Imperfect Competition and Moral Hazard in Financial Markets

Milena Wittwer*

February 5, 2026

Abstract

Moral hazard is traditionally analyzed as a bilateral principal–agent problem. In financial markets, however, intermediaries (agents) investing on behalf of clients (principals) must also compete with one another when trading, creating an interaction between incentives and market clearing. I develop a model to study how frictions from moral hazard and imperfect competition interact. Greater competition can reduce welfare by tightening incentive constraints, contradicting the standard prediction that more traders improve efficiency. Using trade data from Canadian exchanges, I show that the model can be used to assess the relative importance of moral hazard and its interaction with competition in practice, highlighting the empirical relevance of the theoretical welfare results.

Keywords: Moral hazard, financial intermediation, imperfect competition, market microstructure, equity trading

JEL: D4, D44, D61, D8, D82, G1, G24, L1

*Correspondence to Milena Wittwer, Columbia University, NBER, and CEPR: mw3941@columbia.edu. I thank Connor Breck, Carel Chok, Costanza Didonna, Matthew Hagerty, Harrison Lynch, Brendan McLaughlin, Daniel Smith, and Joseph Wagner for excellent research assistance. I thank Pavel Zryumov for his valuable discussions; in addition to Jason Allen, Ana Babus, Eric Budish, Piotr Dworczak, William Fuchs, Simon Gervais, Sergei Glebkin, Zhiguo He, Joshua Mollner, Daniel Neuhann, Christian Opp, Cecilia Parlatore, Indira Puri, Uday Rajan, Michael Sockin, Adi Sunderam, and all participants of the following conferences: 2025 Fall FTG meeting, the 2025 Chicago Market Design Conference, 2025 Germans Abroad Christmas Meeting, and 2025 Workshop for Assistant Professors in Finance.

1 Introduction

Most theoretical work on moral hazard builds on the principal–agent tradition, following Ross (1973), Mirrlees (1976), Holmström (1979), and Holmström and Milgrom (1987), and studies incentives within bilateral relationships. In financial markets, intermediaries (agents) invest on behalf of clients (principals), giving rise to moral hazard concerns (He and Krishnamurthy, 2012, 2013). Yet intermediaries do not operate in isolation. When trading assets, they compete with one another in centralized markets such as exchanges (Kyle, 1989; Vayanos, 1999; Vives, 2011; Rostek and Weretka, 2012). While moral hazard and imperfect competition are each well understood in isolation, we lack a framework to understand their interaction in centrally cleared asset markets.

In this paper, I develop a model that unifies moral hazard and imperfect competition and assess its empirical relevance using data from equity exchanges. Because the two frictions interact, stronger competition can tighten incentive constraints and reduce welfare. This highlights the need for coordinated competition and conduct regulation to avoid unintended consequences.

To motivate and quantify moral hazard and imperfect competition, I use proprietary trade-level data, covering roughly 60% of Canadian equity trading, from two exchanges listing the same securities. The data identify the trading dealer (firm) and whether trades are executed for the dealer's own account or on behalf of a client—granularity rarely available. I show that (holding the dealer, security, and time of execution fixed) prices are systematically worse on the less competitive exchange. Further, client trades receive consistently worse prices than proprietary trades, consistent with moral hazard.

Guided by these empirical patterns, I develop a model that incorporates both frictions. N clients seek exposure to a risky asset with normally distributed returns but lack direct market access. Each client delegates trading to an intermediary, hereafter, referred to as trader. Clients and traders have CARA preferences over wealth, potentially with different risk tolerances. The game unfolds over four periods. First, clients simultaneously offer contracts consisting of a lump-sum payment and a share of realized investment wealth. Second, traders decide whether to accept their contract and, conditional on acceptance, whether to exert effort when acquiring information about the expected asset return relevant for their client. Effort yields a precise signal, whereas shirking yields a private benefit but a noisy signal. Third, traders compete in a centralized market by submitting (net) demand schedules, and the market clears at a single price where aggregate demand equals supply. Fourth, returns and transactions realize.

I characterize a symmetric equilibrium in which clients offer the same contract, traders exert effort, and submit symmetric (linear) demand functions. Under imperfect competition, traders internalize their price impact and strategically scale back demand to improve their terms of trade (Kyle, 1989; Ausubel et al., 2014). Moral hazard, instead, shapes contract choice: larger private benefits from shirking require allocating traders a greater share of wealth to induce effort (Holmström and Tirole, 1997; He and Krishnamurthy, 2012). This feeds back into trading by effectively decreasing traders' risk tolerance and the positions they are willing to take for a given expected return. When incentive constraints are slack contracts coincide with the first best under perfect competition; when they bind contracts depend on the number of traders, and thus on competition, creating an interaction between frictions.

To clarify how these frictions interact, I study whether stronger competition, measured by the number of traders, mitigates or exacerbates moral hazard, as reflected in the required contract wealth shares (that traders need to obtain to exert effort). Three channels drive this interaction: reduced price impact, increased price predictability, and changes in the temptation to shirk driven by the information loss from shirking relative to effort. Together, these effects generate two contract regimes. When shirking noise is sufficiently high, contract shares (weakly) decline in the number of traders, indicating that competition mitigates moral hazard. When shirking noise is low, contract shares instead (weakly) rise, implying that competition amplifies moral hazard.

Whether increased competition through trader entry raises or reduces welfare depends on the prevailing contract regime. When the incentive constraint is slack and contracts do not respond to entry, welfare rises for the standard reason in trading models following Kyle (1989); Vayanos (1999); Vives (2011): improved risk sharing. When contracts instead adjust with trader entry, an additional indirect welfare channel emerges through changes in the equilibrium contract. When the contract increases with the number of traders, this indirect effect (which reflects weaker risk sharing and less efficient use of information per unit of risk) can dominate the direct effect, leading welfare to decline as competition intensifies.

I next illustrate how the model can be used to assess the relative importance of moral hazard and imperfect competition with trade-data (from the Canadian exchanges). Identifying moral hazard is challenging because incentives and contractual arrangements are rarely observed directly and, when they are observable, reflect equilibrium outcomes (Einav et al., 2021). In financial markets, client–intermediary contracts are typically proprietary, institution-specific, and not publicly disclosed, making systematic collection

of contract-level data largely infeasible—especially in settings with competitive market clearing among all active market participants. My approach is to infer the share of wealth traders must receive to exert effort given trade-data. Because trading activity is observable to regulators even when contracts are private, this approach provides a feasible first pass, albeit one that relies on relatively strong assumptions. To better approximate the data-generating process, I introduce a small number of targeted model adjustments, most notably approximating continuous trading with frequent batch auctions (Budish et al., 2015) for subsets of trades, such as client-account transactions on a given exchange.

I estimate two parameters that govern the severity of the two frictions in equilibrium: traders' risk tolerance and the contract share. The estimation is guided by the motivating empirical patterns: Price differences for proprietary trades across exchanges with different levels of competition identify traders' risk tolerance, while systematic price differences between client and proprietary trades pin down the contract. The resulting estimate of traders' risk tolerance is noisy but broadly consistent with existing evidence (Wittwer and Allen, 2026). Estimated contract shares range between 0.6 and 0.8 across specifications, implying that intermediaries retain well over half of the wealth generated on behalf of their clients, consistent with evidence from asset management, hedge funds, and private equity (Berk and van Binsbergen, 2015; Fama and French, 2010; French, 2008; Agarwal et al., 2009; Lan et al., 2013; Metrick and Yasuda, 2010).

To analyze the model's central implication—that contract terms vary with the number of traders when moral hazard interacts with imperfect competition—I estimate contracts separately for days with different levels of trader participation. I find that contracts are more favorable to intermediaries as the number of active traders rises. While this pattern may partly reflect time-varying market conditions or other confounding factors that trade-level data alone cannot fully address, it is consistent with the model's unfavorable-contract regime in which greater competition tightens moral-hazard frictions. In this regime, increased competition can reduce welfare. Empirically, when clients are less risk-tolerant than dealers, adding an additional trader reduces welfare by more than eliminating moral hazard altogether would increase it.

Taken together, my findings highlight that moral hazard and imperfect competition are intertwined: the efficiency of asset-market outcomes depends jointly on the severity of incentive problems and on competition. Policies that promote competition in asset markets can therefore have unintended consequences when incentive problems are severe, as additional entry may tighten incentive constraints and lower welfare. Effective regulation requires a coordinated approach.

Literature. My main contribution is to show how moral hazard and imperfect competition interact in a trading environment and jointly shape welfare. In doing so, I connect to several strands of literature.

At a high level, my paper relates to the principal–agent literature—following Ross (1973), Mirrlees (1976), Holmström (1979) and Holmström and Milgrom (1987). In my setting, the clients are the principals and the traders the agents. In addition, because traders decide whether to acquire information, my model also connects to the costly information acquisition literature—dating back to Stigler (1961), and Grossman and Stiglitz (1980). Lastly, the paper relates to the financial contracting literature, which applies these ideas to settings such as delegated portfolio management (e.g., Bhattacharya and Pfleiderer, 1985; Starks, 1987; Carpenter, 2000; Buffa et al., 2022) and delegated asset management with moral hazard (e.g., DeMarzo and Fishman, 2007; Di Tella and Sannikov, 2021).

A common feature of these three strands of literature is their micro-level focus. Within the principal–agent framework, this typically means analyzing incentives and contract design between a single principal and agent while abstracting from interactions among multiple agents or principals. Although some papers, such as Garcia and Vanden (2009), Buffa et al. (2022), and Kashyap et al. (2023), develop general-equilibrium models, they assume frictionless—and thus perfectly competitive—asset markets.

I instead model strategic interactions among traders in the spirit of Wilson (1979) and Kyle (1989); see Rostek and Yoon (2025) for a literature overview. This allows me to study how imperfect competition and moral hazard interact and jointly shape equilibrium contracts and welfare. This also distinguishes my work from Kyle et al. (2011), who model one informed agent trading against perfectly competitive market makers, and from Baldauf et al. (2024), who analyze contracting under moral hazard with exogenous price impact.[1]

The paper also relates to the banking and corporate finance literature that studies the interaction between imperfect competition and moral hazard. In banking, greater competition is often argued to compress margins and weaken incentives, attracting riskier borrowers and thereby undermining financial stability (Keeley, 1990; Crawford et al., 2018).[2] In corporate finance, product-market competition shapes managerial effort and

[1] More remotely related are papers that study trading and price formation under different frictions but share modeling features with elements of my framework. Examples include Glode and Opp (2020), Chen et al. (2025), and Wittwer and Allen (2026).

[2] Building on Stiglitz and Weiss (1981), subsequent work shows that this relationship is not monotonic: competition can both discipline and destabilize banks, depending on how it affects

incentive design, with competition either mitigating or exacerbating moral hazard depending on its effect on rents and performance sensitivity (Schmidt, 1997; Scharfstein, 1988; Raith, 2003).

My central welfare result—that competition can either mitigate or exacerbate moral hazard in a non-monotonic way—parallels existing findings but arises through different mechanisms. Whereas prior models capture competition through Bertrand or Cournot interactions over prices or quantities, traders in my model compete through demand (or supply) functions, a formulation better suited to trading environments with endogenous price formation. As a result, welfare losses arise from distortions in traders' information–risk trade-off when acting on private information—a channel absent from the existing literature.

2 Motivating empirical evidence

To motivate my focus on imperfect competition and moral hazard in a trading environment, and to illustrate the empirical variation that later informs the model, I draw on evidence from the Canadian stock market. The goal is not to provide causal evidence, and, given that the paper's main contribution is theoretical, I keep the empirical discussions brief. Appendix B.1 provides a more detailed discussion.

Institutional environment. The Canadian equity market provides a useful setting to study how moral hazard and imperfect competition interact in asset pricing. Unlike the United States, Canada—like other major markets such as the European Union and the United Kingdom—does not enforce a strict National Best Bid and Offer (NBBO) regime, granting dealers greater discretion in trade execution.[3] As a result, prices and allocations can vary more strongly with local trading conditions, creating scope for both moral hazard and imperfect competition. Moreover, because Canadian regulation requires equity trades to be executed on a regulated marketplace, exchange trading is representative of a typical equity transaction in Canada.

I focus on the Toronto Stock Exchange (TSX)—the primary exchange—and Alpha, its smaller counterpart trading the same securities. Combined, they represent about 60% of

lending incentives (Boyd and Nicolo, 2005; Martinez-Miera and Rapullo, 2010).

[3]A binding NBBO constrains quoted prices but does not eliminate dealer discretion over execution. Even under an NBBO regime, dealers retain control over order routing, timing, venue choice, queue position, and execution speed, and face heterogeneous local competitive conditions across venues. As a result, both moral hazard and imperfect competition can affect allocations and effective execution outcomes despite compliance with NBBO price protection.

total market activity. Both venues share common microstructure features of modern stock markets. Most trading occurs continuously throughout the day on centralized electronic limit order books, creating a relatively competitive trading environment. Further, only member firms—typically banks, large broker-dealers, and asset managers—are permitted to submit orders directly. All other investors, including large institutional investors, access the market through these members, which actively execute trades both for their own account and on behalf of clients.

In handling client orders, member firms retain considerable discretion. Although their execution decisions are subject to principles-based best-execution requirements under the Canadian Investment Regulatory Organization (CIRO) and the supervisory framework of the Universal Market Integrity Rules (UMIR), these rules largely rely on firms to design, implement, and monitor their own execution policies and procedures.[4]

Data. My data cover all regular trades on TSX and Alpha in 2021 during the continuous trading session (9:30 am to 4:00 pm). For each trade, I observe the time of execution, the traded security (the 'symbol'), trade size, price, and the national best bid and offer (prior trade execution). I also know whether the trade is active—meaning it immediately executes against resting orders in the limit order book and thus consumes liquidity—or passive; throughout the analysis, I focus on active trades.

The data have two features that make them well suited to studying the interaction between imperfect competition and moral hazard. First, I observe the exchange member firm executing each trade, which I refer to as the 'trader', recognizing that multiple individuals may operate under a single firm identifier. Second, I observe whether a trade is executed for the member's own account or on behalf of a client, a distinction that is central to analyzing moral hazard.[5]

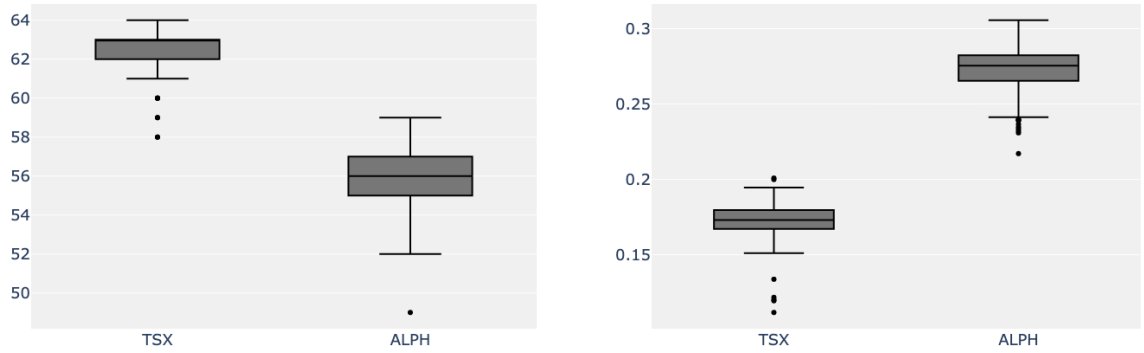
Imperfect competition. To provide suggestive evidence that imperfect competition matters for stock prices, I exploit variation across TSX—the primary Canadian exchange—and Alpha—its smaller counterpart—that list the same set of stocks. Both exchanges are well established, have operated for over a decade, are owned by the same parent company (TMX), and operate under similar market rules and regulatory requirements. TThis shared institutional structure rules out many—though not all—confounding factors un-

[4]Despite this, many institutional investors choose not to become exchange members, because membership is relatively costly: firms must be affiliated with CIRO, have access to the clearing system, and pay a C$65,000 entry fee plus monthly charges.

[5]Other account types exist but are relatively rare and are therefore excluded from most of the analysis.

Figure 1: Market concentration is higher on Alpha

(a) Number of active traders (b) Daily HHI



Notes: Figure 1a shows the distribution of the number of active traders on a day for the TSX and Alpha exchange. Figure 1b shows the distribution of daily HHI indices (which is between 0 and 1) on TSX and Alpha. The $HHI_t = \sum_i (s_{ti})^2$ is computed using the fraction of total daily trade value executed by a dealer i on day t as $s_{ti} \in [0, 1]$.

Table 1: Prices are worse on the more concentrated exchange (Alpha)

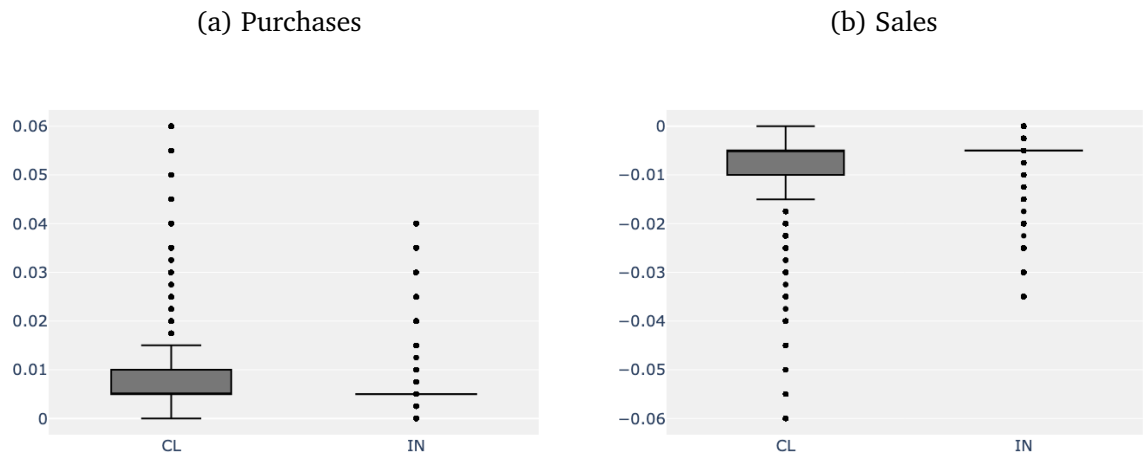
	(Buy)		(Sell)	
Alpha	0.004***	(0.000)	-0.004***	(0.000)
quantity	0.154***	(0.011)	-0.115***	(0.018)
N	136,715,587		131,862,280	

Notes: Table 1 reports estimates from regressions of transaction prices on an indicator for the Alpha exchange (Alpha) and trade size (quantity), controlling for symbol × trade date × hour × trader fixed effects. Appendix Table A2 reports robustness checks. Column (Buy) uses asset purchases, while column (Sell) uses sales. Prices are measured in Canadian dollars, and quantities in millions of shares. In both specifications, the R^2 is close to one due to the rich set of fixed effects, and standard errors are clustered at the daily level. For comparison, the median trade quantity is 0.0001 million C$, the median price is about 22.5C$, the median bid-ask spread using the national best bid and offer is 0.01C$.

related to competition (discussed in Appendix B.1) that could otherwise explain price differences.

As the smaller of the two exchanges, Alpha operates with systematically higher market concentration than TSX (see Figure 1) and exhibits worse execution prices (see Table 1). Controlling for trade size as well as trader–stock–day–hour fixed effects, traders buying on Alpha pay prices that are on average about C$0.004 higher—roughly 2 basis points of the median stock price or 40% of the median bid-ask spread—and conversely receive lower prices when selling. The regression results further confirm early empirical studies

Figure 2: Clients trade at worse prices than traders for their own account—raw data



Notes: These boxplots show the distribution of the difference between the trade-price and the mid-price before the trade executes, excluding upper and lower 2.5% quantiles, using trades on TSX from July 2021 for client (CL) and own inventory (IN) accounts. The mid-price is the average of the national best bid and ask before the trade.

on the NYSE and regional U.S. exchanges (such as Bernhardt et al. (2005)), according to which larger trades are executed at worse prices. This suggests that traders have price impact—a key feature of the model.

Moral hazard. Identifying moral hazard in financial markets is challenging because incentives and contractual arrangements are rarely observed. I provide suggestive evidence based on the idea that, when moral-hazard frictions exist between traders and their clients, trades executed on behalf of clients should occur at systematically worse prices than trades for a trader's own account.

Figure 2 and Table 2 show that, on average, a client-account trade executed on the same exchange, for the same stock, and within the same hour and day occurs at a price about C$0.005 worse—roughly 2 basis points of the median stock price or 50% of the median bid-ask spread—after controlling for trade size. To my knowledge, this type of evidence has not previously been documented, as trade accounts are usually unobserved.

In practice, several forces consistent with moral hazard could lead to worse execution for client accounts. Traders may simply pass through client orders without optimizing execution, or exert less effort in trading on behalf of client. In both cases, the trader does not exert the same effort to improve the client's performance. My model captures these forces in a stylized way: prices are worse when the trader invests on the client's behalf because the trader captures part of the surplus as compensation, pushing prices up when

Table 2: Clients trade at worse prices than traders for their own account

	(Buy)		(Sell)	
client-account	0.005***	(0.000)	-0.005***	(0.000)
quantity	0.156***	(0.014)	-0.114***	(0.018)
N	133,615,215		128,407,743	

Notes: Table 2 reports estimates from regressions of transaction prices on an indicator for client-account trades, traded quantity, and symbol × trade date × hour × exchange fixed effects. Appendix Table A3 reports robustness checks. I do not further interact these fixed effects with trader identifiers because there are insufficient observations of traders operating across both account types and exchanges for the same symbol within the same hour. Own-account trades are the baseline. Column (Buy) uses asset purchases, while column (Sell) uses sales. Prices are measured in Canadian dollars, and quantities in millions of of shares. In both specifications, the R^2 is close to one due to the rich set of fixed effects, and standard errors are clustered at the daily level. For comparison, the median trade quantity is 0.0001 million C$, the median price is about 22.5C$, the median bid-ask spread using the national best bid and offer is 0.01C$.

investors buy and down when they sell.

