Why Trade Over the Counter?*

Tomy Lee[†]
Central European University

Chaojun Wang[‡]
The Wharton School

September 7, 2023

Abstract

Over-the-counter trading thrives despite opposing regulatory pressure. In our model, traders who are less likely to be informed optimally choose the OTC market, where a dealer offers a trader-specific discount. Penalizing OTC trades causes some traders to exit but induces others with larger gains from trade to enter. Overall, aggregate trade volume falls and average bid-ask spread widens, yet welfare strictly improves if adverse selection risk is low. In a benchmark case, closing the OTC market strictly raises welfare specifically where *most* trades occur over the counter. We devise a simple procedure to implement the optimal Pigouvian tax.

JEL-Classification: D47, D61, D82, G14, G18

Keywords: Over-the-counter, exchanges, adverse selection, price discrimination, venue choice

*For invaluable advice and comments, we thank Jordi Mondria, Xianwen Shi, Darrell Duffie, Ettore Damiano, Colin Stewart, Yu An, Daniel Carvalho, Hongseok Choi, Peter Cziraki, Piotr Dworczak, Jim Goldman, Vincent Glode, Botond Kőszegi, Sam Langfield, Katya Malinova, Vincent Maurin, Andreas Park, Cecilia Parlatore, Christine Parlour, Ádám Szeidl, Semih Üslü, Adrian Walton, Liyan Yang, Ádám Zawadowski, and Yao Zeng. We thank Amy Dai, William Fan, Clarise Huang, Dylan Marchlinski, Sifan Tao, and Ceyda Ustun for excellent research assistance. We also thank Jiaqi Zou for excellent professional editing.

[†]Central European University; email: leeso@ceu.edu

[‡]The Wharton School, University of Pennsylvania; email: wangchj@wharton.upenn.edu

1 Introduction

Over-the-counter (OTC) markets are opaque and difficult to access. Yet, they host the vast majority of trades in financial assets, despite a decade of regulatory efforts to shift trades onto exchanges.[1] The EU's MiFID II bans OTC dealers from matching or rerouting equities trades under most circumstances. The US Dodd-Frank Act requires standardized swap trading platforms to offer the option of trading the swaps on an exchange. Nonetheless, 20% of EU equities trades remain over the counter—the same share as before MiFID II (Exton and Healey, 2022)—and so are nearly all trades of standardized swaps in the US (Riggs, Onur, Reiffen and Zhu, 2020). Dealers cite the high and persistent OTC market share and evidence that trading costs are lower over the counter than on exchanges (Collin-Dufresne, Junge and Trolle, 2020; de Roure, Moench, Pelizzon and Schneider, 2021) to argue against regulatory intervention.[2] Do low trading costs and high market shares indicate that OTC trading is efficient? Should policymakers intervene against the OTC market, and under which conditions?

In this paper, we show that trading costs, trade volumes, and market shares can starkly mislead policy. An intervention in the OTC market can widen average bid-ask spread and reduce aggregate trade volume, yet strictly improve welfare. Moreover, the intervention may, under certain conditions, improve welfare *exactly where* the OTC market share is *high*. In contrast to these measures, adverse selection risk provides robust guidance. We show that penalizing or closing the OTC market strictly raises traders' utilitarian welfare whenever adverse selection risk is low, even if the OTC market were frictionless. Beyond penalties and closure, we devise an optimal Pigouvian tax that one can implement using a simple sufficient statistic.

[1] In our context, "over-the-counter (OTC) markets" consist of all financial markets in which trades are executed non-anonymously between a client and a dealer. This definition includes traditional voice markets in which clients contact dealers one by one, and request-for-quote (RFQ) markets in which clients contact multiple dealers at a time. "Exchanges" cover all other markets, including limit order books, batch auctions, dark pools, and all-to-all request-for-quote platforms.

[2] A typical example is the comment of Securities Industry and Financial Markets Association (2018) on a US regulatory proposal to limit the dealers' access to their counterparties' identities (see p. 30). SIFMA suggests regulatory intervention is unwarranted, because (i) the OTC trades have smaller bid-ask spreads and (ii) the swap traders overwhelmingly trade over the counter despite being "free to choose" an exchange.

We develop a model of trade that adopts the choice between an exchange and an OTC market from Seppi (1990) and Desgranges and Foucault (2005) (together SDF). In SDF, dealers can price discriminate among their customers, whereas exchanges cannot. Hence, the dealers can offer a discount to the traders who are less likely to be informed and cream skim them from the exchange. We nest this intuition in the otherwise standard framework of Glosten and Milgrom (1985). Doing so adds one crucial feature for welfare analysis: Traders who differ in their gains from trade and participate endogenously.

Section 2 describes our model. A continuum of traders trade an asset with an uncertain payoff. Uninformed traders have heterogeneous hedging benefits that incentivize them to trade. Informed traders receive signals about the asset payoff and seek profit. Whether a trader is informed or not is her private information. However, each trader is publicly labeled as either *Likely Informed (LI)* or *Likely Uninformed (LU)*, which imperfectly indicates her true type. All traders optimally choose among buying or selling on an exchange, with a dealer over the counter, or exit. The venues differ solely in that only the dealer may condition his prices on each trader's label. In equilibrium, an informed (LI or LU) trader always trades, and an uninformed trader trades if her hedging benefit exceeds her best bid-ask spread. The (informed and uninformed) LU traders endogenously choose the OTC market and the LI traders choose the exchange.

Section 3 presents the main results, which concern utilitarian welfare, aggregate trade volume, and average bid-ask spread. Our first result considers a "penalty" on OTC trades as follows. We reduce the accuracy of traders' labels, which blurs the separation between the informed and the uninformed traders, bringing the two markets' spreads closer together. Penalizing the OTC market this way raises welfare if the ratio β of the mass of informed traders to the mass of uninformed traders is low and reduces welfare if β is high. Yet, the penalty *always* decreases the aggregate volume and widens the average spread in our base setup. Closing the OTC market is equivalent to imposing a large penalty on OTC trades, since the closure pools all traders at the same spread. In sum, (i) how a policy affects aggregate volume and average spread is a poor indicator for its effect on welfare, and (ii) penalizing or closing the OTC market would raise welfare where adverse selection risk is low.

Result (i) is driven by "cheap substitution." Penalizing the OTC market causes exit in the

OTC market by widening the OTC spread, and induces entry on the exchange by narrowing the exchange spread. The aggregate volume decreases whenever the exiters outnumber the entrants. Nonetheless, welfare can increase, since every entrant has a larger hedging benefit than every exiter. In welfare terms, the entrants substitute for the comparably "cheap" exiters.

To illustrate cheap substitution, we consider a concrete example depicted in Figure 1. Initially, all informed traders are correctly labeled as Likely Informed, and hence trade on the exchange. We then relabel a small mass of informed traders as Likely Uninformed. This slight penalty on OTC trades destroys an average hedging benefit $\bar{b}(\text{exiters})$ per exiter and creates $\bar{b}(\text{entrants})$ per entrant. Given the mass of entrants $m(\text{entrants})$, the gross gain in welfare is $\bar{b}(\text{entrants})m(\text{entrants})$. The gross loss is $\bar{b}(\text{exiters})m(\text{exiters})$. On net, welfare increases if and only if $\frac{\bar{b}(\text{entrants})}{\bar{b}(\text{exiters})} > \frac{m(\text{exiters})}{m(\text{entrants})}$. The ratio $\frac{\bar{b}(\text{entrants})}{\bar{b}(\text{exiters})}$ measures the beneficial impact of cheap substitution on welfare. As Figure 1 illustrates, the average entrant gains a hedging benefit slightly smaller than the exchange spread, $\bar{b}(\text{entrants}) \approx S_{\text{E}} > 0$, whereas the average exiter hardly loses, $\bar{b}(\text{exiters}) \approx S_{\text{O}} = 0$. Thus, cheap substitution $\frac{\bar{b}(\text{entrants})}{\bar{b}(\text{exiters})} \approx \frac{S_{\text{E}}}{S_{\text{O}}}$ approaches infinity, and welfare strictly increases no matter how much the aggregate volume declines. Going beyond this example, cheap substitution $\frac{\bar{b}(\text{entrants})}{\bar{b}(\text{exiters})}$ is generally finite. Therefore, in the general case, a penalty or closure can raise or lower welfare depending on whether cheap substitution $\frac{\bar{b}(\text{entrants})}{\bar{b}(\text{exiters})}$ exceeds the effect of any decline in aggregate volume $\frac{m(\text{exiters})}{m(\text{entrants})}$.

Result (ii) follows because cheap substitution dominates where adverse selection risk is low. As a demonstration, we evaluate the impact of closing the OTC market in a numerical example. We set the mass of informed traders to μ and uninformed traders to 1. We suppose 75% of traders are correctly labeled—as LI for the informed traders and LU for the uninformed traders. The distribution of hedging benefits is uniform $\mathbb{U}[0,1]$. With a small μ, the bid-ask spread is locally linear in the ratio β of informed to uninformed traders because β is small in both markets. Then, the exchange spread $S_{\text{E}} \approx c\mu\, 0.75/(1-0.75) = 3c\mu$ and the OTC spread $S_{\text{O}} \approx c\mu\,(1-0.75)/0.75 = c\mu/3$ for some constant c, while closing the OTC market pools both spreads to the No-OTC spread $S_{\text{N}} \approx c\,\mu$. Hence, cheap substitution $\frac{\bar{b}(\text{entrants})}{\bar{b}(\text{exiters})} = \frac{(S_{\text{E}}+S_{\text{N}})/2}{(S_{\text{N}}+S_{\text{O}})/2} \approx 3$. On the other hand, the mass of entrants $m(\text{entrants}) = 0.25(S_{\text{E}} - S_{\text{N}}) \approx c\mu/2$ and likewise $m(\text{exiters}) = 0.75(S_{\text{N}} - S_{\text{O}}) \approx c\mu/2$, thereby their ratio $\frac{m(\text{exiters})}{m(\text{entrants})} \approx$

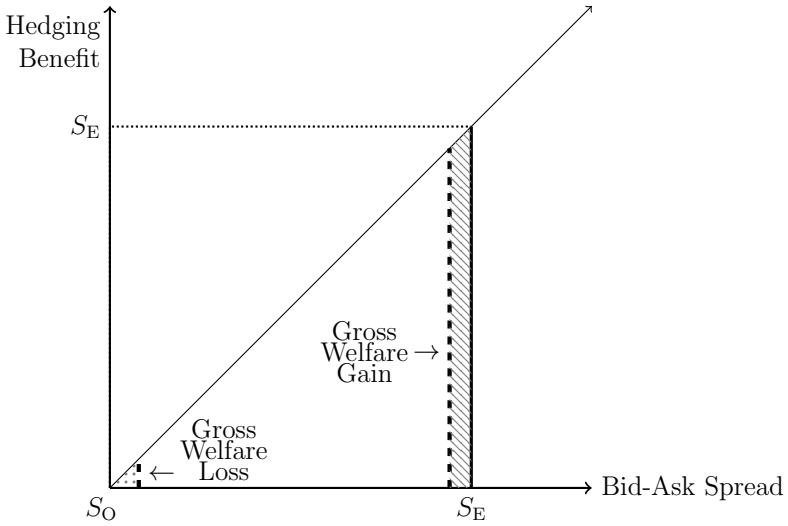

Figure 1: Slightly penalizing the OTC market raises welfare

Effect of reducing the accuracy of traders' labels on welfare. Initially, all informed traders are correctly labeled as Likely Informed. Then, a small mass of them are relabeled as Likely Uninformed.

1. Together, closing the OTC market strictly raises welfare at small μ. Penalizing the OTC market is analogous. Our results are remarkably robust. They hold under every economically relevant distribution, for any initial and any reduced traders' label accuracy, and whether or not the informed traders have hedging benefits themselves (Internet Appendix II).

Dealers frequently invoke the high market share of OTC markets to oppose regulatory intervention. Indeed, exchanges are available for many OTC-dominated assets.[3] Perhaps then the high OTC market share is a competitive equilibrium outcome, which may well be socially optimal. Section 3.4 shows that a high OTC market share is *not* evidence against intervention. A higher ratio of informed to uninformed traders β mechanically reduces the share of Likely Uninformed traders, who choose the OTC market. Hence, under a distributional assumption, β and the OTC market share are negatively related. Proposition 4 combines this negative relationship with our main result: When β drives the variation in

[3] Examples of exchanges for mostly OTC-traded assets include NYSE Bonds for corporate bonds, Saxo Bank SaxoTrader for EU government bonds, Tradeweb Dealerweb for repos, Refinitiv FXall for foreign exchange, and Bloomberg SEF CLOB and GFI Swaps Exchange for swaps. All these exchanges are open to any buy-side trading firm. In addition, several OTC-trading platforms allow clients to anonymously request quotes (e.g., Open Trading on MarketAxess), which fall under "trading on exchange" in our definition.

the OTC market share, penalizing or closing the OTC market *improves* welfare if the OTC market share is sufficiently *high*. High OTC market share is not evidence against intervention; it is not even evidence for keeping the OTC market.

We devise an easily implementable optimal Pigouvian tax on OTC trades in Section 4. The tax optimally trades off cheap substitution against any decline in aggregate volume to attain the highest possible welfare. Implementing the tax does not require policymakers to measure adverse selection risk. Rather, a simple policy experiment suffices. First, one introduces a small tax t on OTC trades. Second, one computes the Weighted Spread Ratio $WSR \coloneqq |S_{\mathrm{E}}/S_{\mathrm{O}} \div (dV_{\mathrm{O}}/dV_{\mathrm{E}})|$. The ratio of spreads $S_{\mathrm{E}}/S_{\mathrm{O}}$ captures the cheap substitution effect, and the ratio of changes in trade volumes $|dV_{\mathrm{O}}/dV_{\mathrm{E}}|$ captures the volume effect. If the WSR exceeds 1, cheap substitution dominates, and raising the tax t would strictly improve welfare. If the WSR is below 1, the decline in aggregate volume dominates, and cutting the tax t would strictly improve welfare.

Our model generates two testable predictions. First, the OTC spread is narrower than the exchange spread (Proposition 0). Internet Appendix IV presents supporting evidence. Second, the exchange's market share and spread are positively correlated through the informed ratio β, because increasing the share of informed traders both raises the exchange market share and widens its spread. Internet Appendix V documents a novel empirical pattern: The total market share of exchanges and their quoted spreads are positively correlated across US equities.

Whether price discrimination in OTC markets is socially beneficial is an increasingly relevant question. Beyond Dodd-Frank and MiFID II, the Commodities and Futures Trading Commission banned "name give-up" in 2020, partly to boost trading on exchanges. Several proposals to implement blockchain technology for recordkeeping of OTC trades would reveal traders' identities to selected dealers. Section 6 discusses these plans and policies through the lens of our model.

We make three contributions to the literature, which we will discuss shortly. First, we introduce cheap substitution, a new mechanism that can overturn the effects of aggregate volume and average spread on welfare. Cheap substitution shows how common proxies for welfare—volume, spread, and market share—are in fact poor indicators of welfare. Second,

we contribute the novel result that pooling improves welfare whenever adverse selection risk is low. Third, we devise a simple sufficient statistic with which one can implement an optimal Pigouvian tax.

1.1 Related Literature

We are most closely related to the literature on venue choice between OTC and centralized markets. One strand in this literature abstracts away from adverse selection and focuses on the presence of search frictions (Pagano, 1989; Rust and Hall, 2003; Vogel, 2019) or limited trading capacity (Dugast, Üslü and Weill, 2022) in OTC markets. Others, like this paper, feature cream skimming driven by price discrimination (Seppi, 1990; Desgranges and Foucault, 2005). Seppi (1990) explains why trade sizes are larger over the counter than on exchanges. Desgranges and Foucault (2005) shows how endogenous dealer-client relationships can concentrate adverse selection risk on exchanges. They do not examine social welfare or aggregate volume. Because the combination of adverse selection risk and heterogeneous gains from trade is absent, the existing papers in this literature cannot generate either of our main results.

We belong to the enduring literature that compares OTC and centralized markets. Benveniste, Marcus and Wilhelm (1992), Pagano and Roell (1996), Biais, Foucault and Salanié (1998), Malinova and Park (2013), and Glode and Opp (2019) compare the case of only having an exchange against only having the OTC market, and study outcomes unrelated to penalizing or closing the OTC market. More distantly related is the literature on how traders choose or split orders across multiple exchanges (Hendershott and Mendelson, 2000; Zhu, 2014; Pagnotta and Philippon, 2018; Lee, 2019; Chao, Yao and Ye, 2019; Babus and Parlatore, 2021; Baldauf and Mollner, 2021).

That the dealer engages in price discrimination links our paper to the literature on third-degree price discrimination (e.g., Pigou, 1920; Aguirre, Cowan and Vickers, 2010; Bergemann, Brooks and Morris, 2015). Pigou (1920) establishes that allowing a monopolist producer to price discriminate can reduce the total surplus. He identifies a "misallocation effect" in which output is inefficiently distributed whenever different consumers are charged different

prices (Aguirre et al., 2010). In their framework, adverse selection is absent and, instead, the distribution of private values determines the effect of price discrimination on welfare. Consequently, there can be *no* guidance over whether price discrimination would raise or lower welfare (Bergemann et al., 2015).[4] We show that adverse selection gives rise to a robust guidance: With minimal assumptions on the distribution of private values, price discrimination lowers welfare whenever adverse selection risk is low.

2 A Model of Venue Choice and Price Discrimination

Section 2.1 sets up a model in which each trader may trade on an exchange, trade with a dealer over the counter, or exit. The dealer observes a label for each trader's type and can use the labels to price discriminate, while the exchange cannot. Section 2.2 interprets our assumptions, and Section 2.3 derives the unique equilibrium.

2.1 Trading Game

A continuum of risk-neutral traders may trade an indivisible asset in a three-stage game. The asset is equally likely to pay $v = 1$ or -1 in the third stage. Each trader either exits without trading, or buys or sells 1 unit in one of two markets. In the OTC market, a dealer acts as the counterparty to the traders and absorbs net demand. A market maker does so on the exchange. The dealer and the market maker set prices such that their expected profit is zero, as we detail below.

A mass μ of the traders are informed and a mass 1 are uninformed. An informed trader has a private binary signal, which equals the true value v with probability $\alpha \in (1/2, 1)$ and $-v$ otherwise. Probability α is the accuracy of the informed traders' signals. Each uninformed trader is equally likely to be a buyer or a seller, and obtains a hedging benefit b_i upon trading in her desired direction. The hedging benefits are independent and uniformly distributed over $[0, 1]$, $b_i \overset{\text{iid}}{\sim} \mathbb{U}[0, 1]$. An uninformed trader's realized hedging benefit b_i and

[4] Specifically, price discrimination can generate any outcome under which the producer is at least as well off as with a uniform price, consumers receive non-negative payoffs, and the allocation is feasible (Bergemann et al., 2015).

whether she is a buyer or a seller are her private information.

An informed trader is labeled $\ell_i = \text{LI}$ ("Likely Informed") with probability θ and LU ("Likely Uninformed") otherwise. An uninformed trader is LU with probability γ and LI otherwise. Hence, there are $\theta\mu$ informed LI traders and γ uninformed LU traders. The labels are informative, $\theta > 1 - \gamma$, in that an LI trader is more likely to be informed than an LU trader. The labels become more informative as their accuracy, θ or γ, increases. We assume the label accuracy $\theta < 1$ or $\gamma < 1$, so that the labels are imperfectly informative.

In Stage 1, the dealer posts a bid to buy and an ask to sell one unit of the asset to every trader i. The dealer's quote is the highest bid and the lowest ask that earn him zero expected profit conditional on the trader's label $\ell_i \in \{\text{LI}, \text{LU}\}$. Simultaneously, the market maker posts the highest bid and the lowest ask that unconditionally earn him zero expected profit (as in Glosten and Milgrom (1985)). That is, the OTC market differs from the exchange in one way, that the dealer observes the label ℓ_i before setting the prices for trader i.

In Stage 2, every trader observes all prices, then makes two decisions: *Whether* to buy, sell, or exit, and *where* to trade. Figure 2 summarizes the timing of our model. All distributions, parameters, and the structure of the game are common knowledge.

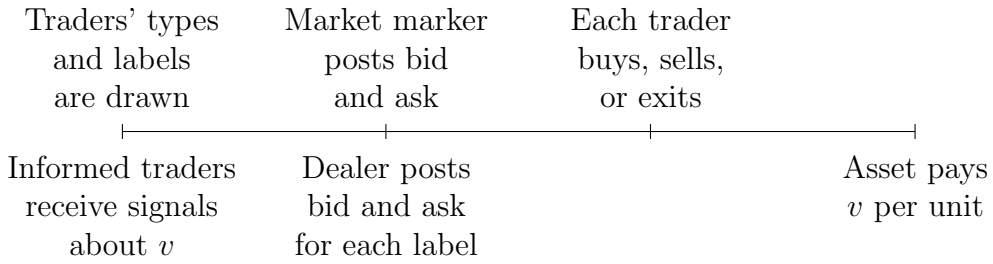

Figure 2: Timeline

We impose a tie-breaking rule to pin down a unique equilibrium.

Assumption 1. If a trader is indifferent between trading over the counter or on the exchange, she trades on the exchange.

Assumption 1 is purely expositional.[5] The rule is equivalent to imposing a small cost on OTC trades, which can represent the inconvenience of soliciting prices that is absent on exchanges.

2.2 Interpretation

Prices. Our setup features competitive prices, as defined by Glosten and Milgrom (1985). That prices on exchanges are competitive is a good proxy of reality. However, search frictions and the dealers' market power limit price competition in OTC markets. We nonetheless assume competitive prices to show that, even when the OTC market is made artificially efficient, penalizing or closing it can improve welfare. Introducing search frictions or market power would further raise the appeal of penalizing OTC trades.

Trading protocol. The dealer spontaneously posts label-dependent prices in the model. An OTC trade in practice occurs in two steps. A trader first requests quotes from her dealers, and then those dealers respond with trader-specific quotes. In Internet Appendix I, we present an extension that incorporates the requests for quote and endogenizes the competition among dealers and among market makers. Every trader requests quotes from all dealers in any extensive-form trembling-hand perfect equilibrium (Selten, 1975). The equilibrium allocation is identical to that of the base model described in the next section. We let the dealer post label-dependent prices to simplify exposition.

