NGEU Grants and Sovereign Auction Demand: Evidence from Euro-Area Primary Markets*

Diogo Sampaio Lima
Maastricht University

Martim Felgueiras
Maastricht University

Last updated: May 27, 2026

Abstract

Does NextGenerationEU (NGEU) raise private sovereign demand at issuance? We study this question using a 14-country panel of sovereign bond auctions from 2018 to 2023, pooling eight euro-area sovereigns exposed to NGEU grants with six non-euro-area placebo issuers. We measure auction demand using the log bid-to-cover ratio and use each country's grants-only National Recovery and Resilience Plan (NRRP) allocation, scaled by 2019 GDP, as the source of treatment intensity after the July 2020 NGEU agreement. Because the Eurosystem is prohibited by Article 123 TFEU from buying sovereign debt at auction, the primary-market setting rules out direct ECB purchase-flow contamination, although indirect monetary-policy channels through secondary-market conditions remain possible. Auction demand rises more strongly for euro-area sovereigns with greater grants exposure: the implied bid-to-cover responses are 7.6% for Italy, 12.6% for Spain, and 13.3% for Portugal. The result is robust to inference designed for few country clusters, including exact wild cluster bootstrap and a country-label permutation test. A grants-versus-loans horse race sharpens the interpretation: the grants effect strengthens when loans are co-estimated, the loans component enters with the opposite sign, and the two coefficients are statistically distinguishable at the 1% level. The effect concentrates in the second half of 2022, consistent with a credibility interpretation in which repeated successful disbursement events built market confidence in the joint-borrowing arrangement.

JEL: E44, E62, F33, G14, H63, H87.
Keywords: NextGenerationEU; bid-to-cover ratio; sovereign auctions; primary-market demand; grants versus loans; few-cluster inference.

*Corresponding author's email address: d.sampaiolima@maastrichtuniversity.nl.

1 Introduction

Next Generation EU (NGEU) is the largest joint fiscal instrument in the history of the European Union. The programme authorizes roughly €750 billion in common EU borrowing to finance grants and loans to Member States through the Recovery and Resilience Facility (RRF), with the grant component distributed according to a formula that concentrates transfers on the countries most affected by the pandemic (European Parliament and Council, 2021; Christie et al., 2021). Whether this new fiscal-transfer architecture changed private sovereign demand is an open empirical question with direct implications for the future design of common European debt instruments.

The empirical literature on NGEU's effect on sovereign markets has operated primarily through secondary-market sovereign prices and closely related financial-market outcomes. Havlik et al. (2022) study fiscal and monetary policy announcements during the pandemic using euro-area sovereign spreads. Corradin and Schwaab (2023) decompose euro-area sovereign yields into risk-premium components and study the effects of monetary and fiscal policy announcements during the COVID-19 period. Pancotto et al. (2023) examine bank-equity returns, sovereign CDS, and bank CDS around the EU recovery-fund agreement. Alessi et al. (2025) use the greenness of national Recovery and Resilience Plans as a treatment in a staggered difference-in-differences design on sovereign yields. The common feature of this literature is that NGEU is studied through secondary-market prices or closely related financial-market outcomes rather than through primary-market auction demand.

The central difficulty in that setting is that NGEU overlapped with the ECB's Pandemic Emergency Purchase Programme (PEPP) and Public Sector Purchase Programme (PSPP), both of which operated directly in the secondary market and were first-order forces in euro-area sovereign-bond pricing during the same period. Distinguishing an NGEU effect from contemporaneous ECB purchase flows, monetary-policy announcements, and secondary-market risk-premium dynamics is therefore the core identification challenge of the secondary-market approach, and the resulting estimates are sensitive to event-window choices, structural decompositions, and identifying assumptions.

This paper moves the outcome from the secondary market to the primary market. We study sovereign bond auctions and use the bid-to-cover ratio (BCR) as a measure of private demand at issuance. Article 123 of the Treaty on the Functioning of the European Union prohibits the Eurosystem from purchasing government debt directly from Member States, so the ECB cannot mechanically participate in sovereign primary auctions. The primary-market design therefore removes direct ECB purchase-flow contamination at issuance, although indirect monetary-policy channels operating through secondary-market yields, volatility, liquidity conditions, and dealer balance sheets remain. The paper should accordingly be read as a cleaner test of private sovereign demand at issuance, not as a complete separation of fiscal and monetary transmission.

The treatment variable is grants-only NRRP intensity — each country's approved grant allocation divided by 2019 nominal GDP — interacted with a post-21-July-2020 indicator. The distinction between grants and loans is economically central. Grants are unconditional transfers from the common EU budget that increase the recipient sovereign's net wealth without creating a repayment obligation. Loans provide financing on favorable terms but also create liabilities. Their implications for sovereign demand are therefore theoretically distinct, and a rational investor pricing sovereign debt should respond more strongly to the grants margin. Italy illustrates the point clearly. Its €68.9 billion grant allocation places it third on grants intensity at 0.038, while its €122.6 billion loan allocation pushes its grants-plus-loans intensity to 0.107, almost tied with Spain.

The empirical analysis uses a panel of 1,471 sovereign bond auctions between January 2018 and July 2023 across fourteen advanced economies. Eight are euro-area sovereigns exposed to NGEU grants (Austria, Belgium, Finland, France, Germany, Italy, Portugal, and Spain), while six are non-euro-area placebos (Australia, Canada, Japan, New Zealand, South Korea, and the United Kingdom). The non-euro-area sovereigns contribute variation to the time fixed effects by helping absorb global shocks that jointly affect sovereign debt markets, while contributing zero to the treatment gradient itself.

The headline result is that sovereigns with greater grants exposure experienced meaningfully stronger auction demand after the July 2020 agreement. The implied responses are around 8 percent for Italy and around 13 percent for Spain and Portugal. The effect survives exact-inference procedures designed for panels with few country clusters and is economically relevant in size.

A grants-versus-loans horse race sharpens the interpretation. When grants and loans are allowed to enter separately, the grants margin becomes stronger, while the loans margin does not show a comparable positive effect. The evidence is therefore not simply that larger overall NGEU envelopes are associated with stronger auction demand. Rather, the data are more consistent with a response to the unconditional-transfer component of NGEU than to the loan component.

The evidence suggests that the composition of supranational fiscal support matters for sovereign-market reception. Common borrowing arrangements structured primarily through unconditional transfers may affect sovereign demand differently from arrangements relying mainly on on-lending facilities. More broadly, the results suggest that fiscal-capacity design and credibility may matter for sovereign funding conditions even when the direct scale of realized transfers remains limited.

The timing of the response also matters. The strongest estimates appear in the second half of 2022 rather than immediately after the July 2020 agreement. Moreover, once approved grants intensity is controlled for, realized cumulative disbursements do not explain the response. We interpret this pattern as consistent with a credibility mechanism in which repeated successful disbursement events gradually strengthened market confidence in the joint-borrowing framework. This interpretation is necessarily reduced-form: the paper cannot separately identify credibility, redenomination-risk, liquidity, or issuance-expectation channels.

The paper also adopts an explicitly conservative approach to inference. The treatment variation is country-level, the effective number of treated clusters is small, and the identifying variation is concentrated in a handful of sovereigns. We therefore complement asymptotic cluster-robust inference with exact-enumeration wild cluster bootstrap, country-label permutation, within-euro-area restricted permutation, and time-placebo calibration exercises. The most demanding within-euro-area permutation weakens the evidence relative to the full sample, indicating that part of the identifying power comes from the contrast between NGEU-exposed euro-area sovereigns and non-euro placebo issuers. The full-sample evidence should therefore be interpreted as combining within-euro-area treatment variation with this broader placebo contrast.

This paper makes four contributions. First, it provides the first primary-market evidence on NGEU's effect on sovereign demand. Second, it introduces the grants-versus-loans distinction into the empirical NGEU literature and shows that the two components operate differently on sovereign auction demand. Third, it applies a comprehensive honest-inference framework to a setting with few clusters where asymptotic approximations are unreliable. Fourth, it combines announcement-intensity and realized-disbursement evidence to motivate a credibility-based interpretation of the timing of the response.

The paper does not claim a welfare effect, does not estimate counterfactual funding costs, and does not decompose the response into specific pricing channels. Nor does it contest the secondary-market evidence in Corradin and Schwaab (2023), Havlik et al. (2022), or Alessi et al. (2025); the primary-market analysis asks a different question on a different outcome.

The remainder of the paper proceeds as follows. Section 2 positions the contribution relative to the existing literature. Section 3 describes the auction panel and treatment construction. Section 4 presents the empirical strategy and identification challenges. Section 5 reports the main results and the grants-versus-loans horse race. Section 6 discusses robustness and pre-trend evidence. Section 7 reports the validation exercises. Section 8 discusses interpretation and mechanism. Section 9 concludes.

2 Related literature

This paper connects three literatures. The first studies NGEU through secondary-market sovereign prices. The second studies sovereign auction demand and establishes the bid-to-cover ratio as an informative primary-market outcome. The third develops inference procedures for settings in which treatment varies at the cluster level and the number of clusters is small. The contribution of this paper is to bring these three strands together: it studies NGEU in the primary market, uses the auction-demand outcome developed in the sovereign-auction literature, and evaluates the evidence with inference tools appropriate for a 14-country panel.

NGEU and euro-area sovereign yields. The empirical NGEU literature has focused primarily on secondary-market sovereign yields, spreads, CDS prices, and financial-sector outcomes. Havlik et al. (2022) estimate event-study regressions of euro-area sovereign spreads around COVID-era monetary and fiscal policy announcements. They find that monetary announcements, especially PEPP, generated larger spread-compressing effects than fiscal announcements, while NGEU-related news was more relevant than pure loan instruments among fiscal crisis tools. Corradin and Schwaab (2023) decompose euro-area sovereign yields into expected short-rate and term-premium components, default risk, redenomination risk, liquidity risk, and segmentation premia. They show that both ECB monetary-policy announcements and EU fiscal-policy announcements affected sovereign yields during the pandemic, mainly through default, redenomination, and segmentation premia. Pancotto et al. (2023) study bank-equity abnormal returns, sovereign CDS, and bank CDS around the July 2020 recovery-fund agreement. They find positive bank-equity responses and significant declines in sovereign and bank CDS spreads, with stronger effects in more indebted countries and in countries that had strongly advocated the recovery fund. Alessi et al. (2025) use the green-investment share of national Recovery and Resilience Plans as a treatment in a staggered difference-in-differences design and study sovereign bond yields. They find that greener recovery-plan investments are associated with lower sovereign yields, with stronger effects at longer residual maturities.

This strand is the natural starting point for the present paper, but it also reveals the core identification difficulty. Secondary-market prices are the object most closely connected to funding costs, but they are also the object most directly exposed to ECB asset-purchase flows. PEPP and PSPP operate in the secondary market, and their scale makes it hard to distinguish fiscal-news effects from monetary-policy purchase effects without relying on narrow event windows or structural decompositions. The present paper does not challenge those studies. Instead, it asks a complementary question: whether NGEU grant exposure is associated with stronger private demand at issuance, measured in a market where direct Eurosystem purchase-flow contamination is absent by institutional design.

Sovereign auction demand and the bid-to-cover ratio. A second literature studies primary-market demand for sovereign debt. Beetsma et al. (2018) show that more successful euro-area sovereign auctions, measured by higher bid-to-cover ratios, are followed by lower secondary-market yields, especially when market volatility is high. Beetsma et al. (2020) study the determinants of euro-area auction BCR and show that secondary-market yields, past auction outcomes, primary-dealer participation, supply, and yield volatility help explain bidding pressure at issuance. Shida (2023) studies more than 600 German federal bond auctions between 2005 and 2022 and identifies robust auction-specific, institutional, regulatory, and financial-market determinants of demand. The paper also finds that central-bank net purchases in the secondary market are positively associated with primary-market demand, at least for short-term bonds.

This literature provides the outcome and institutional setting used here. The bid-to-cover ratio is not a structural demand curve, but it is a standard summary of bidding pressure at issuance. Moving from secondary-market spreads to auction BCR changes the interpretation of the evidence. It does not eliminate all monetary-policy channels: secondary-market yields, liquidity, volatility, and dealer balance sheets can still affect auction demand. It does, however, remove the direct participation of the Eurosystem in the outcome, since Article 123 TFEU prohibits purchases at issuance. The contribution relative to the auction literature is to apply this primary-market framework to NGEU, extend the sample into the post-2020 period, and add non-euro-area placebo issuers to discipline common shocks in the full panel.

Inference with few clusters. The third relevant literature concerns inference when the number of clusters is small and the treatment varies at the cluster level. This is not a technical detail in the present setting: NGEU intensity varies by country, only a small number of euro-area sovereigns have material grant exposure, and the high-exposure variation is concentrated in Italy, Spain, and Portugal. Conventional asymptotic cluster-robust standard errors are therefore not sufficient on their own.

The inference tools used in this paper draw on Cameron et al. (2008) for the wild cluster bootstrap, MacKinnon and Webb (2017) for exact enumeration in small-G settings, Young (2019) and MacKinnon and Webb (2020) for randomization and permutation inference, and Roth (2022) for the cautionary interpretation of event-study pre-trend tests. Lima (2026) assembles these tools into an applied diagnostic framework. In this paper they serve a narrower role: to report the headline result under exact wild cluster bootstrap, country-label permutation, within-euro-area restricted permutation, and time-placebo calibration. These procedures do not solve the identifying assumption, but they discipline inference in a setting where the standard large-cluster approximation is fragile.

3 Data and setup

3.1 Auction panel

The unit of analysis is an individual sovereign bond auction. The panel contains 1,471 competitive primary auctions conducted by fourteen advanced-economy treasuries between 4 January 2018 and 4 July 2023. The auction-level data are hand-collected from official national debt-management-office (DMO) bulletins; auxiliary macro and financial series used in robustness exercises are described in Appendix A. Eight countries are euro-area sovereigns exposed to NGEU grants (Austria, Belgium, Finland, France, Germany, Italy, Portugal, and Spain). Six are non-euro-area sovereigns with zero NGEU exposure (Australia, Canada, Japan, New Zealand, South Korea, and the United Kingdom). The euro-area countries contribute 695 auctions and the non-euro-area placebo issuers contribute 776.

We restrict the sample to conventional fixed-coupon sovereign bonds at the 10-year and 30-year tenors. This restriction keeps the outcome comparable across issuers and excludes instruments with different benchmarks or investor bases, such as bills, linkers, strips, and syndications. Of the retained auctions, 1,021 are 10-year auctions and 450 are 30-year auctions.

The non-euro-area placebo issuers are included for two reasons. First, they provide untreated country trajectories that enter the country-label permutation exercise in section 7. Second, they help the time fixed effects absorb global shocks to sovereign debt demand: shifts in global risk appetite, monetary-policy expectations, or bond-market volatility can affect Canadian, Korean, British, or Japanese auctions as well as euro-area auctions, while only the latter are exposed to NGEU grants. The comparison is not perfect. The placebo countries differ from the euro-area sovereigns in currency regime, monetary-policy framework, auction institutions, and investor base. We therefore treat the non-euro-area comparison as a source of additional discipline and power rather than as a clean control group.

Table 1 reports per-country counts, average bid-to-cover ratios, and average competitive allotments. Two features are useful background for the empirical analysis. First, the high-exposure euro-area sovereigns have lower average BCRs than many of the placebo issuers: Italy and Spain have mean BCRs near 1.55, while Australia and Japan are above 3. Country fixed effects absorb these level differences in the regressions. Second, the within-country standard deviations are moderate relative to the cross-country level differences, which motivates the panel design and the emphasis on changes over time rather than comparisons of unconditional BCR levels.

Table 1: Auction panel summary by country

		Bid-to-cover		Mean comp. alloc.
Country	N	Mean	SD	(€/equiv. mn)
Euro-area treated				
Austria	76	2.36	0.51	577
Belgium	57	1.97	0.53	1,149
Finland	27	1.93	0.96	713
France	155	2.23	0.59	3,126
Germany	126	1.64	0.54	2,110
Italy	96	1.53	0.35	2,282
Portugal	53	2.13	0.45	617
Spain	105	1.55	0.45	1,676
Non-euro-area placebos				
Australia	119	3.48	1.12	929
Canada	91	2.31	0.18	3,469
Japan	126	3.80	0.74	1,309
New Zealand	156	2.83	0.97	187
South Korea	165	2.85	0.34	1,771
United Kingdom	119	2.37	0.34	2,839
Full panel	1,471	2.46	0.93	1,738

Notes: N is the number of 10-year and 30-year conventional sovereign auctions between 4 January 2018 and 4 July 2023. Competitive allotment is in millions of euro or euro-equivalent using contemporaneous spot rates.