Take away. The empirical patterns motivate a model that jointly captures imperfect competition in trading and moral-hazard frictions. The goal is to understand how these frictions interact, how they shape welfare, and under what conditions one dominates the other. To retain tractability, trading occurs in a centralized market that clears at a single price, following Kyle (1989). The model focuses on client-directed investment through intermediaries, abstracting from proprietary trading; this allows me to isolate the interaction between incentives and competition that is central to the empirical analysis.

3 Model with imperfect competition and moral hazard

In the baseline model, traders compete to acquire assets on behalf of clients, consistent with the intermediary asset-pricing literature's focus on investment; extensions to trading are straightforward. I first present the model, then discuss the key tractability assumptions, and finally analyze the interaction between moral hazard and imperfect competition. Proofs are provided in Appendix A.4.

Setting. A continuum of clients seek to invest in one risky asset of fixed supply $A > 0$, but have no direct access to the asset market. Clients are partitioned into $N > 2$ groups, and each group—for simplicity referred to as 'client'—wants to hire a trader to invest on their behalf.

The game unfolds over four periods. In period 1, each client i offers a trader i a two-

part contract, consisting of a fixed amount $K_i \geq 0$, and a fraction $\phi_i \geq 0$ of wealth ω_i that will generated from their investment at the end of the game. Thus, when accepting the job the trader's total compensation is $K_i + \phi_i \omega_i$; when declining it, the trader receives a reservation utility of $\overline{U}$.

Wealth results from purchasing amount a_i at price p of the asset, which generates a per-unit gross return of R_i:

$$\omega_i = a_i(R_i - p).$$

Returns are normally distributed, $R_i \mid \nu_i \sim \mathcal{N}(\nu_i, \Sigma)$, with known covariance matrix $\Sigma > 0$. The mean ν_i captures individual market forecasts, balance-sheet exposures, funding costs, or hedging needs that shape appetite for additional asset holdings. These means are drawn independently from a common prior, $\nu_i \sim \mathcal{N}(\nu, \Sigma_\nu)$, with $\nu \in \mathbb{R}$ and $\Sigma_\nu > 0$. Uncertainty about expected returns is thus driven by idiosyncratic factors, with any common component being absorbed by the known common prior.[6]

In period 2, traders form beliefs about their (and their client's) expected return, ν_i. They can exert effort ($s_i = 0$), in which case they observe $\nu_i \in \mathbb{R}$, and therefore know the conditional return distribution. Alternatively, they can shirk ($s_i = 1$) which implies that they only observe the return distribution with noise. Conditional on their effort-decision, they see the following private return signal in period 3:

$$\mu_i(s_i) = \begin{cases} \nu_i & \text{when trader exerts effort, } s_i = 0 \\ \nu_i + \epsilon_i \text{ with } \epsilon_i \overset{iid}{\sim} \mathcal{N}(0, \sigma_\epsilon) & \text{when trader shirks, } s_i = 1, \end{cases}$$

where $\sigma_\epsilon > 0$. Note that unlike in a standard principal-agent setting (dating back to Ross (1973), and Holmström and Milgrom (1987)) exerting effort does not create larger output, but rather changes the information available to the trader in the subsequent trading stage. In this regard, my model is more similar to the literature on costly information acquisition (dating back to Stigler (1961), and Grossman and Stiglitz (1980)).

Shirking generates a private benefit to the trader, $b \geq 0$, and a loss for the client $\Delta \geq 0$, which is sufficiently high that it is never optimal for the client to allow shirking. In practice, the benefit from shirking may reflect the value of leisure from working less, while the client's loss reflects reduced attention to non-modeled responsibilities, such as reporting, communication, or compliance duties, which can impose administrative or regulatory costs on the client.

[6] Allowing for correlated means or an unknown common prior is feasible but adds algebraic complexity with limited economic insight.

Given their private signal, traders compete against one another to purchase shares of the asset. The asset market clears through a uniform price auction—a common approach to modeling market clearing in trading environments (e.g., Wilson, 1979; Klemperer and Meyer, 1989; Kyle, 1989; Vayanos, 1999; Vives, 2011; Rostek and Yoon, 2021). Each trader submits a decreasing and twice continuous demand function

$$a_i(\cdot, \mu_i(s_i)) : \mathbb{R}_+ \to \mathbb{R}_+$$

that maps prices to quantities, conditional on their private signal (and their effort-decision). The submitted demand functions are aggregated, and the market clears at prices where aggregate demand equals supply.[7] From the perspective of trader i who takes choices of all other traders as given:

$$p^c(s_i) : \sum_i a_i(p^c(s_i), \mu_i(s_i)) = A. \tag{1}$$

There are two possible clearing prices, depending on whether the trader exerts effort or shirks. To highlight this, I label the clearing price as $p^c(s_i)$. Note that this notation omits the fact that the clearing price depends on all contracts and effort choices by all traders, not just trader i.

Clients and traders obtain CARA utility over the wealth they hold at the end of the game. Clients have risk-tolerance $\rho_H > 0$, and traders $\rho_I \geq \rho_H > 0$. Specifically, a trader who chooses their demand function expects to obtain the following utility when choosing demand $a^c(s_i) = a_i(p, \mu_i(s_i))$ at clearing price $p = p^c(s_i)$:

$$\mathbb{E}\left[1 - \exp\left(-\frac{1}{\rho_I}\omega_I^c(s_i;\phi_i)\right)\middle|\mu_i(s_i)\right], \tag{2}$$

$$\text{with } \omega_I^c(s_i;\phi_i) = \phi_i\Big\{a^c(s_i)(R_i - p^c(s_i))\Big\} + s_i b + K_i. \tag{3}$$

Given that all random variables are normally distributed, asset returns conditional on return signals are also normally distributed: $R_i|\mu_i(s_i) \sim \mathcal{N}(\mathbb{E}[R_i \mid \mu_i(s_i)], Var[R_i \mid \mu_i(s_i)])$. Under normal return distributions, maximizing (2) is equivalent to maximizing the cer-

[7] In case of excess demand—which never happens in equilibrium—traders are rationed pro-rata on the margin.

tainty equivalent $\mathcal{U}(s_i;\phi_i)$ point-wise at every clearing price $p^c(s_i)$:

$$\mathcal{U}(s_i;\phi_i)+K_i\ , \tag{4}$$

$$\begin{aligned}\text{where } \mathcal{U}(s_i;\phi_i) &= \phi_i\Big(\mathbb{E}[R_i \mid \mu_i(s_i)]-p^c(s_i)\Big)a_i^c(s_i)\\ &\quad -\phi_i^2\frac{1}{2\rho_I}a_i^c(s_i)Var[R_i \mid \mu_i(s_i)]a_i^c(s_i)+s_i b.\end{aligned} \tag{5}$$

When the trader chooses their effort-level, s_i, they do not know their private signal, $\mu_i(s_i)$. I assume that they take an expectation, integrating out the uncertainty over their signal.

Similarly, each client maximizes their expected certainty equivalent, integrating out the uncertainty about the trader's signal, when choosing the contract. The client's certainty equivalent is defined analogously to the trader's certainty equivalent with four differences: the client's risk tolerance is ρ_H rather than ρ_I, the client obtains fraction $1-\phi_i$ rather than ϕ_i of wealth generated by the asset, and suffers a loss, Δ, when the trader shirks, rather than a benefit, b:

$$\mathcal{W}(\phi_i;s_i)-K_i, \tag{6}$$

$$\begin{aligned}\text{where } \mathcal{W}(\phi_i;s_i) =&(1-\phi_i)\Big(\mathbb{E}[R_i \mid \mu_i(s_i)]-p^c(s_i)\Big)a_i^c(s_i)\\ &-(1-\phi_i)^2\frac{1}{2\rho_H}a_i^c(s_i)Var[R_i \mid \mu_i(s_i)]a_i^c(s_i)-\Delta s_i.\end{aligned} \tag{7}$$

Summarizing, the sequence of events is as follows:

(1) Clients choose their contract, $\{K_i,\phi_i\}$, simultaneously and independently. Each trader accepts their contract or walks away, in which case the game ends.

(2) Traders decide simultaneously and independently whether to exert effort: $s_i\in\{0,1\}$.

(3) Each trader observe their private expected return signal, depending on their chosen effort-level, $\mu_i(s_i)$, and compete to buy assets in a uniform-price auction.

(4) The assets pay their returns, and all transactions take place.

While I introduced the model in a buy-side setting, it is straightforward and common in the finance literature to modify the model to allow for selling or trading. One can simply allow the exogenous net supply to take any real value, $A\in\mathbb{R}$, and interpret traders' submitted schedules as net-demand curves $a_i:\mathbb{R}\to\mathbb{R}$, where positive quantities correspond to purchases and negative ones to sales.

Equilibrium. I focus on symmetric Bayesian Nash equilibria (BNE), in which traders and clients choose identical strategies when observing the same signal. The key equilibrium objects are the contracts $\{K^*, \phi^*\}$, the trader's shirking decisions, $s^* \in \{0,1\}$, and their demand functions, $a^*(\cdot, \mu_i(s_i))$ for $s_i \in \{0,1\}$ for all i.

Proposition 1. *In a symmetric equilibrium:*

(i) All clients offer the following contract $\{K^*, \phi^*\}$:

$$\phi^* = \begin{cases} \Phi(N) = \frac{1}{|A|}\sqrt{\frac{\beta(N)+b\gamma(N)}{\alpha(N)}} & \text{when } \frac{\rho_I}{\rho_I+\rho_H} \leq \Phi(N) \\ \frac{\rho_I}{\rho_I+\rho_H} & \text{otherwise,} \end{cases} \tag{8}$$

and

$$K^* = \bar{U} - \rho_I \frac{N-2}{2(N-1)} \frac{\Sigma_v}{\Sigma} - \frac{A^2\Sigma}{2N(N-2)\rho_I}(\phi^*)^2, \text{ where} \tag{9}$$

$$\alpha(N) = (-1+N)\Sigma^2\left(\tilde{\Sigma} - \Sigma\right) > 0,$$

$$\beta(N) = -(((N-2)N\rho_I^2\sigma_\epsilon\Sigma_v((N-1)\Sigma + (N-2)\Sigma_v))(\sigma_\epsilon + \Sigma_v)^{-1} < 0,$$

and $\gamma(N) = 2(-1+N)N\rho_I\Sigma(2\Sigma + (-2+N)\tilde{\Sigma}) > 0$, *given* $\tilde{\Sigma} = Var(R_i|\mu_i(1))$.

(ii) All traders exert effort, $s^* = 0$, and submit the following demand function:

$$a^*(p, \mu_i(s^*)) = \rho_I(\Lambda^*\rho_I + \phi^* Var[R_i \mid \mu_i(s^*)])^{-1}(\mathbb{E}[R_i \mid \mu_i(s^*)] - p), \tag{10}$$

$$\text{with } \Lambda^* = \frac{\phi^*}{\rho_I(N-2)}\Sigma. \tag{11}$$

where we can insert $Var[R_i \mid \mu_i(s^*)] = \Sigma$ and $\mathbb{E}[R_i \mid \mu_i(s^*)] = v_i$.

This equilibrium exists when both the trader's benefit and the client's loss from shirking are sufficiently large, as characterized in Propositions 3 in Appendix A.1. A sufficiently high trader benefit ($\beta(N) + b\gamma(N) > 0$) implies that the trader's incentive constraint binds for large contract shares. When the trader's benefit from shirking is too small, effort is always optimal for the trader, and the equilibrium contract collapses to $\frac{\rho_I}{\rho_I+\rho_H}$. A sufficiently large client loss rules out profitable deviations in which a single client tolerates shirking while others induce effort. The equilibrium is unique within the class of symmetric equilibria with linear demand functions, a common refinement in the related literature (Rostek and Yoon, 2025) (Proposition 4).

To derive an intuition for the equilibrium consider one client-trader pair, i, and fix the behavior of all clients and traders at the symmetric equilibrium. Consider first the con-

tract choice. Since the client aims at incentivizing effort (to avoid the loss from shirking, Δ), they choose contract $\{\phi_i, K_i\}$ to maximize:

$$\mathbb{E}[\mathcal{W}(\phi_i; 0] - K_i \text{ subject to (IR) and (IC)}, \tag{M}$$

anticipating that the trader will exert effort when

$$\mathbb{E}\left[\mathcal{U}(0; \phi_i)\right] \geq \mathbb{E}\left[\mathcal{U}(1; \phi_i)\right], \tag{IC}$$

and accept the job offer when

$$\mathbb{E}\left[\mathcal{U}(0; \phi_i)\right] + K_i \geq \bar{U}. \tag{IR}$$

Since the trader's incentive constraint is independent of the fixed part of the contract, K_i, it is optimal to set K_i as low as possible, so as to still incentivize the trader to participate. This is the case when

$$K_i \equiv K(\phi_i; 0) = \bar{U} - \mathbb{E}\left[\mathcal{U}(0; \phi_i)\right]. \tag{IR'}$$

Note that $K(\phi_i; 0)$ depends on ϕ_i, because ϕ_i affects the trader's demand function in the trading game. Inserting the binding participation constraint (IR') into the client's objective function (12), the optimization problem simplifies to:

$$\max_{\phi_i} \left\{\mathbb{E}[\mathcal{W}(\phi_i; 0] - K(\phi_i; 0)\right\} = \max_{\phi_i} \left\{\mathbb{E}[\mathcal{W}(\phi_i; 0] + \mathbb{E}[\mathcal{U}(0; \phi_i] - \bar{U}\right\} \tag{12}$$

subject to (IC). There are two solutions, depending on whether the constraint binds or is slack. When it is slack, $\phi_i = \frac{\rho_I}{\rho_I + \rho_H}$, the trader perfectly internalizes that they invest for a client, and correctly chooses demand, depending on their differing degree of risk-tolerance, i.e., maximizes risk-sharing. When the (IC) constraint binds, it pins down the client's best response to other clients choosing ϕ^*. In a symmetric equilibrium, each client i chooses the same contract as all others. Solving for the fixed-point yields a unique positive equilibrium solution, denoted by $\Phi(N)$.

In both the constrained and unconstrained case, the trader exerts effort. They choose demand, $a_i = a_i(p, \mu_i(s_i))$, by point-wise maximizing their certainty equivalent (5) at each possible clearing price. The necessary condition for quantity a_i at the fixed market clearing price p and both $s_i \in \{0, 1\}$ is

$$\mathbb{E}[R_i \mid \mu_i(s_i)] - \phi_i \frac{1}{\rho_I} Var[R_i \mid \mu_i(s_i)] a_i = p + \Lambda_i a_i, \text{ where } \Lambda_i = \left(\frac{dp}{da_i}\right). \tag{13}$$

Λ_i measures the price impact of the trader's own demand on the clearing price. It is the amount by which the market clearing price moves when trader i demands marginally more than quantity a_i. Since all other traders j but i submit the equilibrium guess (11), price impact is

$$\Lambda_i = \frac{\phi^*}{\rho_I(N-2)}\Sigma. \tag{14}$$

The necessary conditions are sufficient, since Σ is positive semi-definite, $N > 2$ and $\phi_i > 0$. Therefore, equation (13) pins down the trader's best reply $a_i(p, \mu_i(s_i))$ for $s_i \in \{0, 1\}$.

Given the contracts, the trader's shirking decisions, and their demand functions, the equilibrium price (15), and the equilibrium allocations are then pinned down via market clearing.

Corollary 1. *In equilibrium, the market clears at price:*

$$p^* = \frac{1}{N}\sum_i \nu_i - \frac{\phi^*}{\rho_I}\Sigma\left(\frac{N-1}{N-2}\right)\frac{A}{N}, \tag{15}$$

and trader i purchases amount

$$a^*(p^*, \nu_i) = \frac{A}{N} + \frac{\rho_I}{\phi^*}\frac{N-2}{N-1}\left(\nu_i - \frac{1}{N}\sum_i \nu_i\right)\Sigma^{-1}, \tag{16}$$

where I used $\mu_i(0) = \nu_i$, $\mathbb{E}[R_i \mid \mu_i(0)] = \nu_i$ *and* $Var[R_i \mid \mu_i(0)] = \Sigma$.

Discussing tractability assumptions. The model relies on a set of assumptions that ensure tractability. Before analyzing the equilibrium and unpacking the underlying economic mechanisms, I discuss the key modeling choices that render the model tractable. The discussion focuses on features that are specific to this paper's contribution, rather than widely used assumptions in the related literatures (such as CARA utility, ex-ante identical traders, normally distributed returns, or linear equilibria).

First, clients offer simple, linear compensation contracts, $\{K_i, \phi_i\}$, consisting of a fixed fee and a constant share of the future wealth generated by the asset.[8] This con-

[8] In the classic principal–agent framework, where the agent's effort directly affects output, Holmström and Milgrom (1987) show that linear contracts are optimal when the agent has CARA preferences and output noise is normally distributed. My environment is more complex along several dimensions—including strategic interaction, price formation, and the indirect mapping from effort to outcomes—which makes solving for the optimal contract difficult. Rather than pursuing optimality in this richer setting, I take institutional features seriously and model contracts that mirror those used in practice.

tract resembles common institutional arrangements, combining a fixed base salary with a performance-based bonus ϕ_i that reflects average outcomes across repeated market clearings, and does not depend on one-time market outcomes. With outcome-contingent contracts traders' demand would no longer be linear in prices—a feature that is known to render uniform-price auctions intractable. However, the main insight that the equilibrium contract ϕ^* depends on the number of competing traders (or investors), once market clearing is taken into account, is robust across different ways of modeling contract choices, and market clearing.

Second, I assume that traders and clients choose their effort levels and contracts by maximizing their expected certainty equivalents—that is, the utility they would obtain from receiving a given level of wealth with certainty, conditional on the realized signal. In doing so, they integrate out the uncertainty that arises from not knowing the signal realization at the time of their decision.[9] This modeling choice preserves tractability and reflects the idea that traders and clients do not make a single, one-shot decision based on a unique signal. Instead, trading occurs repeatedly across many rounds, and decisions are likely to reflect average outcomes across these rounds, which the model captures through expectations.

Third, I restrict attention to equilibria in which clients choose contracts that incentivize effort by imposing a sufficiently large loss when the trader shirks. This assumption rules out deviations in which an individual client tolerates shirking while others induce effort, as well as symmetric equilibria sustained by shirking. It serves as a reduced-form way to capture reputational, institutional, and relational forces that discipline behavior in practice—such as loss of trust, mandate renegotiation, or reputational consequences—that are difficult to model explicitly but plausibly large relative to (one-shot) trading gains.

Fourth, I abstract from common-value concerns, meaning that all traders and clients share the same beliefs about the market-wide sentiment for the asset, ν, so prices convey no information about the value of the asset. Allowing for common values would introduce additional strategic considerations in the trading stage: traders would have an additional incentive to shade their bids to avoid overpaying when they are more optimistic than the market—the classic winner's curse. Crucially, however, effort and contract choices occur before trading and would only be affected insofar as certainty equivalents

[9]Formally, maximizing the expected certainty equivalent, where expectations are taken over signals, $\mathbb{E}[\mathcal{U}(s_i;\phi_i)]$, is not equivalent to maximizing expected CARA utility (2), where expectations are taken jointly over wealth and signals.

incorporate expectations over both the true value and traders' private signals. Thus, abstracting from common values helps isolate the key forces that govern the interaction between competition and moral hazard: a price-impact effect, capturing how traders strategically adjust their quantities, and an information effect, arising from the fact that contracts are determined before trading occurs. Both forces should also arise in a less tractable framework with common values, though their functional forms would differ.

Finally, the model features a single trading stage and abstracts from intertemporal considerations. Even when focusing solely on trading—and setting aside interactions with contract choices—dynamic auction environments are complex (e.g., Vayanos, 1999; Rosket and Weretka, 2015; Du and Zhu, 2017). A promising, though challenging, extension would be to build on this literature by studying a dynamic common-value environment in which prices convey information about asset values and clients can revise their contracts after observing market outcomes.

3.1 How imperfect competition and moral hazard affects choices

My goal is to analyze how the two frictions—imperfect competition and moral hazard—interact in affecting asset allocations, prices, and ultimately welfare.

