ISSN 2704-7407

UNIVERSITÀ CATTOLICA DEL SACRO CUORE
**Dipartimento di Economia e Finanza**

## Working Paper Series

# Does routine supervisory communication move volatility? Evidence from 285 ECB announcements

*Andrea Monticini*

**Working Paper n. 158**

**October 2026**



# Does routine supervisory communication move volatility? Evidence from 285 ECB announcements

**Andrea Monticini**
*Università Cattolica del Sacro Cuore*

Working Paper n. 158
October 2026

Dipartimento di Economia e Finanza
Università Cattolica del Sacro Cuore
Largo Gemelli 1 - 20123 Milano – Italy
tel: +39.02.7234.2976
e-mail: dip.economiaefinanza@unicatt.it

The Working Paper Series promotes the circulation of research results produced by the members and affiliates of the Dipartimento di Economia e Finanza, with the aim of encouraging their dissemination and discussion. Results may be in a preliminary or advanced stage. The Dipartimento di Economia e Finanza is part of the Dipartimenti e Istituti di Scienze Economiche (DISCE) of the Università Cattolica del Sacro Cuore.

# Does routine supervisory communication move volatility?

## Evidence from 285 ECB announcements

by

**Andrea Monticini***

Department of Economics and Finance
Università Cattolica del Sacro Cuore
Via Necchi 5
20123 Milan, Italy

email: **andrea.monticini@unicatt.it**

### Abstract

Supervisory authorities communicate continuously, yet the evidence on market reactions concerns a few landmark events. This paper tests whether the ordinary flow of European Central Bank (ECB) Banking Supervision communication, all 285 press releases published between 2013 and 2025, moves the daily volatility of euro area bank equity. Announcement effects are estimated in an exponential GARCH (EGARCH) model of bank index returns and assessed by randomization inference over placebo dates matched on calendar year and weekday, re-estimating the model on every assignment. The estimated increase in conditional volatility on announcement days is 2.4% ($p = 0.71$). Inverting the same test bounds the effect between $-3.1\%$ and $+4.4\%$, and an equivalence test rejects increases of 4% or more. Applied to 113 monetary policy decisions, the identical design detects a 9.0% increase ($p = 0.023$), although the difference between the two event classes is only marginally significant. No policy regime of the sample rejects the null, and the longest reproduces the bound. Standardized residuals show that announcement days are less often quiet but no more often turbulent. The routine flow of supervisory communication is not a material source of aggregate bank volatility.

Keywords: Banking supervision, ECB, Single Supervisory Mechanism, EGARCH, randomization inference, equivalence testing

JEL codes: C12, C58, G14, G21, G28

October 2026

* I am most grateful to Andrea Tarelli for useful comments.

## 1. Introduction

Prudential supervision is costly for banks and informative about them. Major institutional events have moved bank equity, including the establishment of European banking supervision and the 2014 Comprehensive Assessment: see Fiordelisi, Ricci and Stentella Lopes (2017), Loipersberger (2018) and Sahin and de Haan (2016). Stress test disclosures have also affected bank returns and volatility, as in Petrella and Resti (2013) and Durrani, Ongena and Ponte Marques (2022). These studies examine episodes selected for their importance. This paper asks a different question: whether the ordinary flow of supervisory communication – guidance, consultations, recommendations, sanctions and assessments, taken as issued and without selection – moves the volatility of bank equity.

A rejection is not foreordained. A release can move aggregate bank prices when three conditions hold together: its content is known to the supervisor before it is known to the market, it bears on the banking system rather than on a single institution, and it has not already been priced through an earlier consultation. The flow studied here contains releases of exactly that kind. The aggregate results of supervisory stress tests are computed by the ECB from submissions that no bank or investor observes in full. The recommendations of March and December 2020 on dividend distributions were issued without consultation and constrained payouts across the system; their withdrawal was announced in July 2021. The capital and operational relief of March 2020, the revisions of supervisory expectations on non-performing loans and the outcomes of the annual supervisory review fall in the same sample. No futures market priced their content, and several, the dividend recommendations among them, were not scheduled in advance. A view of supervision as a source of aggregate news therefore predicts that announcement days carry higher conditional variance. The opposite view – that supervisory procedures are deliberately gradual, with consultations preceding guidance and guidance preceding binding acts, so that by the day of any release its content is largely expected – predicts no response.

The two views are not symmetric in what the data must deliver to separate them. A positive estimate would favour the first; a statistically insignificant one favours the second only if the design can exclude effects of the size the first view implies. The paper is built around that requirement. It reports the set of effects consistent with the data, obtained by inverting the randomization test rather than by centring an interval on the estimate, tests formally whether increases of a stated size can be excluded, and benchmarks both against the response to monetary policy decisions, an event class known to move bank equity – see Ricci (2015) – estimated by the same procedure on the same data.

The object of interest is the second moment of bank index returns. Monetary policy studies identify surprises from futures prices – Kuttner (2001), Bernanke and Kuttner (2005), and for the euro area Monticini, Peel and Vaciago (2011) – but no comparable instrument exists for qualitative supervisory news. Since its sign is unidentified, positive and negative reactions can cancel in mean returns, whereas a variance test does not require a predicted direction.

The parameter estimated is the average same-day change in the conditional variance of the bank index on days on which a supervisory release is published. Because the event indicator is switched on for every release, the parameter averages over anticipated and unanticipated content: it measures the effect of the flow as issued, which is the quantity the question concerns, and not the effect of its surprise component, which no available instrument isolates. It excludes intraday adjustment completed before the close and bank-level effects concealed by aggregation.

The estimated change in conditional volatility is +2.4% ($p = 0.71$). The set of effects not rejected by the randomization test at 5% runs from −3.1% to +4.4%, and an equivalence test rejects any increase of 4% or more ($p = 0.041$). The same design detects a 9.0% increase on monetary policy decision days ($p = 0.023$); the difference between the two classes, however, is only marginally significant, with one-sided $p$ values between 0.024 and 0.068 across three tests, so the conclusion rests on the bound rather than on the contrast. Estimated separately in the three policy regimes of the sample – low and negative rates, the pandemic, and the tightening cycle with its aftermath – the supervisory coefficient is never distinguishable from its placebo distribution; the longest regime alone reproduces the full-sample bound, while the two shorter ones are too brief to detect even the monetary policy response. One feature of the data does distinguish announcement days: their standardized residuals are less often close to zero than on other days, while large residuals are no more frequent. Releases are therefore not non-events, but what they add sits in the centre of the return distribution, where it leaves the conditional variance unaffected.

Section 2 describes the data, Section 3 the model and the inference, Section 4 the results and Section 5 concludes. Appendix A documents the construction of the sample, Appendix B the selection of the variance specification, Appendix C the randomization design, Appendix D first-moment evidence and Appendix E the monetary policy surprises.

## 2. Data

Announcements are the 285 ECB Banking Supervision press releases that remain once attachments and periodic statistical publications, identified by the ECB's own publication type rather than by a reading of their content, are removed. Twenty-six administrative releases – appointments, relocations, nominations – are kept as a separate category. Releases before 16:00 CET are assigned to the same trading day and later ones to the next, yielding 242 substantive and 25 administrative event days.

Market data are daily from Bloomberg and comprise the EURO STOXX Banks and FTSE Italia All-Share Banks indices, their market benchmarks and the 10-year BTP–Bund spread. The euro area estimation contains 3,330 observations. The positive control uses the 113 monetary policy decisions of the Governing Council falling in that sample. Appendix A documents sample construction, publication times, overlapping events and alternative event-day conventions.

## 3. Empirical model and inference

Let $r_t$ be the log return on the bank index and $r_t^m$ the return on its market benchmark. The mean equation is

$$r_t = \mu + \rho\, r_{t-1} + \lambda\, r_t^m + \phi\, \Delta s_t + \varepsilon_t, \qquad \varepsilon_t = h_t^{1/2} z_t, \tag{1}$$

where $\Delta s_t$ is the change in the BTP–Bund spread and $z_t$ is a standardized Student-$t$ innovation. The contemporaneous market return isolates the bank-specific component.

The conditional variance follows an EGARCH – Nelson (1991) – augmented with the event dummy:

$$\log h_t = \omega + \alpha z_{t-1} + \gamma\,(|z_{t-1}| - \mathrm{E}|z_{t-1}|) + \beta_1 \log h_{t-1} + \beta_2 \log h_{t-2} + \delta D_t, \tag{2}$$

where $D_t$ marks event days and $\delta$ is the parameter of interest. Since (2) is written in logarithms, $\delta$ is a proportional effect: volatility changes by $\exp(\delta/2) - 1$.

The lag structure, the error distribution and the model family were fixed by a deterministic diagnostic rule set in advance and reported in Appendix B, which returns EGARCH(1,2) with Student-$t$ errors. No candidate fully removes residual ARCH, and the Nyblom (1989) statistic of the selected model, 3.66, exceeds its 1% critical value of 3.27: the parameters of the model are not jointly constant over the sample, although the event coefficient itself is (Section 4.3). Coefficient magnitudes are therefore model-dependent, and Section 4.3 estimates $\delta$ separately within sub-periods. Because one step of the selection rule consults the observed event labels, the entire procedure is repeated within 1,000 placebo assignments; the randomization $p$ value of the studentized event coefficient is then 0.58 (Appendix B).

*Randomization.* Placebo sets preserve the observed number of events within calendar year × weekday strata – monetary policy decisions are Thursdays in 96% of cases – and the true dates, overlaps included, remain in the support, while dates carrying only the other event class are excluded. The model is re-estimated on each assignment, and draws are screened on convergence alone, never on the log-likelihood. Tests are two-sided, obtained by doubling the smaller tail, because a release that resolves uncertainty rather than creating it lowers the conditional variance. Every row that carries an interval uses 10,000 draws. Appendix C gives the implementation.

*Intervals that match the test.* Let $d_b$ denote the placebo estimates and $q_\tau$ their $\tau$ quantile. A value $\delta_0$ is retained when $\hat{\delta} - \delta_0$ falls between $q_{0.025}$ and $q_{0.975}$, so that the 95% interval is

$$\left[\, \hat{\delta} - q_{0.975},\ \hat{\delta} - q_{0.025} \,\right]. \tag{3}$$

This is the inversion of the randomization test under the assumption that a constant effect shifts the distribution of $\hat{\delta}$ without changing its shape. The interval contains zero exactly when the test does not reject, and it inherits the location of the null distribution as well as its width. That location is not zero in the supervisory design: across the 10,000 placebo assignments $\hat{\delta}$ averages 0.032, a property of this sample

under a misspecified variance equation rather than a bias of the estimator, which a Monte Carlo experiment under $\delta = 0$ finds to be centred (Appendix C). An interval centred on $\hat{\delta}$ itself would refer to a different null from the test, and would place the upper bound too high by the same 0.032.

*Equivalence.* To ask whether an increase of a given size $\Delta$ can be excluded, the hypothesis $H_0\colon \delta \geq \delta_\Delta$, with $\delta_\Delta = 2\log(1+\Delta)$, is tested against $\delta < \delta_\Delta$. Under the same shift assumption its $p$ value is the share of placebo draws not exceeding $\hat{\delta} - \delta_\Delta$. Rejection establishes that the volatility response is smaller than $\Delta$; this is the one-sided half of the two one-sided tests procedure of Schuirmann (1987), the other half being uninformative here since a large negative effect is not at issue.

