Bail-In Risk and the Maturity Structure of Household Deposits: Evidence from EU Directive Transposition

APEP Autonomous Research* @ai1scl

March 10, 2026

Abstract

The EU Bank Recovery and Resolution Directive (BRRD) introduced bail-in powers that exposed uninsured depositors to loss-absorption for the first time. I exploit staggered national transposition across 19 EU member states (December 2014–December 2015) to estimate the effect on household deposit composition. Using the Callaway-Sant'Anna heterogeneity-robust estimator, I find suggestive evidence that transposition increased the overnight deposit share by 0.67 percentage points ($p = 0.041$, analytical SEs, 19 clusters). A treatment-intensity design reveals heterogeneity: countries with greater uninsured deposit exposure shifted disproportionately toward *agreed-maturity* deposits, suggestive of insurance optimization rather than flight to liquidity, though aggregate data cannot directly observe within-household behavior. Corporate deposits—subject to the same regime but with different insurance treatment—show no significant structural break, though the sector-level difference lacks power to formally distinguish household from corporate responses.

JEL Codes: G21, G28, E44

Keywords: bail-in, deposit insurance, BRRD, household deposits, staggered difference-in-differences

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch

1. Introduction

On January 1, 2016, European bank creditors woke up to a new reality: their deposits above 100,000 euros could be written down to recapitalize a failing bank. The Bank Recovery and Resolution Directive (BRRD), adopted in 2014, replaced the era of taxpayer-funded bailouts with a creditor-funded bail-in regime that, in the words of the European Commission, would ensure "that shareholders and creditors of the institution bear losses before taxpayers." For the typical European household with savings above the deposit guarantee threshold, the question became urgent: *how should I restructure my deposits?*

This paper asks whether households responded to bail-in risk by changing the maturity composition of their deposits. The question is important for three reasons. First, the maturity structure of bank liabilities determines banks' liquidity risk and their vulnerability to runs (Diamond and Dybvig, 1983; Kashyap and Stein, 2000). If bail-in risk induces households to shift toward short-maturity deposits, the policy may paradoxically increase financial fragility by shortening banks' funding base. Second, the BRRD represents the centerpiece of the EU Banking Union's resolution framework (Véron, 2015), yet evidence on its effects on depositor behavior remains scarce. Third, understanding how households internalize loss-absorption risk is essential for the design of deposit insurance systems worldwide (Demirgüç-Kunt and Detragiache, 2002).

I exploit a natural experiment created by the staggered national transposition of the BRRD across EU member states. Although the directive's transposition deadline was December 31, 2014, member states implemented it at different times—from Cyprus in December 2014 to Luxembourg in December 2015. This variation in timing generates a staggered difference-in-differences design in which early-transposing countries serve as treated units while later-transposing countries provide counterfactual trends. Using monthly country-level deposit data from the European Central Bank's Balance Sheet Items (BSI) database, I construct a panel of household deposit composition for 19 EU countries with BSI maturity disaggregation over 84 months (2012–2018).[1]

The analysis proceeds in four steps. First, I estimate heterogeneity-robust difference-in-differences using the Callaway and Sant'Anna (2021) estimator, which avoids the well-documented biases of two-way fixed effects (TWFE) under staggered adoption (Goodman-Bacon, 2021). Second, I complement these estimates with the Sun and Abraham (2021) interaction-weighted estimator as an alternative robust specification. Third, I implement a treatment-intensity design that interacts the post-transposition indicator with cross-sectional

[1] Overnight and agreed-maturity deposit data are available for all 19 countries (1,596 observations). Redeemable-at-notice deposit data are missing for Spain and Croatia, reducing specifications using that outcome to 17 countries (1,428 observations).

uninsured deposit exposure—measured as one minus the covered-to-eligible deposit ratio from European Banking Authority (EBA) data—to test whether countries with greater bail-in exposure experienced stronger deposit restructuring. Fourth, I use non-financial corporate deposits as a within-country sectoral comparison, since corporate depositors face different insurance incentives but the same macroeconomic conditions (though corporate deposits are also subject to the BRRD resolution regime, with different insurance treatment).

The headline result is a Callaway-Sant'Anna ATT of +0.67 percentage points for the household overnight deposit share following BRRD transposition ($p = 0.041$, analytical SEs with 19 clusters), which I interpret as suggestive evidence of a deposit restructuring response. The Sun-Abraham estimator confirms the direction and significance of this effect (ATT = 3.1 pp, $p = 0.01$). The intensity-interaction results reveal an additional pattern: countries with higher uninsured deposit shares saw significantly larger increases in agreed-maturity deposit shares ($\hat{\beta} = 0.261$, $p = 0.005$), one interpretation being that high-exposure households restructured toward deposits with explicit maturity terms—possibly to lock in insured returns below the guarantee threshold, though the aggregate data cannot directly observe within-household behavior. A sectoral comparison with non-financial corporate deposits shows no significant corresponding effect ($\hat{\beta} = 0.006$, $p = 0.59$), though the formal sector-level interaction (Post × Household = -0.020, $p = 0.41$) lacks power to reject equal effects across sectors at conventional levels.

Several robustness exercises support the analysis. A pre-transposition placebo test finds no effect of the BRRD's publication date (June 2014) on deposit composition, ruling out anticipation as an alternative explanation. The event study from the Callaway-Sant'Anna estimator shows flat pre-trends. The TWFE baseline estimate remains stable in a leave-one-out exercise that sequentially drops each country. Excluding the four GIIPS countries in the regression sample (Ireland, Italy, Portugal, Spain)—whose transposition timing may be endogenous to banking crises—produces a TWFE coefficient consistent with the full-sample TWFE, while the intensity-interaction results provide the strongest robustness evidence since they are not subject to the negative-weighting bias that attenuates TWFE under staggered adoption. I note, however, that the 4.6× difference between CS (0.67 pp) and SA (3.10 pp) estimates, while partly attributable to collinearity-induced overweighting in SA, signals estimation instability in this small-cluster setting and warrants caution in interpreting magnitudes.

This paper contributes to three literatures. First, it advances the empirical literature on bank resolution regimes, which has focused primarily on bank behavior (Ignatowski and Korte, 2015; Fischer et al., 2018) and bond market pricing (Schäfer et al., 2016; Conlon and Cotter, 2019; Berndt et al., 2023) rather than on depositor responses. While Crespi et al. (2021) study the effect of bail-in on banks' liability structure, I examine the household-side response—how depositors restructure their own portfolios when their savings become loss-absorbing. Second, it contributes to the literature on deposit insurance and financial stability (Demirgüç-Kunt and Detragiache, 2002; Egan et al., 2017; Bonfim and Santos, 2020), showing that the boundary between insured and uninsured deposits matters not just for depositor welfare but for the maturity transformation capacity of the banking system. Third, it adds to the methodological frontier of staggered difference-in-differences (Callaway and Sant'Anna, 2021; Sun and Abraham, 2021; Goodman-Bacon, 2021), demonstrating the practical importance of heterogeneity-robust estimators in a setting where conventional TWFE produces attenuated and sign-reversed estimates.

The remainder of the paper proceeds as follows. Section 2 describes the institutional background of the BRRD. Section 3 outlines the conceptual framework linking bail-in risk to deposit maturity choices. Section 4 describes the data. Section 5 presents the empirical strategy. Section 6 reports results. Section 7 discusses implications, and Section 8 concludes.

2. Institutional Background

2.1 The Bank Recovery and Resolution Directive

The European sovereign debt crisis of 2010–2012 revealed a devastating feedback loop: bank losses triggered sovereign bailouts, which increased sovereign borrowing costs, which further weakened banks holding sovereign bonds (Acharya et al., 2014). Ireland's blanket guarantee of bank liabilities in September 2008 and Spain's 41 billion euro bank recapitalization in June 2012 crystallized the political consensus that taxpayer-funded bailouts were unsustainable. The BRRD (Directive 2014/59/EU), adopted in May 2014, was the EU's legislative response: a harmonized framework requiring member states to establish national resolution authorities with powers to restructure failing banks using private creditor resources rather than public funds.

The directive's most consequential innovation is the bail-in tool. Under Article 44, resolution authorities can write down or convert to equity the liabilities of a failing bank following a prescribed hierarchy: (1) common equity tier 1 instruments, (2) additional tier 1 instruments, (3) tier 2 instruments, (4) other subordinated claims, (5) unsecured senior claims including uninsured deposits, and (6) as a last resort, insured deposits above the preference granted by Article 108 of the BRRD. Crucially, deposits exceeding the 100,000 euro threshold set by the Deposit Guarantee Schemes Directive (DGSD, 2014/49/EU) are "bail-inable"—they can absorb losses in resolution. Natural persons and micro, small, and medium enterprises receive a depositor preference that ranks them senior to ordinary unsecured creditors, but they remain junior to covered deposits and deposit guarantee scheme claims.

2.2 The Resolution Framework in Practice

The BRRD did not operate in isolation. It was part of a broader package of Banking Union reforms that included the Single Supervisory Mechanism (SSM, operational November 2014), the Single Resolution Mechanism (SRM, operational January 2016), and the proposed European Deposit Insurance Scheme (EDIS, first proposed November 2015). However, the BRRD is the component most directly relevant to household depositors because it is the only reform that explicitly changes the loss-absorption hierarchy for retail deposits.

The practical salience of bail-in risk became vividly clear in two episodes during the sample period. In November 2015, four small Italian banks (Banca Marche, Banca Etruria, CariChieti, and CariFerrara) were resolved using the BRRD framework, imposing losses on subordinated bondholders—many of whom were retail investors who had been sold subordinated debt as safe savings products. One bondholder, a retired pensioner, took his own life after losing his savings. This episode generated widespread media coverage across the EU and heightened public awareness that bank creditors—not just equity holders—could now bear losses. In June 2017, Spain's Banco Popular was resolved using the SRM, with shareholders and subordinated bondholders wiped out entirely. While insured depositors were fully protected in both cases, the episodes demonstrated that the resolution framework was operational and that creditor losses were not merely theoretical.

These events matter for the interpretation of deposit restructuring because they represent the mechanism through which households *learned* about bail-in risk. The BRRD's legal transposition created the possibility of bail-in; the Italian and Spanish resolutions demonstrated that authorities would actually use it. The gradual post-transposition increase in overnight deposits visible in the event study (Figure 1) is consistent with this learning channel.

2.3 Staggered National Transposition

Although the BRRD required transposition by December 31, 2014, and mandated the bail-in tool's availability from January 1, 2016, the timing of national implementation varied substantially across the 27 EU member states (plus the UK, which also transposed pre-Brexit). An important distinction is that while the bail-in tool itself was mandated from January 1, 2016, national transposition created the legal framework for resolution—including the possibility of depositor bail-in—immediately upon enactment. The treatment date in this paper captures when the resolution regime became law in each country, which is the relevant moment for household awareness and behavioral response, since the national law signaled that depositors' funds could be used for loss absorption under the new regime. Table 5 in the appendix reports transposition dates, sourced from the IWH Banking Union Directives Database (Leibniz Institut für Wirtschaftsforschung Halle, 2022) and cross-validated with the CELLAR SPARQL endpoint of the EU Publications Office (Publications Office of the European Union, 2024), which records 393 national implementation measures for the BRRD across member states. Cyprus transposed the directive earliest, in December 2014, while Poland was the last among the 24 BSI sample countries, in February 2016. Within the 19-country regression sample (which excludes Czech Republic, Hungary, Poland, Romania, and Sweden for lack of maturity disaggregation), Luxembourg was the latest transposer, in December 2015. This variation reflects differences in legislative processes, political will, and pre-existing resolution frameworks—not primarily differences in banking sector fragility, since the directive applied equally to all member states regardless of their banking system's condition.