Traders' labels. One can interpret a trader's label as a summary statistic of her reputation and observable characteristics. Such observables may include the trader's industry (e.g., hedge fund or insurer), marketing or public filings (e.g., active versus passive fund), name (e.g., "Two-Sigma" or "AIG"), and any other public fact that is informative about the trader's motive. We assume imperfectly informative labels, because reputation and observables are noisy signals about the true motive behind a trade. This assumption is consistent with the evidence of Cheng and Xiong (2014) in the commodities futures market. The US Commodity Futures Trading Commission labels traders as "hedgers" if they are commodity

[5] All equilibria that would exist without Assumption 1 are payoff equivalent. We can eliminate Assumption 1 and show that traders who are indifferent between trading on the exchange or in the OTC market must choose the exchange with some probability ρ. Our results only require that $\rho > 0$.

producers and their past trades are not consistently profitable. Cheng and Xiong (2014) finds that the hedgers' positions are far more volatile than their outputs, especially their short positions. These short positions are consistently profitable and uncorrelated with output, which suggests that the so-called hedgers sometimes trade for profit.

Mass of traders. We fix the mass of uninformed traders and will vary the mass of informed traders μ in the welfare analysis. This normalization ensures that the maximum welfare that can be achieved is fixed and equal to the total hedging benefit of all uninformed traders.

2.3 Equilibrium

Definition 1. An equilibrium consists of the dealer's bid-ask spread for the LI traders, his spread for the LU traders, the market maker's spread, and the traders' trading strategies. Each trader's strategy maximizes her expected payoff given her available spreads. Given the traders' strategies, the dealer sets each of his spreads to the lowest spread that earns zero profit in expectation conditional on the label, and the market maker sets his spread to the lowest spread that earns him zero profit in expectation unconditionally.

An equilibrium bid is the negative of the corresponding ask, because the asset value v is symmetric around zero and the uninformed traders are equally likely to be a buyer or a seller. Thus, each pair of equilibrium bid and ask prices can be expressed as the half bid-ask spread s. For brevity, we hereon write "spread" to mean "half bid-ask spread".

The trader's equilibrium strategies are cutoff rules. An uninformed buyer or seller trades 1 unit in her corresponding direction at her smallest available spread if that spread is smaller than her hedging benefit, and exits otherwise. An informed trader buys or sells in the direction of her signal if her smallest spread is below $2\alpha - 1$, the expected value of this long or short position conditional on her signal.

We now pin down equilibrium spreads. If the dealer or the market maker (together "intermediaries") charges a spread s, an uninformed trader who *chooses* the intermediary[6] trades with probability $(1 - s)$. Conditional on such a trade, an intermediary's expected

[6]We say "a trader chooses a market" if the trader would trade in that market were she forced to trade.

profit is s. On the other hand, the intermediary's expected loss per trade with an informed trader is $(2\alpha - 1 - s)^+$. Altogether, the intermediary earns zero profit in expectation if and only if

$$\underbrace{s \cdot (1-s)}_{\substack{\text{Profit from} \\ \text{uninformed traders}}} = \underbrace{(2\alpha - 1 - s)^+ \cdot \beta}_{\text{Loss to informed traders}}, \tag{1}$$

where the *informed ratio* β is the mass of informed traders who choose the intermediary per unit mass of uninformed traders who choose the same. The zero-profit condition (1) has a unique solution, which we denote $S(\beta)$:

$$S(\beta) := \frac{1}{2}\left(\beta + 1 - \sqrt{(\beta-1)^2 + 8(1-\alpha)\beta}\right). \tag{2}$$

Proposition 0. *(a) Without the OTC market, the spread on the exchange is the No-OTC spread $S_N = S(\mu)$. (b) With the OTC market, the exchange spread is $S_E = S\left(\frac{\theta}{1-\gamma}\mu\right)$, and the spreads over the counter are the OTC spread $S_O = S\left(\frac{1-\theta}{\gamma}\mu\right)$ for LU traders and S_E for LI traders. Every LI trader chooses the exchange and receives S_E, and every LU trader chooses the OTC market and receives S_O. (c) The exchange spread is strictly wider than the OTC spread, $S_E > S_O$.*

Part (a) is a standard result of Glosten and Milgrom (1985). Part (b) incorporates venue choice: The LI traders choose the exchange, and the LU traders choose the OTC market in equilibrium. The traders of the same label choose the same market because, if they split, the prices they face in the two markets must be equal. Then, all traders with that label would choose the exchange by Assumption 1. Hence, the LU traders choose the OTC market whose spread is lower for them than the exchange spread, whereas the LI traders are indifferent between the two markets and choose the exchange due to Assumption 1. Our results remain unchanged if the LI traders choose the exchange with a strictly positive probability. Part (c) immediately follows from (2) and part (b). Internet Appendix IV cites recent evidence that trading costs are lower over the counter than on exchanges.

3 Welfare, Volume, and Spread

We analyze utilitarian welfare, aggregate trade volume, and average bid-ask spread upon (A) penalizing or (B) closing the OTC market. We show that (A) or (B) strictly raises welfare if the mass of informed traders μ is small and lowers welfare if μ is large. Yet, doing so may reduce the aggregate volume and widen the average bid-ask spread, irrespective of the value of μ.

Welfare W is the sum of all agents' *ex-ante* payoffs. It measures the total gains from trade in our model. Precisely, welfare W equals the sum of the hedging benefits that uninformed traders obtain through trade.

We let V_{O} denote the equilibrium volume of OTC trades, V_{E} the volume of trades on the exchange, and $V \coloneqq V_{\mathrm{O}} + V_{\mathrm{E}}$ the aggregate trade volume.[7] *Average bid-ask spread* $\bar{S}$ is the volume-weighted average of bid-ask spreads in the OTC market and on the exchange,

$$\bar{S} \coloneqq \frac{V_{\mathrm{O}}}{V} S_{\mathrm{O}} + \frac{V_{\mathrm{E}}}{V} S_{\mathrm{E}}.$$

We say "lower average spread $\bar{S}$" interchangeably with "higher aggregate volume V," because spread $\bar{S} \propto 1/V$ in equilibrium.[8]

The net effect of (A) or (B) on welfare hinges on the tradeoff between the change in trade volume and a "cheap substitution" effect. The aggregate volume of trades may decline on the extensive margin, pulling welfare downwards. However, the gains from each trade rise unambiguously on the intensive margin, pushing welfare upwards—an effect we call cheap substitution. Section 3.1 uses an example to illustrate how cheap substitution arises and can strictly raise welfare, even as the aggregate volume falls. Section 3.2 pinpoints the general conditions under which (A) or (B) raises or lowers welfare.

[7] Explicitly, $V_{\mathrm{O}} = (1-\theta)\mu + \gamma \cdot (1 - S_{\mathrm{O}})$ and $V_{\mathrm{E}} = \theta\mu + (1-\gamma) \cdot (1 - S_{\mathrm{E}})$.

[8] Formally, the zero-profit condition (1) implies

$$\underbrace{V \cdot \bar{S}}_{\text{Revenue}} = \underbrace{(2\alpha - 1) \cdot \mu}_{\text{Gross loss}},$$

which yields the inverse relationship between $\bar{S}$ and V.

3.1 Illustrative Example

We suppose all informed traders are correctly labeled as Likely Informed, $\theta = 1$. We then marginally lower label accuracy θ to $1-\mathrm{d}\theta$, which shifts a small mass of informed traders from the exchange to the OTC market. Marginally lowering θ therefore "marginally penalizes" the OTC market. If we were to continue lowering θ until $\theta = 1 - \gamma$, we would effectively close the OTC market, as the labels become uninformative and the bid-ask spreads of both markets converge.

This example exhibits two properties. Marginally lowering θ always raises welfare. Yet, doing so reduces the aggregate volume V and widens the average bid-ask spread $\bar{S}$.

Cheap substitution drives these properties. When the label accuracy θ falls to $1 - \mathrm{d}\theta$, the exchange spread narrows, $\mathrm{d}S_{\mathrm{E}} < 0$, and induces additional trades by the uninformed LI traders whose hedging benefits are slightly below S_{E} ("entrants"). The OTC spread widens, $\mathrm{d}S_{\mathrm{O}} > 0$, and deters trading by the uninformed LU traders whose hedging benefits are slightly above S_{O} ("exiters"). The net change in welfare equals the entrants' newly attained hedging benefits ("gross welfare gain") less the exiters' hedging benefits ("gross welfare loss"). Letting $\bar{b}$ denote average hedging benefit and m denote a mass of traders,

$$\text{Gross Welfare Gain} = \bar{b}(\text{entrants}) \times m(\text{entrants}),$$
$$\text{Gross Welfare Loss} = \bar{b}(\text{exiters}) \times m(\text{exiters}),$$

and thus

$$\text{welfare increases if and only if } \underbrace{\frac{\bar{b}(\text{entrants})}{\bar{b}(\text{exiters})}}_{\substack{\text{Cheap} \\ \text{Substitution}}} > \underbrace{\frac{m(\text{exiters})}{m(\text{entrants})}}_{\text{Volume Effect}}. \tag{3}$$

The relative hedging benefit of each entrant $\frac{\bar{b}(\text{entrants})}{\bar{b}(\text{exiters})}$ approaches infinity, because the entrant's average hedging benefit $\bar{b}(\text{entrants}) \approx S_{\mathrm{E}} = S(\frac{\theta\mu}{1-\gamma} = \frac{\mu}{1-\gamma}) > 0$ while the exiter's average benefit $\bar{b}(\text{exiters}) \approx S_{\mathrm{O}} = S(\frac{(1-\theta)\mu}{\gamma} = 0) = 0$. In welfare terms, the entrants substitute for "cheap" exiters. Hence, we call the ratio $\frac{\bar{b}(\text{entrants})}{\bar{b}(\text{exiters})}$ 'cheap substitution'.

By contrast, the relative mass of exiters $\frac{m(\text{exiters})}{m(\text{entrants})}$ is finite, since

$$m(\text{exiters}) = \left| \gamma \, \frac{S'(0)\mu}{\gamma} \right| \mathrm{d}\theta = \mu S'(0) \, \mathrm{d}\theta,$$
$$m(\text{entrants}) = \left| (1-\gamma) \, \frac{S'\left(\frac{\mu}{1-\gamma}\right)\mu}{1-\gamma} \right| \mathrm{d}\theta = \mu S'\left(\frac{\mu}{1-\gamma}\right) \mathrm{d}\theta,$$

and both derivatives $S'(0)$ and $S'(\frac{\mu}{1-\gamma})$ are strictly positive. That the bid-ask spread $S(\beta)$ is strictly concave implies $m(\text{exiters}) > m(\text{entrants})$.[9] Altogether, marginally lowering θ strictly raises welfare and strictly reduces the aggregate volume.

Figure 1 illustrates cheap substitution in this example. The exiters' average hedging benefit $\bar{b}(\text{exiters})$ is slightly above $S_{\text{O}} = 0$, whereas $\bar{b}(\text{entrants})$ is slightly beneath $S_{\text{E}} > 0$. The lined area for the gross welfare gain thus greatly exceed the dotted area for the gross welfare loss.

The crucial mechanism is that every entrant has a larger hedging benefit than any of the exiters. In the example, cheap substitution is unboundedly large upon lowering θ such that welfare would rise no matter how steeply aggregate volume were to decline.

Going beyond the example, a generic penalty on OTC trades may raise or lower welfare. Upon lowering either of the label accuracy parameters θ and γ, cheap substitution dominates whenever the mass of informed traders μ is small. The volume effect dominates at large μ. Figure 3a hints at why. When we reduce $\theta < 1$ to a strictly lower value θ_l, the exiters' average hedging benefit $\bar{b}(\text{exiters}) > 0$. Then, cheap substitution $\frac{\bar{b}(\text{entrants})}{\bar{b}(\text{exiters})}$ is no longer infinite. As we next elaborate, it turns out cheap substitution exceeds the volume effect $\frac{m(\text{exiters})}{m(\text{entrants})}$ across small μ.

Closing the OTC market converges all traders' bid-ask spreads to the pooling No-OTC spread S_{N}. Figure 3b shows $\bar{b}(\text{entrants}) \in (S_{\text{N}}, S_{\text{E}})$ and $\bar{b}(\text{exiters}) \in (S_{\text{O}}, S_{\text{N}})$ such that every

[9] The aggregate trade volume falls if and only if $S'(0) > S'\left(\frac{\mu}{1-\gamma}\right)$. This inequality holds if $S(\cdot)$ is strictly concave,

$$\frac{\partial^2}{\partial \beta^2}\left(\beta + 1 - \sqrt{(\beta-1)^2 + 8(1-\alpha)\beta}\right) < 0,$$

which is true whenever $\alpha > 1/2$, as is the case by assumption.

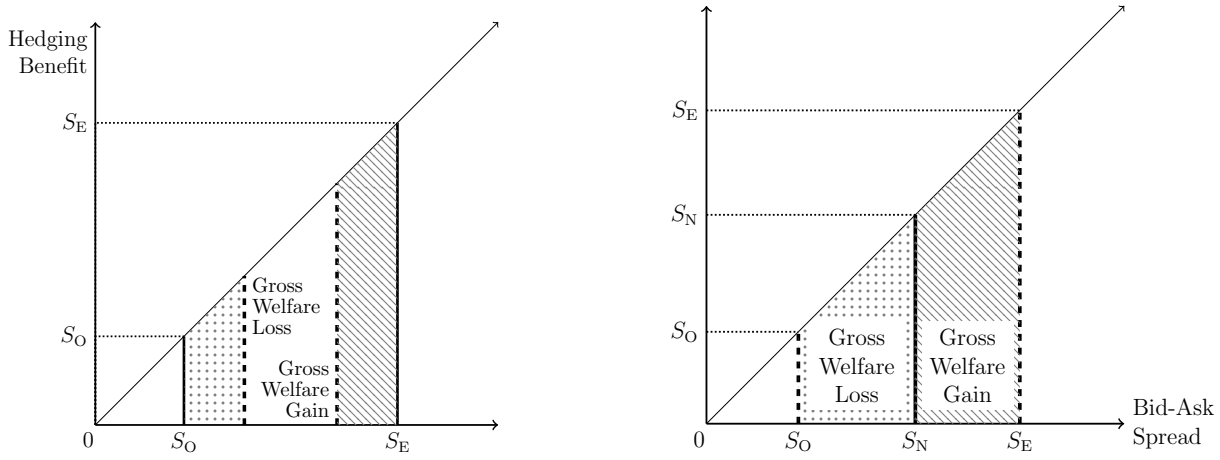

(a) Penalizing the OTC market

(b) Closing the OTC market

Figure 3: Effects on welfare beyond marginally penalizing the OTC market ($\gamma = 1/2$)

Left: The effect of a non-marginal penalty, which reduces the label accuracy θ from an initial level $\theta_h > 1-\gamma$ to any $\theta_l \in (1-\gamma, \theta_h)$. We mark the initial bid-ask spreads S_O and S_E. Right: The effect of closing the OTC market for any label accuracy θ. We mark the No-OTC spread that prevails after the closure S_N.

entrant's hedging benefit remains greater than any exiter's benefit. Comparison to Figure 3a makes clear that closing the OTC market is equivalent to a large penalty. Like the penalty, the closure raises welfare if the mass of informed traders μ is small and lowers welfare if μ is large.

3.2 Main Results

The next proposition states the effects of penalizing or closing the OTC market on welfare, the aggregate trade volume, and the average bid-ask spread. We penalize the OTC market by lowering the traders' label accuracy, θ or γ, which widens its spread. Proofs are in Appendix A.

Proposition 1. *Given any $\alpha > 0$ and any pairs $(\theta_l, \gamma_l) < (\theta_h, \gamma_h)$ that satisfy $\theta_h + \gamma_l > 1$ and $\theta_l + \gamma_h > 1$, there exist cutoffs $\mu', \tilde{\mu}', \hat{\mu}, \tilde{\mu}, \mu^* > 0$ such that:*

(a) Marginally lowering the traders' label accuracy θ from θ_h to $\theta_h - d\theta$ strictly raises W across all mass of informed traders $\mu < \mu'$ and strictly lowers W across all $\mu > \mu'$;

marginally lowering γ from γ_h to $\gamma_h - d\gamma$ strictly raises W across all $\mu < \tilde{\mu}'$ and strictly lowers W across all $\mu > \tilde{\mu}'$.

(b) Lowering the label accuracy θ from θ_h to θ_l strictly raises W across all $\mu < \hat{\mu}$ and strictly lowers W across all $\mu > \hat{\mu}$; lowering γ from γ_h to γ_l strictly raises W across all $\mu < \tilde{\mu}$ and strictly lowers W across all $\mu > \tilde{\mu}$.

(c) Closing the OTC market strictly raises welfare W across all $\mu < \mu^*$ and strictly lowers W across all $\mu > \mu^*$.

(d) Marginally lowering θ or γ, lowering θ or γ to any lower value, or closing the OTC market strictly reduces the aggregate trade volume V and widens the average bid-ask spread $\bar{S}$.

Proposition 1 has two messages. First, penalizing or closing the OTC market improves welfare if the mass of informed traders μ is small and harms welfare if μ is large. Second, in stark contrast, doing so *always* worsens the aggregate trade volume and the average bid-ask spread. Hence, volumes and spreads are poor indicators of welfare in our model.

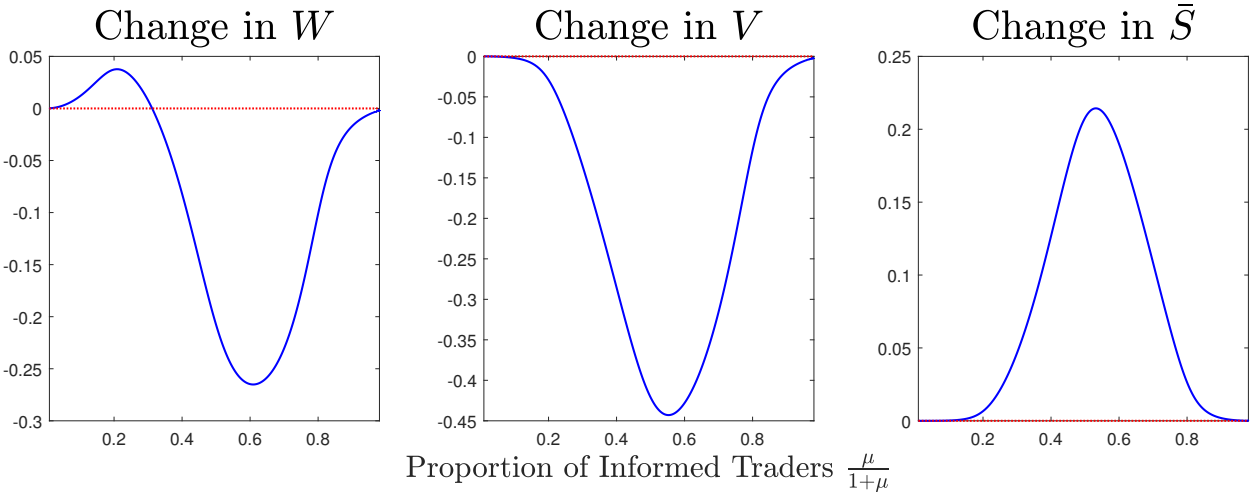

Figure 4: Effects of Closing the OTC Market

Figure 4 illustrates Proposition 1 with parameters $\theta = \gamma = 0.8$ and $\alpha = 0.98$. It plots the changes in welfare W, the aggregate volume V, and the average spread $\bar{S}$ upon closing

the OTC market as the proportion of informed traders $\mu/(1+\mu)$ varies. The changes are positive above the red line. Figure 4 confirms that closing the OTC market raises welfare where μ is low, while it always reduces the aggregate volume and widens the average spread. Adding search frictions or relaxing competitive prices in the OTC market would only expand the range of parameters under which closing the OTC market raises welfare.

Intuition. We explain Proposition 1 for the case of closing the OTC market. The closure widens the OTC spread $S_{\mathrm{O}} = S(\frac{(1-\theta)\mu}{\gamma})$ up to the No-OTC spread $S_{\mathrm{N}} = S(\mu)$, and narrows the exchange spread $S_{\mathrm{E}} = S(\frac{\theta\mu}{1-\gamma})$ down to S_{N}. The average hedging benefit of the exiters is $\bar{b}(\text{exiters}) = (S_{\mathrm{N}} + S_{\mathrm{O}})/2$ and that of the entrants is $\bar{b}(\text{entrants}) = (S_{\mathrm{E}} + S_{\mathrm{N}})/2$. The mass of exiters is $m(\text{exiters}) = \gamma(S_{\mathrm{N}} - S_{\mathrm{O}})$, and the mass of entrants is $m(\text{entrants}) = (1-\gamma)(S_{\mathrm{E}} - S_{\mathrm{N}})$. Altogether, condition (3) becomes

$$\text{welfare increases if and only if } \underbrace{\frac{S_{\mathrm{E}} + S_{\mathrm{N}}}{S_{\mathrm{N}} + S_{\mathrm{O}}}}_{\substack{\text{Cheap} \\ \text{Substitution}}} > \underbrace{\frac{\gamma}{1-\gamma} \cdot \frac{S_{\mathrm{N}} - S_{\mathrm{O}}}{S_{\mathrm{E}} - S_{\mathrm{N}}}}_{\text{Volume Effect}}. \tag{4}$$

For the intuition, we analyze the ratios in (4) at extreme values of the mass of informed traders μ. In the limit as $\mu \to 0$,

$$S_{\mathrm{M}} \sim S(0) + S'(0)\beta_{\mathrm{M}} = S'(0)\beta_{\mathrm{M}}\,, \quad \mathrm{M} \in \{\mathrm{O}, \mathrm{E}, \mathrm{N}\}\,,$$

where β_{M} is the ratio of informed to uninformed traders in market M, and cheap substitution

$$\lim_{\mu\to 0} \frac{S_{\mathrm{E}} + S_{\mathrm{N}}}{S_{\mathrm{N}} + S_{\mathrm{O}}} = \lim_{\mu\to 0} \frac{\beta_{\mathrm{E}} + \beta_{\mathrm{N}}}{\beta_{\mathrm{N}} + \beta_{\mathrm{O}}} = \frac{\frac{\theta}{1-\gamma} + 1}{1 + \frac{1-\theta}{\gamma}} > 1,$$

whereas the volume effect

$$\lim_{\mu\to 0} \frac{\gamma}{1-\gamma} \cdot \frac{S_{\mathrm{N}} - S_{\mathrm{O}}}{S_{\mathrm{E}} - S_{\mathrm{N}}} = \lim_{\mu\to 0} \frac{\gamma}{1-\gamma} \cdot \frac{1 - \frac{1-\theta}{\gamma}}{\frac{\theta}{1-\gamma} - 1} = 1.$$

Thus, at small μ, cheap substitution dominates and welfare improves upon closing the OTC

market. In the opposite limit as $\mu \to \infty$, all spreads approach $2\alpha - 1$ and

$$2\alpha - 1 - S_{\mathrm{M}} \sim \frac{(2\alpha - 1)(2 - 2\alpha)}{\beta_{\mathrm{M}}}.$$

Cheap substitution then converges to 1 as $\lim_{\mu\to\infty} (S_{\mathrm{E}} + S_{\mathrm{N}})/(S_{\mathrm{N}} + S_{\mathrm{O}}) = 1$. The volume effect instead converges to a number greater than 1,

$$\lim_{\mu\to\infty} \frac{\gamma}{1-\gamma} \cdot \frac{S_{\mathrm{N}} - S_{\mathrm{O}}}{S_{\mathrm{E}} - S_{\mathrm{N}}} = \lim_{\mu\to\infty} \frac{\gamma}{1-\gamma} \cdot \frac{(2\alpha - 1 - S_{\mathrm{O}}) - (2\alpha - 1 - S_{\mathrm{N}})}{(2\alpha - 1 - S_{\mathrm{N}}) - (2\alpha - 1 - S_{\mathrm{E}})} = \frac{\gamma}{1-\gamma} \cdot \frac{\frac{\gamma}{1-\theta} - 1}{1 - \frac{1-\gamma}{\theta}} > 1.$$

Then, at high μ, welfare falls upon closing the OTC market. It turns out that cheap substitution crosses the volume effect once, at the single cutoff μ^*, under uniformly distributed hedging benefits. Proposition 3 shows that a broad class of hedging benefit distributions also yields a single cutoff.