3.2 Outcome

The primary outcome is the natural log of the bid-to-cover ratio, $\log(\text{BCR}_{c,t})$. The bid-to-cover ratio is the total amount of competitive bids divided by the competitive amount allotted in auction t of country c. It is a standard measure of auction demand in the sovereign auction literature (Beetsma et al., 2020). We use the logarithm because BCR is right-skewed and because the resulting coefficients can be interpreted approximately as proportional changes in auction demand. All 1,471 auctions have a valid BCR observation, and no winsorising is applied. In the pre-NGEU period (2018–mid-2020), the euro-area subsample has a mean log-BCR of 0.577 and a standard deviation of 0.314.

Two auxiliary outcomes are used in section 7. The first is the log number of primary dealers placing competitive bids, available for 814 auctions across nine countries. The second is the log competitive amount allotted, available for all 1,471 auctions. These outcomes are not alternative measures of the same demand concept. Dealer count captures an extensive-margin participation measure, while competitive allotment captures the issuer's supply choice. They are useful because a demand-side interpretation of the BCR result should not mechanically imply comparable movements in dealer-panel participation or auction size.

Figure 1 provides descriptive evidence on the raw evolution of log-BCR. It plots country-time means by four exposure groups: high-grants, medium-grants, low-grants, and non-euro-area placebo issuers. The high-grants group — Portugal, Spain, and Italy — rises relative to the other groups after the July 2020 agreement. The figure is intentionally unconditional and should not be read as identification evidence; the formal pre-trend and event-time tests are reported in section 6.

Figure 1: Raw country-time means of log bid-to-cover by exposure group, before and after 21 July 2020

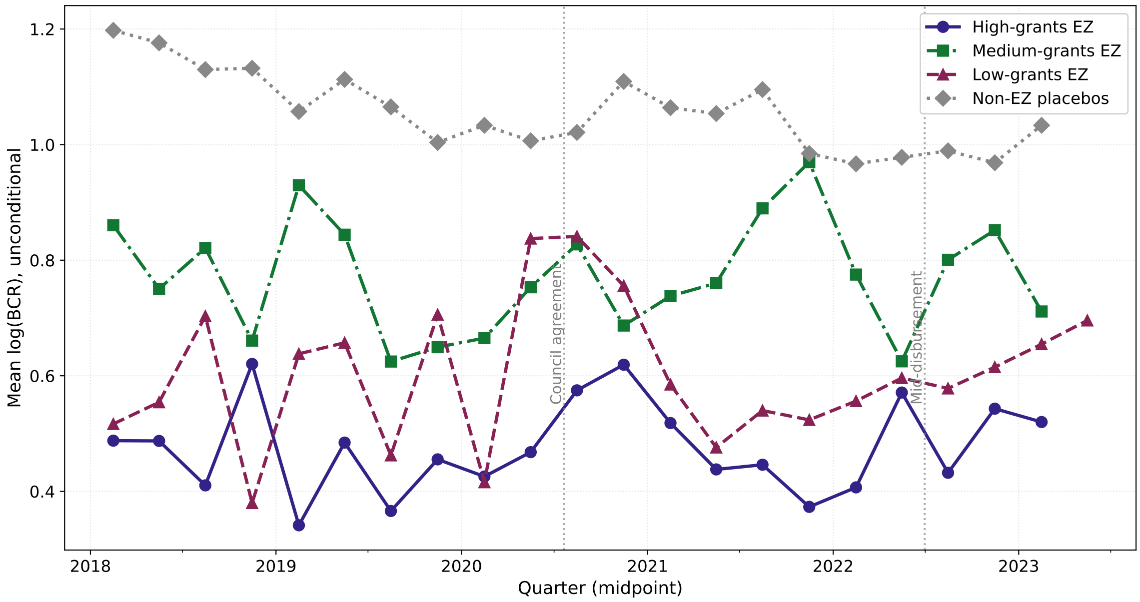

Notes: The figure is descriptive and does not condition on country fixed effects, auction size, or other controls. Formal pre-trend tests are reported in section 6. High-grants: Portugal, Spain, Italy. Medium-grants: France. Low-grants: Austria, Belgium, Finland, Germany. Non-EZ: Australia, Canada, Japan, New Zealand, South Korea, United Kingdom.

3.3 NGEU treatment intensities

The treatment is based on national Recovery and Resilience Plan (NRRP) allocations under the RRF. Let G_c denote approved grants, L_c approved loans, and $Y_{c,2019}$ nominal GDP in 2019. We use three intensity measures:

$$\text{grants}_c \equiv \frac{G_c}{Y_{c,2019}}, \qquad \text{loans}_c \equiv \frac{L_c}{Y_{c,2019}}, \qquad \text{actual}_c \equiv \frac{G_c + L_c}{Y_{c,2019}}.$$

Two further intensity codings appear as alternative regressors in the headline specification of section 5.1. The first is the *formula* allocation key formula$_c$: each country's implied grants share under the 2021 Council allocation formula, which distributes 70 percent of grants according to a function of population, inverse GDP per capita, and the 2015–2019 average unemployment rate, and 30 percent according to the same indicators with realised GDP-loss measures from 2020–2021 replaced by their pre-pandemic projections (Darvas, 2020, 2021). Using projected rather than realised GDP losses produces a country-specific allocation key that is a Bartik-style shift-share instrument: it reflects the formula-implied entitlement that countries would have received under purely pre-pandemic indicators, and so is plausibly exogenous to the auction-period BCR dynamics that the treatment is meant to explain. The second alternative coding is a *binary* high-exposure indicator binary_c equal to one for Italy, Spain, and Portugal — the three peripheral euro-area sovereigns with grants intensity above 0.03 — and zero for all other countries. Non-euro-area placebo issuers take value zero on every intensity measure.

The baseline treatment is grants intensity interacted with an indicator for auctions on or after 21 July 2020. The grants margin is the preferred treatment because grants and loans have different economic content. Grants are transfers from the common EU budget and do not create repayment obligations for the recipient sovereign. Loans provide financing on favorable terms but also create liabilities. The grants-only measure therefore maps more directly into the unconditional-transfer mechanism studied in the paper, while the actual measure is useful as a comparison because it aggregates grants and loans.

Table 2 reports the intensity measures for the eight euro-area sovereigns. Non-euro-area placebo issuers take value zero on all three measures. The distinction between grants and total allocations matters most for Italy. On the grants-only measure, the high-exposure ranking is Portugal (0.065), Spain (0.062), and Italy (0.038). On the grants-plus-loans measure, Spain remains high (0.130), but Italy rises to 0.107 because it drew its full €122.6 billion loan envelope. This difference in treatment rankings is the basis for the grants-versus-loans horse race in section 5.

Table 2: Three NRRP intensity measures for the eight euro-area treated sovereigns

Country	Grants (€bn)	Loans (€bn)	GDP_{2019} (€bn)	Actual intensity	Grants intensity	Loans intensity
Austria	3.5	0.0	397.6	0.009	0.009	0.0000
Belgium	5.0	0.0	480.0	0.010	0.010	0.0000
Finland	2.1	0.0	240.6	0.009	0.009	0.0000
France	39.4	0.0	2437.6	0.016	0.016	0.0000
Germany	28.0	0.0	3473.3	0.008	0.008	0.0000
Italy	68.9	122.6	1796.1	0.107	0.038	0.0683
Portugal	13.9	2.7	214.4	0.077	0.065	0.0126
Spain	77.2	84.0	1244.8	0.130	0.062	0.0675

Notes: Grants and loan allocations come from the approved Recovery and Resilience Plans as ratified by the Council. Non-euro-area sovereigns take value zero on all three. Source: European Commission RRF scoreboard; Eurostat national accounts for 2019 nominal GDP.

3.4 Controls, fixed effects, and sample

The baseline specification includes country fixed effects, year fixed effects, and the log of competitive allotment. Country fixed effects absorb persistent differences across issuers, including average BCR levels, auction institutions, investor bases, and currency regimes. Year fixed effects absorb common annual shocks such as the COVID liquidity episode, the 2022 energy-price shock, and the 2022–2023 monetary-tightening cycle. The log-competitive-allotment control — labelled the *supply control* in Table 3 and throughout the paper — accounts for the relationship between auction size and bid intensity: issuers can adjust auction size in response to market conditions, and BCR mechanically combines bids with allotment. Columns (2)–(4) of Table 3 are estimated without the supply control by design rather than as a sensitivity: those columns reproduce the canonical intensity-variant exercise, which uses a common no-supply specification across alternative intensity codings to keep the codings on a like-for-like footing.

Additional variables enter only in robustness exercises. Maturity is handled through 10-year and 30-year splits in section 6. Secondary-market 10-year yields and 20-day yield volatility are available for 487 auctions in a five-country euro-area subsample and are used to check whether the result is absorbed by contemporaneous secondary-market conditions. The MOVE index of US Treasury volatility is used as a global bond-market volatility control on the full 14-country panel.

The sample ends on 4 July 2023. Extending the panel further would add a period of intensive ECB tightening and APP runoff, during which issuance patterns and secondary-market conditions changed sharply. We therefore keep the baseline sample focused on the period from the pre-NGEU window through the main early implementation years of the RRF.

4 Empirical strategy

4.1 Baseline specification

The empirical design compares changes in auction demand after the July 2020 NGEU agreement across sovereigns with different grants exposure. The baseline specification is

$$\log \mathrm{BCR}_{c,t} = \alpha_c + \gamma_{y(t)} + \beta \left(\mathrm{grants}_c \times \mathbf{1}\{t \geq \text{2020-07-21}\}\right) + \kappa \log \mathrm{CA}_{c,t} + \varepsilon_{c,t},$$

where c indexes countries and t indexes auctions. The outcome is the log bid-to-cover ratio. The term α_c is a country fixed effect, $\gamma_{y(t)}$ is a year fixed effect, grants_c is the grants-only NRRP intensity reported in Table 2, and $\log \mathrm{CA}_{c,t}$ is the log competitive allotment at the auction. The coefficient of interest, β, measures the post-agreement change in log-BCR associated with a one-unit increase in grants intensity. Standard errors are clustered by country.

Three specification choices are central. First, the preferred treatment is grants-only intensity rather than the broader grants-plus-loans envelope. This follows the economic mechanism: grants are unconditional transfers, while loans provide financing but also create repayment obligations. The grants-only specification therefore targets the component of NGEU most likely to affect sovereign net resources.

Second, the specification controls for log competitive allotment. Bid-to-cover is a ratio of bids to allotment, and sovereign issuers can adjust auction quantities in response to market conditions. Controlling for allotment helps separate stronger bidding pressure from changes in the size of the auction itself, a concern emphasized in the sovereign auction literature (Beetsma et al., 2020).

Third, the baseline uses country and year fixed effects. Country fixed effects absorb persistent differences in auction institutions, investor bases, and average bid-to-cover levels. Year fixed effects absorb broad common shocks, including the COVID-19 liquidity window, the energy-price shock, and the beginning of the ECB tightening cycle. Finer time effects are examined in robustness exercises, but the baseline keeps the time structure parsimonious because the auction panel is unbalanced and the treatment varies at the country level.

4.2 Inference with few clusters

The treatment varies at the country level, while the panel contains only fourteen countries and the positive grants exposure is concentrated in Italy, Spain, Portugal, and France. The asymptotic cluster-robust standard error formula is derived under $G \to \infty$ and is known to under-cover (reject too often) when G is small, especially when treatment is concentrated in a few clusters and the cluster-size distribution is uneven (Cameron et al., 2008; MacKinnon and Webb, 2017). Treating the asymptotic p-value as definitive in this setting would overstate the inferential strength of the headline. We therefore complement it with four finite-sample-valid procedures, each chosen to address a specific failure mode of the asymptotic test.

Wild cluster bootstrap with Rademacher weights. The first procedure replaces the asymptotic standard error with a finite-sample reference distribution that respects the cluster structure of the data. Following Cameron et al. (2008), we impose the null on the residuals from the restricted regression, then multiply each cluster's residual vector by a random sign $s_g \in \{-1, +1\}$ and re-estimate. With $G = 14$ clusters there are only $2^{14} = 16{,}384$ possible sign patterns, so we enumerate them exactly rather than sample. The reported p_{WCR} is the share of patterns with absolute pseudo-coefficient at least as large as the observed coefficient. This procedure addresses the failure mode of asymptotic SE under-coverage; it does not address whether the underlying identifying variation is unusual relative to chance.

Country-label permutation. The second procedure asks whether the specific mapping of grants intensities to countries is unusually predictive of post-2020 BCR. Following Young (2019) and MacKinnon and Webb (2020), we randomly reassign the 14-number grants-intensity vector across the 14 countries in the panel over 50,000 draws and re-estimate the baseline regression on each reshuffle. The reported p-value is the share of reshuffles with absolute pseudo-coefficient at least as large as observed. This addresses a failure mode the wild cluster bootstrap does not touch: chance alignment between cross-country heterogeneity and the particular intensity numbers attached to each country. Section 7 reports the corresponding results.

Within-euro-area restricted permutation. The full country-label permutation blends two sources of identifying power — the within-euro-area gradient and the contrast between NGEU-exposed euro-area sovereigns and non-euro placebo issuers — and the resulting p-value is influenced by both. A stricter benchmark fixes the six non-euro-area placebos at zero intensity in every reshuffle and permutes only the eight euro-area intensities among themselves; the full enumeration contains $8! = 40{,}320$ assignments and the procedure is exact. This isolates the within-euro-area gradient as the sole source of identifying variation and is the demanding inference benchmark in the paper. Section 7 reports the corresponding results.

Time-placebo calibration. The fourth procedure asks whether the baseline specification, applied to arbitrary pre-NGEU dates, would routinely deliver coefficients as large as the headline. We re-estimate the regression at 300 random fake event dates drawn from trading days between January 2018 and June 2020, with the grants-intensity vector and the BCR data unchanged. The observed post-21-July-2020 coefficient is compared to the empirical distribution of placebo coefficients. This addresses a third failure mode — specification-induced coincidence between cross-country heterogeneity and the chosen event date — that neither the bootstrap nor the country-label permutation rules out on its own.

Appendix B provides implementation details for all four procedures. The paper reports the 14-country specification alongside the euro-area-only specification throughout, and the contrast between them (in particular, the within-EZ restricted permutation result of $p = 0.089$) is treated as substantive evidence about the design's identifying power.

4.3 Identifying assumption and threats

The identifying assumption is that, conditional on country fixed effects, time effects, and auction-size controls, high-grant and low-grant sovereigns would not have experienced systematically different post-2020 changes in log bid-to-cover ratios in the absence of NGEU. This is a parallel-trends assumption for a continuous country-level treatment.

The assumption is substantive rather than automatic. Grants were not randomly assigned: the allocation formula loads on pandemic exposure and pre-pandemic economic conditions, and the high-grants sovereigns also differ from the low-grants sovereigns in fiscal position, investor base, and exposure to euro-area risk. The empirical strategy therefore uses pre-trend, permutation, placebo, and robustness exercises as discipline on the design, not as proof of exogeneity.

Five threats are particularly relevant. First, grants intensity is correlated with pandemic severity and fiscal vulnerability. If high-grants countries were already on a different recovery path in auction demand, the headline coefficient could partly reflect that trajectory. Second, the euro-area periphery faced differential sovereign-risk dynamics over 2021–2023 for reasons unrelated to NGEU, including the energy shock, the ECB tightening cycle, and the 2023 banking stress episode. Third, PEPP and PSPP may affect primary-market demand indirectly through secondary-market yields, liquidity, volatility, and dealer balance sheets. The primary-market setting removes direct purchase-flow contamination at issuance, not all monetary-policy transmission. Fourth, the non-euro-area placebo issuers operate under different currency regimes, monetary-policy frameworks, auction institutions, and investor bases. They help absorb global shocks, but they are not perfect counterfactuals for euro-area sovereigns. Fifth, because the relevant treatment variation is country-level and concentrated in a few sovereigns, the design is sensitive to leverage and small-cluster inference.