*Comparing the two event classes.* Whether supervisory releases and monetary policy decisions affect volatility differently is tested without the shift assumption. Both indicators enter (2) together, and the statistic is $\hat{\delta}_{\mathrm{mp}} - \hat{\delta}_{\mathrm{sup}}$. Under the null that the two classes have the same effect, the labels "supervisory" and "monetary policy" are exchangeable across event days. They are therefore permuted within calendar year $\times$ weekday cells, holding fixed the two days that carry both, and the model is re-estimated on each permutation. Only cells containing both classes contribute variation, which restricts the permutation set to 144 of the 351 event days and lowers the power of the test; it remains exact for the exchangeability hypothesis.

## 4. Results

### 4.1 *Baseline and positive control*

The baseline estimate is $\hat{\delta} = 0.047$, a 2.4% increase in conditional volatility. It lies close to the centre of the placebo distribution ($p = 0.71$), and the 95% interval obtained by inverting the test is $[-3.1\%, +4.4\%]$. Re-running the selection rule inside each assignment gives no evidence against the null either ($p = 0.58$). Restricting the sample to system-wide releases leaves the interval almost unchanged, $[-3.7\%, +4.8\%]$, and matching placebo dates on the calendar quarter as well narrows it to $[-3.3\%, +2.5\%]$. The one specification with a wider interval removes the market return from the mean equation, which costs about 1,400 points of log-likelihood: the upper bound then rises to 6.7%. Estimates are also small for the Italian index, under alternative event-day conventions and on the post-2018 subsample, for which publication times are known exactly (Appendix B). Mean returns on announcement days are indistinguishable from their placebo distribution ($p = 0.91$), and exactly half of the 242 days close negative (Appendix D).

The positive control gives a different answer. On monetary policy decision days conditional volatility rises by 9.0% ($p = 0.023$), with an interval of $[+1.2\%, +17.4\%]$; without the market return in the mean equation the increase is 12.2% ($p = 0.002$). Matching on the calendar quarter leaves the point estimate unchanged but widens its null distribution, and the control is then no longer significant ($p = 0.13$): monetary policy decisions are almost all Thursdays, and the finer strata leave so few eligible

**Table 1: Effect of announcement days on the conditional variance of bank index returns**

| | $\hat{\delta}$ | $p_{\mathrm{RI}}$ | $\Delta$ vol. | 95% interval |
|---|---|---|---|---|
| *Panel A. Supervisory announcements* | | | | |
| Baseline, euro area index | 0.047 | 0.71 | +2.4% | [−3.1%, +4.4%] |
| No market return in the mean | 0.044 | 0.30 | +2.2% | [−2.1%, +6.7%] |
| System-wide releases only | 0.050 | 0.76 | +2.6% | [−3.7%, +4.8%] |
| Quarter × weekday matched | 0.047 | 0.80 | +2.4% | [−3.3%, +2.5%] |
| *Panel B. Monetary policy decisions (positive control)* | | | | |
| Baseline, euro area index | 0.173 | **0.023** | +9.0% | [+1.2%, +17.4%] |
| No market return in the mean | 0.230 | **0.002** | +12.2% | [+4.5%, +20.9%] |
| Quarter × weekday matched | 0.173 | 0.13 | +9.0% | [−1.9%, +19.1%] |

*Notes.* $\hat{\delta}$ is the coefficient on the event dummy in the log-variance equation (2); $\Delta$ vol. is $\exp(\hat{\delta}/2) - 1$. $p_{\mathrm{RI}}$ is the two-sided randomization $p$ value, obtained by doubling the smaller tail of the placebo distribution, each row using 10,000 draws. The 95% interval inverts the same test, equation (3), and is mapped to volatility. Placebo assignments preserve calendar year × weekday counts; the quarter-matched rows additionally preserve the calendar quarter. Sample: 3,330 observations, 2013–2025; 242 supervisory and 113 monetary policy event days. Further specifications, robust standard errors and asymptotic $p$ values are reported in Appendix B.

dates that placebo sets often coincide with the true ones (Appendix C). Figure 1 plots the two randomization distributions of the baseline.

### 4.2 *How large an effect the data exclude*

Table 2 turns the non-rejection into a bound. The one-sided 95% upper limit for the supervisory effect is 3.8%: an equivalence test rejects increases of 4% or more at the 5% level ($p = 0.041$) and of 5% or more with $p = 0.012$. An effect as large as the one estimated for monetary policy, 9.0%, is rejected with $p < 0.001$. The bound is less than half of the benchmark point estimate, and it is the bound, not the insignificance of $\hat{\delta}$, that carries the result. Matching on the calendar quarter tightens it to 2.0%; only the misspecified fit without the market return widens it, to 6.0%.

Whether the two event classes differ is a separate question, and the answer is less clear-cut (Table 2, Panel B). With both indicators in the variance equation the estimated difference is $\hat{\delta}_{\mathrm{mp}} - \hat{\delta}_{\mathrm{sup}} = 0.136$. Permuting the labels between event days gives a two-sided $p$ value of 0.11 when labels are exchanged within calendar year × weekday cells and 0.13 when they are exchanged across all event days; against the one-sided alternative that monetary policy moves volatility more, the $p$ values are 0.054 and 0.068. Comparing the two separate placebo distributions of Table 1, which relies on the shift assumption, gives 0.047. The data therefore establish that the supervisory effect is small in absolute terms and that the monetary policy effect is distinguishable

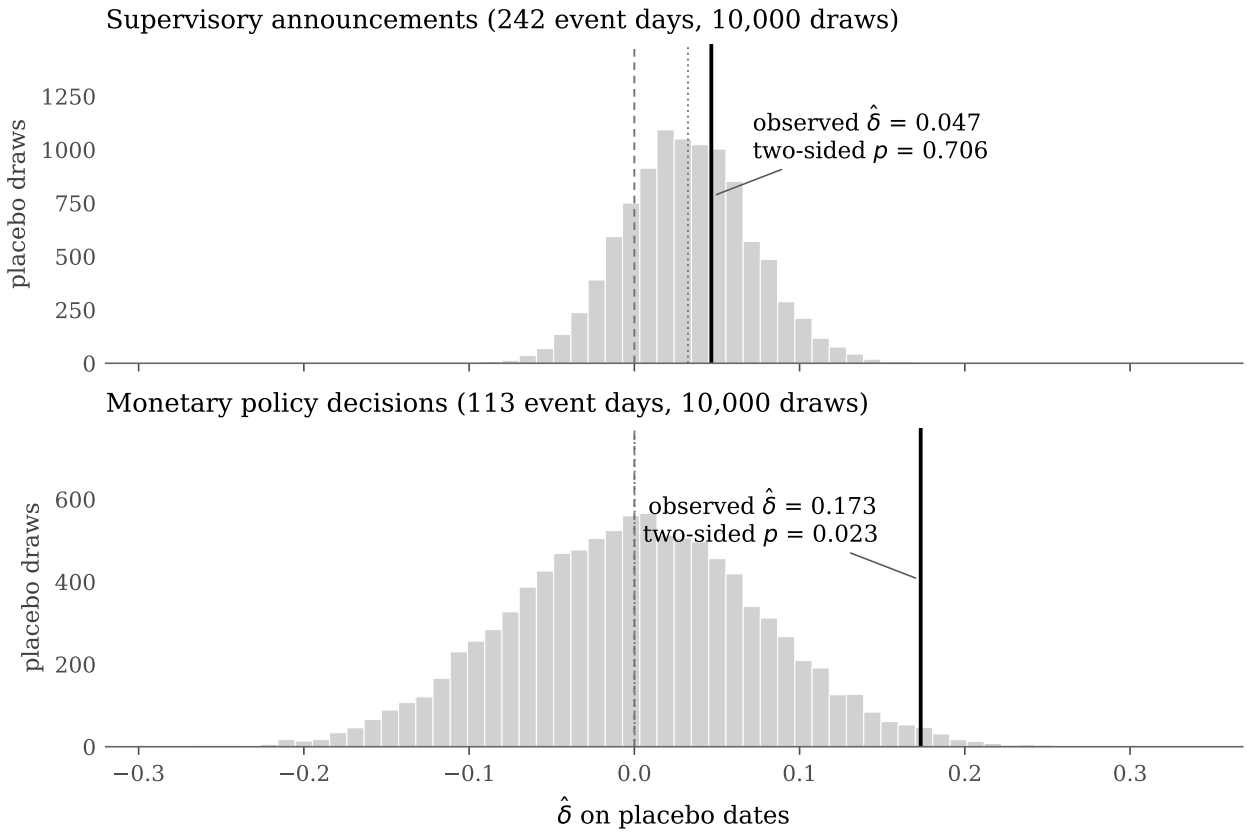


**Figure 1: Randomization distributions of $\hat{\delta}$**

The two panels show the distribution of $\hat{\delta}$ over sets of placebo dates matched on year and weekday, with the observed estimate marked by the solid vertical line and the two-sided $p$ value of Table 1 alongside it. The dashed grey line is zero and the dotted grey line the mean of the placebo distribution; in the lower panel the two coincide. Both panels share the same horizontal scale.

from zero, but they establish only marginally that the two differ. The conclusion of this paper rests on the first statement, the bound, and not on the contrast.

### **4.3** *Stability over time*

The Nyblom statistic rejects parameter constancy over the full sample, and the sample spans three monetary regimes. Table 3 re-estimates the whole model in each, with its own parameters and its own placebo distribution matched within the regime. Within regimes the stability statistic lies between 2.73 and 3.05 for the supervisory fits, below the 1% critical value in every case though above the 5% value in two of the three, against 3.66 on the full sample: much of the instability of the full-sample fit arises between regimes rather than within them. The individual statistics locate it. Over the full sample the statistic for $\delta$ alone is 0.13, well below the 10% critical value of 0.35 for a single parameter; the instability lies in the mean equation – its intercept (1.25), the spread coefficient (0.65) and the market return (0.54) – and in the asymmetry

**Table 2: Bounds, equivalence and the difference between event classes**

| *Panel A. Supervisory announcements: $H_0$ increase in volatility $\geq \Delta$* | | | | |
|---|---|---|---|---|
| | upper bound | $\Delta = 4\%$ | $\Delta = 5\%$ | $\Delta = 9.0\%$ |
| Baseline, euro area index | 3.8% | 0.041 | 0.012 | $< 0.001$ |
| System-wide releases only | 4.1% | 0.056 | 0.021 | $< 0.001$ |
| Quarter × weekday matched | 2.0% | 0.001 | $< 0.001$ | $< 0.001$ |
| No market return in the mean | 6.0% | 0.22 | 0.11 | 0.001 |
| *Panel B. Monetary policy against supervisory: $H_0$ $\delta_{\text{mp}} = \delta_{\text{sup}}$* | | | | |
| | $\hat{\delta}_{\text{mp}} - \hat{\delta}_{\text{sup}}$ | $p$ two-sided | $p$ one-sided | draws |
| Label permutation, year × weekday cells | 0.136 | 0.11 | 0.054 | 2,000 |
| Label permutation, all event days | 0.136 | 0.13 | 0.068 | 2,000 |
| Separate placebo distributions | 0.127 | 0.047 | 0.024 | 10,000 |

*Notes.* Panel A: the upper bound is the one-sided 95% limit for the increase in volatility, $\exp\{(\hat{\delta} - q_{0.05})/2\} - 1$; the remaining columns are $p$ values of the equivalence test of Section 3 for the stated margin, 9.0% being the point estimate for monetary policy decisions; each row uses the 10,000 placebo draws of Table 1. Panel B: in the first two rows both indicators enter the variance equation together and the event labels are permuted between event days, within calendar year × weekday cells or across all event days, holding fixed the two days that carry both; the test is exact for the hypothesis that the labels are exchangeable. The third row compares the two estimates of Table 1, each from its own model, with the difference of their independent placebo distributions, under the shift assumption of equation (3). One-sided $p$ values are against $\delta_{\text{mp}} > \delta_{\text{sup}}$.

term $\alpha$ of the variance equation (0.58). The parameter this paper is about is among the most stable in the model.

