Information and Learning in Heterogeneous Markets

Yaarit Even[*] Alireza Tahbaz-Salehi[†] Xavier Vives[‡]

October 26, 2018

[latest version]

Abstract

This paper studies the implications of leakage of information through prices for the efficient operation of markets with heterogeneous agents. Focusing on uniform-price double auctions, we first characterize how the presence of heterogeneity (e.g., in terms of agents' trading costs, information precision, or risk attitudes) can shape the information content of prices and hence the market's informational efficiency. We find that price informativeness decreases with the extent of heterogeneity in the market. We then establish that such reductions in price informativeness can in turn manifest themselves as an informational externality: in the presence of heterogeneity, agents do not internalize the impact of their trading decisions on the information revealed to others via prices. We also show that the welfare implications of this heterogeneity-induced informational externality depends on the intricate details of the market. Our results thus indicate that accounting for the possibility of information leakage should be an important consideration in designing markets with asymmetric information. We conclude the paper by exploring the welfare implications of market segmentation in the presence of heterogeneous agents and information leakage.

Keywords: information leakage, double auctions, price discovery, information aggregation.

[*]Columbia Business School, Columbia University, yeven18@gsb.columbia.edu.
[†]Kellogg School of Management, Northwestern University, alirezat@kellogg.northwestern.edu.
[‡]IESE Business School, xvives@iese.edu.

1 Introduction

There is, by now, unanimity amongst researchers that the design of markets has an important role in improving market performance. As a result, a growing amount of literature in the operations field has focused on optimal market design in various contexts, such as auctions, two-sided markets, and on-line marketplaces. However, one aspect that has been over-looked is the possibility of information leakage, its interaction with the market design, and the impact it may have on the market outcomes.

With the advance in technology, information leakage has become ever more relevant. The digitalization of information has made the process of observing and processing agents' actions — by the market maker or the market participants — fairly easy and fast, resulting in a high possibility of their private information — on which they rely on when taking their actions — being leaked. This possibility for information leakage, in turn, may affect the actions of market participants – whether they are the agents observing the leaked information or if it is their information that is being leaked to others. Understanding such effects, their interaction with the market structure, and their impact on the outcomes is crucial for policy makers and market designers, who want to guarantee a smooth and efficient operation of the market.

The possibility of information leakage is one of the central features of financial markets. It is by now conventional wisdom that the private information held by various market participants (even anonymous ones for that matter) can be reflected in a security's price. Such a possibility has made the role of information in the "price discovery" process as one of the main concerns of policymakers in designing market regulations. For instance, in its policy report on the emergence of "dark pools" for equity trading, the International Organization of Securities Commissions (2010) argued that "[...] where regulators consider permitting different market structures [...] they should consider the impact of doing so on price discovery [...]." Similarly, Commissioner Troy A. Paredes from the U.S. Securities and Exchange Commission (2010) observed that "price discovery matters because investors would be less willing to invest if the contrarian views of short sellers were not fully incorporated into securities prices" and that "when price discovery is compromised, we run the risk that our securities markets allocate capital inefficiently."

Another prominent example is the class of emissions permits markets in the various "cap-and-trade" systems implemented around the globe. In brief, a cap is set on total amount of permitted pollution, while emission allowances within the cap are distributed to (potential) emitters. The allowances can then be exchanged on a secondary market. One of the key design objectives of such a scheme is for the price in the secondary market to reflect the social costs of emissions, thus inducing firms to internalize the impact of their production decisions. But this means an inefficient price discovery process can lead to arbitrage opportunities for investors, suboptimal budgeting decisions for firms, and inefficient reduction of emissions. Thus, not surprisingly, policymakers consider an accurate price discovery to be an important concern when designing and implementing these systems. For instance, according to the European Commission (2009), "[...]maintaining the functioning and integrity of the secondary markets as lead venues for price discovery and efficient allocation should continue to enjoy highest priority when designing a comprehensive auctioning scheme for the trading period 2013–2020 (Phase III).” Despite this emphasis, studies conducted on data from the European Union Emissions Trading System (the world’s first and so far largest cap-and-trade system) indicate that in the first period of the implementation in 2005–2007, prices failed to aggregate information effectively (Crossland et al. (2013), Slechten and Cantillon (2015)), resulting in suboptimal market operations.

In this paper, we take a step towards understanding how the quality of price discovery in the presence of information leakage can shape the informational and allocative efficiency of the market. We further try to understand how the extent of information leakage, and in turn, the quality of price discovery depends on the market architecture. We pursue this by exploring how heterogeneity in the agent-level (e.g., in terms of agents’ trading costs, information precision, valuations, and risk attitudes) may interact with the market price’s role as an endogenous source of public information. We then explore the implications of price informativeness on allocative efficiency by comparing welfare across various market architectures.

We base our analysis on a standard model of a uniform price double auction. More specifically, following Vives (2011), we focus on a competitive market consisting of finitely many agents who trade a single asset. Agents have interdependent valuations for the asset, but are uncertain about the underlying state that determines the asset’s payoff. Instead, each agent observes a potentially informative private signal about the asset fundamental. As in the rational expectations tradition (e.g., Grossman and Stiglitz (1980) and Diamond and Verrecchia (1981)), the price serves as an endogenous public signal with the ability to (fully or partially) convey each agent’s private information to other market participants. As our key assumption, we allow for ex ante heterogeneity in the precision of private signals and the agents’ preferences, by assuming that agents face potentially heterogeneous trading costs.

Our first set of results, which serves as the basis for the rest of our analysis, establishes that the distribution of trading costs has a direct impact on the informativeness of the price. In particular, we show that, in markets with more than two agents, the market is informationally efficient if and only if all trading costs coincide. This result is a consequence of the fact that introducing a second dimension of heterogeneity (i.e., heterogeneity in trading costs in addition to heterogeneity in information sets) leads to a secondary motive for trade that may be orthogonal to agents’ private information: all else equal, agents with higher trading costs trade less intensely on the same information than those with lower trading costs. This preference-induced heterogeneity, in turn, biases the price towards the private signals of agents with lower trading costs compared to the benchmark with identical traders. In fact, we show that price informativeness decreases in the (weighted) variance of agents’ trading costs, with the weights given by the precision of each agents’ private signal. In summary, higher preference heterogeneity leads to less informative prices.

Given the above observation, our second set of results then establishes that the reduction in price informativeness also manifests itself as an informational externality. More specifically, we show that agents do not internalize the impact of their trading decisions on price informativeness for other traders. Crucially, this informational externality only exists when agents are heterogeneous: we show the equilibrium is constrained efficient when all agents have identical trading costs. In contrast, when agents have heterogeneous trading costs, a subset of agents over-react to their private signals (compared to the constrained efficient benchmark), whereas the remainder of the agents under-react, where the sets of over-reacting and under-reacting agents depend on the underlying model parameters. In particular, whether an agent is over-reacting or under-reacting depends on (i) how his private signal covaries with the asset's payoff estimation error of other traders; and (ii) the slope of agents' demand curves. These two factors, in turn, depend on agents' trading costs and on the dominant role played by the price, respectively. When the informational role of price dominates — i.e., when the main role of the price is as an endogenous public signal in the market — agents with relatively high trading costs under-react to their private information, whereas agents with relatively low trading costs over-react. In contrast, when the main role of the price is as an index of scarcity — i.e., to match supply and demand — agents with high trading costs over-react to their private information, and agents with low trading costs under-react.

With the above results in hand, we then leverage the heterogeneity-induced informational externality identified above to study how the interaction of private information and market architecture determines social welfare. More specifically, we use our framework to compare a centralized market architecture to a segmented market in which agents can only trade with a subset of other individuals. As our main result, we show that, depending on the distribution of trading costs, a segmented market architecture can achieve a higher welfare compared to the centralized market. This is despite the fact that a centralized market provides more trading opportunities and — at least in principle — should lead to higher levels of price informativeness as the price can aggregate the private information of a larger set of individuals. Nonetheless, our results establish that if market centralization leads to sufficiently high levels of heterogeneity, not only price informativeness may decline, but also this decline in the quality of information aggregation and the corresponding informational externality may reduce the welfare in the centralized market below that in the segmented market architecture. Thus, policies that shape the distribution of agents that participate in the market can have a first-order effect on the efficient operations of the market.

Overall, our theoretical findings provide insight on the role of information leakage in shaping market outcomes, and how it may depend on the intricate details of the market. They also suggest that information leakage may have firs-order effect on welfare, and as a result should be an important policy concern when designing markets.

Related Literature Our theoretical framework is related to the literature on rational expectations equilibrium with a Gaussian information structure, such as Hellwig (1980), Diamond and Verrecchia (1981) and Kyle (1989), among many others. Within this literature, our paper is most closely related to Rostek and Weretka (2012), who study a market consisting of traders with identical trading costs but heterogeneous pairwise correlations in valuations. In line with our findings for a model with heterogeneous trading costs, they establish that heterogeneity in pairwise correlations can break informational efficiency. However, unlike our framework, the failure of information aggregation in their model does not translate into an informational externality: even though the market cannot fully reveal the information to all traders, neither can a social planner who has to respect the decentralized information structure of the economy.[1] The disparity between the results of Rostek and Weretka (2012) and our findings is driven by the distinct origins of informational inefficiencies in the two models. More specifically, in the presence of heterogeneous pairwise correlations, the price cannot fully aggregate information because there is no single one-dimensional statistic that can serve as a sufficient statistic for all market participants simultaneously. In contrast, our model admits such a common sufficient statistic. Yet, heterogeneity in trading costs leads to an equilibrium price that does not coincide with this statistic.

Information leakage effects have also been studied by the supply chain management literature. Li (2002) studies the possibility that confidentially-shared information between a retailer and a manufacturer may be leaked to other retailers as they observe the manufacturer's actions. The leaked information, in turn, may affect the strategies of the other retailers, even though they were not part of the information sharing agreement. Relatedly, Anand and Goyal (2009) emphasize the importance of "strategic information management," according to which firms take the possibility of information leakage to competitors into account, while Kong, Rajagopalan, and Zhang (2013) study revenue-sharing contracts that can mitigate the negative effect of information leakage in the supply chain.

Our paper is also related to the literature on optimal information revelation in disclosure policies in the context of platforms and queues. Bimpikis, Ehsani, and Mostagir (2018) focus on the optimal information disclosure policy of a contest designer regarding the competitors' progress. Relatedly, Papanastasiou, Bimpikis, and Savva (2018) study the problem of optimal information provision of on-line platforms that collect and disseminate consumers' experiences, while Candogan and Drakopoulos (2017) study the problem of optimal information revelation in a setting of a social networking platform facing the trade-off between engagement and misinformation. In the context of queues, Guo and Zipkin (2007), Jouini, Aksin, and Dallery (2011), and Allon, Bassamboo, and Gurvich (2011), among others study the effect of different information revelation methods on customers and on the overall performance of the system. In contrast to this literature, where there exists a platform or a service system who controls the nature of the information provision, in our setting, there is no entity who controls for the amount of information revealed, but rather it is determined by the way that prices incorporate and convey information, from and to market participants.

A growing literature has been studying another channel for endogenous public information — ratings and reviews — through which customers can learn about the value of different products and services. For example, Sun (2012), Besbes and Scarsini (2016), Acemoglu, Makhdoumi, Malekian, and Ozdaglar (2017), and Ifrach, Maglaras, Scarsini, and Zseleva (2018) investigate how successful is the learning process in terms of learning the true value of the products. While related, our paper departures from this literature as we allow for learning from prices.

Our paper is also related to the literature on informational efficiency and allocative efficiency of markets. For example, Pesendorfer and Swinkels (2000) studies the general tension between the two notions in a setting of auctions. In a multi-unit auction model with a finite number of bidders, the more sensitive bids are to private information, the more information is aggregated in the price but also the greater is the allocative inefficiency. However, in the limit (of the number of

[1] We formally establish this claim in Appendix A.

items and bidders) Pesendorfer and Swinkels (2000) shows that both are attained – full information aggregation and allocative efficiency. Iyer, Johari, and Moallemi (2014) look at predictions markets in a dynamic setting. They show that when all traders are risk-averse, although prices reflect risk-adjusted probabilities, under some smoothness condition, the allocation is ex-post Pareto efficient. In addition, they show that information is aggregated in the sense that an uninformed observer of the market, sharing only the common knowledge of market participants can infer the true probabilities.

Also related is the literature on efficient use of public versus private information. Morris and Shin (2002) show that in a game with strategic complementarities, agents might over-react to public information, and so releasing more public information can reduce social welfare. Angeletos and Pavan (2007) generalize the model in Morris and Shin (2002) and study for different economies the efficient use of public information. As opposed to our model, they consider exogenous public information, and so the weight agents put on their private information does not affect the content of the public signal.

Finally, our results on the welfare implications of various market architectures are related to the work of Malamud and Rostek (2017), who argue that when agents can exert market power, fragmentation of centralized markets may increase aggregate welfare.In contrast to this paper, we assume that all traders are competitive, but instead allow them to learn from endogenous public signals, i.e., prices. This channel creates an informational externality, whose magnitude is closely tied to the market architecture. As such, a transition from centralized to segmented markets impacts equilibrium price informativeness, thus leading potentially higher aggregate welfare. Also related is the recent work of Iyer, Johari, and Moallemi (2018) who look at the impact of introducing dark pools to financial markets on welfare. In other words, what are the welfare implications of having these ”closed” markets in addition to the open market (i.e., an exchange). They show the answer is ambiguous and depends on the intrinsic value of traders and the mass of speculators. Thus, similar to our paper, they show that a centralized open market may be inferior to a more decentralized market with respect to welfare.

Outline of the Paper The rest of the paper is organized as follows. Section 2 presents the model and the solution concept. In Section 3, we study price informativeness and provide a characterization of the model's informational efficiency as a function of trader characteristics. Section 4 contains our main results, where we identify the informational externality that arises when agents are heterogeneous. In Section 5, we explore the welfare implications of the heterogeneity-induced informational externality in various market architectures. All proofs and some additional technical details are provided in the Appendix.

