5 April 2019

Dear Professors Whited and Ye,

At the suggestion of my advisor Professor Luigi G. Zingales, my coauthors and I would like to submit the attached work-in-progress for consideration for the NBER summer workshop on Big Data and High-Performance Computing for Financial Economics.

In this paper, we leverage a massive and largely comprehensive database of anonymous New York City taxi rides to shed light on private interactions between insiders of the Federal Reserve Bank of New York and commercial bankers. To the best of our knowledge, this is the first work to use these geolocation data in a study either of interactions between institutions outside of the taxi industry or of channels of information flow between them. We have two complementary foci: 1.) An assessment of the value to financial-market participants of contemporaneous knowledge of private-interaction volumes; and 2.) An exploration of changes in opportunities for information flow from the New York Fed to financial markets around major monetary-policy meetings.

The main underlying data are trip-level pick-up and drop-off coordinates and times. We map the coordinates both to circles around relevant buildings and to polygons derived from Manhattan census blocks. Our software for data processing and econometric analyses is a combination of SAS and MATLAB, and we make use of parallel processing. We assess whether information valuable to market participants is captured by the ride data by examining the returns on a trading strategy that is easily implemented by a party with access to contemporaneous ride data. For the study of variation in interaction volumes around monetary-policy meetings, we undertake a formal data-mining exercise and assess the significance of changes using simulation-based data-mining-robust p-values.

We hope that we might have the opportunity to benefit from the lectures and the insights of the workshop's attendees, and we thank you for your consideration of this application.

With kindest regards,

David Andrew Finer

When Bankers Go to Hail: Insights on Fed-Bank Interactions from Taxi Cab Data

Daniel Bradley*, David Andrew Finer†, Matthew Gustafson‡, Jared Williams§

Abstract

We introduce taxi cab ridership between the Federal Reserve Bank of New York and large financial institutions headquartered in New York City as a novel proxy for Fed activity. A market timing strategy that buys the market portfolio (risk-free asset) when lagged Fed-bank ridership is low (high) earns a 50% larger Sharpe ratio than the market portfolio. The strategy's superior performance is especially pronounced around FOMC meetings. We also find highly statistically significant evidence that late-night meetings at the New York Fed and lunchtime offsite interactions increase around FOMC meetings, which suggests increased opportunities for Federal Reserve information to flow to markets along informal or discreet channels.

We are grateful to Marianne Bertrand, Bob DeYoung, Elyas Elyasiani, Zhiguo He, Ivan Ivanov, Bryan Kelly, Elisabeth Kempf, Paul Koch, Randall Kroszner, David Lucca, Felix Meschke, Joe Peek, Jon Scott, Amit Seru, Azeem Shaikh, James Traina, Harald Uhlig, Mihail Velikov, Annette Vissing-Jørgensen, Michael Weber, Jide Wintoki, Linghang Zeng (discussant), Luigi Zingales, Eric Zwick, seminar participants at Chicago Booth, the University of Kansas, and Temple University, as well as participants at Stevanovich Center's Market Microstructure & High Frequency Data conference, the Bank of Canada's Workshop on Data Science for Central Banks, and the 2018 Northern Finance Association (Charlevoix) for their helpful discussions and comments. We also thank individuals who provided insights into the conduct of business in New York City and Weijia Zhao for verifying some of the code and econometric derivations. The George J. Stigler Center for the Study of the Economy and the State and the Lynde and Harry Bradley Foundation provided generous support to David Finer for this work, and The University of Chicago Booth School of Business and the Liew Fama-Miller Fellowship were generous sources of general funding. Jared Williams gratefully acknowledges financial support from the Muma College of Business Center for Analytics and Creativity. We thank Ken French for sharing data on his website. All errors are our own.

*Department of Finance, University of South Florida, Tampa, FL 33620, 813.974.6358, danbradley@usf.edu

†The University of Chicago Booth School of Business, Chicago, IL 60637, dfiner@chicagobooth.edu

‡Department of Finance, Penn State University, State College, PA 16802, 814.867.4042, mtg15@psu.edu

§Department of Finance, University of South Florida, Tampa, FL 33620, 813.974.6316, jwilliams25@usf.edu

1 Introduction

It has been long established that The Federal Reserve Board's monetary policy decisions affect stock market performance (see e.g., Bernanke and Kuttner 2005). More recently, Lucca and Moench (2015), Cieślak, Morse and Vissing-Jørgensen (2019), and others further show that a sizable fraction of total realized equity returns occur at predictable times during the Federal Open Market Committee (FOMC) cycle. Cieślak, Morse and Vissing-Jørgensen (2019) suggest that one explanation for these return patterns is that there are systematic informal interactions between Fed officials and the financial sector, and that monetary policy information is transmitted to investors during these interactions. Yet, the literature lacks a direct link between Fed-bank interactions and aggregate stock market return patterns and there is little evidence on how Fed-bank interactions vary over the FOMC cycle.

In this paper we empirically examine these questions using taxi cab ridership between the New York Fed and the large systemic financial institutions it regularly interacts with as a novel proxy for Fed activity. In doing so, we provide evidence on whether Fed-bank interactions are a likely explanation for the return patterns during the FOMC cycle. Moreover, to the extent that a relation between Fed-bank interactions and stock market returns exists, the direction of the relation sheds new light on the type of information that motivates or is disseminated during these interactions.

Our measure of Fed-bank interactions is quite simple. Using taxi data from the New York City Taxi and Limousine Commission between 2009 and 2016, we pinpoint taxi pick-up and drop-off locations through latitude and longitude coordinates. We match these coordinates with the street address of the NY Fed and financial institutions.[1] We then aggregate daily cab ridership between the Fed and these banks to obtain a measure of Fed-bank interactions.

We proceed to validate that the ride data capture the movements of bankers and FRBNY activity. We begin to assess whether yellow taxi rides capture the movements of commercial bankers with a natural experiment surrounding the relocation of Goldman Sachs's headquarters and trading floor. Ride volumes between the other commercial banks and the old locations exhibit a highly significant decline coincident with the move, while the analogous measure for the new location exhibits roughly the mirror image. No such experiment is available for the New York Fed, but high Fed-bank ride volumes around major events in financial regulation suggest that yellow taxi rides are indicators of interactions between com-

[1]The New York Fed provides an ideal setting to investigate our research question. Many of the largest, most complex, and systemically important financial institutions are headquartered in New York. In addition, the Fed's monetary policy objectives are achieved via the New York Fed's trading desk, where all of the Federal Reserve System's open market operations are executed. Finally, the New York Fed is the only regional Fed branch whose president has a permanent spot on the FOMC. For these reasons, the NY Fed is considered to be the most important of the 12 regional Fed banks.

mercial bankers and the New York Fed. We find elevated ride volumes around the signing of Dodd-Frank, particularly on the business day before the final milestones. The greatest volume over the entire sample occurred the day before its passage through Congress. In addition, the average total daily ridership between the Fed and the financial institutions in our sample is 18.2 on trading days and 6.4 on non-trading days. Placebo tests indicate that this significant difference is not driven by increased traffic flow in the financial district on trading days.

We find that when Fed-bank ridership has recently been high, future stock market returns are lower on average. A tangible example of this relation is the ridership and stock returns surrounding the debt ceiling crisis of 2011. Between August 2 and 9, credit rating agencies downgraded US debt, and the S&P 500 dropped 10%. During the two weeks leading up to these credit downgrades, cab ridership between the Fed and the major financial institutions was at its highest level for the entire year. In a Wall Street Journal article, NY Fed Bank President and Vice Chair of the FOMC, William Dudley commented about private meetings with investors in his office, noting these meetings are "particularly important during periods of market stress" (Pulliam 2011). Our results are consistent with this view—when markets are vulnerable, Fed activity is elevated, presumably because either Fed or bank employees have negative information regarding aggregate returns that has yet to be incorporated into prices.

To formally examine the relation between Fed-bank ridership and stock market returns, we construct a simple market-timing trading strategy. When Fed-bank ridership has been relatively low in the recent past, the strategy invests in the market portfolio. Otherwise, the strategy invests in a risk-free asset. On average, the market returns 14 bps on days that the strategy holds the market portfolio and -0.1 bps on days the strategy holds the risk free asset. Thus, all of the market's return (in excess of the risk-free rate) during our sample period is concentrated in the approximately 50% of days that our strategy owns the market portfolio. This is remarkable, because the stock market's cumulative excess return was 210% during our sample period. This simple strategy generates an annualized Sharpe ratio of 1.45; in comparison, the market portfolio earns an annualized Sharpe ratio of 0.96 during our sample period.[2] Since our trading strategy only invests in extremely liquid assets (the market portfolio and short-term Treasuries), and it averages three trades per month, its superior performance is unlikely to be explained by trading costs. Notably, the strategy's superior performance is not driven by the previously documented pre-FOMC drift (see e.g., Lucca and

[2]The sample period referenced here is February 18, 2009, through June 30, 2016; it does not begin on January 1, 2009, because we require sufficient levels of past ridership to implement the market-timing strategy.

Moench 2015). In fact, our strategy does not consistently hold the market portfolio prior to or during FOMC meeting periods—between FOMC event weeks −3 and +3, the probability of holding the market portfolio is smallest in FOMC event weeks −1 and 0. Moreover, the trading strategy's profitability is robust to reasonable changes to all of the strategy's parameters, including the length of the cab ridership estimation period, or the lag between the ridership estimation period and the holding date. This evidence that the market performs worse when Fed-bank interactions have recently been high corroborates existing anecdotes suggesting that Fed activity increases when the Fed or the banks it interacts with possess negative information that has yet to be impounded into market prices.

We next examine how this trading strategy performs in the days surrounding FOMC meetings, when monetary policy information is most likely to get impounded into stock market prices. We find that the negative relation between Fed-bank ridership and market returns is most significant in FOMC event days −1 through +1, which generally consists of the Monday through Thursday of the FOMC week.[3] Specifically, the strategy earns a statistically significant 42 bps per day during this FOMC-event period, compared to between 0 and 20 bps per day in the three 5-day periods before or after the event period. This suggests that our market-timing trading strategy is picking up signals about the Fed's monetary policy activity.

Motivated by this evidence that our trading strategy's market-timing ability is enhanced in the days leading up to and including FOMC meetings, we construct a dataset with one observation for each of the 59 FOMC meetings in our sample period. We examine the predictability of the market's return over each FOMC meeting (event days –1, 0, and 1) as a function of the trend in Fed-bank ridership levels from the end of the previous FOMC meeting up to the current meeting. We document that stock market returns around FOMC meetings are in fact quite predictable: when daily ridership is higher in the 2 weeks immediately prior to the FOMC meeting relative to the previous 2 weeks, the market's return during the FOMC meeting is negative 24 times and positive 17 times. Conversely, when 2-week pre-FOMC ridership is lower than the previous 2 weeks, the market's return during the FOMC event window is positive 16 out of 18 times.[4] It is extremely unlikely that such patterns are driven by random chance: when a dummy variable for the market rising during the FOMC period is regressed on the dummy for the daily ridership levels increasing in the weeks prior to the FOMC meeting, the t-statistic for the ridership dummy is 3.70.

[3]Scheduled FOMC meetings begin on Tuesdays, and they usually last two days.

[4]These relations are driven by variation in ridership in the two weeks immediately prior to the FOMC meetings. There is similar ridership 3 and 4 weeks prior to FOMC meetings with positive and negative stock market returns returns, but significantly higher ridership in the two weeks before FOMC meetings that are accompanied by negative stock market returns.

Because the market-timing ability of our trading strategy is most pronounced around FOMC meetings, we take a closer look at the nature of Fed-bank interactions around FOMC meetings by examining changes in ride volumes during intraday windows over a roughly two-week span around an FOMC announcement. Strikingly, rides from the major commercial banks to the New York Fed show a highly significant jump almost immediately after the midnight lifting of a communications blackout' imposed on Federal Reserve staff. The increase between 1:00 AM and 4:00 AM is not large at around half a ride per meeting, but it is an increase of approximately 100% relative to non-FOMC period ridership at theses hours. We also find evidence of a significant increase in lunchtime meetings between insiders of the New York Fed and major commercial banks at sites away from those institutions from the day before the FOMC announcement through a week after the announcement. We find slightly over one extra lunchtime coincidental drop-off around each FOMC meeting, an increase of approximately 50% to over 100%, depending on the specification. In line with a hypothesis of informal interaction over lunch, the increase largely occurs in areas associated with dining and shopping.

Given the large increases in percentage terms and the presumably small slice of interactions captured by the taxi rides, our estimates of late-night and lunchtime meetings are likely a conservative representation of the total increase in opportunities for Fed-bank information flow around FOMC meetings. One possible explanation for the late night meetings is that New York Fed staff reduce their interactions with commercial bankers during the blackout, and interactions increase thereafter due to pent-up demand. Federal Reserve staff might also seek information from outside parties to facilitate the implementation of monetary policy. Whatever the purpose, it has been acknowledged both within and outside of the Federal Reserve System that interactions between Federal Reserve staff and outside parties entail the risk of at least accidental information flow (Federal Open Market Committee 2010, Pulliam 2011).

A few caveats are in order when interpreting our results. First, the timing and locations of the roughly 1.5 extra direct rides following the lifting of the blackout, and the coincidental lunchtime drop-offs around FOMC meetings, suggest unofficial or discreet private interactions; however, they certainly do not imply any impropriety. Second, we do not claim that our trading strategy returns our causal. One possible explanation is that the performance of the trading strategy around FOMC meetings reflects special periods of high levels of interaction, including information-gathering by the Fed, not elevated information flow from the Fed to the banks. Finally, our period of study is relatively short (January 2009-June 2016) and unique in that the federal funds rate was effectively zero. Neuhierl and Weber (2018a) exclude this period from their study for this reason and Gilbert, Kurov and Wolfe (2018)

find that the pre-FOMC drift does not exist in the more recent 2011 to 2017 period. Thus, it is not clear how our results would generalize to earlier periods.[5]

Our paper contributes to the literature on Fed activity and stock returns. In addition to Lucca and Moench (2015) and Cieślak, Morse and Vissing-Jørgensen (2019), there is a large recent literature examining stock returns patterns throughout the FOMC cycle. Neuhierl and Weber (2018a) provide evidence that the pre-FOMC drift documented by Lucca and Moench (2015) extends as far as 25 days prior to FOMC announcements. Neuhierl and Weber (2018b) examine weekly changes in 1- and 3-month Fed funds futures to construct a "slope factor" that signals changes in monetary policy. They document that a positive slope factor signals faster than expected monetary policy tightening and is negatively related to future weekly returns. Bernile, Hu and Tang (2016) find evidence of informed trading during embargoes of FOMC scheduled announcements.[6]

To the best of our knowledge, we are the first to employ the New York City taxi data to study either information flow or interactions between firms outside of the transportation sector, although several other recent papers use taxi and Uber data to address non-financial topics.[7]. In doing so, we make a methodological contribution to the economics and finance literature in the presentation of roadmaps from raw taxi data to inferences about meetings (Nolte 1996, Nolte 1997, Center for Cryptologic History 2006). Finally, our study relates to a recent strand of literature that uses aspects of the transportation and travel industries to examine finance-related questions. For instance, Koudijs (2016) examines drivers of eighteenth-century stock-price movements using data on the ships that conveyed news, Yermack (2014) and Lee, Lowry and Shu (2018) employ movements of corporate jets to study business connections and whether corporate jets are value enhancing, and Brown and Huang (2017) use public White House visitor logs to study the relationship between executives' access to policymakers and their firms' activities and outcomes.

2 Interactions between Fed and bank employees

2.1 Cab ridership data

The New York City Taxi & Limousine Commission (TLC, the Commission) has released over a billion yellow taxi trip records back to 2009. Throughout our sample period, this

[5]Note that our sample period also encompasses the tail end of the financial crisis. Our results hold if we eliminate the first half of 2009. (NBER defined June 2009 as the end of the recession.)

[6]See also Thorbecke (1997), Bernanke and Kuttner (2005), Chen (2007), Bomfim (2003), Brusa, Savor and Wilson (2018), Cieślak and Vissing-Jørgensen (2018), Jegadeesh and Wu (2017), and Ozdagli and Velikov (2018) for evidence on the relation between Fed policy and stock returns.

[7]See for example Cramer and Krueger (2016), Mammen and Shim (2018),Farber (2015), Brodeur and Nield (2017), Jackson (2018), Cohen et al. (2016)

dataset includes each trip's pick-up and drop-off times, pick-up and drop-off GPS coordinates, distance, passenger count, fare, tip and manner of payment. We end our sample in June 2016 when the TLC stopped disclosing the exact latitudes and longitudes of pickup and dropoff locations and restrict attention to yellow taxi cab rides because all of the financial institutions we examine are located in a region where only yellow taxis can provide hail service.[8] Although the TLC does not guarantee the completeness and accuracy of the ride data, it audits its ride records and has the authority to take steps to ensure adequate reporting (NYC Taxi & Limousine Commission 2017)[9]. The records were first broadly disseminated in 2014 in response to a request under New York State's Freedom of Information Law (Whong 2014). It is conceivable that for most of the period up to 2014, riders would not have expected that their trip details would ever be published.

2.2 FRBNY and commercial bank offices

The New York Fed's headquarters at 33 Liberty Street occupies its entire block, and the New York Fed has staff at 33 Maiden Lane across the street to the East (Federal Reserve Bank of New York, n.d.). Since the main entrance to 33 Maiden Lane is directly across from 33 Liberty Street, we focus on rides in the vicinity of 33 Liberty. Although some financial and legal institutions have operations in the vicinity of the New York Fed, and a hotel and private residences are located nearby, rides to and from them would have to vary systematically with Fed activity for them to do more than just add noise.

To facilitate the identification of rides with interactions between New York Fed insiders and commercial bankers, We employ only large financial institutions with which New York Fed staff would be expected to have at least a professional relationship. For this set we employ the US-based subset of the Financial Stability Board's Global Systemically Important Banks (G-SIBs) (Financial Stability Board 2014). To maximise the share of rides related to meetings during which information might flow to active participants in financial markets, we consider only important front-office presences, generally defined here as locations that are both listed as properties in 10-K filings and serve as corporate, investment-banking, financial-markets or asset-management headquarters. For most of our sample, we employ nine locations where Bank of America, BNY Mellon, Citigroup, Goldman Sachs, JPMorgan

[8] Green taxis are prohibited from providing hail service to passengers below West 110th Street and East 96th Street.

[9] Additional limitations to the taxi data are the absence of explicit passenger identifiers and affiliations, the absence of pick-up and drop-off addresses, and a lack of information on black cars and other modes of transit.

Chase and Morgan Stanley have offices.[10,11] All of the banks except for BNY Mellon are also primary dealers and are consequently expected to provide the New York Fed with market commentary (Federal Reserve Bank of New York 2010a).

A complication is that Goldman Sachs relocated its headquarters and trading floor during our sample period. Goldman Sachs reported that its primary location was 85 Broad Street at the date it filed its 2008 10-K and 200 West Street when it filed its 2009 10-K, but we do not have an exact date for the change.[12] We treat Goldman Sachs as operating exclusively at its old locations through the end of the third quarter (September 30) of 2009 and operating exclusively at its new location beginning in the fourth quarter (October 1) of 2009. Similarly, Bank of New York Mellon's headquarter location changed from 1 Wall Street to 225 Liberty Street between the filing of its 2014 10-K and its 2015 10-K. Based on our search of news articles, to the best of our knowledge, the change in headquarters occurred sometime in early 2015. Thus, we assign Bank of New York Mellon to 1 Wall Street from January 1, 2009, through March 31, 2015, and we assign it to 225 Liberty Street location from April 1, 2015, through June 30, 2016. Citigroup's headquarter change differed from Goldman Sachs's and Bank of New York Mellon's in that Citigroup moved to a location at which it already had a significant presence: their investment banking operations were housed in their new headquarter location (388-390 Greenwich St) throughout our sample period. We assign Citigroup to two locations (388-390 Greenwich St and 399 Park Ave) from January 1, 2009, through June 30, 2015, and we assign it to the Greenwich St location only from July 1, 2015 through June 30, 2016.

We report the addresses and latitudes/longitudes of the institutions in our sample in Table 1.

[Insert Table 1 here]

2.3 Measuring Fed-Bank Interactions

Since we do not know the identities of the individuals in any of the cab rides, any measure of Fed-bank interactions that we develop will have some measurement error. For example, someone living in an apartment near a bank might take a cab to the Fed to visit the gold

[10] Although non-American G-SIBs such as Deutsche Bank and HSBC appear to have large New York City footprints, an absence of data on staff sizes and building occupancy would make a staff-size-based assessment of which to include highly arbitrary. Conversely, State Street and Wells Fargo are US-based G-SIBs, but they do not report comparably important presences in New York City. Fed insiders' interactions with non-G-SIB asset managers might also be of interest, but it is less clear where there should or would be relationships, and identification would be significantly hampered by asset managers' generally smaller footprints.

[11] The buildings are listed in Table 1.

[12] Carney 2009 reports that the move had begun by November 2009.

vault, and we have no way of distinguishing such trips from actual interactions between Fed and bank employees. We can minimize the risk of false positives by restricting the areas of taxi pickups and dropoffs to extremely small areas around the doors of the Fed and the banks in our sample. However, doing this will exclude some actual Fed-bank interactions from our sample for several reasons.[13] Thus, there is an inherent tradeoff when classifying rides as being between the Fed and a bank or not: if we use too small of an area, we will miss a lot of valid Fed-bank interactions, and if we use too large of an area, we will falsely classify many rides as being Fed-bank interactions. In light of this tradeoff, we develop two measures of Fed-bank interactions: a "circles-based" errs more toward inclusion, and a more targeted "blocks-based" measure that is a cleaner, but less complete, measure of Fed-bank interactions.

