

Latest version.

China Walls

Tomy Lee Daniel Nathan Chaojun Wang*

This version: December 15, 2024

First version: November 20, 2024

Abstract

Conflicts of interest are inherent to banking conglomerates. Regulators increasingly manage these conflicts by enforcing *China Walls*—internal information barriers around key affiliates, in particular, dealers. We map information sharing among dealers and funds using a near universe of foreign exchange transactions involving the Israeli Shekel to evaluate if today’s China Walls are effectively enforced. We employ a difference-in-differences design comparing affiliates to entirely unconnected firms around exceptionally large trades to measure information sharing, and exploit the structure of this market to verify our design. We document islands of informational autarky between dealers and their affiliate funds surrounded by a sea of information sharing: (1) The affiliate dealers and funds never trade and do not share information with each other. (2) The dealers and funds connected via trading relationships systematically share information, including on days when a connected dealer and fund happen to not trade with each other. (3) Affiliate funds without China Walls intensely share information among themselves. (4) Our results hold during crisis and noncrisis periods, and across granular cells of firm and asset characteristics. Our results reveal remarkable regulatory capacity to control information flows within conglomerates.

*Lee is at the Central European University. Nathan is at the Bank of Israel. Wang is at the Wharton School, University of Pennsylvania. Emails: leeso@ceu.edu; daniel.nathan@boi.org.il; wangchj@wharton.upenn.edu.

1 Introduction

Banking conglomerates are rife with conflicts of interest. They manage funds and run broker-dealers that intermediate financial markets, all while investing on their own accounts. Potential for sharing of privileged information among affiliate firms underlies many of these conflicts. In response, regulators in the US and elsewhere increasingly enforce *China Walls*—blunt information barriers around broker-dealers—to preempt information sharing by the affiliates that are the most likely to cause conflicts of interest.¹

Enforcing the China Walls presents a formidable regulatory challenge: Information sharing among affiliates occurs in private, is plausibly deniable, and yields large conglomerate-wide payoffs. More fundamentally, the affiliates have tightly aligned incentives, precluding counterparty litigation that is central to regulatory enforcement in other spheres. On the other hand, today’s China Walls are a product of immense popular support for financial regulation that followed the 2008 crisis. Therefore, effectively enforced China Walls would reveal remarkable regulatory capacity to control information flows when backed by sufficient political will. Are the China Walls effectively enforced within banking conglomerates?

In this paper, we map the bilateral sharing of material private information to evaluate the effectiveness of today’s China Walls.² We exploit regulatory data on the universe of Israeli Shekel trades and the structure of the foreign exchange market, where the China Walls separate dealers from their affiliate funds, but not the affiliate funds from each other. The dealers heavily trade with funds (Figure 4a), yet virtually never with their

¹“China Walls,” or the more common “Chinese Walls,” is a reference to the Great Wall of China (Gozzi, 2003). “Information barriers,” “firewalls,” “ethical screens,” and “insulation walls” are synonymous terms that appear later. We adopt “China Walls,” because it is concise, does not have a common alternative meaning, and is the closest to the original reference.

²Information is material if it would affect the receiving firm’s trading behavior.

affiliate funds (Figure 4b): Between 2019-2024, the dealers trade a total \$5.51 with their affiliate funds. Our empirical design examines gross trade volumes around exceptionally large trades to identify bilateral information sharing from the originator of the large trade to the other firms. This design detects precisely null information sharing between the dealers and their affiliate funds. In stark contrast, we detect systemic information sharing between unaffiliated dealers and funds that often trade with each other, even on days they happen to not trade, verifying that our design reliably detects information sharing. The most intense information sharing occurs between the affiliate funds which are not walled off, consistent with affiliates having large incentives to share information absent the China Walls. All results persist in crisis and noncrisis periods and across a battery of key firm and asset characteristics. In sum, the China Walls sharply reshape the topology of information flows in financial markets.

Section 2 develops our empirical design, in which responses in gross trade volumes identify bilateral sharing of material information, and exceptionally large trades pinpoint when the affiliate dealers and funds are the most likely to violate the China Walls. Figure 1 sketches our test of the hypothesis that the affiliate dealers and funds do not share information. GS Dealer and GS Fund are affiliates. Unrelated Fund is neither an affiliate nor ever trades with the other firms in the figure. We compare the daily gross dollar volume of the GS Fund and the Unrelated Fund around an exceptionally large trade by the GS Dealer (the event). We conclude that the GS Dealer (the event firm) does not share information with the GS Fund (the affiliate firm) if the daily volume of the GS Fund does not change relative to the Unrelated Fund (the control firm) around the event day.

Three pieces of intuition guide our design. First, material information would prompt trading by the receiving firm. Thus, increases in gross volumes can proxy for material




Figure 1: Identification: Information Sharing from Dealers to Affiliate Funds

information flow while filtering out immaterial information. Second, exceptionally large trades would signal the arrivals of especially valuable material information (Kyle, 1985; Easley and O'Hara, 1987), and pinpoint when violating the China Walls would yield the largest conglomerate-wide payoffs from sharing among affiliates.

Third, we conservatively strip away confounding variation by choosing the unrelated firms as controls and confirming parallel trends. Three confounding channels allow the exceptionally large trades of the event firms to coincide with increases in the gross volumes of their affiliates. (I) The large trades may be correlated with public news or aggregate shocks that trigger all firms to trade. (II) The impact of the large trades or the information linked to them may indirectly reach the affiliates through other firms. (III) Sources common to the event firms and their affiliates, say a shared data or trading connection, may simultaneously induce the event trade and the affiliate volume. An event firm would not share information with an unrelated firm, which would nonetheless be

exposed to the aggregate shocks (I) and the indirect effects (II). Therefore, the unrelated firms control for the aggregate and the indirect effects while preserving any bilateral information sharing. Moreover, a common source (III) is likely to prompt trading by some affiliate firms before prompting the event trades, causing the affiliates' daily volumes to increase relative to the unrelated firms leading up to the event day. Thus, we conclude that the common sources do not contaminate our test if the affiliate and the unrelated firms exhibit parallel trends before the event day, giving meaning to the parallel trends assumption in our context.

Our rich setting provides two falsification tests. The first exploits the stylized fact that the dealers and funds with a trading relationship share information (Barbon, Di Maggio, Franzoni, and Landier, 2019; Boyarchenko, Lucca, and Veldkamp, 2021) to test if our design reliably detects information sharing. In Figure 1, BoA Fund frequently trades with the GS Dealer. The BoA Fund is the treated firm in this test: We compare the daily volumes of the BoA Fund (the connected firm) and the Unrelated Fund around the GS Dealer's (the event firm's) exceptionally large trade. To strengthen the test, we exclude any firms from the connected treatment that traded with the event firm on or after the event day. We conclude that our design reliably detects information sharing if the connected firms increases their daily volume relative to the unrelated firms on or after the event day.

The second test considers the necessity of China Wall enforcement. Where even the affiliates without the China Walls refuse to share information, there would be no need for their enforcement. We exploit the presence of affiliate funds—firms that are affiliates yet not walled off—to infer whether the walled-off affiliates would share information if the China Walls were absent. Applying a design similar to Figure 1 to the subsample

of funds, we conclude that the China Wall enforcement is necessary if the affiliate funds increase their daily volumes relative to the unrelated funds on or after the event day.

Section 3 explores the data. We obtain the near universe of nonanonymized *subsidiary-level* transactions in the Israeli Shekel (ILS), totaling 20 million trades from January 2019 to March 2024. The ILS market closely resemble the larger USD market, except that Israeli regulators do not enforce China Walls. Hence, our results reflect enforcement by foreign regulators, especially the US regulators with worldwide jurisdiction. Further, firms do not incriminate themselves when they report data inconsistent with their China Walls to the Israeli regulators, providing assurance that reporting fraud does not drive our results.

We motivate our main analysis using the correlations in daily gross dollar volumes of affiliate dealer-fund pairs, connected dealer-fund pairs, and unrelated dealer-fund pairs. The volumes of the unrelated pairs are correlated over wide leads and lags, attributable to the confounders. The affiliate pairs have correlations that are indistinguishable from the unrelated pairs. By contrast, the connected pairs have sharply higher same-day correlations, although the correlations exclude the days when the pair trades.

Section 4 implements our empirical design as stacked difference-in-differences specifications (Cengiz, Dube, Lindner, and Zipperer, 2019). The daily gross dollar volumes of the affiliate funds are statistically indistinguishable from the unrelated funds before, on, and after the days when a dealer makes an exceptionally large trade. Our estimates rule out increases above 0.53 standard deviations in the daily volume of affiliate funds at the 95% confidence level. In contrast, the daily volumes of the connected funds increase by 0.76 sd (std. error: 0.132 sd) relative to the unrelated funds on the event day. Comparing the dealers around the events of funds, we rule out increases in the daily volumes of the affiliate dealers exceeding 0.082 sd, at the 95% confidence level, relative to the unrelated

dealers on the days when a fund makes an exceptionally large trade. The connected dealers raise their daily volumes by 0.182 sd (std. error: 0.028) relative to the unrelated dealers on the event day. All specifications exhibit parallel trends. We conclude that the affiliate dealers and funds do not share material information, and that our design reliably detects material information sharing.

The remainder of [Section 4](#) applies the design to the subsample of funds. For the funds and event funds whose dealers do not overlap, the daily volumes of the affiliate funds increase by 1.80 sd (std. error: 0.248 sd) relative to the unrelated funds on the days when an event fund makes an exceptionally large trade. For the funds and event funds whose dealers overlap, the affiliate funds increase their daily volumes by a remarkable 3.15 sd (std. error: 0.347 sd) relative to the unrelated funds on the event day, confirming that common sources can contaminate measures of bilateral information sharing. In sum, the affiliate dealers and funds do not share material information, yet would intensely do so if the China Walls were absent. We conclude that the China Walls are, on the whole, effectively enforced.

[Section 5](#) splits firms into granular cells of firm and asset characteristics to scour for breaches of the China Walls. We do not detect information sharing across a battery of affiliate dealer-fund characteristics: for hedge funds or nonhedge funds, for large or small dealers or funds, whether the affiliate dealer and fund have the same or different specializations, and during crisis or noncrisis periods. The connected dealers and funds display robust information sharing, with the hedge fund response double that of the nonhedge funds, 1.92 sd (std. error: 0.831 sd) versus 0.758 sd (std. error: 0.135 sd), and larger responses where the connected dealer and fund specialize in the same currency or asset class. This result is consistent with [Di Maggio, Franzoni, Kermani, and Somnavilla](#)

(2019), and [Kumar, Mullally, Ray, and Tang \(2020\)](#), and our conjecture that material information triggers increases in gross volumes and more so for the more valuable information. Extending this analysis to the subsample of funds, we find universal information sharing among the affiliate funds, with or without overlapping dealers, generating reliably and sharply larger responses than the connected dealers and funds.

