Rillor

METHODS NOTE

Event studies on public communications and commodity moves

How to measure a market's response to public statements: defensible timestamps, windows fixed in advance, random-time baselines, counted tests and the gap between timing and cause.



16 November 2026 Public communications and markets rillor.com/insights/event-studies-public-communications

Contents

01 Summary

02 Background

03 Method

04 Findings

05 Implications

06 Limits of this analysis

07 Sources

Summary

- An event study measures how prices respond to an event by comparing returns around it with a model of normal returns [1, 5]. Applied to public communications, such as central bank statements, OPEC decisions and posts by officials, much of the difficulty lies in time.
- Timestamps need a documented source. Gürkaynak, Sack and Swanson obtained Federal Reserve announcement times from the Fed itself [6]. For OPEC, official times are unavailable and decisions often leak early [9]. Market clocks have their own tolerances, such as the 50-millisecond standard for firms reporting to the Consolidated Audit Trail [18].
- Window choice trades completeness against noise. Published windows range from 30 minutes to several days [9]. Studies of Fed announcements use windows as tight as 10 minutes before to 20 minutes after [6, 7].
- Prices can move before the official time. For nine of 20 market-moving U.S. macroeconomic announcements, prices began moving about 30 minutes early, and that drift averaged about 40% of the total adjustment [12].
- Random-time baselines show whether a response is unusual. In one study, the variance of oil futures price changes on OPEC announcement days was more than three times that on matched control days [9].
- Communications studies generate many tests. Results need a stated correction, such as a familywise or false discovery rate procedure [14], and a higher bar for significance than a single test would need [15, 16].
- Timing is not causation. Confounding releases [6], information revealed by the speaker [7] and early trading [12, 13] can each produce a move after a statement.
- Rillor maintains a time-aligned corpus of public communications with random-baseline controls. This note covers the method.

Background

What an event study is

MacKinlay's survey in the Journal of Economic Literature defines the goal of an event study as measuring "the effects of an economic event on the value of firms," and explains that, if markets are rational, the effect of an event shows up in security prices soon after it, so it can be measured from prices observed over a relatively short period [1]. The approach has a long record in finance, from a 1969 study of how stock prices adjust to new information [2] to Brown and Warner's work on measuring security price performance and on event studies with daily returns [3, 4].

The standard daily design has a few steps. An estimation window before the event, typically 120 days, is used to fit a model of normal returns, most commonly a market model regressing the asset's return on an index. The model predicts normal returns over the event window. The abnormal return is the actual return minus the predicted one, and abnormal returns can be accumulated over the window and averaged across events. A common test divides the abnormal return by the root mean square error of the regression and compares the result with a t-distribution [5]. The same source notes that in volatile periods, such as 2008, this test flags too many returns as abnormal [5], a reminder that volatility changes from period to period.

From days to minutes

High-frequency event studies replaced daily windows with intraday ones. Gürkaynak, Sack and Swanson measured the response of bond yields and stock prices to Federal Reserve announcements using an intraday dataset going back to 1990. They found that two factors are needed to describe the response: one tied to the current federal funds rate target and one tied to the future path of policy, which is closely associated with FOMC statements. Statements had a much greater effect on longer-term Treasury yields than rate changes did [6]. Nakamura and Steinsson used interest rate changes in a 30-minute window around scheduled Fed announcements to estimate the effects of monetary policy [7]. Andersen, Bollerslev, Diebold and Vega built six years of real-time dollar exchange rate quotes, together with expectations and realizations of macroeconomic releases, and found that announcement surprises produce jumps in the conditional mean of exchange rates, with bad news having a greater effect than good news [8].

Communications as events

Three strands of published work are close to the problem of public communications and commodity prices.

- **Central bank statements.** The Fed's statements carry information beyond the rate decision itself [6], and the market's response can reflect what the statement reveals about the economy [7].
- **OPEC announcements.** Känzig identified oil supply news from high-frequency changes in oil futures prices around OPEC announcements, exploiting OPEC's institutional features. Negative supply news raised oil prices immediately [9].