2.3.1 Circles-based measure of Fed-bank interactions

The New York Fed is headquartered at 33 Liberty Street. It only has doors on two (parallel) sides of the building: on Liberty Street and on Maiden Lane.[14] Given that all of the entrances and exits to the building are close together, we use a circle with a 0.05 mile radius centered on the building (i.e., coordinates (40.708366, -74.008712)) to classify rides as being to or from the Fed. The banks in our sample are larger and have more entrances and exits to the buildings. For example, Citigroup has adjacent buildings at 388-390 Greenwich Street. We therefore use larger circles (with radii of 0.10 miles) to classify rides as being to or from these institutions.

2.3.2 Blocks-based measure of Fed-bank interactions

Census blocks largely correspond to city blocks. Where a building or complex of interest is not the only substantial development on its census block, our blocks-based measure restricts attention to the slice of the block occupied by that building or complex (except in robustness checks). Due to the scatter that we observe around roads and the possibility that relevant pick-ups and drop-offs occur across from buildings, we use expanded versions of census blocks or slices thereof.[15] To construct an expanded area, we first obtain points along the expanded

[13] First, the cabs' GPS devices have some measurement error (as described above). Second, cab drivers might not drop employees off right next to the door due to traffic conditions in the immediate vicinity of the buildings—for example, it might be faster for the passenger to be dropped off a block away from the building and walk the remaining distance. Finally, the passenger might choose to walk to a more convenient location to be picked up by the cab—for example, perhaps there is more cab traffic half a block away from the employee's building, so she decides to walk half a block before trying to hail a cab.

[14] The other sides of the building are quite short, and they lie on William Street and Nassau Street. However, there are no doors on these streets.

[15] Figures A.1 and A.2 in the Internet Appendix provide examples.

perimeter by shifting each vertex 100 feet outwards along each of the two perpendiculars to the edges of which that vertex is a part. The intersections of lines parallel to the original edges through those shifted points are the expanded area's vertices.

2.4 Ridership summary statistics

We report statistics from the distribution of ridership levels between the Fed and the financial institutions on trading days in Table 2.[16] Using the circles-based ride classification methodology (Section 2.3.1), Citigroup's 388-390 Greenwich St location exhibits the highest ridership to/from the Fed at 4.45 cab rides per trading day, followed by Goldman Sachs' new headquarter location at 200 West Street with 3.27 rides per trading day. The locations with the least cab ridership to and from the Fed belong to the Bank of New York Mellon with 0.77 and 0.84 rides per trading day at its old headquarter location (1 Wall Street) and its new headquarter location (225 Liberty Street), respectively.

[Insert Table 2 here]

In general, the blocks-based measure indicates similar patterns, with slightly lower levels of ridership. Citi's 388-390 Greenwich St location hast the highest ridership to/from the Fed at 4.14 cab rides per trading day, followed by Goldman Sachs' new headquarter location at 200 West Street with 2.40 rides per trading day.

In Figure 1, we plot the time series of the average daily ridership between the Fed and the banks in our sample. There is a noticeable downward trend in ridership levels during our sample using either our circles- or blocks-based measures (see Panels A and C) as well as more aggregated measures of all rides to and from the Fed (not necessarily to the banks in our sample) and all yellow cab rides in New York City. Presumably, this time trend is partly due to the introduction of Uber.

[Insert Figure 1 here]

As we argued earlier, we undoubtedly misclassify some rides as representing a face-to-face interaction between Fed and bank employees, when in reality, they are unrelated to Fed-bank interactions. We expect this noise component of our measure to be related to idiosyncratic factors (e.g., the weather) that are unlikely to affect the number of cab-facilitated Fed-bank interactions. For example, if cab ridership is high across New York City on a given day (e.g., because it is rainy), we expect to find more rides between the Fed and the banks in our

[16] Since aggregate Manhattan activity reaches a nadir between 4:00 AM and 5:00 AM, and rides very early in the morning are presumably more likely to reflect long days at work than early starts, our activity-based definition of a day is the span from 5:00 AM through 4:59 AM the following day.

sample, but not because we expect there to be more Fed-bank interactions on that day. To account for this, we consider two deflators for our measure: the number of Yellow cab rides across New York City as a whole, and the number of cab rides to or from the New York Fed (but not necessarily to any of the banks in our sample). Notably, these deflated measures, which we plot by month in Panels B and D of Figure 2, exhibit less of a time trend.

2.5 Validating our measures of Fed-bank interactions

2.5.1 Trading-day versus Non-trading-day Ridership

If our measure captures taxicab rides from the New York Fed to the banks, then we expect elevated ridership on trading days compared to non-trading days. In Table 3, Panel A, we report statistics on ridership between the Fed and each of the banks in our sample on trading and non-trading days. Across the entire sample of New York City yellow cab rides, there is relatively little difference between the levels of ridership on trading days versus non-trading days: 467,520 per day on trading days versus 402,537 on non-trading days. In contrast, there is a significant difference between ridership on trading days versus non-trading days among rides that we classify as being between the Fed and one of the five financial institutions: 18.23 per day on trading days versus just 6.36 on non-trading days.

[Insert Table 3 here]

To determine whether the differential ridership to and from the Fed on trading days is driven by cab ridership being generally higher in the financial district on trading days versus non-trading days, we conduct a placebo test where we examine cab ridership inside a "doughnut" surrounding each institution, where the "hole" of the doughnut corresponds to the circle we use to classify rides as being to/from the institutions in our sample. For example, throughout this paper, our circles-based classification considers a rides as being to or from the Fed if its pickup or dropoff location is within 0.05 miles of the Fed. For the placebo tests (only), we classify a ride as being to/from the Fed if its pickup/dropoff location is between 0.05 miles and 0.10 miles from the Fed. Similarly, for the other institutions in our sample, for our placebo tests (only), we classify a ride as being to/from the bank if its pickup/dropoff location is between 0.10 miles and 0.15 miles from the bank. By construction, this test identifies adjacent locations in the financial district to our institutions.

Recall that Fed-institutions ridership was 2.9 (18.23/6.36) times higher on trading versus non-trading days. Our specification tests if this increase in ridership on trading days is significantly different than the placebo classification. Panel B presents the results of this analysis. Note that, as expected, the level of ridership is higher in the placebo sample because

the bandwidth encompasses a much larger footprint. The average increase on trading days is 1.9 (35.56/18.81). This is significantly lower (t-stat=18.27) than the average increase between Fed-bank pairs as shown in Panel C.

2.5.2 Ridership around Goldman Sachs Headquarter Move

We next examine whether taxi rides reflect the movement of business activity surrounding Goldman Sach's headquarters and trading floor relocation from 85 Broad Street and One New York Plaza to 200 West Street from the old (Carney 2009,Craig 2010). Figure 2 Panels A and B present two-deviation confidence bands for monthly means of weekday ride volumes between Goldman Sachs's old locations and new headquarters and the buildings of the other major commercial banks.[17] The decline in the old locations' volume and the increase in the new headquarters' volume are strikingly coincident and begin around November 2009, approximately when the first staff began working at 200 West Street (Craig 2010). The confidence bands after mid-2010 generally do not overlap with any of those before November 2009, which suggests the highly significant structural breaks in the geography of business interactions that would follow a relocation of a major bank's front-office operations. The magnitudes of the changes in volume are different, but usage of taxis, the proximities of pick-ups and drop-offs to the sites and the relocation of Goldman Sachs employees from other locations could be responsible.

2.5.3 Ridership around Dodd-Frank Milestones

To assess whether the taxi data capture the conduct of FRBNY business, we examine ride volumes around milestones in the passage of The Dodd-Frank Wall Street Reform and Consumer Protection Act (Dodd-Frank). Dodd-Frank represents the greatest overhaul of American financial regulation in decades, and one would expect that preparations for its possible passage would increase interactions among FRBNY staff at its main locations, the institutions that the Federal Reserve System regulates and the Fed staff embedded in those institutions (Paletta and Lucchetti 2010). We examine ride volumes on, before and after the dates listed in the United States Congress's overview of actions on Dodd-Frank with the exception of the long list of dates on which only conference committees were held.[18] One observes a striking pattern: from the filing of the conference report onward, the day before

[17]For greater consistency with boundaries of workdays, throughout this paper a day is defined as 5:00 through 4:59 rather than 0:00 through 23:59.

[18]The submission and passage of legislation are presumably more relevant to New York Fed staff than ongoing negotiations, and it is not immediately clear that commercial banks would see the New York Fed as a potentially valuable agent for lobbying Congress.

each milestone is in the top 99.4 per cent of days by ride volume (Table 4 Panel A). The day before the Senate agreed to the conference report and the day before President Obama signed the bill into law saw the two highest volumes of the whole sample.

High ride volumes are not limited to those days: Figure 2 Panel C presents two-deviation confidence bands around the monthly means of weekday ride volumes, and overall ride volumes in June and July 2010 are exceptionally high. This might partially reflect consultations with the major law firm Milbank, Tweed, Hadley & McCloy LLP (Milbank), but the finding of high volumes is robust to the exclusion of rides with pick-ups or drop-offs within 100 feet of Milbank's headquarters (Table 4 Panel B and Figure 2 Panel D).

3 Relationship between ridership and subsequent stock market returns

As noted in the introduction, it is well-known that the Federal Reserve Board's monetary policy decisions affect stock market performance. More recently, researchers have shown that a sizeable fraction of total realized equity returns occur at predictable times during the Federal Open Market Committee (FOMC) cycle (see, e.g., Lucca and Moench 2015; Cieślak, Morse and Vissing-Jørgensen 2019). Cieślak, Morse and Vissing-Jørgensen (2019) argue that the return patterns they document might be driven by systematic informal communications between Fed officials and the financial sector. However, existing literature lacks a link between Fed-bank interactions and stock market return patterns. We seek to fill this void by examining the relation between stock market returns and recent Fed-bank cab ridership.

3.1 Description of trading strategy

We consider a simple trading strategy that either holds the market portfolio or a risk-free asset. If recent cab ridership between the Fed and major financial institutions has been relatively low, the strategy holds the market portfolio; otherwise, it holds a risk-free asset. Specifically, let $Ridership_{10,t}$ denote the (scaled) number of rides between the Fed and the banks in our sample on trading days $t-10$ through $t-1$, inclusive. The strategy holds the market on trading day $t+1$ if $Ridership_{10,t}$ is less than the median value of $Ridership_{10,s}$ between trading days $t-20$ and $t-1$. Otherwise, the strategy owns the risk-free asset on trading day $t+1$.

For our main specification, we utilize the more inclusive circles-based ride measure that we developed in Section 2.3.1, however (as we discuss and show in Appendix A) our main results are qualitatively similar using the blocks-based measure. To account for the possibility that New York City cab ridership varies for a variety of reasons (including weather, the introduction of Uber, etc.), our baseline specification deflates Fed-bank ridership by the

number of cab rides to or from the New York Fed (but not necessarily to any of the banks in our sample).As we show in the following section, our findings are qualitatively similar adjusting this deflator or changing any of the trading strategy's parameters.

3.2 Performance of trading strategy

We analyze the performance of this market timing strategy in several ways. First, we compare the average return of the market (in excess of the risk-free rate) on days that the strategy holds the market with the average market return on days that the strategy holds a risk-free asset. Table 5 shows that on days our strategy holds the market portfolio, the average daily excess return of the market is 13.8 bps.[19] On days our strategy holds a risk-free asset, the average daily excess return of the market is −0.10 bps. This is remarkable, as it reveals that all of the market's excess returns during our 7.5 year sample period—a period when the cumulative excess return of the market was 210%—was entirely concentrated during periods following relatively low Fed-bank cab ridership. The difference in average market returns on days that our strategy owns the market portfolio and days on which it owns the risk-free asset is statistically significant with a t-statistic of 2.69. To our knowledge, this is the first evidence that face-to-face interactions between the Fed and financial institutions predict future stock market returns.

[Insert Table 5]

Next, we examine the Sharpe ratio based on daily returns. When using daily returns, we compute Sharpe ratios as the ratio of the average daily excess return divided by the standard deviation of the daily excess returns, multiplied by the square root of 252. Our market timing strategy generates an annualized Sharpe ratio of 1.45, whereas the strategy that takes the opposite side of our strategy's position (i.e., a risk-free asset if our strategy owns the market, and the market if our strategy owns a risk-free asset) generates an annualized Sharpe ratio of −0.01. For comparison, the annualized Sharpe ratio for the market as a whole over this time period was 0.96.[20]

A notable feature of this trading strategy is that the strategy is not closely related to the previously documented positive returns in the weeks leading up to the FOMC meetings (see e.g., Lucca and Moench 2015). Figure 3 provides evidence of this as our strategy does not always hold the market portfolio in the days leading up to the FOMC meetings cycle. In fact, our strategy tends to invest in a risk-free asset (not the market portfolio) in the week

[19] Excess daily market returns are taken from Ken French's website.

[20] Because we require past levels of ridership to implement our trading strategy, the sample period described here begins on February 18, 2009 (not January 1, 2009).

prior to and the week of the FOMC meeting. Thus, our trading strategy is clearly not a repackaging of the previously documented pre-FOMC drift.

[Insert Figure 3 here]

These findings empirically link Fed-bank interactions and aggregate stock market returns. The fact that the market performs worse when Fed-bank interactions have recently been high suggests that there are more Fed-bank interactions when the Fed has negative private information about the state of the economy. That is, it appears as though the Fed increases activity when it possesses negative information that has yet to be impounded into market prices.

[Insert Figure 4 here]

In Figure 4, we plot the excess return of our strategy versus the excess return of the market, by month. The returns are highly correlated (0.67), have almost identical average monthly excess returns (1.44% for the strategy versus 1.45% for the market). The outperformance of the strategy is evident in the lower standard deviation in the excess returns (3.4% for the strategy versus 4.0% for the market), which causes the strategy to achieve a superior annualized Sharpe ratio (1.46 versus 1.26).[21] Regarding possible crash risk, the strategy does significantly underperform the market in November, 2011, losing 8.0%. However, the market was flat that month (excess return for the month: –32 bps), so the strategy's biggest crash did not occur during a market downturn.

In Figure 5, we examine the trading activity of the strategy over time. Panel A reports the proportion of days each month that the strategy owns the market portfolio. Recall that when the strategy does not hold the market portfolio, it is invested in a risk-free asset, so this figure also reveals the proportion of days that the strategy invests in a risk-free asset each month. In all but one of the 88 months, there is at least one day in the month where the strategy owns the market portfolio, and at least one day that it owns a risk-free asset. The only exception is May 2015, when the strategy owns the market portfolio every day.

[Insert Figure 5 here]

Regarding trading frequency, the strategy makes at least one trade in every month. In the average month, the strategy trades 3.02 times, with a maximum of 12 trades in June,

[21] These annualized Sharpe ratios based on monthly returns do not contradict the Sharpe ratios we reported using daily returns. First, since we are analyzing monthly returns, we begin our sample in March 2009, not February 18, 2009. In addition, annualized Sharpe ratios generally differ when estimated based on daily returns versus monthly returns.

2015. We plot the time series of trading volume in Panel B of Figure 5. The transaction costs of this strategy are likely low because it only invests in liquid assets (the market portfolio and Treasuries).

3.3 Robustness of trading strategy performance to parameter selection

We next examine the robustness of the trading strategy to variations of the strategy, including: (i) the 10 day ridership estimation period (i.e., use of $Ridership_{10,t}$), (ii) our choice of deflator, (iii) our choice of the 50th percentile of past values of $Ridership_{10,t}$, (iv) our choice of a one day lag between the date that $Ridership_{10,t}$ is computed and the day that the strategy takes its position, and (v) our choice of 20 days to estimate a benchmark level of $Ridership_{10,t}$.

In Table 6, we explore the sensitivity of our profits to each of these specification choices. In Panel A, we consider 5 and 15 day estimation periods as opposed to the 10 day estimation period that our baseline specification uses. In Panel B, we deflate our ridership measure by the total yellow taxi ridership across all of New York City as well as not using any deflator. In Panel C, we consider a strategy that decides whether or not to hold the market portfolio by comparing $Ridership_{10,t}$ to the 25th and 75th percentiles in its recent past rather than its past median value. In Panel D, we consider 0 and 5 day lags instead of the 1 day lag that our baseline strategy uses. Finally, in Panel E, we use benchmark periods (i.e., the period we use to look back to compare the current level of $Ridership_{10,t}$ to its past values) of 15 and 25 days instead of the 20 days that our baseline specification utilizes. Our results are qualitatively similar across all these specifications. Recall that our baseline specification produces daily excess returns of 13.8 (–0.1) when we are in (out) of the market. In these alternative models, the average daily excess return of the market on days that our strategy holds the market portfolio ranges from a low of 10.1 bps to a high of 20.6 bps, while the average daily excess return on days the strategy holds a risk-free asset ranges from a low of –2.7 bps to a high of 3.1 bps. The t-statistics for the difference in means range from a low of 1.47 to a high 3.46. The annualized Sharpe ratios for the various trading strategies range from a low of 1.11 to a high of 1.60, while the annualized Sharpe ratios for the strategies that take the opposite position as our strategies (i.e., owns the market when our strategy says to own a risk-free asset, and vice versa) range from a low of –0.27 to a high of 0.30.

[Insert Table 6 here]

In Table **??** and Table **??** we show that the results from this section are qualitatively similar using the blocks-based measure of Fed-bank ridership. Table **??** shows that using the blocks-based measure, the baseline trading strategy yields 12 bps per day when the it

invests in the market, while the market returns only 1 bp per day when the strategy is not invested. This yields a Sharpe Ratio of 1.26.

3.4 Market-timing ability around FOMC meetings

3.4.1 Trading strategy's profits in FOMC event time

The Federal Reserve has two main functions: monetary policy and financial supervision. So far, we have documented that elevated Fed-bank interactions are followed by abnormally low returns in the stock market. This could be due to either of the Fed's primary responsibilities. For example, it is reasonable to expect the Fed to increase its supervisory activity when financial institutions are in trouble. If the market learns about these problems with a lag, then low market returns would be expected to follow elevated Fed-bank interactions.

In this section, we attempt to determine whether our strategy reflects the Fed's monetary policy activity or its supervisory activity. Senior Federal Reserve officials decide monetary-policy matters such as the federal funds target range and changes to quantitative-easing programmes at the eight annual Federal Open Market Committee (FOMC) meetings in Washington, D.C. The New York Fed provides inputs into FOMC deliberations, hosts the trading desk that performs the transactions that implement FOMC decisions, and decides the distribution of bond purchases over maturities. Hence, it is natural to examine the market-timing ability of our trading strategy around these meetings. If the strategy exhibits unusual market-timing ability around FOMC meetings, then its superior performance is likely due (at least in part) to intertemporal variation in the Fed's monetary policy activity.

In Figure 6, we plot the average difference in the expected excess return of the market (in basis points) on days that our trading strategy owns the market portfolio and days that it owns the risk-free asset, aggregated by FOMC-event-week. Specifically, we let FOMC event time 0 denote the returns around the FOMC meeting, namely, the FOMC event days -1, 0, and $+1$. Since scheduled FOMC meetings begin on Tuesday, and they generally last two days, this generally consists of Monday through Thursday of the FOMC week. FOMC event time -1 consists of the week prior to the FOMC announcement period, namely, FOMC event days -6 through -2 (inclusive).[22] Similarly, FOMC event time $+1$ consists of FOMC event days $+2$ through $+6$. The other event times are defined analogously.

[Insert Figure 6 here]

Inspecting the figure, it is apparent that our trading strategy exhibits abnormally good market-timing ability in the immediate vicinity of FOMC meetings. The difference in the

[22] Recall that our analysis is restricted to trading days, hence 5 days generally equals one week.

market's daily excess return when the strategy is in the market versus when it is in risk-free assets is 46 bps, which is more than twice as large as the next largest value (21 bps on FOMC event week −3). Moreover, the market-timing ability in event time 0 is highly statistically significant ($t = 2.72$), while the next largest t-statistic is 1.52 (FOMC event week −3). The fact that our strategy's market-timing ability is strongest around FOMC meetings suggests that the strategy is capturing signals about the Fed's monetary policy activity rather than its supervisory activity.

3.4.2 Predicting returns around FOMC meetings

To delve further into our trading strategy's profitability around FOMC meetings, we construct an event study sample. We define the FOMC event window as event days 1 to +1. This event window is generally 4 days long since scheduled FOMC meetings typically last two days. We use Fed-bank ridership beginning two days after the previous FOMC event window (and ending prior to the event) to predict returns in the event window. Specifically, we compare the average daily ridership in the two weeks prior to the FOMC meeting (FOMC event days –11 through –2), which we label the "Pre-FOMC period," to the average daily ridership in the weeks before, which we label the "benchmark period." The benchmark period begins on the second trading day after the previous FOMC meeting, and it runs through FOMC event day –12 (inclusive) of the FOMC meeting under consideration. We measure the trend in ridership levels prior to the FOMC meeting using a dummy (*Ridership Up*) indicating whether the average daily ridership is higher in the pre-FOMC period than in the benchmark period. We illustrate the timeline of our analysis in Figure 7.[23]

[Insert Figure 7 here]

We consider two measures of market performance during FOMC periods: (i) the cumulative return of the market in excess of the risk-free rate, and (ii) a dummy (*Market Up Dummy*)

[23]Recall that a complication of our analysis is that Goldman Sachs, Citigroup, and Bank of New York Mellon changed the location of their headquarters during our sample period. To ensure that there are an equal number of bank locations on each FOMC event day of a given FOMC meeting, we handle these banks' headquarter changes as follows. For Goldman Sachs, we assign it to its 85 Broad Street and One New York Plaza locations for all FOMC meetings prior to (and including) the FOMC meeting on September 22-23, 2009, and we assign it to its 200 West Street location for all FOMC meetings after (and including) the one on January 26-27, 2010. For Bank of New York Mellon, we assign it to its 1 Wall Street location up to and including the January 27-28, 2015, FOMC meeting, and we assign it to its 225 Liberty Street location for all FOMC meetings after and including the June 16-17, 2015 FOMC meeting. Citigroup's headquarter change differed from the other two in that it already had a significant presence at its new headquarter location (388-390 Greenwich Street). Accordingly, we assign Citigroup to both of its locations up through and including the April 28-29, 2015, FOMC meeting, and we only assign it to its 388-390 Greenwich Street locations for all FOMC meetings thereafter.

indicating whether the cumulative return is positive. In Table 7, we tabulate the number of times that daily ridership rose and fell from the benchmark period to the pre-FOMC period, and how often the market's return exceeded and fell short of the risk-free return. There are 59 FOMC meetings during our sample period, and ridership is higher in the pre-FOMC period than the benchmark period for 41 of the meetings.[24] In these 41 meetings, the market's return is positive (negative) in 17 (24) of them. In the 18 meetings where ridership is lower in the pre-FOMC period than the benchmark period, the market's return is positive (negative) in 16 (2) of them. These counts are reported in Panel A of Table 7.