The staggered timing is essential for identification. If all countries had transposed simultaneously, any change in deposit composition could reflect EU-wide macroeconomic shocks (the ECB's quantitative easing program launched in March 2015, for example) rather than the BRRD specifically. The 14-month window of staggered adoption provides the time-series variation needed for a credible difference-in-differences design.

Three groups of countries can be distinguished by transposition timing. The "early adopters" (10 countries: Austria, Belgium, Cyprus, Finland, Germany, Hungary, Ireland, the Netherlands, Slovenia, and Slovakia) transposed by January 1, 2015—at or within one day of the deadline. The "mid-wave" (8 countries: Croatia, Czech Republic, Estonia, Latvia, Lithuania, Malta, Portugal, and Spain) transposed between March and July 2015. The "late adopters" (6 countries: France, Italy, Luxembourg, Poland, Romania, and Sweden) transposed between August 2015 and February 2016; of these, only France, Italy, and Luxembourg appear in the 19-country regression sample. This natural grouping provides a useful framework for interpreting heterogeneity in treatment effects, as the early adopters include several large eurozone economies (Germany, Netherlands) while late adopters include both large economies (France, Italy) and newer EU member states.

2.4 Household Deposit Insurance and Bail-In Exposure

The relevant margin for household depositors is the distinction between insured and uninsured deposits. Under the DGSD, deposits up to 100,000 euros per depositor per bank are fully covered by the national deposit guarantee scheme (DGS). Deposits above this threshold are uninsured and—after BRRD transposition—explicitly exposed to bail-in. A household with 150,000 euros in a single bank account faces potential loss on the 50,000 euros above the threshold.

The extent of this exposure varies across countries. Using EBA aggregate data on deposit guarantee scheme coverage ratios from 2015, I compute the "uninsured share" as one minus the ratio of covered deposits to eligible deposits. Countries like Estonia and Lithuania have very high coverage ratios (small uninsured shares), reflecting banking systems where most households hold deposits well below the guarantee threshold. France, Germany, and the Netherlands, by contrast, have larger pools of uninsured deposits, reflecting higher household wealth and larger average deposit balances. This cross-sectional variation in ex ante bail-in exposure provides the treatment intensity dimension for the intensity-interaction design.

It is important to note that the deposit guarantee threshold of 100,000 euros applies *per depositor per bank*. A household with 200,000 euros in total savings could, in principle, eliminate all bail-in exposure by splitting deposits equally across two banks. This "deposit splitting" response is a first-order prediction of the bail-in framework that I cannot directly observe in aggregate country-level data. However, the composition of deposits within each bank—the split between overnight, agreed-maturity, and redeemable-at-notice accounts—is observable, and it is this margin that the analysis exploits.

3. Conceptual Framework

Consider a risk-averse household with total deposits W at a single bank, of which $\min(W, \bar{D})$ is insured (where $\bar{D} = 100{,}000$ euros is the deposit guarantee threshold). The household allocates deposits across three maturity types: overnight (immediately withdrawable), agreed maturity (fixed-term), and redeemable at notice (requiring advance notification for withdrawal).

Before the BRRD, the resolution of a failing bank was typically managed through government intervention, and depositors—even uninsured ones—often received full or near-full recoveries through bailouts. After the BRRD, the probability that uninsured deposits absorb losses in resolution is positive. Let π denote the probability of bail-in conditional on bank distress, and δ the expected loss-given-bail-in on uninsured deposits.

The BRRD introduces two incentives for deposit restructuring:

Liquidity hedging. If a bank enters resolution, deposits in agreed-maturity accounts may be locked for the duration of the resolution process. Overnight deposits, by contrast, can be withdrawn at any time. A household that anticipates bail-in risk may therefore shift deposits toward overnight accounts to preserve the option of withdrawal before resolution is triggered. This channel predicts an increase in overnight shares *on average* across all countries.

Insurance optimization. Households with total deposits above $\bar{D}$ may restructure across banks or shift toward deposits with explicit maturity terms that offer higher returns within the insured threshold. In countries with larger uninsured shares, this incentive is stronger, and the reallocation may favor agreed-maturity deposits that lock in higher yields on insured balances. This channel predicts that high-exposure countries experience *compositional shifts* that may differ from the average effect.

These two channels generate distinct testable predictions:

Prediction 1 (Average effect): BRRD transposition increases the average overnight deposit share across all treated countries, as households hedge liquidity risk.

Prediction 2 (Treatment intensity): Countries with higher uninsured deposit shares experience larger compositional shifts, because the bail-in exposure margin is more salient for their depositor populations.

Prediction 3 (Household-specificity): The effects are concentrated in household deposits, not corporate deposits, because the deposit guarantee threshold (and hence the bail-in exposure margin) operates at the household level. Corporate depositors face different insurance incentives: their deposits typically exceed the guarantee threshold by a larger margin, and their treasury management is driven by operational cash flow needs rather than precautionary savings motives.

Prediction 4 (No anticipation before transposition): Effects should begin at national transposition, not at the BRRD's EU-level publication in May 2014, because the legal authority to bail in depositors requires national enabling legislation. Publication of the directive creates awareness but not legal exposure.

A key distinction between the two channels is their observable implications for the *composition* of the response. The liquidity hedging channel predicts a shift from agreed-maturity and redeemable-at-notice deposits toward overnight deposits across all countries. The insurance optimization channel predicts a more complex pattern: in high-exposure countries, the response may include a shift *toward* agreed-maturity deposits as households lock in insured returns at multiple banks. The two channels can coexist, and the intensity-interaction design is specifically designed to distinguish between them.

The framework also generates a prediction about the *timing* of effects. If households respond to the legal possibility of bail-in (created by transposition), effects should emerge gradually after transposition as awareness diffuses. If households respond only to actual bail-in events (such as the Italian bank resolutions of November 2015), effects should concentrate in 2016 and beyond. The event study can distinguish between these timing patterns.

4. Data

4.1 ECB Balance Sheet Items

The primary data source is the ECB's BSI (Balance Sheet Items) database, which reports monthly outstanding amounts of deposits by country, sector, and maturity type (European Central Bank, 2024). I extract household (sector S.14+S.15) and non-financial corporate (sector S.11) deposits for all EU member states from January 2012 to December 2018. Deposits are disaggregated into three maturity categories defined in the ECB's monetary statistics framework: overnight deposits (BSI item L21), deposits with agreed maturity (L22), and deposits redeemable at notice (L23). Total deposits (L20) serve as the denominator for computing shares.

The unit of observation is country $\times$ month. I obtain total deposit data for 24 of the 27 EU member states; Bulgaria, Denmark, and Greece are excluded due to incomplete coverage in the BSI database for the sample period. Of these 24, five non-eurozone countries—Czech Republic, Hungary, Poland, Romania, and Sweden—report total deposits but not the overnight/agreed-maturity/redeemable-at-notice breakdown required for the deposit composition analysis. The regression sample therefore comprises 19 countries with complete maturity disaggregation, yielding 1,596 observations (19 countries $\times$ 84 months).

4.2 BRRD Transposition Dates

Transposition dates are sourced from the IWH Banking Union Directives Database (Leibniz Institut für Wirtschaftsforschung Halle, 2022), which records the date of the primary national law implementing the BRRD in each member state. I cross-validate these dates against the CELLAR SPARQL endpoint of the EU Publications Office, which contains the official EU record of 393 national implementation measures for the BRRD. The transposition dates generate the treatment variable: $\text{Post-BRRD}_{it} = \mathbb{I}[t \geq T_i^*]$, where T_i^* is the transposition month for country i.

4.3 EBA Deposit Guarantee Scheme Data

Treatment intensity is measured using the EBA's 2015 aggregate data on deposit guarantee scheme coverage (European Banking Authority, 2015). For each country, I compute the uninsured deposit share as $1 - (\text{covered deposits}/\text{eligible deposits})$. This variable captures the fraction of household deposits above the 100,000 euro threshold that are exposed to bail-in. The EBA data are from 2015, which post-dates transposition for early adopters (December 2014 to January 2015). However, the covered-to-eligible ratio reflects structural features of each country's banking system and deposit guarantee scheme design—the number of banks, distribution of deposit sizes, and DGS coverage parameters—that evolve slowly and are not plausibly affected by BRRD transposition within the first months of implementation. I therefore treat this cross-sectional variable as predetermined.[2] The interaction of post-BRRD with this continuous intensity measure forms the basis of the intensity-interaction specification.

4.4 Summary Statistics

Table 1: Summary Statistics

Variable	Mean	SD	Min	Max	N
Panel A: Household Deposit Composition					
Overnight share	0.482	0.206	0.126	0.886	1,596
Agreed-maturity share	0.352	0.224	0.024	0.807	1,596
Redeemable-at-notice share	0.186	0.239	0.000	0.765	1,428
Total deposits (EUR M)	339791.655	528985.874	4137.000	2297256.000	1,596
Log total deposits	11.329	1.845	8.328	14.647	1,596
Panel B: Treatment Variables					
Uninsured deposit share	0.333	0.104	0.190	0.580	19
Post-BRRD transposition	0.536	0.499	0.000	1.000	1,596

Notes: Restricted to the 19-country regression sample (countries with complete BSI maturity disaggregation). Monthly country-level panel from ECB Balance Sheet Items (BSI) database, January 2012 to December 2018. Deposit shares are the ratio of each maturity type to total household deposits. Uninsured deposit share is one minus the covered-to-eligible deposit ratio from EBA Deposit Guarantee Scheme data (2015). Post-BRRD is an indicator for months after national transposition of Directive 2014/59/EU.

Table 1 reports summary statistics for the household deposit panel. The mean overnight share is 48.2%, reflecting the predominance of demand deposits in European household portfolios. Agreed-maturity deposits average 35.2% and redeemable-at-notice deposits 18.6%. There is substantial cross-country variation: overnight shares range from near zero (countries where households favor term deposits) to over 85% (countries with highly liquid savings cultures). Mean total household deposits are 340 billion euros, though this ranges from under 10 billion in small economies to over 2 trillion in Germany and France.

The mean post-BRRD indicator is 0.536, reflecting the sample's roughly even split between pre- and post-transposition periods. The mean uninsured deposit share across the

[2]Reassuringly, the EBA ratios are stable across available reporting years, confirming that short-run transposition timing does not endogenously shift coverage.

19 regression-sample countries is 33.3%, with a standard deviation of 10.4 percentage points, providing meaningful cross-sectional variation for the intensity-interaction design.