Imperfect competition. In an imperfectly competitive market, traders have price impact (Kyle, 1989), as characterized (in my model) by expression (11) in equilibrium. Strategic traders take this into account when choosing how to trade: they deliberately reduce their demand for additional units to lower the market-clearing price, thereby improving the terms of their winning bids. This is known as strategic demand reduction (Ausubel et al., 2014). It reduces the market clearing price (compared to the price that would arise if traders were price-takers) and distorts allocations. In particular, traders with above-average return expectations to win too little, and those with below-average expectations to win too much compared to the efficient outcome.[10]

Moral hazard. Moral hazard frictions affect the contract choice, and through the flexible part of the contract, ϕ^*, the trading decisions. Since the fixed part of the contract, K^*, represents a transfer between the client and trader, it does not affect allocative efficiency. I focus on the flexible part, ϕ^*—which I will simply refer to as contract—, for the

[10] To see this formally, compare the market-clearing allocation of trader i when everyone is a price-taker, $a^o(p^o, \nu_i)$, to the equilibrium allocation with strategic traders for a fixed contract: $a^o(p^o, \nu_i) - a(p^*, \nu_i) = \frac{\rho_I}{\phi^*} \frac{1}{N-1} (\nu_i - \frac{1}{N} \sum_i \nu_i) \Sigma^{-1}$, where $\nu_i = \mathbb{E}[R_i \mid \mu_i(0)]$ and $\Sigma = Var(R_i \mid \mu_i(0))$. This expression is positive when ν_i exceeds the average and negative otherwise.

remainder of the paper.

When the benefit of shirking is sufficiently low so that the trader has no incentives to shirk, the contract, $\phi^* = \frac{\rho_I}{\rho_I+\rho_H}$, is unconstraint. This contract achieves the first best-solution when the market is perfectly competitive ($N \to \infty$). When the incentive constraint binds, which is the case when $\frac{\rho_I}{\rho_I+\rho_H}$ is too low, the contract is set to ensure that the trader exerts effort: $\phi^* = \Phi(N)$. While the expression is messy, the key property (we discuss below) is that it depends non-trivially on the number of traders, and therefore the degree of competition.

The contract, in turn, affects trading because it affects the trader's risk-return trade-off. Increasing the contract raises the trader's 'effective' risk-aversion $\frac{\phi^*}{\rho_I}$. Since traders weigh the extra profit they expect from the purchase against the additional risk it brings, traders are less comfortable taking on larger (risky) positions for a fixed expected return. As a result, the price (15) needs to decline to persuade them to absorb extra amounts of the asset, and price impact (11) increases. This ultimately affects asset allocations, and allocative efficiency. Traders with above-average return expectations purchase too little, and traders with below-average expectations too much when the constrained binds. The higher the contract, i.e., the worse the moral hazard friction, the larger this distortion.[11]

How the frictions interact. Intuitively, it is ambiguous whether stronger competition—measured by the number of traders—alleviates or amplifies moral hazard—captured by the size of the contract. This depends on three intertwined forces.

(i) Price-impact effect: The first force is that greater competition reduces each trader's individual price impact,

$$\Lambda^* = \frac{\phi^*}{\rho_I(N-2)}\Sigma, \tag{11}$$

altering their incentive to strategically scale back demand. To be willing to exert effort, traders need to obtain a larger fraction of wealth as soon as the incentive constraint binds.

The price-impact effect is related to, yet somewhat distinct from, the idea that increased competition compresses the rents traders can extract from market power, thereby raising the compensation required to sustain effort. In my trade model, market-power rents depend not only on price impact but also on the information–risk ratio, which measures how informative a trader's signal is relative to the residual uncertainty borne when investing, Σ_v/Σ. The influence of this ratio on market-power rents varies with the num-

[11] To see this formally, compare the market-clearing allocation of trader i when moral hazard is not binding, $a^*_{i,UC}$, to the allocation when the constrained is binding, $a^*_{i,C}$. This gives: $a^*_{i,C} - a^*_{i,UC} = [(\rho_I+\rho_H) - \frac{\rho_I}{\Phi(N)}]\frac{N-2}{N-1}(\frac{1}{N}\sum_i \nu_i - \nu_i)\Sigma^{-1}$.

ber of traders in a way not captured by the price-impact effect.[12]

(ii) Price-information effect: Adding more traders makes prices—and thus allocations—more predictable. In the limit with infinitely many traders, market outcomes become perfectly predictable. Because contracts and effort are chosen prior to trading, changes in predictability affect both decisions.[13] To highlight the role of this effect, I shut down the price-impact effect by assuming that traders behave as price-takers and deriving the resulting equilibrium in Proposition 5 in the appendix. The constrained contract in that equilibrium, $\tilde{\Phi}(N)$, shares similarities to the contract with strategic demand reduction, $\Phi(N)$, emphasizing the central role of the price-information effect in the interaction between moral hazard and imperfect competition.

(iii) Information-risk ratio effect: The third force reflects how tempting it is for a trader to shirk, which depends on how noisy the signal becomes when shirking relative to exerting effort, interacted with how many traders compete for the asset. To isolate this effect from the other two, I consider the limiting case as the number of traders N grows large—an environment approximating perfect competition, where prices are fully predictable and no trader strategically reduces demand.

Corollary 2. *Let* $\gamma_\infty(\sigma_\epsilon) = \frac{\rho_I \sigma_\epsilon \Sigma_\nu(\Sigma+\Sigma_\nu)}{2\Sigma(\sigma_\epsilon\Sigma_\nu+\Sigma(\sigma_\epsilon+\Sigma_\nu))}$*, and note that* $\gamma'_\infty(\sigma_\epsilon) > 0$*.*

(i) When σ_ϵ *is sufficiently low so that* $b > \gamma_\infty(\sigma_\epsilon)$*, there exists no finite contract that incentivizes traders to exert effort in a perfectly competitive market* $(N \to \infty)$*.*

(ii) When σ_ϵ *is sufficiently high so that* $b < \gamma_\infty(\sigma_\epsilon)$*, the incentive constraint never binds in a perfectly competitive market* $(N \to \infty)$*, so that* $\phi^* = \frac{\rho_I}{\rho_I+\rho_H}$*.*

There are two cases: a favorable and an unfavorable one. When the noise from shirking, σ_ϵ, is high relative to the private benefit from shirking, the contract is unconstrained and achieves the first best; when the noise is low, no finite contract can induce effort. To

[12] To show this formally, I consider two alternative measures of market-power rents. The first is from the perspective of an individual trader is to compare the expected payoff when acting strategically to the payoff when behaving as a price-taker while all other traders remain strategic: $\Lambda^* A^2 \phi^* (2(N-1)^2 N)^{-1} + (N-2)\rho_I (2(N-1)^3)^{-1} \Sigma_\nu / \Sigma$. The second measure compares the expected payoff of a trader (who exerts effort) when all traders behave strategically with the payoff when all traders behave as price-takers, holding contracts fixed: $\Lambda^* \phi^* (A/N)^2 - \rho_I (2N(N-1))^{-1} \Sigma_\nu / \Sigma$.

[13] To see this formally, consider a client i choosing a contract ϕ_i, taking all other clients' contracts as given at ϕ^*. As the number of traders N increases, the client can better predict the clearing price—and thus its market-clearing allocation—since, ceteris paribus, the variance of the clearing price declines in N: $Var(p^c(0)) = (2\phi^* + (N-2)\phi_i)^{-2}(\phi^{*2} + (\phi^* + (N-2)\phi_i)^2 (N-1)^{-1})\Sigma_\nu$.

see why, note that in the competitive limit the trader's expected payoff takes a familiar mean–variance structure. It is proportional to the information–risk ratio:

$$\frac{Var(\mathbb{E}[R_i|\mu_i(s_i)])}{Var(R_i|\mu_i(s_i))}.$$

A trader can be incentivized to exert effort if their expected payoff under effort is at least as high as their expected payoff when shirking. When the trader exerts effort, they expect the following payoff:

$$\lim_{N\to\infty} \mathbb{E}[\mathcal{U}(0;\phi_i)] = \frac{\rho_I}{2}\frac{Var(\mathbb{E}[R_i|\mu_i(0)])}{Var(R_i|\mu_i(0))} = \frac{\rho_I}{2}\frac{\Sigma_\nu}{\Sigma}. \tag{17}$$

When they shirk, they expect to get:

$$\lim_{N\to\infty} \mathbb{E}[\mathcal{U}(1;\phi_i)] = \frac{\rho_I}{2}\frac{Var(\mathbb{E}[R_i|\mu_i(1)])}{Var(R_i|\mu_i(1))} + b = \frac{\rho_I}{2}\frac{\Sigma_\nu^2(\Sigma_\nu+\sigma_\epsilon)^{-1}}{\Sigma+(\Sigma_\nu^{-1}+\sigma_\epsilon^{-1})^{-1}} + b. \tag{18}$$

The larger is the noise traders observe when shirking (higher σ_ϵ), the lower is their expected payoff from shirking. Greater noise weakens the signal's ability to predict returns while increasing the residual uncertainty that traders must bear. Both effects reduce the information–risk ratio under shirking and, in turn, lower the expected payoff.

With finitely many traders, the three forces— (i) the price-impact effect, (ii) the price-information effect, and (iiii) the information-risk ratio effect—jointly give rise to a non-trivial relationship between competition and moral hazard. To capture this relationship I analyze how the constraint contract, $\Phi(N)$, changes when we add an additional trader. When the trader needs to obtain a higher fraction of wealth to still exert effort, I say that moral hazard intensifies. When the opposite is true, moral hazard is relaxed.

Corollary 3. *Define* $\gamma(\sigma_\epsilon) \equiv \frac{\rho_I\sigma_\epsilon\Sigma_\nu(2(N-1)^3\Sigma+(N-2)(2+N(2N-3))\Sigma_\nu}{4(N-1)^2\Sigma(N\Sigma\sigma_\epsilon-\sigma_\epsilon\Sigma_\nu+N(\Sigma+\sigma_\epsilon)\Sigma_\nu)}$*; note that* $\gamma'(\sigma_\epsilon) > 0$.

(i) When σ_ϵ *is sufficiently low so that* $b > \gamma(\sigma_\epsilon)$, $\Phi(N)$ *increases in* N.

(ii) When σ_ϵ *is sufficiently high so that* $b < \gamma(\sigma_\epsilon)$, $\Phi(N)$ *decreases in* N.

Corollary 3 underscores the non-linear interaction between moral hazard and imperfect competition, foreshadowed in Corollary 2. When noise from shirking is small, increasing the number of traders worsens the moral hazard problem—the fraction of wealth the trader must obtain to exert effort increases in N (at the margin); otherwise it decreases. Given that the cutoff, $\gamma(\sigma_\epsilon)$ varies in N, however, the relationship is complex.

Figure 3 provides numerical examples of how the contract ϕ^* can behave as the number of traders increases. The examples are chosen to illustrate the two cases presented in Corollary 3, though these cases do not exhaust all possible shapes the contract may take. In the first example—corresponding to statement (i)— noise from shirking is low, $\sigma_\epsilon = 0.5$, and $\Phi(N)$ increases in N for large enough N. In the second example—corresponding to statement (ii)—noise from shirking is high, $\sigma_\epsilon = 100$. Then, $\Phi(N)$ decreases in N for large enough N, and eventually becomes unconstrained.[14]

The finding that contracts vary with the number of traders suggests that policies aimed at promoting competition in asset markets, for example, by lowering entry barriers for intermediaries, may inadvertently intensify moral-hazard problems. The corollary clarifies when this risk is most pronounced: when traders derive substantial private benefits from shirking (or, equivalently, when exerting effort is costly), or when returns can be predicted with relatively high precision even without exerting effort. Such conditions are more likely in markets where information acquisition is costly, such as corporate bonds, structured credit, or OTC derivatives, as well as in environments where returns are largely driven by coarse or public signals, for instance in highly liquid government bond or index markets, or during periods of elevated macroeconomic uncertainty when price movements are dominated by economy-wide rather than asset-specific factors.

3.2 How imperfect competition and moral hazard affect welfare

Intuitively, because allocations depend on the equilibrium contract—and the contract itself varies with the number of traders N—the welfare effects of increased competition should hinge on how the contract responds to N. If stronger competition tightens incentive constraints, welfare may decline even as markets become more competitive. To analyze whether this is the case, I compute the allocative efficiency of the equilibrium allocation:

$$W(\phi^*, N) = \sum_i W_i(\phi^*, N), \tag{19}$$

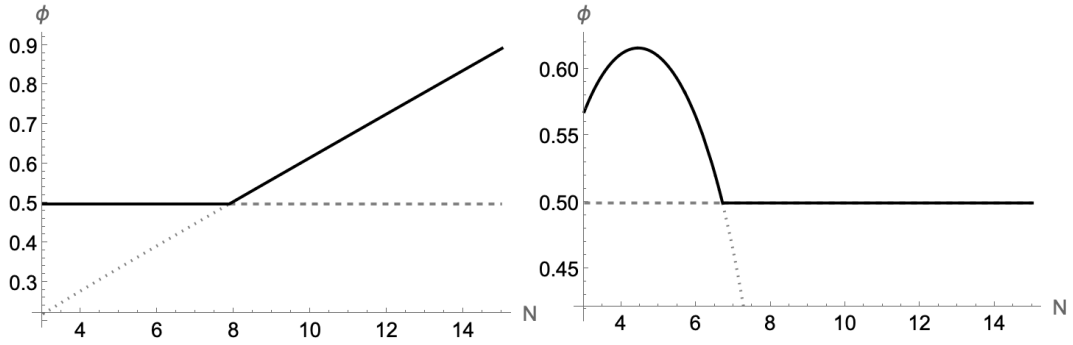
$$W_i(\phi^*, N) = a^*(p^*, \nu_i)\nu_i - \sum_{k\in\{h,i\}} \frac{1}{2}\phi_k^{*2}\frac{1}{\rho_k} a^*(p^*, \nu_i)\Sigma\, a^*(p^*, \nu_i), \tag{20}$$

where $\phi_h = 1 - \phi_i$. Following the market-design literature, I adopt an ex-ante perspective and define 'welfare' as expected allocative efficiency, $\mathbb{E}[W(\phi^*, N)]$. This definition

[14] Note that this example also illustrates the non-linearity of the contract, since the contract is increasing in N for small enough N.

Figure 3: Examples of how the flexible contract, ϕ^*, can change in N

(a) ϕ^*: Low noise from shirking (b) ϕ^*: High noise from shirking



Notes: Figure 3 plot the equilibrium contract ϕ^* (black line) as a function of the number of traders N for two numeric examples. The chosen contract is a piecewise function consisting of the unconstrained contract, $\frac{\rho_I}{\rho_I+\rho_H}$, the dashed gray line, and the constrained contract, $\Phi(N)$, the dotted gray line. In both examples, $\rho_I = \rho_H = 0.5, \Sigma = \Sigma_\nu = 1, b = 0.2$. In Figure 3a shirking-noise is low, $\sigma_\epsilon = 0.5$ and supply is large, $A = 10$, in Figure 3b it is high, $\sigma_\epsilon = 100$ and supply is small $A = 2$. Appendix Figure A1 shows the corresponding fixed part of the contract, K^*.

ensures that welfare comparisons do not depend on a particular realization of return expectations.

Welfare is shaped by three forces: how wealth is split between traders and clients, their relative risk tolerances, and how much useful information traders have relative to the underlying asset's risk (the information–risk ratio). To build intuition, I begin with two simple benchmarks.

In the first benchmark, a single trader ($N = 1$) purchases an asset with normally distributed returns, $\mathcal{N}(\nu, \Sigma)$, and keeps a share ϕ of the payoff. Welfare is

$$\mathbb{E}[W^{B1}(\phi, N)] = A\nu - \frac{1}{2}\sum_{k\in\{i,h\}} \frac{1}{\rho_k}\phi_k^2 \Sigma A^2 = A\nu - \beta(\phi, 1)\Sigma A^2,$$

where $\beta(\phi, N)$ is given by expression (23) in Proposition 2a. The first term reflects expected gains from holding the asset; the second captures risk-bearing costs, which depend on the wealth split and on risk tolerances.

In the second benchmark, $N > 2$ traders have private information about their expected asset return, ν_i, like in my model but they keep all assets ($\phi = 1$). Welfare is

$$\mathbb{E}[W^{B2}(1, N)] = A\nu - \frac{1}{2N\rho_I}\Sigma A^2 + \frac{\rho_I N(N-2)}{2(N-1)}\frac{\Sigma_\nu}{\Sigma} = A\nu - \beta(1, N)\Sigma A^2 + \alpha(1, N)\beta(1, N)\frac{\Sigma_\nu}{\Sigma},$$

with $\alpha(\phi, N)$ given by expression (22) in Proposition 2a. Relative to the first benchmark,

an additional force appears: the information–risk ratio Σ_v/Σ, which raises welfare when information is precise and traders are risk-tolerant. In the full model with moral hazard, welfare is driven by the same three forces.

Proposition 2a. *Welfare is given by*

$$\mathbb{E}[W(\phi^*,N)] = Av - \beta(\phi^*,N)\Sigma A^2 + \alpha(\phi^*,N)\beta(\phi^*,N)\frac{\Sigma_v}{\Sigma}, \tag{21}$$

with

$$\alpha(\phi^*,N) = \frac{(-2+N)N\rho_I^2((2+N(-2+\phi^*)-2\phi^*)\phi^*\rho_H + (-2+N)(-1+\phi)^{*2}\rho_I)}{-(-1+N)\phi^{*2}(\rho_I - 2\phi^*\rho_I + \phi^{*2}(\rho_H+\rho_I))}, \tag{22}$$

$$\beta(\phi^*,N) = \frac{1}{2N}\left(\frac{1}{\rho_I}\phi^{*2} + \frac{1}{\rho_H}(1-\phi^*)^2\right), \tag{23}$$

where $\alpha(\phi^*,N) \geq 0$ *for* $\phi^* < 1$, *and* $\beta(\phi^*,N) \geq 0$.
Note that $\frac{\partial\beta(\phi^*,N)}{\partial N} \leq 0$; $\frac{\partial\alpha(\phi^*,N)\beta(\phi^*,N)}{\partial N} \geq 0$ *for* $\phi^* < 1$; $\frac{\partial\alpha(\phi^*,N)\beta(\phi^*,N)}{\partial\phi} \leq 0$; $\frac{\partial\beta(\phi^*,N)}{\partial\phi} \geq 0$.

Changes in moral hazard (measured by ϕ^*, since an increase in the private benefit of shirking eventually increases ϕ^*) and changes in the degree of competition (captured by N) both affect welfare but in different ways.

Proposition 2b.

(i) Welfare increases when moral hazard becomes less severe, i.e., ϕ^* decreases, because the private benefit from shirking, b, decreases.

(ii) Welfare decreases when the number of traders (N) increases when

$$\underbrace{\frac{\partial\mathbb{E}[W(\Phi(N),N)]}{\partial N}}_{\text{Direct effect}} - \underbrace{\frac{\partial\beta(\Phi(N),N)}{\partial\phi}\frac{\partial\Phi(N)}{\partial N}\Sigma A^2 + \frac{\partial\alpha(\Phi(N),N)\beta(\Phi(N),N)}{\partial\phi}\frac{\partial\Phi(N)}{\partial N}\frac{\Sigma_v}{\Sigma}}_{\text{Indirect effect through } \Phi(N)} < 0, \tag{24}$$

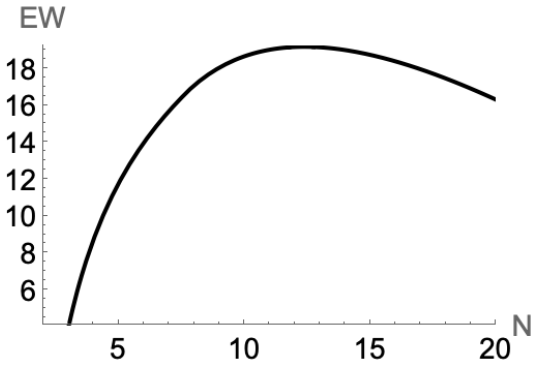
and increases otherwise.

Reducing moral hazard is always welfare enhancing. As the contract share ϕ^* declines, the loss from risk aversion falls ($\frac{\partial\beta(\phi^*,N)}{\partial\phi} \geq 0$). This is because the trader is effectively more risk-averse than the client ($\frac{1-\phi^*}{\rho_H} < \frac{\phi^*}{\rho_I}$) when the incentive constraint is binding ($\frac{\rho_I}{\rho_I+\rho_H} \leq \phi^*$). Therefore, shifting some fraction of the asset from the trader to the client reduces the total risk bearing cost, and improves risk-sharing.[15] In addition, because the trader absorbs a smaller share of the risky position, they react more

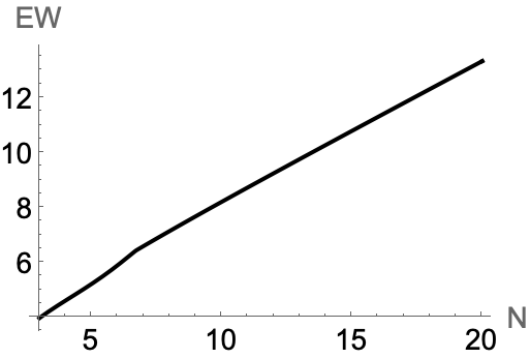
[15]For instance, if trader and client have equal risk tolerance ($\rho_H = \rho_I$), welfare rises when moving from a situation in which the trader receives more than half the surplus toward a more balanced split.

Figure 4: Examples of how equilibrium welfare can change in N

(a) Low noise from shirking



(b) High noise from shirking



Notes: Figure 4 plots equilibrium welfare, $\mathbb{E}[W(\phi^*, N)]$ from Proposition 2a, as a function of the number of traders N for the same two numerical examples as in Figure 3. In both examples, parameters are $\rho_I = \rho_H = 0.5, \Sigma = \Sigma_\nu = 1, b = 0.2$ and $\nu = 2$. Figure 3a considers a low-shirking-noise, high-supply environment ($\sigma_\epsilon = 0.5, A = 10$), while Figure 3b considers a high-shirking-noise, low-supply environment ($\sigma_\epsilon = 100, A = 2$).

aggressively to informative signals, which improves welfare. Formally, this is captured by $\frac{\partial(\alpha(\phi^*,N)\beta(\phi^*,N))}{\partial\phi} \leq 0$, which shows that lowering ϕ^* strengthens the information-driven component of welfare.