We proceed as follows. [Section 1.1](#) summarizes our contributions. [Section 2](#) develops the empirical design. [Section 3](#) describes the data and performs motivating analyses. [Section 4](#) investigates the effectiveness of China Walls. [Section 5](#) contains the heterogeneity analyses.

1.1 Related Literature

We belong to the literature on the capacity of democracies to regulate firm behavior. Rights and electoral turnover constrain democratic governments ([Persson and Svensson, 1989](#); [Alesina and Tabellini, 1990](#)). Nonetheless, their regulators greatly reduce pollution ([Keiser and Shapiro, 2019](#); [Behrer, Glaeser, Ponzetto, and Shleifer, 2021](#)), discrimination in pay ([Bailey, Helgerman, and Stuart, 2024](#)) and access to accommodation ([Cook, Jones, Logan, and Rosé, 2023](#)), insider trading ([Bhattacharya and Daouk, 2002](#)), and misleading financial disclosures ([Greenstone, Oyer, and Vissing-Jorgensen, 2006](#)), among other impacts. In many of these settings, the misaligned interests of parties involved in the regulated activity (e.g., employer vs employee, insider vs outside shareholder, incumbent vs prospective investor) expose violating firms to enforcement via private litigation by the involved counterparties ([Glaeser and Shleifer, 2003](#); [La Porta, Lopez-De-Silanes, and Shleifer, 2006](#)). In our setting, a China Wall violation would involve affiliates under common corporate control, obviating the threat of counterparty litigation. Moreover, the

possibility for plausibly deniable communication limit the scope for effective enforcement of private information sharing (Peluso, 2020). Therefore, our finding that the China Walls are effectively enforced reveal remarkable regulatory capacity in democracies to control firm behavior.

A policy debate rages over the effectiveness of China Walls (Webel, Carpenter, Gnanaiah, Jones, Labonte, Miller, Perkins, Shorter, and Weiss, 2017). Much of the empirical evidence focuses on the period before the surge in rule making since the 2008 crisis. This evidence points to extensive China Wall violations, as legal proceedings would confirm.³ For the period under the post-crisis regulations, Kondor and Pintér (2022) document greater profitability of funds as the proportion of their affiliate dealers' trades with highly connected clients increases. Haselmann, Leuz, and Schreiber (2023) show that bank-affiliated traders earn higher returns on a stock after the affiliate bank lends to the underlying firm than on other stocks.⁴ We contribute identification that isolates bilateral material information sharing and validates the empirical design in conditions where China Walls are absent and information flow is expected. Doing so yields precisely estimated and robust evidence that the China Walls effectively preempt information sharing.

We extend the empirical literature on the diffusion of information in financial markets. The dealers both extract information from their clients' order flow (Hortaçsu and Kastl, 2012) and leak information to their clients (Barbon et al., 2019; Boyarchenko et al., 2021; Chague, Giovannetti, and Herskovic, 2023). More broadly, the dealers act as the conduits

³Lehar and Randl (2006), Irvine, Lipson, and Puckett (2007), Seyhun (2008), Massa and Rehman (2008), Chen and Martin (2011), Ivashina and Sun (2011), Li (2018), and Li, Mukherjee, and Sen (2021) find evidence for China Walls violations in various settings. Griffin, Shu, and Topaloglu (2012) does not find such evidence between affiliate trading desks and investment banks. The latest in-sample year is 2013 among these papers.

⁴A non-methodological reason for the difference in our results with Haselmann et al. (2023) could be that their sample period covers 2012 to 2017, when the Dodd-Frank Act was gradually being implemented.

through which information diffuses throughout each dealer’s trading network (Di Maggio et al., 2019; Hagströmer and Menkveld, 2019; Kumar et al., 2020). Our results identify a stark void in this informational network driven by regulatory intervention, introducing policy impact on information diffusion to this literature. Methodologically, we add the China Walls as a promising source of variation in information flows that is especially relevant today, when the financial sector is highly concentrated.

2 Design

2.1 Context

China Walls refer to a collection of rules and physical barriers that aim to preempt the flow of material private information (MPI) to or from walled-off affiliate firms. An MPI is any information that (a) a reasonable investor would find important for her investment decisions and (b) is not publicly disclosed.⁵ For example, proprietary analysis, inside information, or private trade requests would constitute MPI. Typical China Walls ban require walled-off workplaces be isolated via separate entrances and opaque and soundproof barriers, and the monitoring and recording of walled-off employees’ communications.

New regulations since the 2008 crisis established China Wall requirements on broker-dealers and investment advisers within banking conglomerates.⁶ Today, the failure to

⁵Material non-public information (MNPI) is the more commonly referred type of information in law. The MPI includes analyses based purely on public information, whereas the MNPI expressly excludes such analyses. We use MPI rather than MNPI since proprietary analysis is valuable private information.

⁶A banking conglomerate is a financial conglomerate that includes a deposit-taking affiliate, as [Appendix A](#) elaborates. We write “dealer,” “broker,” and “broker-dealer” interchangeably. A broker matches the orders of clients without trading against the orders itself. A dealer trades against client orders. Every intermediary in the foreign exchange market is a broker-dealer, which brokers or deals at its discretion.

maintain sufficiently strict China Walls around these subsidiaries is a prosecutable offense, even without any evidence of MPI misuse. The US Securities and Exchange Commission (SEC) and other regulators routinely impose large fines for deficiencies in China Walls. [Appendix A](#) details relevant definitions, history and legal precedents, impact of the Dodd-Frank Act, and recent enforcement cases.

Empirical setting. The foreign exchange market is an over-the-counter market, in which trades occur between dealers or a dealer and its client. The dealers are long-lived, trades are non-anonymous, and most firms rely exclusively on one or a few relationship dealers. Hence, reputation concerns preclude behavior frequently seen in centralized markets, such as repeated order submissions without the intent to trade or splitting a large trade quantity into a rapid sequence of small orders. This market operates at high frequency, where news is near-instantaneously incorporated into exchange rates. Therefore, we do not expect private advantage from an MPI to last beyond a few trading days.

Our data covers the universe of Israeli Shekel (ILS) foreign exchange transactions, which we obtain from the Bank of Israel. The ILS market structure is identical to the other foreign exchange markets. More broadly, financial regulations in Israel are largely based on the US. Indeed, 87% of ILS transactions are for the USD-ILS pair and the ILS and the USD markets have the same largest dealers.⁷ A peculiar Israeli law forbids Israeli holding companies from owning both a dealer and a non-dealer investment firm, as the US Glass-Steagall Act did until its 1999 repeal.⁸ As such, the Israeli regulators neither mandate nor enforce the China Walls—the banking conglomerates do not incriminate themselves when reporting data at odds with their China Walls to the Bank of Israel. The

⁷The share of USD in our sample is remarkably close to the 85% of all foreign exchange transactions that involve the USD ([Somogyi, 2022](#)).

⁸Indeed, no Israeli dealer in our sample has an affiliate fund.

enforcers of the China Walls in our setting are the nonIsraeli regulators, especially the US Securities and Exchange Commission whose jurisdiction extends to all activities of the banking conglomerates active in the US (every conglomerate in our sample).

2.2 Empirical Design

We must overcome three challenges to test the hypothesis that the China Walls are effectively enforced. First, the China Wall violations may be violated in circumstances that our test does not examine. In particular, the test may neglect the circumstances when the China Wall violations are the most likely. Second, a proxy for information flows must isolate bilateral MPI sharing, yet also reliably detect them. Third, the walled-off affiliates perhaps would not share MPI even if the Walls were absent, in which case enforcement is unnecessary.

Defining events. To overcome the first challenge, we examine the times when an affiliate receives an especially valuable MPI—if the affiliate does not share MPI in this state, it is unlikely to share the less valuable MPI in the other states. Standard theory shows that a trader submits larger quantities when she holds more valuable private information (Kyle, 1985; Easley and O’Hara, 1987). A trader’s larger trades are indeed more predictive of returns (Kumar et al., 2020; Pinter, Wang, and Zou, 2024). Appendix B presents concurring evidence in our setting. Therefore, we let an event be a firm and a day when the firm makes an exceptionally large trade. These events pinpoint the arrivals of especially valuable MPI.

Identifying bilateral information sharing. To the second challenge, a proxy for bilateral MPI sharing from firm i to firm j must isolate information that is (i) material and (ii) bilaterally shared. The definition of MPI helps resolve (i): An information is material only

if it is important for determining the firms' optimal portfolios. Receiving an MPI would prompt firm j to rebalance its portfolio towards the new optimum. As j does so, its daily gross volume would increase, whereas this volume would not respond to immaterial information. Therefore, we choose the increases in the gross volume of j to proxy for the sharing of material information to j .

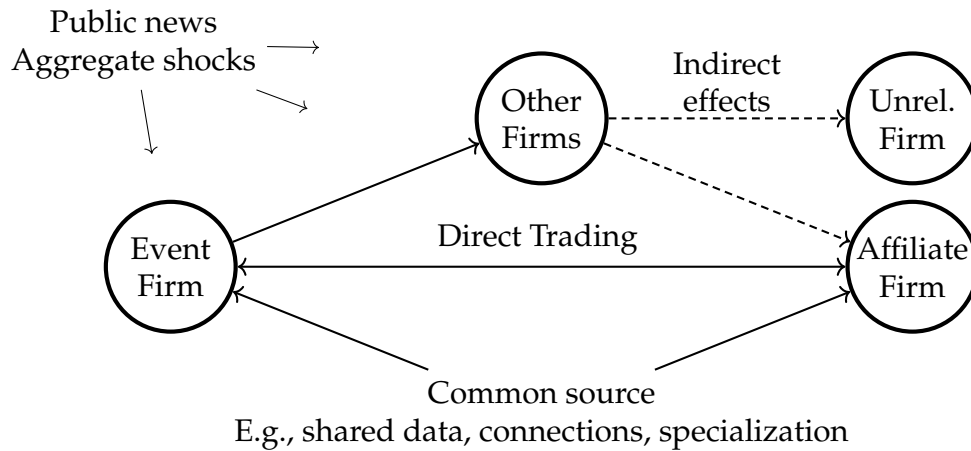


Figure 2: Potential Confounders to Measuring Bilateral Information Sharing

To resolve (ii), we rule out each of the four confounding alternatives to bilateral MPI sharing that can explain the coincidence in the exceptionally large trades of firm i and the heightened trading activity of firm j . **Figure 2** illustrates the alternatives. First, any direct trading between i and j could mechanically induce both, for instance as the large trade itself causes the heightened trading activity. This confounder does not apply when i and j are unrelated firms or are affiliate dealer and fund, since the former do not trade by definition and the latter do not in our data. We shut down this confounder in the case where i and j are connected dealer and fund by excluding j from our comparisons if j traded with i on or after the day (the event day) of the exceptionally large trade (the event trade).

