Definition 2. *A $(g(\cdot), \mu(\cdot), x(\cdot))$ profile forms a perfect Bayesian equilibrium of the interim disclosure game if:*

1. *For every possible signal s, $x(s)$ solves $\max_p \pi(p, s)$, where $\pi(p, s)$ denotes the seller's profit if he quotes a price p and the buyer's valuation is drawn from $\mu(s)$.*

2. *For every $v \in [v_L, v_H]$, $g(v)$ solves $\max_s \max(v - x(s), 0)$, where $v \in g(v)$.*

3. *For every s in the range of g, the seller's belief function $\mu(s)$ is obtained by applying Bayes' rule given the particular signal s.*

Since beliefs are unrestricted following off-equilibrium deviations once we allow for the disclosure decision to happen after the buyer has observed private information, the uninformed seller's market power allows him to drive the buyer's information rents to zero

[6]We use notation $\Delta([v_L, v_H])$ to denote a probability distribution on $[v_L, v_H]$.

following any off-equilibrium deviation in disclosure. This type of behavior then leads to the existence of multiple perfect Bayesian equilibria with various degrees of information revelation, as opposed to the unique equilibrium with full revelation in Grossman (1981) and Milgrom (1981) (see Perez-Richet 2014, for a broader discussion of equilibrium multiplicity when the information designer picks a signal structure after acquiring private information). For instance, either full disclosure, partial disclosure, or no disclosure can be supported in equilibrium if the seller has the following beliefs: if for any s not in the range of g (i.e., whenever s is an off-equilibrium signal), the belief $\mu(s)$ assigns probability 1 to type $\bar{v}(s)$, where $\bar{v}(s) \equiv \max s$.[7]

Given the multiplicity of equilibria, we focus on buyer-preferred equilibria in order to capture the spirit of our earlier setting where the buyer moved first. What it means for the buyer to "prefer" an equilibrium in the current setting is complicated by the fact that he can now be of many types when designing a disclosure plan. Thus, we define as buyer-preferred equilibria the set of equilibria that are not dominated among buyer types (in the Pareto sense) by another equilibrium based on their interim payoffs. As in Riley (1979), we are then treating different informed-agent types as distinct players and looking for equilibria in which no type can be strictly worse off while all other types are equally off than they are in an alternative equilibrium. We conclude this discussion by stating the main result for this section.

Proposition 2. *In any buyer-preferred equilibrium of the interim disclosure game, the buyer's optimal disclosure is partial and results in efficient trade.*

[7] An equilibrium is said to feature full disclosure if $\mu(g(v))$ assigns probability 1 to type v whereas it is said to feature no disclosure if $g(v) = [v_L, v_H]$ for all $v \in [v_L, v_H]$ and thus $\mu([v_L, v_H])$ is equal to $F(v)$, the prior distribution of v.

5 Conclusion

We model a bilateral trading encounter and solve for the information that a privately informed agent would find it optimal to share with a counterparty endowed with market power. While sharing some of his information may reduce the agent's information rents, it can also make it less likely that a counterparty will inefficiently screen him and jeopardize gains to trade. We show that when considering ex post verifiable signals, the privately informed agent always finds it optimal to design a partial disclosure plan that will implement socially efficient trade in equilibrium. Moreover, our analysis shows how his optimal disclosure plan might pool together multiple regions of possible values, sometimes allowing for gaps between them, in order to maximize his information rents while limiting the counterparty's incentives to screen him.

Our paper speaks to the fundamental origins of asymmetric information problems that impede efficient trade under imperfect competition. If information is verifiable, truthfulness is enforced, and information rents only come from the upcoming transaction, trade should be efficient even though the buyer and seller are asymmetrically informed when they first meet. Our insights thus have important implications for regulating information disclosure in bilateral transactions. In our model, a regulator would not need to mandate what information traders must disclose nor would it need to produce additional information for uninformed market participants. The regulator should instead focus on enforcing the truthfulness of disclosures by disciplining traders who send signals that ex post prove to violate their own disclosure standards. Once only verifiably truthful disclosure are possible, traders have incentives to share their private information with counterparties in ways that maximize the efficiency of trade.

