Economics School of Louvain - ESL

Who Pays for Political Risk?

Banks, Borrowers, and the Distribution of Losses in a Monetary Union

Author : Jauder Uribarri Sánchez

Thesis Director : Francesca Monti

Thesis Reader : Rigas Oikonomou

Academic Year 2025 2026

Master in Economics – 120 credits – Econometrics Focus

Université catholique de Louvain - Economics School of Louvain - Place Montesquieu 3, 1348 Louvain-La-Neuve

Acknowledgements.
First, I would like to thank Professor Francesca Monti, my supervisor, for her guidance and feedback throughout the project. I would also like to thank Rigas Oikonomou, my reader. Finally, I would like to thank my family and friends for their support with this project and throughout my Master's degree. I will always remember my LLN experience.

Abstract

Who bears sovereign risk? Losses on government bonds appear highly concentrated among wealthy households, since financial assets sit at the top of the wealth distribution. This paper shows that a channel through bank balance sheets changes that incidence. Because banks also hold government debt, sovereign repricing weakens their balance sheets, and can be transmitted into tighter credit conditions. This reaches borrowers, rather than asset holders, and the two groups can occupy very different parts of the wealth distribution. Because debt is less concentrated among wealthy households than financial assets, the banking channel shifts part of the burden down the wealth distribution. This happens in our event study in Italy, where the bottom two wealth quintiles bear only 4.8% of direct asset losses, but 26.5% of the borrowing-cost component.

The mechanism is quantified with a model that traces sovereign losses from bank balance sheets to household borrowing conditions. A heterogeneous-agent New Keynesian (HANK) model allows the resulting borrower wedge to affect households with different borrowing positions, and tracks their consumption and saving responses, while balance sheet data map the direct and credit components across the wealth distribution. The results show that the aggregate effect is small, but the distributional composition changes meaningfully. Wealthy households still lose the most, but the banking channel shifts part of the burden toward lower-wealth households.

JEL classification: D31, E44, F45, G21, H63

Keywords: sovereign risk; bank net worth; heterogeneous agents; monetary union; household incidence

Contents

CHAPTER 1

Introduction

Government bonds and financial assets sit at the top of the wealth distribution. In Italy the richest fifth of households holds roughly three quarters of the government securities that households own directly, and a similar share of all financial assets. When a sovereign repricing destroys the value of those claims, the loss lands on households who were already wealthy. Read through asset prices alone, sovereign stress looks close to a progressive tax.

That reading leaves out the banks. Domestic banks hold a large share of their own government's debt, so a wider spread lowers their net worth as well, and weaker banks lend on tighter terms. The households who pay for that are the ones who borrow, and borrowers are not bondholders. In the Italian survey data, bank debt is spread almost evenly from the second wealth quintile upward, while securities climb steeply with wealth. The same sovereign movement therefore reaches two different groups of households depending on which channel carries it.

This is the subject of the thesis. Sovereign risk has an incidence and not only an aggregate cost, and the sovereign–bank nexus changes who bears it. The question is the following.

> *When political risk is priced into wider sovereign spreads within a monetary union, how do domestic banks' holdings of sovereign debt amplify the effects on households, and who ultimately bears the losses?*

Italy 2018 supplies the application. After the inconclusive March election, the prospect of a Five Star–League government widened the BTP–Bund spread by 166 basis points at its peak against the pre-event baseline. Mapped through the observed balance sheet of the Italian banking sector, that movement costs the representative bank 14.7% of its equity in economic value and raises the rate faced by borrowers by roughly four to eight basis points.

The distributional answer has two halves and they run in opposite directions. In levels, wealthy households still bear most of the loss, exactly as the asset-price reading predicts. The top quintile loses 104 basis points of its own consumption against 3 for the bottom quintile, and keeps about three quarters of the total even after the credit channel is switched on. In composition, the credit channel moves part of the burden downward. The bottom two quintiles' share of the total rises from 4.80% without it to 5.99–6.08% with it. The aggregate response is small under the baseline closure, about 0.003% of output on impact, so what the exercise measures is a reallocation.

That direction is not a product of the calibration, and the thesis derives it. Adding a bank-credit channel makes a sovereign loss less top-heavy if and only if the borrowers that channel reaches are less concentrated at the top than the holders of the assets that reprice. The condition contains no parameter of any model. It compares two exposure shares that a household balance-sheet survey measures directly, and Italian data satisfy it by a factor of five and a half, with the bottom two quintiles carrying 4.80% of the asset-side loss and 26.51% of the borrower-side loss. What the calibration then supplies is the weight the credit channel receives, and therefore the size of the reallocation rather than its sign.

What makes a sovereign movement dangerous to a banking system is a product rather than a sum. The loss on the sovereign portfolio, as a share of equity, is domestic exposure times portfolio duration times leverage times the spread movement, so a system is badly hit only when several of those line up at once. France 2024 is the useful contrast. French banks carry longer-duration portfolios and slightly higher leverage than Italian ones, yet the same mapping gives a bank loss fourteen times smaller and a borrower wedge of 0.41 basis points, because they hold roughly one fifth as much of their own government's debt. A common currency does not imply common exposure, which is why two systems in the same union can be affected so unequally.