The first regime, 2013–2019, holds 113 of the 242 event days. There the supervisory estimate is +1.0% ($p = 0.51$) and its interval, $[-2.9\%, +4.3\%]$, reproduces the full-sample bound almost exactly. The pandemic and tightening regimes are shorter, 643 and 895 observations, and their intervals are wide. Point estimates there are larger, +10.2% and +4.3%, but neither is distinguishable from its placebo distribution ($p =$ 0.55 and 0.36), and in the same two regimes the monetary policy control is not detected either ($p = 0.74$ and 0.14), with point estimates of +9.1% and +28.2%. The later regimes therefore lack the power to speak about either event class: they do not overturn the full-sample conclusion, nor can they confirm it on their own. What the split does establish is that the bound is not an artefact of pooling regimes with offsetting effects, since the one regime long enough to measure it reproduces it.

### 4.4 *The distribution of standardized residuals*

The conditional variance is estimated by a likelihood that the largest residuals dominate, and the tests above ask whether announcement days carry more of them. A test that does not depend on re-estimation asks a broader question. The model is

**Table 3: The event coefficient within policy regimes**

| | obs. | events | $\hat{\delta}$ | $\Delta$ vol. | $p_{\mathrm{RI}}$ | 95% interval | Nyblom |
|---|---|---|---|---|---|---|---|
| *Panel A. Supervisory announcements* | | | | | | | |
| Full sample | 3,330 | 242 | 0.047 | +2.4% | 0.71 | [−3.1%, +4.4%] | 3.66 |
| 2013–2019 | 1,792 | 113 | 0.019 | +1.0% | 0.51 | [−2.9%, +4.3%] | 2.77 |
| 2020–June 2022 | 643 | 55 | 0.194 | +10.2% | 0.55 | [−14.0%, +34.5%] | 3.05 |
| July 2022–2025 | 895 | 74 | 0.084 | +4.3% | 0.36 | [−6.8%, +23.6%] | 2.73 |
| *Panel B. Monetary policy decisions* | | | | | | | |
| Full sample | 3,330 | 113 | 0.173 | +9.0% | 0.023 | [+1.2%, +17.4%] | 4.25 |
| 2013–2019 | 1,792 | 64 | 0.134 | +6.9% | 0.12 | [−1.2%, +17.5%] | 3.13 |
| 2020–June 2022 | 643 | 21 | 0.174 | +9.1% | 0.74 | [−22.1%, +50.4%] | 2.78 |
| July 2022–2025 | 895 | 28 | 0.498 | +28.2% | 0.14 | [−6.8%, +50.3%] | 2.91 |

*Notes.* Each row re-estimates the model of Section 3 on the stated period, with its own parameters. $p_{\mathrm{RI}}$ and the 95% interval, equation (3), come from 1,000 placebo assignments matched on calendar year and weekday within the period (10,000 for the full sample). Nyblom is the joint parameter-stability statistic of the fitted model, with asymptotic critical values of 2.49, 2.75 and 3.27 at 10%, 5% and 1%; for a single parameter they are 0.35, 0.47 and 0.75, and the full-sample statistic for $\delta$ alone is 0.13. The regimes are fixed on the monetary policy calendar: low and negative policy rates up to the end of 2019, the pandemic until the first rate increase of July 2022, and the tightening cycle with the subsequent easing to the end of the sample.

estimated once with $\delta = 0$ and its standardized residuals $z_t$ are held fixed; only the event labels are then permuted, over the matched draws used throughout, and three statistics are compared between labelled and other days: the mean of $z_t^2$, the quantity the event dummy is meant to capture; the mean of $|z_t|$, less sensitive to the tails; and the mean rank of $z_t^2$, which responds to a shift of the whole distribution. The residuals still depend on the fitted variance model, so the test is lighter than the others, not free of them.

For monetary policy decisions all three statistics agree (Table 4). For supervisory releases they do not: the mean of $z_t^2$ is indistinguishable from its null ($p = 0.26$), the mean of $|z_t|$ is borderline ($p = 0.054$) and the mean rank rejects ($p = 0.012$). The quantiles show why. The ratio of announcement-day to other-day quantiles of $|z_t|$ is 1.44 at the tenth percentile, 1.35 at the quarter, 1.15 at the median, 1.08 at the third quarter and 1.02 at the ninetieth: the lower part of the distribution moves away from zero, while the upper part is almost unchanged. The share of days with $|z_t| < 0.25$ falls from 23.5% to 16.5% ($p = 0.006$); the share with $|z_t| > 2$ is 5.0% against 5.1% ($p = 0.79$). On monetary policy days the pattern is reversed, the upper-tail share rising by 9.4 points ($p = 0.026$) and the lower-tail share not moving significantly.

Supervisory announcement days are therefore less often quiet, but they are not more often turbulent. The finding qualifies the conclusion drawn from the variance equation

**Table 4: Standardized residuals on event days**

| | supervisory | | monetary policy | |
|---|---|---|---|---|
| | stat. | $p$ | stat. | $p$ |
| *Panel A. Difference from the null, in null standard deviations* | | | | |
| Mean of $z_t^2$ | 1.08 | 0.26 | 2.69 | 0.012 |
| Mean of $\lvert z_t \rvert$ | 2.01 | 0.054 | 2.83 | 0.005 |
| Mean rank of $z_t^2$ | 2.52 | 0.012 | 2.24 | 0.024 |
| *Panel B. Change in the share of days, percentage points* | | | | |
| $\lvert z_t \rvert < 0.25$ | −7.0 | 0.006 | −4.6 | 0.41 |
| $\lvert z_t \rvert > 2$ | −0.1 | 0.79 | +9.4 | 0.026 |

*Notes.* Standardized residuals from the model of Section 3 estimated once with $\delta = 0$. Panel A reports the observed difference between event days and other days, standardized by the mean and standard deviation of its permutation distribution; Panel B the difference in the share of days below and above the stated thresholds. Two-sided $p$ values from 20,000 permutations of the event labels matched on calendar year and weekday. Event days: 242 supervisory, 113 monetary policy.

rather than contradicting it: a shift confined to the centre of the distribution raises the mean of $z_t^2$ only slightly, which is why the event coefficient does not register it, and it is consistent with releases that carry some information but not enough to produce the large moves through which conditional volatility rises. The evidence is weaker than the main tests. Three statistics on two event classes make multiplicity real – a Bonferroni adjustment over the three supervisory statistics gives 0.036 for the mean rank – and the result does not survive finer matching intact: adding the calendar quarter to the strata moves the mean-rank $p$ value to 0.072 and the lower-tail one to 0.035.

## 5. Concluding remarks

The ordinary flow of ECB supervisory communication does not raise the same-day volatility of euro area bank equity by more than a few per cent. The data reject increases of 4% or more, while the same design detects a 9% increase on monetary policy decision days, although the contrast between the two classes is only marginally significant. No policy regime rejects the null, and the longest reproduces the full-sample bound, so the result is not an average of offsetting periods; the two shorter regimes are too brief to be informative for either event class. What announcement days do show is fewer quiet sessions, a shift in the centre of the return distribution that the conditional variance does not capture.

The result is consistent with supervisory procedures that are gradual by design, in which consultations precede guidance and guidance precedes binding acts, so that by the day of most releases their content is expected. It is also consistent with releases whose economic content is limited. Daily aggregate data cannot separate these two readings, and the paper does not claim to. What it establishes is narrower and, for

the evaluation of supervisory communication, more useful: the flow as issued is not a material source of aggregate bank volatility, a statement that delimits rather than contradicts the evidence from exceptional institutional and stress test events.

Three things bound the conclusion: aggregation across banks, the daily frequency, and a variance equation that no candidate fits. Because the same estimation procedure is applied to every placebo assignment, inference does not rest on the misspecified model's asymptotic approximation, although magnitudes remain model-dependent; the sub-period estimates and the residual-based test are the two checks that rely least on a single parametric fit.

## References


Altavilla, C., L. Brugnolini, R. S. Gürkaynak, R. Motto, and G. Ragusa (2019). "Measuring Euro Area Monetary Policy", *Journal of Monetary Economics*, **108**, 162–179.

Bernanke, B. S. and K. N. Kuttner (2005). "What Explains the Stock Market's Reaction to Federal Reserve Policy?", *Journal of Finance*, **60**, 1221–1257.

Bollerslev, T. and J. M. Wooldridge (1992). "Quasi-maximum Likelihood Estimation and Inference in Dynamic Models with Time-varying Covariances", *Econometric Reviews*, **11**, 143–172.

Durrani, A., S. Ongena, and A. Ponte Marques (2022). "The Certification Role of the EU-wide Stress Testing Exercises in the Stock Market. What Can We Learn from the Stress Tests (2014–2021)?", *ECB Working Paper* 2711.

Fiordelisi, F., O. Ricci, and F. S. Stentella Lopes (2017). "The Unintended Consequences of the Launch of the Single Supervisory Mechanism in Europe", *Journal of Financial and Quantitative Analysis*, **52**, 2809–2836.

Kuttner, K. N. (2001). "Monetary Policy Surprises and Interest Rates: Evidence from the Fed Funds Futures Market", *Journal of Monetary Economics*, **47**, 523–544.

Loipersberger, F. (2018). "The Effect of Supranational Banking Supervision on the Financial Sector: Event Study Evidence from Europe", *Journal of Banking and Finance*, **91**, 34–48.

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

Monticini, A., D. Peel, and G. Vaciago (2011). "The Impact of ECB and FED Announcements on the Euro Interest Rates", *Economics Letters*, **113**, 139–142.

Nelson, D. B. (1991). "Conditional Heteroskedasticity in Asset Returns: A New Approach", *Econometrica*, **59**, 347–370.

Nyblom, J. (1989). "Testing for the Constancy of Parameters over Time", *Journal of the American Statistical Association*, **84**, 223–230.

Petrella, G. and A. Resti (2013). "Supervisors as Information Producers: Do Stress Tests Reduce Bank Opaqueness?", *Journal of Banking and Finance*, **37**, 5406–5420.

Ricci, O. (2015). "The Impact of Monetary Policy Announcements on the Stock Price of Large European Banks during the Financial Crisis", *Journal of Banking and Finance*, **52**, 245–255.