Appendix

Proof of Proposition 1: By contradiction, suppose that the buyer's optimal disclosure plan is represented by $g(\cdot)$, which does not implement efficient trade. We show that there exists another disclosure plan that yields a higher profit for the buyer. Denote by $x(s)$ the price the seller would quote upon receiving a signal of s. Then there exists $s_0 \in S$, such that if the signal is s_0, the seller quotes a price $x(s_0) > \hat{v}$ and $x(s_0) > \inf_v\{v : g(v) = s_0\}$. A buyer whose valuation belongs to $\{v : g(v) = s_0\} \cap (\hat{v}, x(s_0))$ would refuse to pay the seller's quoted price $x(s_0)$, leading to inefficient trade.

Now, consider the following disclosure plan where $S' = S \cup \{s'\}$ for some $s' \notin S$ and

$$\tilde{g}(v) = \begin{cases} g(v) & \text{if } g(v) \neq s_0 \\ s_0 & \text{else if } g(v) = s_0, v > x(s_0) \\ s' & \text{otherwise.} \end{cases} \tag{A1}$$

By definition, the disclosure plan $\tilde{g}(\cdot)$ would also be ex-post verifiable. We now show that $\tilde{g}(\cdot)$ would give the buyer a strictly higher ex-ante expected profit. First, note that if $s \neq s_0$, the seller would still quote a price $x(s)$. Second, if $s = s_0$, the seller would also quote $x(s_0)$ under the alternative disclosure plan $\tilde{g}(\cdot)$ as long as he receives a signal of s_0. To see this, it is sufficient to establish the following lemma.

Lemma 1. *Suppose that the seller would quote a price x if the buyer's valuation was drawn from a random distribution with CDF $G(v)$. Let $G_0(v)$ denote the distribution $G(v)$ truncated from below at x, i.e., $G(v|v \geq x)$. Then the seller would also quote a price x if the buyer's valuation was drawn from $G_0(v)$.*

Proof. We argue by contradiction. Suppose the seller would instead quote a price y if the buyer's valuation was drawn from $G_0(\cdot)$. Then $y > x$ since the support of $G_0(\cdot)$ is bounded below at x. Thus,

$$(1 - G_0(y))y + G_0(y)\mathbb{E}_G[c(v)|x \leq v < y] > x \tag{A2}$$

where the subscript G in $\mathbb{E}_G$ reminds the distribution of v. Note that $G_0(y) = \frac{G(y)-G(x)}{1-G(x)}$ and we thus have:

$$(1 - G(y))y + (G(y) - G(x))\mathbb{E}_G[c(v)|x \leq v < y] > (1 - G(x))x. \tag{A3}$$

Since $(G(y)-G(x))\mathbb{E}_G[c(v)|x \leq v < y] = \int_x^y c(v)dG(v) = \int_{v_L}^y c(v)dG(v) - \int_{v_L}^x c(v)dG(v)$, we can rewrite the inequality as:

$$(1 - G(y))y + \int_{v_L}^{y} c(v)dG(v) > (1 - G(x))x + \int_{v_L}^{x} c(v)dG(v) \tag{A4}$$

or equivalently,

$$(1 - G(y))y + G(y)\mathbb{E}_G[c(v)|v < y] > (1 - G(x))x + G(x)\mathbb{E}_G[c(v)|v < x]. \tag{A5}$$

This inequality contradicts our initial statement that the seller would quote a price x if the buyer's valuation was drawn from $G(\cdot)$. □

Given this lemma, we know that the seller would quote a price at $x(s_0)$ upon receiving a signal of s_0 under the alternative disclosure plan $\tilde{g}(\cdot)$.

Finally, suppose the seller would quote a price at z if he receives a signal of s'. Since quoting $x(s_0)$ yields zero profit in this case, it must be that $z \in [\inf g^{-1}(s_0), x(s_0))$. As a result, the buyer's ex-ante expected profit under the alternative disclosure plan $\tilde{g}(\cdot)$ is given by:

$$\sum_{s\in S}\underbrace{\int_{g^{-1}(s)\cap[x(s),v_H]}(v-x(s))1_{v\geq\hat{v}}dF(v)}_{\text{Profit from } s\in S}+\underbrace{\int_{g^{-1}(s_0)\cap[z,x(s_0))}(v-z)1_{v\geq\hat{v}}dF(v)}_{\text{Profit from } s'} \tag{A6}$$

while the profit under the disclosure plan $g(\cdot)$ is only the first term. Since $x(s_0) > z$ and $x(s_0) > \hat{v}$, the second term is strictly positive. So the buyer earns a strictly higher profit under the disclosure plan $\tilde{g}(\cdot)$ than that under $g(\cdot)$. This is a contradiction to the optimality of $g(\cdot)$.

