China Walls

Tomy Lee Daniel Nathan Chaojun Wang*

March 27, 2026

Abstract

We evaluate the enforcement of information barriers—China Walls—within conglomerates. Our setting is the 25 million trades in 2018–2024 in the Israeli Shekel market, where the SEC tightened the China Walls around dealers from July 2018. Our difference-in-differences design compares the trade volumes and profits of funds that are affiliated with, clients of, or entirely unrelated to a dealer around the days when the dealer is especially likely to hold valuable information. Before July 2018, dealers extensively share information with both their affiliate funds and their clients. After July 2018, the dealers cease sharing information with their affiliate funds, while continuing to do so with their clients unabated. A back-of-the-envelope calculation shows that the China Walls eliminated 86% of profits by the affiliate funds on event days.

JEL classification: G21, G28, G14, G15

Keywords: Banking conglomerates, trading networks, information barriers, information sharing, foreign exchange

*Lee is at the Central European University. Nathan is at Hong Kong Polytechnic University and is affiliated with the Bank of Israel. Wang is at the Wharton School, University of Pennsylvania. We especially thank Adam Szeidl for valuable comments. We thank Jamie Coen, Philip Bond, Andras Danis, Xavier D'Haultfoeuille, Efrat Dressler, Felix Fattinger, Jian Feng, Thomas Gehrig, Vincent Glode, Itay Goldstein, Terrence Hendershott, Simon Jurkatis, Holger Mueller, Florian Nagler, Gabor Pinter, Christo Pirinsky, Robert Richmond, Antonino Emanuele Rizzo, Clara Vega, Toni Whited, Adam Zawadowski, and Josef Zechner for helpful comments. Emails: leeso@ceu.edu; daniel.nathan@polyu.edu.hk; wangchj@wharton.upenn.edu.

1 Introduction

Banking conglomerates are rife with conflicts of interest. They analyze securities while underwriting them, manage their own investment funds while acting as dealers for the funds of others, and much else. Acting on these conflicts contributed to the last financial crisis (Griffin, 2021), and involves affiliates sharing confidential information with each other. In response, regulators tightened information barriers within banking conglomerates—China Walls[1]—to preempt information sharing between conflicted affiliates.

Enforcing China Walls is a formidable regulatory challenge, because it is difficult to detect private information sharing and to investigate every setting where information might be shared. Many policymakers and legal scholars have been skeptical for these reasons (Avci, Schipani, and Seyhun, 2018), with some calling for banking conglomerates to be broken up as an alternative way to remedy conflicts of interest.[2] To address their concerns, the Dodd-Frank Act authorized the US Securities and Exchange Commission (SEC) to prosecute banking conglomerates for failing to maintain China Walls around their dealer arms, even if the SEC cannot prove a misuse of confidential information. The SEC began to fully exercise this authority in July 2018, when it fined Mizuho Securities $1.25 million for lacking effective China Walls. Unlike in prior cases, the SEC did not cite *any* evidence that information was misused. It is unknown whether today's heightened enforcement effectively walls off dealers from their affiliates.

We document effective enforcement of China Walls in today's foreign exchange market.[3] Dealers intermediate all foreign exchange trades. Each fund is affiliated to a dealer via

[1]"China Walls," or more commonly "Chinese Walls," "information barriers,'' "firewalls," or "ethical screens," was originally a reference to the Great Wall of China (Gozzi, 2003). We adopt "China Walls," because it is concise and the closest to the original reference.

[2]The former Federal Reserve Chairman Paul Volcker was among those advocating breaking them up, testifying that he is "[not] so naive as to think that [...] so-called Chinese walls can remain impermeable against the pressures to seek maximum profit and personal remuneration" (Reuters, 2010).

[3]Private information is widespread in the foreign exchange market. Order flows predict short-run currency returns (Evans and Lyons, 2002; Froot and Ramadorai, 2005). The flows and trade volumes of informed traders, such as hedge funds, are especially predictive (Menkhoff, Sarno, Schmeling, and Schrimpf, 2016; Ranaldo and Somogyi, 2021; Cespa, Gargano, Riddiough, and Sarno, 2022; Hacioglu Hoke, Ostry, Rey, Rousset Planat, Stavrakeva, and Tang, 2026).

common control, connected via a trading relationship, or entirely unrelated to the dealer. China Walls in this setting isolate dealers from their affiliate funds to prevent the sharing of confidential information that the dealers glean from client order flow. Our difference-in-differences design compares the trading activities and profits of funds that are affiliates of, connected to, and unrelated to a dealer around the days when the dealer holds especially valuable private information, which we pinpoint using exceptionally large trades. Heightened trading or profit by the affiliate funds, relative to the unrelated funds, around those days detects violations of China Walls. A failure to detect violations plausibly implies compliance in other times, when the dealer has less valuable information to share. The connected funds and a tightening of China Wall enforcement around July 2018 allow us to potentially falsify our design.

We obtain the near universe of foreign exchange trades involving the Israeli Shekel, covering 25 million trades from 2018 to 2024. The Shekel market is large and liquid, with average daily trades worth USD 13 billion and 83% of them in the US dollar-Shekel currency pair. The largest dealers in the Shekel market are identical to those in the broader US dollar market, and the Israeli financial regulations are mainly based on the US regulations. An exception is that Israel does not impose China Walls, leaving the SEC (which regulates all registered broker-dealers, see Section 2.1) as the main enforcer of China Walls in our setting.

Figure 1 illustrates our design. An event is a trade by a dealer in the top 0.1 percentile by dollar value among this dealer's trades. We compare the daily gross dollar volumes, net volumes, and one-week future profit-and-loss of the dealer's affiliate funds and the funds unrelated to this dealer around the event day. This design detects extensive violations before July 2018, and *none* thereafter.

We can falsify whether our design reliably detects information sharing, using the connected funds. The China Walls are absent between nonaffiliated dealers and funds, and dealers are well known to share information with their clients (Barbon, Di Maggio, Franzoni, and Landier, 2019; Boyarchenko, Lucca, and Veldkamp, 2021). We thus expect dealers

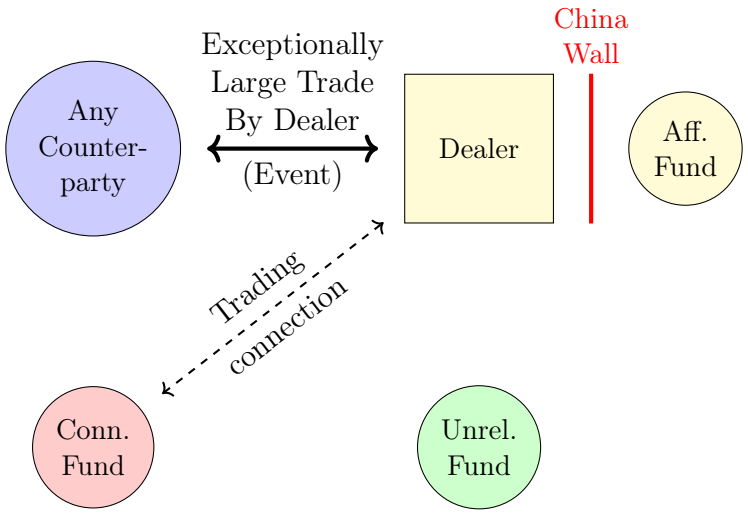

Figure 1: Identifying Information Sharing from Dealers to Affiliate Funds

to share information with their nonaffiliate client funds throughout our sample period. Our falsification test (illustrated in Figure 1), compares the daily volumes and profits of a dealer's nonaffiliate client funds (connected funds) and its unrelated funds around the day when the dealer made an exceptionally large trade. We consistently detect information sharing between connected dealers and funds both before and after July 2018.

Section 3 explains our design. We face the same two challenges as the regulators. First, we must detect information sharing where it exists without falsely detecting it where it does not exist. Second, we need settings in which ruling out violations allows us to rule them out elsewhere. Absent such settings, ruling out violations would require examining every possible setting, an infeasible task.

Using the unrelated funds as our control lets us cleanly identify information sharing. Under the plausible assumption that dealers do not share private information with funds that are neither their affiliates nor clients, our control preserves any increase in volume or profit due to information sharing while removing spurious increases due to common shocks. Our control partials out macroeconomic news, fluctuations in currency demand, and other common shocks because the unrelated funds are exposed to those shocks like other funds. Separately, the event trades may move prices, exhaust liquidity, or prompt additional trading

as the dealer manages resulting inventory risk. Our control partials out those movements in prices or liquidity, because they affect all funds. To eliminate the remaining spurious variation from the event dealer's trades around its event day, we exclude any fund–event pair in which the fund traded with the event dealer on the event day or within the following five days.

Our events are the days when each dealer has the strongest incentive to violate its China Walls. Under the assumption that especially valuable information prompts exceptionally large trades, such days occur around the dealer's exceptionally large trades. Detecting no information sharing around these events rules out sharing in other settings, with weaker incentives to violate China Walls. Our assumption is supported by standard theory (Kyle, 1985; Easley and O'Hara, 1987) and evidence in other markets (Kumar, Mullally, Ray, and Tang, 2020; Pinter, Wang, and Zou, 2024).

Appendix C tests and rejects two threats to this assumption. First, dealers may split orders to disguise private information. We find that the direction of event trades predicts exchange-rate changes for up to five days, whereas smaller trades are not predictive, and we only detect information sharing between connected dealers and funds around the event trades. These findings are inconsistent with order splitting, and consistent with dealers making the event trades to exploit especially valuable information. Second, regulators might focus attention on large trades (DeMarzo, Fishman, and Hagerty, 1998), pushing dealers to hide violations among smaller trades. Affiliate dealers and funds do not share information around any size of the dealer's trades after July 2018, inconsistent with smaller trades hiding violations.

Section 4 describes the data. Dealers frequently trade with their affiliate funds before July 2018, soon after which they virtually *never* do so. We observe no such discontinuity between dealers and their nonaffiliate connected funds. Unrelated dealers and funds exhibit highly correlated trade volumes and profits, verifying the need for a control group to partial out common shocks.

Section 5 implements our design as stacked difference-in-differences specifications with never-treated controls (Cengiz, Dube, Lindner, and Zipperer, 2019) after the SEC escalated enforcement, covering 2019 to 2024. In the 11 trading days around an exceptionally large trade by an event dealer, the daily gross dollar volumes of the funds affiliated to this dealer differ by -0.001 standard deviation on the event day (clustered std. error: 0.003 sd) relative to the unrelated funds. In contrast, the funds connected to the event dealer increase their volumes by 2.0 sd (std. error: 0.004 sd) on the event day relative to the unrelated funds. All results remain when we replace gross volumes with one-week future profits or net volumes signed in the direction of the event trade.

Section 6 looks for heterogeneity in these results. We never detect information sharing between affiliate dealers and funds across crisis and noncrisis periods, asset classes, currencies, and fund types.[4] Connected dealers and funds, by contrast, share information across every characteristic. The patterns in their responses are consistent with the connected funds acting on the information shared by the dealers, rather than inventory or liquidity impacts of the event trades.

Section 7 investigates the quarters around the escalation of enforcement in July 2018. The volumes and profits of affiliate funds sharply increase on the event day relative to the unrelated funds before 2018Q3, with this response dramatically shrinking in 2018Q3, and then falling to zero by 2018Q4. Responses of the connected funds stay nearly identical over the year and into 2019. In sum, we fail to detect information sharing exactly where and when the China Walls are tightly enforced, and consistent detect them elsewhere.

Prior work. Existing evidence on China Walls exploits samples that predate 2018. This evidence identifies extensive breaches, as we replicate.[5] We further evaluate the China

[4]This finding rejects the possibility that affiliate dealers and funds use thirdparties to share information with each other. If that were true, we would expect the affiliate funds to respond more intensely than the unrelated funds for at least some fund, asset, or event characteristic.

[5]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.

Walls during the recent period of tightened enforcement.[6] 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 information sharing has been documented in the literature. Applying this design to a large and granular dataset yields precise estimates and robust evidence that today's China Walls effectively preempt information sharing.

Broader contributions. We belong to the literature on information diffusion in financial markets. Dealers extract information from their clients' order flow (Hortaçsu and Kastl, 2012), leak information to certain clients (Barbon et al., 2019; Boyarchenko et al., 2021; Chague, Giovannetti, and Herskovic, 2023), and generally act as the conduits through which information diffuses throughout their trading networks (Di Maggio, Franzoni, Kermani, and Sommavilla, 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 in banking conglomerates.

Roadmap. Section 2 details the regulatory context and the empirical setting. Section 3 develops the empirical design. Section 4 describes the data and provides preliminary evidence. Section 5 investigates whether the China Walls are effectively enforced in the Israeli Shekel market. Section 6 examines heterogeneity across fund, asset, and event characteristics. Section 7 investigates how enforcement reshaped information sharing. Section 8 concludes.

[6] Garrett (2024) and Beck, Silva-Buston, and Wagner (2025) also document the effectiveness of recent regulations on banking conglomerates. Garrett (2024) finds that a Dodd-Frank ban on concurrent advising and bond underwriting for the same municipality lowers the financing costs of municipalities. Beck et al. (2025) find that supervisory cooperation agreements between national regulators cause banking conglomerates to shift lending to subsidiaries in countries without such agreements. We examine limits on information sharing within banking conglomerates, as opposed to limits on concurrent services or geographic coverage of regulation.

2 Context

Section 2.1 defines the China Walls and describes their enforcement by the SEC, including the Dodd-Frank Act and the Mizuho settlement. Section 2.2 introduces the foreign exchange market in the Israeli Shekel.

2.1 Regulatory Context

China Walls are collections of rules and physical barriers that aim to block the flow of material private information (MPI) between a walled-off subsidiary and its affiliates. MPI is any information that (a) a reasonable investor would find important for her investment decisions and (b) is not publicly disclosed.[7] MPI includes, for example, proprietary analysis and clients' order flows. 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.

The SEC is responsible for the China Walls through its authority over the internal organization of registered broker-dealers (Exchange Act §15(g)). This authority covers any financial firm that trades securities with US persons (§15(a)), including every dealer in our sample. While the Commodity Futures Trading Commission (CFTC) enforces anti-fraud rules involving spot foreign exchange trades, the CFTC does not regulate broker-dealers' internal organization.

Legal precedent. Before the 2010 Dodd-Frank Act, regulators did not enforce China Walls. The Act granted financial regulators risk-based enforcement powers, under which they may penalize conduct that raises the risk of insider trading without having to prove that insider trading occurred. Today, the SEC monitors the China Walls around dealers as a part of its risk-based enforcement against insider trading.

On July 23, 2018, the SEC signaled the tightening of China Wall requirements in its

[7] Material non-public information (MNPI) is more commonly mentioned in law. MNPI excludes analyses purely based on public information, whereas MPI expressly includes such analyses. We focus on MPI, because we are agnostic about the source of private information, whether proprietary analysis or inside information.

settlement with Mizuho Securities, a dealer. The settlement accused Mizuho of failing to maintain and enforce China Walls around a trading desk, without citing any trade that could have possibly occurred due to the information this desk shared (US Securities and Exchange Commission, 2018; Barrack, Moskowitz-Hesse, Richards, and Cox, 2020). From then onward, neither establishing China Wall rules nor the absence of any trade that could have been based on illicitly shared information would be sufficient to avoid penalties. A recent example is the SEC's case against Virtu Financial, whose key database was accessible to both their investment fund and their broker-dealer employees (US Securities and Exchange Commission, 2024). The SEC did not mention any trades in its case against Virtu, building on the precedent of Mizuho case. In December 2025, Virtu agreed to a $2.5 million settlement for its China Wall violations under the new SEC commissioner. Appendix A details regulatory history, other rules on banking conglomerates, and recent cases.

2.2 Empirical Setting

The *foreign exchange market* is an over-the-counter market, in which trades occur between dealers or a dealer and its client. Trades are nonanonymous, and most firms rely exclusively on one or a few relationship dealers. This market operates at high frequency, where news is rapidly incorporated into exchange rates (Menkhoff et al., 2016). We thus expect private advantage from an MPI to dissipate in a few trading days or faster.