3.3 General Distributions

Proposition 1 is remarkably sharp: Penalizing or closing the OTC market always reduces aggregate volume, while their impact on welfare hinges on a single cutoff over the mass of informed traders μ. Crucial to the results is the rate of endogenous entry versus exit, which plainly depends on the distribution of hedging benefits. In this section, we generalize our results to most economically relevant distributions.

We suppose that the hedging benefits follow a generic distribution F, $b_i \overset{iid}{\sim} F$, whose support is $[0, 1]$ with a pdf $f(s)$ that is analytic in some neighborhoods of 0 and $2\alpha - 1$ (the

lower and upper bounds of the bid-ask spread).[10] The zero-profit condition (1) becomes

$$\underbrace{s \cdot (1 - F(s))}_{\substack{\text{Profit from} \\ \text{uninformed traders}}} = \underbrace{(2\alpha - 1 - s)^+ \cdot \beta}_{\text{Loss to informed traders}}. \tag{5}$$

The right-hand side is unchanged and the left-hand side changes from $s \cdot (1-s)$ to $s \cdot (1-F(s))$.

Two statistics characterize the effects of penalizing or closing the OTC market on welfare and volume. *Marginal volume* Δ_V is the decrease in the proportion of uninformed trades upon a marginal increase in the ratio of informed to uninformed traders β,

$$\Delta_V(\beta) \coloneqq -\left(\int_{S(\beta)}^1 f(s)\,\mathrm{d}s\right)' = S'(\beta) \cdot f(S(\beta)). \tag{6}$$

Marginal welfare Δ_W is the decrease in welfare per unit mass of uninformed traders upon a marginal increase in β. It is equivalent to the marginal volume Δ_V weighted by the hedging benefit $b = S(\beta)$ of the marginal uninformed trader:

$$\Delta_W(\beta) \coloneqq -\left(\int_{S(\beta)}^1 s f(s)\,\mathrm{d}s\right)' = \underbrace{S(\beta)}_{\substack{\text{Hedging} \\ \text{benefit} \\ \text{lost}}} \cdot \underbrace{S'(\beta) f(S(\beta))}_{\Delta_V}. \tag{7}$$

We show in Appendix A.2 that both Δ_V and Δ_W are well-defined.

The following proposition says our welfare result holds under any distribution f whose marginal welfare Δ_W is strictly quasiconcave.[11] For brevity, we write "strictly quasiconcave" and "quasiconcave" interchangeably. We sketch an example in Figure 5.

Proposition 2. *Given any α, we let f be a pdf analytic in some neighborhoods of 0 and*

[10] A function f is analytic at some x_0 if f is equal to its Taylor series at x_0 in a neighborhood of x_0,

$$f(x) = \sum_{n=0}^{\infty} a_n (x - x_0)^n \qquad \forall x \in (x_0 - \varepsilon, x_0 + \varepsilon) \text{ for some } \varepsilon > 0.$$

Analytic functions include the pdf of any beta distribution, all polynomials, the exponential function, all trigonometric functions, logarithms, and the power function, among many others.

[11] A strictly quasiconcave Δ_W is strictly increasing up to a cutoff, then strictly decreasing thereafter. If Δ_W is weakly quasiconcave, Proposition 2 remains exactly the same except that "strictly" becomes "weakly."

$(2\alpha - 1)$ *with support* $[0, 1]$*, and each uninformed trader's hedging benefit* $b_i \overset{iid}{\sim} f$*. Given any pairs* $(\theta_l, \gamma_l) < (\theta_h, \gamma_h) < (1, 1)$ *that satisfy* $\theta_h + \gamma_l > 1$ *and* $\theta_l + \gamma_h > 1$*, there exists cutoffs* $\mu'_h > \mu'_l$*,* $\tilde{\mu}'_h > \tilde{\mu}'_l$*,* $\mu_h > \mu_l$*,* $\tilde{\mu}_h > \tilde{\mu}_l$*, and* $\bar{\mu} > \underline{\mu}$*, all strictly positive, such that:*

(a) Marginally lowering the traders' label accuracy θ *from* θ'_h *to* $\theta'_h - d\theta$ *strictly raises welfare* W *across all mass of informed traders* $\mu < \mu'_l$ *and strictly lowers* W *across all* $\mu > \mu'_h$*; marginally lowering* γ *from* γ'_h *to* $\gamma'_h - d\gamma$ *strictly raises* W *across all* $\mu < \tilde{\mu}'_l$ *and strictly lowers* W *across all* $\mu > \tilde{\mu}'_h$*.*

(b) Lowering the label accuracy θ *from* θ_h *to* θ_l *strictly raises* W *across all* $\mu < \mu_l$ *and strictly lowers* W *across all* $\mu > \mu_h$*; lowering* γ *from* γ_h *to* γ_l *strictly raises* W *across all* $\mu < \tilde{\mu}_l$ *and strictly lowers* W *across all* $\mu > \tilde{\mu}_h$*.*

(c) Closing the OTC market strictly raises W *across all* $\mu < \underline{\mu}$ *and strictly lowers* W *across all* $\mu > \bar{\mu}$*.*

Moreover:

(d) If the marginal welfare Δ_W *(defined in (7)) is strictly quasiconcave, the cutoffs* $\mu'_h = \mu'_l$*,* $\tilde{\mu}'_h = \tilde{\mu}'_l$*,* $\mu_h = \mu_l$*,* $\tilde{\mu}_h = \tilde{\mu}_l$*, and* $\bar{\mu} = \underline{\mu}$*.*

(e) If the marginal volume Δ_V *(defined in (6)) is strictly decreasing, marginally lowering* θ *or* γ*, lowering* θ *or* γ *to any lower value, or closing the OTC market strictly reduces the aggregate trade volume* V *and strictly widens the average bid-ask spread* $\bar{S}$ *across all* $\mu > 0$*.*

Proposition 2 parts (a)–(c) say, under *any* economically relevant distributions, penalizing or closing the OTC market raises welfare if the informed mass μ is sufficiently small and reduces welfare if μ is sufficiently large. Part (d) sharpens this result to a single cutoff on μ for distributions whose marginal welfares Δ_W are quasiconcave. Part (e) implies that the aggregate volume and the average spread are poor indicators of welfare: Wherever the marginal volume Δ_V is decreasing, penalizing or closing the OTC market reduces the aggregate volume and widens the average spread, no matter how they affect welfare. Later,

Proposition 3 shows that most relevant distributions admit a quasiconcave Δ_W, and lists a class of distributions whose Δ_V are decreasing.

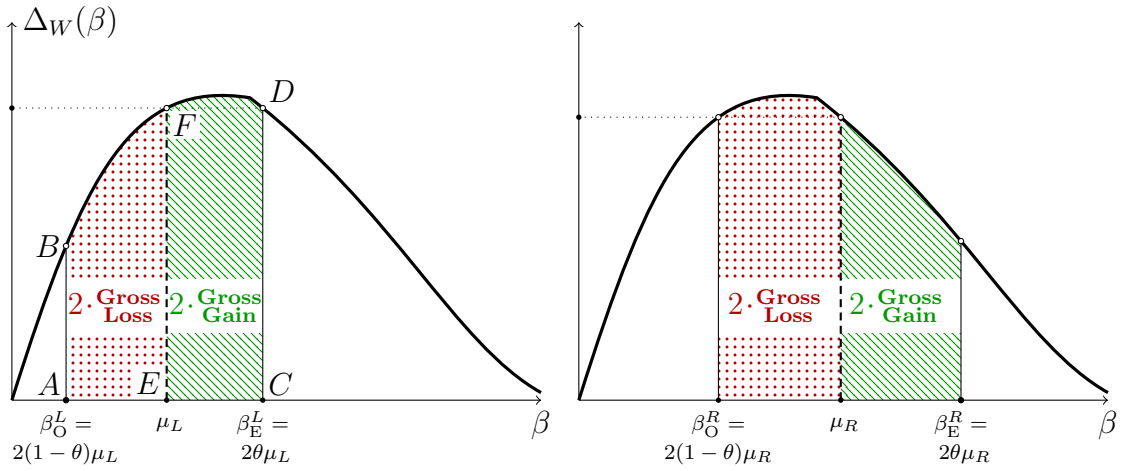

Figure 5: The Effect of Closing the OTC Market on Welfare ($\gamma = 1/2$, $\theta > 1/2$)

We now outline the proof of the single cutoff result (Proposition 2 part (d)) upon closing the OTC market. Appendix A.2 generalizes this proof for the other parts. To simplify exposition, we consider the example where the label accuracy $\gamma = 1/2$ and the marginal welfare Δ_W is continuous. Figure 5 plots a strictly quasiconcave Δ_W. On the left panel, the mass of informed traders is μ_L. The initial ratio of informed to uninformed traders β in the OTC market is $\beta_{\mathrm{O}}^L = 2(1-\theta)\mu_L < \mu_L$ and on the exchange is $\beta_{\mathrm{E}}^L = 2\theta\mu_L > \mu_L$. Closing the OTC market pools traders on the exchange and results in $\beta = \mu_L \in (\beta_{\mathrm{O}}^L, \beta_{\mathrm{E}}^L)$, which yields the LU traders a higher spread and the LI traders a lower spread relative to their initial spreads. Some uninformed LU traders consequently exit and cause a gross welfare loss marked by $\gamma = 1/2$ times the red-dotted area between the points $\beta = \mu_L$ and $\beta = \beta_{\mathrm{O}}^L$. Some uninformed LI traders enter and generate a gross welfare gain marked by $1 - \gamma = 1/2$ times the green-lined area between the points $\beta = \mu_L$ and $\beta = \beta_{\mathrm{E}}^L$.

With the aid of Figure 5, we prove Proposition 2 part (d) in two steps. First, we set constants μ_L and μ_R such that the marginal welfares $\Delta_W(\mu_L) = \Delta_W(\beta_{\mathrm{E}}^L)$ and $\Delta_W(\mu_R) = \Delta_W(\beta_{\mathrm{O}}^R)$. The left panel of Figure 5 shows that the areas of the gross welfare gain and the gross welfare loss have equal base lengths $(2\theta - 1)\mu$, while the gross welfare gain is

uniformly taller than the gross welfare loss. This remains true for all informed mass $\mu \leq \mu_L$. Symmetrically, on the right panel, the gross welfare loss is strictly larger than the gross welfare gain, and remains so for all $\mu \geq \mu_R$.

Second, we prove that the net gain in welfare is strictly decreasing in the informed mass μ between $\mu = \mu_L$ and $\mu = \mu_R$. Figure 5 shows that as we marginally increase μ, the net gain in welfare changes by

$$\underbrace{\left(\theta \cdot ||\overline{CD}|| - \frac{1}{2} \cdot ||\overline{EF}||\right)}_{\text{Derivative of gross welfare gain}} - \underbrace{\left(\frac{1}{2} \cdot ||\overline{EF}|| - (1-\theta) \cdot ||\overline{AB}||\right)}_{\text{Derivative of gross welfare loss}}.$$

The derivative is strictly negative for any $\mu \in (\mu_L, \mu_R)$, as both segments $\overline{AB}$ and $\overline{CD}$ are strictly shorter than $\overline{EF}$. Hence, increasing μ in the interval (μ_L, μ_R) monotonically reduces the net gain in welfare. There must then be a unique cutoff in (μ_L, μ_R) at which the net gain in welfare crosses 0 from above. Part (d) follows.

In the case where $\gamma \neq 1/2$, the base lengths of the gross welfare gain and loss in Figure 5 differ by a factor of $\gamma/(1-\gamma)$. The masses of uninformed traders scale the gross gain and loss with an inverse factor of $(1-\gamma)/\gamma$, which precisely offsets the difference in their base lengths. Therefore, the same arguments as above establish part (d) for all values of γ. We generalize this basic proof to establish all other parts in Appendix A.2.

Economic relevance. Most economically relevant distributions F have quasiconcave marginal welfare Δ_W. Proposition 3 states the necessary and sufficient conditions on the distribution for its Δ_W to be quasiconcave and its marginal volume Δ_V to be decreasing.

Proposition 3. *(i) Marginal welfare Δ_W is strictly quasiconcave if and only if the function*

$$\frac{(2\alpha-1)(1-F(x))}{xf(x)(2\alpha-1-x)^2} - \frac{1}{2\alpha-1-x} \quad \textit{is strictly quasiconvex in } x \in (0, 2\alpha-1). \tag{8}$$

Any beta distribution Beta(a,b) for all $a,b>0$ satisfies condition (8).

(ii) Marginal volume Δ_V is decreasing if and only if the function

$$\frac{(2\alpha-1)(1-F(x))}{f(x)(2\alpha-1-x)^2}-\frac{x}{2\alpha-1-x} \quad \textit{is increasing in } x\in(0,2\alpha-1). \tag{9}$$

Any beta distribution Beta(a,b) for all $a\leq 1, b\leq 1$ satisfies condition (9).

Taylor expansions around $2\alpha-1$ verify that the uniform distribution $\mathbb{U}[0,1]$ and any beta distribution—for all parameter values—satisfy condition (8) and thereby have quasiconcave marginal welfare Δ_W.[12] To see what distributions are excluded by (8), we note that (8) holds at the extreme ends under *any* pdf f analytic in some neighborhoods of 0 and $2\alpha-1$: The expression in (8) approaches infinity at the limit as $x\to 0$ and as $x\to 2\alpha-1$. Condition (8) only excludes pdf f that oscillate over moderate values in the support $(0,2\alpha-1)$.

We further verify that (9) is satisfied for any uniform distribution and for beta distributions Beta(a,b) with parameters $a\leq 1, b\leq 1$. That a wide range of distributions satisfy (9) shows how the aggregate volume and the average bid-ask spread are poor indicators of policies' impact on welfare.

3.4 Market Shares

Our results so far link welfare to adverse selection risk. Because regulatory debates often cite market shares, we analyze whether and how the OTC market share relates to welfare. Proposition 4 formalizes one message: Penalizing OTC trading can strictly raise welfare specifically where the OTC market share is high.

Proposition 4. *Given any α, we let pdf f and its cdf F have support $[0,1]$ and satisfy*

$$x\cdot\left(\frac{1}{2\alpha-1-x}-\frac{f(x)}{1-F(x)}\right) \quad \textit{is strictly increasing in} \quad x\in(0,2\alpha-1),\text{[13]} \tag{10}$$

and each uninformed trader's hedging benefit $b_i\overset{iid}{\sim} f$. Given any pairs $(\theta_l,\gamma_l)<(\theta_h,\gamma_h)<(1,1)$ that satisfy $\theta_h+\gamma_l>1$ and $\theta_l+\gamma_h>1$, there exist cutoffs $M,\widehat{M},M^*>0$ such that:

[12] The beta distribution Beta(a,b) (pdf $f(x)=\frac{x^{a-1}(1-x)^{b-1}}{B(a,b)}$) is a general class of bounded distributions that embeds the uniform distribution when $a=b=1$. We can numerically verify that common unbounded distributions—Normal, Chi-squared, and Gamma—satisfy condition (8) when truncated to $[0,1]$.

(a) OTC market share $\frac{V_O}{V}$ *is strictly decreasing in* μ.

(b) Lowering the label accuracy θ *from* θ_h *to* θ_l *strictly raises welfare* W *across all* $\frac{V_O}{V} > M$ *and strictly reduces* W *across all* $\frac{V_O}{V} < M$*; lowering* γ *from* γ_h *to* γ_l *strictly raises welfare* W *across all* $\frac{V_O}{V} > \widehat{M}$ *and strictly reduces* W *across all* $\frac{V_O}{V} < \widehat{M}$.

(c) Closing the OTC market strictly raises W *across all* $\frac{V_O}{V} > M^*$ *and strictly reduces* W *across all* $\frac{V_O}{V} < M^*$.

Intuitively, a larger μ mechanically raises the share of Likely Informed traders, who trade on the exchange. Under condition (10), this mechanical effect causes the OTC market share $\frac{V_O}{V}$ to strictly decrease in μ (part (a)). Parts (b) and (c) are corollaries of part (a) and Proposition 2: Keeping all parameters except μ equal, penalizing or closing the OTC market would raise welfare *if and only if* its market share $\frac{V_O}{V}$ is sufficiently high.

Raising μ, all else equal, simultaneously reduces $\frac{V_O}{V}$ and widens the spread on the exchange S_E. Stated empirically, adverse selection risk induces a positive correlation between the exchange spread S_E and its market share $\frac{V_E}{V}$. We document a positive correlation between quoted spreads on exchanges and their market share across US-listed equities (Internet Appendix V.3).

4 Optimal Pigouvian Tax

Section 3 provides policy guidance where one knows that adverse selection risk is low (penalize or close the OTC market) or high (keep the OTC market). We now introduce an optimal Pigouvian tax that one can implement without knowing the severity of adverse selection risk. Instead, a simple sufficient statistic derived from a local experiment determines whether the current tax is too low or too high. Computing this statistic, Weighted Spread Ratio (WSR), only requires trade volumes and bid-ask spreads.

[13] The uniform distribution, $f = \mathbb{U}[0,1]$, trivially meets (10). We numerically verify that, given any α, the beta distribution satisfies (10) for wide ranges of its parameters $a, b > 0$.

4.1 Characterizing the Optimal Pigouvian Tax

We analyze the Pigouvian tax under the uniform distribution, then generalize to any distribution (p. 27). Gross tax revenue t is transferred from the dealer to the market maker.[14] The tax t is common knowledge from the beginning of the game. We fix all parameters $\{\mu, \theta, \gamma, \alpha\}$ except the tax t.

All implementations of a Pigouvian tax levied on the dealer that raise the same gross revenue t are equivalent to each other, as they all lead to the same zero-profit conditions.[15] Precisely, the dealer's condition becomes

$$S_{\mathrm{O}}(t) \cdot \underbrace{\big(1 - S_{\mathrm{O}}(t) + \beta_{\mathrm{O}}\big)\gamma}_{V_{\mathrm{O}}(t)} = (2\alpha - 1)\beta_{\mathrm{O}}\gamma + t, \tag{11}$$

and the market maker's condition becomes

$$S_{\mathrm{E}}(t) \cdot \underbrace{\big(1 - S_{\mathrm{E}}(t) + \beta_{\mathrm{E}}\big) \cdot (1 - \gamma)}_{V_{\mathrm{E}}(t)} = (2\alpha - 1)\beta_{\mathrm{E}} \cdot (1 - \gamma) - t. \tag{12}$$

Increasing the tax t raises the OTC spread $S_{\mathrm{O}}(t)$ and lowers the exchange spread $S_{\mathrm{E}}(t)$.

We define the *optimal tax* as follows.

$$\text{The } \textit{optimal Pigouvian tax } t^* \text{ is the tax } t \text{ that maximizes welfare } W. \tag{13}$$

Levying this optimal tax t^* attains a strictly higher welfare than closing the OTC market.

Proposition 5. *For $f = \mathbb{U}[0,1]$ and given any triple (α, θ, γ): (a) There exists a unique optimal Pigouvian tax t^* (defined in (13)). (b) Imposing the optimal tax t^* attains a strictly higher welfare W than closing the OTC market.*

The optimal Pigouvian tax *always* outperforms closing the OTC market. This is because cheap substitution vanishes under the extreme tax that effectively closes the OTC market.

[14] Taxing the market maker is ineffective, as every LI trader would choose the OTC market and no trade would occur on the exchange.

[15] Examples include flat per-trade transaction tax, lump-sum tax, and non-linear levy on OTC volume.

We define this *Zero-OTC* tax $\bar{t}$ as follows:

Zero-OTC tax $\bar{t}$ is the smallest tax such that the OTC trade volume is zero.

At the Zero-OTC tax $\bar{t}$, each trader is indifferent between trading over the counter or on the exchange. In particular,

$$S_{\mathrm{O}}(\bar{t}) = S_{\mathrm{E}}(\bar{t}).$$

Thereby, the Zero-OTC tax $\bar{t}$ equalizes the hedging benefits of the marginal uninformed traders in both markets. A marginal tax cut from $\bar{t}$ would create entrants and exiters with the same hedging benefits, eliminating cheap substitution. Because a tax cut from $\bar{t}$ strictly increases the aggregate trade volume with the uniform distribution, the optimal tax here is strictly below the Zero-OTC tax, $t^* < \bar{t}$.

4.2 Implementing the Optimal Pigouvian Tax

We use the trade-off between the cheap substitution versus the volume effect (defined in (3)) to develop a simple precedure to implement the optimal Pigouvian tax t^*. Marginally increasing the tax t raises welfare W if and only if

$$\underbrace{\left|\frac{S_{\mathrm{E}}}{S_{\mathrm{O}}}\right|}_{\substack{\text{Cheap}\\ \text{Substitution}}} > \underbrace{\left|\frac{dV_{\mathrm{O}}}{dV_{\mathrm{E}}}\right|}_{\substack{\text{Volume}\\ \text{Effect}}}.$$

The spread ratio $S_{\mathrm{E}}/S_{\mathrm{O}}$ captures cheap substitution, which raises welfare upon the tax increase. The volume ratio $|dV_{\mathrm{O}}/dV_{\mathrm{E}}|$ captures the volume effect, which reduces welfare.

