Arms Sales in Financial Markets

Vincent Glode Xingtan Zhang *

August 8, 2023

Abstract

Many financial transactions are of a fixed-sum nature, meaning that improvements in the terms of trade for one party come at the expense of another party. We model how the sales of trading advantages (e.g., data or co-location services) affect traders' endogenous participation in financial markets and vice-versa. Sellers of trading advantages (e.g., data providers or securities exchanges) maximize profits by charging prices that may lead to inefficiently low levels of market participation and liquidity in equilibrium. Optimal sales of trading advantages lead less sophisticated investors to conclude that financial markets are too "rigged" and to exit these markets.

JEL Codes: G20, G14, D82, D42.

*We thank Markus Baldauf, Jeremy Bertomeu, Itay Goldstein, Jonathan Lewellen, Joshua Mollner, Greg Nini, Liyan Yang, Yao Zeng, and seminar participants at the SFS-Cavalcade (Asia), PHBS Workshop in Macroeconomics and Finance, HKUST, Rice, UT-Austin, Wash-U, and Wharton for their helpful comments.
Glode is at the Wharton School, University of Pennsylvania, 3620 Locust Walk, Philadelphia, PA 19104 and can be reached by e-mail at vglode@wharton.upenn.edu. Zhang is at the Leeds School of Business, University of Colorado at Boulder, 995 Regent Drive, Boulder, CO 80309 and can be reached by e-mail at xingtan.zhang@colorado.edu.

1 Introduction

> *By the summer of 2013, the world's financial markets were designed to maximize the number of collisions between ordinary investors and high-frequency traders — at the expense of ordinary investors, and for the benefit of high-frequency traders, exchanges, Wall Street banks, and online brokerage firms. Around those collisions an entire ecosystem had arisen.*
> – Michael Lewis, Flash Boys (p.179)

Many financial transactions are of a fixed-sum nature, meaning that any improvement in the terms of trade for one party comes at the expense of another party. This feature of financial markets has been shown to promote "arms races", that is, the inefficient acquisition of resources aimed at gaining a relative advantage over rivals and appropriating their surplus.[1] While large financial institutions spend astronomical sums on data and co-location services hoping to take advantage of their counterparties, 41% of individuals who do not participate in financial markets blame the fact that these markets are "rigged" against them.[2] Accordingly, the ecosystem that collects revenues by providing goods and services that benefit a subset of traders at the expense of their counterparties must account for the fixed-sum nature of trading. By providing a substantial advantage to too many traders, a data provider or a securities exchange might push unsophisticated investors to exit the market, thereby reducing the value of the advantage being purchased.

We propose a model to study how the sales of goods and services that impose negative externalities on counterparties affect financial market outcomes, including market participation and volume. In our model, agents differ in their probability of being able to supply liquidity to counterparties. The likelier a market participant is to be asked to supply liquidity, the more valuable gaining a "trading advantage" through superior data and co-location

[1] See, e.g., Glode, Green, and Lowery (2012), Biais, Foucault, and Moinas (2015), Budish, Cramton, and Shim (2015), Glode and Lowery (2016), and Glode and Ordoñez (2022).

[2] See Royal, James. (March 24, 2021.) "Survey: More than half of investors think the stock market is rigged against individuals." Bankrate.

is. Each agent chooses whether to participate in the market and whether to acquire a trading advantage at a price chosen by a monopolist. We focus our analysis on the interactions among agents' endogenous participation, the optimal pricing of the trading advantage, and the provision of liquidity.

We show how a monopolist may maximize its profit by setting the price of the trading advantage higher than in the standard monopolist problem without externalities. The endogenous market participation of traders that choose not to acquire an advantage creates a second elasticity that the monopolist must consider, in addition to the demand elasticity typically featured in the classic monopoly pricing problem. When lowering the price of the advantage, a monopolist increases the quantity demanded but also makes financial markets appear more "rigged" to traders that do not acquire this advantage. As a result, these unsophisticated traders might decide to exit the market, thereby reducing the frequency at which those acquiring this advantage get to trade. With fewer traders that demand liquidity participating in the market, an advantage that can be used when supplying liquidity becomes less valuable.

Yet, we show that the monopolist's optimal pricing strategy can result in socially excessive sales of goods and services that benefit a subset of market participants, leaving a fraction of the potential surplus from trade unrealized due to the non-participation of their disadvantaged counterparties. The welfare losses caused by these excessive sales increase in the magnitude of the negative externalities associated with the trading advantage as well as in the share of the surplus that liquidity suppliers can extract when trading. Moreover, the excessive sales of goods and services such as data and co-location services affect agents differently. When the negative externalities are small, the most likely liquidity suppliers benefit from the sales of trading advantages whereas the most likely liquidity demanders are harmed by them. When the negative externalities are large, however, all traders are made worse off by the resulting high exit rates of market participants in equilibrium. This case arises because the monopolist maximizes its profit by pricing the advantage low enough to have socially costly exit levels that destroy the surplus collected by *all* traders in equilibrium. Altogether, our model highlights how the magnitude of the pecuniary externality a trading advantage imposes on counterparties is an important determinant of various market outcomes.

Literature review. Our paper contributes to the literature on arms races in financial markets. Glode, Green, and Lowery (2012), Biais, Foucault, and Moinas (2015), Budish, Cramton, and Shim (2015), Glode and Lowery (2016), and Glode and Ordoñez (2022) show that financial firms may have incentives to overinvest (from a social standpoint) in goods and services such as information, expertise, and fast-trading technology that allow their acquirers to take advantage of counterparties. Our paper instead focuses on the pricing of resources that provide a trading advantage to a subset of market participants, from the perspective of a monopolist (e.g., a data provider or securities exchange) that must account for the impact of its decisions on market participation and liquidity in order to mazimize its profit.

Our insights on the decreasing returns to a trading advantage relate our paper to the literature on information acquisition in financial markets. In models with noise traders or a noisy asset supply such as Grossman and Stiglitz (1980), Admati and Pfleiderer (1986), and Admati and Pfleiderer (1990) just to name a few, the value of information decreases as more traders acquire it and as asset prices become more informative.[3] In those models, the volume of transactions and the available surplus from trade are considered exogenous by traders deciding to acquire information. Our paper instead focuses on how the sales

[3]Other models of information sales that assume exogenously noisy trading include Garcia and Sangiorgi (2011) who show that how traders compete for financial assets affects the optimal design of the signals offered for sales, Han and Yang (2013) who show that social interactions reduce the incentives to acquire information (due to a free-riding problem), and Huang, Xiong, and Yang (2022) who show that investors' ability to acquire skills needed to analyze data affects optimal data sales.

of trading advantages endogenously reduce trader participation in financial markets — the more rigged a market appears to be, the less attractive it is for unsophisticated investors (i.e., traders that do not possess these trading advantages). Moreover, the fewer investors agree to participate in the market, the less valuable acquiring a trading advantage is, thereby resulting in a feedback loop between market participation and the sales of trading advantages. We show how the seller of a trading advantage chooses its optimal pricing strategy in light of its effect on market participation and liquidity, which drive how much traders are willing to pay for such advantage.

This focus on optimal pricing strategies also distinguishes our paper from the recent literature aimed at explaining trends in stock market participation and active/index investing. When traders are exogenously endowed with heterogenous skill levels, Chaderina and Green (2014) show that the fixed-sum nature of trading can lead to participation cycles where exogenous negative shocks to the available trading profits can result in low-skill traders exiting financial markets. Stambaugh (2014) shows how exogenous time-variations in the amount of noise trading (as proxied by individual equity ownership) can explain, in an equilibrium model, the observed shift of assets under management from active to passive investing. Also using noisy rational expectations models, Peress (2005) and Bond and Garcia (2022) study how exogenous declines in participation and indexing costs impact market participation, the choice to invest passively or actively, and price informativeness. Our paper instead focuses on how market participation and liquidity respond to the price a profit-maximizing entity chooses to charge for a trading advantage whose value depends on market participation and liquidity.

Our paper also contributes to our understanding of the role of exchanges in the good functioning of financial markets. The early literature treats exchanges as a passive entity where trading takes place (see, e.g., Arrow 1951, Debreu 1951). Recently, many papers have studied the consequences of exchanges offering new products that solve market incompletness or that facilitate speculative activities (see, e.g., Allen and Gale 1994, Simsek 2013). Budish, Lee, and Shim (2022) consider a market design adoption game where exchanges choose whether to allow for continuous- or discrete-time trading. By allowing for continuous trading, exchanges gain market power in the sales of speed advantages. Budish, Lee, and Shim (2022) show that inefficient market designs can persist in equilibrium, due to exchanges benefitting from the arms race in speed. Our paper differs from this literature through its focus on how an entity's business model of selling trading advantages affects the market participation of the trading counterparties that are being taken advantage of, which then feeds back into the entity's profit-maximization decision.

Our paper also contributes to the literature on the optimal design of auctions for goods with externalities (e.g., Jéhiel, Moldovanu, and Stacchetti 1996, Eső, Nocke, and White 2010). Unlike in these papers, the monopolist's pricing decision we study affects the quantity of goods or services produced and, as a result, the magnitude of the externalities imposed on the subset of agents that do not acquire them. Furthermore, we show how the endogenous participation channel affects the optimal pricing, which is a channel that matters in financial markets according to the survey evidence mentioned earlier, but arguably not in many of the settings considered by the existing literature — for example, a country troubled by a neighbor arming up cannot simply "opt out" of a war if attacked. Similar to our approach with quantities in mind, the mechanism design literature (see Segal 1999, Segal 2003, among others) considers the general contracting problem between a principal and multiple agents when allocations can impose externalities. Our setting instead focuses on the interaction between a seller with market power and multiple potential buyers. The full characterization of the monopolist's profit function allows our model to generate sharp predictions about the roles played by the underlying distribution of traders' types, the monopolist's knowledge of agents' types, and the degree of competition that are all missing from more abstract mechanism-design frameworks.

Roadmap. In the next section, we introduce an environment with heterogenous prospective traders and a monopolist offering them access to a trading advantage. In Section 3, we analyze agents' decisions whether to participate in a financial market and, if so, whether to acquire a trading advantage. In Section 4, we analyze the monopolist's optimal pricing of the trading advantage given prospective traders' equilibrium responses. In Section 5, we study how varying the market structure can affect equilibrium behaviors by prospective traders and by the entity selling the trading advantage. The last section concludes. Proofs of formal results are relegated to the Appendix.

2 Model

We develop a simple model to study how a monopolist (e.g., a data provider or securities exchange) maximizes its profits from providing a good or service (e.g., data or co-location services) that improves the position of a subset of traders at the expense of their counterparties.[4] While most of the literature with information sales/acquisition analyzes trading environments with risk-averse traders and exogenous noise trading, we abstract from these frictions in order to make the fixed-sum nature of financial markets and its impact on endogenous trader participation as transparent as possible.