2 Model

Consider a market consisting of n agents, denoted by $\{1, 2, \ldots, n\}$, who trade a divisible good. These agents may correspond to firms trading emissions permits in a secondary market in a cap-and-trade scheme or traders buying and selling assets in financial markets. The realized payoff of agent i who obtains x_i units of the good is given by

$$\pi_i(x_i) = \theta_i x_i - \frac{1}{2}\lambda_i x_i^2 - p x_i, \tag{1}$$

where p denotes the price of the good and θ_i, which we refer to as i's valuation, is a random variable that is drawn from the standard normal distribution. We allow for interdependence in traders' valuations by assuming that $\mathrm{corr}(\theta_i, \theta_j) = \rho$ for all pairs of agents $i \neq j$, where $\rho \in [0, 1)$. This formulation thus nests the cases with independent ($\rho = 0$) and common ($\rho \to 1$) valuations as special cases. We refer to parameter λ_i in (1) as agent i's trading cost and treat the collection of parameters $(\lambda_1, \ldots, \lambda_n)$ as a primitive of the model, which we assume to be commonly known to all agents.

In the context of financial markets, θ_i represents the dividend of the traded asset, x_i is the quantity of the asset purchased by trader i, and the trading cost λ_i can arise due to transaction taxes, inventory costs for holding the asset, or other costs incurred as a consequence of trade. Alternatively, to interpret equation (1) in the context of the emissions permits market, suppose each polluter i can produce one unit of output per one unit of pollution permit. Thus, to produce x_i units of output, which results in a revenue of $\theta_i x_i$, polluter i incurs a cost $p x_i$ to obtain the required permits as well as a quadratic production cost $\lambda_i x_i^2/2$. Regardless of the interpretation, equation (1) represents a market consisting of agents with interdependent valuations and potentially heterogeneous costs.

Prior to trading, each agent i observes a noisy private signal $s_i = \theta_i + \epsilon_i$ about her valuation, where $\epsilon_i \sim N(0, \sigma_i^2)$ are mutually independent and σ_i parametrizes i's uncertainty about θ_i. Under this specification, all signals $(s_1, \ldots, s_n)$ are informative about agent i's valuation as long as $\rho \neq 0$ and $\sigma_i > 0$.

The good/asset is supplied by a competitive market of outside agents, represented by the inverse aggregate supply function $p = \alpha + \beta \sum_{i=1}^n x_i$, where α and β are non-negative constants and $\sum_{i=1}^n x_i$ is the (inside) agents' aggregate demand for the good. Such an inverse supply function can arise by assuming that, in addition to the n traders discussed above, the market contains a representative outside agent, indexed 0, with payoff

$$\pi_0(y) = \alpha y - \beta y^2/2 - p y, \tag{2}$$

where y is the total units of the good purchased by the outside agent.[2] In the context of the emissions permits market, the outside agent can be though of as the government or regulator supplying the asset, with $\alpha y - \beta y^2/2$ capturing the social cost of y units of emissions. Market clearing requires that the traders' aggregate demand and the demand of the outside agent satisfy $y + \sum_{i=1}^n x_i = 0$.

Trade occurs via a one-shot, uniform-price double auction mechanism, according to which all agents simultaneously submit demand schedules that specify their demand for the asset as a function of the price p. Under such a trading mechanism, the strategy of trader $i \in \{1, \ldots, n\}$ is a mapping $x_i(s_i, p)$ from her private information and the price to a quantity, whereas the strategy of the representative outside agent is a function $y(p)$ that specifies his demand at any given price p. The price is then determined by the submitted demand functions and the market-clearing condition.

[2] In parallel with traders indexed 1 through n, one can interpret α and β as the outside agent's valuation and trading cost, respectively.

The competitive equilibrium of this market is defined in the usual way: it consists of a collection of demand schedules $x_i(s_i, p)$ and $y(p)$ such that (i) each trader $i \in \{1, \ldots, n\}$ maximizes her expected payoff conditional on her information set $\{s_i, p\}$ while taking the price as given, (ii) the representative outside agent maximizes his payoff given the price, and (iii) the market clears. Throughout, we restrict our attention to equilibria in linear strategies, according to which each agent i's demand schedule is an affine function of her private signal s_i and the market price p.

Before presenting our results, a few remarks are in order. First, note that the assumption that agents submit price-contingent demand schedules enables them to take the information content of the price into account, thus paving the way for the possibility of information leakage in the market: the price can serve as an endogenous public signal with the ability to (fully or partially) convey agents' private information to one another. Second, the absence of noise traders in our framework enables us to perform a well-defined welfare analysis. Such an analysis will be instrumental in disentangling the market's informational inefficiency from its allocative inefficiency. Finally, our assumption that agents are price takers ensures that the inefficiencies identified by the welfare analysis are not driven by market power or other departures from the competitive benchmark. Our main results on information leakage and the market's informational inefficiency extent to the settings where agents exert market power.

We have the following result:

Proposition 1. *There exists an equilibrium in linear strategies $x_i = a_i s_i + b_i - c_i p$, where the coefficients corresponding to trader i's strategy depend on the price via*

$$\lambda_i a_i = \frac{\mathrm{var}(p) - \mathbb{E}[ps_i]\mathbb{E}[p\theta_i]}{\mathbb{E}[s_i^2]\,\mathrm{var}(p) - \mathbb{E}^2[ps_i]} \tag{3}$$

$$\lambda_i b_i = \frac{\mathbb{E}[ps_i] - \mathbb{E}[s_i^2]\mathbb{E}[p\theta_i]}{\mathbb{E}[s_i^2]\,\mathrm{var}(p) - \mathbb{E}^2[ps_i]}\mathbb{E}[p] \tag{4}$$

$$\lambda_i c_i = 1 + \frac{\mathbb{E}[ps_i] - \mathbb{E}[s_i^2]\mathbb{E}[p\theta_i]}{\mathbb{E}[s_i^2]\,\mathrm{var}(p) - \mathbb{E}^2[ps_i]} \tag{5}$$

and the price depends on the equilibrium strategies via

$$p = \frac{\alpha + \beta \sum_{k=1}^{n}(a_k s_k + b_k)}{1 + \beta \sum_{k=1}^{n} c_k}. \tag{6}$$

Furthermore, coefficients $(a_1, \ldots, a_n)$ are independent of parameters α and β.

The above result, which will serve as the basis for the rest of our analysis, provides an implicit characterization of agents' equilibrium strategies and market-clearing price as a function of trading costs and signal precisions.[3] Despite the implicit nature of Proposition 1, a few observations are immediate. First, equation (6) establishes that the price is an affine function of all traders' private signals, thus formalizing the idea that the equilibrium price is an endogenous public signal, with the weighted average $\sum_{k=1}^{n} a_k s_k$ serving as a sufficient statistic for the information content of the price. Second, the fact that coefficients $(a_1, \ldots, a_n)$ are independent of α and β implies that even though

[3] Even though not explicit, the traders' signal precisions are reflected in the various variance and covariance terms between θ_i, s_i, and p. We explore these relationships in detail in subsequent sections.

the price level depends on the characteristics of the outside agent, its information content does not. Third and most importantly, the expression in (6) illustrates that various agents' private signals do not impact the information content of the price symmetrically. Rather, the price is biased towards the private signals of agents who assign larger weights on their signals in equilibrium. In view of equations (3)–(5), this observation implies that, in general, equilibrium price informativeness may depend on the entire profile of trading costs $(\lambda_1, \ldots, \lambda_n)$ and signal precisions $(\sigma_1^{-1}, \ldots, \sigma_n^{-1})$, an issue which will be the main focus of Section 3.

As a final remark, we note that the expression in (5) underscores the trade-off between the two roles played by the price: (i) as a measure of the opportunity cost of obtaining an extra unit of the asset and (ii) as a potentially informative endogenous public signal about the asset's underlying payoff. In particular, when the price is uninformative about the underlying state (e.g., when $\sigma_i = 0$), equation (5) implies that $c_i = 1/\lambda_i$. On the other hand, the informational role of the price is captured by the second term on the right-hand side of (5): if the price contains some information about θ_i above and beyond agent i's private information, she infers that a higher p reflects a higher payoff, thus reducing her opportunity cost of obtaining the asset. This reduction in opportunity cost is reflected as a smaller coefficient c_i in equilibrium. Put differently, the slope of the demand curve submitted by agent i not only reflects i's opportunity cost of trade, but also her desire to utilize the information contained in the price in her demand. Importantly, the relative importance of the two roles played by the price depends on the slope of the inverse aggregate supply function β. For small values of β, the price level is insensitive to the aggregate demand $\sum_{k=1}^{n} x_k$. Thus, while a small increase in the price does not change the marginal cost of acquiring the asset, such an increase is interpreted by the market as a strong positive signal about the asset's underlying value. As a result, the informational role of the price dominates, inducing the agents to submit upward-sloping demand curves ($c_i < 0$). In contrast, when β is large, the price is very sensitive to the aggregate demand $\sum_{k=1}^{n} x_k$. As a result, an increase in demand by an agent in the market — say, due to a positive signal —results in a sharp increase in the price, which induces other agents to purchase less of the asset. In other words, the role of the price as a measure of opportunity cost of the asset dominates its informational role, inducing downward-sloping demand curves ($c_i > 0$).

3 Information Leakage and Informational Efficiency

With the equilibrium characterization in Proposition 1 in hand, we now turn to studying how model primitives, and in particular, the profile of trading costs and signal precisions, shape the informational content of the price and hence the market's informational efficiency. Throughout, we rely on the following notion of informational efficiency:

Definition 1. The equilibrium is fully privately revealing to trader i if $\mathbb{E}[\theta_i | s_i, p] = \mathbb{E}[\theta_i | s_1, ..., s_n]$.

In other words, under full private revelation, the price coupled with agent i's private signal serve as a sufficient statistic for all the information dispersed throughout the market. We say the market is *informationally efficient* if the equilibrium is fully privately revealing to all agents simultaneously.

Thus, in an informationally efficient market, the leakage of information via the price is complete. We have the following result:

Proposition 2. *Suppose $\rho \neq 0$ and $\sigma_i > 0$ for all traders i. The market is informationally efficient if and only if either*

(i) there are only two agents in the market (i.e., $n = 2$); or

(ii) all trading costs coincide (i.e., $\lambda_1 = \cdots = \lambda_n$).

The above result thus establishes that when all agents have identical trading costs, the leakage of information is complete, in the sense that all agents behave as if they had access to all the private information held by other agents in the market. This is the case regardless of the profile of signal precisions $(\sigma_1, \ldots, \sigma_n)$ and hence how private information is initially distributed among the agents. Conversely, Proposition 2 also shows that in a market consisting of $n \geq 3$ agent, any heterogeneity in trading costs would make the equilibrium price to be less than fully revealing to at least one market participant.

To see the intuition underlying this result, consider the special case in which all signals are of equal precision, that is, $\sigma_i = \sigma$ for all i. In such an environment, it is immediate that full private revelation requires the equilibrium price to be a sufficient statistic for the unweighted average of all traders' private signals, i.e., $p = d_0 + d_1 \sum_{k=1}^{n} s_k$ for some constants d_0 and d_1. Yet, as we established in Proposition 1, the equilibrium price reflects $\sum_{k=1}^{n} a_k s_k$, where the coefficient a_k depends on the entire profile of trading costs $(\lambda_1, \ldots, \lambda_n)$. Thus, as long as there are two traders i and j with non-identical trading costs, the equilibrium price would reflect a weighted average of private signals, making the extraction of the unweighted average of the signals and hence full revelation impossible. Note, however, that this argument is no longer valid if $n = 2$. In that case, each trader can back out the private signal of the other trader from the price (which reveals $a_1 s_1 + a_2 s_2$) and her own signal, irrespective of the coefficients a_1 and a_2.

We remark that the failure of information aggregation established in Proposition 2 is distinct from the reasons behind partial revelation in Jordan (1983) and Rostek and Weretka (2012). Jordan (1983) illustrates that equilibrium is generically inconsistent with the efficient market hypothesis when "the dimension of the signal space is larger than the number of assets," or in other words, when the dimension of payoff-relevant variables exceeds the number of prices. On the other hand, Rostek and Weretka (2012) construct a class of models under which there is no single statistic that can simultaneously serve as a sufficient statistic for all agents in the market. In contrast to these papers, in our environment, the (single-dimensional) linear combination $\sum_{k=1}^{n} s_k/(1 - \rho + \sigma_k^2)$ is a sufficient statistic for all the information in the market for all traders simultaneously. Yet, the failure of information aggregation is a consequence of agents' equilibrium actions: the heterogeneity in trading costs induces a dispersion in agents' trading intensity that is orthogonal to their private signals, thus biasing the information content of the price towards the private information of agents with lower trading costs.

Proposition 2 thus illustrates that the nature and extent of information leakage in the market is highly sensitive to the distribution of agents' trading costs. Our next result provides a refinement of this observation by relating the extent of informational inefficiency in the market to the distribution of agents' trading costs. For each trader i, define *the information revelation gap* as

$$\phi_i = \frac{\mathrm{var}(\theta_i|s_i, p) - \mathrm{var}(\theta_i|s_1, \dots, s_n)}{\mathrm{var}(\theta_i|s_i) - \mathrm{var}(\theta_i|s_1, \dots, s_n)}. \tag{7}$$

This index, which is always a number between 0 and 1 measures the extent to which the price reduces agent i's uncertainty relative to a benchmark with no informational asymmetry. More specifically, $\phi_i = 0$ whenever the equilibrium is fully privately revealing to agent i, whereas $\phi_i = 1$ if the price does not provide agent i with any new information above and beyond her private signal s_i.[4]

Our next result relates each agent's information revelation gap to the distribution of trading costs in the market.

Proposition 3. *The information revelation gap of trader i satisfies*

$$\phi_i = \frac{\Sigma_i^2}{\Sigma_i^2 + (1/\Lambda_i)^2} + o(\rho), \tag{8}$$

where

$$\Lambda_i = \left(\frac{\sum_{k\neq i} w_k \lambda_k^{-1}}{\sum_{k\neq i} w_k}\right)^{-1} \qquad , \qquad \Sigma_i^2 = \frac{\sum_{k\neq i} w_k \left(1/\lambda_k - 1/\Lambda_i\right)^2}{\sum_{k\neq i} w_k}$$

are, respectively, the weighted harmonic mean and weighted variance of the reciprocal of trading costs of agents $k \neq i$ with weights $w_k = 1/\mathrm{var}(s_k)$.

The above result thus provides a refinement of Proposition 2 (for small values of ρ) by linking the information content of the price to the distribution of trading costs in the market. More specifically, it illustrates that, keeping the harmonic mean of trading costs Λ_i constant, an increase in the heterogeneity Σ_i in the trading costs of agents $k \neq i$ widens i's information revelation gap. In contrast, when agents $k \neq i$ have identical trading costs, equation (8) implies that the equilibrium is fully privately revealing to agent i ($\phi_i = 0$), thus recovering Proposition 2(ii) as a special case. Also note that, in line with condition (i) of Proposition 2, the above result implies that $\phi_i = 0$ when there are only two agents in the market, irrespective of their trading costs.