Our empirical setting is the foreign exchange market in the Israeli Shekel (ILS), whose data we obtain from the Bank of Israel (BOI).[8] The ILS market structure is identical to other major foreign exchange markets. Some 83% of ILS transactions involve the USD, similar to the USD-share across all foreign exchange transactions (Somogyi, 2022). The ILS and the USD markets have the same largest dealers. Financial regulations in Israel and in the US are similar, except that Israel forbids its banks from controlling investment funds (Shemesh and

[8]The BOI can enforce its worldwide reporting requirements (detailed in Footnote 9) through Israeli banks, which are firmly under BOI jurisdiction. Israeli banks must verify that their counterparties are not in breach of BOI reporting requirements. Any nonIsraeli firm trading ILS have a strong incentive to abide by those requirements, because Israeli banks maintain the largest holdings of ILS.

Abir, 2024, p. 191), as the US Glass-Steagall Act did until 1999. No Israeli fund is affiliated with a dealer, because banks own all Israeli dealers. The Israeli regulators do not enforce China Walls.

Two implications follow. (a) Israeli dealers and funds cannot drive our results on the information sharing between affiliate dealers and funds. (b) Foreign banking conglomerates do not incriminate themselves when reporting data at odds with their China Walls to the BOI. Hence, ours is the dataset in which we would most expect to observe any violations of the China Walls.

3 Design

Section 3.1 describes and explains our empirical design. Section 3.2 provides a test of our central identification assumption.

3.1 Empirical Design

We overcome three challenges to test the hypothesis that the China Walls are effectively enforced. First, it is infeasible to evaluate enforcement in all circumstances. Second, any proxy for bilateral sharing of MPI can be contaminated by public news and other common shocks. Third, our design may fail to reliably detect bilateral MPI sharing, in which case we cannot establish that China Walls are effectively enforced.

Defining events. We seek events that pinpoint the days when a dealer receives especially valuable MPI, under the assumption that China Wall violations are most likely to occur when gains from sharing information with affiliates are largest. We let an event be a dealer and a day (event day) when the event dealer makes a trade (event trade) that is exceptionally large compared to the dealer's other trades. Precisely, an event trade is a trade whose dollar value is in the top 0.1 percentile among all trades by the event dealer. Appendix E examines events by funds and shows qualitatively identical results.

Standard theory shows that an informed trader requests to trade larger quantities when she holds more valuable private information (Kyle, 1985; Easley and O'Hara, 1987). Empirically, unusually large trades by the same trader are particularly predictive of returns (Kumar et al., 2020; Pinter et al., 2024). Section 3.2 develops a test of our assumption. Appendix C implements the test and concludes that the assumption holds in our setting.

Isolating information sharing. We consider a dealer d and one of its affiliate funds f in the days surrounding an exceptionally large trade by d. A proxy for MPI sharing between d and f must isolate information that is (i) material and (ii) shared bilaterally, rather than indirectly through the market. Receiving material information would prompt f to rebalance its portfolio and earn greater trading profits. More specifically, f would trade more in the direction predicted by the MPI to earn positive returns, which would be the direction of the event trade if the same MPI prompted d to make the exceptionally large trade. Taken together, we look for increases in the daily gross dollar volume, net dollar volume in the direction of the event trade, and one-week future profit and loss of f around an event trade by d to proxy for MPI sharing between d and f.

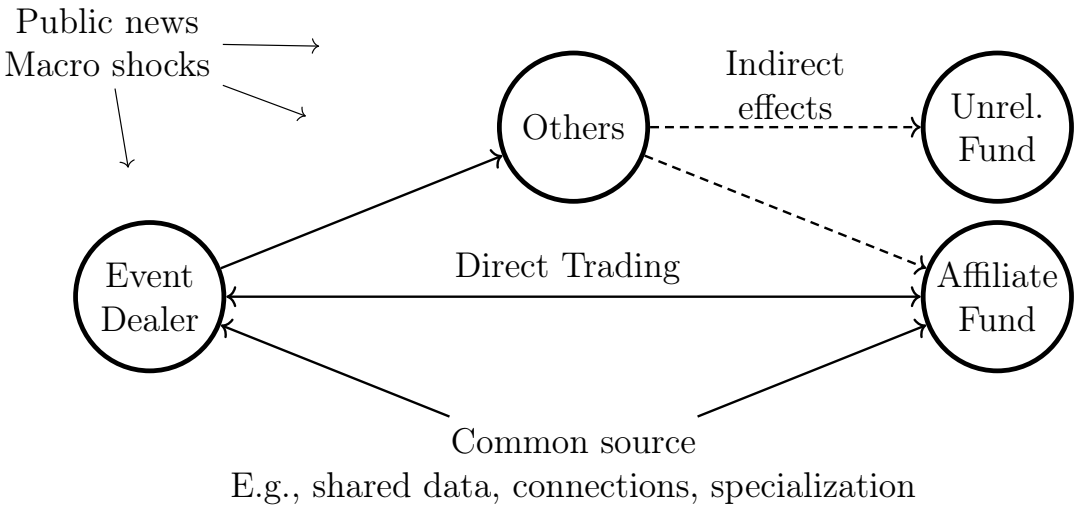

Figure 2: Potential Confounders to Measuring Bilateral Information Sharing

We remove the four confounders that could simultaneously induce an exceptionally large trade by dealer d and raise the trade volumes or profit of its affiliate fund f. Figure 2

illustrates the confounders. First, any direct trade between d and f would cause a mechanical correlation in their gross trade volumes. We avoid mechanical correlation by excluding any f from our analyses whenever it traded with d on or in the five trading days after the event day.

Second, common shocks, such as public news, can simultaneously affect the trade sizes of d and the trade volumes and profits of f. Third, the event trade may have price impact or deplete marketwide liquidity, and thus indirectly affect the volumes or profits of f. For example, the event trade could be a trade between d and another dealer, which in turn offloads the newly gained inventory to f, raising the fund's volume.

We partial out the common shocks and indirect effects by comparing the affiliate fund f to the *unrelated funds*, which are neither affiliated with nor ever trades with the event dealer d. A dealer would never share private information with an unrelated fund. At the same time, such funds as affected as f by the common shocks and the indirect effects of the event trade. Therefore, comparing f to the unrelated funds preserves all variation in volumes and profits due to MPI sharing between d and f, and removes the variation due to the common shocks and indirect effects. Section 4.5 verifies the presence of these confounders.

Fourth, shocks specific to a banking conglomerate could trigger the event trade by d while increasing the volumes and profits of f. Conglomerate-specific shocks could arise if, say, d and f share more data sources or more thirdparty dealers than a random dealer-fund pair. If this confounder were important, we would see an increase in the volumes and profits of f relative to the unrelated funds around the event trade by d. Our precise null results after July 2018 (Section 5) indicate that this confounder is unimportant in practice.

Falsification test. We would falsely conclude that the China Walls are effectively enforced if our design fails to reliably detect MPI sharing although it exists. We test our design's reliability by exploiting the lack of China Walls between nonaffiliated dealers and funds and the stylized fact that dealers extensively share information with their clients (Barbon et al., 2019; Kumar et al., 2020; Chague et al., 2023). Combining those facts, we

expect unaffiliated dealers and funds connected via trading relationships to consistently share MPI with each other. Therefore, if our design were reliable, we would consistently detect MPI sharing between those *connected* dealers and funds. We falsify our design whenever we fail to detect an increase in the volumes or profits of the connected funds relative to the unrelated funds around the event day. We exclude any connected fund that trades with the event dealer on or in the five days after the event day to avoid mechanical correlations.

3.2 Identification Tests

We assume that a dealer makes an exceptionally large trade when the dealer learns especially valuable private information. Appendix C adjudicates the stronger assumption that dealers trade larger sizes when they have more valuable private information. Specifically, we test *(I)* this stronger assumption, *(II)* the claim that our design detects MPI sharing if and only if such sharing exists, and *(III)* the threat that the affiliate dealers and funds hide their MPI sharing among their smaller trades.

We define placebo events as a dealer and a day when the dealer makes a trade in the X to $X + 0.1$ percentile of the dealer's trades, where X is each decile, $X \in \{10, 20, \ldots, 90, 99.9\}$. Our test consists of two parts. First, we separately compute the price impacts of the exceptionally large trades (top 0.1 percentile) and the placebo event trades for each decile (down to bottom 99.9 percentile). We verify (I) if the dealers' exceptionally large trades predict price movements whereas their smaller trades do not. Second, we compare the volumes and profits of the connected and the unrelated funds around each decile of trades by dealers. We verify (II) if we detect information sharing between connected dealers and funds only around the exceptionally large trades, and none around the smaller trades.

The last threat is the possibility that the SEC focuses its enforcement on the largest trades (DeMarzo et al., 1998), (III) pushing China Wall violation towards the dealers' smaller trades. We compare the volumes and profits of the affiliate and the unrelated funds around

each decile of trades by dealers. We reject (III) if we fail to detect information sharing between affiliate dealers and funds around any decile of event trades.

4 Data

Section 4.1 describes the raw data and trading volumes. Section 4.2 determines the affiliations, connections, domiciles, and fund types. Section 4.3 constructs the analytical sample. Section 4.4 defines the dependent variables. Section 4.5 documents spurious correlations in trading volumes between unrelated dealers and funds, verifying the need to include the unrelated funds as controls.

4.1 Raw Data

We obtain the near universe of foreign exchange trades involving the Israeli Shekel from the BOI for the sample period spanning January 2018 and March 2024.[9] Each trade specifies the currency pair (ILS and another currency), contemporaneous exchange rate, quantity, date and time, asset class (spot, forward, swap, or option), the legal names of counterparties, their trade directions, whether each counterparty is a financial firm (i.e., dealers and investment funds), and whether it is domiciled in Israel. We do not use the time stamps, as the BOI does not verify them and an improbably high 6.6% of trades report 00:00:00.

Table 1 describes financial firms in the raw data. *Dealers* are financial firms on the BOI's list of foreign exchange dealers. *Funds* are all other financial firms. *Conglomerates* are holding companies that each controls two or more financial firms. The plurality of dealers and most funds are in the US with the others mainly in Europe or the UK. There

[9] All Israeli banks, including the Israeli branches of banking conglomerates, must report all their ILS foreign exchange trades to the BOI. Non-Israeli banks fall under the same reporting requirement if their foreign exchange trades (in any currency pair) in the previous year exceed $15 million per day on average. Practically all dealers fall under the reporting requirement. Exceptions are the three dealers that are not affiliated with a bank whose trades we do not observe. (We do observe their trades with the other dealers, which must report to the BOI.) 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/.

Table 1: Unique Financial Firms in the Raw Data

	Conglomerates	Dealers		Funds		
		All	No affiliate fund	All	No affiliate dealer	Hedge fund
US	13	94	0	4824	4307	364
Europe & UK	13	101	0	1380	968	88
Israel	11	18	18	192	192	2
Elsewhere	9	22	0	1379	1193	178
Total	46	235	18	7775	6660	632

Dealers are financial firms on the BOI's list of foreign exchange dealers. *Funds* are all other financial firms. *Conglomerates* are holding companies that each controls two or more financial firms.

are 46 conglomerates, each of which controls one or more dealers and altogether control all dealers. Most conglomerates control several funds, though most funds are independent. All conglomerates without a fund are Israeli.

Figure 3 plots the total daily gross dollar volume of trades among the dealers and funds. Regular falloffs in volume correspond to end-of-year holidays. Nonaffiliate dealers and funds trade USD2.8 billion each day with each other on average (Figure 3a), and this volume is gradually increasing throughout the sample period. Affiliate dealers and funds made regular trades before the announcement of the Mizuho settlement on July 23, 2018. They cease trading soon thereafter, making a mere four trades worth USD5.51 million over 2019 to 2024. Trading between affiliate dealers and funds would require communication, which effective China Walls would prevent. (They could share information without trading, which we look for.) We can attribute the delay in ceasing of trades to the time needed to interpret the settlement and update dealer workflows. The first relevant legal analysis we found is dated July 31th (Morrison Mahoney LLP, 2018), and the first by a national law firm is dated August 7th (Patterson Belknap Webb & Tyler LLP, 2018).

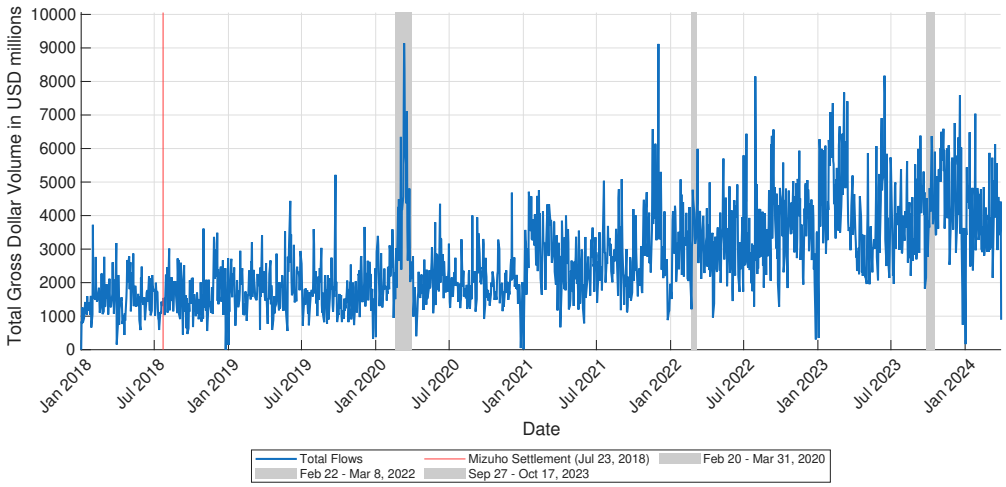

(a) Unaffiliated Dealers and Funds

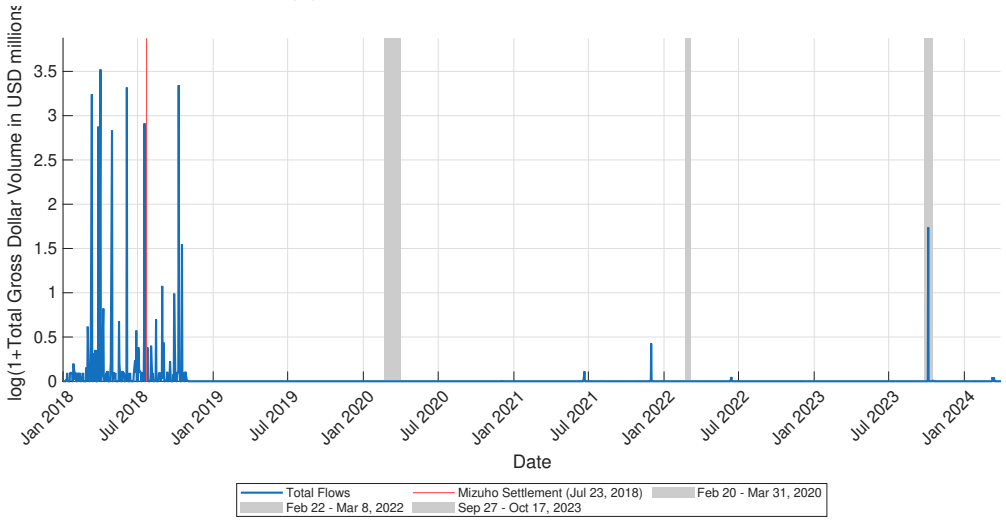

(b) Affiliate Dealers and Funds

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.

4.2 Affiliations, Connections, Domiciles, and Fund Types

A *connected* dealer-fund pair are unaffiliated and trade at least ten times with each other in our sample. An *unrelated* dealer-fund pair are unaffiliated and never trade with each other in our sample.

