China Walls

Tomy Lee Daniel Nathan Chaojun Wang*

March 17, 2025

Abstract

Conflicts of interest are inherent to banking conglomerates. Regulators manage these conflicts by enforcing *China Walls*—internal information barriers around 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. Our design compares trading activities of affiliates and entirely unrelated firms around exceptionally large trades to detect information sharing. We document islands of informational autarky between dealers and their affiliate funds surrounded by a sea of information sharing: (i) Dealers never trade nor share information with affiliated funds. (ii) Dealers consistently share information with their client funds, including when they do not trade with each other. (iii) Affiliated funds, which are free to share information with each other, intensely share information among themselves. From a back-of-the-envelope calculation, establishing China Walls between affiliated funds would eliminate $16.1 billion in trades, comprising 37% of their trades on the event dates. Our results hold during crisis and noncrisis periods, and across granular cells of firm and asset characteristics. We reveal remarkable regulatory capacity to control information flows.

JEL classification: G21, G28, G14, G15

Keywords: Banking conglomerates, financial networks, information barriers, information sharing, regulatory capacity

*First version: November 20, 2024. Lee is at the Central European University. Nathan is at the Bank of Israel. Wang is at the Wharton School, University of Pennsylvania. We thank Markus Bak-Hansen, David Card, Andras Danis, Xavier D'Haultfoeuille, Felix Fattinger, Thomas Gehrig, Sasha Indarte, Simon Jurkatis, Attila Lindner, Florian Nagler, Gabor Pinter, Adam Szeidl, Toni Whited, Adam Zawadowski, and Josef Zechner for valuable comments. 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 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, affiliates have tightly aligned incentives, precluding counterparty litigation that is often central to regulatory enforcement. Therefore, effectively enforced China Walls would reveal remarkable regulatory capacity to control information flows—an especially relevant finding today, when concerns over privacy are widespread. Are China Walls effectively enforced within banking conglomerates?

We document that they are. Our empirical challenges mirror that of the regulators: Information sharing is infeasible to directly observe, and effective enforcement in one circumstance does not rule out violations at other times. We compare trading activities of dealers and funds around exceptionally large trades to overcome these challenges. Our difference-in-differences design detects information sharing between a dealer and its affiliate funds if the funds increase trading on days that the dealer makes an exceptionally

[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.

large trade, (i) compared to funds that are entirely unconnected to the dealer and (ii) controlling for granular fixed effects. Three plausible assumptions underpin our design. First, exceptionally large trades pinpoint arrivals of especially valuable private information, when there is the greatest incentive to violate China Walls. Second, funds increase trading upon receiving valuable information. Third, dealers never share private information with funds that are neither affiliates nor its clients.

We implement this design on the near universe of foreign exchange trades involving the Israeli Shekel that covers 21 million transactions from 2019 to 2024. Among them, 87% are transactions in the US dollar and Shekel currency pair. Moreover, the largest dealers in the Shekel market is identical to those in the broader US dollar market and Israeli financial regulations are mainly based on US regulations. An exception is that Israel does not impose China Walls, leaving US regulators (whose jurisdiction reaches worldwide) as the main enforcers of China Walls in our setting. Their rules wall off dealers from their affiliate funds, while leaving funds affiliated to each other free to share information among themselves.

Figure 1 illustrates our main comparison. GS Dealer and GS Fund are affiliates. (GS, MS, and BoA are illustrative names.) Unrelated Fund is unaffiliated and never trades with the other firms in the figure. An event is an exceptionally large trade by the GS Dealer (event firm) that belongs in the top 0.1 percentile of the GS Dealer's trades (event trade). We compare the daily gross dollar volumes of the GS Fund (affiliate firm) and the Unrelated Fund (control) around the event day. We conclude that the event dealers share information with their affiliate funds if the daily volumes of the affiliate funds increase relative to the unrelated funds around the event day.

This approach detects no information sharing from dealers to their affiliate funds nor,

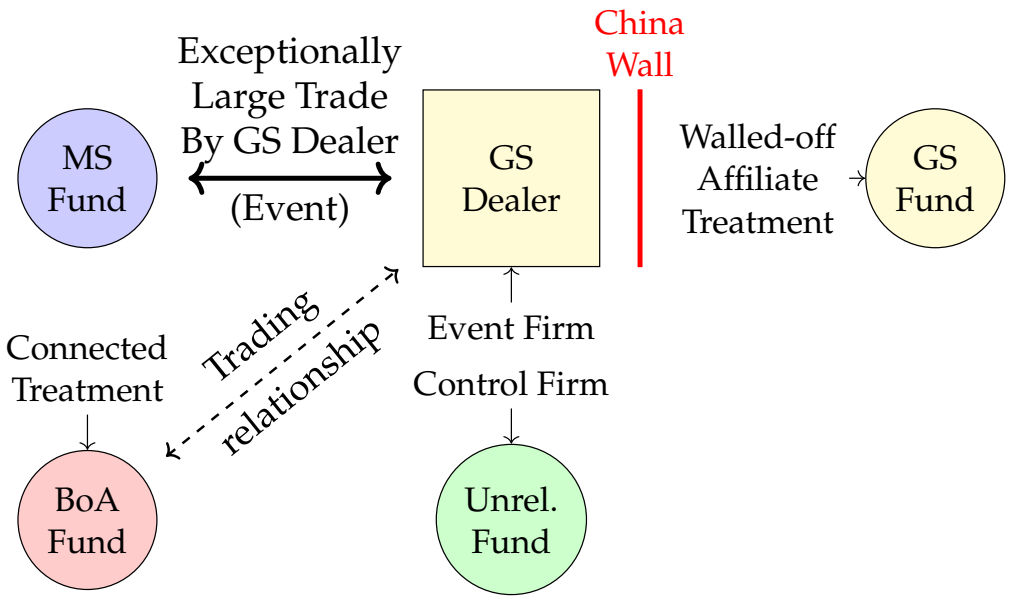


Figure 1: Identifying Information Sharing from Dealers to Affiliate Funds

reversing their roles, from funds to their affiliate dealers. Richness of our setting provides two falsification tests. First, we verify whether our design reliably detects information sharing where it exists. Since dealers are well known to share information with their client funds (Barbon, Di Maggio, Franzoni, and Landier, 2019; Boyarchenko, Lucca, and Veldkamp, 2021), a reliable design must detect information sharing between such connected dealers and funds. In Figure 1, the BoA Fund is a client of the GS Dealer. Our verifying test compares the daily volumes of the BoA Fund (connected fund) and the Unrelated Fund around the day of the GS Dealer's exceptionally large trade. We consistently detect information sharing between connected dealers and funds. Second, we exploit funds that are affiliates but not walled off from each other to infer whether affiliate dealers and funds would share information with each other in the absence of China Walls. Affiliate funds intensely share information among themselves, and thus we infer that affiliate dealers and funds would share information if China Walls were absent.

Section 2 develops the design. Our key identifying assumption is that especially valuable information prompts exceptionally large trades. This assumption is consistent with standard theory (Kyle, 1985; Easley and O'Hara, 1987) and empirical evidence in other settings (Kumar, Mullally, Ray, and Tang, 2020; Pinter, Wang, and Zou, 2024). A theoretical threat is the possibility that firms would split orders to disguise their private information. However, finely splitting orders is likely impractical and impose prohibitive reputational costs, as foreign exchange markets are non-anonymous and intermediated by a small number of dealers (46 in our sample). Appendix B jointly tests this assumption and the claim that our design detect information sharing if and only if it exists. Consistent with these claims, exceptionally large trades predict returns for up to three days following the trade, smaller trades do not, and we only detect information sharing between connected dealers and funds around the large trades.

We use dealers or funds that are unrelated to the event dealer or fund to strip away two sources of contamination when measuring information sharing. First, public news or aggregate shocks would trigger all dealers and funds increase trading and in larger trade sizes. Second, the price and liquidity impacts of event trades may induce other dealers and funds to trade. Because the unrelated funds are as exposed to aggregate shocks and the impacts of event trades as other funds, yet the event dealer would not share information with funds unrelated to it, using the unrelated funds as controls removes the two contaminants while preserving any increases in trading due to information sharing.

Section 3 describes the data and motivating analyses. We compute correlations in daily gross volumes within dealer-fund pairs. Volumes of the unrelated dealer-fund pairs are highly correlated over wide leads and lags, emphasizing the importance of using the unrelated firms to filter out aggregate shocks and trade impacts. Foreshadowing our

results, the affiliate dealer-fund pairs show correlations that are indistinguishable from the unrelated pairs, whereas the connected dealer-fund pairs have dramatically higher contemporaneous correlation.