We define the ratio of cheap substitution to the volume effect as the *weighted spread ratio*:

$$\mathrm{WSR}(t) \coloneqq \left|\frac{S_{\mathrm{E}}(t)dV_{\mathrm{E}}(t)}{S_{\mathrm{O}}(t)dV_{\mathrm{O}}(t)}\right|. \tag{14}$$

The WSR allows us to derive a simple "WSR Rule" for tax setting: Raise the tax t when $\mathrm{WSR}(t) > 1$, cut t when $\mathrm{WSR}(t) < 1$, and keep the current tax when $\mathrm{WSR}(t) = 1$. Welfare strictly increases whenever one changes the tax following the WSR Rule. Moreover, under

the uniform distribution $F = \mathbb{U}[0,1]$, iteratively following the WSR Rule obtains the optimal tax t^*.

Proposition 6. *We suppose $F = \mathbb{U}[0,1]$ and the current Pigouvian tax is $t \geq 0$. If the tax marginally increases by $dt > 0$ and the consequent weighted spread ratio is* $\mathrm{WSR}(t)$ *(defined in (14)):*

(a) Welfare W increases ($dW/dt > 0$) if and only if $\mathrm{WSR}(t) > 1$*, and W decreases ($dW/dt < 0$) if and only if* $\mathrm{WSR}(t) < 1$.

(b) The current tax is strictly lower than the optimal tax $t < t^*$ if and only if $\mathrm{WSR}(t) > 1$*, and $t > t^*$ if and only if* $\mathrm{WSR}(t) < 1$.

(c) $\mathrm{WSR}(t)$ *is strictly decreasing in t.*

Proposition 6 spells out a procedure to implement the optimal Pigouvian tax t^*:

Procedure

1. Introduce a tax near zero.
2. Calculate the weighted spread ratio WSR.
3. If WSR is above 1, raise the tax; if WSR is below 1, lower the tax; otherwise, do nothing and end this procedure (the WSR Rule).
4. Repeat steps 2 and 3 until the procedure ends.

Pigouvian tax under general distributions

Remark. We let f be a pdf with support $[0,1]$, each uninformed trader's hedging benefit $b_i \overset{iid}{\sim} f$, and the current Pigouvian tax be $t > 0$. If the tax marginally increases by $dt > 0$, welfare W increases ($dW/dt > 0$) if and only if $\mathrm{WSR}(t) > 1$ and decreases ($dW/dt < 0$) if and only if $\mathrm{WSR}(t) < 1$.

The WSR Rule holds under any pdf f with support $[0,1]$. Shifting the tax t following the WSR Rule would strictly raise welfare. When $\mathrm{WSR}(t) = 1$, welfare is at a local maximum.

Hence, following our procedure implements the lowest tax that attains a locally maximal welfare. Further, we learn whether one should tax OTC trades from the first WSR computed in the procedure $\text{WSR}(t = 0)$—the optimal tax t^* is strictly positive whenever $\text{WSR}(0) > 1$.

Estimating the Weighted Spread Ratio

The volume term in the WSR, $|dV_{\text{O}}/dV_{\text{E}}|$, is easily computed using trade volumes just before and just after the change in the tax t. Already in the US and the EU, most trades must be reported in nearly real time and specify whether it was executed on an exchange or over the counter. The spread term $S_{\text{E}}/S_{\text{O}}$ is empirically the ratio of the average bid-ask spread on exchanges to the average spread over the counter just before the tax change. To estimate it, the exchange spread S_{E} can be the best quoted ask minus the best quoted bid, divided by the midpoint price. If no exchange has a bid or an ask, the effective spread from recent exchange trades can substitute for the quoted spread. In the OTC market, one can approximate the spread S_{O} with the effective spreads of similar assets (e.g., corporate bonds at the same firm with similar maturity) or adopt benchmarks widely used among traders (e.g., MarketAxess' Composite+ or Bloomberg's BVAL).

Applicability of the optimal Pigouvian tax

Our Pigouvian tax results apply to OTC-traded assets that are also traded via pre-trade anonymous methods. Such methods include limit order books, batch auctions, dark pools, and all-to-all requests for quote, which we describe in Section 6. Most important asset classes are actively traded via these methods, including equities, their futures and options, treasuries, corporate bonds, and repurchase agreements. Footnote 3 lists examples.

5 Related Mechanisms

We find that price discrimination can harm social welfare via cheap substitution. That price discrimination can be inefficient is well-known. One might wonder if cheap substitution repackages a known mechanism. Below, we discuss whether prior mechanisms can generate our two main results: *(i)* Welfare can decline while the aggregate volume increases, and *(ii)* penalizing or closing the OTC market raises welfare where adverse selection risk is low.

The original Akerlof (1970) framework

In Akerlof (1970), assets have varying common values, and an uninformed trader's private value from owning an asset is proportional to that asset's common value. Therefore, private values are heterogeneous across assets. This setting does not allow Akerlof (1970) to obtain our main results, for two reasons: (I) The private and the common values are perfectly correlated. (II) Each uninformed trader is unaware of her own private value; otherwise, she would learn the common value and face no adverse selection risk. Internet Appendix III adds imperfect labels to Akerlof (1970) and shows how the combination of these two features *reverses* cheap substitution: Upon pooling, every entrant has a *smaller* private value than any exiter. As the section shows, heterogeneity in private values that are (I') decoupled from the common value and (II') known to the uninformed traders is the minimum extension to the Akerlof (1970) framework necessary to generate our results.

Cream skimming

Bolton, Santos and Scheinkman (2016) also features cream skimming by dealers. Nonetheless, Bolton et al. (2016) cannot generate our results, because cream skimming in their framework affects welfare through a tradeoff orthogonal to adverse selection risk. On the one hand, two exogenous frictions in the OTC market lower welfare: Their dealers (a) hold market power and (b) incur a deadweight cost to separate high- and low-quality assets. On the other hand, the endogenous opportunity to sell to an informed dealer incentivizes effort in origination, which raises welfare. Therefore, cream skimming necessarily raises welfare in absence of the two exogenous frictions, (a) and (b). We strip away effort in origination, and separate the private values of traders from the common value of the asset. Under this setup, we show that cream skimming strictly lowers welfare whenever adverse selection risk is low, even with a competitive dealer and zero deadweight costs.

Non-anonymity in financial markets

Cheap substitution requires different uninformed traders to be affected in opposite directions. In our case, price discrimination lowers the trading cost of uninformed LU traders and drives up that of uninformed LI traders. In Röell (1990), revealing the orders of certain uninformed traders benefits those traders and leaves others worse off. A large literature

shows analogous effects (Admati and Pfleiderer, 1991; Forster and George, 1992; Fishman and Longstaff, 1992; Foucault, Moinas and Theissen, 2007; Rindi, 2008). All such models include noise or liquidity traders who trade an exogenous quantity at any price. Consequently, these models cannot produce cheap substitution or our main results. Indeed, the literature focuses on measures of liquidity or price discovery, rather than social welfare.[16]

Binding minimum wage

Policies we examine lead the uninformed traders with the smaller gains from trade to exit and those with larger gains to enter. Likewise, a higher minimum wage forces out the workers with the least surplus from employment, and can strictly raise social welfare under redistributive preferences (Allen, 1987; Guesnerie and Roberts, 1987; Boadway and Cuff, 2001; Lee and Saez, 2012). However, a higher minimum wage does not induce any worker into employment absent additional features such as endogenous search or effort on the job (Clemens, 2021; Manning, 2021). Without such features, the minimum wage can only reduce aggregate employment and utilitarian welfare.

6 Policy Implications

This section applies our results to recent policy debates.

Post-trade name give-up

The US Dodd-Frank Act forces most swaps trades to occur on Swap Execution Facilities (SEFs), a kind of electronic trading platform. The SEFs offer two ways to trade: Request for quote (RFQ) or all to all (A2A). In an RFQ, a trader simultaneously submits a non-anonymous trade request to multiple dealers. Because an RFQ is non-anonymous pre-trade, it is an OTC trade in our model. An A2A trade occurs either on the limit order book or through a first-price auction open to all traders, both of which are anonymous pre-trade. The A2A protocol thus represents the exchange in our model.

Most SEFs traditionally practiced post-trade name give-up (NGU). With NGU, each counterparty to an A2A trade learn the other party's identity after the trade. Buy-side

[16]The only exception is Admati and Pfleiderer (1991), whose result on welfare (their Proposition 1 (d)) is the opposite of ours.

firms intensely oppose NGU, due to the "information leakage associated with sharing [their] trading activity" (Citadel LLC, 2020). Dealers, on the other hand, claim that NGU helps tailor their pricing on RFQs and increases the overall liquidity they provide (JPMorgan Chase & Co., 2018). They also cite research that RFQs have narrower spreads and generate larger trade volumes than A2A to argue in favor of NGU (Securities Industry and Financial Markets Association, 2018). Despite the dealers' claims, the Commodity Futures Trading Commission (CFTC) banned NGU in 2020.

Our results undercut the dealers' argument for keeping NGU—trade volumes, bid-ask spreads, and market shares are *not* indicators of welfare. Moreover, banning NGU would reduce the information each dealer has about their clients, which may correspond to a lower label accuracy, θ or γ, in our model. Evidence is mixed on whether swaps trades have low adverse selection risk. If swaps trades are low risk, as several studies find,[17] our results suggest that the ban is welfare-improving.

Permissioned blockchain

Blockchain is an electronic recordkeeping procedure for a network of members, called "nodes." Each transaction between the nodes is broadcast to the other nodes in the network, and each node maintains a ledger of all transactions. An algorithm periodically reconciles these ledgers, and the sequence of reconciled ledgers forms the 'official' record. Transactions on a blockchain include pseudonymized identifiers of the counterparties involved in the trade. The reconciled record provides the complete trade history of every account on the blockchain. Since pseudonymity violates anti-money-laundering rules (Elwell, Murphy and Seitzinger, 2013), most proposed blockchains for financial markets fully reveal traders' identities to their nodes.

The Depository Trust & Clearing Corporation (DTCC)[18] proposed a plan to move the ownership records of most credit derivatives onto a blockchain (Irrera, 2017). The DTCC envisions a "permissioned" blockchain, wherein 15 major dealers act as the nodes (DTCC,

[17] As examples, credit default swap spreads tend to follow stock price and not vice versa (Narayan, Sharma and Thuraisamy, 2014; Hilscher, Pollet and Wilson, 2015), and interest rate swap spreads tend to follow futures prices (Poskitt, 2007). Augustin, Subrahmanyam, Tang and Wang (2014) surveys the asset-pricing literature on swaps.

[18] DTCC is the dominant clearinghouse for most securities.

2018). Soon thereafter, ICAP (the dominant interdealer broker) initiated a pilot of a permissioned blockchain for foreign exchange, and the DTCC proposed a permissioned blockchain for equities (DTCC, 2020).

All three proposals would reveal each trader's transaction history to the major dealers. The trade histories may aid the dealers to more accurately separate traders by their trading motives. If so, the blockchain proposals would increase the label accuracy, θ or γ, in our model. Then, our model predicts a decrease in welfare for the assets with low adverse selection risk and a widening of the bid-ask spreads on exchanges.

7 Conclusion

We show that penalizing OTC trades can improve utilitarian welfare, under the conservative setup of competitive pricing in the OTC market. In practice, search frictions and the dealers' market power hamper price competition in OTC markets. As OTC trading moves onto electronic platforms, such frictions are dissipating (Hendershott and Madhavan, 2015; O'Hara and Alex Zhou, 2021; Hau, Hoffmann, Langfield and Timmer, 2021). Price discrimination by the dealers remains a fundamental feature of OTC trading.

Our model abstracts away from some important sources of inefficiency on exchanges. We do not consider how price impact (Vives, 2011) or sniping (Budish, Cramton and Shim, 2015) might affect welfare. Several papers already propose improvements to the design of exchanges that address such frictions.[19] Our focus on resolving the inefficiency of OTC trading complements this literature.

Previous work show that price discovery in secondary markets affects corporate investment decisions (for example, Goldstein and Guembel, 2008). We leave for future research the analysis of price discovery in the presence of an exchange and an OTC market for two reasons. First, in our model, price discovery *within* each market is uninteresting—price discovery is monotonically increasing in the informed ratio β in the market. Second, analyzing

[19]For example, Malamud and Rostek (2017) and Chen and Duffie (2021) show that optimal market fragmentation can address price impact, and Budish et al. (2015) proposes frequent batch auctions to resolve sniping by fast traders.

the *aggregate* price discovery requires a stance on exactly how the quotes and the transaction prices are aggregated across the two markets. Any effect on aggregate price discovery would be driven by, for instance, the content and the timing of disclosures.

Appendix

A Proofs

A.1 Proofs for Section 2.3

Proof of Proposition 0. Part (a): It suffices to show that there exists at least one solution to the zero-profit condition (1) so that $S(\beta)$ is well-defined. We use Figure A.1, which plots the market maker's payoff $s \cdot (1-s) - (2\alpha - 1 - s)^+\beta$ over the spread s. (With a general distribution F as in Section 3.3, the payoff is $s \cdot [1 - F(s)] - (2\alpha - 1 - s)^+\beta$ and the proof is identical otherwise.) The payoff curve is continuous. Her payoff is negative at $s = 0$, as she is adversely selected yet has no revenue. It is positive at $s = 2\alpha - 1$, as she breaks-even on the trades against the informed and profits on the uninformed. The Intermediate Value Theorem implies that there exists at least one solution to the zero-profit condition. The smallest solution $S(\beta)$ is thus well-defined.

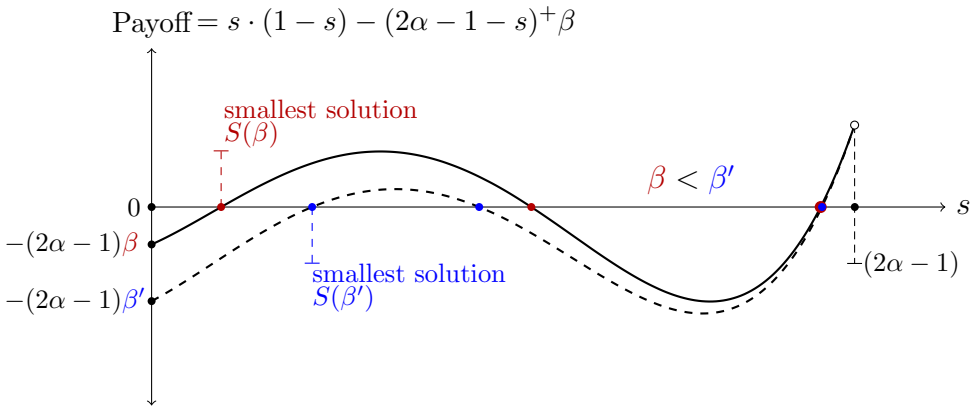

Figure A.1: Finding the equilibrium spread

Part (b): We proceed in three steps. First, we show that the spread function $S(\beta)$ is increasing in the informed ratio $\beta \in [0, \infty)$. We see this easily in Figure A.1: increasing β (to β') shifts the entire payoff curve downwards and the crossing point $S(\beta)$ to the right. Intuitively, as the informed traders impose losses on the market maker, he requires a wider spread to break-even when there are more informed traders.

Second, we solve for the spreads in the OTC market. All traders with the same label share the same OTC spread, because they are indistinguishable to the dealer. For an LU trader, the equilibrium OTC spread is $S(\beta_{\text{LU}})$, where β_{LU} is the informed ratio of the LU traders. As LU traders consist of $(1-\theta)\mu$ informed and γ uninformed traders, their OTC spread is $S\left(\frac{1-\theta}{\gamma}\,\mu\right)$. Similarly, the LI traders' OTC spread is $S\left(\frac{\theta}{1-\gamma}\,\mu\right)$.

Third, we turn to the exchange spread S_{E}. If the market maker sets $S_{\text{E}} \in (S(\beta_{\text{LU}}), S(\beta_{\text{LI}})]$, all LU traders choose the OTC market, whereas all LI traders choose the exchange. Then the informed ratio on the exchange $\beta_{\text{E}} = \beta_{\text{LI}}$. The market maker thus earns zero profit if and only if she sets $S_{\text{E}} = S(\beta_{\text{LI}})$ in this case. If the market maker sets $S_{\text{E}} \leq S(\beta_{\text{LU}})$, then every trader chooses the exchange, implying $\beta_{\text{E}} = \mu > \beta_{\text{LU}}$, and thus the market maker earns a non-zero profit. Therefore, in equilibrium, *(i)* the exchange spread is $S_{\text{E}} = S(\beta_{\text{LI}}) = S\left(\frac{\theta}{1-\gamma}\,\mu\right)$, the lowest spread that earns the market maker a zero profit, and *(ii)* all LU traders choose the OTC market, whereas all LI traders choose the exchange.

Part (c): Since $\beta_{\text{O}} < \beta_{\text{E}}$ and the spread function $S(\beta)$ is strictly increasing in the informed ratio β, then $S_{\text{O}} < S_{\text{E}}$. □

A.2 Proofs for Section 3

We proceed with the proofs in reverse order. We first prove Proposition 2. Applying Proposition 2 to the case $F = \mathbb{U}[0,1]$ proves Proposition 1. The section ends with a proof of Proposition 4.

The proof of Proposition 0 shows that the spread function $S(\beta)$ is increasing. As $S(\beta)$ is also left-continuous in β, then $S(\beta)$ is left-differentiable. We let $S'(\beta)$ be the left derivative. Marginal volume Δ_V and marginal welfare Δ_W as in (6) and (7) are thus well-defined.

Proof of Proposition 2. We first show that Δ_W begins at $\Delta_W(0) = 0$, from which Δ_W strictly increases before eventually strictly decreasing to zero. For this, we establish three properties: (i) $\lim_{s\downarrow 0} sf(s) = 0$, *(ii)* $\lim_{s\downarrow 0} (sf)'(s) \in \mathbb{R}^+ \cup \{\infty\}$, and *(iii)* $\lim_{s\downarrow 0} [\ln(sf)]'(s) = \infty$. (i) Since $sf(s)$ is analytic in a neighborhood $(0,\varepsilon)$ of 0, $\lim_{s\downarrow 0} sf(s)$ exists (the limit can but need not be infinite). Otherwise, $sf(s)$ would cross some strictly positive constant $c > 0$ infinitely often as $s \downarrow 0$, giving rise to an accumulation point of roots for the analytic function

$sf(s) - c$. Then the Identity Theorem would imply that $sf(s) - c \equiv 0$ in the neighborhood $(0, \varepsilon)$ of 0, implying that $f(s) \sim c/s$ as $s \downarrow 0$. This contradicts the integrability of the pdf f in $(0, \varepsilon)$. Further, it must be that $\lim_{s\downarrow 0} sf(s) = 0$ since f is integrable. (ii) Since $(sf)'$ is analytic in a neighborhood of 0, $\lim_{s\downarrow 0}(sf)'(s)$ exists. The limit must be either non-negative or infinity, since $\lim_{s\downarrow 0} sf(s) = 0$. (iii) Since $[\ln(sf)]'$ is analytic in a neighborhood of 0, $\lim_{s\downarrow 0}[\ln(sf)]'(s)$ exists. Since $\lim_{s\downarrow 0} sf(s) = 0$, it must be that $\lim_{s\downarrow 0}[\ln(sf)]'(s) = \infty$. As β approaches 0, $S(\beta)$ approaches 0. Then,

$$S'(\beta) = \frac{1}{\beta'(S(\beta))} = \frac{2\alpha - 1 - S(\beta)}{\frac{(2\alpha-1)[1-F(S(\beta))]}{2\alpha-1-S(\beta)} - S(\beta)f(S(\beta))} \xrightarrow{\beta\downarrow 0} 2\alpha - 1.$$

and thus property (i) implies that $\Delta_W(\beta) \to 0$ as $\beta \to 0$. One can verify that $\ln(\Delta_W)$ is differentiable in a neighborhood of $\beta = 0$ and properties (i)–(iii) imply that

$$(\ln \Delta_W)'(\beta) = [\ln(S')]'(\beta) + S'(\beta)\,[\ln(sf)]'(S(\beta)) \xrightarrow{\beta\downarrow 0} \infty.$$

Thus, $\ln(\Delta_W)$ is strictly increasing in a neighborhood of $\beta = 0$. That is, there exists some $\beta_l > 0$ below which Δ_W is strictly increasing. By a similar argument, there exists some $\beta_h > \beta_l$ above which $\Delta_W(\beta)$ is strictly decreasing to 0. Altogether, Δ_W begins at $\Delta_W(0) = 0$, from which Δ_W strictly increases before eventually strictly decreasing towards the lower limit of zero.

Closing the OTC market is equivalent to lowering label accuracy θ from the current level $\theta_h > 1-\gamma$ to the uninformative level $\theta_l := 1-\gamma$. Hence, it suffices to establish Proposition 2 for when the label accuracy θ or γ falls.