[Insert Table 7 here]

In Panel B of Table 7, we test for the statistical significance of this relation by running a regression where our unit of observation is FOMC meeting, our dependent variable is the dummy indicating that the market rose during the FOMC meeting, and our independent variable is the dummy indicating that ridership was higher in the pre-FOMC period than in the benchmark period. The coefficient of the ridership dummy is negative (–0.474) and highly significant (t=3.70), revealing that the relationship documented in Panel A of Table 7 is highly significant.[25] In the second column of Panel B, we use the market's excess return (taken from Ken French's website) over the FOMC event window as our dependent variable. From this regression, we see that the average return of the market in the FOMC window is 1.32% when ridership is lower in the pre-FOMC period than in the benchmark period, and –62 bps (= 1.32% − 1.94%) when ridership is higher in the pre-FOMC period than in the benchmark period. This difference in returns (1.94%) is highly statistically significant (t = 3.42). Table 7 was produced using Fed-deflated ridership, but the results are qualitatively similar with NYC-deflated ridership and with no deflator.

In the preceding analysis, we examined returns around FOMC meetings as a function of Fed-bank ridership trends before the FOMC meeting. An alternative approach is to examine the pre-FOMC ridership trends as a function of the market's performance during the FOMC meeting. In Figure 8, we plot the daily ridership levels in FOMC event week separately for the 33 FOMC meetings where the market rises and the 26 meetings where the market falls during the FOMC meeting. Event week 0 is defined as the FOMC meeting period (FOMC

[24] We exclude the January 27-28, 2009, FOMC meeting from this analysis because there is insufficient data to estimate ridership in the benchmark period.

[25] Of course, the coefficient and intercept in the regression can be directly computed from the counts that are reported in Panel A. The intercept (0.889) is the probability of the market rising conditional on ridership falling, which from Panel A equals 16/18. The intercept plus the coefficient of the ridership dummy ($0.889 - 0.474 = 0.415$) represents the probability of the market rising conditional on ridership rising, which from Panel A equals 17/41.

event days −1 to +1), event week −1 is the five trading days before this period (FOMC event days −6 to −2), and the other FOMC event weeks are defined analogously.

[Insert Figure 8 here]

In both samples, daily ridership is higher immediately prior to (and during) the FOMC meetings than it is 3-4 weeks prior to the meeting. However, the increase in ridership is far greater when the market declines during the FOMC meeting than it is when the market rises during the FOMC meeting. In terms of economic significance, daily ridership rises from 16.7 in event week −4 to 20.3 in event week −1 among the FOMC meetings where the market declines, whereas it only rises from 17.3 in event week −4 to 17.9 in event week −1 among the FOMC meetings where the market rises. Importantly, the evidence in Figure 8 suggests that ridership in the two weeks immediately prior to the FOMC meetings, but not ridership in the benchmark period two weeks earlier, predicts FOMC period returns.

4 Anomalous Ridership Activity around the FOMC Cycle

Having established that the market-timing ability of our trading strategy is pronounced around FOMC meetings, we seek a more detailed understanding of patterns of communication themselves. The times at which significant changes in meetings occur and their proximities to FOMC meetings may provide insights into the nature of meetings that the broad aggregates underlying the strategy cannot. We therefore proceed to examine patterns of rides at a finer temporal resolution than in the previous analyses to see if there is any evidence of unusual Fed-bank interactions around FOMC meetings. The significantly lower degree of temporal aggregation implies that significantly less noise will be aggregated away. Under the presumption that the signal decreases with distance from the buildings of interest, we switch to more precise and generally smaller areas around the buildings of interest and utilize the block-based ride classification measure that we developed in Section 2.3.2. Anomalies are inferred from an estimation over the full sample period; thus, to avoid a structural break in geography, we employ a static set of locations that omits Goldman Sachs's old headquarters and trading floor. In addition, we restrict our sample to the 2009-2014 period. 2014 saw a large increase in the use of rideshare apps by businesses and the first public release of the taxi data under New York State's Freedom of Information Law. Significant patterns in rides may reflect the movements of small sets of individuals, and substitution by any of them away from taxis due to changes in business practice or due to concerns about privacy could result in a fatal loss of signal (Rao 2015,Saitto 2014, White 2015,Whong 2014).

The period around FOMC meetings that we examine is motivated by Federal Reserve policy and evidence of leakage. During a blackout period around these meetings (Board of

Governors 2005, Federal Reserve Bank of New York 2010b), Federal Reserve staff are not permitted to discuss with the public monetary-policy and economic matters that have not already been cleared and widely disseminated. Over the sample employed in this paper, the blackout began a week before the FOMC meeting's first day – typically eight days before the announcement – and ended a day after the announcement of policy decisions (Board of Governors 2014).[26]

Despite the rules, FOMC information has reached the public before its official release. In the late 1980s and early 1990s, for instance, FOMC decisions were only to be released after a substantial lag, but the Wall Street Journal frequently wrote about them within a week of the meetings (Cieślak, Morse and Vissing-Jørgensen 2019). More recently, Medley Global Advisors circulated information on FOMC deliberations prior to their authorised release (Appelbaum 2017). Cieślak, Morse and Vissing-Jørgensen (2019) use patterns in futures around monetary-policy meetings and narrative evidence to argue that market-moving information systematically flows from Federal Reserve insiders to outside parties through informal communication. Bernile, Hu and Tang (2016) present market data that they interpret as informed trading in the minutes prior to FOMC surprises. The possibility of systematic leakage is controversial, and Lucca and Moench (2015) present counter-arguments.

Information can flow even accidentally whenever Fed insiders and commercial bankers interact, and evidence of leakage need not suggest any improper intent by anyone associated with the Federal Reserve. Laurence Fink, the CEO of BlackRock, is reported as stating that insights into the Federal Reserve could be obtained even from questions that the Fed posed to members of the private sector (Pulliam 2011). Then-Governor Yellen noted during the 2-3 November 2010 FOMC meeting that external parties could induce problematic disclosures (Federal Open Market Committee 2010).[27] Indeed, the Richmond Fed president resigned in April 2017 over inappropriate communication with an analyst at Medley Global Advisors (Appelbaum 2017). Despite such risks, some Fed employees might find it challenging to impose strict limitations on the volume of their interactions with external parties lest they should harm relationships that are important both for the conduct of Fed business and for those employees' external employment options (Zingales 2013 and Lucca, Seru and Trebbi 2014 for the intuition).

The key identifying assumption in this section is that FOMC meetings in Washington,

[26] 2017 revisions to Federal Reserve policies further restricted permissible communication between Federal Reserve insiders and outsider parties around FOMC meetings (Board of Governors 2017).

[27] "We're obliged to maintain the confidentiality of FOMC information – period, full stop. And that includes documents that we look at in the FOMC and information on who said what. It's obvious that these guidelines have been breached. I also know from personal experience over the years that it's easy for this to happen – it can happen pretty innocently when an experienced reporter lures one into revealing things that end up crossing the line." (Federal Open Market Committee 2010)

D.C., have a much more significant impact on Manhattan interactions between the New York Fed's and commercial banks' insiders than on Manhattan interactions involving unrelated parties in their neighbourhoods. Even though we cannot infer levels of interaction between New York Fed insiders and commercial-bank insiders around FOMC meetings, we can identify significant changes in rides with significant changes in interaction if this assumption holds. Since every Fed-commercial-banker interaction entails the risk of leakage, evidence of an increase in the volume of interactions is evidence of an increased probability of leakage.

4.1 Summary statistics over intraday windows

Figure 9 presents summary statistics for aggregate Manhattan taxi activity–the mean of pick-up and drop-off volumes for intra-Manhattan rides–and for rides between the New York Fed and the major commercial banks. For consistency with later work, we employ overlapping two-hour windows that begin at the top of an hour. Since aggregate Manhattan activity reaches a nadir between 4:00 and 5:00, and rides very early in the morning are presumably more likely to reflect long days at work than early starts, our activity-based definition of a day is the span from 5:00 through 4:59 (Panel A).[28] Activity at 1:30 on a given date, for example, is mapped to the previous date. Trips between the New York Fed and the major commercial banks are not frequent, with medians below 1 for almost all hours (Panels B and C). Anecdotally, meetings at financial institutions with external parties are not likely before 8:00, so the relatively high counts of rides from the vicinity of the New York Fed to the vicinities of the major commercial banks in the morning might be largely driven by the commutes of residents of the condominium near the New York Fed or of guests of the hotel across from it. The relatively high volumes in the opposite direction in the evening are also consistent with the capture of commutes. Intraday ride volumes from the FRBNY's vicinity to those of the major commercial banks and *vice versa* are generally positively skewed. As the maximum count over these two ride classes over all hourly windows and days is 6, there are no extreme outliers.

[Insert Figure 9 here]

4.2 Ride variation during intraday and FOMC windows

A natural starting point is the first day of the FOMC communications blackout. During the sample period, this is seven days before the first day of an FOMC meeting and typically eight

[28] Times of the form XX:00:00 and XX:59:59 are rendered as XX:00 and XX:59.

days before the announcement.[29] Guided by the timing of market movements in Cieślak, Morse and Vissing-Jørgensen (2019) and symmetry, we choose as an endpoint the seventh day after the FOMC announcement.

We map a calendar day to an event day by obtaining the offset in calendar days from the nearest FOMC announcement, with a negative integer indicating a day prior to the announcement and a positive integer indicating a day subsequent to the announcement.[30] In the case of the Wednesday FOMC announcements which dominate the sample, the weekdays during the span that we examine are the twelve event days −8, −7, −6, −5, −2, −1, 0, +1, +2, +5, +6 and +7. Tuesday FOMC announcements and the single Thursday FOMC announcement also permit the estimation of changes for event days −4, −3, +3 and +4, but we do not examine those event days individually as the paucity of observations would make their estimates unreliable.[31]

We employ all two- and three-hour spans that begin at the top of an hour so that *i.*) We have windows short enough to be informative about potential activities during them; *ii.*) any cluster of rides around the boundary between two windows will be entirely captured by other windows; and *iii.*) the windows are long enough that the choice of whether to map rides to windows by their pick-up or drop-off times is not of first-order importance. We arbitrarily choose to map by drop-off time.

Given that ride counts are discrete and non-negative, we employ Poisson regressions in their analysis. The prototypical regression for rides during intraday window f employs a conditional mean, or intensity, λ_t^f of the following form:

$$\lambda_t^f = exp\Big(\iota_t\beta_\iota + z_t^f\gamma + \alpha_{ym(t)} + \theta_{wd(t)}\Big) \tag{1}$$

where t indexes the date; ι_t is an indicator of whether date t falls in a specified set of event days (an FOMC window); z_t^f is aggregate Manhattan taxi activity during window f on date t (the mean of all pick-ups and drop-offs of intra-Manhattan rides); $ym(t)$ is the year-month associated with date t; α_{ym} is a fixed effect for year-month ym; $wd(t)$ is the weekday associated with date t, and θ_{wd} is a weekday indicator. FOMC meetings are not randomly distributed over weekdays and months, and the period controls will control for intraweek cyclicality and lower-frequency seasonality and trends. Some FOMC windows might just coincide with unrelated high levels of Manhattan taxi rides, and z_t^f should reduce the danger of serious omitted-variable bias. Camerer et al. (1997) provide evidence that New York City

[29] The blackout begins at midnight, but for convenience we begin at the start of the day as defined in this paper, 5:00.

[30] We define a day as the span 5:00-4:59, and windows beginning at 0:00 through 4:00 on one calendar day are mapped to the previous calendar day and the associated event day.

[31] Section A.2 of the Internet Appendix discusses the handling of sparse data.

taxi drivers make labour-supply decisions daily, so the control for Manhattan activity might control for some supply-side variation and tighten inference.

The QMLE fixed-effects Poisson estimator used in this paper can be consistent for the parameters in λ_t^f even when the ride data do not resemble Poisson draws (Wooldridge 2010).[32,33] We cluster at the year-month level to accomodate arbitrary heteroskedasticity and within-month serial correlation. The cluster-robust covariance matrix is derived from M-estimation. Our null hypothesis for each FOMC window is a non-positive change, and our alternative is a positive change. To minimise the risk of a false rejection of the null, we test against zero change.[34]

4.3 Changes in direct rides

We first examine direct rides between the New York Fed and the major commercial banks. Figure 10 Panel A presents the estimate of the change in rides from the New York Fed to the major commercial banks during each bihourly window of each of the twelve event days.[35] We find some variation that is individually significant at the 1% confidence level, and there is a small cluster of increases that are individually significant at the 5% confidence level during mornings around the end of the blackout, but the lowest p-value in a set of hundreds is only 0.2%. The large and individually significant increase associated with that minimum p-value occurs between 7:00 and 9:00 on event day +5. It is largely offset by subsequent drops and thus might reflect a change in intraday scheduling.

[Insert Figure 10 here]

An identical analysis of rides from the major commercial banks to the New York Fed shows substantially more significant variation, particularly at night (Panel B). Strikingly, the most statistically significant increases in rides occur back to back in the hours after the end of the FOMC blackout. Rides increase by 0.32 with a p-value of less than 0.1% during the

[32] Cameron and Trivedi (2009) provide the estimator specialised to a conditional mean that is exponential in a term that is linear in parameters.

[33] Cameron and Trivedi (2009) note that the consistency depends on strong exogeneity, and it is not obvious that the exclusive use of period controls and contemporaneous aggregate taxi activity as regressors would necessarily lead to any important violation.

[34] We obtain one-sided p-values from pairs bootstrapping with asymptotic refinement (see, for example, Cameron, Gelbach and Miller 2008 and Cameron and Trivedi 2009). Each bootstrap simulation entails 72 year-month draws with replacement. We employ 1×10^4 simulations when assessing significance up to the 1% confidence level and 1×10^5 simulations when assessing significance up to the 0.1% confidence level. Section A.2 of the Internet Appendix provides additional technical details.

[35] The Poisson regressions yield percentage changes. To obtain a change over an FOMC window in units of rides, we calculate the mean partial effect over those event days, *i.e.* the mean estimated difference in rides of having the indicator on relative to the counterfactual of having it off for the calendar days in the sample mapped to those event days.

1:00-2:59 window and by 0.28 with a p-value of less than 0.1% during the 2:00-3:59 window. These are also the fourth- and sixth-largest changes in ride count despite their occurrence during typically low-volume intraday windows. The increase between 3:00 and 5:00 is also individually significant at the 1% confidence level but is less statistically and economically significant with a magnitude of 0.12 and a p-value of 0.4%. The small and insignificant change between 4:00 and 6:00 suggests that the increase is concentrated between 1:00 and 4:00.

A formal data-mining exercise bolsters the salience and significance of the post-blackout increase in rides from the major commercial banks to the New York Fed. The questions that we seek to answer are *i.*) what the most significant intraday-window-event-day pair is; and *ii.*) whether the associated change is statistically significant conditional on the examination of the changes over the full set of pairs. Our determination of the most significant change balances statistical and economic significance. We separately rank each coefficient by its t-statistic and by the associated change in units of rides and then pick the pair of windows that minimise the mean of the two ranks. In light of the individually large and significant back-to-back increases, we add to the set of intraday windows all three-hour windows that start at the top of an hour. Significant increases might also occur during other windows, so focusing on a single pair yields a lower bound on ride growth.

To assess significance conditional on data-mining, we employ a Monte-Carlo variant of the Romano and Wolf (2005) one-sided StepM approach with studentisation. This approach seeks to maximise the number of rejected null hypotheses for a given type-I error rate and accounts for the correlations among coefficients across different specifications. We test for significance at the 0.1%, 1%, 5%, and 10% confidence levels.[36]

We find that the most significant pair of windows is the span from 1:00 through 3:59 after the lifting of the FOMC blackout, and the associated increase is significant at the 1% confidence level even when we account for the data mining over 572 models. Rides increase by 118.8%, and the average partial effect is 0.43 rides (Table 8 Column 1). This increase is specific to the immediate vicinity of the New York Fed's headquarters and not its neighbourhood. For the same pair of windows, we examine the variation in drop-offs that are mapped to the full set of blocks bordering the New York Fed's blocks and are more than 100 feet from New York Fed's headquarters and its annex across Maiden Lane.[37] Rides over this much larger area increase by only 0.23 (+6.5%) during the 1:00-3:59 window on event day +1, and the variation is far from significant at any conventional level (Column 2).

[36]Details on our implementation are presented in Section A.3 of the Internet Appendix.

[37]The block across Nassau Street from 33 Liberty Street is only a small sliver, so we also include the block between that sliver and Broadway.

To obtain a more complete picture of what underlies these results, we examine the volume of rides unexplained by year-month, weekday and Manhattan taxi-activity controls. Figure 11 Panel A presents for each of 2009 through 2014 that year's mean post-blackout ride residual for the 1:00-3:59 window.[38] The mean residual is generally non-trivially positive and never non-trivially negative, but the change in rides estimated over the full sample is clearly driven to a large degree by 2012. In 2012, there were on average roughly 1.18 unexplained rides to the New York Fed in the hours after the end of the FOMC blackout period. 2012 saw the Fed's adoption of an explicit inflation target and its initiation of a third round of quantitative easing, QE3. We observe only one ride in the wake of the blackout period around the QE3 announcement, but the second-highest volume over all 1:00-3:59 windows, 4 rides, occurred after the subsequent blackout period, and the third-highest, 3 rides, after the previous. Neither discreet information-seeking nor clarification of Fed communication is inconceivable.

[Insert Figure 11 here]

Late-night meetings might be scheduled to gather information about bond-market conditions. While the FOMC decides the scale of asset purchases, the distribution of purchases across maturities is the responsibility of the New York Fed. Bond markets are more opaque than stock markets, and the New York Fed might gather information about the supply of and demand for debt securities. Primary dealers are expected to provide market commentary to the New York Fed, and the New York Fed might seek information pertinent to the implementation of just-announced policy decisions or to mooted plans as soon as it is permitted (Federal Reserve Bank of New York 2010a). All of the major commercial banks in the sample are also primary dealers except for BNY Mellon, and the primary dealers drive the results.

4.4 Changes in coincidental drop-offs

Not all meetings between FRBNY insiders and commercial bankers need to occur at the New York Fed or at a major bank. People can, for example, chat informally at a park, in a café or over lunch at a restaurant. We employ coincidental drop-offs away from the New York Fed and the major commercial banks as noisy indicators of such informal meetings. "Coincidental drop-offs" will henceforth imply coincidental drop-offs away from the New York Fed and the major commercial banks' buildings.

[38] We perform a Poisson regression of 1:00-3:59 ride volumes only on year-month, weekday and Manhattan activity controls and then subtract the fit levels from those observed.

For drop-offs to be considered coincidences, they must satisfy three spatial criteria and one temporal criterion. Since coincidences are used to capture meetings away from the financial institutions, the first spatial criterion is that neither ride's drop-off be mapped to the New York Fed or any of the major commercial banks. The second is that the drop-offs be mapped to the same census block. With roads in New York City largely following an approximately NS-EW grid arrangement, the final spatial criterion is that the drop-offs be within a certain distance of each other along the NS axis and along the EW axis. The baseline distance is a quarter of a typical Midtown Manhattan block along its shorter edge, 66 feet.[39] This makes some allowance for GPS noise and for dispersion in where individuals are dropped off. The baseline temporal criterion is that the drop-offs be no more than 10 minutes apart. Counts only include coincidences for which both drop-offs are in the specified intraday window, precluding, for example, a drop-off at 11:06 only coincidental with a drop-off at 10:58 from contributing to the count for the 11:00-12:59 window.[40] We ignore coincidental drop-offs at Pennsylvania Station and Grand Central Terminal as well as their adjacent blocks due to an expectation that coincidences around those transit hubs rarely reflect meetings. Coincidences are not common, with the mean over a bihourly window never exceeding 0.24. The maximum over a bihourly window is 3, so there are no extreme outliers (Figure 9 Panel D).

We again begin by examining variation during overlapping bihourly windows over the twelve event days around an FOMC meeting (Figure 10 Panel C). The 2:00-3:59 through 4:00-5:59 windows are omitted due to a paucity of observations. Coincidental drop-offs around noon are consistently and strikingly elevated from the day before the FOMC announcement onward. That starting point typically corresponds to the first day of an FOMC meeting. Moreoever, changes within a day of the FOMC announcement are statistically significant: an increase of 0.19 the prior day between 11:00 and 13:00 is individually significant at the 5% confidence level; an increase of 0.20 the day after the announcement between 11:00 and 13:00 is individually significant at the 5% confidence level, and an increase of 0.25 the day after the announcement between 12:00 and 14:00 is individually significant at the 1% confidence level. A third individually significant lunchtime change is an increase of 0.14 on event day +5 between 12:00 and 14:00. While the individual p-values for the noon increases are not very

[39] There are approximately 20 blocks to a mile along a Manhattan avenue, so a quarter of a block is approximately 66 feet (Pollak 2006).

[40] One complication is that a single New York Fed drop-off might be coincidental with more than one drop-off from the major commercial banks or *vice versa*. Since the interest is in meetings and not individuals, we calculate the number of coincidental drop-offs over an intraday window as the minimum of *i.*) the number of rides originating at the New York Fed with drop-offs that are coincidental with rides originating at the major commercial banks; and *ii.*) the number of rides originating at the major commercial banks with drop-offs that are coincidental with rides originating at the New York Fed.

low, particularly given the number of changes examined, a potential increase in lunchtime meetings during a period for which Cieślak, Morse and Vissing-Jørgensen (2019) provide evidence of informal Federal Reserve communication motivates further examination.