We determine affiliations in four steps. First, the remaining dealers and funds with obviously indicative names (e.g., "`[major bank]` Luxembourg S.A.") are assigned to the

corresponding conglomerate. Second, we obtain the affiliations of most US-based dealers and funds from the National Information Center database (https://www.ffiec.gov/npw/). We assign those affiliations as of the fourth quarter in 2023, because dealers and funds rarely change their affiliations and typically change their legal names when they do. Third, the remaining dealer and fund legal names are entered into `ChatGPT 4.0o` as prompts in the form, "as of `[date legal name last appears in the sample]`, does `[legal name]` belong to a financial conglomerate or is it independent? Which holding company if any does `[legal name]` belong to?" Fourth, we verify each answer generated in step three, by searching for the legal name paired with "independent" or the `ChatGPT`-suggested holding company name.

We dealer and fund domiciles and whether each fund is a hedge fund by querying and verifying, "as of `[date legal name last appears in the sample]`, which country is the legal domicile of `[legal name]`?" and "as of `[date legal name last appears in the sample]`, is `[legal name]` a hedge fund?"

4.3 Sample Construction

We construct our analytical sample in three steps. First, we drop all options trades (for insufficient observations) and all second legs of swaps trades (to avoid double counting).[10] Second, we consolidate the dealers up to the conglomerate-level by dropping all trades between affiliated dealers and combining the dealers under conglomerate-level labels. Doing so treats affiliated dealers as a single economic entity, given their ability to split incoming or-

[10]Foreign exchange swap is a spot trade with the commitment to reverse the spot trade at a predetermined exchange rate on a future date. The first leg is the initial spot trade and the second leg is the reverse trade. For example, consider a swap trade today to buy one USD using ILS at 3.03 USD/ILS spot rate (first leg; USD/ILS means ILS per USD by convention) then promise to sell one USD for ILS at 2.86 USD/ILS in five days (second leg). This trade would yield a profit if ILS appreciates below 2.86 USD/ILS in five days. To see this, the trader can immediately sell the USD from the first leg for 3.03 ILS; then buy back one USD just before the second leg using, say, 2.78 ILS; then deliver the USD for 2.86 ILS. Overall profit is $-3.03 + 3.03 - 2.78 + 2.86 = 0.08$ ILS.

ders and transfer assets and capital among themselves for tax and balance sheet purposes.[11] Third, we aggregate the trades of each fund up to the daily frequency.

Table 2: Each Step in Sample Construction

	Obs.	Unique dealers	Unique funds	Mean value (USD millions)		Share of trades (%)		
				Dealer obs.	Fund obs.	USD	Forward	Swap
Raw data	25,525,488	235	7,775	3.69	1.73	82.8	15.1	14.7
Drop								
options	25,315,204	235	7,775	3.70	1.73	82.8	15.2	14.8
2nd legs	23,372,145	235	7,775	2.45	1.32	82.7	16.5	7.7
Consolidate								
aff. dealers	20,733,381	46	7,775	2.12	1.32	82.6	17.1	7.8
Aggregate to								
fund-by-day	12,688,800	46	7,775	—	0.32	93.1	21.5	2.8
dealer-by-day	75,072	46	7,775	584.17	—	82.6	17.1	7.8

Mean value is in USD millions and is a simple average across observations. *Share of trades* is the percentage of observations with the corresponding characteristic (i.e., is in USD-ILS currency pair, a forward trade, or a swap trade). *Raw data* includes all foreign exchange trades reported to the Bank of Israel between January 2018 and March 2024. We *drop* all options trades and all second legs of swap trades (see Footnote 10). *Consolidate aff. dealers* involves combining the dealers affiliated to each other up to the conglomerate-level, then dropping all trades between them. *Fund-by-day* is the analytical sample, in which an observation is an aggregate across all trades and asset classes by a fund on a given trading day in USD terms. *Dealer-by-day* is the sample in which an observation is the aggregate up to a dealer and a day, used for the analysis in Appendix E.

Table 2 summarizes each of the three steps. The final sample for analysis comprise of 12 million fund-by-day observations. The funds trade more USD and less swaps than the dealers. Appendix E examines the dealer-by-day sample and finds qualitatively identical results.

Table 3 describes the network structure in the analytical sample. Most dealers have at least one affiliate fund and have traded ten or more times with (are connected to) several funds. The overwhelming majority of funds are not affiliated with any dealer and are connected to no more than a single dealer. A quarter of dealers neither ever trade nor are

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

Table 3: Summary Statistics of the Trading Network

	Mean	Std. Dev.	Min	25%	50%	75%	Max
Dealers (N = 46)							
Affiliate funds	24.35	39.81	0.00	1.00	10.50	30.00	182.00
Connected funds	181.89	329.24	0.00	0.00	8.00	178.00	1248.00
Unrelated funds	7251.24	869.94	4747.00	7040.00	7756.50	7775.00	7775.00
Funds (N = 7775)							
Affiliate dealers	0.14	0.35	0.00	0.00	0.00	0.00	1.00
Connected dealers	1.08	1.82	0.00	0.00	0.00	1.00	16.00
Unrelated dealers	42.59	3.15	28.00	41.00	44.00	45.00	45.00

Number of affiliate, connected, and unrelated counterparties per dealer and per fund. Each dealer-fund pair is exclusively affiliate, connected, unrelated, or none (if the pair trades 1 to 9 times with each other in our sample).

affiliated with (are unrelated to) any fund, because they either only trade with nonfinancial firms or are one of the three nonbank dealers that are not required to report their trades to the BOI. (We do observe the nonbank dealers' trades with the other dealers.) Dropping the nonbank dealers does not affect our results. Each dealer-fund pair is exclusively affiliate, connected, unrelated, or none (if the pair trades 1 to 9 times with each other in our sample).

4.4 Events and Dependent Variables

We define *events* in three steps. First, we convert all dealer trades into USD using the contemporaneous exchange rates. Second, an *event trade* is a trade that belongs in the top 0.1 percentile in USD terms among all trades made by a given dealer. Third, an event is a dealer and a day on which the dealer made one or more event trades. Each event takes on the characteristics of its event trade. If an event includes multiple event trades, the event takes on the characteristics of the largest trade among them.

Figure 4 plots the distribution of events over time. No specific period is likely to be driving our results. The event frequency is about constant over time, even as total volumes have steadily increased, because events are defined at the dealer level and each dealer's distribution of trade sizes have stayed constant.

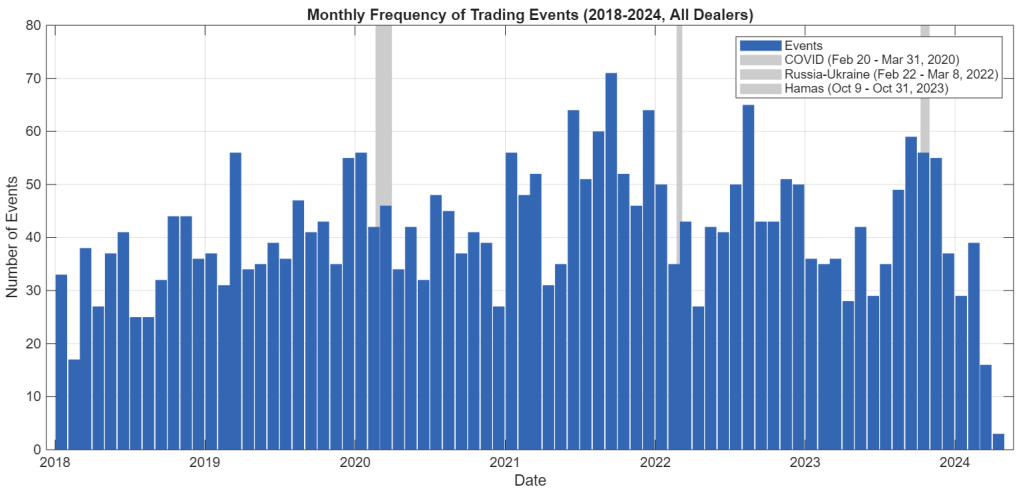

Figure 4: Distribution of Events over Time

Gross dollar volume of fund f on day t aggregates the fund's daily spot, forward, and swap trades. For each (f, t)-pair, we add together the notional amounts of the spot, forward, swap (first legs only) trades in USD terms. We convert all trade values to USD at the contemporaneous exchange rate.

Net dollar volume in the direction of the event trade by fund f on day t in event e is computed in two steps. For each (f, t, e, asset class, currency pair)-quintuplet, we compute a separate net notional amount (i.e., buys less sales of the non-ILS currency for spot and forward and the opposite for the first legs of swaps; see Footnote 10) for each non-ILS currency (CCY) in USD terms. We multiply the net notional amount of a given CCY by -1 if the event dealer sold the CCY in the event trade. The net dollar volume is the sum of the net notional amounts across CCY by f on t in e.

One-week future profit-and-loss is the dollar net profit from trades by fund f on day t if those positions were held for five trading days. For each (f, t, asset class, currency pair)-quadruplet, we compute the net notional amount for each CCY in CCY terms. The five-trading day realized return for a given CCY is the percent change in the exchange rate, $R_t^{sf}(CCY) := [(CCY/ILS)_{t+5}/(CCY/ILS)_t - 1]$, using the rates at 17:00 EST from

Bloomberg.[12] We multiply each CCY net notional amount by the corresponding realized return, then convert the resulting amounts into USD using the 17:00 EST exchange rate on $t+5$ to obtain the net P&L in the CCY by f on t. The one-week future P&L is the sum of the converted net P&L across CCY by f on t.

Table 4: Summary Statistics of the Analytical Sample

	Mean	Std. Dev.	Min	25%	50%	75%	Max
Events ($N = 3096$)							
Event trade value	253.00	143.56	1.53	199.40	233.61	296.73	1750.00
Has multiple event trades	0.11	0.31	0.00	0.00	0.00	0.00	1.00
Crisis	0.06	0.23	0.00	0.00	0.00	0.00	1.00
USD	0.89	0.31	0.00	1.00	1.00	1.00	1.00
JPY	0.03	0.17	0.00	0.00	0.00	0.00	1.00
EUR	0.04	0.20	0.00	0.00	0.00	0.00	1.00
Spot	0.18	0.38	0.00	0.00	0.00	0.00	1.00
Forward	0.13	0.33	0.00	0.00	0.00	0.00	1.00
Swap	0.70	0.46	0.00	0.00	1.00	1.00	1.00
Fund-by-Day ($N = 12,688,800$)							
Has trade	0.04	0.21	0.00	0.00	0.00	0.00	1.00
Gross volume	0.36	8.36	0.00	0.00	0.00	0.00	4210.10
Net volume	0.00	2.66	-1201.00	0.00	0.00	0.00	1046.75
1-week future P&L	0.00	0.07	-32.25	0.00	0.00	0.00	35.65

All volumes, P&L, and values are in USD millions. Each event takes on the characteristics of its event trade. If an event includes multiple event trades, the event takes on the characteristics of the largest trade among them.

Table 4 summarizes the events and the dependent variables in the analytical sample. Event trade values, volumes, and P&L are in USD millions. Event trades widely vary in size, and are mostly swap trade and for the USD-ILS pair. Only a small proportion occur during crisis periods, including the Covid pandemic (February 1 to March 31, 2020), the 2022 Russian invasion of Ukraine (February 16 to March 8, 2022), and the 2023 Hamas attack (September 27 to October 17, 2023). Some 95% of fund-by-day observations do not have a trade, because most funds do not trade on a given day. Funds on average do not earn a trading profit nor accumulate a net position in ILS, as one would expect.

[12] Although the foreign exchange (FX) market is open around the clock, 17:00 EST is the beginning of the sole interval in a trading day when all major FX markets are closed. It also avoids the various FX fixes throughout the day, around which the USD predictably appreciates then reverses (Krohn, Mueller, and Whelan, 2024).

4.5 Correlations within Unrelated Dealer-Fund Pairs

We examine correlations in trading activity within unrelated dealer-fund pairs to verify whether a control group is necessary in our empirical design. An unrelated pair is a dealer and a fund that are unaffiliated and never trade with each other.

For each unrelated pair of dealer d and fund f and each lag $l = -10 \ldots +10$, we compute the correlation between the day-t gross volume of d and the day-$(t+l)$ gross volume of f. We present the residuals from regressing these pairwise correlations on day-of-week, dealer, and fund fixed effects, with robust standard errors.

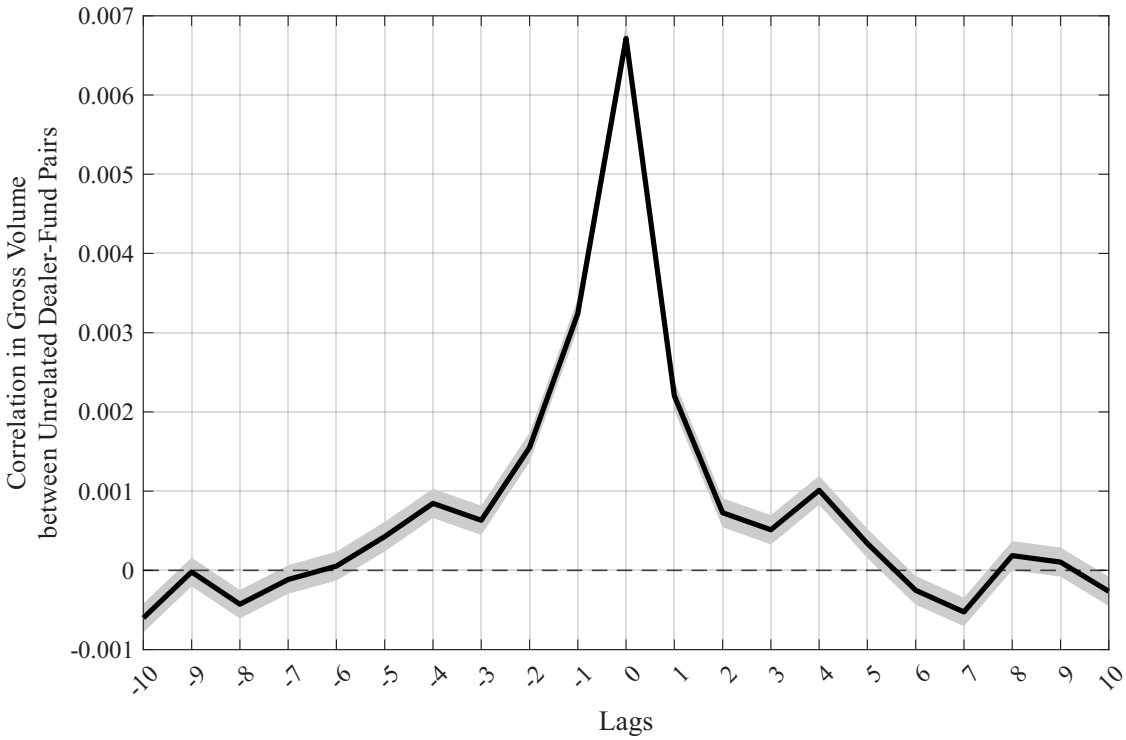

Figure 5: Correlations in Daily Gross Volumes within Unrelated Dealer–Fund Pairs

Residuals from regressing raw within-unrelated-pair correlations on day-of-week, dealer, and fund fixed effects. Shaded region shows the 95% confidence intervals based on robust standard errors.

Figure 5 plots the results. There are positive and significant correlations in the daily gross volumes between dealers and their unrelated funds. The correlations peak on the same day (lag 0), likely reflecting common exposure to macroeconomic news and exchange rate

movements. Absent a control group, these confounding comovements in trading activities would contaminate any proxy for bilateral information sharing.

5 Are the China Walls Effectively Enforced?

Section 5.1 specifies the stacked difference-in-differences regression that implements our design. Section 5.2 presents the results for 2019–2024, today's period of tightened China Wall enforcement by the SEC.

5.1 Implementation

We adopt the stacked difference-in-differences specification with never-treated controls of Cengiz et al. (2019).[13] Each event creates a stack, a subsample in which an observation is a treated or control fund on a given day from the analytical sample in the 11-trading day window (event window) around the event day. We append the stacks across all events to obtain the sample for regression analysis.