Increasing competition by adding traders can either raise or reduce welfare. When the contract does not adjust with N—that is, when incentive constraints are slack—more competition unambiguously increases welfare, for two reasons. First, as the number of traders grows, each trader (and thus each client) holds a smaller share of the risky asset, A/N, so the total loss from risk aversion declines. Formally, the term $-\frac{\partial\beta(\phi^*,N)}{\partial N}\Sigma A^2 \geq 0$ contributes positively to welfare. Second, the clearing-price aggregates information across a larger set of signals, which improves welfare through the information–risk channel, captured by $\frac{\partial(\alpha(\phi^*,N)\beta(\phi^*,N))}{\partial N}\frac{\Sigma_\nu}{\Sigma} \geq 0$.

When the contract is constrained, welfare can decline because changes in the number of traders affect welfare indirectly through the equilibrium contract. This indirect effect is negative when the contract becomes more generous to traders as N rises, that is, in the unfavorable contract regime. Welfare declines when this negative indirect effect—reflecting reduced risk sharing and suboptimal use of signal information per unit of risk (the second line of (24))—outweighs the positive direct effect of competition (the first line of (24)).

To illustrate this, I return to the two numerical examples. In the low–σ_ϵ example of Figure 3a, the contract is constrained at $N = 15$ (with $\Phi(15) = 0.89$) and is increasing

in N. At this point, the indirect marginal effect of competition on welfare is negative: $\frac{\partial \mathbb{E}[W(\Phi(15),15)]}{\partial N} = 0.667 - 0.977 < 0$. Although the direct effect $\frac{\partial \mathbb{E}[W(\phi^*_{UC},N)]}{\partial N}$ is positive, the contract's unfavorable response to competition leads to an overall decline in welfare. By contrast, in the high–σ_ϵ example of Figure 3b, the contract is decreasing at $N = 6$ (with $\Phi(6) = 0.56$). The indirect effect is therefore positive, reinforcing the direct effect, and the overall welfare effect of increasing N is positive.

Take away. The finding implies that policies aimed at promoting competition in asset trade settings may, paradoxically, reduce efficiency by exacerbating moral hazard problems. In such cases, an increase in the number of traders worsens incentive misalignments and thereby diminishes welfare. This highlights a potential negative side-effect of competition that runs counter to the standard intuition that increase competition promotes efficiency (especially in Kyle-models where traders have price impact and distort allocations because of it). Appendix A.3 compares the welfare impact of moral hazard and imperfect competition and identifies when each dominates. The results inform whether policy should focus on promoting competition or mitigating moral hazard, recognizing their interaction.

4 Quantification

In the remainder of the paper, I link the theoretical framework back to the empirical patterns motivating the two frictions, in particular Tables 1 and 2. The goal is to provide a rough assessment of the relative importance of moral hazard and imperfect competition using trade-level data alone, which are observable to regulators, rather than contractual terms or direct measures of effort, which are typically private. While such information may be available for some market participants, observing it for all traders and clients whose interactions jointly determine prices is likely infeasible.

Because neither contracts nor effort are observed, the analysis necessarily relies on relatively strong assumptions, for instance, that market participants behave according to the effort-inducing equilibrium characterized above (subject to the modifications discussed below). The exercise is therefore not intended as a definitive empirical test. Rather, it aims to illustrate the empirical relevance of the interaction between these frictions and to motivate future work that can examine them more directly.

Model-data mapping. I introduce a few targeted model adjustments that preserve tractability while better approximating the data-generating process and allowing the

model to capture the empirical patterns that motivate it.

First, the model is static, whereas trading in practice unfolds continuously throughout the day. To address this, I approximate continuous trading with frequent, periodically held batch auctions (Budish et al., 2015)—effectively discretizing trading into a sequence of uniform-price auctions that clear at Γ periods during the day. In each auction, I treat active orders as net demand and passive resting orders as fixed supply.

Second, in the model, all trades clear at a single uniform price, whereas in practice trading is subject to bid–ask spreads, so buys and sells occur at different prices. To address this, I model separate auctions for the (active) buy- versus sell-side. Treating the two sides separately avoids a mechanical bias from averaging transaction prices over short intervals—namely, higher averages when active buy orders exceed active sell orders, and lower ones when the reverse holds.

To reflect the observed price differences across accounts and exchanges, I also model separate auctions for client-account and own-account trades on each exchange. Institutionally, client and proprietary books within a dealer are typically ring-fenced by mandate, client-priority rules, and distinct risk budgets, making joint optimization across accounts operationally costly and often infeasible in real time. Conceptually, allowing for optimal bidding in an auction with multiple asymmetric account types would introduce substantial analytical complexity. It is therefore plausible—both operationally and theoretically—that traders themselves rely on a simpler, account-by-account decision rule.

Finally, I treat each security separately so that prices and trade sizes are not averaged across securities with different price levels, market capitalizations, or liquidity. Each auction for a given security clears independently: demand and supply are aggregated only within that security, not across securities. This reflects how stock markets operate, and implies that traders choose their demand for one security independently of all other securities.

Data sample. I focus on liquid S&P/TSX 60 securities to ensure sufficient trading activity over short intervals. In my baseline specification, I divide the trading day into 15-minute intervals—the shortest window that still provides sufficient activity on both exchanges, with Alpha being the limiting factor. As robustness, I also estimate the parameters using 10-, 20-, and 30-minute intervals.

Given this discretization, I collapse the raw data per exchange $e \in \{\text{TSX, Alpha}\}$, period τ in a day t, security s, account-type $k \in \{\text{CL, IN}\}$, and side $l \in \{\text{buy, sell}\}$. I average over prices with these average prices reflecting the (hypothetical) auction prices, $p^{el}_{tsk\tau}$;

and I aggregate traded quantities to obtain the total amount traded in each auction, $A^{el}_{tsk\tau}$; see Appendix Figures A2a and A2b. For each day–security–account-type–side combination, I count the number of distinct trader IDs executing a trade and use this count as the number of potential bidders in the corresponding auction, N^{el}_{tsk}; see Appendix Figure A2c.[16]

I approximate the return variance, Σ_{ts}, using the intraday variance of the national mid-quote for each security and day; and, as robustness, also construct weekly and monthly mid-quote variances. Quote-based volatility measures are common in the related literature, because they avoid microstructure noise that inflates transaction-price-based estimates (see Andersen et al., 2001; Aït-Sahalia et al., 2005; Hendershott et al., 2011). The median variance is approximately C$0.038, about 8 basis points of the median stock price (C$50.33); the full distribution of Σ_{ts} is shown in Appendix Figure A2d.

Estimation. The goal is to learn about the trader's risk-tolerance, ρ_I, the contract size, ϕ^*, and the clients' risk-tolerance, ρ_H.

The first two parameters can be point-identified. The price difference between own-account trades on TSX and Alpha identifies the trader's risk tolerance, ρ_I (similar to Table 1); while the price difference between client- and own-account trades on a given exchange identifies the contract size, ϕ^* (similar to Table 2). I rely on TSX trades for this second step, as TSX is Canada's largest and most liquid exchange and therefore the natural venue for identifying contract terms.

To see why these moments identify ρ_I, ϕ^*, consider an l-side auction of security s on exchange $e \in \{\text{TSX, Alpha}\}$ on day t in period τ for account $k \in \{IN, CL\}$. This auction clears at price:

$$p^{el}_{t\tau sk} = \frac{1}{N^{el}_{tsk}} \sum_i v_{ti\tau s} - \frac{\phi^*_k}{\rho_I} \Sigma_{ts} \left(\frac{N^{el}_{tsk} - 1}{N^{el}_{tsk} - 2} \right) \frac{A^{el}_{t\tau sk}}{N^{el}_{tsk}}, \tag{15'}$$

with $\phi^*_{IN} = 1$ for own-accounts, and $\phi^*_{CL} = \phi^* \geq 0$ for client-accounts. Since $v_{ti\tau s}$'s are

[16] This measure is an approximation: In the model, at least three active traders each receive a signal and submit one bid, whereas in the data I observe executed trades (not all bids) within each auction window, and the associated trader (firm) IDs. One alternative would count each executed trade within a period as a separate bidder, implicitly assuming that multiple trades by the same ID reflect independent private signals; this would overstate competition, particularly since traders within the same firm may not compete independently. Another approach would aggregate trades by trader ID within each auction window and count distinct IDs. I do not pursue this approach because it would likely overstate coordination within firms and understate competition, and because there are too often fewer than three trader IDs, especially for the inventory account on Alpha.

drawn iid within a day-security from $\mathcal{N}(\nu_{ts}, \Sigma_{ts})$, averaging over intraday periods yields the following average auction price:[17]

$$p^{el}_{tsk} = \frac{1}{\Gamma}\sum_{\tau} p^{el}_{t\tau sk} \approx \nu_{ts} - \frac{\phi^*_k}{\rho_I}\Sigma_{ts}\left(\frac{N^{el}_{tsk}-1}{N^{el}_{tsk}-2}\right)\frac{1}{N^{el}_{tsk}}\frac{1}{\Gamma}\sum_{\tau} A^{el}_{t\tau sk}. \tag{15"}$$

Since p^{el}_{tsk}, N^{el}_{tsk}, A^{el}_{tsk}, and Σ_{ts} are observed (or estimated), the only unknown parameters in expression (15")—after differencing prices across accounts or exchanges—are the trader's risk tolerance ρ_I and the contract size ϕ^*. For this, I use moments from both buy- and sell-side auctions, and drop observations when a security is not bought (or sold) by at least three traders during all periods in a day to ensure that taking daily averages provides a reasonable approximation to the expected auction price for a given account type and trade side.

Having identified ρ_I and ϕ^*, I exploit the equilibrium condition from Proposition 1 that the contract must satisfy $\frac{\rho_I}{\rho_H+\rho_I} \leq \phi^*$. This inequality reflects the fact that clients must offer sufficiently strong incentives to deter traders from shirking, and tighter incentives are required when clients are more risk-averse. Rearranging therefore yields a lower bound on the client's risk-tolerance parameter, ρ^{lb}_H, consistent with the observed contract.

For inference, I cluster standard errors at the daily level by resampling entire trading days with replacement and re-estimating both parameters in each bootstrap replication. I construct confidence intervals from the empirical percentiles of the bootstrap draws, which provide more reliable coverage for median-based estimators with skewed sampling distributions. For more details see Appendix B.2.

Median risk-tolerance and contract. The median estimate of risk tolerance, ρ_I, is about 670 (see Table 3). Although imprecisely estimated, the estimate is broadly consistent with Wittwer and Allen (2026), who infer slightly lower risk tolerance ($\rho_I \approx 250$) for Canadian primary dealers, which account for the largest stock trading volume in Canada (Wittwer and Uthemann, 2024). While numerical values are not directly comparable across the two settings, the qualitative conclusion is the same.

The median estimated contract share, ϕ^*, is measured with high precision and is consistently around 0.6 across specifications. This implies that intermediaries capture more than half of the surplus generated by a trade, indicating economically meaningful moral-hazard frictions. Related work in asset management infers intermediaries' surplus

[17] In the limit, when either the number of periods or the number of traders grows to infinity, this expression holds with equality.

shares from management fees and performance-based compensation (e.g., Berk and van Binsbergen, 2015; Fama and French, 2010; French, 2008; Agarwal et al., 2009; Lan et al., 2013). Although these estimates are not directly comparable to ϕ^* and apply to specific intermediaries such as mutual funds or hedge funds, they fall in a broadly similar range and likewise suggest that intermediaries capture more than half of total surplus.

Finally, the implied lower bound on clients' risk tolerance is around 385, which is about 60% of the trader's risk-tolerance. Many of the active clients in my setting are large institutional investors (such as pension funds, endowments, foundations, or fully managed investment funds). These investors typically hold diversified portfolios, making them relatively risk-tolerant.

Discussion. Because traders' risk tolerance is identified from price differences across exchanges, any unobserved factors that systematically affect execution prices for otherwise comparable inventory trades (liquidity-consuming, same security, same trade side, same time window) can bias the estimate of ρ_I. For example, if more difficult–to–execute inventory trades are disproportionately routed to Alpha while simpler orders are sent to the more liquid TSX, ρ_I would be biased upward. Similarly, if inventory trades differ in their information content across venues, observed price differences may partly reflect adverse selection.

In contrast, unobserved factors that systematically differentiate execution prices (for comparable trades) between client and inventory trades on TSX would affect the estimated contract share ϕ^*. For instance, client orders may be more urgent or require faster completion, leading to worse execution prices relative to inventory trades for reasons outside the model. In that case, contract shares would be biased upward.

While transaction data alone cannot fully rule out such concerns, identification relies on comparisons of average execution prices, which helps smooth out trade–specific effects. Moreover, the analysis focuses on trades executed within narrow time windows on two exchanges that share a common owner (TMX) and operate under similar market rules. This institutional similarity reduces the likelihood that unobserved factors generate quantitatively important biases. Reassuringly, the resulting estimates are comparable in magnitude to those reported in the existing literature.

How the contract varies with the number of traders. My model highlights that moral hazard and imperfect competition interact, implying that the contract should vary with the number of traders, unless incentive constraints do not bind (Corollary 3). To test this, I estimate $\phi^*(N)$ separately for each number of active TSX client-account traders

Table 3: Parameter estimates

Parameter	Estimate	SE (SD)	SE (Pct)	95% CI (lower)	95% CI (upper)
ρ_I	670.562	138.219	127.743	404.247	904.998
ρ_H^{lb}	385.241	82.914	80.283	221.820	536.528
ϕ^*	0.635	0.024	0.022	0.599	0.686

Notes: Table 3 reports parameter estimates from the main specification (2021, 15-minute intervals, with return variance measured daily at the security level). Appendix Table A4 reports robustness checks. Standard errors and confidence intervals are computed using a cluster bootstrap with 999 draws, clustering at the daily level (228 clusters). I report both standard-deviation bootstrap SEs and percentile-based SEs, the latter being more reliable for median-based estimators. The table also reports the implied lower bound on client risk tolerance, ρ_H^{lb}, constructed jointly with the bootstrap distribution of $(\hat{\rho}_I, \hat{\phi}^*)$; see Appendix B.2 for details.

per security and day, $N_{ts}^{TSX,l} = N$; see Appendix B.2 for details.

Figure 5a shows an upward trend but with substantial noise at the lowest and highest trader counts, where data are sparse. Excluding these thinly populated trader-count groups in Figure 5b reveals a pronounced pattern: the contract increases from about 0.45 to 0.76 (by 70%) as the number of active traders rises from 22 to 31 (by 40%). This empirical pattern is consistent with the unfavorable contract regime, in which moral hazard frictions intensify in stronger competition.

With trade-data alone, I cannot rule out that this empirical pattern is at least partially driven by confounding factors, such as reflect time-varying market conditions.[18] Cleanly accounting for time-varying market conditions would require jointly estimating all model primitives and allowing key parameters—such as the noise and the private benefit from shirking, σ_ϵ and b—to vary in market conditions. This is not feasible with trade-level data alone, because key inputs are unobserved: actual contract terms and the behavioral primitives that determine the costs and benefits of shirking. Since shirking does not occur in equilibrium, these objects leave no direct imprint in transaction data and cannot be recovered without additional information.

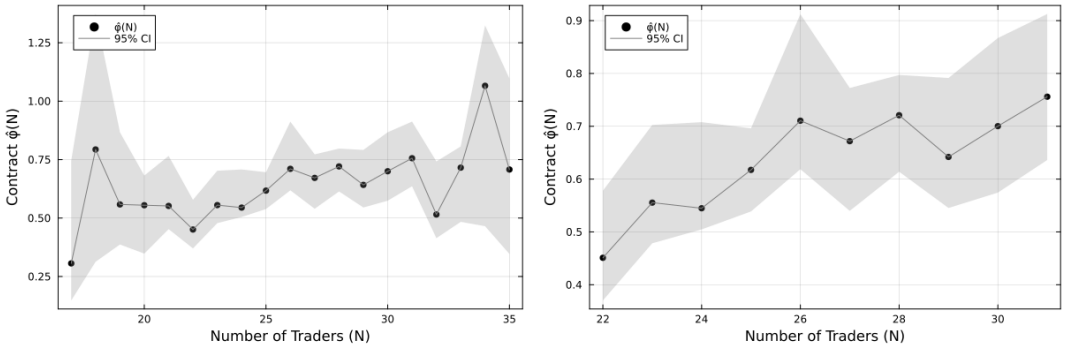
Welfare effects. If we are willing to take the model at face-value, Proposition 2b offers a way to gauge the relative importance of moral hazard and imperfect competition.

To illustrate this, I consider a representative TSX client-account auction constructed using the median return variance $\overline{\Sigma} = \Sigma_{ts}$ and the median daily average net supply

[18]For example, the model predicts that $\Phi(N)$ declines in the total amount traded, $|A|$; thus, a systematic decrease in A as N rises could generate the observed pattern. Appendix Figure A3 shows that the opposite is true: $|A|$ tends to increase with N. I additionally examine whether return variance differs systematically across trader-count groups and find no evidence of such differences.

Figure 5: How ϕ^* changes in N

(a) No restrictions (b) At least 50 observations per group



Notes: Figure 5a reports the contract estimates and their 95% confidence intervals for each trader-count group using all available observations. Figure 5b repeats the exercise but excludes groups with fewer than 50 observations. This restriction narrows the support of trader counts from 17–35 to 22–31 and removes much of the noise in the extremes. Appendix B.2 provides further details on the estimation procedure.

(C$1,604) from TSX client-account auctions with 25 traders, which lie at the center of the empirical distribution (Figure 5b).

Because I do not know the exact values of the client's risk tolerance or the information–risk ratio, $\kappa = \overline{\Sigma}/\Sigma_\nu$, I compute welfare effects over a range of values. I vary the client's risk tolerance, ρ_H, from its lower bound (415) to values exceeding the trader's risk tolerance;[19] and the information-risk ratio from 0 to 2, with 1 as the baseline. To express welfare changes in percentage terms, I set the trader's expected signal ν equal to the median mid-quote (C$57.16); because this level is large relative to the return variance (0.038), percentage changes are small, so I focus on relative welfare differences across scenarios.

Table 4 displays the welfare effects of (i) removing moral hazard, while keeping competition constant; and (ii) of increasing competition by adding one more trader for all parameter constellations. I focus on the theoretically ambiguous welfare effect of increasing competition. Adding traders has a positive direct effect through improved risk sharing, but also an indirect effect operating through endogenous contract adjustments as N changes. While the direct effect is always positive, the indirect effect can be positive or negative depending on how contracts respond to greater competition.

[19]This lower bound exceeds the median lower bound in Table 3 because the contract at $N = 25$ (0.617) is slightly smaller than the median estimate (0.635). For lower values of ρ_H, the unconstrained contract would exceed the observed one, which cannot occur in equilibrium and would mechanically reduce welfare.

Table 4: Welfare effects for median client auction

ρ		ϕ			Total		Decomp.	
$\hat{\rho}_I$	ρ_H	ϕ_{UC}	$\hat{\phi}^*$	$\hat{\phi}^*_+$	ΔW_ϕ	ΔW_N	ΔW_N^{d}	$\Delta W_N^{\mathrm{ind}}$
670	415	0.62	0.62	0.71	0.00	-0.15	0.52	-0.67
670	500	0.57	0.62	0.71	0.23	-0.55	0.55	-1.10
670	670	0.50	0.62	0.71	1.23	-1.04	0.59	-1.63
670	800	0.46	0.62	0.71	2.25	-1.27	0.60	-1.88
670	1000	0.40	0.62	0.71	4.03	-1.51	0.62	-2.13

Notes: Table 4 reports welfare effects (in percent) for different values of the client's risk tolerance ρ_H in a representative TSX client-account auction ($N = 25$, $A =$ C$1,604, $\nu =$ C$57.16, and $\Sigma_\nu = \overline{\Sigma} = 0.038$). Appendix Table A1 reports the corresponding welfare effects in dollar terms. Columns 1–2 list the estimated trader risk tolerance $\hat{\rho}_I$ and the values of ρ_H. Columns 3–5 report the unconstrained contract $\phi_{UC} = \rho_I/(\rho_I + \rho_H)$, the estimated contract $\hat{\phi}^*(N)$ at $N = 25$, and the estimated contract $\hat{\phi}^*(N{+}1)$ at $N{+}1 = 26$. Columns 6–7 report total welfare changes: ΔW_ϕ is the gain from eliminating moral hazard at $N = 25$, while ΔW_N is the total effect of increasing competition from N to $N{+}1$ with endogenous contract adjustment. Columns 8–9 decompose ΔW_N into a direct effect, ΔW_N^{d}, holding the contract fixed at $\hat{\phi}^*(N)$, and an indirect effect, $\Delta W_N^{\mathrm{ind}}$, arising from the contract change to $\hat{\phi}^*(N{+}1)$.

In the data, the contract rises when moving from 25 to 26 traders, generating a negative indirect effect that dominates the positive direct effect and leads to a net welfare loss from additional entry. When the contract instead falls the indirect effect becomes positive. Over wider ranges, such as from 22 to 31 traders, where the contract predominantly rises, the indirect effect is negative and grows in magnitude. The resulting welfare loss is amplified at lower information–risk ratios; Appendix Figure A4 for how welfare effects (i) and (ii) vary with this ratio.