- **Posts by officials.** Bianchi, Kind and Kung studied presidential posts criticizing the Federal Reserve, identifying their effect from a short window around the precise timestamp of each post. They found a negative and statistically significant average effect on the expected federal funds rate, growing with the horizon [11].

Commodities add a further question: whether prices respond to general macroeconomic news at all. Kilian and Vega regressed daily energy price changes on U.S. macroeconomic news and found no compelling evidence of feedback at daily or monthly horizons [10].

Method

This section sets out a protocol for event studies on public communications and commodity prices. Each step is tied to the published work that motivates it.

1. Define the event set before measuring anything

An event set is the list of communications under study and the rule that selected them. The rule is written down before any price data is examined. It states the speakers, channels and topics included, and whether non-events count. Gürkaynak, Sack and Swanson included FOMC meetings that ended with no change in policy, because a decision not to act sometimes surprised markets [6]. A study that keeps only the statements that seemed to matter has already selected on the outcome.

2. Record where each timestamp came from

Every event carries a timestamp and a record of its source. For Fed announcements, Gürkaynak, Sack and Swanson obtained dates and times from the office of the secretary of the Federal Reserve Board. They noted that press releases after scheduled meetings came at about 2:15 p.m. from 1995, and that before 1994, when the FOMC did not announce changes explicitly, the announcement time was that of the next open market operation [6]. For OPEC, Känzig notes that official announcement times are unavailable and that information about decisions often leaks before the official announcement [9]. Posts on social platforms carry platform timestamps, which allow the short windows used by Bianchi, Kind and Kung [11].

The source matters because it sets the uncertainty. A timestamp from the issuing institution, a platform timestamp and a time inferred from the first news report can differ, and the difference decides whether a tight window is meaningful.

3. Align every clock

Communications and prices come from different systems. Each needs to be mapped to one reference clock, with its tolerance stated. For U.S. trading, FINRA Rule 6820 requires firms to synchronize business clocks used for Consolidated Audit Trail reporting to within 50 milliseconds of the time kept by the NIST atomic clock, or one second for manual order events. The tolerance includes the difference from NIST, the transmission delay from the source and the drift of the firm's clock, and clocks must be synchronized every business day before market open [18].

A study whose windows are measured in seconds needs timestamps accurate to well under a second on both sides. A study with daily windows does not. A sound study records the tolerance for each source and uses no window narrower than the combined tolerance.

4. Build the market calendar

Scheduled releases and market hours shape every window.

- **Scheduled releases.** The BLS releases the Employment Situation at 8:30 a.m. [20]. The EIA's Weekly Petroleum Status Report comes out at 10:30 a.m. Eastern on Wednesdays, except in the holiday weeks the EIA lists, when it comes out on Thursday at 12:00 p.m. In the week of Veterans Day 2026, for example, the EIA's schedule moved the release of data for the week ending 6 November to Thursday, 12 November [19]. A study that assumes the regular schedule misplaces those events.
- **Market hours.** CME crude oil futures trade on Globex from Sunday to Friday, 5:00 p.m. to 4:00 p.m. Central Time, with a 60-minute break each day [21]. A statement made during the break or on a weekend meets the futures market only when it reopens.
- **Markets that do not close.** On-chain perpetuals on crude oil trade through weekends. On Saturday, 28 February 2026, oil-linked perpetuals on Hyperliquid rose more than 5% after strikes on Iran while traditional markets were closed [23]. Castle Labs found that over that weekend and the next, one such market captured part of the gap that CME later opened with, limited by a ±5% price bound [22]. Where such markets exist, a weekend statement has a measurable first response, with its own design limits.

Each event should be tagged with the scheduled releases and market sessions that overlap its window.