Figure 1.1 sets out the stages of the exercise and the evidence behind each. Political chronology fixes the episode and the window, and market data measure the sovereign movement inside it. Conditional on that movement, observed exposure, duration and leverage give the loss on banks' government-bond portfolios. Italy-2011 lending evidence calibrates the step from bank net worth to borrowing costs. A one-asset heterogeneous-agent New Keynesian (HANK) model, solved with the sequence-space Jacobian method of Auclert et al. (2021), turns the resulting wedge into aggregate and household responses. Fixed balance-sheet vectors from the Survey on Household Income and Wealth (SHIW) then allocate the losses across net-wealth quintiles. Bank equity returns provide a separate check on the first half of the chain, asking whether more exposed institutions did worse over the event window.

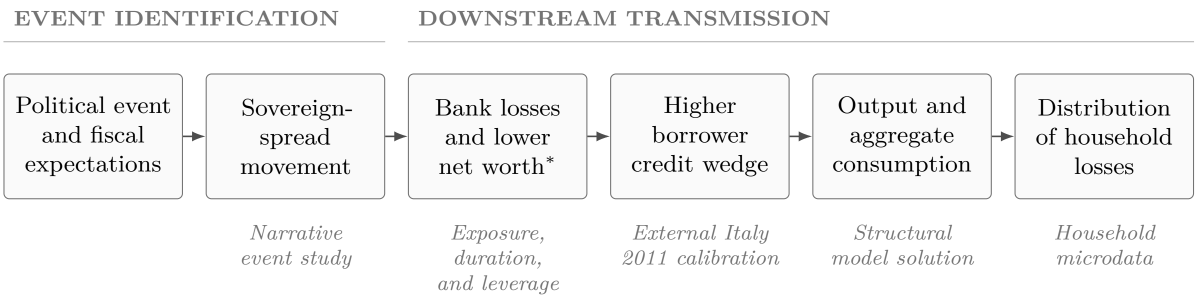

*Supporting bank-market evidence assesses whether institutions with greater sovereign exposure experience larger equity and credit default swap (CDS) losses.

Figure 1.1. Identification and transmission framework. Market data measure the event-associated repricing, observed balance sheets determine the bank loss, the bank-to-credit mapping is calibrated externally on Italy 2011, and fixed household vectors allocate the resulting incidence across wealth quintiles.

The four episodes follow a strict hierarchy, so that no episode both supplies a calibrated parameter and receives the headline result. The bank-to-credit step is calibrated outside the main application. Its structure follows Bocola (2016), and its scale combines the Italy-2011 lending-rate differential of Bofondi et al. (2018) with the Italy-2011 bank-loss mapping constructed here. That parameter is then held fixed. Italy 2018 is the main application with high exposure, France 2024 the comparison with low exposure, and Italy 2022 a policy contrast that estimates nothing about the Transmission Protection Instrument. Calibrating on 2011 and applying to 2018 keeps the exercise non-circular.

The exercise is conditional, not causal. Three steps make it so. The sovereign repricing is taken as observed and never identified as a political shock. The bank-to-credit elasticity is disciplined externally on Italy-2011 firm-lending evidence, with nothing estimated for Italian households in 2018. And incidence is evaluated on fixed SHIW balance sheets, never on observed contracts. Two things nevertheless survive that conditionality. The France comparison rests only on identically constructed inputs and an accounting identity, and every incidence exhibit changes one object at a time against a fixed household set, common weights and a common denominator. What the calibration puts at risk is therefore the size of the reallocation rather than its existence, and Chapters 4 to 6 report how far that size moves with timing and calibration.

The thesis sits between two literatures. Work on the sovereign–bank nexus studies how bank balance sheets amplify sovereign stress, and mostly stops at financial stability, firm borrowing, or aggregate lending. Work on heterogeneous households studies who bears a banking-sector loss, and mostly starts from a generic shock rather than a dated national episode. Neither asks the incidence question that sits between them. Once a sovereign movement reaches households through asset prices and bank credit at the same time, the relative strength of the two channels decides who pays. This thesis states the condition under which the credit channel shifts the burden downward, shows that Italian microdata satisfy it, and quantifies the shift for one dated episode using that country's own measured exposures and a record of when each balance-sheet input was observed. It also checks its own aggregation, replacing the representative intermediary with twenty-five named banks.

The thesis proceeds as follows. Chapter 2 reviews the literature on political risk, the sovereign-bank nexus, credit supply, and household heterogeneity. Chapter 3 sets out the sovereign-bank block and the borrower-credit HANK model. Chapter 4 describes the episodes, the measurement of sovereign repricing, and the calibration. Chapter 5 reports the dated-bridge results, the market check, the aggregate and distributional findings, and a bank-aggregation sensitivity. Chapter 6 discusses what they mean, what they imply for policy, and where the exercise stops. Appendix A documents numerical validation, Appendix B records data provenance, Appendix C reports calibration and allocation robustness, and Appendix D describes the two-asset portfolio-rebalancing extension.