Section 4 implements our design in stacked difference-in-differences specifications with never-treated controls of Cengiz, Dube, Lindner, and Zipperer (2019).[2] The daily gross dollar volumes of the affiliate funds are statistically indistinguishable from the unrelated funds in each of the five trading days before, the day of, and the five days after an exceptionally large trade by a dealer. On the event day, the affiliate funds change their daily volumes by a statistically insignificant -0.018 standard deviation (std. error: 0.044 sd) relative to the unrelated funds. In stark contrast, the daily volumes of the connected funds increase by 0.97 sd (std. error: 0.0042 sd) on the event day.[3] On days when a fund makes an exceptionally large trade, the daily volumes of affiliate dealers increase by an insignificant 0.016 sd (std. error: 0.049 sd) relative to unrelated dealers, whereas connected dealers raise theirs by a significant 0.15 sd (std. error: 0.018 sd). Taken together, we reject information sharing between affiliate dealers and funds, and conclude that we would have detected information sharing if they were doing so.

We address three key threats to identification. First, our results are not an artifact of weak power. Every affiliate estimate yields a tight confidence interval around zero, while each connected estimates has an interval far above zero. Second, trades between affiliate dealers and funds might generate mechanical increases in trading around event trades. Such mechanical effects cannot confound our affiliate estimates, because dealers

[2]Our estimates are average treatment-on-the-treated (ATT) effects, unlike the uninterpretable estimates generated by traditional staggered two-way fixed-effects specifications (Gardner, 2022; Roth, Sant'Anna, Bilinski, and Poe, 2023).

[3]The extreme precision arises from large sample sizes. This specification has 42 million observations across 7,710 events. We cluster standard errors by event-and-firm and by calendar date.

virtually *never* trade with their affiliate funds, itself consistent with the effectiveness of China Walls. For the connected treatment, our specification excludes any fund that trades with the event dealer on or after the event day, and the same of dealers with respect to the event funds, precluding the mechanical effects. Third, affiliate dealers and funds may be more likely to share trading connections or information sources, which would introduce simultaneity bias. We establish that an affiliate dealer and fund are in fact less likely to share a connection than unrelated dealers and funds. Moreover, we test for parallel trends to determine if these shared connections are biasing our estimates. Because shared connections would reach some affiliate firms before prompting the event trade, these connections would imply an upward trend leading up to the event dates, violating parallel trends. All of our estimates exhibit parallel trends, ruling out shared connections as a material bias in our estimates.

Section 5 restricts attention to funds affiliated to each other. A fund increases its daily volume by 1.7 sd (std. error: 0.22 sd) relative to unrelated funds on a day when an affiliated fund makes an exceptionally large trade. We exclude funds from the affiliate treatment if it has a dealer connection that overlaps with the event fund. Including the funds with a dealer overlap, the estimate soars to 2.6 sd (std. error: 0.30 sd), demonstrating that shared connections could introduce severe bias. In sum, affiliate dealers and funds do not share information and would do so if the China Walls were absent. We conclude that China Walls are effectively enforced.

Section 6 scours granular cells of asset and firm characteristics for China Walls violations. We do not detect information sharing between affiliate dealers and funds across all characteristic cells while consistently detecting sharing between connected dealers and funds. Among funds affiliated to each other, information sharing is universal and the es-

timates are reliably and sharply larger than for the connected treatment. In particular, on days when a hedge fund makes an exceptionally large trade, its affiliate hedge funds raise their volumes by 5.1 sd (std. error: 1.3 sd), whereas other affiliate funds do not respond, -0.032 sd (std. error: 0.24 sd). This finding echos the evidence that hedge funds are especially sensitive to material information (Di Maggio, Franzoni, Kermani, and Sommavilla, 2019; Kumar et al., 2020), lending further credence to our assumption that exceptionally large trades pinpoint arrivals of valuable information.

Related literature. We belong to the literature on the capacity to regulate firm behavior in states governed by rule-of-law. Despite legal constraints, regulators manage to 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). In these settings, the misaligned interests of parties involved in the regulated activity (e.g., employer vs employee, insider vs outside shareholder) 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 findings reveal a remarkable capacity to regulate firms beyond what is established in prior work.

A policy debate rages over the effectiveness of China Walls (Webel, Carpenter, Gnanarajah, Jones, Labonte, Miller, Perkins, Shorter, and Weiss, 2017). Much of the empiri-

cal 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.[4] 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.[5] We contribute identification that isolates bilateral information sharing and validates the empirical design in conditions where China Walls are absent and information sharing 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. 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, dealers act as conduits through which information diffuses throughout their trading networks (Di Maggio et al., 2019; Hagströmer and Menkveld, 2019; Kumar et al., 2020). We identify a stark void in this informational network driven by regulatory intervention, introducing influence of policy on information diffusion to this literature. Methodologically, we add China Walls as a promising source of variation in information flows that is especially relevant today, when the financial sector is highly concentrated.

[4] Lehar 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 Wall violations in various settings. Their latest in-sample year is 2013, well before the start of worldwide proactive enforcement of the Walls by the US Securities and Exchange Commission in 2018, on the basis of new powers granted under the Dodd-Frank Act (Appendix A).

[5] A 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.

Section 2 develops the empirical design. Section 3 describes the data and performs motivating analyses. Section 4 investigates the effectiveness of China Walls. Section 6 contains the heterogeneity analyses.

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.[6] 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. Today, the failure to 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.

[6]Material 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.

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 rapidly 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 near 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. Indeed, 87% of ILS transactions are for the USD-ILS pair and the ILS and the USD markets have the same largest dealers.[7] More broadly, financial regulations in Israel are largely based on the US. 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. 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 enforcers of the China Walls in our setting are the nonIsraeli regulators, especially the US SEC whose jurisdiction extends to all banking conglomerates active in the US (every conglomerate in our sample).

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

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 may not share MPI even if the Walls were absent, in which case enforcement is moot.

Defining events. We examine events 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 et al., 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. 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, increasing its daily gross volume. Therefore, we choose the increases in the gross volume of firm j to proxy for the flow of material information to j.

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

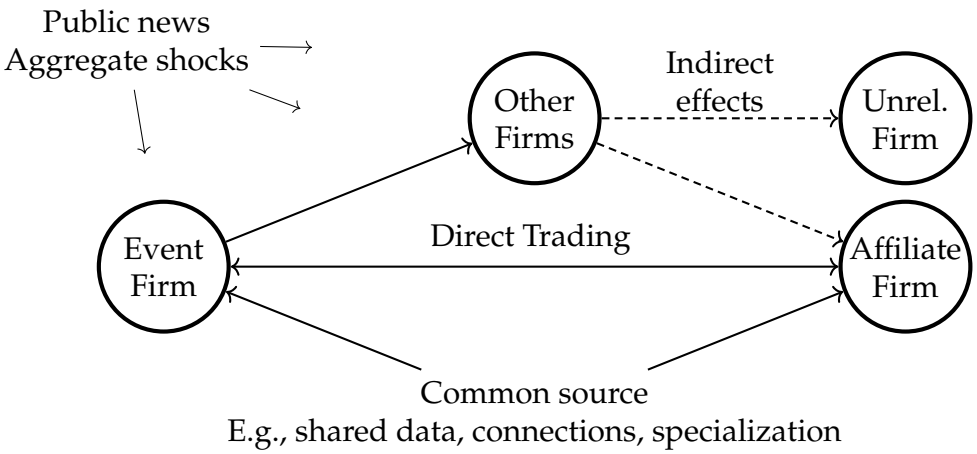


Figure 2: Potential Confounders to Measuring Bilateral Information Sharing

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.[8] 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 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

[8] Any 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.

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.

2.3 Implementation

We adopt the stacked difference-in-differences specification with never-treated controls of Cengiz et al. (2019).[9] 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.[10] 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

[9]Two-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 et al., 2023). Our ATT estimates alway assign positive weights to all individual treatment effects (Gardner, 2022).

[10]For 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.

the event day, because exchange rates fully incorporate private information in about a trading week (Menkhoff, Sarno, Schmeling, and Schrimpf, 2016).

