

Regulating Over-the-Counter Markets

TOMY LEE and CHAOJUN WANG*

ABSTRACT

Over-the-counter (OTC) trading thrives despite competition from exchanges. We let OTC dealers cream skim from exchanges in an otherwise standard Glosten and Milgrom framework. Restricting the dealer's ability to cream skim induces "cheap substitution": some traders exit while others with larger gains from trade enter. Cheap substitution implies trading costs, trade volumes, and market shares are poor policy indicators. In a benchmark case, restricting the dealer raises welfare only if trading cost increases, volume falls, and OTC market share is high. By contrast, the restriction improves welfare when adverse selection risk is low. A simple procedure implements the optimal Pigouvian tax.

OVER-THE-COUNTER (OTC) MARKETS ARE OPAQUE and difficult to access,¹ yet they host the vast majority of financial trades, as most assets are seldom traded on centralized exchanges. Dealers argue against regulatory intervention, citing high OTC market shares and evidence of lower trading costs compared to exchanges.² Do the high market shares and low trading costs

*Tomy Lee is at the Central European University. Chaojun Wang is at the Wharton School, University of Pennsylvania. This paper was previously titled "Why Trade Over the Counter?" We thank Philip Bond (Editor), Associate Editor, two referees, Jordi Mondria, Xianwen Shi, Darrell Duffie, Ettore Damiano, Colin Stewart, Adam Szeidl, Adam Zawadowski, Yu An, Daniel Carvalho, Hongseok Choi, Peter Cziraki, Piotr Dworzak, Jim Goldman, Vincent Glode, Botond Kőszegi, Sam Langfield, Katya Malinova, Vincent Maurin, Andreas Park, Cecilia Parlatore, Christine Parlour, Semih Üslü, Adrian Walton, Liyan Yang, and Yao Zeng for invaluable advice and guidance. Amy Dai, William Fan, Clarise Huang, Dylan Marchlinski, Sifan Tao, and Ceyda Ustun provided excellent research assistance. Jiaqi Zou provided excellent copy editing. We have read *The Journal of Finance* disclosure policy and have no conflict of interest to disclose.

Correspondence: Tomy Lee, Central European University, Department of Economics, B514-Quellenstrasse 51, Wien, Austria; e-mail: leeso@ceu.edu

This is an open access article under the terms of the [Creative Commons Attribution](#) License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

¹ In our context, "over-the-counter markets" consist of all financial markets in which trades are executed nonanonymously 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 RFQ platforms.

² A typical example is the comment of the Securities Industry and Financial Markets Association (2018) on a U.S. regulatory proposal to limit dealers' access to their counterparties' identities. DOI: 10.1111/jofi.13461

© 2025 The Author(s). *The Journal of Finance* published by Wiley Periodicals LLC on behalf of American Finance Association.

indicate that having OTC markets improve welfare? Should policymakers restrict OTC trading, and under which conditions?

In this paper, we show that trading cost, trade volume, and market share can starkly mislead policy. Restricting OTC trading can widen the average bid-ask spread and reduce aggregate trade volume while strictly raising welfare. Moreover, under certain conditions, such restrictions improve welfare *exactly* where the OTC market share is *high*. In contrast to these measures, adverse selection risk provides robust guidance. We show that restrictions on OTC trading strictly raise traders' utilitarian welfare whenever adverse selection risk is low, even if the OTC market is frictionless. Beyond trade restrictions, we devise an optimal Pigouvian tax that can be implemented using a simple sufficient statistic.

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). OTC dealers can price discriminate among their customers, whereas exchanges cannot. As a result, the dealers can offer a discount to those 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: endogenous participation of traders who differ in their gains from trade.

Section I 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 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 to buy or sell on an exchange, with a dealer over the counter, or to exit. The venues differ solely in that only the dealer may condition his prices on each trader's label. In equilibrium, the (informed and uninformed) LI traders endogenously choose the exchange, the LU traders choose the OTC market, and the exchange spread is wider than the LU traders' OTC spread. An informed (LI or LU) trader always trades and an uninformed trader trades if her hedging benefit exceeds her best half bid-ask spread.

Section II examines how restricting the dealer's ability to price discriminate affects utilitarian welfare, aggregate trade volume, and the average bid-ask spread. Specifically, we reduce the accuracy of the traders' labels, which further pools together the informed and the uninformed traders in the eyes of the OTC dealer. Restricting the dealer this way until the LI and the LU traders are equally likely to be informed is equivalent to closing the OTC market. We show that restricting the dealer strictly raises welfare if the ratio β of the mass of informed traders to the mass of uninformed traders is low and reduces welfare if

ties (called posttrade name give-up). SIFMA suggests that regulatory intervention is unwarranted because (i) the OTC trades have smaller bid-ask spreads (citing Riggs et al. (2020) and Collin-Dufresne, Junge, and Trolle (2020)) and (ii) the swap traders overwhelmingly trade over the counter despite being "free to choose" an exchange. de Roure et al. (2021) provide further evidence for (i).

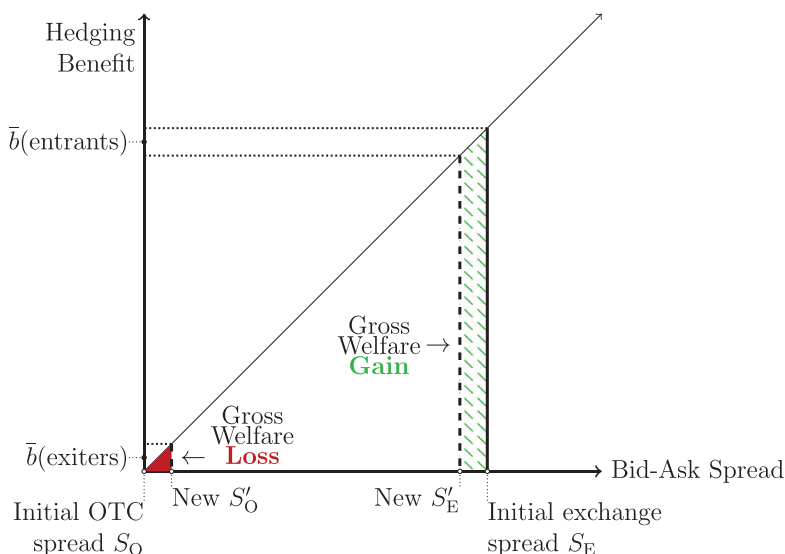


Figure 1. Restricting the OTC dealer can raise welfare. Initially, all informed traders are correctly labeled as Likely Informed. We then mislabel a small mass of informed traders as Likely Uninformed. The average hedging benefit of the resulting entrants is $\bar{b}(\text{entrants})$ and of the exiters is $\bar{b}(\text{exitors})$. (Color figure can be viewed at wileyonlinelibrary.com)

β is high. Yet the restriction can *always* decrease aggregate volume and widen the average spread. In sum, (i) the effect of a policy on aggregate volume and the average spread is a poor indicator of its effect on welfare, and (ii) restricting the OTC dealer improves welfare whenever adverse selection risk is low.

Result (i) is driven by “cheap substitution.” The exchange spread always exceeds the OTC spread, and thus the marginal uninformed traders on the exchange have larger hedging benefits than the marginal uninformed traders in the OTC market. Restricting the OTC dealer pulls the two spreads toward each other, narrowing the exchange spread and widening the OTC spread. As a result, the marginal uninformed traders with the larger hedging benefits enter the exchange while those with the smaller hedging benefits exit the OTC market. In welfare terms, the entrants substitute for the comparably “cheap” exiters. Due to this cheap substitution, welfare can increase even if the exiters outnumber the entrants.

Figure 1 depicts a striking example of cheap substitution. We let $m(\text{entrants})$ and $\bar{b}(\text{entrants})$ denote the mass and the average hedging benefit of entrants upon restricting the OTC dealer, and we define $m(\text{exitors})$ and $\bar{b}(\text{exitors})$ analogously. The restriction generates a gross welfare gain worth $\bar{b}(\text{entrants}) \times m(\text{entrants})$, and destroys $\bar{b}(\text{exitors}) \times m(\text{exitors})$. On net,

$$\text{welfare increases if and only if } \underbrace{\frac{\bar{b}(\text{entrants})}{\bar{b}(\text{exitors})}}_{\text{Cheap Substitution}} > \underbrace{\frac{m(\text{exitors})}{m(\text{entrants})}}_{\text{Volume Effect}}. \quad (1)$$

The ratio $\frac{\bar{b}(\text{entrants})}{\bar{b}(\text{exitors})}$ measures the beneficial impact of cheap substitution on welfare. In Figure 1, all informed traders are initially correctly labeled as LI, and hence trade on the exchange. We then restrict the dealer's ability to price discriminate by mislabeling a small mass of informed traders as LU. The resulting entrants have an average hedging benefit between the initial exchange spread S_E and the new wider spread S'_E , $\bar{b}(\text{entrants}) \in (S'_E, S_E) > 0$, whereas the average exiter hardly loses, $\bar{b}(\text{exitors}) \in (S_O, S'_O) \approx 0$. Thus, cheap substitution $\frac{\bar{b}(\text{entrants})}{\bar{b}(\text{exitors})} \approx \frac{S_E}{S_O}$ approaches infinity, and welfare strictly increases no matter how much aggregate volume declines.³ Going beyond this example, cheap substitution $\frac{\bar{b}(\text{entrants})}{\bar{b}(\text{exitors})}$ is generally finite, and therefore the restriction can raise or lower welfare.