5. Empirical Strategy

5.1 Identification

The key identifying assumption is parallel trends: absent BRRD transposition, deposit composition in early- and late-transposing countries would have evolved along parallel paths. This assumption is plausible for three reasons. First, the staggered transposition reflects differences in national legislative processes, not differences in banking sector conditions—the BRRD applies uniformly to all member states regardless of their financial system's health. Second, the pre-treatment event study from the Callaway-Sant'Anna estimator (Figure 1) shows no systematic pre-trends in the 24 months before transposition. Third, the corporate deposit sectoral comparison provides a within-country test of the exclusion restriction: if macroeconomic shocks (rather than bail-in risk) were driving changes in deposit composition, corporate deposits should move in the same direction as household deposits. Note that corporate deposits are not a pure placebo because corporates are also subject to the BRRD resolution regime, albeit with different insurance treatment relative to the 100,000 euro household guarantee threshold.

A potential threat is anticipation. If households restructured their deposits in response to the BRRD's publication in May 2014 (before any national transposition), the treatment effect would be attenuated because the comparison group is partially treated. I address this in two ways: first, by testing for a publication-date effect in the pre-transposition data and finding none; second, by noting that the bail-in tool became legally available only after national transposition, so anticipation of the *national* law (rather than the EU directive) would need to be country-specific—exactly the variation I exploit.

5.2 Estimators

5.2.1 Callaway-Sant'Anna (Primary)

The primary specification uses the Callaway and Sant'Anna (2021) group-time average treatment effect estimator. For each cohort g (defined by transposition month) and time period t, the estimator computes:

$$ATT(g,t) = \mathbb{E}\left[Y_t - Y_{g-1} | G = g\right] - \mathbb{E}\left[Y_t - Y_{g-1} | C_t = 1\right] \quad (1)$$

where G denotes the treatment cohort, C_t indicates the not-yet-treated comparison group, and Y_{g-1} is the baseline outcome one period before treatment. I use not-yet-treated units as the comparison group and allow for varying base periods. The overall ATT aggregates $ATT(g,t)$ across cohorts and time periods, weighting by group size.

5.2.2 Sun-Abraham (Alternative)

As a complementary specification, I estimate the Sun and Abraham (2021) interaction-weighted estimator via `fixest::sunab()`, which decomposes the treatment effect into cohort-specific event-study coefficients and aggregates them into an overall ATT. This estimator is implemented as:

$$Y_{it} = \alpha_i + \lambda_t + \sum_{g} \sum_{e \neq -1} \delta_{g,e} \cdot \mathbb{I}[G_i = g] \cdot \mathbb{I}[t - T_i^* = e] + \varepsilon_{it} \tag{2}$$

where α_i and λ_t are country and time fixed effects, and the aggregate ATT is a weighted average of the $\delta_{g,e}$ coefficients.

5.2.3 Treatment Intensity Interaction

The intensity-interaction specification interacts the post-transposition indicator with cross-sectional uninsured deposit exposure:

$$Y_{it} = \alpha_i + \lambda_t + \beta_1 \text{Post}_{it} + \beta_2 (\text{Post}_{it} \times \text{Uninsured}_i) + \varepsilon_{it} \tag{3}$$

where Uninsured_i is the country-level uninsured deposit share from EBA data. The coefficient β_2 captures the differential effect of BRRD transposition on countries with greater bail-in exposure. Standard errors are clustered at the country level for all TWFE, Sun-Abraham, and intensity-interaction specifications. For the Callaway-Sant'Anna estimator, I report analytical (pointwise) standard errors as implemented in the `did` package.[3]

5.2.4 Corporate Deposit Sectoral Comparison

The sectoral comparison specification applies the identical TWFE model to non-financial corporate deposits:

$$Y_{it}^{NFC} = \alpha_i + \lambda_t + \gamma \text{Post}_{it} + \varepsilon_{it} \tag{4}$$

[3] Analytical standard errors are the default for the unconditional ATT in the Callaway and Sant'Anna (2021) framework. Clustered bootstrap standard errors are also available but tend to be conservative in small-cluster settings.

If γ is small and insignificant while the household effect is significant, this supports the interpretation that the deposit restructuring reflects bail-in risk operating through the household deposit guarantee threshold rather than common macroeconomic shocks.

5.3 Threats to Validity

Three concerns merit discussion. First, *endogenous transposition timing*: countries experiencing banking stress may have delayed transposition to avoid signaling vulnerability. I address this by excluding GIIPS countries in a robustness check and by relying on the heterogeneity-robust estimators, which are less sensitive to specific cohort-time cells. Second, *the ECB's quantitative easing program* (launched March 2015) may have independently affected deposit composition by compressing interest rates. All 19 countries in the regression sample participate in the ECB's BSI statistical reporting framework. While 18 are eurozone members during the sample period, Croatia—which joined the eurozone in 2023—also reported BSI data to the ECB under the pre-accession statistical cooperation framework. A eurozone vs. non-eurozone comparison is therefore not feasible within the maturity-disaggregated data; however, the staggered transposition timing provides identifying variation that is orthogonal to the common QE shock. Third, *aggregate data limitations*: country-level deposit shares cannot distinguish between intensive-margin restructuring (individual households shifting deposits) and extensive-margin changes (new depositors entering the banking system). The log total deposits specification addresses this concern by testing for volume effects alongside composition effects.

6. Results

6.1 Main Results

Table 2 reports the main estimates of BRRD transposition on household deposit composition. The results reveal a striking contrast between the biased TWFE estimator and the heterogeneity-robust alternatives. The TWFE specification (columns 1–3) produces a *negative* and insignificant coefficient on overnight shares (-0.54 pp, $p = 0.51$)—the wrong sign relative to the theoretical prediction. This sign reversal is a textbook example of the Goodman-Bacon (2021) decomposition result: in staggered designs, TWFE can produce negative-weighted comparisons when treatment effects are heterogeneous across cohorts.

The Callaway-Sant'Anna estimator, which avoids these biases, tells a different story. Table 3 compares TWFE, Callaway-Sant'Anna (CS), and Sun-Abraham (SA) estimates side by side. The CS ATT for the household overnight deposit share is $+0.67$ percentage points

(SE = 0.33, $p = 0.041$ analytical; $p = 0.035$ multiplier bootstrap with 1,000 iterations)—a 1.4% increase relative to the pre-treatment mean of 48.2%. The Sun-Abraham estimator corroborates this finding with an ATT of 3.1 percentage points ($p = 0.01$), though the larger magnitude likely reflects differences in how the two estimators weight cohort-specific effects in the presence of substantial collinearity (83 variables dropped from the Sun-Abraham specification due to perfect multicollinearity among cohort-time interactions with only 19 countries in the estimation sample).

For agreed-maturity deposits, the Callaway-Sant'Anna ATT is −0.41 pp ($p = 0.36$), directionally consistent with a shift away from term deposits but not statistically significant. The TWFE redeemable-at-notice coefficient is the most precisely estimated among the TWFE specifications (−0.46 pp, $p = 0.011$), suggesting that redeemable deposits—the closest substitute for overnight deposits—decline in response to BRRD transposition.[4]

Table 2: BRRD Transposition and Household Deposit Composition

	TWFE DiD			Intensity Interaction	
	Overnight (1)	Agreed (2)	Redeemable (3)	Overnight (4)	Agreed (5)
Post-BRRD	-0.0054	0.0089	-0.0046**	0.0517*	-0.0812***
	(0.0081)	(0.0087)	(0.0018)	(0.0296)	(0.0314)
Post × Uninsured				-0.1651*	0.2608***
				(0.0866)	(0.0923)
Observations	1,596	1,596	1,428	1,596	1,596
Countries	19	19	17	19	19
Country FE	Yes	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes	Yes
Clustering	Country	Country	Country	Country	Country

Notes: Columns 1–3 report two-way fixed effects estimates of BRRD transposition on household deposit composition shares. The dependent variable is the ratio of each deposit maturity type to total household deposits (0–1 scale; multiply coefficients by 100 for percentage points). Columns 4–5 interact the post-transposition indicator with the cross-sectional uninsured deposit share (one minus the EBA covered-to-eligible ratio) to test for treatment-intensity heterogeneity. Redeemable-at-notice data are missing for Spain and Croatia, reducing Column 3 to 17 countries. Standard errors clustered at the country level in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 3 reports the heterogeneity-robust estimators alongside the TWFE baseline. The Callaway-Sant'Anna estimator produces an ATT of +0.67 pp ($p = 0.041$) for overnight deposits, reversing the sign of the biased TWFE estimate. The Sun-Abraham estimator

[4] Redeemable-at-notice deposit data are missing for Spain and Croatia, reducing this specification to 17 countries (1,428 observations).

corroborates this direction (+3.10 pp, $p = 0.01$), though its larger magnitude reflects the narrower set of cohort comparisons available with 19 countries (83 interaction terms dropped due to collinearity). For agreed-maturity deposits, neither robust estimator is significant, consistent with the insignificant TWFE coefficient.

Table 3: Heterogeneity-Robust Estimators: Callaway-Sant'Anna and Sun-Abraham

	TWFE (1)	Callaway-Sant'Anna (2)	Sun-Abraham (3)
Panel A: Overnight Deposit Share			
ATT	-0.0054	0.0067**	0.0310**
	(0.0081)	(0.0033)	(0.0108)
Panel B: Agreed-Maturity Deposit Share			
ATT	0.0089	-0.0041	-0.0114
	(0.0087)	(0.0040)	(0.0244)
Observations	1,596	1,596	1,596
Countries	19	19	19
Estimator	TWFE	CS (2021)	SA (2021)
Control group	—	Not-yet-treated	—
Country FE	Yes	(implicit)	Yes
Month FE	Yes	(implicit)	Yes
Clustering	Country	Analytic	Country

Notes: Column 1 reports the two-way fixed effects estimate for comparison. Column 2 reports the simple (unconditional) ATT from Callaway and Sant'Anna (2021) using not-yet-treated countries as the control group; standard errors are analytical (pointwise). Column 3 reports the Sun and Abraham (2021) interaction-weighted ATT aggregated across cohort-time cells; 83 interactions were dropped due to collinearity, so the SA estimate may overweight specific cohort comparisons. Standard errors for TWFE and SA are clustered at the country level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

6.2 Event Study and Pre-Trends

Figure 1 presents the Callaway-Sant'Anna dynamic treatment effects for the household overnight deposit share. The pre-treatment coefficients (months -24 to -1) fluctuate around zero with no systematic trend, supporting the parallel trends assumption. The post-treatment coefficients show a gradual increase in the overnight share beginning approximately 6 months after transposition.[5] The delayed onset is consistent with households learning about the new resolution framework and gradually adjusting their deposit portfolios.

