China Walls

Tomy Lee Daniel Nathan Chaojun Wang*

May 17, 2025

Abstract

Regulators manage conflicts of interest within banking conglomerates by enforcing *China Walls*—internal information barriers around dealers. To evaluate if today's China Walls are effectively enforced, we map information sharing between dealers and funds using the universe of Israeli Shekel foreign exchange trades. Our design compares trading activities of affiliates against 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)* A dealer never trades nor shares information with its affiliated funds. *(ii)* Dealers consistently share information with their client funds, including on days when they do not trade with each other. *(iii)* Affiliated funds, which are free to share information with each other, intensely do so 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. To limit these conflicts, regulators in the US increasingly enforce *China Walls*—blunt information barriers around broker-dealers, which are particularly exposed to conflicts of interest.[1]

Enforcing China Walls is a formidable 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 crucial to regulatory enforcement. As such, effectively enforced China Walls would reveal remarkable regulatory capacity to control information flows—especially relevant today, when concerns over privacy are widespread. Are today's China Walls effectively enforced within banking conglomerates?

We document that they are. Our empirical challenges mirror that of the regulators: information sharing is not directly observable, and compliance 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 large trade relative to funds that are entirely unrelated to the dealer. Three plausible assumptions underpin our design. First, exceptionally large trades pinpoint arrivals of especially valuable

[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.

private information, when there would be the strongest incentive to violate China Walls. Second, funds increase trading activity upon receiving valuable information. Third, dealers never share private information with unrelated funds that are neither their affiliates nor clients.

We implement this design on the near universe of foreign exchange trades involving the Israeli Shekel covering 21 million trades from 2019 to 2024. Among them, 87% are trades in the US dollar-Shekel currency pair. Moreover, the largest dealers in the Shekel market are 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 the 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. Appendix A details the legal context.

Figure 1 illustrates our design. 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 (event trade) by the GS Dealer (event firm) that belongs in the top 0.1 percentile of the GS Dealer's trades. 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, 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

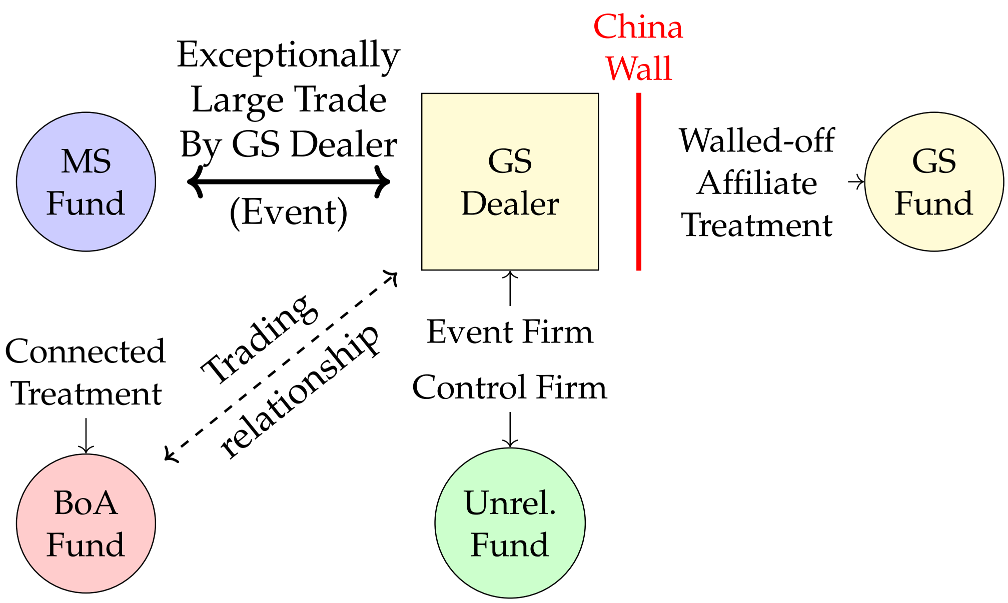

Figure 1: Identifying Information Sharing from Dealers to Affiliate Funds

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 first falsification test compares the daily volumes of the BoA Fund (connected firm) and the Unrelated Fund (control) 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 determine whether affiliates do 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 the 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 empirically holds in other markets (Kumar, Mullally, Ray, and Tang, 2020; Pinter, Wang, and Zou, 2024). A threat is the possibility that firms would split orders to disguise their private information. Appendix B jointly tests this assumption and the claim that our design isolates information sharing. Consistent with these claims, exceptionally large trades predict future returns, smaller trades do not, and we detect information sharing between connected dealers and funds only around the large trades.

We then strip away three sources of confounding variation in trade volumes. First, public news or aggregate shocks can simultaneously trigger the funds to increase trading and the dealers to make exceptionally large trades. Second, the liquidity and price impacts of the event trades, rather than information sharing, can induce funds to increase trading. Because no dealer would share information with a fund that is neither an affiliate nor a client—and yet these unrelated funds are as exposed to the aggregate shocks and the impacts of event trades as other funds—using the unrelated funds as controls removes the two confounders while preserving any variation due to information sharing. And third, we may be omitting relevant characteristics of events and funds. Our calendar date, days-relative-to-event, and event-by-fund fixed effects eliminate any confounding variation that is common across funds over time, or specific to a fund as long as it is invariant over the two weeks (the event window) around the event.

Section 3 describes the data. There are 7,700 funds, 46 conglomerates that control dealers, and 17,000 events in our sample. The dealers virtually *never* trade with their affiliate funds, perhaps due to the onerous constraints of the dealers' China Walls. Our main analyses test whether the China Walls preempt information sharing in addition to barring trades between affiliate dealers and funds.

Section 4 implements our design in stacked difference-in-differences specifications with never-treated controls of Cengiz, Dube, Lindner, and Zipperer (2019). Our analytical samples contain millions of observations across thousands of events and firms, providing the power to detect even tiny differences between treated and control groups. Despite the high power, in the 11 trading days around an exceptionally large trade by an event dealer, the daily gross dollar volumes of funds affiliated to this dealer are statistically indistinguishable from those of unrelated funds, differing by -0.02 standard deviation on the event day (clustered std. error: 0.04 sd). In stark contrast, the funds connected to the event dealer increase their volumes by a precisely estimated 1.9 sd (std. error: 0.007 sd) on the event day relative to the unrelated funds. Likewise, around a day when an event fund makes an exceptionally large trade, the gross volumes of its affiliate and unrelated dealers are indistinguishable from each other, whereas its connected dealers sharply increase their volumes on the event day relative to the unrelated dealers. All results remain when we replace gross volumes with net volumes signed in the direction of the event trade.[2]

Section 5 applies this design to funds affiliated to each other. On a day when an event fund makes an exceptionally large trade, the funds affiliated to the event fund increase their volumes by 1.7 sd (std. error: 0.2 sd) relative to the unrelated funds. Taken together, we reject information sharing between dealers and their affiliate funds—exactly where China Walls are present—and detect extensive information sharing elsewhere, both among affiliated funds and between dealers and their clients. We conclude that China Walls are effectively enforced.

[2]We do not observe who initiated each trade. To proxy the direction of each event trade, we assume that (i) any trade between a dealer and a fund is initiated by the fund, and (ii) for event trades between two dealers, the event dealer initiated the trade. We focus on gross volume, because these assumptions add noise to our net-volume estimates.

We address two key threats to this conclusion. First, trades between dealers and their client funds might generate mechanical increases in trading around event trades. We exclude, from each event, any fund or dealer that trades with the event firm on or up to five days after the event day, precluding such mechanical effects. Second, funds affiliated to each other may be exposed to common shocks through shared dealer connections. To remove these shocks, we flexibly control for overlaps in the sets of dealer connections between the event fund and its affiliate funds.

Section 6 scours granular cells of key event, asset, and firm characteristics for China Wall violations. We never detect information sharing between affiliate dealers and funds across crisis and noncrisis periods, asset classes, currency pairs, and fund types. We consistently detect information sharing among affiliated funds and between dealers and their clients. Moreover, hedge funds respond more intensely to event trades than other funds, and particularly so when the event trade was by another hedge fund (Table 4), echoing evidence that hedge funds are more informed and more sensitive to information than other funds (Di Maggio, Franzoni, Kermani, and Sommavilla, 2019; Kumar et al., 2020). Events are more likely during crisis periods, and yet treated firms' responses to crisis and noncrises events are precisely equal (Table 5)—our design fully strips away aggregate shocks and any variation that correlates with such shocks. Last, an event prompts the largest responses by connected firms that specialize in the currency pair and asset class of the event trade (Tables 6 and 7): homophily one would expect if the event trades indicate the type of information they embody. It is hard to imagine a confounder that can plausibly explain this rich combination of results.