The characterization in Proposition 3 also underscores that the extent of information leakage depends on the joint distribution of agents' trading costs and signal precisions. In particular, equation (8) establishes that the information revelation gap ϕ_i depends not on the dispersion of trading costs, but rather on a *weighted* variance of the reciprocal of trading costs of agents $k \neq i$ with weights $w_k = 1/\mathrm{var}(s_k)$. This expression captures the idea that agent k's trading cost matters for revelation only to the extent that she posses informative signals, with the trading cost of agents with uninformative signals assigned a weight $w_k = 0$.

[4]Our notion of information revelation gap as a measure of price informativeness is distinct, but closely related to what Rostek and Weretka (2012) refer to as the index of price informativeness. More specifically, their index, ψ_i, measures the contribution of the price signal to the reduction of i's uncertainty relative to the *complete information* benchmark with no uncertainty. In contrast, ϕ_i in equation (7) measures i's residual uncertainty relative to the benchmark of *full revelation*. Formally, the two indices are related to one another via $\psi_i = (1-\phi_i)(1-\mathrm{var}(\theta_i|s_1,\dots,s_n)/\mathrm{var}(\theta_i|s_i))$.

3.1 Large Markets

Our results thus far focused on a market consisting of finitely many agents. We conclude this section by studying the model's behavior at its continuum limit and illustrating that, as long as all agents are informationally small, the market is informationally inefficient unless the distribution of trading costs is degenerate.

Formally, we consider a sequence of markets indexed by the number of agents n and focus on the limit as $n \to \infty$. Let λ_{in} and σ_{in} respectively denote the trading cost and the standard deviation of noise in agent i's signal in the market consisting of n agents, with their joint empirical distribution denoted by $\mathbf{F}_n(\lambda, \sigma)$. We use $\mathbf{F}_n(\lambda)$ and $\mathbf{F}_n(\sigma)$ to denote the corresponding marginals. Furthermore, we assume that

$$\lim_{n\to\infty} \mathbf{F}_n(\lambda, \sigma/\sqrt{n}) = \mathbf{F}(\lambda, \sigma) \tag{9}$$

for all λ and σ, with $\mathbf{F}(0, \sigma) = 0$ for all σ. This assumption serves a dual purpose. First, it ensures that the limiting market is well-defined, with almost all agents exhibiting a non-zero trading cost. Second and more importantly, as in Bergemann and Välimäki (1997), the normalization constant $1/\sqrt{n}$ on the left-hand side of (9) guarantees that each agent is informationally small as $n \to \infty$: the variance of noise σ_{in}^2 in each agent i's signal grows linearly in n. Intuitively, this normalization implies that the aggregate amount of information dispersed among all agents remains bounded even as $n \to \infty$. More specifically, it guarantees that $\liminf_{n\to\infty} \mathrm{var}(\bar{\theta}_n | s_{1n}, \dots, s_{nn}) > 0$, where $\bar{\theta}_n = \frac{1}{n}\sum_{i=1}^{n} \theta_{in}$ is the average of agents' valuations and s_{kn} denotes the private signal of agent k in the market with n agents. We have the following result:

Proposition 4. *Let ϕ_{in} denote the information revelation gap of agent i in the market consisting of n traders. Then, $\phi^* = \lim_{n\to\infty} \phi_{in} = 0$ if and only if the marginal distribution $\mathbf{F}(\lambda)$ is degenerate.*

4 Informational Externality

Propositions 2–4 in the previous section illustrate that, as long as $n \geq 3$, the equilibrium is not fully privately revealing to all market participants simultaneously unless all agents have identical trading costs. These results, however, are silent about the (in)efficiency of the equilibrium *allocation*. In this section, we study the welfare implications of the market's failure to aggregate information and show that heterogeneity in the market can lead to the emergence of an informational externality, whereby traders do not internalize how their actions shape the information content of the price. This analysis will serve as the basis for our results in Section 5 on the welfare implications of various market architectures with endogenous public signals.

We consider the constrained efficiency benchmark of Angeletos and Pavan (2007, 2009), according to which the social planner maximizes total expected surplus in the market

$$\mathbb{E}[W] = \mathbb{E}[\pi_0] + \sum_{i=1}^{n} \mathbb{E}[\pi_i]$$

subject to the same informational constraints faced by the agents in equilibrium, where recall that π_i denotes trader i's payoff and π_0 is the payoff of the outside agent. Under this specification, the action prescribed to each agent cannot depend on the private information of other agents. Thus, this formulation ensures that the planner internalizes any potential externality that agents may impose on one another while respecting the decentralized information structure of the market.[5] As in the equilibrium, we restrict the planner to affine strategies in the form of $x_i = a_i s_i + b_i - c_i p$, while imposing the market-clearing condition $y + \sum_{i=1}^{n} x_i = 0$.

We have the following result:

Proposition 5. *Suppose $\sigma_i^2 > 0$ for all i. The equilibrium is constrained efficient if either*

(i) there are only two agents in the market (i.e., $n = 2$);

(ii) agents have private valuations (i.e., $\rho = 0$); or

(iii) all traders have identical trading costs (i.e., $\lambda_1 = \cdots = \lambda_n$).

If the above conditions are violated, then the equilibrium is constrained inefficient for almost all β.

The above result thus provides a necessary and sufficient condition for constrained efficiency of the equilibrium allocation. More specifically, Proposition 5 establishes that if either condition (i)–(iii) is satisfied, then the social planner cannot improve on the equilibrium allocation without violating the decentralized information structure of the market. Contrasting this observation with Proposition 2 illustrates that these conditions are identical to the conditions that guarantee the market's informational efficiency: in a market with $n \geq 3$ agents, the equilibrium attains both informational and allocative efficiency when all agents have identical trading costs.[6]

More importantly for our purposes, however, the juxtaposition of Propositions 2 and 5 also establishes a converse implication: within our environment, any heterogeneity in trading costs not only leads to an informational inefficiency, but also a constrained inefficient allocation. In other words, as long as there is a pair of agents i and j with $\lambda_i \neq \lambda_j$, the market exhibits an externality that is not fully internalized by the market participants in equilibrium. Crucially, this externality is absent if agents have either perfect information ($\sigma_i = 0$) or private valuations ($\rho = 0$). Under either scenario, the price cannot provide the agents with any useful information. This simple observation thus implies that the heterogeneity-induced externality identified in Proposition 5 is an *informational externality*: agents do not internalize how their actions shape the information content of the endogenous public signal.

To see the intuition for the relationship between heterogeneity and the emergence of the informational externality, first consider the case in which all agents have identical trading costs.

[5] This concept bypasses the details of specific policy instruments and instead directly identifies the strategy that maximizes welfare under the restriction that information cannot be centralized.

[6] The equivalence between informational and allocative efficiency does not hold in general. See Appendix A for a slight variation of the model along the lines of Rostek and Weretka (2012), in which the equilibrium is constrained efficient even though the price is not fully privately revealing to any of the market participants, i.e., $\phi_i > 0$ for all i. In other words, even though informational efficiency in a competitive market implies allocative efficiency (as argued by Grossman (1981)), the converse is not generally true. This means that taking informational efficiency as a proxy for allocative efficiency — without performing a proper welfare analysis — may result in misleading conclusions.

Since such a market is informationally efficient, any deviation from the equilibrium strategies can only reduce price informativeness, thus implying that the social planner cannot improve on the equilibrium allocation. In contrast, when the distribution of trading costs is non-degenerate, a marginal deviation by agent i away from her equilibrium strategy results in a second-order loss in i's payoff, but potentially a first-order gain in price informativeness for other market participants and hence a first-order increase in aggregate welfare. Our next result captures how this informational externality manifests itself:

Proposition 6. *A marginal deviation by agent i away from the equilibrium weight she assigns on her private signal leads to a first-order change in aggregate welfare given by*

$$\left.\frac{d\mathbb{E}[W]}{da_i}\right|_{\text{eq}} = \gamma \sum_{k=1}^{n} \frac{\partial x_k}{\partial p} \operatorname{cov}\big(s_i, \underbrace{\theta_k - \mathbb{E}[\theta_k|s_k, p]}_{e_k}\big) \tag{10}$$

where $\gamma > 0$ is some positive constant.

The above result provides a characterization for how changes in the equilibrium strategy of agent i shapes the total surplus in the market as a function of slope of demand curves submitted by any given agent k ($\partial x_k/\partial p$) and the covariance between agent i's private signal and k's estimation error $e_k = \theta_k - \mathbb{E}[\theta_k|s_k, p]$. Before exploring the intuition underlying equation (10), we first note that the right-hand side of this equation is equal to zero whenever the market is informationally efficient. This is a consequence of the fact that when the price is fully privately revealing to agent k, then there cannot be a systematic relationship between k's estimation error and i's private signal (as otherwise the equilibrium could not have been privately revealing to k). To see this formally, note that, under full private revelation to agent k, the covariance between i's private signal and k's estimation error is given by

$$\operatorname{cov}(s_i, e_k) = \mathbb{E}[s_i\theta_k] - \mathbb{E}[s_i\mathbb{E}[\theta_k|s_k, p]] = \mathbb{E}[s_i\theta_k] - \mathbb{E}[\mathbb{E}[s_i\theta_k|s_k, p]] = 0,$$

where the second equality is a consequence of the fact that the price is fully revealing to agent k (and hence already reflects agent i's private signal), whereas the last equality is a consequence of the law of iterated expectations. Thus, Proposition 6 substantiates the relationship between Propositions 2 and 5 discussed earlier in this section: the same conditions that guarantee informational efficiency also guarantee an efficient allocation in the market.

More importantly, the characterization in Proposition 6 also illustrates that when the equilibrium is not fully privately revealing, the nature of the informational externality depends on the *interaction* between the above covariance and the slopes of the equilibrium demand curves. To see this in the most transparent manner, consider a scenario with incomplete information leakage and suppose that $\operatorname{cov}(s_i, e_k) > 0$ between a pair of agents $i \neq k$. This means that whenever agent i has better signals, agent k tends to underestimate the true underlying value of the asset. Thus, a change in the trading strategy of agent i can improve agent k's utility. Importantly, the exact nature of this change, depends on the slope of the equilibrium demand curve submitted by agent k. If agent k submits upward-sloping demand curves (so that $\partial x_k/\partial p > 0$), then a marginal increase in agent i's trading intensity would raise the price, thus inducing agent k to acquire more of the asset exactly in the states of the world in which k was underestimating its value. This increases k's utility. In contrast, if agent k submits downward-sloping demand curves (i.e., $\partial x_k/\partial p < 0$), then by putting a marginally higher weight on her private signal, agent i increases the price and induces agent k to acquire less of the good exactly when the latter was underestimating the asset's value. This reduces k's utility. An analogous argument shows that when $\operatorname{cov}(s_i, e_k) < 0$, a marginal increase in a_i decreases k's utility when k's demand curve is upward sloping, whereas it increases k's utility when k's demand is downward sloping. Finally, note that the equilibrium is constrained efficient as long as traders strategies are not indexed to the price (i.e., $\partial x_k/\partial p = 0$), in which case the model reduces to a competition in quantities as opposed to schedules.

Our next result explores the implications of Proposition 6 by relating the market's allocative efficiency to the distribution of agents' trading costs in the market.

Proposition 7. *Let a_i^{eq} and a_i^{eff} denote the weights that i assigns to her private signal in equilibrium and constrained efficient allocations, respectively. There exist $\bar{\rho} > 0$ and functions $\underline{\beta}(\rho) < \overline{\beta}(\rho)$ such that*

(a) if $\rho < \bar{\rho}$ and $\beta < \underline{\beta}$, then $a_i^{\text{eq}} < a_i^{\text{eff}}$ if and only if

$$\frac{1}{\lambda_i} < \frac{\displaystyle\sum_{k \neq i} \frac{1 - w_k}{\lambda_k} \left(\frac{\sum_{j \neq k} w_j/\lambda_j}{\sum_{j \neq k} w_j/\lambda_j^2} \right)}{\displaystyle\sum_{k \neq i} \frac{1 - w_k}{\lambda_k} \left(\frac{\sum_{j \neq k} w_j/\lambda_j}{\sum_{j \neq k} w_j/\lambda_j^2} \right)^2};$$

(b) if $\rho < \bar{\rho}$ and $\beta > \overline{\beta}$, then $a_i^{\text{eq}} < a_i^{\text{eff}}$ if and only if

$$\frac{1}{\lambda_i} > \frac{\displaystyle\sum_{k \neq i} \frac{1 - w_k}{\lambda_k}}{\displaystyle\sum_{k \neq i} \frac{1 - w_k}{\lambda_k} \left(\frac{\sum_{j \neq k} w_j/\lambda_j}{\sum_{j \neq k} w_j/\lambda_j^2} \right)},$$

where $w_k = 1/(1 + \sigma_k^2)$.

The above result therefore characterizes the set of traders that over- and under-react to their private signals. It shows that not all departures from the efficient strategy profile are in the same direction: while some traders over-react to their private signals in equilibrium, others under-react relative to the constrained efficient benchmark. Importantly, proposition 7 also illustrates that whether any given agent i over- or under-reacts to her private signal also depends on the value of β. For example, agents with large trading costs under-react to their private signals when β is small, the same agents over-react to their signals when β is large.

It is instructive to interpret Proposition 7 through the prism of Proposition 6. To this end, suppose β is small and consider an agent i with the largest trading cost. Our discussion in Section 3 indicates that the private signal of such an agent is reflected in the price with a small weight a_i relative to the sufficient statistic that would have resulted in full private revelation to all agents. This under-reflection means that when agent i has a strong positive signal, other agents tend to underestimate the value of the asset, thus implying that $\mathrm{cov}(s_i, e_k) > 0$. On the other hand, recall from the discussion following Proposition 1 that a small β means that the informational role of the price dominates its role as an index of scarcity and as a result leads to upward-sloping demand curves, as agents interpret higher prices as strong signals in favor of the asset's underlying value. Thus, by Proposition 6, an increase in a_i induces all other agents to acquire more of the asset when they underestimate its value, thus increasing the overall welfare in the market. This is indeed the statement of Proposition 7(a). In contrast, for large values of β, the informational role of the price is weakened, resulting in downward-sloping demand curves. Hence, equation (10) indicates that a marginal increase in a_i would reduce the welfare of all other agents, consistent with Proposition 7(b).