Parts (a)–(c): As the label accuracy θ falls from θ_h to θ_l, the change in aggregate trade volume V is

$$\underbrace{(1-\gamma)\int_{S\left(\frac{\theta_l}{1-\gamma}\mu\right)}^{S\left(\frac{\theta_h}{1-\gamma}\mu\right)} f(s)\,\mathrm{d}s}_{\substack{\text{Entry by uninformed} \\ \text{LI-traders}}} - \underbrace{\gamma\int_{S\left(\frac{1-\theta_h}{\gamma}\mu\right)}^{S\left(\frac{1-\theta_l}{\gamma}\mu\right)} f(s)\,\mathrm{d}s}_{\substack{\text{Exit by uninformed} \\ \text{LU traders}}},$$

which is equal to

$$(1-\gamma)\int_{\frac{\theta_l}{1-\gamma}\mu}^{\frac{\theta_h}{1-\gamma}\mu}\Delta_V(\beta)\,\mathrm{d}\beta-\gamma\int_{\frac{1-\theta_h}{\gamma}\mu}^{\frac{1-\theta_l}{\gamma}\mu}\Delta_V(\beta)\,\mathrm{d}\beta. \tag{A.1}$$

Similarly, the change in welfare is

$$\underbrace{(1-\gamma)\int_{\frac{\theta_l}{1-\gamma}\mu}^{\frac{\theta_h}{1-\gamma}\mu}\Delta_W(\beta)\,\mathrm{d}\beta}_{\text{Gross welfare gain}}-\underbrace{\gamma\int_{\frac{1-\theta_h}{\gamma}\mu}^{\frac{1-\theta_l}{\gamma}\mu}\Delta_W(\beta)\,\mathrm{d}\beta}_{\text{Gross welfare loss}}. \tag{A.2}$$

The proofs are intuitive with the aid of graphs. Figure A.2 plots a generic Δ_W. We

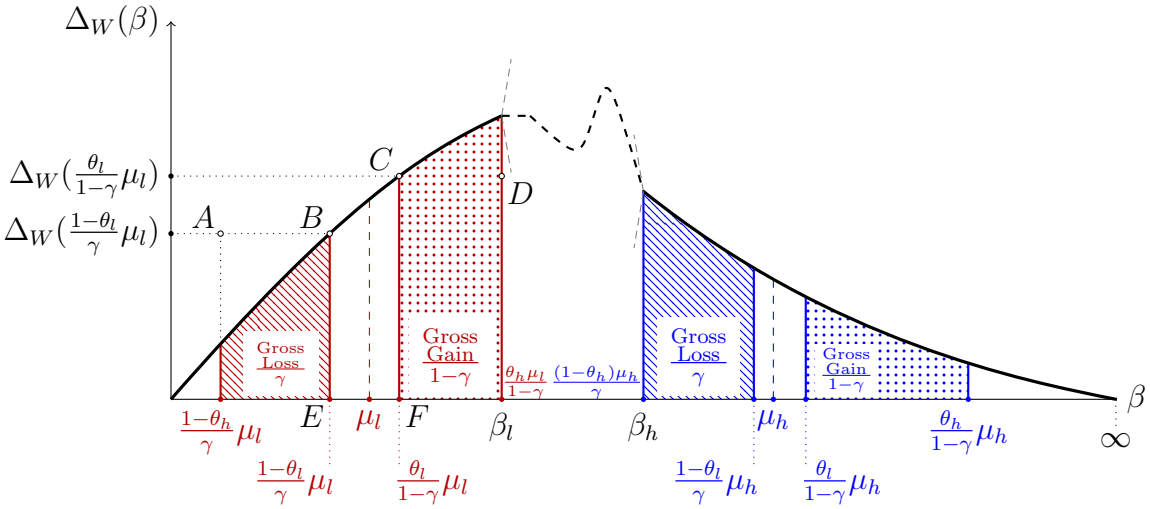

Figure A.2: Generic Δ_W

define μ_l such that $\frac{\theta_h}{1-\gamma}\mu_l=\beta_l$ (marked in Figure A.2). If $\mu=\mu_l$, the second integral in (A.2) (marked "Gross Loss/γ" in red) has a strict upper bound

$$\left(\frac{1-\theta_l}{\gamma}\mu_l-\frac{1-\theta_h}{\gamma}\mu_l\right)\cdot\Delta_W\left(\frac{1-\theta_l}{\gamma}\mu_l\right)=\frac{\theta_h-\theta_l}{\gamma}\mu_l\cdot||\overline{BE}||,$$

which corresponds to the area ABE. The first integral in (A.2) (marked "Gross Gain/$(1-\gamma)$" in red) has a strict lower bound

$$\frac{\theta_h-\theta_l}{1-\gamma}\mu_l\cdot||\overline{CF}||,$$

marked by the area DCF. Since the segment $\overline{CF}$ is longer than $\overline{BE}$ (because Δ_W is strictly increasing in $\beta \in [0, \beta_l]$), the Gross Gain in welfare is strictly larger than the Gross Loss. The same argument applies to any $\mu < \mu_l$, so that welfare rises if the mass of informed traders μ is small. Likewise, we choose μ_h such that $\frac{1-\theta_h}{\gamma}\mu_h = \beta_h$ and follow analogous steps to show that the Gross Loss in welfare (in blue) is larger than the Gross Gain if μ is large $\mu \geq \mu_h$. The above continues to hold in the limit where $\theta_l \uparrow \theta_h$, which corresponds to marginally lowering the label accuracy θ (part (a)).

As the label accuracy γ falls from γ_h to γ_l, the change in aggregate trade volume V is

$$\underbrace{(1-\gamma_h)\int_{\frac{\theta}{1-\gamma_l}\mu}^{\frac{\theta}{1-\gamma_h}\mu}\Delta_V(\beta)\,\mathrm{d}\beta}_{\substack{\text{Entry by uninformed}\\ \text{LI-traders}}} - \underbrace{\gamma_h\int_{\frac{1-\theta}{\gamma_h}\mu}^{\frac{1-\theta}{\gamma_l}\mu}\Delta_V(\beta)\,\mathrm{d}\beta}_{\substack{\text{Exit by uninformed}\\ \text{LU traders}}} - \underbrace{(\gamma_h-\gamma_l)\int_{\frac{1-\theta}{\gamma_l}\mu}^{\frac{\theta}{1-\gamma_l}\mu}\Delta_V(\beta)\,\mathrm{d}\beta}_{\text{Exit by relabeled traders}}. \tag{A.3}$$

Similarly, the change in welfare is

$$(1-\gamma_h)\int_{\frac{\theta}{1-\gamma_l}\mu}^{\frac{\theta}{1-\gamma_h}\mu}\Delta_W(\beta)\,\mathrm{d}\beta - \gamma_h\int_{\frac{1-\theta}{\gamma_h}\mu}^{\frac{1-\theta}{\gamma_l}\mu}\Delta_W(\beta)\,\mathrm{d}\beta - (\gamma_h-\gamma_l)\int_{\frac{1-\theta}{\gamma_l}\mu}^{\frac{\theta}{1-\gamma_l}\mu}\Delta_W(\beta)\,\mathrm{d}\beta. \tag{A.4}$$

We define $\tilde{\mu}_l$ such that $\frac{\theta}{1-\gamma_h}\tilde{\mu}_l = \beta_l$ (marked in Figure A.3). If $\mu = \tilde{\mu}_l$, the second integral

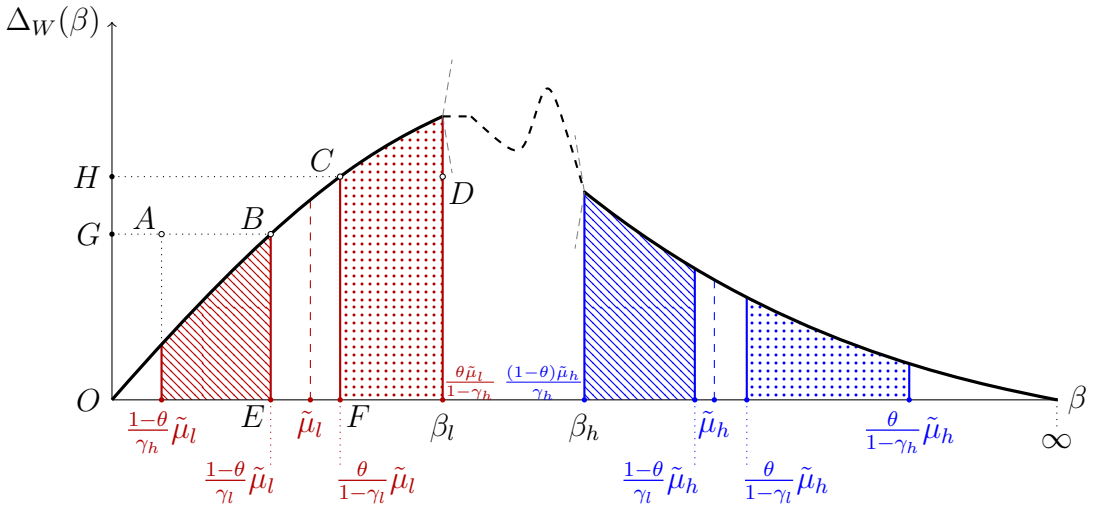

Figure A.3: Generic Δ_W

in (A.4) (marked in shaded red) has a strict upper bound

$$\left(\frac{1-\theta}{\gamma_l}\,\tilde{\mu}_l - \frac{1-\theta}{\gamma_h}\,\tilde{\mu}_l\right)\cdot\Delta_W\left(\frac{1-\theta}{\gamma_l}\,\tilde{\mu}_l\right) = (\gamma_h-\gamma_l)\,\frac{1-\theta}{\gamma_h\gamma_l}\,\tilde{\mu}_l\cdot||\overline{BE}||,$$

which corresponds to the area ABE. Thus, the second term has a strict upper bound

$$(\gamma_h-\gamma_l)\cdot\frac{1-\theta}{\gamma_l}\,\tilde{\mu}_l\cdot||\overline{BE}|| = (\gamma_h-\gamma_l)\cdot||\overline{GB}||\cdot||\overline{BE}|| = (\gamma_h-\gamma_l)\cdot||\overline{GBEO}||,$$

Similarly, the first term in (A.4) has a strict lower bound

$$(\gamma_h-\gamma_l)\cdot||HCFO||,$$

The third term in (A.4) equals $(\gamma_h-\gamma_l)\cdot||BCFE||$. Then (A.4) is strictly more than

$$(\gamma_h-\gamma_l)\left(||HCFO|| - ||GBEO|| - ||BCFE||\right) > 0.$$

The same argument applies to any $\mu < \mu_l$, so that welfare rises if the mass of informed traders μ is small. Likewise, we choose μ_h such that $\frac{1-\theta}{\gamma_h}\,\tilde{\mu}_h = \beta_h$ and follow analogous steps to show that the welfare declines if μ is large $\mu \geq \mu_h$. The above continues to hold in the limit where $\gamma_l \uparrow \gamma_h$, which corresponds to marginally lowering the label accuracy γ (part (a)).

Part (d): Figure A.4 plots a quasiconcave Δ_W. As the label accuracy θ falls from θ_h to θ_l, we choose two constants, μ_L and μ_R, as shown in Figure A.4. We set μ_L to be the highest μ such that $\Delta_W\left(\frac{1-\theta_l}{\gamma}\,\mu\right) \leq \Delta_W\left(\frac{\theta_h}{1-\gamma}\,\mu\right)$, and set μ_R to be the highest μ such that $\Delta_W\left(\frac{1-\theta_h}{\gamma}\,\mu\right) \leq \Delta_W\left(\frac{\theta_l}{1-\gamma}\,\mu\right)$. As Δ_W is quasiconcave and $\lim_{\beta\downarrow 0}\Delta_W(\beta) = \lim_{\beta\uparrow\infty}\Delta_W(\beta) = 0$, $0 < \mu_L < \mu_R < \infty$. For illustration only, Figure A.4 plots the case where Δ_W is continuous so that $\Delta_W\left(\frac{1-\theta_l}{\gamma}\,\mu_L\right) = \Delta_W\left(\frac{\theta_h}{1-\gamma}\,\mu_L\right)$ (line $\overline{BD}$ in Figure A.4a) and $\Delta_W\left(\frac{1-\theta_h}{\gamma}\,\mu_R\right) = \Delta_W\left(\frac{\theta_l}{1-\gamma}\,\mu_R\right)$ (line $\overline{IG}$ in Figure A.4b). The proof works whether or not Δ_W is continuous. We show that *(i)* the change in welfare (A.2) is strictly positive for all $\mu < \mu_L$ and strictly negative for all $\mu > \mu_R$; and *(ii)* (A.2) is strictly decreasing between μ_L and μ_R. Together, *(i)* and *(ii)* establish the existence of a single cutoff.

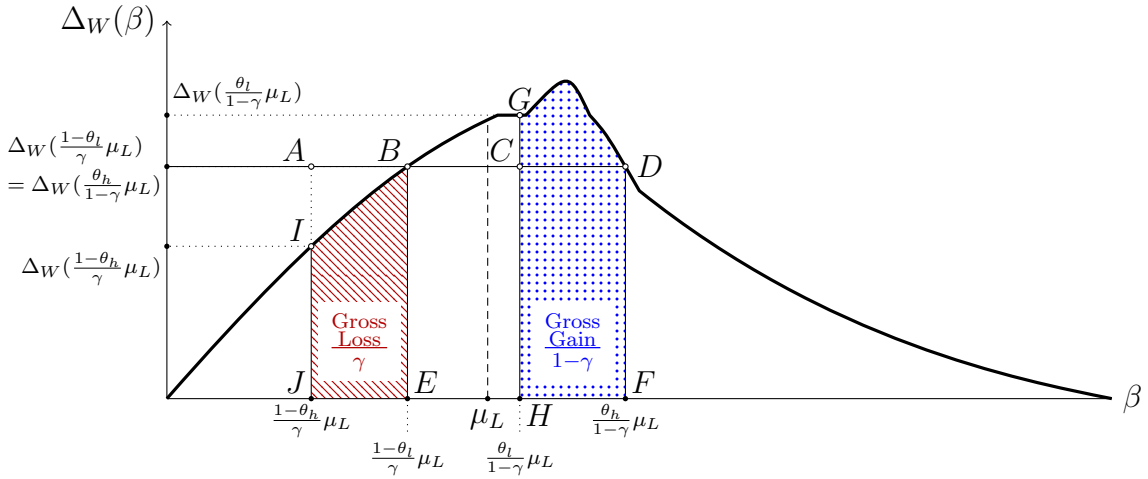

(a) Small mass of informed traders $\mu = \mu_L$

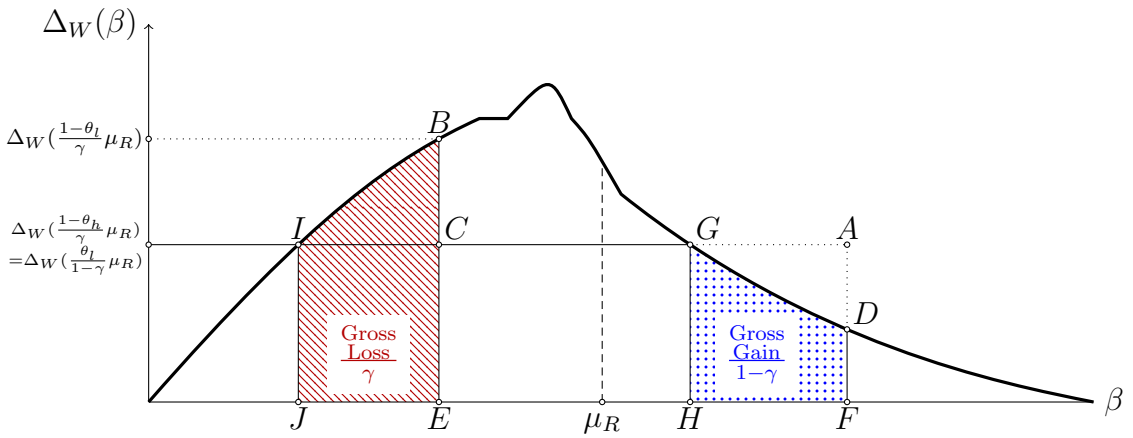

(b) Large $\mu = \mu_R > \mu_L$

Figure A.4: Quasiconcave Δ_W

To prove *(i)*, we set $\mu = \mu_L$. A strict upper bound of the second integral in (A.2) ("Gross Loss/γ" in red) is

$$\frac{\theta_h - \theta_l}{\gamma}\,\mu_L \cdot \Delta_W\left(\frac{1-\theta_l}{\gamma}\,\mu_L\right) = ||\overline{AB}|| \cdot ||\overline{BE}|| = ||ABEJ||,$$

A strict lower bound of the first integral in (A.2) ("Gross Gain/$(1-\gamma)$" in red) is $||CDFH||$. As $||\overline{BE}|| \leq ||\overline{DF}||$, the Gross Gain in welfare is strictly larger than the Gross Loss. The same argument applies to any $\mu < \mu_L$, so that (A.2) is strictly positive for all $\mu \leq \mu_L$. Similarly, (A.2) is strictly negative for all $\mu \geq \mu_R$.

To prove *(ii)* that (A.2) is strictly decreasing over $\mu \in (\mu_L, \mu_R)$, we take the derivative of (A.2) with respect to μ,

$$\underbrace{\left(\theta_h \cdot ||\overline{DF}|| - \theta_l \cdot ||\overline{GH}||\right)}_{\text{Derivative of the gross welfare gain}} - \underbrace{\left((1-\theta_l) \cdot ||\overline{BE}|| - (1-\theta_h) \cdot ||\overline{IJ}||\right)}_{\text{Derivative of the gross welfare loss}}. \tag{A.5}$$

Due to Δ_W being quasiconcave and how μ_L and μ_R are chosen, both $||\overline{BE}||$ and $||\overline{GH}||$ are strictly greater than $||\overline{DF}||$ and $||\overline{IJ}||$ when $\mu_L < \mu < \mu_R$. Then, (A.5) is strictly negative. In sum, as θ decreases from θ_h to θ_l, the change in welfare is strictly positive if $\mu \leq \mu_L$, strictly negative if $\mu \geq \mu_R$, and strictly decreasing across $\mu \in (\mu_L, \mu_R)$, which together imply that a single cutoff exists. The above continues to hold in the limit where $\theta_l \uparrow \theta_h$, which corresponds to marginally lowering the label accuracy θ.

We follow similar steps when the label accuracy γ falls from γ_h to γ_l. We choose two constants, $\tilde{\mu}_L < \tilde{\mu}_R$, as shown in Figure A.5. We set $\tilde{\mu}_L$ to be the highest μ such that $\Delta_W\left(\frac{\theta}{1-\gamma_l}\,\mu\right) \leq \Delta_W\left(\frac{\theta}{1-\gamma_h}\,\mu\right)$ (geometrically, $||\overline{GH}|| \leq ||\overline{CD}||$ in the left panel of Figure A.5), and set $\tilde{\mu}_R$ to be the highest μ such that $\Delta_W\left(\frac{1-\theta}{\gamma_h}\,\mu\right) \leq \Delta_W\left(\frac{1-\theta}{\gamma_l}\,\mu\right)$ ($||\overline{AB}|| \leq ||\overline{EF}||$ in the right panel).

For $\mu \leq \tilde{\mu}_L$, the change in welfare (A.4) is strictly more than

$$(\gamma_h - \gamma_l)\left(||O'GHO|| - ||O'E'FO|| - ||EGHF||\right) > 0.$$

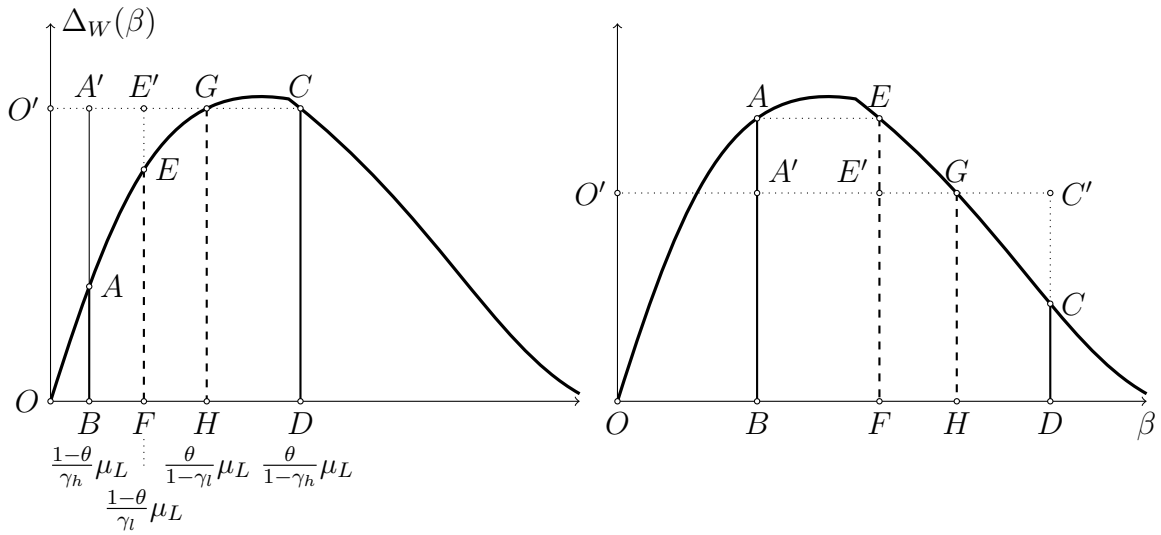

Figure A.5: Decrease in label accuracy γ

For $\mu \geq \tilde{\mu}_R$, (A.4) is strictly less than

$$(\gamma_h - \gamma_l)\left(||O'GHO|| - ||O'E'FO|| - ||EGHF||\right) < 0.$$

For $\mu \in (\tilde{\mu}_L, \tilde{\mu}_R)$, the derivative of (A.4) is

$$(1-\theta)\cdot||\overline{AB}|| + \theta\cdot||\overline{CD}|| - (1-\theta)\cdot||\overline{EF}|| - \theta\cdot||\overline{GH}|| < 0.$$

Therefore, a single cutoff exists. The above continues to hold in the limit where $\gamma_l \uparrow \gamma_h$, which corresponds to marginally lowering the label accuracy γ.

Parts (e): Figure A.6 plots a decreasing Δ_V. As the label accuracy θ falls from θ_h to θ_l, the second term of (A.1) has a strict lower bound

$$(\theta_h - \theta_l)\,\mu\cdot||\overline{BE}||,$$

which is larger than the first term's strict upper bound

$$(\theta_h - \theta_l)\,\mu\cdot||\overline{DF}||,$$

and thus the change (A.1) in aggregate volume V is strictly negative. The above continues to hold in the limit where $\theta_l \uparrow \theta_h$, which corresponds to marginally lowering the label accuracy θ.