Context and previous work on China Walls. The US regulators did not enforce China Walls before 2018. Previously, banking conglomerates voluntarily adopted China Walls to protect against corporate liability from insider trading by their employees. (The employees themselves would remain liable.) The 2010 Dodd-Frank Act allowed US regulators to conduct "risk-based" enforcement, under which they can prosecute firms for practices that substantially raise the risk of a crime, even without evidence that the crime has actually occurred. The US Securities and Exchange Commission (SEC) began to exercise this power to enforce China Walls in 2018: insufficiently maintaining China Walls itself, not only insider trading, is now a prosecutable offense. Appendix A provides further detail.

Existing evidence on China Walls exploit samples that predate 2018. This evidence identifies extensive violations, as legal proceedings would eventually confirm.[3] We instead evaluate the China Walls during the recent period of their active enforcement. As importantly, we contribute a novel identification strategy that uses unrelated funds as controls to isolate the effects of information sharing. We validate our strategy in conditions where the China Walls are absent and information sharing is expected. Applying this design to a large and granular dataset yields precisely estimated and robust evidence that today's China Walls effectively preempt information sharing.

Broader contributions. We belong to the literature on the capacity of states to regulate firms.[4] In their settings, a large extent of regulatory enforcement occurs through private litigation by parties involved in the regulated activity (e.g., employer vs employee, insider vs outside shareholder; Glaeser and Shleifer (2003), La Porta, Lopez-De-Silanes, and Shleifer (2006)). In our setting, a China Wall violation involves affiliates under common

[3] 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), Li, Mukherjee, and Sen (2021), Kondor and Pintér (2022), and Haselmann, Leuz, and Schreiber (2023) find evidence of China Wall violations in various settings. The latest in-sample year among them is 2017.

[4] Regulators have greatly reduced pollution (Keiser and Shapiro, 2019; Behrer, Glaeser, Ponzetto, and Shleifer, 2021), insider trading (Bhattacharya and Daouk, 2002), misleading financial disclosures (Greenstone, Oyer, and Vissing-Jorgensen, 2006), and discrimination in pay (Bailey, Helgerman, and Stuart, 2024) and access to accommodation (Cook, Jones, Logan, and Rosé, 2023).

corporate control, eliminating the threat of counterparty litigation. Moreover, bankers often communicate in plausibly deniable ways (Peluso, 2020). Therefore, our results reveal a remarkable regulatory capacity to control information flows beyond what is established in prior work.

We extend the empirical literature on information diffusion in financial markets. Dealers extract information from their clients' order flow (Hortaçsu and Kastl, 2012) and leak information to certain 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, thereby adding China Walls as a promising source of variation in information flows that is especially relevant today, when the financial sector is highly concentrated.

Roadmap. Section 2 develops the empirical design. Section 3 describes the data and performs motivating analyses. Sections 4 and 5 investigate the effectiveness of China Walls. Section 6 performs 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 the walled-off subsidiaries of banking conglomerates and their affiliates. An MPI is any information that (a) a reasonable investor would find important for her investment decisions and (b) is not publicly disclosed.[5] For example, proprietary analysis, inside information, or private trade requests would constitute MPI. Typical China Walls require walled-off subsidiaries to be isolated via separate entrances, opaque and soundproof barriers, and the monitoring and recording of their employees' communications.

New regulations since the 2008 financial crisis established today's China Walls around broker-dealers within banking conglomerates (and bank-owned investment advisers, which we do not examine). Today, the US SEC routinely imposes large fines for deficiencies in the dealers' China Walls. Appendix A details relevant definitions, history and legal precedents, impacts of the Dodd-Frank Act, and recent enforcement cases.

Empirical setting. The foreign exchange market is an over-the-counter market, in which trades occur between dealers or a dealer and its client. The dealers are long-lived, trades are nonanonymous, 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.[6] More broadly, financial

[5]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.

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

regulations in Israel are largely based on the US. A peculiar Israeli law forbids Israeli holding companies from owning both a dealer and a nondealer 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).

2.2 Empirical Design

We must overcome three challenges to test the hypothesis that the China Walls are effectively enforced. First, the China Walls may be violated in circumstances that we do not examine. In particular, the test may neglect the circumstances where China Wall violations are the most likely to occur. Second, we need a proxy that isolates the variation specifically due to bilateral MPI sharing, all while reliably detecting any such information sharing. Third, the bank-owned dealers may not share MPI with their affiliates even if their China Walls were absent, in which case enforcement is moot.

Defining events. We seek events that pinpoint when a dealer or a fund receives especially valuable MPI. Under the plausible assumption that China Wall violations are the most likely to occur when there are the largest gains from sharing information with affiliates, rejecting violations during such events would also rule out violations at other times. Standard theory shows that an informed trader submits larger trades when she holds more valuable private information (Kyle, 1985; Easley and O'Hara, 1987). Empirically, the trades that are unusually large compared to its trader's other trades are particularly 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 dealer or a fund (a firm) and a day (event day) when the event firm makes a trade (event trade) that is exceptionally large compared to the its other trades.

Isolating information sharing. We consider an event firm i and a treated firm j such that, if i is a dealer, then j is either an affiliate fund or a client (i.e., connected) fund. A proxy for MPI sharing from the event firm i to the treated firm j must isolate information that is (i) material and (ii) bilaterally shared. 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. Alternatively, j may become more likely to trade in the direction of the price change that the MPI predicts. Therefore, we choose increases in the gross volume of firm j to proxy for the sharing of MPI by the event firm i to j, and confirm that our results are robust to using net volumes signed in the direction of the event trades.

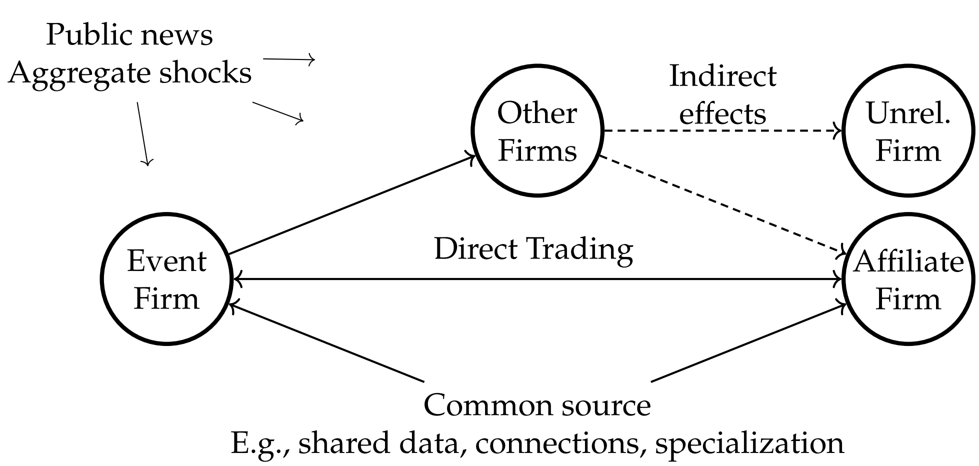

Figure 2: Potential Confounders to Measuring Bilateral Information Sharing

To isolate bilateral MPI sharing, we remove each of the four confounders that can also generate the coincidence of firm i's exceptionally large trades and increases in firm j's gross volume. Figure 2 illustrates the confounders. First, any direct trade between i and j could mechanically induce both the event trade and the increased gross volume, for instance as the large trade itself causes the increase in gross volume. This confounder does not apply when firms i and j are affiliated or unrelated (i.e., neither affiliated nor ever trades with i in our sample), since the former do not trade in our data and the latter do not by definition. In the case where i and j are connected, we shut down the confounder by excluding firm j for the events in which j traded with i on or after the event day up to the end of the event window (five days after the event day).

Second, arrivals of public news or other aggregate shocks may trigger all firms to trade, including the event trades. Third, event firm i's MPI that corresponds to the event trade may indirectly induce firm j to increase its trading activity. Either the event trade itself or any sharing of the MPI by i to firms other than j could affect liquidity or prices throughout the market. As these liquidity or price impacts reach firm j, they may prompt j to increase trading. For example, if an event trade is between dealer i and another dealer, the second dealer might contact fund j to offload the newly gained inventory. Fund j's gross volume would then increase if it agrees to this trade, or if the contact reveals information to j.

We filter out the aggregate-shock and the indirect-impact channels by comparing the gross volume of firm j and those of the firms that are unrelated to the event firm i. The gross volumes of the unrelated firms would capture the aggregate shocks and the indirect impacts of firm i's event trade around the event day. At the same time, the event firm i would not share MPI with unrelated firms. Hence, our design detects bilateral MPI sharing from i to j if and only if the gross volume of j increases relative to unrelated firms on or after the event day.