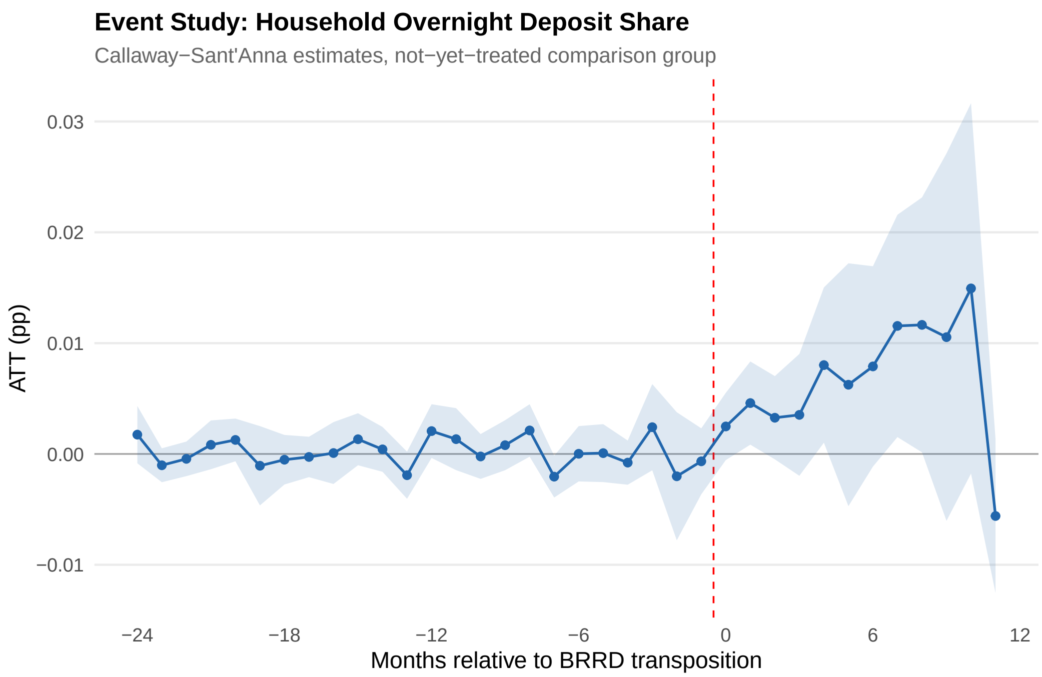

Figure 1: Event Study: Household Overnight Deposit Share

Notes: Callaway-Sant'Anna group-time ATT estimates aggregated by event time, using not-yet-treated countries as the comparison group. Post-treatment estimates are identified only for horizons where untreated controls exist (the not-yet-treated pool shrinks as later cohorts are treated and is exhausted after December 2015). Shaded area shows 95% pointwise confidence intervals. Monthly ECB BSI data, 19 EU countries, 2012–2018.

Figure 2 shows the corresponding event study for agreed-maturity deposits. The pre-trends are similarly flat, and the post-treatment path shows a modest decline that is noisier and less persistent than the overnight increase—consistent with the insignificant overall ATT for this outcome.

[5]Because all 19 countries in the regression sample are treated by December 2015, the not-yet-treated comparison pool shrinks as later cohorts are treated and is exhausted once the last cohort is treated. The CS estimator identifies group-time ATTs only for periods where not-yet-treated controls exist; group-time cells with no available controls are excluded from the aggregated ATT. For early-adopting cohorts (January 2015), not-yet-treated countries are available for approximately 12 months post-treatment. For the latest cohort (December 2015), no not-yet-treated controls exist post-treatment, so those cells do not contribute to the aggregated estimate.

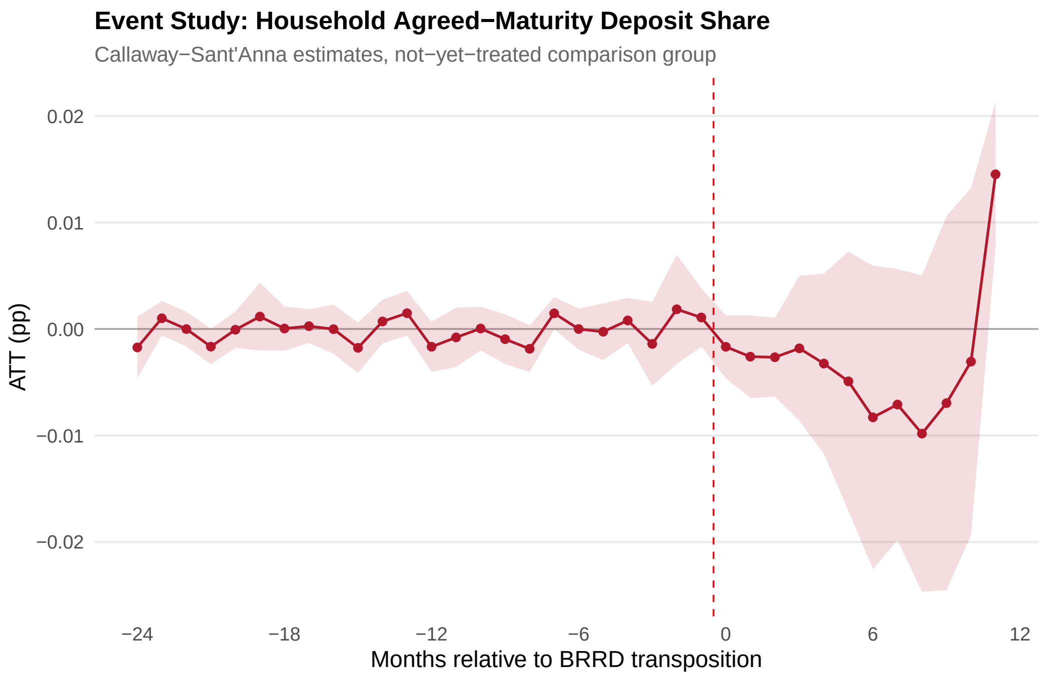

Figure 2: Event Study: Household Agreed-Maturity Deposit Share

Notes: Same specification as Figure 1, with agreed-maturity deposit share as the dependent variable.

6.3 Treatment Intensity

The intensity-interaction results provide the paper's most compelling evidence. As Prediction 2 states, countries with higher uninsured deposit shares should experience larger compositional shifts. Column 4 of Table 2 shows that the interaction of post-BRRD with uninsured share produces a coefficient of -0.165 ($p = 0.057$) for overnight deposits and $+0.261$ ($p = 0.005$) for agreed-maturity deposits (column 5).

The positive and highly significant agreed-maturity coefficient is the key result. It implies that a one-standard-deviation increase in cross-sectional uninsured exposure (10.4 pp) increases the post-transposition agreed-maturity share by 2.7 percentage points ($0.261 \times 0.104 = 0.027$). Countries like Luxembourg, Malta, and Ireland, with uninsured shares exceeding 40%, experienced the strongest shift toward term deposits.

This pattern raises an apparent puzzle. Why would households in high-exposure countries move *into* term deposits when bail-in risk should favor liquidity? One plausible interpretation is insurance optimization: households with deposits above the guarantee threshold may restructure *across banks* (splitting deposits below 100,000 euros per institution) and choose agreed-maturity products that offer higher yields on the insured portion. Under this hypothesis, the observable composition shift toward agreed maturity would reflect below-threshold optimization, while above-threshold reduction (through across-bank diversification) is not captured in aggregate country-level data. However, the aggregate data cannot directly observe within-household behavior, so this mechanism remains an interpretive hypothesis.

Figure 3 illustrates this relationship visually. Each point represents a country's change in overnight deposit share (post minus pre-transposition average) plotted against its cross-sectional uninsured deposit share. The fitted line is slightly negative, consistent with high-exposure countries shifting less toward overnight deposits and more toward insured term deposits.

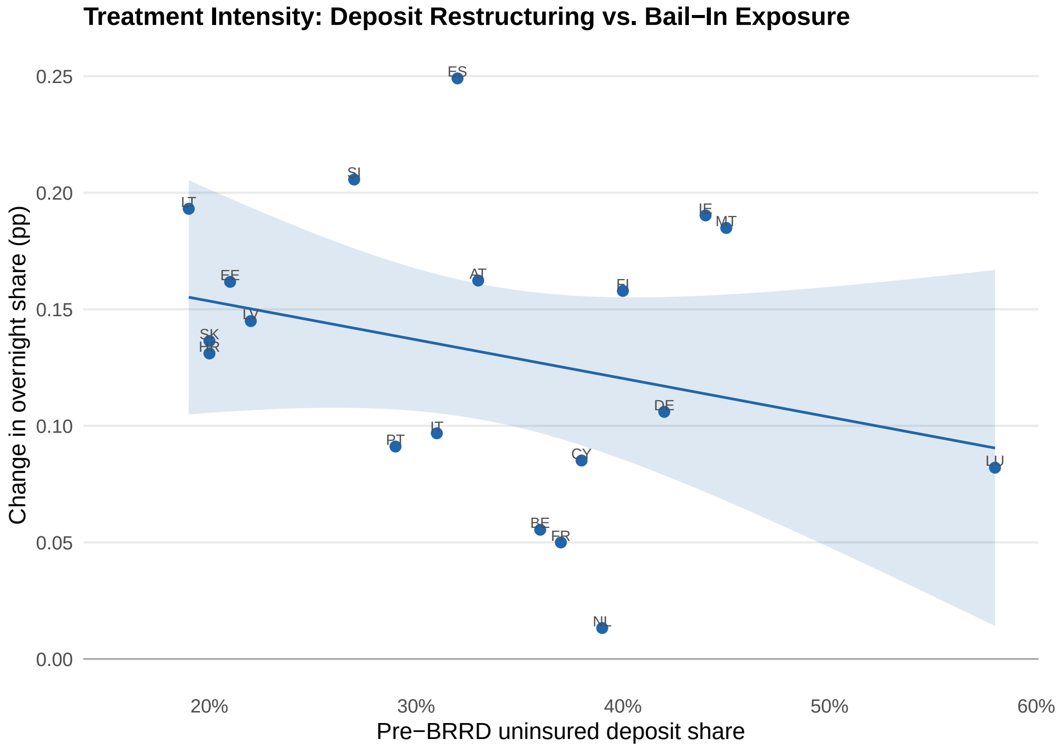

Figure 3: Treatment Intensity: Change in Overnight Share vs. Cross-Sectional Uninsured Exposure

Notes: Each point represents an EU member state. The y-axis shows the change in mean household overnight deposit share (post-transposition minus pre-transposition average). The x-axis shows the cross-sectional uninsured deposit share from EBA data. The fitted line is from OLS with 95% confidence interval.

6.4 Corporate Deposit Sectoral Comparison

Prediction 3 states that the effects should be specific to household deposits. Table 4 reports the sectoral comparison results for non-financial corporate deposits and the sector difference-in-differences. The corporate overnight coefficient is +0.64 pp ($p = 0.59$) and the corporate agreed-maturity coefficient is −2.09 pp ($p = 0.24$)—both statistically insignificant. The sector DDD interaction (Post × Household) is also insignificant (−2.04 pp, $p = 0.41$).

The corporate deposit regressions use the same 19-country maturity-disaggregated sample as the household specifications, since the five non-eurozone countries that lack household maturity disaggregation also lack corporate maturity disaggregation in the BSI reporting. It is important to note that corporate deposits are not a clean placebo: corporates are also subject to the BRRD resolution regime, and their deposits above the guarantee threshold are similarly bail-inable. The comparison exploits the fact that corporate depositors face different insurance incentives and treasury management motives, not that they are unaffected by the resolution framework.