CHAPTER 2

Literature Review

This chapter reviews the four areas of literature that the thesis draws on. The links in the chain studied here are established by political risk and sovereign pricing, the sovereign–bank nexus, sovereign stress and credit supply, and household heterogeneity. The literature reviewed here does not jointly connect these links in a single, dated episode, which is the gap that the thesis addresses.

2.1 Political Risk and Sovereign Repricing

Three papers provide the necessary information for a political date to be considered an economic input. Pástor and Veronesi (2013) explain why such a date is priced at all: investors demand compensation for exposure to political news. Hazell and Hobler (2025) offer a high-frequency narrative benchmark: they identify an event that released information about future US deficits, quantify its narrative content, analyse high-frequency asset-price responses, and conclude with a heterogeneous-agent model that reproduces its magnitude. This thesis does not reproduce that causal design. It instead uses an event-associated Italian sovereign repricing as the input to a conditional exercise on household borrowing conditions.

Handler and Jankowitsch (2025) apply this approach to the correct asset and country, documenting the effects on prices and liquidity around Italian summits and elections. Prices fall before these events amid illiquidity and sell-side pressure, and event returns are positive afterwards. The window construction in Chapter 4 is informed by this evidence: it uses a pre-specified pre-event baseline and a peak over a broader horizon, rather than an announcement-day return.

Nevertheless, sovereign spreads are composite, reflecting default risk, redenomination risk, liquidity, safe-haven demand and common-rate expectations (Afonso et al., 2015; Corradin and Schwaab, 2023). As no narrow window separates these components, the measured object is an event-associated spread movement throughout.

2.2 The Sovereign–Bank Nexus and Domestic Home Bias

Three papers establish the transmission this thesis takes as given. Gennaioli et al. (2014) show that sovereign trouble is costly because it destroys the balance sheets of the domestic banks holding the debt, and that private credit falls further where those banks hold more of it. Their trigger is outright default, so what transports here is the direction. Bocola (2016) carries that direction into a structural transmission from bank net worth to borrowing conditions, which the bank block of Chapter 3 follows in reduced form. De Marco (2019) measures the real effects of banks' mark-to-market sovereign losses and shows that they bite through short-term funding and not through recognised equity. The economic-value loss measure of Section 3.2 is taken from that finding.

Why banks hold the debt bears on interpretation. Acharya and Steffen (2015) find that euro-area bank equity returns are correlated with peripheral sovereign bonds in proportion to holdings, most strongly among large, thinly capitalised banks. The bank-market check in Section 5.3 borrows that specification and its capitalisation controls. Battistini et al. (2014) find peripheral banks *increased* domestic exposure as their own sovereign's risk rose, so exposure is chosen, not assigned. Two consequences follow for this thesis.

First, home bias creates a usable exposure gradient. Banks holding more domestic public debt will lose more when the domestic spread widens, making cross-bank outcomes an observable implication of the mechanism. However, this gradient is not randomly assigned. Since banks may differ in terms of franchise value, expected support, and business model, a cross-bank return relationship cannot be interpreted as an identified credit-supply elasticity. The thesis uses the gradient as an upstream diagnostic only.

Second, the analysis keeps the valuation-to-credit leg and omits the reverse leg, from bank distress back to sovereign risk. In full models of the nexus, feedback amplifies both the sovereign price and the bank loss (Farhi and Tirole, 2018). Taking the observed spread path as given sacrifices that amplification. However, this approach provides a direct link to exposure, duration and leverage at a given point in time. Therefore, the bank block should be interpreted as a conditional transmission map, and it is no theory of sovereign pricing.

2.3 Sovereign Stress and Credit Supply

Bocola (2016) provides the structural precedent. Sovereign risk weakens banks holding domestic debt, tightens their constraints, and raises borrowing costs. Bofondi et al. (2018) supply the empirical scale. By comparing Italian and foreign banks lending to the same firms during the 2011 crisis, they found that Italian banks charged 15–20 basis points more and lent about three percentage points less. Holding the borrower fixed across lenders isolates a supply-side difference, which is why this estimate anchors the calibration.

It also identifies a firm-credit differential rather than a household one, and the model wedge applies to households. Neri (2013) bridges the two. By estimating the effect of sovereign tensions on euro-area lending rates, he found that the rates on new corporate loans and residential mortgages in peripheral countries would have been approximately 130 and 60 basis points lower at the end of 2011, respectively, had the spreads remained at their April 2010 levels. However, these levels are not comparable with the same-firm differential since they combine the effects of funding, capital, risk and demand without a control group. The ratio between the two is more informative and is used in Appendix C to scale the firm target down to households. This serves as an illustrative same-episode attenuation factor with no claim to identification, because the two lending categories differ in terms of contract structure, borrower composition and transmission channels. Theobald and Tober (2020) reach a similar conclusion for private-sector borrowing costs more generally, finding significant but small long-run effects with wide cross-country variation. Given that dispersion, this thesis reports a range of calibrated wedges.