Fourth, a source common to firms i and j, but not to the unrelated firms, may simultaneously trigger i's event trade and heighten j's gross volume. We present three examples, in which we suppose that i and j are affiliated or connected with each other. Firms i and j may be more likely than two unrelated firms to have common data or research sources. They may also be more likely to specialize in the same assets, or be connected to the same third firm. In these examples, any information from a common data source would sometimes reach j before i, and likewise for those from asset-specific shocks or the shared connection. This noise in timing would generate pretrends before the event dates. We reject the presence of the common-sources channel if j and unrelated firms show parallel trends prior to the event day.

Connected treatment. 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).[7] 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 or its largest trade if the firm made fewer than 1,000 trades in our sample. All event trades by the same firm on the same day are combined into a single event. A firm is treated on or after the event day within the event window if the firm is an affiliate or a connection of the event firm. A firm is a control if it is unrelated to the event firm and not treated in any other event during the event window. Our event window is the 11 trading days around the event day, because exchange rates fully incorporate private information in about a trading week (Menkhoff, Sarno, Schmeling, and Schrimpf, 2016).

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 21-day panel around event e. 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

[7]This implementation yields average treatment-on-the-treated (ATT) effect estimates that always place positive weights on all groups (Gardner, 2022), unlike those of traditional staggered two-way fixed-effects difference-in-differences specifications (Roth, Sant'Anna, Bilinski, and Poe, 2023).

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,

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

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.

2.4 Identification Tests

We assume that a firm trades an exceptionally large size when especially valuable material private information arrives at the firm. The underlying claim is that firms submit 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 in over-the-counter markets that larger trades by each firm are more predictive of returns than its smaller trades (Kumar et al., 2020; Pinter et al., 2024). However, the theory on order splitting (Bernhardt and Hughson, 1997) and the lack of similar evidence specifically on the foreign exchange market question our claim.

Appendix B adjudicates our assumptions in the data. Placebo tests using small and medium trades as event trades jointly falsify two claims. (i) Exceptionally large (0.1 percentile) trades indicate arrivals of especially valuable MPI. (ii) Our design isolates bilateral MPI sharing, in that it yields significantly positive treatment coefficients if and only if the event firms bilaterally share MPI with the treated firms. 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. If informed firms do not split orders, and rather trade exceptionally large sizes to exploit especially valuable MPI, the large trades would predict returns and smaller trades would not. Moreover, if our design isolates information sharing, connected firms would increase their volumes relative to unrelated firms around the dates of trades that predict returns, and not for nonpredictive trades.

We find that the exceptionally large trades predict returns up to three days following the trade. The small and medium trades do not predict returns. Moreover, we find zero evidence of an increase in the gross volume of connected firms relative to unrelated firms around the small-event or the median-event days. We conclude that our design isolates the sharing of especially valuable MPI.

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.[8] Each observation specifies the currency pair (ILS and another currency), price, quantity, date and time,[9] 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 analyze. A three-step process generates samples whose observations are firm-by-date. First, we consolidate dealers up to the conglomerate by dropping all trades between dealers affiliated to each other and combining them under conglomerate-level labels. We do so, because a group of affiliate dealers are free to split incoming orders and transfer assets and capital among themselves, and are thus effectively a single economic entity. As such, consolidating affiliate dealers minimizes

[8] 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/.

[9] We do not exploit intraday time stamps, because a large proportion of trades report 00:00:00 rather than the actual trade time. (The Bank only requires that firms report the correct date, not time.)

noise from nonmarket transactions that shift cash and inventory for tax or balance sheet purposes.[10] Second, we aggregate the daily gross volumes of each dealer and fund across asset classes (i.e., spot, forward, and swap). While aggregating, we keep the notional amount from each swap trade's first leg and ignore its second leg to avoid double counting. Third, we winsorize observations in the top 0.5 percentile by daily gross volume, calculated separately for dealers and for funds, because their daily volume distributions differ dramatically.

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/). We assign affiliations to firms as of last quarter, 2023, for the whole sample period, because financial firms rarely change their affiliations and typically change their legal names when they do. Second, all firms with obviously indicative names are linked to the corresponding conglomerate (e.g., "Deutsche Bank Luxembourg S.A."). Third, the remaining firm names are entered into `ChatGPT 4.0` as a query in the form, "as of `[date the firm last appears in the sample]`, is `[firm name]` independent? If not, which holding company does `[firm name]` belong to?" Fourth, we manually verify each answer generated in step three, by searching for the firm name paired with "independent" or the `ChatGPT`-suggested holding company name.

[10]Some 8% of foreign exchange spot trades are "back-to-back" trades between affiliate dealers for accounting or inventory rebalancing reasons (Bank for International Settlements, 2022). All trades by affiliate funds are market-based, since they only trade with nonaffiliate dealers.

Table 1: Sample Characteristics

	All trades	Fund trades	Final Sample: Dealer day	Final Sample: 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 nonfinancial 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.

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, due to Israeli law that forbids common ownership of banks and dealers.

3.2 Motivating Analyses

Three analyses motivate our main empirical design. First, Figure 3 plots the total daily dollar volume of transactions. The dealers trade USD2.8 billion with the funds daily, of which near precisely zero is with their affiliate funds—there are four trades between a dealer and an affiliate fund, worth a mere USD5.51, in our sample.

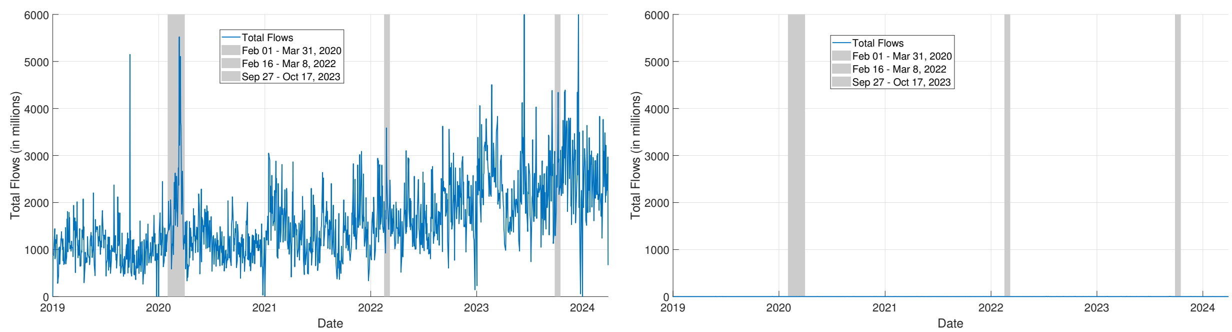

(a) Unaffiliated Dealers and Funds

(b) Affiliate Dealers and Funds

Figure 3: Daily Gross Dollar Volumes Traded Between Dealers and 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.

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 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

$$CorrGV_{ijl} = a_l Affiliate_{ij} + b_l Connected_{ij} + c_i + d_j + \varepsilon_{ijl}. \tag{3}$$

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 c_i and d_j control for time-invariant factors specific to each dealer and each fund.

Figure 4b plots the coefficients c_l and d_l across $l = -10 \ldots 10$. Daily gross volumes of the affiliate dealer-fund pairs are no more correlated than those of the unrelated pairs across all lags l. In contrast, the connected dealer-fund pairs are sharply more correlated contemporaneously than the unrelated pairs. These correlations suggest that China Walls effectively block material information flows between walled-off firms, while information flows freely among connected firms. Our main design isolates bilateral information sharing and focuses on the dates when there is the greatest incentive to violate the China Walls.