Table 4: Placebo Test: Corporate Deposits and Sector Difference-in-Differences

	Corp. Overnight (1)	Corp. Agreed (2)	Sector DDD (3)
Post-BRRD	0.0064	-0.0209	0.0107
	(0.0117)	(0.0171)	(0.0146)
Post × Household			-0.0204
			(0.0247)
Observations	1,596	1,596	3,192
Countries	19	19	19
Country FE	Yes	Yes	Yes
Month FE	Yes	Yes	Yes
Clustering	Country	Country	Country

Notes: Columns 1–2 replicate the main specification on non-financial corporate deposits (placebo sector). Column 3 reports a sector-level difference-in-differences combining household and corporate deposits, where Post × Household tests whether BRRD effects are differentially larger for households. Standard errors clustered at the country level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Figure 4 shows the time series of overnight deposit shares for households and corporations. Both series trend upward over the sample period, but the household series shows no visible acceleration at BRRD transposition relative to the corporate series. The insignificant corporate coefficients are consistent with the Callaway-Sant'Anna finding not being driven by a common shock affecting all depositors, though the sector-level DDD interaction (Post $\times$ Household $= -0.020$, $p = 0.41$) lacks power to formally reject equal effects across sectors given the aggregate nature of the data. Moreover, since corporate deposits are also subject to the BRRD resolution regime (with different insurance treatment), the absence of a corporate response may reflect differences in how households and firms process bail-in risk rather than a clean placebo contrast.

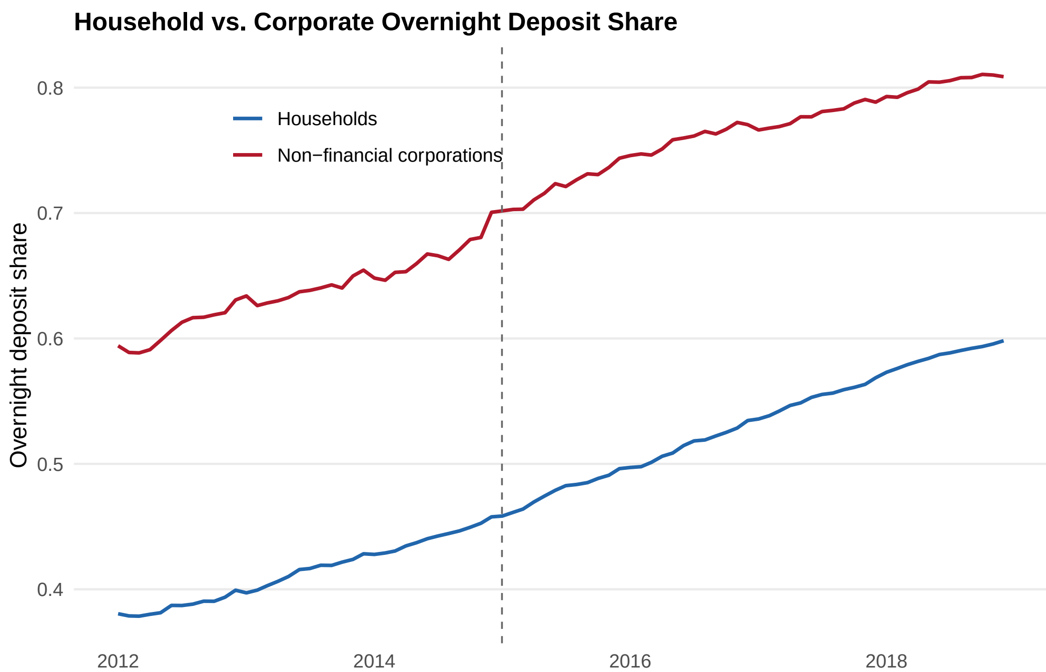

Figure 4: Household vs. Corporate Overnight Deposit Share

Notes: Cross-country average overnight deposit share by sector. Dashed vertical line at December 2014 indicates the BRRD transposition deadline. Monthly ECB BSI data, 19 EU countries with maturity disaggregation, 2012–2018.

6.5 Early vs. Late Transposers

Figure 5 displays the evolution of overnight and agreed-maturity deposit shares separately for early transposers (before the median transposition date) and late transposers (after the median). The overnight share is higher for early transposers throughout the sample, consistent with level differences that are absorbed by country fixed effects. The key pattern is the *relative* movement: early transposers show a visible uptick in overnight shares beginning in 2015, while late transposers' series remains flatter until late 2015. This visual evidence corroborates the event study results and illustrates the staggered treatment timing that identifies the effect.

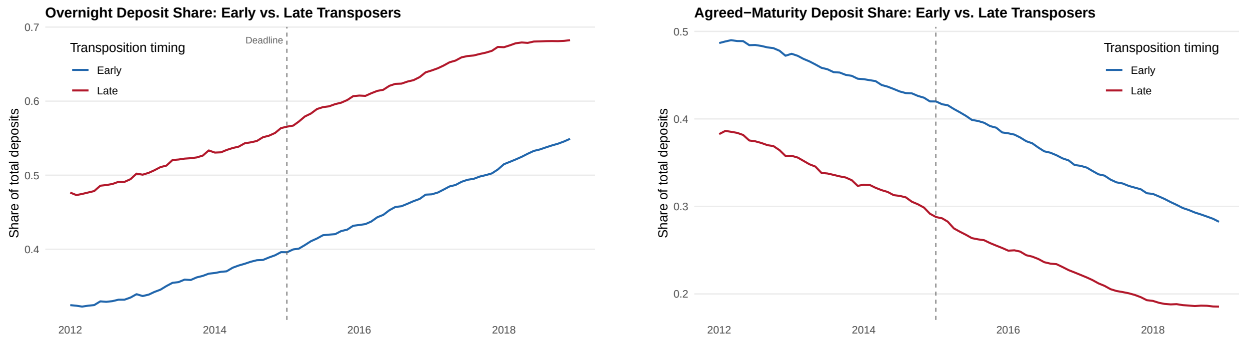

(a) Overnight deposits

(b) Agreed maturity

Figure 5: Deposit Composition by Transposition Timing Group

Notes: Countries are split by median transposition date into "Early" and "Late" groups. Monthly cross-country averages within each group. ECB BSI data, 19 EU countries with maturity disaggregation, 2012–2018.

6.6 Robustness

A battery of specification checks probes the sensitivity of the TWFE estimates; details are reported in Appendix C.[6]

Pre-treatment placebo. If the BRRD's publication in May 2014 triggered anticipatory deposit restructuring, interacting the publication-date indicator with uninsured exposure should be significant in the pre-transposition period. The coefficient is small and insignificant (-0.026, $p = 0.63$), ruling out an EU-wide anticipation channel operating before national transposition.

Leave-one-out. Sequentially dropping each of the 19 countries produces TWFE estimates ranging from -0.009 to $+0.001$, demonstrating that the baseline result is not driven by any single country (Figure 6). Figure 6 also shows that 95% confidence intervals overlap across all exclusions.

Excluding GIIPS countries. Five countries (Greece, Ireland, Italy, Portugal, Spain) experienced severe banking crises that may have endogenized their transposition timing. Greece is already excluded from the regression sample due to BSI data gaps; excluding the remaining four GIIPS members (Ireland, Italy, Portugal, Spain) produces a TWFE coefficient of -0.012 ($p = 0.16$), directionally consistent with the full sample but with reduced precision due to the smaller sample (15 countries).

[6]The robustness exercises below use the TWFE estimator, whose staggered-adoption bias produces attenuated or sign-reversed overnight coefficients relative to the Callaway-Sant'Anna preferred estimate. These checks therefore test the stability of the TWFE specification rather than directly validating the CS-DiD headline. The consistency of the intensity-interaction results (columns 4–5 of Table 2), which do not suffer from the same negative-weighting bias, provides the strongest robustness evidence.

Bail-in tool activation. The BRRD mandated that the bail-in tool be available from January 1, 2016—a common shock to all countries. Interacting this date with uninsured exposure produces a coefficient of −0.181 (SE = 0.085), suggesting that the bail-in tool's formal activation had an additional effect on high-exposure countries, operating through the same channel as transposition itself.

TWFE vs. robust estimators. The divergence between the TWFE baseline (−0.54 pp, insignificant) and the Callaway-Sant'Anna (+0.67 pp, $p = 0.041$) and Sun-Abraham (+3.10 pp, $p = 0.01$) estimates is itself informative. It illustrates the Goodman-Bacon (2021) decomposition in practice: with staggered adoption and heterogeneous treatment effects, TWFE produces attenuated or sign-reversed estimates. In this setting, the bias is substantial—the TWFE coefficient has the wrong sign.

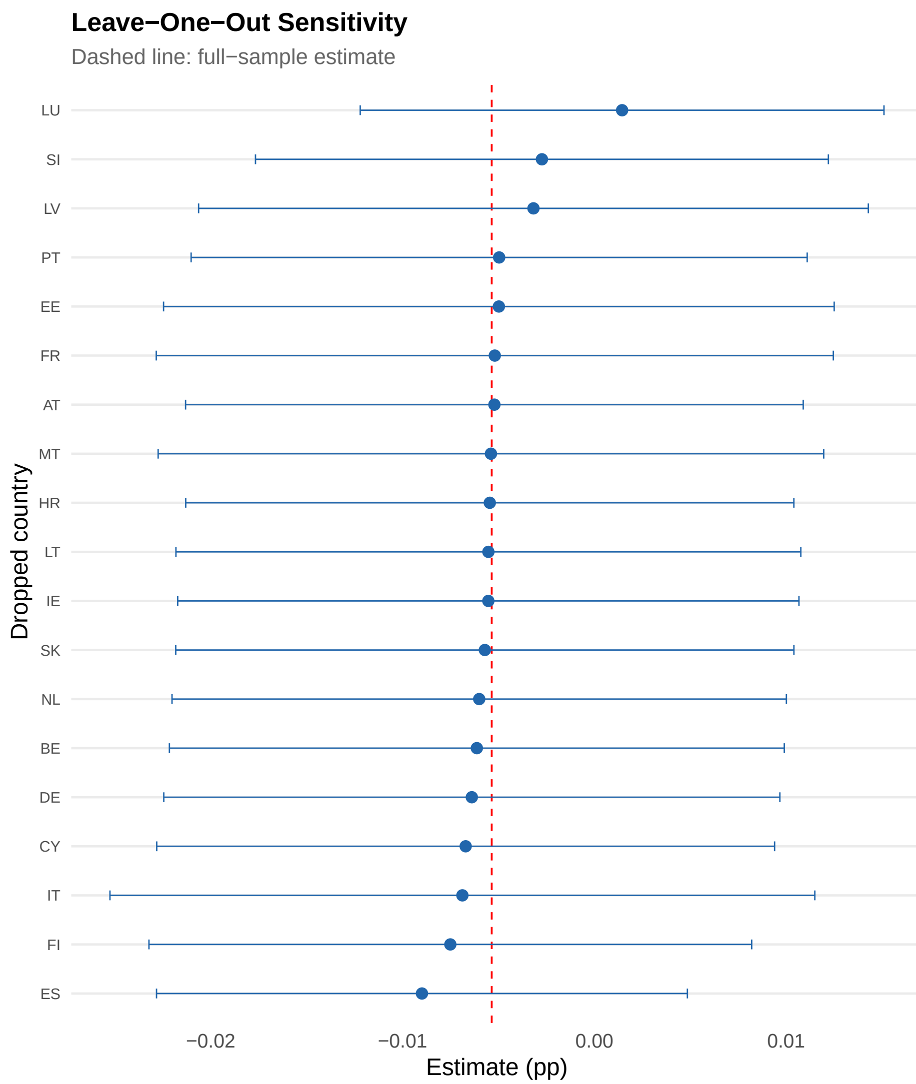

Figure 6: Leave-One-Out Sensitivity of TWFE Overnight Share Estimate

Notes: Each point represents the TWFE coefficient for overnight deposit share when the labeled country is excluded. Horizontal bars show 95% confidence intervals. Dashed red line shows the full-sample estimate.