Taken together, Propositions 6 and 7 illustrate that, in the presence of information leakage, the efficient operation of the market is highly sensitive to (i) the extent of market's informational efficiency and (ii) the relative importance of the price's informational and allocative roles, parameterized by parameter β in our setting.

5 Information Leakage and Market Architecture

Our results in Section 4 illustrate that the price's role as an endogenous public signal leads to the emergence of an informational externality whenever the distribution of trading costs is non-degenerate. In this section, we study how this externality can lead to non-trivial implications by comparing welfare across various market architectures. More specifically, we consider two market architectures, one centralized and one segmented, and show that market centralization may reduce price informativeness by strengthening the informational externality, thus resulting in a reduction of aggregate welfare compared to a segmented market architecture.

To this end, fix the set of traders $\{1, \dots, n\}$ with profile of trading costs $(\lambda_1, \dots, \lambda_n)$ and signal precisions $(\sigma_1^{-1}, \dots, \sigma_n^{-1})$ and consider two different market architectures: a centralized architecture in which all trade occurs on the same exchange with a single market-clearing price — as in the model studied thus far — and a segmented market architecture in which each trader can only trade in one of the multiple exchanges with a specific subset of other market participants. Formally, the segmented market architecture is defined as a partition $\mathcal{S} = \{S_1, \dots, S_m\}$ of the set of traders $\{1, \dots, n\}$ for some $m \geq 2$, with trader $i \in S_k$ only capable of trading with other traders $j \in S_k$. Thus, as in Malamud and Rostek (2017), each segment S_k in the segmented market architecture has a separate market-clearing price. To ensure consistency between the centralized and decentralized architectures, we also assume that a fraction $\zeta_k \in [0, 1]$ of outside traders are also active in segment S_k, with $\sum_{k=1}^{m} \zeta_k = 1$.

We start with the following benchmark result:

Proposition 8. *Expected welfare in the centralized market architecture is higher than the segmented architecture, if either of the following conditions are satisfied:*

(i) all traders have complete information about their valuations ($\sigma_i = 0$);

(ii) traders' valuations are independent ($\rho = 0$);

(iii) all trading costs are identical.

Proposition 8 provides a benchmark for the comparison between various market architectures. In particular, it establishes that as long as the market is informationally efficient — which occurs either because market participants have no use for the private information of other traders or when they all have identical trading costs — then market centralization leads to a higher aggregate welfare. This increase in welfare operates via two distinct channels. First, market centralization enables each agent to trade in a market consisting of a larger number of participants, thus leading to further realization of gains from trade. This is the channel that underlies the gains from centralization in cases (i) and (ii) above. Second, in the case that traders can benefit from other market participants' private information — as in case (iii) above — market centralization means that the price aggregates the private signals of a larger number of traders, thus increasing price informativeness for all agents and hence welfare.

With Proposition 8 as the benchmark, our next result provides a comparison between the two market architectures in the presence of trader heterogeneity. For any trader i in the centralized architecture, define $m_i^{\text{cen}} = \sum_{j \neq i} 1/\text{var}(s_j)$. Similarly for the segmented architecture, let $m_i^{\text{seg}} = \sum_{j \in S(i) \setminus \{i\}} 1/\text{var}(s_j)$, where $S(i)$ denotes the segment that agent i can trade in. We have the following result:

Proposition 9. *There exist constants $\underline{\rho} > 0$ and $\underline{\beta} > 0$ such that if $\rho < \underline{\rho}$ and $\beta < \underline{\beta}$, then expected welfare in the centralized market structure is higher than the segmented market structure only if*

$$\sum_{i=1}^{n} \frac{1}{\lambda_i} \left(\frac{\sigma_i^2}{1 + \sigma_i^2} \right)^2 \left((m_i^{\text{cen}} - m_i^{\text{seg}}) - (\phi_i^{\text{cen}} m_i^{\text{cen}} - \phi_i^{\text{seg}} m_i^{\text{seg}}) \right) \geq 0, \tag{11}$$

where ϕ_i^{cen} and ϕ_i^{seg} are trader i's information revelation gaps in the centralized and segmented markets, respectively.

The term $m_i^{\text{cen}} - m_i^{\text{seg}}$ measures how much more information is available to the market participants post centralization. But then there is a penalty term which reduces gains from centralization if the information revelation gap increases. The second term on the left-hand side of (11) creates a countervailing force that may reduce and even reverse the welfare gains from centralization. The juxtaposition of the above result with Proposition 3 relates welfare gains in centralizing the markets to the heterogeneity in each segment and the market in general. Furthermore, the above result reduces to part (iii) of Proposition 8 when all trading costs are identical, in which case $\phi_i^{\text{cen}} = \phi_i^{\text{seg}} = 0$ for all i, implying that $\mathbb{E}[W^{\text{cen}}] \geq \mathbb{E}[W^{\text{seg}}]$. On the other hand, when (11) is violated, we get the opposite. Finally, each term is also weighed by trader i's trading cost: clearly traders with high trading costs will trade less and hence matter less for aggregate welfare.

The result above thus implies that policies that shape the distribution of agents that participate in the market, which in turn, shapes price informativeness can have a first-order effect on the efficient operations of the market.

6 Conclusions

In this paper, we investigate how heterogeneity of market participants can shape the information content of the price, in the presence of information leakage. We find that price informativeness is highly sensitive to the characteristics of market participants. In particular, we find that the price is less informative the more heterogeneous are the agents. This is a consequence of the fact that agents' characteristics determine the intensity of their market activity and hence the extent to which their private information will be leaked via prices.

Moreover, we find that the market's *informational* inefficiency translates into an informational externality, resulting in an *allocative* inefficient market. In a heterogeneous market, agents do not internalize how their actions shape the information that is leaked via prices and hence the information available to other agents. The underlying of this informational externality is twofold: One, the dominant role played by the price — informational or allocative — which determines the slope of agents' demand curves. Second, how the private information of each agent covaries with the payoff estimation of other agents, where a positive covariance characterizes relatively high trading cost agents, and implies that other agents under-estimate the asset's value, and a negative covariance characterizes relatively low trading cost agents, and implies that other agents over-estimate the value of the asset. Taken together, we find that when the informational role of price dominates, agents with high trading costs are under-reacting to their private information, whereas agents with low trading costs are over-reacting. However, when the allocative role of price dominates, the opposite is true.

We further conclude that the extent of information leakage and its effect on market performance is tightly related to the market architecture. As opposed to conventional belief, we find that welfare in a centralized market — where, potentially, there are higher realizations of gain and more information can be aggregated — may be low compared to a segmented market (i.e., agents are allowed to trade only within one segment of the market). This result emphasizes the potential impact that the heterogeneity-induced informational externality may have on market welfare.

Our findings suggest that the extent of information leakage via prices may vary with the intricate details of the market structure. Policies that shape the distribution of agents that participate in the market can have a first-order effect on the efficient operations of the market. Thus, accounting for the possibility and extent of information leakage should be a central pillar of optimal market design, specially in environments with highly dispersed information.

Appendix

A Informationally-Inefficient Efficient Markets

Our results in the main body of the paper establish that when trading costs are heterogeneous, (i) the price signal does not fully aggregate the information in the market and (ii) the equilibrium is constrained inefficient, as traders do not internalize the impact of their trading decisions on the information content of the price. In other words, the equilibrium is both *informational* and *allocative* inefficient. In this appendix, we show that, in general, incomplete aggregation of information does not necessarily imply allocative inefficiency. We illustrate this by contrasting our results to an extension of the model of Rostek and Weretka (2012), where traders have homogeneous trading costs but are asymmetric in the correlation between their private valuations. More specifically, we show that even though such heterogeneity leads to an incomplete aggregation of information, the equilibrium is constrained efficient, in the sense that the social planner cannot improve on the allocation.

As in the baseline model in Section 2, consider a market consisting of n price-taking traders with payoffs given by (1) and private signals $s_i = \theta_i + \epsilon_i$, where $\epsilon_i \sim \mathcal{N}(0, \sigma_i^2)$. As in our baseline model, the assumption that traders take the price as given guarantees that any potential inefficiency is not driven by traders' market power. In a departure from the baseline model, however, suppose that the interdependencies in private valuations can be heterogeneous among different pairs of traders. More specifically, suppose $\mathrm{corr}(\theta_i, \theta_j) = \rho_{ij}$, with the assumption that

$$\frac{1}{n-1}\sum_{j \neq i} \rho_{ij} = \bar{\rho} \tag{12}$$

for some $\bar{\rho} \in (0,1)$ and all traders i. This assumption ensures that all traders face the same average interdependencies in the market. We have the following result:

Proposition A.1. *Suppose pairwise correlations satisfy* (12). *Also suppose all trading costs and signal precisions coincide. Then,*

(a) The equilibrium is fully privately revealing to all traders if and only if $\rho_{ij} = \bar{\rho}$ for all $i \neq j$.

(b) The equilibrium is constrained efficient, regardless of the pairwise correlations.

The first statement of the above proposition, which generalizes Proposition 3 of Rostek and Weretka (2012), illustrates that heterogeneity in pairwise correlations prevents full private revelation in the sense of Definition 1. This is a consequence of the fact that full private revelation for trader i requires the price to be equal to a specific weighted average of private signals. But the presence of heterogeneous correlations means that this weight average may be different for different traders, implying that at least one trader cannot fully extract the sufficient statistic of other traders' private signals by observing the price.

More importantly for our purposes however, part (b) of Proposition A.1 illustrates that the failure of informational efficiency highlighted in part (a) may not translate into allocative inefficiency: no matter what the pairwise correlations are, all traders internalize the impact of their actions on others and no policy can improve upon the equilibrium allocation. This result thus underscores that equating informational efficiency with allocative efficiency — without performing a proper welfare analysis — can lead to misleading conclusions.

B Proofs

Proof of Proposition 1

Recall that the (ex ante) expected profit of trader i is given by $\mathbb{E}[\pi_i] = \mathbb{E}[\theta_i x_i] - \frac{1}{2}\lambda_i \mathbb{E}[x_i^2] - \mathbb{E}[px_i]$ and suppose trader i follows a linear strategy given by $x_i = a_i s_i + b_i - c_i p$, where a_i, b_i, and c_i are coefficients that only depend on model parameters. Plugging this expression into i's expected profit implies that

$$\mathbb{E}[\pi_i] = a_i - c_i\mathbb{E}[\theta_i p] - \frac{1}{2}\lambda_i(1+\sigma_i^2)a_i^2 - \frac{1}{2}(\lambda_i c_i^2 - 2c_i)\mathbb{E}[p^2] + a_i(\lambda_i c_i - 1)\mathbb{E}[ps_i] - \frac{1}{2}\lambda_i b_i^2 + b_i(\lambda_i c_i - 1)\mathbb{E}[p].$$

Trader i's objective is to maximize her expected profit while taking the price as given. As a first observation, note that i's objective function is jointly concave in (a_i, b_i, c_i). Therefore, the first-order conditions with respect to these parameters are both necessary and sufficient for optimality. Hence, the best-response strategy of trader i satisfies the following relationships:

$$1 - \lambda_i(1+\sigma_i^2)a_i + (\lambda_i c_i - 1)\mathbb{E}[ps_i] = 0 \tag{13}$$

$$-\lambda_i b_i + (\lambda_i c_i - 1)\mathbb{E}[p] = 0 \tag{14}$$

$$-\mathbb{E}[\theta_i p] - (\lambda_i c_i - 1)\mathbb{E}[p^2] + \lambda_i a_i \mathbb{E}[ps_i] + \lambda_i b_i \mathbb{E}[p] = 0. \tag{15}$$

On the other hand, market clearing requires that $y + \sum_{i=1}^n x_i = 0$, where y is the quantity demanded by the outside trader. Hence,

$$\alpha - p + \beta \sum_{i=1}^n (a_i s_i + b_i - c_i p) = 0,$$

where we are using the fact that the first-order condition of the outside trader is given by $\alpha - p + \beta y = 0$. Rearranging the above terms therefore implies that the equilibrium price is given by (6).

Equations (13)–(6) provide a system of equations that relate traders' equilibrium strategies to the model fundamentals. Plugging in the expression for the price (6) in equations (13)–(15) (followed by some tedious calculations) then implies that

$$\lambda_i a_i = \frac{\sum_{k\neq i}(1-\rho+\sigma_k^2)a_k^2 + \rho(1-\rho)(\sum_{k\neq i} a_k)^2 - \rho\sigma_i^2 a_i \sum_{k\neq i} a_k}{(1+\sigma_i^2)\sum_{k\neq i} a_k^2(1-\rho+\sigma_k^2) + \rho(1-\rho+\sigma_i^2)(\sum_{k\neq i} a_k)^2} \tag{16}$$

$$\beta\lambda_i b_i = -\frac{\rho\sigma_i^2(\sum_{k\neq i} a_k)(\alpha + \beta\sum_{k=1}^n b_k)}{(1+\sigma_i^2)\sum_{k\neq i} a_k^2(1-\rho+\sigma_k^2) + \rho(1-\rho+\sigma_i^2)(\sum_{k\neq i} a_k)^2} \tag{17}$$

$$\beta(1-\lambda_i c_i) = \frac{\rho\sigma_i^2(\sum_{k\neq i} a_k)(1 + \beta\sum_{k=1}^n c_k)}{(1+\sigma_i^2)\sum_{k\neq i} a_k^2(1-\rho+\sigma_k^2) + \rho(1-\rho+\sigma_i^2)(\sum_{k\neq i} a_k)^2}. \tag{18}$$

The proof is complete once we show that the system of equations (16)–(18) has a solution $(a_i, b_i, c_i)_{i=1}^n$. We first establish that there exists a vector $a = (a_1, \dots, a_n)$ that satisfies (16) for all i. To this end, define the mapping $\Phi : \mathbb{R}^n_{++} \to \mathbb{R}^n_{++}$ as

$$\Phi_i(a) = \frac{\sum_{k \neq i}(1-\rho+\sigma_k^2)a_k^2 + \rho(1-\rho)\left(\sum_{k\neq i} a_k\right)^2}{\lambda_i\left((1+\sigma_i^2)\sum_{k\neq i} a_k^2(1-\rho+\sigma_k^2) + \rho(1-\rho+\sigma_i^2)(\sum_{k\neq i} a_k)^2\right) + \rho\sigma_i^2\sum_{k\neq i} a_k}.$$

Note that a satisfies equilibrium condition (16) if and only $\Phi(a) = a$. Define the set $A = \prod_{i=1}^n [\underline{a}_i, \overline{a}_i]$, where $\underline{a}_i = \lambda_{\max}^{-1}(1-\rho)/(1-\rho+\sigma_i^2)$ and $\overline{a}_i = \lambda_{\min}^{-1}(1-\rho)/(1-\rho+\sigma_i^2)$, with $\lambda_{\max}$ and $\lambda_{\min}$ denoting the largest and smallest trading costs, respectively. It is easy to verify that $\Phi_i(a) \geq \underline{a}_i$ whenever $\rho\sigma_i^2\sum_{k\neq i} a_k(\lambda_{\max}(1-\rho+\sigma_k^2)a_k - (1-\rho)) \geq 0$, which holds trivially as long as $a_k \geq \underline{a}_k$ for all $k \neq i$. Similarly, $\Phi_i(a) \leq \overline{a}_i$ as long as $\rho\sigma_i^2\sum_{k\neq i} a_k(\lambda_{\min}(1-\rho+\sigma_k^2)a_k - (1-\rho)) \leq 0$, an inequality that is satisfied when $a_k \leq \overline{a}_k$ for all $k \neq i$. These observations therefore imply that Φ maps the compact and convex set A to itself. Thus, by Brouwer's fixed point theorem, there exists $a \in A$ such that $\Phi(a) = a$, hence guaranteeing that there exist coefficients $a_1, \dots, a_n$ that satisfy equation (16) for all i simultaneously.