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 gross volumes on the event day by -0.02 standard deviation (std. error: 0.04 sd), whereas the connected funds increase theirs by 1.9 sd (std. error: 0.007 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

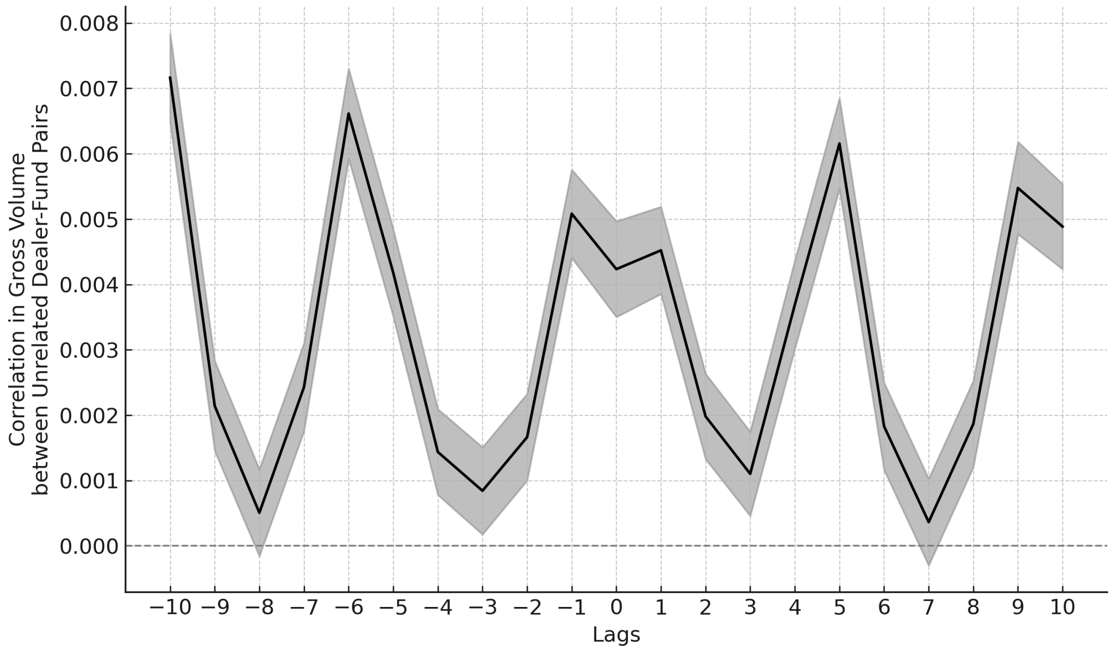

(a) Unrelated Pairs

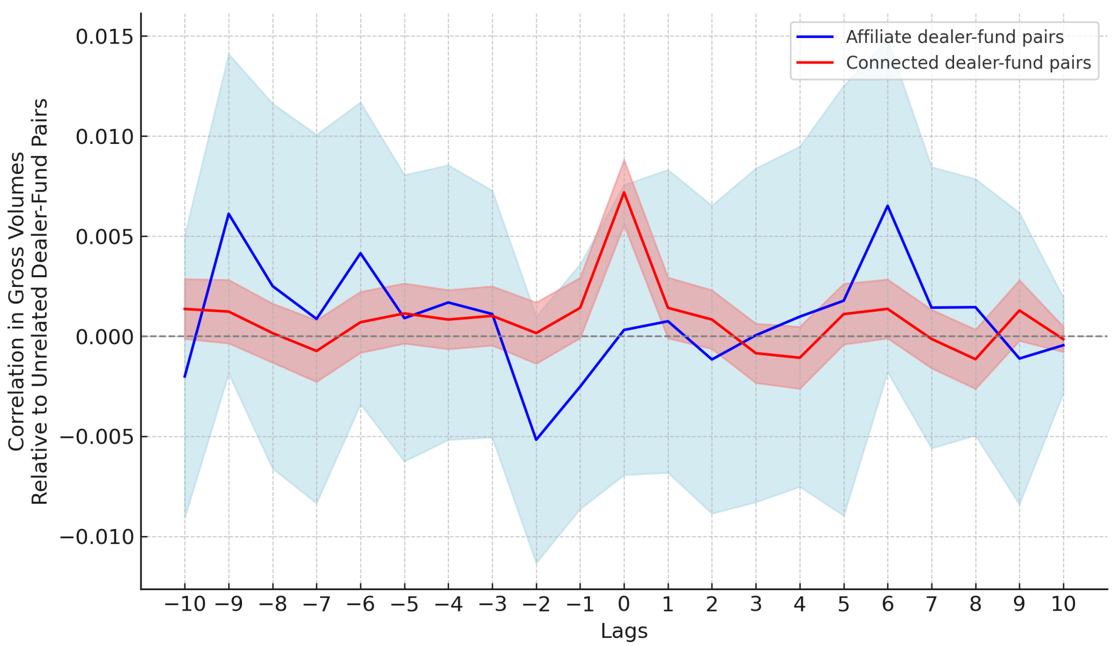

(b) Affiliate and Connected Pairs using Unrelated Pairs as Controls

Figure 4: Correlations in Daily Gross Volumes within Dealer-Fund Pairs

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 the estimated increase in their gross volumes is precisely nil. The connected dealers increase their gross volumes by 0.26 sd (std. error: 0.02 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.[11] In contrast, the connected dealers and funds respond strongly to each other's information, with estimated coefficients in the multiples of the affiliate coefficients. By far the most responsive are the funds to the information from their affiliate funds. Altogether, the pooled results confirm that the China Walls are effective enforced.

5 Do Affiliates Without China Walls Share Information?

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 affiliate funds bilaterally share MPI with each other, we infer that dealers and their affiliate funds would also share MPI in the absence of China Walls.

[11] The fund-to-dealer specification has far fewer events and observations than the dealer-to-fund specification, because there are many more funds than dealers. Each event fund has no more than one affiliate dealer and a few connected dealers, whereas each event dealer has several affiliate funds and numerous connected funds.

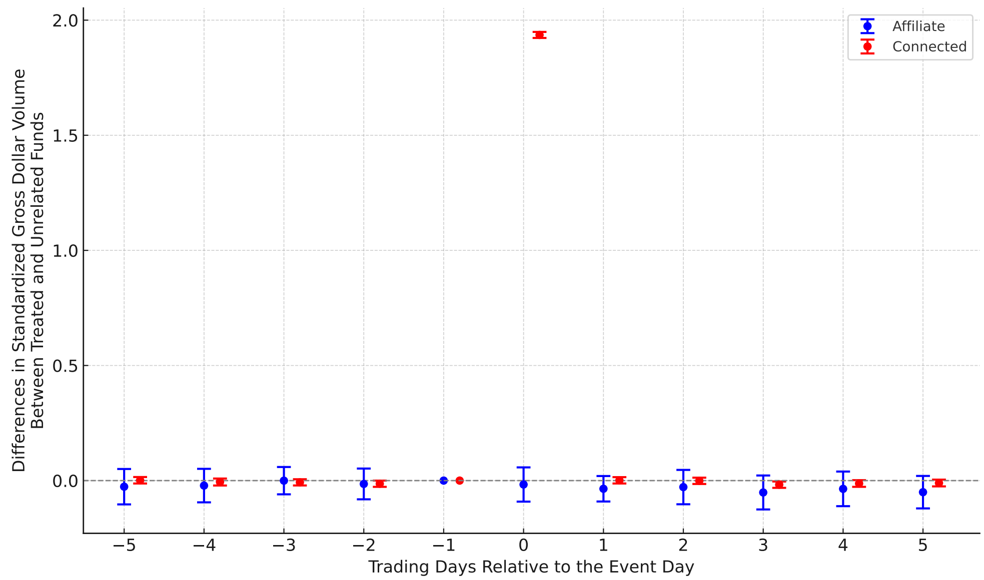

(a) Fund Responses to Dealer Information

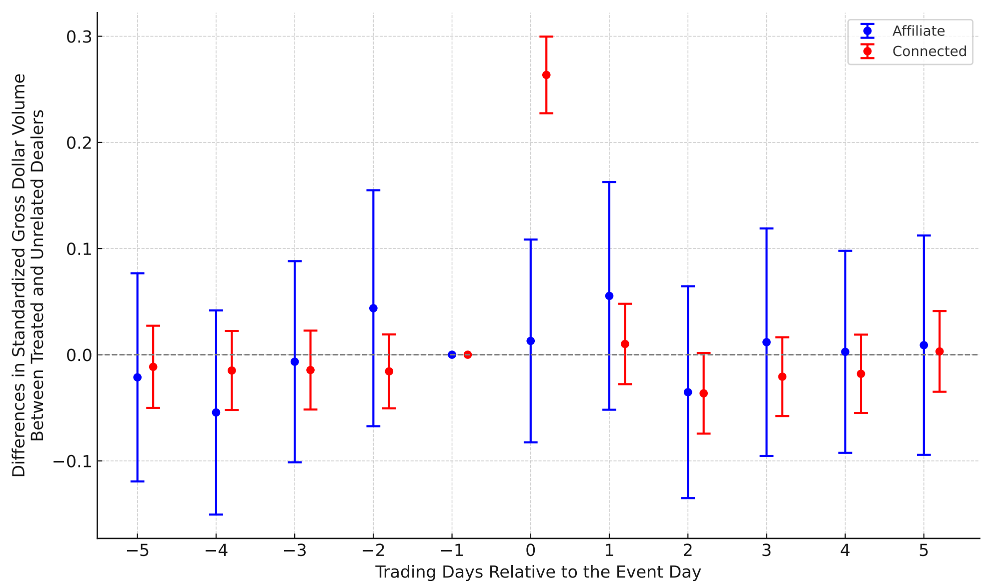

(b) Dealer Responses to Fund Information

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

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

	D2F Affiliate	F2D Affiliate	D2F Connected	F2D Connected	F2F Affiliate
$Post \times Affiliate$	-0.024	0.012			0.23***
	[0.017]	[0.028]			[0.021]
$Post \times Connected$			0.32***	0.034**	
			[0.0047]	[0.014]	
$Post \times DealerOverlap$					0.01***
					[0.004]
$Post \times Affiliate$					0.25***
$\times DealerOverlap$					[0.026]
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	0.022	0.41	-0.007	0.45	0.068
Within R-squared	0	0.0001	0.0005	0.0001	0.0001
Events	7,710	7,894	7,710	7,894	7,894
Observations	89,005,179	4,156,128	42,150,672	3,614,383	12,664,366