For all regressions, we standardize each dependent variable of fund f by dividing its value by its standard deviation computed over the observations of f. We winsorize these variables at the top 0.5 percent for gross volume, or at both the top and the bottom 0.5 percent for the other variables, after stacking observations and before estimating each regression. All results remain without winsorization.

Our main regression specification is

$$Y_{e(d)ft} = \sum_{\tau=-5}^{5} \alpha_\tau \mathbb{1}_{t=\ell_{e(d)}+\tau} Affiliate_{e(d)f} + \delta_{e(d)f} + \varphi_t + \sum_{\tau=-5}^{5} \gamma_\tau \mathbb{1}_{t=\ell_{e(d)}+\tau} + \varepsilon_{e(d)ft}. \quad (1)$$

The dependent variable $Y_{e(d)ft}$ is the gross dollar volume, the net dollar volume in the direction of the event trade, or the one-week future P&L of fund f (all defined in Section 4.4) for

[13] 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).

event $e(d)$ by event dealer d on calendar date t. The affiliate treatment dummy $Affiliate_{e(d)f}$ equals 1 if f is affiliated with d. The dummy $Affiliate_{e(d)f} = 0$ if f is *(i)* not affiliated with and never trades with d and *(ii)* not affiliated with any event dealer of other events within the 21 trading days around the event day $\ell_{e(d)}$ ($t = \ell_{e(d)} - 10, \ldots, \ell_{e(d)} + 10$). We control for event-by-fund, calendar date, and event date fixed effects, $\delta_{e(d)f}$, φ_t, and γ_τ. These effects embed all possible event-and-fund-specific controls as well as common trends over calendar and event times. We cluster standard errors by event-and-fund and by calendar date, because our treatments are assigned event-by-fund and the incidence of events varies over time. Our clustered variances likely approximates the true variances, since we observe the near universe of trades in the Israeli Shekel, implying a high sampling probability (Abadie, Athey, Imbens, and Wooldridge, 2023).

The second specification considers the connected funds:

$$Y_{e(d)ft} = \sum_{\tau=-5}^{5} \beta_\tau \mathbb{1}_{t=\ell_{e(d)}+\tau} Connected_{e(d)f} + \delta_{e(d)f} + \varphi_t + \sum_{\tau=-5}^{5} \gamma_\tau \mathbb{1}_{t=\ell_{e(d)}+\tau} + \varepsilon_{e(d)ft}. \tag{2}$$

The connected treatment dummy $Connected_{e(d)f}$ equals 1 if fund f *(a)* is not affiliated with event dealer d, *(b)* trades 10 or more times with d in the full 2018–2024 analytical sample, and *(c)* does not trade with d on or in the five days after the event day ($t = \ell_{e(d)}, \ldots, \ell_{e(d)}+5$). Condition (a) keeps the connected and the affiliate treatments mutually exclusive. Condition (b) fixes connections constant across the full sample period, because dealer-to-client trading relationships are highly persistent (Hendershott, Li, Livdan, and Schürhoff, 2020). Condition (c) excludes funds whose trades with the event dealer would add mechanical variations in volumes or profits around the event day. The dummy $Connected_{e(d)f} = 0$ if and only if $Affiliate_{e(d)f} = 0$, identifying the same control group consisting of the unrelated funds. All remaining aspects of equation (2) are identical to those in equation (1).

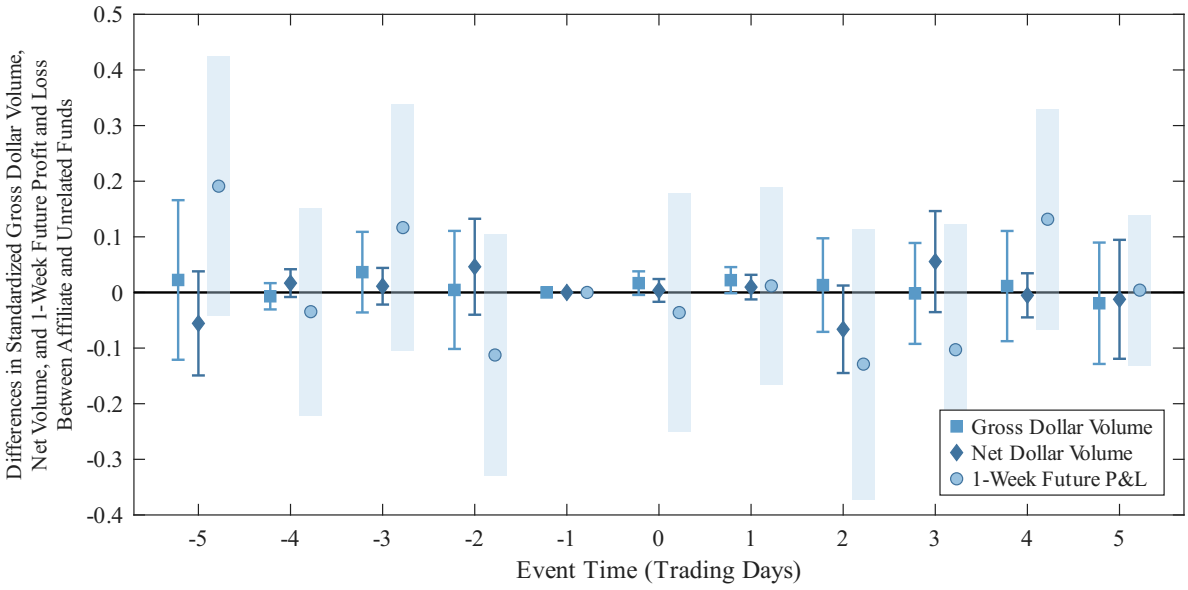

(a) Affiliate and Unrelated Funds: Differences in Standardized Volumes and P&L

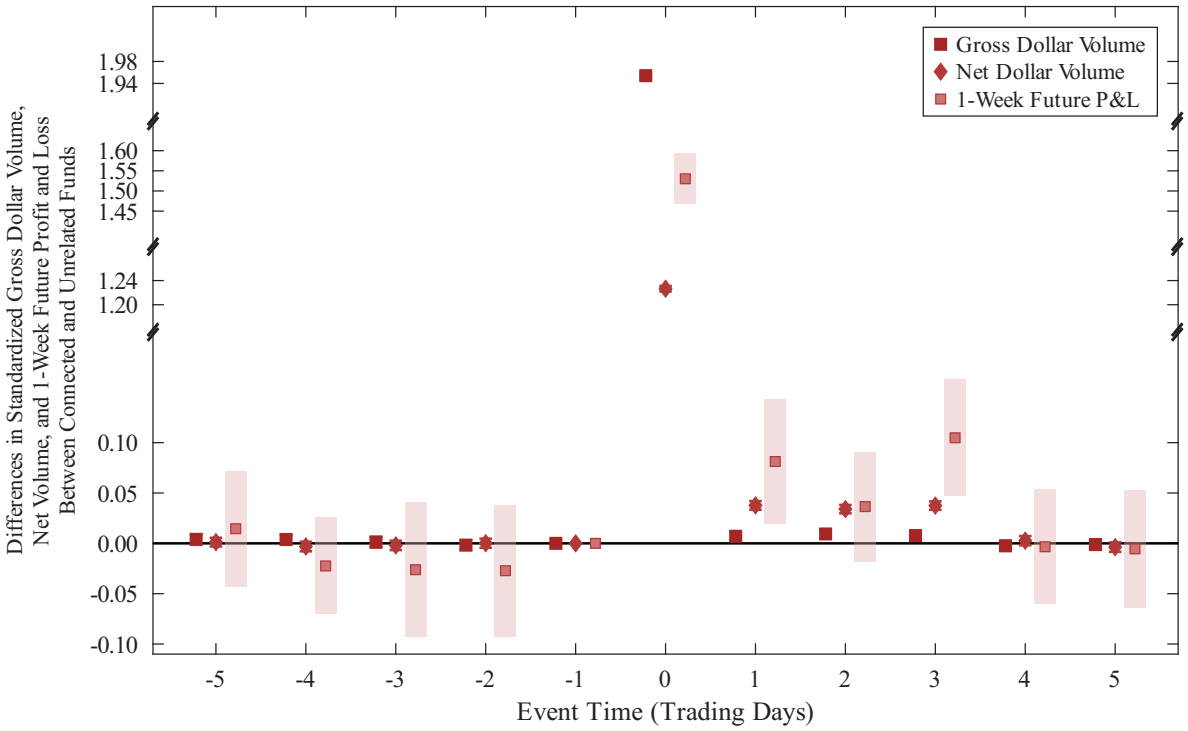

(b) Connected and Unrelated Funds: Differences in Standardized Volumes and P&L

Figure 6: Coefficient Estimates from (1) and (2) in the 2019–2014 Sample

5.2 Results

Figure 6a plots the differences α_τ (in equation (1)) in standardized gross volume, net volume, and one-week future profit and loss between affiliate and unrelated funds around the days of exceptionally large trades by dealers in the 2019–2024 period. Figure 6b plots the corresponding differences β_τ (in equation (2)) between the connected and the unrelated funds. Each panel displays gross dollar volume (capped confidence intervals), net dollar volume (uncapped), and one-week future P&L (translucent bars). The affiliate funds are indistinguishable from the unrelated funds across all outcomes over the entire event window. The connected funds show zero pretrend, then a large positive response in every outcomes on the event day and a weaker positive response for up to three days thereafter.[14]

More precisely, on a day when a dealer makes an exceptionally large trade, its affiliate funds increase their gross volumes by −0.001 sd (std. error: 0.003 sd) relative to the unrelated funds, and this response remains nil on the following days. The connected funds increase theirs by 2.0 sd (std. error: 0.004 sd) on the event day. Likewise, the affiliate funds exhibit precisely nil increase in their net volumes and one-week future P&L on the event day and onward, whereas the connected funds increases their net volumes by 1.0 sd (std. error: 0.003 sd) and P&L by 0.6 sd (std. error: 0.02 sd), both of which stay elevated for up to three days afterwards.

Taken together, the affiliate funds are indistinguishable from the unrelated funds around each dealer's exceptionally large trade, even as the connected funds trade dramatically more, in the same direction as the event trade, and earn larger profits on those trades. We conclude that our design reliably detects information sharing, and that the China Walls effectively preempt information sharing between affiliate dealers and funds. Appendix E examines the responses of dealers around exceptionally large trades by funds, and finds qualitatively identical results.

[14] That dealers share information with their connected funds does not necessarily indicate illegal activity. For example, dealers can legally share proprietary analysis of public information with their connected funds.

Table 5: Responses by Funds on and after Event Days

	Affiliate Gross Volume	Affiliate Net Volume	Affiliate P&L	Connected Gross Volume	Connected Net Volume	Connected P&L
$Post \times Affiliate$	-0.00	-0.01	0.01			
	[0.032]	[0.021]	[0.021]			
$Post \times Connected$				0.33***	0.22***	0.30***
				[0.006]	[0.004]	[0.013]
Event×Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Calendar Date FE	Yes	Yes	Yes	Yes	Yes	Yes
Days-since-Event FE	Yes	Yes	Yes	Yes	Yes	Yes
Adjusted R-squared	-0.001	-0.001	0.061	0.079	0.008	0.046
Within R-squared	0.0000	0.0000	0.0000	0.0072	0.0011	0.0028
Events	1,284	1,284	1,284	2,637	2,637	2,637
Observations	184,356,710	184,356,710	184,014,987	94,802,495	94,802,495	94,450,117

Coefficient estimates from the pooled counterparts to equations (1) and (2). *Gross Volume*: Dependent variable is a fund's standardized daily gross US dollar volume winsorized at the top 0.5 percentile. *Net Volume*: Dependent variable is a fund's standardized daily net US dollar volume in the direction of the event trade winsorized at the top and bottom 0.5 percentile. *P&L*: Dependent variable is a fund's standardized daily one-week future P&L winsorized at the top and bottom 0.5 percentiles. An event is a dealer and a day when the dealer made a trade in the 0.1 percentile among its trades. Each event window is the 11 trading days around the event day. *Affiliate*: Fund is controlled by the same conglomerate as the event dealer. *Connected*: Fund trades at least 10 times with the event dealer in 2019–2024, and does not trade with the event dealer on or after the event day. Control funds are unaffiliated and never trades with the event dealer in 2019–2024, and are not treated in another event during the event window. We include event-by-firm, calendar date, and days-since-event fixed effects. Standard errors in square brackets are clustered at the event-by-firm and date levels. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 5 details the pooled regression counterpart to Figure 6. It crystallizes our power to detect even tiny treatment effects. Each stacked sample includes many million observations corresponding to hundreds of events and numerous affiliate, connected, and unrelated funds per event. The affiliate-treatment regressions contain fewer events than the connected regressions, because the Israeli and the nonbank dealers have no affiliate funds, and thus their events are dropped from the affiliate regressions. The connected regressions have fewer observations than the affiliate ones, because dealers have far more connected funds than affiliate funds (Table 4), leading to many more unrelated funds being dropped to keep the control group never-treated.

6 Heterogeneity

We examine whether the China Walls hold uniformly across fund, asset, and event characteristics in the 2019–2024 sample. We find that affiliate dealers and funds never share information, regardless of the characteristic. Connected dealers and funds share information across all characteristics, with patterns consistent with the existing literature on foreign exchange trading.

Section 6.1 details our procedure. Section 6.2 presents and discusses the results.

6.1 Implementation

We examine cells of fund types, currency pair, and asset class (i.e., spot, forward, or swap), and events during crisis and noncrisis periods. Rather than split our sample, we add interaction terms to the pooled counterparts of equations (1) and (2) and estimate over the full 2019–2024 sample, maximizing our power to detect China Wall violations. We define the fund and event characteristics as follows.

Fund characteristics. The dummy $Small_f = 1$ if and only if the total gross dollar volume of the treated or control fund f is smaller than the median across all funds. A fund is "Large" if $Small_f = 0$. Other fund-specific dummies indicate whether a fund is a hedge fund or whether the share of its trades in a currency pair or asset class is greater than the median across funds.

Event-trade characteristics. The dummy $SmallEventTrade_e = 1$ if and only if the dollar value of the event trade is smaller than the median across all event trades. The event is "Large" if $SmallEventTrade_e = 0$. The dummy $Crisis_e = 1$ if and only if the event trade occurred during the initial Covid panic, the 2022 Russian invasion of Ukraine, or the 2023 Hamas attack. Other event-specific dummies indicate whether the event trade (or the largest one if the event has multiple) was in a given currency pair or asset class.

Estimation. We illustrate the estimation procedure using currency-pair characteristics.

First, we define a dummy $USDFund_f = 1$ if and only if USD-ILS trades as the share of all trades by fund f exceeds the median across all funds, and a dummy $USDEvent_e = 1$ if and only if the event trade is in USD-ILS. Analogous dummies are created for all other currency pairs. Second, we add these currency dummies and all possible interactions among them and the affiliate treatment dummy into the pooled counterpart to equation (1). Third, to compute the point estimate for the affiliate treatment when both the fund's specialization and the event trade are in USD-ILS, we sum over the estimated coefficients of $Affiliate_{ef}$, $Affiliate_{ef} \times USDFund_f$, $Affiliate_{ef} \times USDEvent_e$, and $Affiliate_{ef} \times USDFund_f \times USDEvent_e$ terms. The delta method obtains the standard errors.