Overall, the welfare loss generated by the interaction between moral hazard and imperfect competition is sizable relative to the gains from eliminating moral hazard in isolation and is robust across specifications. The effect is strongest when clients are less risk-tolerant than traders, an empirically plausible configuration. As a result, analyses that abstract from limited competition can substantially understate the severity of moral-hazard frictions.

Take away. Overall, the quantification suggests that inefficiency is driven primarily by the interaction between moral hazard and imperfect competition rather than by either friction alone. Understanding how this interaction varies with market conditions requires richer data, pointing to a natural direction for future research.

5 Conclusion

I develop a model in which imperfect competition and moral hazard jointly shape trading outcomes and welfare. Using proprietary trade data from Canadian exchanges, I show how the model helps interpret observed price differences and quantify the importance of each friction. The results imply that effective policy must address competition and incentives jointly.

References

Agarwal, V., N. D. Daniel, and N. Y. Naik (2009). Role of managerial incentives and discretion in hedge fund performance. *Journal of Finance 64*(5), 2221–2256.

Aït-Sahalia, Y., P. A. Mykland, and L. Zhang (2005). How often to sample a continuous-time process in the presence of market microstructure noise. *Journal of Finance 60*(2), 1207–1252.

Andersen, T. G., T. Bollerslev, F. X. Diebold, and P. Labys (2001). The distribution of realized exchange rate volatility. *Journal of the American Statistical Association 96*(453), 42–55.

Ausubel, L. M., P. Cramton, M. Pycia, M. Rostek, and M. Weretka (2014). Demand reduction and inefficiency in multi-unit auctions. *Review of Economic Studies 81*, 1366–1400.

Baldauf, M., C. Frei, and J. Mollner (2024). Block trade contracting. *Journal of Financial Economics 160*, 1–29.

Berk, J. B. and J. H. van Binsbergen (2015). Measuring skill in the mutual fund industry. *Journal of Financial Economics 118*(1), 1–20.

Bernhardt, D., V. Dvoracek, E. Hughson, and I. M. Werner (2005). Why do larger orders receive discounts on the london stock exchange? *The Review of Financial Studies 18*(4), 1343–1368.

Bhattacharya, S. and P. Pfleiderer (1985). Delegated portfolio management. *Journal of Economic Theory 36*(1), 1–25.

Boyd, J. H. and G. D. Nicolo (2005). The theory of bank risk taking and competition revisited. *The Journal of Finance 60*(3), 1329–1343.

Budish, E., P. Cramton, and J. Shim (2015). The high-frequency trading arms race: Frequent batch auctions as a market design response. *The Quarterly Journal of Economics 130*(4), 1547–1621.

Buffa, A. M., D. Vayanos, and P. Woolley (2022). Asset management contract and equilibrium prices. *Journal of Political Economy 130*(12), 3025–3342.

Carpenter, J. N. (2000). Does option compensation increase managerial risk appetite? *Journal of Finance 55*(5), 2311–2311.

Chen, S., R. Kaniel, and C. C. Opp (2025). Market power in the securities lending market. Working paper.

Crawford, G., N. Pavanini, and F. Schivardi (2018). Asymmetric information and imperfect competition in lending markets. *American Economic Review 108*(7), 1659–1701.

DeMarzo, P. M. and M. J. Fishman (2007). Optimal long-term financial contracting. *Review of Economic Studies 20*(6), 2079–2128.

Di Tella, S. and Y. Sannikov (2021). Optimal asset management contracts with hidden savings. *Econometrica 89*(3), 1099–1139.

Du, S. and H. Zhu (2017). What is the optimal trading frequency in financial markets? *Review of Economic Studies 84*, 1606–1651.

Einav, L., A. Fikelstein, and N. Mahoney (2021). Chapter 14: The IO of selection markets. In: Handbook of Industrial Organization, Volume 5, Issue 1: 389-429.

Fama, E. F. and K. R. French (2010). Luck versus skill in the cross-section of mutual fund returns. *Journal of Finance 65*(5), 1915–1947.

French, K. R. (2008). Presidential address: The cost of active investing. *Journal of Finance 63*(4), 1537–1573.

Garcia, D. and J. M. Vanden (2009). Information acquisition and mutual funds. *Journal of Economic Theory 144*(5), 1965–1995.

Glode, V. and C. C. Opp (2020). Over-the-counter versus limit-order-markets: The role of traders' expertise. *The Review of Financial Studies 33*(2), 866–915.

Grossman, S. J. and J. E. Stiglitz (1980). On the impossibility of informationally efficient markets. *The American Economic Review 70*(3), 393–408.

He, Z. and A. Krishnamurthy (2012). A model of capital and crises. *The Review of Economic Studies 79*(2), 735–777.

He, Z. and A. Krishnamurthy (2013). Intermediary asset pricing. *The American Economic Review 103*(2), 732–770.

Hendershott, T., C. Jones, and A. J. Menkveld (2011). Does algorithmic trading improve liquidity? *Journal of Finance 66*(1), 1–33.

Holmström, B. (1979). Moral hazard and observability. *The Bell Journal of Economics 10*(1), 74–91.

Holmström, B. and P. Milgrom (1987). Aggregation and linearity in the provision of intertemporal incentives. *Econometrica 55*(2), 303–328.

Holmström, B. and J. Tirole (1997). Financial intermediation, loanable funds, and the real sector. *The Quarterly Journal of Economics 112*(3), 663–691.

Kashyap, A. K., N. Kovrijnykh, J. Li, and A. Pavlova (2023). Is there too much benchmarking in asset managemetn? *American Economic Review 113*(4), 1112–1141.

Keeley, M. (1990). Deposit insurance, risk, and market power in banking. *American Economic Review 80*(5), 1183–1200.

Klemperer, P. D. and M. A. Meyer (1989). Supply function equilibria in oligopoly under uncertainty. *Econometrica 57*(6), 1243–1277.

Kyle, A. S. (1989). Informed speculation with imperfect competition. *Review of Economic Studies 56*(3), 317–355.

Kyle, A. S., H. Ou-Yang, and B. Wei (2011). A model of portfolio delegation and strategic trading. *The Review of Financial Studies 24*(11), 3778–3812.

Lan, Y., C. Wang, and L. Yang (2013). The economics of hedge funds. *Journal of Financial Economics 110*(2), 300–323.

Martinez-Miera, D. and R. Rapullo (2010). Does competition reduce the risk of bank failure? *The Review of Financial Studies 23*(10), 3638–3664.

Metrick, A. and A. Yasuda (2010). The economics of private equity funds. *Review of Financial Studies 23*(6), 2303–2341.

Mirrlees, J. A. (1976). The optimal structure of incentives and authority within an organization. *The Bell Journal of Economics 7*(1), 105–131.

Newey, W. K. and D. McFadden (1994). Large sample estimation and hypothesis testing. In R. F. Engle and D. McFadden (Eds.), *Handbook of Econometrics, Volume IV*, Chapter 36, pp. 2111–2245. Elsevier.

Raith, M. (2003). Competition, risk, and managerial incentives. *American Economic Review 93*(4), 1425–1436.

Rosket, M. and M. Weretka (2015). Dynamic thin markets. *Review of Financial Studies 28*, 2946–2992.

Ross, S. (1973). The economic theory of agency: The principal's problem. *American Economic Review 63*(2), 134–139.

Rostek, M. and M. Weretka (2012). Price inference in small markets. *Econometrica 80*(2), 687–711.

Rostek, M. and J. Yoon (2025). Equilibrium theory of financial markets: Recent developments. *The Journal of Economic Literature*, forthcoming.

Rostek, M. J. and J. H. Yoon (2021). Exchange design and efficiency. *Econometrica 89*(6), 2887–2928.

Scharfstein, D. S. (1988). Product-market competition and managerial slack. *RAND Journal of Economics 19*(1), 147–155.

Schmidt, K. M. (1997). Managerial incentives and product market competition. *Review of Economic Studies 64*(2), 191–213.

Starks, L. T. (1987). Performance incentive fees: An agency theoretic approach. *Journal of Financial and Quantitative Analysis 22*(1), 17–32.

Stigler, G. J. (1961). The economics of information. *Journal of Political Economy 69*(3), 213–225.

Stiglitz, J. E. and A. Weiss (1981). Credit rationing in markets with imperfect information. *American Economic Review 71*(3), 393–410.

Vayanos, D. (1999). Strategic trading and welfare in a dynamic market. *The Review of Economic Studies 66*(2), 219–254.

Vives, X. (2011). Strategic supply function competition with private information. *Econometrica 79*(6), 1919–1966.

Wilson, R. (1979). Auctions of shares. *The Quarterly Journal of Economics 93*(4), 675–689.

Wittwer, M. and J. Allen (2026). Market power and capital constraints. American Economic Review.

Wittwer, M. and A. Uthemann (2024). Specialization in financial markets. Working paper.

APPENDIX

Imperfect Competition and Moral Hazard in Financial Markets

by Milena Wittwer

Section A shows additional theoretical results, and proofs.
Section B presents estimation details and additional evidence.
Appendix Tables and Figures are displayed at the end.

A Additional theoretical evidence

Section A.1 provides existence and uniqueness conditions.
Section A.2 solves for an equilibrium with price-taking traders.
Section A.3 analyses when imperfect competition is worse than moral hazard.
Section A.4 presents mathematical auxiliary statements and all proofs.

A.1 Existence and uniqueness

Throughout the paper, I assume that the trader's private benefit from shirking is sufficiently large and that the client incurs a sufficiently large loss when the trader shirks. The trader's benefit captures private gains from reduced effort, such as lower monitoring costs, time savings, or diverted attention, while the client's loss reflects the idea that shirking can trigger persistent reputational, institutional, or relational costs. Both can be large relative to one–shot trading gains.

These two forces jointly determine whether effort can be sustained in equilibrium. The following proposition characterizes how large the trader's benefit and the client's loss must be to guarantee existence of the equilibrium described in Proposition 1. Given the complexity of the equilibrium contract and the endogenous interactions among traders through market clearing, the conditions are expressed in terms of expected equilibrium payoffs. Crucially, these conditions can be expressed entirely in model primitives by simply substituting in the equilibrium contract.

Proposition 3 (Existence)**.** *Define, $f(\phi_i, \phi)$ is the payoff the client expects (net of $\bar{U}$) when choosing some contract ϕ_i in response to everyone else choosing some contract, ϕ, and the trader exerts effort; while the clients expects payoff $g(\phi_i, \phi) + b - \Delta$ when the trader shirks instead. Formally:*

$$f(\phi_i, \phi) \equiv \mathbb{E}[\mathcal{W}(\phi_i; 0)] + \mathbb{E}[\mathcal{U}(0; \phi_i)], \tag{25}$$

$$g(\phi_i, \phi) + b - \Delta \equiv \mathbb{E}[\mathcal{W}(\phi_i; 1)] + \mathbb{E}[\mathcal{U}(1; \phi_i)]. \tag{26}$$

The effort-inducing equilibrium characterized in Proposition 1 exists when (i) the trader's benefit from shirking is sufficiently high

$$b > \frac{\beta(N)}{\gamma(N)},$$

with $\beta(N)$ and $\gamma(N)$ defined in Proposition 1; and (ii) the client's loss from shirking is sufficiently high:

$$\Delta \;\geq\; b + g\big(\phi_i^S(\phi^*), \phi^*\big) - f(\phi^*, \phi^*). \tag{27}$$

Here ϕ^* denotes the equilibrium contract in Proposition 1, and $\tilde{\phi}_i^S(\phi^*)$ denotes the client's best shirking-inducing deviation against ϕ^*. Formally:

$$\phi_i^S(\phi^*) = \begin{cases} \tilde{\phi}_i^S(\phi^*) & \text{if } \tilde{\phi}_i^S(\phi^*) < \dfrac{\rho_I}{\rho_I + \rho_H} \quad (\textit{shirking-IC binds}), \\ \dfrac{\rho_I}{\rho_I + \rho_H} & \text{otherwise} \quad (\textit{shirking-IC slack}), \end{cases} \tag{28}$$

where $\tilde{\phi}_i^S(\phi^*)$ denote a contract that makes the trader indifferent between shirking and effort when other clients choose ϕ^*, i.e., $\mathbb{E}[\mathcal{U}(1; \tilde{\phi}_i^S(\phi^*))] = \mathbb{E}[\mathcal{U}(0; \tilde{\phi}_i^S(\phi^*))]$.

Condition (i) ensures that the trader has incentives to shirk for sufficiently large contract shares. If the private benefit from shirking, b, is too small, the trader would instead exert effort even under a zero contract. To see this, note that when all clients choose ϕ^*, the trader's expected payoff from exerting effort relative to shirking is

$$\mathbb{E}[\mathcal{U}(0; \phi^*)] - \mathbb{E}[\mathcal{U}(1; \phi^*)] = \frac{\alpha(N)}{\gamma(N)}(\phi^* A)^2 - \frac{\beta(N)}{\gamma(N)} - b.$$

If $b < \beta(N)/\gamma(N)$, this difference is positive even at $\phi^* = 0$. In that case, exerting effort strictly dominates shirking for all feasible contracts, so the incentive constraint never binds and the equilibrium contract is always unconstrained.

Condition (ii) compares the client's expected payoff from the equilibrium contract ϕ^*—under which the trader exerts effort—to the highest payoff the client could obtain by deviating to a shirking-inducing contract when all other clients choose ϕ^*. The condition requires that the loss from shirking, Δ, is large enough to offset the maximal gain

from such a deviation. When this holds, no client has a profitable deviation to inducing shirking given ϕ^*, and the effort-inducing equilibrium of Proposition 1 exists.

Proposition 4 (Uniqueness within symmetric linear-demand equilibria)**.** *There is no other symmetric equilibrium with linear demand functions.*

There are only two symmetric equilibrium candidates with linear demand: one that induces effort by all traders (Proposition 1) and one that induces universal shirking. In the effort-inducing equilibrium, the contract share must be sufficiently high when the effort incentive constraint binds, $\frac{\rho_I}{\rho_I+\rho_H} < \Phi(N)$, where $\Phi(N)$ denotes the symmetric contract that renders traders indifferent between exerting effort and shirking. In contrast, in the shirking-inducing case the contract share must be sufficiently low when the corresponding constraint binds, $\Phi(N) < \frac{\rho_I}{\rho_I+\rho_H}$. In this region, each client has a profitable deviation to the higher, unconstrained contract share, which improves risk sharing. As a result, the shirking-inducing candidate does not constitute an equilibrium.

A.2 Price-takers

To isolate the role of the price-impact effect, I shut it down by assuming that traders behave as price takers and derive the resulting equilibrium (Proposition 5). Comparing this equilibrium to the one with strategic traders (Proposition 1) reveals that the constrained contract, $\tilde{\Phi}(N)$, closely resembles the contract under strategic demand reduction, $\Phi(N)$. This similarity highlights the central role of the information effect in shaping the interaction between moral hazard and imperfect competition.

In this appendix, I study the price-taking equilibrium in more detail.

Proposition 5. *In a symmetric equilibrium:*

(i) All clients offer the following contract $\{K^*, \phi^*\}$:

$$\phi^* = \begin{cases} \Phi_C(N) & \text{when the incentive constraint binds, i.e., when } \Phi_U(N) \leq \Phi_C(N) \\ \Phi_U(N) & \text{otherwise,} \end{cases}$$

where

$$\Phi_C(N) = \frac{1}{|A|}\sqrt{\frac{\tilde{\beta}(N) + b\tilde{\gamma}(N)}{\tilde{\alpha}(N)}} \tag{29}$$

$$\Phi_U(N) = \frac{1}{2}\left(\frac{\sqrt{(\rho_H+\rho_I)((N-2)^2\rho_H + N^2\rho_I)}}{\rho_H+\rho_I} - (N-2)\right) > 0, \tag{30}$$

and

$$K^* = \bar{U} - \rho_I \frac{N-1}{2N} \frac{\Sigma_\nu}{\Sigma} - \frac{A^2 \Sigma}{2N^2 \rho_I} (\phi^*)^2, \text{ where} \tag{31}$$

$\tilde{\alpha}(N) = \Sigma^2(\Sigma - \tilde{\Sigma})(\Sigma - (N-1)^2\tilde{\Sigma}) > 0,$
$\tilde{\beta}(N) = -(N-1)N\rho_I^2(\Sigma^2\Sigma_\nu + \Sigma((N-2)\Sigma_\nu - \bar{m}^2(N-1)N(\sigma_\epsilon + \Sigma_\nu))\tilde{\Sigma} + (N-1)^2\Sigma_\nu\tilde{\Sigma}^2) < 0,$
$\tilde{\gamma}(N) = 2N^2\rho_I\Sigma(\Sigma + (N-1)\tilde{\Sigma})^2 > 0$, *given* $\bar{m} = \Sigma_\nu(\Sigma_\nu + \sigma_\epsilon)^{-1}$ *and* $\tilde{\Sigma} \equiv Var[R_i|\mu_i(1)]$.

(ii) All traders exert effort, and submit the following demand function:

$$a^*(p, \mu_i(s_i)) = \rho_I(\phi^* Var[R_i|\mu_i(s_i)])^{-1}(\mathbb{E}[R_i \mid \mu_i(s_i)] - p). \tag{32}$$

This equilibrium exists if $b > \frac{\tilde{\beta}(N)}{\tilde{\gamma}(N)}$.

The equilibrium with price-taking traders is closely related to the equilibrium with strategic traders characterized in Proposition 1, but differs in two ways. First, by assumption, traders submit demand functions as in equation (11) while setting $\Lambda_i = 0$, effectively ignoring their own price impact. Relative to the strategic benchmark, traders therefore overestimate how aggressively they can trade without affecting prices, leading them to behave as if competition were stronger than it truly is. This misperception alters the functional form of individual demands, and changes how those demands aggregate and, in equilibrium, how prices respond to order flow.

Second, and crucially, clients correctly anticipate this behavior and adjust the contracts they offer accordingly. Even though traders act as if they were price takers, clients understand that this belief is incorrect and take into account how traders' misperceived incentives translate into equilibrium prices and allocations. As a result, contract choice differs from the strategic-trading case both when the incentive constraint binds and when it does not. In particular, clients internalize that traders' failure to account for price impact changes incentives, and they respond by adjusting the strength of the contract to partially offset this mistake.

Importantly, clients' ability to correct traders' misperception is limited. Contracts are set ex ante and are not auction-specific; in particular, they do not condition on realized clearing prices or on the trader's contribution to price formation. As a consequence, clients cannot directly discipline traders for ignoring their price impact, nor can they fully restore the strategic-trading benchmark through contract design alone. This limitation is reflected in the structure of the unconstrained contract $\Phi_U(N)$, which depends explicitly on the number of traders. Unlike in the strategic equilibrium, $\Phi_U(N)$ does not coincide

with the second-best benchmark $\rho_I/(\rho_I+\rho_H)$ for finite N. Instead, it converges to this benchmark only as $N \to \infty$, when traders truly become price takers and the misperception underlying their behavior vanishes.

A.3 When is imperfect competition worse than moral hazard?

In this Appendix I investigate which friction—when considered 'in isolation'—has a larger impact on welfare and identify the conditions under which one dominates the other.

Specifically, to assess the welfare loss from imperfect competition, which arises due to demand reduction, I remove frictions from moral hazard (by setting $\phi^* = \frac{\rho_I}{\rho_I+\rho_H}$). I compare the expected welfare when traders are price-takers (i.e., do not reduce their demand strategically) to the corresponding welfare when traders act strategically and reduce their demand (as in Proposition 1).

To analyze the effect of moral hazard, I, instead, remove the loss due to strategic demand reduction. I assume traders act as price-takers and compute the welfare loss from moral hazard as the difference between the expected welfare that would arise under unconstrained contracting—that is, with $\phi^* = \frac{\rho_I}{\rho_I+\rho_H}$ for all traders—and the welfare when moral hazard constraints bind for all traders, $\phi^* = \Phi(N)$.

Proposition 6. *Consider the scenario where moral hazard constraints bind, $\frac{\rho_I}{\rho_H+\rho_I} < \Phi(N)$.*

The welfare loss from imperfect competition is, for any Σ, A and distribution of v_i, smaller than the welfare loss from moral hazard if

$$f(\phi) \leq N \text{ , where } f(\phi) = 1 + \frac{\phi\sqrt{\rho_H(\rho_H+\rho_I)}}{-\rho_I + \phi(\rho_H+\rho_I)} \text{ at } \phi = \Phi(N). \tag{33}$$

At $\phi = \Phi(N)$ this expression rearranges to:

$$\Phi(N) \geq \frac{\rho_I(N-1)}{(N-1)(\rho_H+\rho_I) - \sqrt{\rho_H(\rho_H+\rho_I)}}, \tag{34}$$

with this cutoff being larger than $\frac{\rho_I}{\rho_I+\rho_H}$.

The first part of the corollary (33) identifies the minimal number of traders necessary for imperfect competition to distort asset allocations less than moral hazard for a given contract ϕ.[20] As moral hazard becomes less sever (b declines which decreases ϕ), $f(\phi)$

[20]The relationship is visualized in Appendix Figure A5 for a simplified environment in which traders and clients have the same risk tolerance ($\rho_I = \rho_H$).

increases. This means the number of traders such that the loss from imperfect competition is smaller than the loss from moral hazard increases. When moral hazard disappears ($\phi \to \frac{\rho_I}{\rho_I+\rho_H}$), there is no finite number of traders such that that imperfect competition is less welfare harming than moral hazard.