Coefficient estimates from the pooled counterparts to Equations (1), (2) and (4). The dependent variable is the standardized daily gross US dollar volume of a firm winsorized at the top 0.5 percentile. An event is a firm and a day when the firm made a trade in the 0.1 percentile among its trades. Each event window is 11 days around the event day. 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 in our sample, and do not trade with the event firm on or after the event day. Affiliate and Connected treatments are mutually exclusive, because no dealer trades with an affiliate fund in our sample. Controls includes firms that are unaffiliated and never trades with the event firm, and are not treated in another event throughout the event window. We include event-by-firm, calendar date, and days-since-event fixed effects. *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. *F2F*: All firms are funds. *DealerOverlap* indicates a treated or control fund whose set of connected dealers overlaps with that of the event fund. Standard errors in square brackets are clustered at the event-by-firm and date levels. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

5.1 Design and Implementation

Figure 6 depicts the 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.

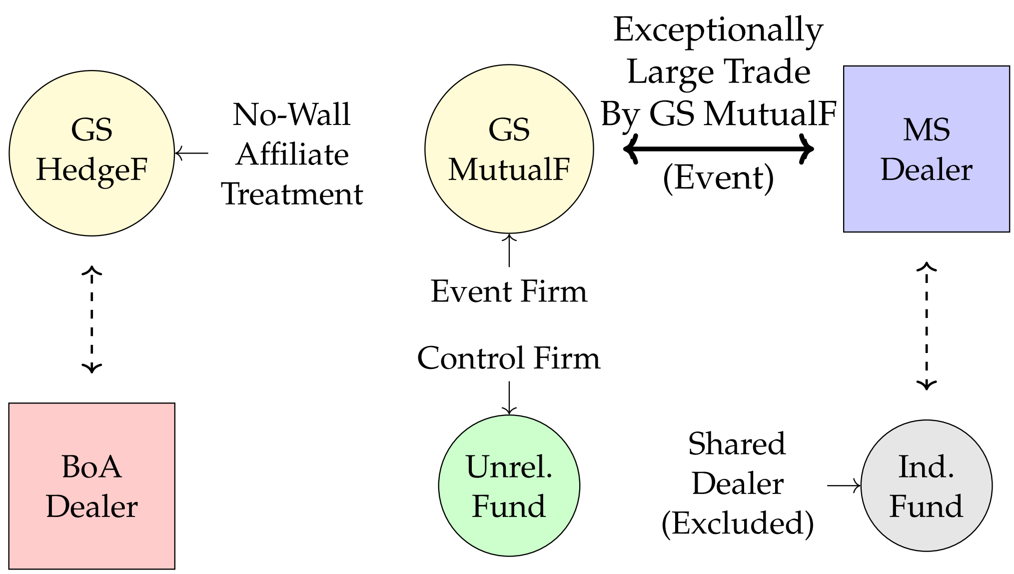

Figure 6: Identification: Information Sharing Between Affiliate Funds

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.

We apply the following specification to implement this design on the subsample of funds:

$$\begin{aligned} Y_{ejt} = & \sum_{\tau=-5}^{5} \nu_\tau \mathbb{1}_{t=\ell_e+\tau} Affiliate_{ej} + \delta_{ej} + \varphi_t + \sum_{\tau=-5}^{5} \gamma_\tau \mathbb{1}_{t=\ell_e+\tau} \\ & + \sum_{\tau=-5}^{5} \kappa_\tau \mathbb{1}_{t=\ell_e+\tau} Affiliate_{ej} DealerOverlap_{ej} \\ & + \sum_{\tau=-5}^{5} \eta_\tau \mathbb{1}_{t=\ell_e+\tau} DealerOverlap_{ej} + \varepsilon_{ejt}. \end{aligned} \tag{4}$$

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.

5.2 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 (4) 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 their gross volumes by a precisely estimated 1.7 standard deviations (std. error: 0.2 sd) on the event date. The large size of this response is consistent with affiliated funds being eager to share information among themselves. 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 2.6 sd (std. error: 0.3 sd).

6 Heterogeneity

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

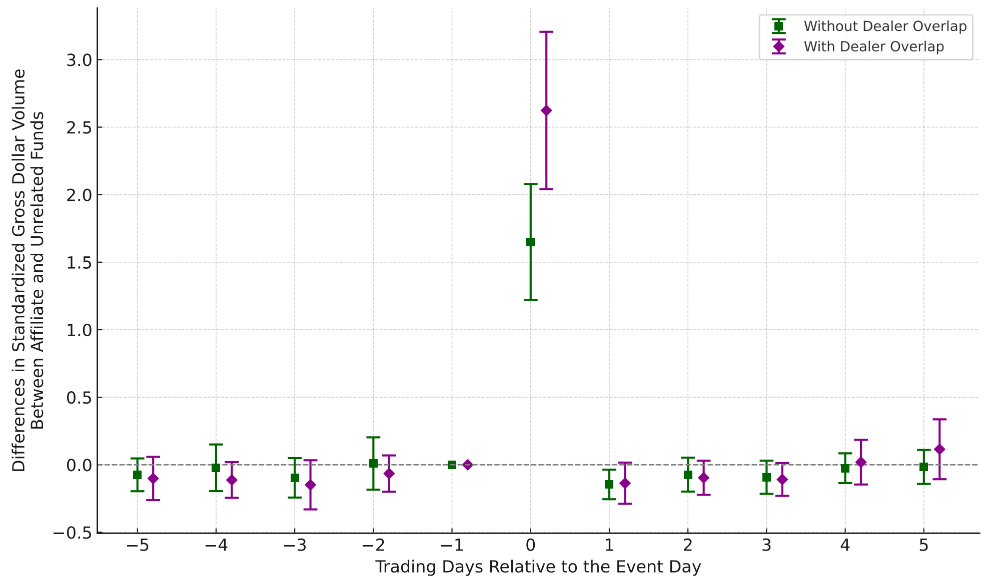

Figure 7: Affiliate Fund Response to Event Fund Information

affiliate dealers and funds would not share MPI absent the China Walls.

To do so, we repeat the analyses of Section 4 across cells of fund types, currency pair, and asset class (i.e., spot, forward, or swap), and for event during crisis and noncrisis periods. Since our main specification yields a null result, we maximize the power to detect deviations from the null by interacting dummy variables corresponding to each characteristic with the complete set of terms in the pooled counterparts of Equations (1), (2) and (4). (Rather than splitting our sample across those characteristics.) We examine both event-level and firm-level characteristics. The dummy $HedgeFund = 1$ if the treated or control firm is a hedge fund, and $HedgeFundEvent = 1$ if the event firm is a hedge fund. Other firm-level dummies indicate whether a firm's share of trades in a currency pair or asset class is greater than the median across firms, separately for dealers and for funds. Event-level dummies indicate whether the event trade was in a given currency pair or asset class, and whether the event occurred during the crisis periods following the

2020 Covid shock, the 2022 Russian invasion of Ukraine, and the 2023 Hamas attack.

6.1 Hedge Funds versus Other Funds

Table 4 separates the responses of hedge funds and nonhedge funds to events by hedge funds and nonhedge funds.

6.2 Crisis Periods

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

6.3 Currency Pairs and Asset Classes

Tables 6 and 7 present the treated firms' responses split by currency pair and asset class. Each table cell is the increase in the treated firms' daily gross volumes relative to the unrelated firms on and after the event day, where the firms' specializations and the events belong to the currency pair or asset class specified for the row.