Thus, the optimal disclosure plan must result in efficient trade. We also know that the optimal disclosure plan must reveal the buyer's information only partially. Otherwise, the seller quotes the buyer a price $p = v$ for all realizations of v and the buyer obtains no surplus. A full disclosure plan is therefore weakly dominated by a no-disclosure plan that leads to inefficient trade, which is then strictly dominated by a partial disclosure plan that leads to efficient trade, consistent with the arguments above. □

Proof of Proposition 2: To show that trade is efficient in any buyer-preferred equilibrium $(g(\cdot), \mu(\cdot), x(\cdot))$, we argue by contradiction. Suppose there exists a signal $s_0 = g(v)$ for some $v \in [v_L, v_H]$ such that $x(s_0) > \hat{v}$ and $x(s_0) > \inf_v\{v : g(v) = s_0\}$. A buyer whose valuation belongs to $\{v : g(v) = s_0\} \cap (\hat{v}, x(s_0))$ would refuse to pay the seller's quoted price $x(s_0)$, leading to inefficient trade. Let $s' \equiv \{v \in s_0 : \bar{v} \leq v < x(s)\}$. Consider the following candidate equilibrium $(\tilde{g}(\cdot), \tilde{\mu}(\cdot), \tilde{x}(\cdot))$, where

$$\tilde{g}(v) = \begin{cases} g(v) & \text{if } v \notin s_0 \\ s' & \text{else if } v \in s' \\ s_0 \backslash s' & \text{otherwise } v \in s_0 \backslash s' \end{cases}$$

We obtain $\tilde{\mu}(s')$ and $\tilde{\mu}(s_0\backslash s')$ using Bayes' rule at s' and $s_0\backslash s'$, respectively. For any signal outside the range of s_0, $\tilde{\mu} = \mu$. For any other signal s, $\tilde{\mu}(s)$ assigns probability 1 to $\bar{v}(s)$. Let $\tilde{x}(s_0)$ solves $\max_p \pi(p, s_0)$, where $\pi(p, s_0)$ denotes the seller's profit if she quotes a price p and the buyer's valuation is drawn from $\tilde{\mu}(s_0)$. It is clear that we are indeed in an equilibrium, since deviating to any other disclosure yields a profit of 0 for the buyer. Now consider the buyer's interim payoffs in this alternative equilibrium. For buyer types $v \notin s_0$ and $v \in s_0\backslash s'$, they receive payoffs identical to those from the original equilibrium $(g(\cdot), \mu(\cdot), x(\cdot))$. However, for buyer types in s', they receive weakly higher payoffs. Moreover, a buyer type $x(s_0) - \epsilon$, where ϵ is a small positive number, receives a strictly higher payoff, since he made zero profit in the original equilibrium. In all, if trade is not efficient in an equilibrium, then it is Pareto dominated among buyer types by a more efficient equilibrium. Consequently, in any buyer-preferred equilibrium of the interim disclosure game, trade must be efficient.

To show that a buyer-preferred equilibrium does not feature full disclosure, where each buyer type is quoted $p = v$ and makes zero profit, it is sufficient to construct an equilibrium where some buyer types receive positive payoffs (as no buyer type can do worse than zero profit given their right to reject a price quote). Consider the equilibrium induced by the ex ante disclosure plan we solved for in Proposition 1 from Section 3. Formally, suppose $g(\cdot)$ is the ex post verifiable disclosure plan chosen by the buyer in the ex ante disclosure game and let the interim disclosure plan follow $g(v)$ for all $v \in [v_L.v_H]$. Now, let $\mu(\cdot)$ be a belief function obtained using Bayes' rule on the equilibrium path and that assigns probability 1 to the highest type for any signal off the equilibrium path. Lastly, $x(s)$ maximizes the seller's profit based on the belief $\mu(s)$. The profile $(g(\cdot), \mu(\cdot), x(\cdot))$ is clearly an equilibrium of the interim disclosure game. In this equilibrium, the buyer receives profits identical to those in the ex ante disclosure. Thus, this equilibrium featuring partial disclosure Pareto dominates among buyer types any equilibrium with full disclosure. □

References

Acharya, Viral V., Peter DeMarzo, and Ilan Kremer. 2011. “Endogenous Information Flows and the Clustering of Announcements.” *American Economic Review* 101: 2955-2979.