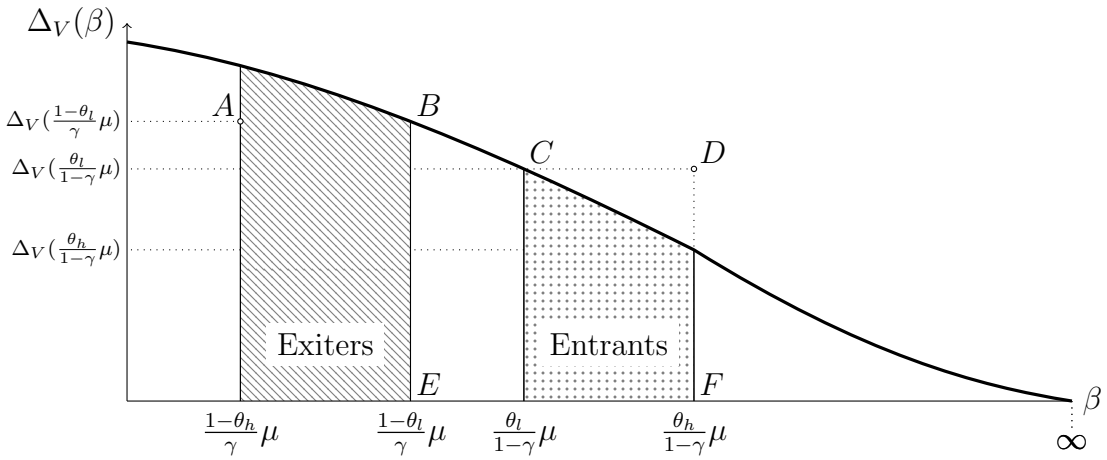

Figure A.6: Decreasing Δ_V

As the label accuracy γ falls from γ_h to γ_l, the change (A.3) in aggregate volume V is strictly less than

$$(\gamma_h - \gamma_l)(||O'GHO|| - ||O'E'FO|| - ||EGHF||) < 0,$$

marked in Figure A.7. The above continues to hold in the limit where $\gamma_l \uparrow \gamma_h$, which corresponds to marginally lowering the label accuracy γ. □

Proof of Proposition 3. The marginal welfare Δ_W can be written as

$$\Delta_W(\beta) = S'(\beta)S(\beta)f(S(\beta)) = \frac{S(\beta)f(S(\beta))}{\beta'(S(\beta))}.$$

Because the spread function $S(\beta)$ is strictly increasing in β, then $\Delta_W(\beta)$ is quasiconcave if and only if $\beta'(x)/(xf(x))$ is quasiconvex in $x \in (0, 2\alpha - 1)$. Differentiating (5) with respect to x yields

$$1 - F(x) + \beta - xf(x) = (2\alpha - 1 - x)\beta'(x),$$

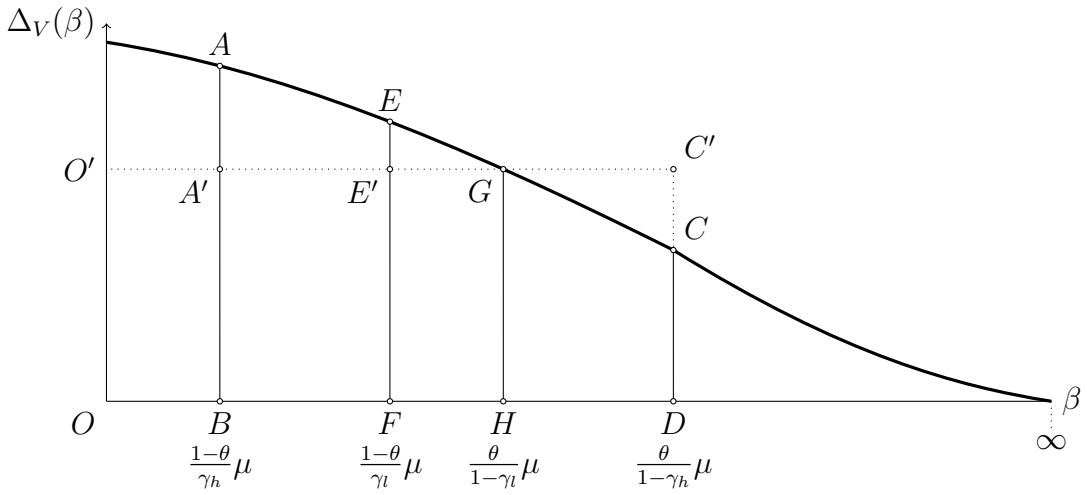

Figure A.7: Decrease in label accuracy γ

which can be rearranged to

$$\beta'(x) = \frac{1 - F(x) + \beta - xf(x)}{2\alpha - 1 - x}. \tag{A.6}$$

From (5), we can express $\beta(x)$ as function of x:

$$\beta(x) = \frac{(1 - F(x))x}{2\alpha - 1 - x}. \tag{A.7}$$

Then, marginal welfare $\Delta_W(\beta)$ is quasiconcave if and only if condition (8) holds. Likewise, marginal volume $\Delta_V(\beta)$ is decreasing if and only if (9) is true. □

Proof of Proposition 4. Part (a): The ratio $V_{\mathrm{E}}/V_{\mathrm{O}}$ equals

$$\frac{V_{\mathrm{E}}}{V_{\mathrm{O}}} := \frac{(1-\gamma)[1 - F(S(\beta_{\mathrm{E}}))] + \theta\mu}{\gamma[1 - F(S(\beta_{\mathrm{O}}))] + (1-\theta)\mu}.$$

The derivative of $V_{\mathrm{E}}/V_{\mathrm{O}}$ with respect to μ is strictly positive if and only if

$$\frac{-f(S(\beta_{\mathrm{O}}))S'(\beta_{\mathrm{O}}) + 1}{\frac{1-F(S(\beta_{\mathrm{O}}))}{\beta_{\mathrm{O}}} + 1} < \frac{-f(S(\beta_{\mathrm{E}}))S'(\beta_{\mathrm{E}}) + 1}{\frac{1-F(S(\beta_{\mathrm{E}}))}{\beta_{\mathrm{E}}} + 1}.$$

The above inequality holds for every μ, θ, and γ if and only if

$$\frac{-f(S(\beta))S'(\beta)+1}{\frac{1-F(S(\beta))}{\beta}+1}$$

is strictly increasing in β. After a change of variable using (A.6) and (A.7), the above is true if and only if

$$\frac{-\frac{(2\alpha-1-x)f(x)}{\frac{(2\alpha-1)[1-F(x)]}{2\alpha-1-x}-xf(x)}+1}{\frac{2\alpha-1-x}{x}+1}$$

is strictly increasing in $x \in (0, 2\alpha-1)$. After simplifying, the above is true if and only if

$$x\left(\frac{1}{2\alpha-1-x}-\frac{f(x)}{1-F(x)}\right)$$

is strictly increasing in $x \in (0, 2\alpha-1)$.

When $F = \mathbb{U}[0,1]$, the above expression equals

$$2(1-\alpha)\,\frac{s}{(2\alpha-1-s)(1-s)}$$

which is strictly increasing in $x \in (0, 2\alpha-1)$.

Parts (b) and (c) are corollaries of part (a) and Proposition 2. □

A.3 Proofs for Section 4

Proof of Propositions 5 and 6. We prove Propositions 5 and 6 together. First, imposing some Pigouvian tax t obtains a strictly higher welfare than closing the OTC market (Proposition 5 (b)). To show why, we proceed in three steps.

Step 1: We show that when $t = \bar{t}$, both the OTC and the exchange spreads are equal to the No-OTC spread S_{N}, $S_{\mathrm{O}}(\bar{t}) = S_{\mathrm{E}}(\bar{t}) = S_{\mathrm{N}}$. The definition of $\bar{t}$ implies $S_{\mathrm{O}}(\bar{t}) = S_{\mathrm{E}}(\bar{t})$, wherein the dealer's and the market maker's zero-profit conditions are

$$S(\bar{t})\big(1-S(\bar{t})+\beta_{\mathrm{O}}\big)\gamma = (2\alpha-1)\beta_{\mathrm{O}}\gamma+\bar{t},$$

$$S(\bar{t})\big(1-S(\bar{t})+\beta_{\mathrm{E}}\big)\cdot(1-\gamma) = (2\alpha-1)\beta_{\mathrm{E}}\cdot(1-\gamma)-\bar{t}.$$

Summing the two equations gives

$$S(\bar{t})\big(1 - S(\bar{t}) + \mu\big) = (2\alpha - 1)\mu,$$

which shows $S_{\mathrm{O}}(\bar{t}) = S_{\mathrm{E}}(\bar{t}) = S_{\mathrm{N}}$.

Step 2: We show that a marginal tax cut $-dt$ from the tax $\bar{t}$ increases the aggregate trade volume V. The volume V increases if and only if $|dV_{\mathrm{O}}(\bar{t})/dV_{\mathrm{E}}(\bar{t})| > 1$, where

$$dV_{\mathrm{O}}(\bar{t}) = S'_{\mathrm{O}}(\bar{t})\gamma\, dt \text{ and } dV_{\mathrm{E}}(\bar{t}) = S'_{\mathrm{E}}(\bar{t}) \cdot (1-\gamma)\, dt.$$

Taking the derivative of the zero-profit conditions (11)-(12) with respect to t:

$$S'_{\mathrm{O}}(t)\big(1 - 2S_{\mathrm{O}}(t) + \beta_{\mathrm{O}}\big)\gamma = 1,$$

$$S'_{\mathrm{E}}(t)\big(1 - 2S_{\mathrm{E}}(t) + \beta_{\mathrm{E}}\big) \cdot (1-\gamma) = -1.$$

Then, using that $S_{\mathrm{O}}(\bar{t}) = S_{\mathrm{E}}(\bar{t}) = S_N$,

$$\left|\frac{dV_{\mathrm{O}}(\bar{t})}{dV_{\mathrm{E}}(\bar{t})}\right| = \left|\frac{S'_{\mathrm{O}}(\bar{t})\gamma}{S'_{\mathrm{E}}(\bar{t}) \cdot (1-\gamma)}\right| = \frac{1 - 2S_{\mathrm{N}} + \beta_{\mathrm{E}}}{1 - 2S_{\mathrm{N}} + \beta_{\mathrm{O}}}.$$

The ratio $|dV_{\mathrm{O}}(\bar{t})/dV_{\mathrm{E}}(\bar{t})| > 1$ because $\beta_{\mathrm{E}} > \beta_{\mathrm{O}}$.

Step 3: We show that a strictly larger volume V implies a strictly higher welfare W. The entrants and the exiters upon a marginal change in the tax around $\bar{t}$ have the same hedging benefit. Since the entrants outnumber the exiters, welfare W is strictly higher.

Altogether, Steps 1-3 imply that the tax cut $-dt$ from $\bar{t}$ strictly raises welfare.

Second, we establish Proposition 6 (a) and (c). Upon a marginal increase dt in the tax $t < \bar{t}$, the gross welfare loss among traders over the counter is $|S_{\mathrm{O}}(t)dV_{\mathrm{O}}(t)|$. The gross welfare gain among traders on the exchange is $|S_{\mathrm{E}}(t)dV_{\mathrm{E}}(t)|$. Thus, welfare W increases ($dW/dt > 0$) if and only if $\mathrm{WSR}(t) > 1$ (Proposition 6 (a)). From the first half of this proof,

$$|S_{\mathrm{O}}(t)dV_{\mathrm{O}}(t)| = \frac{S_{\mathrm{O}}(t)}{1 - 2S_{\mathrm{O}}(t) + \beta_{\mathrm{O}}}\, dt = \frac{1}{\frac{1+\beta_{\mathrm{O}}}{S_{\mathrm{O}}(t)} - 2}\, dt,$$

and

$$|S_{\mathrm{E}}(t)dV_{\mathrm{E}}(t)| = \frac{S_{\mathrm{E}}(t)}{1 - 2S_{\mathrm{E}}(t) + \beta_{\mathrm{E}}}\,dt = \frac{1}{\frac{1+\beta_{\mathrm{E}}}{S_{\mathrm{E}}(t)} - 2}\,dt.$$

Because $S_O(t)$ is strictly increasing in t while $S_{\mathrm{E}}(t)$ is strictly decreasing, $|S_O(t)dV_O(t)|$ is strictly increasing in t and $|S_E(t)dV_E(t)|$ is strictly decreasing. Hence,

$$\mathrm{WSR}(t) = \left|\frac{S_{\mathrm{E}}(t)dV_{\mathrm{E}}(t)}{S_{\mathrm{O}}(t)dV_{\mathrm{O}}(t)}\right| = \frac{\frac{1+\beta_{\mathrm{O}}}{S_{\mathrm{O}}(t)} - 2}{\frac{1+\beta_E}{S_E(t)} - 2}$$

is strictly decreasing in t (Proposition 6 (c)).

Therefore, there exists a unique $t^* \in [0, \bar{t})$ such that $\mathrm{WSR}(t) > 1$ for $t \in [0, t^*)$ and $\mathrm{WSR}(t) < 1$ for $t > (t^*, \bar{t})$. Proposition 6 (a) then implies that $dW/dt > 0$ for $t < t^*$ and $dW/dt < 0$ for $t > t^*$. That is, t^* is the unique optimal tax that maximizes welfare W. Proposition 5 (a) and Proposition 6 (b) follow. □

References

Admati, Anat R. and Paul Pfleiderer, "Sunshine Trading and Financial Market Equilibrium," *The Review of Financial Studies*, 1991, *4* (3), 443–481.

Aguirre, Iñaki, Simon Cowan, and John Vickers, "Monopoly Price Discrimination and Demand Curvature," *American Economic Review*, September 2010, *100* (4), 1601–1615.

Akerlof, George A., "The Market for "Lemons": Quality Uncertainty and the Market Mechanism," *The Quarterly Journal of Economics*, 1970, *84* (3), 488–500.

Allen, Stephen P., "Taxes, Redistribution, and the Minimum Wage: A Theoretical Analysis*," *The Quarterly Journal of Economics*, August 1987, *102* (3), 477–489.

Augustin, Patrick, Marti G. Subrahmanyam, Dragon Yongjun Tang, and Sarah Qian Wang, "Credit Default Swaps: A Survey," *Foundations and Trends® in Finance*, December 2014, *9* (1–2), 1–196.

Babus, Ana and Cecilia Parlatore, "Strategic Fragmented Markets," *Journal of Financial Economics*, September 2021.

Baldauf, Markus and Joshua Mollner, "Trading in Fragmented Markets," *Journal of Financial and Quantitative Analysis*, February 2021, *56* (1), 93–121.

Benveniste, Lawrence M., Alan J. Marcus, and William J. Wilhelm, "What's Special about the Specialist?," *Journal of Financial Economics*, 1992, *32* (1), 61–86.

Bergemann, Dirk, Benjamin Brooks, and Stephen Morris, "The Limits of Price Discrimination," *American Economic Review*, March 2015, *105* (3), 921–957.

Bernhardt, Dan, Vladimir Dvoracek, Eric Hughson, and Ingrid M. Werner, "Why Do Larger Orders Receive Discounts on the London Stock Exchange?," *Review of Financial Studies*, 2005, *18* (4), 1343–1368.

Bessembinder, Hendrik and Kumar Venkataraman, "Does an Electronic Stock Exchange Need an Upstairs Market?," *Journal of Financial Economics*, July 2004, *73* (1), 3–36.

Biais, Bruno and Richard C. Green, "The Microstructure of the Bond Market in the 20th Century," *GSIA Working Papers*, August 2007.

_ , Thierry Foucault, and François Salanié, "Floors, Dealer Markets and Limit Order Markets," *Journal of Financial Markets*, September 1998, *1* (3), 253–284.

Boadway, Robin and Katherine Cuff, "A Minimum Wage Can Be Welfare-Improving and Employment-Enhancing," *European Economic Review*, March 2001, *45* (3), 553–576.

Bogousslavsky, Vincent and Pierre Collin-Dufresne, "Liquidity, Volume, and Order Imbalance Volatility," *Journal of Finance, Forthcoming*, February 2022.

Bolton, Patrick, Tano Santos, and Jose A. Scheinkman, "Cream-Skimming in Financial Markets," *The Journal of Finance*, 2016, *71* (2), 709–736.

Booth, G. Geoffrey, Ji-Chai Lin, Teppo Martikainen, and Yiuman Tse, "Trading and Pricing in Upstairs and Downstairs Stock Markets," *The Review of Financial Studies*, July 2002, *15* (4), 1111–1135.

Budish, Eric, Peter Cramton, and John Shim, "The High-Frequency Trading Arms Race: Frequent Batch Auctions as a Market Design Response," *The Quarterly Journal of Economics*, November 2015, *130* (4), 1547–1621.

Chao, Yong, Chen Yao, and Mao Ye, "Why Discrete Price Fragments U.S. Stock Exchanges and Disperses Their Fee Structures," *The Review of Financial Studies*, March 2019, *32* (3), 1068–1101.

Chen, Daniel and Darrell Duffie, "Market Fragmentation," *American Economic Review*, 2021, *111*, 2247–2274.

Cheng, Ing-Haw and Wei Xiong, "Why Do Hedgers Trade so Much?," *The Journal of Legal Studies*, 2014, *43* (S2), S183–S207.

Citadel LLC, "Re: Post-Trade Name Give-Up on Swap Execution Facilities (RIN 3038–AE79)," March 2020.

Clemens, Jeffrey, "How Do Firms Respond to Minimum Wage Increases? Understanding the Relevance of Non-Employment Margins," *Journal of Economic Perspectives*, February 2021, *35* (1), 51–72.

Collin-Dufresne, Pierre, Benjamin Junge, and Anders B. Trolle, "Market Structure and Transaction Costs of Index CDSs," *The Journal of Finance*, 2020, *75* (5), 2719–2763.

de Roure, Calebe, Emanuel Moench, Loriana Pelizzon, and Michael Schneider, "OTC Discount," September 2021.

Desgranges, Gabriel and Thierry Foucault, "Reputation-Based Pricing and Price Improvements," *Journal of Economics and Business*, 2005, *57* (6), 493–527.

DTCC, "DTCC Enters Test Phase on Distributed Ledger Project for Credit Derivatives with MarkitSERV & 15 Leading Global Banks," Technical Report, Depository Trust & Clearing Corporation November 2018.

_ , "Project ION Case Study," Technical Report, Depository Trust & Clearing Corporation May 2020.

Dugast, Jérôme, Semih Üslü, and Pierre-Olivier Weill, "A Theory of Participation in OTC and Centralized Markets," *The Review of Economic Studies*, February 2022, p. rdac010.

Elwell, Craig K, M Maureen Murphy, and Michael V Seitzinger, "Bitcoin : Questions , Answers , and Analysis of Legal Issues," Technical Report 7-5700, Congressional Research Service October 2013.

Exton, Gareth and Rebecca Healey, "Liquidity Landscape H1 2022," Technical Report, Liquidnet August 2022.

Fishman, Michael J. and Francis A. Longstaff, "Dual Trading in Futures Markets," *The Journal of Finance*, 1992, *47* (2), 643–671.

Forster, Margaret M. and Thomas J. George, "Anonymity in Securities Markets," *Journal of Financial Intermediation*, 1992, *2* (2), 168–206.

Foucault, Thierry, Sophie Moinas, and Erik Theissen, "Does Anonymity Matter in Electronic Limit Order Markets?," *Review of Financial Studies*, 2007, *20* (5), 1707–1747.

GFI LLC, "GFI Swaps Exchange LLC Rulebook," Technical Report November 2022.

Glode, Vincent and Christian C. Opp, "Over-the-Counter vs. Limit-Order Markets: The Role of Traders' Expertise," *Review of Financial Studies*, 2019, *33* (2), 866–915.

Glosten, Lawrence R. and Paul R. Milgrom, "Bid, Ask and Transaction Prices in a Specialist Market with Heterogeneously Informed Traders," *Journal of Financial Economics*, March 1985, *14* (1), 71–100.

Goldstein, Itay and Alexander Guembel, "Manipulation and the Allocational Role of Prices," *The Review of Economic Studies*, January 2008, *75* (1), 133–164.

Guesnerie, Roger and Kevin Roberts, "Minimum Wage Legislation as a Second Best Policy," *European Economic Review*, February 1987, *31* (1), 490–498.

Hau, Harald, Peter Hoffmann, Sam Langfield, and Yannick Timmer, "Discriminatory Pricing of Over-the-Counter Derivatives," *Management Science*, March 2021.

Hendershott, Terrence and Ananth Madhavan, "Click or Call? Auction versus Search in the over-the-Counter Market," *Journal of Finance*, 2015, *70* (1), 419–447.

_ **and Haim Mendelson**, "Crossing Networks and Dealer Markets: Competition and Performance," *The Journal of Finance*, October 2000, *55* (5), 2071–2115.

Hilscher, Jens, Joshua M. Pollet, and Mungo Wilson, "Are Credit Default Swaps a Sideshow? Evidence That Information Flows from Equity to CDS Markets," *Journal of Financial and Quantitative Analysis*, June 2015, *50* (3), 543–567.

Homer, Sidney, "The Historical Evolution of Today's Bond Market," in "Explorations in Economic Research," Vol. 2 of *NBER Books*, Cambridge, MA: National Bureau of Economic Research, Inc, 1975, pp. 378–389.

Irrera, Anna, "DTCC Completes Blockchain Repo Test," https://www.reuters.com/article/us-dtcc-blockchain-repos/dtcc-completes-blockchain-repo-test-idUSKBN1661L9 February 2017.

JPMorgan Chase & Co., "Re: Post-Trade Name Give-Up on Swap Execution Facilities: Proposed Rule – RIN 3038-AE79, 84 Fed. Reg. 72262," December 2018.

Lee, David and Emmanuel Saez, "Optimal Minimum Wage Policy in Competitive Labor Markets," *Journal of Public Economics*, October 2012, *96* (9), 739–749.

Lee, Tomy, "Latency in Fragmented Markets," *Review of Economic Dynamics*, July 2019, *33*, 128–153.

Madhavan, Ananth and Minder Cheng, "In Search of Liquidity: Block Trades in the Upstairs and Downstairs Markets," *The Review of Financial Studies*, January 1997, *10* (1), 175–203.

Malamud, Semyon and Marzena Rostek, "Decentralized Exchange," *American Economic Review*, November 2017, *107* (11), 3320–3362.

Malinova, Katya and Andreas Park, "Liquidity, Volume and Price Efficiency: The Impact of Order vs. Quote Driven Trading," *Journal of Financial Markets*, February 2013, *16* (1), 104–126.

Manning, Alan, "The Elusive Employment Effect of the Minimum Wage," *Journal of Economic Perspectives*, February 2021, *35* (1), 3–26.

Narayan, Paresh Kumar, Susan Sunila Sharma, and Kannan Sivananthan Thuraisamy, "An Analysis of Price Discovery from Panel Data Models of CDS and Equity Returns," *Journal of Banking & Finance*, April 2014, *41*, 167–177.

O'Hara, Maureen and Xing Alex Zhou, "The Electronic Evolution of Corporate Bond Dealers," *Journal of Financial Economics*, May 2021, *140* (2), 368–390.

Pagano, Marco, "Trading Volume and Asset Liquidity," *The Quarterly Journal of Economics*, 1989, *104* (2), 255–274.

_ **and Ailsa Roell**, "Transparency and Liquidity: A Comparison of Auction and Dealer Markets with Informed Trading," *Journal of Finance*, 1996, *51* (2), 579–611.

Pagnotta, Emiliano S and Thomas Philippon, "Competing on Speed," *Econometrica*, May 2018, *86* (3), 1067–1115.

Pigou, A., *The Economics of Welfare*, St. Martin's Street, London: MacMillan and Co., Limited, 1920.