Our first regression specification is

$$Y_{ejt} = \sum_{\tau=-5}^{5} \alpha_\tau \mathbb{1}_{t=\ell_e+\tau} Affiliate_{ej} + \delta_{ej} + \varphi_t + \sum_{\tau=-5}^{5} \gamma_\tau \mathbb{1}_{t=\ell_e+\tau} + \varepsilon_{ejt}. \tag{1}$$

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 $Affiliate_{ej}$ equals 1 if firm j is an affiliate of the event firm. The dummy $Affiliate_{ej} = 0$ if j is unrelated to the event firm and is not treated in any other event within the window of event e.[11] 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,

$$Y_{ejt} = \sum_{\tau=-5}^{5} \beta_\tau \mathbb{1}_{t=\ell_e+\tau} Connected_{ej} + \delta_{ej} + \varphi_t + \sum_{\tau=-5}^{5} \gamma_\tau \mathbb{1}_{t=\ell_e+\tau} + \varepsilon_{ejt}. \tag{2}$$

[11] To 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.

The connected treatment dummy $Connected_{ej}$ equals 1 if (a) firm j trades 10 or more times with the event firm in the sample, and (b) does not trade with the event firm on the event day and five days afterwards, $t = \ell_e, \ldots, \ell_e + 5$. 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.[12] 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.[13] 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 top 0.5 percentile of observations by daily gross volume

[12] All 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/.

[13] Some 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.

separately for the dealer and the fund subsamples, since their distributions dramatically differ.

Table 1: Sample Characteristics

			Final Sample	
	All trades	Fund trades	Dealer day	Fund day
Mean daily volume (USD millions)	29,510	2,843	19,940	2,843
Mean daily volume per firm (USD millions)	642	3.7	433	469
Dollar value per observation (USD millions)	2.7	1.7	635	0.34
Currency				
USD	0.87	0.76	0.94	0.76
JPY	0.07	0.22	0.004	0.22
EUR	0.02	0.02	0.03	0.02
Asset class				
Spot	0.36	0.50	0.32	0.50
Forward	0.13	0.40	0.11	0.40
Swap	0.50	0.10	0.58	0.10
Observations	20,832,686	2,762,406	62,974	10,643,975

All trades: Raw data set containing the near universe of Israeli Shekel transactions. *Fund trades*: Transactions involving a fund. *Dealer day*: Dealer transactions aggregated to the daily gross dollar volume in USD; excludes trades between dealers affiliated to each other and trades with non-financial firms. *Fund day*: Fund transactions aggregated to the daily gross dollar volume in USD. *Mean daily volume* is the average daily total dollar volume in USD billions. *Mean daily volume per firm* is the mean daily volume divided by the number of firms in the sample. All Currency and Asset class figures are weighted by dollar volume.

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 en-

Table 2: Number of Unique Entities

	Conglomerates	Dealers	Funds
US	15	92	4,826
Israeli	11	15	192
Independent	–	11	6,660
Hedge funds	–	–	632
Total	46	229	7,775

A conglomerate is a holding company and the group of firms that the holding company ultimately controls. "Dealers" also include brokers and broker-dealers. All dealers in our sample are broker-dealers, which match client orders or trade on their own accounts at the their discretion. "Independent" denotes entities that do not belong to a conglomerate. All independent dealers are Israeli, perhaps due to Israeli law that forbids common ownership of banks and dealers.

tered 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?"[14] Fourth, we manually check each answer generated in step three.[15]

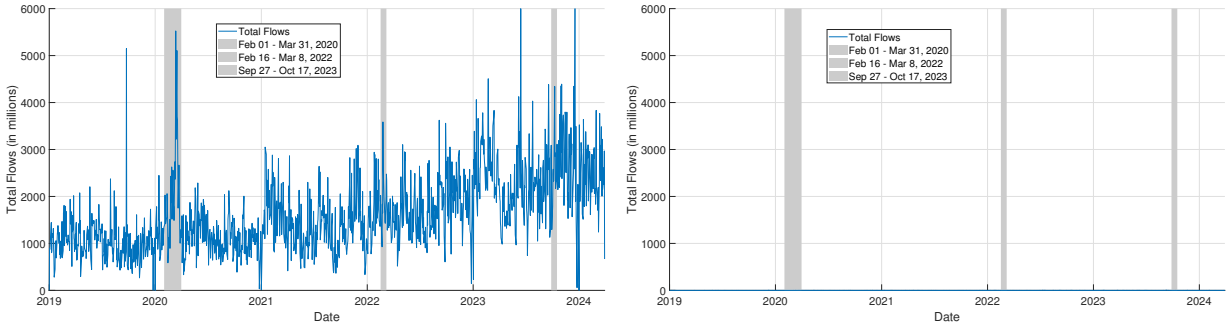
3.2 Motivating Analyses

Three analyses motivate our main empirical design. First, Figure 3 plots the total daily dollar volume of transactions across all paris 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 4 computes the correlation in daily gross volumes within unrelated dealer-fund pairs. For each lag $l = -10 \ldots + 10$ and a pair of dealer i and fund j that

[14] We 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.

[15] Verification is typically immediate upon searching for the firm name paired with either "independent" or the `ChatGPT`-suggested holding company name.



(a) Pairs of Nonaffiliate Dealer and Fund

(b) Pairs of Affiliate Dealer and Fund

Figure 3: Daily Gross Dollar Volumes Between Dealers and their Nonaffiliate Funds

Figure 3a: 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 3b: 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.

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

Figure 4a 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

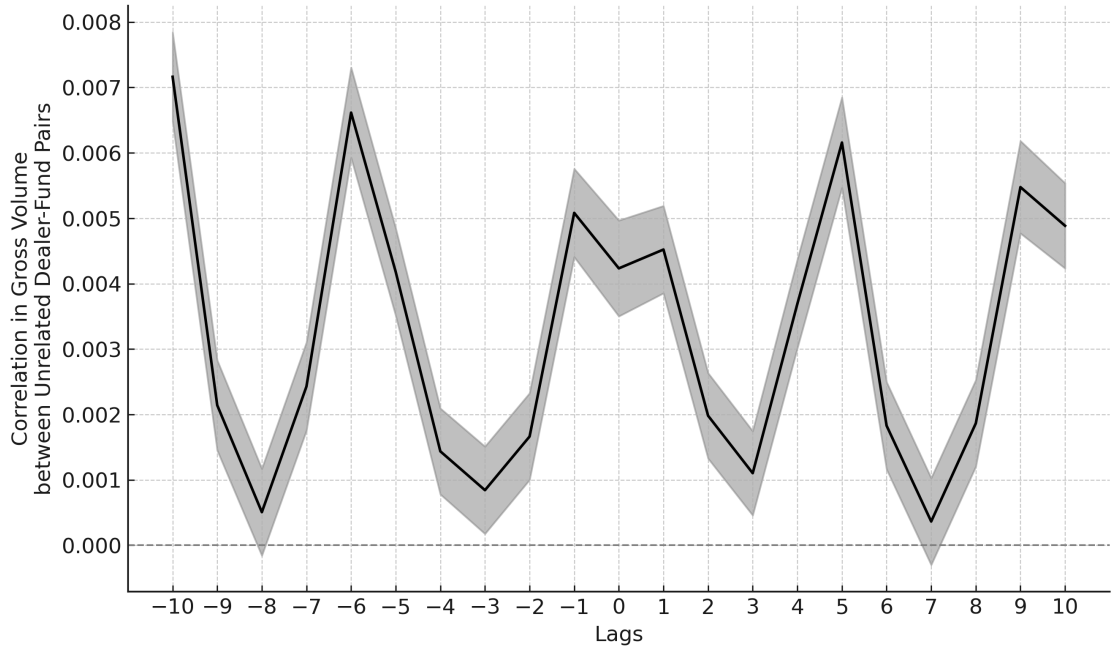
$$CorrGV_{ijl} = a_i + b_j + c_l Affiliate_{ij} + d_l Connected_{ij} + \varepsilon_{ijl}. \tag{4}$$

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.