Second, arrivals of public news or other aggregate shocks may trigger all firms to trade, sometimes triggering large trades. Third, the MPI of i corresponding to the event trade may indirectly induce the trading activity of j . Either the event trade itself or any sharing of the MPI by i to the firms other than j could affect the liquidity in the market or the information percolating among the other firms. These liquidity or informational impacts may trigger trading activity once they reach j . For example, a dealer that is the counterparty to the event trade might contact j to offload the newly gained inventory. Firm j may then trade more if this contact partly reveals the MPI to j or if j agrees to trade with the contacting dealer.

We filter out the aggregate-shock and the indirect-impact channels by comparing the gross volumes of j and the firms unrelated to i (i.e., those neither affiliated nor ever trade with i). The gross volumes of the unrelated firms would respond to the aggregate shocks and to the indirect impacts of i 's actions around the event day. Moreover, firm i would not bilaterally share the MPI with the unrelated firms. Hence, under the plausible assumption that directly receiving the MPI from i would induce greater trading activity than the MPI's indirect impacts, we conclude that i bilaterally shares the MPI with j only if the average difference between the gross volumes of j and an unrelated firm is positive on or after the event day. Our proxy for M_{ij} is the increase in the daily gross dollar volume of firm j relative to the unrelated firms around the event trade.

Fourth, a source common to i and j , and not to the unrelated firms, may simultaneously trigger i 's event trade and heighten j 's gross volume.⁹ We present three examples, in which we suppose that i and j are affiliated or connected with each other. If i and j are more likely than two unrelated firms to share a data or research subscription, then the

⁹Any effect of a source common to i , j , and the unrelated firms would be partialled out by our use of the unrelated firms as controls.

MPI that originates from the shared subscription could drive the correlation in i 's event trades and j 's gross volume relative to the unrelated firms. If i and j are more likely to specialize in the same asset, then the asset-specific shocks could generate the correlation by inducing i and j to trade around the same days. If i and j are more likely trade with the same third firm, then the correlation could be due to the third firm contacting i and j in response to its own inventory or information shocks. In each example, j would sometimes exhibit heightened gross volume before i makes the event trade: j might begin trading based on the MPI from the shared subscription or the asset-specific shock before i finds a counterparty for an exceptionally large trade, and the third firm might contact j before i . Therefore, we reject the presence of the common-sources channel if j and the unrelated firms show parallel trends in the days before the event day.

A key remaining threat is the possibility that our design does not reliably detect bilateral MPI sharing where it exists. We exploit the stylized fact that a connected dealer and fund extensively share information (Barbon et al., 2019; Kumar et al., 2020; Chague et al., 2023) to falsify the reliability of our design to detect bilateral MPI sharing. If our design is reliable, then we will detect the sharing whenever i and j are connected to each other. Thus, we falsify the reliability of our design if the daily gross volumes of connected firms do not increase relative to the unrelated firms on or after the event day. We strengthen this falsification test by excluding the connected firms that trades with the event firm on or after the event day.

Are China Walls necessary at all? To the third challenge, we exploit that each banking conglomerate owns multiple funds to infer whether the China Walls are necessary. A pair of affiliate funds belong to the same entity, yet are not walled off. Where the affiliate fund pairs bilaterally share MPI, we infer that the affiliate dealer-fund pairs would share MPI

if the China Walls were absent.

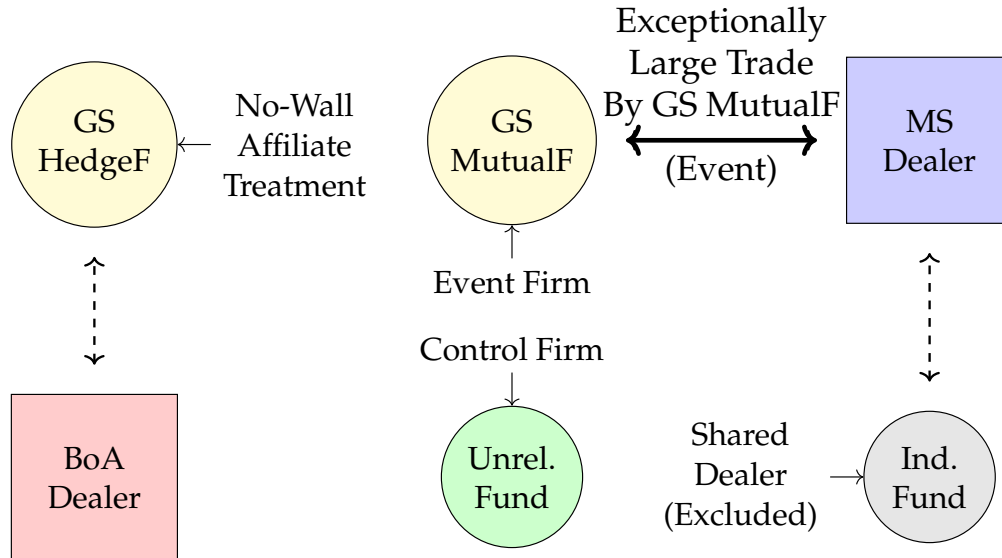


Figure 3: Identification: Information Sharing Between Affiliate Funds

Figure 3 depicts the corresponding design. Dotted arrows indicate trading relationships. GS Hedge Fund's sole dealer connection is BoA Dealer. GS Mutual Fund and the GS Hedge Fund are affiliate funds whose dealer connections do not overlap. All funds that trade with the MS Dealer are dropped, such as Independent Fund, to remove any confounding variation due to overlapping dealer connections. We compare the daily gross dollar volume of the GS Hedge Fund (the affiliate fund) to the Unrelated Fund around an exceptionally large trade by the GS Mutual Fund (the fund event). We conclude that the enforcement of China Walls are necessary if the daily volumes of the affiliate funds increase relative to the unrelated funds on or after the fund event day.

Several conspicuous differences between the affiliate fund pairs and the dealer-fund pairs threaten the validity of this inference. Specifically, a dealer and a fund are likely farther apart in size and in trading strategy than two funds. We partition the affiliate fund

pairs into granular cells of similar or greatly differing sizes and trading strategies to help address this threat to inference. We reject that the China Walls are unnecessary if and only if the gross volumes of the affiliate funds increase relative to the unrelated funds on or after the day when a fund makes an exceptionally large trade consistently across the cells of fund-event fund characteristics. We exclude the affiliate funds that frequently trade with a dealer with whom the event fund also frequently trades. Removing the effects of overlapping dealers this way prevents confounding variation due to common dealer connections, strengthening our inference.

2.3 Implementation

We adopt the stacked difference-in-differences specification with never-treated controls of [Cengiz et al. \(2019\)](#).¹⁰ An event is a firm and a date on which the firm made a trade in the 0.1 percentile of its trades by dollar value.¹¹ A firm is treated on or after the event day within the event window if the firm is an affiliate or a connection of the event firm. A firm is a control if it is unrelated to the event firm and not treated in any other event during the event window. Our event window is the 11 trading days around the event day, because exchange rates fully incorporate private information in about a trading week ([Menkhoff, Sarno, Schmeling, and Schrimpf, 2016](#)).

¹⁰Two-way fixed-effect DiD designs may assign negative weights to individual treatment effects in their ATT estimates if the treatments are staggered and have dynamic effects ([Roth, Sant’Anna, Bilinski, and Poe, 2023](#)). Our ATT estimates always assign positive weights to all individual treatment effects ([Gardner, 2022](#)).

¹¹For a firm with fewer than 1000 trades, an event is the day of the firm’s largest trade. If a firm makes multiple trades in its top 0.1 percentile on the same day, those trades are combined into a single event.

Our first regression specification is

$$Y_{ejt} = \sum_{\tau=-5}^5 \alpha_{\tau} \mathbb{1}_{t=\ell_e+\tau} \textit{Affiliate}_{ej} + \delta_{ej} + \varphi_t + \sum_{\tau=-5}^5 \gamma_{\tau} \mathbb{1}_{t=\ell_e+\tau} + \varepsilon_{ejt}. \quad (1)$$

The dependent variable Y_{ejt} is the gross dollar volume of firm j on calendar date t and event e , standardized at the firm level. The affiliate treatment dummy $\textit{Affiliate}_{ej}$ equals 1 if firm j is an affiliate of the event firm. The dummy $\textit{Affiliate}_{ej} = 0$ if j is unrelated to the event firm and is not treated in any other event within the window of event e .¹² The indicator variable $\mathbb{1}_{t=\ell_e+\tau}$ equals 1 when t equals the event day ℓ_e shifted by τ days, and 0 otherwise. We control for event-by-firm, calendar date, and event date fixed effects δ_{ej} , φ_t , and γ_{τ} . These effects control for event-and-firm-specific factors as well as common trends over calendar and event times. We cluster standard errors by event-and-firm and by calendar date, because our treatments are assigned event-by-firm and the incidence of events varies over time. Our data contains the near universe of transactions in the currency pairs we examine, as detailed in [Section 3](#), implying a high sampling probability. Therefore, the clustered variances likely approximates the true variances ([Abadie, Athey, Imbens, and Wooldridge, 2023](#)).