Environment. Consider a financial market with a continuum of potential participants (i.e., traders). Each agent i has a type $\theta_i \in [0, 1]$ that denotes the probability that this agent would be in position to supply liquidity to the market at $t = 1$ if it decided to participate in the market at $t = 0$. With probability $(1 - \theta_i)$, however, this agent would be hit by a liquidity

[4] Trading advantages that improve the position of a subset of traders at the expense of their counterparties are prevalent in financial markets. While we use data and co-location services as primary examples, other examples include providing an early peek to the Michigan Index of Consumer Sentiment (Hu, Pan, and Wang (2017)), sharing information with a selected set of clients within a broker's network (Di Maggio, Franzoni, Kermani, and Sommavilla (2019)), and giving trading priority within a dark pool (see https://www.sec.gov/news/press-release/2016-16 for the anecdotal evidence).

shock after entering the market and would demand liquidity from market participants in position to supply it (e.g., by exchanging an illiquid asset for cash). We assume that types θ_i are independent and identically distributed across agents. The cumulative distribution function (CDF) of types θ_i is denoted $F : [\underline{\theta}, \overline{\theta}] \to [0, 1]$, where $\underline{\theta} \geq 0$ and $\overline{\theta} \leq 1$, and the probability density function (PDF) is denoted f. We use $\mu \equiv \int_{\underline{\theta}}^{\overline{\theta}} \theta dF(\theta)$ to denote the population mean of θ_i.

When an agent participates in the financial market, it generates a surplus of Δ (thanks to unmodeled diversification benefits, for example). If this agent is not hit by a liquidity shock, it retains the entire surplus Δ. However, when hit by a liquidity shock, this agent might need to share its surplus with a liquidity supplier. Indeed, we assume that being asked to supply liquidity by a counterparty desperate for a transaction allows a liquidity supplier i to extract a payoff $\omega \cdot \Delta + \sigma \cdot a_i$ from the liquidity demander if a transaction occurs at $t = 1$. As we further explain below, the term $a_i \in \{0, 1\}$ denotes whether agent i acquired a trading advantage at $t = 0$, which boosts the payoff from supplying liquidity at $t = 1$. As a result of having to share surplus, agent i's liquidity demander only retains a payoff $\Delta - [\omega \cdot \Delta + \sigma \cdot a_i] = (1 - \omega) \cdot \Delta - \sigma \cdot a_i$ at $t = 1$.[5]

At $t = 0$, the monopolist chooses a price p to charge for a trading advantage and then each agent chooses whether to participate in the financial market and, if so, whether to acquire the trading advantage at that price. For simplicity, we assume that the *marginal* cost of producing a trading advantage such as sharing data already collected or providing co-location access to one more trader is zero.

While we interpret $(1 - \theta_i)$ as a probability of needing liquidity at $t = 1$, our insights generalize to other factors driving the need for some agents to transact with counterparties. For example, a low θ_i could capture an inflexible balance sheet that forces trader i to search

[5]See Glode, Green, and Lowery (2012), Glode and Lowery (2016), and Glode and Ordoñez (2022) for various micro-foundations for these types of surplus-sharing functions.

for a counterparty willing to bargain over a more efficient allocation of inventory risk. In fact, one can think of agents with relatively high θ_i as market makers and hedge funds, which tend to be able to take advantage of profitable transactions when they arise, whereas agents with relatively low θ_i can be thought of retail investors and pension plans, whose inflexibility and trading needs tend to create profitable opportunities for their counterparties. As a result of these interpretations, we assume that agents choose at $t = 0$ whether to participate in the financial market and whether to acquire a trading advantage fully knowing their own types. Yet, the uncertainty embedded in an agent's type about whether it will supply or demand liquidity at $t = 1$ captures the possibility that financial market conditions might change between when agents decide to participate and acquire trading advantages such as data and co-location services and when they actually engage in financial transactions.[6]

Matching of traders. We assume that the matching of liquidity suppliers and demanders is random and that, as a whole, liquidity suppliers can fulfill the needs of as many liquidity demanders as there are. Specifically, each liquidity demander is randomly matched to a liquidity supplier with probability 1 and each liquidity supplier ends up participating in a measure η of financial transactions, which is determined in equilibrium based on all agents' decisions at $t = 0$ on whether to participate in the market.

Traders' payoff. Prior to knowing whether it will supply or demand liquidity at $t = 1$, agent i expects to collect a payoff:

$$(1 - \theta_i)\mathbb{E}[(1 - \omega)\Delta - \sigma a_{-i}] + \theta_i[\Delta + \eta(\omega\Delta + \sigma a_i)], \tag{1}$$

where a_{-i} denotes the equilibrium strategies of agent i's potential liquidity suppliers participating in the market (more on this later). The first term in the payoff function (1) captures

[6] Without such uncertainty, the distribution f simply collapses to featuring mass-points at $\theta_i = 0$ and $\theta_i = 1$.

the expected surplus a market participant can retain when demanding liquidity from a random counterparty. The second term combines the benefit from participating in the market, which agent i fully retains when it has liquidity, with the payoff that can be extracted by supplying liquidity to the measure η of liquidity demanders.

Equilibrium. Since the benefit of acquiring a trading advantage increases with the probability of being able to supply liquidity, we conjecture that an agent i accepts to pay p in exchange for a trading advantage that boosts its per-transaction payoff by σ whenever $\theta_i \geq \theta_A$. This threshold θ_A will later be determined in equilibrium based on all agents' decisions. We also conjecture that agents decide at $t = 0$ to stay out of the market and collect a normalized payoff of 0 whenever their liquidity type is below θ_P, which again will be determined in equilibrium. For the ease of exposition, the subscripts A and P stand for "Advantage" and "Participation," respectively. Given a price p, the resulting pair (θ_P, θ_A) chosen by traders is defined as the *equilibrium of the traders' subgame*.

Finally, the monopolist charges the price that maximizes its expected profit. The *global equilibrium* is characterized by (p, θ_P, θ_A), where (θ_P, θ_A) is an equilibrium of the traders' subgame when the price of the trading advantage is p.

Welfare. Since a surplus is generated whenever an agent participates in the financial market, an equilibrium (p, θ_P, θ_A) produces a level of total welfare given by:

$$\Delta \cdot \int_{\theta_P}^{\overline{\theta}} dF(\theta) = \Delta \left[1 - F(\theta_P)\right], \tag{2}$$

which is a decreasing function of θ_P. The more agents agree to participate in the financial market, the more surplus is available to be split among agents through trade. Since liquidity matching solely transfers the surplus Δ among liquidity demanders and suppliers, the parameters ω and σ do not affect the level of welfare in the economy, aside from their impact on the determination of traders' participation threshold θ_P. Similarly, since the purchase of a trading advantage results in a transfer p from traders to the monopolist, the total welfare is unaffected by this transaction, except through its impact on θ_P. While the production of trading advantages such as access to satellite imaging and microwave transmission services might very well represent a socially wasteful use of resources in reality, we abstract away from this type of inefficiency already featured in Glode, Green, and Lowery (2012), Biais, Foucault, and Moinas (2015), and Budish, Cramton, and Shim (2015), among many others, in order to focus our analysis on how the sales of trading advantages affect market participation and the demand and supply of liquidity.

3 Traders' Decisions

In this section, we take the price of the trading advantage, p, as given and analyze how traders choose (θ_P, θ_A) as a result. Solving for these quantities involves a fixed-point problem: the participation threshold θ_P depends on the advantage-acquisition threshold θ_A and the advantage-acquisition threshold θ_A depends on the matching of liquidity demanders with suppliers, which is captured by η and depends on the participation threshold θ_P.

3.1 Solving for Equilibria of the Traders' Subgame

In order to solve for an equilibrium of the traders' subgame at a given price p, we need to (i) characterize the matching of liquidity demanders and suppliers for a given participation threshold θ_P, (ii) analyze traders' willingness to pay p in exchange for the trading advantage, which then dictates the advantage-acquisition threshold θ_A, and (iii) verify that θ_A and θ_P are optimal responses to each other and to p.

Matching of liquidity demanders and suppliers. Taking the equilibrium value of θ_P as given, we can derive the transaction volume that each liquidity supplier receives as:

$$\eta = \frac{\int_{\theta_P}^{\overline{\theta}} (1-\theta) dF(\theta)}{\int_{\theta_P}^{\overline{\theta}} \theta dF(\theta)}. \tag{3}$$

The numerator in the ratio represents the measure of agents demanding liquidity in the market whereas the denominator represents the measure of agents that can supply it. Their ratio thus represents the measure of transactions that each liquidity supplier is randomly matched to (and this ratio can be smaller or larger than 1). A liquidity supplier's transaction volume η is then decreasing in the participation threshold θ_P:

$$\begin{aligned} \frac{\partial \eta}{\partial \theta_P} &= \frac{-(1-\theta_P) f(\theta_P) \cdot \int_{\theta_P}^{\overline{\theta}} \theta dF(\theta) + \theta_P f(\theta_P) \cdot \int_{\theta_P}^{\overline{\theta}} (1-\theta) dF(\theta)}{\left[\int_{\theta_P}^{\overline{\theta}} \theta dF(\theta)\right]^2} \\ &= \frac{-f(\theta_P) \cdot \int_{\theta_P}^{\overline{\theta}} (\theta - \theta_P) dF(\theta)}{\left[\int_{\theta_P}^{\overline{\theta}} \theta dF(\theta)\right]^2} < 0. \end{aligned} \tag{4}$$

A higher participation threshold θ_P implies that more of the agents likely to need liquidity if they were to participate in the financial market opt for not participating in this market. Thus, the aggregate demand for liquidity goes down, meaning that each trader in position to supply liquidity ends up being involved in fewer transactions.

When agent i makes decisions to participate and to acquire a trading advantage, it is based on the expected equilibrium decisions of its potential counterparties (see equation (1)). These equilibrium concerns are captured in our model through the matching function η. Our model's matching process can be interpreted as originating from either centralized or decentralized markets. On one hand, the matching protocol assumed above can be interpreted as each liquidity demander randomly meeting one liquidity supplier as part of a bilateral transaction typical of over-the-counter markets, thereby implying that each liquidity supplier expects to meet a measure η of liquidity demanders. On the other hand, η can be interpreted as the volume of trade orders that each liquidity supplier with high-frequency trading capacities could either cream skim or scalp on an electronic stock exchange. In both cases, what matters to an agent making decisions at $t = 0$ is how much surplus its counterparties are expected to appropriate when this agent demands liquidity (i.e., $\sigma \cdot \mathbb{E}[a_{-i}]$) and the volume of transactions on which this agent could use its trading advantage when supplying liquidity (i.e., η).

Acquisition of a trading advantage. At $t = 0$, each agent i must decide whether to pay p to the monopolist in order to gain a trading advantage yielding an extra payoff σ in all transactions where it supplies liquidity. Using equations (1) and (3), we know that agent i agrees to pay p whenever:

$$\begin{aligned} p &\leq \theta_i \eta \sigma \\ &= \theta_i \sigma \cdot \frac{\int_{\theta_P}^{\overline{\theta}} (1-\theta) dF(\theta)}{\int_{\theta_P}^{\overline{\theta}} \theta dF(\theta)}. \end{aligned} \tag{5}$$

Since a trading advantage can be beneficial only if one participates in the market, agent i pays to acquire a trading advantage whenever:

$$\theta_i \geq \theta_A = \max\left\{\theta_P, \frac{p}{\sigma} \frac{\int_{\theta_P}^{\overline{\theta}} \theta dF(\theta)}{\int_{\theta_P}^{\overline{\theta}} (1-\theta) dF(\theta)}\right\}. \tag{6}$$

This condition shows that more trader types are willing to acquire the trading advantage when its price is low relative to the payoff boost it provides. Also, it shows that more trader types are willing to acquire the advantage when the measure of liquidity suppliers relative to liquidity demanders among market participants is low, since each liquidity supplier receives more transaction volume and can thereby use the acquired trading advantage against a larger measure of liquidity demanders.