The consequences extend beyond lending terms. Acharya et al. (2018) show that impairment of banks' sovereign exposures, together with risk-shifting by weakly capitalised banks, cut the probability of a firm receiving a new syndicated loan by up to 53%. They also show that this credit crunch accounts for between 44% and 66% of the fall in investment, employment and sales growth at affected firms. This is the point at which this thesis stops short. It considers household borrowing costs, but does not follow the mechanism through to real outcomes.

Evidence from 2018 itself is supportive but distinct. Sabatini (2022) exploits the same May-2018 yield increase and finds that more exposed banks contracted credit, mainly through capitalisation and not through collateral or liquidity. That result concerns quantities, not the rate the model represents, so it motivates an illustrative attenuation without replacing the rate calibration. Giannetti et al. (2026) show that securities losses impair lending, whether or not they are recognised by accounting rules. This supports the use of an economic-value loss measure here. Grandi (2019) finds that banks with larger sovereign exposures transmit monetary easing less strongly, which matters here because it locates the same exposure gradient outside a crisis window and in the rate margin the model represents.

2.4 Household Heterogeneity and Distributional Consequences

Kaplan et al. (2018) show that liquid wealth, debt and marginal propensities to consume influence both the propagation and incidence of aggregate effects, with transmission mostly occurring through indirect general-equilibrium effects on labour income.

Faccini et al. (2026) are the closest precedent for the mechanism itself. Using Danish consumer-credit evidence, they incorporate consumer credit spreads into a HANK model with household credit supplied by banks and demonstrate that such spreads are distributed unevenly across the wealth spectrum. The borrower wedge of Chapter 3 is that object, with its origin moved upstream to sovereign-induced bank losses. This is a mechanism precedent, not evidence on Italian household pass-through.

Two papers address the question of who bears banking-sector losses. Mendicino et al. (2025) study a two-asset model in which banking losses hit low-income and indebted households through borrowing costs and labour income, while households with sufficient liquidity can purchase assets at depressed prices and benefit from the recovery. In their US calibration, one fifth of households benefit from banking-sector losses. This model-based result motivates the portfolio-rebalancing restriction of the baseline, but it is not comparable to the unpromoted Italian diagnostic in Appendix D.

Roldán (2025) works the saver side. Sovereign risk raises precautionary savings among heterogeneous domestic savers and imposes unequal welfare costs even when default never occurs. The baseline here sets saver-side pass-through to zero on the Italian deposit evidence of Section 3.3, so that channel is switched off and incidence runs entirely through borrowers. Neither paper observes Italian household contracts in 2018.

The empirical counterpart is the survey-based decomposition of Slacalek et al. (2020), who combine euro-area household microdata with estimated aggregate responses to quantify separate transmission channels. Their approach is methodologically closest to the incidence step used in Chapter 5, and their finding that indirect labour-income effects dominate the aggregate consumption response is a caution about the fixed-wage closure adopted here.

Auclert (2019) names what the incidence exercise here does and does not capture. It considers households' unhedged interest-rate exposure in conjunction with heterogeneous marginal propensities to consume, meaning that a rate change is redistributed according to balance-sheet position. The allocation in Chapter 5 keeps the exposure half of that logic and drops the behavioural half. It holds the SHIW vectors fixed and lets no quintile's consumption response feed back.

The measured heterogeneity is also narrower than the model heterogeneity. A household may be asset-rich yet indebted, liquid-poor yet not borrowing, or exposed through contract features that are not visible in a cross-sectional wealth table. Survey aggregates only capture selected balance sheet margins. These are applied to the model after it has been solved using common population weights, quintile boundaries and denominators. This makes the counterfactual internally comparable, though it does not imply that the model recovers the joint distribution of assets, debt, contracts and lenders.

2.5 Model Architecture and Computational Methods

The baseline is a one-liquid-asset heterogeneous-household model with a reduced monetary-union closure. It is not a model of lender choice or portfolio rebalancing, and the fuller two-country structure of Bayer et al. (2024) is deliberately not attempted. Their monetary-union framework motivates treating national exposure as a relevant margin, but it is not a calibration or direct test of this thesis's Italian transmission map.

The stationary household problem is solved by the endogenous-gridpoint method (Carroll, 2006) and the transition by sequence-space Jacobians (Auclert et al., 2021). The income process is discretised with the Rouwenhorst method, whose accuracy for persistent processes is analysed by Kopecky and Suen (2010).

These methods determine which aspects can and cannot be checked numerically. The endogenous-gridpoint method focuses on policy accuracy in the stationary problem, while the sequence-space method focuses on the local derivative around it. Appendix A therefore reports Euler and budget residuals, finite-difference Jacobian checks, grid and horizon comparisons, and an independently coded value-function cross-check. These verify the implementation. However, they do not address whether the calibration is correct, which is a separate issue that requires evidence, not arithmetic.