The second specification repurposes [Equation \(1\)](#) to measure the MPI sharing between connected dealers and funds,

$$Y_{ejt} = \sum_{\tau=-5}^5 \beta_{\tau} \mathbb{1}_{t=\ell_e+\tau} \textit{Connected}_{ej} + \delta_{ej} + \varphi_t + \sum_{\tau=-5}^5 \gamma_{\tau} \mathbb{1}_{t=\ell_e+\tau} + \varepsilon_{ejt}. \quad (2)$$

The connected treatment dummy $\textit{Connected}_{ej}$ equals 1 if (a) firm j trades 10 or more times

¹²To do so, we exclude from the control group any firm that is treated in another event that occurs in the 21-day panel around event e .

with the event firm in the sample, and (b) does not trade with the event firm on the event day and five days afterwards, $t = \ell_e, \dots, \ell_e + 5$. Condition (a) restricts the connected firms to nonaffiliates, because exactly zero pair of affiliate dealer and fund trades 10 or more times. Condition (b) removes any mechanical increase in the gross volumes of the connected firms relative to the unrelated firms due to trades with the event firm. The conditions for $Connected_{ej} = 0$ and $Affiliate_{ej} = 0$ are identical, and the other elements in Equation (2) are the same as the corresponding elements in Equation (1).

We estimate each of Equations (1) and (2) twice. Either the dealers are the event firms and we examine the daily gross volumes of the funds, or the funds are the event firms and we examine the volumes of the dealers.

The third specification, applied solely to the subsample of funds, is

$$\begin{aligned}
Y_{ejt} = & \sum_{\tau=-5}^5 \nu_{\tau} \mathbb{1}_{t=\ell_e+\tau} Affiliate_{ej} + \delta_{ej} + \varphi_t + \sum_{\tau=-5}^5 \gamma_{\tau} \mathbb{1}_{t=\ell_e+\tau} \\
& + \sum_{\tau=-5}^5 \kappa_{\tau} \mathbb{1}_{t=\ell_e+\tau} Affiliate_{ej} DealerOverlap_{ej} \\
& + \sum_{\tau=-5}^5 \eta_{\tau} \mathbb{1}_{t=\ell_e+\tau} DealerOverlap_{ej} + \varepsilon_{ejt}.
\end{aligned} \tag{3}$$

The control dummy $DealerOverlap_{ej}$ equals 1 if the set of dealers with whom fund j trades at least 10 times in the sample overlaps with the event fund's analogous set of dealers, and equals 0 otherwise. Our focus is on the coefficients ν_{τ} , which measure the MPI sharing from the event funds to their affiliate funds without an overlapping dealer. Separate event-date effects, γ_{τ} and η_{τ} , flexibly control for any trend over event time specific to the funds with or without an overlapping dealer.

2.4 Identification Tests

We assume that a trade that is exceptionally large for the responsible firm indicates an arrival of especially valuable material private information at the firm. The underlying claim is that a firm submits larger orders when it has more valuable private information. This claim is consistent with standard theory (Kyle, 1985; Easley and O'Hara, 1987) and recent empirical evidence that the larger trades by a firm are more predictive of returns than the firm's smaller trades in over-the-counter markets (Kumar et al., 2020; Pinter et al., 2024). On the other hand, the theory on order splitting (Bernhardt and Hughson, 1997) and the lack of similar evidence specifically on the foreign exchange market question the claim.

Appendix B adjudicates these claims in the data. Placebo tests using small and medium trades to define events can falsify our assumption. We define a small event as a firm and a day when the firm makes a trade in the 99.9 to 100th percentile of its trades by dollar volume, and a median event as the same except in the 50 to 50.1st percentile. Any increase in the gross volume of the connected firms relative to the unrelated firms on or after the small-event or the median-event days would suggest that ordinary trades contain MPI. This finding would falsify our assumption that the large trades are especially indicative of MPI arrival. We find zero evidence of such increase across all specifications.

3 Data and Descriptive Results

3.1 Data

We obtain the near universe of foreign exchange transactions involving the Israeli Shekel from the Bank of Israel (the Bank) in the sample period January 2019 to March 2024, spanning 1,368 trading days.¹³ Each observation specifies the currency pair (ILS and another currency), price, date and time, asset class (spot, forward, swap, or option), and the counterparty names. We exclude options due to insufficient observations and convert all nonUSD transaction values into USD at the contemporaneous official exchange rate published by the Bank.

Table 1 summarizes the samples we use in our analyses. A two-step process generates these samples. First, we consolidate the dealers up to the conglomerate by dropping all transactions between affiliate dealers and combining the affiliate dealers into conglomerate-level labels. The dealers are free to split incoming orders or transfer assets and capital with affiliate dealers. Consolidating the affiliate dealers minimizes noise from nonmarket transactions among them.¹⁴ Second, we aggregate the transactions of each dealer and each fund into daily gross dollar volumes, summing across asset classes. For each swap transaction, we only keep the notional amount from the first leg to avoid double counting. We winsorize the largest 0.5 percentile of observations by daily gross

¹³All Israeli firms, including the Israeli branches of conglomerates, must report each ILS transaction to the Bank. Non-Israeli firms fall under the same reporting requirement if their foreign exchange transactions in the previous year exceed \$15 million per day on average, whether on their own accounts or on behalf of clients. This reporting requirement applies to practically all significant financial firms, because any foreign currency spot or derivative transaction is included in the reporting threshold, even if the firm rarely trades ILS. Rules can be retrieved from <https://www.boi.org.il/en/economic-roles/statistics/reports-to-bank-of-israel/reporting-on-activity-in-the-foreign-currency-derivative/>.

¹⁴Some 8% of foreign exchange spot trades are “back-to-back” trades between affiliate dealers for purely accounting or inventory rebalancing reasons (Bank for International Settlements, 2022). All transactions by affiliate funds would be market-based, since they only trade with nonaffiliate dealers.

volume separately for the dealer and the fund subsamples.¹⁵

Affiliations. A four-step procedure identifies the affiliations of all firms. First, we determine the affiliations of most US-based firms using the quarterly organizational hierarchy data accessible via the National Information Center (<https://www.ffiec.gov/npw/>). Second, all firms with obviously indicative names are linked to the indicated conglomerate (e.g., “Deutsche Bank Luxembourg S.A.”). Third, the remaining firm names are entered into ChatGPT 4.0 as a query in the form, “as of [date the firm last appears in the sample], is [firm name] independent? If not, which holding company does [firm name] belong to?”¹⁶ Fourth, we manually check each answer generated in step three.¹⁷

3.2 Motivating Analyses

Three analyses motivate our main empirical design. First, **Figure 4** plots the total daily dollar volume of transactions across all pairs of a dealer and a fund. The dealers trade USD2.8 billion with the funds daily, while almost never trading with their affiliate funds—in total there are 4 affiliate dealer-fund trades worth USD5.51.

Second, **Figure 5** computes the correlation in daily gross volumes within unrelated dealer-fund pairs. For each lag $l = -10 \dots +10$ and a pair of dealer i and fund j that are nonaffiliates and do not trade in the sample, we compute the correlation $CorrGV_{ijl}$ between the date t gross volume of i and date $t + l$ gross volume of j . We average this

¹⁵The dealer or the fund subsamples have dramatically different distributions. The smallest observations are not winsorized, since gross volume is bounded below at zero.

¹⁶We keep affiliations of a given firm name constant throughout the sample period, because financial firms rarely change their affiliations and do so under new names when they do. Our results do not change after excluding the firms affiliated to Credit Suisse—the largest group to switch affiliations during our sample period.

¹⁷Verification is typically immediate upon searching for the firm name paired with either “independent” or the ChatGPT-suggested holding company name.

Table 1: Sample Characteristics

	All trades	Fund trades	Final Sample	
			Dealer day	Fund day
Mean daily volume (USD millions)	29,510	2,843		2,843
Mean daily volume across firms (USD millions)	641.52	3.66		3.66
Mean trade size (USD millions)	2.72	1.68	2.70	1.68
Currency				
USD	0.87	0.76		0.76
JPY	0.07	0.22		0.22
EUR	0.02	0.02		0.02
Asset class				
Spot	0.36	0.50		0.50
Forward	0.13	0.40		0.40
Swap	0.50	0.10		0.10
Unique firms				
Conglomerate	46	46	46	46
US	15	15	15	15
Israeli	11	11	11	11
Dealer	229	229	46	46
US	92	92	15	15
Israeli	15	15	11	11
Fund	7,775	7,775	7,775	7,775
Independent	6,660	6,660	6,660	6,660
US	4,826	4,826	4,826	4,826
Israeli	192	192	192	192
Hedge funds	632	632	632	632
Observations	20,832,686	2,762,406	50,570	490,408

All trades: Raw data set containing the near universe of Israeli Shekel transactions. *Dealer day*: Daily total dollar volumes of dealers in USD, excluding trades between affiliate dealers. *Fund day*: Daily total dollar volumes of funds in USD. *Mean daily total* is the average daily total dollar volume in USD billions. *Mean trade size* is the average dollar volume per transaction. *Mean daily volume* is the average across the daily dollar volume across firms. All Currency and Asset class figures are weighted by dollar volume.




Figure 4: Daily Gross Dollar Volumes Between Dealers and their Nonaffiliate Funds

Figure 4a: The sum of daily gross dollar volume in USD millions across pairs of dealer and fund that are not affiliated with the same banking conglomerate. **Figure 4b:** The sum of daily gross dollar volume in USD millions across pairs of affiliate dealer and fund. Shaded regions mark the onsets of the Covid pandemic, the Russian Invasion of Ukraine, and the Hamas attack on Israel.

correlation across the unrelated dealer-fund pairs for each l .

Figure 5a plots the results. There are strongly positive and significant correlations in trading activity among the unrelated dealers and funds. Absent a control group, the common shocks driving comovement among the unrelated firms may severely contaminate measures of bilateral information sharing.

Third, we estimate a simplified version of our main specifications (1)-(2). We compare the correlations $CorrGV_{ijl}$ within the affiliate and the connected dealer-fund pairs against the unrelated pairs. Doing so tests whether the trading activities of affiliates and connected firms correlate once stripped of common shocks. Our implementation uses the regression specification

$$CorrGV_{ijl} = a_i + b_j + c_l Affiliate_{ij} + d_l Connected_{ij} + \varepsilon_{ijl}. \quad (4)$$


The dummy variable $Affiliate_{ij}$ equals 1 if dealer i and fund j are affiliates and 0 if they are unrelated. The dummy $Connected_{ij}$ equals 1 if i and j trades 10 or more times in the

sample and 0 if they are unrelated. We exclude the trades between i and j to compute $CorrGV_{ijl}$, which avoids mechanical correlations due to within-pair trades. The dealer and the fund effects a_i and b_j control for time-invariant factors specific to each dealer and each fund.