6.2 Results

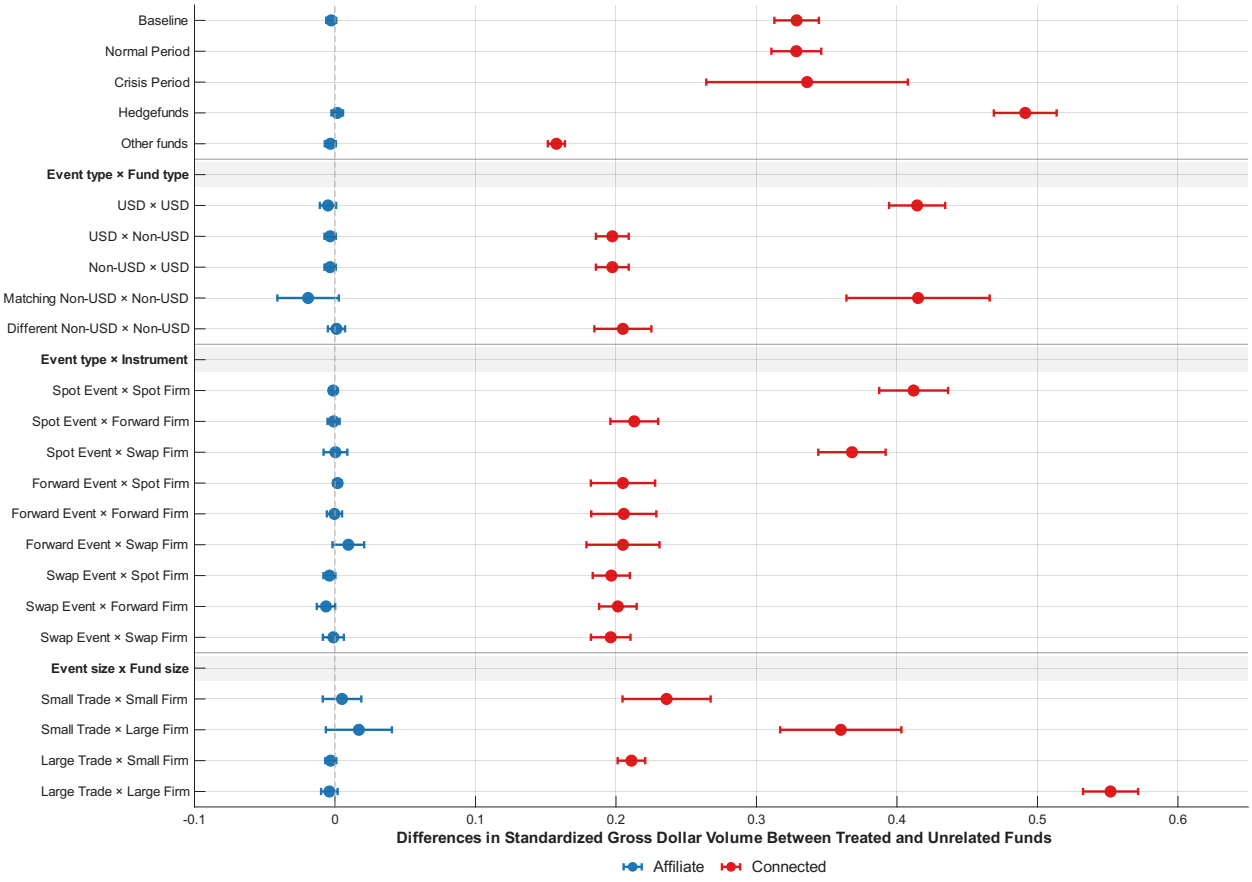

Figure 7: Responses of Funds to Dealer Events by Fund, Asset, and Event Characteristics: Standardized Daily Gross Dollar Volume and 95% Confidence Intervals

Figure 7 presents the estimates for daily gross dollar volume as the dependent variable. We find qualitatively similar results for the other dependent variables.

We never detect information sharing between affiliate dealers and funds across the event, fund, and asset characteristics. Meanwhile, connected dealers and funds share information across all characteristics. A connected fund responds more strongly to event trades in its the fund's specialty currency ('Event types × Fund type' rows), consistent with connected dealers and funds acting on the same information that prompted the dealers' event trades.

Relation to literature. Cespa et al. (2022, Table 9) find that spot trades are more predictive of currency movements than forwards and swaps. Figure 7 indeed shows that the connected funds have the strongest responses to the event trades that are spot trades ('Event type × Instrument' rows). Hacioglu Hoke et al. (2026) find that hedge funds are more likely to speculate on currency movements, and their trades are more predictive of exchange rates, than other funds. In line with this finding, we estimate a far stronger response by connected hedge funds than by other connected funds ('Hedge funds' and 'Other funds' rows) to the event trades.

7 How Did China Walls Reshape Information Sharing?

We investigate the impacts of tighter enforcement by the SEC using the 2018 sample period around the Mizuho settlement on July 23, 2018 (Section 2.1). We find that dealers cease sharing information with their affiliate funds immediately upon the Mizuho settlement, while their information sharing with their connected funds remains unchanged.

Section 7.1 develops the hypotheses and decision rules. Section 7.2 presents the results.

7.1 Hypotheses and Decision Rules

We so far infer that the SEC is the regulator responsible for enforcing the China Walls in our setting, based on its authority over registered broker-dealers (Section 2.1) and that

most ILS trades involve the same dealers as in the USD market. We tighten this inference by looking for a sharp decline in information sharing between affiliate dealers and funds around the Mizuho settlement in July 2018. We conclude that the SEC is responsible for enforcing the China Walls if the affiliate dealers and funds cease sharing information after July 2018.

We employ the connected funds as a placebo and to test for spillover effects of China Walls. First, the dealers would share less information with their connected funds if the dealers became systemically less willing to share information around July 2018, perhaps due to regulatory changes other than the Mizuho settlement. We conclude that effectively enforced China Walls drive our results if we fail to detect a decline in information sharing between connected dealers and funds around July 2018.

Second, the China Walls between affiliate dealers and funds might have spillover effects on how other dealers and funds share information. For example, after the dealers cease sharing MPI with their affiliate funds, other funds trading on the same MPI could earn a larger profit than before. The larger potential profit may induce the dealers to share more MPI with their connected funds. We reject the presence of such spillover effects if we do not detect an increase in information sharing between connected dealers and funds around July 2018.

7.2 Results

Implementation. We separately estimate equations (1) and (2) for each quarter of 2018 and the first quarter of 2019. We focus on the event-day coefficient estimates, because any response by treated funds concentrate on the event day as in Figure 6b. (Pooled estimates yield qualitatively identical results.) Events, the affiliate and the connected treatments, and the unrelated-fund control group are defined as in Section 5. We truncate the events whose 11-trading-day windows extend beyond the quarter at the end of the quarter.

Results. Figure 8 plots the resulting event-day coefficient estimates. The affiliate funds increase their gross volumes on the event day by 0.13 sd relative to the unrelated funds

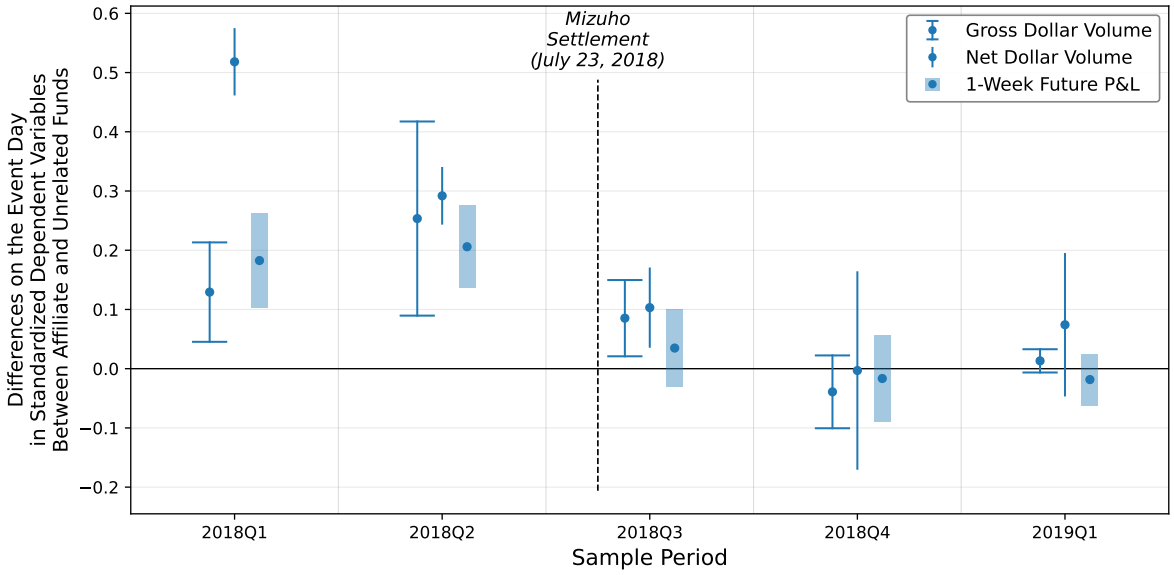

(a) Affiliate Funds

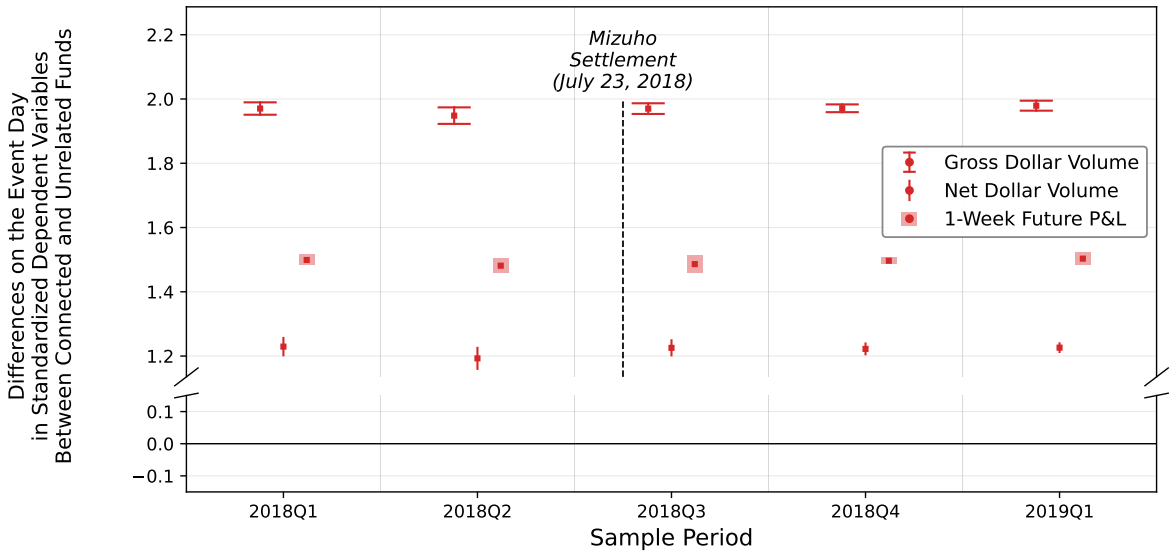

(b) Connected Funds

Figure 8: Event-Day Coefficient Estimates from (1) and (2) by Sample Period

in 2018Q1 (std. error: 0.04 sd) and by 0.25 sd in 2018Q2 (std. error: 0.08 sd). Their net volumes and P&L likewise greatly increase on the event day relative to the unrelated funds. Upon the Mizuho settlement in July 2018, the affiliate coefficient for both volumes and P&L dramatically shrink, then become economically and statistically insignificant from 2018Q4 onwards. On the other hand, the connected funds exhibit a large and stable increases in

volumes and P&L on the event day across all quarters around July 2018 (Figure 8b), with magnitudes comparable to the estimates for the 2019–2024 period (Figure 6b).

We come to three conclusions. The sharp decline in information sharing between affiliate dealers and funds from 2018Q1–Q2 through 2018Q3–2019Q1, coinciding with the Mizuho settlement, points to the SEC as the regulator responsible for enforcing China Walls in our setting. That information sharing between the connected dealers and funds do not decline around the settlement establishes the China Walls, rather than other regulatory changes, as the driver of our results. Last, that the connected dealers and funds do not increase their information sharing after the settlement rules out spillover effects due to the escalating enforcement of China Walls.

8 Conclusion

We document that today's China Walls effectively block information sharing between affiliate dealers and funds in the foreign exchange market. We identify the SEC's risk-based enforcement as the cause. Given the ease of sharing information in plausibly deniable ways (Peluso, 2020) and the difficulty in monitoring all circumstances where information might be shared, our findings reveal a remarkable regulatory capacity to control information flows within conglomerates.

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). Figure 9 summarizes the components of a banking conglomerate. A conglomerate partitions their services into insurers, commercial banks (deposits, loans), investment banks (underwriting, corporate advising), investment funds (asset management), broker-dealers (brokering, dealing, analysis, proprietary trading), and investment advisers.

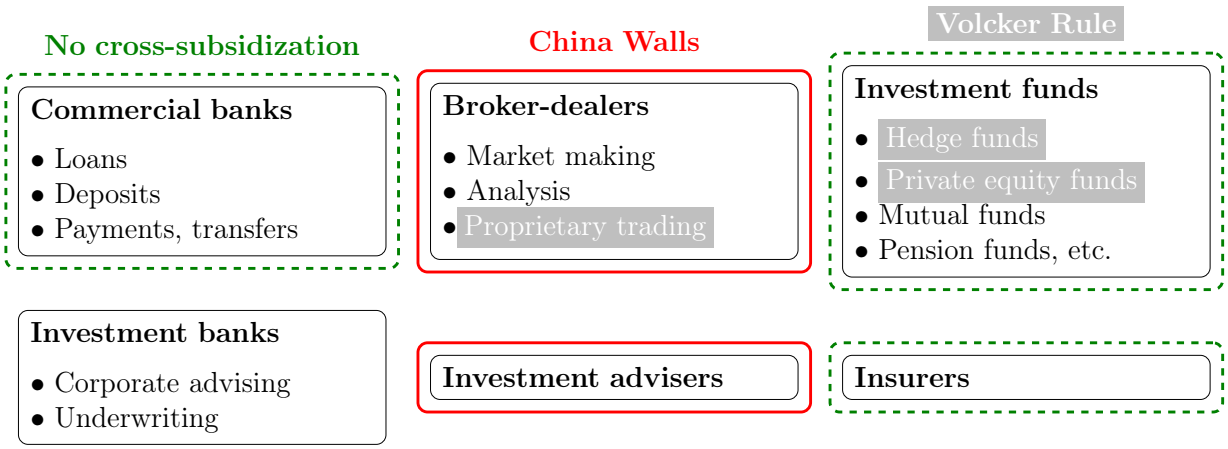

Figure 9: Stylized Banking Conglomerate and Relevant Legal Restrictions

Green dotted lines indicate restrictions on cross-subsidization: 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 China Walls: broker-dealers and investment advisers are required to prevent their employees from interacting with the employees of affiliates. Black-on-white fonts highlight the Volcker Rule restrictions on proprietary trading and ownership of hedge funds and private equity funds.

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

affect market prices. In practice, common-law courts treat as MNPI any non-publicly disclosed information that reasonable investors in the relevant securities would find important for their investment decisions. For example, insider earnings information or outstanding order flows of clients are MNPI.[15] Possessing, sharing, or acting on MNPI is not generically illegal. However, financial intermediaries owe legal duties over MNPIs, as we soon elaborate.

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

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

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

A.2 Key Regulations on Banking Conglomerates

The markings in Figure 9 indicate each key regulation on banking conglomerates. Two concerns underlie the regulations. First, the conglomerates may divert publicly insured de-

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

posits or insurance premiums towards risky trades or to cross-subsidize affiliates by 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.[16] 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 permits any trade held for more than 60 days and trading by broker-dealers that is necessary to support customer-related services (market making, hedging, etc.). 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 gather private information through market making.

Third, as we elaborate next, the China Walls around broker-dealers and around in-

[16] 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.)

vestment advisers seek to minimize information leakage surrounding these firms. Statutes single out investment advisers for their large potential impact on investment decisions. The broker-dealers are singled out, because their role as intermediaries provide constant stream of privileged information gleaned from their clients' orders. Under the argument that broker-dealers leaking this information to affiliate funds or receiving inside information from affiliates would place the investing public at a sharp disadvantage, preventing such information flows is necessary to maintain trust and participation in financial markets.