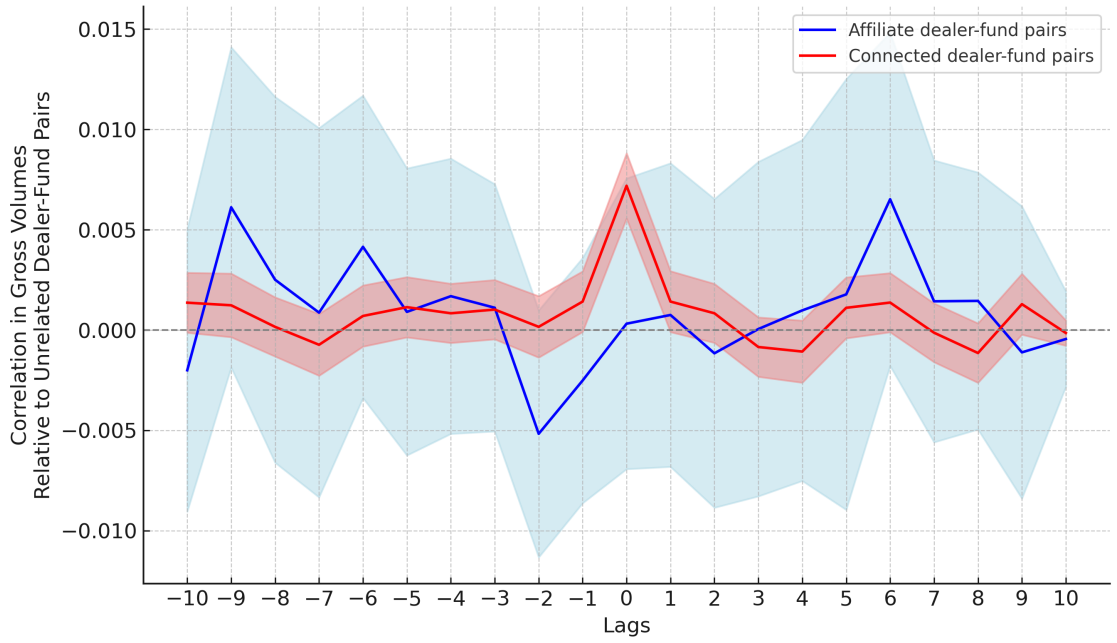
Figure 4b plots the coefficients c_l and d_l across $l = -10 \ldots 10$. Trading activities of the affiliate dealers and funds are no more correlated than the unrelated dealers and funds across l. In stark contrast, the connected dealers and funds are significantly more correlated than the unrelated 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 5a 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



(a) Correlation Within Unrelated Pairs



(b) Affiliate and Connected Pairs

Figure 4: Correlations in Gross Volume

no more than 0.53 standard deviations with 95% confidence, whereas the connected funds increase theirs by highly statistically significant 0.76 sd.

We interpret Figure 5a 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.

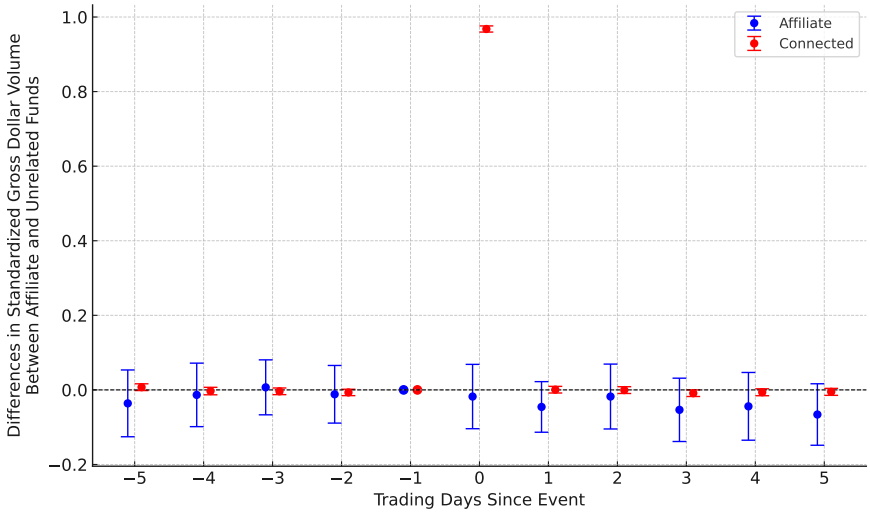
The mechanical channel is ruled out by the exclusion of funds that traded with the event dealer on or after the event day for each event. The aggregate-shock and the indirect-effect channels are stripped away by the unrelated fund control group, since the unrelated funds would be exposed to the aggregate 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.

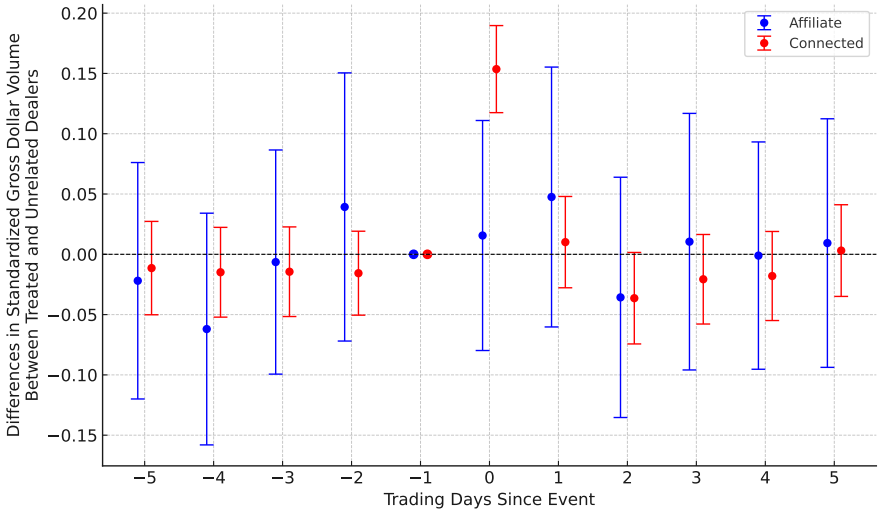
Figure 5b presents the coefficient estimates of Equations (1) and (2) where we examine the standardized daily gross volumes of the dealers around the days when a fund makes an exceptionally large trade. In blue are the differences α_τ in the volumes between the affiliate and the unrelated dealers around 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.

Based on the results of Figures 5 and 7, we conclude that the China Walls are effective on the whole. Table 3 details the pooled regression counterparts to Figures 5 and 7. The affiliate funds have precisely null response to the arrival of especially valuable information at the dealers and the converse for the affiliate dealers to the funds.[16] 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.

[16]The 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.



(a) Fund Responses to Dealer Information



(b) Dealer Responses to Fund Information

Figure 5: Coefficient Estimates from Equations (1) and (2)

Table 3: Responses in Daily Volumes by Affiliate Firms and Connected Firms on and after the Event Day

	D2F Affiliate	F2D Affiliate	D2F Connected	F2D Connected	F2F Affiliate
$Post \times Affiliate$	-0.020	0.012			0.24***
	[0.015]	[0.028]			[0.021]
$Post \times Connected$			0.32***	0.034**	
			[0.0047]	[0.015]	
$Post \times DealerOverlap$					0.017***
					[0.0036]
$Post \times Affiliate$					0.21***
$\times DealerOverlap$					[0.025]
Event×Firm FE	Yes	Yes	Yes	Yes	Yes
Calendar Date FE	Yes	Yes	Yes	Yes	Yes
Days-since-Event FE	Yes	Yes	Yes	Yes	Yes
Adjusted R-squared	0.023	0.41	-0.007	0.48	0.051
Within R-squared	0	0.0001	0.0005	0.0001	0.0002
Events	7,710	9,472	7,710	9,472	9,472
Observations	79,671,983	4,057,372	42,150,672	3,614,383	12,664,366

Coefficient estimates from Equations (1), (2) and (5). The dependent variable is the standardized daily gross US dollar volume of a firm winsorized at the top 0.5 percentile. We include event-by-firm, calendar date, and days-relative-to-event-date (Days-since-Event) fixed effects. 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. *F2F*: All firms are funds. The F2F estimates are for the funds whose dealer connections do not overlap with the event fund. Affiliate treatment includes firms that belong to the same conglomerate as the event firm. Connected treatment includes firms that trade at least 10 times with the event firm during the sample period, and do not trade with the event firm on the event day nor on the five subsequent trading days. The Affiliate and Connected treatments are mutually exclusive, because no dealer trades 10 or more times with an affiliate fund in our sample. Control firms are unaffiliated and never trades with the event firm, and are not treated in another event on any of the 11 trading days around the event day.

5 Would Affiliates Share Information Absent the China Walls?

We exploit that each banking conglomerate owns multiple funds to infer whether the affiliate dealers and funds would share MPI if their China Walls were absent. 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.