Sahin, C. and J. de Haan (2016). "Market Reactions to the ECB's Comprehensive Assessment", *Economics Letters*, **140**, 1–5.

Savor, P., and M. Wilson (2013). "How Much Do Investors Care About Macroeconomic Risk? Evidence from Scheduled Economic Announcements", *Journal of Financial and Quantitative Analysis*, **48**, 343–375.

Schuirmann, D. J. (1987). "A Comparison of the Two One-Sided Tests Procedure and the Power Approach for Assessing the Equivalence of Average Bioavailability", *Journal of Pharmacokinetics and Biopharmaceutics*, **15**, 657–680.

# Appendices

## A. Construction of the announcement sample

Announcements were assembled twice and independently. The first extraction scraped the press release archive of ECB Banking Supervision page by page, taking each date from the annual index and from the file name. The second queried the publication database that the website loads to render its index, whose records carry a `pub_timestamp` field; both extractions are reproduced by the script accompanying this paper.

The two agree completely. Restricted to press releases published up to 19 December 2025 they return 297 and 295 documents, *agree on every one of the 294 dates they share*, and differ only over four attachments: three FAQ and annex documents that the scrape collects and the database does not list as publications, and one FAQ the database carries separately. Every one of the 285 releases in the sample appears in both. For the 203 releases whose publication time is recorded the two sources give *the same time to the minute, in all 203 cases*, and the remaining 82 are flagged as untimed by both. A third route, reading the date declared on each release page itself, agrees with the scrape on every shared date as well.

The database also settles two questions that would otherwise rest on judgement. The nine periodic statistical publications excluded from the sample are exactly the nine press releases that the ECB's own classification assigns to the publication type reserved for supervisory banking statistics – the correspondence is exact in both directions, no release of that type is kept and no release of another type is dropped. The three attachments excluded are not listed in the database as publications at all. Neither exclusion required a judgement about content.

Publication times come from this database alone: the release pages carry a date but no time. Of the 82 releases treated as having no recorded time, 70 are stored at midnight and the remaining twelve between 00:10 and 02:00, which is why any stamp before 04:00 is read as a placeholder rather than as a publication time.

From 297 documents three attachments are dropped, since they duplicate the event date of a parent release, together with nine periodic statistical publications ("Supervisory banking statistics" releases), leaving 285.

*Event days.* The website records publication times for 203 of the 285 releases, essentially all of them from mid-2017 onwards. A release published before 16:00 CET is assigned to the same trading day, a later release to the next trading day, and a weekend release to the following Monday. Of the 203 timed releases two fall before 08:00, 113 between 08:00 and 12:00, 26 between 12:00 and 16:00, and 62 at or after 16:00, thirty-two of them at 18:00. Those 62 are all moved forward, but one falls on a Saturday and is counted under the weekend rule, so 61 carry the late-release rule. The 82 releases whose time is not recorded are assigned to the same trading day, which is an assumption and not a measurement: had they followed the timed distribution,

roughly a quarter of them would belong to $t+1$ instead. Two estimates in Table B.3 bound the consequence: the specification that moves every event to $t+1$, and the post-2018 subsample, in which publication times are known and no assumption is needed. Twelve days carry more than one release, so the dummy is constructed per day rather than per release. The 285 releases map to 242 substantive and 25 administrative event days: the twenty-six administrative releases fall on twenty-six days, one of which (15 December 2016) also carries a substantive release and is therefore counted with the substantive days. The 16:00 cut-off is the one used throughout the paper; the robustness row labelled "cut-off at 17:30" re-assigns to the same trading day the ten releases published between 16:00 and 17:30, and changes the estimate from 0.0466 to 0.0467.

*Administrative releases.* Twenty-six releases concern appointments, office relocations and board nominations and carry no prudential content. They are kept as a separate category rather than pooled.

*Scope.* Releases are classified as institution-specific or system-wide from the titles alone, never from returns, by two independent criteria: verbs identifying action on a single intermediary (*sanctions*, *failing or likely to fail*, *asset quality review of*, *withdraws banking licence*, *temporary administrators*), and the presence of a bank's proper name in the title. Fifty of the 259 substantive releases (19%) are institution-specific. They fall on 46 distinct dates, 45 of them trading days; four of those also carry a system-wide release and are assigned to the system-wide group, leaving the 41 institution-specific and 201 system-wide event days.

## B. Selection of the variance specification

The lag structure, the error distribution and the model family were selected by a rule fixed in advance and applied in this order: among specifications that converge, that respond to the event regressor and that are numerically stable, take the one with the lowest Nyblom (1989) joint stability statistic among those not rejected by an ARCH LM test at 5%; if no candidate survives that test, take the lowest Nyblom among the stable ones. The rule never looks at the sign of $\hat{\delta}$, at its magnitude or at its $p$ value. It is not, however, blind to the event dates: the screen described below establishes whether a family responds to a variance regressor at all, and to do so it estimates each family once with the true dummy.

The second branch is the one that fires. No candidate passes the ARCH LM test, a fact recorded in the selection output rather than discovered afterwards, and the rule therefore selects on stability alone. This is stated here because it is the weakest link in the design: the variance equation is misspecified for every family examined, which is why the paper treats $\hat{\delta}$ as provisional in magnitude and relies on randomization rather than asymptotics for inference.

*A prior screen.* Each family is estimated twice, once with the true dummy and once with the dummy shifted forward by one day. The gap $\Delta \log L$ between the two fits

**Table B.1: Candidate specifications**

| | $\hat{\delta}$ | AIC | ARCH LM | Nyblom | $\Delta \log L$ | usable |
|---|---|---|---|---|---|---|
| eGARCH(2,1)-*t* | 0.033 | 2.6494 | 0.029 | 4.20 | 8605.07 | no |
| NAGARCH(1,1)-*t* | 0.000 | 2.6520 | 0.005 | 5.10 | 0.00 | no |
| **eGARCH(1,2)-*t*** | **0.047** | **2.6522** | **0.003** | **3.66** | **0.29** | **yes** |
| gjrGARCH(1,1)-*t* | 0.000 | 2.6530 | 0.003 | 4.60 | 0.00 | no |
| apARCH(1,1)-*t* | 0.000 | 2.6531 | 0.002 | 4.43 | 0.00 | no |
| sGARCH(1,1)-*t* | 0.000 | 2.6535 | 0.002 | 3.96 | 0.00 | no |
| TGARCH(1,1)-*t* | 0.000 | 2.6543 | 0.001 | 4.09 | 0.00 | no |
| AVGARCH(1,1)-*t* | 0.012 | 2.6544 | 0.000 | 3.67 | 0.27 | yes |
| eGARCH(1,1)-skew*t* | 0.018 | 2.6853 | 0.001 | 17.84 | 14017.35 | no |
| eGARCH(1,1)-*t* | 0.164 | 2.6935 | 0.000 | 15.93 | 528.16 | no |

*Notes.* All ten candidates converge. ARCH LM is the $p$ value of the test at five lags; none exceeds 0.05. Nyblom is the joint parameter-stability statistic, the criterion on which the rule selects: lower is better, and for the selected eleven-parameter model the asymptotic critical values are 2.49, 2.75 and 3.27 at the 10%, 5% and 1% levels. $\Delta \log L$ is the gap in log-likelihood between the fit with the true dummy and the fit with the dummy shifted by one day; a gap of zero means the regressor does not enter the fit, a large gap means the optimum is not stable. "Usable" is true when $\Delta \log L < 50$ and $\hat{\delta}$ is not zero to four decimals; the second condition is applied here rather than by the estimation code, as explained above. Reported $\hat{\delta}$ is rounded to three decimals; the five entries shown as 0.000 are zero to four decimals and pair with $\Delta \log L$ of zero. The selected specification is in bold.

separates the two ways a candidate fails. A gap in the hundreds or thousands means the optimiser lands in different regions depending on which dummy it is given, which matters because the randomization test refits the model a thousand times; the screen discards those, and they are the three specifications marked unstable in Table B.1. A gap of exactly zero means the opposite: the fit does not respond to the dummy at all. This happens in five of the six non-logarithmic families, where the optimum for the coefficient on a variance regressor lies on the boundary that the positivity of the conditional variance imposes. A boundary optimum is not a failure of identification; it means that under those constraints the maximum is at the edge of the parameter space, where symmetric inference does not apply, and that the family cannot deliver an unconstrained estimate of $\delta$.

One point of implementation concerns the replication files. The screen tests for a non-zero coefficient as $|\hat{\delta}| > 10^{-8}$, and the five degenerate families return values that are zero to four decimals but not below that threshold, so they passed the test and entered the pool from which the Nyblom criterion selected. They are excluded here on the stronger evidence of $\Delta \log L = 0$, which is decisive and which the threshold should have encoded. The selection is unaffected: eGARCH(1,2)-*t* has the lowest Nyblom statistic among all stable candidates, degenerate ones included.

*The screen repeated on a placebo.* Because the screen uses the true event dates, it

was repeated with a placebo dummy of the same size drawn from event-free days. Restricted to the branch that fires on the true data – lowest Nyblom among the usable candidates – the placebo screen selects eGARCH(1,2)-$t$ again, with a Nyblom statistic of 3.77 against 15.34 for the only other usable family. Applied in full, however, the rule returns eGARCH(2,1)-$t$: under the placebo dummy that family passes the ARCH LM test ($p = 0.14$), the first branch of the rule therefore fires, and it selects on ARCH LM rather than on stability. Two things follow. The criterion actually used here is invariant to the event dates. The other branch is not robust, since whether it fires at all turns on a single test crossing 0.05, and a rule for a future paper should not leave that hinge in place.

What the screen cannot do is manufacture the null. It keeps the families whose fit *responds* to the true dummy and discards those whose optimum ignores it, so the step that looks at the event dates selects towards a non-zero $\hat{\delta}$, not towards zero. The main estimate is indistinguishable from its placebo distribution all the same. Contamination from selection therefore works against the reported result rather than for it, and the same holds for the alternative usable candidate: AVGARCH(1,1)-$t$ would have returned 0.012 in place of 0.047.

*The rule re-run inside every draw.* The argument above bounds the problem; it does not remove it. A fully valid randomization test re-runs the selection under every label assignment, so that the observed statistic enjoys no conditioning the placebo statistics lack. For each of 1,000 matched placebo assignments the entire rule is applied afresh – all ten families, each fitted with the placebo dummy and with it shifted by one day, Nyblom and ARCH LM recomputed – and the coefficient is taken from whichever specification the rule selects *for those labels*.

The statistic is the robust $t$ of the event coefficient, $\hat{\delta}$ over its standard error, and not $\hat{\delta}$ itself. The candidates are not all logarithmic: in eGARCH $\delta$ is a proportional effect, while in AVGARCH the dummy enters the conditional standard deviation, so a null distribution of $\hat{\delta}$ would mix parameters that are not the same quantity. Two draws illustrate the point, returning $\hat{\delta}$ of 1.66 and $-32.2$ with standard errors of 10.6 and 67.1; as $t$ they are 0.16 and $-0.48$, which is what they are worth. This is therefore a check on the validity of the design and not a restatement of Table 1, whose statistic is different.