The subsequent sections map these threats into diagnostics. The binned event-study examines pre-trends. The exact wild cluster bootstrap and country-label permutation address few-cluster inference. The within-euro-area restricted permutation separates the within-euro-area treatment gradient from the broader euro-versus-placebo contrast. Leave-one-country-out estimates assess country leverage. The MOVE-control specification checks whether global bond-market volatility drives the result. The cumulative-disbursement horse race asks whether the timing is better explained by realized cash flows or by approved grant exposure. Taken together, these exercises define the scope of the causal interpretation: the paper provides disciplined reduced-form evidence on auction demand, not a fully structural decomposition of the NGEU mechanism.

5 Main results

5.1 Baseline result

Table 3 reports the baseline specification and three alternative ways of measuring NGEU exposure. The preferred measure is grants-only intensity, because it corresponds to the unconditional-transfer component of the RRF. In column (1), higher grants exposure is associated with a statistically reliable increase in auction demand after the July 2020 agreement. The coefficient is positive, economically large, and lies outside the country-label permutation interval. The exact wild cluster bootstrap and the country-label permutation both reject the null at conventional levels.

Table 3: Baseline two-way fixed-effects difference-in-differences

Intensity measure:	(1) Grants-only	(2) Actual	(3) Formula	(4) Binary
$\hat{\beta}$ (log-BCR per unit)	1.921	0.723	1.142	+0.075
Cluster-robust SE	(0.485)	(0.246)	(0.408)	(0.027)
Exact WCR p-value	0.036	—	—	—
Permutation p-value	0.021	0.118	0.113	0.116
Observed inside 95% perm CI?	No	Yes	Yes	Yes
N	1,471	1,471	1,471	1,471
Countries	14	14	14	14
Supply control	Yes	No	No	No

Notes: Column (1) is the headline specification with grants-only intensity, country and year fixed effects, and log competitive allotment as supply control. Columns (2)–(4) replace the intensity variable with the actual (grants-plus- loans) intensity, the 2021 Council formula, and a binary high-intensity dummy. Standard errors are cluster-robust on country. The exact wild cluster bootstrap enumerates all 2^{14} sign patterns. The permutation test reshuffles the intensity vector across the fourteen countries 50,000 times. (See section 4.2)

The contrast across columns is informative. The broader grants-plus-loans measure, the formula measure, and the binary high-exposure indicator all deliver positive point estimates, but none passes the country-label permutation test. The data therefore do not simply say that any measure of NGEU exposure predicts stronger auction demand. The evidence is strongest when exposure is measured on the grants margin, which is the component most closely aligned with the transfer mechanism.

5.2 Economic magnitudes

The coefficient in column (1) is a semi-elasticity with respect to a country-level intensity measure. To express the result in the scale of the outcome, we multiply the coefficient by each country's grants intensity and exponentiate: $\exp(\hat{\beta} \times \text{grants}_c) - 1$. Table 4 reports the implied responses.

Table 4: Country-level implied BCR responses

Country	Grants intensity	$\hat{\beta} \times \text{grants}_c$	Implied ΔBCR
Portugal	0.065	0.125	13.3%
Spain	0.062	0.119	12.6%
Italy	0.038	0.073	7.6%
France	0.016	0.031	3.1%
Belgium	0.010	0.019	1.9%
Austria	0.009	0.017	1.7%
Finland	0.009	0.017	1.7%
Germany	0.008	0.015	1.5%

Notes: The third column is the linear implied log-change; the fourth is the exponentiated percent change in the level of BCR.

Figure 2: Country-level implied post-2020 BCR responses

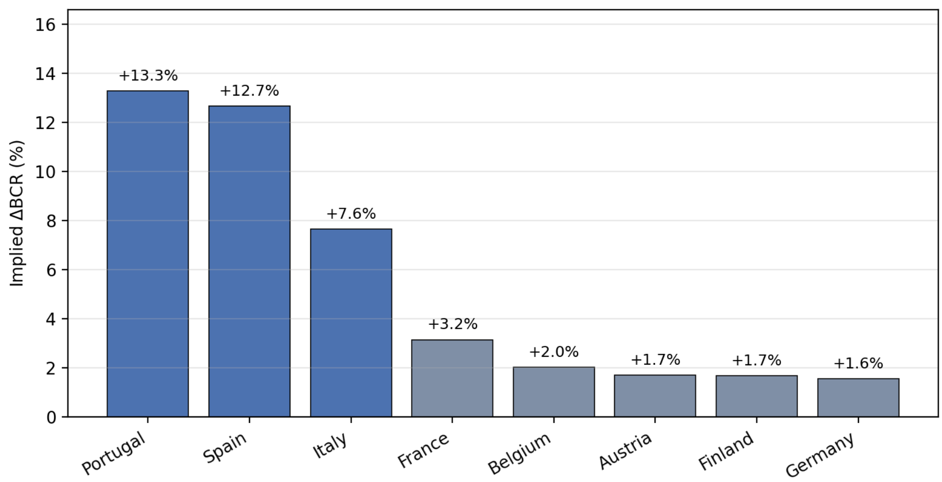

Notes: Each bar is $\exp(\hat{\beta} \times \text{grants}_c) - 1$ in percent. Portugal, Spain, and Italy (dark blue) are the three peripheral euro-area sovereigns with material grants exposure.

The magnitudes are meaningful but not implausibly large. The estimated responses are about 13% for Portugal and Spain and about 8% for Italy, while the low-exposure core countries have implied responses below 2%. For Spain and Portugal, the implied log-change is roughly 0.4 pre-period standard deviations of euro-area log BCR. This is large enough to matter for auction demand, but still within the range of normal auction-day variation documented in the pre-NGEU auction literature.

The country pattern is also informative. Italy receives a smaller implied response than Spain or Portugal even though it received a large absolute NGEU envelope. This follows directly from the grants-only treatment: Italy's loan envelope is large, but its grant intensity is lower than those of Spain and Portugal. The result therefore lines up with the interpretation that investors respond more strongly to the transfer component than to the total nominal envelope.

5.3 Grants versus loans

The baseline result uses grants-only intensity. A natural next question is whether the loan component carries similar information for auction demand. We therefore estimate a horse-race specification in which grants intensity and loans intensity enter jointly, each interacted with the post-21-July-2020 indicator. This specification asks whether the two components of NGEU are priced symmetrically by primary-market investors.

Table 5: Grants-versus-loans horse race

	C1 grants-only	H3b grants	H3b loans
$\hat{\beta}$ (log-BCR per unit)	1.921	2.690	-0.844
Cluster-robust SE	(0.485)	(0.846)	(0.580)
95% CI	$[0.97, 2.87]$	$[1.03, 4.35]$	$[-1.98, 0.29]$
Exact WCR p-value	0.036	0.003	0.073
Permutation p-value	0.021	0.002	0.271
N	1,471	1,471	1,471
Countries	14	14	14

Notes: Column 1 reproduces the grants-only headline from Table 3. Columns 2–3 report the two coefficients from a single regression that includes grants intensity and loans intensity, each interacted with the post-21-July-2020 indicator, alongside country and year fixed effects and the log-competitive-allotment supply control. Joint F-test of $\beta_G = \beta_L = 0$: $F = 7.69$, $p = 0.0005$; test of equality $\beta_G = \beta_L$: $F = 6.64$, $p = 0.010$.

Table 5 shows that the grants and loans margins behave differently. Once loans intensity is included, the grants coefficient becomes larger than in the grants-only baseline. The loan coefficient enters with the opposite sign and is statistically distinguishable from the grants coefficient. The loan estimate itself is not the main claim: it is suggestive under the wild cluster bootstrap and not significant under the country-label permutation. The stronger conclusion is the asymmetry between the two components.

This asymmetry is central to the economic interpretation. If the result were driven only by a broad "large NGEU recipients did better" pattern, grants and loans should carry similar information. Instead, the response is concentrated on the grants margin. That pattern is consistent with the idea that unconditional transfers improve perceived sovereign resources in a way that on-lending does not. Loans may still matter through favorable funding terms, rollover-risk relief, or conditionality, but they are not equivalent to grants from the perspective of private auction demand.

Section 8 returns to this mechanism and connects it to the timing evidence and the realized-disbursement horse race.

6 Pre-trend, maturity, and sample robustness

This section reports the main stress tests for the headline specification. The purpose is not to present every diagnostic as an independent confirmation, but to separate three pieces of evidence: whether the pre-period is compatible with the identifying assumption, whether the result is concentrated in a particular maturity or sample composition, and where the evidence becomes weaker under more demanding specifications. The overall pattern is that the coefficient remains positive across the main alternatives, while exact inference is strongest in the full 14-country panel and weaker when the design is restricted to the euro-area countries alone or saturated with country-specific trends.

6.1 Binned event-study pre-trend check

The identifying assumption is a continuous-treatment version of parallel trends: absent NGEU, high-grants and low-grants sovereigns would not have experienced systematically different changes in log-BCR after conditioning on the fixed effects and auction-size control. A direct way to probe this assumption is to replace the post-treatment interaction with event-time bins and inspect the pre-event coefficients (Athey and Imbens, 2017; Roth, 2022).

Table 6 reports the binned coefficients and Figure 3 plots them. The reference window is the twelve months before the 21 July 2020 agreement. The two pre-event bins cover the period twelve to twenty-four months before the agreement (pre_{2y}) and twenty-four months or more before (pre_{2+}). Four post-event bins cover the immediate announcement window (months 0–6), the first disbursement year (months 6–18), the second disbursement year (months 18–30), and the tail (thirty months or more).

Table 6: Binned event-study coefficients on grants intensity

Bin	Period (relative to July 2020)	$\hat{\beta}$	p_{asymp}	p_{WCR}
pre_{2+}	$\geq$ 24 months before	1.260	0.082	0.216
pre_{2y}	12 to 24 months before	0.948	0.331	0.380
(reference)	1 to 12 months before	0	—	—
post_{h1}	0 to 6 months after	2.663	0.003	0.259
post_{1y}	6 to 18 months after	1.471	0.146	0.138
post_{2y}	18 to 30 months after	3.487	0.002	0.039
post_{3y+}	$\geq$ 30 months after	2.928	0.082	0.072

Notes: The dependent variable is log-BCR; each reported coefficient is the interaction of grants intensity with an event-time bin. Coefficients are in log-BCR units per unit grants intensity. The pre-event bins are not significant under the exact wild cluster bootstrap; the disbursement-phase bin (post_{2y}) delivers the largest coefficient and the tightest honest p-value.

Figure 3 shows the pattern visually: the two pre-event bins are positive but lie well inside their confidence intervals, while the post-event bins rise from a modest announcement-window response to a peak in the second-disbursement-year bin. Appendix Figure 8 and Appendix Table 23 report the same event-study coefficients under three alternative binning schemes (annual, six-monthly, and the current scheme); the location of the mass in the disbursement window is robust across them.

Two patterns are useful for interpretation. First, the pre-event coefficients are positive but imprecise. The two pre-period bins do not reject the parallel-trends null under the exact wild cluster bootstrap: pre_{2y} has an exact WCR p-value of 0.38, and pre_{2+} has an exact WCR p-value of 0.22. This does not prove parallel trends, and the positive pre-coefficients caution against an overly strong reading. It does show that the pre-period differences are not statistically

Figure 3: Binned event-study coefficients on grants intensity

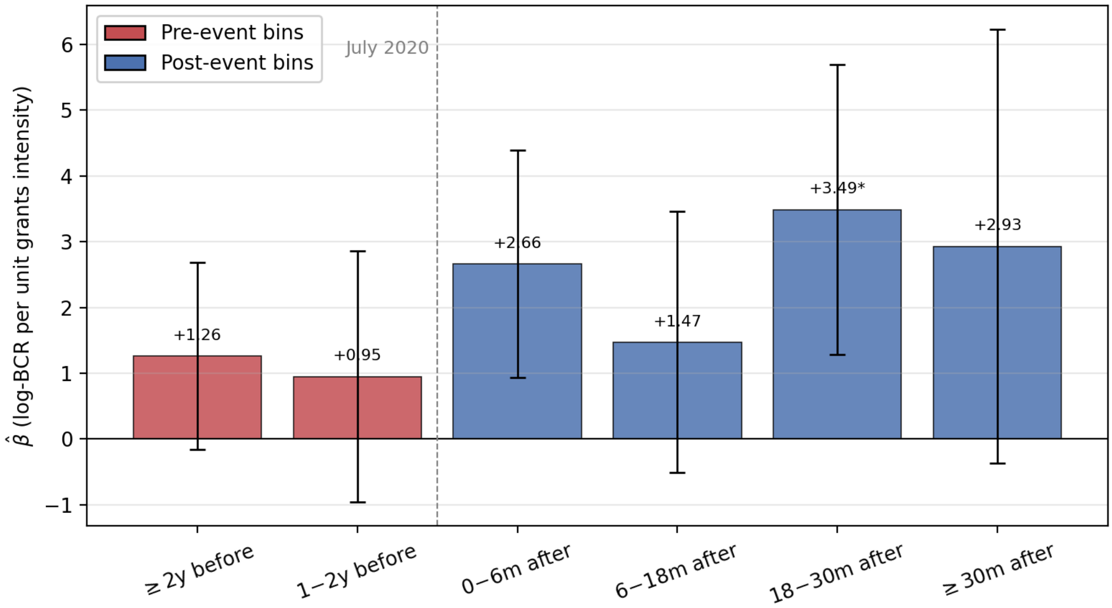

Notes: Red bars are pre-event bins (reference window dropped); blue bars are post-event bins. Whiskers show the $\pm 1.96 \cdot$ SE asymptotic interval. A star marks exact-enumeration wild cluster bootstrap significance at the 5% level. The second-disbursement-year bin (post_{2y}) delivers the largest coefficient and the tightest honest p-value.

distinguishable from zero under the same few-cluster inference used for the headline result.

Second, the post-event coefficients are largest after the initial announcement period. The immediate 0–6 month window has a large point estimate but is not significant under exact inference. The strongest post-period estimate appears in the 18–30 month bin, with an exact WCR p-value of 0.039. Splitting this bin further shows that the peak is concentrated in July–December 2022: the 18–24 month sub-bin has an exact WCR p-value of 0.073, while the 24–30 month sub-bin has a coefficient of 3.951 and an exact WCR p-value of 0.046. The timing is therefore more consistent with a gradual credibility-building interpretation than with a pure announcement effect. Section 8 returns to this point.

6.2 Maturity split

A first robustness concern is that the pooled estimate may be driven by a single maturity segment. Table 7 therefore reports the headline specification separately for the 10-year and 30-year auctions.

The two maturity-specific point estimates are positive and close to the pooled estimate. The 10-year coefficient is 1.59 and the 30-year coefficient is 1.67. Neither subsample, however, clears the exact wild-cluster threshold at 5%. This weakening is unsurprising given the smaller samples, especially the 30-year subsample with 450 auctions. The maturity evidence therefore supports the sign and approximate magnitude of the pooled estimate, but the precision comes from pooling maturities. Appendix Table 20 reports the full five-specification maturity grid, including a specification with a maturity-bucket dummy and one with the three-way interaction grants ×

Table 7: Baseline headline specification estimated on the full panel, the 10-year subsample, and the 30-year subsample

Subsample	$\hat{\beta}$	SE	p_{asymp}	p_{WCR}
All maturities (baseline)	1.921	0.485	0.00008	0.036
10-year only	1.587	0.462	0.001	0.074
30-year only	1.675	1.149	0.146	0.161

Notes: The full-panel specification is the column (1) headline from Table 3. The 10-year and 30-year splits use identical fixed effects and supply controls.

post $\times$ $\mathbf{1}\{\text{maturity} = 30\}$.

6.3 Sample-composition robustness

The second concern is sample composition, especially the role of the non-euro-area placebo countries. Table 8 reports the headline grants-only specification on the full panel, the full panel without the supply control, the euro-area-only panel, and several euro-area restrictions.

Table 8: Sample-composition robustness of the headline grants- only coefficient

Sample	$\hat{\beta}$	SE	p_{asymp}	p_{perm}
14-country (baseline)	1.921	0.485	< 0.001	0.021
14-country, no supply ctrl	1.790	0.481	0.00021	0.026
8-country EZ only	1.578	0.510	0.002	0.079
EZ7 (drop Finland)	1.687	0.508	0.001	0.070
EZ7 (drop Austria)	1.331	0.535	0.013	0.195
EZ6 (drop AT & FI)	1.460	0.533	0.006	0.139

Notes: Each row is the same headline specification estimated on a different slice of the panel. The baseline 14-country row matches column (1) of Table 3. Dropping the non-euro-area placebos reduces the effective variation in the time fixed effects and widens the permutation p-value into the 0.05–0.10 band.