A.3 China Wall Enforcement Over Time

Origins. Under common-law tradition, insider trading on behalf of clients was encouraged. Brokers and dealers were expected to use all information that came into their possession, and further solicit inside information, to fulfill their fiduciary duty. This expectation was upended in 1961, when a landmark judgement held each conglomerate liable for damages incurred by the investing public due to trades based on its MNPI. The ruling demands that the intermediaries holding MNPI either publicly disclose or take no action whatsoever related to the MNPI. Subsequent court rulings placed the full burden of avoiding incompatible duties onto the conglomerates.[17]

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

[17] 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.

the conglomerate to liability if the mutual fund trading on the MNPI cause losses to the investing public. 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 began offering a safe harbor from liability for conglomerates that implement sufficiently strict China Walls, as determined by the SEC.[18] 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 the dealer were walled off from this fund. 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 against less regulated foreign competitors. Because the duty to disclose or abstain might render full-service conglomerates nonviable, new statutes explicitly incorporated China Walls as a safe harbor and broadened their legal protections (Brooke,

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

Burrows, Faber, Harpum, and Silber, 1995, p. 98).[19] 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.[20]

Pre-2008 crisis regulatory regime. The China Walls were initially an legal benefit available to the banking conglomerates—not a regulatory requirement. As such, the China Walls enforcement was purely reactive, occurring in the course of assigning liability upon the discovery of fraud or breach of duty. Indeed, no US regulator proactively evaluated the China Walls between 1990 and 2012, the years when the SEC reviewed the Walls within broker-dealers as a research exercise (Office of Compliance Inspections and Examinations and US Securities and Exchange Commission, 2012).[21] 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 offense in itself.

Further, financial regulators had more limited enforcement powers. Imposition of large penalties or punishment of individuals required court judgment, with 5-year statute of limitations. A firm that aided a violator could only be prosecuted if the firm knowingly assisted in the violation, a high legal bar. Most importantly, regulatory action required 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, was not themselves actionable by regulators.

[19]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 a safe harbor from a wide range of liabilities to the financial conglomerates that adopt China Walls.

[20]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).

[21]The 1990 review was in response to the 1988 ITSFEA Act that explicitly gave a safe harbor to walled-off broker-dealers. The 2012 review was in response to the Dodd-Frank Act.

Current Regulatory Regime. The US Dodd-Frank Act 2010, and partly coordinated laws elsewhere, dramatically reshaped the enforcement of China Walls. 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 sufficiently enforce information barriers among its trading desks (US Securities and Exchange Commission, 2018; Barrack et al., 2020).[22] The SEC subsequently intensified regulatory action against failures to establish, maintain, or enforce China Walls. As an ongoing example, in 2021, the SEC began a proactive sweep of monitoring and retention of business-related communication among employees across all broker-dealers and investment advisors. The first resulting settlement included a $125 million fine on Morgan Stanley for its 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 fined Virtu Financial $2.5 million in 2025 for merely having a database accessible to both broker-dealer and non-broker-dealer employees—despite producing no evidence that any MNPI was leaked (US Securities and Exchange Commission, 2024; Alston & Bird LLP, 2025). Therefore, following Dodd-Frank, the regulatory regime over China Walls morphed from reactive to proactive.

[22]The Mizuho settlement was novel in two aspects. (i) The SEC did not cite any trades potentially based on inappropriately shared MNPI. (ii) The SEC acknowledged that Mizuho had established procedures that were adequate on paper. Mizuho was penalized solely due to its failure to maintain and enforce its those policies. Closest other SEC cases involved trading while in possession of MNPI (Dechert LLP, 2017), information sharing between analysts and traders (US Securities and Exchange Commission, 2016), or prolonged recordkeeping failures (US Securities and Exchange Commission, 2015).

B Do Affiliates Without China Walls Share Information?

We exploit the conglomerates that own multiple funds to infer whether our results isolate the effects of China Walls and are not contaminated by broader regulations on banking conglomerates. The affiliate funds in the same conglomerate are not walled off from each other, while being subject the broader rules on banking conglomerates. If the affiliate funds share information among themselves in the post-2019 period of active China Wall enforcement, we infer that our results are due to the China Walls, rather than the broader regulations.

B.1 Design and Implementation

Figure 10 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. We compare the daily gross dollar volume and one-week future P&L 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 control for whether the event fund and an affiliate or unrelated fund are connected to a common dealer, removing the confounding variation from overlapping dealer connections in our estimates of interest. We conclude that our results isolate the effect of China Walls if the volumes and the P&L of the affiliate funds increase relative to the unrelated funds on or after the fund event day.

A threat to the validity of this inference is that two funds tend to be closer together in size and trading strategies than a dealer and a fund. Section B.3 partitions the funds into cells of similar and differing sizes and trading strategies to evaluate this threat.

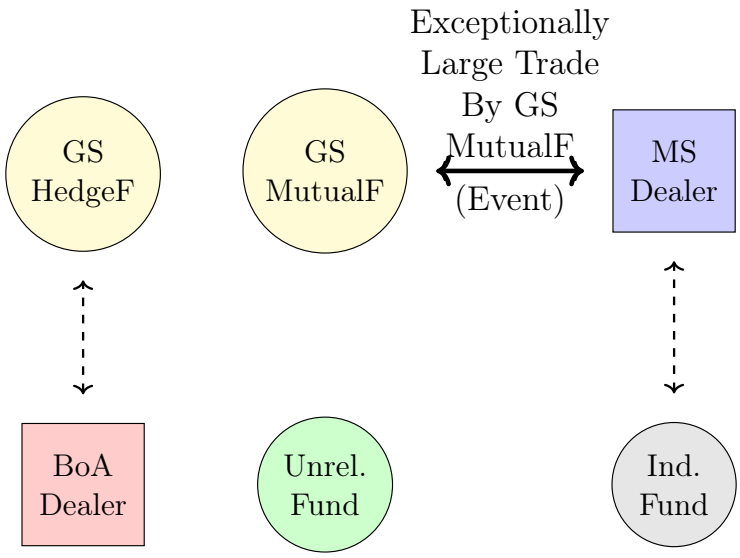

Figure 10: Identification: Information Sharing Between Affiliate Funds

We apply the following specification to the subsample of funds:

$$
\begin{aligned}
Y_{e(i)jt} = &\sum_{\tau=-5}^{5} \nu_\tau \mathbb{1}_{t=\ell_{e(i)}+\tau} Affiliate_{e(i)j} + \delta_{e(i)j} + \varphi_t + \sum_{\tau=-5}^{5} \gamma_\tau \mathbb{1}_{t=\ell_{e(i)}+\tau} \\
&+ \sum_{\tau=-5}^{5} \kappa_\tau \mathbb{1}_{t=\ell_{e(i)}+\tau} Affiliate_{e(i)j} DealerOverlap_{e(i)j} \\
&+ \sum_{\tau=-5}^{5} \eta_\tau \mathbb{1}_{t=\ell_{e(i)}+\tau} DealerOverlap_{e(i)j} + \varepsilon_{e(i)jt}.
\end{aligned} \tag{3}
$$

The control dummy $DealerOverlap_{e(i)j}$ equals 1 if fund j's set of connected dealers overlaps with event fund i's set of connected dealers, and equals 0 otherwise. We focus 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.

B.2 Results

Figure 11 presents the results from equation (3) estimated on the subsample of funds. In green are the differences ν_τ in standardized gross volume or one-week P&L between the

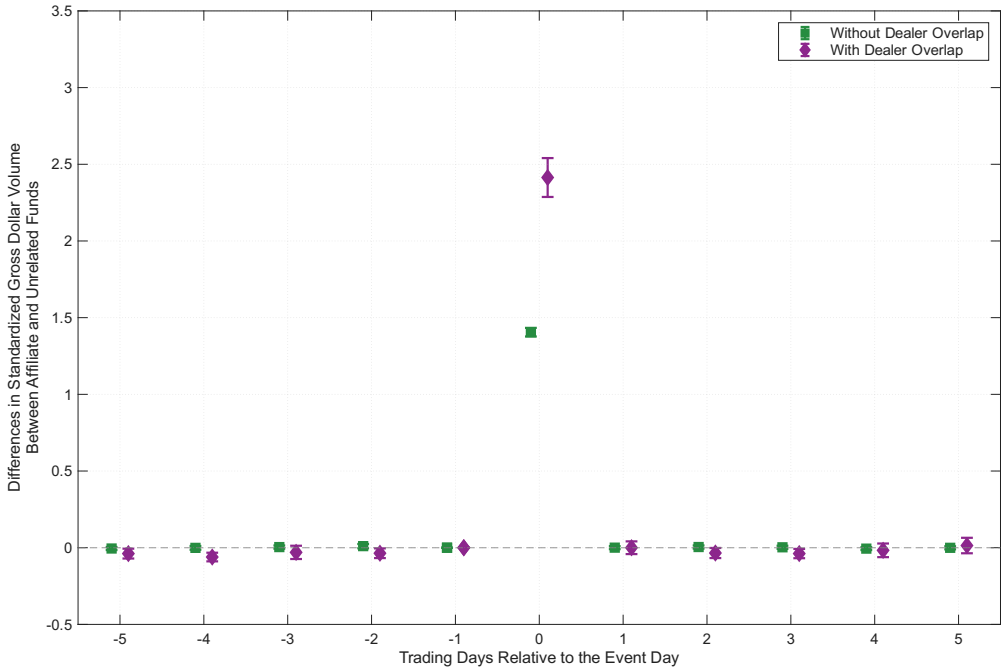

(a) Standardized Daily Gross Dollar Volume

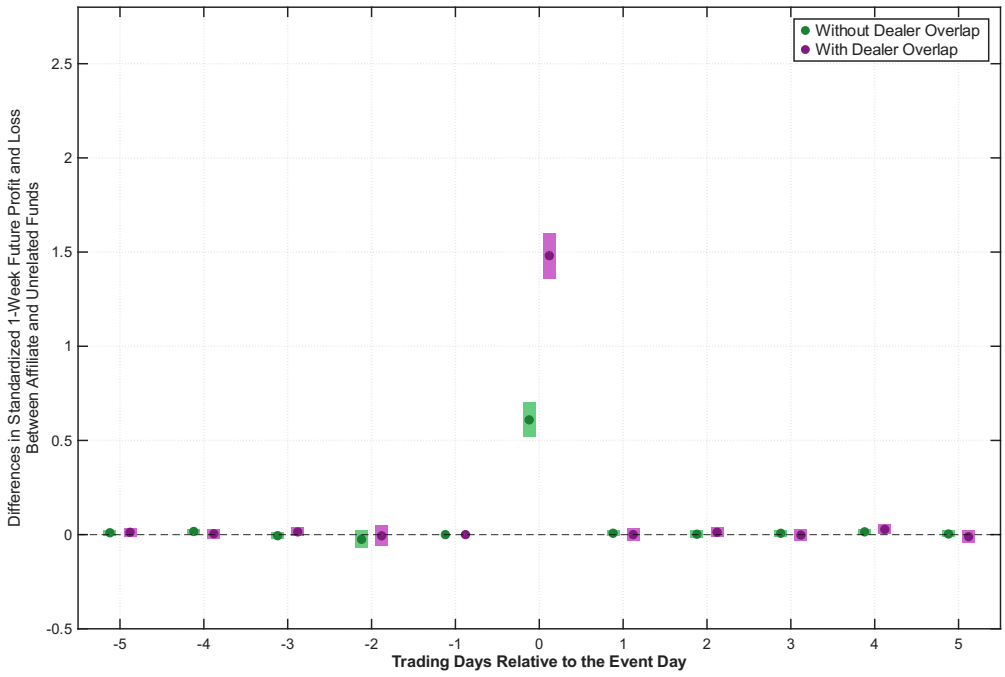

(b) Standardized Daily 1-Week Future P&L

Figure 11: Coefficient Estimates from (3): Fund Responses to Fund Events and 95% Confidence Intervals

affiliate funds and the unrelated funds whose dealer connections do not overlap with the event fund around an exceptionally large trade by the event fund. Despite removing the common shocks through any overlapping dealers, the affiliate funds increase their gross volumes by 1.4 standard deviations (std. error: 0.01 sd; Figure 11a) and their one-week P&L by 0.6 sd (std. error: 0.05 sd; Figure 11b) 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 unrelated nonoverlapping funds. As one might expect, incorporating overlapping dealer effects dramatically raises the event-date responses further.

Rightmost columns in Table 5 presents the pooled counterpart to Figure 11. They present the large sample sizes and hundreds of events that lead to our tight standard errors. There are much fewer fund events than dealer events, because most funds are independent (so have no affiliates) and do not trade frequently (so cannot have more than one event). Among the remaining fund events, many have zero never-treated affiliate funds within the event window, explaining the lower number of fund events in Table 5 compared to Table 4.

B.3 Heterogeneity in Fund Responses to Fund Events

We test the possibility that affiliate funds only respond to an event if the event fund is similar in size and other characteristics. If so, we would be limited in our ability to infer about the dealers' willingness to share information with their affiliate funds absent the China Walls.

We estimate the pooled counterpart to (3) whose treatment variables are interacted with the rich set of dummies described in Section 6. The dependent variable is the standardized daily gross dollar volume. Figure 12 plots the results. Its estimates and their standard errors are computed as we do for Figure 7, except that we only use the coefficients that are not interacted with the $DealerOverlap_{ej}$ dummy. The estimates are consistently positive and significant, and generally large, across asset, event, and fund characteristics.

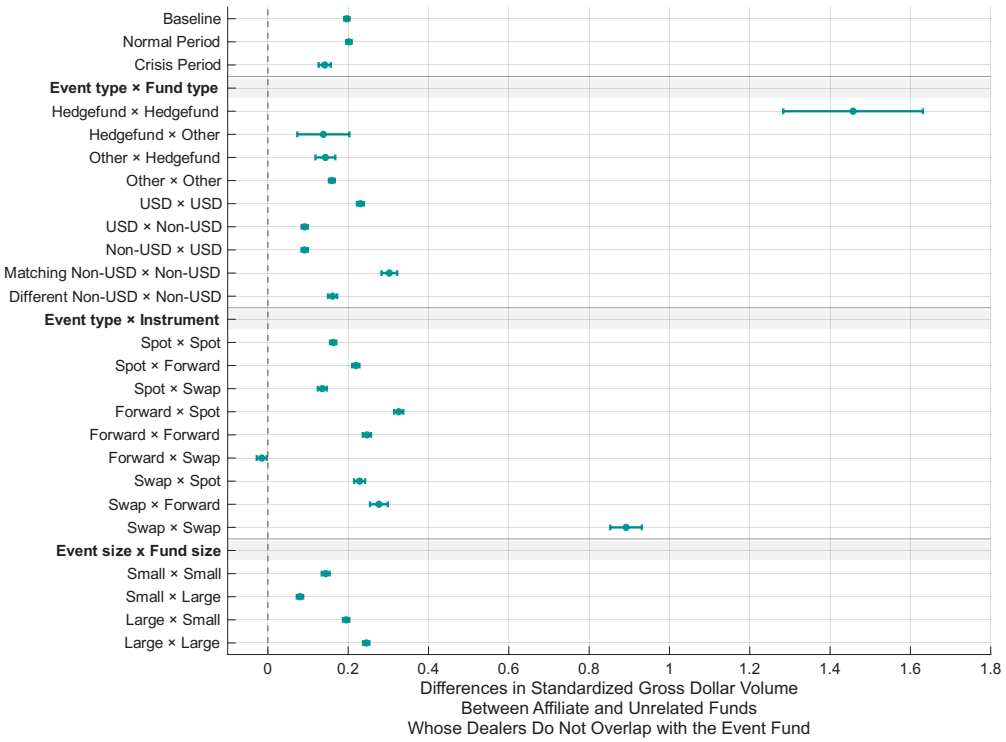

Figure 12: Fund Responses to Fund Events: Standardized Daily Gross Dollar Volume and 95% Confidence Intervals

C Identification Tests

Two exercises jointly test our two assumptions that: (I) Dealers and funds trade large sizes when they have more valuable private information; and (II) our design detects MPI sharing if and only if such sharing exists. We further rule out the possibility that dealers only maintain China Walls around exceptionally large trades, which might attract heightened attention.