By assuming that a trading advantage only serves a trader when supplying liquidity, the probability of supplying liquidity acts as a scaling factor for the benefit of acquiring a trading advantage. As a result, agent type θ_i generates the heterogeneity required to have some traders (e.g., market makers and hedge funds) being willing to purchase the trading advantage while others (e.g., retail investors and endowment funds) are not. Alternative sources of heterogeneity that also scale up or down the benefit of acquiring the advantage for a given trader (e.g., assets under management, trading volume, flexibility of balance sheet) could play a similar role in our analysis of the optimal sales of trading advantages with endogenous market participation. But as will be clear below, our framework benefits from the tractability of having the same source of heterogeneity driving both the advantage-acquisition and market-participation decisions.

Market participation. At $t = 0$, each agent i must decide whether to enter the financial market. If planning to pursue strategy $a_i \in \{0, 1\}$ upon entering, agent i's expected payoff from entering the market is:

$$\begin{aligned} & (1-\theta_i)\mathbb{E}[(1-\omega)\Delta - \sigma a_{-i}] + \theta_i[\Delta + \eta(\omega\Delta + \sigma a_i)] - pa_i \\ = \ & (1-\theta_i)\left[(1-\omega)\Delta - \sigma \cdot \frac{\int_{\theta_A}^{\overline{\theta}} \theta dF(\theta)}{\int_{\theta_P}^{\overline{\theta}} \theta dF(\theta)}\right] + \theta_i\Delta + \theta_i(\omega\Delta + \sigma a_i) \cdot \frac{\int_{\theta_P}^{\overline{\theta}} (1-\theta) dF(\theta)}{\int_{\theta_P}^{\overline{\theta}} \theta dF(\theta)} - pa_i, \end{aligned} \tag{7}$$

where we use the fact that:

$$\mathbb{E}[a_{-i}] = \frac{\int_{\theta_A}^{\overline{\theta}} \theta dF(\theta)}{\int_{\theta_P}^{\overline{\theta}} \theta dF(\theta)}. \tag{8}$$

Thus, agent i finds it optimal to participate in the market whenever this expected payoff is greater than zero:

$$(1-\theta_i)\mathbb{E}[(1-\omega)\Delta - \sigma a_{-i}] + \theta_i[\Delta + \eta(\omega\Delta + \sigma a_i)] - pa_i \geq 0. \tag{9}$$

We know that the expected payoff (7) is increasing in agent i's liquidity type θ_i since:

$$\Delta + \eta[\omega\Delta + \sigma a_i] \geq \Delta \geq \mathbb{E}[(1-\omega)\Delta - \sigma a_{-i}]. \tag{10}$$

Equation (10) thus confirms our equilibrium conjecture that only agents with $\theta_i \geq \theta_P$ agree to participate in the market. Put simply, agents who would be most likely to be hit by a liquidity shock if they were to participate in the financial market (i.e., agents with high $(1-\theta_i)$) are also those who find participation least attractive.

3.2 Properties of Equilibria of the Traders' Subgame

Two cases must be considered as potential equilibrium outcomes for the traders' subgame, based on whether the marginal market participant whose $\theta_i = \theta_P$ acquires the trading advantage (i.e., $a_i = 1$) or not (i.e., $a_i = 0$).

The case with $\theta_A > \theta_P$. In such case, not all agents who decide to participate in the financial market acquire a trading advantage in equilibrium. Thus, the marginal participant whose $\theta_i = \theta_P$ opts for $a_i = 0$. Adjusting the participation condition (9) by setting $a_i = 0$ yields:

$$\theta_i \geq \frac{\sigma\mathbb{E}[a_{-i}] - (1-\omega)\Delta}{\sigma\mathbb{E}[a_{-i}] + (1+\eta)\omega\Delta}. \tag{11}$$

When the above inequality becomes an equality, it identifies the agent type that is indifferent between participating in the market or not. By substituting $\mathbb{E}[a_{-i}]$ and η into the condition above and recalling that $\theta_P \geq \underline{\theta}$, we obtain an expression that determines θ_P:

$$\theta_P = \max\left(\underline{\theta}, \frac{\sigma \int_{\theta_A}^{\overline{\theta}} \theta dF(\theta) - (1-\omega)\Delta \int_{\theta_P}^{\overline{\theta}} \theta dF(\theta)}{\sigma \int_{\theta_A}^{\overline{\theta}} \theta dF(\theta) + \omega\Delta\left[1 - F(\theta_P)\right]}\right). \tag{12}$$

To better understand the above equation, consider two possibilities. First, suppose that the first term in the maximand is smaller than the second term. As a result, θ_P is equal to the second term. Rearranging the equation implies:

$$\frac{\sigma}{\Delta}\int_{\theta_A}^{\overline{\theta}} \theta dF(\theta) = \frac{\int_{\theta_P}^{\overline{\theta}} \left[(1-\omega)\theta + \omega\theta_P\right] dF(\theta)}{1 - \theta_P}. \tag{13}$$

This condition shows that the participation threshold, which affects the right-hand side, and the acquisition threshold, which affects the left-hand side, are mutually related. Intuitively, agents' participation decisions depend on how many agents acquire the trading advantage and vice-versa. Moreover, since the left-hand side of equation (13) is monotonic in θ_A, for a given θ_P, there is at most one θ_A that satisfies (13).

Second, suppose instead that the second term in the maximand (12) is weakly smaller than the first term, in which case the equilibrium $\theta_P = \underline{\theta}$ and we have full participation. Plugging $\theta_P = \underline{\theta}$ into the second term in the maximand implies that θ_A must satisfy the following condition:

$$\frac{\sigma}{\Delta}\int_{\theta_A}^{\overline{\theta}} \theta \, dF(\theta) \leq \frac{(1-\omega)\mu + \omega\underline{\theta}}{1 - \underline{\theta}}. \tag{14}$$

Since the left-hand side is decreasing in θ_A, this condition implies a lower bound on θ_A.

Combining the two possible scenarios, we can obtain all possible pairs (θ_P, θ_A) that may arise in equilibrium. Panel (a) of Figure I plots, for a simple numerical example, the possible (θ_P, θ_A) pairs induced by the participation constraint. We denote the set of the possible equilibrium pairs (θ_P, θ_A) as Θ and investigate the price that must be charged by

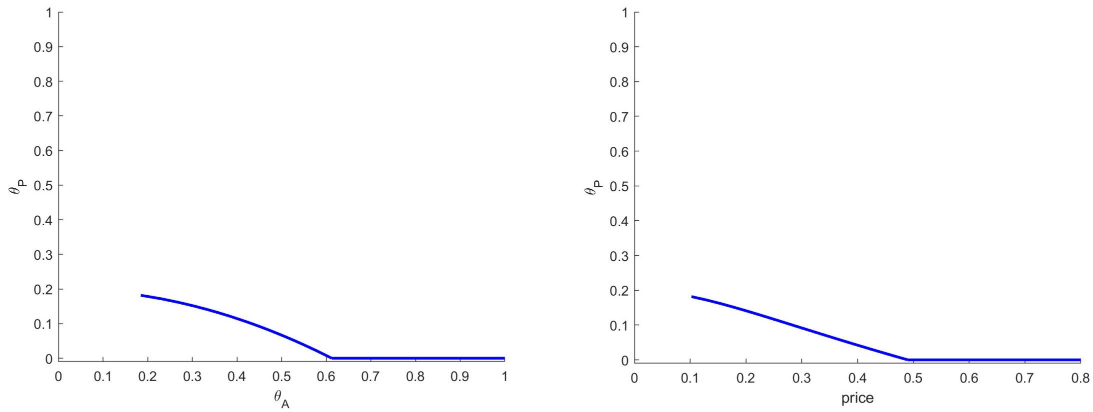

(a) (θ_A, θ_P) pairs.

(b) Implied price.

FIGURE I

This figure plots the possible (θ_P, θ_A) pairs and the implied prices p that trigger them, assuming $\omega = 0.5$, $\Delta = 1$, $\sigma = 0.8$, and $\theta_i \sim U[0,1]$.

the monopolist in order to induce each given pair in Θ. Given that agents would choose $\theta_A = \frac{p}{\sigma\eta}$ and η is determined by θ_P, we can compute the associated price for a given element of Θ as $p = \sigma\eta\theta_A$.

The case with $\theta_A = \theta_P$. In such case, agents either participate in the market while acquiring a trading advantage or they do not participate at all. Similar to the previous discussion, we can pin down the possible pairs (θ_P, θ_A) and their implied prices. To avoid repetition, we defer the formal analysis of this case to Appendix B. Figure II illustrates the implied prices when $\theta_A = \theta_P$ using the same numerical example as above.

Combining both cases. The analyses of cases $\theta_A > \theta_P$ and $\theta_A = \theta_P$ suggest that the participation threshold can be non-monotonic in the monopolist's price. For example, combining Panel (b) of Figure I with Panel (a) of Figure II yields Panel (a) of Figure III where the blue curve corresponds to the case with $\theta_A > \theta_P$ and the red curve corresponds to the case with $\theta_A = \theta_P$. Using this figure, we can see that the participation threshold is increasing in p for low values of p and then decreasing in p for high values of p.

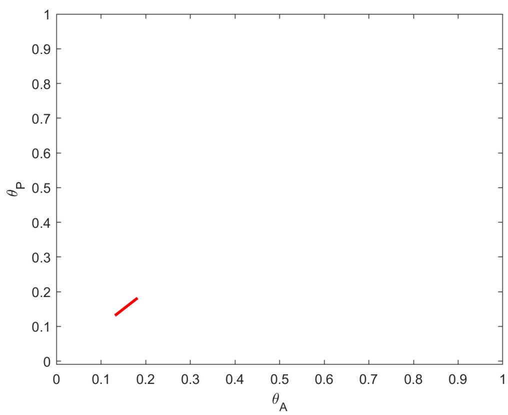

(a) (θ_A, θ_P) pairs.

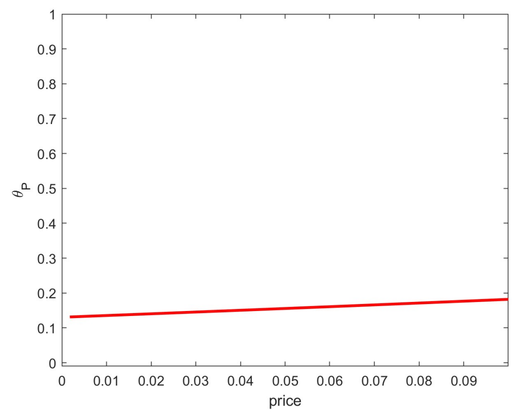

(b) Implied price.

FIGURE II

This figure plots the possible (θ_P, θ_A) pairs and the implied prices p that trigger them, assuming $\theta_A = \theta_P$, $\omega = 0.5$, $\Delta = 1$, $\sigma = 0.8$, and $\theta_i \sim U[0, 1]$.

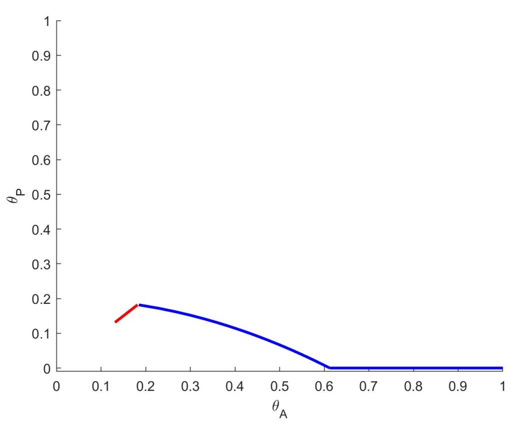

(a) (θ_A, θ_P) pairs.