To arrive at each estimate in Table 6, we first run the pooled counterparts to Equations (1), (2) and (4) augmented with the complete set of interactions involving four dummies: $USDFirm_f = 1$ if the firm's dollar-value share of trades in USD-ILS pair over our sample exceeds the median across firms (i.e., the firm "specializes in USD"), $USDEvent_e = 1$ if the event trade (or any event trade for every event with multiple event trades) was for USD-ILS, and $NonUSDFirm_f$ and $NonUSDEvent_e$ are defined

Table 4: Responses by Fund Type

	D2F Affiliate	F2D Affiliate	D2F Connected	F2D Connected	F2F Affiliate
$Post \times Affiliate$	-0.023	0.026			0.157***
	[0.018]	[0.030]			[0.022]
$Post \times Affiliate \times Hedge\ Fund$	-0.0068				0.176
	[0.045]				[0.32]
$Post \times Affiliate \times HF\ Event\ Trade$		0.024			0.047
		[0.14]			[0.14]
$Post \times Connected$			0.122***	0.029*	
			[0.0036]	[0.016]	
$Post \times Connected \times Hedge\ Fund$			0.364***		
			[0.012]		
$Post \times Connected \times HF\ Event\ Trade$				0.112***	
				[0.033]	
$Post \times Affiliate \times HF\ Event\ Trade \times Hedge\ Fund$					2.21***
					[0.39]
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	0.022	0.457	-0.007	0.449	0.049
Within R-squared	0	0.0001	0.0003	0.0002	0.0001
Events	7,710	7,894	7,710	7,894	7,894
Observations	89,005,179	4,156,128	42,150,672	3,614,383	12,664,366

Coefficient estimates from the pooled counterparts to Equations (1), (2) and (4). The dependent variable is the standardized daily gross US dollar volume of a firm winsorized at the top 0.5 percentile. An event is a firm and a day when the firm made a trade in the 0.1 percentile among its trades. Each event window is 11 days around the event day. 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 in our sample, and do not trade with the event firm on or after the event day. Affiliate and Connected treatments are mutually exclusive, because no dealer trades with an affiliate fund in our sample. Controls includes firms that are unaffiliated and never trades with the event firm, and are not treated in another event throughout the event window. We include event-by-firm, calendar date, and days-since-event fixed effects. *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. *F2F*: All firms are funds. Below: F2F estimates exclude treated and control funds whose dealer connections overlap with the event fund. HF Event Trade is an event whose event trade was by a hedge fund. Standard errors in square brackets are clustered at the event-by-firm and date levels. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 5: Responses to Events During Crisis and Noncrisis Periods

	D2F Affiliate	F2D Affiliate	D2F Connected	F2D Connected	F2F Affiliate
$Post \times Affiliate$	-0.021	0.012			0.226***
	[0.018]	[0.028]			[0.062]
$Post \times Affiliate$	-0.055	-0.032			-0.022
$\times Crisis$	[0.076]	[0.095]			[0.201]
$Post \times Connected$			0.319***	0.033**	
			[0.005]	[0.015]	
$Post \times Connected$			0.012	-0.006	
$\times Crisis$			[0.026]	[0.044]	
$Post \times DealerOverlap$					0.014
					[0.012]
$Post \times Affiliate$					0.264***
$\times DealerOverlap$					[0.073]
Event$\times$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	0.022	0.457	-0.007	0.449	0.033
Within R-squared	0	0.0001	0.0005	0.0002	0.0002
Crisis Events	440	3,303	440	3,303	3,303
Total Events	7,710	7,894	7,710	7,894	7,894
Observations	89,005,179	4,156,128	42,150,672	3,614,383	12,664,366

Coefficient estimates from the pooled counterparts to Equations (1), (2) and (4). The dependent variable is the standardized daily gross US dollar volume of a firm winsorized at the top 0.5 percentile. An event is a firm and a day when the firm made a trade in the 0.1 percentile among its trades. Each event window is 11 days around the event day. 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 in our sample, and do not trade with the event firm on or after the event day. Affiliate and Connected treatments are mutually exclusive, because no dealer trades with an affiliate fund in our sample. Controls includes firms that are unaffiliated and never trades with the event firm, and are not treated in another event throughout the event window. We include event-by-firm, calendar date, and days-since-event fixed effects. *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. *F2F*: All firms are funds. $DealerOverlap$ indicates a treated or control fund whose set of connected dealers overlaps with that of the event fund. *Crisis*: Event occured during the beginnings of Covid pandemic, the Russian Invasion of Ukraine, or the Hamas-Israeli War. Standard errors in square brackets are clustered at the event-by-firm and date levels. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

analogously. We separately compute the medians for dealers and for funds when assigning $USDFirm_f$ and $NonUSDFirm_f$. We further interact the NonUSD dummies with $MatchingNonUSD_{ef}$, which equals one if $NonUSDFirm_f = NonUSDEvent_e = 1$ and the firm's currency specialization is in the currency of the event trade. Table 6 displays the appropriate sums of the estimated coefficients, and uses the coefficients' covariance matrices to obtain the standard errors of those sums. The estimate in the USD Event Trade-USD Firm by D2F Connected cell, for example, is the sum of coefficients that remain when we set $USDFirm_f = USDEvent_e = 1$ and $NonUSDFirm_f = NonUSDEvent_e = 0$.

Table 7 estimates are calculated in the same way as Table 6, except using dummies corresponding to spot, forward, and swap asset classes, rather than currency pairs.

Table 6: Responses by Currency

	D2F Affiliate	F2D Affiliate	D2F Connected	F2D Connected	F2F Affiliate
USD Event Trade -USD Firm	-0.019	0.021	0.376***	0.034*	0.561***
	[0.035]	[0.073]	[0.006]	[0.019]	[0.219]
USD Event Trade -NonUSD Firm	-0.029	-0.077	0.280***	0.038*	0.194***
	[0.021]	[0.080]	[0.004]	[0.022]	[0.067]
NonUSD Event Trade -USD Firm	0.041	0.122	0.121***	-0.062	0.313***
	[0.033]	[0.108]	[0.007]	[0.052]	[0.097]
NonUSD Event Trade -NonUSD Firm (Matching Currency)	-0.074	0.029	0.391***	-0.010	0.462***
	[0.257]	[0.213]	[0.019]	[0.045]	[0.167]
NonUSD Event Trade -NonUSD Firm (Nonmatching)	-0.028	-0.033	0.137***	0.007	0.225***
	[0.030]	[0.063]	[0.005]	[0.070]	[0.087]
Adjusted R-squared	0.022	0.290	-0.016	0.449	0.033
Within R-squared	0	0.0001	0.0011	0.0002	0.0002
Events	7,710	7,894	7,710	7,894	7,894
Observations	89,005,179	4,156,128	42,150,672	3,614,383	12,664,366

Coefficient estimates from the pooled counterparts to Equations (1), (2) and (4). The dependent variable is the standardized daily gross US dollar volume of a firm winsorized at the top 0.5 percentile. An event is a firm and a day when the firm made a trade in the 0.1 percentile among its trades. Each event window is 11 days around the event day. 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 in our sample, and do not trade with the event firm on or after the event day. Affiliate and Connected treatments are mutually exclusive, because no dealer trades with an affiliate fund in our sample. Controls includes firms that are unaffiliated and never trades with the event firm, and are not treated in another event throughout the event window. We include event-by-firm, calendar date, and days-since-event fixed effects. *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. *F2F*: All firms are funds. Below: F2F estimates exclude treated and control funds whose dealer connections overlap with the event fund. USD Event Trade is an event whose event trade was a USD trade. USD Firm is a treated or control fund (dealer) whose share of trades by dollar volume involving the USD is above the median across all funds (dealers). Standard errors in square brackets are clustered at the event-by-firm and date levels. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 7: Responses by Asset Class

	D2F Affiliate	F2D Affiliate	D2F Connected	F2D Connected	F2F Affiliate
Spot Event Trade	0.013	-0.031	0.520***	0.068	0.203***
-Spot Firm	[0.021]	[0.049]	[0.008]	[0.049]	[0.029]
Spot Event Trade	0.044	0.026	0.323***	0.049**	0.250***
-Forward Firm	[0.035]	[0.055]	[0.009]	[0.023]	[0.044]
Spot Event Trade	0.007	0.041	0.318***	0.028	0.174**
-Swap Firm	[0.044]	[0.097]	[0.015]	[0.054]	[0.088]
Forward Event Trade	-0.052	-0.104	0.226***	0.118**	0.248***
-Spot Firm	[0.040]	[0.081]	[0.010]	[0.050]	[0.038]
Forward Event Trade	-0.082	0.004	0.241***	0.099***	0.300***
-Forward Firm	[0.066]	[0.066]	[0.012]	[0.027]	[0.056]
Forward Event Trade	0.026	-0.200	0.134***	0.079	0.301**
-Swap Firm	[0.112]	[0.157]	[0.019]	[0.055]	[0.134]
Swap Event Trade	-0.039	0.149	0.219***	0.019	0.246***
-Spot Firm	[0.026]	[0.147]	[0.005]	[0.054]	[0.073]
Swap Event Trade	-0.029	0.240	0.237***	-0.000	0.224*
-Forward Firm	[0.027]	[0.258]	[0.006]	[0.031]	[0.137]
Swap Event Trade	-0.003	0.280	0.485***	-0.020	0.645*
-Swap Firm	[0.030]	[0.412]	[0.009]	[0.059]	[0.338]
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	0.030	0.489	-0.007	0.449	0.051
Within R-squared	0	0.0001	0.0007	0.0002	0.0002
Events	7,710	7,894	7,710	7,894	7,894
Observations	89,005,179	4,156,128	42,150,672	3,614,383	12,664,366