For a more rigorous sense of where coincidences are elevated, we proceed to undertake a similar data-mining exercise to that for the direct rides. Given the evidence of an increase in coincidental drop-offs over lunch across multiple event days, we expand the candidate set to include all contiguous spans of length 1 through 16 calendar days between event day −8 and +7. We again examine two- and three-hour intraday spans but we only employ those between 9:00 and 17:00 in order to exclude a large share of the coincidences that reflect commuting.

The data-mining exercise yields that the most significant windows are 10:00-12:59 from event days −2 through +7. We find an extra 1.41 coincidences over this span, an increase of 49.8%, and that the increase is significant at the 1% confidence level when data mining over the 1,690 specifications is accounted for (Table 8 Column 3). Also significant at the 5% confidence level is the increase during a subset of this span: 11:00-12:59 from event day −1 through event day +7. The increase over this subset is 1.10 extra coincidences (+67.4%) (Column 4). Given the evidence of a significant increase in lunchtime rides and the Cieślak, Morse and Vissing-Jørgensen (2019) evidence of informal information flow, we focus on this narrow pair of windows for the rest of this analysis. The increased level of coincidences is not a general phenomenon in the neighbourhood of the New York Fed. With pick-ups at blocks adjacent to the New York Fed's buildings but more than 100 feet from them, the change is 0.24 coincidences, 2.4%, and is far from statistically significant even before data mining is addressed (Column 5).

Further supporting an identification of the coincidences with planned meetings over lunch are the spatial distribution of drop-offs and the degree to which our results are driven by coincidences of cabs with single passengers. Financial firms are concentrated in Midtown and the Financial District, and the omission both of these areas where drop-offs might be related to official duties and of hospitals where clusters are also observed yields an increase of 0.71 rides (+96.1%) and significance at the 5% confidence level when we account for data mining (Column 6). The lunchtime coincidences are largely concentrated in areas associated with dining and shopping like the Meatpacking District and TriBeCa (Figure 12). Where cabs contain multiple passengers, coincidental drop-offs might simply reflect common tastes in lunch destinations rather than planned meetings. Fed staff might conceivably go for lunch with each other once their FOMC-related duties end. We find, though, that the growth in coincidences is primarily attributable to the coincidental drop-offs of single passengers. They increase by 0.85 coincidences (+108.8%), and the change is significant at the 1% confidence

level conditional on data mining (Column 7).

[Insert Figure 12 here]

As with post-blackout rides to the New York Fed, 2012 contributes importantly to the estimated increase in lunchtime coincidental drop-offs. Figure 11 Panel B plots the mean per-FOMC-meeting count of unexplained lunchtime coincidences during for each year from 2009 through 2014.[41] The mean volume of unexplained coincidences per FOMC meeting is reasonably large and positive from 2010 through 2013 and is only negative, but trivially so, in 2009. While scarcely exceeding 1 extra ride in any other year, it reaches 1.85 in 2012. The synchronicity of high residuals suggests information-seeking linked to changes in monetary policy such as the adoption of a formal inflation target or the third round of quantitative easing, but nothing can be concluded about the cause without collateral data.

Meetings could be linked to Fed information-gathering regarding bond-market conditions pertinent to the implementation of monetary policy. Alternatively, Federal Reserve staff who broadly restrict their interactions with outside parties during the blackout might address pent-up demand by scheduling an above-average volume of meetings in the days after its end. Whatever the intent might be, sensitive information could flow even accidentally. Although the blackout might reduce the risk of leakage of preparations for FOMC meetings, it might also increase opportunities for outside parties to gain insights into their still-confidential proceedings and support material over the following days.

5 Concluding remarks

Motivated by prior studies documenting the importance of the monetary decisions of the Federal Reserve Board on aggregate stock market returns, we investigated the predictability of future market returns based on past levels of ridership between the New York Fed and the major, systemically important financial institutions headquartered in New York City. We documented a strong negative association between past ridership levels and future returns. A simple strategy that invests in the market portfolio (a risk-free asset) when past ridership levels have been low (high) earns a Sharpe ratio that is 50% larger than the market portfolio. The abnormal performance of this strategy is especially pronounced around FOMC meetings, when the Fed's monetary policy decisions are made. Overall, our evidence suggests that the Fed engages in more face-to-face interactions with systemically important financial institutions when it has negative private information about the economy. The fact that our results are strongest around FOMC meetings suggests that the private information is related to monetary policy, and this information gets released to the public around FOMC meetings.

[41] The specification is that of Table 8 Column 4.

Cieślak, Morse and Vissing-Jørgensen (2019) hypothesize that markets learn about the Federal Reserve around monetary-policy meetings through private, informal channels. We thus examine anomalous interactions between Fed and bank employees, including late-night visits and offsite lunch meetings. We find evidence suggestive of an increase in both types of meetings around FOMC meetings, suggesting that these meetings might be a subset of those channels.

References

Appelbaum, Binyamin. 2017. "Richmond Fed President Resigns, Admitting He Violated Confidentiality". *The New York Times* April 4, 2017.

Bernanke, Ben, and Kenneth Kuttner. 2005. "What Explains the Stock Market's Reaction to Federal Reserve Policy?" *Journal of Finance* 60 (3): 1221–1257.

Bernile, Gennaro, Jianfeng Hu and Yuehua Tang. 2016. "Can information be locked up? Informed trading ahead of macro-news announcements". *Journal of Financial Economics* 121 (3): 496–520.

Board of Governors. 2014. "Federal Open Market Committee Rules and Authorizations". https://www.federalreserve.gov/monetarypolicy/files/FOMC_RulesAuthPamphlet_201401.pdf.

— . 2017. "Federal Open Market Committee Rules and Authorizations". https://www.federalreserve.gov/monetarypolicy/files/FOMC_RulesAuthPamphlet_201701.pdf.

— . 2005. *The Federal Reserve System: Purposes & Functions*. Washington, D.C.: Board of Governors of the Federal Reserve System.

Bomfim, Antulio N. 2003. "Pre-announcement effects, news effects, and volatility: Monetary policy and the stock market". *Journal of Banking and Finance* 27 (1): 133–151.

Brodeur, Abel, and Kerry Nield. 2017. "Has Uber Made it Easier to Get a Ride in the Rain?" *Working paper*.

Brown, Jeffrey R., and Jiekun Huang. 2017. "All the President's Friends: Political Access and Firm Value". *NBER Working Paper No. 23356*.

Brusa, Francesca, Pavel Savor and Mungo Wilson. 2018. "One Central Bank to Rule Them All". *Working paper*.

Camerer, Colin, et al. 1997. "Labor Supply of New York City Cabdrivers: One Day at a Time". *Quarterly Journal of Economics*: 407–441.

Cameron, A. Colin, Jonah B. Gelbach and Douglas L. Miller. 2008. "Bootstrap-Based Improvements for Inference with Clustered Errors". *The Review of Economics and Statistics* 90 (3): 414–427.

Cameron, A. Colin, and Pravin K. Trivedi. 2009. *Microeconometrics: Methods and Applications*. Cambridge University Press.

Carney, John. 2009. "Goldman Sachs Kills The Express Elevator To The Equity Trading Floor At One New York Plaza". *Business Insider* December 23, 2009.

Center for Cryptologic History. 2006. *The Friedman Legacy: A Tribute to William and Elizebeth Friedman*. Sources in Cryptologic History 3. National Security Agency.

Chen, Shiu-Sheng. 2007. "Does Monetary Policy Have Asymmetric Effects on Stock Returns?" *Journal of Money, Credit and Banking* 39 (2-3): 667–688.

Cieślak, Anna, Adair Morse and Annette Vissing-Jørgensen. 2019. "Stock Returns Over the FOMC Cycle". *Journal of Finance (forthcoming)*.

Cieślak, Anna, and Annette Vissing-Jørgensen. 2018. "The Economics of the Fed Put". *Working paper*.

Cohen, Peter, et al. 2016. "Using Big Data to Estimate Consumer Surplus: The Case of Uber". *NBER Working Paper No. 22627*.

Craig, Susanne. 2010. "Goldman Sachs's New Palace Creates Princes, Serfs". *The Wall Street Journal* April 16, 2010.

Cramer, Judd, and Alan B. Krueger. 2016. "Disruptive Change in the Taxi Business: The Case of Uber". *American Economic Review* 106 (5): 177–182.

Farber, Henry S. 2015. "Why you Can't Find a Taxi in the Rain and Other Labor Supply Lessons from Cab Drivers". *The Quarterly Journal of Economics* 130 (4): 1975–2026.

Federal Open Market Committee. 2010. *Meeting of the Federal Open Market Committee on November 2-3, 2010*. https://www.federalreserve.gov/monetarypolicy/files/FOMC20101103meeting.pdf.

Federal Reserve Bank of New York. n.d. "About the Building". https://www.newyorkfed.org/aboutthefed/about_the_building.html.

— . 2010a. *OPERATING POLICY Administration of Relationships with Primary Dealers*. https://www.newyorkfed.org/markets/pridealers_policies_100111.html.

— . 2010b. *Statement Regarding Purchases of Treasury Securities*. https://www.newyorkfed.org/markets/opolicy/operating_policy_101103.html.

Financial Stability Board. 2014. "2014 update of list of global systemically important banks (G-SIBs)". http://www.fsb.org/wp-content/uploads/r_141106b.pdf.

Gilbert, Thomas, Alexander Kurov and Marketa Wolfe. 2018. "The Disappearing Pre-FOMC Announcement Drift". *Working paper*.

Jackson, Osborne. 2018. "The Supply Side of Discrimination: Evidence from the Labor Supply of Boston Taxi Drivers". *Working paper*.

Jegadeesh, Narasimhan, and Di Wu. 2017. "Deciphering Fedspeak: The Information Content of FOMC Meetings". *Working paper*.

Koudijs, Peter. 2016. "The Boats That Did Not Sail: Asset Price Volatility in a Natural Experiment". *The Journal of Finance* 71 (3): 1185–1226.

Lee, Lian Fen, Michelle Lowry and Susan Shu. 2018. "Are all perks solely perks? Evidence from corporate jets". *Journal of Corporate Finance* 48:460–473.

Lucca, David, Amit Seru and Francesco Trebbi. 2014. "The Revolving Door and Worker Flows in Banking Regulation". *Journal of Monetary Economics* 65:17–32.

Lucca, David O., and Emanuel Moench. 2015. "The Pre-FOMC Announcement Drift". *The Journal of Finance* 70 (1): 329–371.

Mammen, Kristin, and Hyoung Suk Shim. 2018. "New York City Taxis in an Uber World". *Journal of Marketing Development and Competitiveness* 12 (2): 10–25.

Neuhierl, Andreas, and Michael Weber. 2018a. "Monetary Momentum". *NBER Working paper*.

— . 2018b. "Monetary Policy Slope and the Stock Market". *Working paper*.

Nolte, William. 1997. "Interview with "Ski" [REDACTED]". *Cryptolog* 23 (2): 1–9.

— . 1996. "Louis W. Tordella and the Making of NSA". *Cryptolog* 22 (1): iii–iv.

NYC Taxi & Limousine Commission. 2017. *TLC Trip Record Data*. http://www.nyc.gov/html/tlc/html/about/trip_record_data.shtml.

Ozdagli, Ali K., and Mihail Velikov. 2018. "Show Me the Money: The Monetary Policy Risk Premium". *Journal of Financial Economics (forthcoming)*.

Paletta, Damian, and Aaron Lucchetti. 2010. "Law Remakes U.S. Financial Landscape". *The Wall Street Journal* July 16, 2010.

Pollak, Michael. 2006. "Knowing the Distance". *The New York Times* September 17, 2006.

Pulliam, Susan. 2011. "Investors Bullish on Fed Tips". *The Wall Street Journal* November 23, 2011.

Rao, Leena. 2015. "Uber has signed up thousands of business for rides". *Fortune* September 2, 2015.

Romano, Joseph P., Azeem M. Shaikh and Michael Wolf. 2008. "Formalized Data Snooping Based on Generalized Error Rates". *Econometric Theory* 24:404–447.

Romano, Joseph P., and Michael Wolf. 2005. "Stepwise Multiple Testing as Formalized Data Snooping". *Econometrica* 73 (4): 1237–1282.

Saitto, Serena. 2014. "Uber to Debut Car-Booking Services for Business Travelers". *Bloomberg* July 28, 2014.

Thorbecke, Willem. 1997. "On stock market returns and monetary policy". *Journal of Finance* 52 (2): 635–654.

U.S. Congress. n.d. *Actions Overview: H.R.4173 — 111th Congress (2009-2010)*. https://www.congress.gov/bill/111th-congress/house-bill/4173/actions.

White, Martha C. 2015. "Companies Adapt to Uber and Lyft Rides on the Expense Report". *The New York Times* March 23, 2015.

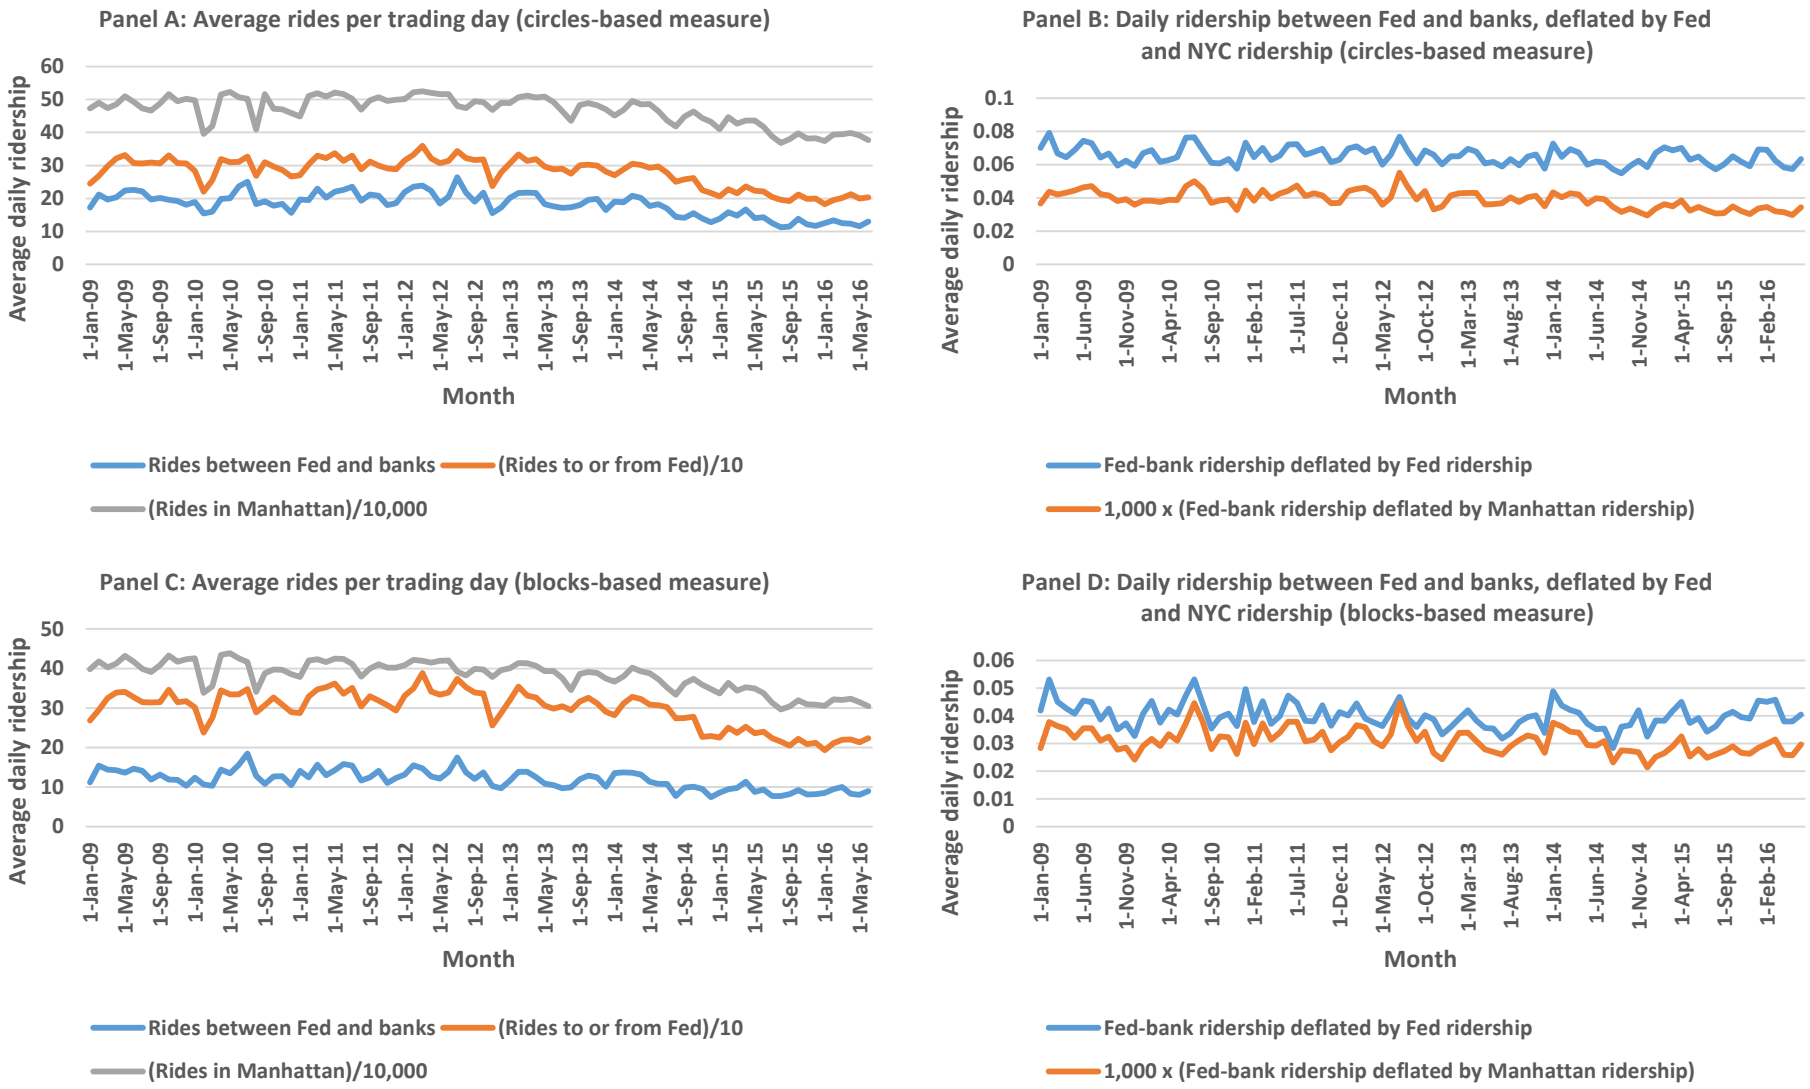
Whong, Chris. 2014. *FOILing NYCs Taxi Trip Data*. https://web.archive.org/web/20140618074624/http://chriswhong.com/open-data/foil_nyc_taxi/.

Wooldridge, Jeffrey M. 2010. *Econometric Analysis of Cross Section and Panel Data*. 2nd ed. The MIT Press.

Yermack, David. 2014. "Tailspotting: Identifying and profiting from CEO vacation trips". *Journal of Financial Economics* 113 (2): 252–269.

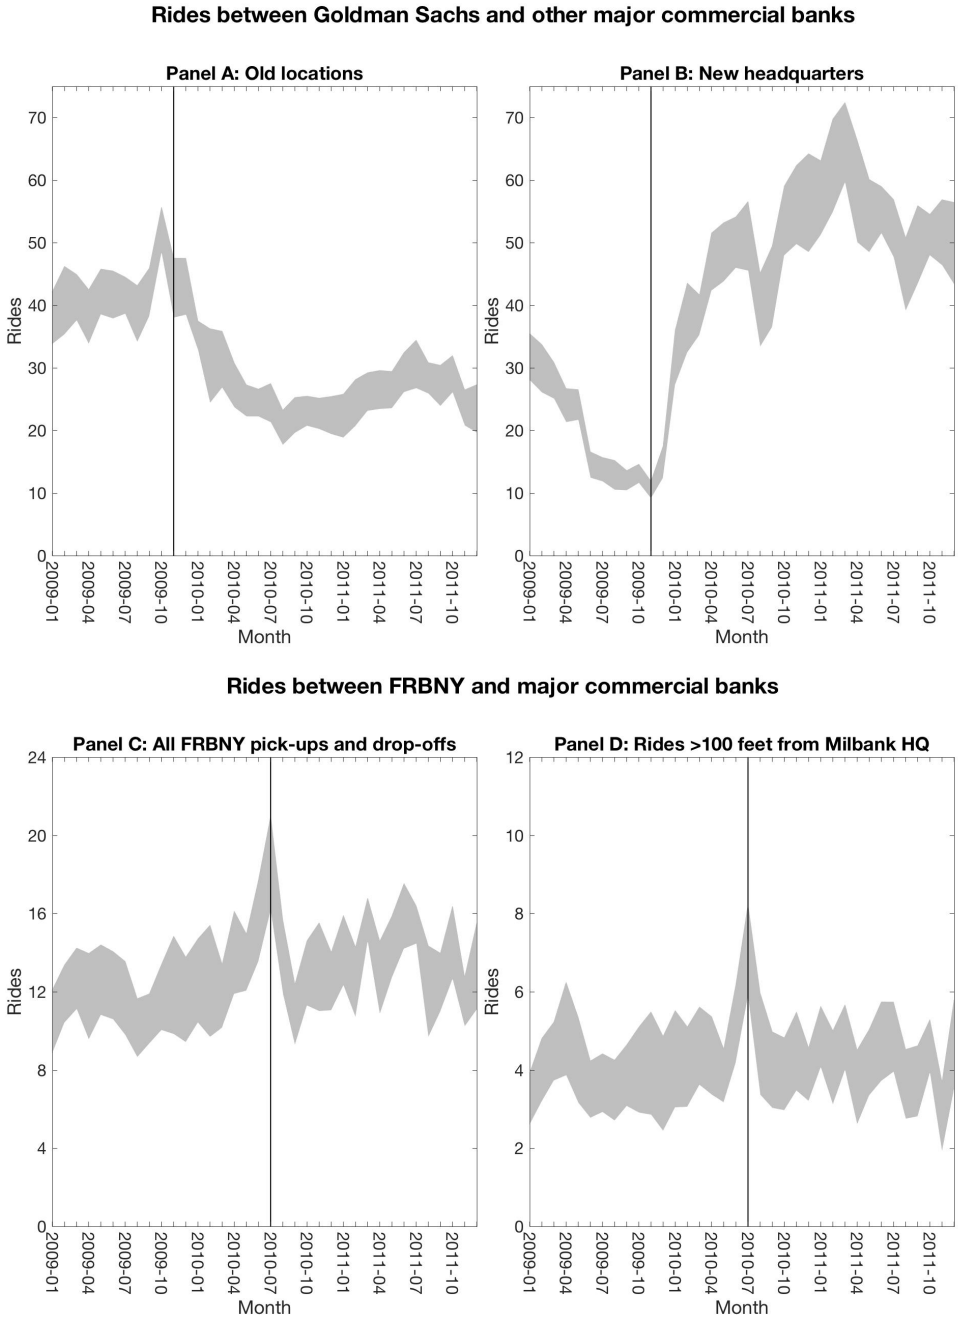
Zingales, Luigi. 2013. "Preventing Economists' Capture". Booth Working Paper No. 13-81.