7. Discussion

7.1 Economic Magnitude

The central finding is suggestive evidence that BRRD transposition induced a modest restructuring of household deposit portfolios toward more liquid instruments. The magnitude—0.67 percentage points on a base of 48.2%—is small in absolute terms, but its economic significance should be evaluated in the context of EU-wide deposit volumes. Total household deposits across the 19 sample countries averaged approximately 6.5 trillion euros during the sample period; a 0.67 pp shift in overnight share translates to roughly 43 billion euros in additional overnight deposits across the EU. For individual banks, a shift of this magnitude can meaningfully change the Liquidity Coverage Ratio (LCR) and Net Stable Funding Ratio (NSFR) calculations that regulators use to assess funding stability.

To put this in comparative perspective, the effect is substantially smaller than the deposit outflows documented during acute banking crises—Iyer and Puri (2012) document 10–15% deposit withdrawals during bank runs—but comparable in magnitude to the effects of deposit insurance changes studied by Demirgüç-Kunt and Detragiache (2002). The modest size is consistent with the BRRD being a *prospective* change in the resolution framework (creating the possibility of bail-in) rather than an *acute* event (an actual bank failure). The gradual onset visible in the event study—effects emerging 6–12 months after transposition—further supports the interpretation that the BRRD operates through a slow-moving awareness and portfolio adjustment channel rather than an immediate fear response.

7.2 Interpreting the Intensity-Interaction Results

The intensity-interaction results illuminate a more nuanced story than the simple average effect suggests. Countries with higher uninsured exposure experienced a disproportionate shift toward *agreed-maturity* deposits, not overnight deposits. This pattern is suggestive of the insurance optimization channel described in Section 3, under which households restructured to maximize insured returns rather than to maximize liquidity. However, the aggregate country-level data cannot directly observe within-household portfolio decisions, so this interpretation remains a hypothesis rather than established fact. The finding echoes Egan et al. (2017), who show that deposit insurance fundamentally shapes the competitive landscape for retail deposits; here, the introduction of bail-in risk may amplify the importance of the insurance boundary.

Why would high-exposure households shift *into* term deposits when bail-in risk should favor liquidity? One plausible interpretation is deposit splitting combined with maturity choice. Consider a household with 200,000 euros at a single bank. After learning about bail-in risk, this household could eliminate exposure by splitting deposits across two banks. Once the exposure is eliminated through splitting, the household faces no bail-in risk on either account and can freely choose the maturity structure that maximizes returns. Agreed-maturity deposits typically offer higher interest rates than overnight deposits, so the optimal post-splitting portfolio would favor term deposits. In aggregate country-level data, this would appear as a shift toward agreed maturity in high-exposure countries—reflecting below-threshold optimization by households who have diversified across banks. This interpretation is consistent with the data but cannot be directly verified with aggregate country-level deposit statistics; micro-level data on individual household deposit positions across banks would be needed to confirm this mechanism.

This interpretation generates an additional prediction: the agreed-maturity shift should be larger in countries where the yield spread between term and overnight deposits is wider. While testing this prediction requires bank-level interest rate data that is beyond the scope of the current analysis, it offers a productive direction for future research.

7.3 Methodological Implications

The divergence between TWFE (−0.54 pp, insignificant) and the heterogeneity-robust estimators (CS: +0.67 pp, $p = 0.041$; SA: +3.10 pp, $p = 0.01$) is not merely a technical curiosity. It demonstrates that TWFE bias can be severe enough to reverse the sign of the treatment effect, potentially leading applied researchers to conclude that a policy has *no effect* (or even the opposite effect) when it has a real and significant impact. This finding has implications beyond the BRRD context: any staggered policy evaluation in Europe—and EU regulatory harmonization typically proceeds through staggered national transposition—should be wary of TWFE estimates and prefer heterogeneity-robust alternatives.

The 4.6× magnitude gap between CS (+0.67 pp) and SA (+3.10 pp) also deserves comment. The Sun-Abraham estimator dropped 83 cohort-time interaction terms due to perfect multicollinearity, meaning the SA estimate is driven by a narrower set of cohort comparisons among the 19 regression-sample countries and may overweight specific cohort-time cells where treatment effects happen to be large. The CS estimator, which uses a different aggregation procedure and retains all cohort-time cells, produces a more conservative estimate that I treat as the preferred specification. However, the scale of this divergence—even after accounting for the collinearity explanation—signals estimation instability inherent in this small-cluster setting and warrants caution in interpreting the precise magnitude of either estimate.

7.4 Limitations

Five limitations warrant discussion. First, the aggregate country-level data cannot identify individual household responses. The composition shifts I observe could reflect heterogeneous responses across the deposit distribution—for example, wealthy households (with deposits above the guarantee threshold) may shift to term deposits at multiple banks, while middle-class households shift to overnight accounts for precautionary reasons. Micro-level deposit data, such as the ECB's Household Finance and Consumption Survey, could shed light on this heterogeneity.

Second, the 19-country sample is modest for a staggered DiD design. While the Callaway-Sant'Anna and Sun-Abraham estimators are designed for this setting, inference with small numbers of clusters remains a concern. To address this, I supplement the analytical standard errors with a multiplier bootstrap for the CS-DiD estimator (1,000 iterations), yielding a bootstrap p-value of 0.035—slightly more significant than the analytical $p = 0.041$—which provides reassurance that the headline result is not an artifact of small-cluster inference. The TWFE wild cluster bootstrap is also reported in Appendix C.

Third, I cannot fully separate the BRRD's effect from other Banking Union reforms implemented over the same period, including the Single Supervisory Mechanism (SSM) and the Single Resolution Mechanism (SRM). The BRRD is the component most directly relevant to household depositors, but the overall regulatory package may have contributed to the observed effects. The corporate deposit sectoral comparison provides some reassurance—if the Banking Union broadly (rather than bail-in specifically) were driving deposit restructuring, we would expect corporate deposits to move as well.

Fourth, the paper cannot cleanly distinguish the effects of national transposition from the common January 1, 2016 bail-in tool activation date, which applied simultaneously to all member states. The robustness check showing a significant bail-in activation × uninsured interaction (-0.181, SE = 0.085) supports this concern: the activation date may independently drive part of the observed deposit restructuring, and the staggered transposition design cannot fully separate the two channels. The gradual event-study dynamics are consistent with either national transposition effects or anticipation of the common activation date.

Fifth, the sample period (2012–2018) may be too short to capture the full long-run adjustment of household portfolios. Deposit restructuring is a slow process, particularly for less financially sophisticated households who may not learn about bail-in risk until an actual resolution event occurs in their country. The Italian bank resolutions of 2015 and the Banco Popular resolution of 2017 may have triggered additional adjustment that is partially captured in the later years of the sample, but a longer post-treatment window would provide more definitive evidence.

7.5 Policy Implications

If confirmed by future research with more granular data, these results would have implications for the design of resolution frameworks. If bail-in risk shortens the maturity of bank funding, it may create a tension between the policy goals of reducing moral hazard (by imposing losses on creditors) and maintaining financial stability (by preserving banks' ability to engage in maturity transformation). The BRRD was designed to make the financial system safer by eliminating the implicit guarantee that taxpayers would rescue failing banks. But if the consequence is that depositors restructure toward shorter-maturity instruments—making banks more dependent on volatile overnight funding—the net effect on financial stability is ambiguous.

Policymakers may need to consider how the minimum requirement for own funds and eligible liabilities (MREL) and the total loss-absorbing capacity (TLAC) standards interact with depositor behavior to affect overall bank funding stability. If banks must hold more bail-in-able liabilities (MREL), and if households shift away from these liabilities toward insured short-term deposits, banks may face higher funding costs and shorter average funding maturities—precisely the conditions that increase vulnerability to liquidity crises.

The insurance optimization channel suggested by the intensity-interaction results, if operative, would also have implications for the design of deposit guarantee schemes. If households respond to bail-in risk by splitting deposits across banks to stay below the 100,000 euro threshold, the effective coverage of the DGS increases (more deposits become insured), which increases the DGS's contingent liability. European policymakers debating the proposed EDIS should consider how the BRRD's bail-in framework has already altered the distribution of deposits relative to the guarantee threshold.

8. Conclusion

This paper presents suggestive evidence that the introduction of bail-in risk through the BRRD reshaped the maturity structure of European household deposits. While the CS-DiD estimate of 0.67 pp ($p = 0.041$) with 19 clusters and analytical standard errors warrants interpretive caution, the finding points to a fundamental trade-off in resolution policy: the credibility of bail-in as a disciplining device depends on depositors believing their funds are at risk, but the behavioral response to that risk—shortening deposit maturities and optimizing across insurance boundaries—may weaken precisely the funding stability that resolution frameworks aim to protect.

The most important insight is not the average effect but the heterogeneity: countries with greater bail-in exposure experienced a shift toward term deposits, not away from them. This is suggestive of households responding to bail-in risk not primarily through flight to liquidity but through optimization within the deposit guarantee framework, though the aggregate data cannot directly observe within-household behavior. If this interpretation holds, it would challenge the assumption, common in banking regulation, that household depositors are informationally inert. Understanding how different depositor segments respond to resolution risk may prove important for calibrating bail-in thresholds, deposit preference rules, and the scope of deposit guarantee schemes in the next generation of bank resolution reforms.

Three avenues for future research are particularly promising. First, the advent of granular supervisory data—such as the ECB's AnaCredit database and the Household Finance and Consumption Survey—would enable a household-level analysis that distinguishes between deposit splitting (across banks) and maturity restructuring (within banks), the two margins that aggregate country data conflates, and would allow direct testing of the insurance optimization hypothesis that remains conjectural in this paper. Second, the temporal dynamics of the effect merit further investigation: does the gradual post-transposition increase reflect slow diffusion of awareness, or is it driven by specific salient events (the Italian resolutions of 2015, Banco Popular in 2017) that made bail-in risk concrete? Third, the interaction between bail-in risk and the ECB's negative interest rate policy—which compressed the yield advantage of term deposits—deserves attention, as it may have moderated or amplified the BRRD's effect on deposit composition through its impact on the relative attractiveness of overnight versus agreed-maturity deposits.