Result (ii), which holds that restricting the OTC dealer raises welfare when adverse selection risk is low, follows because cheap substitution persists as the risk becomes small, while the effect on volume vanishes. Our results are remarkably robust. They hold under any commonly used distribution, for any change in the accuracy of traders' labels, and regardless of whether the informed traders have hedging benefits themselves (see Section III of the [Internet Appendix](#)).⁴

Dealers frequently invoke the high market share of OTC markets to oppose regulatory intervention. Exchanges are indeed available for many OTC-dominated assets.⁵ Perhaps, then, the high OTC market share is a competitive equilibrium outcome, which may well be socially optimal. Section II.C 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 LU traders, who choose the OTC market. Therefore, β and the OTC market share are negatively related under many distributions. Proposition 4 combines this negative relationship with our main result: when β drives the variation in the OTC market share, restricting or closing the OTC market *strictly improves* welfare if the OTC market share is sufficiently *high*. Thus, high OTC market share is not evidence against intervention; it is not even evidence for keeping the OTC market.

Venue choice generates the key externality underlying our results. The traders who opt to trade over the counter do not internalize that their choice worsens the spread on the exchange. Section III devises a Pigouvian tax that optimally penalizes this externality. A simple policy experiment suffices to

³ Proposition 1 states that the volume effect $m(\text{exitors})/m(\text{entrants})$ is finite for any distribution of hedging benefits whose probability density function (pdf) f is continuous and strictly positive in $[0,1]$.

⁴ The [Internet Appendix](#) may be found in the online version of this article.

⁵ 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 of 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.

implement the tax. First, one imposes a small tax T on OTC trades. Second, one computes the weighted spread ratio $WSR := |S_E/S_O \div (dV_O/dV_E)|$. The ratio of spreads S_E/S_O captures the cheap substitution effect, and the ratio of changes in trade volumes $|dV_O/dV_E|$ captures the volume effect. If the weighted spread ratio (WSR) exceeds one, cheap substitution dominates and raising the tax T would strictly improve welfare. If the WSR is below one, 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 1). Section V of the [Internet Appendix](#) 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. Section VI of the [Internet Appendix](#) documents a novel empirical pattern: the total market share of exchanges and their quoted spreads are positively correlated across U.S. equities.

Whether price discrimination in OTC markets is socially beneficial is an increasingly relevant question. Electronic OTC trading platforms (e.g., MarketAxess, Bloomberg) have the capability to fine tune the information revealed to dealers. Moreover, several proposals to implement blockchain technology for record-keeping of OTC trades are under consideration. Most such proposals would disseminate the traders' identities to selected dealers. Section I of the [Internet Appendix](#) discusses these proposals through the lens of our model.

Section IV compares our mechanism to those in the literature. We make three contributions. First, we introduce cheap substitution, a new mechanism that can overturn the effects of aggregate volume and average spread on welfare. Second, we contribute the novel result that pooling improves welfare whenever adverse selection risk is low. And third, we devise a simple sufficient statistic with which one can implement the optimal Pigouvian tax.

The remainder of the paper proceeds as follows. Section I describes our model. Section II states and explains our main results. Section III devises the optimal Pigouvian tax on OTC trades. Section IV compares our mechanism and results to existing work. Finally, Section V concludes with a discussion of frictions that are not captured by our model.

I. A Model of Venue Choice and Price Discrimination

Section I.A sets up a model in which each trader may trade over the counter, on an exchange, or exit. The OTC dealer can price discriminate across the traders' public labels, whereas the exchange dealer cannot. Section I.B interprets our assumptions. Section I.C derives the unique equilibrium.

A. 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 one unit in one of two markets. In the OTC market, a dealer acts as the counterparty to the traders and absorbs net demand. Another dealer does so on the exchange. The OTC dealer and the exchange dealer set prices such that their expected profit is zero, as we detail below.

A mass μ of the traders are informed and a mass one 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 independently drawn from distribution F , $b_i \stackrel{\text{iid}}{\sim} F$, with support $[0,1]$. An uninformed trader's realized hedging benefit b_i and 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. We assume that $\theta < 1$ or $\gamma < 1$, and that a trader's odds \mathcal{O}_{LI} of being informed conditional on being labeled LI is strictly higher than the unconditional odds μ and the odds \mathcal{O}_{LU} conditional on being LU . More precisely, $\mathcal{O}_{\text{LI}} = \frac{\theta\mu}{1-\gamma} > \mu > \frac{(1-\theta)\mu}{\gamma} = \mathcal{O}_{\text{LU}}$ or, equivalently, $\theta > 1 - \gamma$.

In Stage 1, the OTC dealer posts a bid to buy and an ask to sell one unit of the asset to every trader i . The OTC 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 exchange dealer 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—the OTC dealer observes the label ℓ_i before setting the prices for trader i .

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

We impose a regularity condition on the distribution F of hedging benefits.

ASSUMPTION 1: *There exist neighborhoods of $x = 0$ and $x = 2\alpha - 1$ such that the hedging benefit distribution F has a pdf f that is analytic over these two neighborhoods.*⁶

Assumption 1 precludes distributions whose pdfs oscillate between extreme values around the lower and upper bounds of the equilibrium bid-ask spread

⁶ 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 ,

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

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.

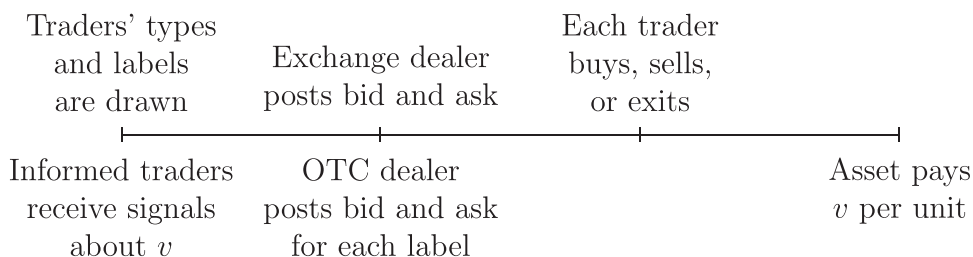


Figure 2. Timeline.

(which we derive in Section I.C). In practice, any commonly used distribution in economics satisfies Assumption 1.

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

ASSUMPTION 2: *If a trader is indifferent between trading over the counter or on the exchange, she trades on the exchange.*

Assumption 2 is purely expositional.⁷ 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.

B. Interpretation

Prices. Our setup features competitive prices, as defined in Glosten and Milgrom (1985). That prices on exchanges are competitive is a good proxy for reality. However, search frictions and 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, restricting or taxing it can still improve welfare. Introducing search frictions or market power would further raise the appeal of restricting the OTC dealer's ability to price discriminate.

Trading protocol. The OTC dealer spontaneously posts label-dependent prices in the model. In practice, an OTC trade occurs in two steps. A trader first requests quotes from her dealers, and then those dealers respond with trader-specific quotes. Section II of the Internet Appendix presents an extension that incorporates RFQ and endogenizes the competition among dealers. 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 OTC 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 in-

⁷ All equilibria that would exist without Assumption 2 are payoff equivalent. We can eliminate Assumption 2 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 $\rho > 0$.

clude the trader's industry (e.g., hedge fund or insurer), marketing or public filings (e.g., active vs. 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 in Cheng and Xiong (2014) from the commodities futures market. The U.S. Commodity Futures Trading Commission labels traders as "hedgers" if they are commodity producers and their past trades are not consistently profitable. Cheng and Xiong (2014) find that the hedgers' positions are far more volatile than their output, 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 at one and 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.

C. Equilibrium

DEFINITION 1: An equilibrium consists of the OTC dealer's bid-ask spread for the LI traders, his spread for the LU traders, the exchange dealer's spread, and traders' venue choice and trading strategies. Each trader chooses the market that offers the lowest spread available to her, and her trading strategy maximizes her expected payoff given this spread. Given the traders' strategies, the OTC dealer sets each of his spreads to the lowest spread that earns zero profit in expectation conditional on the label, and the exchange dealer 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 hereafter use "spread" to refer to "half bid-ask spread."