The rule reselects eGARCH(1,2)-$t$ in 67.8% of draws. It picks eGARCH(1,1)-$t$ in 16.1%, AVGARCH(1,1)-$t$ in 10.8%, eGARCH(2,1)-$t$ in 4.7% and eGARCH(1,1)-skew$t$ in 0.6%: the specification does depend on the labels, but in two draws out of three it does not change. The observed $t$ is 1.093 against a null distribution with mean 0.701 and standard deviation 0.773, so the estimate lies half a standard deviation above the centre and 290 draws out of 1,000 exceed it. The two-sided $p$ value is 0.58, with a Monte Carlo error of 0.03. As in the fixed-specification test the null is not centred on zero: the procedure returns a positive $t$ on any matched set of 242 days, which is the same property recorded for $\hat{\delta}$ above.

The comparison measures the size of the contamination. Holding the specification fixed the same 1,000 draws give $p = 0.69$; re-selecting inside each draw gives 0.58.

The conditioning on the true dates is therefore worth about a tenth of a $p$ value: the conditioning matters, but the conclusion does not turn on it, since 0.58 is as far from any conventional threshold as 0.69. The exercise covers the selection of the variance specification. The set of ten candidates and the mean equation were not chosen by a rule, and no re-randomization can repair that.

Three features of Table B.1 bear on Section 4. The information criteria are reported as `rugarch` computes them, divided by the number of observations; in ordinary units they span 147 points across the ten candidates and 17 across the best eight, so they do discriminate. They do not, however, select: AIC ranks eGARCH(2,1)-$t$ first and the Schwarz criterion sGARCH(1,1)-$t$, and both are excluded before any criterion is consulted, the first as numerically unstable and the second because it does not respond to the variance regressor. Between the two usable candidates the two criteria agree, and by the same margin since both have eleven parameters, favouring eGARCH(1,2)-$t$ over AVGARCH(1,1)-$t$ by 7.3 points. Of the two usable candidates, eGARCH(1,2)-$t$ has the lower Nyblom statistic, 3.66 against 3.67 – a margin so thin that the alternative deserves stating: AVGARCH(1,1)-$t$ would have given $\hat{\delta} = 0.012$ instead of 0.047, that is a smaller estimated effect, so the selection rule did not pick the specification most favourable to a null. And no candidate passes the ARCH LM test, which establishes only that none of the families examined describes the residual dynamics adequately – not that a family outside this set would fail as well.

*The mean equation.* The selection rule above concerns the variance equation. The mean equation was not selected by a rule, and one of its terms deserves scrutiny of its own: the contemporaneous market return $r_t^m$. If a supervisory release moved the entire market, including $r_t^m$ would remove that movement from $\varepsilon_t$ before the conditional variance is estimated, and $\hat{\delta}$ would be biased towards zero by construction. Table B.2 estimates the same variance equation under five mean equations, with the supervisory dummy and with the monetary policy dummy.

Three things follow. The market return is worth about 1,437 log-likelihood points, so the rows that drop it are deliberately misspecified checks and not competing specifications; they are reported because the objection they answer is about bias, not about fit. Dropping it does not raise the supervisory estimate, which is what a market-wide supervisory effect concealed by the control would require: the estimate moves from 0.047 to 0.044, and the randomization interval does not widen, so the check costs no precision either. And the same removal raises the monetary policy estimate from 0.173 to 0.230, and to 0.268 once $\Delta s_t$ goes as well. Dropping the spread alone – which the market return leaves in place, and which is itself a possible outcome of the news – changes neither series much, 0.035 and 0.174. A common component on event days is therefore present and measurable in these data; it is in the monetary policy decisions. The autoregressive term changes nothing anywhere, 0.047 against 0.048.

*Robustness of the event coefficient.* Table B.3 collects the specifications summarised in Section 4. Every row re-estimates the variance equation on the stated sample or dummy, holding the selection rule fixed.

**Table B.2: The event coefficient under alternative mean equations**

| | supervisory | | monetary policy | | |
|---|---|---|---|---|---|
| mean equation | $\hat{\delta}$ | s.e. | $\hat{\delta}$ | s.e. | $\log L$ |
| $\mu + \rho r_{t-1} + \lambda r_t^m + \phi \Delta s_t$ (baseline) | 0.047 | 0.043 | 0.173 | 0.122 | −4405.0 |
| without $\lambda r_t^m$ | 0.044 | 0.040 | 0.230 | 0.101 | −5842.1 |
| without $\phi \Delta s_t$ | 0.035 | 0.043 | 0.174 | 0.122 | −4408.0 |
| without $\lambda r_t^m$ and $\phi \Delta s_t$ | 0.015 | 0.043 | 0.268 | 0.107 | −6119.6 |
| without $\rho r_{t-1}$ | 0.048 | 0.044 | 0.179 | 0.122 | −4407.9 |
| constant only | 0.015 | 0.044 | 0.273 | 0.107 | −6120.1 |

*Notes.* The variance equation is the one throughout, estimated on the euro area index over the same 3,330 observations; only the mean equation changes. Standard errors are Bollerslev–Wooldridge robust. $\log L$ is that of the supervisory fit; the monetary policy fits differ from it by less than three points in every row. The two rows without $r_t^m$ were also subjected to the full randomization test of Appendix C: the supervisory estimate gives, on 10,000 draws, a two-sided $p$ value of 0.30 with a 95% interval of $[-2.1\%, +6.7\%]$ in volatility, and the monetary policy estimate a $p$ value of 0.002 with $[+4.5\%, +20.9\%]$, as reported in Table 1.

## C. The randomization distribution and the estimator

*Matching.* Placebo dates are matched on calendar year *and* weekday: each placebo set contains, for every year and every day of the week, as many dates as the true event set does. Year alone is not enough. Monetary policy decisions are announced on a published calendar and 96.5% of them fall on a Thursday, while supervisory releases lean towards Monday (26.9% against 20.0% of event-free days) and away from Thursday (14.0%). Whether this matters is an empirical question and the answer is that it does not: on days free of known events the variance of market-adjusted returns is 1.040 on Thursdays against 1.012 overall, a log ratio of 0.033 with a randomization $p$ value of 0.71. A calendar regularity of that size cannot generate a $\hat{\delta}$ of 0.17. The matching is retained because it costs nothing, not because it rescues anything.

*The support.* The true event dates are eligible as placebo dates. The alternative is to draw only from days free of every known event; the true dates are then never drawn and sit in the control group of all 1,000 re-estimations. If event days carry higher variance this depresses the placebo estimates and inflates the contrast with the observed one, and the $p$ value is too small: under that alternative the monetary policy control gives 0.008, against 0.023 here. For the supervisory test the same bias works against the null, which survives it. Days carrying *only* the other event class are excluded from the support. Overlapping dates remain eligible, because they belong to the target event set: two of the 242 supervisory days are also monetary policy days and are drawn from, while the 111 monetary policy days that carry no release are not.

*Why not match more finely.* Adding the calendar quarter to the strata is the obvious

**Table B.3: The event coefficient across specifications**

| | $\hat{\delta}$ | s.e. | $p_{\text{as}}$ | $\Delta$ vol. |
|---|---|---|---|---|
| *Panel A. Supervisory announcements* | | | | |
| Baseline, euro area index | 0.047 | 0.043 | 0.27 | +2.4% |
| Italian bank index | 0.016 | 0.054 | 0.76 | +0.8% |
| Cut-off at 17:30 | 0.047 | 0.043 | 0.27 | +2.4% |
| Event day $t+1$ | 0.034 | 0.044 | 0.44 | +1.7% |
| Day before the release | 0.036 | 0.046 | 0.43 | +1.8% |
| Normal errors | 0.040 | 0.050 | 0.42 | +2.0% |
| Post-2018 subsample | 0.017 | 0.054 | 0.75 | +0.8% |
| From November 2014 | 0.039 | 0.047 | 0.40 | +2.0% |
| No market return in the mean | 0.044 | 0.040 | 0.27 | +2.2% |
| System-wide releases only | 0.050 | 0.049 | 0.30 | +2.6% |
| Institution-specific only | 0.058 | 0.097 | 0.55 | +2.9% |
| Administrative releases (placebo) | −0.126 | 0.204 | 0.54 | −6.1% |
| *Panel B. Monetary policy decisions* | | | | |
| Baseline, euro area index | 0.173 | 0.122 | 0.16 | +9.0% |
| No market return in the mean | 0.230 | 0.101 | 0.02 | +12.2% |
| Both dummies: monetary policy | 0.193 | 0.122 | 0.11 | +10.1% |
| Both dummies: supervision | 0.057 | 0.044 | 0.19 | +2.9% |

*Notes.* Standard errors are Bollerslev–Wooldridge robust, see Bollerslev and Wooldridge (1992), and $p_{\text{as}}$ is the associated asymptotic $p$ value, reported here rather than in Section 4 because it disagrees systematically with the randomization $p$ value: the baseline and the 17:30 cut-off have the same coefficient and the same standard error and give 0.71 against 0.27, which is why the paper relies on $p_{\text{RI}}$. The rows carrying a randomization $p$ value and an interval are in Table 1. About a fifth of the sample concerns individual institutions, which cannot move an aggregate index except through contagion, so pooling might dilute any effect mechanically; it does not detectably, 201 system-wide event days giving 0.050 against 0.058 on 41 institution-specific days, though 41 days is imprecise enough that only a large dilution would show.

refinement, and it was run on both event classes. It does not move the supervisory estimate: $\hat{\delta}$ is unchanged at 0.047 and its $p$ value goes from 0.71 to 0.80. But it substantially reduces the power of the positive control. On the monetary policy decisions the same refinement takes the $p$ value from 0.023 to 0.13 and opens the interval to $[-1.9\%, +19.1\%]$, so an effect of 9% that the coarser design detects is no longer distinguishable from zero. Both are run on 10,000 draws, as are the central specifications.

The mechanism is capacity. Monetary policy decisions are almost all Thursdays; adding the quarter spreads 113 of them over 54 strata of nine to twelve eligible days each, from which two must be drawn. Each draw therefore takes a large share of a small pool and lands on the true dates often – an expected overlap of 18% – and those draws return a large $\hat{\delta}$. The null distribution is contaminated towards the alternative:

its standard deviation rises from 0.076 to 0.098 and its 95th percentile from 0.117 to 0.177, which is almost the observed 0.173. The supervisory design suffers the same effect more mildly, 242 events spread over 154 strata and an overlap rising from 10% to 16%.

The conclusion is not that the supervisory null survives a more demanding test. It is that this particular refinement buys stringency by reducing the number of admissible placebo assignments, and with it the power that the positive control measures. The matching was fixed before these results, on the calendar regularities set out above; what the quarter-matched rows add is a measure of what a finer design would cost.

*Implementation.* Every re-estimation starts from the coefficients of the baseline fit, without which a substantial share converges to degenerate optima: without starting values, 39% of draws returned $|\hat{\delta}| > 0.5$ against a standard error of 0.095, and 9% piled up on a single spurious optimum near $\hat{\delta} = -32$.