The coefficient remains positive in every sample, but precision depends on the full panel. In the euro-area-only sample, the point estimate is still sizable and close to the baseline, but the permutation p-value rises to 0.079. Dropping Austria or Finland further weakens the evidence. This is an important qualification: the non-euro placebo countries are not merely decorative. They contribute to common-shock absorption and to the precision of the full-sample design. The appropriate conclusion is therefore not that the result is driven by one country, but that the most decisive inference uses the combination of within-euro-area treatment variation and the broader placebo contrast. Appendix Table 13 reports the full 29-row sample-composition grid (14 leave-one-country-out specifications, 7 high-exposure exclusions including Italy/Spain/Portugal combinations, and 8 euro-area-only restrictions).

6.4 Compact robustness summary

Table 9 summarizes the main robustness patterns. This table is meant to discipline the interpretation of the paper. The result is strongest in specifications that preserve the full 14-country panel and allow time effects to absorb common shocks without over-saturating the country-level treatment variation. It weakens when the sample is restricted to the euro area alone, split by maturity, or saturated with country-specific trends.

Table 9: Compact robustness summary

Specification	$\hat{\beta}$	Honest p	Interpretation
Baseline, 14-country panel	1.921	0.036	Headline grants-only specification, exact WCR.
Half-year fixed effects	1.620	0.034	Finer time fixed effects saturate within-year shocks while preserving cross-country variation.
+ MOVE volatility control	1.796	0.039	Adds global Treasury-volatility control to the full panel.
Equal-country weights	1.813	0.049	Reweights to give each country equal influence.
10-year only	1.589	0.074	Single-tenor sub-panel, $N = 1{,}021$.
30-year only	1.675	0.161	Single-tenor sub-panel, $N = 450$; smaller power.
Euro-area only ($G = 8$)	1.221	0.164	Drops the six non-euro placebos.
Within-EZ restricted permutation	1.921	0.089	Full 8! enumeration; within-EZ gradient alone.
Country-specific linear trends	1.981	0.346	Saturates pre-event trends; identification power falls.

Notes: Each row reports the headline grants-only coefficient on a different specification. Honest p-value is the exact wild cluster bootstrap p-value (p_{WCR}) for all rows except the within-EZ row, which uses the full 8!-enumeration restricted permutation p-value. Sources by row: baseline reproduces Table 3; half-year row uses section 6.5 (full grid in Appendix Table 16); MOVE-control row uses Appendix Table 17; equal-country-weights row uses Appendix Table 19; 10-year and 30-year rows reproduce Table 7 (full maturity grid in Appendix Table 20); euro-area-only row reproduces row of Table 8 (full sample-composition grid in Appendix Table 13); within-EZ permutation row uses section 7 and Appendix Table 24; country-specific linear trends row uses Appendix Table 15.

The compact summary gives a balanced message. The positive coefficient is not fragile in sign: it remains positive under the listed alternatives. Nor is it driven by high-frequency issuers, since equal-country weights leave the estimate close to the baseline. At the same time, exact inference is not uniformly strong. The euro-area-only, within-EZ permutation, and country-trends rows show where the design is least powerful. These rows are central to the interpretation of the paper and motivate the cautious wording used throughout.

6.5 Half-year fixed effects

A natural concern is that annual fixed effects are too coarse to absorb within-year shocks that may correlate with the post-21-July-2020 indicator. Replacing year fixed effects with calendar-half fixed effects absorbs more time variation while preserving enough cross-country treatment variation to estimate the dose-response. Under this specification the coefficient is 1.620, with an exact wild cluster bootstrap p-value of 0.034. The finer time controls therefore leave the headline result qualitatively unchanged. Appendix Table 16 reports the full eight-specification grid spanning year, half-year, quarter-year, month-year, country-by-maturity, country-specific linear trends, and the pre-event detrended outcome; Appendix Table 15 adds a complementary grid that includes the country-trends and raw intensity-coding variants. Appendix Table 18 reports the same headline under seven alternative transformations of the BCR outcome (log, levels, country-standardised log, country-demeaned log, winsorised log, inverse hyperbolic sine, and the residualised log); the sign and significance of the headline coefficient are preserved across transformations.

6.6 Leave-one-country-out diagnostics

Appendix Figure 4 reports the leave-one-country-out estimates. No single-country exclusion reverses the sign of the coefficient. The fourteen single-country-drop estimates range from 1.609 when New Zealand is excluded to 2.264 when the United Kingdom is excluded, bracketing the baseline estimate of 1.921. This check is especially useful because the treatment intensity is concentrated in a small number of countries: the headline result is not mechanically driven by removing or retaining one particular country.

The more demanding restrictions are the euro-area-only and high-exposure drop samples. Dropping all six non-euro-area placebos leaves an EZ-only panel of eight countries and 695 auctions. The coefficient remains positive, but the exact wild cluster bootstrap p-value rises to 0.164. This is consistent with the within-EZ restricted permutation in Section 7, which yields a p-value of 0.089. The result should therefore be read as a full-panel finding that combines a within-euro-area grants gradient with the contrast between NGEU-exposed euro-area issuers and non-euro placebo issuers, not as a result identified solely within the euro-area sample.

7 Validation: permutation, time placebo, and auxiliary outcomes

This section reports three validation exercises designed to probe whether the headline result is unusually strong relative to plausible placebo assignments, placebo timing, and auxiliary auction outcomes. The exercises do not replace the identifying assumption in Section 4; rather, they discipline the interpretation of the baseline estimate in a setting with few country clusters and concentrated treatment exposure.

7.1 Country-label permutation

The country-label permutation test asks one specific question: if NGEU grants intensity had been randomly assigned to countries – i.e., the mapping (Spain $\rightarrow$ 0.062, Italy $\rightarrow$ 0.038, Australia $\rightarrow 0, \ldots$) had been a coin toss – how often would we have gotten a coefficient as large as the one we actually observe? If the observed assignment looks unusual relative to that distribution of random reassignments, the data are saying the actual country-to-intensity mapping really is doing the work, not chance. We reassign the grants-intensity vector uniformly at random across the fourteen countries, re-estimate the headline specification for each reshuffle, and compare the observed coefficient to the empirical distribution of pseudo-coefficients (Young, 2019; MacKinnon and Webb, 2020). Under the sharp null of no treatment effect, the observed assignment should not look exceptional relative to these placebo assignments.

The reported p-value is two-sided, computed as the fraction of the 50,000 reshuffles whose absolute pseudo-coefficient exceeds the absolute observed coefficient. For the headline specification, the permutation p-value is 0.021 (Monte Carlo SE = 0.0006), and the observed coefficient sits outside the empirical 95% interval $[-1.087, 1.889]$. Appendix Figure 5 plots the full permutation distribution.

The permutation exercise also helps distinguish the grants margin from broader NGEU exposure. The actual-intensity (grants-plus-loans) specification has a permutation p-value of 0.118 in Table 3, and its observed coefficient falls inside the corresponding permutation interval. This contrast supports the grants-only specification: the randomization evidence is stronger for the transfer component of NGEU than for the total envelope that combines grants and loans. The grants-versus-loans horse race in section 5.3 gives the same message from a different angle: once grants and loans enter separately, the grants coefficient strengthens, the loans coefficient enters with the opposite sign, and equality of the two coefficients is rejected at the 1% level.

Within-EZ restricted permutation

The full country-label permutation uses both within-euro-area variation in grants intensity and the contrast between NGEU-exposed euro-area sovereigns and non-euro placebo issuers. The non-euro placebo contrast contributes meaningful identifying power because the placebo countries are distinct from the treated countries on currency regime, monetary stance, and investor base, and several of them carry small but non-zero grants intensities in the permuted distribution. A stricter benchmark fixes the six non-euro-area countries at zero intensity and permutes only the eight euro-area intensities among themselves. The full enumeration over these within-EZ assignments gives a two-sided p-value of 0.089, reported in the within-EZ row of Table 9 and as procedure M2 of the inference-suite consolidation in Appendix Table 24.

This result is an important qualification. The within-euro-area gradient alone does not reject at the 5% level under exact restricted permutation. The sign and magnitude remain positive, but exact inference is less decisive once the non-euro placebo contrast is removed. The full-sample evidence should therefore be read as combining within-euro-area treatment variation with the additional information provided by the non-euro placebo issuers.

7.2 Time-placebo calibration

The time-placebo exercise asks whether the same specification would generate similarly large estimates if the event date were moved to dates before NGEU existed. The intensity vector and fixed effects are held fixed, while the treatment date is varied. We draw 300 random fake event dates from the pre-NGEU window (2018-01 through 2020-06) and re-estimate the headline specification with each fake date as the treatment date. The resulting coefficients form an empirical distribution of estimates that the specification would generate under placebo timing.

Table 10 reports the placebo distribution summary. The observed post-July-2020 coefficient is 1.921 (reproduced from column (1) of Table 3). The placebo distribution has mean 0.485 and standard deviation 0.403, with an empirical 2.5th percentile of -0.218 and an empirical 97.5th percentile of 1.143. The observed coefficient lies outside this placebo 95 percent interval: no random pre-NGEU date in the calibration produced a coefficient of comparable magnitude. The time-placebo false-positive rate at $\alpha = 0.05$ is 6.3 percent, close to nominal. Appendix Figure 6 plots the full time-placebo distribution underlying the summary in Table 10.

Table 10: Time-placebo calibration summary

Quantity	Value
Observed coefficient (true event date 21 July 2020)	1.921
Number of fake event dates (K)	300
Placebo-distribution mean	0.485
Placebo-distribution SD	0.403
Empirical 2.5th percentile of placebo distribution	-0.218
Empirical 97.5th percentile of placebo distribution	1.143
Empirical false-positive rate at $\alpha = 0.05$	6.3%

Notes: The headline grants-only specification is re-estimated at 300 random fake event dates drawn uniformly from trading days between January 2018 and June 2020, holding the grants-intensity vector and BCR data fixed. The observed coefficient lies outside the empirical 95 percent interval of the placebo distribution. The empirical false-positive rate at nominal $\alpha = 0.05$ is 6.3 percent, close to nominal.

This exercise addresses a concrete concern about small country panels: a flexible difference-in-differences specification might produce large estimates for arbitrary dates if cross-country heterogeneity happens to line up with the treatment gradient. In this application, the placebo distribution is much tighter than the observed estimate. The result therefore weighs against the interpretation that the headline coefficient is a generic artifact of the timing choice, while still leaving the identification subject to the economic threats discussed in Section 4.

7.3 Auxiliary-outcome checks

The auxiliary-outcome checks ask whether the grants-intensity pattern appears on outcomes other than the bid-to-cover ratio. We re-estimate the headline specification using two alternative auction-level outcomes: the log number of primary dealers placing competitive bids and the log competitive amount allotted. These outcomes are useful because they capture different margins of the auction process. Dealer count is an extensive-margin participation measure, while competitive allotment is an issuer-side quantity choice. A demand-pressure interpretation predicts the clearest response on BCR, not necessarily on these auxiliary margins.

Table 11 reports the results. Log primary-dealer count runs on the 9-country sub-panel where dealer counts are published, delivering a coefficient of 0.109 with a cluster-robust standard error of 0.545 and a 95% confidence interval of $[-0.96, 1.18]$. The exact wild cluster bootstrap p-value is 0.801, and the country-label permutation p-value is 0.870. Log competitive allotment runs on the full 14-country panel, delivering a coefficient of 2.136 with a cluster-robust standard error of 2.371 and a 95% confidence interval of $[-2.51, 6.78]$. The asymptotic p-value is 0.408 and the permutation p-value is 0.497.

The two auxiliary outcomes carry different information.

Dealer-count outcome. The dealer-count estimate is small relative to the BCR response. The upper bound of its 95% confidence interval implies an effect of roughly 7.6% at Spain's grants

Table 11: Auxiliary-outcome checks

Outcome	Sample	$\hat{\beta}$	p_{WCR}	p_{perm}	N
log-BCR (headline)	14c	1.921	0.036	0.021	1,471
log primary-dealer count	9c	0.109	0.801	0.870	814
log competitive allotment	14c	2.136	0.408	0.497	1,471

Notes: Same headline specification, same sample definitions, different outcomes. The log primary-dealer count regression drops the log-competitive- allotment supply control when log-CA is itself the outcome.

intensity (0.062), below the estimated BCR response for Spain. This suggests that the headline result is not simply an extensive-margin increase in the number of participating dealers.

Competitive-allotment outcome. The competitive-allotment estimate is an imprecise null. Its confidence interval is wide enough to include economically large effects of either sign at the peripheral grants intensities. The result therefore does not provide tight evidence that issuer quantity choices were unaffected; it only shows that this margin does not display a statistically comparable response in the available specification.

Taken together, the auxiliary outcomes weigh against a broad joint-movement story in which all auction-level margins moved similarly for high-grants sovereigns after 2020. The evidence is stronger on the dealer-count margin, where the estimated response is economically small, and weaker on the allotment margin, where the confidence interval is wide. The checks therefore support a demand-pressure reading of the BCR result, but they do not by themselves rule out all supply-side or auction-design channels. The complementary evidence is the grants-versus-loans horse race in Section 5.3: a generic "peripheral-auctions-in-2022" story would have less reason to distinguish sharply between grants and loans, whereas the data do.

Appendix Table 21 reports the full falsification grid, combining the three auxiliary outcomes shown here with three placebo-treatment specifications (random uniform intensities, time-placebo dates, and a loans-only horse race). Appendix Table 24 consolidates the four few-cluster-robust inference procedures (exact wild cluster bootstrap, country-label permutation, within-EZ restricted permutation, and binary high-grants subset enumeration) into a single comparison for the headline specification.

8 Interpretation

This section interprets the empirical results in economic terms. The evidence points to three related messages. First, the response is tied more closely to grants than to loans. Second, the timing is more consistent with a credibility-building interpretation than with a pure announcement effect or a simple realized-cash-flow channel. Third, the estimated magnitudes are economically meaningful but not large enough to support welfare or funding-cost claims without additional structure.

8.1 Grants, not loans

The grants-versus-loans horse race reported in section 5.3 clarifies what the headline coefficient is capturing. Once loans intensity is co-estimated with grants intensity in the headline specification, the grants coefficient strengthens from 1.921 to 2.690, while the loans coefficient enters negatively at -0.844. The two coefficients are statistically distinguishable from one another: the equality test rejects at the 1% level.

The economic reading is straightforward. Grants increase public resources without creating a repayment obligation for the recipient sovereign. Loans provide financing on favorable terms but also create a liability. Their effect on private sovereign demand is therefore theoretically weaker and may depend on the spread between EU funding costs and market funding costs, on rollover-risk relief, and on investor beliefs about seniority or conditionality.

The negative point estimate on loans should not be overinterpreted: it is not significant under the country-label permutation. The stronger and more robust result is the asymmetry between the two components. The data are inconsistent with a symmetric market response to grants and loans, which supports the paper's decision to focus on the unconditional transfer component of NGEU rather than on the total grants-plus-loans envelope.

8.2 Timing is consistent with a credibility interpretation

The timing of the response is also informative. The binned event-study in Table 6 locates the strongest primary-market demand response in the second disbursement year, not at the announcement. The half-year decomposition reported in section 6 concentrates the peak further in the second half of 2022.

This pattern is difficult to read as a pure July 2020 announcement effect. A natural alternative is that market confidence in the joint-borrowing arrangement increased as the facility moved from political agreement to operational implementation. Under this interpretation, repeated successful disbursement events gradually made the NGEU commitment more credible to private investors.

A related question is whether the late-2022 concentration reflects actual cash arriving in national treasuries, rather than credibility built through implementation. To separate these interpretations, We estimate a horse-race specification that includes cumulative grants disbursed, scaled by 2019 GDP, alongside the approved-grants-times-post interaction. Table 12 reports the result; Appendix Figure 7 plots cumulative RRF grants disbursed (scaled by 2019 GDP) by country.