Figure 5b plots the coefficients c_l and d_l across $l = -10 \dots 10$. Trading activities of the affiliate dealers and funds are no more correlated than the unrelated dealers and funds across l . In stark contrast, the connected dealers and funds are significantly more correlated than the unrelated firms for every l . These correlational results suggest that the China Walls effectively block material information flows among walled-off firms, even while information freely flows among connected firms. Our main empirical design isolates bilateral sharing of material information and focuses on the dates with the greatest incentive to share information.

4 Are China Walls Effectively Enforced?

We first estimate **Equations (1) and (2)** selecting the dealers as the event firms and the funds as the treated and the control firms. **Figure 6a** plots in blue the differences α_τ in standardized gross volume between affiliate and unrelated funds around the days of exceptionally large trades by dealers, and in red the differences β_τ between the connected and the unrelated funds. The affiliate funds exhibit neither pretrends nor posttrends. The connected funds show no pretrends and a positive estimate on the event day. The event-day estimates are far apart: The affiliate funds increase their event-day gross volume by no more than 0.53 standard deviations with 95% confidence, whereas the connected funds increase theirs by highly statistically significant 0.76 sd.



(a) Correlation Within Unrelated Pairs



(b) Affiliate and Connected Pairs

Figure 5: Correlations in Gross Volume

We interpret [Figure 6a](#) as follows. The exceptionally large trades pinpoint the arrivals of especially valuable MPI at the dealers, and receiving MPI would prompt increases in trading activity. The null posttrend of the affiliate funds implies that the dealers do not share the especially valuable MPI with their affiliate funds. The positive posttrend of the connected funds means that the dealers obtain the MPI on the days that their connected funds exhibit heightened trading activity. We partition how this coincidence of MPI and increase in gross volume could arise into the four channels other than bilateral MPI sharing. The MPI may induce the dealers to trade with the connected funds, in which case the coincidence would be mechanical. An aggregate shock affecting all firms could simultaneously cause both the event trade and the increase in gross volume. The MPI, the event trade, and related trading or information sharing by the event dealers may indirectly affect the connected funds as the dealers' actions percolate throughout the market. There may be common shocks specific to the connected dealers and funds, perhaps because they tend to share sources of information or common thirdparty connections.


The mechanical channel is ruled out by the exclusion of funds that traded with the event dealer on or after the event day for each event. The aggregate-shock and the indirect-effect channels are stripped away by the unrelated fund control group, since the unrelated funds would be exposed to the aggregate shocks and the indirect effects of the dealers' actions. This control would preserve any increase in gross volume due to bilateral MPI sharing, because the dealers would not share MPI with the unrelated funds. The common-shocks channel is rejected by the parallel pretrend, as the shocks common to the connected dealers and funds would sometimes cause the connected funds' gross volumes to increase before the event dealers make their exceptionally large trades. Altogether, only the bilateral sharing channel remains. We conclude that the dealers do not

bilaterally share MPI to their affiliate funds.


Figure 6b presents the coefficient estimates of **Equations (1) and (2)** where we examine the standardized daily gross volumes of the dealers around the days when a fund makes an exceptionally large trade. In blue are the differences α_τ in the volumes between the affiliate and the unrelated dealers around the event days. In red are the differences β_τ between the connected and the unrelated dealers. Neither the affiliate nor the connected dealers exhibit pretrends. The affiliate dealers do not show posttrends, and we can reject an increase in their gross volumes exceeding 0.082 sd on the event day at the 95% confidence level. The connected dealers increase their gross volumes by a highly statistically significant 0.189 sd on the event day. We conclude the funds do not bilaterally share MPI to their affiliate dealers.

Affiliates without China Walls. **Figure 6** establishes that the affiliate dealers and funds do not share material information. (And that, if they did, our design would reliably detect it.) One interpretation is that the China Walls are effectively enforced. The alternative is that the affiliate dealers and funds would not share MPI even if the China Walls were absent, rendering their enforcement unnecessary. We exploit that affiliate funds are not walled off from one another to infer whether the walled-off affiliates would share information absent the China Walls.

Figure 7 presents the results from **Equation (3)** estimated on the subsample of funds. In green are the differences ν_τ in standardized gross volume between the affiliate funds and the other funds whose dealer connections do not overlap with the event fund around exceptionally large trades by a fund. Despite removing the common shocks through any overlapping dealers, the affiliate funds increase trading activity by 1.803 standard deviations on the event date. This response is precisely estimated (standard error 0.248) and



(a) Fund Responses to Dealer Information



(b) Dealer Responses to Fund Information

Figure 6: Coefficient Estimates from Equations (1) and (2)

more than twice the response of the connected funds to dealer events in Figure 6a, consistent with greater willingness to share material information with affiliates than with outside firms. In magenta are the differences $\nu_\tau + \kappa_\tau + \eta_\tau$ between the affiliate funds whose dealer connections do overlap with the event fund and the nonaffiliate nonoverlapping funds. As one might expect, incorporating overlapping dealer effects dramatically raises the event date response, to 3.154 sd.




Figure 7: Affiliate Fund Response to Event Fund Information

Based on the results of Figures 6 and 7, we conclude that the China Walls are effective on the whole. Table 2 details the pooled regression counterparts to Figures 6 and 7. The affiliate funds have precisely null response to the arrival of especially valuable informa-

tion at the dealers and the converse for the affiliate dealers to the funds.¹⁸ In contrast, the connected dealers and funds respond strongly to each other’s information, with estimated coefficients in the multiples of the affiliate coefficients. By far the most responsive are the funds to the information from their affiliate funds. Altogether, the pooled results confirm that the China Walls are effective enforced.

5 Heterogeneity and Robustness

Our heterogeneity exercises aim to test the robustness of the China Walls. It is particularly important to test the robustness of our affiliate fund-to-fund results: Where even the affiliate funds only share MPI under special contexts, there is high likelihood that the affiliate dealers and funds would not share MPI absent the China Walls.

To do so, we repeat the analyses of [Section 4](#) across granular cells of jurisdiction, firm size, fund types, and currency and asset class specializations. We also explore the possibility that the affiliate dealers and funds violate their China Walls specifically during crisis period, when information is especially valuable. We create dummy variables corresponding to each characteristic, add the complete set of interaction terms to [Equations \(1\) to \(3\)](#), then reestimate the expanded regression specifications. Of the characteristic dummies, half are for the event firms and the others are for the treated and the control firms. Precisely, the dummy *HedgeFund* equals 1 if the treated or the control firm is a hedge fund, and *HedgeFundEvent* = 1 if the event firm is a hedge fund. Other dummy variables indi-

¹⁸The dealer-to-fund specification has far fewer events and observations than the fund-to-dealer specification because most funds have zero volume on most days—the median number of trades by a fund is 18 over the whole sample. Thus, majority of dealer events get dropped due to missing a treated or a control fund. Every dealer trades every day so few fund events are dropped. We verify that our filters do not materially affect the gap in the numbers of events and observations.

Table 2: Pooled Regression Estimates

	D2F Affiliate	F2D Affiliate	D2F Connected	F2D Connected	F2F Affiliate
$Post \times Affiliate$	0.0163 [0.2011]	0.0221 [0.0311]			0.2360*** [0.0216]
$Post \times Connected$			0.1515*** [0.0453]	0.0520*** [0.0156]	
$Post \times DealerOverlap$					0.0075** [0.0039]
$Post \times Affiliate$ $\times DealerOverlap$					0.3158*** [0.0333]
Event*Firm FE	Yes	Yes	Yes	Yes	Yes
Date and Event Date FE	Yes	Yes	Yes	Yes	Yes
Within R-squared	0.0023	0.0001	0.0001	0.0002	0.0002
Events	32,831	264,985	28,538	269,288	992,728
Observations	361,141	2,878,748	313,055	2,925,568	10,904,439

Coefficient estimates from [Equations \(1\) to \(3\)](#). The dependent variable is the standardized daily gross US dollar volume of a firm winsorized at the top 0.5 percentile. Standard errors in square brackets are clustered at the event-by-firm and date levels. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

D2F: Dealers are the event firms and funds are the treated and the control firms. *F2D*: Funds are the event firms and dealers are the treated and the control firms. An event is a firm and a day when the firm made a trade in the 0.1 percentile among its trades. The treatment $Affiliate = 1$ if the firm is an affiliate of the event firm. The treatment $Connected = 1$ if the firm trades at least 10 times with the event firm in the sample, and does not trade with the event firm on the event day nor on the five subsequent trading days. The two treatments are mutually exclusive, because affiliate dealer-fund pairs do not trade. All control firms, whose $Affiliate = Connected = 0$, are neither affiliated nor ever trades with the event firm.


cate whether a firm is located in the US, or has a higher share of USD or swap trades than the median. We call the high USD share firms simply “USD firms” and similarly “Swap firms.”

Figure 8 depicts the separate effects of hedge funds and nonhedge funds as our leading heterogeneity exercise. **Figure 8a** reports the responses of hedge funds and nonhedge funds to dealer events. Its nonhedge fund estimate is the coefficient on $\mathbb{1}_{t=\ell_e+\tau}Treated_{ej}$, where $Treated_{ej}$ is $Affiliate_{ej}$ or $Connected_{ej}$ depending on the specification. The hedge fund estimate is the sum of this coefficient and the those on $\mathbb{1}_{t=\ell_e+\tau}Treated_{ej}HedgeFund_{ej}$ and $\mathbb{1}_{t=\ell_e+\tau}HedgeFund_{ej}$.¹⁹ **Figure 8b** does so symmetrically for the dealer responses to hedge fund and nonhedge fund events, and **Figure 9** partitions the responses to and from hedge funds and nonhedge funds without dealer overlap into four cells.²⁰


Table 5 compares the coefficient estimates for the events during crisis and noncrisis periods. The crisis periods span Covid (February 1st to March 31, 2020), the Russian Invasion of Ukraine (February 16 to March 8, 2022), and the Hamas Attack (September 27 to October 17, 2023).

¹⁹We use the constituent coefficients’ covariance matrices to compute the standard errors of their sums.