Acknowledgements

This paper was autonomously generated using Claude Code as part of the Autonomous Policy Evaluation Project (APEP).

Project Repository: https://github.com/SocialCatalystLab/ape-papers

Contributors: @ai1scl

First Contributor: https://github.com/ai1scl

References

Acharya, Viral V., Itamar Drechsler, and Philipp Schnabl, "A Pyrrhic Victory? Bank Bailouts and Sovereign Credit Risk," *Journal of Finance*, 2014, *69* (6), 2689–2739.

Berndt, Antje, Darrell Duffie, and Yichao Zhu, "The Effects of the Bail-in Tool on Bank Funding Costs," *Journal of Financial and Quantitative Analysis*, 2023.

Bonfim, Diana and João A.C. Santos, "The Composition Effect of Deposit Insurance on Bank Deposits," *Journal of Financial Intermediation*, 2020, *43*, 100817.

Callaway, Brantly and Pedro H.C. Sant'Anna, "Difference-in-Differences with Multiple Time Periods," *Journal of Econometrics*, 2021, *225* (2), 200–230.

Conlon, Thomas and John Cotter, "The Effect of Bail-in on Bank Bond Yields and Market Discipline," *Journal of Financial Stability*, 2019, *42*, 100707.

Crespi, Fabrizio, Emanuela Giacomini, and Danilo Valerio Mascia, "Bail-in vs Bail-out: Bank Resolution and Liability Structure," *International Review of Financial Analysis*, 2021, *73*, 101642.

Demirgüç-Kunt, Asli and Enrica Detragiache, "Does Deposit Insurance Increase Banking System Stability? An Empirical Investigation," *Journal of Monetary Economics*, 2002, *49* (7), 1373–1406.

Diamond, Douglas W. and Philip H. Dybvig, "Bank Runs, Deposit Insurance, and Liquidity," *Journal of Political Economy*, 1983, *91* (3), 401–419.

Egan, Mark, Ali Hortacsu, and Gregor Matvos, "Deposit Competition and Financial Fragility: Evidence from the US Banking Sector," *American Economic Review*, 2017, *107* (1), 169–216.

European Banking Authority, "Report on Deposit Guarantee Scheme Coverage Level and Practices," Technical Report, EBA 2015.

European Central Bank, "Statistical Data Warehouse: Balance Sheet Items," Technical Report, ECB 2024. Accessed February 2026.

Fischer, Marcel, Christa Hainz, Jörg Rocholl, and Sascha Steffen, "The Effect of the EU Bank Recovery and Resolution Directive on Equity and Bond Prices," *Research in International Business and Finance*, 2018, *46*, 205–225.

Goodman-Bacon, Andrew, "Difference-in-Differences with Variation in Treatment Timing," *Journal of Econometrics*, 2021, *225* (2), 254–277.

Ignatowski, Magdalena and Josef Korte, "Wishful Thinking or Effective Threat? Tightening Bank Resolution Regimes and Bank Risk-Taking," *Journal of Financial Stability*, 2015, *15*, 264–281.

Iyer, Rajkamal and Manju Puri, "Understanding Bank Runs: The Importance of Depositor-Bank Relationships and Networks," *American Economic Review*, 2012, *102* (4), 1414–1445.

Kashyap, Anil K. and Jeremy C. Stein, "What Do a Million Observations on Banks Say about the Transmission of Monetary Policy?," *American Economic Review*, 2000, *90* (3), 407–428.

Leibniz Institut für Wirtschaftsforschung Halle, "IWH Banking Union Directives Database," *IWH*, 2022.

Publications Office of the European Union, "CELLAR SPARQL Endpoint for EU Legislation," Technical Report, EU Publications Office 2024. Accessed February 2026.

Schäfer, Alexander, Isabel Schnabel, and Beatrice Weder di Mauro, "Bail-in Expectations for European Banks: Actions Speak Louder than Words," *Journal of Financial Economics*, 2016.

Sun, Liyang and Sarah Abraham, "Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects," *Journal of Econometrics*, 2021, *225* (2), 175–199.

Véron, Nicolas, "Europe's Radical Banking Union," *Bruegel Essay and Lecture Series*, 2015.

A. Data Appendix

A.1 ECB BSI Data Retrieval

Monthly deposit data are retrieved from the ECB Statistical Data Warehouse (SDW) via the SDMX-REST API. The BSI key format is:

```
BSI/M.{country}.N.A.{item}.A.1.U2.{sector}.Z01.E
```

where `item` $\in$ {L20 (total), L21 (overnight), L22 (agreed maturity), L23 (redeemable at notice)} and `sector` $\in$ {2250 (households + NPISH), 2240 (non-financial corporations)}. Data are reported in all currencies (Z01 dimension) for deposit items, which differs from the EUR denomination used for asset items.

Three countries are excluded from the panel entirely: Bulgaria (non-euro area with incomplete BSI coverage), Denmark (non-euro area with limited maturity disaggregation), and Greece (BSI data gaps during the capital controls period of 2015–2017). The resulting 24-country panel covers Austria, Belgium, Cyprus, Czech Republic, Germany, Estonia, Spain, Finland, France, Croatia, Hungary, Ireland, Italy, Lithuania, Luxembourg, Latvia, Malta, Netherlands, Poland, Portugal, Romania, Sweden, Slovenia, and Slovakia. Of these, five non-eurozone countries—Czech Republic, Hungary, Poland, Romania, and Sweden—report total deposits but not the overnight/agreed-maturity/redeemable-at-notice breakdown. These five countries are excluded from all regression specifications, yielding a 19-country regression sample.

A.2 BRRD Transposition Dates

Transposition dates are sourced from the IWH Banking Union Directives Database (Leibniz Institut für Wirtschaftsforschung Halle, 2022), which systematically codes the primary national law implementing each Banking Union directive across EU member states. Cross-validation against the CELLAR SPARQL endpoint of the EU Publications Office yields 393 national implementation measures (NIMs) for the BRRD (CELEX number 32014L0059), covering all 27 member states plus the United Kingdom.

A.3 EBA Deposit Guarantee Scheme Data

The treatment intensity variable is constructed from EBA aggregate data on deposit guarantee scheme coverage (European Banking Authority, 2015). For each country, the covered-to-eligible deposit ratio represents the proportion of deposits below the 100,000 euro threshold. The uninsured share is computed as $1 - (\text{covered}/\text{eligible})$. These data are cross-sectional (2015 vintage) and are treated as pre-determined with respect to BRRD transposition.

Table 5: BRRD National Transposition Dates and Deposit Insurance Coverage

Country	Transposition Date	Days Relative to Deadline	Uninsured Share
Austria	January 01, 2015	1	33.0%
Belgium	January 01, 2015	1	36.0%
Cyprus	December 19, 2014	-12	38.0%
Czechia	June 01, 2015	152	22.0%
Germany	January 01, 2015	1	42.0%
Estonia	July 01, 2015	182	21.0%
Spain	June 19, 2015	170	32.0%
Finland	December 31, 2014	0	40.0%
France	August 20, 2015	232	37.0%
Croatia	April 24, 2015	114	20.0%
Hungary	December 31, 2014	0	17.0%
Ireland	January 01, 2015	1	44.0%
Italy	November 16, 2015	320	31.0%
Lithuania	March 26, 2015	85	19.0%
Luxembourg	December 18, 2015	352	58.0%
Latvia	July 01, 2015	182	22.0%
Malta	July 31, 2015	212	45.0%
Netherlands	January 01, 2015	1	39.0%
Poland	February 03, 2016	399	18.0%
Portugal	March 31, 2015	90	29.0%
Romania	October 23, 2015	296	16.0%
Sweden	February 01, 2016	397	43.0%
Slovenia	January 01, 2015	1	27.0%
Slovakia	January 01, 2015	1	20.0%

Notes: Transposition dates from the IWH Banking Union Directives Database, cross-validated with CELLAR SPARQL national implementation measures. Uninsured share computed as one minus the covered-to-eligible deposit ratio from EBA Deposit Guarantee Scheme aggregate data (2015). The BRRD transposition deadline was December 31, 2014. Three countries (Bulgaria, Denmark, Greece) are excluded due to incomplete BSI coverage. Five countries (Czech Republic, Hungary, Poland, Romania, Sweden) are included in this table for completeness but lack maturity-disaggregated BSI data and are excluded from all regression analyses.

A.4 Variable Definitions

- **Overnight share:** $L21_{it}/L20_{it}$, the ratio of household overnight deposits to total household deposits.
- **Agreed-maturity share:** $L22_{it}/L20_{it}$.
- **Redeemable-at-notice share:** $L23_{it}/L20_{it}$.
- **Post-BRRD:** $\mathbb{I}[t \geq T_i^*]$, where T_i^* is the YYYY-MM of national BRRD transposition.
- **Uninsured share:** $1 - (\text{covered}_i/\text{eligible}_i)$ from EBA DGS data (2015).
- **Post × Uninsured:** $\text{Post-BRRD}_{it} \times \text{Uninsured}_i$.
- **Log total deposits:** $\log(L20_{it})$, the natural logarithm of total household deposits in millions of euros.

B. Identification Appendix

B.1 BRRD Treatment Rollout

Figure 7 displays the complete transposition timeline for the 24 EU member states in the BSI sample. The transposition deadline (December 31, 2014) and the bail-in tool mandatory date (January 1, 2016) are marked. Three countries (Cyprus, Finland, Hungary) transposed on or before the deadline, seven more (Austria, Belgium, Germany, Ireland, Netherlands, Slovenia, Slovakia) transposed on January 1, 2015, and the remaining 14 transposed between March 2015 and February 2016.