Figure 1: Average Daily Ridership on Trading Days, by Month



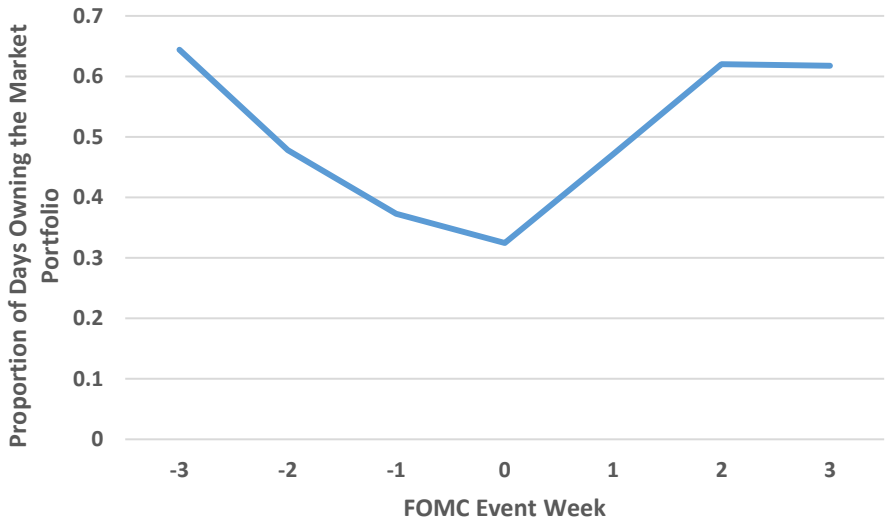
Notes: We report the average daily cab ridership per trading day, by month. In Panels A and C, we report the ridership (i) between the Fed and the systemically important banks in our sample, (ii) between the Fed and any location in NYC (divided by 10), and (iii) across all NYC (divided by 10,000). In Panels B and D, we deflate (i.e., scale) the average daily cab ridership per trading day between the Fed and the systemically important banks in our sample. We consider two deflators: all rides to or from the Fed, and all rides in NYC. We multiply the NYC-deflated average daily ridership by 1,000 so that its units are similar in magnitude to the Fed-deflated values. In Panels A and B, we use the circles-based ride assignment methodology developed in Section 2.3.1. In Panels C and D, we report ridership using the blocks-based ride assignment methodology developed in Section 2.3.2. Taxi cab ridership comes from the New York City Taxi and Limousine Commission.

Figure 2: Two-deviation confidence bands around monthly means of weekday ride volumes between major financial institutions



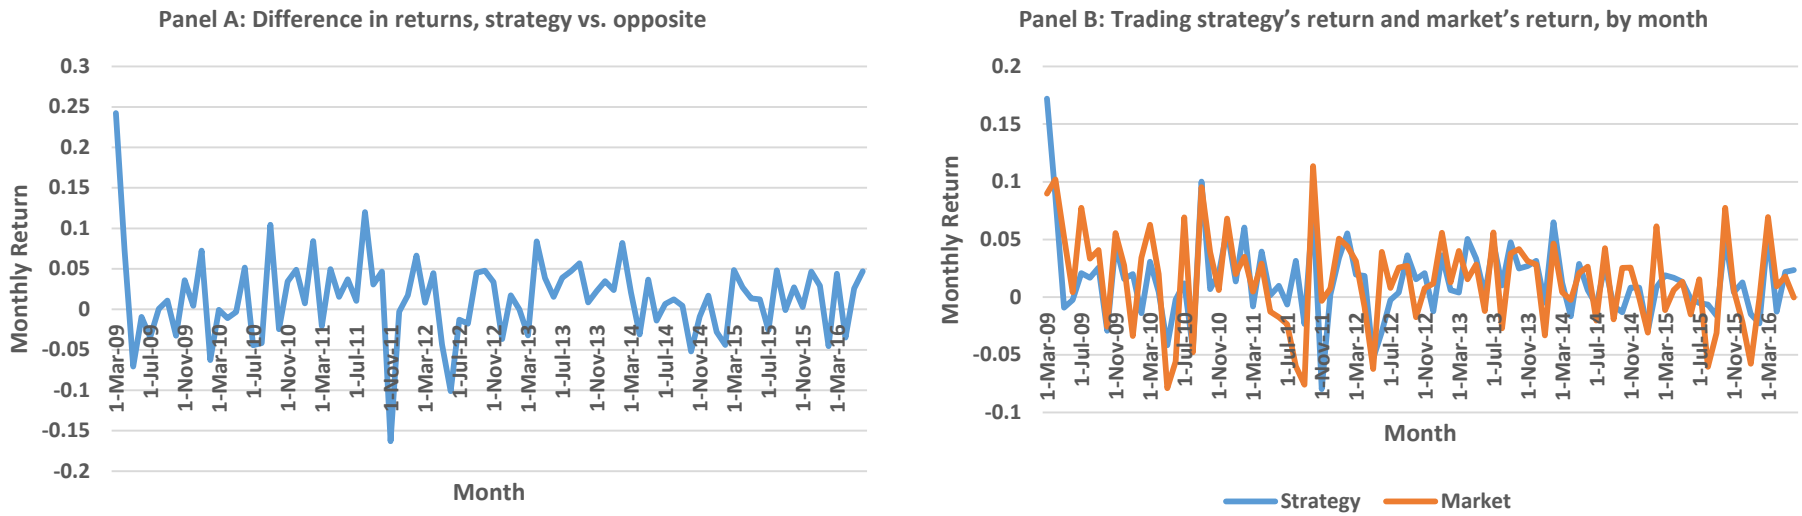
Notes: The vertical bars in Panels A and B indicate the commencement of Goldman Sachs's staff relocation by November 2009. The vertical bars in Panels C and D mark the signing of Dodd-Frank into law in July 2010. A day spans the period from 5:00 to 4:59 the next morning.

Figure 3: Trading Strategy Holding by FOMC Event Week



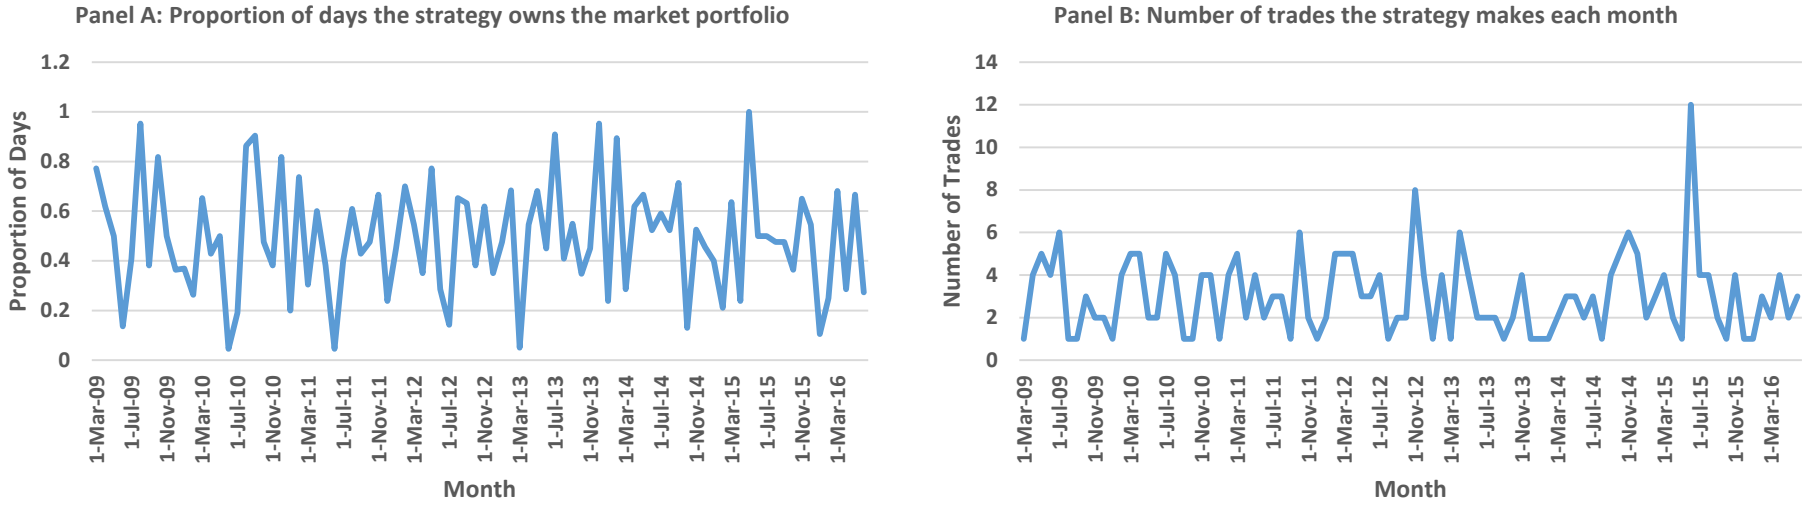
Notes: We plot the average proportion of days that the trading strategy described in Section 3.1 holds the market portfolio (as opposed to a risk-free security) across the FOMC cycle.

Figure 4: Time Series of Trading Strategy's Monthly Profits



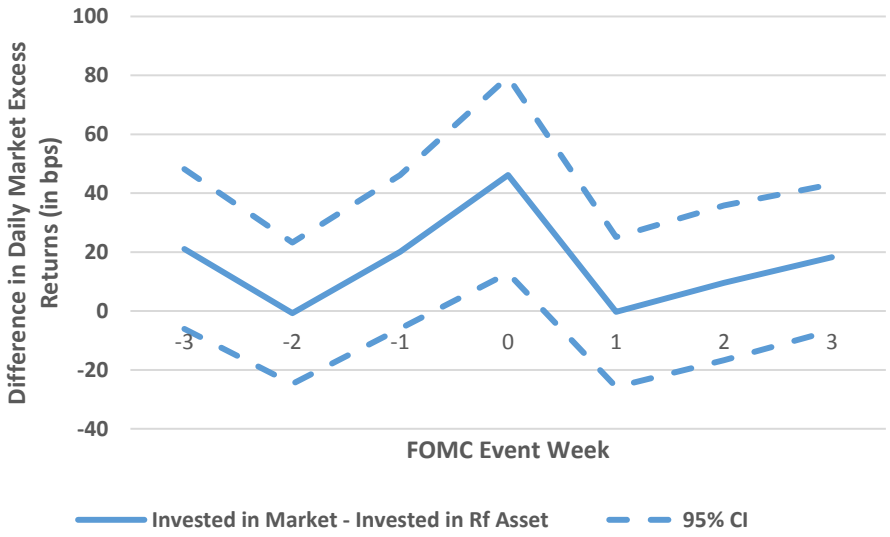
Notes: We plot the time series of monthly profits from the trading strategy described in Section 3.1. Each day, the trading strategy either holds the market portfolio or a risk-free asset. It holds the market portfolio if the ridership between the Fed and the systemically important banks in our sample has been low in the previous 10 trading days relative to the 20 trading days prior; otherwise, it holds the risk-free asset. In Panel A, we report the difference between the cumulative excess return of the market (in excess of the risk-free rate) on the days on which the strategy owns the market and the days on which the strategy owns a risk-free asset. This represents the return differential between our strategy and an "opposite" strategy that holds the market portfolio when our strategy owns a risk free asset, and holds a risk free asset when our strategy owns the market portfolio. Taxi cab ridership comes from the New York City Taxi and Limousine Commission. Our daily excess returns are taken from Ken French's website. In Panel B, we report the monthly excess returns of our trading strategy and the market portfolio.

Figure 5: Holdings and Trading Volume of Strategy



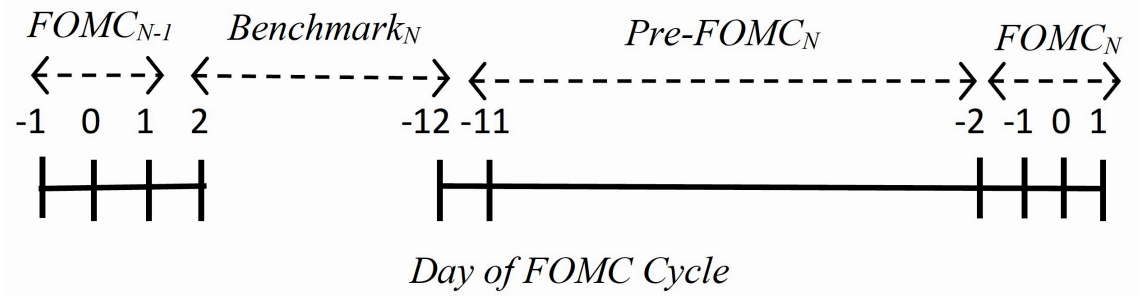
Notes: We plot the proportion of days the trading strategy owns the market portfolio (Panel A) and the number of trades the strategy makes (Panel B) by month. Each day, the trading strategy either holds the market portfolio or a risk-free asset. It holds the market portfolio if the ridership between the Fed and the systemically important banks in our sample has been low in the previous 10 trading days relative to the 20 trading days prior; otherwise, it holds the risk-free asset. The trading strategy is described in further detail in Section 3.1. Taxi cab ridership comes from the New York City Taxi and Limousine Commission.

Figure 6: Market Timing Ability in FOMC Event Time



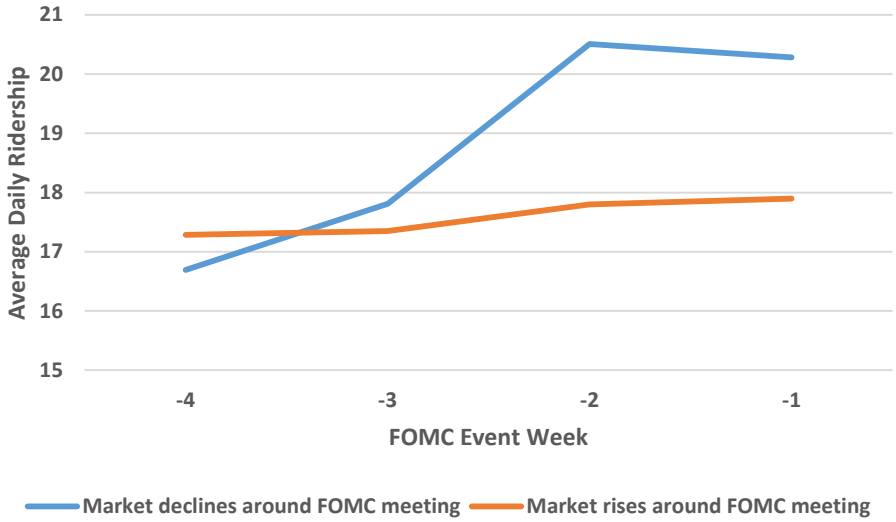
Notes: We plot the average difference in average daily return of the market (in basis points) on days that our trading strategy owns the market portfolio and days that it owns a risk-free asset, by FOMC event week. The 95% confidence interval is plotted in dashed lines. FOMC event week 0 corresponds to the week of an FOMC meeting. Taxi cab ridership comes from the New York City Taxi and Limousine Commission. Our daily excess returns are taken from Ken French's website.

Figure 7: Illustration of FOMC event time analysis



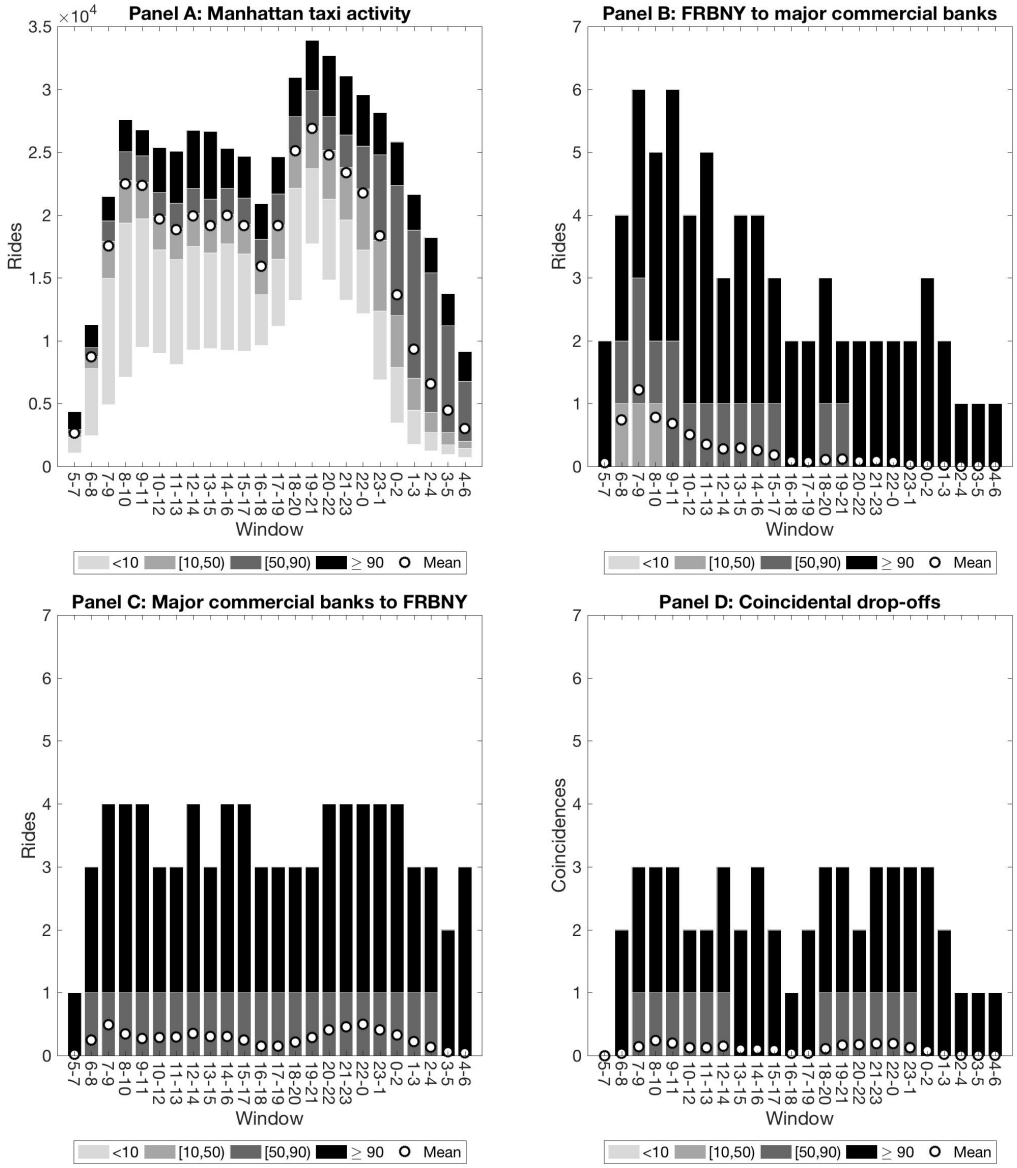
Notes: This figure plots the relevant time periods for our FOMC event time analysis. We estimate returns around FOMC meetings as the cumulative excess return from FOMC event day −1 through FOMC event day +1, where the sample is restricted to trading days. The "benchmark period" for a given pre-FOMC period consists of the second trading day after the previous FOMC meeting through FOMC event day −12 of the current FOMC meeting. The pre-FOMC period consists of the two weeks before the present FOMC meeting window, namely, FOMC event days −11 through −2 (inclusive). We compare the average daily ridership in the benchmark period to the pre-FOMC period, and we examine the predictability of returns around FOMC meetings based on these ridership levels.