Coefficient estimates from the pooled counterparts to Equations (1), (2) and (4). The dependent variable is the standardized daily gross US dollar volume of a firm winsorized at the top 0.5 percentile. An event is a firm and a day when the firm made a trade in the 0.1 percentile among its trades. Each event window is 11 days around the event day. 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 in our sample, and do not trade with the event firm on or after the event day. Affiliate and Connected treatments are mutually exclusive, because no dealer trades with an affiliate fund in our sample. Controls includes firms that are unaffiliated and never trades with the event firm, and are not treated in another event throughout the event window. We include event-by-firm, calendar date, and days-since-event fixed effects. *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. *F2F*: All firms are funds. Below: F2F estimates exclude treated and control funds whose dealer connections overlap with the event fund. Event Trades are separated by asset class. Spot Firm is a treated or control fund (dealer) whose share of trades by dollar value involving spot trades is above the median across all funds (dealers). Forward Firm and Swap Firm are defined analogously. Standard errors in square brackets are clustered at the event-by-firm and date levels. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

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 8 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.

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

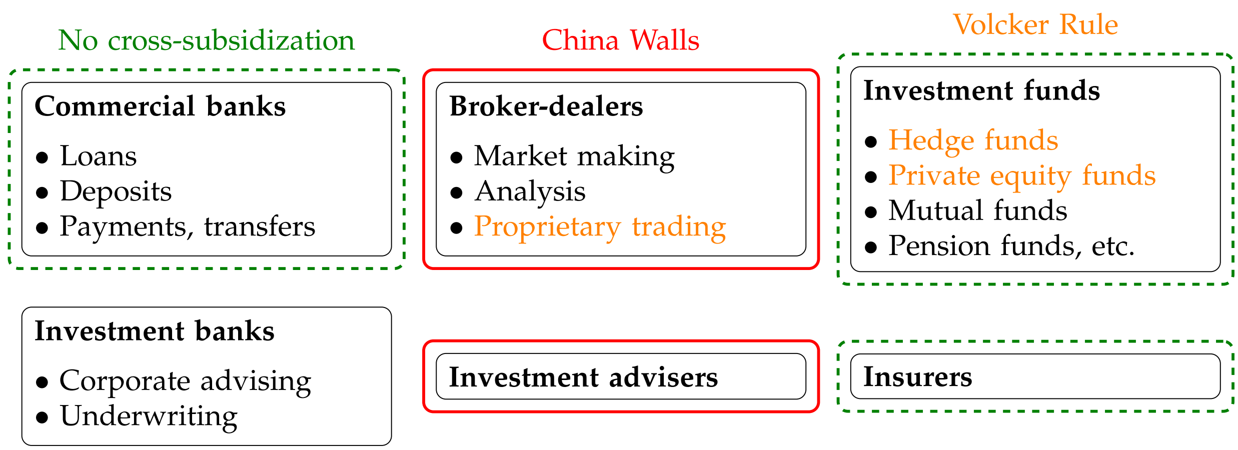

Figure 8: Stylized Banking Conglomerate and Relevant Legal Restrictions

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

preciably affect market prices. In practice, common-law courts treat as MNPI any non-publicly disclosed information that reasonable investors in the relevant securities would find important for their investment decisions. For example, insider earnings information or outstanding order flows of clients are MNPI.[12] 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.

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

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

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

A.2 Key Regulations on Banking Conglomerates

The markings in Figure 8 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.[13] 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 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

[13] 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.)

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.[14]

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

[14] A typical case is *Black and Shearson v. Hammill Co.* (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.

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.[15] The logic is that walled-off subsidiaries can be considered separate entities for the purpose of determining whether a legal duty has been breached. Continuing the example, the dealer would not owe fiduciary duty to the investors of the affiliate mutual fund if this fund were walled off from the dealer. The US financial conglomerates widely adopted the China Walls, which became broadly standardized according to SEC guidelines. Financial conglomerates in other jurisdictions followed, whether through their US operations or regulatory standardization (in Australia, Canada, France, Germany, Japan, Switzerland, and the UK).

Pre-2008 crisis legal status. A 1980 US Supreme Court case (*Chiarella v. United States*) 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.

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

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

statutes explicitly incorporated the China Walls as safe harbor and broadened their legal protections (Brooke, Burrows, Faber, Harpum, and Silber, 1995, p. 98).[16] Suppose a fund consistently earns large profits whenever an affiliate dealer receives large order flows. Under the new statutes, presence of a China Wall between the dealer and the fund would protect the conglomerate against liabilities to the dealer's clients and to the fund's counterparties.[17]

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.[18] The prosecutions over the LIBOR scandal highlights the nonobligatory status of China Walls precrisis: 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

[16] 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.

[17] 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).

[18] 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.

statute of limitations. A firm that aided a violator could only be prosecuted if the firm knowingly assisted in the violation, a high legal bar. Most importantly, regulatory action required the evidence of actual fraud or breach of duty. Engaging in transactions with a high risk of fraud or duty breach, or failing to maintain China Walls that could greatly suppress the misuse of MNPI were not themselves actionable by regulators.

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

Today's China Walls form 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 to proactive.

B Placebo Results

Two exercises jointly test two identifying assumptions that: (a) Exceptionally large trades pinpoint the arrivals of especially valuable MPI; and (b) our design yields a significant and positive coefficient at $t = 0$ if and only if the event firms bilaterally share MPI with the treated firms.

The first exercise is to 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. To sign each event trade, we assume, for each dealer-fund trade, that the fund was the initiator. We assume that the event dealer was the initiator for each interdealer trade.

A three-step procedure obtains the price impact of firm type k, firm's trade-size percentile p, and cumulative return horizon ℓ. First, 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, $d_{t,k,p} = 1$ if its positive, and zero otherwise. Second, we calculate the cumulative returns $R_{t,t+\ell}$ between t and $t+\ell$, $\ell \in \{0,\ldots,9\}$, using Bloomberg benchmark exchange rates. Third, the price impact is the coefficient $\rho_{k,p,\ell}$ in the time-series regression (5):

$$R_{t,t+\ell} = \alpha_{k,p,\ell} + \rho_{k,p,\ell} \cdot d_{t,k,p} + \varepsilon_{t,k,p,\ell}. \tag{5}$$

Figure 9 plots the price impact estimates. The net volumes from exceptionally large trades predict future returns, whereas the median and the small trades do not.

The second exercise replicates Figures 5 and 7, except redefining an event to be a day when a firm makes a median or a small trade. As Figure 10 depicts, across all specifications, every coefficient estimate is insignificant at the 95% confidence level. Combined with Figures 5 and 9, 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.