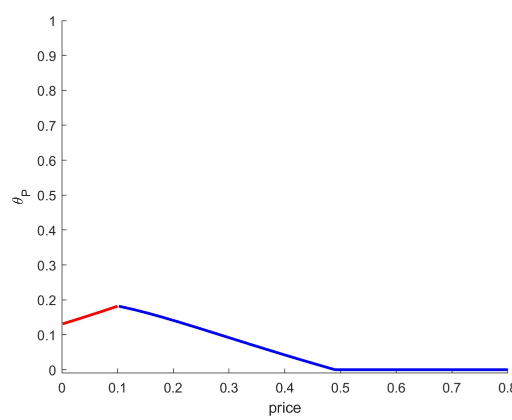

(b) Implied price.

FIGURE III

This figure combines both cases with $\theta_A > \theta_P$ and $\theta_A = \theta_P$ and shows how to find an equilibrium of the traders' subgame for a given price, assuming $\omega = 0.5$, $\Delta = 1$, $\sigma = 0.8$, and $\theta_i \sim U[0, 1]$.

4 Monopolist's Pricing Decision

Having characterized the possible equilibria of the traders' subgame, we now analyze how a monopolist chooses a price that maximizes the profit from selling a trading advantage while taking into account prospective traders' equilibrium responses.

As stated above, the monopolist's cost of providing an *additional* trader with access to a trading advantage (e.g., by sharing access to already collected data or to an existing high-speed communication network) is assumed to be zero. The monopolist thus chooses a price p to maximize revenues $p \cdot \int_{\theta_A}^{\overline{\theta}} dF(\theta) = p[1 - F(\theta_A)]$, subject to all agents' optimal decisions regarding market participation (i.e., condition (9)) and the costly acquisition of a trading advantage (i.e., condition (5)).

Before we can solve for the equilibrium pricing strategy, we need to revisit the two cases we considered in Section 3 with an emphasis on pricing. We first consider the case where the monopolist targets a strict subset of market participants in equilibrium, that is, $\theta_A > \theta_P$.

The case with $\theta_A > \theta_P$. Recall from subsection 3.2 that, for each possible pair $(\theta_P, \theta_A) \in \Theta$, we can find the price charged by the monopolist, which then allows us to compute the monopolist's profit. To compare this optimization problem with the classic monopolist's problem, we write the monopolist's profit as a function of θ_A, which can be interpreted as the marginal buyer type. In our model, the monopolist's objective is to maximize: $p[1 - F(\theta_A)] = \sigma\eta\theta_A[1 - F(\theta_A)]$. The first-order condition is then:

$$\frac{\partial \eta}{\partial \theta_A}\theta_A[1 - F(\theta_A)] + \eta\left[1 - F(\theta_A) - \theta_A f(\theta_A)\right] = 0, \tag{15}$$

which can be rearranged as:

$$\frac{\partial\eta/\eta}{\partial\theta_A/\theta_A} + 1 - H(\theta_A) = 0, \tag{16}$$

where $H(\theta_A) \equiv \frac{\theta_A f(\theta_A)}{1-F(\theta_A)}$. Note that in the classic monopoly pricing problem (i.e., when η is assumed to be constant), the standard FOC is $1-H(\theta_A) = 0$. However, our pricing problem is different because the level of θ_A affects η. Intuitively, a change in the fraction of agents that acquire the trading advantage affects which buyer types are willing to participate in the market, which in turn affects the matching of liquidity suppliers and demanders (i.e., η).

We write ϵ to denote $\frac{\partial\eta/\eta}{\partial\theta_A/\theta_A}$, which can be interpreted as the elasticity of agents' non-participation to the quantity of advantages acquired. This elasticity captures how a market's liquidity demand is impacted by how "rigged" a market appears to be for unsophisticated agents — if $\epsilon > 0$, decreasing θ_A leads to more traders acquiring the trading advantage, which may convince agents with low θ_i to stay away from the financial market and thereby increase θ_P. For ease of exposition, we make the following assumption:

Assumption 1. *$H(\theta)$ is strictly increasing in θ for $\theta \in [0, 1]$.*

Assumption 1 resembles the definition of a strictly regular environment by Fuchs and Skrzypacz (2015) and the standard assumption in auction theory that bidders' virtual valuation functions are strictly increasing. This regularity condition is satisfied by many well-known distributions, including the uniform distribution.

Our pricing problem can be simplified as setting $H(\theta_A) = 1 + \epsilon$. If the elasticity $\epsilon > 0$, then the optimal θ_A is higher than that in the classic monopolist problem, meaning there are fewer agents that end up acquiring the trading advantage. Alternatively, if the elasticity $\epsilon < 0$, then the optimal marginal buyer type θ_A is smaller than what it would be in the classic problem. We remark that ϵ is a function of θ_A, thus the solution to $H(\theta_A) = 1 + \epsilon$

may be challenging to solve for — there may be multiple solutions. Yet, we can show that at the optimum the elasticity ϵ is weakly positive. This is because the trading advantage cannot be a Giffen good in a neighborhood around the monopolist's optimal price.[7] We summarize the result in the following proposition:

Proposition 1. *In an equilibrium where $\theta_A > \theta_P$, the monopolist optimally quotes a price that induces $\theta_A \geq \theta_0$, where θ_0 denotes the smallest $\theta_i \in [\underline{\theta}, \overline{\theta}]$ such that $H(\theta_i) \geq 1$.*

Note that θ_0 is the marginal buyer type targeted in the classic monopolistic setting. Proposition 1 states that in our current setting with an endogenous η, the monopolist targets a weakly smaller mass of agents to purchase the trading advantage than would be targeted in the classic monopolistic setting where η is constant.

The case with $\theta_A = \theta_P$. We now explore the pricing of the advantage in the second case we considered above. As discussed in subsection 3.2 and Appendix B, an equilibrium with $\theta_A = \theta_P$ would require the price p charged by the monopolist to be solely a function of θ_P (see equation (B4) in Appendix B). The monopolist's profit would then be given by $p[1 - F(\theta_A)] = p[1 - F(\theta_P)]$, which is a function of θ_P. The pricing problem is thus a one-dimensional maximization problem subject to constraint (5) which captures traders' optimal purchasing behavior.

Remark 1 (Equilibrium Uniqueness)**.** *Since the monopolist's problem is a univariate maximization problem, there is a unique price that solves the problem for generic parameters. Global equilibrium uniqueness is determined by whether there are multiple equilibrium of the traders' subgame. In Appendix B, we show that the global equilibrium is generically*

[7] If $\epsilon < 0$ for a given p, then decreasing θ_A increases the relative liquidity demand η, which increases traders' willingness to pay for the trading advantage and the monopolist's profit, contradicting that p can be the monopolist's optimal price.

unique under mild conditions. Moreover, in our numerical exercises below, we verify that the equilibrium is unique.

4.1 Efficient/Full-Participation Equilibrium

Having analyzed the monopolist's profit in two possible and mutually exclusive cases, we now compare the monopolist's profit when triggering an equilibrium of the traders' subgame with $\theta_A = \theta_P$ versus one with $\theta_A > \theta_P$ to identify what the monopolist will want to achieve with its pricing strategy. In this subsection, we derive a sufficient condition under which the pecuniary externality parameter σ is small enough to allow for the optimal sales of trading advantages that do not prevent efficient market participation in equilibrium. This small-externality "benchmark" highlights how the negative externalities associated with the trading advantage interact with traders' participation decisions.

To start, consider the following arguments. Fixing θ_A, the monopolist would prefer an equilibrium with a lower θ_P as it leads to a higher η and a higher willingness to pay for the trading advantage. However, how low θ_P can go is constrained by agents' unwillingness to participate in "rigged" markets, as captured by equation (12). If the negative externality associated with a trading advantage is sufficiently small, however, we can identify a condition that ensures that all agents are willing to participate in the market and that the pricing problem reduces to the classic monopoly pricing problem (i.e., since η remains constant over a range of θ_A).

When all agents participate and $\theta_P = \underline{\theta}$, no welfare loss is incurred. If we ignore the participation constraint (12) for the time being, it follows that η is a constant when $\theta_P = \underline{\theta}$. The first-order condition of the profit maximization problem with respect to θ_A is precisely $H(\theta_A) = 1$. Given Assumption 1, the monopolist chooses $\theta_A = \theta_0 \equiv \max(\underline{\theta}, H^{-1}(1))$. Note that if $H(\underline{\theta}) \geq 1$, then $\theta_0 = \underline{\theta}$; otherwise $\theta_0 > \underline{\theta}$.

Bringing back the participation constraint (12) as part of the analysis means that we need to make sure that all agents are willing to participate even when the monopolist sells the trading advantage to traders whose θ_i is higher than θ_0. Equation (11) becomes:

$$\underline{\theta} \geq \frac{\sigma \mathbb{E}[a_{-i}] - (1-\omega)\Delta}{\sigma \mathbb{E}[a_{-i}] + (1+\eta)\omega\Delta}. \tag{17}$$

The above condition must hold for $\theta_P = \underline{\theta}$ and $\theta_A = \theta_0$. As we show in the proof of Proposition 2, substituting in θ_P and θ_A allows to rearrange condition (17) as:

$$\frac{\sigma}{\Delta} \leq \frac{(1-\omega)\mu + \omega\underline{\theta}}{(1-\underline{\theta}) \int_{\theta_0}^{\bar{\theta}} \theta dF(\theta)} \equiv k. \tag{18}$$

As a result, we can establish the following proposition:

Proposition 2. *In equilibrium, we have $(\theta_P, \theta_A) = (\underline{\theta}, \theta_0)$ if and only if $\frac{\sigma}{\Delta} \leq k$, in which case the equilibrium is socially efficient.*

We refer to any equilibrium that features $(\theta_P, \theta_A) = (\underline{\theta}, \theta_0)$ as a small-externality equilibrium. In such equilibrium, the monopolist's pricing decision is locally unaffected by agents' endogenous participation. To understand Proposition 2, recall that σ/Δ can be thought of as how powerful a trading advantage is since it measures the fraction of the surplus that a liquidity supplier can extract from a liquidity demander thanks to this advantage. Proposition 2 shows that the equilibrium outcome is socially efficient (i.e., $\theta_P = \underline{\theta}$) when $\sigma/\Delta \leq k$. In that case, the monopolist prefers to induce an equilibrium of the traders' subgame where η is as large as possible, that is, where $\theta_P = \underline{\theta}$.

The condition $\sigma/\Delta \leq k$ is sufficient for the existence of an efficient equilibrium, but it is not a necessary condition. In some cases with $\sigma/\Delta > k$, the monopolist might still find it optimal to induce full participation (i.e., $\theta_P = \underline{\theta}$) by raising θ_A slightly, in which case there is no welfare loss but fewer traders acquire the trading advantage than in a small-externality equilibrium. In fact, when $\sigma/\Delta > k$, either $\theta_P \neq \underline{\theta}$ or $\theta_A \neq \theta_0$. In the former case the equilibrium is inefficient, and in the latter case the monopolist can no longer induce what would have been induced in the classic monopoly pricing problem. As will become clear in the subsection below, as the trading advantage becomes more powerful, inducing an equilibrium of the traders' subgame with full participation requires that too few traders acquire a trading advantage and yields a suboptimal level of profit for the monopolist.