Figure 8: Daily ridership in FOMC event time as a function of future FOMC returns



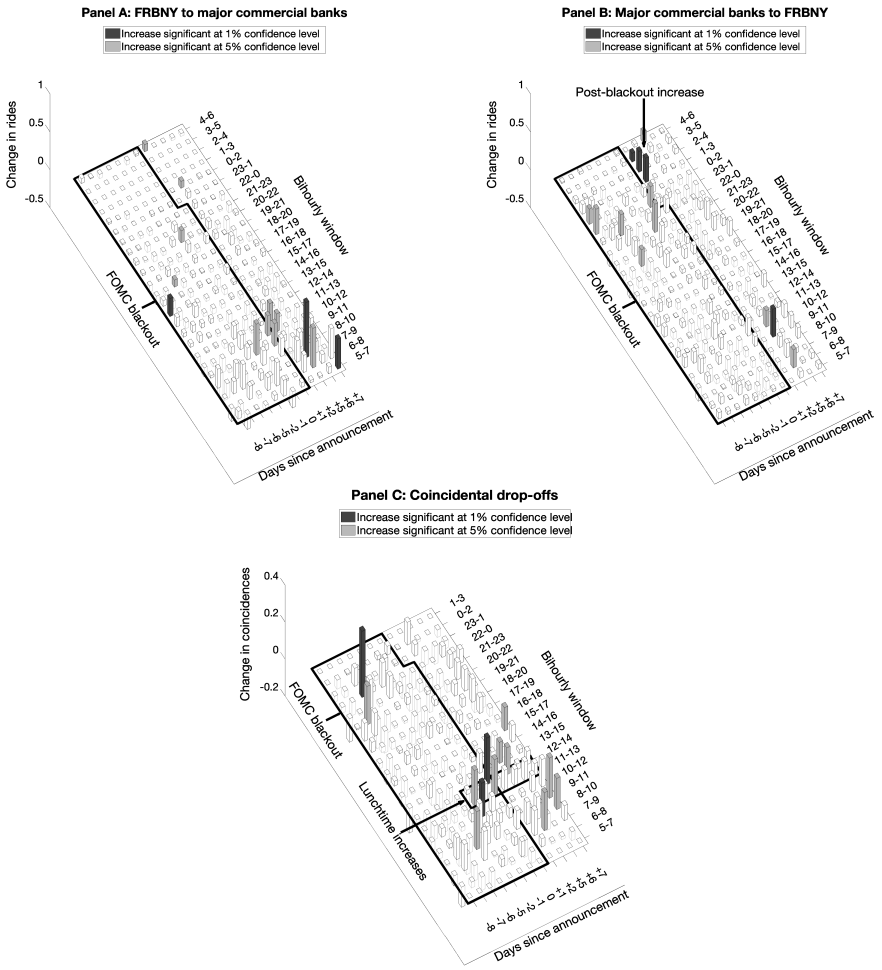
Notes: We plot average daily ridership between the Fed and the banks in our sample in FOMC event time. We plot this separately for FOMC meetings in which the market had positive and negative excess returns during the FOMC window (event days -1 to $+1$). FOMC event week 0 corresponds to FOMC event days 1 to $+1$. FOMC event week 1 corresponds to FOMC event days 6 through 2 (inclusive), and the other weeks are defined analogously. Taxi cab ridership comes from the New York City Taxi and Limousine Commission. Our daily excess returns are taken from Ken French's website.

Figure 9: Means and quantiles of counts of weekday rides and coincidental drop-offs



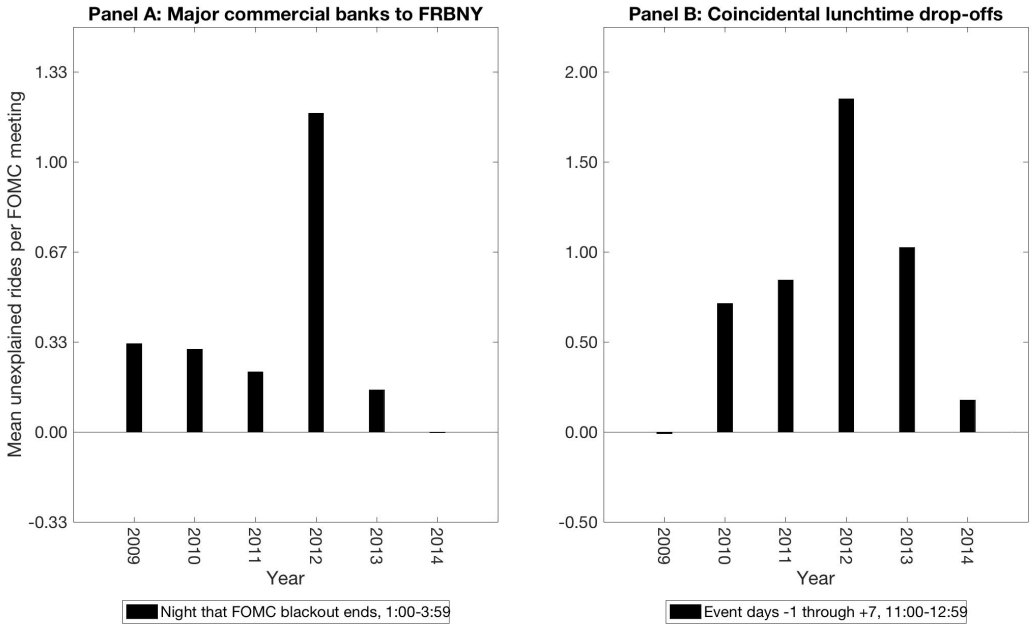
Notes: The sample spans 1425 weekdays from the beginning of 2009 through the end of 2014. For greater consistency with a typical workday, an intraday window starting during the span from 0:00 up to and including 4:00 is mapped to the previous calendar day. Panel A: the mean of bihourly pick-up and drop-off volumes of intra-Manhattan yellow taxi rides. Panel B: bihourly volumes of yellow taxi rides from FRBNY to the sample of locations where major commercial banks have substantial front-office presences. Panel C: bihourly volumes of yellow taxi rides from the sample of locations where major commercial banks have substantial front-office presences to FRBNY. Panel D: bihourly volumes of coincidental drop-offs of passengers picked up around FRBNY and passengers picked up around locations where major commercial banks have substantial front-office presences.

Figure 10: Changes in direct rides and coincidental drop-offs around an FOMC meeting



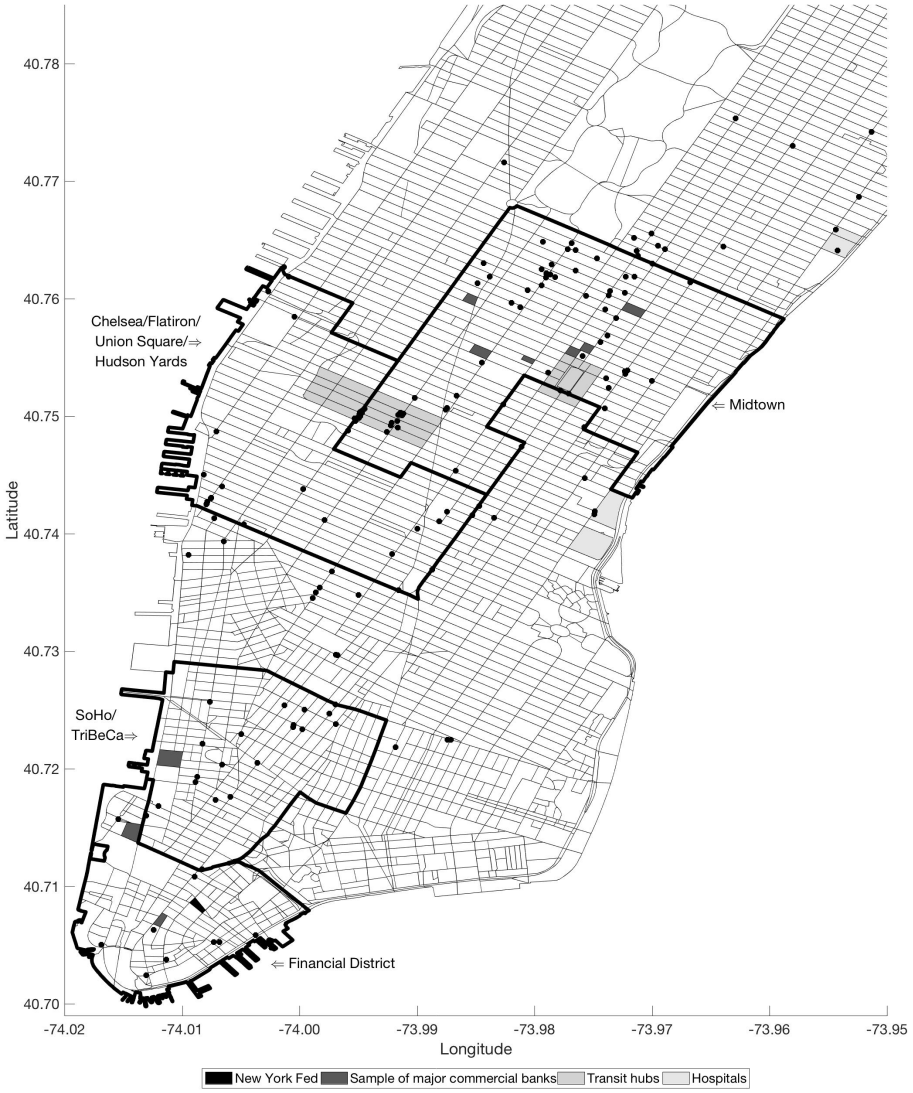
Notes: A separate Poisson regression is run for each intraday-window-event-day pair, where event day t is an offset of t calendar days from an FOMC announcement. For greater consistency with a typical workday, an intraday window starting during the span from 0:00 up to and including 4:00 is mapped to the previous calendar day. The intensity of rides or coincidences is given by Eq. 1. Year-month fixed effects, weekday indicators and overall Manhattan taxi activity are used as controls. To be deemed coincidental, drop-offs must be mapped to the same block; be separated by no more than 1/4 block; be separated by no more than 10 minutes; and not be mapped to the vicinity of a transit hub or any of the FRBNY or commercial-bank buildings. One-sided p-values for coefficients with t-statistics greater than 1.25 are obtained from pairs bootstrapping of year-month observations with at least 1×10^4 repetitions and employ asymptotic refinement. The sample spans 1425 weekdays from the beginning of 2009 through the end of 2014.

Figure 11: Assessment of individual years' contributions to the estimated increases in direct rides and coincidental drop-offs



Notes: Intraday counts are regressed only on year-month fixed effects, weekday controls and overall Manhattan taxi activity, and the mean of the unexplained counts over each year's FOMC windows is plotted. More precisely, we obtain for each year the mean residual over days in the specified FOMC window and multiply that mean by the typical number of weekdays over that window. A higher mean is suggestive of a greater contribution to the increase over the FOMC window estimated using the full sample. The sample spans 1425 weekdays from the beginning of 2009 through the end of 2014.

Figure 12: Locations of lunchtime coincidental drop-offs around FOMC announcements that underlie regression results



Notes: The lunchtime period spans 11:00 through 12:59. The window around an FOMC meeting spans the day before the announcement through a week after the announcement. To be deemed coincidental, drop-offs must be mapped to the same block; be separated by no more than 1/4 block along both the North-South axis and the East-West axis; be separated by no more than 10 minutes; and not be mapped to any of the FRBNY or major-commercial-bank buildings. The outlined areas are based on New York City Neighborhood Tabulation Areas (NTAs). Midtown comprises the Midtown-Midtown South and Turtle Bay-East Midtown NTAs; Chelsea/Flatiron/Union Square/Hudson Yards comprises the Hudson Yards-Chelsea-Flat Iron [*sic*]-Union Square NTA; SoHo/TriBeCa comprises the SoHo-TriBeCa-Civic Center-Little Italy NTA, and the Financial District comprises Battery Park City-Lower Manhattan and parkland. The sample period is 2009 through 2014.

Table 1: Locations of the institutions in our sample, and the distances between them

Panel A: Locations of the Financial Institutions in our Sample

Federal Reserve Bank of New York	33 Liberty Street	40.7084	–74.0087
Bank of America (Merrill Lynch)	One Bryant Park	40.7557	–73.9849
BNY Mellon Original Headquarters	1 Wall Street	40.7075	–74.0116
BNY Mellon New Headquarters	225 Liberty Street	40.712	–74.0153
Citigroup I-banking and New HQ	388 Greenwich Street	40.7207	–74.0111
Citigroup Old Headquarters	399 Park Avenue	40.7591	–73.9717
Goldman Sachs Old Headquarters	85 Broad Street	40.7041	–74.0111
Goldman Sachs Old Trading Floor	One New York Plaza	40.7022	–74.0118
Goldman Sachs New Headquarters	200 West Street	40.7149	–74.0145
JPMorgan Chase Headquarters	270 Park Avenue	40.7558	–73.9754
JPMorgan Investment Banking	383 Madison Avenue	40.7555	–73.9766
Morgan Stanley Headquarters	1585 Broadway	40.7604	–73.9857
Morgan Stanley Investment Management	522 Fifth Avenue	40.7548	–73.9805

Panel B: Distances (in miles) between the Financial Institutions in our Sample

	Institution	1	2	3	4	5	6	7	8	9	10	11	12	13
1	Fed	0	3.5	0.16	0.43	0.86	4.01	0.32	0.45	0.54	3.71	3.66	3.79	3.53
2	Bank of America	3.5	0	3.61	3.41	2.78	0.73	3.82	3.95	3.22	0.5	0.43	0.33	0.24
3	BNY Mellon Orig HQ	0.16	3.61	0	0.37	0.91	4.13	0.24	0.36	0.53	3.84	3.79	3.9	3.65
4	BNY Mellon New HQ	0.43	3.41	0.37	0	0.64	3.98	0.59	0.7	0.2	3.68	3.62	3.68	3.47
5	Citigroup IB and New	0.86	2.78	0.91	0.64	0	3.36	1.15	1.28	0.44	3.07	3.01	3.05	2.85
6	Citigroup Old HQ	4.01	0.73	4.13	3.98	3.36	0	4.33	4.46	3.79	0.3	0.36	0.74	0.55
7	Goldman Sachs Old HQ	0.32	3.82	0.24	0.59	1.15	4.33	0	0.13	0.77	4.03	3.98	4.11	3.85
8	Goldman Sachs Old TF	0.45	3.95	0.36	0.7	1.28	4.46	0.13	0	0.89	4.16	4.11	4.24	3.98
9	Goldman Sachs New HQ	0.54	3.22	0.53	0.2	0.44	3.79	0.77	0.89	0	3.49	3.44	3.49	3.28
10	JPMorgan Chase HQ	3.71	0.5	3.84	3.68	3.07	0.3	4.03	4.16	3.49	0	0.07	0.63	0.28
11	JPMorgan IB	3.66	0.43	3.79	3.62	3.01	0.36	3.98	4.11	3.44	0.07	0	0.58	0.21
12	Morgan Stanley HQ	3.79	0.33	3.9	3.68	3.05	0.74	4.11	4.24	3.49	0.63	0.58	0	0.48
13	Morgan Stanley Inv M	3.53	0.24	3.65	3.47	2.85	0.55	3.85	3.98	3.28	0.28	0.21	0.48	0

Notes: In Panel A, we report the address and locations (latitudes and longitudes) of the financial institutions in our sample. In Panel B, we report the distance (in miles) between the New York Fed and the banks in our sample. Distances are computed based on the latitudes and longitudes of the institutions. Taxi cab ridership comes from the New York City Taxi and Limousine Commission.

Table 2: Distribution of daily ridership between the New York Fed and banks

Panel A: Circles-based ridership measure

			Percentiles				
	Mean	Std Dev	10	25	50	75	90
Fed ridership	279.24	59.32	198	238	288	320	348
Ridership between the Fed and:							
Any bank location in our sample	18.23	6.10	11	14	18	22	26
BNY Mellon Original Headquarters	0.77	0.91	0	0	1	1	2
BNY Mellon New Headquarters	0.84	0.98	0	0	1	1	2
Bank of America (Merrill Lynch)	2.02	1.48	0	1	2	3	4
Citigroup I-banking and New HQ	4.45	2.35	2	3	4	6	8
Citigroup Old Headquarters	2.34	1.63	0	1	2	3	5
Goldman Sachs Old Headquarters	2.12	1.57	0	1	2	3	4
Goldman Sachs Old Trading Floor	2.46	1.55	1	1	2	3	4
Goldman Sachs New Headquarters	3.27	2.04	1	2	3	4	6
JPMorgan Chase Headquarters	2.36	1.78	0	1	2	3	5
JPMorgan Investment Banking	1.80	1.52	0	1	2	3	4
Morgan Stanley Headquarters	1.93	1.48	0	1	2	3	4
Morgan Stanley Investment Management	0.92	1.00	0	0	1	1	2

Panel B: Blocks-based ridership measure

			Percentiles				
	Mean	Std Dev	10	25	50	75	90
Fed ridership	296.67	62.93	210	250	305	342	370
Ridership between the Fed and:							
Any bank location in our sample	11.84	4.45	6	9	12	15	18
BNY Mellon Original Headquarters	0.36	0.62	0	0	0	1	1
BNY Mellon New Headquarters	0.52	0.74	0	0	0	1	1
Bank of America (Merrill Lynch)	1.34	1.14	0	0	1	2	3
Citigroup I-banking and New HQ	4.14	2.23	1	3	4	5	7
Citigroup Old Headquarters	1.23	1.14	0	0	1	2	3
Goldman Sachs Old Headquarters	0.85	0.93	0	0	1	1	2
Goldman Sachs Old Trading Floor	2.04	1.48	0	1	2	3	4
Goldman Sachs New Headquarters	2.40	1.73	0	1	2	3	5
JPMorgan Chase (HQ and I-banking)	1.38	1.32	0	0	1	2	3
Morgan Stanley Headquarters	0.86	0.94	0	0	1	1	2
Morgan Stanley Investment Management	0.22	0.47	0	0	0	0	1

Notes: We report statistics on the distribution of daily cab ridership between the New York Fed and the bank locations in our sample. In Panel A, we use the circles-based ridership classification methodology described in Section 2.3.1. In Panel B, we use the blocks-based methodology described in Section 2.3.2. The sum of the average daily ridership to the individual bank locations is not equal to the average daily ridership between the Fed and any of the bank locations in our sample. This arises for two reasons. First, many of these locations are very close to one another (see panel B of Table 1). When counting rides between the Fed and any bank location, we only count the ride once, but the same ride might be matched to more than one bank location, which would inflate the sum of the individual bank location ridership levels. Second, recall that the Goldman Sachs, Bank of New York Mellon, and Citigroup old headquarters locations do not span the entire sample period; rather, each of these locations is only valid for a subperiod, and the averages are only computed over the period of time when they are occupied by the bank.

Table 3: Taxi Ridership, Trading Days versus Non-Trading Days

Panel A: Full sample	Rides per day (1)	Rides per trading day (2)	Rides per non-trading day (3)
All NYC cab rides (yellow cab)	447,322.37	467,519.61	402,537.20
Cab rides to or from the NY Fed	252.93	279.24	194.59
Cab rides between the NY fed and banks	14.54	18.23	6.36
Panel B: Placebo test			
All NYC cab rides (yellow cab)	447,322.37	467,519.61	402,537.2
Cab rides to or from the NY Fed	766.85	812.33	666
Cab rides between the NY fed and banks	30.36	35.56	18.81
Panel C: T-test trading/non-trading day ratio	Full sample (Columns 2/3)	Placebo sample (Columns 2/3)	t-statistic
Any of the bank locations in our sample	18.23/6.36	35.56/18.81	18.27

Notes: We report the average number of daily taxi rides to and from various locations. In Panel A, rides are considered to (from) the Fed if the dropoff (pickup) location is within 0.05 miles from the Fed. We use a larger radius (0.10 miles) for the financial institutions since they are generally larger than the Fed. These are the cutoffs we use throughout our paper. In Panel B, we conduct placebo tests by forming "doughnuts" around the Fed and each of the institutions in our sample. Namely, rides are classified to (from) the Fed if the dropoff (pickup) location is between 0.05 miles and 0.10 miles from the Fed. For the banks in our sample, we require that the pickup/dropoff location be between 0.10 miles and 0.15 miles. In Panel C, we test whether the ratio of ridership on trading days to non-trading days is significantly higher using our baseline classification than for our placebo classification. We conduct this test by computing the ratio of daily ridership on trading to non-trading days each month for both our baseline classification and our placebo classification. We then compute the difference in these ratios each month, and we conduct a t-test on this sample of 90 differences (where our unit of observation is month). Taxi cab ridership comes from the New York City Taxi and Limousine Commission.

Table 4: Quantiles and counts of daily rides between the New York Fed and major commercial banks around Dodd-Frank milestones

Panel A: All direct rides – quantile (count)

Dodd-Frank Milestone	Date	$t-1$	t	$t+1$
Introduced in House	2/12/2009	86.9 (17)	3.3 (5)	68.5 (14)
Passed in House	11/12/2009	16.6 (8)	81.9 (16)	60.6 (13)
Passed with amendment in Senate	20/5/2010	68.5 (14)	33.3 (10)	51.6 (12)
Conference report filed	29/6/2010	99.4 (24)	99.8 (27)	60.6 (13)
Conference report agreed in House	30/6/2010	99.8 (27)	60.6 (13)	51.6 (12)
Conference report agreed in Senate	15/7/2010	100.0 (32)	75.9 (15)	96.4 (20)
Signed by President	21/7/2010	99.9 (29)	96.4 (20)	86.9 (17)

Panel B: Pick-ups and drop-offs >100 feet from Milbank HQ – quantile (count)

Dodd-Frank Milestone	Date	$t-1$	t	$t+1$
Introduced in House	2/12/2009	93.9 (7)	2.7 (0)	88.4 (6)
Passed in House	11/12/2009	26.6 (2)	77.5 (5)	10.4 (1)
Passed with amendment in Senate	20/5/2010	44.4 (3)	26.6 (2)	44.4 (3)
Conference report filed	29/6/2010	97.2 (8)	98.8 (9)	26.6 (2)
Conference report agreed in House	30/6/2010	98.8 (9)	26.6 (2)	44.4 (3)
Conference report agreed in Senate	15/7/2010	100.0 (12)	88.4 (6)	97.2 (8)
Signed by President	21/7/2010	93.9 (7)	93.9 (7)	62.9 (4)

Notes: Quantiles as percentages are presented first and then counts in parentheses. The milestones are the events in the passage of Dodd-Frank listed in U.S. Congress (n.d.) with the exception of conference committees. t indicates the date of the event, while $t-1$ and $t+1$ are respectively the previous weekday and the subsequent weekday. For greater consistency with workdays, daily ride volume is calculated as the count of rides from 5:00 through 4:59. The quantiles are calculated over the 1425 weekdays in the filtered sample spanning 2009 through 2014.