²⁰To obtain the HF Event-HF Firm estimates in **Figure 9**, we sum across the coefficients on $\mathbb{1}_{t=\ell_e+\tau}Affiliate_{ej}$, $\mathbb{1}_{t=\ell_e+\tau}Affiliate_{ej}HedgeFund_{ej}$, $\mathbb{1}_{t=\ell_e+\tau}Affiliate_{ej}HedgeFundEvent_{ej}$, and $\mathbb{1}_{t=\ell_e+\tau}Affiliate_{ej}HedgeFund_{ej}HedgeFundEvent_{ej}$. Similar steps yield all heterogeneity estimates. We omit the coefficients on the terms interacted with $DealerOverlap_{ej}$.



(a) Dealer-to-Fund Responses



(b) Fund-to-Dealer Responses

Figure 8: Dealer-and-Fund Responses: Hedge Funds vs Nonhedge Funds




Figure 9: Affiliate Fund-to-Fund Responses: Hedge Funds vs Nonhedge Funds

Table 3: Pooled Regression Estimates: USD vs NonUSD Firms

	D2F Affiliate	D2F Connected	F2D Affiliate	F2D Connected	F2F Affiliate
NonUSD Event Firm -NonUSD Firm	0.1453 [0.1950]	0.1220* [0.0752]	-0.0023 [0.0548]	0.0800 [0.1122]	0.2696*** [0.0267]
NonUSD Event Firm -USD Firm	0.0651 [0.1969]	0.0893 [0.0822]	-0.0009 [0.0562]	0.0825 [0.1130]	0.2745*** [0.0270]
USD Event Firm -NonUSD Firm	0.2914 [0.2632]	0.1142 [0.1088]	-0.0101 [0.0570]	0.0717 [0.1128]	0.2653*** [0.0268]
USD Event Firm -USD Firm	0.0874 [0.7148]	-0.0470 [0.1104]	-0.0840 [0.0696]	0.1544** [0.0720]	0.3096*** [0.0449]
Event*Firm FE	Yes	Yes	Yes	Yes	Yes
Date and Event Date FE	Yes	Yes	Yes	Yes	Yes
Within R-squared	0.0009	0.0002	0.0001	0.0001	0.0002
Events					
Observations	361,141	313,055	2,878,748	2,925,568	10,904,439

Coefficient estimates from [Equations \(1\) to \(3\)](#). The dependent variable is the standardized daily gross US dollar volume of a firm winsorized at the top 0.5 percentile. Standard errors in square brackets are clustered at the event-by-firm and date levels. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

USD Event Firm indicates an event by a firm with above-median share of USD-ILS transactions. USD Firm indicates a treated or a control firm with above-median share of USD-ILS transactions. *D2F*: Dealers are the event firms and funds are the treated and the control firms. *F2D*: Funds are the event firms and dealers are the treated and the control firms. An event is a firm and a day when the firm made a trade in the 0.1 percentile among its trades. The treatment *Affiliate* = 1 if the firm is an affiliate of the event firm. The treatment *Connected* = 1 if the firm trades at least 10 times with the event firm in the sample, and does not trade with the event firm on the event day nor on the five subsequent trading days. The two treatments are mutually exclusive, because affiliate dealer-fund pairs do not trade. All control firms, whose *Affiliate* = *Connected* = 0, are neither affiliated nor ever trades with the event firm.

Table 4: Pooled Regression Estimates: Swap vs NonSwap Firms

	D2F Affiliate	D2F Connected	F2D Affiliate	F2D Connected	F2F Affiliate
NonSwap Event Firm -NonSwap Firm	0.1216 [0.1700]	0.0699** [0.0313]	0.0058 [0.0312]	0.0746** [0.0293]	0.2332*** [0.0227]
NonSwap Event Firm -Swap Firm	0.0665 [0.1752]	0.1595 [0.0685]	0.0173 [0.0325]	0.0860** [0.0301]	0.2360*** [0.0237]
Swap Event Firm -NonSwap Firm	-0.0624 [0.2016]	-0.0967** [0.0433]	0.0265 [0.0372]	0.0952** [0.0351]	0.2271*** [0.0232]
Swap Event Firm -Swap Firm	1.953*** [0.7148]	-0.7186 [0.6750]	0.4491** [0.1567]	0.1020** [0.0498]	0.1458 [0.1074]
Event*Firm FE	Yes	Yes	Yes	Yes	Yes
Date and Event Date FE	Yes	Yes	Yes	Yes	Yes
Within R-squared	0.0009	0.0002	0.0001	0.0001	0.0002
Events					
Observations	361,141	313,055	2,878,748	2,925,568	10,904,439

Coefficient estimates from [Equations \(1\) to \(3\)](#). The dependent variable is the standardized daily gross US dollar volume of a firm winsorized at the top 0.5 percentile. Standard errors in square brackets are clustered at the event-by-firm and date levels. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Swap Event Firm indicates an event by a firm with above-median share of swap transactions. Swap Firm indicates a treated or a control firm with above-median share of swap transactions. *D2F*: Dealers are the event firms and funds are the treated and the control firms. *F2D*: Funds are the event firms and dealers are the treated and the control firms. An event is a firm and a day when the firm made a trade in the 0.1 percentile among its trades. The treatment *Affiliate* = 1 if the firm is an affiliate of the event firm. The treatment *Connected* = 1 if the firm trades at least 10 times with the event firm in the sample, and does not trade with the event firm on the event day nor on the five subsequent trading days. The two treatments are mutually exclusive, because affiliate dealer-fund pairs do not trade. All control firms, whose *Affiliate* = *Connected* = 0, are neither affiliated nor ever trades with the event firm.

Table 5: Pooled Regression Estimates: Crisis Periods

	D2F Affiliate	F2D Affiliate	D2F Connected	F2D Connected	F2F Affiliate
<i>Post</i> \times <i>Affiliate</i>	-0.1586 [0.2339]	0.0316 [0.0323]			0.3657*** [0.0290]
<i>Post</i> \times <i>Connected</i>			0.1312*** [0.0464]	0.0521*** [0.0159]	
<i>Post</i> \times <i>Affiliate</i> \times <i>Crisis</i>	0.2297 [0.6945]	-0.0115 [0.0802]			0.0833 [0.0943]
<i>Post</i> \times <i>Connected</i> \times <i>Crisis</i>			0.1410 [0.1128]	0.0001 [0.0309]	
Event*Firm FE	Yes	Yes	Yes	Yes	Yes
Date and Event Date FE	Yes	Yes	Yes	Yes	Yes
Within R-squared	0.0023	0.0001	0.0002	0.0001	0.0002
Crisis Events	3,325	13,497	2,966	13,579	51,350
Events	32,831	264,979	28,538	269,288	331,803
Observations	361,141	2,878,748	313,055	2,925,568	10,904,439

Coefficient estimates from [Equations \(1\) to \(3\)](#). The dependent variable is the standardized daily gross US dollar volume of a firm winsorized at the top 0.5 percentile. Standard errors in square brackets are clustered at the event-by-firm and date levels. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Crisis: Event occurred during the start of Covid, the Russian Invasion of Ukraine, or the Hamas-Israeli War. *D2F*: Dealers are the event firms and funds are the treated and the control firms. *F2D*: Funds are the event firms and dealers are the treated and the control firms. An event is a firm and a day when the firm made a trade in the 0.1 percentile among its trades. The treatment *Affiliate* = 1 if the firm is an affiliate of the event firm. The treatment *Connected* = 1 if the firm trades at least 10 times with the event firm in the sample, and does not trade with the event firm on the event day nor on the five subsequent trading days. The two treatments are mutually exclusive, because affiliate dealer-fund pairs do not trade. All control firms, whose *Affiliate* = *Connected* = 0, are neither affiliated nor ever trades with the event firm.

Appendix

A Detailed Context

This section provides detailed institutional context with a focus on the US.

A.1 Definitions

A *banking conglomerate* is a group of firms controlled by the same holding company and that includes a depository institution (i.e., a bank). A *financial conglomerate* is a broader term encompassing any such groups that includes firms offering financial services as its primary activity. We write “financial conglomerate” when discussing the period up to the 2000s, when most financial conglomerates became banking conglomerates, and “banking conglomerates” elsewhere.

Figure 10 summarizes the components of a banking conglomerate. Their services include deposits, lending, insurance, asset management (i.e., investing clients’ capital), proprietary trading (investing own capital), brokering (matching client orders) and dealing (absorbing client orders onto inventory), investment analysis and advising, underwriting (asset issuance), corporate advising (on mergers and acquisitions and other strategic decisions), and payments and trade finance. A conglomerate partitions these services into insurers, commercial banks (deposits, loans), investment banks (underwriting, corporate advising), investment funds (asset management), broker-dealers (brokering, dealing, analysis, proprietary trading), and investment advisers.

All regulations against the misuse or leakage of financial information target *material non-public information* (MNPI). Information is MNPI if its public disclosure would ap-




Figure 10: Stylized Banking Conglomerate and Relevant Legal Restrictions

Green dotted lines indicate restrictions on transactions and transfers: Banking laws, fiduciary duty to investors, and state-level insurance laws bar commercial banks, investment funds, and insurers from transferring capital to affiliates or trading with them at unfavorable terms. Red solid lines indicate the China Walls that aim to block the flow of information around subsidiaries in which conflicts of interest concentrate: Broker-dealers and investment advisers are required to prevent their employees interacting with the employees of affiliates. Orange fonts highlight the Volcker Rule restrictions on proprietary trading and ownership of hedge funds and private equity funds by banking conglomerates.

preciably affect market prices. In practice, common-law courts treat as MNPI any non-publicly disclosed information that reasonable investors in the relevant securities would find important for their investment decisions. For example, insider earnings information or outstanding order flows of clients are MNPI.²¹ Possessing, sharing, or acting on MNPI is not generically illegal. However, financial intermediaries owe legal duties over MNPIs, as we soon elaborate.

The *China Walls* are blunt internal barriers set around subsidiaries with especially high risk of MNPI misuse. The Walls include both physical barriers and rules, typically:

- Separate offices, elevators, and entry ways for walled-off affiliates, with opaque and soundproof physical barriers when located on the same floor.

²¹Analyses of MNPI are MNPI, whereas analyses of publicly available information are not.

- Cool-down periods for employees transferring between walled-off affiliates.
- Watch lists that prohibit employees from trading or advising on the listed securities.
- Records of every instance where an “over-the-wall” executive (who oversees multiple affiliates walled off from each other) receives MNPI from any subsidiary, and requirement that the executive recuse themselves from any business related to the MNPI.
- Monitor and retain all business-related emails and messages sent by employees, and review those containing MNPI.
- Contingency plans when MNPI leaks through the China Walls, and the appointment of officers responsible for enforcing the Walls and handling the contingencies.