Next, consider (17). This system of equations has a trivial solution of $b_i = 0$ for all i when $\alpha = 0$. We therefore consider the case that $\alpha \neq 0$. Dividing both sides of the equation by λ_i and summing over all i leads to

$$\beta\sum_{i=1}^n b_i = -\left(\alpha + \beta\sum_{i=1}^n b_i\right)\sum_{i=1}^n \frac{\rho\sigma_i^2}{\lambda_i\delta_i}\sum_{k\neq i} a_k, \tag{19}$$

where

$$\delta_k = (1+\sigma_k^2)\sum_{j\neq k} a_j^2(1-\rho+\sigma_j^2) + \rho(1-\rho+\sigma_k^2)(\sum_{j\neq k} a_j)^2. \tag{20}$$

Since $a_i > 0$ for all i, it must be the case that $\sum_{i=1}^n \frac{\rho\sigma_i^2}{\lambda_i\delta_i}\sum_{k\neq i} a_k \neq -1$. Therefore, given coefficients $a_1, \dots, a_n$, there exists a unique $\sum_{i=1}^n b_i$ that satisfies (19). Plugging back this solution into (17) then implies that there exists a collection of constants $(b_1, \dots, b_n)$ that satisfy the equilibrium condition.

Finally, consider (18). This equation implies that

$$\beta\sum_{i=1}^n c_i = \beta\sum_{i=1}^n \frac{1}{\lambda_i} - \left(1 + \beta\sum_{i=1}^n c_i\right)\sum_{i=1}^n \frac{\rho\sigma_i^2}{\lambda_i\delta_i}\sum_{k\neq i} a_k.$$

Once again, the fact that $\sum_{i=1}^n \frac{\rho\sigma_i^2}{\lambda_i\delta_i}\sum_{k\neq i} a_k \neq -1$ guarantees that the exists a unique $\sum_{i=1}^n c_i$ that satisfies the above equation. Plugging back this solution into (18) then implies that there exists a collection $(c_1, \dots, c_n)$ that satisfies the equilibrium conditions. $\square$

Proof of Proposition 2

Lemma B.1. *$a_i(1-\rho+\sigma_i^2) = a_k(1-\rho+\sigma_k^2)$ for all pairs i and k if and only if all trading costs coincide.*

Proof. First suppose all trading costs are identical, i.e., $\lambda_i = \lambda$ for all i. Under such an assumption, it is immediate to verify that $\lambda a_i = (1-\rho)/(1-\rho+\sigma_i^2)$, thus implying that $a_i(1-\rho+\sigma_i^2) = a_k(1-\rho+\sigma_k^2)$ for all pairs of traders i and k.

To prove the converse implication, suppose $a_i(1-\rho+\sigma_i^2) = a_k(1-\rho+\sigma_k^2)$ for all pairs $i \neq k$. This means that there exists a constant $S > 0$ such that $(1-\rho+\sigma_k^2)a_k = S$ for all k. Plugging this expression into equilibrium condition (16) leads to

$$S\lambda_i\left((1+\sigma_i^2)+\rho(1-\rho+\sigma_i^2)\sum_{k\neq i}(1-\rho+\sigma_k^2)^{-1}\right)+\rho\sigma_i^2 = (1-\rho+\sigma_i^2)\left(1+\rho(1-\rho)\sum_{k\neq i}(1-\rho+\sigma_k^2)^{-1}\right).$$

Solving for the constant S from the above expression implies that $S = (1-\rho)/\lambda_i$ for all i, which can hold only if $\lambda_i = \lambda_k$ for all i and k. □

We now turn to the proof of Proposition 2. As a first observation, note that when $n = 2$, it is immediate that the equilibrium is fully privately revealing to both traders. Hence, in the rest of the proof we assume that there are at least three traders in the market. Suppose that the equilibrium is fully privately revealing to all traders, where recall from Definition 1 that this is equivalent to assuming that $\mathbb{E}[\theta_i|s_i, p] = \mathbb{E}[\theta_i|s_1, ..., s_n]$ for all i, where

$$\mathbb{E}[\theta_i|s_1,\dots,s_n] = \left(\frac{1-\rho}{1-\rho+\sigma_i^2}\right)s_i + \frac{\rho\sigma_i^2}{(1-\rho+\sigma_i^2)\left(1+\rho\sum_{j=1}^n(1-\rho+\sigma_j^2)^{-1}\right)}\sum_{k=1}^n\left(\frac{1}{1-\rho+\sigma_k^2}\right)s_k.$$

On the other hand, the fact that market-clearing price satisfies (6) means that

$$\mathbb{E}[\theta_i|s_i,p] = \frac{1}{\delta_i}\left(\sum_{k\neq i}(1-\rho+\sigma_k^2)a_k^2+\rho(1-\rho)\Big(\sum_{k\neq i}a_k\Big)^2-\rho\sigma_i^2a_i\sum_{k\neq i}a_k\right)s_i+\frac{1}{\delta_i}\left(\rho\sigma_i^2\sum_{j\neq i}a_j\right)\sum_{k=1}^n a_ks_k,$$

where δ_i is given by (20). Hence, full private revelation requires that the coefficient on signal s_k in the above two expressions coincide for all k. Hence, as long as there are at least three traders in the market, full private revelation to all traders i implies that $a_j/a_k = (1-\rho+\sigma_k^2)/(1-\rho+\sigma_j^2)$ for all $j, k \neq i$. Consequently, by Lemma B.1, all trading costs have to coincide. □

Proof of Proposition 3

As a first observation, note that

$$\mathrm{var}(\theta_i|s_1,\dots,s_n) = \frac{\sigma_i^2}{1-\rho+\sigma_i^2}\left(1-\frac{\rho(1-\rho+\sigma_i^2)^{-1}+\rho^2\sum_{k\neq i}(1-\rho+\sigma_k^2)^{-1}}{1+\rho\sum_{k=1}^n(1-\rho+\sigma_k^2)^{-1}}\right). \tag{21}$$

Furthermore, recall from (6) that $\mathrm{var}(\theta_i|s_i,p) = \mathrm{var}(\theta_i|s_i, \sum_{k\neq i}a_ks_k)$. Therefore,

$$\mathrm{var}(\theta_i|s_i,p) = \frac{\sigma_i^2}{1-\rho+\sigma_i^2}\left(1-\frac{\rho(1-\rho+\sigma_i^2)^{-1}\sum_{k\neq i}a_k^2(1-\rho+\sigma_k^2)+\rho^2(\sum_{k\neq i}a_k)^2}{\left(1+\rho(1-\rho+\sigma_i^2)^{-1}\right)\sum_{k\neq i}a_k^2(1-\rho+\sigma_k^2)+\rho(\sum_{k\neq i}a_k)^2}\right).$$

Finally, note that $\mathrm{var}(\theta_i|s_i) = \sigma_i^2/(1+\sigma_i^2)$. Combining the above expressions implies that trader i's information revelation gap, defined in (7), is given by

$$\phi_i = \left(\frac{1+\sigma_i^2}{1-\rho+\sigma_i^2}\right)\frac{\sum_{k\neq i}a_k^2(1-\rho+\sigma_k^2)-(\sum_{k\neq i}a_k)^2(\sum_{k\neq i}(1-\rho+\sigma_k^2)^{-1})^{-1}}{\left(1+\rho(1-\rho+\sigma_i^2)^{-1}\right)\sum_{k\neq i}a_k^2(1-\rho+\sigma_k^2)+\rho(\sum_{k\neq i}a_k)^2}. \tag{22}$$

On the other hand, equation (16) implies that $\lim_{\rho\to 0} a_i = w_i/\lambda_i$ for all traders i, where $w_i = 1/(1+\sigma_i^2)$. Taking the limit as $\rho \to 0$ from both sides of the above equation implies that

$$\lim_{\rho\to 0} \phi_i = \frac{\left(\sum_{k\neq i} w_k/\lambda_k^2\right) - \left(\sum_{k\neq i} w_k/\lambda_k\right)^2 \Big/ \left(\sum_{k\neq i} w_k\right)}{\left(\sum_{k\neq i} w_k/\lambda_k^2\right)}.$$

Dividing both the numerator and the denominator by $\sum_{k\neq i} w_k$ then complete the proof. □

Proof of Proposition 4

The implication that $\phi_i^* = \lim_{n\to\infty} \phi_{in} = 0$ whenever $\mathbf{F}(\lambda)$ is degenerate is trivial. We therefore only provide the proof of the converse implication. In particular, suppose that $\phi_i^* = 0$. Recall from the proof of Proposition 3 that trader i's information revelation gap satisfies equation (22). Therefore,

$$\phi_i^* = \frac{\int a^2\sigma^2 d\mathbf{G} - \left(\int a d\mathbf{G}\right)^2 \Big/ \int \sigma^{-2} d\mathbf{G}}{\int a^2\sigma^2 d\mathbf{G} + \rho\left(\int a d\mathbf{G}\right)^2}, \tag{23}$$

where $\mathbf{G}(a,\lambda,\sigma) = \lim_{n\to\infty} \mathbf{G}_n(na,\lambda,\sigma/\sqrt{n})$ and $\mathbf{G}_n(a,\lambda,\sigma)$ denotes the joint empirical distribution of the weight that traders assign to their private signals, their trading costs, and the signal precisions. Note that whereas the joint distribution of λ_i and σ_i, denoted by $\mathbf{F}_n(\lambda,\sigma)$ is exogenous, the weights that traders assign to their private signals are equilibrium objects that are determined endogenously. Nonetheless, $\mathbf{G}_n$ can always be expressed in terms of the model primitive $\mathbf{F}_n$ using equation (16).[7]

Since $\phi_i^* = 0$, (23) implies that

$$\int a^2\sigma^2 d\mathbf{G} \int \sigma^{-2} d\mathbf{G} = \left(\int a d\mathbf{G}\right)^2.$$

But by he Cauchy-Schwarz inequality, the above equality can hold only if $a_i\sigma_i^2 = a_j\sigma_j^2$ for almost all pairs i and j. Hence, by equation (16), it must be the case that $\lambda_i = \lambda_j$, which means that $\mathbf{F}(\lambda)$ is degenerate. □

Proof of Proposition 5

Lemma B.2. *Let $x_i = a_i s_i + b_i - c_i p$ denote traders' equilibrium strategies. Then,*

$$\frac{\beta c_k}{1+\beta\sum_{j=1}^n c_j} = \frac{\beta Q_k - M_k}{\lambda_k(1+\beta\sum_{j=1}^n 1/\lambda_j)}, \tag{24}$$

where Q_k and M_k are independent of the value of β.

[7] The normalization constant n in the definition $\mathbf{G}$ is a consequence of the assumption that all traders are informationally small as $n\to\infty$. More specifically, the fact that σ_{in} grows at rate $\sqrt{n}$ implies that the weight a_{in} has to decay to zero at rate n.

Proof. Recall that equilibrium coefficients (a_i, b_i, c_i) satisfy equations (16)–(18). Summing both sides of (18) over all traders i and solving for $1+\beta\sum_{i=1}^n c_i$ implies that

$$1+\beta\sum_{j=1}^n c_j = \frac{1+\beta\sum_{j=1}^n 1/\lambda_j}{1+\rho\sum_{j=1}^n\sum_{r\neq j} a_r\sigma_j^2/(\lambda_j\delta_j)}, \tag{25}$$

where δ_k is given by (20). Plugging the above back into the expression for c_k in (18) then establishes (24), where Q_k and M_k are given by $Q_k = 1+\rho\sum_{j\neq k}\left(\sum_{r\neq j}\frac{\sigma_j^2 a_r}{\lambda_j\delta_j} - \sum_{r\neq k}\frac{\sigma_k^2 a_r}{\delta_k\lambda_j}\right)$ and $M_k = \frac{\rho\sigma_k^2}{\delta_k}\sum_{j\neq k} a_j$. Finally, to establish that Q_k and M_k are independent of β, recall that the coefficients $(a_1,\ldots,a_n)$ are solutions to the system of equations given by (16), which does not depend on β. Hence, Q_k and M_k are independent of β. □

With the above lemma in hand, we now return to the proof of Proposition 5. We prove this result by determining the conditions under which the equilibrium strategies identified in Proposition 1 satisfy the optimality conditions of the planner's problem.