Admati, Anat R., and Paul Pfleiderer. 2000. “Forcing Firms to Talk: Financial Disclosure Regulation and Externalities.” *Review of Financial Studies* 13: 479-519.

Akerlof, George A. 1970. “The Market for “Lemons”: Quality Uncertainty and the Market Mechanism.” *Quarterly Journal of Economics* 84: 488-500.

Ambrose, Brent W., Michael J. Highfield, and Peter D. Linneman. 2005. “Real Estate and Economies of Scale: The Case of REITs.” *Real Estate Economics* 33: 323-350.

Atkeson, Andrew G., Andrea L. Eisfeldt, and Pierre-Olivier Weill. 2013. “The Market for OTC Credit Derivatives.” Unpublished.

Begenau, Juliane, Monika Piazzesi, and Martin Schneider. 2015. “Banks’ Risk Exposures.” Unpublished.

Bergemann, Dirk, Benjamin Brooks, and Stephen Morris. 2015. “The Limits of Price Discrimination.” *American Economic Review* 105: 921-957.

Boone, Audra L. and J. Harold Mulherin. 2007. “How are Firms Sold?” *Journal of Finance* 62: 847-875.

Crawford, Vincent P., and Joel Sobel. 1982. “Strategic Information Transmission.” *Econometrica* 50: 1431-1451.

Diamond, Douglas W. 1985. “Optimal Release of Information by Firms.” *Journal of Finance* 40: 1071-1094.

Di Maggio, Marco, Amir Kermani, and Zhaogang Song. 2016. “The Value of Trading Relationships in Turbulent Times.” *Journal of Financial Economics*, forthcoming.

Eckbo, Espen B., Ronald M. Giammarino, and Robert L. Heinkel. 1990. “Asymmetric Information and the Medium of Exchange in Takeovers: Theory and Tests.” *Review of Financial Studies* 3: 651-675.

Ely, Jeffrey C.. 2015. “Beeps.” *American Economic Review*, forthcoming.

Eső, Peter, and Balázs Szentes. 2007. “Optimal Information Disclosure in Auctions and the Handicap Auction.” *Review of Economic Studies* 74: 705-731.

Fuchs, William, and Andrzej Skrzypacz. 2015. “Government Interventions in a Dynamic Market with Adverse Selection.” *Journal of Economic Theory* 158: 371-406.

Gal-Or, Esther. 1985. “Information Sharing in Oligopoly.” *Econometrica* 53: 329-343.

Garmaise, Mark J., and Tobias J. Moskowitz. 2004. “Confronting Information Asymmetries: Evidence from Real Estate Markets.” *Review of Financial Studies* 17: 405-437.

Glaeser, Edward L., Joseph Gyourko, and Raven Saks. 2005. “Why is Manhattan so Expensive? Regulation and the Rise in Housing Prices.” *Journal of Law and Economics* 48: 331-369.

Glode, Vincent, Richard C. Green, and Richard Lowery. 2012. “Financial Expertise as an Arms Race.” *Journal of Finance* 67: 1723-1759.

Glode, Vincent, and Christian Opp. 2016. “Asymmetric Information and Intermediation Chains.” *American Economic Review* 106: 2699-2721.

Glode, Vincent, Christian Opp, and Xingtan Zhang. 2016. “On the Efficiency of Long Intermediation Chains.” Unpublished.

Glosten, Lawrence R., and Paul R. Milgrom. 1985. "Bid, Ask and Transaction Prices in a Specialist Market with Heterogeneously Informed Traders." *Journal of Financial Economics* 14: 71-100.

Goldstein, Itay, and Yaron Leitner. 2015. "Stress Tests and Information Disclosure." Unpublished.