The second condition (34) states that the welfare loss from imperfect competition is smaller than the loss from moral hazard when $\Phi(N)$ is sufficiently high. In the unfavorable contract regime, where $\Phi(N)$ increases in N and moral hazard frictions worsen as the number of traders rises, the left-hand side of the condition increases while the right-hand side declines in N, making the condition more likely to hold. Moreover, Corollary 2 implies that, in the limit, when $b > \gamma(\sigma_\epsilon)$, the condition must be satisfied, so that the welfare loss from imperfect competition is dominated by the loss from moral hazard.

Proposition 6 offers a tool for evaluating whether policy should primarily target imperfect competition or moral hazard, while recognizing that the two frictions interact. For instance, enhancing competition could involve lowering exchange membership fees to attract additional traders. Addressing moral hazard, by contrast, would require redesigning the contractual relationship between dealer institutions and traders—such as by structuring compensation to more directly reward effort that improves returns.

A.4 Additional mathematical results and proofs

In this appendix I provide additional mathematical statements, and all proofs. I start with deriving the unconditional and conditions distribution of returns (1). Then I provide the full proof of Proposition 1, after which I characterize how high the client's loss from shirking must be to rule out deviations to shirking. All other proofs follow, first for the propositions and then for the corollaries.

Lemma 1. *(i)* $R_i \sim \mathcal{N}(\nu, \Sigma+\Sigma_\nu)$. *(ii)* $R_i|\mu_i(s_i) \sim \mathcal{N}(\mathbb{E}[R_i \mid \mu_i(s_i)], Var[R_i \mid \mu_i(s_i)])$, *where*

$$\mathbb{E}[R_i \mid \mu_i(s_i)] = \begin{cases} \nu + \Sigma_\nu(\Sigma_\nu + \sigma_\epsilon)^{-1}\big((\nu_i + \epsilon_i) - \nu\big) & when s_i = 1 \\ \nu_i & when\ s_i = 0. \end{cases} \tag{35}$$

$$Var[R_i \mid \mu_i(s_i)] = \begin{cases} \Sigma + \left(\Sigma_\nu^{-1} + \sigma_\epsilon^{-1}\right)^{-1} & when\ s_i = 1 \\ \Sigma & when\ s_i = 0. \end{cases} \tag{36}$$

Proof of Lemma 1. (i) Because $\nu_i \sim \mathcal{N}(\nu, \Sigma_\nu)$ and $R_i \mid \nu_i \sim \mathcal{N}(\nu_i, \Sigma)$, the marginal distribution of R_i is given by: $R_i \sim \mathcal{N}(\nu, \Sigma_\nu + \Sigma)$, because $\mathbb{E}[R_i] = \mathbb{E}[\mathbb{E}[R_i|\nu_i]] = \mathbb{E}[\nu_i] =$

ν and $Var(R_i) = \mathbb{E}[Var(R_i|\nu_i)] + Var(\mathbb{E}[R_i|\nu_i]) = \Sigma + \Sigma_\nu$.

(ii) To derive the conditional expectation, note:

$$\mathbb{E}[R_i \mid \mu_i(0)] = \nu_i,$$
$$\mathbb{E}[R_i \mid \mu_i(1)] = \mathbb{E}[\nu_i \mid \nu_i + \epsilon_i] = \nu + Cov(\nu_i, \nu_i + \epsilon_i)Var(\nu_i + \epsilon_i)^{-1}\Big((\nu_i + \epsilon_i) - \nu\Big),$$

where

$$Cov(\nu_i, \nu_i + \epsilon_i) = \Sigma_\nu, \qquad Var(\nu_i + \epsilon_i) = \Sigma_\nu + \sigma_\epsilon.$$

Therefore,

$$\mathbb{E}[\nu_i \mid \nu_i + \epsilon_i] = \nu + \Sigma_\nu(\Sigma_\nu + \sigma_\epsilon)^{-1}\Big((\nu_i + \epsilon_i) - \nu\Big).$$

To derive the conditional variance, I use the law of total variance: $Var(R_i \mid \mu_i(s_i)) = \mathbb{E}[Var(R_i \mid \nu_i) \mid \mu_i(s_i)] + Var(\mathbb{E}[R_i \mid \nu_i] \mid \mu_i(s_i))$. Since $R_i \mid \nu_i \sim \mathcal{N}(\nu_i, \Sigma)$, we have:

$$Var(R_i \mid \mu_i(s_i)) = \Sigma + Var(\nu_i \mid \mu_i(s_i)).$$

To compute $Var(\nu_i \mid \mu_i(s_i))$, observe that this is a standard Bayesian updating problem with Gaussian prior and linear Gaussian signal. The prior is $\nu_i \sim \mathcal{N}(\nu, \Sigma_\nu)$, and the signal model is: $\mu_i(s_i) = \nu_i + \epsilon_i, \epsilon_i \sim \mathcal{N}(0, \sigma_\epsilon(s_i))$ with $\sigma_\epsilon(0) = 0$ and $\sigma_\epsilon(1) = \sigma_\epsilon > 0$. Therefore:

$$Var(\nu_i \mid \mu_i(1)) = \left(\Sigma_\nu^{-1} + \sigma_\epsilon^{-1}\right)^{-1} \text{ and } Var(\nu_i \mid \mu_i(0)) = 0. \quad \square$$

Proof of Proposition 1. I follow a guess-and verify strategy, and solve the game backwards. Starting in the last period, I consider the choice of trader i and assume all other traders play the equilibrium strategy. The trader observes signal $\mu_i(s_i)$, depending on their previous choice of effort, and chooses their demand schedule. Their objective is to maximize the expected utility from the money they will receive at the end of the game, which is:

$$\omega^c(s_i; \phi_i) = \phi_i\Big\{a_i(p, \mu_i(s_i))(R_i - p)\Big\} + s_i b + K_i.$$

Fo this, the trader chooses quantities $a_i(p, \mu_i(s_i))$ for $s_i \in \{0, 1\}$, point-wise for each clearing price p to maximize their expected payoff from purchasing assets:

$$\mathbb{E}\left[1 - \exp\left(-\frac{1}{\rho_I}\omega^c(s_i; \phi_i)\right)\middle|\mu_i(s_i)\right]. \tag{2}$$

Since $R_i|\mu_i(s_i) \sim \mathcal{N}(\mathbb{E}[R_i \mid \mu_i(s_i)], Var[R_i \mid \mu_i(s_i)])$, point-wise maximizing (2) is equivalent to standard mean-variance optimization:

$$\phi_i(\mathbb{E}[R_i \mid \mu_i(s_i)] - p)a_i - \frac{1}{2}\frac{\phi_i^2}{\rho_I}a_i Var[R_i \mid \mu_i(s_i)]a_i + s_i b + K_i,$$

where $a_i = a_i(p, \mu_i(s_i))$. The necessary condition for quantity a_i at the fixed market clearing price p and both $s_i \in \{0, 1\}$ is

$$\mathbb{E}[R_i \mid \mu_i(s_i)] - \phi_i \frac{1}{\rho_I} Var[R_i \mid \mu_i(s_i)]a_i = p + \Lambda_i a_i, \text{ where } \Lambda_i = \left(\frac{dp}{da_i}\right). \tag{13}$$

Λ_i measures the price impact of the trader's own demand on the clearing price. Since all other traders j but i submit the equilibrium guess (11), price impact is $\Lambda_i = \frac{\phi^*}{\rho_I(N-2)}\Sigma,$. Moreover, the necessary conditions are sufficient, since $Var[R_i \mid \mu_i(s_i)]$ is positive semi-definite, $N > 2$ and $\phi_i > 0$.

Therefore, equation (13) pins down the trader's best reply $a_i(p, \mu_i(s_i))$ for $s_i \in \{0, 1\}$. When trader i submits this demand curve and everyone else the equilibrium demand, the market clears at

$$p^c(s_i) : A = a_i^c(s_i) + \sum_{j \neq i} a^*(p^c, \mu_j(0)),$$

where $a_i^c(s_i)$ abbreviates the clearing quantity, $a_i(p^c, \mu_i(s_i))$ and $a^*(p^c, \mu_j(0))$ is the equilibrium candidate (11).

Before trading and observing their private signal, the trader chooses their effort level in period 2. They exert effort when

$$\mathbb{E}[\mathcal{U}(0; \phi_i)] + K_i \geq \mathbb{E}[\mathcal{U}(1; \phi_i)] + K_i, \tag{IC}$$

where

$$\mathcal{U}(s_i; \phi_i) = \phi_i(\mathbb{E}[R_i \mid \mu_i(s_i)] - p^c(s_i))a_i^c(s_i) - \frac{1}{2}\frac{\phi_i^2}{\rho_I}a_i^c(s_i) Var[R_i \mid \mu_i(s_i)]a_i^c(s_i) + bs_i. \tag{5}$$

In the first period client i chooses the contract so as to maximize the expected certainty

equivalent $\mathbb{E}[\mathcal{W}(\phi_i;s_i)]-K_i$, with

$$\begin{aligned}
&\mathcal{W}(\phi_i;s_i)-K_i,\\
\text{where } \mathcal{W}(\phi_i;s_i) =&(1-\phi_i)\Big(\mathbb{E}[R_i \mid \mu_i(s_i)]-p^c(s_i)\Big)a_i^c(s_i)\\
&-(1-\phi_i)^2\frac{1}{2\rho_H}a_i^c(s_i)Var[R_i \mid \mu_i(s_i)]a_i^c(s_i)-\Delta s_i. \qquad (19)
\end{aligned}$$

The client faces two constraints:

$$\mathbb{E}\left[\mathcal{U}(0;\phi_i)\right] \geq \mathbb{E}\left[\mathcal{U}(1;\phi_i)\right] \qquad \text{(IC)}$$

$$\mathbb{E}\left[\mathcal{U}(0;\phi_i)\right] + K_i \geq \bar{U} \qquad \text{(IR)}$$

Since the incentive constraint is independent of K_i, we can choose K_i to ensure participation:

$$K_i \equiv K(\phi_i;0) = \bar{U} - \mathbb{E}\left[\mathcal{U}(0;\phi_i)\right]. \qquad \text{(IR')}$$

Note that $K(\phi_i;0)$ depends on ϕ_i, because ϕ_i affects the trader's demand function in the trading game. Inserting the binding participation constraint (IR') into the client's objective function (12), the optimization problem simplifies to:

$$\max_{\phi_i}\Big\{\mathbb{E}[\mathcal{W}(\phi_i;0]-K_i(\phi_i;0)\Big\} = \max_{\phi_i}\Big\{\mathbb{E}[\mathcal{W}(\phi_i;0]+\mathbb{E}[\mathcal{U}(0;\phi_i]-\bar{U}\Big\} \qquad (12)$$

subject to (IC), assuming Δ is sufficiently high that the client always wants to incentivize effort. To preview, there are two possible solutions for ϕ_i: an unconstrained contract, $\phi_i = \frac{\rho_I}{\rho_I+\rho_H}$, and a constrained contract, $\phi_i = \Phi(N)$. The unconstrained contract maximizes $\mathbb{E}[\mathcal{W}(\phi_i;0)]+\mathbb{E}[\mathcal{U}(0;\phi_i)]-\bar{U}$. The constrained contract is pinned down by the binding incentive constraint: $\mathbb{E}[\mathcal{U}(0;\phi_i)] = \mathbb{E}[\mathcal{U}(1;\phi_i)]$. The solution is unconstrained when $\Phi(N) \leq \frac{\rho_I}{\rho_I+\rho_H}$ and constrained otherwise.

To see why this is the case, consider first the constrained scenario. The constrained contract is given by the binding incentive constraint, defined below, in expressions (37) and (38), respectively. Since all clients must choose the same contract in the symmetric equilibrium, I can solve for ϕ $\mathbb{E}[\mathcal{U}(0;\phi)] = \mathbb{E}[\mathcal{U}(1;\phi)]$ for ϕ. This gives a unique positive solution, which I call $\Phi(N)$.

Consider now the unconstrained scenario, which arises when $\Phi(N) \leq \frac{\rho_I}{\rho_I+\rho_H}$. First note that $\mathbb{E}[\mathcal{W}(\phi_i;0]+\mathbb{E}[\mathcal{U}(0;\phi_i]-\bar{U}$ is maximized at $\phi_i = \frac{\rho_I}{\rho_I+\rho_H}$, independently of other clients' contracts. Since clients are ex-ante identical, we know that other clients

must be choosing the unconstrained contract as well if this is optimal for one client.

Next, I verify that when the solution is unconstrained, the trader exerts effort, i.e., the incentive constraint is satisfied. To check this, compute the trader's expected certainty equivalents for an arbitrary ϕ_i for some given ϕ chosen by the other clients, by relying on market clearing and equilibrium trading behavior in period 3:

$$\mathbb{E}[\mathcal{U}(0;\phi_i)] = \phi_i \frac{A^2(N-1)\phi^2\Sigma^2 + (N-2)^2 N\rho_I^2\Sigma_\nu)}{2(N-2)(N-1)(2\phi+(N-2)\phi_i)\rho_I\Sigma}, \tag{37}$$

$$\mathbb{E}[\mathcal{U}(1;\phi_i)] = b + \phi_i \frac{A^2(N-1)\phi^2\Sigma^2(\sigma_\epsilon+\Sigma_\nu) + (N-2)^2\rho_I^2\Sigma_\nu(\sigma_\epsilon+N\Sigma_\nu)}{2(N-2)(N-1)\rho_I(2\Sigma(\sigma_\epsilon+\Sigma_\nu)+(N-2)\phi_i(\Sigma\sigma_\epsilon+(\Sigma+\sigma_\epsilon)\Sigma_\nu))} \tag{38}$$

To ensure that the trader has no incentive to deviate to shirking when the contract is unconstraint, $\phi_i = \frac{\rho_I}{\rho_H+\rho_I}$, I need to show that $\mathbb{E}[\mathcal{U}(0;\phi_i)] \geq \mathbb{E}[\mathcal{U}(1;\phi_i)]$ when all clients choose the unconstrained contract: $\phi_i = \phi = \frac{\rho_I}{\rho_H+\rho_I}$. To show this, let's first impose $\phi_i = \phi$ for some ϕ and simplify the incentive constraint:

$$\mathbb{E}[\mathcal{U}(0;\phi)] \geq \mathbb{E}[\mathcal{U}(1;\phi)] \Leftrightarrow b \leq \mathcal{C}(\phi) = \frac{A^2\phi^2\alpha(N)-\beta(N)}{\gamma(N)},$$

given $\alpha(N), \beta(N), \gamma(N)$ as defined in the proposition. When $b \leq \mathcal{C}\left(\frac{\rho_I}{\rho_I+\rho_H}\right)$ the incentive constraint is fulfilled. By construction, we know that $\mathcal{C}(\Phi(N)) = b$. Since I consider the case where $\Phi(N) \leq \frac{\rho_I}{\rho_I+\rho_H}$, and $\mathcal{C}'(\phi) \geq 0$, the incentive constraint is satisfied.

Finally, inserting the optimal ϕ^* into the expression for the fixed part of the contract (IR') completes the proof. □

Proof of Proposition 2a. I first compute realized welfare (19) in the equilibrium. Then I take expectations over return expectations, and rely on the simplifying assumptions that these expectations are independent across traders i. Then to show that $\alpha(\phi,N)$ and $\beta(\phi,N)$ are non-negative for $\phi^* < 1$, I rely on the assumptions that $0 < \rho_H \leq \rho_I$, and $N > 2$, and the property that $\phi^* > 0$. □

Proof of Proposition 2b. To show statement (i) I take the partial derivative of welfare, characterize in Corollary 2a, with respect to ϕ. Given $0 < \rho_H \leq \rho_I$, $N > 2$, and $\phi^* > 0$: $\frac{\partial\alpha(\phi^*,N)\beta(\phi^*,N)}{\partial\phi} \leq 0$, $\frac{\partial\beta(\phi^*,N)}{\partial\phi} \geq 0$. Therefore, welfare always decreases in ϕ, provided that $\mathbb{E}[\nu_i\Sigma^{-1}\nu_i] \geq \mathbb{E}[\nu_i]\Sigma^{-1}\mathbb{E}[\nu_i]$ and $A\Sigma A \geq 0$, which is always the case (in a one-asset setting). Similarly, taking the partial derivative with respect to the number of traders, N, shows statement (ii). □

Proof of Proposition 3. Existence requires that both (i) the trader's private benefit from shirking, b, and (ii) the client's loss from shirking, Δ, are sufficiently large.

To establish condition (i), note that the trader must have an incentive to shirk for sufficiently large contract shares for there two be two cases—an unconstrained contract and a constrained one. Suppose instead that b is small, and let everyone play according to the equilibrium. The trader's expected payoff from exerting effort relative to shirking is given by

$$\mathbb{E}[\mathcal{U}(0;\phi^*)]-\mathbb{E}[\mathcal{U}(1;\phi^*)]=\frac{\alpha(N)}{\gamma(N)}(\phi^* A)^2-\frac{\beta(N)}{\gamma(N)}-b.$$

If $b<\beta(N)/\gamma(N)$, this expression is strictly positive even at $\phi^*=0$. Hence, exerting effort strictly dominates shirking for all feasible contracts. In this case, the incentive constraint never binds and the equilibrium contract is unconstrained, ruling out the equilibrium of interest.

To derive condition (ii), consider a client who seeks to induce shirking while all other clients choose the equilibrium contract ϕ^*. Such a client solves

$$\max_{\phi_i}\left\{\mathbb{E}[\mathcal{W}(\phi_i;1)]-K_i(\phi_i;1)\right\}=\max_{\phi_i}\left\{\mathbb{E}[\mathcal{W}(\phi_i;1)]+\mathbb{E}[\mathcal{U}(1;\phi_i)]-\bar{U}\right\} \tag{39}$$

$$\text{subject to } \mathbb{E}[\mathcal{U}(0;\phi_i)]\leq\mathbb{E}[\mathcal{U}(1;\phi_i)]. \tag{40}$$

The best reply is:

$$\phi_i^S(\phi^*)=\begin{cases}\tilde{\phi}_i^S(\phi^*) & \text{if } \tilde{\phi}_i^S(\phi^*)<\dfrac{\rho_I}{\rho_I+\rho_H} \quad \text{(shirking-IC binds)},\\ \dfrac{\rho_I}{\rho_I+\rho_H} & \text{otherwise} \quad \text{(shirking-IC slack)},\end{cases} \tag{28}$$

where $\tilde{\phi}_i^S(\phi^*)$ makes the trader indifferent between shirking and effort, i.e., $\mathbb{E}[\mathcal{U}(1;\tilde{\phi}_i^S(\phi^*))]=\mathbb{E}[\mathcal{U}(0;\tilde{\phi}_i^S(\phi^*))]$. Therefore, the deviation is not profitable when

$$\mathbb{E}[\mathcal{W}(\phi_i^S(\phi^*);1)]+\mathbb{E}[\mathcal{U}(1;\phi_i^S(\phi^*))]\leq\mathbb{E}[\mathcal{W}(\phi^*;0)]+\mathbb{E}[\mathcal{U}(0;\phi^*)],$$

which is equivalent to condition (27).