Table 5: Profitability of Trading Strategy

Daily market excess return, strategy in market	13.83
Daily market excess return, strategy out of market	–0.10
T-stat (difference of means)	2.69
Sharpe ratio of trading strategy	1.45
Sharpe ratio of opposite strategy	–0.01

Notes: We consider the profitability of a trading strategy that holds the market only if there has been relatively low recent ridership between the Fed and the major, systemically important banks in our sample. Each day, the strategy either holds the market or a risk-free asset. We report the average market return (in excess of the risk-free rate) on the days that the strategy is in and out of the market. In addition, we report the t-statistic for the difference in means between these statistics, as well as the Sharpe ratio for (i) our strategy, and (ii) the strategy that does the opposite of our market (i.e., invests in a risk-free asset when our strategy invests in the market portfolio, and invests in the market portfolio when our strategy invests in a risk-free asset). To determine whether to hold the market, we compare the level of ridership between the Fed and the banks in our sample in the past 10 days (ridership period = 10 days), i.e., $Ridership_{10}$, to the distribution of $Ridership_{10}$ in the 20 days prior to the given date (benchmark period = 20 days). Our strategy holds the market if and only if $Ridership_{10}$ is less than the median value (percentile cutoff = 50) of $Ridership_{10}$ during the benchmark period. We deflate (i.e., scale) the ridership level each day by the total number of rides to and from the New York Fed, and we use a one day lag between the estimation period and the trading strategy (lag = 1 day). In other words, the estimation period that runs through Tuesday at 4:00 PM determines the trading strategy's position on Thursday, not Wednesday. Our baseline strategy has five parameters: ridership period, choice of deflator, a cutoff value for what constitutes high or low levels of recent ridership, the lag between the benchmark period and the ridership estimation period, and the length of the ridership estimation period. In Table 6, we report the results for alternative choices for the parameter values. Taxi cab ridership comes from the New York City Taxi and Limousine Commission. Our daily excess returns are taken from Ken French's website.

Table 6: Robustness of Trading Strategy Profits to Parameter Selection

Panel A: Alternate ridership periods

	5 day ridership	15 day ridership
Daily market excess return, strategy in market	15.34	11.97
Daily market excess return, strategy out of market	–2.65	2.30
T-stat (difference of means)	3.46	1.87
Sharpe ratio of trading strategy	1.48	1.22
Sharpe ratio of opposite strategy	–0.27	0.23

Panel B: Alternate deflators

	NYC deflator	No deflator
Daily market excess return, strategy in market	10.66	12.00
Daily market excess return, strategy out of market	3.06	1.84
T-stat (difference of means)	1.47	1.96
Sharpe ratio of trading strategy	1.11	1.22
Sharpe ratio of opposite strategy	0.30	0.18

Panel C: Alternate cutoff values for high/low ridership

	25th pctile	75th pctile
Daily market excess return, strategy in market	20.58	10.06
Daily market excess return, strategy out of market	0.95	–0.47
T-stat (difference of means)	3.46	1.89
Sharpe ratio of trading strategy	1.57	1.21
Sharpe ratio of opposite strategy	0.11	–0.04

Panel D: Alternate lags between benchmark and estimation periods

	0 day lag	5 day lag
Daily market excess return, strategy in market	15.18	14.16
Daily market excess return, strategy out of market	–1.86	0.23
T-stat (difference of means)	3.28	2.69
Sharpe ratio of trading strategy	1.60	1.53
Sharpe ratio of opposite strategy	–0.18	0.02

Panel E: Alternate benchmark period lengths

	15 day benchmark	25 day benchmark
Daily market excess return, strategy in market	14.89	13.52
Daily market excess return, strategy out of market	–2.34	0.39
T-stat (difference of means)	3.31	2.55
Sharpe ratio of trading strategy	1.54	1.42
Sharpe ratio of opposite strategy	–0.23	0.04

Notes: Each panel considers permutation of the parameter choices in our baseline trading strategy (see Table 5). Our baseline strategy has five parameters: ridership period, choice of deflator, a cutoff value for what constitutes high or low levels of recent ridership, the lag between the benchmark period and the ridership estimation period, and the length of the ridership estimation period. In Panel A we vary the number of days over which Fed-bank ridership is calculated, in Panel B we consider no deflator as well as deflating each day's ridership level by the total number of Yellow cab rides in New York City as a whole, in Panel C, we compare $Ridership_{10}$ to the 25th and 75th percentiles of $Ridership_{10}$ during the benchmark period, as opposed to our baseline specification (the median value of $Ridership_{10}$), and in Panel E we consider different benchmark period lengths.

Table 7: Stock Market Predictability around FOMC meetings

Panel A: Raw Counts

	Market up	Market down
Ridership up	17	24
Ridership down	16	2

Panel B: Regressions

	Market up dummy	Market return
Ridership up dummy	–0.474***	–0.0194***
	[0.128]	[0.00567]
Constant	0.889***	0.0132***
	[0.107]	[0.00473]
Observations	59	59
R-squared	0.19	0.17

Notes: We sort the 59 FOMC meetings during our sample period based on whether (Fed-deflated) daily cab ridership between the Fed and the banks in our sample increased or decreased from the benchmark period (the second trading day after the previous FOMC meeting through FOMC event day 12, inclusive) and the pre-FOMC period (event days -11 through -2, inclusive). Our measure of returns around FOMC meetings is the cumulative excess return of the market over FOMC event days 1 through $+1$, inclusive. "Market up" refers to FOMC meetings in which the market return exceeded the risk-free rate, and "Market down" is defined analogously. Taxi cab ridership comes from the New York City Taxi and Limousine Commission. Our daily excess returns are taken from Ken French's website.

Table 8: Poisson regression analysis of taxi trips around FOMC meetings. Trip types are *i.*) rides from major commercial banks to the New York Fed; and *ii.*) coincidental drop-offs of passengers picked up near the New York Fed and passengers picked up near major commercial banks.

Trip type:	*Commercial bank to FRBNY*		*Coincidental drop-off of FRBNY and commercial-bank pick-ups*				
Candidate windows: (FOMC × intraday)	*Single day × all intraday*		*All spans × all intraday between 9:00 and 17:00*				
	(1)	*(2)*	*(3)*	*(4)*	*(5)*	*(6)*	*(7)*
			Change in trips per meeting				
Growth (%)	118.8	6.5	49.8	67.4	2.4	96.1	108.8
t-statistic	(4.39)	(0.75)	(4.35)	(3.98)	(0.36)	(3.83)	(4.30)
Extra trips	0.43	0.23	1.41	1.10	0.24	0.71	0.85
Indiv. p-value (%)	<0.1	23.5	<0.1	<0.1	35.7	<0.1	<0.1
Robust signif. (%)	1.0		1.0	5.0		5.0	1.0
			Selected FOMC and intraday windows				
Max. overall significance	Yes	No	Yes	No	No	No	No
First event day	+1	+1	−2	−1	−1	−1	−1
Last event day	+1	+1	+7	+7	+7	+7	+7
Start of intraday window	1:00	1:00	10:00	11:00	11:00	11:00	11:00
End of intraday window	3:59	3:59	12:59	12:59	12:59	12:59	12:59
			Geographic and temporal restrictions				
Only FRBNY neighbours	No	Yes	No	No	Yes	No	No
Extra restrictions*	No	No	No	No	No	Yes	No
Only single passengers	No	No	No	No	No	No	Yes
Pseudo-R^2	0.19	0.23	0.08	0.09	0.15	0.07	0.11
Observations	1425	1425	1425	1425	1425	1425	1425
Candidate models	572		1690	1690		1685	1690

Notes: Trips are modelled as Poisson processes with intensity given by Eq. 1, and the count of extra trips over the FOMC window is the associated average partial effect over dates mapped to that window times the typical number of weekdays during that window. Coincidental drop-offs do not include trips where either drop-off is mapped to the New York Fed, a major commercial bank or a transit hub. Event day X refers to a date that is offset by X calendar days from an FOMC announcement, with negative values indicating dates before an announcement. Windows beginning during the span from 0:00 through 4:00 are treated as part of the preceding calendar day. For direct rides, the set of candidate FOMC windows is the set of single event days from −8 through +7, and the set of intraday windows are all two- and three-hour spans beginning at the top of an hour. For coincidental drop-offs, the set of candidate FOMC windows is the set of all contiguous spans during the period from event day −8 through event day +7, and the set of intraday windows are all two- and three-hour spans beginning at the top of an hour within the span from 9:00 until 17:00. All specifications include controls for overall Manhattan taxi activity, weekday controls and year-month fixed effects, and all standard errors employ clustering at the year-month level. Individual p-values are right-tail quantiles, are obtained from pairs bootstrapping at the year-month level with 1×10^5 simulations and employ asymptotic refinement. Sections A.2 and A.3 of the Internet Appendix provides additional details on the estimation and the variant of the Romano and Wolf (2005) StepM procedure employed in the assessment of data-mining-robust significance. Robust significance of changes for the sample of FRBNY neighbours was not calculated given their individual insignificance. Pseudo-R^2 is $1 - SSR/TSS$. Counts of candidate models can vary across specifications of the same ride type due to variation in data sparsity. Regressions span a filtered set of weekdays from 2009 through 2014. ∗: The extra restrictions entail the omission of drop-offs mapped to Midtown, the Financial District and a set of hospitals.

A Internet Appendix

A.1 Trading Strategy Results using the Blocks-Based Ride Classification Methodology

In Section 3.1, we analyzed the profitability of our trading strategy using the circles-based methodology developed in Section 2.3.1. In Figures A.3-A.7 and Tables A.1-A.3, we reproduce Figures 3-6, Figure 8, and Tables 5-7, except that we use the blocks-based ride classification methodology developed in Section 2.3.2.

[Insert Figures A.3-A.7 and Tables A.1-A.3 here]

A.2 Estimation details

The potential for sparse rides necessitates a filtering of coefficients. In the calculation of standard errors, we cluster at the year-month level. The effective number of clusters is equal to the number of months with non-zero ride or coincidence volumes, so we omit particularly sparse windows to avoid unreliable standard errors. Specifically, for a given ride type and FOMC window, we drop any intraday windows for which there are positive rides in fewer than 36 of the 72 months. In general, an absence of observations during an FOMC window theoretically yields a coefficient estimate of $-\infty$ and an undefined standard error. As the numerical estimation might return a highly negative but finite value instead, we drop coefficients corresponding to a decrease by a factor of at least 0.9999. As a guard against unrealistic estimates from sparse data, we also perform the mirror-image filtration and drop any coefficients corresponding to an increase by a factor of at least 1×10^4. As a precaution against unreliable estimates due to near-singularity, we omit every coefficient for which the matrix to be inverted in the calculation of a standard error has a condition number of at least 1×10^6.

We employ bootstrapped p-values with asymptotic refinement. Each simulation is a draw of 72 year-month observations with replacement. I follow Cameron, Gelbach and Miller (2008) and Cameron and Trivedi (2009) in employing the empirical distribution of the standardised deviations from the coefficient obtained from the unsimulated data. Non-positivity of a coefficient is rejected at the $100\alpha\%$ confidence level when its t-statistic exceeds the $1 - \alpha$ quantile of the empirical distribution.

To reduce the computational burden of the production of Figure 10, the number of bootstrap simulations employed may vary from coefficient to coefficient. Our interest there is primarily in whether changes are individually significant at the 5% or 1% confidence level, and we initially obtain bootstrapped p-values for all coefficients with t-statistics greater than

1.25 using 1×10^4 simulations. When we obtain a p-value of 0.5% or below, we subsequently employ 1×10^5 simulations to assess significance at the 0.1% confidence level.

A.3 Data-mining-robust significance

To assess the significance of increases conditional on data mining, we follow the StepM procedure with studentisation described in Romano and Wolf (2005) and Romano, Shaikh and Wolf (2008). This is a step-down procedure that seeks to maximise the number of rejected null hypotheses for a given type-I error rate α. To fix concepts, let $\{\beta_i\}, i \in \{1, 2, ..., N\}$ be a set of true parameters, and let $\{\hat{\beta}_i\}$ and $\{\hat{\sigma}_i\}$ be respectively their point estimates and standard errors. we seek to answer which $\beta_j \in \{\beta_i\}$ are significantly greater than 0.

The Romano and Wolf (2005) StepM procedure is based on the order statistics of simulated estimates' deviations from those obtained from the unsimulated data. When one employs studentisation, one first simulates M realisations of the N estimates' standardised deviations from the true values. For each of the M simulations, one obtains the maximum standardised deviation from the corresponding estimate obtained from the unsimulated data. Let q_1^α denote the $1 - \alpha$ quantile of the distribution of maximum standardised deviations. One rejects the null hypothesis of non-positivity for any β_i for which $\hat{\beta}_i/\hat{\sigma}_i$ is greater than q_1^α. Using the same simulations, one proceeds to calculate the $1 - \alpha$ quantile of the distribution of maximum standardised deviations for the random variables for which the null was not rejected, henceforth q_2^α. One then rejects the null for any β_i for which $\hat{\beta}_i/\hat{\sigma}_i$ is greater than q_2^α. This process continues until no additional null hypotheses are rejected.

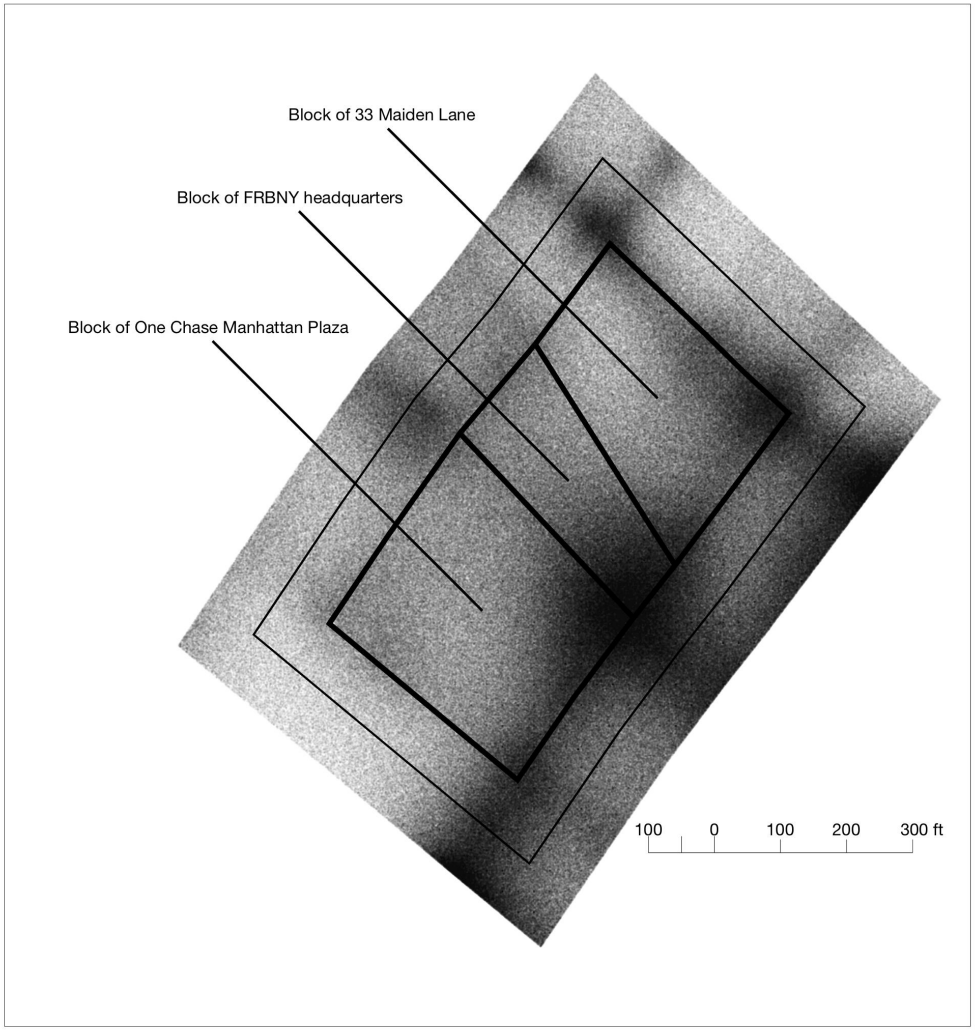
This assessment of significance is conservative in that it asymptotically provides an upper bound (Romano and Wolf 2005). This approach does, however, provide power benefits over a simple Bonferroni correction in that the Bonferroni assumption that estimates are independent imposes the highest bar for the rejection of statistical insignificance (Romano and Wolf 2005). The overlap of the intraday windows in this study guarantees at least some dependence. Additional dependence will be introduced if the decision to meet during a certain window is a function of the timing of proximate interactions.

Romano and Wolf (2005) present a bootstrap-based approach applicable to a wide variety of random variables, but bootstrapping here is problematic. Given that we must estimate the right tail of the distribution of the maximum standardised deviation from a true value over hundreds of coefficients, we require an extremely large number of simulations. It is not feasible to obtain a sufficient number of bootstrap simulations for such a large number of coefficients that lack closed-form estimators. Moreover, the block bootstrapping at the year-month level oversamples outliers in what are already leptokurtic data. For example,

counts with a true probability of less than 1 in 1425 will be drawn with probability 1 in 72. With hundreds of coefficients per simulation, there is a danger that the maxima for a non-trivial number of bootstrap simulations might be driven by the oversampling of outliers.

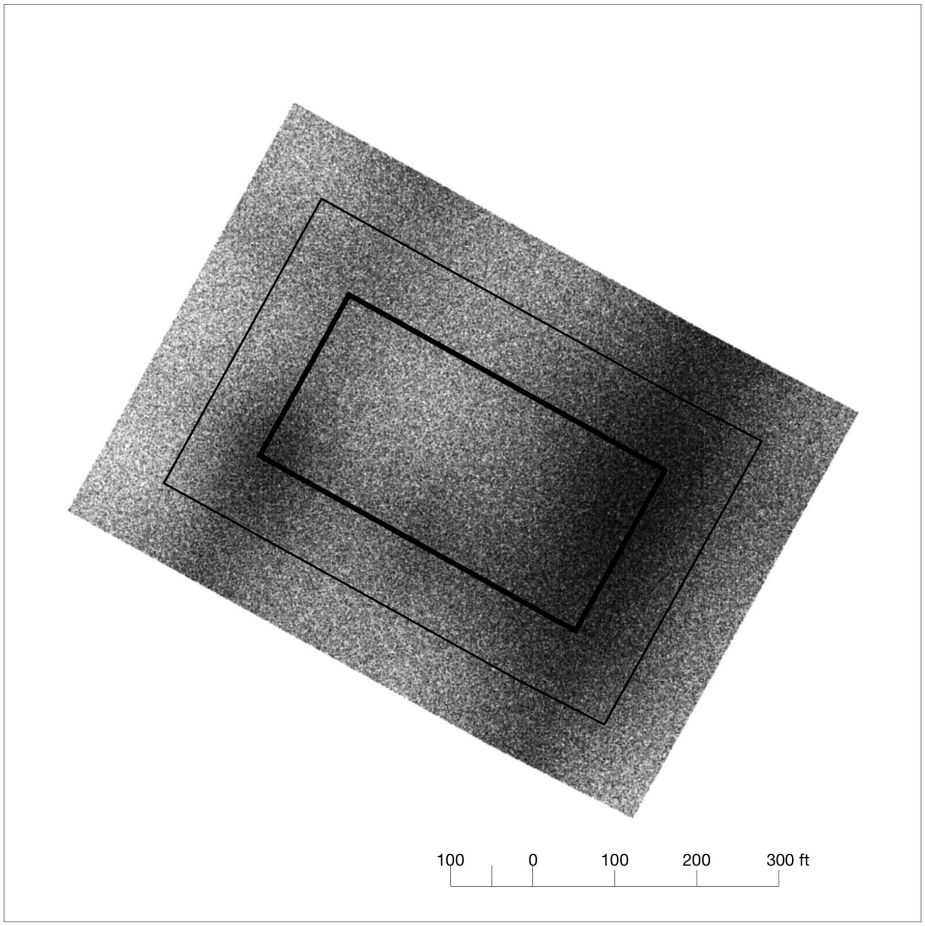
Instead of bootstrapping, we exploit the properties of GMM estimates and employ Monte Carlo simulation. For the GMM estimation, we collect all of the models' Poisson-regression moment conditions. The point estimates are the same as when the regressions are performed individually. In the estimation of the covariance matrix, we impose no restrictions on the covariances of errors on the days in year-month i both within and across models, and we impose that all errors in year-month i be independent of all errors in year-month j for $i \neq j$. The covariance matrices for the individual regressions are along the diagonal. Asymptotically, joint draws of standardised deviations from the true coefficients on the FOMC-window indicators are distributed normally with mean 0 and with a covariance matrix equal to those coefficients' correlation matrix. To obtain a reasonably good picture of the right tail of the distribution of maxima, we take 1×10^6 draws from a multivariate normal distribution with mean 0 and a covariance matrix equal to the estimated correlation matrix of the coefficients on the FOMC-window indicators. Each step of the StepM procedure uses the same set of simulations and simply removes from the selection of the maxima the contributions of the coefficients for which non-positivity has been rejected.