The trader's equilibrium strategies are cutoff rules. An uninformed buyer or seller trades one 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 a dealer charges a spread s , an uninformed trader who chooses the dealer trades with probability $1 - F(s)$. Conditional on such a trade, the dealer's expected profit is s . However, the dealer's expected loss per trade with an informed trader is $(2\alpha - 1 - s)^+$. Taken

together, the dealer earns zero profit in expectation if and only if

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

where the *informed ratio* β is the mass of informed traders who choose the dealer per unit mass of uninformed traders who choose the same. The zero-profit condition (2) has a unique solution, which we denote by $S(\beta)$.

PROPOSITION 1: (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 OTC spreads for LU and LI traders, respectively, are $S_O = S\left(\frac{1-\theta}{\gamma} \mu\right)$ and S_E . 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 and the No-OTC spread is strictly between the two, $S_E > S_N > S_O$.*

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. Intuitively, traders of the same label choose the same market because, if they split, the prices they face in the two markets must be equal, and thus all traders with that label would choose the exchange by Assumption 2. 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 2. Our results remain unchanged if the LI traders choose the exchange with strictly positive probability. Part (c) immediately follows from (2) and part (b). Section V of the [Internet Appendix](#) cites recent evidence that trading costs are lower over the counter than on exchanges.

II. Welfare, Volume, and Spread

We analyze utilitarian welfare, aggregate trade volume, and the average bid-ask spread upon restricting the OTC dealer's ability to price discriminate. To restrict the OTC dealer, we either lower the accuracy of the traders' labels, θ or γ , or close the OTC market altogether. The closure is equivalent to lowering θ or γ until $\theta = 1 - \gamma$, where the LU and LI traders are equally likely to be informed.

Section II.A defines key terms. Section II.B details the general conditions under which the restriction raises or lowers welfare. Section II.C explains how adverse selection risk jointly determines market shares and welfare.

A. Definitions

Welfare W is the sum of all agents' *ex ante* payoffs. It measures the total gains from trade in our model. More 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 $V := V_O + V_E$ the aggregate trade volume.⁸ The average bid-ask spread \bar{S} is the volume-weighted average of bid-ask spreads in the OTC market and on the exchange,

$$\bar{S} := \frac{V_O}{V} S_O + \frac{V_E}{V} S_E.$$

We refer to “lower average spread \bar{S} ” interchangeably with “higher aggregate volume V ,” because spread $\bar{S} \propto 1/V$ in equilibrium.⁹

Two statistics characterize the effects of restricting the OTC dealer on welfare and volume. First, *marginal volume* Δ_V is the decrease $m(\text{exiters})$ in the mass of uninformed trades upon a marginal increase in the ratio of informed to uninformed traders β :

$$\Delta_V(\beta) := -\left(\int_{S(\beta)}^1 f(s) ds\right)' = \underbrace{S'(\beta) \cdot f(S(\beta))}_{\frac{dm(\text{exiters})}{d\beta}}. \tag{3}$$

Second, *marginal welfare* Δ_W is the decrease in welfare, $m(\text{exiters}) \times \bar{b}(\text{exiters})$, upon a marginal increase in β . It is equivalent to $\Delta_V(\beta)$ times the marginal exiters' hedging benefit $\bar{b}(\text{exiters})$:

$$\Delta_W(\beta) := -\left(\int_{S(\beta)}^1 s f(s) ds\right)' = \underbrace{S'(\beta) f(S(\beta))}_{\Delta_V(\beta)} \cdot \underbrace{S(\beta)}_{\bar{b}(\text{exiters})}. \tag{4}$$

We show in the [Appendix](#) that both Δ_V and Δ_W are well-defined.

B. Main Results

The next proposition states the effects of restricting the OTC dealer on welfare, aggregate trade volume, and the average bid-ask spread. Proofs are in the [Appendix](#).

PROPOSITION 2: *Given any α and 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 $\mu'_h > \mu'_l, \tilde{\mu}'_h > \tilde{\mu}'_l, \mu_h > \mu_l,$*

⁸ Explicitly, $V_O = (1 - \theta)\mu + \gamma \cdot (1 - S_O)$ and $V_E = \theta\mu + (1 - \gamma) \cdot (1 - S_E)$.

⁹ Formally, the zero-profit condition (2) 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 .

$\tilde{\mu}_h > \tilde{\mu}_l$, and $\tilde{\mu} > \underline{\mu}$, all strictly positive, such that:

- (a) Marginally lowering traders' label accuracy θ from θ_h to $\theta_h - d\theta$ strictly raises welfare W for all mass of informed traders $\mu < \mu'_l$ and strictly lowers W for all $\mu > \mu'_h$; marginally lowering γ from γ_h to $\gamma_h - d\gamma$ strictly raises W for all $\mu < \tilde{\mu}'_l$ and strictly lowers W for all $\mu > \tilde{\mu}'_h$.
- (b) Lowering the label accuracy θ from θ_h to θ_l strictly raises W for all $\mu < \mu_l$ and strictly lowers W for all $\mu > \mu_h$; lowering γ from γ_h to γ_l strictly raises W for all $\mu < \tilde{\mu}_l$ and strictly lowers W for all $\mu > \tilde{\mu}_h$.
- (c) Closing the OTC market strictly raises W for all $\mu < \underline{\mu}$ and strictly lowers W for all $\mu > \underline{\mu}$.

Moreover:

- (d) If the marginal welfare Δ_W (defined in Definition 4) is strictly quasiconcave, then $\mu'_h = \mu'_l$, $\tilde{\mu}'_h = \tilde{\mu}'_l$, $\mu_h = \mu_l$, $\tilde{\mu}_h = \tilde{\mu}_l$, and $\tilde{\mu} = \underline{\mu}$.¹⁰
- (e) If the marginal volume Δ_V (defined in Definition 3) is strictly decreasing, lowering θ or γ (marginally or otherwise) or closing the OTC market strictly reduces aggregate trade volume V and strictly widens the average bid-ask spread \bar{S} for all $\mu > 0$.

Proposition 2 has two messages. First, restricting the OTC dealer improves welfare if the mass of informed traders μ is small and harms welfare if μ is large. Second, in stark contrast, these interventions can *always* worsen aggregate trade volume and the average bid-ask spread.

More precisely, Proposition 2 parts (a) to (c) say that, under any commonly used distribution, lowering label accuracy or closing the OTC market raises welfare if 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) says that the restrictions reduce aggregate volume and widen the average spread—no matter how they affect welfare—whenever the marginal volume Δ_V is decreasing. Thus, the aggregate measures of volume and spread are poor indicators of welfare. Proposition 3 states the conditions on model primitives that are necessary and sufficient for parts (d) and (e), and lists broad classes of distributions that satisfy each condition.

Figure 3 illustrates Proposition 2 under the uniform distribution, $F = \mathbb{U}[0, 1]$. The initial parameters are $\theta = \gamma = 0.9$ and $\alpha = 0.98$. We plot the changes in welfare W , the aggregate volume V , and the average spread \bar{S} upon lowering the label accuracy θ to 0.6 as the proportion of informed traders $\mu/(1 + \mu)$ varies. The changes are positive above the red line. Figure 3 confirms that restricting the OTC dealer raises welfare, while it always reduces the aggregate volume and widens the average spread, where μ is low. Adding search frictions or relaxing competitive prices in the OTC market would only expand the range of parameters under which the restriction raises welfare.

¹⁰ 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 the same except that “strictly” becomes “weakly.”




Figure 3. Effects of restricting the OTC dealer's ability to price discriminate. (Color figure can be viewed at wileyonlinelibrary.com)

Intuition. Marginally lowering the label accuracy θ by $d\theta$ raises welfare when the mass μ of informed traders is low because cheap substitution persists yet the volume effect vanishes. (The intuition is analogous for marginally lowering γ , nonmarginally lowering θ or γ , and closing the OTC market.) The key to the intuition is that cheap substitution depends on the relative *levels* of the spreads, while the volume effect depends on the relative *changes* in the spreads of the two markets.

To see that cheap substitution persists when μ is small, we note that cheap substitution equals $\bar{b}(\text{entrants})/\bar{b}(\text{exitors}) = S_E/S_O$. Here, the two spreads S_E and S_O are each proportional to the informed ratio β in their respective markets.¹¹ These β are equal to μ scaled by a constant $\frac{\theta}{1-\gamma}$ on the exchange and by $\frac{1-\theta}{\gamma}$ in the OTC market. The constant is strictly larger on the exchange because the uninformed traders are more likely to be labeled LU than the informed traders, $\gamma > 1 - \theta$. Therefore, cheap substitution S_E/S_O is bounded away from one as μ becomes small.