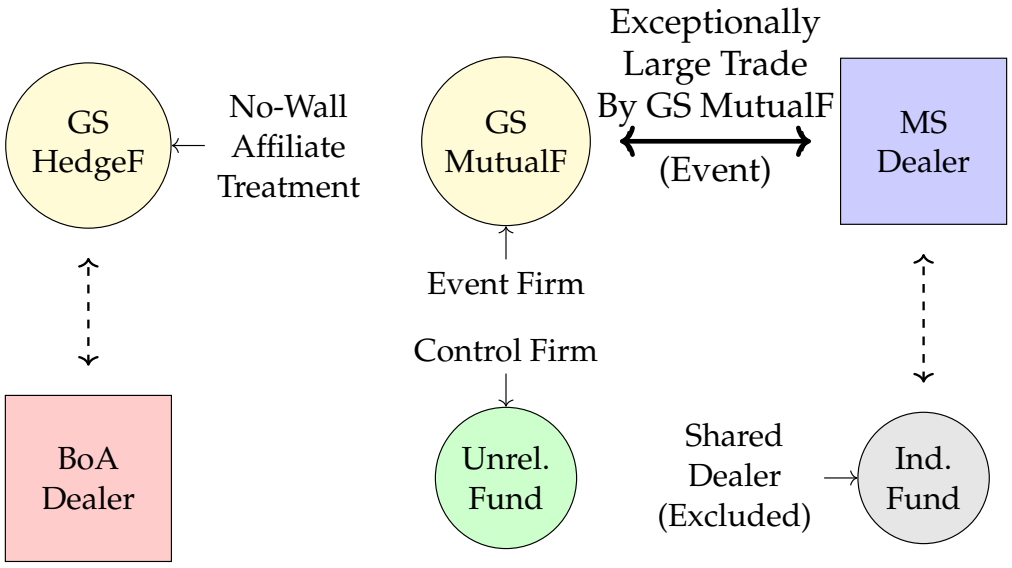


Figure 6: Identification: Information Sharing Between Affiliate Funds

Figure 6 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.

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

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

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.

Results. Figure 5 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 (5) 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 more than twice the response of the connected funds to dealer events in Figure 5a, 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.

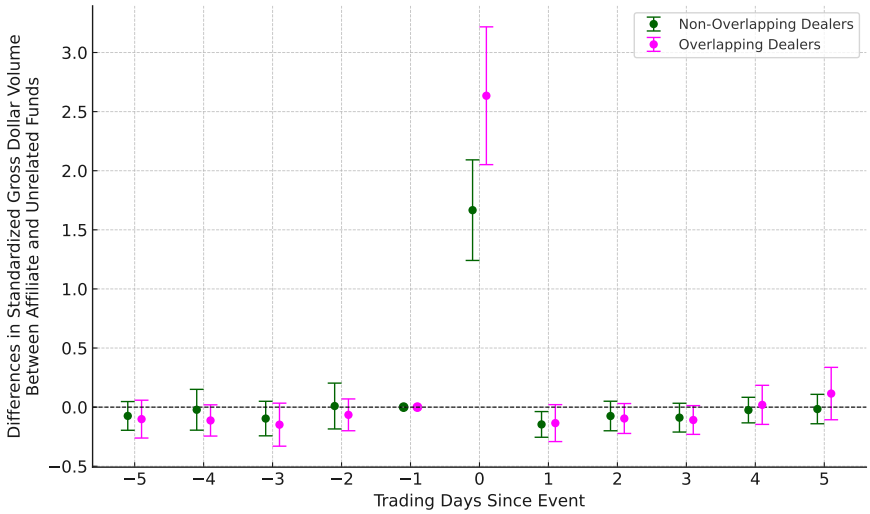


Figure 7: Affiliate Fund Response to Event Fund Information

6 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), (2) and (5), 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

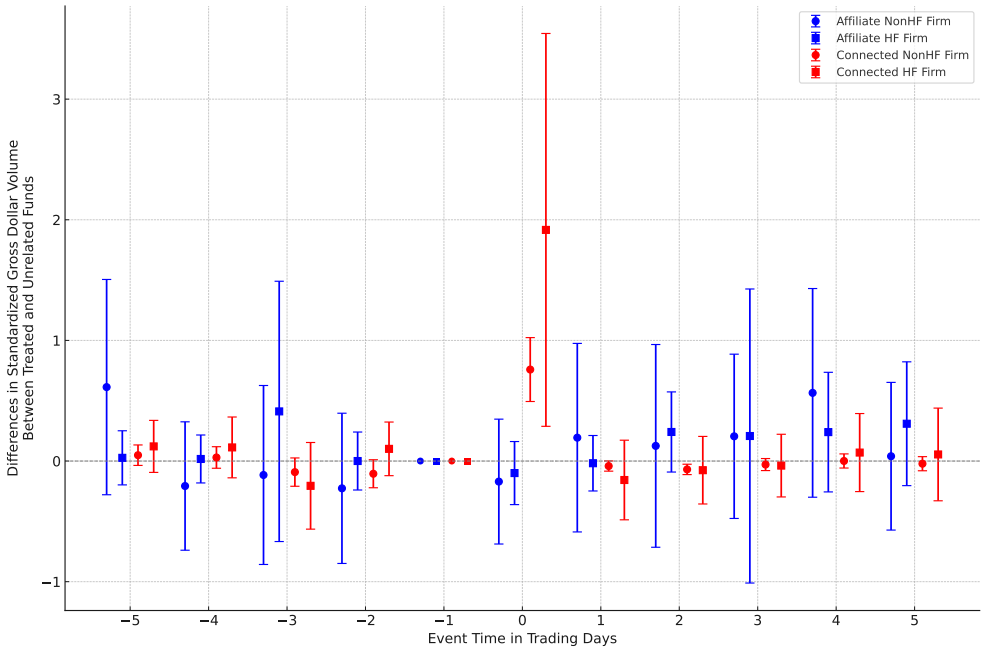
firms. Precisely, the dummy $HedgeFund$ equals 1 if the treated or the control firm is a hedge fund, and $HedgeFundEvent = 1$ if the event firm is a hedge fund. Other dummy variables indicate whether a firm is located in the US, or has a higher share of USD or swap trades than the median. We call the 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 $\mathbb{1}_{t=\ell_e+\tau} HedgeFund_{ej}$.[17] Figure 8b does so symmetrically for the dealer responses to hedge fund and nonhedge fund events, and Figure 9 partitions the reponses to and from hedge funds and nonhedge funds without dealer overlap into four cells.[18]

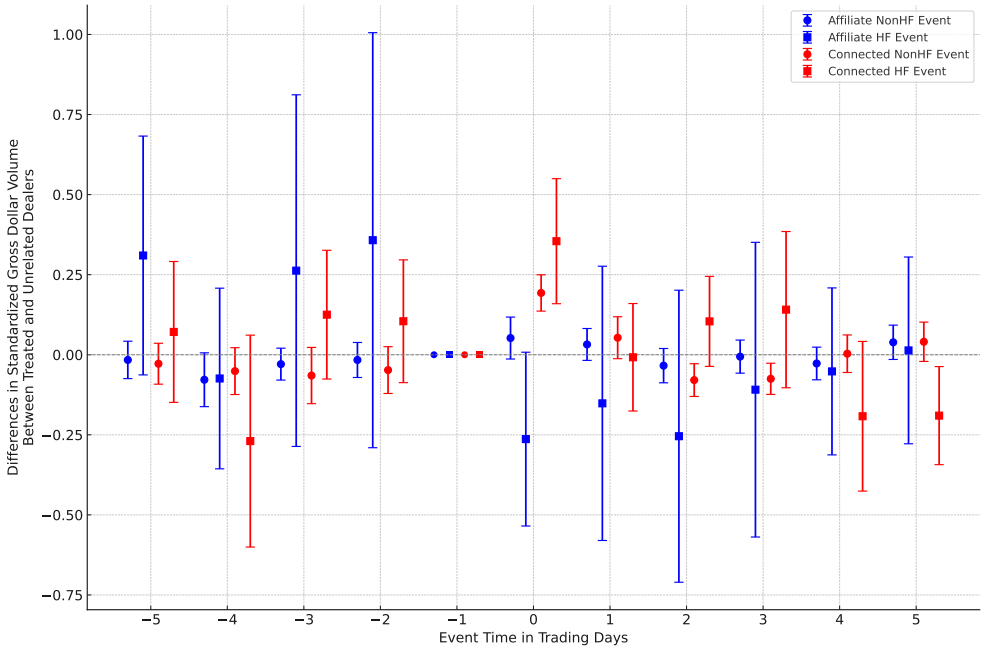
Table 6 compares the coefficient estimates for the events during crisis and noncrisis periods. The crisis periods span Covid (February 1st to March 31, 2020), the Russian Invasion of Ukraine (February 16 to March 8, 2022), and the Hamas Attack (September 27 to October 17, 2023).