5. Fix windows and outcomes in advance

The window must be long enough to capture the response and short enough to exclude other news. Känzig sets out this trade-off and notes that common choices range from 30 minutes to multiple days [9]. Published choices include:

Study	Events	Window
Gürkaynak, Sack and Swanson [6]	FOMC announcements	Tight: 10 minutes before to 20 minutes after. Wide: 15 minutes before to 45 minutes after. Daily: close to close
Nakamura and Steinsson [7]	Scheduled Fed announcements	30 minutes
Känzig [9]	OPEC announcements	Daily, using futures settlement prices
Bianchi, Kind and Kung [11]	Presidential posts	A short window around each post's timestamp

Gürkaynak, Sack and Swanson give the case for tight windows: a narrow window around the announcement ensures the decision was not influenced by asset prices or other news over that interval, and shrinking it to an hour or less makes it much less likely that other significant events fall inside [6]. Känzig gives the case for a

wider one when timing is uncertain: OPEC communicates less clearly than a central bank, markets need time to process an announcement, official times are unavailable, and intraday data covers only the later part of his sample [9].

Three things should be fixed before any returns are examined: the window, the outcome measure and the price used at each edge of the window. The last matters in thin markets. Gürkaynak, Sack and Swanson used the most recent trade before the start of a window and the next available trade after its end when no trade fell exactly on the boundary [6].

6. Open the window early enough

Prices can move before the official time. Kurov, Sancetta, Strasser and Wolfe studied stock index and Treasury futures around 30 U.S. macroeconomic announcements. Of the 20 that moved markets, nine showed evidence of substantial informed trading: prices began moving in the direction later justified by the news about 30 minutes before the release, and the drift accounted for about 40% of the total price adjustment on average [12]. Bernile, Hu and Tang found evidence consistent with informed trading in E-mini S&P 500 futures during the embargo windows before FOMC announcements, with order imbalances in the direction of the later surprise, and no such evidence before the embargoes began or during lockups before nonfarm payroll, PPI and GDP releases [13].

The pre-event segment should therefore be measured separately from the post-event segment, and a response that begins before the timestamp reported as such.

7. Compare with random-time baselines

A move after a statement means little until it is compared with moves at times without one. Känzig compared daily oil futures price changes on OPEC announcement days with control days that contained no announcement but matched on weekday and week of the month in the months before each announcement. Price changes were more volatile on announcement days, with variance more than three times higher, and a Brown-Forsythe test confirmed that the difference was statistically significant. Känzig also found non-negligible background noise in the daily window and cites Nakamura and Steinsson's finding that background noise can lead to unreliable inference and overstate precision, especially with longer windows [9].

Rillor's communications corpus is time-aligned and built with random-baseline controls: windows drawn at random times that contain no event, against which each event's response is compared. The comparison answers the question any reader of an event study should ask first: how often does the market do this with no statement at all?

8. Correct for every test

A communications study runs many tests: speakers times assets times windows times outcome definitions. Each test is another chance for noise to pass a threshold.

- **Familywise error rate.** Controlling the probability of at least one false rejection across the family is the common approach [14].
- **False discovery rate.** Benjamini and Hochberg proposed controlling the expected proportion of false rejections among all rejections. It equals the familywise rate when every null hypothesis is true and is less strict otherwise, which gives more power when that is the appropriate goal [14].
- **Higher thresholds.** Harvey, Liu and Zhu argued that, after decades of factor testing, a new factor in finance should clear a t-statistic above 3.0, and that most claimed research findings in financial economics are likely false [15].
- **Data snooping.** White proposed a reality check for data snooping [17].

Ioannidis's list of conditions that make published findings less likely to be true applies directly: small effects, many tested relationships with little preselection, and flexibility in designs, definitions, outcomes and analysis [16]. Communications studies can have all of them. A sound study counts every test, including the ones that are dropped, and reports the correction used.

The same rule covers patterns found along the way. In Rillor's study of crude oil futures and an on-chain perpetual, 108 tests were corrected together with Bonferroni, and a concentration of returns around 08:30 Eastern release times that surfaced in the test data was recorded as a hypothesis for new data [24].