Poskitt, Russell, "Benchmark Tipping and the Role of the Swap Market in Price Discovery," *Journal of Futures Markets*, 2007, *27* (10), 981–1001.

Riggs, Lynn, Esen Onur, David Reiffen, and Haoxiang Zhu, "Swap Trading after Dodd-Frank: Evidence from Index CDS," *Journal of Financial Economics*, September 2020, *137* (3), 857–886.

Rindi, Barbara, "Informed Traders as Liquidity Providers: Anonymity, Liquidity and Price Formation," *Review of Finance*, January 2008, *12* (3), 497–532.

Röell, Ailsa, "Dual-Capacity Trading and the Quality of the Market," *Journal of Financial Intermediation*, June 1990, *1* (2), 105–124.

Rose, Annica, "The Informational Effect and Market Quality Impact of Upstairs Trading and Fleeting Orders on the Australian Securities Exchange," *Journal of Empirical Finance*, September 2014, *28* (Supplement C), 171–184.

Rust, John and George Hall, "Middlemen versus Market Makers: A Theory of Competitive Exchange," *Journal of Political Economy*, 2003, *111* (2), 353–403.

Securities Industry and Financial Markets Association, "Re: Post-Trade Name Give-Up on Swap Execution Facilities; Request for Comment – RIN 3038-AE79, 83 Fed. Reg. 61751," November 2018.

Selten, R., "Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games," *International Journal of Game Theory*, March 1975, *4* (1), 25–55.

Seppi, Duane J., "Equilibrium Block Trading and Asymmetric Information," *The Journal of Finance*, 1990, *45* (1), 73–94.

Smith, Brian F., D. Alasdair S. Turnbull, and Robert W. White, "Upstairs Market for Principal and Agency Trades: Analysis of Adverse Information and Price Effects," *The Journal of Finance*, 2001, *56* (5), 1723–1746.

Vives, Xavier, "Strategic Supply Function Competition With Private Information," *Econometrica*, November 2011, *79* (6), 1919–1966.

Vogel, Sebastian, "When to Introduce Electronic Trading Platforms in Over-the-Counter Markets?," *Working Paper*, November 2019.

Wang, Chaojun, "The Limits of Multi-Dealer Platforms," *Journal of Financial Economics*, September 2023, *149* (3), 434–450.

Westerholm, P. Joakim, "Do Uninformed Crossed and Internalized Trades Tap into Unexpressed Liquidity? The Case of Nokia," *Accounting & Finance*, June 2009, *49* (2), 407–424.

Zhu, Haoxiang, "Do Dark Pools Harm Price Discovery?," *Review of Financial Studies*, 2014, *27* (3), 747–789.

Internet Appendices

I Request for Quote

We microfound the zero-profit condition and the trading protocol in the OTC market. We add two features to the model of Section 2: There are $n \geq 2$ dealers and $m \geq 2$ market markers ($n + m$ intermediaries), and traders must request quotes from the dealers to trade over the counter. We show that these additions do not change the equilibrium allocation.

Strategies. In Stage 1, each trader chooses (i) whether to simultaneously request two-way quotes and (ii) from whom among the n dealers (possibly all n). (The equilibrium allocation is the same whether the traders' choices (i) and (ii) are disclosed to the intermediaries or not.) In Stage 2, each market maker posts one bid and one ask available to all traders and, for each quote request received, each dealer offers a bid and an ask conditional on the requester's label. In Stage 3, each trader observes her prices then makes two decisions: (a) buy, sell, or exit and (b) against which price to trade. We impose the same tie-breaking rule Assumption 1 as in Section 2.1. Figure I.1 summarizes the timing and Figure I.2 the traders' choices.

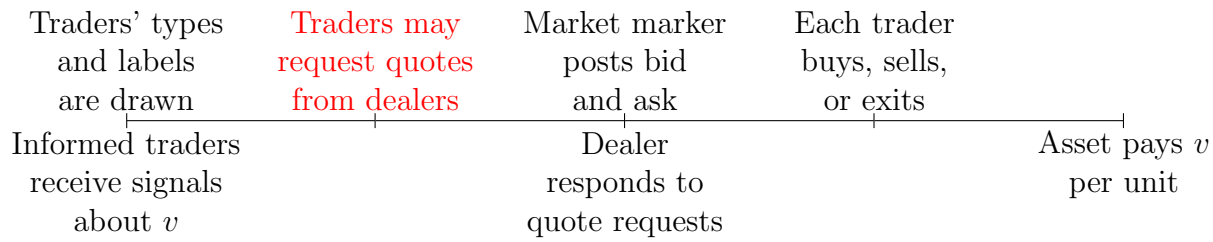

Figure I.1: Timeline

Equilibrium. We study the properties of an extensive-form trembling-hand perfect (THP) equilibrium (Selten, 1975). The next proposition shows that any extensive-form THP equilibrium must induce an identical allocation as the equilibrium of the original model.

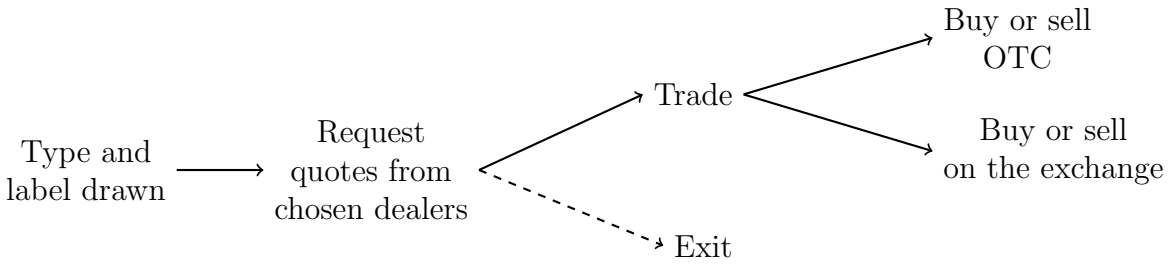

Figure I.2: Traders' Choices

Proposition I.1. *Given any extensive-form trembling-hand perfect equilibrium,*[20] *with probability* 1*:*

- *In Stage 1, each trader sends a quote request to all n dealers.*
- *In Stage 2, every market maker posts the exchange spread S_E, and every dealer offers the OTC spread S_O to a Likely Uninformed trader and S_E to a Likely Informed trader. (The spreads S_O and S_E are given in Proposition 0.)*
- *In Stage 3, every LI trader chooses the exchange and receives S_E, and every LU trader chooses the OTC market and receives S_O.*

Proof. We first show that in Stage 1, each trader sends an RFQ to all the n dealers[21] with probability 1 under a THP equilibrium E. We fix a converging sequence of completely mixed strategy profiles E_n that gives rise to the equilibrium E. Each strategy profile E_n is a Nash equilibrium of a perturbed game. In Stage 1, each trader optimally submits an RFQ to every dealer with the maximum probability that is allowed in the perturbed game for two reasons. First, sending an RFQ is cost-free. Second, every dealer and market maker follows a completely mixed pricing strategy under the Nash equilibrium E_n, so any given dealer

[20] There may be multiple THP equilibria. The proof of Proposition I.1 establishes that all such equilibria must have the same allocation. For example, the THP equilibria may differ in how a dealer responds at an information set where the dealer does not receive quote requests from all traders. Such an information set is reached with probability 0 under any THP equilibrium.

[21] When it is costly for a dealer to respond to an RFQ, Wang (2023) shows that a trader submits an RFQ to only two dealers, and prices become noncompetitive. Adding such an OTC friction to our model would strengthen the benefit of penalizing OTC trades. We abstract away from any OTC friction to focus on the effect of cheap substitution.

offers the overall best price with a strictly positive probability. Hence, each trader submits an RFQ to all the dealers with probability 1 under the limiting equilibrium E.

Since every trader submits an RFQ to all the dealers with probability 1, the extended trading game reduces to the original game with price posting instead of RFQ. Therefore, a THP equilbrium of the extended trading game induces an identical equilibrium allocation as the equilibrium of Section 2.3. □

II Informed Traders with Hedging Benefits

We modify the model of Section 3.3 to endow the informed traders with hedging benefits. Precisely, each trader's payoff V_i from owning an asset (or $-V_i$ from shorting an asset) is the sum of a private value $\pm b_i$ (a hedging need) and the common value v of the asset, $V_i = \pm b_i + v$. The trader's private value is equally likely to be b_i and $-b_i$. The hedging benefits $b_i \overset{iid}{\sim} F$, whose pdf f is analytic in some neighborhoods of 0 and $(2\alpha - 1)$ with support $[0, 1]$. We normalize the total mass of traders to 1. A fraction μ of the traders are informed in that they each receive a private binary signal about v, as do the informed traders of the baseline model. This setup guarantees that the total attainable hedging benefit is invariant to μ. A trader's private value $\pm b_i$ and whether she is informed are her private information. The remaining setup is identical to the model in Section 3.3. In this extension, the informed traders have "dual intent" of earning speculative profit and attaining their hedging benefits.

Proposition II.1. *Given any α, we let f be a pdf analytic in some neighborhoods of 0 and $(2\alpha - 1)$ with support $[0, 1]$, and each uninformed trader's hedging benefit $b_i \overset{iid}{\sim} f$. Given any pairs $(\theta_l, \gamma_l) < (\theta_h, \gamma_h) < (1, 1)$ that satisfy $\theta_h + \gamma_l > 1$ and $\theta_l + \gamma_h > 1$, there exists a cutoff $\mu^d > 0$ such that:*

(a) (1) Marginally lowering traders' label accuracy θ from θ_h to $\theta_h - d\theta$ or γ from γ_h to $\gamma_h - d\gamma$, (2) lowering the label accuracy θ from θ_h to θ_l or γ from γ_h to γ_l, or (3) closing the OTC market strictly raises welfare W across all $\mu < \mu^d$.

(b) For $f = \mathbb{U}[0, 1]$, (1), (2), or (3) strictly reduces the aggregate trade volume V and widens the average bid-ask spread $\bar{S}$.

Penalizing or closing the OTC market raises welfare if the mass of informed traders μ is sufficiently small. Welfare may rise or decline where μ is large.

The intuition is that cheap substitution remains dominant with small μ while, with large μ, other effects with ambiguous signs may dominate. In the base model, cheap substitution arises among (A) the uninformed traders, because the uninformed LI traders enter and the uninformed LU traders exit. Adding dual intent, two further instances of cheap substitution arise. Among (B) the informed traders with large hedging benefits, who only trade in the direction of their benefits, cheap substitution takes the same form as for (A). Among (C) the informed traders with small benefits, who only trade against the direction of their benefits to profit on their private signals, cheap substitution is "double flipped." (First flip) Trading by (C) destroys welfare as each such trade is in the *opposite* direction of its underlying hedging benefit. (Second flip) Each exiter among (C) has a *larger* hedging benefit than each entrant, since a trader in (C) trades only if her hedging benefit is *small*. As a result of the double flip, cheap substitution in (C) raises welfare because each exiter would otherwise inflict a larger welfare loss than each entrant. All three cheap substitutions raise welfare and dominate if μ is low. When μ is high, all cheap substitutions vanish, which leaves the volume effect (present in the base model) and several other effects (absent in the base model), whose net effect depends on distributional assumptions.

Proof of Proposition II.1. If an intermediary charges a spread s, the intermediary earns zero profit in expectation if and only if

$$s \cdot \underbrace{[1 - F(s)]}_{(A)} = \Bigg[\quad (2\alpha - 1 - s) \left(\frac{1}{2} + \underbrace{\frac{F(2\alpha - 1 - s)}{2}}_{(C)} \right) - (2\alpha - 1 + s) \underbrace{\frac{1 - F(2\alpha - 1 + s)}{2}}_{(B)} \Bigg] \cdot \beta. \tag{II.1}$$

The left-hand side is the intermediary's expected profit when facing an uninformed trader, and is the same as the LHS of (1). On the right-hand side, an intermediary who faces an informed trader j expects either (i) a loss of $2\alpha - 1 - s$ if j trades in the direction of her

signal, or (ii) a gain of $2\alpha - 1 + s$ if j trades against her signal. Thus, any zero-profit spread must be less than $2\alpha - 1$. Given a spread $s < 2\alpha - 1$, event (i) occurs if and only if the private value $\pm b_j$ of informed trader j and her signal are in the same direction, or (C) $\pm b_j$ is in the opposite direction of her signal and $b_j < 2\alpha - 1 - s$. Event (i) occurs with probability $[1 + F(2\alpha - 1 - s)]/2$. Event (ii) occurs if and only if (B) $\pm b_j$ is in the opposite direction of her signal and $b_j > 2\alpha - 1 + s$. Event (ii) occurs with probability $[1 - F(2\alpha - 1 + s)]/2$. Last, β is the mass of informed traders per unit mass of uninformed traders.

The Intermediate Value Theorem implies that the zero-profit condition (1) admits at least one solution in $(0, 2\alpha - 1)$. Its smallest solution is the equilibrium spread, denoted as $\tilde{S}(\beta)$.

Part (a): Welfare is equal to

$$
\begin{aligned}
(1-\mu)\Bigg(\quad (1-\gamma)\Bigg[\quad & \int_{\tilde{S}(\beta_{\mathrm{E}})}^{1} s f(s)\,\mathrm{d}s + \\
& \frac{\beta_{\mathrm{E}}}{2}\left(\int_{2\alpha-1-\tilde{S}(\beta_{\mathrm{E}})}^{1} s f(s)\,\mathrm{d}s + \int_{2\alpha+1+\tilde{S}(\beta_{\mathrm{E}})}^{1} s f(s)\,\mathrm{d}s\right)\Bigg] \\
+\gamma\Bigg[\quad & \int_{\tilde{S}(\beta_{\mathrm{O}})}^{1} s f(s)\,\mathrm{d}s + \\
& \frac{\beta_{\mathrm{O}}}{2}\left(\int_{2\alpha-1-\tilde{S}(\beta_{\mathrm{O}})}^{1} s f(s)\,\mathrm{d}s + \int_{2\alpha+1+\tilde{S}(\beta_{\mathrm{O}})}^{1} s f(s)\,\mathrm{d}s\right)\Bigg]\Bigg),
\end{aligned}
$$

where

$$
\beta_{\mathrm{E}} = \frac{\theta\mu}{(1-\gamma)(1-\mu)}, \qquad \beta_{\mathrm{O}} = \frac{(1-\theta)\mu}{\gamma(1-\mu)}.
$$

Expanding the terms, for $M \in \{\text{O,E}\}$, we have

$$\begin{aligned}
&\frac{\beta_M}{2}\int_{2\alpha-1-\tilde{S}(\beta_M)}^{1} sf(s)\,\mathrm{d}s \\
&= \frac{1}{2}\int_0^{\beta_M}\Bigg[\int_{2\alpha-1-\tilde{S}(\beta)}^{1} sf(s)\,\mathrm{d}s + \\
&\qquad\qquad \beta\tilde{S}'(\beta)\left(2\alpha-1-\tilde{S}(\beta)\right) f\left(2\alpha-1-\tilde{S}(\beta)\right)\Bigg]\mathrm{d}\beta,
\end{aligned}$$

$$\begin{aligned}
&\frac{\beta_M}{2}\int_{2\alpha-1+\tilde{S}(\beta_M)}^{1} sf(s)\,\mathrm{d}s \\
&= \frac{1}{2}\int_0^{\beta_M}\Bigg[\int_{2\alpha-1+\tilde{S}(\beta)}^{1} sf(s)\,\mathrm{d}s - \\
&\qquad\qquad \beta\tilde{S}'(\beta)\left(2\alpha-1+\tilde{S}(\beta)\right) f\left(2\alpha-1+\tilde{S}(\beta)\right)\Bigg]\mathrm{d}\beta.
\end{aligned}$$

As the label accuracy θ falls from θ_h to θ_l, the change in welfare W is

$$(1-\mu)\left[(1-\gamma)\int_{\frac{\theta_l\mu}{(1-\gamma)(1-\mu)}}^{\frac{\theta_h\mu}{(1-\gamma)(1-\mu)}} \Delta_W^d(\beta)\,\mathrm{d}\beta - \gamma\int_{\frac{(1-\theta_h)\mu}{\gamma(1-\mu)}}^{\frac{(1-\theta_l)\mu}{\gamma(1-\mu)}} \Delta_W^d(\beta)\,\mathrm{d}\beta\right],$$

where

$$\begin{aligned}
\Delta_W^d(\beta) = {}& \tilde{S}(\beta)\tilde{S}'(\beta)f(\tilde{S}(\beta)) \\
&-\frac{1}{2}\Bigg[\int_{2\alpha-1-\tilde{S}(\beta)}^{1} sf(s)\,\mathrm{d}s + \beta\tilde{S}'(\beta)\left(2\alpha-1-\tilde{S}(\beta)\right) f\left(2\alpha-1-\tilde{S}(\beta)\right) + \\
&\qquad \int_{2\alpha-1+\tilde{S}(\beta)}^{1} sf(s)\,\mathrm{d}s - \beta\tilde{S}'(\beta)\left(2\alpha-1+\tilde{S}(\beta)\right) f\left(2\alpha-1+\tilde{S}(\beta)\right)\Bigg].
\end{aligned}$$

As the label accuracy γ falls from γ_h to γ_l, the change in welfare W is

$$(1-\mu)\left[(1-\gamma_h)\int_{\frac{\theta\mu}{(1-\gamma_l)(1-\mu)}}^{\frac{\theta\mu}{(1-\gamma_h)(1-\mu)}}\Delta_W^d(\beta)\,\mathrm{d}\beta-\gamma_h\int_{\frac{(1-\theta)\mu}{\gamma_h(1-\mu)}}^{\frac{(1-\theta)\mu}{\gamma_l(1-\mu)}}\Delta_W^d(\beta)\,\mathrm{d}\beta\right.$$
$$\left.-(\gamma_h-\gamma_l)\int_{\frac{(1-\theta)\mu}{\gamma_l(1-\mu)}}^{\frac{\theta\mu}{(1-\gamma_l)(1-\mu)}}\Delta_W^d(\beta)\,\mathrm{d}\beta\right].$$

Next, we show that Δ_W^d is strictly increasing in a neighborhood of $\beta=0$. Then, Proposition II.1 Part (a) follows from identical steps as in the proof of Proposition 2 (a)–(c).

First, one can show that $\lim_{\beta\downarrow 0}[\tilde{S}(\beta)\tilde{S}'(\beta)f(\tilde{S}(\beta))]'(\beta)$ exists and is strictly positive. Thus, $\tilde{S}(\beta)\tilde{S}'(\beta)f(\tilde{S}(\beta))$ is increasing in a neighborhood of $\beta=0$ with a strictly positive slope. Second,

$$\begin{aligned}
&\left(\int_{2\alpha-1-\tilde{S}(\beta)}^{1}sf(s)\,\mathrm{d}s+\int_{2\alpha-1+\tilde{S}(\beta)}^{1}sf(s)\,\mathrm{d}s\right)'(\beta)\\
=\;&\tilde{S}'(\beta)\big[\;(2\alpha-1-\tilde{S}(\beta))f(2\alpha-1-\tilde{S}(\beta))\\
&\quad-(2\alpha-1+\tilde{S}(\beta))f(2\alpha-1+\tilde{S}(\beta))\big]\\
\xrightarrow{\beta\downarrow 0}\;&0\qquad\qquad\left(\text{since }\tilde{S}(\beta)\xrightarrow{\beta\downarrow 0}0\text{ and }f\text{ is analytic in a neighborhood of }2\alpha-1\right)
\end{aligned}$$

Third,

$$\frac{1}{\beta}\big[\;\beta\tilde{S}'(\beta)(2\alpha-1-\tilde{S}(\beta))f(2\alpha-1-\tilde{S}(\beta))-$$
$$\beta\tilde{S}'(\beta)(2\alpha-1+\tilde{S}(\beta))f(2\alpha-1+\tilde{S}(\beta))\big]\xrightarrow{\beta\downarrow 0}\quad 0$$

Therefore, Δ_W^d is strictly increasing in a neighborhood of $\beta=0$.

Part (b): Under the uniform distribution $F=\mathbb{U}[0,1]$, as θ falls from θ_h to θ_l, the change in aggregate volume V is

$$(1-\mu)\left[(1-\gamma)\int_{\frac{\theta_l\mu}{(1-\gamma)(1-\mu)}}^{\frac{\theta_h\mu}{(1-\gamma)(1-\mu)}}\Delta_V^d(\beta)\,\mathrm{d}\beta-\gamma\int_{\frac{(1-\theta_h)\mu}{\gamma(1-\mu)}}^{\frac{(1-\theta_l)\mu}{\gamma(1-\mu)}}\Delta_V^d(\beta)\,\mathrm{d}\beta\right],$$

where

$$\Delta_V^d(\beta)=\tilde{S}'(\beta)-2\,(1-\alpha)\,.$$

As γ falls from γ_h to γ_l, the change in V is

$$(1-\mu)\left[(1-\gamma_h)\int_{\frac{\theta\mu}{(1-\gamma_l)(1-\mu)}}^{\frac{\theta\mu}{(1-\gamma_h)(1-\mu)}}\Delta_V^d(\beta)\,\mathrm{d}\beta-\gamma_h\int_{\frac{(1-\theta)\mu}{\gamma_h(1-\mu)}}^{\frac{(1-\theta)\mu}{\gamma_l(1-\mu)}}\Delta_V^d(\beta)\,\mathrm{d}\beta\right.$$
$$\left.-(\gamma_h-\gamma_l)\int_{\frac{(1-\theta)\mu}{\gamma_l(1-\mu)}}^{\frac{\theta\mu}{(1-\gamma_l)(1-\mu)}}\Delta_V^d(\beta)\,\mathrm{d}\beta\right].$$

One can show that $\tilde{S}'$ is strictly decreasing in $\beta \geq 0$ and so is Δ_V^d. Then, Proposition II.1 Part (b) follows from identical steps as in the proof of Proposition 2 (e). □

III Relation to Akerlof (1970)

The setup of Akerlof (1970) differs from ours in two ways: (I) the private values are perfectly correlated with the common value, and (II) each uninformed trader does not know her own private value. In the presence of assumption (I), assumption (II) is necessary to preserve information asymmetry. This section adds imperfect labels to Akerlof (1970) under every coherent combination of (I) and (II). First, assumptions (I) and (II) together *reverse* the effect of cheap substitution on welfare. Second, assumption (II) alone turns off cheap substitution and is equivalent to a setting with homogeneous hedging benefits. Under the first two settings, pooling can raise social welfare only if it also raises the aggregate trade volume. Third, removing both assumptions is necessary and sufficient to obtain our main results.