The volume effect $m(\text{exitors})/m(\text{entrants})$ vanishes as μ becomes small because each dealer's adverse selection cost (gross loss to the informed traders) equals his gross revenue $V_M S_M$ in his market M. For small μ , V_M is approximately equal to the mass of uninformed traders who choose market M whether we lower θ or not. Lowering θ moves a fraction of informed traders from the exchange to the OTC market,¹² which transfers the corresponding adverse selection cost from the exchange to the OTC dealer. The resulting change in the dealer's gross loss $V_M S_M$ is approximately the same in the two markets and equal to $V_M dS_M$. Therefore, the masses of exitors and entrants are about equal at small μ : $m(\text{exitors}) = \gamma |dS_O| \approx V_O |dS_O| \approx V_E |dS_E| \approx$

¹¹ More precisely, $S_M \sim S'(0)\beta_M$, for $M \in \{E, O\}$ as $\mu \rightarrow 0$, where $\beta_E = \frac{\theta}{1-\gamma}\mu$ and $\beta_O = \frac{1-\theta}{\gamma}\mu$.

¹² As θ decreases, the number of informed traders in the exchange $\theta\mu$ decreases and the number of informed traders in the OTC market $(1-\theta)\mu$ increases.

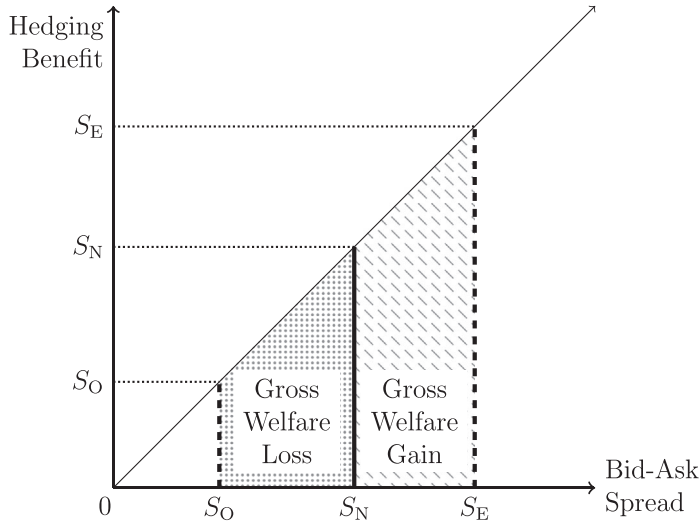


Figure 4. Effect of closing the OTC market on welfare. We mark the initial bid-ask spreads S_O and S_E , and the No-OTC spread S_N that prevails after the closure.

$(1 - \gamma)dS_E = m(\text{entrants})$.¹³ The volume effect $m(\text{exiters})/m(\text{entrants})$ is thus close to one.

Equivalence between label accuracy and closure. Closing the OTC market converges all traders’ bid-ask spreads to the No-OTC spread S_N . Hence, the closure is equivalent to lowering the traders’ label accuracy until the LI and the LU traders are equally likely to be informed, $\theta = 1 - \gamma$. Figure 4 shows $\bar{b}(\text{entrants}) \in (S_N, S_E)$ and $\bar{b}(\text{exiters}) \in (S_O, S_N)$ such that every entrant’s hedging benefit is greater than any exiter’s benefit.

Economic relevance. Most commonly used distributions have strictly quasiconcave marginal welfares Δ_W . Proposition 3 states the necessary and sufficient conditions on the pdf f for its Δ_W to be strictly 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*

$$\frac{(2\alpha - 1)(1 - F(x))}{xf(x)(2\alpha - 1 - x)^2} - \frac{1}{2\alpha - 1 - x} \text{ is strictly quasiconvex in } x \in (0, 2\alpha - 1). \tag{5}$$

Any beta distribution $Beta(a, b)$ for all $a, b > 0$ satisfies condition (5).

¹³ This sentence uses the case of the uniform distribution. For analytical distributions whose $\lim_{x \rightarrow 0} f'(x) > 0$, the approximation is $m(\text{exiters}) = \lim_{x \rightarrow 0} f(x)\gamma|dS_O| \approx \lim_{x \rightarrow 0} f(x)V_O|dS_O| \approx \lim_{x \rightarrow 0} f(x)V_E|dS_E| \approx \lim_{x \rightarrow 0} f(x)(1 - \gamma)dS_E = m(\text{entrants})$. If the distribution has $\lim_{x \rightarrow 0} f'(x) = 0$ or ∞ , the intuition is more involved. Our proof in the Appendix allows all such cases.

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

$$\frac{(2\alpha - 1)(1 - F(x))}{f(x)(2\alpha - 1 - x)^2} - \frac{x}{2\alpha - 1 - x} \text{ is increasing in } x \in (0, 2\alpha - 1). \quad (6)$$

Any beta distribution $Beta(a, b)$ for all $a \leq 1, b \leq 1$ satisfies condition (6).

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 (5) and thereby have quasiconcave marginal welfare Δ_W .¹⁴ To see what distributions are excluded by (5), we note that (5) holds at the extreme ends under any pdf f analytic in some neighborhoods of zero and $2\alpha - 1$: The expression in (5) approaches infinity at the limit as $x \rightarrow 0$ and as $x \rightarrow 2\alpha - 1$. Condition (5) only excludes pdfs that oscillate over moderate values in the support $(0, 2\alpha - 1)$.

We further verify that (6) 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 (6) shows how the aggregate volume and the average bid-ask spread are poor indicators of policies' impact on welfare.

C. 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: restricting the OTC dealer can strictly raise welfare specifically where the OTC market share is high.

PROPOSITION 4: Given any α , we let F and its pdf f be such that¹⁵

$$x \cdot \left(\frac{1}{2\alpha - 1 - x} - \frac{f(x)}{1 - F(x)} \right) \text{ is strictly increasing in } x \in (0, 2\alpha - 1). \quad (7)$$

Given any pairs $(\theta_l, \gamma_l) \leq (\theta_h, \gamma_h) < (1, 1)$ that satisfy $\theta_h + \gamma_l > 1$ and $\theta_l + \gamma_h > 1$, there exist cutoffs $M, \hat{M}, M^* > 0$ such that:

- (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 for all $\frac{V_O}{V} > M$ and strictly reduces W for all $\frac{V_O}{V} < M$; lowering γ from γ_h to γ_l strictly raises welfare W for all $\frac{V_O}{V} > \hat{M}$ and strictly reduces W for all $\frac{V_O}{V} < \hat{M}$.
- (c) Closing the OTC market strictly raises W for all $\frac{V_O}{V} > M^*$ and strictly reduces W for all $\frac{V_O}{V} < M^*$.

¹⁴ 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 (5) when truncated to $[0, 1]$.

¹⁵ The uniform distribution, $F = \mathbb{U}[0, 1]$, trivially meets (7). We numerically verify that given any α , the beta distribution satisfies (7) for wide ranges of its parameters $a, b > 0$.

Intuitively, a larger μ mechanically raises the share of LI traders, who trade on the exchange. Under condition (7), 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: Holding all parameters constant except for μ , restricting the OTC dealer 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 exchange spread 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 U.S.-listed equities (see Section VI.C of the [Internet Appendix](#)).

III. Optimal Pigouvian Tax

Section II provides policy guidance when one knows that adverse selection risk is low (restrict the OTC dealer) or high (keep as is). We now turn to an optimal Pigouvian tax that can be implemented without knowing the severity of adverse selection risk. Instead, a simple sufficient statistic derived from a local experiment can determine whether the current tax is too low or too high. Computing this statistic, the WSR, requires only trade volumes and bid-ask spreads.

A. Characterizing the Optimal Pigouvian Tax

Gross tax revenue T is transferred from the OTC dealer to the exchange dealer.¹⁶ The lump-sum 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.¹⁷ More precisely, the OTC dealer's condition becomes

$$S_O(T) \cdot \underbrace{[1 - F(S_O(T)) + \beta_O]}_{V_O(T)} \gamma = (2\alpha - 1)\beta_O \gamma + T, \quad (8)$$

and the exchange dealer's condition becomes

$$S_E(T) \cdot \underbrace{[1 - F(S_E(T)) + \beta_E]}_{V_E(T)} \cdot (1 - \gamma) = (2\alpha - 1)\beta_E \cdot (1 - \gamma) - T. \quad (9)$$

Increasing tax T raises OTC spread $S_O(T)$ and lowers exchange spread $S_E(T)$.

¹⁶ Taxing the exchange dealer is ineffective, as every LI trader would choose the OTC market and no trade would occur on the exchange.

¹⁷ Examples include a flat per-trade transaction tax, a lump-sum tax, and a nonlinear levy on OTC volume. The same is true with multiple dealers in each market as long as the tax T is split between the dealers in proportion to their trade volumes.

We define the *optimal Pigouvian tax* as follows.

DEFINITION 2: The optimal Pigouvian tax T^* is the lump-sum tax T that maximizes welfare W .

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

PROPOSITION 5: Given any triple (α, θ, γ) : (a) Imposing the optimal tax T^* attains a strictly higher welfare W than closing the OTC market. (b) If $F = \mathcal{U}[0, 1]$, the optimal Pigouvian tax T^* (defined in Definition 2) is unique.

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. We define this *zero-OTC tax* \bar{T} as follows.