2.6 Backstops and Monetary-Union Policy

Central bank backstops act upstream of everything described above by limiting the extent to which sovereign spreads can move in the first place. Krishnamurthy et al. (2018) show that OMT and related interventions lowered sovereign yields and bank stress, with effects depending on the distribution of exposures. Rojas and Thaler (2024) argue that domestic sovereign exposure also strengthens governments' incentives to repay, so restricting it involves a trade-off and not a clear gain.

This literature frames the policy discussion in Chapter 6 without identifying anything. A backstop that compresses a spread mechanically reduces every downstream term in the mapping used here. Whether it improves welfare depends on who ends up holding the risk, which is outside the model. The Italy-2022 episode illustrates the logic but cannot estimate the effect of the Transmission Protection Instrument, whose announcement coincided with a 50-basis-point tightening and a domestic political crisis.

2.7 Research Gap and Positioning

Together, the four literatures leave a specific gap because each establishes one link but does not jointly close the chain. Political risk and sovereign pricing stop at the price, saying nothing about whose balance sheet absorbs it. The sovereign–bank nexus carries the price into bank net worth and stops there. Sovereign stress and credit supply carry bank losses into borrowing conditions and stop at the firm, since their sharpest estimates are, by construction, about firms. Household heterogeneity shows who bears a banking-sector loss once it occurs, but it typically starts from a generic shock rather than a specific national event with measurable exposures. The division is between amplification and incidence: sovereign-bank work asks how much a shock is magnified, while household work asks who bears one. The literature reviewed here does not jointly map a dated sovereign repricing through bank balance sheets into the distribution of household exposure.

This thesis constructs a conditional mapping for Italy in 2018. Measured repricing and observed exposure provide a path for bank losses, while an externally calibrated wedge provides the modelled borrower response. Fixed household vectors then allocate this response across wealth quintiles. Each stage retains its own source, date, and permitted interpretation, so a weak diagnostic narrows the claim instead of disappearing into a single number. The contribution is therefore integration and auditability, with the weakest link named. The anchor is a firm-lending differential, and the Neri (2013) ratio is the only same-episode rate-comparison evidence used in this thesis to scale that anchor to the household margin.

CHAPTER 3

From Sovereign Repricing to Household Credit, A Model

This chapter presents a modular model for studying the transmission of sovereign repricing to household credit.[1] When sovereign yields rise, the market value of outstanding government bonds falls. If domestic banks hold this debt, the resulting valuation losses weaken bank net worth and can tighten financial intermediation. This transmission logic builds on Bocola (2016), who models how sovereign risk impairs banks holding domestic government debt through tighter funding constraints and precautionary deleveraging. The framework here is deliberately narrower. Rather than solving the bank's optimisation problem, it retains only the downstream balance sheet and credit transmission mapping required for the present exercise.

The resulting credit wedge enters a heterogeneous household economy and applies specifically to households with negative liquid-asset positions. Faccini et al. (2026) show that consumer-credit spreads have heterogeneous effects on consumption and asset accumulation and can amplify aggregate shocks in a HANK economy with bank-supplied household credit. The mechanism developed here combines this borrower-side transmission with the sovereign-bank origin emphasized by Bocola (2016). As borrower status is endogenous, the model tracks households moving into and out of debt, their responses in terms of consumption and savings, and the resulting evolution of the liquid-asset distribution.

The household block is embedded in a New Keynesian economy representing a small member state of a monetary union. This takes into account country-specific sovereign repricing and represents the rest of the union through a common policy anchor, stopping short of the full two-country HANK setting solved by Bayer et al. (2024) (Section 3.4). The aggregate transition is computed using the sequence-space Jacobian method of Auclert et al. (2021), allowing the borrower-wedge path to propagate through household policies, the evolving distribution, and the remaining equilibrium blocks.

The one-asset heterogeneous-household benchmark is the only specification behind the headline results. It isolates the borrower-credit mechanism under a reduced monetary-union closure. Two richer architectures are reported alongside it but carry no headline number. One is a 25-bank aggregation of the bank side in Chapter 5, the other a two-asset extension with portfolio rebalancing in Appendix D. Each is a check on one restriction of the baseline, and Section 3.6 states what each is and is not allowed to settle.

[1] It combines the sovereign-to-bank balance sheet transmission mechanism (Bocola, 2016) with the HANK framework (Kaplan et al., 2018).

3.1 Model Overview and Scope

Formally, the model starts with an event-related sovereign spread movement measured outside the structural economy. For an event e, let $\Delta s_e^{\text{model}}$ denote the constructed spread input described in Chapter 4. For the subsequent transition, suppressing the event subscript, the model can be summarised as follows.

$$\underbrace{\Delta s_e^{\text{model}}}_{\text{measured event input}} \longrightarrow \underbrace{\{\ell_e^B, N_0^B\}}_{\substack{\text{impact valuation}\\ \text{and bank loss}}} \longrightarrow \underbrace{\{N_t^B, \phi_t\}_{t\geq 0}}_{\substack{\text{bank-net-worth recovery}\\ \text{and credit pass-through}}} \longrightarrow \underbrace{\{c_t(a,z),\, \mu_t,\, C_t,\, Y_t,\, \pi_t\}}_{\substack{\text{endogenous HANK}\\ \text{equilibrium}}}. \tag{3.1}$$

The event input and the upstream-mapping parameters are measured, calibrated, or assigned outside the heterogeneous-agent equilibrium. The reduced-form banking block then generates the impact bank loss and the paths of bank net worth and the borrower wedge. Household policies, the asset-income distribution, and aggregate consumption, output and inflation are determined endogenously. This separation lets the model quantify the downstream consequences of a sovereign repricing without needing to explain the political event or price sovereign risk.