9. Separate timing from cause

Timing evidence and causal interpretation should be reported separately, with four alternative explanations checked for each result:

1. **Coincident releases.** Gürkaynak, Sack and Swanson found that eight FOMC announcements in their sample fell on the day of an employment report, seven on a GDP release, nine on a CPI release and five on a PPI release [6]. The market calendar from step 4 flags these overlaps.
2. **Information effects.** Nakamura and Steinsson found responses consistent with Fed announcements changing beliefs about economic fundamentals as well as about policy, and their model suggests these information effects are an important part of how monetary policy shocks affect output [7]. A market may respond to what a statement reveals about the speaker's information as well as to the stated intent.
3. **Early information.** Pre-release drift [12, 13] means some of the response can precede the event.
4. **Asymmetry.** Andersen and coauthors found that bad news has a greater effect than good news [8]. Pooling positive and negative events can hide a response that exists on one side only.

10. Report everything needed to check the result

A complete report states the event selection rule, the timestamp source and tolerance per event type, the market calendar used, the windows and outcomes fixed in advance, the random-baseline construction, the number of tests and the correction, and the null results alongside the positive ones.

Findings

This is a methods note. Its findings are what the published record shows about how these studies behave.

Window choice matters on crowded days

Gürkaynak, Sack and Swanson compared tight, wide and daily measures of policy surprises. The average absolute difference between the daily and intraday measures was only about 2 basis points, but on a handful of days, almost all of them employment report days, the gap was large: on 4 September 1992 the intraday surprise was 0 basis points and the daily one was -22, because the daily window also captured the response to a weak employment report. The two intraday windows agreed closely on every date [6]. Känzig made the opposite choice for OPEC, for documented reasons, and then tested how much background noise the daily window let in [9]. The right window depends on how precisely the event can be timed and how fast the market absorbs it.

Communications move prices measurably

FOMC statements had a larger effect on longer-term Treasury yields than rate decisions did [6]. Presidential posts about the Fed had a statistically significant average effect on the expected federal funds rate [11]. Oil futures price changes on OPEC announcement days were more than three times as variable as on matched control days [9]. Macroeconomic surprises produced jumps in high-frequency exchange rate data [8].

Some responses start early

Nine of 20 market-moving U.S. macroeconomic announcements showed substantial informed trading before release, with the early drift averaging about 40% of the total adjustment [12]. Evidence consistent with informed trading also appeared during FOMC embargo windows [13]. Studies that start their windows at the official timestamp can understate the total response and misattribute when it happened.

Null results are part of the record

Kilian and Vega found no compelling evidence that daily energy prices respond to U.S. macroeconomic news [10]. A commodity study that finds no response to a class of communications is consistent with published evidence and worth reporting.

Multiple testing is the common failure

Harvey, Liu and Zhu's review of hundreds of published factors concluded that most claimed findings in financial economics are likely false and that the bar for a new finding should be a t-statistic above 3.0 [15]. Communications studies, which can test many speakers, assets and windows, are exposed to the same problem.

Implications

For researchers

Write down the event set, timestamp sources, windows and outcomes before examining prices. Pair every event with random-time baselines from the same market session. Count every test, report the correction and keep patterns found along the way as hypotheses until new data confirms them [14, 15, 24].

For buyers of event data and analysis

Ask four questions of any claim that a statement moved a market. Where did the timestamp come from, and how accurate is it? Was the window fixed in advance? How does the move compare with random times? How many other tests were run? A claim that cannot answer them is an anecdote.

For builders of forecasting systems

Systems that read communications and act on them, including superintelligence (SI) agents, inherit every problem in this note. A model trained on statements and price moves without clock alignment, random baselines and a count of tests can learn noise as readily as a real effect. The controls that make a study credible also make a training set more reliable.

For risk managers in commodity markets

Communications can arrive when the main futures market is closed [21, 23]. Where 24/7 markets exist, they give a first response, bounded by their own design rules [22]. A risk process that waits for the futures open sees the move after the fact.

Limits of this analysis

- This note describes a method and the published evidence. It reports no results from Rillor's communications corpus.
- Most of the cited literature studies interest rates, equities and exchange rates. Commodity-specific evidence cited here comes mainly from OPEC announcements and energy price responses to macroeconomic news [9, 10].
- For several papers, this note relies on the published abstract and bibliographic record without the full text [1, 7, 8, 10, 11, 12, 13, 15]. Claims drawn from them are limited to what those records state.
- Clock standards and release schedules change. The figures cited are those published by FINRA, the EIA and the BLS when read [18, 19, 20].
- Nothing in this note is investment advice or a recommendation to trade on any communication.