DEFINITION 3: The 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,

$$S_O(\bar{T}) = S_E(\bar{T}).$$

The zero-OTC tax \bar{T} therefore 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. Given that 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, $T^* < \bar{T}$.

B. Implementing the Optimal Pigouvian Tax

We use the trade-off between the cheap substitution effect and the volume effect (defined in (1)) to develop a simple procedure to implement the optimal Pigouvian tax T^* . Marginally increasing the tax T raises welfare W if and only if

$$\underbrace{\left| \frac{S_E}{S_O} \right|}_{\text{Cheap Substitution}} > \underbrace{\left| \frac{dV_O}{dV_E} \right|}_{\text{Volume Effect}}.$$

The optimal tax T^* , where $T^* > 0$, equalizes cheap substitution and the volume effect.¹⁸

$$\frac{S_E(T^*)}{S_O(T^*)} = \left| \frac{dV_O(T^*)}{dV_E(T^*)} \right|.$$

¹⁸ We differentiate (8) and (9) to compute dV_O and dV_E . Doing so, we obtain

$$\begin{aligned} |dV_O(t)| &= |\gamma S'_O(t)| dt = \frac{dt}{1 - 2S_O(t) + \beta_O} \\ |dV_E(t)| &= |(1 - \gamma) S'_E(t)| dt = \frac{dt}{1 - 2S_E(t) + \beta_E}. \end{aligned}$$

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

$$\text{WSR}(T) := \left| \frac{S_E(T)dV_E(T)}{S_O(T)dV_O(T)} \right|. \quad (10)$$

PROPOSITION 6: *Let the current Pigouvian tax be $T \geq 0$. If the tax marginally increases by $dT > 0$ and the consequent weighted spread ratio is $\text{WSR}(T)$ (defined in Definition 10):*

- (a) *Welfare W increases ($dW/dT > 0$) if and only if $\text{WSR}(T) > 1$, and W decreases ($dW/dT < 0$) if and only if $\text{WSR}(T) < 1$.*
- (b) *If $F = \mathcal{U}[0, 1]$, the current tax is strictly lower than the optimal tax $T < T^*$ if and only if $\text{WSR}(T) > 1$, and $T > T^*$ if and only if $\text{WSR}(T) < 1$.*
- (c) *If $F = \mathcal{U}[0, 1]$, $\text{WSR}(T)$ is strictly decreasing in T .*

Proposition 6 spells out a simple “WSR Rule” for tax setting: raise the tax T when $\text{WSR}(T) > 1$, cut T when $\text{WSR}(T) < 1$, and keep the current tax when $\text{WSR}(T) = 1$. Shifting the tax T following the WSR Rule would strictly raise welfare. When $\text{WSR}(T) = 1$, welfare is at a local maximum. The WSR Rule holds under any distribution F with support $[0,1]$. Moreover, under the uniform distribution $F = \mathcal{U}[0, 1]$, iteratively following the WSR Rule obtains the optimal tax T^* .

Estimating the WSR. The volume term in the WSR, $|dV_O/dV_E|$, is easily computed using trade volumes just before and just after the change in the tax T . In the United States and the European Union, most trades must already be reported in nearly real time and specify whether it was executed on an exchange or over the counter. The spread term S_E/S_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 pretrade anonymous methods. Such methods include limit order books, batch auctions, dark pools, and all-to-all requests for quote. Most important asset classes are actively traded via these methods, including equities, their futures and options, treasuries, corporate bonds, and repurchase agreements. Footnote lists examples.

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 aggregate volume increases, and (ii) closing the OTC market increases welfare where adverse selection risk is low.

OTC versus exchange. Our analysis is most closely related to the literature on venue choice between OTC and centralized markets. One strand in this literature abstracts from adverse selection and focuses on the presence of search frictions (Pagano (1989), Vogel (2019), Rust and Hall (2003)) 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) show 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.

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. Akerlof (1970) cannot obtain our main results for two reasons: (I) the private and the common values are perfectly correlated and (II) each uninformed trader is unaware of her own private value; otherwise, she would learn the common value and face no adverse selection risk. Section IV of the [Internet Appendix](#) 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, the minimum extension to the Akerlof (1970) framework necessary to generate our results is heterogeneity in private values that are (I') decoupled from the common value and (II') known to the uninformed traders.

Cream skimming. Bolton, Santos, and Scheinkman (2016) also feature cream skimming by dealers. Nonetheless, Bolton, Santos, and Scheinkman (2016) cannot generate our results because cream skimming in their framework affects welfare through a trade-off that is orthogonal to adverse selection risk. On the one hand, two exogenous frictions in the OTC market reduce welfare: their dealers (i) hold market power and (ii) incur a deadweight cost to separate high-quality 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 the absence of the two exogenous frictions, (i) and (ii). 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 reduces welfare whenever adverse selection risk is low, even with a competitive dealer and zero deadweight costs.

Nonanonymity 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 trade orders of certain uninformed traders benefits those traders and leaves others worse off. A large literature shows analogous effects (Admati and Pfleiderer (1991), Fishman and Longstaff (1992), Forster and George (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.¹⁹

Binding minimum wage. Policies that we examine lead the uninformed traders with 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, raising the minimum wage does not increase employment without 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.

Third-degree price discrimination. 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 total surplus. He identifies a “misallocation effect” whereby output is inefficiently distributed whenever different consumers are charged different prices (Aguirre, Cowan, and Vickers (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 increase or decrease welfare (Bergemann, Brooks, and Morris (2015)).²⁰ 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.

Broadly related literature. This paper belongs to the large 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 restricting price discrimination. More distantly related is the literature on

¹⁹ The only exception is Admati and Pfleiderer (1991), whose result on welfare (their Proposition 1(d)) is the opposite of ours.

²⁰ Specifically, price discrimination can generate any outcome under which the producer is at least as well off as with a uniform price, consumers receive nonnegative payoffs, and the allocation is feasible (Bergemann, Brooks, and Morris (2015)).

how traders choose or split orders across multiple venues that do not feature price discrimination (Hendershott and Mendelson (2000), Zhu (2014), Pagnotta and Philippon (2018), Lee (2019), Chao, Yao, and Ye (2019), Baldauf and Mollner (2021), Babus and Parlato (2024)).

V. Conclusion

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

Our model abstracts from some important sources of inefficiency on exchanges. We do not consider how welfare might be affected by price impact (Vives (2011)) or sniping (Budish, Cramton, and Shim (2015)). Several papers already propose improvements to the design of exchanges that address such frictions.²¹ Our focus on resolving the inefficiency of OTC trading complements this literature.

Previous work shows that price discovery in secondary markets affects corporate investment decisions (e.g., 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 ratio β of informed to uninformed traders. Second, analyzing the *aggregate* price discovery requires a stance on how exactly 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.

Acknowledgments

Open access funding provided by Central European University Private University - CEU GmbH/KEMÖ.

Initial submission: April 1, 2022; Accepted: November 8, 2023
Editors: Stefan Nagel, Philip Bond, Amit Seru, and Wei Xiong

²¹ For example, Malamud and Rostek (2017) and Chen and Duffie (2021) show that optimal market fragmentation can address price impact, and Budish, Cramton, and Shim (2015) proposes frequent batch auctions to resolve sniping by fast traders.

Appendix: Proofs

A. Proofs for Section I.C

PROOF OF PROPOSITION 1: *Part (a)*: It suffices to show that there exists at least one solution to the zero-profit condition (2) so that $S(\beta)$ is well-defined. We use Figure A.1, which plots the exchange dealer’s payoff $s \cdot [1 - F(s)] - (2\alpha - 1 - s)^+\beta$ over the spread s . The payoff curve is continuous. Her payoff is negative at $s = 0$, as she is adversely selected yet has no revenue, while it is positive at $s = 2\alpha - 1$, as she breaks even on the trades against the informed traders and profits on the uninformed traders. The Intermediate Value Theorem implies that there exists at least one solution to the zero-profit condition. *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 clearly in Figure A.1: increasing β (to β') shifts the entire payoff curve downward and the crossing point $S(\beta)$ to the right. Intuitively, as the informed traders impose losses on the exchange dealer, 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 OTC dealer. For an LU trader, the equilibrium OTC spread is $S(\beta_{LU})$, 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, LI traders’ OTC spread is $S\left(\frac{\theta}{1-\gamma}\mu\right)$.