A draw is retained if the optimizer converges and its log-likelihood is not below that of the nested fit with $\delta = 0$, which adding a parameter cannot lower. No draw fails either condition in any test reported. Nothing else is screened, and in particular the log-likelihood is not used as a measure of the quality of a draw. It cannot be: across ten thousand placebo draws it correlates with $|\hat{\delta}|$ at 0.94, a placebo that falls on turbulent days raising the likelihood and producing a large coefficient by the same mechanism, so a likelihood screen is a screen on the statistic and would truncate the null distribution exactly in the tails where the $p$ value is read. Mean $\hat{\delta}$ rises monotonically across deciles of the log-likelihood, from 0.000 to 0.098. A tolerance set relative to the largest log-likelihood in the batch is worse again, since the maximum of $B$ draws grows with $B$: on the baseline the same rule discards nothing at a thousand draws and three quarters of them at ten thousand, and what it keeps is the upper tail of $|\hat{\delta}|$ – enough to move the $p$ value from 0.71 to 0.06. None of the estimates reported here uses such a rule. Every row of Table 1 that reports an interval uses 10,000 draws rather than 1,000. Two considerations motivate the larger number. For the positive control the $p$ value rested on twelve draws in the tail at a thousand; at ten thousand it rests on 116 and moves from 0.026 to 0.023. The same control without the market return rested on a single draw, which is a resolution floor and not a measurement: at ten thousand it rests on eight and the $p$ value is 0.002. For the supervisory rows the $p$ value was never in doubt – 0.71 is far from any threshold, and its Monte Carlo error at a thousand draws is about 0.03 – but the interval endpoints are quantiles estimated from some twenty-five draws per tail, with a Monte Carlo standard error of 0.17 percentage points on the upper endpoint. At ten thousand that falls to 0.05, which is what the reported decimal requires. Moving from one thousand draws to ten thousand changes the baseline interval from $[-3.02\%, +4.33\%]$ to $[-3.08\%, +4.38\%]$.

*The p value.* Reported $p$ values are two-sided, obtained by doubling the smaller tail of the null distribution rather than by referring $|\hat{\delta}|$ to it: the null distribution is not centred on zero, so the absolute-value convention would compare the observed statistic with the wrong reference. Two-sidedness is substantive and not a formality, since a

**Table C.1: Monte Carlo distribution of $\hat{\delta}$ under $\delta = 0$**

| event days | mean | s.d. | $t$ |
|---|---|---|---|
| 52 | −0.008 | 0.115 | −1.65 |
| 117 | 0.003 | 0.077 | 0.92 |
| 246 | −0.001 | 0.050 | −0.45 |

release that resolves uncertainty rather than creating it lowers the conditional variance; the administrative placebo of Table B.3 is the one estimate with that sign.

*The interval.* The interval (3) uses the quantiles of the placebo distribution as they are, not centred on its mean. The alternative, common in practice, writes $\bar{d}$ for the placebo mean and reports $\left[\hat{\delta} - q_{0.975}(d_b - \bar{d}),\ \hat{\delta} - q_{0.025}(d_b - \bar{d})\right]$. It has the same width but is centred on $\hat{\delta}$, and when the null distribution is not centred on zero it no longer corresponds to the test: it can exclude zero while the test does not reject, or the reverse. On the supervisory baseline, whose null mean is 0.032, it gives $[-1.5\%, +6.1\%]$ against $[-3.1\%, +4.4\%]$ for (3); on the monetary policy baseline, whose null mean is 0.000, the two coincide. Both rest on the same approximation, a constant effect that shifts the null distribution without changing its shape; an exact inversion would recompute the placebo distribution on a grid of values of $\delta_0$, at thousands of re-estimations per grid point, and is not attempted.

*No detectable bias in the estimator.* The randomization distribution is not centred on zero: its mean is 0.032 in the supervisory design and 0.000 in the monetary policy design. Because a finite-sample bias could depend on event frequency and on the structure of the assignments, this was checked directly. Five hundred series were simulated from the fitted model with $\delta = 0$, holding the mean-equation regressors at their observed values, and a dummy on randomly chosen dates was estimated on each.

The estimator is unbiased at every event frequency. The simulated dispersion matches the randomization distribution at the monetary policy frequency (0.077 against 0.072 at 117 and 113 event days) but exceeds it at the supervisory one (0.050 against 0.039 at 246 and 242 event days), the simulated series being drawn from a model that is itself misspecified for the data; the comparison bears on scale, not on the location that the table establishes. Where the null distribution sits on one realization of the data is therefore a property of that sample and carries no interpretation; the test does not require it to, comparing the observed statistic with the distribution of the same statistic on the same data under randomly assigned dates.

This is also why the asymptotic $p$ values of Table B.3 diverge from the randomization $p$ values of Table 1. Under a variance equation that no candidate specification fits, asymptotic inference is unreliable in both location and scale, whereas the null distribution used here is generated under the same specification as the estimate compared with it.

*The residual-based test.* The test of Section 4.4 estimates the model once with $\delta = 0$ –

**Table D.1: Returns by event class, percentage points**

| | $n$ | mean | median | % negative | $p$ |
|---|---|---|---|---|---|
| *Panel A. Market-adjusted returns, bank index* | | | | | |
| Supervisory announcements | 242 | −0.003 | 0.018 | 50.0 | 0.91 |
| system-wide | 201 | −0.015 | −0.002 | 50.7 | 0.99 |
| institution-specific | 41 | 0.055 | 0.179 | 46.3 | 0.76 |
| Administrative releases | 25 | 0.205 | 0.191 | 44.0 | 0.20 |
| Monetary policy decisions | 113 | 0.059 | 0.172 | 44.2 | 0.60 |
| Days without known events | 2,962 | −0.005 | −0.031 | 51.6 | – |
| *Panel B. Raw returns, bank index* | | | | | |
| Supervisory announcements | 242 | −0.030 | 0.061 | 48.3 | 0.77 |
| Administrative releases | 25 | 0.128 | 0.128 | 44.0 | 0.62 |
| Monetary policy decisions | 113 | 0.157 | 0.446 | 38.1 | 0.31 |
| Days without known events | 2,962 | 0.019 | 0.048 | 47.9 | – |
| *Panel C. Raw returns, market benchmark* | | | | | |
| Supervisory announcements | 242 | −0.024 | 0.029 | 48.3 | 0.59 |
| Administrative releases | 25 | −0.070 | −0.137 | 56.0 | 0.66 |
| Monetary policy decisions | 113 | 0.089 | 0.291 | 39.8 | 0.34 |
| Days without known events | 2,962 | 0.022 | 0.053 | 47.0 | – |

*Notes.* Two-sided randomization $p$ values on the difference in means against all other days, 20,000 draws stratified by calendar year and weekday, the same matching used for the variance tests. Panel A is the bank index net of 1.103 times the market, 1.103 being the estimate of $\lambda$ in the mean equation; Panel B is the same index unadjusted; Panel C is the EURO STOXX 50 on the same dates. The 2,962 days without known events are the 2,963 event-free trading days of Appendix C less the first day of the sample, for which no return is defined.

no dummy in the variance, log-likelihood −4405.576 on ten parameters – and keeps its standardized residuals fixed ($\mathrm{var}(z) = 0.9996$ on all 3,330 observations). Each statistic is the difference between its mean on labelled days and on the rest, and $p$ values use 20,000 matched permutations. A fourth statistic, the mean of $\log z_t^2$, is computed by the accompanying script only to show why it must be discarded: six observations have $|z_t| < 0.001$ and the smallest $\log z_t^2$ is −21.4, so a handful of days on which the model happened to fit almost exactly dominate the average.

## D. First-moment evidence

Table D.1 reports returns by event class, first market-adjusted – the bank index net of 1.10 times the market – and then raw, both for the bank index and for the market benchmark alone.

Exactly half of the announcement days close negative, 121 of 242. Monetary policy days are not distinguishable in the mean either, which is the empirical counterpart of

**Table E.1: Dispersion of the surprise and estimated effects**

| *Panel A. Standard deviation of the surprise, basis points* | | | |
|---|---|---|---|
| | $n$ | 3-month | 10-year |
| 2013–2015 | 32 | 1.87 | 4.06 |
| 2016–2021 (lower bound) | 48 | 1.81 | 2.09 |
| 2022–October 2023 | 15 | 6.16 | 6.28 |
| *Panel B. Estimates* | | | |
| | $\hat{\beta}$ | s.e. | $p$ |
| 3-month surprise, mean equation | 0.015 | 0.047 | 0.75 |
| 3-month surprise, with long end | 0.012 | 0.040 | 0.76 |
| 10-year surprise, mean equation | 0.089 | 0.043 | 0.04 |
| \|10-year surprise\|, variance equation | 0.013 | 0.033 | 0.70 |

the identification argument of Section 1: the sign of the surprise is not identified ex ante for either event class, so pooled mean tests are uninformative for both.

Panels B and C qualify that statement for monetary policy, without overturning it. Raw bank index returns average +0.157 on decision days against +0.019 on event-free days, with a median of +0.446 and only 38 per cent of days closing negative; the market benchmark alone averages +0.089 with a median of +0.291. This is the pattern documented for FOMC days on United States data, where the announcement-day return is a market-wide phenomenon rather than a sectoral one – Lucca and Moench (2015), Savor and Wilson (2013) – and it is the component that the market adjustment of Panel A removes. It is also, with 113 decisions and a daily standard deviation near 1.7 percentage points, not distinguishable from zero: $p = 0.40$ raw against $p = 0.52$ adjusted. The gap between mean and median says where the difficulty lies – what signal there is sits in the centre of the distribution and is swamped by the tails – and it is a further reason to read the second moment rather than the first. For supervisory announcements nothing changes between panels: −0.003 adjusted, −0.030 raw, neither distinguishable from zero.

## E. Monetary policy surprises

The intraday OIS surprises of Altavilla, Brugnolini, Gürkaynak, Motto and Ragusa (2019) are matched to 95 of the 113 monetary policy decisions; 13 of these carry a surprise of exactly zero. Of the eighteen unmatched decisions, seventeen fall beyond the database's coverage, which ends on 26 October 2023, and the eighteenth is the emergency decision of 18 March 2020, announced at 23:45 and therefore outside the intraday window the database is built on. Surprises are dated to the announcement day and mapped to the trading day on which the news reaches the market.

The short-end surprise has a standard deviation below two basis points for eighty of the ninety-five matched decisions, the policy rate being at its lower bound over most of the sample, so its insignificance is a statement about the instrument rather than

about the design. The long-end surprise, through which forward guidance and asset purchases operate, enters the mean equation with a coefficient of 0.089: the design recovers a signed effect where a measure of surprise exists. The estimate provides no evidence that the conditional variance varies in proportion to the magnitude of the measured rate surprise. That is consistent with the reading adopted in the body of the paper, and weaker than it: an absence of evidence over ninety-five decisions, most of them at the lower bound, does not establish that the variance responds to the arrival of information rather than to its size.