[17] We use the constituent coefficients' covariance matrices to compute the standard errors of their sums.

[18] To 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

Table 4: Responses to USD and NonUSD Event Trades in Daily Volumes by Firms Specialized in USD and NonUSD Assets on and after the Event Day

	D2F Affiliate	D2F Connected	F2D Affiliate	F2D Connected	F2F Affiliate
NonUSD Event Trade	-0.047	0.32***	-0.019	0.030	0.15*
-NonUSD Firm	[0.033]	[0.011]	[0.046]	[0.091]	[0.088]
NonUSD Event Trade	-0.047	0.33***	-0.021	0.019	0.24***
-USD Firm	[0.033]	[0.014]	[0.047]	[0.092]	[0.080]
USD Event Trade	-0.050	0.31***	-0.025	0.020	0.22***
-NonUSD Firm	[0.033]	[0.015]	[0.047]	[0.092]	[0.080]
USD Event Trade	-0.060	0.34***	-0.089	0.098***	0.24***
-USD Firm	[0.049]	[0.015]	[0.067]	[0.032]	[0.080]
Event×Firm FE	Yes	Yes	Yes	Yes	Yes
Calendar Date FE	Yes	Yes	Yes	Yes	Yes
Days-since-Event FE	Yes	Yes	Yes	Yes	Yes
Adjusted R-squared	0.016	-0.007	0.50	0.48	0.052
Within R-squared	0.0000	0.0005	0.0001	0.0002	0.0002
Events	7,710	7,710	9,472	9,472	9,472
Observations	79,671,983	42,150,672	42,150,672	3,614,383	12,664,366

Coefficient estimates from Equations (1), (2) and (5). The dependent variable is the standardized daily gross US dollar volume of a firm winsorized at the top 0.5 percentile. We include event-by-firm, calendar date, and days-relative-to-event-date (Days-since-Event) fixed effects. 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 Trade indicates an event whose corresponding event trade is a USD-ILS transaction. USD Firm indicates a treated or a control firm whose USD-ILS share of trades by dollar volume exceeds the median across firms. *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. *F2F*: All firms are funds. The F2F estimates are for the funds whose dealer connections do not overlap with the event fund. Affiliate treatment includes firms that belong to the same conglomerate as the event firm. Connected treatment includes firms that trade at least 10 times with the event firm during the sample period, and do not trade with the event firm on the event day nor on the five subsequent trading days. The Affiliate and Connected treatments are mutually exclusive, because no dealer trades 10 or more times with an affiliate fund in our sample. Control firms are unaffiliated and never trades with the event firm, and are not treated in another event on any of the 11 trading days around the event day.

Table 5: Responses to Swap and NonSwap Event Trades in Daily Volumes by Firms Specialized in Swap and NonSwap Assets on and after the Event Day

	D2F Affiliate	D2F Connected	F2D Affiliate	F2D Connected	F2F Affiliate
NonSwap Event Trade	-0.0013	0.31***	0.027	0.083***	0.23***
-NonSwap Firm	[0.015]	[0.0064]	[0.031]	[0.025]	[0.022]
NonSwap Event Trade	-0.0038	0.31***	-0.037	0.065**	0.23***
-Swap Firm	[0.015]	[0.0088]	[0.033]	[0.027]	[0.023]
Swap Event Trade	-0.00097	0.30***	-0.022	0.088***	0.24***
-NonSwap Firm	[0.015]	[0.0078]	[0.033]	[0.027]	[0.023]
Swap Event Trade	-0.0019	0.32***	0.10	0.035	0.42***
-Swap Firm	[0.025]	[0.011]	[0.13]	[0.044]	[0.16]
Event×Firm FE	Yes	Yes	Yes	Yes	Yes
Calendar Date FE	Yes	Yes	Yes	Yes	Yes
Days-since-Event FE	Yes	Yes	Yes	Yes	Yes
Adjusted R-squared	0.023	-0.012	0.49	0.49	0.052
Within R-squared	0.0000	0.0007	0.0001	0.0002	0.0002
Events	7,710	7,710	9,472	9,472	9,472
Observations	79,671,983	42,150,672	42,150,672	3,614,383	12,664,366

Coefficient estimates from Equations (1), (2) and (5). The dependent variable is the standardized daily gross US dollar volume of a firm winsorized at the top 0.5 percentile. We include event-by-firm, calendar date, and days-relative-to-event-date (Days-since-Event) fixed effects. 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 Trade indicates an event whose corresponding event trade is a swap transaction. Swap Firm indicates a treated or a control firm whose swaps share of trades by dollar value exceeds the median across firms. *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. *F2F*: All firms are funds. The F2F estimates are for the funds whose dealer connections do not overlap with the event fund. Affiliate treatment includes firms that belong to the same conglomerate as the event firm. Connected treatment includes firms that trade at least 10 times with the event firm during the sample period, and do not trade with the event firm on the event day nor on the five subsequent trading days. The Affiliate and Connected treatments are mutually exclusive, because no dealer trades 10 or more times with an affiliate fund in our sample. Control firms are unaffiliated and never trades with the event firm, and are not treated in another event on any of the 11 trading days around the event day.

Table 6: Responses to Events During Crisis and NonCrisis Periods in Daily Volumes by Affiliate Firms and Connected Firms on and after the Event Day

	D2F Affiliate	D2F Connected	F2D Affiliate	F2D Connected	F2F Affiliate
Post × *Treated*	-0.022	0.32***	0.0086	0.035**	0.35***
	[0.019]	[0.0049]	[0.031]	[0.014]	[0.028]
Post × *Treated*	-0.055	0.012	-0.075	-0.0063	-0.0029
× *Crisis*	[0.081]	[0.026]	[0.12]	[0.040]	[0.077]
Event×Firm FE	Yes	Yes	Yes	Yes	Yes
Calendar Date FE	Yes	Yes	Yes	Yes	Yes
Days-since-Event FE	Yes	Yes	Yes	Yes	Yes
Adjusted R-squared	0.016	-0.007	0.20	0.48	0.043
Within R-squared	0.0000	0.0005	0.0001	0.0002	0.0002
Crisis Events	440	440	3,303	3,303	3,303
Events	7,710	7,710	9,472	9,472	9,472
Observations	79,671,983	42,150,672	4,057,372	3,614,383	12,664,366

Coefficient estimates from Equations (1), (2) and (5). The dependent variable is the standardized daily gross US dollar volume of a firm winsorized at the top 0.5 percentile. We include event-by-firm, calendar date, and days-relative-to-event-date (Days-since-Event) fixed effects. 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 occured 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. *F2F*: All firms are funds. The F2F estimates are for the funds whose dealer connections do not overlap with the event fund. Affiliate treatment includes firms that belong to the same conglomerate as the event firm. Connected treatment includes firms that trade at least 10 times with the event firm during the sample period, and do not trade with the event firm on the event day nor on the five subsequent trading days. The Affiliate and Connected treatments are mutually exclusive, because no dealer trades 10 or more times with an affiliate fund in our sample. Control firms are unaffiliated and never trades with the event firm, and are not treated in another event on any of the 11 trading days around the event day.

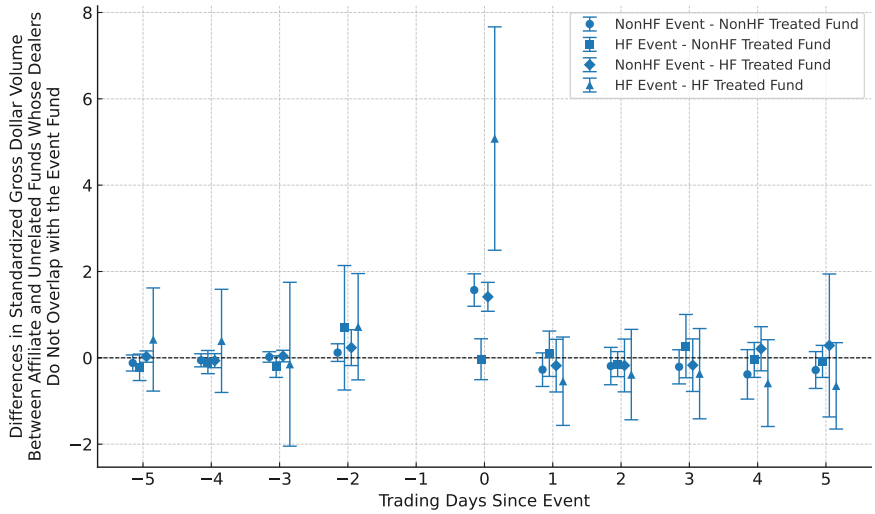


Figure 9: Affiliate Fund-to-Fund Responses: Hedge Funds vs Nonhedge Funds

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 includes 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.