Green, Richard C., Burton Hollifield, and Norman Schürhoff. 2007. "Financial Intermediation and the Costs of Trading in an Opaque Market." *Review of Financial Studies* 20: 275-314.

Grossman, Sanford J. 1981. "The Informational Role of Warranties and Private Disclosure about Product Quality." *Journal of Law and Economics* 24: 461-483.

Grossman, Sanford J., and Oliver D. Hart. 1980. "Disclosure Lawa and Takeover Bids." *Journal of Finance* 35: 323-334.

Guttman, Ilan, Ilan Kremer, and Andrzej Skrzypacz. 2014. "Not Only What but also When: A Theory of Dynamic Voluntary Disclosure." *American Economic Review* 104: 2400-2420.

Hendershott, Terrence, Dan Li, Dmitry Livdan, and Norman Schürhoff. 2015. "Relationship Trading in OTC Markets." Unpublished.

Hirshleifer, Jack. 1971. "The Private and Social Value of Information and the Reward to Inventive Activity." *American Economic Review* 61: 561-574.

Hollifield, Burton, Artem Neklyudov, and Chester Spatt. 2014. "Bid-Ask Spreads and the Pricing of Securitizations: 144a vs. Registered Securitizations." Unpublished.

Jiang, Hao, and Zheng Sun. 2015. "News and Corporate Bond Liquidity." Unpublished.

Kamenica, Emir, and Matthew Gentzkow. 2011. "Bayesian Persuasion." *American Economic Review*, 101: 2590-2615.

King, Michael R., Carol Osler, and Dagfinn Rime. 2012. "Foreign Exchange Market Structure, Players, and Evolution." in *Handbook of Exchange Rates* (eds J. James, I.W. Marsh, and L. Sarno). Hoboken: John Wiley & Sons.

Kurlat, Pablo, and Laura Veldkamp. 2015. "Should we Regulate Financial Information?" *Journal of Economic Theory* 158: 697-720.

Lewis, Tracy R., and David E. M. Sappington. 1994."Supplying Information to Facilitate Price Discrimination." *International Economic Review* 35: 309-327.

Li, Dan, and Norman Schürhoff. 2014. "Dealer Networks." Unpublished.

Li, Di, Lucian A. Taylor, and Wenyu Wang. 2016. "Inefficiencies and Externalities from Opportunistic Acquirers." Unpublished.

Menkhoff, Lukas, Lucio Sarno, Maik Schmeling, and Andreas Schrimpf. 2016. "Information Flows in Foreign Exchange Markets: Dissecting Customer Currency Trades." *Journal of Finance* 71: 601-633.

Milgrom, Paul R. 1981. "Good News and Bad News: Representation Theorems and Applications." *Bell Journal of Economics* 12: 380-391.

Milgrom, Paul R. 2008. "What the Seller Won't Tell You: Persuasion and Disclosure in Markets." *Journal of Economic Perspectives* 22: 115-131.

Milgrom, Paul R., and John Roberts. 1986. "Relying on the Information of Interested Parties." *Rand Journal of Economics* 17: 18-32.

Myerson, Roger B. 1981. “Optimal Auction Design.” *Mathematics of Operations Research* 6: 58-73.

Myerson, Roger B., and Mark A. Satterthwaite. 1983. “Efficient Mechanisms for Bilateral Trading.” *Journal of Economic Theory* 29: 265-281.

Perez-Richet, Eduardo. 2014. “Interim Bayesian Persuasion: First Steps.” *American Economic Review* 104: 469-474.

Rayo, Luis, and Ilya Segal. 2010. “Optimal Information Disclosure.” *Journal of Political Economy* 118: 949-987.

Riley, John G. 1979. “Informational Equilibrium.” *Econometrica* 47: 331-359.

Roesler, Anne-Katrin, and Balázs Szentes. 2016. “Buyer-Optimal Learning and Monopoly Pricing.” Unpublished.

Shin, Hyun Song. 2003. “Disclosures and Asset Returns.” *Econometrica* 71: 105-133.

Siriwardane, Emil. 2016. “Concentrated Capital Losses and the Pricing of Corporate Credit Risk.” Unpublished.

Stroebel, Johannes. 2016. “Asymmetric Information about Collateral Values.” *Journal of Finance* 71: 1071-1111.