These restrictions on employee interactions effectively ban transactions between walled-off affiliates.

A.2 Key Regulations on Banking Conglomerates

The markings in **Figure 10** indicate each key regulation on the banking conglomerates. Two concerns underlie the regulations. First, the conglomerates may divert publicly insured deposits or insurance premiums towards risky trades or to cross-subsidize affiliates, thereby shifting risk onto the state or the insureds. Second, the conflicts of interest inherent in combining intermediation, advisory, and trading functions could disadvantage retail investors and undermine trust in financial markets.

Three constraints on banking conglomerates address these concerns. First, a bank or an insurer cannot cross-subsidize affiliates. The US Regulation W (and similar rules

elsewhere) limit the outstanding value of bank-to-affiliate transactions to 20 percent of the bank's capital and 10 percent with any single affiliate.²² These trades must occur at prevailing market prices and under punitive collateral requirements. Moreover, banks cannot trade securities issued by its affiliates, accept them as collateral, nor guarantee a trade, loan, or securities issuance that involves an affiliate. Analogous rules on insurers, which are harmonized across the US yet enforced by state authorities, prevent their capital being used to subsidize affiliates.²³

Second, the Volcker Rule restricts banking conglomerates from proprietary trading and owning risky investment funds. Specifically, a banking conglomerate cannot use its own capital to make short-term profit-seeking trades. The Rule also limits its ownership stake and exposure to hedge funds and private equity funds. Broad exemptions apply. The Rule exempts the trades linked to market making by broker-dealers and any trade held for more than 60 days. Further, hedge funds and private equity funds active entirely outside the US are exempt and, within the US, a conglomerate may sponsor and control such funds if it holds less than 3 percent of the funds' assets. Therefore, most banking conglomerates contain hedge funds and considerable scope remains for bank-affiliated broker-dealers to trade on private information using own capital.

Third, as we elaborate next, the China Walls around broker-dealers and around investment advisers seek to minimize information leakage surrounding these firms. Statutes single out investment advisers for their large potential impact on investment decisions. The broker-dealers are singled out, because their role as intermediaries provide constant

²²Outstanding transaction value include loans, face value of guaranteed assets or liabilities, and gross purchases from affiliates. For example, purchasing \$1 million of an asset from an affiliate would raise the outstanding value by \$1 million until the bank sells \$1 million of the same asset back to that affiliate. (Sales to other affiliates or of other assets do not affect the outstanding value generated by this purchase.)

²³[NAICS Model Law; implementation]

stream of privileged information gleaned from their clients' orders. Under the argument that broker-dealers leaking this information to affiliate funds or receiving inside information from affiliates would place the investing public at a sharp disadvantage, preventing such information flows is necessary to maintain trust and participation in financial markets.

A.3 China Wall Enforcement Over Time

Origins. Under common-law tradition, insider trading on behalf of clients was encouraged. Brokers and dealers were expected to use all information that came into their possession, and further solicit inside information, to fulfill their fiduciary duty. This expectation was upended in 1961, when a landmark judgement held each conglomerate liable for damages incurred by the investing public due to trades based on its MNPI. The ruling demands that the intermediaries holding MNPI either publicly disclose or take no action whatsoever related to the MNPI. Subsequent court rulings placed the full burden of avoiding incompatible duties onto the conglomerates.²⁴

Financial conglomerates were in an impossible legal jeopardy. Beyond fiduciary duty and the new duty to the investing public, the agency principle requires the firms acting as agents to safeguard the private information of their principal (Tuch, 2014). Suppose a conglomerate owns a dealer and a mutual fund, and the dealer receives a large trade request from a client hedge fund—an MNPI. By fiduciary duty, the dealer ought to share

²⁴A typical case is [Black and Shearson, Hammill Co. \(1968\)](#) which rules, “conflict in duties is the classic problem encountered by one who serves two masters. It should not be resolved by weighing the conflicting duties; it should be avoided in advance [...] or terminated when it appears.” The judgement upheld awards of \$25 thousand (1968 dollars) each to two customers of a dealer, which sold debentures of a failing firm whose board included a partner at the dealer’s parent company. The conflicting duties were the dealer’s fiduciary duty to its customers and the partner’s duty to keep the inside information of the failing firm confidential.

this MNPI with the mutual fund for the benefit of the fund's investors. Yet, doing so would expose the conglomerate to liability if the mutual fund trading on the MNPI cause losses to some traders. This liability can be avoided only by publicly disclosing the hedge fund's trade request, in violation of the agency principle. These incompatible duties left financial conglomerates in near-permanent state of legal liability.

The China Walls provided a way out. In 1968, the US Securities and Exchange Commission (SEC) began offering safe harbor from liability for the conglomerates that implement sufficiently strict China Walls, as determined by the SEC.²⁵ The logic is that walled-off subsidiaries can be considered separate entities for the purpose of determining whether a legal duty has been breached. Continuing the example, the dealer would not owe fiduciary duty to the investors of the affiliate mutual fund if this fund were walled off from the dealer. The US financial conglomerates widely adopted the China Walls, which became broadly standardized according to SEC guidelines. Financial conglomerates in other jurisdictions followed, whether through their US operations or regulatory standardization (in Australia, Canada, France, Germany, Japan, Switzerland, and the UK).

Pre-2008 crisis legal status. A 1980 US Supreme Court case replaced the constellation of duties with one overarching duty to "disclose or abstain." A person has the duty to disclose or abstain from acting on an MNPI when: (a) she owes fiduciary duty to the source of the MNPI; and (b) the action would give her a personal benefit.²⁶ Today's legal standard thereby narrows where the China Walls are valuable to the subsidiaries whose

²⁵Alternative means to avoid incompatible-duty liabilities, such as obtaining client consent to waive fiduciary duties, are likely ineffective under most circumstances (Tuch, 2014).

²⁶A recipient of an MNPI inherits any duty to disclose or abstain. Moreover, a sender is liable for any breach of duty by all recipients whose MNPI traces back to the sender. Suppose a client shares an MNPI with her dealer, which shares it with a non-financial company, which does so with a fund. The company, then the fund, inherits the duty to disclose or abstain. If the fund trades on the MNPI without disclosing it, the fund, then the company, and so the dealer become liable.

(i) clients routinely provide MNPI and (ii) affiliates would personally benefit from the MNPI.

The 1980s also saw the deregulation of financial conglomeration in the US and the UK. The arguments were that full-service financial conglomerates would generate economies of scope and be more competitive versus less regulated foreign competitors. Because the duty to disclose or abstain might render full-service conglomerates nonviable, new statutes explicitly incorporated the China Walls as safe harbor and broadened their legal protections (Brooke, Burrows, Faber, Harpum, and Silber, 1995, p. 98).²⁷ Suppose a fund consistently earns large profits whenever an affiliate dealer receives large order flows. Under the new statutes, presence of a China Wall between the dealer and the fund would protect the conglomerate against liabilities to the dealer's clients and to the fund's counterparties.²⁸

Pre-2008 crisis regulatory regime. The China Walls were initially an legal benefit available to the banking conglomerates—not a regulatory requirement. As such, the China Walls enforcement was purely reactive, occurring in the course of assigning liability upon the discovery of fraud or breach of duty. Indeed, no US regulator proactively evaluated the China Walls between 1990 and 2012, the years when the SEC reviewed the Walls within broker-dealers as a research exercise.²⁹ The prosecutions over the LIBOR scandal

²⁷The UK removed most restrictions on financial conglomeration in 1986. The US gradually weakened the Glass-Steagall Act provisions throughout the 1980s and 90s, until largely repealing the Act in 1999. The UK Financial Services Act 1986 (FSA) and the US Insider Trading and Securities Fraud Enforcement Act 1988 (ITSFEA) explicitly provide safe harbor from a wide range of liabilities to the financial conglomerates that adopt China Walls.

²⁸The China Walls grant similar protection elsewhere. For instance, in a landmark Australian case, *ASIC v. Citigroup* (2007), Citigroup's trading arm purchased one million shares of a target firm one day before its acquisition announcement, in a deal where Citigroup's investment bank was advising the acquirer. The judge dismissed the case, on the basis that the China Wall between Citigroup's trading and investment bank arms was sufficient to preclude conflict of interest (Hanrahan, 2007).

²⁹The 1990 review was in response to the 1998 ITSFEA Act that explicitly gave safe harbor to walled-off broker-dealers. The 2012 review was in response to the Dodd-Frank Act.

highlights the non-obligatory status of China Walls pre-crisis: While each settlement with an implicated banking conglomerate often delves into its China Walls, the sole purpose of doing so were to determine the degree of the conglomerate's legal liability for fraud and insider trading. Lacking sufficient China Walls was not an offence in itself.

Further, financial regulators had more limited enforcement powers. Imposition of large penalties or punishment of individuals required court judgement, with 5-year statute of limitations. A firm that aided a violator could only be prosecuted if the firm knowingly assisted in the violation, a high legal bar. Most importantly, regulatory action required the evidence of actual fraud or breach of duty. Engaging in transactions with a high risk of fraud or duty breach, or failing to maintain China Walls that could greatly suppress the misuse of MNPI were not themselves actionable by regulators.


Current Regulatory Regime. The US Dodd-Frank Act 2010, and partly coordinated laws elsewhere, dramatically reshape the enforcement of China Walls today. The key change is the "risk-based" enforcement powers granted to financial regulators. Rather than requiring actual illegality before the regulators can act, Dodd-Frank gave them the ability to prosecute behavior that raises the risk of fraud or duty breaches. Moreover, a regulator can now prescribe corporate organization and internal rules that the regulator believes necessary to cap the risk of illegality to a reasonable level.

The China Walls is today a heavily enforced risk-based regulatory prescription. The landmark case is the SEC's 2018 settlement with Mizuho Securities in which Mizuho paid \$1.25 million partly for failing to maintain information barriers between its broker-dealer and hedge fund trading desks ([US Securities and Exchange Commission, 2018](#)). This case began a series of prosecutions by the SEC where the key issue was the effectiveness of the China Walls itself ([Barrack, Moskowitz-Hesse, Richards, and Cox, 2020](#)). As an on-


going example, in 2021, the SEC began a proactive sweep of monitoring and retention of business-related communication among employees across all broker-dealers and investment advisors. The first consequent settlement included a \$125 million fine on Morgan Stanley for their failure to retain all business-related messages sent by its broker-dealer employees *on their private devices* (US Securities and Exchange Commission, 2021). As of early 2024, over \$2 billion in fines have been meted out to dozens of broker-dealers and investment advisors over similar failures. Similarly, the SEC charged Virtu Financial in 2024 merely for having a database accessible to both broker-dealer and nonbroker-dealer employees—despite producing no evidence that any MNPI was leaked (US Securities and Exchange Commission, 2024). Therefore, following Dodd-Frank, the regulatory regime over China Walls morphed from reactive and indirect to proactive and direct.