**Working Paper del Dipartimento di Economia e Finanza**

1. L. Colombo, H. Dawid, *Strategic Location Choice under Dynamic Oligopolistic Competition and Spillovers*, novembre 2013.
2. M. Bordignon, M. Gamalerio, G. Turati, *Decentralization, Vertical Fiscal Imbalance, and Political Selection*, novembre 2013.
3. M. Guerini, *Is the Friedman Rule Stabilizing? Some Unpleasant Results in a Heterogeneous Expectations Framework*, novembre 2013.
4. E. Brenna, C. Di Novi, *Is caring for elderly parents detrimental to women's mental health? The influence of the European North-South gradient*, novembre 2013.
5. F. Sobbrio, *Citizen-Editors' Endogenous Information Acquisition and News Accuracy*, novembre 2013.
6. P. Bingley, L. Cappellari, *Correlation of Brothers Earnings and Intergenerational Transmission*, novembre 2013.
7. T. Assenza, W. A. Brock, C. H. Hommes, *Animal Spirits, Heterogeneous Expectations and the Emergence of Booms and Busts*, dicembre 2013.
8. D. Parisi, *Is There Room for 'Fear' as a Human Passion in the Work by Adam Smith?*, gennaio 2014.
9. E. Brenna, F. Spandonaro, *Does federalism induce patients' mobility across regions? Evidence from the Italian experience*, febbraio 2014.
10. A. Monticini, F. Ravazzolo, *Forecasting the intraday market price of money*, febbraio 2014.
11. Tiziana Assenza, Jakob Grazzini, Cars Hommes, Domenico Massaro, *PQ Strategies in Monopolistic Competition: Some Insights from the Lab*, marzo 2014.
12. R. Davidson, A. Monticini, *Heteroskedasticity-and-Autocorrelation-Consistent Bootstrapping*, marzo 2014.
13. C. Lucifora, S. Moriconi, *Policy Myopia and Labour Market Institutions*, giugno 2014.
14. N. Pecora, A. Spelta, *Shareholding Network in the Euro Area Banking Market*, giugno 2014.
15. G. Mazzolini, *The economic consequences of accidents at work*, giugno 2014.
16. M. Ambrosanio, P. Balduzzi, M. Bordignon, *Economic crisis and fiscal federalism in Italy*, settembre 2014.
17. P. Bingley, L. Cappellari, K. Tatsiramos, *Family, Community and Long-Term Earnings Inequality*, ottobre 2014.
18. S. Frazzoni, M. L. Mancusi, Z. Rotondi, M. Sobrero, A. Vezzulli, *Innovation and export in SMEs: the role of relationship banking*, novembre 2014.
19. H. Gnutzmann, *Price Discrimination in Asymmetric Industries: Implications for Competition and Welfare*, novembre 2014.
20. A. Baglioni, A. Boitani, M. Bordignon, *Labor mobility and fiscal policy in a currency union*, novembre 2014.
21. C. Nielsen, *Rational Overconfidence and Social Security*, dicembre 2014.
22. M. Kurz, M. Motolese, G. Piccillo, H. Wu, *Monetary Policy with Diverse Private Expectations*, febbraio 2015.
23. S. Piccolo, P. Tedeschi, G. Ursino, *How Limiting Deceptive Practices Harms Consumers*, maggio 2015.
24. A.K.S. Chand, S. Currarini, G. Ursino, *Cheap Talk with Correlated Signals*, maggio 2015.
25. S. Piccolo, P. Tedeschi, G. Ursino, *Deceptive Advertising with Rational Buyers*, giugno 2015.

26. S. Piccolo, E. Tarantino, G. Ursino, *The Value of Transparency in Multidivisional Firms*, giugno 2015.
27. G. Ursino, *Supply Chain Control: a Theory of Vertical Integration*, giugno 2015.
28. I. Aldasoro, D. Delli Gatti, E. Faia, *Bank Networks: Contagion, Systemic Risk and Prudential Policy*, luglio 2015.
29. S. Moriconi, G. Peri, *Country-Specific Preferences and Employment Rates in Europe*, settembre 2015.
30. R. Crinò, L. Ogliari, *Financial Frictions, Product Quality, and International Trade*, settembre 2015.
31. J. Grazzini, A. Spelta, *An empirical analysis of the global input-output network and its evolution*, ottobre 2015.
32. L. Cappellari, A. Di Paolo, *Bilingual Schooling and Earnings: Evidence from a Language-in-Education Reform*, novembre 2015.
33. A. Litina, S. Moriconi, S. Zanaj, *The Cultural Transmission of Environmental Preferences: Evidence from International Migration*, novembre 2015.
34. S. Moriconi, P. M. Picard, S. Zanaj, *Commodity Taxation and Regulatory Competition*, novembre 2015.
35. M. Bordignon, V. Grembi, S. Piazza, *Who do you blame in local finance? An analysis of municipal financing in Italy*, dicembre 2015.
36. A. Spelta, *A unified view of systemic risk: detecting SIFIs and forecasting the financial cycle via EWSs*, gennaio 2016.
37. N. Pecora, A. Spelta, *Discovering SIFIs in interbank communities*, febbraio 2016.
38. M. Botta, L. Colombo, *Macroeconomic and Institutional Determinants of Capital Structure Decisions,* aprile 2016.
39. A. Gamba, G. Immordino, S. Piccolo, *Organized Crime and the Bright Side of Subversion of Law*, maggio 2016.
40. L. Corno, N. Hildebrandt, A. Voena, *Weather Shocks, Age of Marriage and the Direction of Marriage Payments*, maggio 2016.
41. A. Spelta, *Stock prices prediction via tensor decomposition and links forecast*, maggio 2016.
42. T. Assenza, D. Delli Gatti, J. Grazzini, G. Ricchiuti, *Heterogeneous Firms and International Trade: The role of productivity and financial fragility*, giugno 2016.
43. S. Moriconi, *Taxation, industry integration and production efficiency*, giugno 2016.
44. L. Fiorito, C. Orsi, *Survival Value and a Robust, Practical, Joyless Individualism: Thomas Nixon Carver, Social Justice, and Eugenics*, luglio 2016.
45. E. Cottini, P. Ghinetti, *Employment insecurity and employees' health in Denmark*, settembre 2016.
46. G. Cecere, N. Corrocher, M. L. Mancusi, *Financial constraints and public funding for eco-innovation: Empirical evidence on European SMEs*, settembre 2016.
47. E. Brenna, L. Gitto, *Financing elderly care in Italy and Europe. Is there a common vision?*, settembre 2016.
48. D. G. C. Britto, *Unemployment Insurance and the Duration of Employment: Theory and Evidence from a Regression Kink Design*, settembre 2016.
49. E. Caroli, C.Lucifora, D. Vigani, *Is there a Retirement-Health Care utilization puzzle? Evidence from SHARE data in Europe*, ottobre 2016.
50. G. Femminis, *From simple growth to numerical simulations: A primer in dynamic programming*, ottobre 2016.
51. C. Lucifora, M. Tonello, *Monitoring and sanctioning cheating at school: What works? Evidence from a national evaluation program*, ottobre 2016.

52. A. Baglioni, M. Esposito, *Modigliani-Miller Doesn't Hold in a "Bailinable" World: A New Capital Structure to Reduce the Banks' Funding Cost*, novembre 2016.
53. L. Cappellari, P. Castelnovo, D. Checchi, M. Leonardi, *Skilled or educated? Educational reforms, human capital and earnings*, novembre 2016.
54. D. Britto, S. Fiorin, *Corruption and Legislature Size: Evidence from Brazil*, dicembre 2016.
55. F. Andreoli, E. Peluso, *So close yet so unequal: Reconsidering spatial inequality in U.S. cities*, febbraio 2017.
56. E. Cottini, P. Ghinetti, *Is it the way you live or the job you have? Health effects of lifestyles and working conditions*, marzo 2017.
57. A. Albanese, L. Cappellari, M. Leonardi, *The Effects of Youth Labor Market Reforms: Evidence from Italian Apprenticeships*; maggio 2017.
58. S. Perdichizzi, *Estimating Fiscal multipliers in the Eurozone. A Nonlinear Panel Data Approach*, maggio 2017.
59. S. Perdichizzi, *The impact of ECBs conventional and unconventional monetary policies on European banking indexes returns*, maggio 2017.
60. E. Brenna, *Healthcare tax credits: financial help to taxpayers or support to higher income and better educated patients? Evidence from Italy*, giugno 2017.
61. G. Gokmen, T. Nannicini, M. G. Onorato, C. Papageorgiou, *Policies in Hard Times: Assessing the Impact of Financial Crises on Structural Reforms*, settembre 2017.
62. M. Tettamanzi, *E Many Pluribus Unum: A Behavioural Macro-Economic Agent Based Model*, novembre 2017.
63. A. Boitani, C. Punzo, *Banks' leverage behaviour in a two-agent New Keynesian model*, gennaio 2018.
64. M. Bertoni, G. Brunello, L. Cappellari, *Parents, Siblings and Schoolmates. The Effects of Family-School Interactions on Educational Achievement and Long-term Labor Market Outcomes*, gennaio 2018.
65. G. P. Barbetta, G. Sorrenti, G. Turati, *Multigrading and Child Achievement*, gennaio 2018.
66. S. Gagliarducci, M. G. Onorato, F. Sobbrio, G. Tabellini, *War of the Waves: Radio and Resistance During World War II*, febbraio 2018.
67. P. Bingley, L. Cappellari, *Workers, Firms and Life-Cycle Wage Dynamics*, marzo 2018.
68. A. Boitani, S. Perdichizzi, *Public Expenditure Multipliers in recessions. Evidence from the Eurozone*, marzo 2018.
69. M. Le Moglie, G. Turati, *Electoral Cycle Bias in the Media Coverage of Corruption News*, aprile 2018.
70. R. Davidson, A. Monticini, *Improvements in Bootstrap Inference*, aprile 2018.
71. R. Crinò, G. Immordino, S. Piccolo, *Fighting Mobile Crime*, giugno 2018.
72. R. Caminal, L. Cappellari, A. Di Paolo, *Linguistic skills and the intergenerational transmission of language*, agosto 2018.
73. E. Brenna, L. Gitto, *Adult education, the use of Information and Communication Technologies and the impact on quality of life: a case study*, settembre 2018.
74. M. Bordignon, Y. Deng, J. Huang, J. Yang, *Plunging into the Sea: Ideological Change, Institutional Environments and Private Entrepreneurship in China*, settembre 2018.
75. M. Bordignon, D. Xiang, L. Zhan, *Predicting the Effects of a Sugar Sweetened Beverage Tax in a Household Production Model*, settembre 2018.
76. C. Punzo, L. Rossi, *The Redistributive Effects of a Money-Financed Fiscal Stimulus*, gennaio 2019.
77. A. Baglioni, L. Colombo, P. Rossi, *Debt restructuring with multiple bank relationships*, gennaio 2019.