Upon impact, the measured sovereign repricing reduces the market value of banks' domestic government bond portfolios and establishes the initial bank net worth state, N_0^B. Thereafter, at each monthly date t, the inherited state N_t^B determines the borrower wedge ϕ_t. Households choose consumption and savings based on this wedge, while prices and aggregate quantities are jointly determined in general equilibrium. At the end of the period, retained intermediary income and mean reversion determine the next state, N_{t+1}^B through the recovery law of Section 3.2. The structural transition is therefore solved at a monthly frequency. However, borrower wedges are reported on an annualised basis for comparability with the lending-rate evidence used in Section 4.3, while bank net-worth losses are reported as percentages of equity.

Four boundaries follow from this design. The spread movement Δs_t enters as a measured input, so the political event stays outside the model. It enters purely as a valuation shock to existing bond holdings, so default probability plays no role. Transmission runs one way, from sovereign prices to bank net worth to credit, which leaves out the reverse leg of the doom loop discussed in Chapter 2 and the amplification it would add. And households face a single economy-wide wedge ϕ_t and never the balance sheet of a named lender, which puts household–bank matching beyond what the model can speak to.

3.2 Sovereign Repricing, Bank Net Worth, and the Borrower Wedge

Valuation Losses on Sovereign Holdings

Let Δs_e^{bp}, the peak model-input measure $\Delta s_e^{\text{model}}$ of Chapter 4 expressed in basis points, denote the event-associated sovereign-spread increase used as the model input for episode e, and define $\Delta s_e^{dec} = \Delta s_e^{bp}/10{,}000$. Domestic banks hold domestic sovereign securities equal to a fraction ω_e^{dom} of total assets. Let D_e denote the maturity-based proxy for the portfolio's spread duration and

L_e sector leverage, defined as assets over equity. Treating the measured spread widening as an increase in the sovereign-specific discount spread, conditional on the common benchmark yield curve, the first-order economic loss on the sovereign portfolio can be approximated as follows.

$$\ell_e^B = \omega_e^{dom} D_e L_e \, \Delta s_e^{dec}. \tag{3.2}$$

The corresponding impact bank-net-worth ratio is

$$\frac{N_{e,0}^B}{N_{ss}^B} = 1 - \ell_e^B. \tag{3.3}$$

Bank-Net-Worth Dynamics

Equation (3.3) gives only the impact loss. Bank net worth then evolves according to a reduced-form monthly recovery law,

$$N_{t+1}^B = \min\left\{N_{ss}^B,\ N_t^B + \Delta t\left[\theta_B \phi_t^{\text{ann}} \bar{K}^B + (\delta_B + \iota)(N_{ss}^B - N_t^B)\right]\right\}, \tag{3.4}$$

where $\Delta t = 1/12$, ϕ_t^{ann} is the annualised borrower wedge, and $\bar{K}^B$ is the steady-state annual intermediation scale. Thus, the bracket is an annual net-worth flow and Δt converts it into a monthly increment. The term $\theta_B \phi_t^{\text{ann}} \bar{K}^B$ is a reduced-form retained-income term, allowing part of the borrower wedge earned on the steady-state intermediation scale to rebuild bank net worth.[2] The term $(\delta_B + \iota)(N_{ss}^B - N_t^B)$ captures mean reversion and recapitalisation toward steady-state net worth. Because δ_B and ι enter only through their sum, they are not separately identified by this law. The minimum operator prevents mechanical overshooting above the steady-state bank-capital level.

Pass-Through to Household Credit

The impairment of bank capital tightens credit conditions for borrowers. Banks can absorb a capital loss through higher rates, smaller volumes, or tighter non-price terms. The model represents only the rate margin, since that is what enters household intertemporal choice. It captures this margin through an asymmetric credit wedge,

$$\phi_t^{\text{ann}} = \phi_0 \max\left[\left(\frac{N_{ss}^B}{N_t^B}\right)^{\phi_1} - 1,\ 0\right], \tag{3.5}$$

The annualised wedge is calibrated to match the Italy-2011 bank-lending evidence of Bofondi et al. (2018) and held fixed when the model is applied to later events. The curvature parameter ϕ_1 is set to one in the baseline model. This is a tractability assumption, imposed and not estimated, with no literature source behind it. The max operator makes the pass-through one-sided. A net-worth shortfall raises borrowing costs, whereas the model never generates a corresponding borrower subsidy.