4.2 Inefficient/Partial-Participation Equilibrium

We now write the monopolist's profit as a function of θ_P and study its properties when θ_P approaches $\underline{\theta}$. A sufficient condition for inefficient participation in equilibrium is that this profit function is strictly increasing when $\theta_P = \underline{\theta}$. We can thus derive a sufficient condition for the case where the monopolist optimally quotes a price that induces low-θ agents to not enter the financial market, i.e., $\theta_P > \underline{\theta}$.

We assume that $\theta_A > \theta_P$ for the main text and defer the analysis of $\theta_A = \theta_P$ to Appendix B. Recall the participation constraint (13) for the case with $\theta_P > \underline{\theta}$:

$$\frac{\sigma}{\Delta}\int_{\theta_A}^{\overline{\theta}} \theta dF(\theta) = \frac{\int_{\theta_P}^{\overline{\theta}} [(1-\omega)\theta + \omega\theta_P]\, dF(\theta)}{1-\theta_P}. \tag{19}$$

In order to derive our inefficiency condition, a relevant variable we need to keep track of is the solution to θ_A from equation (19) when $\theta_P \to \underline{\theta}$. We denote the solution as θ_1 (assuming it exists). For the equilibrium to be inefficient, θ_1 has to be greater than θ_0, otherwise $(\theta_P, \theta_A) = (\underline{\theta}, \theta_0)$ can be supported, and as argued in the previous subsection, it is then in the monopolist's best interest to quote a price that induces efficient participation. From the proof of Proposition 2, the condition $\theta_1 > \theta_0$ is precisely equivalent to $\sigma/\Delta > k$. Since the left-hand side of equation (19) approaches zero as $\theta_A \to \overline{\theta}$, it follows that, under the condition $\sigma/\Delta > k$, there is a unique θ_1.

A sufficient condition for the equilibrium to be inefficient is that the monopolist's profit function is locally increasing at $\theta_P = \underline{\theta}$. Note that:

$$\frac{\partial[\sigma\eta\theta_A(1-F(\theta_A)]}{\partial\theta_P}\Big|_{\underline{\theta}} = \sigma\left(\frac{\partial\eta}{\partial\theta_P}\theta_A[1-F(\theta_A)] + \eta[1-F(\theta_A)-f(\theta_A)\theta_A]\frac{\partial\theta_A}{\partial\theta_P}\right)\Big|_{\underline{\theta}}. \tag{20}$$

Directly evaluating the right-hand side of the above equation (see the proof of Propostiion 3), we show that it is strictly positive whenever the following condition is satisfied:

$$\frac{\sigma}{\Delta} < \frac{(1-\mu)\mu[1-H(\theta_1)][\underline{\theta}(1-\underline{\theta})f(\underline{\theta}) - \omega - (1-\omega)\mu]}{f(\underline{\theta})(\mu-\underline{\theta})\theta_1{}^2 f(\theta_1)(1-\underline{\theta})^2}. \tag{21}$$

This condition ensures that the monopolist's profit function does not peak at $\theta_P = \underline{\theta}$. In particular, locally increasing θ_P above $\underline{\theta}$ can strictly increase the monopolist's profit. The following proposition summarizes the result:

Proposition 3. *Suppose $\theta_A > \theta_P$, $\sigma/\Delta > k$, and condition* (21) *holds, then the monopolist optimally quotes a price p that induces $\theta_P > \underline{\theta}$, in which case the equilibrium is socially inefficient.*

Proposition 3 identifies sufficient conditions for the equilibrium to feature excessive sales of trading advantages — here, the term "excessive" refers to the fact that too many traders acquire the advantage to allow the maximum levels of participation and social welfare. Condition (21) should, however, *not* be interpreted as requiring that σ/Δ is sufficiently small as θ_1 generally depends on σ/Δ. For example, when the distribution F of agent types is a standard uniform distribution, condition (21) simplifies to:

$$\frac{\sigma}{\Delta} < \frac{(2\theta_1 - 1)(1+\omega)}{4(1-\theta_1)\theta_1^2}, \tag{22}$$

where $\theta_1 = \sqrt{1 - \frac{1-\omega}{\sigma/\Delta}}$. We can then show that condition (22) is satisfied when either σ/Δ or ω is sufficiently large. We discuss these relationships in detail in the next subsection with the help of numerical examples.

4.3 Parameterization

We now focus on presenting numerical results that further illustrate our paper's main insights. To do so, we assume that F, the distribution of θ_i, is a standard uniform distribution. While there are still two possible cases to keep track of (i.e., $\theta_A > \theta_P$ or $\theta_A = \theta_P$), throughout the numerical analysis we made sure that the highest profit is always achieved in the case with $\theta_A > \theta_P$ (i.e., the case featured in most of our analysis above). For example, Figure IV plots the monopolist's profit as a function of the price of the trading advantage when $\omega = 0.5$ and $\omega = 0.8$, respectively. As we can see from the figures, the monopolist would optimally choose a price under which the equilibrium features $\theta_A > \theta_P$. Moreover, the "kink" on the blue curve corresponds to the case that the equilibrium of the traders' subgame features $(\theta_P, \theta_A) = (0, \theta_1)$. In the left panel, the monopolist induces $\theta_P = 0$ so the equilibrium is efficient, while in the right panel, the equilibrium features excessive "arms" sales.

We now plot various comparative statics. First, we study how exogenous parameters ω and σ/Δ affect the participation threshold θ_P. In our model, the payoff parameters Δ and σ impact equilibrium trader decisions only through their ratio σ/Δ. However, to be able to compare equilibrium prices and payoffs, we normalize $\Delta = 1$. Figures V and VI show that, as we increase ω or σ, the participation threshold θ_P weakly increases, indicating more welfare losses. Put differently, as liquidity suppliers can extract a higher fraction of the liquidity demanders' surplus or as the externality is stronger, the equilibrium is more likely to feature inefficiently low levels of trader participation.

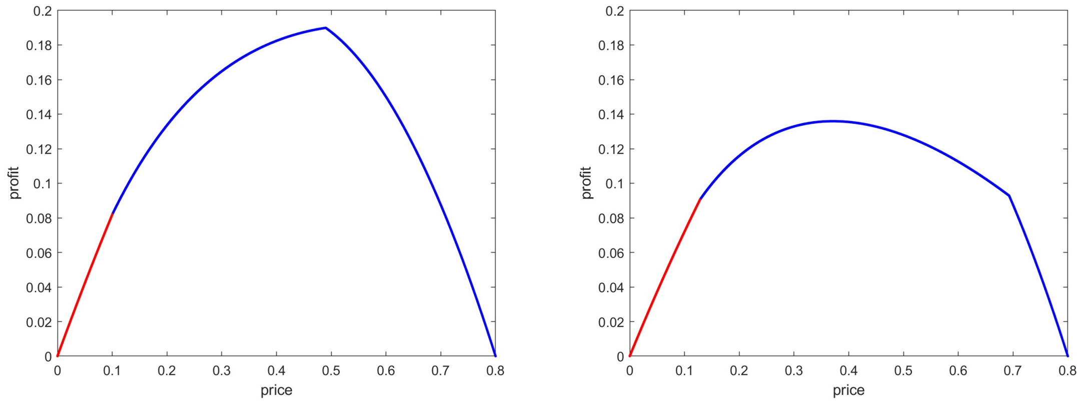

(a) $\omega = 0.5$.

(b) $\omega = 0.8$.

FIGURE IV

This figure plots the monopolist's profit as a function of price, assuming $\Delta = 1$, $\sigma = 0.8$, and $\theta_i \sim U[0, 1]$. The blue curve corresponds to the case with $\theta_P < \theta_A$, while the red curve corresponds to the case with $\theta_P = \theta_A$.

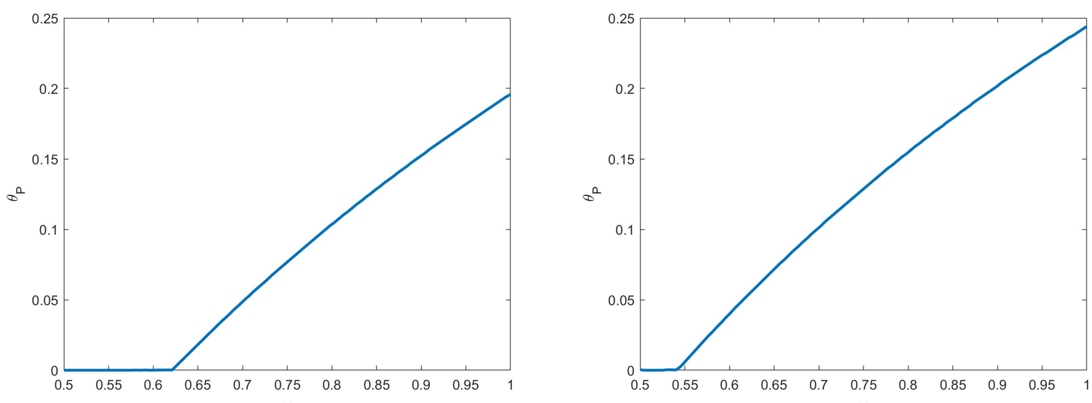

(a) $\sigma = 0.6$.

(b) $\sigma = 0.8$.

FIGURE V

This figure plots the value of θ_P when ω changes, assuming $\Delta = 1$ and $\theta_i \sim U[0, 1]$.

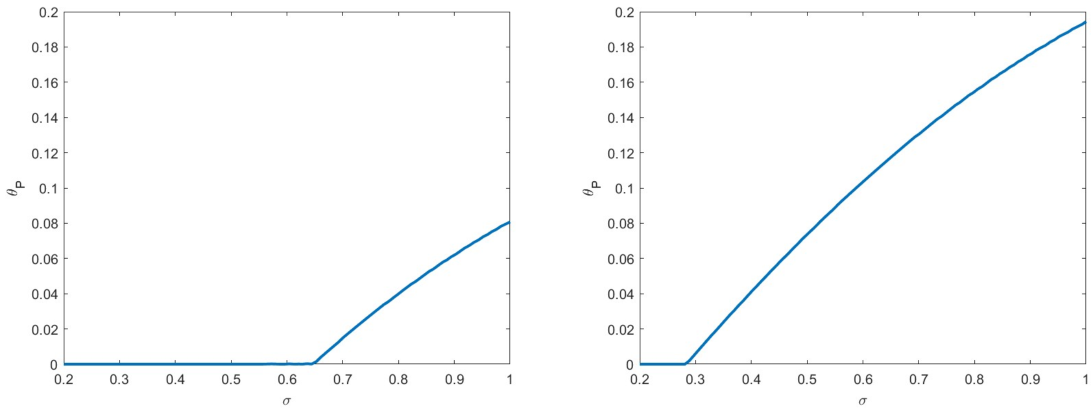

(a) $\omega = 0.6$. **(b)** $\omega = 0.8$.

FIGURE VI

This figure plots the value of θ_P when σ/Δ changes, assuming $\Delta = 1$ and $\theta_i \sim U[0, 1]$.