The realized-cash-flow interpretation receives little support from this horse race. The cumulative-disbursement coefficient is imprecisely estimated and not significant under any of the inference procedures. The approved-intensity coefficient remains positive but becomes less precise once cumulative disbursements are included, as expected when two closely related regressors are entered together. The important point is that the horse race does not shift the evidence toward realized cash flows as the independent driver of the BCR response.

We therefore interpret the late-2022 concentration as consistent with a credibility-building

Table 12: Approved-intensity versus realized-disbursement specifications

	E4 (approved × post)	E5 horse race approved × post	cum. grants disb.
$\hat{\beta}$	1.921	1.423	4.441
Cluster-robust SE	(0.485)	(0.675)	(4.757)
Asymptotic p-value	< 0.001	0.035	0.351
Exact WCR p-value	0.036	0.221	0.353
Permutation p-value	0.018	0.116	0.435
N	1,471	1,471	1,471

Notes: Column 1 is the headline specification repeated for reference. Columns 2–3 report the two coefficients from a single regression that includes approved grants intensity times post and the country's cumulative grants disbursed (scaled by 2019 GDP) as a step function in time. Joint F-test of both coefficients zero: $F = 6.29$, $p = 0.002$.

channel. This is a reduced-form interpretation, not a separately identified mechanism. The same timing could also reflect changing expectations about EU fiscal backing, reduced redenomination or fragmentation risk, improved market sentiment toward high-grants sovereigns, or shifts in expected issuance needs. The evidence supports the narrower claim that the response is not explained simply by the amount of cash already disbursed.

8.3 Magnitudes in context

The implied country-level effects in Table 4 are meaningful but not extreme. The estimates correspond to roughly 12.6% for Spain, 13.3% for Portugal, and 7.6% for Italy, or about 0.4 pre-period log-BCR standard deviations for Spain and Portugal. These are economically relevant movements in auction demand, but they remain within the range of routine auction-day variation documented in the pre-NGEU auction literature. (e.g. (Beetsma et al., 2018, 2020))

This scale is important for interpretation. The paper does not claim that NGEU eliminated peripheral sovereign risk or generated a large funding-cost effect. It shows that NGEU grant exposure is associated with a detectable increase in primary-market demand at issuance. The Italian magnitude is also informative. Italy received a large absolute grant allocation, but its grants intensity is lower than those of Spain and Portugal because its economy is larger. A grants-channel reading therefore predicts a smaller Italian response, which is what the estimated country-level magnitudes imply.

8.4 What the paper does not claim

The coefficient reported in Table 3 is an estimate of the semi-elasticity of primary-market log-BCR with respect to grants-only NRRP intensity. It is not a welfare number, not an estimate of counterfactual funding costs, and not a decomposition into specific pricing channels. An improved-creditworthiness channel, a safe-asset-supply channel (Christie et al., 2021), a liquidity channel, and a credibility-of-commitment channel are all consistent with the sign and magnitude of the observed coefficient. Distinguishing among these mechanisms would require richer auction-level allocation data, yield-tail information, or a structural demand model that separates price and quantity effects at issuance.

The paper also does not contest the secondary-market results of Corradin and Schwaab (2023), Havlik et al. (2022), or Alessi et al. (2025). Those papers ask whether NGEU changed secondary-market yields and spreads. This paper asks whether NGEU grant exposure is associated with stronger primary-market demand at issuance. The primary-market setting removes direct Eurosystem purchase-flow contamination from the auction outcome, but ECB policy can still affect bidding indirectly through secondary-market conditions.

Finally, the grants-channel and credibility interpretations are reduced-form readings of the evidence. Whether the underlying transmission ran through fiscal-multiplier expectations, reduced redenomination risk, perceived joint liability, improved market sentiment, or some combination of these channels remains outside the scope of the paper. The contribution is to document a primary-market demand response and to show that it is more closely aligned with grants than with loans, not to provide a full structural account of NGEU's macroeconomic effects.

9 Conclusion

This paper provides primary-market evidence on whether NextGenerationEU changed private sovereign demand at issuance. Using a 14-country panel of 1,471 sovereign bond auctions between 2018 and 2023, we show that sovereigns with greater exposure to NGEU grants experienced stronger bid-to-cover ratios after the July 2020 agreement. The implied magnitudes are economically meaningful: the estimated response is close to 8% for Italy and around 13% for Spain and Portugal. The result survives exact-inference procedures designed for settings with few country clusters, although the interpretation remains subject to the identification limits discussed in the paper.

The evidence points to three main conclusions. First, the composition of fiscal support matters. In specifications that separate grants from loans, the grants margin is associated with stronger auction demand, while the loan margin does not show a comparable positive effect. This pattern is consistent with the economic distinction between unconditional transfers and on-lending: grants increase the recipient sovereign's net resources without creating a repayment obligation, whereas loans provide financing while adding liabilities.

Second, the timing of the response is consistent with a credibility interpretation. The strongest effects appear in the second half of 2022 rather than immediately after the July 2020 agreement. Moreover, realized cumulative disbursements do not explain the response once approved grants intensity is controlled for. This suggests that markets may have responded not simply to cash arriving in national treasuries, but to the accumulating evidence that the joint-borrowing arrangement was operational and politically credible. This interpretation is necessarily reduced-form. The paper cannot separately identify credibility, redenomination-risk, liquidity, safe-asset-supply, or issuance-expectation channels.

Third, the primary-market setting changes the nature of the NGEU identification problem. Because Article 123 TFEU prohibits the Eurosystem from bidding at sovereign auctions, the bid-to-cover ratio is not directly contaminated by ECB purchase flows at issuance. This does not eliminate monetary-policy transmission: PEPP and PSPP can still affect auction demand indirectly through secondary-market yields, volatility, liquidity, and dealer balance sheets. The contribution is therefore not to fully separate fiscal and monetary transmission, but to study a market outcome in which direct Eurosystem purchase-flow contamination is mechanically absent.

The main limitation is that the identifying variation is country-level and concentrated in a small number of high-exposure sovereigns. The most demanding within-euro-area restricted permutation does not reject at conventional levels, so the full-sample evidence draws meaningful power from the contrast between NGEU-exposed euro-area sovereigns and non-euro placebo issuers. Those placebo issuers are useful for absorbing global shocks, but they are not perfect controls: they differ in currency regime, monetary policy, auction institutions, and investor base. The results should therefore be read as disciplined reduced-form evidence, not as a definitive causal decomposition of the NGEU mechanism.

The policy implication is correspondingly narrow. The evidence suggests that the design of supranational fiscal support matters for sovereign-market reception. In this setting, private auction demand appears more responsive to the transfer component of common borrowing than to the on-lending component. That does not imply that NGEU generated welfare gains, lowered funding costs by a specific amount, or establishes a case for permanent fiscal union. It does imply that markets may distinguish between grants and loans when assessing the sovereign-demand consequences of common fiscal instruments.

Two directions for future research follow naturally. The first is to combine auction-level data with bidder-level allocations or yield-tail information in order to separate demand, price, and quantity channels more directly. The second is to embed a joint-borrowing instrument like NGEU in a structural sovereign-debt model with rollover risk, default risk, or redenomination risk, and to use the reduced-form moments documented here to discipline the model. Both extensions would help move from the demand response documented in this paper toward a fuller account of how common European fiscal instruments affect sovereign debt markets.

During the preparation of this work the author(s) used ChatGPT and Claude code in order to assist with language editing, organization of manuscript revisions, code review, and robustness-check documentation. After using this tool/service, the author(s) reviewed and edited the content as needed and take(s) full responsibility for the content of the published article.

References

Lucia Alessi, Amela Duranovic, Virmantas Kvedaras, and Irene Monasterolo. From risks to opportunities: the impact of green public investment programs on sovereign yields. *Research in International Business and Finance*, 2025. doi: 10.1016/j.ribaf.2025.102866.

Susan Athey and Guido W. Imbens. The state of applied econometrics: causality and policy evaluation. *Journal of Economic Perspectives*, 31(2):3–32, 2017. doi: 10.1257/jep.31.2.3.

Roel Beetsma, Massimo Giuliodori, Jesper Hanson, and Frank de Jong. Bid-to-cover and yield changes around public debt auctions in the euro area. *Journal of Banking and Finance*, 87: 118–134, 2018.

Roel Beetsma, Massimo Giuliodori, Jesper Hanson, and Frank de Jong. Determinants of the bid-to-cover ratio in Eurozone sovereign debt auctions. *Journal of Empirical Finance*, 58:96–120, 2020.

A. Colin Cameron, Jonah B. Gelbach, and Douglas L. Miller. Bootstrap-based improvements for inference with clustered errors. *Review of Economics and Statistics*, 90(3):414–427, 2008. doi: 10.1162/rest.90.3.414.

Rebecca Christie, Grégory Claeys, and Pauline Weil. Next Generation EU borrowing: a first assessment. Bruegel Policy Contribution 22/2021, Bruegel, 2021.

Stefano Corradin and Bernd Schwaab. Euro area sovereign bond risk premia during the Covid-19 pandemic. *European Economic Review*, 153:104402, 2023. doi: 10.1016/j.euroecorev.2023.104402. Earlier version: ECB Working Paper 2561 (2021), with Grimm as third author.

Zsolt Darvas. Next Generation EU: 75% of grants will have to wait until 2023, 2020. Bruegel Blog Post, 10 June 2020.

Zsolt Darvas. The nonsense of Next Generation EU net balance calculations, 2021. Bruegel Blog Post.

European Parliament and Council. Regulation (EU) 2021/241 of the European Parliament and of the Council of 12 February 2021 establishing the Recovery and Resilience Facility, 2021. Official Journal of the European Union L57/17, including Annex II allocation formula.

Annika Havlik, Friedrich Heinemann, Samuel Helbig, and Justus Nover. Dispelling the shadow of fiscal dominance? Fiscal and monetary announcement effects for euro area sovereign spreads in the corona pandemic. *Journal of International Money and Finance*, 122:102578, 2022. doi: 10.1016/j.jimonfin.2021.102578.

Diogo Sampaio Lima. Identification limits in sovereign spread event studies: four diagnostics from the NGEU debate. Working paper, available from the author on request, 2026.

James G. MacKinnon and Matthew D. Webb. Wild bootstrap inference for wildly different cluster sizes. *Journal of Applied Econometrics*, 32(2):233–254, 2017. doi: 10.1002/jae.2508.

James G. MacKinnon and Matthew D. Webb. Randomization inference for difference-in-differences with few treated clusters. *Journal of Econometrics*, 218(2):435–450, 2020. doi: 10.1016/j.jeconom.2020.04.024.

Livia Pancotto, Owain ap Gwilym, and Philip Molyneux. Deal! Market reactions to the agreement on the European Union Covid-19 recovery fund. *Journal of Financial Stability*, page 101157, 2023. doi: 10.1016/j.jfs.2023.101157.

Jonathan Roth. Pretest with caution: event-study estimates after testing for parallel trends. *American Economic Review: Insights*, 4(3):305–322, 2022. doi: 10.1257/aeri.20210236.

Jun Shida. Primary market demand for German government bonds. *Journal of International Money and Finance*, 137:102909, 2023. doi: 10.1016/j.jimonfin.2023.102909.

Alwyn Young. Channeling Fisher: randomization tests and the statistical insignificance of seemingly significant experimental results. *Quarterly Journal of Economics*, 134(2):557–598, 2019. doi: 10.1093/qje/qjy029.

A Data construction and auction comparability

This appendix collects detailed data-construction and auction- comparability material referenced from §3.

Sources and sample. The auction panel is hand-collected from official national debt-management-office (DMO) bulletins, one DMO at a time, and consolidated into a single auction-level panel. The sample runs from 4 January 2018 to 4 July 2023 at the daily auction level, covering eight euro-area and six non-euro-area sovereigns. The 1,471 retained auctions are restricted to conventional fixed-coupon sovereign bonds at the 10-year and 30-year tenors; linkers, strips, and short-dated bills are excluded because they price against different benchmarks and attract different investor clienteles.

Per-DMO source URLs. The 14 DMO sources are documented country by country: Austria, `oebfa.at`; Belgium, Belgian Debt Agency (`debtagency.be`); Finland, State Treasury (`treasuryfinland.fi`); France, Agence France Trésor (`aft.gouv.fr`); Germany, German Finance Agency (`deutsche-finanzagentur.de`) and Bundesbank auction archive; Italy, Ministero dell'Economia e delle Finanze (`mef.gov.it`); Portugal, IGCP (`igcp.pt`); Spain, Tesoro Público (`tesoro.es`); Australia, AOFM (`aofm.gov.au`); Canada, Bank of Canada auction results (`bankofcanada.ca`); Japan, Ministry of Finance (`mof.go.jp`); New Zealand, NZ Debt Management (`nzdmo.govt.nz`); South Korea, Ministry of Economy and Finance (`moef.go.kr`); United Kingdom, UK DMO (`dmo.gov.uk`).

Auxiliary macro and financial data. Two auxiliary data layers are used in the robustness exercises and the ECB asset-purchase controls. Secondary-market 10-year benchmark yields for the macro-financial-controls regression in Appendix C come from FactSet Workstation (constant-maturity 10Y series; FactSet tickers of the form `TRY[XX]10Y-FDS`, where `XX` is the ISO country code). The MOVE index of US Treasury-volatility is from ICE BofA; the VIX is from FRED (series `VIXCLS`). The Eurosystem capital keys applied to construct the deviation are the post-Brexit shares effective 1 February 2020, from the ECB's capital-subscription publications. NRRP grant and loan allocations, and cumulative RRF disbursements, come from the European Commission RRF scoreboard. 2019 nominal GDP is from Eurostat national accounts.

Treatment construction. The three NRRP intensity measures are computed from the approved Recovery and Resilience Plans as ratified by the Council. Grants and loan allocations are read from the European Commission's RRF scoreboard; 2019 nominal GDP is from Eurostat national accounts. Non-euro-area sovereigns take the value zero on all three intensity measures.

Auction-format differences. The placebo set is institutionally heterogeneous. Italian sovereign auctions are run under a uniform-price (marginal) format, while most other countries in the panel use multiple-price discriminatory auctions. The literature suggests that the choice of auction format shifts mean BCR levels but does not systematically bias the slope of BCR on a continuous treatment once country fixed effects are partialled out.

Bundesbank retention. German competitive allotment is computed net of the Bundesbank-retention share, which has historically averaged around 20% of the announced amount; this shifts the level of competitive allotment used as the supply control in Germany but is absorbed by the country fixed effect.

Primary-dealer coverage. Japan and South Korea operate primary-dealer panels, but neither debt-management office publishes the number of dealers placing competitive bids at each auction. The five countries that do not publish dealer counts are Australia, Canada, Japan, New Zealand, and South Korea. They are missing from the dealer-count sub-panel used for the auxiliary-outcome falsification check in section 7, but contribute to the main BCR panel without restriction.

Data cleaning. Auction records were checked for duplicates by (country, auction date, ISIN, tenor); duplicate rows were collapsed to one observation. Auctions with missing bid-to-cover were dropped; no winsorising of BCR is applied. The dedup'd panel has $N = 1{,}471$ auctions across 14 countries.

B Inference procedures

B.1 Exact-enumeration wild cluster bootstrap

The wild cluster bootstrap is implemented exactly as in Lima (2026). For $G = 14$ clusters the number of Rademacher sign patterns is $2^{14} = 16{,}384$ using Frisch-Waugh-Lovell partialling. The residualisation step absorbs country and year fixed effects and the log-competitive-allotment control, leaves a residualised outcome and a residualised treatment, and then bootstraps the residuals under the null by multiplying each cluster's residual vector by $s_g \in \{-1, 1\}$. The enumeration floor on the two-sided p-value is $1/2^{G-1} = 1/8{,}192 \approx 0.00012$, well below any conventional cutoff. See Cameron et al. (2008); MacKinnon and Webb (2017) for the underlying theory and Lima (2026) for the exact implementation used here.