Equation (3.3) measures a sovereign-specific spread-induced valuation component, conditional

[2] The logic is that of the macroeconomic banking literature (e.g., Gerali et al., 2010), where banks rebuild capital from retained earnings and wider loan margins when capital is scarce. The formulation here abstracts from a full dynamic dividend-payout problem.

on the common benchmark yield curve. It is a counterfactual country-specific revaluation, not the observed change in the bond's own market price, which also moves with the benchmark leg (Appendix C quantifies the difference for France 2024). Separately, whether any such valuation change is translated into reported accounting income or regulatory capital depends on the classification of sovereign holdings (e.g. amortised cost versus fair value) and the application of relevant prudential filters. These factors can delay or reduce the recognition of the same underlying loss. This thesis uses the spread-induced economic-value measure because it can affect funding and balance-sheet capacity even when accounting recognition is delayed.

3.3 Heterogeneous Households

Preferences, Income Risk, and Asset Positions

The household side is a continuum of infinitely lived households facing uninsurable idiosyncratic labour-income risk. Preferences are constant-relative-risk-aversion,

$$u(c) = \frac{c^{1-\gamma}}{1-\gamma}, \tag{3.6}$$

where γ is the coefficient of relative risk aversion. Each household draws an idiosyncratic productivity state z_t from a finite-state Markov chain obtained by discretising a persistent income process with the Rouwenhorst method. Kopecky and Suen (2010) provide the relevant finite-state approximation analysis. Labour income is $y_t = w_t z_t$. The discount factor β is internally calibrated, not fixed exogenously. It is solved jointly with steady-state output so that the stationary equilibrium clears the goods market and matches the targeted aggregate liquid-asset stock. The exact calibration is detailed in Section 3.5.

Households save or borrow using a single uncontingent liquid asset a_t, subject to the borrowing constraint $a_{t+1} \geq \underline{a}$. Whether a household is a saver ($a_t \geq 0$) or a borrower ($a_t < 0$) is an endogenous outcome of its history of income shocks and consumption choices, not a fixed household type. The same household can therefore move between the two regions over time. This endogenous composition of the borrower pool is precisely what the fixed borrower share of the two-agent benchmark in Section 3.6 abstracts from.

The Household Problem

The gross return on the household's liquid position depends on the sign of its current holdings, a_t, not on its next-period choice. The annualised wedge becomes the monthly wedge $\phi_t^m = (1+\phi_t^{\mathrm{ann}})^{1/12} - 1$ before it enters the household problem. Borrowers therefore pay $r_t + \phi_t^m$. Savers earn $r_t - \chi\phi_t^m$, where $\chi \in [0, 1]$ is the share of the wedge passed through to the rate on positive liquid balances. The baseline sets $\chi = 0$ on direct evidence, and not for tractability. Albertazzi et al. (2014) estimate the effect of the BTP–Bund spread on Italian banks' own funding rates during sovereign stress and find a sizeable effect on term deposits and newly issued bonds, but "virtually no effect on overnight deposits," the closest empirical counterpart to this model's single liquid asset. The sign of that limited pass-through is a rise in what banks pay, not a fall. A wedge that lowers what savers earn is therefore not supported by the evidence for this specific margin and this specific episode. $\chi = 0.25$ is retained as a declared closure sensitivity in Appendix C, not the baseline. Expressed as $R_t(a_t) = 1 + r_t + \mathbf{1}_{\{a_t<0\}}\phi_t^m - \mathbf{1}_{\{a_t\geq 0\}}\chi\phi_t^m$, the household budget constraint is

$$c_t + a_{t+1} = R_t(a_t)\, a_t + w_t z_t. \tag{3.7}$$

The household problem carries no transfer or lump-sum tax term. What the bank extracts from borrowers stays inside the bank block through the $\theta_B \phi_t^{\text{ann}} \bar{K}^B$ term in the recovery law (3.4), so wedge income leaves the modelled household sector entirely. This fits a small member state whose banking sector is partly foreign-owned and whose intermediation income partly accrues abroad. Because a_t is negative for borrowers and non-negative for savers, the wedge enters resources with the same sign as the position it applies to.

$$\frac{\partial\, [R_t(a_t)a_t]}{\partial \phi_t^m} = \mathbf{1}_{\{a_t<0\}}\, a_t - \mathbf{1}_{\{a_t\geq 0\}}\, \chi\, a_t, \tag{3.8}$$

which is negative for borrowers and, at the baseline $\chi = 0$, exactly zero for savers. A larger wedge reduces resources only for households with negative liquid assets. This borrower-only exposure mirrors the consumer-credit margin emphasised by Faccini et al. (2026), with ϕ_t here originating in sovereign-induced bank-net-worth impairment. The household-incidence allocation in Section 5.1 applies the wedge to household debt only, which follows directly from this budget constraint.