Second, Figure VII plots, for different levels of σ, the price charged by the monopolist in Panel (a), the equilibrium values for θ_A and θ_P in Panel (b), and the equilibrium payoffs for three types of agents ($\theta_i = 0.1, 0.5, 0.9$) in Panel (c). We can see that as the strength of the externality increases, fewer traders acquire the trading advantage, yet the equilibrium becomes more inefficient in the sense that θ_P is weakly increasing. However, the difference in agents' payoffs between the $\theta_i = 0.9$ type and the $\theta_i = 0.1$ type is non-monotonic in σ/Δ. Intuitively, when σ/Δ is relatively small, an increase in σ benefits high-θ_i agents at the expense of low-θ_i agents. Yet, as the trading advantage gets stronger, high-θ_i agents are made worse off in equilibrium while low-θ_i agents are either only slightly affected or unaffected once they decide to exit the market. When σ/Δ is large, the monopolist optimally induces fewer traders to acquire the trading advantage by charging a higher price for it. For example, the agent with $\theta_i = 0.9$ purchases the trading advantage in equilibrium for all levels of σ/Δ that we plot. Since having fewer agents purchasing the trading advantage benefits those that purchase it, it also allows the monopolist to charge a higher price. The latter channel may dominate, in which case some high-θ_i agents are made worse off by the trading advantage being stronger.

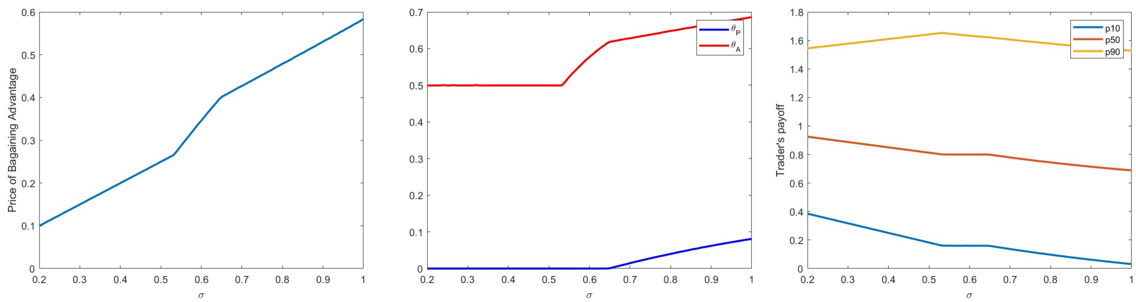

(a) Price charged by the monopolist.

(b) Equilibrium thresholds (θ_P, θ_A).

(c) Traders' expected payoffs.

FIGURE VII

This figure plots equilibrium objects as functions of σ, assuming $\omega = 0.6$ and $\Delta = 1$.

Third, Figure VIII identifies the parameter regions where the global equilibrium is inefficient. By shifting the distribution of θ_i to the right (from $U[0, 1]$ in Panel (a) to $U[0.1, 1]$ in Panel (b)), the inefficient region shrinks, indicating that it is easier to sustain the efficient-participation equilibrium. By shifting the distribution to the left (from $U[0, 1]$ in Panel (a) to $U[0, 0.9]$ in Panel (c)), the inefficient region expands slightly, indicating that it is slightly harder to sustain the efficient-participation equilibrium. Reducing the dispersion of types but preserving the mean (from $U[0, 1]$ in Panel (a) to $U[0.1, 0.9]$ in Panel (d)) makes the inefficient region shrink. Overall, dispersion in traders' liquidity types leads to more heterogeneity in the value of acquiring an advantage and a higher propensity for the inefficient non-participation of a subset of agents. Moreover, lowering the upper bound of F's support does not impact the inefficiency region as much as raising the lower bound of F's support does. These results emphasize the market-participation benefits of policies that reduce retail investors' liquidity needs and the frequency at which more sophisticated traders can take advantage of them.

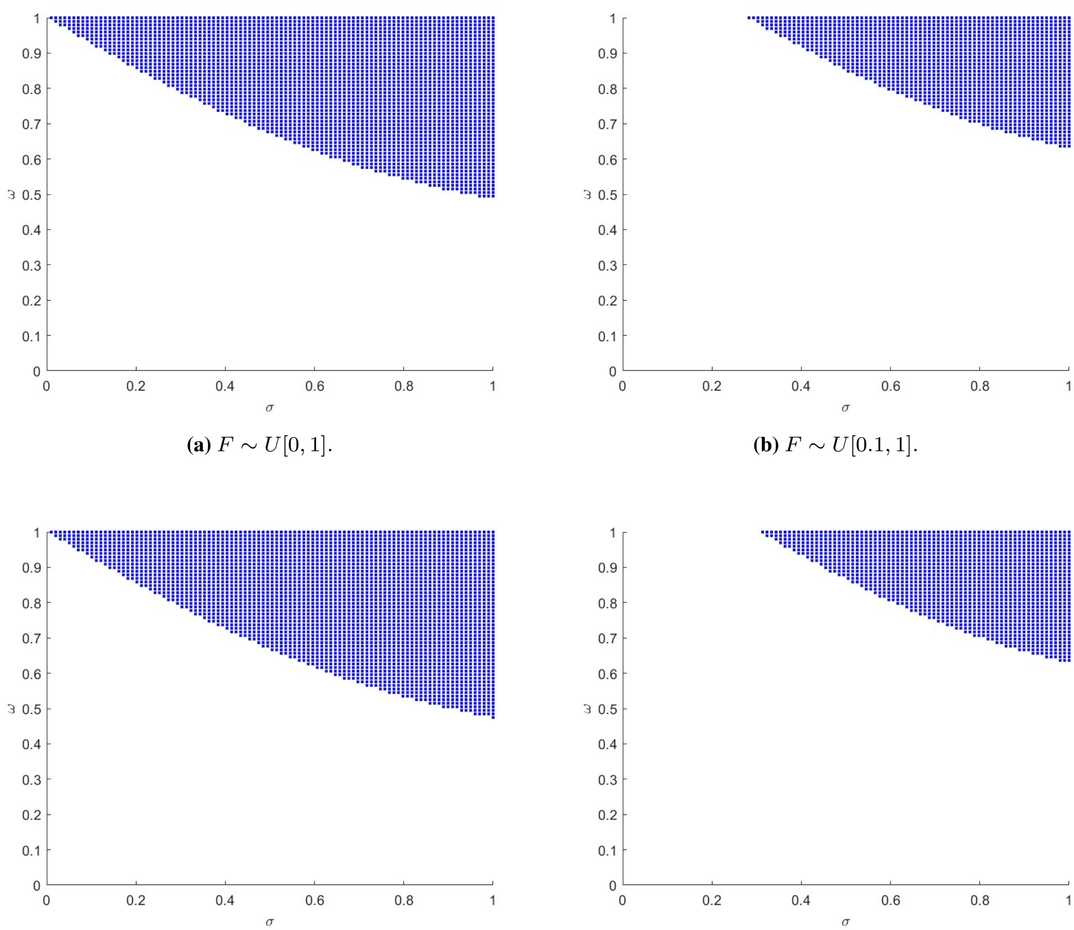

(a) $F \sim U[0, 1]$.

(b) $F \sim U[0.1, 1]$.

(c) $F \sim U[0, 0.9]$.

(d) $F \sim U[0.1, 0.9]$.

FIGURE VIII

This figure identifies, in blue, the parameter regions where the equilibrium is inefficient for various levels of trader-type heterogeneity.

5 Discussion

In this section, we discuss how varying the market structure can affect equilibrium behaviors by traders and by the entity selling the trading advantage.

5.1 Segmenting Traders Across Multiple Trading Venues

We have shown that the magnitude of the externality a trading advantage imposes on counterparties is a key determinant of various financial market outcomes. When full efficiency is not achieved, it is natural to ask what policy makers can do to improve market participation and liquidity. In this subsection, we consider the impact of interventions that segment traders of various types across multiple trading venues.

To fix ideas, suppose that the optimal price charged by the monopolist induces inefficient participation, i.e., $\theta_P > \underline{\theta}$, in our baseline environment. Now consider the introduction of a second trading venue that attracts agents whose type is above θ_P while the existing venue retains agents whose type is below θ_P (or vice-versa). We study the implications of this market-design adjustment in two scenarios that differ in whether the monopolist can quote different prices for the trading advantage based on the trading venue a trader participates in.

Let's first assume that the monopolist can charge differentiated prices across venues (say, if the trading advantage is co-location access to specific exchanges). As in Glode, Opp, and Zhang (2018), an interesting observation is that the monopolist's problem is unchanged in the first venue, because the optimal participation and acquisition conditions are unaffected by the addition of a second venue that caters to non-participants. Thus, all agents who were willing to participate when there was only one venue still participate in the first venue. However, the introduction of a second venue that only targets non-participants in the first venue results in some of those agents now being willing to participate. It immediately follows that the efficiency with the two venues is strictly higher than in the baseline environment with only one venue. Moreover, if we could further segment traders into any number of venues by repeating the same argument, then we should achieve full participation and maximum social welfare.

Let's now assume that at the time of the sale of a trading advantage, the monopolist must charge the same price to all agents, irrespective of which trading venue they will trade in (say, if the trading advantage is access to data). Since the monopolist could always charge the equilibrium price from the baseline environment, the monopolist is made weakly better off by the introduction of a second trading venue. In the meantime, our previous argument still applies, as introducing a second venue attracts more participation, thereby improving social welfare. Similarly as above, if we could further segment traders into any number of venues, then we should achieve full participation and maximum social welfare.

Summarizing our discussion above, we have the following proposition:

Proposition 4. *If a planner can segment traders across multiple trading venues, there exists a combination of trading venues such that all agents participate, in which case full efficiency is achieved, regardless of whether the monopolist is able to charge differentiated prices across segments/venues.*

A somewhat surprising allocation of trading advantage can arise in an environment with multiple trading venues. Suppose there are two trading venues and the first venue attracts agents whose type is above θ_P while the second venue attracts agents whose type is below θ_P. It is then possible that agents whose type is slightly below θ_P are willing to purchase the trading advantage while agents whose type is slightly above θ_P are not. This discrepancy is due to agents whose type is slightly below θ_P being in the second venue where counterparties are more likely to demand liquidity, thereby making the trading advantage more valuable.

5.2 Giving the Monopolist Less Power through Competition

Our baseline environment featured a monopolist selling a trading advantage to prospective traders. We now discuss the potential impact of adding competition for the sales of trading advantages. Specifically, we consider the case where there are at least two sellers engaging in Bertrand competition. Since the marginal cost of the trading advantage is assumed to be zero, the price charged by all sellers is zero in equilibrium. All agents, if they agree to participate, acquire the advantage for free. Does this imply that we have a full-participation equilibrium? Since trader's expected payoff is increasing in its liquidity type θ_i, we only need to check the lowest type trader's participation constraint:

$$(1-\underline{\theta})[(1-\omega)\Delta-\sigma]+\underline{\theta}[\Delta+\eta(\omega\Delta+\sigma)]\geq 0. \tag{23}$$

Equivalently,

$$\frac{\sigma}{\Delta}\leq\frac{(1-\omega)\mu+\omega\underline{\theta}}{\mu-\underline{\theta}}\equiv k_2. \tag{24}$$

If $\sigma/\Delta\leq k_2$, seller competition can ensure that all agents participate in equilibrium, in which case no welfare loss is incurred. However, if $\sigma/\Delta>k_2$, seller competition cannot induce a full-participation equilibrium. Proposition 2 stated that the equilibrium is efficient if $\sigma/\Delta<k$, yet we can show that k can be bigger than k_2. In that case, when $k_2<\sigma/\Delta<k$, a monopolist can induce an efficient equilibrium, while competing sellers cannot. For example, when F is a uniform over $[0,1]$, we find that $k=\frac{4}{3}(1-\omega)$ and $k_2=1-\omega$, in which case $k>k_2$.