Recall that the total ex ante surplus in the market is given by $\mathbb{E}[W] = \mathbb{E}[\pi_0] + \sum_{i=1}^n \mathbb{E}[\pi_i]$, where π_0 is the surplus of the outside trader and π_i is the profit of trader i. Therefore, the market-clearing condition $y+\sum_{i=1}^n x_i = 0$ implies that

$$\mathbb{E}[W] = \sum_{i=1}^n \mathbb{E}[\theta_i x_i] - \frac{1}{2}\sum_{i=1}^n \lambda_i\mathbb{E}[x_i^2] + \alpha\mathbb{E}[y] - \frac{\beta}{2}\mathbb{E}[y^2]. \tag{26}$$

When agents follow linear strategies in the form of $x_i = a_i s_i + b_i - c_i p$, the expected total surplus is given by

$$\mathbb{E}[W] = \sum_{i=1}^n \mathbb{E}[(\theta_i-\alpha)(a_i s_i+b_i-c_i p)] - \frac{1}{2}\sum_{i=1}^n \lambda_i\mathbb{E}[(a_i s_i+b_i-c_i p)^2] - \frac{\beta}{2}\mathbb{E}\left[\sum_{i=1}^n (a_i s_i+b_i-c_i p)\right]^2, \tag{27}$$

where once again we are using the market-clearing condition. Thus the social planner chooses the constants a_i, b_i, and c_i to maximize the total expected surplus in (27). We now determine the conditions under which the equilibrium strategies identified in Proposition 1 satisfy the first-order conditions corresponding to the planner's problem.

First, consider the planner's first-order condition with respect to coefficients $(b_1,\ldots,b_n)$. Differentiating (27) with respect to b_i and using the fact that the market-clearing price satisfies (6) implies that

$$\frac{d\mathbb{E}[W]}{db_i} = -\lambda_i b_i + \lambda_i c_i\mathbb{E}[p] - \frac{\alpha+\beta\sum_{k=1}^n b_k}{1+\beta\sum_{k=1}^n c_k} + \frac{\beta}{1+\beta\sum_{k=1}^n c_k}\sum_{k=1}^n c_k\left(\lambda_k b_k + (1-\lambda_k c_k)\mathbb{E}[p]\right). \tag{28}$$

On the other hand, recall from equation (14) that equilibrium coefficients satisfy $(\lambda_i c_i - 1)\mathbb{E}[p] = \lambda_i b_i$. Consequently, the first-order condition of the planner's problem with respect to b_i evaluated at the equilibrium strategies is given by

$$\left.\frac{d\mathbb{E}[W]}{db_i}\right|_{\text{eq}} = \mathbb{E}[p] - \frac{\alpha+\beta\sum_{k=1}^n b_k}{1+\beta\sum_{k=1}^n c_k}.$$

But note that (6) implies that the right-hand side of the above expression is equal to zero, thus implying that equilibrium strategies always satisfy the planner's first-order conditions with respect to b_i for all parameter values.

Next, consider the social planner's first-order condition with respect to coefficients $(c_1, \ldots, c_n)$. Differentiating (27) with respect to c_i leads to

$$\begin{aligned}\frac{d\mathbb{E}[W]}{dc_i} &= \lambda_i a_i \mathbb{E}[ps_i] + \lambda_i b_i \mathbb{E}[p] - \lambda_i c_i \mathbb{E}[p^2] - \mathbb{E}[\theta_i p] + \frac{\alpha + \beta\sum_{k=1}^n b_k}{1+\beta\sum_{k=1}^n c_k}\mathbb{E}[p] + \frac{\beta\sum_{k=1}^n a_k \mathbb{E}[s_k p]}{1+\beta\sum_{k=1}^n c_k} \\ &+ \frac{\beta}{1+\beta\sum_{k=1}^n c_k}\sum_{k=1}^n c_k\left(\mathbb{E}[\theta_k p] + (\lambda_k c_k - 1)\mathbb{E}[p^2] - \lambda_k a_k \mathbb{E}[ps_k] - \lambda_k b_k \mathbb{E}[p]\right),\end{aligned} \tag{29}$$

where once again we are using the fact that the market-clearing price satisfies (6). On the other hand, recall that equilibrium coefficients satisfy equation (15). Therefore, the first-order condition of the planner's problem with respect to c_i evaluated at equilibrium strategies is equal to

$$\left.\frac{d\mathbb{E}[W]}{dc_i}\right|_{\text{eq}} = -\mathbb{E}[p^2] + \frac{\alpha + \beta\sum_{k=1}^n b_k}{1+\beta\sum_{k=1}^n c_k}\mathbb{E}[p] + \frac{\beta\sum_{k=1}^n a_k \mathbb{E}[s_k p]}{1+\beta\sum_{k=1}^n c_k}.$$

Equation (6) then implies that right-hand side of the above equation is equal to zero. In other words, no matter the parameter values, the equilibrium strategies always satisfy the planner's first-order conditions with respect to $(c_1, \ldots, c_n)$.

Finally, we consider the planner's first-order condition with respect to $(a_1, \ldots, a_n)$. Differentiating (27) with respect to a_i and using the fact that the market-clearing price satisfies (6) implies that

$$\frac{d\mathbb{E}[W]}{da_i} = 1 - \lambda_i(1+\sigma_i^2)a_i + (\lambda_i c_i - 1)\mathbb{E}[ps_i] + \frac{\beta}{1+\beta\sum_{j=1}^n c_j}\sum_{k=1}^n c_k\left(\lambda_k a_k \mathbb{E}[s_i s_k] - \mathbb{E}[\theta_k s_i] + (1-\lambda_k c_k)\mathbb{E}[s_i p]\right). \tag{30}$$

Recall that we have already established that $d\mathbb{E}[W]/db_i = d\mathbb{E}[W]/dc_i = 0$ at the equilibrium strategies. Therefore, the equilibrium is constrained efficient if only if the above expression is equal to zero when evaluated at the equilibrium strategies. Furthermore, recall that equilibrium strategies satisfy equations (13)–(15). Hence, by (13), it is immediate that

$$\left.\frac{d\mathbb{E}[W]}{da_i}\right|_{\text{eq}} = \frac{\beta}{1+\beta\sum_{j=1}^n c_j}\sum_{k=1}^n c_k\left(\lambda_k a_k \mathbb{E}[s_i s_k] - \mathbb{E}[\theta_k s_i] + (1-\lambda_k c_k)\mathbb{E}[s_i p]\right),$$

which by using (13) one more time further simplifies to

$$\left.\frac{d\mathbb{E}[W]}{da_i}\right|_{\text{eq}} = \frac{\beta}{1+\beta\sum_{j=1}^n c_j}\sum_{k\neq i} c_k\left(\rho(\lambda_k a_k - 1) + (1-\lambda_k c_k)\mathbb{E}[s_i p]\right).$$

Replacing for coefficients a_i and c_i from equations (16) and (18) and using the fact that the market-clearing price satisfies (6) implies that

$$\left.\frac{d\mathbb{E}[W]}{da_i}\right|_{\text{eq}} = \frac{\beta\rho}{1+\beta\sum_{j=1}^n c_j}\sum_{k\neq i}\frac{c_k\sigma_k^2}{\delta_k}\sum_{j\neq k} a_j\Big(a_i(1-\rho+\sigma_i^2) - a_j(1-\rho+\sigma_j^2)\Big),$$

where δ_k is defined in (20). Thus, by Lemma B.2,

$$\left.\frac{d\mathbb{E}[W]}{da_i}\right|_{\text{eq}} = \frac{\rho}{1+\beta\sum_{r=1}^{n} 1/\lambda_r}\sum_{k\neq i}\frac{\sigma_k^2}{\delta_k\lambda_k}(\beta Q_k - M_k)\sum_{j\neq k} a_j\Big(a_i(1-\rho+\sigma_i^2) - a_j(1-\rho+\sigma_j^2)\Big). \tag{31}$$

We now use (31) to prove Proposition 5. As a first observation, note that when $\rho = 0$, the right-hand side of the above equation is equal to zero, thus implying that the equilibrium is constrained efficient for all profiles of trading costs. Next, consider the case that $n = 2$. With only two traders, it is immediate that the right-hand side of (31) is also equal to zero for all parameter values, thus once again implying constrained efficiency. To establish that the equilibrium is constrained efficient when all trading costs coincide, recall from Lemma B.1 that $\lambda_i = \lambda$ guarantees that $a_i(1-\rho+\sigma_i^2) = a_j(1-\rho+\sigma_j^2)$ for all i and j. Therefore, when all trading costs are identical, the right-hand side of (31) is equal to zero, thus guaranteeing constrained efficiency.

Finally, we show that as long as $n \geq 3$, trading costs are heterogeneous, and $\rho > 0$, the equilibrium is constrained inefficient for almost all values of β. We establish this by contradiction. Suppose there exist $\beta \neq \tilde{\beta}$ for which the equilibrium is constrained efficient. Hence, the right-hand side of (31) is equal to zero for both β and $\tilde{\beta}$ and all traders i. Since $\rho \neq 0$, this implies that

$$\sum_{k\neq i}\frac{\sigma_k^2}{\delta_k\lambda_k}(\beta Q_k - M_k)\sum_{j\neq k} a_j\Big(a_j(1-\rho+\sigma_j^2) - a_i(1-\rho+\sigma_i^2)\Big) = 0$$
$$\sum_{k\neq i}\frac{\sigma_k^2}{\delta_k\lambda_k}(\tilde{\beta} Q_k - M_k)\sum_{j\neq k} a_j\Big(a_j(1-\rho+\sigma_j^2) - a_i(1-\rho+\sigma_i^2)\Big) = 0,$$

where recall that the coefficients $(a_1, \ldots, a_n)$ are the solution to the fixed point equation (16) and hence are independent of the value of β. Subtracting the above two equations from one another and using the fact that $\beta \neq \tilde{\beta}$ leads to

$$\sum_{k\neq i}\frac{\sigma_k^2 M_k}{\delta_k\lambda_k}\sum_{j\neq k} a_j\Big(a_j(1-\rho+\sigma_j^2) - a_i(1-\rho+\sigma_i^2)\Big) = 0 \tag{32}$$

for all traders i. Since not all trading costs are identical, Lemma B.1 in the proof of Proposition 2 guarantees that there exists a i such that $a_i(1-\rho+\sigma_i^2) \leq a_j(1-\rho+\sigma_j^2)$ for all j, with at least one inequality holding strictly. But since $M_k > 0$, this means that the left-hand side of (32) has to be strictly negative, leading to a contradiction. □

Proof of Proposition 6

Recall that the total ex ante surplus in the market is given by $\mathbb{E}[W] = \mathbb{E}[\pi_0] + \sum_{i=1}^{n}\mathbb{E}[\pi_i]$, where π_0 is the surplus of the outside trader and π_i is the profit of trader i. For ease of notation, denote $\pi_i(x_i) = u_i(x_i) - px_i$ and $\pi_0 = u_0(y) - py$. Differentiating the total surplus with respect to a_i implies

$$\frac{d}{da_i}\mathbb{E}[W] = \sum_{k=1}^{n}\mathbb{E}\left[\left(\frac{\partial u_k}{\partial x_k} - p\right)\left(\frac{dx_k}{da_i} + \frac{\partial x_k}{\partial p}\frac{dp}{da_i}\right) - x_k\frac{dp}{da_i}\right] + \mathbb{E}\left[\left(\frac{\partial u_0}{\partial y} - p\right)\frac{dy}{da_i}\right] - \mathbb{E}\left[y\frac{dp}{da_i}\right]$$

Therefore, the market-clearing condition $y + \sum_{i=1}^n x_i = 0$ implies that

$$\frac{d}{da_i}\mathbb{E}[W] = \mathbb{E}\left[\left(\frac{\partial u_i}{\partial x_i} - p\right)\left(\frac{dx_i}{da_i}\right)\right] + \sum_{k=1}^n \mathbb{E}\left[\left(\frac{\partial u_k}{\partial x_k} - p\right)\left(\frac{\partial x_k}{\partial p}\frac{dp}{da_i}\right)\right] \tag{33}$$

Recall that in equilibrium the first-order condition is given by

$$\mathbb{E}\left[\left(\frac{du_i}{dx_i} - p\right)\frac{dx_i}{da_i}\right] = 0$$

Writing the same expression in the ex post form, results in

$$\mathbb{E}\left[\frac{\partial u_k}{\partial x_k}|s_k, p\right] - p = 0.$$

Therefor, plugging the equilibrium action in equation (33) results in

$$\left.\frac{d\mathbb{E}[W]}{da_i}\right|_{\text{eq}} = \sum_{k=1}^n \mathbb{E}\left[\left(\frac{\partial u_k}{\partial x_k} - \mathbb{E}\left[\left.\frac{\partial u_k}{\partial x_k}\right| s_k, p\right]\right)\left(\frac{\partial x_k}{\partial p}\frac{dp}{da_i}\right)\right]$$

By the law of iterated expectations,

$$\left.\frac{d\mathbb{E}[W]}{da_i}\right|_{\text{eq}} = \sum_{k=1}^n \mathbb{E}\left[\mathbb{E}\left[\left.\left(\frac{\partial u_k}{\partial x_k} - \mathbb{E}\left[\left.\frac{\partial u_k}{\partial x_k}\right| s_k, p\right]\right)\left(\frac{\partial x_k}{\partial p}\frac{dp}{da_i}\right)\right]\right| s_1, \ldots, s_n\right]$$

Note that all equilibrium variables have to be measurable with respect to the collection of all the signals in the market. Consequently, we have,

$$\left.\frac{d\mathbb{E}[W]}{da_i}\right|_{\text{eq}} = \sum_{k=1}^n \mathbb{E}\left[\frac{\partial x_k}{\partial p}\frac{dp}{da_i}\left(\mathbb{E}\left[\left.\frac{\partial u_k}{\partial x_k}\right| s_1, \ldots, s_n\right] - \mathbb{E}\left[\left.\frac{\partial u_k}{\partial x_k}\right| s_k, p\right]\right)\right]$$

Now, plugging back the marginal utility functions $u_k' = \theta_k - \lambda_k x_k$ and the linear strategies $x_i = a_i s_i - c_i p$, implies that

$$\left.\frac{d\mathbb{E}[W]}{da_i}\right|_{\text{eq}} = \sum_{k=1}^n \frac{\partial x_k}{\partial p}\mathbb{E}\left[\frac{dp}{da_i}\Big(\mathbb{E}\left[\theta_k | s_1, \ldots, s_n\right] - \mathbb{E}\left[\theta_k | s_k, p\right]\Big)\right]$$

From equation (6) we have $\partial p/\partial a_i = \gamma s_i$, where $\gamma = \beta/(1 + \beta\sum_{k=1}^n c_k)$. On the other hand, recall from equation (25) that $\gamma > 0$. Thus, by the law of iterated expectations,

$$\left.\frac{d\mathbb{E}[W]}{da_i}\right|_{\text{eq}} = \gamma \sum_{k=1}^n \frac{\partial x_k}{\partial p}\operatorname{cov}\big(s_i, \mathbb{E}\left[\theta_k | s_1, \ldots, s_n\right] - \mathbb{E}\left[\theta_k | s_k, p\right]\big).$$