Third, we turn to the exchange spread S_E . If the exchange dealer sets $S_E \in (S(\beta_{LU}), S(\beta_{LI})]$, all LU traders choose the OTC market, whereas all LI traders choose the exchange. Then the informed ratio on the exchange is $\beta_E = \beta_{LI}$. The exchange dealer thus earns zero profit if and only if she sets

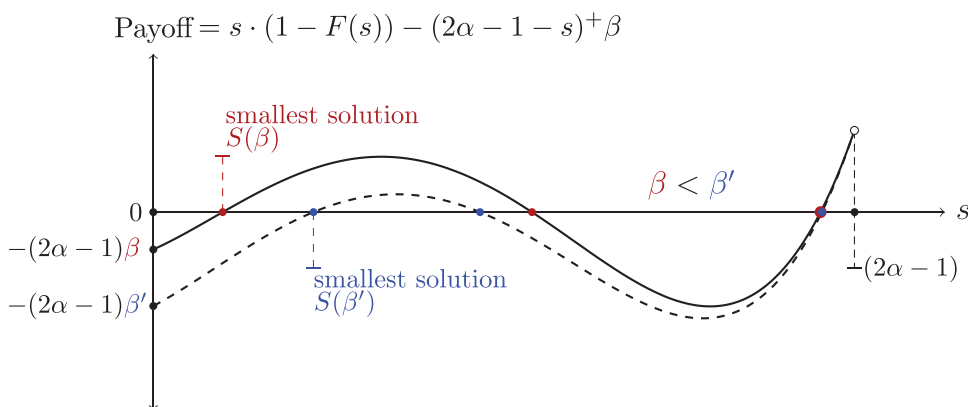


Figure A.1. Finding the equilibrium spread. (Color figure can be viewed at wileyonlinelibrary.com)

$S_E = S(\beta_{LI})$ in this case. If the exchange dealer sets $S_E \leq S(\beta_{LU})$, then every trader chooses the exchange, implying $\beta_E = \mu > \beta_{LU}$, and thus the exchange dealer earns a nonzero profit. Therefore, in equilibrium, (i) the exchange spread is $S_E = S(\beta_{LI}) = S\left(\frac{\theta}{1-\gamma} \mu\right)$, the lowest spread that earns the exchange dealer a zero profit, and (ii) all LU traders choose the OTC market, whereas all LI traders choose the exchange.

Part (c): Since $\beta_O < \beta_E$ and the spread function $S(\beta)$ is strictly increasing in the informed ratio β , it follows that $S_O < S_E$.

PROPOSITION A1: *If the pdf f of the hedging benefit distribution F is continuous and strictly positive in $[0,1]$, then the volume effect $\frac{m(\text{exitors})}{m(\text{entrants})}$ associated with marginally reducing θ from $\theta = 1$ is finite.*

PROOF OF PROPOSITION A1: The equilibrium half bid-ask spread $S(\beta)$ is the smallest solution to the dealers' zero-profit condition,

$$s \cdot [1 - F(s)] - (2\alpha - 1 - s)^+ \beta = 0, \quad (\text{A1})$$

which Figure A.1 illustrates.

The masses of entrants and exitors are

$$m(\text{exitors}) = f(S(\beta_O))S'(\beta_O^+) \mu \, d\theta = f(0)S'(0^+) \mu \, d\theta,$$

$$m(\text{entrants}) = f(S(\beta_E))S'(\beta_E^-) \mu \, d\theta,$$

in which $S'(\beta^\pm)$ are the left and right derivatives of $S(\beta)$ at β . Applying the Implicit Function Theorem to (A1), for any $\beta > 0$, we have

$$S'(\beta^-) = \frac{2\alpha - 1 - S(\beta)}{1 - F(S(\beta)) - S(\beta)f(S(\beta)) + \beta}.$$

Since a dealer's expected payoff, $s \cdot [1 - F(s)] - (2\alpha - 1 - s)^+ \beta$, crosses zero for the first time at $s = S(\beta)$ from below, its derivative with respect to s at $s = S(\beta)$ must be nonnegative,

$$1 - F(S(\beta)) - S(\beta)f(S(\beta)) + \beta \geq 0.$$

Thus, $S'(\beta^-)$ is either infinite or strictly positive for any $\beta > 0$. Similarly, $S'(0^+) > 0$. Therefore, $m(\text{exitors})/m(\text{entrants})$ is finite.

B. Proofs for Section II

The proof of Proposition 1 shows that the spread function $S(\beta)$ is increasing. Given that $S(\beta)$ is also left-continuous in β , $S(\beta)$ is left-differentiable. We let $S'(\beta)$ be the left derivative. Marginal volume Δ_V and marginal welfare Δ_W as in Definitions 3 and 4 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) $\lim_{s \downarrow 0} s f(s) = 0$, (ii) $\lim_{s \downarrow 0} (s f)'(s) \in$

$\mathbb{R}^+ \cup \{\infty\}$, 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 $sf(s) - c$. The Identity Theorem would then imply that $sf(s) - c \equiv 0$ in the neighborhood $(0, \varepsilon)$ of zero, implying in turn that $f(s) \sim c/s$ as $s \downarrow 0$. This contradicts the integrability of the pdf f in $(0, \varepsilon)$. Furthermore, it must be the case 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 nonnegative or infinity, since $\lim_{s \downarrow 0} sf(s) = 0$. (iii) Since $[\ln(sf)]'$ is analytic in a neighborhood of zero, $\lim_{s \downarrow 0} [\ln(sf)]'(s)$ exists. Since $\lim_{s \downarrow 0} sf(s) = 0$, it must be that $\lim_{s \downarrow 0} [\ln(sf)]'(s) = \infty$. As β approaches zero, $S(\beta)$ approaches zero. 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 $\Delta_W(\beta) \rightarrow 0$ as $\beta \rightarrow 0$. One can verify that $\ln(\Delta_W)$ is differentiable in a neighborhood of $\beta = 0$ and properties (i) to (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 zero. Taken together, Δ_W begins at $\Delta_W(0) = 0$, from which Δ_W strictly increases before eventually strictly decreasing toward 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) to (c): As the label accuracy θ falls from θ_h to θ_l , the change in aggregate trade volume V is

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

which is equal to

$$(1 - \gamma) \int_{\frac{\theta_l}{1-\gamma}\mu}^{\frac{\theta_h}{1-\gamma}\mu} \Delta_V(\beta) d\beta - \gamma \int_{\frac{1-\theta_l}{\gamma}\mu}^{\frac{1-\theta_h}{\gamma}\mu} \Delta_V(\beta) d\beta. \tag{A2}$$




Figure A.2. Generic Δ_W . (Color figure can be viewed at wileyonlinelibrary.com)

Similarly, the change in welfare is

$$(1 - \gamma) \underbrace{\int_{\frac{\theta_l}{1-\gamma} \mu}^{\frac{\theta_h}{1-\gamma} \mu} \Delta_W(\beta) d\beta}_{\text{Gross welfare gain}} - \gamma \underbrace{\int_{\frac{1-\theta_h}{\gamma} \mu}^{\frac{1-\theta_l}{\gamma} \mu} \Delta_W(\beta) d\beta}_{\text{Gross welfare loss}}. \tag{A3}$$

The proofs are intuitive with the aid of graphs. Figure A.2 plots a generic Δ_W . We 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 (A3) (marked “Gross Loss/ γ ” in red) has 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 (A3) (marked “Gross Gain/ $(1 - \gamma)$ ” in red) has 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)).




Figure A.3. Generic Δ_W . (Color figure can be viewed at wileyonlinelibrary.com)

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

$$\underbrace{(1 - \gamma_h) \int_{\frac{\theta}{1-\gamma_l} \mu}^{\frac{\theta}{1-\gamma_h} \mu} \Delta_V(\beta) d\beta}_{\text{Entry by uninformed LI traders}} - \underbrace{\gamma_h \int_{\frac{1-\theta}{\gamma_h} \mu}^{\frac{1-\theta}{\gamma_l} \mu} \Delta_V(\beta) d\beta}_{\text{Exit by uninformed LU traders}} - \underbrace{(\gamma_h - \gamma_l) \int_{\frac{1-\theta}{\gamma_l} \mu}^{\frac{\theta}{1-\gamma_l} \mu} \Delta_V(\beta) d\beta}_{\text{Exit by relabeled traders}}. \tag{A4}$$

Similarly, the change in welfare is

$$(1 - \gamma_h) \int_{\frac{\theta}{1-\gamma_l} \mu}^{\frac{\theta}{1-\gamma_h} \mu} \Delta_W(\beta) d\beta - \gamma_h \int_{\frac{1-\theta}{\gamma_h} \mu}^{\frac{1-\theta}{\gamma_l} \mu} \Delta_W(\beta) d\beta - (\gamma_h - \gamma_l) \int_{\frac{1-\theta}{\gamma_l} \mu}^{\frac{\theta}{1-\gamma_l} \mu} \Delta_W(\beta) d\beta. \tag{A5}$$

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 in (A5) (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 (A5) has a strict lower bound

$$(\gamma_h - \gamma_l) \cdot \|\overline{HCFO}\|.$$




Figure A.4. Quasiconcave Δ_W . (Color figure can be viewed at wileyonlinelibrary.com)

The third term in (A5) equals $(\gamma_h - \gamma_l) \cdot \|BCFE\|$. Then, (A5) is strictly greater than

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

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 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 in which 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.4 Panel (a)) 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.4 Panel (b)). The proof works whether or not Δ_W is continuous. We show that (i) the change in welfare (A3) is strictly positive for all $\mu < \mu_L$ and strictly negative for all $\mu > \mu_R$; and (ii) (A3) is strictly decreasing between μ_L and μ_R . Together, (i) and (ii) establish the existence of a single cut-off.