Overall, a market with a monopolist can result in more efficient trader participation than a market with competing sellers. The simple intuition is that the trading advantage can be socially harmful and the losses associated with non-participation are partially internalized by a monopolist. Adding competition eliminates how sellers internalize non-participation in their pricing decisions, and can result in so many traders acquiring the advantage that it pushes traders with low θ out of financial markets. While monopoly power has its benefits, we show below that giving the monopolist even more power through the possibility of price discrimination can lead to a less efficient level of market participation.

5.3 Giving the Monopolist More Power through Price Discrimination

In subsection 5.2, we showed that reducing market power by adding a second seller can make market participation more inefficient in equilibrium. In this subsection, we instead boost market power by allowing the monopolist to observe a trader's type before quoting a price and thereby engage in first-degree price discrimination.

We now assume that the monopolist sells the trading advantage to all traders whose type is above θ_A, but unlike in the baseline model, the price quoted to each trader θ_i equals its willingness to pay $p_i = \sigma\eta\theta_i$ for all $\theta_i \geq \theta_A$. The monopolist's profit is then: $\sigma\eta \int_{\theta_A}^{\overline{\theta}} \theta dF(\theta)$. Expressing the profit as a function of θ_P and taking the derivative with respect to θ_P yields:

$$\frac{\partial[\sigma\eta \int_{\theta_A}^{\overline{\theta}} \theta dF(\theta)]}{\partial\theta_P} = \sigma\left(\frac{\partial\eta}{\partial\theta_P}\int_{\theta_A}^{\overline{\theta}} \theta dF(\theta) - \eta\theta_A f(\theta_A)\frac{\partial\theta_A}{\partial\theta_P}\right). \tag{25}$$

Setting this expression to be strictly positive, and substituting $\frac{\partial\eta}{\partial\theta_P}$ with (4) and $\frac{\partial\theta_A}{\partial\theta_P}$ with (A7) yields:

$$(1-\mu)\frac{\omega + (1-\omega)\mu - \underline{\theta}(1-\underline{\theta})f(\underline{\theta})}{(1-\underline{\theta})^2} > \frac{\sigma}{\Delta} f(\underline{\theta})\frac{\mu - \underline{\theta}}{\mu}\int_{\theta_1}^{\overline{\theta}} \theta dF(\theta). \tag{26}$$

To understand this condition, we suppose that F is a uniform distribution over $[0, 1]$. Condition (26) then reduces to:

$$\frac{\sigma}{\Delta} < \frac{2}{3}(1+\omega). \tag{27}$$

From Proposition 2, we know that when $\sigma/\Delta < k$, a non-discriminating monopolist would optimally induce an *efficient* equilibrium. When $\theta_i \sim U[0,1]$, $k = \frac{4}{3}(1-\omega)$ if F is $U[0,1]$. If $\omega > \frac{1}{3}$, then $\frac{4}{3}(1-\omega) < \frac{2}{3}(1+\omega)$, which means that when $\sigma/\Delta < k$, (27) holds, and a discriminating monopolist would induce an *inefficient* equilibrium unlike a non-discriminating monopolist.

6 Conclusion

Many financial transactions are of a fixed-sum nature, meaning that any improvement in the terms of trade for one party must come at the expense of another party. We model how the sales of trading advantages (e.g., data or co-location services) by a monopolist (e.g., data provider or securities exchange) affect traders' endogenous participation in a market and vice-versa. We show how the magnitude of the externality a trading advantage imposes on counterparties impacts financial market conditions. While the monopolist accounts for how its pricing strategy affects market participation, its optimal pricing strategy can result in socially excessive sales of goods and services that benefit a subset of market participants, leaving a fraction of the potential surplus from trade unrealized due to the non-participation of their disadvantaged counterparties. We show how different types of market participants are heterogenously impacted by this behavior and we study how changing the market structure reduces or sometimes amplifies the excessive sales of trading advantages.

References

Admati, Anat R., and Paul Pfleiderer. 1986. "A Monopolistic Market for Information." *Journal of Economic Theory* 39, 400-38.

Admati, Anat R., and Paul Pfleiderer. 1990. "Direct and Indirect Sale of Information." *Econometrica* 58, 901-28.

Allen, F. and Gale, D. 1994. "Financial Innovation and Risk Sharing." *MIT press.*

Arrow, Kenneth. 1951. "Alternative Approaches to the Theory of Choice in Risk-Taking Situations." *Econometrica* 19, 404–37.

Biais, Bruno, Thierry Foucault, and Sophie Moinas. 2015. "Equilibrium Fast Trading." *Journal of Financial Economics* 116, 292-313.

Bond, Philip, and Diego García. 2022. "The Equilibrium Consequences of Indexing" *Review of Financial Studies* 35: 3175-230.

Budish, Eric, Peter Cramton, and John Shim. 2015. "The High-Frequency Trading Arms Race: Frequent Batch Auctions as a Market Design Response." *Quarterly Journal of Economics* 130, 1547-621.

Budish, Eric, Robin Lee, and John Shim. 2022. "A Theory of Stock Exchange Competition and Innovation: Will the Market Fix the Market?" *National Bureau of Economic Research* working paper No. w25855.

Chaderina, Maria, and Richard C. Green. 2014. "Predators and Prey on Wall Street." *Review of Asset Pricing Studies* 4: 1-38.

Eső, Péter, Volker Nocke, and Lucy White. 2010. "Competition for Scarce Resources." *RAND Journal of Economics* 41, 524-48.

Debreu, Gerard. 1951. "The Coefficient of Resource Utilization." *Econometrica* 19, 273–92.

Di Maggio, Marco and Franzoni, Francesco and Kermani, Amir and Sommavilla, Carlo. 2019. "The Relevance of Broker Networks for Information Diffusion in the Stock Market." *Journal of Financial Economics* 134: 419–446.

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

Garcia, Diego, and Francesco Sangiorgi. 2011. "Information Sales and Strategic Trading." *Review of Financial Studies* 24, 3069-104.

Glode, Vincent, Richard C. Green, and J. Richard Lowery. 2012. "Financial Expertise as an Arms Race." *Journal of Finance* 67, 1723-59.

Glode, Vincent, and J. Richard Lowery. 2016. "Compensating Financial Experts." *Journal of Finance* 71, 2781-808.

Glode, Vincent, and Guillermo Ordoñez. 2022. "Technological Progress and Rent Seeking." Unpublished working paper.

Glode, Vincent, Christian Opp, and Xingtan Zhang 2018. "Voluntary Disclosure in Bilateral Transactions." *Journal of Economic Theory* 175, 652-688.

Grossman, Sanford J,, and Joseph E. Stiglitz. 1980. "On the Impossibility of Informationally Efficient Markets." *American Economic Review* 70, 393-408.

Han, Bing, and Liyan Yang. 2013. "Social Networks, Information Acquisition, and Asset Prices." *Management Science* 59, 1444-57.

Hu, Grace Xing, Jun Pan, and Jiang Wang. 2017. "Early Peek Advantage? Efficient Price Discovery with Tiered Information Disclosure." *Journal of Financial Economics* 126, 399-421.

Huang, Shiyang, Yan Xiong, and Liyan Yang. 2022. "Skill Acquisition and Data Sales." *Management Science* 68, 5557-6354.

Jéhiel, Philippe, Benny Moldovanu, and Ennio Stacchetti. 1996. "How (Not) to Sell Nuclear Weapons." *American Economic Review* 86, 814-29.

Peress, Joel. 2005. "Information vs. Entry Costs: What Explains U.S. Stock Market Evolution?" *Journal of Financial and Quantitative Analysis* 40: 563-94.

Segal, Ilya. 1999. "Contracting with Externalities." *The Quarterly Journal of Economics* 114(2), pp.337-88.

Segal, Ilya. 2003. "Coordination and Discrimination in Contracting with Externalities: Divide and Conquer?" *Journal of Economic Theory* 113: 147-81.

Simsek, Alp. 2013. "Speculation and Risk Sharing with New Financial Assets" *Quarterly journal of economics* 128: 1365-96.

Stambaugh, Robert F. 2014. "Presidential Address: Investment Noise and Trends." *Journal of Finance* 69: 1415-53.

A Proofs of Formal Results

Proof of Proposition 1: We argue this result by contradiction. Suppose that $\theta_A < \theta_0$. Given the monotonically of H, it implies that $H(\theta_A) < 1$. As in the main text, there are two cases we need to consider. First, we conjecture an equilibrium where $\theta_P = \underline{\theta}$. Recall from equation (14) that the participation constraint implies that there is a lower bound on θ_A. So marginally increasing θ_A is still feasible (i.e., the participation constraint is satisfied). Given that $H(\theta_A) < 1$, marginally increasing θ_A leads to a higher profit. This is a contradiction to the optimality of θ_A.

Second, we conjecture an equilibrium where $\theta_P > \underline{\theta}$. In this case, θ_A is pinned down by equation (13). We first show that $\frac{\partial \theta_A(\theta_P)}{\partial \theta_P}|_{\theta_P} \leq 0$, again by contradiction. Suppose instead that $\frac{\partial \theta_A}{\partial \theta_P}|_{\theta_P} > 0$. Then marginally decreasing θ_P would decrease θ_A as well given that decreasing θ_P always increases η. So in this case, not only agents are willing to pay more for the trading advantage, the monopolist also sells this advantage to more traders. The monopolist can enjoy a higher profit by decreasing θ_P, contradicted with the optimality of θ_P. Thus, $\frac{\partial \theta_A(\theta_P)}{\partial \theta_P}|_{\theta_P} \leq 0$. Returning to the proof of Proposition 1, together with $\frac{\partial \eta}{\partial \theta_P} \leq 0$, it then follows that $\frac{\partial \eta}{\partial \theta_A} = \frac{\partial \eta}{\partial \theta_P}\frac{\partial \theta_P}{\partial \theta_A} \geq 0$. Thus, the elasticity $\frac{\partial \eta / \eta}{\partial \theta_A / \theta_A} \geq 0$. From the monopolist's FOC $H(\theta_A) = 1 + \epsilon$. If the elasticity $\epsilon \geq 0$, then the optimal θ_A is higher than θ_0. Again, this is a contradiction. $\square$

Proof of Proposition 2: For full participation to be an equilibrium outcome, the fraction of traders that acquire the trading advantage cannot be too large. Specifically, equation (12) implies that for $\theta_P = \underline{\theta}$ we need:

$$\underline{\theta} \geq \frac{\sigma \mathbb{E}[a_{-i}] - (1 - \omega)\Delta}{\sigma \mathbb{E}[a_{-i}] + (1 + \eta)\omega\Delta}.$$

Substituting in $\mathbb{E}[a_{-i}]$ from (8),

$$\underline{\theta} \geq \frac{\sigma \cdot \frac{\int_{\theta_A}^{\overline{\theta}} \theta dF(\theta)}{\int_{\theta_P}^{\overline{\theta}} \theta dF(\theta)} - (1-\omega)\Delta}{\sigma \cdot \frac{\int_{\theta_A}^{\overline{\theta}} \theta dF(\theta)}{\int_{\theta_P}^{\overline{\theta}} \theta dF(\theta)} + (\omega\Delta) \cdot \frac{\int_{\theta_P}^{\overline{\theta}} 1 dF(\theta)}{\int_{\theta_P}^{\overline{\theta}} \theta dF(\theta)}}. \tag{A1}$$