The first exercise is to compute the ability of trades to predict price movements. We consider the exceptionally large, various decile (10 to 10.1st percentile and so on), and exceptionally small (99.9 to 100th percentile) trades of each dealer or fund. We call the trades in the given percentile "event trades" in this section. Under the intuition that net volumes contain information (Kyle, 1985), we net all event trades in each day separately for dealers and funds. We do not observe who initiated each trade. To sign each trade, we assume that the fund was the initiator for each dealer-fund trade, and that the event dealer

was the initiator for each interdealer trade. We limit to trades in the USD-ILS pair to avoid aggregating across currency pairs. For the fund trades, we exclude the funds with fewer than 1000 trades in our sample to keep meaningful variation between different percentiles.

For firm type $k \in \{\text{dealer}, \text{fund}\}$, firm's trade-size percentile $p \in \{[0, 0.1], [10, 10.1], \ldots, [90, 90.1], [99.9, 100]\}$, and cumulative return horizon $\ell \in \{0, \ldots, 9\}$, a three-step procedure obtains its ability to predict price movement. First, we convert the net dollar volumes on day t into trade-direction dummies $d_{t,k,p} \in \{-1, 0, 1\}$, for k and percentile p. 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 $t + \ell$ using Bloomberg USD/ILS exchange rate at 17:00 EST. Third, the ability to predict price movement is the coefficient $\rho_{k,p,\ell}$ in the time-series regression (4):

$$R_{t,t+\ell} = \rho_{k,p,\ell} \cdot d_{t,k,p} + \alpha_{k,p,\ell} + \varepsilon_{t,k,p,\ell}. \tag{4}$$

We estimate this regression for each triple (k, p, ℓ) using OLS with Newey-West standard errors that correct for serial correlation up to ℓ lags. Figure 13 plots the estimates. The net volumes from exceptionally large trades predict price movements, whereas all smaller trades do not, consistent with assumption (I).

The second exercise replicates Figure 6, except redefining event trades to be those in the various percentiles p. Figure 14a depicts the corresponding coefficient estimates for the connected funds on the event day and across p. The estimate is close to zero for every percentile except the exceptionally large trade baseline of Figure 6b. Combined with Figure 13, the daily gross volumes of the connected funds only increase relative to the unrelated funds around the event trades that are predictive of price movement, consistent with assumption (II).

Figure 14b depicts the coefficient estimates for the affiliate funds on the event day and across p. The estimate is nearly exactly zero, and statistically insignificant, for every per-

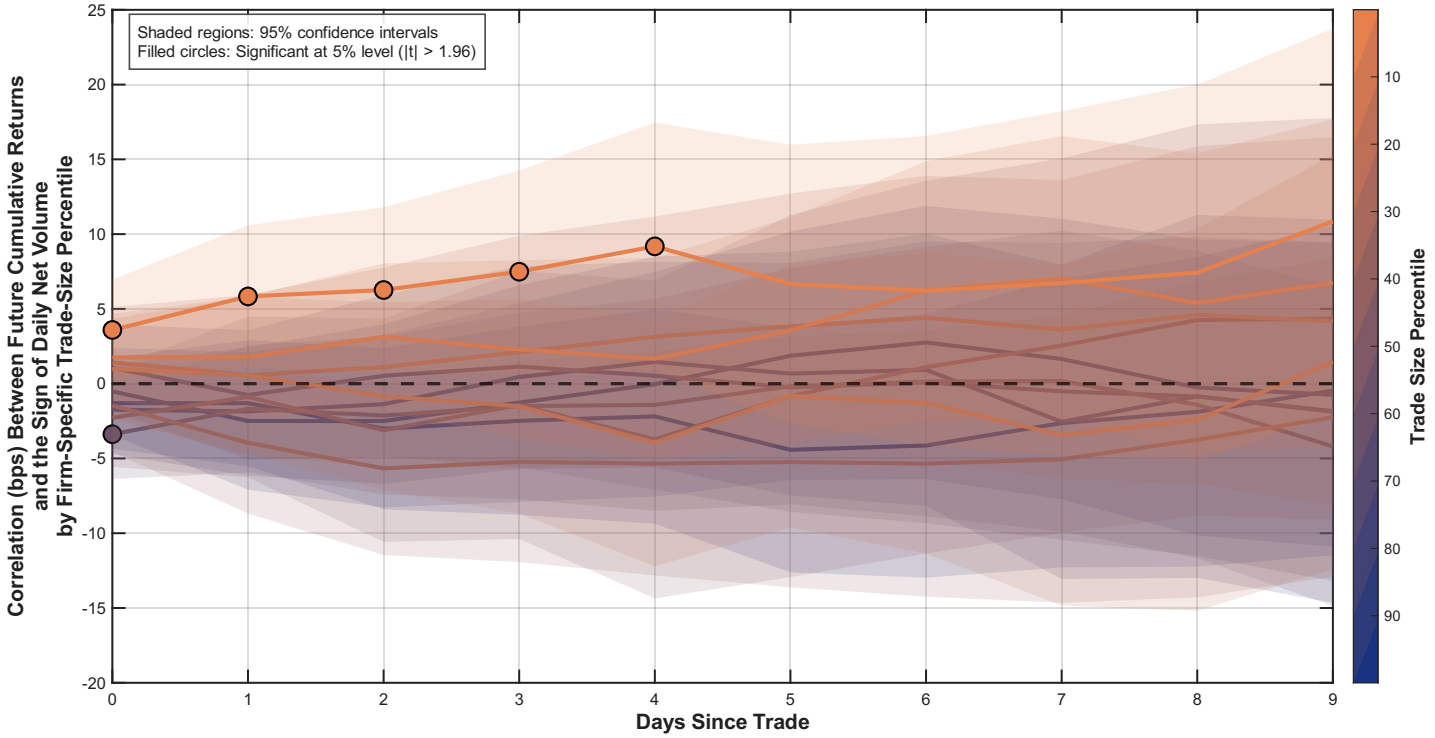

(a) Trades by Dealers

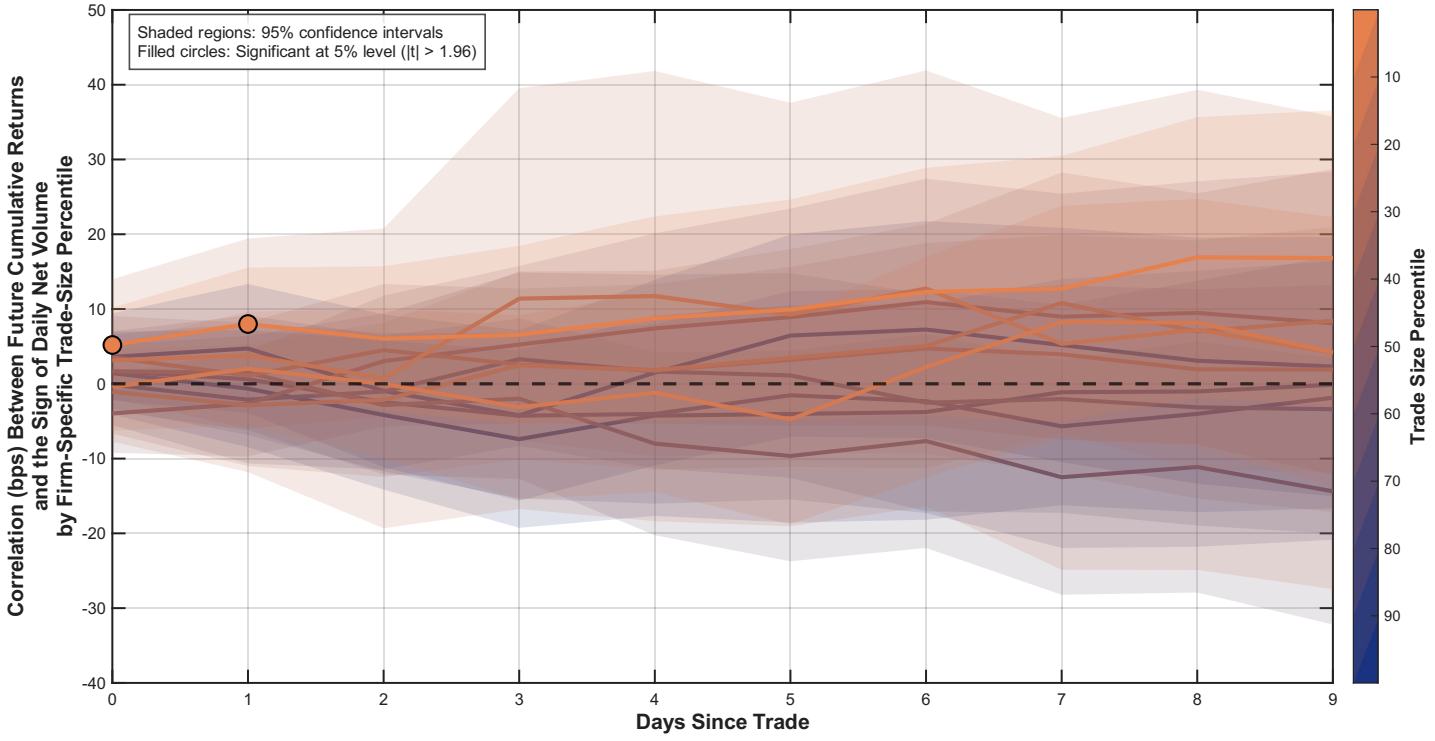

(b) Trades by Funds

Figure 13: Ability of Event Trades to Predict Price Movement and 95% Confidence Intervals

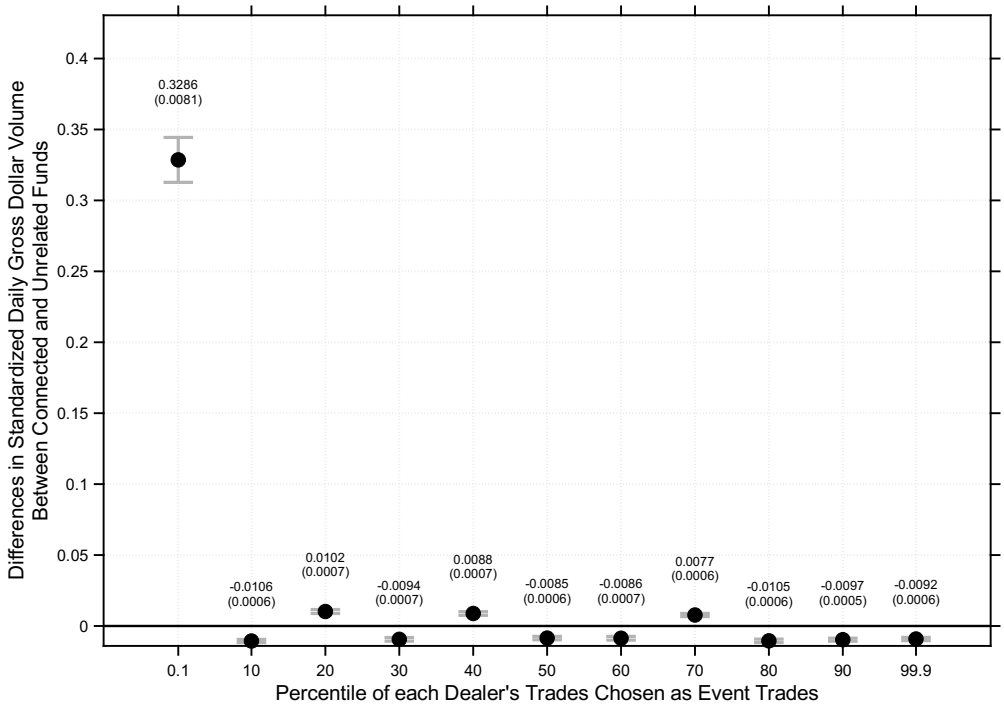

(a) Connected-Fund Treatment Effect on the Event Day

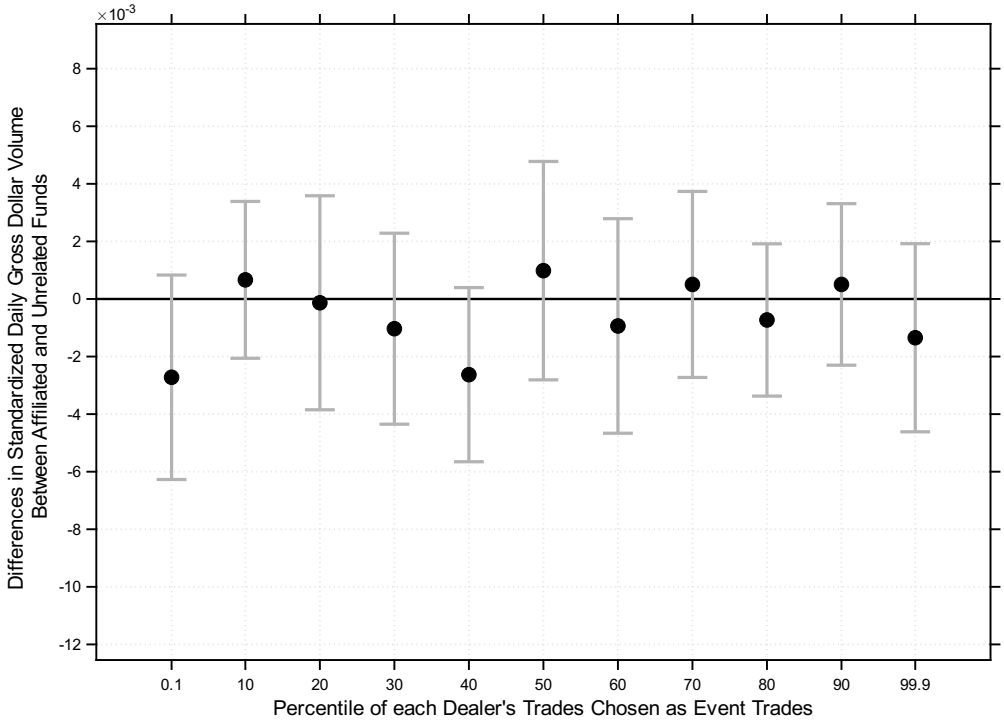

(b) Affiliate-Fund Treatment Effect on the Event Day

Figure 14: Placebo Estimates and 95% Confidence Intervals

centile. This finding is inconsistent with the possibility that the dealers only maintain China Walls around the exceptionally large trades.

D Net Dollar Volumes

Figure 15a plots the responses in the daily net dollar volumes of the affiliate funds relative to the unrelated funds around exceptionally large trades by the event dealers. Figure 15b is an analogous plot for the responses of the affiliate dealers to the unrelated dealers around the fund events. Figure 16 is the corresponding plot for the subsample of funds.

E Information Sharing from Funds to Dealers

Table 6: Summary Statistics of the Fund-to-Dealer Analytical Sample

	Mean	Std. Dev.	Min	25%	50%	75%	Max
Dealer-Day ($N = 62974$)							
Has trade	0.82	0.38	0.00	1.00	1.00	1.00	1.00
Is an event	0.04	0.20	0.00	0.00	0.00	0.00	1.00
Gross volume	774.83	1298.78	0.00	19.51	239.01	897.63	12214.47
Net volume	36.65	213.19	-2170.21	-5.76	0.04	19.99	2093.41
1-week future P&L	-0.01	2.72	-119.03	-0.01	0.00	0.01	128.06
Fund Events ($N = 1315$)							
Event trade value	15.90	53.89	0.01	0.25	1.37	7.80	1349.84
Has multiple event trades	0.10	0.30	0.00	0.00	0.00	0.00	1.00
Crisis	0.08	0.28	0.00	0.00	0.00	0.00	1.00
USD	0.88	0.33	0.00	1.00	1.00	1.00	1.00
JPY	0.01	0.10	0.00	0.00	0.00	0.00	1.00
EUR	0.05	0.22	0.00	0.00	0.00	0.00	1.00
Spot	0.62	0.49	0.00	0.00	1.00	1.00	1.00
Forward	0.29	0.45	0.00	0.00	0.00	1.00	1.00
Swap	0.09	0.29	0.00	0.00	0.00	0.00	1.00

All volumes, P&L, and values are in USD millions. No variables are standardized in this table. For the events that combine multiple event trades, we assign the characteristics of the largest trade among them to the whole event.