Sources

1. A. Craig MacKinlay, Event Studies in Economics and Finance. Journal of Economic Literature 35(1), 13-39, March 1997, IDEAS record, viewed 15 November 2026. https://ideas.repec.org/a/aea/jeclit/v35y1997i1p13-39.html

2. Eugene F. Fama et al., The Adjustment of Stock Prices to New Information. International Economic Review 10(1), 1-21, February 1969, IDEAS record, viewed 15 November 2026. https://ideas.repec.org/a/ier/iecrev/v10y1969i1p1-21.html

3. Stephen J. Brown and Jerold B. Warner, Measuring security price performance. Journal of Financial Economics 8(3), 205-258, September 1980, IDEAS record, viewed 15 November 2026. https://ideas.repec.org/a/eee/jfinec/v8y1980i3p205-258.html

4. Stephen J. Brown and Jerold B. Warner, Using daily stock returns: The case of event studies. Journal of Financial Economics 14(1), 3-31, March 1985, IDEAS record, viewed 15 November 2026. https://ideas.repec.org/a/eee/jfinec/v14y1985i1p3-31.html

5. Wikipedia, Event study, last edited 22 August 2026, viewed 15 November 2026. https://en.wikipedia.org/wiki/Event_study

6. Refet S. Gürkaynak, Brian Sack and Eric T. Swanson, Do Actions Speak Louder Than Words? The Response of Asset Prices to Monetary Policy Actions and Statements. International Journal of Central Banking 1(1), May 2005, viewed 15 November 2026. https://www.ijcb.org/sites/default/files/journal/v1n1/ijcb-v1n1-do-actions-speak-louder-words-response-asset-prices-monetary-policy-actions-and-statements.pdf

7. Emi Nakamura and Jón Steinsson, High-Frequency Identification of Monetary Non-Neutrality: The Information Effect. Quarterly Journal of Economics 133(3), 1283-1330, 2018, IDEAS record, viewed 15 November 2026. https://ideas.repec.org/a/oup/qjecon/v133y2018i3p1283-1330..html

8. Torben G. Andersen, Tim Bollerslev, Francis X. Diebold and Clara Vega, Micro Effects of Macro Announcements: Real-Time Price Discovery in Foreign Exchange. American Economic Review 93(1), 38-62, March 2003, IDEAS record, viewed 15 November 2026. https://ideas.repec.org/a/aea/aecrev/v93y2003i1p38-62.html

9. Diego R. Känzig, The Macroeconomic Effects of Oil Supply News: Evidence from OPEC Announcements. American Economic Review 111(4), 1092-1125, April 2021, IDEAS record, viewed 15 November 2026. https://ideas.repec.org/a/aea/aecrev/v111y2021i4p1092-1125.html. Working paper version, MPRA Paper No. 106249, 2020, viewed 15 November 2026. https://mpra.ub.uni-muenchen.de/106249/

10. Lutz Kilian and Clara Vega, Do Energy Prices Respond to U.S. Macroeconomic News? A Test of the Hypothesis of Predetermined Energy Prices. Review of Economics and Statistics 93(2), 660-671, May 2011, IDEAS record, viewed 15 November 2026. https://ideas.repec.org/a/tpr/restat/v93y2011i2p660-671.html

11. Francesco Bianchi, Thilo Kind and Howard Kung, Threats to Central Bank Independence: High-Frequency Identification with Twitter. NBER Working Paper 26308, September 2019, revised September 2022, viewed 15 November 2026. https://www.nber.org/papers/w26308

12. Alexander Kurov, Alessio Sancetta, Georg Strasser and Marketa Halova Wolfe, Price Drift Before U.S. Macroeconomic News: Private Information about Public Announcements? Journal of Financial and Quantitative Analysis 54(1), 449-479, February 2019, IDEAS record, viewed 15 November 2026. https://ideas.repec.org/a/cup/jfinqa/v54y2019i01p449-