B.2 Country-label permutation procedure

The country-label permutation test is implemented as in Young (2019) and MacKinnon and Webb (2020). For each of $K = 50{,}000$ iterations, the grants-intensity vector $(\text{grants}_1, \ldots, \text{grants}_{14})$ is randomly permuted across the fourteen countries in the panel. The interaction $\text{grants}_{\sigma(c)} \times \mathbf{1}\{t \geq \text{2020-07-21}\}$ is re-computed under the permuted assignment, and the baseline two-way fixed-effects specification is re-estimated to obtain a pseudo-coefficient $\hat{\beta}^{(k)}$. The two-sided permutation p-value is $(1/K) \sum_{k=1}^{K} \mathbf{1}\{|\hat{\beta}^{(k)}| \geq |\hat{\beta}^{\text{obs}}|\}$. The empirical 2.5th and 97.5th percentiles of the pseudo-distribution define the permutation 95% interval reported in the body of the paper. The implementation uses the same FWL precomputation trick described in B.1: residualise the outcome and each country-times-post indicator on the rest of the design once, then each draw is a pure-numpy linear combination of the precomputed residual basis.

B.3 Within-EZ restricted permutation

The within-EZ restricted permutation fixes the six non-euro-area countries at zero intensity and permutes only the eight EZ countries' intensities among themselves. The full enumeration over $8! = 40{,}320$ permutations gives a two-sided p-value of 0.089. This test isolates the within-EZ gradient channel from the EZ-versus-placebo contrast and provides the most demanding inference benchmark reported in the paper.

B.4 Time-placebo calibration

The time-placebo test draws 300 fake event dates uniformly at random from trading days between 2 January 2018 and 30 June 2020. For each fake date, the dose_post regressor is re-computed using the same grants-intensity vector but with the fake date in place of 21 July 2020, the baseline specification is re-estimated, and the fake coefficient is stored. The empirical distribution of 300 fake coefficients gives the placebo 95% interval reported in section 7. The false-positive rate at nominal $\alpha = 0.05$ on the placebo distribution is 6.3%, close to nominal, so the specification is well-calibrated against the time-placebo null.

Cluster-robust standard errors. Throughout the paper, asymptotic standard errors are clustered at the country level ($G = 14$) using the linearmodels PanelOLS `clustered` covariance estimator with country as the cluster variable. The small-cluster correction is the standard $G/(G-1)$ Liang-Zeger factor. Asymptotic p-values are computed with $G - 1$ degrees of freedom for the t-statistic. The exact wild cluster bootstrap in B.1 is the preferred small-cluster inference; the asymptotic p-values are reported alongside for transparency only.

Saved permutation distributions. The 50,000-draw country-label permutation distribution and the 300-draw time-placebo distribution are saved as part of the project output.

C Robustness tables

This appendix reports the full set of robustness tables summarised in section 6. Each table reports the indicated alternative specification with the same outcome (log BCR), the same treatment timing (post $\equiv \mathbf{1}\{t \geq 21 \text{ July } 2020\}$), and the same cluster-robust inference at the country level. The canonical headline coefficient ($\hat{\beta} = 1.921$, $N = 1{,}471$, $G = 14$) is reproduced for reference inside each table.

C.1 Sample-composition robustness

Table 13 reports the full set of 29 sample-composition specifications: leave-one-country-out (Panel A, 14 rows), high-exposure exclusions targeting Italy, Spain, and Portugal (Panel B, 7 rows), and core/periphery splits restricting attention to the euro-area treated countries (Panel C, 8 rows).

Table 13: Sample-composition robustness

Specification	$\hat{\beta}$	SE	p_{asymp}	p_{WCR}	p_{perm}	N
Panel A — Leave-one-country-out						
Drop Australia	1.855	0.563	0.001	0.047	0.028	1,352
Drop Austria	1.903	0.500	0.000	0.034	0.021	1,395
Drop Belgium	1.867	0.483	0.000	0.036	0.025	1,414
Drop Canada	2.027	0.495	0.000	0.034	0.023	1,380
Drop Finland	1.932	0.478	0.000	0.035	0.015	1,444
Drop France	1.885	0.448	0.000	0.056	0.045	1,316
Drop Germany	1.882	0.488	0.000	0.045	0.019	1,345
Drop Italy	1.988	0.496	0.000	0.043	0.022	1,375
Drop Japan	1.743	0.521	0.001	0.060	0.038	1,345
Drop New Zealand	1.609	0.494	0.001	0.063	0.043	1,315
Drop Portugal	1.882	0.576	0.001	0.066	0.082	1,418
Drop South Korea	1.945	0.576	0.001	0.045	0.024	1,306
Drop Spain	2.241	0.676	0.001	0.037	0.029	1,366
Drop UK	2.264	0.376	0.000	0.012	0.005	1,352
Panel B — High-exposure exclusions						
Drop Italy	1.988	0.496	0.000	0.043	0.022	1,375
Drop Spain	2.241	0.676	0.001	0.037	0.029	1,366
Drop Portugal	1.882	0.576	0.001	0.066	0.082	1,418
Drop Spain + Portugal	2.596	1.351	0.055	0.093	0.175	1,313
Drop Italy + Spain	2.606	0.840	0.002	0.022	0.034	1,270
Drop Italy + Portugal	1.963	0.652	0.003	0.017	0.092	1,322
Drop Italy + Spain + Portugal	5.524	2.439	0.024	0.178	0.158	1,217
Panel C — Core/periphery splits (euro-area only)						
EZ only (drop 6 placebos)	1.221	0.569	0.032	0.164	0.093	695
EZ only, drop Germany	0.982	0.730	0.179	0.313	0.168	569
EZ only, drop France	1.782	0.298	0.000	0.063	0.116	540
EZ only, drop Italy	1.405	0.552	0.011	0.172	0.026	599
EZ only, drop Spain	1.102	0.715	0.124	0.328	0.254	590
EZ only, drop Portugal	1.193	0.659	0.071	0.313	0.212	642
IT/ES/PT vs AT/BE/FI/FR/DE	1.221	0.569	0.032	0.164	0.093	695
Exclude non-EZ placebos	1.221	0.569	0.032	0.164	0.093	695

Notes: Specification mirrors section 5: country fixed effects, year fixed effects, log competitive allotment, country-clustered standard errors. p_{asymp} is the asymptotic cluster-robust p-value, p_{WCR} is the exact wild cluster bootstrap p-value over 2^G Rademacher sign patterns where G is the number of countries remaining in the sample, and p_{perm} uses 10,000 country-label random reassignments

C.2 Intensity-variant comparison (canonical permutation)

Table 14 reports the full set of 45 intensity-variant specifications: five intensity measures (grants, actual including loans, formula allocation key, binary high-exposure indicator, and the dominance-preserving tiebreak continuous variant), each estimated on five sample configurations (full 14-country panel, the eight euro-area treated countries, two leave-one-out variants of EZ8, and the six largest euro-area economies), and under three control sets (none, supply, supply plus maturity bucket dummies). All rows use the canonical permutation procedure: 50,000 country-label random reassignments, on the post-dedup panel ($N = 1{,}471$ for full-14c specifications). The headline row (Grants, 14c, supply) reproduces the body's $\hat{\beta} = 1.921$ with $p_{\text{perm}} = 0.021$.

Table 14: Intensity-variant comparison (canonical permutation: 50,000 draws, post-dedup panel).

Intensity	Sample	Controls	$\hat{\beta}$	p_{asymp}	p_{perm}	N
Grants only (preferred)						
Grants	14c	none	1.787	0.000	0.032	1,471
Grants	14c	supply	1.921	0.000	0.021	1,471
Grants	14c	supply + mat	1.853	0.000	0.051	1,471
Grants	EZ8	none	1.574	0.002	0.101	695
Grants	EZ8	supply	1.221	0.032	0.096	695
Grants	EZ8	supply + mat	1.263	0.000	0.100	695
Grants	EZ7_noFI	none	1.684	0.001	0.078	668
Grants	EZ7_noAT	none	1.328	0.013	0.183	619
Grants	EZ6	none	1.458	0.006	0.168	592
Actual (grants + loans)						
Actual	14c	none	0.721	0.004	0.122	1,471
Actual	14c	supply	0.794	0.001	0.104	1,471
Actual	14c	supply + mat	0.779	0.001	0.131	1,471
Actual	EZ8	none	0.548	0.057	0.310	695
Actual	EZ8	supply	0.471	0.116	0.262	695
Actual	EZ8	supply + mat	0.531	0.009	0.184	695
Actual	EZ7_noFI	none	0.591	0.045	0.280	668
Actual	EZ7_noAT	none	0.415	0.142	0.474	619
Actual	EZ6	none	0.467	0.111	0.517	592
Formula intensity (RRF allocation key)						
Formula	14c	none	1.141	0.005	0.109	1,471
Formula	14c	supply	1.245	0.002	0.082	1,471
Formula	14c	supply + mat	1.198	0.003	0.127	1,471
Formula	EZ8	none	0.889	0.014	0.239	695
Formula	EZ8	supply	0.804	0.014	0.148	695
Formula	EZ8	supply + mat	0.829	0.000	0.132	695
Formula	EZ7_noFI	none	0.944	0.010	0.189	668
Formula	EZ7_noAT	none	0.692	0.034	0.350	619
Formula	EZ6	none	0.760	0.024	0.334	592
Binary high vs not (IT/ES/PT vs rest)						
Binary	14c	none	0.074	0.006	0.121	1,471
Binary	14c	supply	0.080	0.003	0.124	1,471
Binary	14c	supply + mat	0.079	0.001	0.131	1,471
Binary	EZ8	none	0.056	0.091	0.270	695
Binary	EZ8	supply	0.043	0.183	0.306	695
Binary	EZ8	supply + mat	0.049	0.027	0.216	695
Binary	EZ7_noFI	none	0.060	0.081	0.200	668
Binary	EZ7_noAT	none	0.042	0.191	0.343	619
Binary	EZ6	none	0.047	0.166	0.398	592
Tiebreak (continuous, dominance-preserving)						
Tiebreak	14c	none	0.721	0.004	0.122	1,471
Tiebreak	14c	supply	0.794	0.001	0.104	1,471
Tiebreak	14c	supply + mat	0.779	0.001	0.131	1,471
Tiebreak	EZ8	none	0.548	0.057	0.310	695
Tiebreak	EZ8	supply	0.471	0.116	0.262	695
Tiebreak	EZ8	supply + mat	0.531	0.009	0.184	695
Tiebreak	EZ7_noFI	none	0.591	0.045	0.280	668
Tiebreak	EZ7_noAT	none	0.415	0.142	0.474	619
Tiebreak	EZ6	none	0.467	0.111	0.517	592

Notes: Each row is one fit of the headline specification with the indicated intensity vector as the treatment variable. "14c" = full 14-country panel; "EZ8" = the eight euro-area treated countries only; "EZ7_noFI" = EZ8 with Finland dropped; "EZ7_noAT" = EZ8 with Austria dropped; "EZ6" = the six largest euro-area economies (BE, DE, ES, FR, IT, PT). Control sets: "none" = country and year fixed effects only; "supply" = adds log competitive allotment; "supply + mat" = adds maturity bucket dummies. p_{perm} is the country-label permutation p-value over 50,000 random reassignments with seed 20,260,415, on the post-dedup panel (so the full 14-country fits show $N = 1{,}471$).

C.3 Baseline specifications (alternative FE and intensity coding)

Table 15 reports the eight baseline specifications: the C0 anchor without supply control, the C1 headline, two finer time fixed effects (quarter-year C2 and month-year C3), country-specific linear trends (C4), maturity fixed effects (C5), country×maturity fixed effects (C6), and the raw four-decimal intensity coding (C7).

Table 15: Baseline specifications: log(BCR) on grants intensity × `post`. Post-dedup panel.

	(C0) Headline anchor (no supply)	(C1) Headline (paper)	(C2) Quarter-year FE	(C3) Month-year FE	(C4) Country trends	(C5) + Maturity FE	(C6) Country x Maturity FE	(C7) Raw grants intensity
$\hat{\beta}$	1.787	1.921	1.570	1.568	1.981	1.853	0.555	1.921
(cluster SE)	0.481	0.484	0.450	0.498	1.120	0.487	0.690	0.485
p_{asymp}	0.000	0.000	0.001	0.002	0.077	0.000	0.421	0.000
p_{WCR}	0.033	0.036	0.035	0.039	0.346	0.041	0.469	0.035
p_{perm}	0.030	0.018	0.049	0.064	0.281	0.050	0.664	0.019
FE	Cty + Year	Cty + Year	Cty + QY	Cty + MY	Cty + Year + Cty trends	Cty + Year + Mat	Cty×Mat + Year	Cty + Year
Supply ctrl	none	`log_supply`	`log_supply`	`log_supply`	`log_supply`	`log_supply`	`log_supply`	`log_supply`
Intensity	rounded	rounded	rounded	rounded	rounded	rounded	rounded	raw
N	1471	1471	1471	1471	1471	1471	1471	1471
# countries	14	14	14	14	14	14	14	14
Within R^2	0.024	0.045	0.065	0.091	0.078	0.117	-0.178	0.045

Outcome is log(BCR) for the 1,471-auction post-dedup panel (14 countries; 8 EZ treated + 6 non-EZ placebos). Treatment is grants intensity × `post` where `post = 1{auction_date` $\geq$ `2020-07-21}`; the intensity vector uses 3-decimal rounded values from `GRANTS_14` except in column (C7) which uses raw 4-decimal values. Standard errors clustered on country. p_{WCR} is the exact Rademacher wild-cluster bootstrap p-value (2^{14} =16,384 patterns where the cluster count $G = 14$). p_{perm} is the two-sided country-label permutation p-value across 10,000 draws. Cells reported as *absorbed* indicate that the treatment is collinear with the listed fixed-effect structure under the saturation policy.

C.4 Fixed-effects sensitivity

Table 16 reports eight fixed-effects specifications K1–K8 spanning year, half-year, quarter-year, month-year, country×maturity, country-specific linear trends, and the pre-event detrended outcome. The half-year FE row (K2) is the source for the corresponding row in the body's compact robustness Table 9.

Table 16: Fixed-effects sensitivity: log(BCR) on grants intensity × `post`. Post-dedup panel.

	(K1) Anchor (C1 recap)	(K2) Half-year FE (NEW)	(K3) Quarter-year FE (C2 recap)	(K4) Month-year FE (C3 recap)	(K5) + Maturity FE (C5 recap)	(K6) Country x Maturity FE (C6)	(K7) Country trends (C4 recap)	(K8) Pre-trend de-trended Y (NEW)
$\hat{\beta}$	1.921	1.620	1.570	1.568	1.853	0.555	1.981	1.690
(cluster SE)	0.484	0.428	0.450	0.498	0.487	0.690	1.120	1.243
95% CI	[0.972, 2.869]	[0.780, 2.460]	[0.688, 2.453]	[0.592, 2.545]	[0.899, 2.807]	[-0.797, 1.907]	[-0.214, 4.176]	[-0.747, 4.127]
p_{asymp}	0.000	0.000	0.001	0.002	0.000	0.421	0.077	0.174
p_{WCR}	0.036	0.034	0.035	0.039	0.041	0.469	0.346	0.346
p_{perm}	0.018	0.038	0.049	0.064	0.050	0.664	0.281	0.394
FE	Cty + Year	Cty + HY	Cty + QY	Cty + MY	Cty + Year + Mat	Cty×Mat + Year	Cty + Year + Cty trends	Cty + Year (de-tr. Y)
N	1471	1471	1471	1471	1471	1471	1471	1471
# countries	14	14	14	14	14	14	14	14
Within R^2	0.045	0.050	0.065	0.091	0.117	-0.178	0.078	0.038
C.3 ref. β	1.9206	–	1.5704	1.5685	1.8529	0.5548	1.9810	–

Outcome is log(BCR) for the 1,471-auction post-dedup panel (14 countries; 8 EZ treated + 6 non-EZ placebos). Treatment is grants intensity × `post`; intensity uses 3-decimal `GRANTS_14`. All specifications include `log_supply` and cluster-robust SE at country level (G=14). p_{WCR} is the exact Rademacher wild-cluster bootstrap p-value (2^{14} =16,384 patterns). p_{perm} is the two-sided country-label permutation p-value across 10,000 draws.

C.5 Secondary-market and macro-financial controls

Table 17 adds secondary-market and macro-financial controls (MOVE, VIX, secondary-market yield, lagged BCR) to the C1 headline specification. Most columns restrict to the five euro-area countries with consistent secondary-market data (BE, FR, DE, IT, PT); F5 fits the MOVE control on the full 14-country panel. The MOVE-control row is the source for the corresponding row in the body's compact robustness Table 9.

Table 17: Secondary-market and macro-financial controls. log(BCR) on grants intensity $\times$ `post`.