Figure 17 plots the responses in the daily gross dollar volumes and one-week future P&L of affiliate dealers relative to unrelated dealers around exceptionally large trades by event funds. Neither the affiliate dealers' volumes nor P&L vary relative to the unrelated dealers

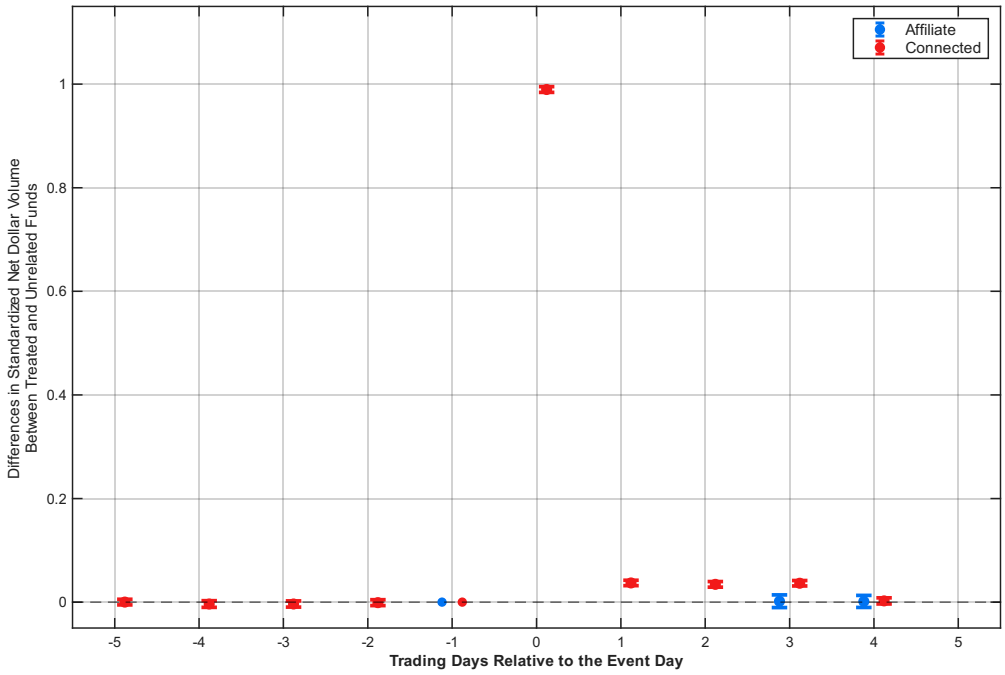

(a) Fund Responses to Dealer Events

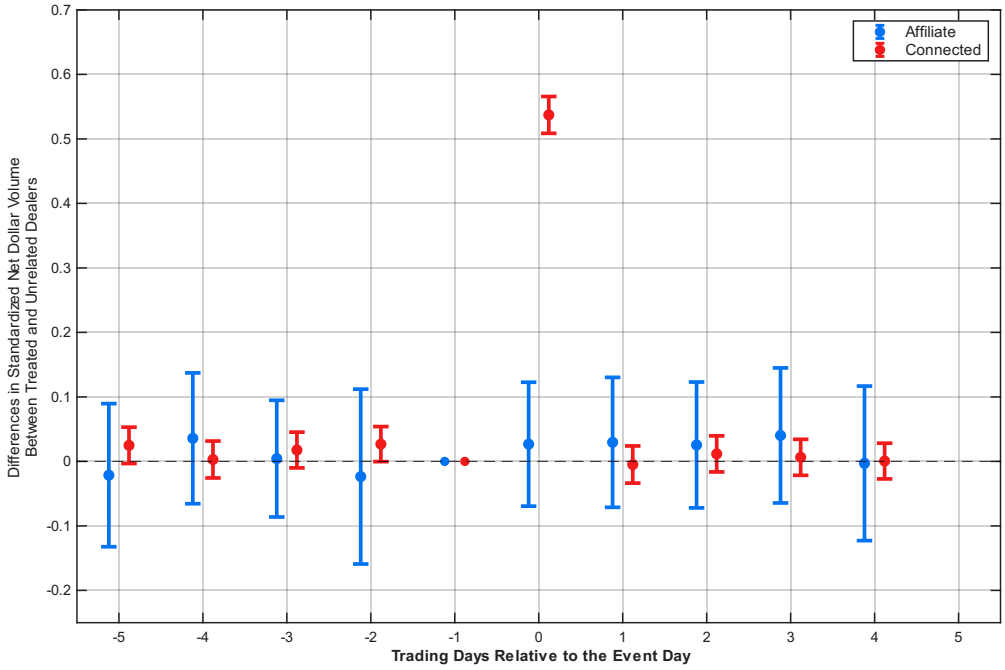

(b) Dealer Responses to Fund Events

Figure 15: Fund and Dealer Responses in Daily Net Dollar Volume and 95% Confidence Intervals

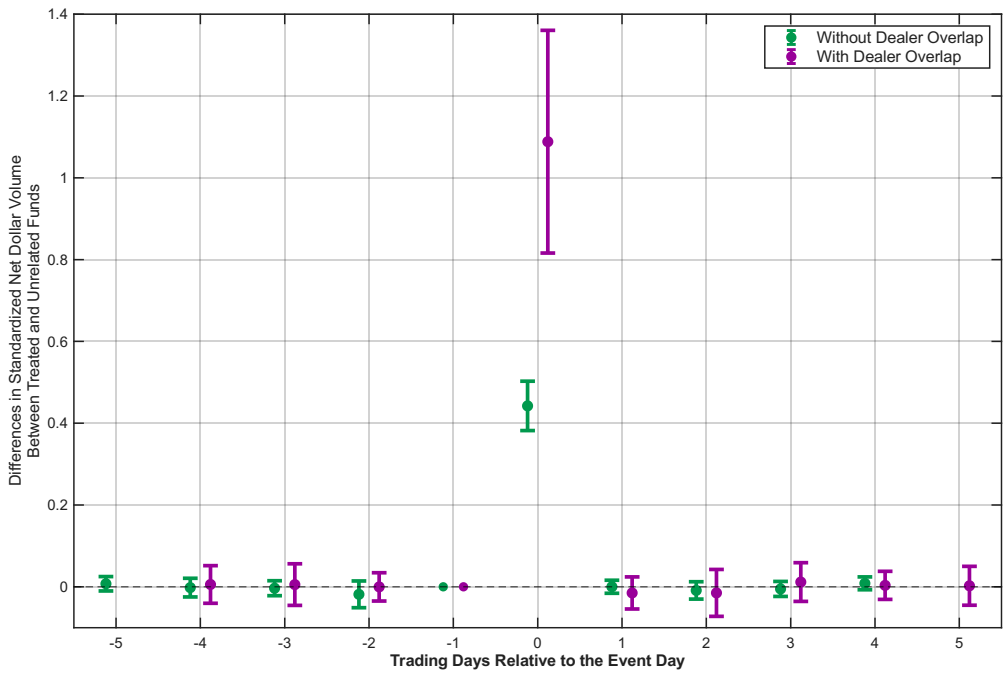

Figure 16: Fund Responses to Fund Events: Standardized Daily Net Dollar Volume and 95% Confidence Intervals

around the event day. Both the connected dealers' volumes and P&L sharply increase on the event day relative to the unrelated dealers. We conclude that the China Walls effectively block funds from sharing MPI with their affiliate dealers.

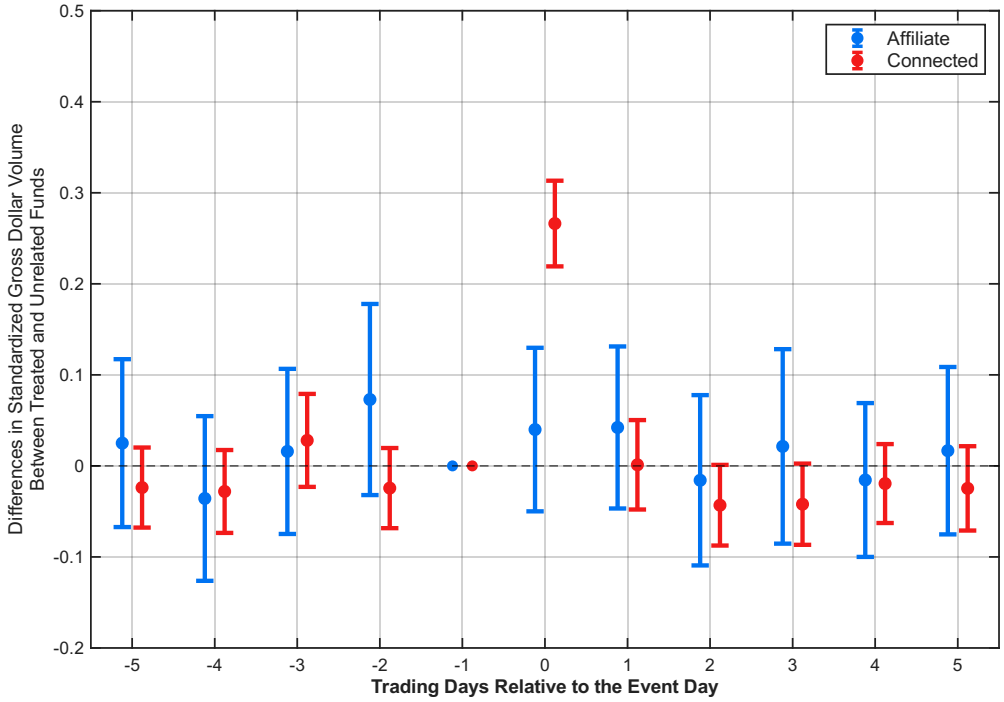

(a) Standardized Daily Gross Dollar Volume

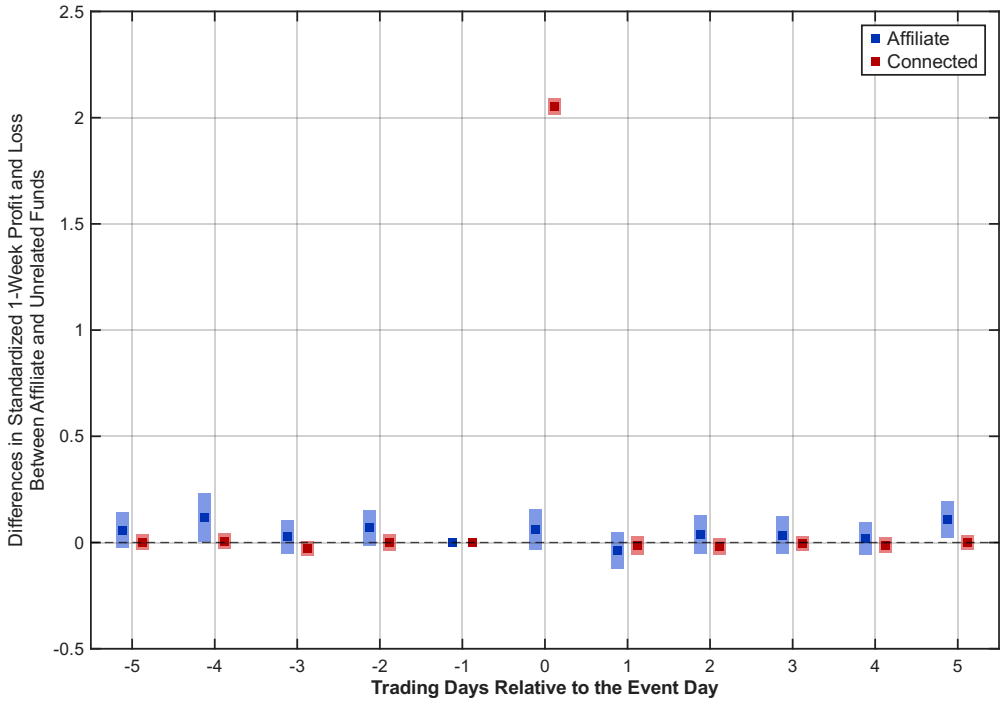

(b) Standardized Daily One-Week Future P&L

Figure 17: Dealer Responses to Fund Events and 95% Confidence Intervals

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.

Alston & Bird LLP (2025, December). SEC Settles with Virtu over Customer MNPI Controls.

Avci, S. B., C. A. Schipani, and H. N. Seyhun (2018). Eliminating Conflicts of Interests in Banks: The Significance of the Volcker Rule. *Yale Journal on Regulation 35*(2), 343–382.

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*.

Beck, T., C. Silva-Buston, and W. Wagner (2025, March). Supervisory cooperation and regulatory arbitrage. *Review of Finance 29*(2), 381–413.

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.

Cespa, G., A. Gargano, S. J. Riddiough, and L. Sarno (2022, May). Foreign Exchange Volume. *The Review of Financial Studies 35*(5), 2386–2427.

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.

Dechert LLP (2017). SEC fines broker-dealer for inadequate information barriers.

DeMarzo, P. M., M. J. Fishman, and K. M. Hagerty (1998, June). The Optimal Enforcement of Insider Trading Regulations. *Journal of Political Economy 106*(3), 602–632.

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.

Evans, M. D. D. and R. K. Lyons (2002, February). Order Flow and Exchange Rate Dynamics. *Journal of Political Economy 110*(1), 170–180.

Froot, K. A. and T. Ramadorai (2005). Currency Returns, Intrinsic Value, and Institutional-Investor Flows. *The Journal of Finance 60*(3), 1535–1566. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1540-6261.2005.00769.x.

Gardner, J. (2022, July). Two-stage differences in differences.

Garrett, D. G. (2024, December). Conflicts of Interest in Municipal Bond Advising and Underwriting. *Review of Financial Studies 37*(12), 3835–3876.

Gozzi, R. (2003). The Chinese Wall Metaphor. *ETC: A Review of General Semantics 60*(2), 171–174.

Griffin, J. M. (2021, December). Ten Years of Evidence: Was Fraud a Force in the Financial Crisis? *Journal of Economic Literature 59*(4), 1293–1321.

Hacioglu Hoke, S., D. A. Ostry, H. Rey, A. Rousset Planat, V. Stavrakeva, and J. Tang (2026, January). Topography of the FX Derivatives Market: A View from London.

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.

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.

Hendershott, T., D. Li, D. Livdan, and N. Schürhoff (2020). Relationship Trading in Over-the-Counter Markets. *Journal of Finance 75*(2), 683–734. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/jofi.12864 tex.ids= hendershottRelationshipTradingOTC2015.

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.

Kondor, P. and G. Pintér (2022). Clients' Connections: Measuring the Role of Private Information in Decentralized Markets. *Journal of Finance 77*(1), 505–544.

Krohn, I., P. Mueller, and P. Whelan (2024). Foreign Exchange Fixings and Returns around the Clock. *The Journal of Finance 79*(1), 541–578. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/jofi.13306.

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.

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.

Morrison Mahoney LLP (2018, July). With Announcement of $1.25 Million Penalty, SEC Again Places Companies on Notice of Duty to Protect Customer Information.

Office of Compliance Inspections and Examinations and US Securities and Exchange Commission (2012, September). Staff Summary Report on Examinations of Information Barriers: Broker-Dealer Practices under Section 15(g) of the Securities Exchange Act of 1934. Technical report, US Securities and Exchange Commission.

Patterson Belknap Webb & Tyler LLP (2018, July). SEC Fines Mizuho for Failing to Protect Customer Data.

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.

Ranaldo, A. and F. Somogyi (2021, May). Asymmetric information risk in FX markets. *Journal of Financial Economics 140*(2), 391–411.

Reuters (2010, February). Big banks' risky trading should be curbed-Volcker.

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.

Shemesh, M. and A. Abir (2024, March). The changing nature of the financial system in Israel in the last two decades. *BIS Papers 148*, 187–201.

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 (2015). In the Matter of Citigroup Global Markets, Inc.

US Securities and Exchange Commission (2016). In the Matter of Deutsche Bank Securities Inc.

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.