Recall that we are looking for the condition such that the equilibrium features $\theta_A = \theta_0$ and $\theta_P = \underline{\theta}$, it follows that

$$\underline{\theta} \geq \frac{\sigma \cdot \int_{\theta_0}^{\overline{\theta}} \theta dF(\theta) - (1-\omega)\Delta\mu}{\sigma \cdot \int_{\theta_0}^{\overline{\theta}} \theta dF(\theta) + (\omega\Delta)}. \tag{A2}$$

Simplifying yields:

$$\underline{\theta}\omega\Delta + (1-\omega)\Delta\mu \geq (1-\underline{\theta})\sigma \int_{\theta_0}^{\overline{\theta}} \theta dF(\theta), \tag{A3}$$

which can be rearranged as:

$$\frac{\sigma}{\Delta} \leq \frac{(1-\omega)\mu + \omega\underline{\theta}}{(1-\underline{\theta}) \int_{\theta_0}^{\overline{\theta}} \theta dF(\theta)}. \tag{A4}$$

□

Proof of Proposition 3: Recall that the participation constraint (12) when $\theta_P > \underline{\theta}$ is given by:

$$\frac{\sigma}{\Delta} \int_{\theta_A}^{\overline{\theta}} \theta \, dF(\theta) = \frac{\omega\theta_P(1-F(\theta_P)) + (1-\omega) \int_{\theta_P}^{\overline{\theta}} \theta \, dF(\theta)}{1-\theta_P}. \tag{A5}$$

Taking $\frac{\partial}{\partial \theta_P}$ on both sides of (A5) yields:

$$-\frac{\sigma}{\Delta}\theta_A f(\theta_A)\frac{\partial \theta_A}{\partial \theta_P} = \frac{[\omega(1-F(\theta_P)) - \theta_P f(\theta_P)](1-\theta_P) + \omega\theta_P(1-F(\theta_P)) + (1-\omega) \int_{\theta_P}^{\overline{\theta}} \theta \, dF(\theta)}{(1-\theta_P)^2} \tag{A6}$$

Evaluating the above equation at $\theta_P = \underline{\theta}$, it follows that:

$$-\frac{\sigma}{\Delta}\theta_1 f(\theta_1)\frac{\partial\theta_A}{\partial\theta_P}\Big|_{\underline{\theta}} = \frac{\omega + (1-\omega)\mu - \underline{\theta}(1-\underline{\theta})f(\underline{\theta})}{(1-\underline{\theta})^2}. \tag{A7}$$

Now returning to the monopolist's profit as a function of θ_P and taking the derivative with respect to θ_P and evaluating at $\theta_P = \underline{\theta}$, we get:

$$\frac{\partial[\sigma\eta\theta_A(1-F(\theta_A)]}{\partial\theta_P}\Big|_{\underline{\theta}} = \sigma\left(\frac{\partial\eta}{\partial\theta_P}\theta_A(1-F(\theta_A)) + \eta(1-F(\theta_A) - f(\theta_A)\theta_A)\frac{\partial\theta_A}{\partial\theta_P}\right)\Big|_{\underline{\theta}}. \tag{A8}$$

A sufficient condition for the equilibrium to be inefficient is that the monopolist's profit function is locally increasing at $\theta_P = \underline{\theta}$.

Substituting $\frac{\partial\eta}{\partial\theta_P}$ from equation (4), the above condition becomes:

$$\frac{-f(\theta_P)\cdot\int_{\theta_P}^{\overline{\theta}}(\theta-\theta_P)dF(\theta)}{\left[\int_{\theta_P}^{\overline{\theta}}\theta dF(\theta)\right]^2}\theta_A(1-F(\theta_A))\Big|_{\underline{\theta}} + \eta(1-F(\theta_A) - f(\theta_A)\theta_A)\frac{\partial\theta_A}{\partial\theta_P}\Big|_{\underline{\theta}} > 0. \tag{A9}$$

Plugging $\theta_P = \underline{\theta}$ and substituting $\frac{\partial\theta_A}{\partial\theta_P}$ from equation (A7) yields:

$$\frac{-f(\underline{\theta})(\mu-\underline{\theta})}{\mu^2}\theta_1(1-F(\theta_1)) + \frac{1-\mu}{\mu}(1-F(\theta_1) - f(\theta_1)\theta_1)\frac{\frac{\omega+(1-\omega)\mu-\underline{\theta}(1-\underline{\theta})f(\underline{\theta})}{(1-\underline{\theta})^2}}{-\frac{\sigma}{\Delta}\theta_1 f(\theta_1)} > 0. \tag{A10}$$

Rearranging, we obtain the following condition:

$$\frac{\sigma}{\Delta} < \frac{\frac{1-\mu}{\mu}(1-F(\theta_1) - f(\theta_1)\theta_1)}{\frac{f(\underline{\theta})(\mu-\underline{\theta})}{\mu^2}\theta_1(1-F(\theta_1))\theta_1 f(\theta_1)}\frac{\underline{\theta}(1-\underline{\theta})f(\underline{\theta}) - \omega - (1-\omega)\mu}{(1-\underline{\theta})^2}, \tag{A11}$$

equivalently,

$$\frac{\sigma}{\Delta} < \frac{(1-\mu)\mu(1-H(\theta_1))}{f(\underline{\theta})(\mu-\underline{\theta})\theta_1{}^2 f(\theta_1)} \frac{\underline{\theta}(1-\underline{\theta})f(\underline{\theta}) - \omega - (1-\omega)\mu}{(1-\underline{\theta})^2}, \tag{A12}$$

which is condition (21) in the main text. □

Proof of Proposition 3: The proposition follows from the arguments preceding the proposition in the main text. □

Proof of Proposition 4: The proposition follows from the arguments preceding the proposition in the main text. □

B Additional Analysis Omitted from the Main Text

Deriving the Implied Price when $\theta_A = \theta_P$ in the Equilibrium of the Traders' Subgame

The participation condition implies that:

$$\theta_P = \max\left(\underline{\theta}, \frac{p-(1-\omega)\Delta+\sigma\cdot\mathbb{E}[a_{-i}]}{\Delta+(\omega\Delta+\sigma)\cdot\eta-(1-\omega)\Delta+\sigma\cdot\mathbb{E}[a_{-i}]}\right), \tag{B1}$$

Substituting in $\mathbb{E}[a_{-i}]$ and η,

$$\theta_P = \max\left(\underline{\theta}, \frac{p-(1-\omega)\Delta+\sigma\cdot\frac{\int_{\theta_A}^{\overline{\theta}}\theta dF(\theta)}{\int_{\theta_P}^{\overline{\theta}}\theta dF(\theta)}}{\Delta+(\omega\Delta+\sigma)\cdot\frac{\int_{\theta_P}^{\overline{\theta}}(1-\theta)dF(\theta)}{\int_{\theta_P}^{\overline{\theta}}\theta dF(\theta)}-(1-\omega)\Delta+\sigma\cdot\frac{\int_{\theta_A}^{\overline{\theta}}\theta dF(\theta)}{\int_{\theta_P}^{\overline{\theta}}\theta dF(\theta)}}\right), \tag{B2}$$

which means that:

$$\frac{p-(1-\omega)\Delta+\sigma\cdot\frac{\int_{\theta_A}^{\overline{\theta}}\theta dF(\theta)}{\int_{\theta_P}^{\overline{\theta}}\theta dF(\theta)}}{(\omega\Delta+\sigma)\cdot\frac{\int_{\theta_P}^{\overline{\theta}}(1-\theta)dF(\theta)}{\int_{\theta_P}^{\overline{\theta}}\theta dF(\theta)}+\omega\Delta+\sigma\cdot\frac{\int_{\theta_A}^{\overline{\theta}}\theta dF(\theta)}{\int_{\theta_P}^{\overline{\theta}}\theta dF(\theta)}}\leq\theta_P, \tag{B3}$$

and the weak inequality becomes an equality when $\theta_P \neq \underline{\theta}$. Let's focus on the equality case, i.e., $\theta_P > \underline{\theta}$. Rearranging equation (B3) and setting $\theta_A = \theta_P$ implies that the price charged by the monopolist must satisfy the following condition:

$$p=\theta_P\left(\frac{(\omega\Delta+\sigma)(1-F(\theta_P))}{\int_{\theta_P}^{\overline{\theta}}\theta dF(\theta)}\right)+(1-\omega)\Delta-\sigma. \tag{B4}$$

We also need the last participating agent, whose $\theta_i = \theta_A$, to be willing to pay p to acquire the advantage, that is, $p \leq \eta\theta_A\sigma = \eta\theta_P\sigma$ which is equivalent to imposing that:

$$\omega\frac{\theta_P(1-F(\theta_P))}{\int_{\theta_P}^{\overline{\theta}}\theta dF(\theta)}\leq(1-\theta_P)\frac{\sigma}{\Delta}-(1-\omega). \tag{B5}$$

Solving condition (B5) leads to a range of θ_P. For any θ_P in this range, there exists a price p, given by equation (B4), such that there exists an equilibrium in which $\theta_A = \theta_P$. We remark that when condition (B5) is slack, the marginal type of buyer, i.e., θ_A, can enjoy a strictly positive rent.

Proof of the Claim in Remark 1

In Remark 1, we claim that the equilibrium is unique under generic parameters. We provide a formal proof here.

For generic parameters, there is a unique price that maximizes the monopolist's profit. We are left to show that there is a unique equilibrium of the traders' subgame that is consistent with the unique price. We consider the following possibilities.

First, suppose $\theta_A > \theta_P$. Then (θ_P, θ_A) are pinned down by the participation constraint (12). If the implied relationship $\frac{\partial \theta_P}{\partial \theta_A}$ from (12) is negative, there is a unique equilibrium. This is because in this case, η is decreasing in θ_P, which is decreasing θ_A. Thus, η is increasing in θ_A, while in turns implies that p is increasing θ_A. From the unique price, we can pin down the unique θ_A, while in turns helps pin down the unique θ_P.

We note that while the condition $\frac{\partial \theta_P}{\partial \theta_A} < 0$ is generally hard to check due to non parametrized distributions F, we argue that the condition is very mild. Economically, the condition says that more agents participate when less traders purchase the trading advantage. Moreover, we can verify this condition is satisfied for all the numerical exercises in the main text.

Second, suppose $\theta_P = \theta_A$. To show unique, we argue by contradiction. Otherwise, there are two pairs (θ_P, θ_A) and (θ'_P, θ'_A) that are consistent with the price charged by the monopolist. Without loss, suppose $\theta_P = \theta_A < \theta'_P = \theta'_A$. Then the monopolist gets a strictly higher profit in the the equilibrium with (θ_P, θ_A) because there are more buyers of the trading technology than in the equilibrium with (θ'_P, θ'_A).

Finally, we need to show that for the optimal price p, there does not exist one equilibrium of the traders' subgame in the form of $\theta_P < \theta_A$ and one equilibrium in the form of $\theta_P = \theta_A$. We argue by contradiction. Because the monopolist's profit is maximized, then the thresholds for the agents who purchase the trading advantage must be the same. Suppose the two equilibrium are (θ_P, θ_A) and (θ_A, θ_A). Recall that η is decreasing in θ_P, so the market liquidity is strictly higher in the equilibrium with (θ_P, θ_A) than in (θ_A, θ_A). As a result, the monopolist earns a strictly higher profit in the in the equilibrium with (θ_P, θ_A) than in (θ_A, θ_A), contradicted to that the monopolist's profits are the same in the two equilibria.