Thus far, I have ruled out deviations to contracts that induce shirking. To complete the proof, it remains to rule out deviations from ϕ^* that leave the trader's effort choice unchanged. When the incentive constraint is slack, the client has no profitable deviation, since ϕ^* maximizes expected payoff. When the incentive constraint binds, the only de-

viation that preserves effort is to offer a higher contract share. Such a deviation lowers the client's payoff and is therefore unprofitable. □

Note that $\tilde{\phi}_i^S(\phi^*)$ and the f and g functions can be expressed explicitly. Using expressions (37) and (38), this contract is given by

$$\begin{aligned}\tilde{\phi}_i^S(\phi^*) = &\Big(8b(N-1)\phi^{*2}\rho_I\Sigma^2\Big)\Big[-(N-2)(N-1)\phi^*\rho_I\Sigma\Big(-\rho_I\Sigma_\nu+\tilde{m}^2\rho_I(\sigma_\epsilon+\Sigma_\nu)+2b(\Sigma+\tilde{\Sigma})\Big)\\ &\pm\Big\{\phi^{*2}\rho_I\Sigma^2\Big((2-3N+N^2)^2\rho_I(-\rho_I\Sigma_\nu+\tilde{m}^2\rho_I(\sigma_\epsilon+\Sigma_\nu)+2b(\Sigma+\tilde{\Sigma}))^2\\ &-8b(-1+N)\Big(A^2(-1+N)\phi^{*2}\Sigma^2(\Sigma-\tilde{\Sigma})+(-2+N)^2\rho_I\Big(\tilde{m}^2(-1+N)\rho_I\Sigma(\sigma_\epsilon+\Sigma_\nu)\\ &+2b(-1+N)\Sigma\tilde{\Sigma}+\rho_I\Sigma_\nu(\Sigma-N\tilde{\Sigma})\Big)\Big)\Big)\Big\}^{1/2}\Big]^{-1},\end{aligned} \tag{41}$$

where $\tilde{m}=\Sigma_\nu(\Sigma_\nu+\sigma_\epsilon)^{-1}$, $\tilde{\Sigma}=Var(R_i|\mu_i(1))$. The f function is:

$$\begin{aligned}f(\phi_i,\phi)\equiv&\mathbb{E}[\mathcal{W}(\phi_i;0)]+\mathbb{E}[\mathcal{U}(0;\phi_i)]\\ =&\Big((2\phi\rho_H+(-2+N)(-\rho_I-(-2+\phi_i)\phi_i(\rho_H+\rho_I)))\\ &(A^2(-1+N)\phi^2\Sigma^2+(-2+N)^2N\rho_I^2\Sigma_\nu)\Big)\\ &\Big(2(-2+N)(-1+N)(2\phi+(-2+N)\phi_i)^2\rho_H\rho_I\Sigma\Big)^{-1}.\end{aligned} \tag{25}$$

In comparison, the client expected payoff when allowing shirking and choosing some ϕ_i in responds to all other traders choosing contract ϕ is $g(\phi_i,\phi)+b-\Delta$, where:

$$\begin{aligned}g(\phi_i,\phi)\equiv&\mathbb{E}[\mathcal{W}(\phi_i;1)]+\mathbb{E}[\mathcal{U}(1;\phi_i)]-b+\Delta\\ =&\Big((A^2(-1+N)\phi^2\Sigma^2+(-2+N)^2\rho_I^2(\Sigma_\nu+\tilde{m}^2(-1+N)\\ &(\sigma_\epsilon+\Sigma_\nu)))(2\phi\rho_H\Sigma-(-2+N)(\rho_I+(-2+\phi_i)\phi_i(\rho_H+\rho_I))\tilde{\Sigma})\Big)\\ &\Big(2(-2+N)(-1+N)\rho_H\rho_I(2\phi\Sigma+(-2+N)\phi_i\tilde{\Sigma})^2\Big)^{-1}.\end{aligned} \tag{26}$$

Proof of Proposition 4. There is only one other symmetric equilibrium candidate with linear demand: one that induces shirking by all traders, and it takes the following form:

$$\phi^{*S}=\begin{cases}\Phi(N) & \text{if } \Phi(N)<\dfrac{\rho_I}{\rho_I+\rho_H} \quad \text{(shirking-IC binds)},\\ \phi_{UC}^S & \text{otherwise.}\end{cases} \tag{42}$$

for some ϕ_{UC}^S. I show that when $\Phi(N)<\frac{\rho_I}{\rho_I+\rho_H}$, a client has a profitable unilateral deviation to the unconstrained contract, regardless of whether this deviation induces the trader to continue shirking or to switch to exerting effort. The reason is that the client's

objective function,

$$\begin{cases} \mathbb{E}[\mathcal{W}(\phi_i;1)] + \mathbb{E}[\mathcal{U}(1;\phi_i)] - \bar{U}, & \text{if the trader shirks,} \\ \mathbb{E}[\mathcal{W}(\phi_i;0)] + \mathbb{E}[\mathcal{U}(0;\phi_i)] - \bar{U}, & \text{if the trader exerts effort,} \end{cases} \tag{43}$$

is, for any fixed contract chosen by all other clients $\phi_{-i} = \phi$, strictly concave in ϕ_i and maximized at $\frac{\rho_I}{\rho_H+\rho_I}$ over the region

$$\phi_i < \underbrace{\frac{\phi}{N-2}\frac{\Sigma}{Var(R_i \mid \mu_i(s_i))} + \frac{3\rho_I}{2(\rho_H+\rho_I)}}_{\geq \frac{\rho_I}{\rho_I+\rho_H}}.$$

Since $\Phi(N) < \frac{\rho_I}{\rho_I+\rho_H}$ lies within this region, deviating to the unconstrained contract is strictly profitable. □

Proof of Proposition 5. The proof is analogous to the proof of Proposition 1. The only difference is that we now set $\Lambda_i = 0$. □

Proof of Proposition 6. To assess the welfare loss from imperfect competition, which arises due to demand reduction, I remove frictions from moral hazard (by setting $\phi = \frac{\rho_I}{\rho_I+\rho_H}$). I compare the welfare when traders are price-takers to the corresponding welfare when traders act strategically and reduce their demand (as in Proposition 1).

To analyze the effect of moral hazard, I assume traders act as price-takers and compute the loss as the difference between the expected welfare that would arise under unconstrained contracting—that is, with $\phi^* = \frac{\rho_I}{\rho_I+\rho_H}$ for all traders—and the welfare when moral hazard constraints bind for all traders, $\phi^* = \Phi(N)$.

The welfare loss due to imperfect competition (when shutting off frictions arising from moral hazard) is

$$LIC = +\underline{\gamma}\left(\mathbb{E}[\nu_i\Sigma^{-1}\nu_i] - \mathbb{E}[\nu_i]\Sigma^{-1}\mathbb{E}[\nu_i]\right) - \underline{\delta}A\Sigma A, \text{ with} \tag{44}$$

$$\underline{\gamma} = \frac{\rho_H+\rho_I}{2(N-1)} \geq 0 \tag{45}$$

$$\underline{\delta} = 0. \tag{46}$$

As expected, the welfare loss converges to zero when $N \to \infty$.

The welfare loss due to moral hazard (when shutting off frictions arising from imper-

fect competition) is

$$LMH(N,\phi) = \tilde{\gamma}\left(\mathbb{E}[\nu_i\Sigma^{-1}\nu_i] - \mathbb{E}[\nu_i]\Sigma^{-1}\mathbb{E}[\nu_i]\right) - \tilde{\delta}A\Sigma A, \text{ with} \tag{47}$$

$$\tilde{\gamma} = \frac{(\rho_I - \phi(\rho_H+\rho_I))^2(N-1)N\rho_I(\rho_H+\rho_I)}{2N\phi^2\rho_H\rho_I(\rho_H+\rho_I)} \geq 0 \tag{48}$$

$$\tilde{\delta} = -\frac{(\rho_I - \phi(\rho_H+\rho_I))^2\phi^2}{2N\phi^2\rho_H\rho_I(\rho_H+\rho_I)} \leq 0 \tag{49}$$

where $\phi = \Phi(N) \geq \frac{\rho_I}{\rho_I+\rho_H}$; $LMH(N, \frac{\rho_I}{\rho_H+\rho_I}) = 0$ when moral hazard constraints do not bind.

Imperfect competition leads to a smaller welfare loss (in allocative efficiency) for any A, Σ and distribution of ν_i when

$$\underline{\gamma} \leq \tilde{\gamma} \text{ and } \tilde{\delta} \leq \underline{\delta},$$

given $A\Sigma A \geq 0$ and $\left(\mathbb{E}[\nu_i\Sigma^{-1}\nu_i] - \mathbb{E}[\nu_i']\Sigma^{-1}\mathbb{E}[\nu_i]\right) \geq 0$ by assumption. Since $N > 2$, $\frac{\rho_I}{\rho_I+\rho_H} < \phi$, this is the case when $f(\phi) \leq N$ with $f(\phi)$ given in Proposition 6.

To derive the condition when ϕ depends on N, we insert $\Phi(N)$ and rearrange the condition leveraging the fact that $\Phi(N) \geq \frac{\rho_I}{\rho_H+\rho_I}$ and $N \geq 2$:

$$1 + \frac{\Phi(N)\sqrt{\rho_H(\rho_H+\rho_I)}}{-\rho_I + \Phi(N)(\rho_H+\rho_I)} \leq N \iff \Phi(N)\left(\sqrt{\rho_H(\rho_H+\rho_I)} - (N-1)(\rho_H+\rho_I)\right) \leq -(N-1)\rho_I$$

$$\iff \Phi(N) \geq \frac{(N-1)\rho_I}{(N-1)(\rho_H+\rho_I) - \sqrt{(\rho_H+\rho_I)\rho_H}}. \quad \square$$

Proof of Corollary 1. To derive the market clearing price in equilibrium it suffices to insert the equilibrium demand functions of all traders into the market clearing condition, $\sum_i a^*(p^*, \mu_i(0)) = A$, and solve for the price. The equilibrium allocations can then be found by inserting the clearing price into the equilibrium demand. $\square$

Proof of Corollary 2. To see this, I analyze the incentive constraint as the market becomes perfectly competitive:

$$\lim_{N\to\infty} \mathbb{E}[\mathcal{U}(\mu_i(0);0)] = \frac{\rho_I}{2}\frac{\Sigma_\nu}{\Sigma} \tag{17}$$

$$\lim_{N\to\infty} \mathbb{E}[\mathcal{U}(\mu_i(1);1)] = \frac{\rho_I}{2}\frac{\Sigma_\nu^2(\Sigma_\nu+\sigma_\epsilon)^{-1}}{\Sigma + (\Sigma_\nu^{-1}+\sigma_\epsilon^{-1})^{-1}} + b. \tag{18}$$

From here we see that the incentive constrained cannot be fulfilled for any ϕ_i when

$$b > \gamma_\infty(\sigma_\epsilon) \equiv \frac{\rho_I \sigma_\epsilon \Sigma_\nu (\Sigma + \Sigma_\nu)}{2\Sigma(\sigma_\epsilon \Sigma_\nu + \Sigma(\sigma_\epsilon + \Sigma_\nu))},$$

and is always fulfilled otherwise. □

Proof of Corollary 3. $\Phi'(N) < 0$ when $(\beta(N) + b\gamma(N))\alpha'(N) > (\beta'(N) + b\gamma'(N))\alpha(N)$; else $\Phi'(N) \geq 0$. Inserting for $\alpha(N), \beta(N)$, and $\gamma(N)$ and rearranging gives the cutoff of the corollary. □

B Additional empirical evidence

Appendix B.1 provides additional motivating evidence.
Appendix B.2 describes the estimation procedure used to take the model to the data.
Appendix Tables and Figures are presented at the end.

B.1 Additional motivating evidence and discussion

Imperfect competition. This appendix provides additional discussion and robustness evidence related to the price differences between TSX and Alpha documented in Table 1.

One potential concern is that the observed price gaps reflect execution-related liquidity frictions. A smaller venue may exhibit thinner order-book depth or slower resiliency, increasing execution costs. Moreover, orders routed to Alpha—whether client or proprietary—may differ in urgency or aggressiveness, for example by walking further into the book or demanding faster execution, which would mechanically worsen prices.

Importantly, not all such liquidity differences constitute confounding factors. Many of them translate into higher price impact, and are therefore conceptually similar to a reduction in competition. What cannot be fully ruled out with transaction data alone are execution-related frictions orthogonal to competition-related changes in price impact.

Several features of the evidence suggest that these (unobservable) execution-related frictions are unlikely to be the sole drivers of the observed price difference. TSX and Alpha are owned by the same parent company and operate under similar rule books and execution protocols, limiting the scope for institutional differences. The Alpha price gap is symmetric across trade direction—higher prices for purchases and lower prices for sales—consistent with a general price-impact mechanism rather than one-sided urgency or execution failures. Moreover, as shown in Appendix Table A2, Panel D, the Alpha effect

persists—and if anything becomes larger—when attention is restricted to small trades, which are least likely to mechanically walk up a thinner the limit order book.

Another concern is strategic venue selection. Traders may route orders across exchanges based on order characteristics that are imperfectly observed by the econometrician—for example, sending more difficult or information-sensitive orders to Alpha while routing simpler orders to TSX. To address this, the regressions condition ontrader × symbol × trade date × hour fixed effects, so that identification comes from comparing the same trader executing the same stock, on the same side of the market, within the same narrow time window, across the two venues. Any remaining selection would therefore need to operate at this level to explain the results.

Appendix Table A2 further shows that the estimated Alpha coefficient is stable across increasingly restrictive fixed-effects specifications. The similarity of the estimates when moving from symbol–day fixed effects (Panel A), to hour-level controls (Panel B), and to trader-level fixed effects (Panel C) indicates that the price gap persists within narrowly defined trader–asset–time cells and is therefore unlikely to be driven by coarse liquidity differences, intraday market conditions, or persistent trader-specific routing strategies.

Taken together, the evidence is consistent with imperfect competition affecting prices, though—as is typical with transaction-level data—it does not establish a causal relationship.

Moral hazard. This appendix provides additional discussion and robustness considerations related to the price differences between client-account and own-account trades documented in Table 2.

One concern is that client orders differ from own-account trades in ways that are imperfectly observed by the econometrician—for example, they may be more urgent, less strategically timed, or subject to constraints imposed by clients or internal execution protocols. If such differences were pervasive, worse prices for client-account trades could reflect higher execution difficulty rather than distorted incentives. Importantly, to the extent that traders exert less effort, monitor less closely, or optimize less aggressively when trading on behalf of clients—because they do not internalize the full surplus—these frictions are part of the moral-hazard mechanism itself.

To mitigate concerns about unobservable differences across account-types that are orthogonal to moral hazard frictions, regressions condition on symbol × trade date × hour × exchange fixed effects, so that client and own-account trades are compared within the same security, on the same venue, and within the same narrow time window, absorbing

systematic differences in asset-level liquidity and intraday market conditions.[21] The estimated price gap is symmetric across trade direction: client purchases occur at higher prices and client sales at lower prices than own-account trades. This pattern is consistent with a general deterioration in execution quality rather than one-sided urgency or directional information effects. Moreover, Appendix Table A3, Panel C shows that the price gap persists when attention is restricted to small trades, which may be more comparable across account types in terms of trade motive and informational content and less likely to differ mechanically in how they interact with the limit order book.

Another concern is selection across account types. Traders may choose whether to execute a given order on behalf of a client or for their own account based on private information about expected execution quality or short-term price dynamics. If traders systematically internalized easier or more profitable opportunities while allocating more difficult trades to client accounts, the observed price gap could reflect strategic selection rather than incentive distortions. According to the rules of the exchange, however, such selection is not permitted.

Taken together, while the evidence does not establish moral hazard in a causal sense, it documents a robust—and previously unobserved—empirical pattern that is consistent with incentive distortions in delegated trading (Appendix Table A3). In the model, moral hazard manifests as reduced effort or suboptimal execution when traders act on behalf of clients, leading to systematically worse prices for client trades. The empirical findings align closely with this mechanism.

B.2 Estimation procedure

I estimate ρ_I and ϕ^* parameters using median-based moment conditions in a two-stage GMM procedure, and then compute the implied lower bound on client risk tolerance. I use a two-stage rather than a joint estimator because the first-stage estimate of ρ_I is relatively noisy, making sequential estimation more stable and numerically robust, while remaining consistent and asymptotically equivalent to the joint estimator under regularity conditions (Newey and McFadden, 1994).[22] I rely on medians rather than

[21] Including trader fixed effects is not feasible in this setting because relatively few traders operate on both account types within the same security and hour, leaving insufficient within-trader variation to support reliable estimation. Appendix Table A3, Panel D shows that when trader fixed effects are included without hour interactions, the estimated effects remain in the same direction as in the baseline specifications and are statistically significant on the sell side.

[22] The first parameter, ρ_I, is identified from price differences between inventory–account trades on TSX and Alpha, which provide only a noisy signal of risk tolerance because the data sample is

means to limit the influence of outliers.

Estimating ρ_I and ϕ^*. Recall the average daily price per auction type (15"):

$$p^{el}_{tsk} = \nu_{ts} - \frac{\phi^*_k}{\rho_I}\Sigma_{ts}\left(\frac{N^{el}_{tsk}-1}{N^{el}_{tsk}-2}\right)\frac{A^{el}_{tsk}}{N^{el}_{tsk}}, \qquad (e,l)\in\{\text{TSX},\text{Alpha}\}\times\{b,q\}, \tag{15"}$$

with $\phi^*_k = 1$ for $k = IN$ and $\phi^*_k > 0$ for $k = CL$, and where p^{el}_{tsk}, N^{el}_{tsk}, $A^{el}_{tsk} = \frac{1}{\Gamma}\sum_\tau A_{t\tau sk}$, and Σ_{ts} are observed (or estimated), while (ρ_I, ϕ^*) are unknown and ν_{ts} cancels in price differences. Define

$$m^{el}_{tsk} = \Sigma_{ts}\left(\frac{N^{el}_{tsk}-1}{N^{el}_{tsk}-2}\right)\frac{A^{el}_{tsk}}{N^{el}_{tsk}}, \tag{50}$$

and denote an observation, i.e., a (t,s,e,l,k) cell by ι.

To obtain the first moment for estimation, take the difference between Alpha inventory trades and TSX inventory prices using expression (15") and rearrange for ρ_I:

$$R^{(1)}_\iota \equiv \frac{m^{\text{TSX},l}_{ts,IN} - m^{\text{Alpha},l}_{ts,IN}}{p^{\text{Alpha},l}_{ts,IN} - p^{\text{TSX},l}_{ts,IN}}. \tag{51}$$

Similarly, take the difference between TSX inventory and TSX client prices using expression (15") and rearrange for ϕ^* to obtain the second moment:

$$R^{(2)}_\iota(\rho) \equiv \frac{m^{\text{TSX},l}_{ts,IN} - \rho\left(p^{\text{TSX},l}_{ts,CL} - p^{\text{TSX},l}_{ts,IN}\right)}{m^{\text{TSX},l}_{ts,CL}}. \tag{52}$$

Let $\psi(u) = \frac{1}{2} - \mathbf{1}\{u < 0\}$ denote the median ($\tau = \frac{1}{2}$) score function. The associated sample moment conditions are

$$\bar{g}^{(1)}_n(\rho) = \frac{1}{n}\sum_{\iota=1}^{n}\psi\big(R^{(1)}_\iota - \rho\big), \qquad \bar{g}^{(2)}_n(\phi^*;\rho) = \frac{1}{n}\sum_{\iota=1}^{n}\psi\big(R^{(2)}_\iota(\rho) - \phi^*\big). \tag{53}$$

The parameters (ρ_I, ϕ^*) are estimated in two stages as just-identified quantile (Z-) esti-

relatively small, and the differences in execution prices across exchanges are small and affected by venue-specific microstructure variation. In contrast, ϕ^* is identified from the price difference between client and inventory trades on the same exchange, which yields a much sharper signal. Estimating both parameters jointly would therefore force the contract–size estimate to absorb noise from the weakly identified first moment, reducing efficiency and stability.

mators. The first-stage estimator $\hat{\rho}_I$ solves

$$\bar{g}_n^{(1)}(\hat{\rho}_I) = 0 \quad \Longleftrightarrow \quad \hat{\rho}_I = \text{median}\big\{R_\iota^{(1)}\big\}, \tag{54}$$

i.e., the sample median of $R_\iota^{(1)}$ across all valid observations. Conditional on $\hat{\rho}_I$, the second-stage estimator $\hat{\phi}^*$ solves

$$\bar{g}_n^{(2)}(\hat{\phi}^*; \hat{\rho}_I) = 0 \quad \Longleftrightarrow \quad \hat{\phi}^* = \text{median}\big\{R_\iota^{(2)}(\hat{\rho}_I)\big\}, \tag{55}$$

the median of the adjusted ratios computed using the plug-in estimate $\hat{\rho}_I$.

To obtain standard errors and confidence intervals, I implement a nonparametric cluster bootstrap at the daily level. Let days $t \in \mathcal{T}$ denote the clustering units, allowing arbitrary dependence within a day but independence across days. For each bootstrap replication $b = 1, \ldots, B$:

1. Resample $|\mathcal{T}|$ days from $\mathcal{T}$ with replacement.
2. Reconstruct the bootstrap sample containing all observations belonging to the resampled days.
3. Recompute $(\hat{\rho}_I^{(b)}, \hat{\phi}^{*(b)})$ using the same two–step procedure.

Bootstrap standard errors are computed as the empirical standard deviations of the bootstrap estimates,

$$\widehat{\text{SE}}(\hat{\theta}) = \text{sd}\big(\hat{\theta}^{(1)}, \ldots, \hat{\theta}^{(B)}\big), \qquad \hat{\theta} \in \{\hat{\rho}_I, \hat{\phi}^*\}. \tag{56}$$

Normal-approximation confidence intervals are given by

$$\text{CI}_{95}^{\text{norm}}(\hat{\theta}) = \big[\hat{\theta} \pm 1.96\, \widehat{\text{SE}}(\hat{\theta})\big]. \tag{57}$$

I also report percentile-based bootstrap confidence intervals,

$$\text{CI}_{95}^{\text{pct}}(\hat{\theta}) = \big[Q_{0.025}(\hat{\theta}^{(1)}, \ldots, \hat{\theta}^{(B)}),\ Q_{0.975}(\hat{\theta}^{(1)}, \ldots, \hat{\theta}^{(B)})\big], \tag{58}$$

and the corresponding implied standard errors,

$$\widehat{\text{SE}}^{\text{pct}}(\hat{\theta}) = \frac{Q_{0.975}(\hat{\theta}^{(1)}, \ldots, \hat{\theta}^{(B)}) - Q_{0.025}(\hat{\theta}^{(1)}, \ldots, \hat{\theta}^{(B)})}{2 \times 1.96}. \tag{59}$$

Percentile intervals are my preferred specification, as they make no symmetry or normality assumption about the sampling distribution of the estimator and are therefore more reliable when the bootstrap distribution is skewed or heavy–tailed, as is typical for ratio– and median–based estimators. By re–estimating ρ_I in each bootstrap replication, the inference procedure correctly accounts for first–stage estimation uncertainty. In practice, the percentile and normal–approximation intervals yield very similar results.

Implied lower bound on client risk tolerance. In equilibrium $\frac{\rho_I}{\rho_I+\rho_H}$ is weakly below the contract, ϕ^*. Thus the model implies a lower bound on client risk tolerance,

$$\rho_H^{\text{lb}} = \rho_I\Big(\frac{1}{\phi^*} - 1\Big). \tag{60}$$

Plugging in my estimates $(\hat{\rho}_I, \hat{\phi}^*)$, I can compute the implied $\hat{\rho}_H^{\text{lb}}$.