B Placebo Results

Figure 11 replicates Figures 6 and 7 using median (50 to 50.1st percentile) and exceptionally small (99.9 to 100th percentile) trades to define the events. There are no pretrends nor posttrends across all specifications. We conclude that exceptionally large trades pinpoint the arrival of material information while the smaller trades do not.



(a) Event Trade in 50 to 50.1st Percentile



(b) Event Trade in 99.9 to 100th Percentile

Figure 11: Placebo Estimates

References

- Abadie, A., S. Athey, G. W. Imbens, and J. M. Wooldridge (2023, February). When Should You Adjust Standard Errors for Clustering? *The Quarterly Journal of Economics* 138(1), 1–35.
- Acemoglu, D., G. Egorov, and K. Sonin (2013). A Political Theory of Populism. *The Quarterly Journal of Economics* 128(2), 771–805.
- Aghion, P., Y. Algan, P. Cahuc, and A. Shleifer (2010). Regulation and Distrust. *The Quarterly Journal of Economics* 125(3), 1015–1049.
- Alesina, A. and G. Tabellini (1990). A Positive Theory of Fiscal Deficits and Government Debt. *The Review of Economic Studies* 57(3), 403–414.
- Bailey, M. J., T. Helgerman, and B. A. Stuart (2024, August). How the 1963 Equal Pay Act and 1964 Civil Rights Act Shaped the Gender Gap in Pay. *The Quarterly Journal of Economics* 139(3), 1827–1878.
- Bank for International Settlements (2022, October). OTC foreign exchange turnover in April 2022. Technical report, Bank for International Settlements, Basel, Switzerland.
- Barbon, A., M. Di Maggio, F. Franzoni, and A. Landier (2019). Brokers and Order Flow Leakage: Evidence from Fire Sales. *The Journal of Finance* 74(6), 2707–2749.
- Barrack, C., M. Moskowitz-Hesse, L. Richards, and P. Cox (2020, March). Protecting Firm and Client Information: MNPI and Client Confidentiality. In *Securities Industry and Financial Markets Association*.
- Behrer, A. P., E. L. Glaeser, G. A. M. Ponzetto, and A. Shleifer (2021, April). Securing Property Rights. *Journal of Political Economy* 129(4), 1157–1192.
- Bernhardt, D. and E. Hughson (1997). Splitting orders. *Review of Financial Studies* 10(1), 69–101.
- Bhattacharya, U. and H. Daouk (2002). The World Price of Insider Trading. *The Journal of Finance* 57(1), 75–108.
- Black and Shearson, Hammill Co. (1968, October). Black v. Shearson, Hammill Co.
- Boyarchenko, N., D. O. Lucca, and L. Veldkamp (2021, February). Taking Orders and Taking Notes: Dealer Information Sharing in Treasury Auctions. *Journal of Political Economy* 129(2), 607–645.

- Brooke, J., A. Burrows, D. Faber, C. Harpum, and S. Silber (1995, December). Fiduciary Duties and Regulatory Rules. Technical Report LAW COM No. 236, The Law Commission, London.
- Cengiz, D., A. Dube, A. Lindner, and B. Zipperer (2019, August). The Effect of Minimum Wages on Low-Wage Jobs. *The Quarterly Journal of Economics* 134(3), 1405–1454.
- Chague, F., B. Giovannetti, and B. Herskovic (2023, February). Information Leakage from Short Sellers.
- Chen, T. and X. Martin (2011). Do Bank-Affiliated Analysts Benefit from Lending Relationships? *Journal of Accounting Research* 49(3), 633–675.
- Cook, L. D., M. E. C. Jones, T. D. Logan, and D. Rosé (2023, February). The Evolution of Access to Public Accommodations in the United States. *The Quarterly Journal of Economics* 138(1), 37–102.
- Di Maggio, M., F. Franzoni, A. Kermani, and C. Somnavilla (2019, November). The relevance of broker networks for information diffusion in the stock market. *Journal of Financial Economics* 134(2), 419–446.
- Duffie, D., S. Malamud, and G. Manso (2015). Information percolation in segmented markets. *Journal of Economic Theory* 158(PB), 838–869.
- Easley, D. and M. O’Hara (1987, September). Price, trade size, and information in securities markets. *Journal of Financial Economics* 19(1), 69–90.
- Gardner, J. (2022, July). Two-stage differences in differences.
- Glaeser, E. L. and A. Shleifer (2003, June). The Rise of the Regulatory State. *Journal of Economic Literature* 41(2), 401–425.
- Gozzi, R. (2003). The Chinese Wall Metaphor. *ETC: A Review of General Semantics* 60(2), 171–174.
- Greenstone, M., P. Oyer, and A. Vissing-Jorgensen (2006). Mandated Disclosure, Stock Returns, and the 1964 Securities Acts Amendments. *The Quarterly Journal of Economics* 121(2), 399–460.
- Griffin, J. M., T. Shu, and S. Topaloglu (2012, July). Examining the Dark Side of Financial Markets: Do Institutions Trade on Information from Investment Bank Connections? *The Review of Financial Studies* 25(7), 2155–2188.

- Hagströmer, B. and A. J. Menkveld (2019). Information Revelation in Decentralized Markets. *The Journal of Finance* 74(6), 2751–2787.
- Hanrahan, P. F. (2007, December). *ASIC v Citigroup: Investment Banks, Conflicts of Interest, and Chinese Walls*, pp. 117–142. London: IMPERIAL COLLEGE PRESS.
- Haselmann, R. F. H., C. Leuz, and S. Schreiber (2023, March). Know Your Customer: Informed Trading by Banks.
- Hortaçsu, A. and J. Kastl (2012). Valuing Dealers’ Informational Advantage: A Study of Canadian Treasury Auctions. *Econometrica* 80(6), 2511–2542.
- Irvine, P., M. Lipson, and A. Puckett (2007, May). Tipping. *The Review of Financial Studies* 20(3), 741–768.
- Ivashina, V. and Z. Sun (2011, May). Institutional stock trading on loan market information. *Journal of Financial Economics* 100(2), 284–303.
- Keiser, D. A. and J. S. Shapiro (2019, February). Consequences of the Clean Water Act and the Demand for Water Quality. *The Quarterly Journal of Economics* 134(1), 349–396.
- Kondor, P. and G. Pintér (2022). Clients’ Connections: Measuring the Role of Private Information in Decentralized Markets. *The Journal of Finance* 77(1), 505–544.
- Kumar, N., K. Mullally, S. Ray, and Y. Tang (2020, August). Prime (information) brokerage. *Journal of Financial Economics* 137(2), 371–391.
- Kyle, A. S. (1985). Continuous Auctions and Insider Trading. *Econometrica* 53(6), 1315–1335.
- La Porta, R., F. Lopez-De-Silanes, and A. Shleifer (2006). What Works in Securities Laws? *The Journal of Finance* 61(1), 1–32.
- Lehar, A. and O. Randl (2006, January). Chinese Walls in German Banks. *Review of Finance* 10(2), 301–320.
- Li, F. W., A. Mukherjee, and R. Sen (2021, September). Inside brokers. *Journal of Financial Economics* 141(3), 1096–1118.
- Li, T. (2018, June). Outsourcing Corporate Governance: Conflicts of Interest Within the Proxy Advisory Industry. *Management Science* 64(6), 2951–2971.
- Maskin, E. and J. Tirole (2004, September). The Politician and the Judge: Accountability in Government. *American Economic Review* 94(4), 1034–1054.

- Massa, M. and Z. Rehman (2008, August). Information flows within financial conglomerates: Evidence from the banks–mutual funds relation. *Journal of Financial Economics* 89(2), 288–306.
- Menkhoff, L., L. Sarno, M. Schmeling, and A. Schrimpf (2016, April). Information Flows in Foreign Exchange Markets: Dissecting Customer Currency Trades. *The Journal of Finance* 71(2), 601–634.
- Peluso, D. (2020, December). Turning a blind eye: The complicit trespassing of ‘Chinese walls’ in financial institutions in New York. *Critique of Anthropology* 40(4), 438–454.
- Persson, T. and L. E. O. Svensson (1989). Why a Stubborn Conservative would Run a Deficit: Policy with Time- Inconsistent Preferences. *The Quarterly Journal of Economics* 104(2), 325–345.
- Pinter, G., C. Wang, and J. Zou (2024, July). Size Discount and Size Penalty: Trading Costs in Bond Markets. *The Review of Financial Studies* 37(7), 2156–2190.
- Roth, J., P. H. C. Sant’Anna, A. Bilinski, and J. Poe (2023, August). What’s trending in difference-in-differences? A synthesis of the recent econometrics literature. *Journal of Econometrics* 235(2), 2218–2244.
- Seyhun, H. N. (2007/2008). Insider Trading and the Effectiveness of Chinese Walls in Securities Firms Insider Trading Symposium - January 27, 2007. *Journal of Law, Economics & Policy* 4(2), 369–408.
- Somogyi, F. (2022, August). Dollar Dominance in FX Trading.
- Tuch, A. F. (2014). Financial Conglomerates and Information Barriers. *Journal of Corporation Law* 39(3), 563–616.
- US Securities and Exchange Commission (2018, July). In the Matter of Mizuho Securities USA LLC.
- US Securities and Exchange Commission (2021, December). In the Matter of J.P. Morgan Securities LLC.
- US Securities and Exchange Commission (2024, January). Securities and Exchange Commission v. Virtu Financial Inc. and Virtu Americas LLC.
- Webel, B., D. H. Carpenter, R. Gnanarajah, K. Jones, M. Labonte, R. S. Miller, D. W. Perkins, G. Shorter, and N. E. Weiss (2017, April). The Dodd-Frank Wall Street Reform and Consumer Protection Act: Background and Summary. CRS Report R41350, Congressional Research Service, Washington, D.C.