78. E. Cottini, P. Ghinetti, S. Moriconi, *Higher Education Supply, Neighbourhood effects and Economic Welfare,* febbraio 2019.
79. S. Della Lena, F. Panebianco, *Cultural Transmission with Incomplete Information: Parental Perceived Efficacy and Group Misrepresentation,* marzo 2019.
80. T. Colussi, Ingo E. Isphording, Nico Pestel, *Minority Salience and Political Extremism*, marzo 2019.
81. G. P. Barbetta, P. Canino, S. Cima, *Let's tweet again? The impact of social networks on literature achievement in high school students: Evidence from a randomized controlled trial*, maggio 2019.
82. Y. Brilli, C. Lucifora, A. Russo, M. Tonello, *Vaccination take-up and health: evidence from a flu vaccination program for the elderly*, giugno 2019.
83. C. Di Novi, M. Piacenza, S. Robone, G. Turati, *Does fiscal decentralization affect regional disparities in health? Quasi-experimental evidence from Italy,* luglio 2019.
84. L. Abrardi, L. Colombo, P. Tedeschi, *The Gains of Ignoring Risk: Insurance with Better Informed Principals,* luglio 2019.
85. A. Garnero, C. Lucifora, *Turning a Blind Eye? Compliance to Minimum Wages and Employment*, gennaio 2020.
86. M. Bordignon, M. Gamalerio, E. Slerca, G. Turati, *Stop invasion! The electoral tipping point in anti-immigrant voting*, marzo 2020.
87. D. Vigani, C. Lucifora, *Losing control? Unions' Representativeness, "Pirate" Collective Agreements and Wages,* marzo 2020.
88. S. L. Comi, E. Cottini, C. Lucifora, *The effect of retirement on social relationships: new evidence from SHARE*, maggio 2020.
89. A. Boitani, S. Perdichizzi, C. Punzo, *Nonlinearities and expenditure multipliers in the Eurozone*, giugno 2020.
90. R. A. Ramos, F. Bassi, D. Lang, *Bet against the trend and cash in profits*, ottobre 2020.
91. F. Bassi, *Chronic Excess Capacity and Unemployment Hysteresis in EU Countries. A Structural Approach*, ottobre 2020.
92. M. Bordignon, T. Colussi, *Dancing with the Populist. New Parties, Electoral Rules and Italian Municipal Elections*, ottobre 2020.
93. E. Cottini, C. Lucifora, G. Turati, D. Vigani, *Children Use of Emergency Care: Differences Between Natives and Migrants in Italy,* ottobre 2020.
94. B. Fanfani, *Tastes for Discrimination in Monopsonistic Labour Markets*, ottobre 2020.
95. B. Fanfani, *The Employment Effects of Collective Bargaining*, ottobre 2020.
96. O. Giuntella, J. Lonsky, F. Mazzonna, L. Stella, *Immigration Policy and Immigrants' Sleep. Evidence from DACA,* dicembre 2020.
97. E. Cottini, P. Ghinetti, E. Iossa, P. Sacco, *Stress and Incentives at Work*, gennaio 2021.
98. L. Pieroni, M. R. Roig, L. Salmasi, *Italy: immigration and the evolution of populism*, gennaio 2021.
99. L. Corno, E. La Ferrara, A. Voena, *Female Genital Cutting and the Slave Trade*, febbraio 2021.
100. O. Giuntella, L. Rotunno, L. Stella, *Trade Shocks, Fertility, and Marital Behavior* , marzo 2021.
101. P. Bingley, L. Cappellari, K. Tatsiramos, *Parental Assortative Mating and the Intergenerational Transmission of Human Capital*, aprile 2021.
102. F. Devicienti, B. Fanfani, *Firms' Margins of Adjustment to Wage Growth. The Case of Italian Collective Bargaining*; aprile 2021.
103. C. Lucifora, A. Russo, D. Vigani, *Does prescribing appropriateness reduce health expenditures? Main e_ects and unintended outcomes*, maggio 2021.

104. T. Colussi, *The Political Effects of Threats to the Nation: Evidence from the Cuban Missile Crisis*, giugno 2021.
105. M. Bordignon, N. Gatti, M. G. Onorato, *Getting closer or falling apart? Euro countries after the Euro crisis*, giugno 2021.
106. E. Battistin, M. Ovidi, *Rising Stars*, giugno 2021.
107. D. Checchi, A. Fenizia, C. Lucifora, PUBLIC SECTOR JOBS: *Working in the public sector in Europe and the US*, giugno 2021.
108. K. Aktas, G. Argentin, G. P. Barbetta, G. Barbieri, L. V. A. Colombo, *High School Choices by Immigrant Students in Italy: Evidence from Administrative Data*, luglio 2021.
109. B. Fanfani, C. Lucifora, D. Vigani, *Employer Association in Italy. Trends and Economic Outcomes*, luglio 2021.
110. F. Bassi, A. Boitani, *Monetary and macroprudential policy: The multiplier effects of cooperation,* settembre 2021.
111. S. Basiglio, A. Foresta, G. Turati, *Impatience and crime. Evidence from the NLSY97*, settembre 2021.
112. A. Baglioni, A. Monticini, D. Peel, *The Impact of the ECB Banking Supervision Announcements on the EU Stock Market*, novembre 2021.
113. E. Facchetti, L. Neri, M. Ovidi, *Should you Meet The Parents? The impact of information on non-test score attributes on school choice*, dicembre 2021.
114. M. Bratti, E. Cottini, P. Ghinetti, *Education, health and health-related behaviors: Evidence from higher education expansion*, febbraio 2022.
115. A. Boitani, C. Dragomirescu-Gaina, *News and narratives: A cointegration analysis of Russian economic policy uncertainty*, aprile 2022.
116. D. Delli Gatti, J. Grazzini, D. Massaro, F. Panebianco, *The Impact of Growth on the Transmission of Patience*, luglio 2022.
117. I. Torrini, C. Lucifora, A. Russo, *The Long-Term Effects of Hospitalization on Health Care Expenditures: An Empirical Analysis for the Young-Old Population*, luglio 2022.
118. T. Colussi, M. Romagnoli, E. Villar, *The Intended and Unintended Consequences of Taxing Waste*, settembre 2022.
119. D. Delli Gatti, G. Iannotta, *Behavioural credit cycles*, settembre 2022.
120. C. Punzo, G. Rivolta, *Money versus debt financed regime: Evidence from an estimated DSGE model*, novembre 2022.
121. M. Ovidi, *Parents Know Better: Sorting on Match Effects in Primary School*, novembre 2022.
122. L. Cappellari, D. Checchi, M. Ovidi, *The effects of schooling on cognitive skills: evidence from education expansions*, dicembre 2022.
123. M. Bertoni, G. Brunello, L. Cappellari, M. De Paola, *The long-run earnings effects of winning a mayoral election*, gennaio 2023.
124. M. Bordignon, F. Franzoni, M. Gamalerio, *Is Populism reversible? Evidence from Italian local elections during the pandemic*, gennaio 2023.
125. L. Colombo, G. Femminis, A. Pavan, *Subsidies to Technology Adoption when Firms' Information is Endogenous*, gennaio 2023.
126. L. Pieroni, M. Rosselló Roig, L. Salmasi, G. Turati, *Legal status and voluntary abortions by immigrants*, gennaio 2023.
127. F. Di Pace, G. Mangiante, R. Masolo, *Do firm expectations respond to Monetary Policy announcements?*, febbraio 2023.
128. R. Masolo, *Heterogeneity and the Equitable Rate of Interest*, febbraio 2023.
129. D. Delli Gatti, R.Terranova, E. M. Turco, *Mind the knowledge gap! The origins of declining business dynamism in a macro agent-based model*, ottobre 2023.

130. R. Davidson, A. Monticini, *Bootstrap Performance with Heteroskedasticity*, novembre 2023.
131. V. Colombo, A. Paccagnini, *Uncertainty and the Federal Reserve's Balance Sheet Monetary Policy*, febbraio 2024.
132. M. Bordignon, T. Colussi, F. Porcelli, *Populists at work. Italian municipal finance under M5s governments*, febbraio 2024.
133. A. Boitani, L. Di Domenico, S. Ricchiuti, *Monetary policy and inequality: an heterogenous agents' approach*, settembre 2024.
134. M. Di Giacomo, M. Piacenza, L. Salmasi, G. Turati, *Understanding productivity in maternity wards*, ottobre 2024.
135. D. Bonciani, R. Masolo, S. Sarpietro, *How Food Prices Shape Inflation Expectations and the Monetary Policy Response*, ottobre 2024.
136. G. Battiston, L. Corno, E. La Ferrara, *Informing Risky Migration: Evidence from a field experiment in Guinea*, dicembre 2024.
137. F. Di Pace, G. Mangiante, R. Masolo, *Monetary Policy Rules: the Market's View*, dicembre 2024.
138. L. Corno, E. La Ferrara, A. Voena, *Female Genital Cutting and the Slave Trade*, maggio 2025.
139. M. Carlana, L. Corno, *Peer influence in Educational Choices: Social Image Concerns and Same-Gender Interactions*, maggio 2025.
140. R. Disney, J. Gathergood, S. Machin, M. Sandi, *Human Capital from Childhood Exposure to Homeownership:Evidence from Right-to-Buy*, giugno 2025.
141. J. Boshoff, S. Machin, M. Sandi, *Youth Crime and Delinquency In And Out Of School*, giugno 2025.
142. A. Boitani, C. Dragomirescu-Gaina, A. Monticini, *Managing the chaos: policy challenges in a hyperinflationary environment*, giugno 2025.
143. E. Cottini, L. Popoescu, L. Salmasi, G. Turati, *Poisoned Air, Shortened Lives: PM2.5 Exposure and Premature Mortality in Southern European Cities*, luglio 2025.
144. M. Bordignon, N. Gatti, G. Turati, *Are citizens willing to reduce public debt? Beliefs, information and policy preferences*, luglio 2025.
145. J. Costa-Fon, N. Gatti, G. Turati, D. Wiesen, *A prosocial legacy of COVID-19 among healthcare professionals?*, ottobre 2025.
146. A. Ariu, G. Rivolta, *The Micro and Macro of a Large and Sudden Devaluation*, ottobre 2025.
147. F. Di Pace, G. Mangiante, R. Masolo, *Brexit and the cost of living: a tale of two phases*, dicembre 2025.
148. A. Albanese, L. Cappellari, M. Ovidi, *Earning ability over the life cycle*, dicembre 2025.
149. L. Cappellari, M. Morelli, *Backward Spillovers and Equalized Access: The Effects of the Bologna Process across Italy's Education System*, febbraio 2026.
150. M. Bordignon, T. Colussi, *Populism in the Italian Municipal Elections: the M5S experience*, febbraio 2026.
151. L. Cappellari, A. di Paolo, T. Ogajah Tawiah, *Linguistic Distance and Job Quality in a Bilingual Labour Market*, febbraio 2026.
152. P. Bingley, L. Cappellari, K. Tatsiramos, *Assortative Mating and Education Gradients within and across Generations*, marzo 2026.
153. G. Letta, M. C. Nurchis, L. Salmasi, G. Turati, *Do vaccination-based access restriction policies curb hesitancy? A comparison between Italy and Spain*, marzo 2026.
154. L. Cappellari, T. Colussi, M. Ovidi, *Gendered Honors: Graduation Committees, Academic Distinction, and Early-Career Outcomes*, luglio 2026.

155. L. Cappellari, B. Fanfani, *Collective Bargaining and the Wage Structure* ,luglio 2026.
156. A. Boitani, M. Franzini, *Inequality, growth and welfare*, settembre 2026.
157. G. Letta, L. Salmasi, G. Turati, *Household Income Uncertainty and Pregnancy Decisions: Evidence from Temporary Work Suspensions*, settembre 2026.
158. A. Monticini, *Does routine supervisory communication move volatility? Evidence from 285 ECB announcements*, ottobre 2026.