For inference, I construct the bootstrap distribution of the bound jointly with $(\hat{\rho}_I, \hat{\phi}^*)$: in each bootstrap replication b, I compute the corresponding $\hat{\rho}_H^{\text{lb},(b)}$, whenever the ratio is well defined (which is always the case in the data). The bootstrap standard error and percentile 95% confidence interval are then computed analogously to those for $\hat{\rho}_I$ and $\hat{\phi}^*$. I also report the corresponding percentile–implied standard error, constructed using the same mapping from percentile bounds as above. By re–estimating all underlying objects in each bootstrap replication, this procedure ensures that uncertainty in the bound $\hat{\rho}_H^{\text{lb}}$ reflects sampling variation in both $\hat{\rho}_I$ and $\hat{\phi}^*$.

Estimating the contract as a function of N. To estimate ϕ^* as a function of the number of traders N, I apply the same two–stage procedure as in the baseline estimation, restricting attention to subsamples with a fixed trader count. Specifically, I first estimate ρ_I using the full sample, as described above. Conditional on $\hat{\rho}_I$, I estimate $\phi^*(N)$ as the median of the ratio in (52), which combines information from TSX inventory and client auctions, restricting attention to observations for which the corresponding TSX client auction involves exactly $N_{ts,CL}^{\text{TSX},,l} = N$ active traders. Inference follows the same clustered bootstrap by trading day as in the baseline case.

Appendix Table A1: Welfare effects for median client auction in dollar values

ρ		ϕ			Total		Decomp.	
$\hat{\rho}_I$	ρ_H	ϕ_{UC}	$\hat{\phi}^*$	$\hat{\phi}^*_+$	ΔW_ϕ	ΔW_N	ΔW_N^{d}	$\Delta W_N^{\mathrm{ind}}$
670	415	0.62	0.62	0.71	0.00	-0.15	0.54	-0.70
670	450	0.60	0.62	0.71	0.06	-0.35	0.56	-0.91
670	500	0.57	0.62	0.71	0.24	-0.58	0.58	-1.16
670	670	0.50	0.62	0.71	1.31	-1.11	0.62	-1.73
670	800	0.46	0.62	0.71	2.40	-1.36	0.64	-2.01
670	1000	0.40	0.62	0.71	4.33	-1.62	0.67	-2.29

Notes: Appendix Table A1 is analogous to Table 4 but displays welfare effects in dollar values, in thousand C$. Percentage and level welfare changes are similar because baseline welfare is large (approximately C$105,000), dominated by the level term Av, while risk-adjustment components are small. The table reports welfare effects in levels (thousand C$) for different values of the client's risk tolerance ρ_H in a representative TSX client-account auction ($N = 25$, $A =$ C$1,604, $v =$ C$57.16, and $\Sigma_v = \overline{\Sigma} = 0.038$). Columns 1–2 list the estimated trader risk tolerance $\hat{\rho}_I$ and the values of ρ_H. Columns 3–5 report the unconstrained contract $\phi_{UC} = \rho_I/(\rho_I + \rho_H)$, the estimated contract $\hat{\phi}^*(N)$ at $N = 25$, and the estimated contract $\hat{\phi}^*(N{+}1)$ at $N{+}1 = 26$. Columns 6–7 report total welfare changes: ΔW_ϕ is the welfare gain from eliminating moral hazard at $N = 25$, while ΔW_N is the total effect of increasing competition from N to $N{+}1$ with endogenous contract adjustment. Columns 8–9 decompose ΔW_N into a direct effect, ΔW_N^{d}, holding the contract fixed at $\hat{\phi}^*(N)$, and an indirect effect, $\Delta W_N^{\mathrm{ind}}$, arising from the contract change to $\hat{\phi}^*(N{+}1)$.

Appendix Table A2: Robustness—Prices are worse on the more concentrated exchange (Alpha)

	(Buy)		(Sell)	
Panel A: Symbol × Trade Date Fixed Effects				
Alpha	0.005***	(0.001)	-0.003*	(0.001)
quantity	0.219***	(0.028)	-0.203***	(0.032)
N	139,626,867		134,387,016	
Panel B: Symbol × Trade Date × Hour Fixed Effects				
Alpha	0.006***	(0.000)	-0.006***	(0.000)
quantity	0.129***	(0.013)	-0.104***	(0.016)
N	139,302,602		134,088,687	
Panel C: Symbol × Trade Date × Hour × Trader Fixed Effects				
Alpha	0.004***	(0.000)	-0.004***	(0.000)
quantity	0.154***	(0.011)	-0.115***	(0.018)
N	136,715,587		131,862,280	
Panel D: Symbol × Trade Date × Hour × Trader Fixed Effects (Small Trades)				
Alpha	0.006***	(0.001)	-0.005***	(0.000)
quantity	67.345	(37.727)	-68.919	(48.785)
N	96,753,660		91,755,344	

Notes: Appendix Table A2 reports robustness regressions of transaction prices on an indicator for the Alpha exchange and trade size. Panels differ by the granularity of fixed effects. Panel A controls for symbol × trade date fixed effects; Panel B adds symbol × trade date × hour fixed effects; Panel C further includes trader fixed effects, conditioning on execution prices within narrowly defined symbol–time–trader cells. Panel D replicates the Panel C specification but restricts the sample to small trades, defined as trades below the 25th percentile of the trade-size distribution. Column (Buy) uses asset purchases, while column (Sell) uses sales. Prices are measured in Canadian dollars, and quantities in millions of shares. Standard errors are clustered at the daily level. Across all specifications, trades executed on Alpha occur at worse prices than those on TSX.

Appendix Table A3: Robustness—Client-Account Price Differences Across Specifications

	(Buy)		(Sell)	
Panel A: Symbol × Trade Date × Exchange Fixed Effects				
Client-account	0.010***	(0.001)	-0.013***	(0.001)
quantity	0.253***	(0.030)	-0.213***	(0.034)
N	134,116,675		128,870,237	
Panel B: Symbol × Trade Date × Exchange × Hour Fixed Effects				
Client-account	0.005***	(0.000)	-0.005***	(0.000)
quantity	0.156***	(0.014)	-0.114***	(0.018)
N	133,615,215		128,407,743	
Panel C: Symbol × Trade Date × Exchange × Hour Fixed Effects (Small Trades)				
Client-account	0.005***	(0.000)	-0.005***	(0.000)
quantity	85.586*	(41.785)	-98.354*	(45.592)
N	95,083,035		89,801,803	
Panel D: Symbol × Trade Date × Exchange × Trader Fixed Effects				
Client-account	0.001	(0.002)	-0.008***	(0.002)
quantity	0.187***	(0.026)	-0.198***	(0.032)
N	132,993,000		127,838,651	

Notes: Appendix Table A3 reports robustness regressions of transaction prices on an indicator for client-account trades and traded quantity. Panels differ by the granularity of fixed effects and sample restrictions. Panel A includes symbol × trade date × exchange fixed effects. Panel B further adds hour-of-day fixed effects, comparing client and own-account trades executed in the same security, on the same exchange, and within the same hour. Panel C restricts the sample to small trades, defined as trades below the 25th percentile of the trade-size distribution. Panel D includes trader fixed effects interacted with symbol, trade date, and exchange, conditioning on price differences within narrowly defined trader–asset–date cells. Own-account trades serve as the baseline category. Column (Buy) uses asset purchases, while column (Sell) uses sales. Prices are measured in Canadian dollars, and quantities in millions of shares. Standard errors are clustered at the daily level.

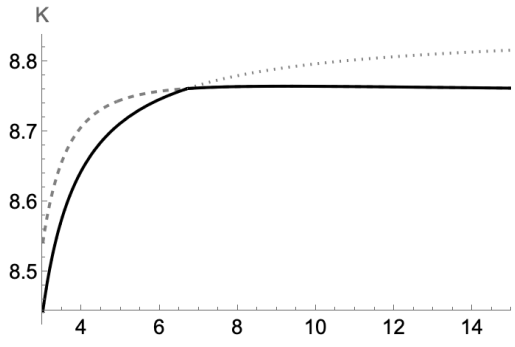
Appendix Table A4: Robustness—Parameter estimates across specifications

Parameter	Estimate	SE (SD)	SE (Pct)	95% CI Low	95% CI High
10 minutes					
ρ_I	805.528	461.094	429.936	216.588	1901.94
ϕ^*	0.666	0.042	0.040	0.602	0.761
ρ_H^{lb}	403.689	220.855	218.851	112.004	969.902
20 minutes					
ρ_I	506.907	65.4669	64.7832	390.245	644.195
ϕ^*	0.608	0.014	0.014	0.581	0.637
ρ_H^{lb}	326.405	47.5308	48.0952	234.604	423.137
30 minutes					
ρ_I	455.981	44.642	42.6444	366.303	533.469
ϕ^*	0.576	0.009	0.009	0.558	0.594
ρ_H^{lb}	334.799	34.9132	33.2562	262.411	392.775
Weekly return variance					
ρ_I	3610.320	875.799	930.790	1999.730	5648.430
ϕ^*	0.616	0.019	0.018	0.582	0.657
ρ_H^{lb}	2243.720	566.188	563.248	1212.500	3420.430
Monthly return variance					
ρ_I	18689.300	4838.380	4862.650	11276.200	30337.700
ϕ^*	0.634	0.0222	0.021	0.591	0.674
ρ_H^{lb}	10754.500	2820.610	2737.740	6539.260	17271.200

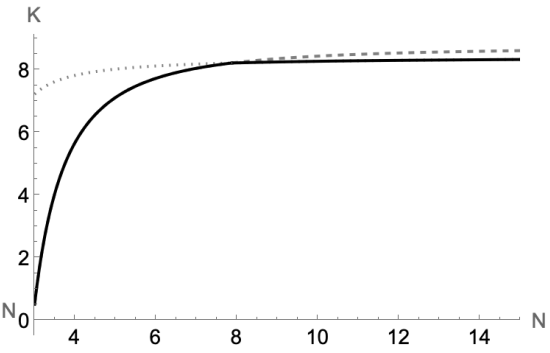
Notes: Appendix Table A4 reports robustness results for the estimated trader risk tolerance ρ_I, the contract share ϕ^*, and the implied lower bound on client risk tolerance ρ_H^{lb}. In the three upper sections, I vary the time interval used to aggregate trades into hypothetical auctions. Because Alpha is relatively illiquid, intervals shorter than 10 minutes yield too few observations to cluster standard errors reliably. As the aggregation window becomes coarser, the estimated trader risk tolerance declines, but confidence intervals overlap across specifications. In the two lowest sections, I use alternative measures of return variance. The baseline uses within-day mid-quote variation at the symbol level. Using weekly or monthly return variances mechanically shifts the level of estimated risk tolerance, since risk tolerance is defined relative to the unit of risk, which changes in the time-horizon. Formally, prices depend on the ratio Σ/ρ_I, implying that an increase in Σ must be offset by a proportional increase in ρ_I. Across all specifications, the estimated contract share remains close to 0.6.

Appendix Figure A1: Examples of how the fixed part of the contract, K^*, can change in N

(a) K^*: Low noise from shirking



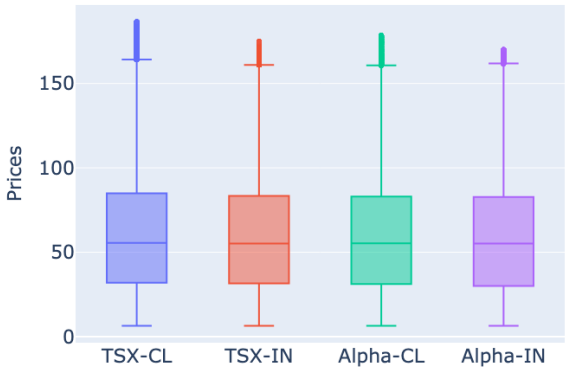
(b) K^*: High noise from shirking



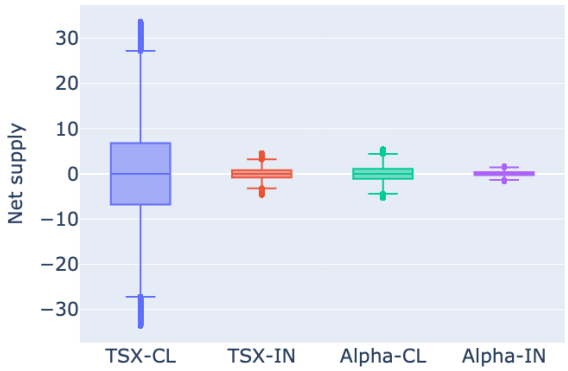
Notes: Appendix Figure A1 plot the equilibrium contract K^* (black line) as a function of the number of traders N for two numeric examples. The chosen contract is a piecewise function implied by the unconstrained contract (dashed gray line) and the constrained contract (dotted gray line). In both examples, $\rho_I = \rho_H = 0.5, \Sigma = \Sigma_v = 1, b = 0.2$, $\bar{U} = 9$. In Figure A1b shirking-noise is low, $\sigma_\epsilon = 0.5$ and supply is large, $A = 10$, in Figure A1a it is high, $\sigma_\epsilon = 100$ and supply is small $A = 2$.

Appendix Figure A2: Distributions of prices, net supply, trader counts, and return variance

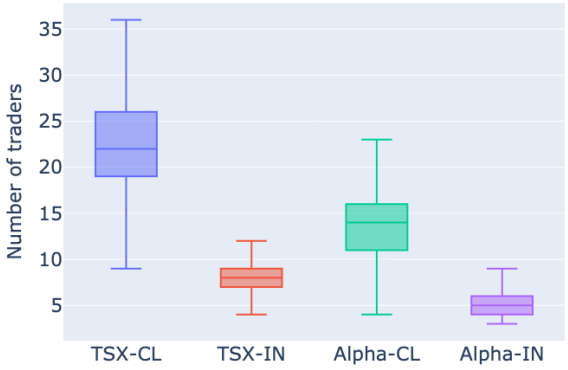
(a) Prices by auction type



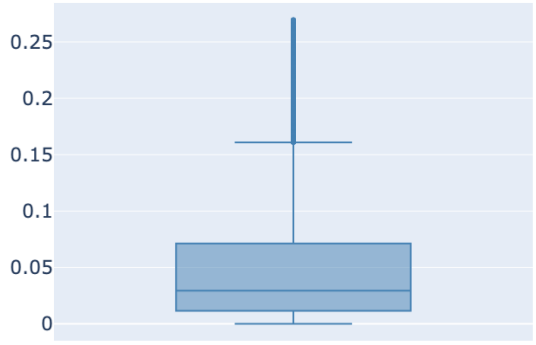
(b) Net supply by auction type



(c) Number of active traders



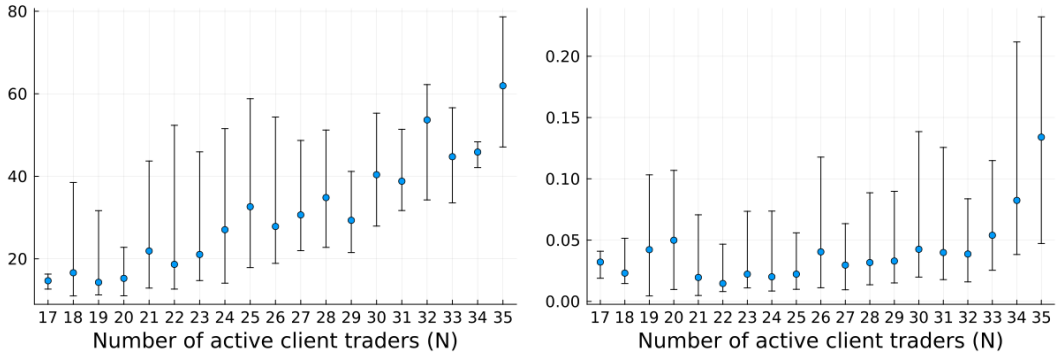
(d) Return variance, Σ_{ts}



Notes: Appendix Figure A2 shows the distribution of transaction prices (in C$), net supply (in thousand C$), the number of active traders, and my proxy for the variance of daily returns (Σ_{ts}) across all discretized "auctions". Each box plot reports TSX inventory (TSX–IN), TSX client (TSX–CL), Alpha inventory (Alpha–IN), and Alpha client (Alpha–CL) auctions separately. Boxes report the interquartile range, whiskers exclude outliers, and individual points indicate outliers.

Appendix Figure A3: Amount traded and return variance by number of traders

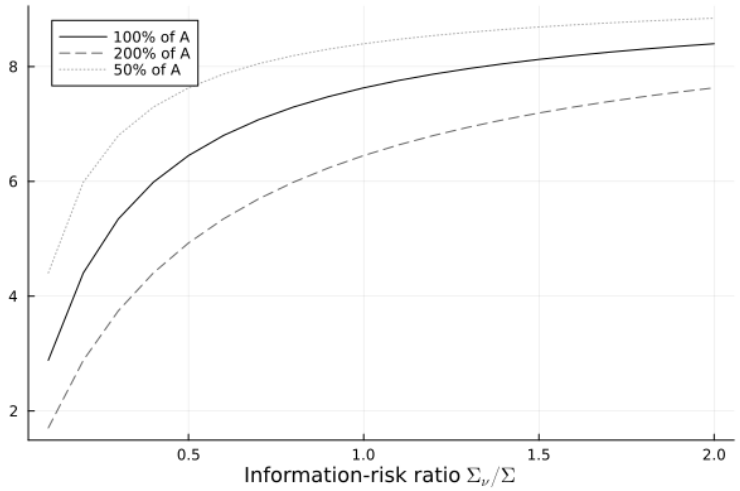
(a) Total amount traded $|A|$ (b) Return variance Σ



Notes: Appendix Figure A3a shows the distribution of the average daily amount traded in a client auction (either on the buy- or sell-side) on TSX per number of traders, N, active in that auction. The distribution is over days, securities, and sides of the auction. Figure A3b shows the distribution of return variances, Σ_{ts}, per N over days and securities. Supply is expressed in 1,000 shares (unites), and return variance in Canadian dollars.

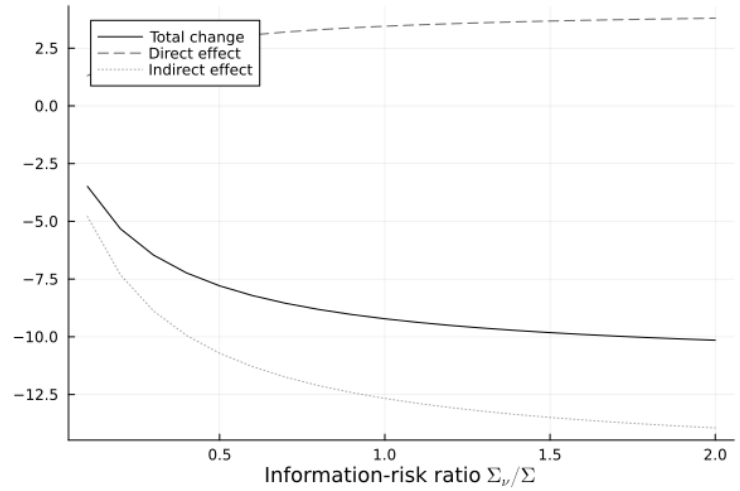
Appendix Figure A4: Welfare effects across information risk and supply

(a) Eliminating moral hazard



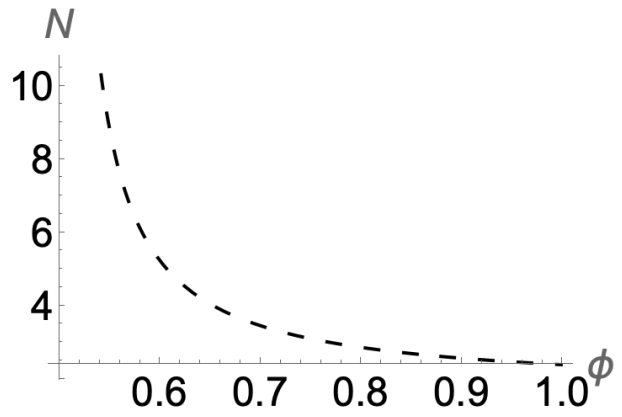
Notes: Appendix Figure A4a plots welfare effects—measured in percentage points relative to total welfare before relaxing the respective friction—for a representative TSX client-account auction, where clients and traders share the same risk tolerance (670), as a function of the information–risk ratio, $\kappa = \Sigma_\nu/\Sigma$. The curves correspond to different auction sizes, scaled relative to the median net supply: $A/2$, A, and $2A$. The percentage gain declines in supply, because the improvement from stronger incentives grows much more slowly with supply, A, than baseline welfare, which rises roughly linearly in $A\nu$.

(b) Adding one trader



Notes: Appendix Figure A4b shows the percentage welfare loss from increasing the number of active traders from N to $N+1$, decomposed into the direct effect of adding a trader (holding the contract fixed) and the indirect effect operating through the change in the contract, $\hat{\phi}^*(N+1) - \hat{\phi}^*(N)$. Because welfare in the model is affine in the information–risk ratio, both experiments produce similarly shaped curves in κ, and the informative differences lie in their levels, signs, and how they scale with auction size and contract changes.

Appendix Figure A5: When imperfect competition is worse than moral hazard ($\rho_I = \rho_H$)



Notes: Appendix Figure A5 shows cutoff function, $f(\phi)$, of Corollary 6 when $\rho_I = \rho_H$ in the dashed gray line. The x-axis shows ϕ from $\frac{1}{2}$, when moral hazard frictions are not binding, to 1, when all assets go to the trader to incentivize effort.