To prove (i), we set $\mu = \mu_L$. A strict upper bound of the second integral in (A3) (“Gross Loss/ γ ” in red) is

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

A strict lower bound of the first integral in (A3) (“Gross Gain/(1 - γ)” in red) is $\|\overline{CDFH}\|$. 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 (A3) is strictly positive for all $\mu \leq \mu_L$. Similarly, (A3) is strictly negative for all $\mu \geq \mu_R$.

To prove (ii), that is, that (A3) is strictly decreasing over $\mu \in (\mu_L, \mu_R)$, we take the derivative of (A3) with respect to μ ,

$$\underbrace{\left(\theta_h \cdot \|\overline{DF}\| - \theta_l \cdot \|\overline{GH}\|\right)}_{\text{Derivative of the gross welfare gain}} - \underbrace{\left((1 - \theta_l) \cdot \|\overline{BE}\| - (1 - \theta_h) \cdot \|\overline{IJ}\|\right)}_{\text{Derivative of the gross welfare loss}}. \tag{A6}$$

Due to Δ_W being quasiconcave and given how μ_L and μ_R are chosen, both $\|\overline{BE}\|$ and $\|\overline{GH}\|$ are strictly greater than $\|\overline{DF}\|$ and $\|\overline{IJ}\|$ when $\mu_L < \mu < \mu_R$, and thus (A6) 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 (A5) is strictly more than

$$(\gamma_h - \gamma_l)(\|\overline{O'GHO}\| - \|\overline{O'E'FO}\| - \|\overline{EGHF}\|) > 0.$$

For $\mu \geq \tilde{\mu}_R$, (A5) is strictly less than

$$(\gamma_h - \gamma_l)(\|\overline{O'GHO}\| - \|\overline{O'E'FO}\| - \|\overline{EGHF}\|) < 0.$$




Figure A.5. Decrease in label accuracy γ .

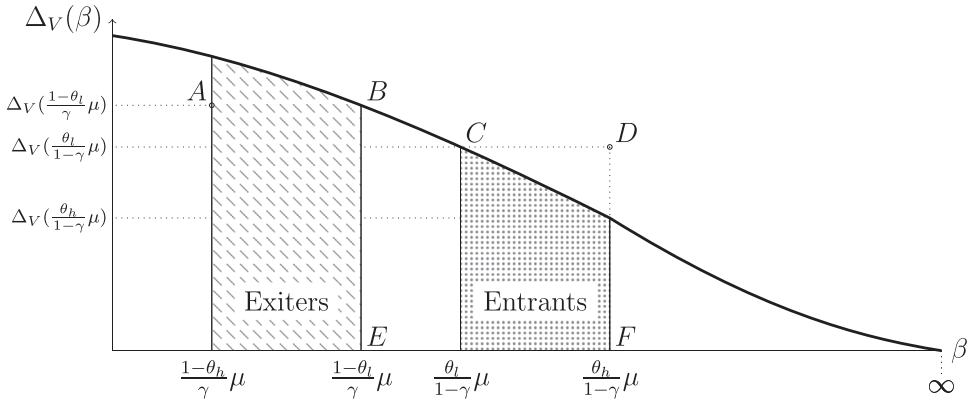


Figure A.6. Decreasing Δ_V .

For $\mu \in (\tilde{\mu}_L, \tilde{\mu}_R)$, the derivative of (A5) 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 .

Part (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 (A2) has a strict lower bound

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

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

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




Figure A.7. Decrease in label accuracy γ .

and thus the change (A2) 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 .

As the label accuracy γ falls from γ_h to γ_l , the change (A4) in aggregate volume V is strictly less than

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

as 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 , it follows that $\Delta_W(\beta)$ is quasiconcave if and only if $\beta'(x)/(xf(x))$ is quasiconvex in $x \in (0, 2\alpha - 1)$. Differentiating (2) with respect to x yields

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

which can be rearranged to

$$\beta'(x) = \frac{1 - F(x) + \beta - xf(x)}{2\alpha - 1 - x}. \tag{A7}$$

From (2), we can express $\beta(x)$ as function of x :

$$\beta(x) = \frac{(1 - F(x))x}{2\alpha - 1 - x}. \tag{A8}$$

Therefore, marginal welfare $\Delta_W(\beta)$ is quasiconcave if and only if condition (5) holds. Likewise, marginal volume $\Delta_V(\beta)$ is decreasing if and only if (6) is true.

PROOF OF PROPOSITION 4: *Part (a)*: The ratio V_E/V_O equals

$$\frac{V_E}{V_O} := \frac{(1 - \gamma)[1 - F(S(\beta_E))] + \theta\mu}{\gamma[1 - F(S(\beta_O))] + (1 - \theta)\mu}.$$

The derivative of V_E/V_O with respect to μ is strictly positive if and only if

$$\frac{-f(S(\beta_O))S'(\beta_O) + 1}{\frac{1 - F(S(\beta_O))}{\beta_O} + 1} < \frac{-f(S(\beta_E))S'(\beta_E) + 1}{\frac{1 - F(S(\beta_E))}{\beta_E} + 1}.$$

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 (A7) and (A8), the above is true if and only if

$$\frac{-\frac{(2\alpha - 1 - x)f(x)}{\frac{2\alpha - 1 - x}{2\alpha - 1} - x} + 1}{\frac{2\alpha - 1 - x}{x} + 1}$$

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.

C. Proofs for Section III

PROOF OF PROPOSITIONS 5 and 6: We prove Propositions 5 and 6 together. First, imposing some Pigouvian tax T obtains strictly higher welfare than closing the OTC market (Proposition 5(a)). 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 , $S_O(\bar{T}) = S_E(\bar{T}) = S_N$. The definition of \bar{T} implies $S_O(\bar{T}) = S_E(\bar{T})$, where the OTC dealer's and the exchange dealer's

zero-profit conditions are

$$S(\bar{T})\left[1 - F(S(\bar{T})) + \beta_O\right]\gamma = (2\alpha - 1)\beta_O\gamma + \bar{T},$$

$$S(\bar{T})\left[1 - F(S(\bar{T})) + \beta_E\right] \cdot (1 - \gamma) = (2\alpha - 1)\beta_E \cdot (1 - \gamma) - \bar{T}.$$

Summing the two equations gives

$$S(\bar{T})\left[1 - F(S(\bar{T})) + \mu\right] = (2\alpha - 1)\mu.$$

Since the zero-OTC tax \bar{T} is unique, it must be the case that $S_O(\bar{T}) = S_E(\bar{T}) = S_N$.

Step 2: We show that a marginal tax cut $-dT$ from the tax \bar{T} increases aggregate trade volume V . Volume V increases if and only if $|dV_O(\bar{T})/dV_E(\bar{T})| > 1$, where

$$dV_O(\bar{T}) = f(S_N)S'_O(\bar{T})\gamma dT \text{ and } dV_E(\bar{T}) = f(S_N)S'_E(\bar{T}) \cdot (1 - \gamma) dT.$$

Taking the derivative of the zero-profit conditions (8) and (9) with respect to T ,

$$S'_O(T)\left[1 - F(S_O(T)) - f(S_O(T))S_O(T) + \beta_O\right]\gamma = 1,$$

$$[5em]S'_E(T)\left[1 - F(S_E(T)) - f(S_E(T))S_E(T) + \beta_E\right] \cdot (1 - \gamma) = -1.$$

Then, using $S_O(\bar{T}) = S_E(\bar{T}) = S_N$, we have

$$\left|\frac{dV_O(\bar{T})}{dV_E(\bar{T})}\right| = \left|\frac{S'_O(\bar{T})\gamma}{S'_E(\bar{T}) \cdot (1 - \gamma)}\right| = \frac{1 - F(S_N) - f(S_N)S_N + \beta_E}{1 - F(S_N) - f(S_N)S_N + \beta_O}.$$

The ratio $|dV_O(\bar{T})/dV_E(\bar{T})| > 1$ because $\beta_E > \beta_O$.

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

Taken together, Steps 1 to 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 $|S_O(T)dV_O(T)|$. The gross welfare gain among traders on the exchange is $|S_E(T)dV_E(T)|$. Thus, welfare W increases ($dW/dT > 0$) if and only if $WSR(T) > 1$ (Proposition 6 (a)).