479_00.html

13. Gennaro Bernile, Jianfeng Hu and Yuehua Tang, Can information be locked up? Informed trading ahead of macro-news announcements. Journal of Financial Economics 121(3), 496-520, 2016, IDEAS record, viewed 15 November 2026. https://ideas.repec.org/a/eee/jfinec/v121y2016i3p496-520.html

14. Yoav Benjamini and Yosef Hochberg, Controlling the False Discovery Rate: A Practical and Powerful Approach to Multiple Testing. Journal of the Royal Statistical Society, Series B 57(1), 289-300, 1995, viewed 15 November 2026. https://academic.oup.com/jrsssb/article/57/1/289/7035855

15. Campbell R. Harvey, Yan Liu and Heqing Zhu, ... and the Cross-Section of Expected Returns. Review of Financial Studies 29(1), 5-68, 2016, IDEAS record, viewed 15 November 2026. https://ideas.repec.org/a/oup/rfinst/v29y2016i1p5-68..html

16. John P. A. Ioannidis, Why Most Published Research Findings Are False. PLOS Medicine 2(8), e124, 30 August 2005, viewed 15 November 2026. https://journals.plos.org/plosmedicine/article?id=10.1371/journal.pmed.0020124

17. Halbert White, A Reality Check for Data Snooping. Econometrica 68(5), 1097-1126, September 2000, IDEAS record, viewed 15 November 2026. https://ideas.repec.org/a/ecm/emetrp/v68y2000i5p1097-1126.html

18. FINRA, Rule 6820: Clock Synchronization. FINRA rulebook, viewed 15 November 2026. https://www.finra.org/rules-guidance/rulebooks/finra-rules/6820

19. U.S. Energy Information Administration, Weekly Petroleum Status Report: Release Schedule, viewed 15 November 2026. https://www.eia.gov/petroleum/supply/weekly/schedule.php

20. U.S. Bureau of Labor Statistics, Schedule of Releases for The Employment Situation, viewed 15 November 2026. https://www.bls.gov/schedule/news_release/empsit.htm

21. CME Group, Crude Oil Futures: Contract Specs. CME Group website, viewed 15 November 2026. https://www.cmegroup.com/markets/energy/crude-oil/light-sweet-crude.contractSpecs.html

22. Castle Labs (@noveleader and @francescoweb3), 432 Hours of Hyperliquid Oil Market Data: A Microstructure Comparison with CME WTI Futures. Castle Labs Research, 1 April 2026, viewed 15 November 2026. https://research.castlelabs.io/p/432-hours-of-hyperliquid-oil-market

23. Omkar Godbole, Oil-linked futures on Hyperliquid surge 5% after U.S.-Israel strike on Iran. CoinDesk, 28 February 2026, viewed 15 November 2026. https://www.coindesk.com/markets/2026/02/28/oil-linked-futures-on-hyperliquid-surge-5-after-u-s-israel-strike-on-iran

24. Rillor, Cross-venue price discovery: regulated crude futures and on-chain perpetuals. Rillor Insights, 9 October 2026, viewed 15 November 2026. https://rillor.com/insights/cross-venue-price-discovery-crude-perps

About this report

Published by Rillor on 16 November 2026. Superintelligence (SI) tools assist research and drafting, and every figure cites its source. Rillor builds and secures SI systems, supplies the data they learn from and delivers the infrastructure they run on. We apply SI to financial and compute markets.

Cite as: Rillor. Event studies on public communications and commodity moves. 16 November 2026. https://rillor.com/insights/event-studies-public-communications

Notices

Rillor is not a registered investment adviser or commodity trading advisor and does not provide investment or trading advice. Research, forecasts, data and software described on this site are for research and engineering use. Nothing here is an offer or recommendation to buy or sell any security, commodity interest or digital asset. Past or simulated results do not indicate future results.

Work with Rillor

Send the sources, assets and time windows your study needs, and we set out the records and baselines that cover them. Contact us at rillor.com/contact?type=communications or contact@rillor.com. More on public communications and markets: rillor.com/markets/forecasting#events