Finally, using the law of iterated expectations one more time to establish that $\operatorname{cov}\left(s_i, \mathbb{E}\left[\theta_k | s_1, \ldots, s_n\right] - \mathbb{E}\left[\theta_k | s_k, p\right]\right) = \operatorname{cov}(s_i, \theta_k - \mathbb{E}\left[\theta_k | s_k, p\right])$ then completes the proof. □

Proof of Proposition 7

Proof of part (a) Recall from equation (31) that

$$\lim_{\beta \to 0}\left.\frac{d\mathbb{E}[W]}{da_i}\right|_{\text{eq}} = \rho^2 \sum_{k \neq i}\frac{\sigma_k^4}{\delta_k^2 \lambda_k}\left(\sum_{j \neq k} a_j\right)\left(\sum_{j \neq k} a_j\Big(a_j(1 - \rho + \sigma_j^2) - a_i(1 - \rho + \sigma_i^2)\Big)\right),$$

where δ_k is given by (20). The above expression therefore implies that

$$\lim_{\rho\to 0}\lim_{\beta\to 0}\frac{1}{\rho^2}\left.\frac{d\mathbb{E}[W]}{da_i}\right|_{\text{eq}} > 0 \tag{34}$$

if and only if

$$\lim_{\rho\to 0} a_i(1-\rho+\sigma_i^2) < \lim_{\rho\to 0}\frac{\sum_{k\neq i}\frac{\sigma_k^4}{\delta_k^2\lambda_k}(\sum_{j\neq k} a_j)(\sum_{j\neq k} a_j^2(1-\rho+\sigma_j^2))}{\sum_{k\neq i}\frac{\sigma_k^4}{\delta_k^2\lambda_k}(\sum_{j\neq k} a_j)^2}.$$

On the other hand, equation (13) implies that $\lim_{\rho\to 0}\lambda_i a_i = 1/(1+\sigma_i^2)$. Consequently, replacing for a_i in the above equation implies that inequality (34) holds if and only if

$$\frac{1}{\lambda_i} < \frac{\displaystyle\sum_{k\neq i}\frac{\sigma_k^4}{\lambda_k(1+\sigma_k^2)^2}\frac{\sum_{j\neq k}\frac{1}{\lambda_j(1+\sigma_j^2)}}{\sum_{j\neq k}\frac{1}{\lambda_j^2(1+\sigma_j^2)}}}{\displaystyle\sum_{k\neq i}\frac{\sigma_k^4}{\lambda_k(1+\sigma_k^2)^2}\left(\frac{\sum_{j\neq k}\frac{1}{\lambda_j(1+\sigma_j^2)}}{\sum_{j\neq k}\frac{1}{\lambda_j^2(1+\sigma_j^2)}}\right)^2}.$$

□

Proof of part (b) Next, consider the case that $\beta\to\infty$. In this case, we have

$$\lim_{\rho\to 0}\lim_{\beta\to \infty}\frac{1}{\rho}\left.\frac{d\mathbb{E}[W]}{da_i}\right|_{\text{eq}} = \frac{1}{\sum_{r=1}^n 1/\lambda_r}\sum_{k\neq i}\frac{\sigma_k^2}{\lambda_k(1+\sigma_k^2)}\left(\frac{1}{\lambda_i}\frac{\sum_{j\neq k}\frac{1}{\lambda_j(1+\sigma_j^2)}}{\sum_{j\neq k}\frac{1}{\lambda_j^2(1+\sigma_j^2)}}-1\right).$$

Therefore,

$$\lim_{\rho\to 0}\lim_{\beta\to \infty}\frac{1}{\rho}\left.\frac{d\mathbb{E}[W]}{da_i}\right|_{\text{eq}} > 0$$

if and only if

$$\frac{1}{\lambda_i} > \frac{\displaystyle\sum_{k\neq i}\frac{\sigma_k^2}{\lambda_k(1+\sigma_k^2)}}{\displaystyle\sum_{k\neq i}\frac{\sigma_k^2}{\lambda_k(1+\sigma_k^2)}\frac{\sum_{j\neq k}\frac{1}{\lambda_j(1+\sigma_j^2)}}{\sum_{j\neq k}\frac{1}{\lambda_j^2(1+\sigma_j^2)}}}$$

□

Proof of Proposition 8

Before presenting the proof, we state and prove two simple lemmas.

Lemma B.3. *Suppose* $\zeta_1,\dots,\zeta_m \geq 0$ *and* $\sum_{k=1}^m \zeta_k = 1$. *Then,*

$$\sum_{k=1}^m \frac{y_k}{\zeta_k+z_k} \geq \frac{(\sum_{k=1}^m\sqrt{y_k})^2}{1+\sum_{k=1}^m z_k} \tag{35}$$

for any collection of non-negative numbers $y_1,\dots,y_m$ *and* $z_1,\dots,z_m$.

Proof. We establish the lemma by showing that $\min_\zeta f(\zeta)$ subject to the constraint that $\sum_{k=1}^n \zeta_k = 1$ is equal to the right-hand side of (35), where $f(\zeta) = \sum_{k=1}^m y_k/(\zeta_k + z_k)$. First, note that $f(\zeta)$ is convex in ζ, thus implying that the first-order condition is a sufficient for optimality. This implies that $\eta = y_k/(\zeta_k + z_k)^2$, where η is the Lagrange multiplier corresponding to the constraint. Plugging this expression into the constraint implies that the optimal value of ζ_k is given by

$$\zeta_k = \left(\frac{1 + \sum_{j=1}^m z_j}{\sum_{j=1}^m \sqrt{y_j}}\right)\sqrt{y_k} - z_k.$$

Evaluating $f(\zeta)$ at the above values leads to the right-hand side of (35), thus completing the proof. □

Lemma B.4. *Suppose $\alpha = 0$. The equilibrium welfare in a market consisting of n traders is*

$$\mathbb{E}[W] = \sum_{i=1}^n \frac{1}{2\lambda_i}(1 - \mathrm{var}(\theta_i|s_i, p)) - \left(\frac{1}{2\beta} + \sum_{i=1}^n \frac{1}{2\lambda_i}\right)\mathbb{E}[p^2].$$

Proof. Recall from the proof of Proposition 5 that the expected welfare in the market is given by (26). Furthermore, note that the first-order condition of trader i is given by $x_i = (\mathbb{E}[\theta_i|s_i, p] - p)/\lambda_i$, whereas that of the outside trader is given by $y = (\alpha - p)/\beta$. Plugging these expressions into (26) therefore implies that

$$\mathbb{E}[W] = \sum_{i=1}^n \frac{1}{\lambda_i}\mathbb{E}[\mathbb{E}^2[\theta_i|s_i,p]] - \sum_{i=1}^n \frac{1}{\lambda_i}\mathbb{E}[\theta_i p] - \sum_{i=1}^n \frac{1}{2\lambda_i}\mathbb{E}[(\mathbb{E}[\theta_i|s_i,p] - p)^2] + \frac{\alpha}{\beta}\mathbb{E}[\alpha - p] - \frac{1}{2\beta}\mathbb{E}[(\alpha - p)^2].$$

Consequently,

$$\begin{aligned}
\mathbb{E}[W] &= \sum_{i=1}^n \frac{1}{2\lambda_i}\mathbb{E}[\mathbb{E}^2[\theta_i|s_i,p]] - \left(\frac{1}{2\beta} + \sum_{i=1}^n \frac{1}{2\lambda_i}\right)\mathbb{E}[p^2] + \frac{\alpha^2}{2\beta} \\
&= \sum_{i=1}^n \frac{1}{2\lambda_i}\left(\mathrm{var}(\theta_i) - \mathbb{E}[\mathrm{var}(\theta_i|s_i,p)]\right) - \left(\frac{1}{2\beta} + \sum_{i=1}^n \frac{1}{2\lambda_i}\right)\mathbb{E}[p^2] + \frac{\alpha^2}{2\beta},
\end{aligned}$$

where the second equality is a consequence of the fact that $\mathbb{E}[\mathbb{E}[\theta_i|s_i,p]] = \mathbb{E}[\theta_i] = 0$ and the law of total variance. Noting that $\mathrm{var}(\theta_i) = 1$ and $\mathbb{E}[\mathrm{var}(\theta_i|s_i,p)] = \mathrm{var}(\theta_i|s_i,p)$, which is a consequence of normality, and setting $\alpha = 0$ completes the proof. □

With the above lemmas in hand, we now proceed to proving Proposition 8.

Proof of part (a). Suppose traders face no uncertainties about their private valuations, i.e., $\sigma_i = 0$ for all i. This means that $\mathrm{var}(\theta_i|s_i,p) = 0$ for all traders regardless of the market structure. Thus, by Lemma B.4, expected welfare in the centralized market is given by

$$\mathbb{E}[W^{\mathrm{cen}}] = \sum_{i=1}^n \frac{1}{2\lambda_i} - \left(\frac{1}{2\beta} + \sum_{i=1}^n \frac{1}{2\lambda_i}\right)\mathbb{E}[p^2].$$

On the other hand, equations (16)–(18) imply that when $\sigma_i = 0$, equilibrium strategies satisfy $b_i = 0$ and $a_i = c_i = \lambda_i^{-1}$. Thus, by equation (6), the market clearing price in the centralized market is equal to $p = \beta \sum_{i=1}^n s_i \lambda_i^{-1}/(1+\beta\sum_{i=1}^n \lambda_i^{-1})$. Therefore,

$$\mathbb{E}[W^{\text{cen}}] = \sum_{i=1}^n \frac{1}{2\lambda_i} - \frac{\beta}{2}\left(\frac{(1-\rho)\sum_{i=1}^n 1/\lambda_i^2 + \rho(\sum_{i=1}^n 1/\lambda_i)^2}{1+\beta\sum_{i=1}^n 1/\lambda_i}\right). \tag{36}$$

Following similar steps implies that expected welfare in the segmented architecture is given by

$$\mathbb{E}[W^{\text{seg}}] = \sum_{i=1}^n \frac{1}{2\lambda_i} - \frac{\beta}{2}\sum_{S_k\in\mathcal{S}} \frac{(1-\rho)\sum_{i\in S_k} 1/\lambda_i^2 + \rho(\sum_{i\in S_k} 1/\lambda_i)^2}{\zeta_k + \beta\sum_{i\in S_k} 1/\lambda_i},$$

where S_k denotes the set of traders in the k-th segment and ζ_k is the fraction of outside traders that are active in that segment. Applying Lemma B.3 to the second term on the right-hand side above and noting that $\sum_{S_k\in\mathcal{S}}\zeta_k = 1$ leads to

$$\mathbb{E}[W^{\text{seg}}] \leq \sum_{i=1}^n \frac{1}{2\lambda_i} - \frac{\beta/2}{1+\beta\sum_{i=1}^n 1/\lambda_i}\left(\sum_{S_k\in\mathcal{S}}\sqrt{(1-\rho)\sum_{i\in S_k} 1/\lambda_i^2 + \rho\Big(\sum_{i\in S_k} 1/\lambda_i\Big)^2}\right)^2,$$

which in turn implies that

$$\mathbb{E}[W^{\text{seg}}] \leq \sum_{i=1}^n \frac{1}{2\lambda_i} - \frac{\beta/2}{1+\beta\sum_{i=1}^n 1/\lambda_i}\left((1-\rho)\sum_{i=1}^n \frac{1}{\lambda_i^2} + \rho\sum_{S_k\in\mathcal{S}}\Big(\sum_{i\in S_k}\frac{1}{\lambda_i}\Big)^2 + \rho\sum_{S_k\neq S_j}\Big(\sum_{i\in S_k}\frac{1}{\lambda_i}\Big)\Big(\sum_{i\in S_j}\frac{1}{\lambda_i}\Big)\right).$$

Equation (36) implies that the right-hand side of the above inequality coincides with $\mathbb{E}[W^{\text{cen}}]$, thus establishing that expected welfare is weakly higher in the centralized market structure. □

Proof of part (b). Suppose $\rho = 0$. This means that $\text{var}(\theta_i|s_i,p) = \text{var}(\theta_i|s_i) = \sigma_i^2/(1+\sigma_i^2)$ regardless of the market structure. Consequently, Lemma B.4 implies that expected welfare in the centralized architecture is equal to

$$\mathbb{E}[W^{\text{cen}}] = \sum_{i=1}^n \frac{1}{2\lambda_i(1+\sigma_i^2)} - \left(\frac{1}{2\beta} + \sum_{i=1}^n \frac{1}{2\lambda_i}\right)\mathbb{E}[p^2].$$

Equations (16)–(18) imply that when $\rho = 0$, the coefficients corresponding to equilibrium strategies satisfy $a_i = \lambda_i^{-1}/(1+\sigma_i^2)$, $b_i = 0$, and $c_i = \lambda_i^{-1}$. Replacing for p from (6) leads to

$$\mathbb{E}[W^{\text{cen}}] = \sum_{i=1}^n \frac{1}{2\lambda_i(1+\sigma_i^2)} - \frac{\beta\sum_{i=1}^n \frac{1}{\lambda_i^2(1+\sigma_i^2)}}{2(1+\beta\sum_{i=1}^n \lambda_i^{-1})}. \tag{37}$$

Following similar steps for the segmented market structure implies that

$$\mathbb{E}[W^{\text{seg}}] = \sum_{i=1}^n \frac{1}{2\lambda_i(1+\sigma_i^2)} - \sum_{S_k\in\mathcal{S}} \frac{\beta\sum_{i\in S_k} \frac{1}{\lambda_i^2(1+\sigma_i^2)}}{2(\zeta_k+\beta\sum_{i\in S_k} \lambda_i^{-1})},$$

and as a result,

$$\mathbb{E}[W^{\text{seg}}] \leq \sum_{i=1}^n \frac{1}{2\lambda_i(1+\sigma_i^2)} - \sum_{S_k\in\mathcal{S}} \frac{\beta\sum_{i\in S_k} \frac{1}{\lambda_i^2(1+\sigma_i^2)}}{2(1+\beta\sum_{i=1}^n \lambda_i^{-1})}.$$

Note that, by (37), the right-hand side of the above inequality is equal to $\mathbb{E}[W^{\text{cen}}]$, thus implying that expected welfare in the centralized architecture is higher than the segmented architecture. □