If $F = \mathbb{U}[0, 1]$, from the first half of this proof,

$$|S_O(T)dV_O(T)| = \frac{S_O(T)}{1 - 2S_O(T) + \beta_O} dT = \frac{1}{\frac{1+\beta_O}{S_O(T)} - 2} dT,$$

and

$$|S_E(T)dV_E(T)| = \frac{S_E(T)}{1 - 2S_E(T) + \beta_E} dT = \frac{1}{\frac{1+\beta_E}{S_E(T)} - 2} dT.$$

Because $S_O(T)$ is strictly increasing in T while $S_E(T)$ 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,

$$\text{WSR}(T) = \left| \frac{S_E(T)dV_E(T)}{S_O(T)dV_O(T)} \right| = \frac{\frac{1+\beta_O}{S_O(T)} - 2}{\frac{1+\beta_E}{S_E(T)} - 2}$$

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

Therefore, there exists a unique $T^* \in [0, \bar{T}]$ such that $\text{WSR}(T) > 1$ for $T \in [0, T^*)$ and $\text{WSR}(T) < 1$ for $T \in (T^*, \bar{T}]$. 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 . Propositions 5(b) and 6(b) follow.

REFERENCES

- Admati, Anat R., and Paul Pfleiderer, 1991, Sunshine trading and financial market equilibrium, *Review of Financial Studies* 4, 443–481.
- Aguirre, Iñaki, Simon Cowan, and John Vickers, 2010, Monopoly price discrimination and demand curvature, *American Economic Review* 100, 1601–1615.
- Akerlof, George A., 1970, The market for “Lemons”: Quality uncertainty and the market mechanism, *Quarterly Journal of Economics* 84, 488–500.
- Allen, Stephen P., 1987, Taxes, redistribution, and the minimum wage: A theoretical analysis, *Quarterly Journal of Economics* 102, 477–489.
- Babus, Ana, and Cecilia Parlatore, 2024, Strategic fragmented markets, *Journal of Financial Economics* 145, 876–908.
- Baldauf, Markus, and Joshua Mollner, 2021, Trading in fragmented markets, *Journal of Financial and Quantitative Analysis* 56, 93–121.
- Benveniste, Lawrence M., Alan J. Marcus, and William J. Wilhelm, 1992, What’s special about the specialist? *Journal of Financial Economics* 32, 61–86.
- Bergemann, Dirk, Benjamin Brooks, and Stephen Morris, 2015, The limits of price discrimination, *American Economic Review* 105, 921–957.
- Biais, Bruno, Thierry Foucault, and François Salanié, 1998, Floors, dealer markets and limit order markets, *Journal of Financial Markets* 1, 253–284.
- Boadway, Robin, and Katherine Cuff, 2001, A minimum wage can be welfare-improving and employment-enhancing, *European Economic Review* 45, 553–576.
- Bolton, Patrick, Tano Santos, and Jose A. Scheinkman, 2016, Cream-skimming in financial markets, *Journal of Finance* 71, 709–736.
- Budish, Eric, Peter Cramton, and John Shim, 2015, The high-frequency trading arms race: Frequent batch auctions as a market design response, *Quarterly Journal of Economics* 130, 1547–1621.
- Chao, Yong, Chen Yao, and Mao Ye, 2019, Why discrete price fragments U.S. stock exchanges and disperses their fee structures, *Review of Financial Studies* 32, 1068–1101.
- Chen, Daniel, and Darrell Duffie, 2021, Market fragmentation, *American Economic Review* 111, 2247–2274.
- Cheng, Ing-Haw, and Wei Xiong, 2014, Why do hedgers trade so much? *Journal of Legal Studies* 43, S183–S207.

- Clemens, Jeffrey, 2021, How do firms respond to minimum wage increases? Understanding the relevance of non-employment margins, *Journal of Economic Perspectives* 35, 51–72.
- Collin-Dufresne, Pierre, Benjamin Junge, and Anders B. Trolle, 2020, Market structure and transaction costs of index CDSs, *Journal of Finance* 75, 2719–2763.
- de Roure, Calebe, Emanuel Moench, Loriana Pelizzon, and Michael Schneider, 2021, OTC discount. <https://papers.ssrn.com/abstract=3744758>
- Desgranges, Gabriel, and Thierry Foucault, 2005, Reputation-based pricing and price improvements, *Journal of Economics and Business* 57, 493–527.
- Dugast, Jérôme, Semih Üslü, and Pierre-Olivier Weill, 2022, A theory of participation in OTC and centralized markets, *Review of Economic Studies* 89, 3223–3266.
- Fishman, Michael J., and Francis A. Longstaff, 1992, Dual trading in futures markets, *Journal of Finance* 47, 643–671.
- Forster, Margaret M., and Thomas J. George, 1992, Anonymity in securities markets, *Journal of Financial Intermediation* 2, 168–206.
- Foucault, Thierry, Sophie Moinas, and Erik Theissen, 2007, Does anonymity matter in electronic limit order markets? *Review of Financial Studies* 20, 1707–1747.
- Glode, Vincent, and Christian C. Opp, 2019, Over-the-counter vs. limit-order markets: The role of traders' expertise, *Review of Financial Studies* 33, 866–915.
- Glosten, Lawrence R., and Paul R. Milgrom, 1985, Bid, ask and transaction prices in a specialist market with heterogeneously informed traders, *Journal of Financial Economics* 14, 71–100.
- Goldstein, Itay, and Alexander Guembel, 2008, Manipulation and the allocational role of prices, *Review of Economic Studies* 75, 133–164.
- Guesnerie, Roger, and Kevin Roberts, 1987, Minimum wage legislation as a second best policy, *European Economic Review* 31, 490–498.
- Hau, Harald, Peter Hoffmann, Sam Langfield, and Yannick Timmer, 2021, Discriminatory pricing of over-the-counter derivatives, *Management Science* 67, 6660–6677.
- Hendershott, Terrence, and Ananth Madhavan, 2015, Click or call? Auction versus search in the over-the-counter market, *Journal of Finance* 70, 419–447.
- Hendershott, Terrence, and Haim Mendelson, 2000, Crossing networks and dealer markets: Competition and performance, *Journal of Finance* 55, 2071–2115.
- Lee, David, and Emmanuel Saez, 2012, Optimal minimum wage policy in competitive labor markets, *Journal of Public Economics* 96, 739–749.
- Lee, Tomy, 2019, Latency in fragmented markets, *Review of Economic Dynamics* 33, 128–153.
- Malamud, Semyon, and Marzena Rostek, 2017, Decentralized exchange, *American Economic Review* 107, 3320–3362.
- Malinova, Katya, and Andreas Park, 2013, Liquidity, volume and price efficiency: The impact of order vs. quote driven trading, *Journal of Financial Markets* 16, 104–126.
- Manning, Alan, 2021, The elusive employment effect of the minimum wage, *Journal of Economic Perspectives* 35, 3–26.
- O'Hara, Maureen, and Xing Alex Zhou, 2021, The electronic evolution of corporate bond dealers, *Journal of Financial Economics* 140, 368–390.
- Pagano, Marco, 1989, Trading volume and asset liquidity, *Quarterly Journal of Economics* 104, 255–274.
- Pagano, Marco, and Ailsa Roell, 1996, Transparency and liquidity: A comparison of auction and dealer markets with informed trading, *Journal of Finance* 51, 579–611.
- Pagnotta, Emiliano S, and Thomas Philippon, 2018, Competing on speed, *Econometrica* 86, 1067–1115.
- Pigou, A., 1920, *The Economics of Welfare* (MacMillan and Co., London).
- Riggs, Lynn, Esen Onur, David Reiffen, and Haoxiang Zhu, 2020, Swap trading after Dodd-Frank: Evidence from index CDS, *Journal of Financial Economics* 137, 857–886.
- Rindi, Barbara, 2008, Informed traders as liquidity providers: Anonymity, liquidity and price formation, *Review of Finance* 12, 497–532.
- Röell, Ailsa, 1990, Dual-capacity trading and the quality of the market, *Journal of Financial Intermediation* 1, 105–124.

- Rust, John, and George Hall, 2003, Middlemen versus market makers: A theory of competitive exchange, *Journal of Political Economy* 111, 353–403.
- Securities Industry and Financial Markets Association, 2018, Re: Post-trade name give-up on swap execution facilities; Request for Comment - RIN 3038-AE79, 83 Fed. Reg. 61751.
- Selten, R., 1975, Reexamination of the perfectness concept for equilibrium points in extensive games, *International Journal of Game Theory* 4, 25–55.
- Seppi, Duane J., 1990, Equilibrium block trading and asymmetric information, *Journal of Finance* 45, 73–94.
- Vives, Xavier, 2011, Strategic supply function competition with private information, *Econometrica* 79, 1919–1966.
- Vogel, Sebastian, 2019, When to introduce electronic trading platforms in over-the-counter markets? <http://www.ssrn.com/abstract=2895222>
- Zhu, Haoxiang, 2014, Do dark pools harm price discovery? *Review of Financial Studies* 27, 747–789.

Supporting Information

Additional Supporting Information may be found in the online version of this article at the publisher's website:

**Appendix S1: Internet Appendix.
Replication Code.**