	(F0) Anchor on 5-country subsample	(F1) + secondary yield level	(F2) + yield volatility	(F3) + yield + volatility	(F4) + spread to Bund	(F5) Full 14c panel + MOVE	(F6) All controls (5c)	(F7) All controls (5c, dup of F6)
$\hat{\beta}$	0.985	0.545	0.705	0.410	0.775	1.796	-0.478	-0.478
(cluster SE)	0.844	1.007	0.854	0.986	0.612	0.466	1.015	1.015
p_{asymp}	0.244	0.588	0.410	0.678	0.207	0.000	0.638	0.638
p_{WCR}	0.438	0.625	0.438	0.750	0.438	0.039	0.688	0.688
p_{perm}	0.308	0.558	0.400	0.667	0.775	0.020	0.675	0.675
Controls	supply	supply, y	supply, vol	supply, y, vol	supply, spr	supply, MOVE	supply, y, vol, cbf, cbs, ngeu, MOVE	supply, y, vol, cbf, cbs, ngeu, MOVE
N	487	487	487	487	305	1471	487	487
# countries	5	5	5	5	5	14	5	5
Within R^2	0.141	0.154	0.146	0.156	0.087	0.050	0.196	0.196

Outcome is log(BCR). Treatment is `GRANTS_14` $\times$ `post` with `post = 1{date` $\geq$ `2020-07-21}`. All specifications include country FE, year FE, and `log_supply`. Columns F0–F4 and F6–F7 are restricted to the five euro-area countries with secondary-market controls observed (BE, FR, DE, IT, PT; $G = 5$). Column F5 fits on the full 14-country panel ($G = 14$). p_{WCR}: exact Rademacher wild-cluster bootstrap p-value (2^G enumeration). p_{perm}: country-label permutation p-value. For $G = 5$ all $5! = 120$ unique permutations are enumerated (deterministic). For $G = 14$ (F5), 10,000 random draws are used.

C.6 Outcome transformations

Table 18 re-runs the C1 anchor under seven alternative transformations of the BCR outcome: log, levels, country-standardised log, country-demeaned log, winsorised log, the inverse hyperbolic sine, and the residualised log (FWL-projected onto entity FE plus year FE plus log-supply). The sign and significance of the headline coefficient are preserved across transformations.

Table 18: Outcome transformations of the C1 anchor (grants×`post` on the post-dedup panel).

	(I1) log_bcr	(I2) bid_to_cover_ratio (level)	(I3) std_log_bcr_country	(I4) demeaned_log_bcr_country	(I5) winsorized_log_bcr_1_99	(I6) ihs_bcr	(I7) residualized_log_bcr
$\hat{\beta}$	1.921	4.371	8.396	1.921	1.861	2.761	1.921
(cluster SE)	0.484	1.256	2.591	0.484	0.505	0.759	0.484
95% CI	[0.97, 2.87]	[1.91, 6.83]	[3.32, 13.47]	[0.97, 2.87]	[0.87, 2.85]	[1.27, 4.25]	[0.97, 2.87]
p_{asymp}	0.000	0.001	0.001	0.000	0.000	0.000	0.000
p_{WCR}	0.036	0.037	0.054	0.036	0.036	0.036	0.036
p_{perm}	0.018	0.035	0.058	0.018	0.023	0.027	0.018
Implied @ Italy 0.038	7.57%	0.166	0.319	7.57%	7.33%	10.51%	7.57%
(units)	pct_BCR	BCR_points	country_SD_units	pct_BCR	pct_BCR	pct_BCR_approx	pct_BCR
N	1471	1471	1471	1471	1471	1471	1471
# countries	14	14	14	14	14	14	14
Within R^2	0.045	0.037	0.053	0.045	0.046	0.044	0.007

Outcome transformations of the C1 anchor (country FE + year FE + `log_supply` control + grants × `post`). I1 reproduces C1 on `log_bcr`. I2 uses BCR in levels. I3 standardizes `log_bcr` by within-country mean and SD. I4 demeans `log_bcr` by country mean. I5 winsorizes `log_bcr` at pooled 1st and 99th percentiles. I6 uses the inverse hyperbolic sine transformation, $2 \cdot \text{arcsinh}(\text{BCR}/2)$ (Bellemare & Wichman 2020). I7 residualizes `log_bcr` on `log_supply`+entity FE+year FE; by FWL, the coefficient must equal I1's. p_{WCR} is the exact Rademacher wild-cluster bootstrap (2^{14} =16,384 patterns; $G = 14$). p_{perm} is the country-label permutation across 10,000 draws with seed 20260415. Implied effect at Italy uses grants intensity = 0.038. Units *pct_BCR* = $\exp(\beta \cdot 0.038) - 1$; *BCR_points* = $\beta \cdot 0.038$; *country_SD_units* = $\beta \cdot 0.038$; *pct_BCR_approx* = $\sinh(\beta \cdot 0.038)$ (IHS column).

C.7 Weighting and aggregation

Table 19 reports six weighting and aggregation variants: equal observation weight (J1), equal-country weights (J2), competitive-allotment weights (J3), and three aggregation granularities (country×year J4, country× half-year J5, country×quarter-year J6). The equal-country weights row (J2) is the source for the corresponding row in the body's compact robustness Table 9.

Table 19: Weighting and aggregation sensitivity. Outcome: $\log(\text{BCR})$ (auction-level) or $\overline{\log(\text{BCR})}$ per cell (aggregated). All specs: country FE + year FE; cluster SE at country (G=14).

Panel A: Auction-level weighted regressions			
	(J1) Unweighted (anchor)	(J2) Equal-country weights	(J3) Allotment-weighted
$\hat{\beta}$	1.921	1.813	1.106
(cluster SE)	0.484	0.446	0.602
95% CI	[0.972, 2.869]	[0.939, 2.687]	[-0.075, 2.286]
p_{asymp}	0.000	0.000	0.067
p_{WCR}	0.036	0.049	0.188
p_{perm}	0.018	0.018	0.086
Within R^2	0.045	0.082	0.048
N	1471	1471	1471
# countries	14	14	14

Panel B: Aggregated to coarser time cells			
	(J4) Country-quarter	(J5) Country-month	(J6) Country-year
$\hat{\beta}$	1.557	1.582	2.074
(cluster SE)	0.602	0.552	0.878
95% CI	[0.376, 2.738]	[0.500, 2.663]	[0.353, 3.795]
p_{asymp}	0.010	0.004	0.021
p_{WCR}	0.143	0.110	0.192
p_{perm}	0.109	0.127	0.158
Within R^2	0.063	0.059	0.057
N	289	730	84
# countries	14	14	14

Panel A re-fits the C1 spec (country FE + year FE + `log_supply` + grants × `post`) on the 1,471-auction post-dedup panel under three weighting schemes: (J1) unweighted, (J2) equal-country weights $w_i = (N/G)/n_c \approx 105.07/n_c$, (J3) competitive-allotment weights $w_i = \text{competitive\_allotment}_i$. Panel B collapses the auction panel to (country, cell) cells of three granularities; per cell, $\overline{\log(\text{BCR})}$ is the unweighted mean over auctions, log TotalAllotment replaces `log_supply`, and `post` is the fraction of cell days on/after 2020-07-21 (cells straddling the cutoff get `post` $\in (0,1)$). p_{WCR}: exact Rademacher wild-cluster bootstrap (2^{14} =16,384 patterns). p_{perm}: country-label permutation, 10,000 draws.

C.8 Maturity and tenor splits

Table 20 reports five maturity-split specifications L0–L4: the C1 anchor (L0), 10-year-only (L1), 30-year-only (L2), with maturity-bucket dummy (L3), and with the three-way interaction grants $\times$ post $\times$ $\mathbf{1}\{\text{maturity} = 30\}$ (L4). The 10-year and 30-year subsamples reported in the body's Table 7 are the L1 and L2 rows of this table.

Table 20: Maturity-split robustness for the headline auction-demand specification: $\log(\text{BCR})$ on grants intensity $\times$ `post`

	(L0) Anchor	(L1) 10-year only	(L2) 30-year only	(L3) Pooled + maturity FE	(L4_main) 3-way interaction: beta at 10y	(L4_diff) 3-way interaction: 30y differential (delta)	(L4_30y) 3-way interaction: beta at 30y
Treatment term	`grants_post`	`grants_post`	`grants_post`	`grants_post`	`grants_post`	`grants_post_x_mb30`	`grants_post + grants_post_x_mb30`
Sample	all 1,471	maturity == 10	maturity == 30	all 1,471	all 1,471	all 1,471	all 1,471
$\hat{\beta}$	1.921	1.589	1.675	1.853	2.366	-1.962	0.403
(cluster SE)	0.484	0.461	1.149	0.487	0.538	1.505	1.331
95% CI	[0.97, 2.87]	[0.69, 2.49]	[-0.58, 3.93]	[0.90, 2.81]	[1.31, 3.42]	[-4.91, 0.99]	[-2.20, 3.01]
p_{asymp}	0.000	0.001	0.146	0.000	0.000	0.193	0.762
p_{WCR}	0.036	0.074	0.161	0.041	0.032	0.562	–
p_{perm}	0.018	0.087	0.290	0.050	0.002	0.600	–
N	1471	1021	450	1471	1471	1471	1471
# countries	14	14	14	14	14	14	14
Within R^2	0.045	0.110	0.148	0.117	0.047	0.047	0.047

Outcome is $\log(\text{BCR})$ on the 1,471-auction post-dedup panel (14 countries; 8 EZ treated + 6 non-EZ placebos). All specs include country FE, year FE, `log_supply`; cluster-robust SE at country level. Treatment is $\text{grants}_c \times \text{post}_t$ from `GRANTS_14`. L0 reproduces the C1 anchor; L1 / L2 restrict to maturity $= 10$ / $= 30$; L3 adds the maturity-bucket dummy `mb_10` (reproduces C5 / K5); L4 adds the three-way interaction grants $\times$ post $\times$ $\mathbf{1}\{\text{maturity} = 30\}$. For L4 the table reports the 10y effect (`L4_main`), the 30y differential δ (`L4_diff`), and the derived 30y effect $\beta_{10y} + \delta$ (`L4_30y`, delta-method SE). p_{WCR} is the exact Rademacher wild-cluster bootstrap (16,384 patterns when $G = 14$); p_{perm} is the country-label permutation across 10,000 draws. F-test of $\delta = 0$ (L4): $F = 1.699$, $\text{df} = (1, 1449)$, $p_F = 0.193$.

C.9 Falsification and placebo grid

Table 21 reports the auxiliary-outcome falsification grid (H1: log primary-dealer count, log competitive amount, log secondary yield) together with the placebo-treatment grid (H3: random uniform placebo intensities, time-placebo dates, and the loans-only horse race). The H1 rows reproduce the body's Table 11; the H3 rows are new to the appendix.

Table 21: Falsification outcomes (H1) and placebo treatments (H3). Post-dedup panel; country and year FE; `log_supply` control (except H1c). Cluster-robust SE at country level.

Spec	Description	$\hat{\beta}$	95% CI	p_{asymp}	p_{WCR}	p_{perm}	N	# c
Panel A: H1 falsification outcomes (real grants × post; varies outcome)								
H1a	log_bcr	1.921	$[0.97, 2.87]$	0.000	0.036	0.018	1471	14
H1b	log_primary_dealers	0.109	$[-0.96, 1.18]$	0.842	0.801	0.870	814	9
H1c	log_competitive_allotment	2.136	$[-2.51, 6.78]$	0.368	0.408	0.496	1471	14
H1d	log_total_bids	1.921	$[0.97, 2.87]$	0.000	0.036	0.018	1471	14
Panel B: H3 placebo treatments (outcome = `log_bcr`)								
H3a	loans × post	1.060	$[0.30, 1.82]$	0.006	0.164	0.257	1471	14
H3b (β_G)	grants horse race	2.690	$[1.03, 4.35]$	0.002	0.003	0.002	1471	14
H3b (β_L)	loans horse race	-0.844	$[-1.98, 0.29]$	0.146	0.073	0.271	1471	14
H3c	random non-EZ (mean β)	-0.782	$[-2.35, 0.90]$	–	–	0.063	–	6
H3d	fake post 2019-07-21	0.897	$[-0.21, 2.00]$	0.113	0.243	0.377	1471	14

p_{WCR}: exact Rademacher wild-cluster bootstrap ($2^{14} = 16{,}384$ patterns for $G = 14$; $2^9 = 512$ for H1b at $G = 9$). p_{perm}: country-label permutation, 10,000 draws. 95% CI $= \hat{\beta} \pm 1.96 \times \text{SE}_{\text{cluster}}$. H1c drops `log_supply` because the outcome equals the control by construction. H3c reports the mean coefficient across 1,000 draws of random uniform$[0, 0.05]$ intensity on the six non-EZ countries (EZ countries set to 0); p_{perm} there is the fraction of fake draws with $|\hat{\beta}_{\text{fake}}| \geq |\hat{\beta}_{\text{C1}}| = 1.92$. Joint F-test of $\beta_G = \beta_L = 0$ (H3b): $F = 7.686$, df $= (2, 1449)$, $p = 0.000$. Test $\beta_G = \beta_L$: $F = 6.642$, df $= (1, 1449)$, $p = 0.010$.

C.10 Disbursement specifications

Table 22 reports the six disbursement specifications E1–E6: cumulative grants disbursed (E1), total disbursed (E2), loans disbursed (E3), the C1 approved-grants anchor (E4), the horse race (E5), and the event-window amount specification (E6). The E4 and E5 specifications are reported in the body's Table 12.

Table 22: Actual RRF disbursement specifications. Post-dedup panel; country + year FE; log_supply control.

	(E1) Cumulative grants disbursed / GDP	(E2) Cumulative total disbursed / GDP	(E3) Cumulative loans disbursed / GDP	(E4) Approved grants x post (C1 anchor)	(E5a) Horse race: approved grants x post	(E5b) Horse race: cum_grants_disb	(E6) Event window amount ($\pm$30 days)
$\hat{\beta}$	7.604	5.423	9.276	1.921	1.423	4.441	5.444
(cluster SE)	4.023	1.695	2.108	0.484	0.675	4.757	3.587
p_{asymp}	0.059	0.001	0.000	0.000	0.035	0.351	0.129
p_{WCR}	0.096	0.091	0.005	0.036	0.221	0.353	0.200
p_{perm}	0.169	0.154	0.247	0.018	0.116	0.435	0.302
N	1471	1471	1471	1471	1471	1471	1471
Within R^2	0.043	0.044	0.042	0.045	0.046	0.046	0.038

E5 joint F-test of $\beta_1 = \beta_2 = 0$: $F = 6.286$, df $= (2, 1449)$, $p_F = 0.002$.

Outcome is log(BCR) on the 1,471-auction post-dedup panel. All specs include country FE, year FE, `log_supply`; cluster-robust SE at country level (G=14). Treatments: E1 cumulative grants disbursed/GDP, E2 cumulative total disbursed/GDP, E3 cumulative loans disbursed/GDP, E4 the C1 anchor (approved grants $\times$ post, GRANTS_14), E5a/E5b the horse-race specification with both treatments together, E6 an event-window amount ($\pm$30 days). Cumulative disbursement amounts are from EC press releases). p_{WCR} is the exact Rademacher wild-cluster bootstrap (2^{14} =16,384 patterns). p_{perm} is the country-label permutation across 10,000 draws.

C.11 Event-study binning schemes

Table 23 reports the binned event study under three binning schemes (G_curr matching the body's pre-trend table; G_ann with annual bins; G_6m with six-month bins). The G_curr scheme is the one shown in the body's Table 6 and Figure 3.

Table 23: Binned event study: log(BCR) on grants intensity $\times$ bin dummy. Post-dedup panel (1,471 obs, 14 countries). Country FE + Year FE + log_supply. Cluster: country.