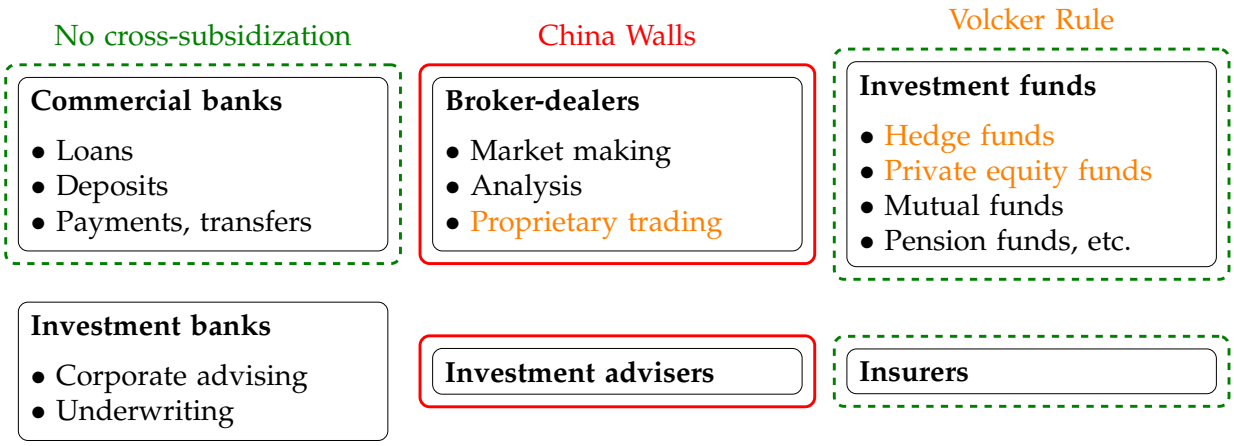


Figure 10: Stylized Banking Conglomerate and Relevant Legal Restrictions

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.

All regulations against the misuse or leakage of financial information target *material non-public information* (MNPI). Information is MNPI if its public disclosure would appreciably 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.[19] 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.
- 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.

[19] Analyses of MNPI are MNPI, whereas analyses of publicly available information are not.

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.[20] 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 (Hamilton, 2011).

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

[20]Outstanding 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.)

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 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.[21]

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 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.[22] 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

[21] A 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.

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

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.[23] Today's legal standard thereby narrows where the China Walls are valuable to the subsidiaries whose (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).[24] 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

[23] A 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.

[24] The 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.

counterparties.[25]

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, occuring 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.[26] The prosecutions over the LIBOR scandal 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

[25]The 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).

[26]The 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.

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 ongoing example, in 2021, the SEC began a proactive sweep of monitoring and retension 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

Two exercises jointly test two identifying assumptions that: (a) Exceptionally large trades pinpoint the arrivals of especially valuable MPI; and (b) bilateral MPI sharing drives our connected (2) and fund-to-fund affiliate (5) treatment results.

First, we compute the price impacts of exceptionally large, median (50 to 50.1st percentile among the event firm's trades by dollar value), and exceptionally small (99.9 to 100th percentile) trades. We do not observe who initiated each trade. Instead, under the intuition that net volumes determine prices (Kyle, 1985), we net all trades in the given percentile in each day separately for funds and dealers. We sign the funds' daily net volumes from the perspective of funds, and the dealers' net volumes from the dealers' perspective. For example, on a certain day, if the GS Fund and the BoA Fund made sales in the top 0.1-percentile among their respective trades, and the two sales sums up to \$1 million, then the net fund volume in the 0.1-percentile on that day would be $-$\$1 million. We net out each interdealer trade. Since doing so removes the largest dealer trades in the 0.1 percentile, we are conservatively estimating the price impact of the dealers' exceptionally large trades.

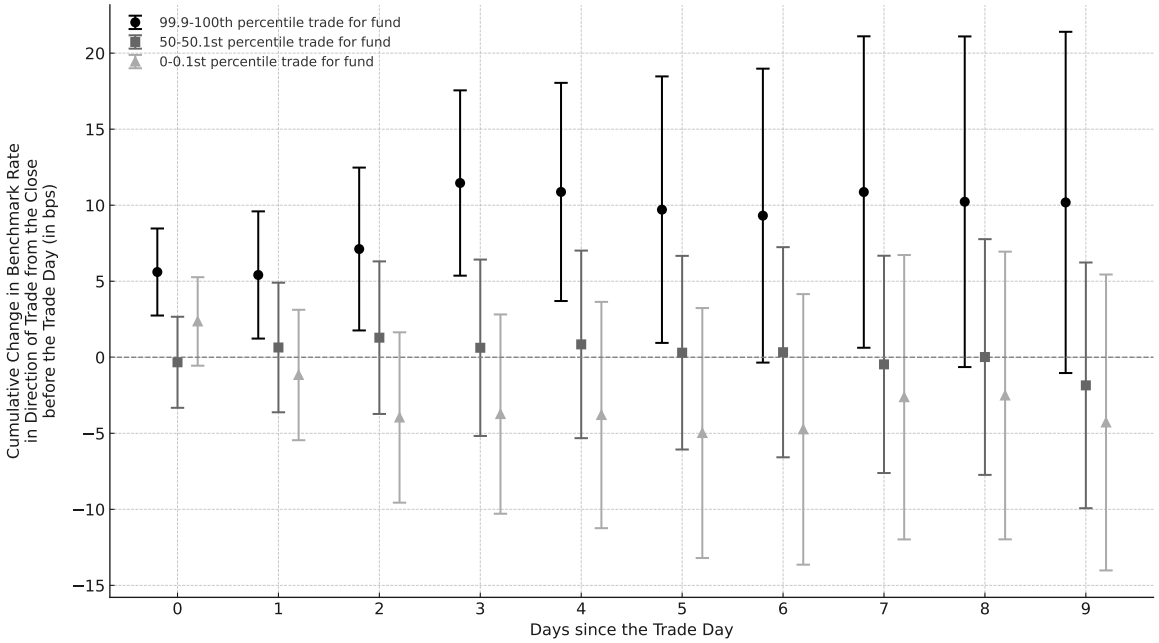
To compute the price impact, we convert the net dollar volumes on day t into trade-direction dummies $d_{t,k,p} \in \{-1, 0, 1\}$, for $k \in \{\text{fund}, \text{dealer}\}$ and percentile $p \in \{[0, 0.1], [50, 50.1], [99.9, 100]\}$. The dummy $d_{t,k,p} = -1$ if the day's net volume is negative and so on. We then import the Bloomberg benchmark exchange rates for each day, and compute the cumulative returns $R_{t,t+\ell}$ between t and $t + \ell$, $\ell \in \{0, \ldots, 9\}$. The price impact for firm type k, percentile p, and cumulative return horizon ℓ is the coefficient $\rho_{k,p,\ell}$

in the time-series regression (6):