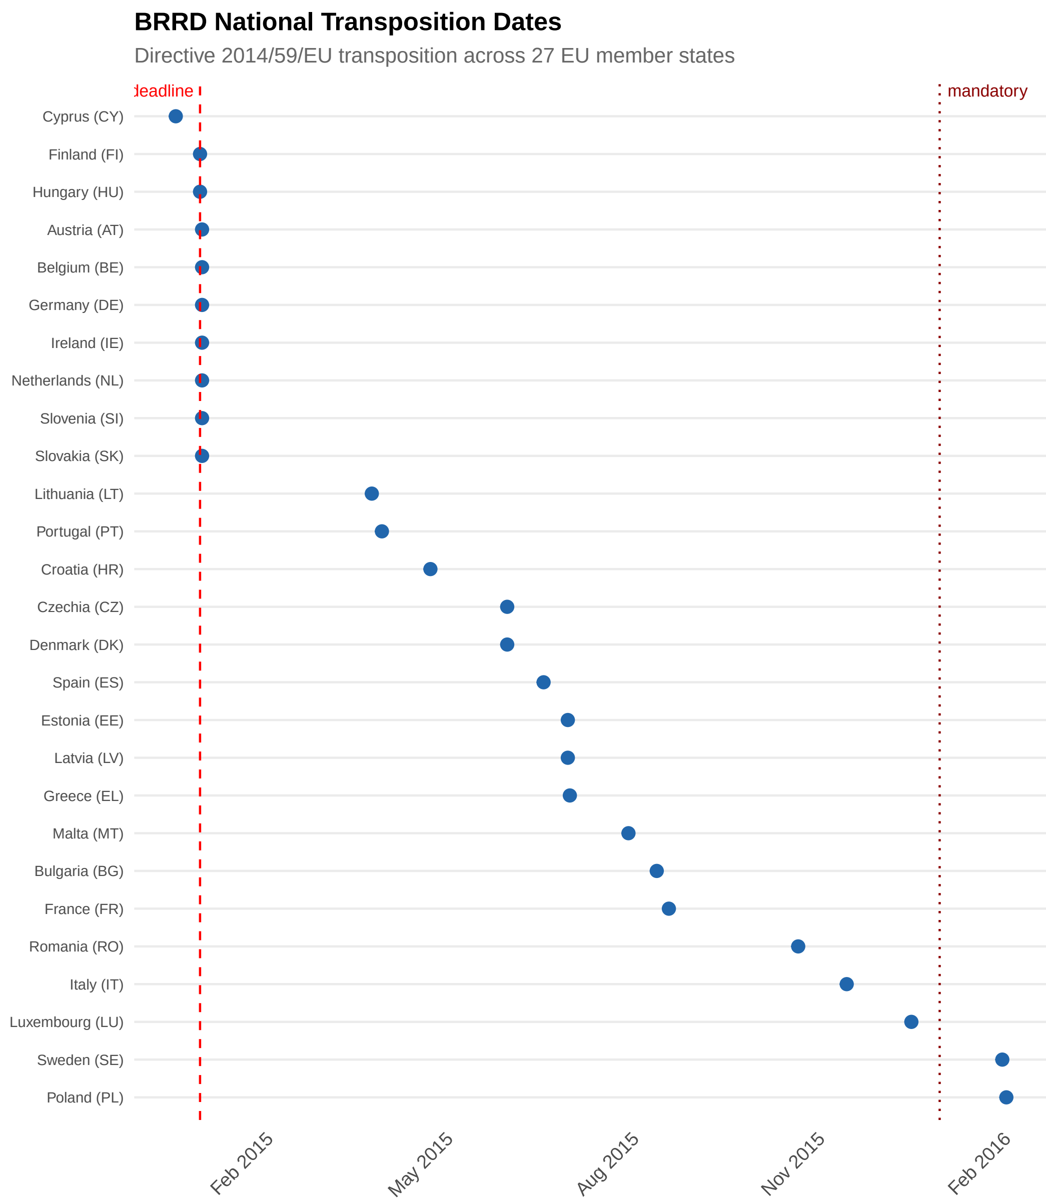

Figure 7: BRRD National Transposition Timeline

Notes: Each point represents the transposition date of the primary national law implementing the BRRD. Red dashed line: transposition deadline (December 31, 2014). Dark red dotted line: bail-in tool mandatory date (January 1, 2016). Sources: IWH Banking Union Directives Database, CELLAR SPARQL.

B.2 Pre-Trend Validation

The Callaway-Sant'Anna event study in Figure 1 provides the primary pre-trend diagnostic. All 24 pre-treatment coefficients (months -24 to -1) are individually insignificant at the 5% level, and there is no monotonic trend in the pre-treatment path. The corresponding event study for agreed-maturity deposits (Figure 2) similarly shows no pre-trends.

B.3 Randomization Inference

To assess the sharpness of the treatment assignment, I conduct a randomization inference exercise that permutes transposition dates across countries 1,000 times and re-estimates the TWFE specification for each permutation. The two-sided RI p-value for the TWFE overnight coefficient is 0.463, reflecting the fact that the TWFE estimator is biased and underpowered in this staggered setting. The distribution of permuted estimates (Figure 8) is centered near zero with standard deviation 0.007, confirming that the observed TWFE coefficient (-0.005) is well within the range of null estimates. This underscores the importance of using heterogeneity-robust estimators rather than TWFE for inference.

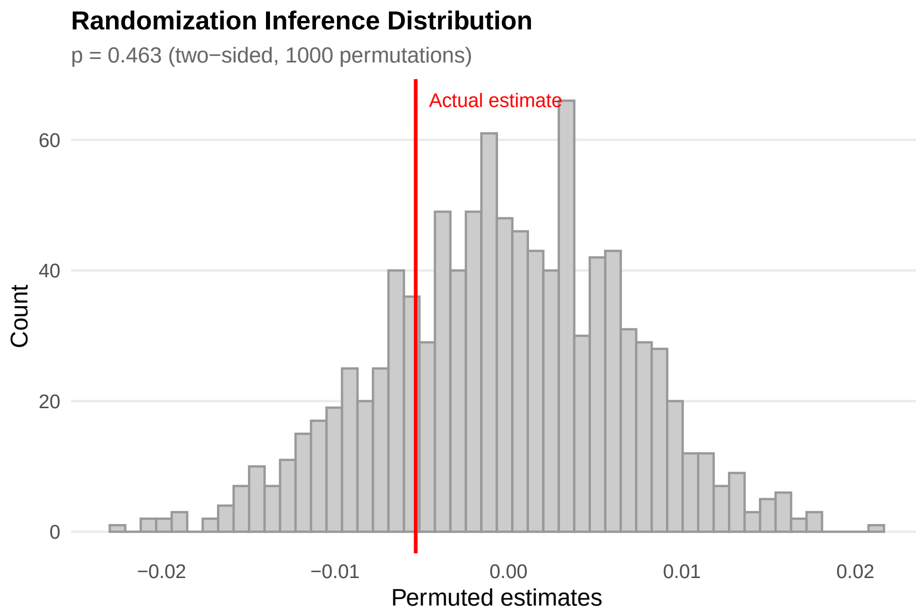

Figure 8: Randomization Inference: Distribution of Permuted TWFE Estimates

Notes: Histogram of TWFE overnight share coefficients from 1,000 random permutations of transposition dates. Red vertical line: actual TWFE estimate. The RI p-value (0.463) reflects the TWFE's attenuation bias; the Callaway-Sant'Anna estimate (+0.0067) corrects for this bias.

C. Robustness Appendix

The following additional robustness checks are referenced in the main text. Table 6 summarizes all robustness checks.

Table 6: Summary of Robustness Checks

Check	Estimate	SE	N	Note
Leave-one-out range	[-0.009, 0.0014]	—	19 countries	Stable across all country exclusions
RI p-value (TWFE)	0.463	—	1000 permutations	Two-sided, 1000 permutations
Placebo: publication x uninsured	-0.0255	0.0514	740	Pre-transposition period only
Bail-in tool x uninsured	-0.1811	0.0848	1596	Jan 1, 2016 bail-in tool x cross-sectional exposure
Excluding GIIPS	-0.0117	0.0078	1260	N=15 countries
Log total deposits	-0.0056	0.0097	1596	Extensive margin: total deposit growth
Log deposits x intensity	-0.3469	0.2426	1596	Intensive margin: treatment intensity on total deposits

Notes: Summary of robustness checks from the main TWFE and intensity-interaction specifications. Leave-one-out reports the range of TWFE overnight coefficients when each country is dropped. RI p-value tests the TWFE specification (which is biased under staggered adoption). All specifications include country and time fixed effects with standard errors clustered at the country level.

Placebo timing test. Restricting the sample to the pre-transposition period (2012 to each country's transposition date), I interact a post-publication indicator (June 2014) with the uninsured share. The coefficient is -0.026 (SE $= 0.051$, $p = 0.63$), indicating no anticipatory response to the BRRD's announcement.

Bail-in tool activation. The interaction of the January 2016 bail-in tool date with uninsured exposure produces a coefficient of -0.181 (SE $= 0.085$). This suggests that the formal activation of bail-in powers had an additional compression effect on overnight shares among high-exposure countries, a pattern suggestive of the insurance optimization mechanism. However, this result also raises a concern about treatment timing: the common January 1, 2016 activation date cannot be separated from national transposition effects for countries that transposed close to that date.

Alternative outcomes. Log total household deposits show no significant change following BRRD transposition (-0.006, SE $= 0.010$), indicating that the effects operate through the *composition* of deposits (intensive margin) rather than total deposit volumes (extensive margin). The deposit growth-intensity interaction is also insignificant (-0.347, SE $= 0.243$).

D. Heterogeneity Appendix

Figure 5 in the main text presents the primary heterogeneity analysis: early vs. late transposers. Additional cross-sectional heterogeneity (by eurozone membership, by banking sector concentration, by sovereign credit rating) is infeasible with the current sample size of 19 countries—the number of clusters in each subgroup is too small for reliable inference.

E. Additional Figures and Tables

All main figures and tables are presented in the body of the paper and in the preceding appendix sections. No additional exhibits are needed.

F. Standardized Effect Sizes

Table 7: Standardized Effect Sizes for Main Outcomes

Outcome	Specification	$\hat{\beta}$	SD(X)	SD(Y)	SDE	Classification
Overnight share	CS-DiD, ATT	+0.0067	—	0.206	+0.033	Null
Overnight share (intensity)	Post × Uninsured, Table 2 Col. 4	−0.165	0.104	0.206	−0.083	Small negative
Agreed-maturity share (intensity)	Post × Uninsured, Table 2 Col. 5	+0.261	0.104	0.224	+0.121	Large positive
Redeemable share	TWFE, Table 2 Col. 3	−0.0046	—	0.239	−0.019	Null

Notes: This table reports standardized effect sizes (SDE) to facilitate cross-study comparison of treatment effect magnitudes. For binary (0/1) treatments (Post-BRRD indicator), SDE $= \hat{\beta}/\text{SD}(Y)$ and the SD(X) column is marked "—". For continuous treatments (Post × Uninsured interaction), SDE $= \hat{\beta} \times \text{SD}(X)/\text{SD}(Y)$, which gives the effect of a one-standard-deviation change in the treatment variable, measured in standard deviations of the outcome. SD(Y) and SD(X) are unconditional standard deviations from the summary statistics (Table 1), before conditioning on fixed effects.

Research question: Does the introduction of bail-in risk through the EU Bank Recovery and Resolution Directive (BRRD) affect the maturity composition of household deposits? **Treatment:** Post-BRRD transposition (binary 0/1, staggered across 19 EU member states with maturity-disaggregated BSI data) and its interaction with cross-sectional uninsured deposit share (continuous, range 0–1). **Data:** ECB Balance Sheet Items (BSI) monthly data, January 2012 to December 2018, country-level; N = 1,596 observations for overnight and agreed-maturity outcomes (19 × 84 balanced panel); N = 1,428 for redeemable-at-notice (Spain and Croatia lack this series). **Method:** Callaway-Sant'Anna heterogeneity-robust DiD for binary treatment (analytical SEs); two-way fixed effects for intensity interaction (country-clustered SEs). **Sample:** 19 EU member states with maturity-disaggregated BSI data. Excludes Bulgaria, Denmark, Greece (incomplete BSI coverage) and Czech Republic, Hungary, Poland, Romania, Sweden (no maturity breakdown).

Classification thresholds: large negative (< -0.10), small negative (-0.10 to -0.05), null (-0.05 to 0.05), small positive (0.05 to 0.10), large positive (> 0.10).