The baseline model. We recall the seminal Akerlof (1970) model. Each of a unit mass of sellers owns one unit of an asset with value $x_i \overset{\text{iid}}{\sim} \mathbb{U}[0,2]$, which is private information to seller i. One buyer values an asset at ax_i, with $a > 1$ (in Akerlof (1970), $a = 3/2$), and has no information about the realization of x_i. The buyer has Y dollars to purchase assets. In a pooled market with price p, the sellers supply $S(p) = p/2$ of assets, and the expected value of those assets is $\mu = p/2$. The demand of the buyer equals $D(p) = Y/p$ if $a\mu > p$ (or equivalently, $a > 2$) and $D(p) = 0$ if $a\mu < p$ (i.e., $a < 2$). If $a < 2$, no trade can exist in the pooled market. If $a > 2$, then $D(p^*) = S(p^*)$ in equilibrium. Therefore, $p^* = \sqrt{2Y}$, and only

the assets with common values below p^* get traded.

Adding imperfect labels. The first step of our analysis extends Akerlof (1970) with imperfect labels to compare pooling and separating outcomes. The buyer receives an imperfect label of x_i, which indicates $x_i \sim G_1$ with probability $\beta > 0$ and $x_i \sim G_2$ with probability $1-\beta$ without revealing the realization of x_i. The distributions satisfy $G_1 \neq G_2$ and $\beta G_1 + (1-\beta)G_2 = \mathbb{U}[0,2]$. One can similarly solve for the market clearing prices p_1^* and p_2^* in each of the two markets. If market m breaks down ($m = 1, 2$), we set $p_m^* = 0$ by convention. We suppose that $p_1^* < p_2^*$ without loss of generality. Then, one must have $p_1^* < p^* < p_2^*$.[22] In market m, only the assets with common values below p_m^* get traded. Once pooled together at the price p^*, those market-1 assets with values between $[p_1^*, p^*]$ "enter," while those market-2 assets with values between $[p^*, p_2^*]$ "exit." Each entering asset attains a private value between $[(a-1)p_1^*, (a-1)p^*]$, and each exiting asset destroys a private value between $[(a-1)p_1^*, (a-1)p_2^*]$. Thus, every entrant has a lower private value than every exiter, which is precisely the opposite of cheap substitution. Therefore, pooling the two markets can raise social welfare only if it also raises the aggregate trade volume.

The reversal of cheap substitution is caused by assumption (I), the perfect correlation between private and common values. Intuitively, because only the assets with relatively lower common values get traded, only lower private values get realized under the perfect correlation.

Removing the correlation between private and common values. Second, we break assumption (I), the perfect correlation between common and private values, while maintaining assumption (II) that the buyer does not know her own private value. To do so, we make one modification to the first extension: The buyer now values each asset at $x_i + b_i$, where $b_i \overset{\text{iid}}{\sim} F$. The private values b_i and the common values x_i are independently distributed. The buyer remains uninformed of the private values b_i. This extension is equivalent to the one with a homogeneous private value in that the private values can be replaced by their

[22] If $p^* \leq p_1^* < p_2^*$, then

$$D(p^*) = D_1(p^*) + D_2(p^*) > D_1(p_1^*) + D_2(p_2^*) = S_1(p_1^*) + S_2(p_2^*) > S_1(p^*) + S_2(p^*) = S(p^*),$$

which contradicts $D(p^*) = S(p^*)$. Likewise, one can rule out $p_1^* < p_2^* \leq p^*$.

expectation $\mathbb{E}b_i$, because not knowing the private values makes $\mathbb{E}b_i$ a sufficient statistic for the buyer.

We denote the pooling equilibrium price as $\tilde{p}$. The separating equilibrium prices are $\tilde{p}_1$ and $\tilde{p}_2$, and we suppose $\tilde{p}_1 < \tilde{p}_2$ without loss of generality. By the same arguments as before, $\tilde{p}_1 < \tilde{p} < \tilde{p}_2$. Upon pooling, the market-1 assets with common values in $[\tilde{p}_1, \tilde{p}]$ "enter," while those market-2 assets valued between $[\tilde{p}, \tilde{p}_2]$ "exit." The aggregate volume changes by $\Delta V = \beta[G_1(\tilde{p}) - G_1(\tilde{p}_1)] - (1-\beta)[G_2(\tilde{p}_2) - G_2(\tilde{p})]$. Each entering asset attains the expected private value $\mathbb{E}b_i$, and each exiting asset destroys $\mathbb{E}b_i$. Thus, every entrant has the same expected private value as any exiter. The net change in welfare is $\mathbb{E}b_i \Delta V$. As such, pooling raises welfare if and only if the aggregate volume also rises.

Revealing private value to the buyer. Third, we reveal the private value b_i to the buyer. We denote the market-clearing prices as $\hat{p}$, $\hat{p}_1$, and $\hat{p}_2$. In the pooled market, an asset gets traded if and only if $\hat{p}/2 + b_i > \hat{p}$, or $b_i > \hat{p}/2$. That is, only private values *above* $\underline{b} \coloneqq \hat{p}/2$ get traded. Likewise, when separated, only the assets with private values *above* $\underline{b}_m \coloneqq \hat{p}_m - \mathbb{E}_{G_m}(x \,|\, x < \hat{p}_m)$ get traded in market m $(m = 1, 2)$. We suppose $\underline{b}_1 < \underline{b}_2$ without loss of generality. Then, one must have $\underline{b}_1 < \underline{b} < \underline{b}_2$.[23] Upon pooling, the market-1

[23] One can show that $\underline{b}$ is some weighted average between $\underline{b}_1$ and $\underline{b}_2$:

$$
\begin{aligned}
\underline{b}_m &= \hat{p}_m - \mathbb{E}_{G_m}(x \,|\, x < \hat{p}_m) \\
&= \mathbb{E}_{G_m}(\hat{p}_m - x \,|\, x < \hat{p}_m) \\
&= \mathbb{E}_{\widetilde{G}_m}(\Delta_m \,|\, \Delta_m > 0) \text{ where } \Delta_m \coloneqq \hat{p}_m - x \text{ and } \widetilde{G}_m \text{ is the distribution of } \Delta_m \\
&= \frac{\mathbb{E}_{\widetilde{G}_m}(\Delta_m^+)}{\mathbb{P}_{\widetilde{G}_m}(\Delta_m > 0)} \\
&= \frac{\mathbb{E}_{\widehat{G}_m}(\Delta_m^+)}{\mathbb{P}_{\widetilde{G}_m}(\Delta_m > 0)} \text{ where } \widehat{G}_m \text{ is the distribution of } \Delta_m^+.
\end{aligned}
$$

We define G, Δ, $\widetilde{G}$, and $\widehat{G}$ in a similar manner in the pooled market: $G \coloneqq \mathbb{U}[0,2]$, $\Delta \coloneqq \hat{p} - x$, $\widetilde{G}$ is the distribution of Δ, and $\widehat{G}$ is the distribution of Δ^+. Since $\beta G_1 + (1-\beta) G_2 = G$, then

$$
\beta \widetilde{G}_1 + (1-\beta)\widetilde{G}_2 = \widetilde{G}\,, \qquad \beta \widehat{G}_1 + (1-\beta)\widehat{G}_2 = \widehat{G}.
$$

Therefore,

$$
\begin{aligned}
&\beta \mathbb{E}_{\widehat{G}_1}\left(\Delta_1^+\right) + (1-\beta)\mathbb{E}_{\widehat{G}_2}\left(\Delta_2^+\right) = \mathbb{E}_{\widehat{G}}\left(\Delta^+\right), \\
&\beta \mathbb{P}_{\widetilde{G}_1}(\Delta_1 > 0) + (1-\beta)\mathbb{P}_{\widetilde{G}_2}(\Delta_2 > 0) = \mathbb{P}_{\widetilde{G}}(\Delta > 0).
\end{aligned}
$$

assets with private values between $[\underline{b}_1, \underline{b}]$ "exit," and the market-2 assets with private values between $[\underline{b}, \underline{b}_2]$ "enter." Therefore, every entrant has a higher private value than any exiter, which is precisely cheap substitution. This extension generates our result that pooling can raise welfare while reducing the aggregate trade volume.

IV Evidence

OTC discount

A prediction of our model is that the bid-ask spread over the counter is lower than the spread on the exchange. Recent evidence in de Roure et al. (2021) is consistent with this prediction. They examine data on German Bunds, which are traded both in the OTC market and on exchanges. Their headline result is that 87% of OTC trades execute at a better price (for the initiator of the trade) than the contemporaneous price on the exchange. A typical OTC discount is substantial: The median OTC trade saves 65.7% of the half spread it would have cost on the exchange. Moreover, the OTC discount is not an artifact of larger trade sizes. The median trade size over the counter (EUR 5.4 million) is nearly identical to that on the exchange (EUR 5.0M). On average, the OTC trades are larger (EUR 12.74M vs 7.88M), solely because the biggest 5 percentile of trades are far larger over the counter (EUR 50M vs 13M). Likewise, Collin-Dufresne et al. (2020) finds that both dealer-to-client (D2C) and dealer-to-dealer (D2D) trades of index CDSs incur smaller trading costs than the contemporaneous bid-ask spread on the largest limit order book.[24]

Analogous evidence is found in the literature on the upstairs market for equities. (Upstairs trades are the OTC trades of stocks.) Rose (2014) compares upstairs and limit order book (LOB) trades on the Australian Stock Exchange. While OTC trades make a loss on average, LOB trades earn a profit. Further, traders who consistently make losses are more likely to trade upstairs and receive a discount relative to the price on the LOB. The opposite holds for consistently profit-earning traders. Others find that upstairs trades are less

[24] Collin-Dufresne et al. (2020) focuses on one limit order book, GFI Swaps Exchange ("D2D SEF" in their parlance). GFI Swaps Exchange provides non-discriminatory access to both dealers and buyside firms (GFI LLC, 2022, Sections 301, 302).

informed and less costly (Madhavan and Cheng, 1997; Smith, Turnbull and White, 2001; Booth, Lin, Martikainen and Tse, 2002; Bessembinder and Venkataraman, 2004; Bernhardt, Dvoracek, Hughson and Werner, 2005; Westerholm, 2009). Large trades do not drive the upstairs discount. In fact, the discount is minimized among large orders (Bernhardt et al., 2005) and can be monotonically shrinking in trade size (Westerholm, 2009). The findings echo our mechanism that dealers offer a discount to the less informed traders and cream skim them into the OTC market.

Migration of corporate bonds to the OTC market

Cream skimming in our model pulls a larger share of trades into the OTC market when a smaller share of traders are informed. The Depression-era migration of corporate bond trades is an episode consistent with this prediction. Most corporate bond trades moved from exchanges to the OTC market in the 1930s and 40s (Biais and Green, 2007). This movement coincided with a sudden, dramatic rise in the proportion of corporate bonds held by insurance companies and pension funds. Homer (1975) delivers the author's first-hand account of this transition as a brokerage president in the 1930s: Pre-1930s, the corporate bond traders primarily traded on NYSE and were "small, country investors or big-city investors" (p. 379). Particularly active were sophisticated traders who "grabbed up" new issues, which "if well priced they sold at quick premiums" (p. 379). During the Great Depression, distressed investors sold out to dealers, who in turn resold the bonds to life insurance companies (p. 381). The life insurers preferred to trade over the counter although "[t]he exchange tried hard to retain its bond business" (p. 381). Because insurers typically trade with a hedging motive, the history of corporate bonds mirrors our prediction that OTC trading dominates where most traders are uninformed.

V Empirics

We use US equities data to document (1) substantial OTC market share in US-listed equities and (2) a positive correlation between the exchange market share and the exchange quoted spread, which Proposition V.1 below predicts.

Proposition V.1. *For $f = \mathbb{U}[0,1]$ and given any triple (α, θ, γ), the exchange spread S_E and the exchange market share V_E/V are each strictly increasing in the mass of informed traders μ.*

V.1 Data, Variables, and Summary Statistics

We combine millisecond Trade and Quote (TAQ) data with weekly OTC trade volumes from the Financial Industry Regulatory Agency (FINRA).[25] In the FINRA data, OTC volumes are separated into Alternative Trading Systems (ATS) versus Non-ATS OTC volumes. The ATS consists of dark pools, batch auctions, and limit order books that are not designated as "national securities exchanges" by the US Securities and Exchange Commission. Because trading on the ATS is anonymous, the ATS corresponds to the exchange in our model. The Non-ATS OTC refers to traditional bilateral and request-for-quote trades. Therefore, only the Non-ATS trades are counted as over the counter in our analysis.

The sample period is January 2, 2017–March 5, 2021, the available range of FINRA OTC data at the time of this analysis. We exclude all trades outside of market hours. Our sample consists of 3,210 US-listed non-ETF tickers that exist in both TAQ and FINRA data on both the first and the last weeks of the 218 weeks in the sample period. Only the trades during market hours are included to avoid an upward bias for OTC market share. (Results remain nearly identical if all trades are included).

Exchange market share for ticker i and week w is 1 minus the ratio of week w OTC dollar volume of trades from FINRA to the week w aggregate dollar volume from TAQ. To compute *percent quoted spread* for ticker i and week w, we use the millisecond TAQ quotes to calculate the time-weighted best quoted spread in percentage, *(best ask - best bid) /midpoint*, for each day, and then we take the simple average across the number of days observed for ticker i in week w. The total *number of trades* and total *dollar volume* of trades are computed for each ticker i and week w from millisecond TAQ trade data. Table V.1 provides the summary statistics, first for all observations, and then by average weekly dollar volume of trades in quintiles.

[25] FINRA data used here is available publicly at http://www.finra.org/industry/otc-transparency.

Table V.1: Summary Statistics

Each observation is one ticker for one week.

	All Observations							
	Obs	Mean	SD	Min	25%	50%	75%	Max
Exchange Market Share	694,305	0.750	0.198	0.000	0.597	0.828	0.914	1.000
Percent Quoted Spread	694,305	0.027	0.053	0.000	0.004	0.010	0.026	6.995
Dollar Volume of Trades	694,305	384.4M	2,752.6M	0.0M	3.3M	27.5M	187.2M	376,000.0M
Number of Trades	694,305	38,401	107,820	2	1,714	9,434	36,594	13,458,361
	By Average Weekly Dollar Volume of Trades:							
	Quintile 1 Observations							
	Obs	Mean	SD	Min	25%	50%	75%	Max
Exchange Market Share	137,455	0.615	0.185	0.000	0.484	0.609	0.758	1.000
Percent Quoted Spread	137,455	0.067	0.086	0.001	0.016	0.042	0.089	6.995
Dollar Volume of Trades	137,455	1.6M	3.7M	0.0M	0.4M	0.9M	1.9M	543.8M
Number of Trades	137,455	997	3,495	2	189	474	1,027	491,416
	Quintile 2 Observations							
	Obs	Mean	SD	Min	25%	50%	75%	Max
Exchange Market Share	139,184	0.659	0.193	0.001	0.506	0.652	0.839	1.000
Percent Quoted Spread	139,184	0.037	0.055	0.001	0.011	0.022	0.044	4.943
Dollar Volume of Trades	139,184	8.3M	18.8M	0.0M	2.7M	5.4M	9.8M	1,570.1M
Number of Trades	139,184	5,240	14,348	7	1,345	2,815	5,745	950,784
	Quintile 3 Observations							
	Obs	Mean	SD	Min	25%	50%	75%	Max
Exchange Market Share	139,010	0.775	0.194	0.005	0.667	0.859	0.922	1.000
Percent Quoted Spread	139,010	0.018	0.033	0.000	0.007	0.012	0.019	3.687
Dollar Volume of Trades	139,010	37.6M	50.9M	0.0M	16.2M	27.8M	46.3M	4,236.7M
Number of Trades	139,010	14,577	25,721	5	5,120	9,755	17,342	1,986,649
	Quintile 4 Observations							
	Obs	Mean	SD	Min	25%	50%	75%	Max
Exchange Market Share	139,273	0.830	0.164	0.008	0.800	0.896	0.935	1.000
Percent Quoted Spread	139,273	0.008	0.012	0.000	0.003	0.005	0.009	0.487
Dollar Volume of Trades	139,273	153.5M	153.4M	0.0M	72.2M	119.7M	194.3M	14,100.0M
Number of Trades	139,273	32,783	36,820	59	14,249	25,524	41,738	3,485,211
	Quintile 5 Observations							
	Obs	Mean	SD	Min	25%	50%	75%	Max
Exchange Market Share	139,383	0.867	0.108	0.013	0.851	0.902	0.931	1.000
Percent Quoted Spread	139,383	0.003	0.005	0.000	0.001	0.002	0.004	0.226
Dollar Volume of Trades	139,383	1,714.1M	5,957.4M	0.3M	427.1M	726.2M	1,409.3M	376,000.0M
Number of Trades	139,383	137,773	206,702	1,161	51,324	86,649	153,925	13,458,361

Table V.2: Summary Statistics by Exchange Market Share

Each observation is one ticker for one week. We present simple averages across observations with exchange market share less than 50% on the left column, and across observations with exchange share more than 50% on the right.

	Exchange Market Share $< 50\%$	Exchange Market Share $> 50\%$
Observations	80,746	613,559
Exchange Market Share	0.443	0.79
Percent Quoted Spread	0.031	0.026
Dollar Volume of Trades	40,283,547	429,702,482
Number of Trades	6,227	42,635

V.2 Exchange Market Share for US-listed Equities

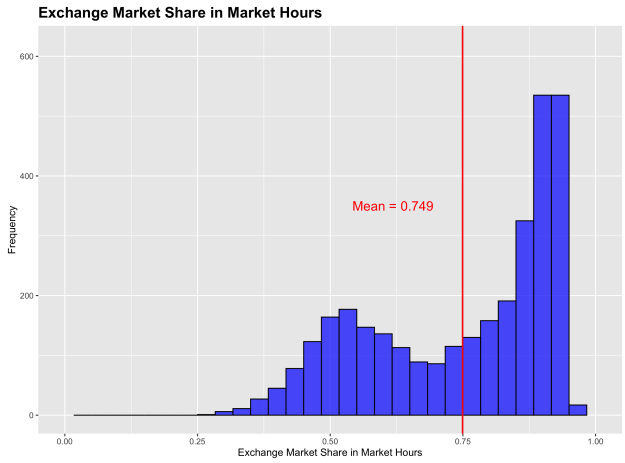

Figure V.1: Exchange Market Share of Dollar Value of Trades in US-listed Equities

Does OTC trading dominate only for historically OTC-traded assets such as bonds? Figure V.1 plots the average weekly exchange market shares of tickers in our sample. It shows that even among US-listed equities, many are OTC-dominated, and much of the remainder exhibit large OTC market shares.

V.3 Correlation between Exchange Market Share and Spread

Proposition V.1 predicts a positive correlation between exchange market share and the exchange quoted spread *if* the observations vary substantively in the mass of informed traders μ but not in signal or label accuracy (α, θ, γ). We estimate the correlation within a narrow asset class, controlling for time (week) fixed effects. Specifically, we partition US exchange-listed equities into quintiles by average weekly dollar volume and estimate the correlation within each quintile. Our assumption is that across tickers within each dollar volume quintile in a given week, (i) similar amounts of information about individual traders (e.g., type of firm, reputation, past disclosures) are available, so the tickers have similar label accuracies (θ, γ), (ii) similar speculators are present, thus similar signal accuracies α, and (iii) there is significant variation in adverse selection risk among the tickers (due to, say, ticker-specific news).

Table V.3 presents the regression estimates for log exchange market share on log quoted spreads.[26] All regressions control for week fixed effects, hence our estimates capture cross-sectional variation. Standard errors are clustered at the ticker level. Under each quintile, the left-most regression has no controls. Consistent with Proposition V.1, the correlation between log quoted spread and log exchange market share is positive within every quintile. COVID or other time variation common across tickers do not drive our results, as we control for week fixed effects.

A typical explanation for OTC trading is that illiquidity renders trading certain assets on exchanges impractical. We evaluate whether our results are genuinely separate from this liquidity mechanism and, if so, the relative magnitudes of the two mechanisms on the exchange market share. To this end, Table V.3 presents regressions whose controls include one of two proxies for liquidity, log weekly total dollar volume ("dollar volume") or log weekly total number of trades. It shows that the quoted spread is positively correlated with the exchange market share after controlling for liquidity—in fact, the coefficient estimates

[26] Bogousslavsky and Collin-Dufresne (2022) reports a positive correlation between aggregate turnover and bid-ask spread. We examine the intensive margin of *where* trades occur and necessarily distinguish between the trades on exchanges versus over the counter. They examine the extensive margin of *how much* trading occurs in the aggregate.

Table V.3: Dependent Variable: log(Exchange Market Share)

All regressions control for week fixed effects. Observations are weekly and include 3,210 non-ETF US-listed tickers that exist in both the first and last weeks of the 218-week sample, from January 2, 2017 to March 5, 2021. Trades outside of market hours are excluded. Standard errors are clustered at the ticker level. Corresponding t-statistics are shown in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels.

Independent Variables	Quintile 1			Quintile 2			Quintile 3		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
log(percent quoted spread)	0.074***	0.076***	0.072***	0.046***	0.144***	0.069***	0.061***	0.143***	0.077***
	(11.61)	(11.05)	(11.66)	(4.32)	(15.57)	(7.52)	(3.62)	(10.10)	(5.71)
log(dollar volume)		0.005			0.129***			0.125***	
		(0.89)			(24.39)			(18.90)	
log(number of trades)			0.029***			0.123***			0.150***
			(5.74)			(19.27)			(14.49)
R^2	0.059	0.059	0.068	0.027	0.168	0.182	0.029	0.140	0.208
N		137,453			139,184			139,010	

Independent Variables	Quintile 4			Quintile 5		
	(10)	(11)	(12)	(13)	(14)	(15)
log(percent quoted spread)	0.087***	0.131***	0.095***	0.028***	0.038***	0.044***
	(5.74)	(9.65)	(8.16)	(2.61)	(2.98)	(3.20)
log(dollar volume)		0.095***			0.014	
		(10.30)			(1.57)	
log(number of trades)			0.146***			0.034***
			(13.08)			(2.83)
R^2	0.074	0.146	0.273	0.051	0.056	0.078
N		139,273			139,383	

seemingly become larger. Therefore, our findings are not explained by more liquid tickers being easier to trade on exchanges. Further, the coefficient estimates for the quoted spread and for the liquidity proxies are in the same magnitude, which suggests that our mechanism is not second-order to the illiquidity explanation.