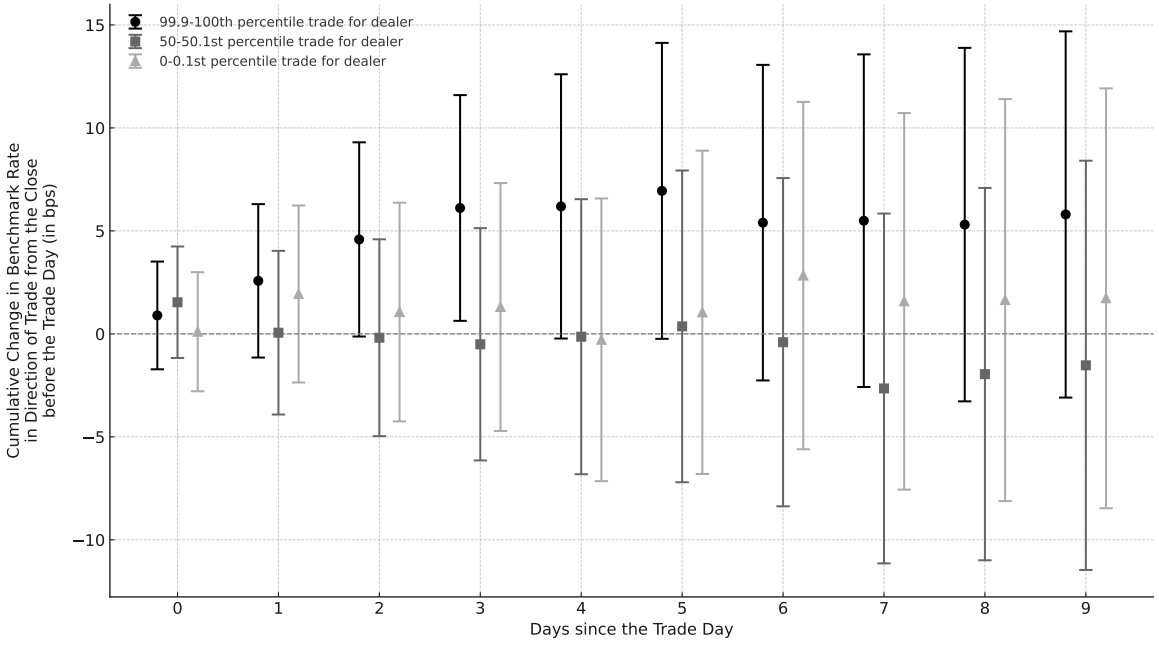
$$R_{t,t+\ell} = \alpha_{k,p,\ell} + \rho_{k,p,\ell} \cdot d_{t,k,p} + \varepsilon_{t,k,p,\ell}. \tag{6}$$

Figure 11 plots the price impact estimates. The net volumes from exceptionally large trades predict future returns, whereas the median and the exceptionally small trades do not.

Second, Figure 12 replicates Figures 5 and 7, except redefining an event to be a day when a firm makes a median or an exceptionally small trade. Across all specifications, there are no pretrend nor posttrend, and the event-day coefficient estimate is insignificant at the 95% confidence level. Combined with Figures 5 and 11, these results show that the daily gross volumes of connected firms and non-walled-off affiliate funds increase only in response to the trades that are predictive of returns. We conclude that the exceptionally large trades pinpoint the arrivals of valuable MPI, and that the bilateral sharing of the valuable MPI drives our results.

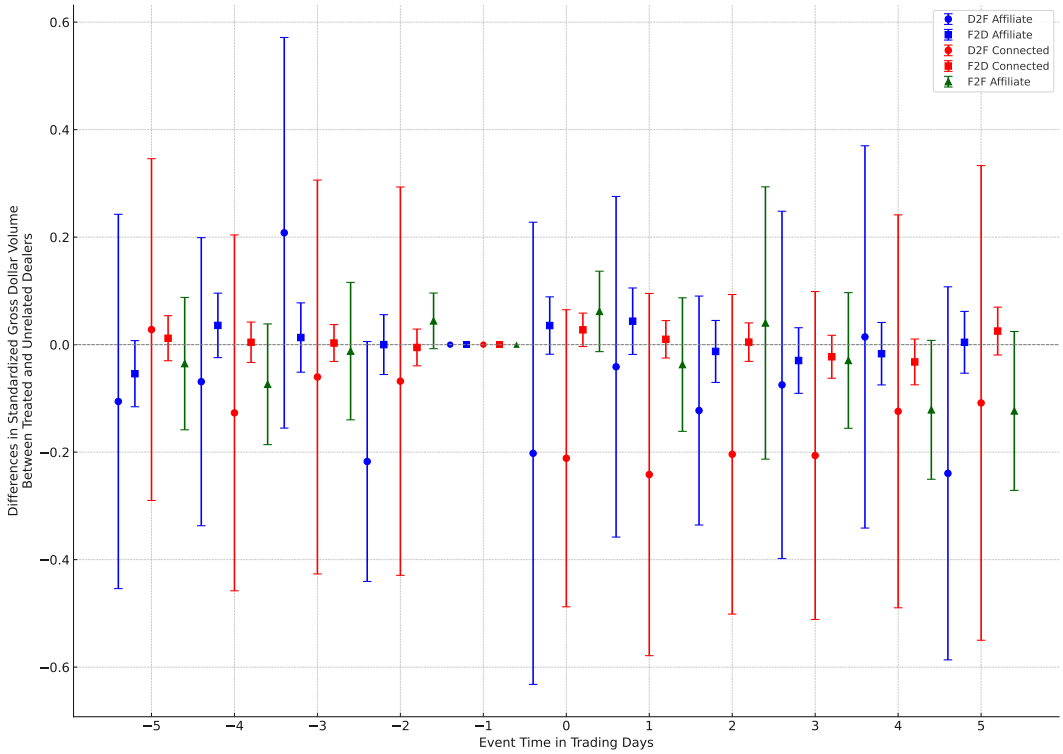


(a) Trades by Funds

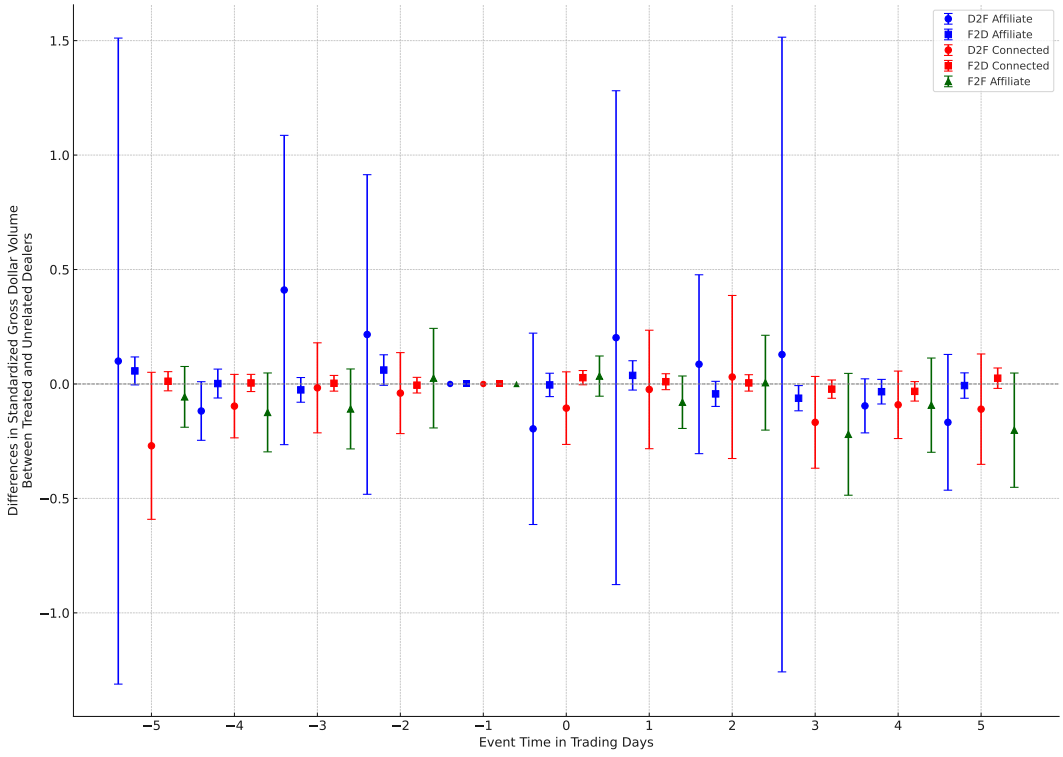


(b) Trades by Dealers

Figure 11: Price Impact Estimates



(a) Event Trade in 50 to 50.1st Percentile



(b) Event Trade in 99.9 to 100th Percentile

Figure 12: 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.

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. Sommavilla (2019, November). The relevance of broker networks for information diffusion in the stock market. *Journal of Financial Economics 134*(2), 419–446.

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.

Hagströmer, B. and A. J. Menkveld (2019). Information Revelation in Decentralized Markets. *The Journal of Finance 74*(6), 2751–2787.

Hamilton, L. (2011, February). US - NAIC Adopts Modified Insurance Holding Company System model Act and regulation. *Global Corporate Insurance & Regulatory Bulletin*.

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.

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.

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.