Scheme	Bin (period)	$\hat{\beta}$	SE	95% CI	p_{asymp}	p_{WCR}	p_{perm}	N_{bin}	#cty	joint pre F (p)
G_curr	ref_pre_12m (1-12 months before)	0.000	–	–	–	–	–	–	–	1.53 (0.218)
	pre_2plus (>= 24 months before)	1.260	0.724	[-0.16, 2.68]	0.082	0.215	0.354	130	14	
	pre_2y (12-24 months before)	0.948	0.975	[-0.96, 2.86]	0.331	0.380	0.362	242	14	
	post_h1 (0-6 months after)	2.663	0.882	[0.93, 4.39]	0.003	0.260	0.167	152	14	
	post_1y (6-18 months after)	1.461	1.012	[-0.52, 3.45]	0.149	0.143	0.370	281	14	
	post_2y (18-30 months after)	3.487	1.122	[1.29, 5.69]	0.002	0.039	0.030	324	14	
	post_3yplus (>= 30 months after)	2.927	1.683	[-0.37, 6.23]	0.082	0.072	0.138	72	14	
G_ann	ref_pre_12m (0-12 months before agreement (Aug 2019 - Jul 2020))	0.000	–	–	–	–	–	–	–	1.54 (0.214)
	pre_2 (>= 24 months before)	1.280	0.730	[-0.15, 2.71]	0.080	0.218	0.349	130	14	
	pre_1 (12-24 months before)	0.962	0.969	[-0.94, 2.86]	0.321	0.366	0.353	242	14	
	post_1 (0-12 months after)	2.381	0.773	[0.87, 3.90]	0.002	0.219	0.079	299	14	
	post_2 (12-24 months after)	2.251	0.978	[0.33, 4.17]	0.022	0.034	0.018	297	14	
	post_3plus (>= 24 months after)	3.330	1.176	[1.03, 5.64]	0.005	0.035	0.005	233	14	
G_6m	ref_pre_0_6m (0-6 months before)	0.000	–	–	–	–	–	–	–	1.46 (0.211)
	pre_30plus (>= 30 months before)	+nan	nan	[nan, nan]	–	–	–	0	0	
	pre_24_30 (24-30 months before)	1.740	0.897	[-0.02, 3.50]	0.053	0.265	0.230	130	14	
	pre_18_24 (18-24 months before)	2.099	0.912	[0.31, 3.89]	0.021	0.218	0.077	120	14	
	pre_12_18 (12-18 months before)	0.571	1.494	[-2.36, 3.50]	0.702	0.742	0.730	122	14	
	pre_6_12 (6-12 months before)	0.538	0.806	[-1.04, 2.12]	0.505	0.563	0.685	125	14	
	post_0_6 (0-6 months after)	2.918	0.919	[1.12, 4.72]	0.002	0.298	0.136	152	14	
	post_6_12 (6-12 months after)	2.197	0.829	[0.57, 3.82]	0.008	0.218	0.054	147	14	
	post_12_18 (12-18 months after)	1.365	0.952	[-0.50, 3.23]	0.152	0.104	0.525	134	14	
	post_18_24 (18-24 months after)	3.680	1.340	[1.05, 6.31]	0.006	0.073	0.001	163	14	
	post_24_30 (24-30 months after)	3.951	1.176	[1.65, 6.26]	0.001	0.046	0.000	161	14	
	post_30plus (>= 30 months after)	3.255	1.589	[0.14, 6.37]	0.041	0.084	0.107	72	14	

Outcome: log(BCR). Treatment: grants intensity (GRANTS_14 rounded to 3 d.p.) interacted with bin dummies, omitting the reference bin. Three bin schemes: `G_curr` matches the current paper's `tab:pretrend`; `G_ann` uses annual bins; `G_6m` uses six-month bins. Errors clustered on country (G=14). 95% CI = $\hat{\beta} \pm 1.96\times$SE. p_{WCR}: exact wild-cluster Rademacher enumeration over $2^{14} = 16{,}384$ patterns (null imposed via FWL). p_{perm}: country-label permutation over 10,000 draws (FWL precomputation). Joint pre-trend F-test: Wald test that all pre-event bin coefficients equal zero, asymptotic cluster-robust covariance; with G=14 the finite-sample concerns.

C.12 Inference-suite summary

Table 24 consolidates the four honest-inference procedures applied to the C1 headline: the exact wild cluster bootstrap (M0, $2^{14} = 16{,}384$ Rademacher enumeration), the canonical country-label permutation (M1, 50,000 draws, the within-EZ restricted permutation (M2, full $8! = 40{,}320$ enumeration), and the binary high-grants subset test (M3, full $\binom{14}{3} = 364$ enumeration). The first two procedures reject at the 5 percent level, while the within-EZ and binary high-grants restrictions provide more demanding benchmarks under which the evidence is weaker.

Table 24: WCR + three permutation layers for the C1 headline. All four methods test the grants-intensity $\times$ `post` coefficient.

Method	$\hat{\beta}$	$p_{\text{two-sided}}$	MC SE(p)	$\text{CI}_{2.5\%}$	$\text{CI}_{97.5\%}$	$n_{\text{assignments}}$
M0: WCR exact enumeration (2^{14})	1.9206	0.0355	0.0000	–	–	16,384
M1: Country-label permutation (50K draws)	1.9206	0.0206	0.0006	-1.0872	1.8891	50,000
M2: Within-EZ permutation (full 8!)	1.9206	0.0885	0.0000	0.5031	2.1475	40,320
M3: Binary high-vs-not (full $\binom{14}{3}$)	0.0796	0.1209	0.0000	-0.0693	0.0980	364

Outcome is log(BCR) on the 1,471-auction post-dedup panel (14 countries; 8 EZ + 6 non-EZ placebos). Headline spec C1 = country FE + year FE + `log_supply` + grants $\times$ `post`. M0 enumerates all $2^{14} = 16{,}384$ Rademacher sign patterns; finite floor $= 1/16{,}384$. M1 uses seed 20,260,415 and 50,000 random country-label shuffles. M2 enumerates all $8! = 40{,}320$ permutations of intensities across the 8 EZ countries, holding non-EZ placebos at zero. M3 enumerates all $\binom{14}{3} = 364$ subsets of *high-grants* labels; observed subset $= \{$Italy, Spain, Portugal$\}$. MC SE $= 0$ where indicated indicates full enumeration (no Monte Carlo error).

D Additional figures

Figure 4: Full leave-one-country-out coefficient plot

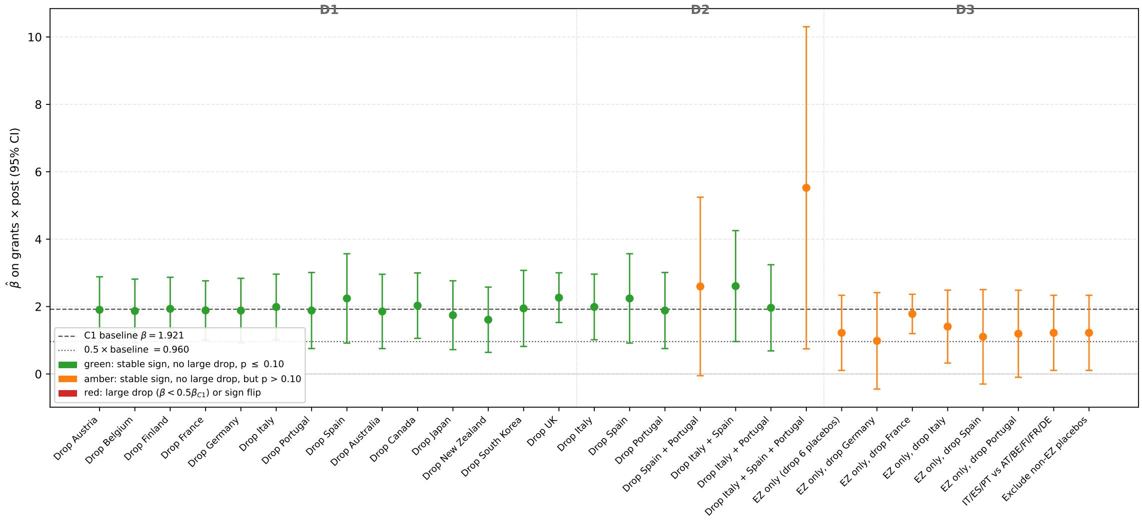

Notes: Each marker is the headline grants-only coefficient estimated on the 13-country panel that drops the indicated country. Bars show cluster-robust 95% confidence intervals. The dashed horizontal line marks the 14-country baseline $\hat{\beta} = 1.921$.

Figure 5: Country-label permutation distribution of the headline coefficient

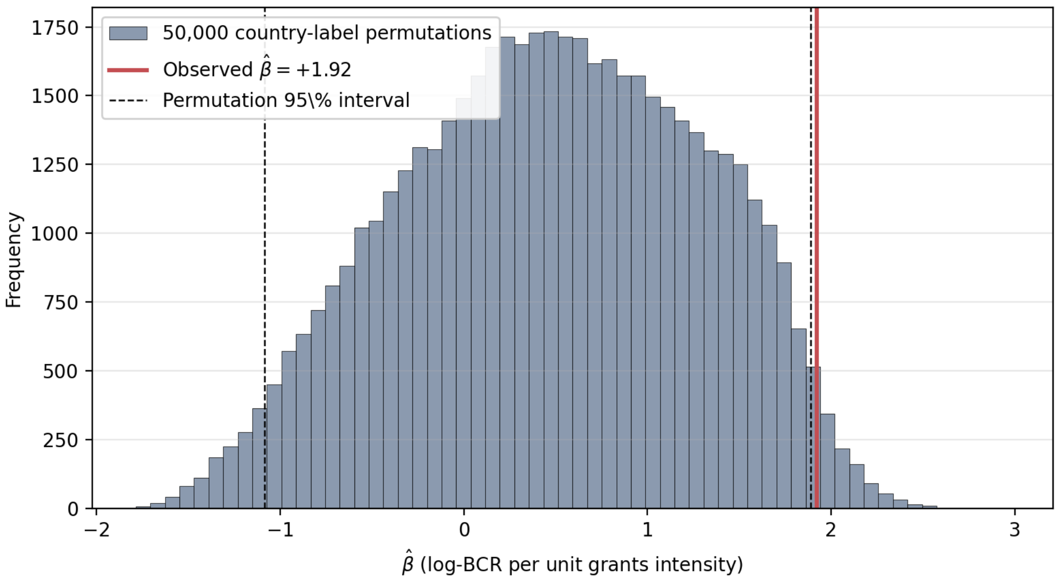

Notes: Grey histogram is over 50,000 random reshuffles of the grants-intensity vector across the fourteen countries; the red line marks the observed coefficient under the true assignment. The observed coefficient sits to the right of the empirical 97.5th percentile of the permutation distribution.

Figure 6: Time-placebo distribution of the headline coefficient

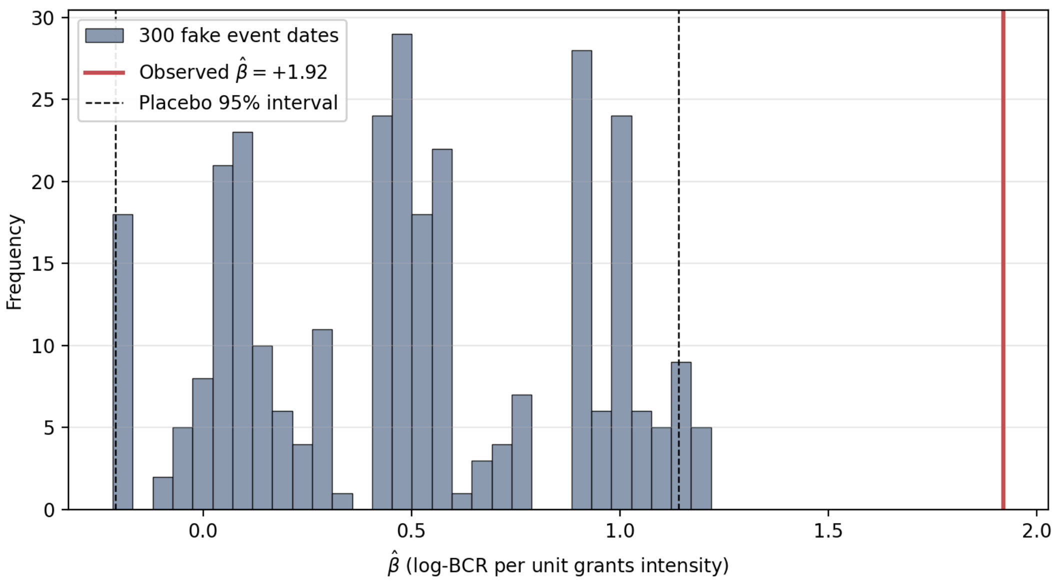

Notes: Grey histogram is over 300 random fake event dates drawn uniformly from business days between June 2018 and December 2019; the red vertical line marks the observed coefficient under the true 21 July 2020 event date. The observed coefficient sits outside the empirical 95% interval of the placebo distribution.

Figure 7: Cumulative RRF grants disbursed (scaled by 2019 GDP) by country

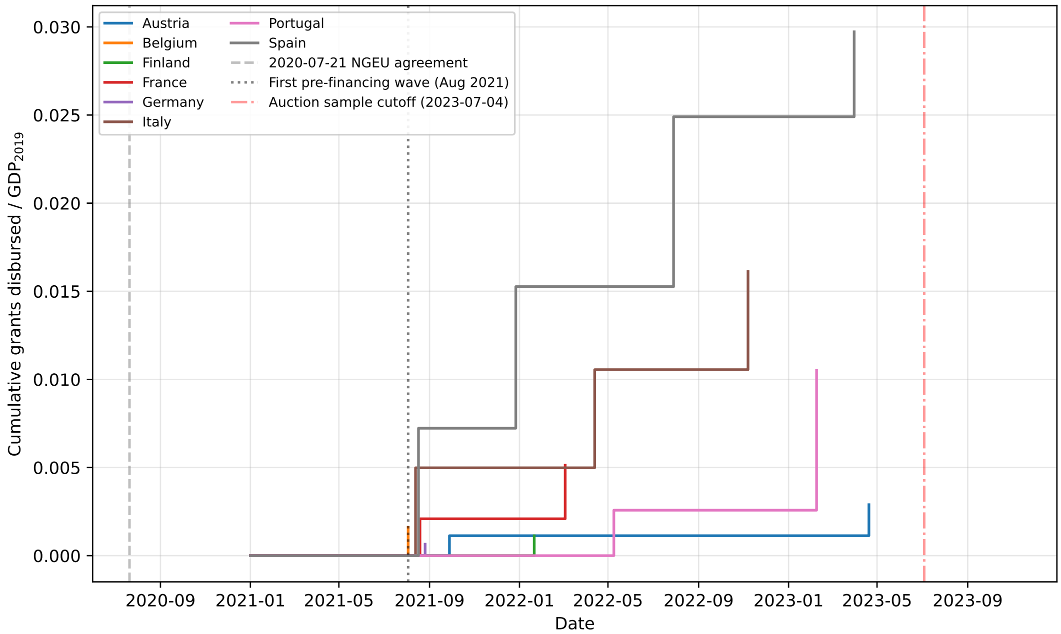

Notes: used in the realized-cash horse race of section 8.2. Source: European Commission RRF scoreboard.

Figure 8: Event-study coefficients by event-time bin across the three intensity coding schemes (G_curr, G_ann, G_6m)

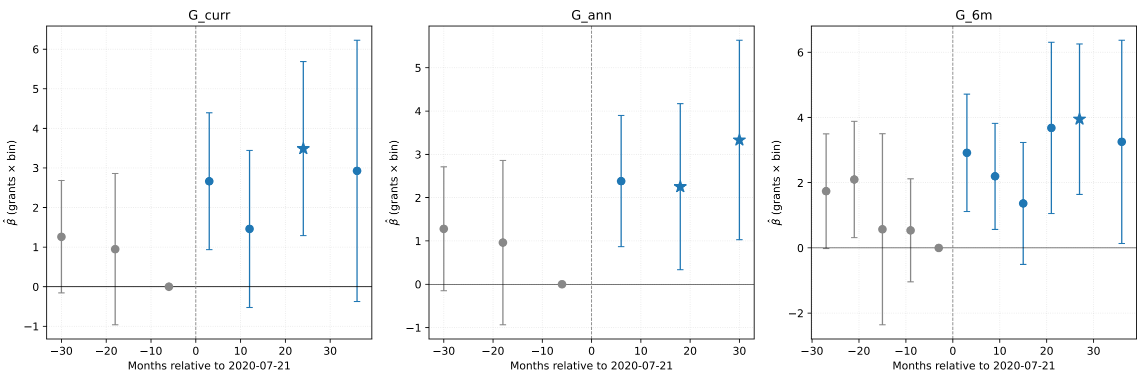

Notes: Vertical bars show $\pm 1.96 \cdot$ SE asymptotic intervals.