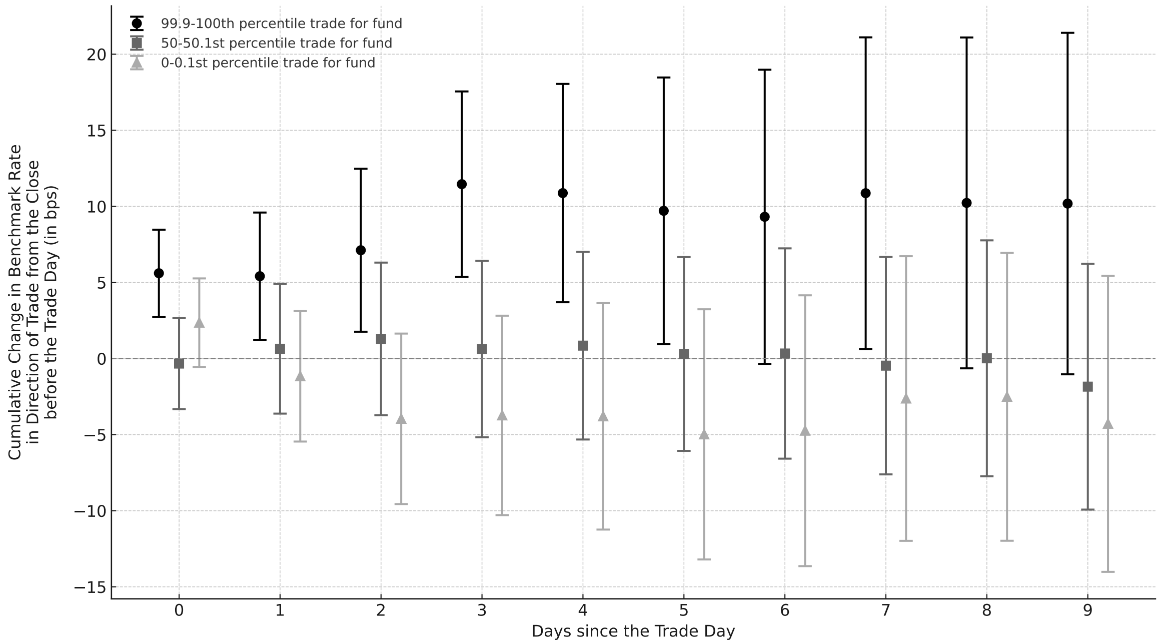

(a) Trades by Funds

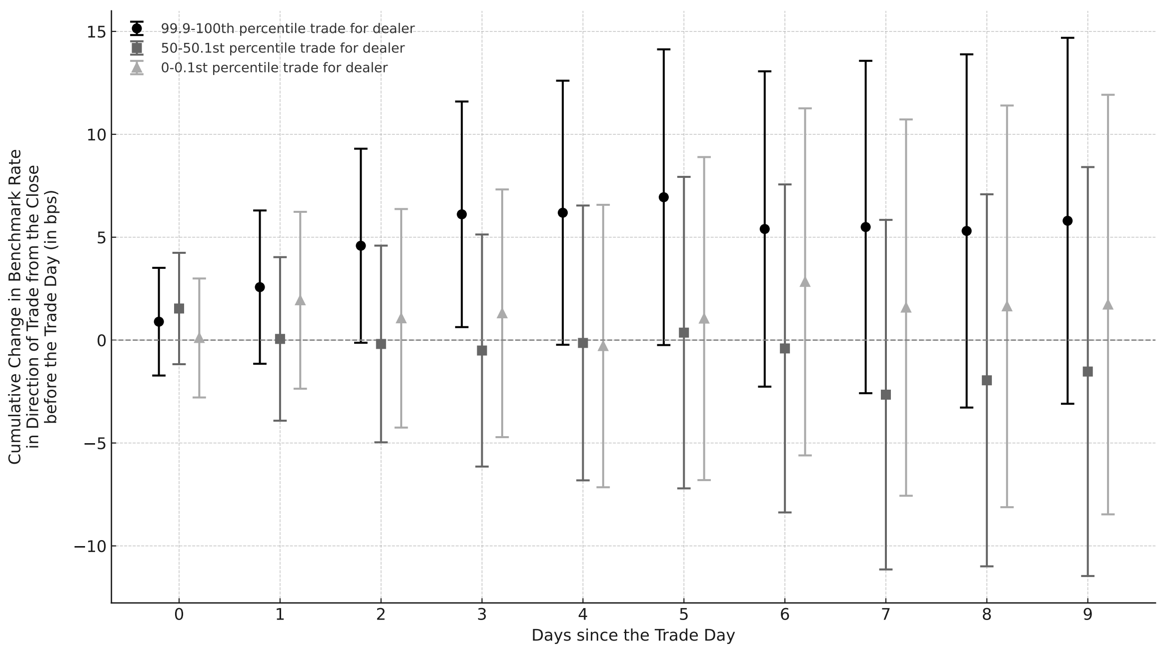

(b) Trades by Dealers

Figure 9: Price Impact Estimates

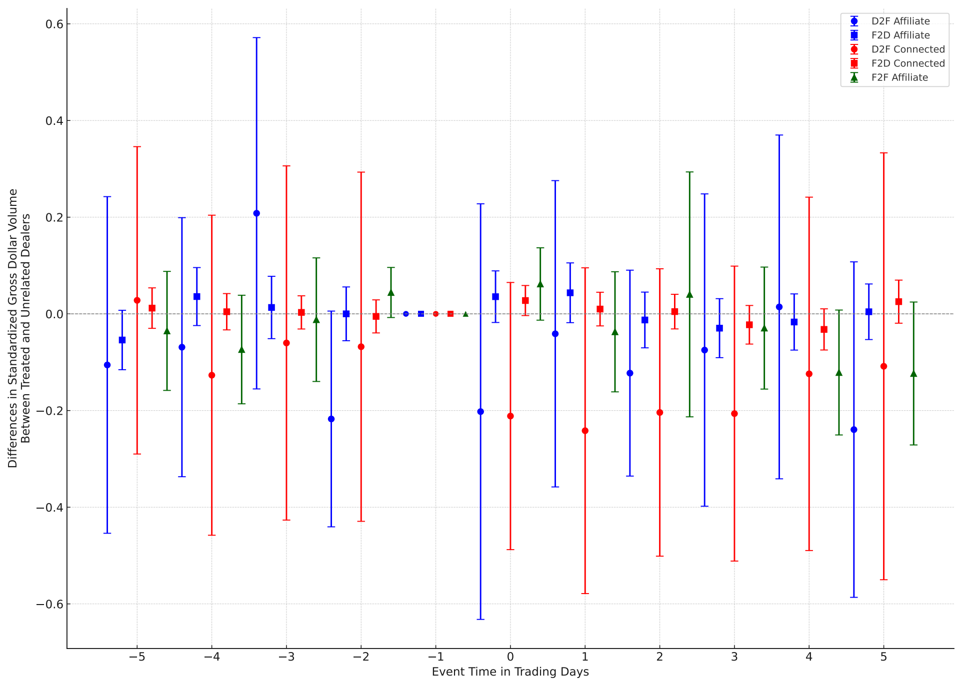

(a) Event Trade in 50 to 50.1st Percentile

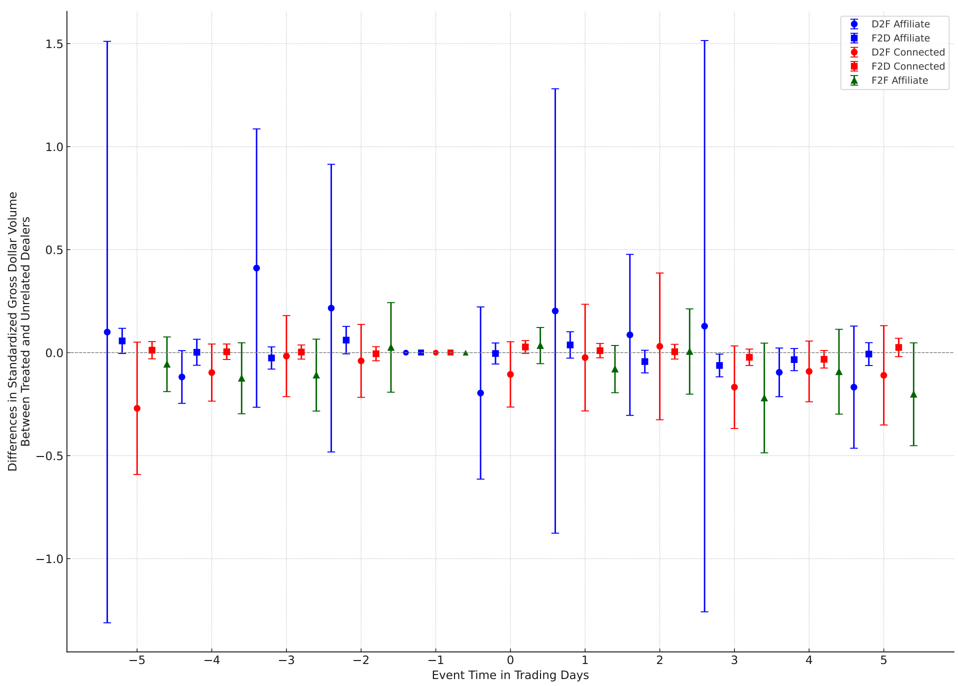

(b) Event Trade in 99.9 to 100th Percentile

Figure 10: Placebo Estimates

References

Abadie, A., S. Athey, G. W. Imbens, and J. M. Wooldridge (2023, February). When Should You Adjust Standard Errors for Clustering? *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. *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. *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. *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. *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. *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. *Quarterly Journal of Economics 121*(2), 399–460.

Hagströmer, B. and A. J. Menkveld (2019). Information Revelation in Decentralized Markets. *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. *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. *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. *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? *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. *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. *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. *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.