Proof of Proposition 9

First consider the centralized market architecture. By Lemma B.4,

$$\lim_{\beta\to 0}\mathbb{E}[W^{\text{cen}}]=\sum_{i=1}^{n}\frac{1}{2\lambda_i}(1-\operatorname{var}(\theta_i|s_i,p)),$$

where we are using the fact that, by equation (6), $\lim_{\beta\to 0}p^2/\beta=0$. Replacing for $\operatorname{var}(\theta_i|s_i,p)$ in terms of the information revelation gap defined in (7) leads to

$$\lim_{\beta\to 0}\mathbb{E}[W^{\text{cen}}]=\sum_{i=1}^{n}\frac{1}{2\lambda_i}\left(1-\phi_i^{\text{cen}}\operatorname{var}(\theta_i|s_i)-(1-\phi_i^{\text{cen}})\operatorname{var}(\theta_i|s_1,\dots,s_n)\right).$$

On the other hand, recall that $\operatorname{var}(\theta_i|s_i)=\sigma_i^2/(1+\sigma_i^2)$, whereas (21) implies that $\operatorname{var}(\theta_i|s_1,\dots,s_n)=\frac{\sigma_i^2}{1+\sigma_i^2}-\frac{\rho^2\sigma_i^4}{(1+\sigma_i^2)^2}\sum_{j\neq i}(1+\sigma_j^2)^{-1}+o(\rho^2)$. Consequently,

$$\lim_{\beta\to 0}\mathbb{E}[W^{\text{cen}}]=\sum_{i=1}^{n}\frac{1}{2\lambda_i}\left(1-\phi_i^{\text{cen}}\frac{\sigma_i^2}{1+\sigma_i^2}-(1-\phi_i^{\text{cen}})\left(\frac{\sigma_i^2}{1+\sigma_i^2}-\frac{\rho^2\sigma_i^4}{(1+\sigma_i^2)^2}\sum_{j\neq i}(1+\sigma_j^2)^{-1}\right)\right)+o(\rho^2).$$

Following similar steps for the segmented market structure implies that

$$\lim_{\beta\to 0}\mathbb{E}[W^{\text{seg}}]=\sum_{S_k\in\mathcal{S}}\sum_{i\in S_k}\frac{1}{2\lambda_i}\left(1-\phi_i^{\text{seg}}\frac{\sigma_i^2}{1+\sigma_i^2}-(1-\phi_i^{\text{seg}})\left(\frac{\sigma_i^2}{1+\sigma_i^2}-\frac{\rho^2\sigma_i^4}{(1+\sigma_i^2)^2}\sum_{\substack{j\in S_k\\ j\neq i}}(1+\sigma_j^2)^{-1}\right)\right)+o(\rho^2),$$

where ϕ_i^{seg} is trader i's information revelation gap in the segmented market structure. Subtracting the above two equations from one another implies that

$$\lim_{\beta\to 0}\left(\mathbb{E}[W^{\text{cen}}]-\mathbb{E}[W^{\text{seg}}]\right)=\rho^2\sum_{S_k\in\mathcal{S}}\sum_{i\in S_k}\frac{\sigma_i^4}{2\lambda_i(1+\sigma_i^2)^2}\left((1-\phi_i^{\text{cen}})\sum_{j\neq i}\frac{1}{1+\sigma_j^2}-(1-\phi_i^{\text{seg}})\sum_{\substack{j\in S_k\\ j\neq i}}\frac{1}{1+\sigma_j^2}\right)+o(\rho^2),$$

thus completing the proof. □

Proof of Proposition A.1

Proof of part (a) The proof of part (a) is similar to that of Proposition 3 of Rostek and Weretka (2012). First suppose that $\rho_{ij}=\rho$ for all $i\neq j$. Since all trading costs coincide, then Proposition 2 guarantees that the equilibrium is fully privately revealing to all traders simultaneously.

To prove the converse implication, suppose the price is fully privately revealing to all traders. That is, $\mathbb{E}[\theta_i|s_i,p]=\mathbb{E}[\theta_i|s_1,\dots,s_n]$ for all i. In addition, recall that when traders follow linear strategies in the form of $x_i=a_is_i+b_i-c_ip$, the corresponding coefficients satisfy (13)–(15). Consequently,

$$\lambda a_i=\frac{\operatorname{var}(p)-\mathbb{E}[ps_i]\mathbb{E}[p\theta_i]}{(1+\sigma^2)\operatorname{var}(p)-\mathbb{E}^2[ps_i]},$$

where we are using the fact that all traders have identical trading costs and signal precisions. Also recall that the market-clearing price satisfies (6). Replacing for the price in the above expression therefore implies that coefficients $(a_1, \dots, a_n)$ are the solution to the following system of equations:

$$\lambda a_i = \frac{\sum_{k\neq i} a_k^2(1+\sigma^2) + \sum_{j,k\neq i} \rho_{kj} a_k a_j - (\sum_{k\neq i} \rho_{ik} a_k)^2 - a_i\sigma^2 \sum_{k\neq i} \rho_{ik} a_k}{(1+\sigma^2)\sum_{k\neq i} a_k^2(1+\sigma^2) + (1+\sigma^2)\sum_{j,k\neq i} \rho_{kj} a_k a_j - (\sum_{k\neq i} \rho_{ik} a_k)^2}. \tag{38}$$

It is easy to verify that the solution to the above system of equations is given by

$$a_i = \frac{1-\bar{\rho}}{\lambda(1-\bar{\rho}+\sigma^2)}, \tag{39}$$

where $\bar{\rho}$ is defined (12). Since $a_i = a_j$ for all pairs of traders i and j, equation (6) implies that the price is a sufficient statistic for the unweighted average of traders' signals, namely, $(1/n)\sum_{k=1}^n s_k$. Therefore,

$$\mathbb{E}[\theta_i|s_1,\dots,s_n] = \mathbb{E}[\theta_i|s_i,p] = \left(\frac{1-\bar{\rho}}{1-\bar{\rho}+\sigma^2}\right) s_i + \frac{\bar{\rho}\sigma^2}{(1-\bar{\rho}+\sigma^2)(1+\sigma^2+\bar{\rho}(n-1))}\sum_{k=1}^n s_k,$$

where we are using the fact that the equilibrium is fully privately revealing to trader i. Consequently,

$$\mathbb{E}[\theta_i s_j] = \left(\frac{1-\bar{\rho}}{1-\bar{\rho}+\sigma^2}\right)\rho_{ij} + \frac{\bar{\rho}\sigma^2}{(1-\bar{\rho}+\sigma^2)(1+\sigma^2+\bar{\rho}(n-1))}\left(1+\sigma^2+\sum_{k\neq j}\rho_{jk}\right)$$

for any $j \neq i$. Replacing the left-hand side of the above equation with ρ_{ij} and noting that $\sum_{k\neq j}\rho_{jk} = (n-1)\bar{\rho}$ implies that the above equality is satisfied for all $i \neq j$ only if $\rho_{ij} = \bar{\rho}$ for all pairs of traders $i \neq j$. ☐

Proof of part (b) Recall from the proof of Proposition 1 that equilibrium strategies satisfy equations (13)–(15). Furthermore, recall from the proof of Proposition 5 that the first-order conditions of the planner's problem with respect to coefficients a_i, b_i, and c_i are given by (30), (28), and (29), respectively. As in the proof of Proposition 5, it is immediate to verify that, as long as (14) is satisfied, the right-hand side of (28) is equal to zero, thus implying that equilibrium strategies satisfy the planner's first-order condition with respect to b_i. Similarly, using (15) to simplify (29) implies that the right-hand side of the latter equation is also equal to zero for all parameter values, which establishes that equilibrium strategies satisfy the planner's first-order condition with respect to c_i.

Having established $d\mathbb{E}[W]/db_i = d\mathbb{E}[W]/dc_i = 0$ for all i, it is therefore sufficient to verify that the right-hand side of (30), when evaluated at equilibrium strategies, is equal to zero. The fact that equilibrium strategies satisfy (13) implies that

$$\left.\frac{d\mathbb{E}[W]}{da_i}\right|_{\text{eq}} = \frac{\beta}{1+\beta\sum_{j=1}^n c_j}\sum_{k\neq i} c_k\left(\rho_{ik}(\lambda a_k - 1) + (1-\lambda(1+\sigma^2)a_k)\frac{\mathbb{E}[s_i p]}{\mathbb{E}[s_k p]}\right),$$

where we are using the fact that all traders have identical trading costs and signal precisions. Plugging for equilibrium actions from (39) and noting that equilibrium strategies are symmetric lead to

$$\left.\frac{d\mathbb{E}[W]}{da_i}\right|_{\text{eq}} = \frac{\beta}{1+n\beta c}\frac{c\sigma^2}{1-\bar{\rho}+\sigma^2}\sum_{k\neq i}\left(\bar{\rho}\left(\frac{1+\sigma^2+\sum_{j\neq i}\rho_{ij}}{1+\sigma^2+\sum_{j\neq i}\rho_{jk}}\right) - \rho_{ik}\right) = \frac{\beta}{1+n\beta c}\frac{c\sigma^2}{1-\bar{\rho}+\sigma^2}\sum_{k\neq i}(\bar{\rho}-\rho_{ik}).$$

The definition of $\bar{\rho}$ in (12) now guarantees that the right-hand side of the above equality is equal to zero, thus completing the proof. ☐

References

Acemoglu, Daron, Ali Makhdoumi, Azarakhsh Malekian, and Asu Ozdaglar (2017), "Fast and slow learning from reviews." Working paper.

Allon, Gad, Achal Bassamboo, and Itai Gurvich (2011), ""We will be right with you": Managing customer expectations with vague promises and cheap talk." *Operations Research*, 59, 1382–1394.

Anand, Krishnan S. and Manu Goyal (2009), "Strategic information management under leakage in a supply chain." *Management Science*, 55, 438–452.

Angeletos, George-Marios and Alessandro Pavan (2007), "Efficient use of information and social value of information." *Econometrica*, 75, 1103–1142.

Angeletos, George-Marios and Alessandro Pavan (2009), "Policy with dispersed information." *Journal of European Economic Association*, 7, 11–60.

Bergemann, Dirk and Juuso Välimäki (1997), "Market diffusion with two-sided learning." *The RAND Journal of Economics*, 773–795.

Besbes, Omar and Marco Scarsini (2016), "On information distortions in online ratings." Working paper.

Bimpikis, Kostas, Shayan Ehsani, and Mohamed Mostagir (2018), "Designing dynamic contests." *Operations Research*. Forthcoming.

Candogan, Ozan and Kimon Drakopoulos (2017), "Optimal signaling of content accuracy: Engagement vs. misinformation." Working paper.

Crossland, Jarrod, Bin Li, and Eduardo Roca (2013), "Is the European Union Emissions Trading Scheme (EU ETS) informationally efficient? Evidence from momentum-based trading strategies." *Applied Energy*, 109, 10–23.

Diamond, Douglas W. and R. E. Verrecchia (1981), "Information aggregation in a noisy rational expectations economy." *Journal of Financial Economics*, 9, 221–235.

European Commission (2009), "EU Emissions Trading Scheme (ETS) – consultation on design and organisation of emissions allowance auctions.", URL https://ec.europa.eu/clima/sites/clima/files/docs/0002/trader/edf_trading_en.pdf.

Grossman, Sanford J. (1981), "An introduction to the theory of rational expectations under asymmetric information." *Review of Economic Studies*, 48, 541–559.

Grossman, Sanford J. and Joseph E. Stiglitz (1980), "On the impossibility of informationally efficient markets." *American Economic Review*, 70, 393–408.

Guo, Pengfei and Paul Zipkin (2007), "Analysis and comparison of queues with different levels of delay information." *Manufacturing and Service Operations Management*, 53, 962–970.

Hellwig, Martin F. (1980), “On the aggregation of information in competitive markets.” *Journal of Economic Theory*, 22, 477–498.

Ifrach, Bar, Costis Maglaras, Marco Scarsini, and Anna Zseleva (2018), “Bayesian social learning from consumer reviews.” Working paper.

International Organization of Securities Commissions (2010), “Issues raised by dark liquidity.”, URL https://www.iosco.org/library/pubdocs/pdf/IOSCOPD336.pdf. A consultation report by the Technical Committee of the International Organization of Securities Commissions, CR05/10.

Iyer, Krishnamurthy, Ramesh Johari, and Ciamac C. Moallemi (2014), “Information aggregation and allocative efficiency in smooth markets.” *Management Science*, 60, 2509–2524.

Iyer, Krishnamurthy, Ramesh Johari, and Ciamac C. Moallemi (2018), “Welfare analysis of dark pools.” Working paper.

Jordan, James S. (1983), “On the efficient markets hypothesis.” *Econometrica*, 51, 1325–1343.

Jouini, Oualid, Zeynep Aksin, and Yves Dallery (2011), “Call centers with delay information: Models and insights.” *Management Science*, 13, 534–548.

Kong, Guangwen, Sampath Rajagopalan, and Hau Zhang (2013), “Revenue sharing and information leakage in a supply chain.” *Management Science*, 59, 556–572.

Kyle, Albert S (1989), “Informed speculation with imperfect competition.” *Review of Economic Studies*, 56, 317–355.

Li, Lode (2002), “Information sharing in a supply chain with horizontal competition.” *Management Science*, 48, 1196–1212.

Malamud, Semyon and Marzena Rostek (2017), “Decentralized exchange.” *American Economic Review*, 107, 3320–3362.

Morris, Stephen and Hyun Song Shin (2002), “Social value of public information.” *American Economic Review*, 92, 1521–1534.

Papanastasiou, Yiangos, Kostas Bimpikis, and Nicos Savva (2018), “Crowdsourcing exploration.” *Management Science*, 64, 1727–1746.

Pesendorfer, Wolfgang and Jeroen M. Swinkels (2000), “Efficiency and information aggregation in auctions.” *American Economic Review*, 90, 499–525.

Rostek, Marzena and Marek Weretka (2012), “Price inference in small markets.” *Econometrica*, 80, 687–711.

Slechten, Aurélie and Estelle Cantillon (2015), “Price formation in the european carbon market: The role of firm participation and market structure.” Working paper.

Sun, Monic (2012), “How does the variance of product ratings matter?” *Management Science*, 58, 696–707.

U.S. Securities and Exchange Commission (2010), “Statement at open meeting and dissent regarding the adoption of amendments to regulation SHO (the ”alternative uptick rule”).”, URL https://www.sec.gov/news/speech/2010/spch022410tap-shortsales.htm. Speech by SEC Commissioner Troy A. Paredes.

Vives, Xavier (2011), “Strategic supply function competition with private information.” *Econometrica*, 79, 1919–1966.