Figure A.1: Yellow taxi pick-up and drop-off coordinates in the neighbourhood of the New York Fed



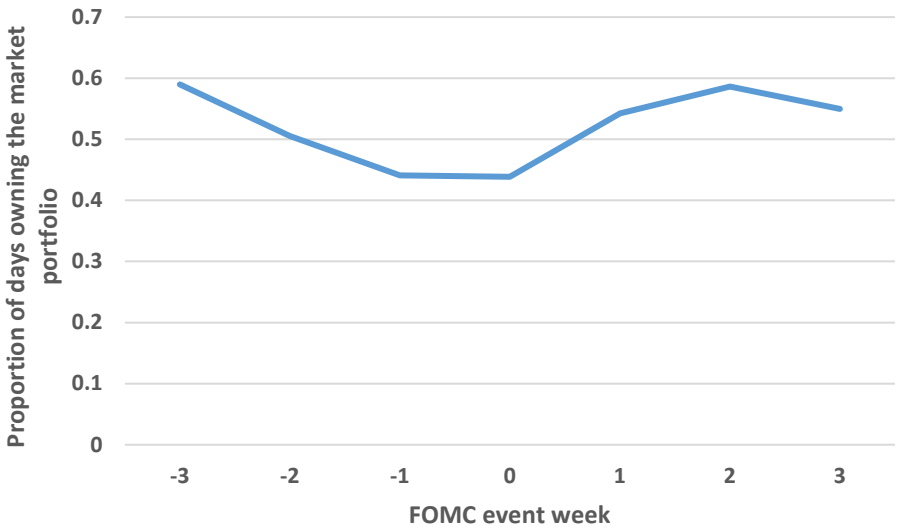
Notes: The borders of the blocks are superimposed and approximately indicate the locations of roads (thick line segments). The thin line segments are an outward shift of the boundary of the three blocks by 100 feet. The data span all rides from 2009 through 2014.

Figure A.2: Yellow taxi pick-up and drop-off coordinates around the block of JPMorgan Chase's corporate headquarters



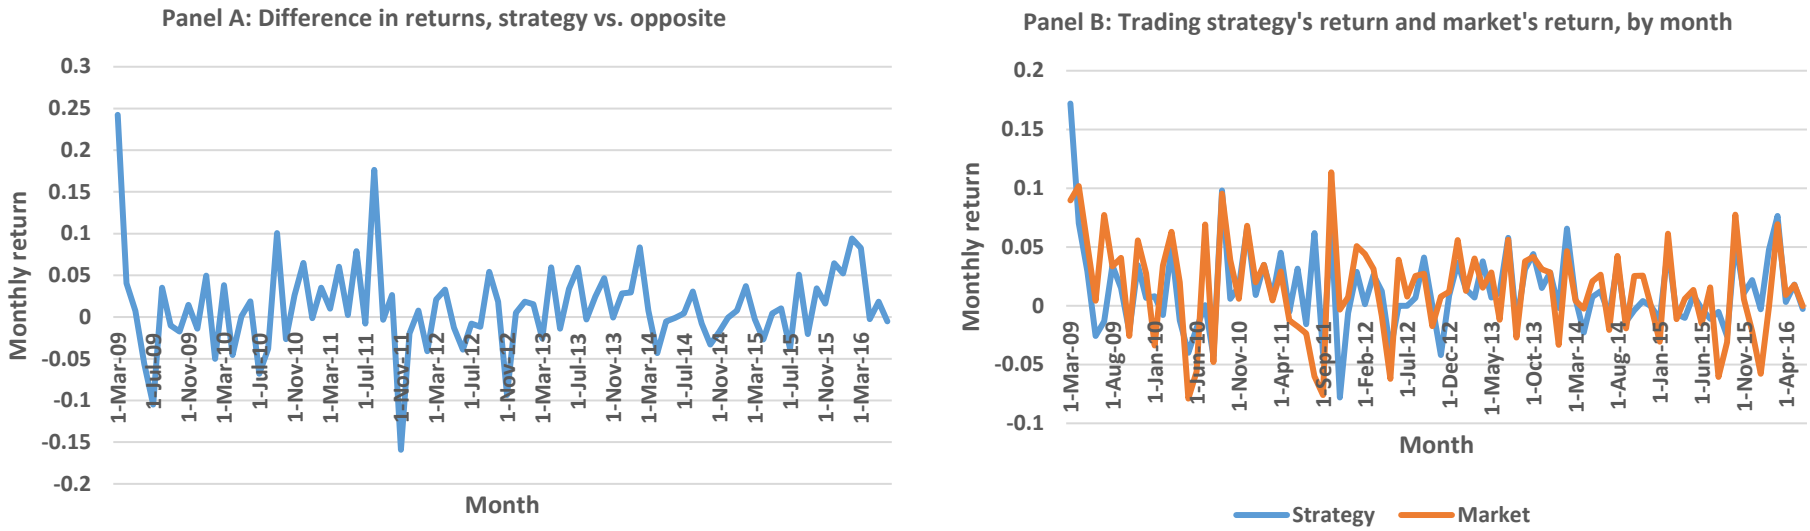
Notes: The border of the block is superimposed and approximately indicates the locations of roads (thick line segments). The thin line segments are an outward shift of the boundary of the block by 100 feet. For greater clarity given the very large volume of rides, the data are random draws of one-tenth of the rides from 2009 through 2014 rather than the full set.

Figure A.3: Trading Strategy Holding by FOMC Event Week



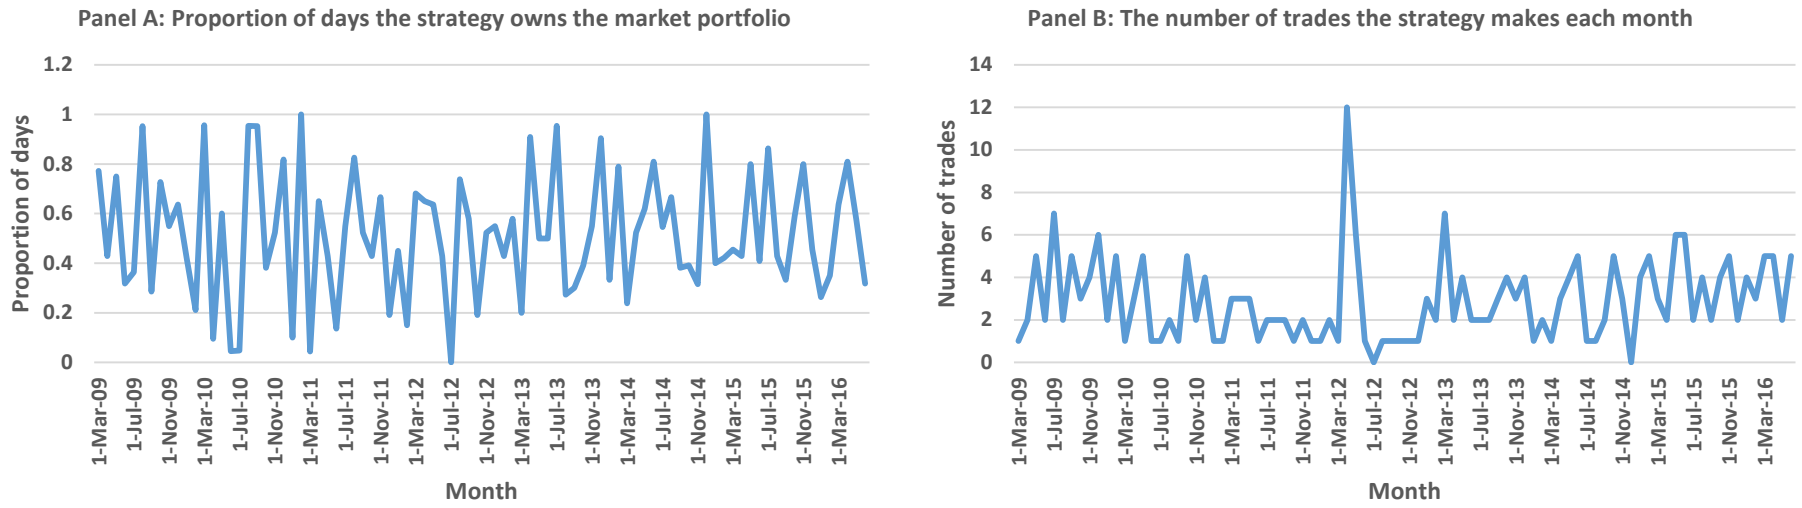
Notes: This figure reproduces Figure 3, except that we use the blocks-based ride classification methodology developed in Section 2.3.2 instead of the circles-based methodology developed in Section 2.3.1. We plot the average frequency with which the trading strategy described in Section 3.1 holds the market portfolio (as opposed to a risk-free security) across the FOMC cycle.

Figure A.4: Time Series of Trading Strategy's Monthly Profits



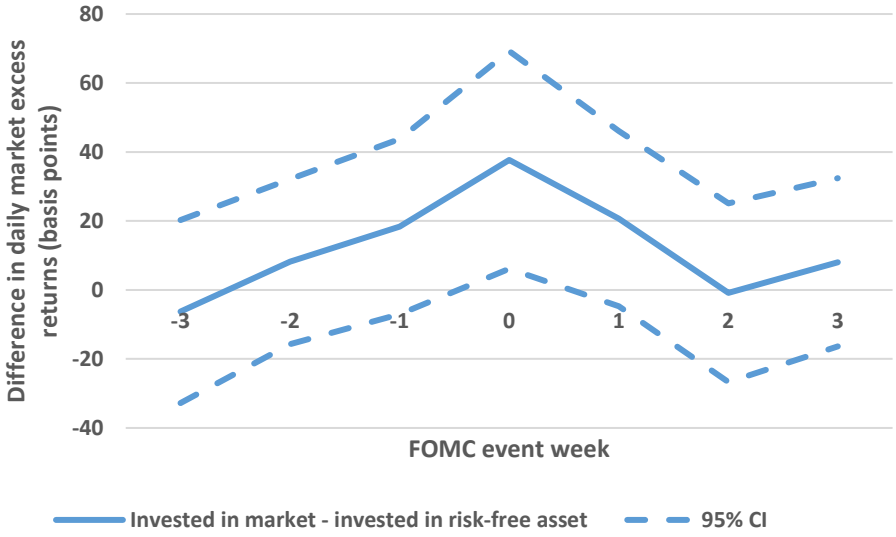
Notes: This figure reproduces Figure 4, except that we use the blocks-based ride classification methodology developed in Section 2.3.2 instead of the circles-based methodology developed in Section 2.3.1. We plot the time series of monthly profits from the trading strategy described in Section 3.1. Each day, the trading strategy either holds the market portfolio or a risk-free asset. It holds the market portfolio if the ridership between the Fed and the systemically important banks in our sample has been low in the previous 10 trading days relative to the 20 trading days prior; otherwise, it holds the risk-free asset. In Panel A, we report the difference between the cumulative excess return of the market (in excess of the risk-free rate) on the days on which the strategy owns the market and the days on which the strategy owns a risk-free asset. This represents the return differential between our strategy and an "opposite" strategy that holds the market portfolio when our strategy owns a risk free asset, and holds a risk free asset when our strategy owns the market portfolio. Taxi cab ridership comes from the New York City Taxi and Limousine Commission. Our daily excess returns are taken from Ken French's website. In Panel B, we report the monthly excess returns of our trading strategy and the market portfolio.

Figure A.5: Holdings and Trading Volume of Strategy



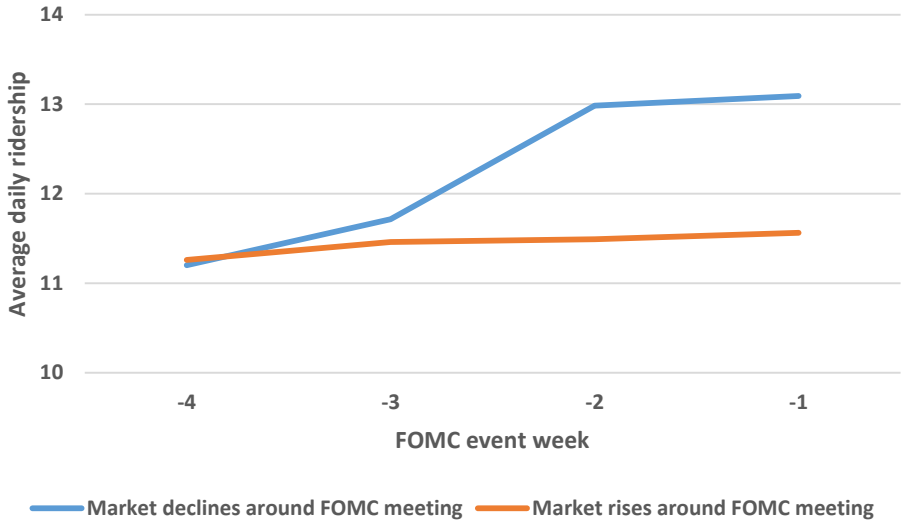
Notes: This figure reproduces Figure 5, except that we use the blocks-based ride classification methodology developed in Section 2.3.2 instead of the circles-based methodology developed in Section 2.3.1. We plot the proportion of days the trading strategy owns the market portfolio (Panel A) and the number of trades the strategy makes (Panel B) by month. Each day, the trading strategy either holds the market portfolio or a risk-free asset. It holds the market portfolio if the ridership between the Fed and the systemically important banks in our sample has been low in the previous 10 trading days relative to the 20 trading days prior; otherwise, it holds the risk-free asset. The trading strategy is described in further detail in Section 3.1. Taxi cab ridership comes from the New York City Taxi and Limousine Commission.

Figure A.6: Market Timing Ability in FOMC Event Time



Notes: This figure reproduces Figure 6, except that we use the blocks-based ride classification methodology developed in Section 2.3.2 instead of the circles-based methodology developed in Section 2.3.1. We plot the average difference in average daily return of the market (in basis points) on days that our trading strategy owns the market portfolio and days that it owns a risk-free asset, by FOMC event week. The 95% confidence interval is plotted in dashed lines. FOMC event week 0 corresponds to the week of an FOMC meeting. Taxi cab ridership comes from the New York City Taxi and Limousine Commission. Our daily excess returns are taken from Ken French's website.

Figure A.7: Daily ridership in FOMC event time as a function of future FOMC returns



Notes: This figure reproduces Figure 8, except that we use the blocks-based ride classification methodology developed in Section 2.3.2 instead of the circles-based methodology developed in Section 2.3.1. We plot average daily ridership between the Fed and the banks in our sample in FOMC event time. We plot this separately for FOMC meetings in which the market had positive and negative excess returns during the FOMC window (event days -1 to $+1$). FOMC event week 0 corresponds to FOMC event days 1 to $+1$. FOMC event week 1 corresponds to FOMC event days 6 through 2 (inclusive), and the other weeks are defined analogously. Taxi cab ridership comes from the New York City Taxi and Limousine Commission. Our daily excess returns are taken from Ken French's website.

Table A.1: Profitability of Trading Strategy

Daily market excess return, strategy in market	12.08
Daily market excess return, strategy out of market	1.05
T-stat (difference of means)	2.13
Sharpe ratio of trading strategy	1.26
Sharpe ratio of opposite strategy	0.10

Notes: This table reproduces Table 5, except that we use the blocks-based ride classification methodology developed in Section 2.3.2 instead of the circles-based methodology developed in Section 2.3.1. We consider the profitability of a trading strategy that holds the market only if there has been relatively low levels of recent ridership between the Fed and the major, systemically important banks in our sample. Each day, the strategy either holds the market or a risk-free asset. We report the average market return (in excess of the risk-free rate) on the days that the strategy is in and out of the market. In addition, we report the t-statistic for the difference in means between these statistics, as well as the Sharpe ratio for (i) our strategy, and (ii) the strategy that does the opposite of our market (i.e., invests in a risk-free asset when our strategy invests in the market portfolio, and invests in the market portfolio when our strategy invests in a risk-free asset). To determine whether to hold the market, we compare the level of ridership between the Fed and the banks in our sample in the past 10 days (ridership period = 10 days), i.e., $Ridership_{10}$, to the distribution of $Ridership_{10}$ in the 20 days prior to the given date (benchmark period = 20 days). Our strategy holds the market if and only if $Ridership_{10}$ is less than the median value (percentile cutoff = 50) of $Ridership_{10}$ during the benchmark period. We deflate (i.e., scale) the ridership level each day by the total number of rides to and from the New York Fed, and we use a one day lag between the estimation period and the trading strategy (lag = 1 day). In other words, the estimation period that runs through Tuesday at 4:00 PM determines the trading strategy's position on Thursday, not Wednesday. Our baseline strategy has five parameters: ridership period, choice of deflator, a cutoff value for what constitutes high or low levels of recent ridership, the lag between the benchmark period and the ridership estimation period, and the length of the ridership estimation period. In Table 6, we report the results for alternative choices for the parameter values. Taxi cab ridership comes from the New York City Taxi and Limousine Commission. Our daily excess returns are taken from Ken French's website.

Table A.2: Robustness of Trading Strategy Profits to Parameter Selection

Panel A: Alternate ridership periods

	5 day ridership	15 day ridership
Daily market excess return, strategy in market	8.28	10.87
Daily market excess return, strategy out of market	4.13	2.55
T-stat (difference of means)	0.79	1.61
Sharpe ratio of trading strategy	0.82	1.11
Sharpe ratio of opposite strategy	0.42	0.26

Panel B: Alternate deflators

	NYC deflator	No deflator
Daily market excess return, strategy in market	11.69	11.44
Daily market excess return, strategy out of market	1.46	2.14
T-stat (difference of means)	1.98	1.79
Sharpe ratio of trading strategy	1.20	1.13
Sharpe ratio of opposite strategy	0.14	0.22

Panel C: Alternate cutoff values for high/low ridership

	25th pctile	75th pctile
Daily market excess return, strategy in market	12.26	9.77
Daily market excess return, strategy out of market	4.16	-0.03
T-stat (difference of means)	1.45	1.75
Sharpe ratio of trading strategy	0.96	1.16
Sharpe ratio of opposite strategy	0.49	0.00

Panel D: Alternate lags between benchmark and estimation periods

	0 day lag	5 day lag
Daily market excess return, strategy in market	9.16	7.90
Daily market excess return, strategy out of market	3.65	6.18
T-stat (difference of means)	1.06	0.33
Sharpe ratio of trading strategy	0.93	0.84
Sharpe ratio of opposite strategy	0.36	0.59

Panel E: Alternate benchmark period lengths

	15 day benchmark	25 day benchmark
Daily market excess return, strategy in market	11.61	11.48
Daily market excess return, strategy out of market	0.71	2.27
T-stat (difference of means)	2.09	1.79
Sharpe ratio of trading strategy	1.18	1.14
Sharpe ratio of opposite strategy	0.07	0.23

Notes: This table reproduces Table 6, except that we use the blocks-based ride classification methodology developed in Section 2.3.2 instead of the circles-based methodology developed in Section 2.3.1. Each panel considers permutation of the parameter choices in our baseline trading strategy (see Table 5). Our baseline strategy has five parameters: ridership period, choice of deflator, a cutoff value for what constitutes high or low levels of recent ridership, the lag between the benchmark period and the ridership estimation period, and the length of the ridership estimation period. In Panel A we vary the number of days over which Fed-bank ridership is calculated, in Panel B we consider no deflator as well as deflating each day's ridership level by the total number of Yellow cab rides in New York City as a whole, in Panel C, we compare $Ridership_{10}$ to the 25th and 75th percentiles of $Ridership_{10}$ during the benchmark period, as opposed to our baseline specification (the median value of $Ridership_{10}$), and in Panel E we consider different benchmark period lengths.

Table A.3: Stock Market Predictability around FOMC meetings

Panel A: Raw Counts

	Market up	Market down
Ridership up (NYC deflator, circles measure)	17	24
Ridership down (NYC deflator, circles measure)	16	2
Ridership up (no deflator, circles measure)	17	23
Ridership down (no deflator, circles measure)	16	3
Ridership up (Fed deflator, blocks measure)	18	17
Ridership down (Fed deflator, blocks measure)	15	9
Ridership up (NYC deflator, blocks measure)	17	20
Ridership down (NYC deflator, blocks measure)	16	6
Ridership up (no deflator, blocks measure)	19	20
Ridership down (no deflator, blocks measure)	14	6

Panel B: Regressions, dummy for positive returns around FOMC meeting

	(1)	(2)	(3)	(4)	(5)
Ridership up dummy	–0.474***	–0.417***	-0.111	–0.268**	-0.213
	[0.128]	[0.129]	[0.133]	[0.131]	[0.136]
Constant	0.889***	0.842***	0.625***	0.727***	0.700***
	[0.107]	[0.107]	[0.102]	[0.104]	[0.111]
Deflator	NYC	None	Fed	NYC	None
Ride measure	Circles	Circles	Blocks	Blocks	Blocks
Observations	59	59	59	59	59
R-squared	0.19	0.15	0.01	0.07	0.04

Panel C: Regressions, Excess return around FOMC meeting

	(1)	(2)	(3)	(4)	(5)
Ridership up dummy	–0.022***	–0.019***	–0.007	–0.011*	–0.009
	[0.006]	[0.006]	0.006	0.006	0.006
Constant	0.015***	0.012***	0.004	0.007	0.005
	[0.005]	[0.005]	0.004	0.005	0.005
Deflator	NYC	None	Fed	NYC	None
Ride measure	Circles	Circles	Blocks	Blocks	Blocks
Observations	59	59	59	59	59
R-squared	0.21	0.16	0.02	0.07	0.03

Notes: This table reproduces Table 7, except that we consider the blocks-based ride classification methodology developed in Section 2.3.2 in addition to the circles-based methodology developed in Section 2.3.1. We also consider alternative deflators besides Fed ridership, namely, all yellow cab ridership in Manhattan as well as no deflator. We sort the 59 FOMC meetings during our sample period based on whether (Fed-deflated) daily cab ridership between the Fed and the banks in our sample increased or decreased from the benchmark period (the second trading day after the previous FOMC meeting through FOMC event day 12, inclusive) and the pre-FOMC period (event days −11 through −2, inclusive). Our measure of returns around FOMC meetings is the cumulative excess return of the market over FOMC event days 1 through +1, inclusive. "Market up" refers to FOMC meetings in which the market return exceeded the risk-free rate, and "Market down" is defined analogously. Taxi cab ridership comes from the New York City Taxi and Limousine Commission. Our daily excess returns are taken from Ken French's website.