Households choose consumption and next-period assets to maximise

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t u(c_t) \tag{3.9}$$

subject to equation (3.7) and $a_{t+1} \geq \underline{a}$. Away from the borrowing constraint and the zero-asset threshold, the associated Euler equation is

$$u'(c_t) = \beta\, \mathbb{E}_t \left[R_{t+1}(a_{t+1})\, u'(c_{t+1}) \right]. \tag{3.10}$$

At $a_{t+1} = \underline{a}$ and $a_{t+1} = 0$, the corresponding one-sided optimality conditions apply. Because borrowers and savers face different marginal rates, the payoff schedule $R_{t+1}(a_{t+1})a_{t+1}$ is continuous but non-differentiable at zero. Uninsurable income risk and the borrowing constraint generate the usual buffer-stock saving motive. As ϕ_t rises, current borrowers face an immediate resource loss, while households close to zero have an additional incentive to deleverage or remain on the saving side of that threshold.

Distribution Dynamics and Aggregation

The optimal policies $c_t(a, z)$ and $a_{t+1}(a, z)$ determine the evolution of the joint distribution of asset and income states $\mu_t(a, z)$. Applying today's asset policy to today's distribution yields tomorrow's.

$$\mu_{t+1}(A, Z) = \int \mathbf{1}_{\{a_{t+1}(a,z)\in A\}}\, P(z' \in Z \mid z)\, d\mu_t(a, z). \tag{3.11}$$

Aggregate consumption and assets are $C_t = \int c_t(a, z)\, d\mu_t(a, z)$ and $A_t = \int a\, d\mu_t(a, z)$, while the borrower mass is $\lambda_t^B = \int \mathbf{1}_{\{a<0\}}\, d\mu_t(a, z)$. A representative-agent economy can track the first two aggregates, but not the borrower mass or the distribution of exposure to ϕ_t. The evolution of μ_t

therefore determines who faces the credit shock, how borrower status shifts over the transition, and how the resulting losses are ultimately distributed.

3.4 Aggregate Economy and Common Monetary Policy

New Keynesian Supply and Goods-Market Block

The household block is embedded in a standard New Keynesian aggregate closure with Rotemberg price adjustment and a common monetary policy anchor. The economy is modelled as a small member state. Domestic household demand, output and inflation are calculated for one member state, while the rest of the currency union is represented by the common policy anchor, with no second country solved. Inflation follows a New Keynesian Phillips curve,

$$0 = \kappa \left(\frac{w_t}{Z_t} - \frac{1}{\mu} \right) + \frac{Y_{t+1}}{Y_t} \frac{\ln(1 + \pi_{t+1})}{1 + r_{t+1}} - \ln(1 + \pi_t), \tag{3.12}$$

with a companion goods-market condition, $\mathcal{G}_t = 0$, that embeds the standard Rotemberg-style quadratic price adjustment cost,

$$\mathcal{G}_t \equiv Y_t - C_t - \frac{\mu}{\mu - 1} \frac{1}{2\kappa} [\ln(1 + \pi_t)]^2 Y_t. \tag{3.13}$$

The sovereign disturbance does not affect either condition directly. Its only input into the aggregate economy is the borrower wedge ϕ_t.

Common Monetary Policy

The monetary block is a simplified domestic rule tied to a union-wide policy rate, not a literal description of how the ECB actually sets interest rates.[3] Let all rates be net monthly rates, and let $i_t = r_t^* + \phi_\pi \pi_t$ denote the stylised domestic nominal-rate component. Combining its lagged implementation with the Fisher relation gives

$$1 + r_t = \frac{1 + r_{t-1}^* + \phi_\pi \pi_{t-1}}{1 + \pi_t}. \tag{3.14}$$

The anchor r_t^* is taken as union wide and is not directly affected by the domestic sovereign event. Domestic inflation nevertheless affects the real rate and, under the maintained reduced-form feedback rule, the nominal-rate component as well.

Single Country Aggregate Closure

The two variables, aggregate output and inflation $\{Y_t, \pi_t\}$, are solved jointly against the goods-market and Phillips-curve residuals. Throughout the transition, the baseline holds the real wage fixed at its steady-state value, $w_{ss} = Z/\mu$, and technology is fixed as well. Consequently, the Phillips-curve cost-push term is zero along the baseline path. Inflation and the real rate remain at their steady-state values, and the reported aggregate response is generated by the direct borrower-cost wedge alone, with no wage, price, or monetary-policy feedback behind it.

[3] A reaction function specifies how a central bank sets its policy rate in response to conditions such as inflation and output. The Taylor rule is the standard example. What follows is a stylised domestic analogue of one, not an estimate of actual ECB policy.

This deliberately isolates the borrower channel. It differs from the indirect wage-income channel emphasised by Kaplan et al. (2018).

This single-country closure is motivated by Equiza-Goñi et al. (2023), who demonstrate that government debt portfolios in a currency union absorb aggregate fiscal risk far more effectively than country-specific fiscal risk. This motivates a national balance-sheet channel that currency-union membership does not automatically neutralise.

3.5 Equilibrium and Solution Method

Stationary Equilibrium

The stationary equilibrium prior to any political event consists of household policies $c_{ss}(a, z)$ and $a'_{ss}(a, z)$, a stationary distribution $\mu_{ss}(a, z)$, a zero borrower wedge, and aggregate prices and quantities consistent with household optimisation and market clearing. Because $\phi_{ss} = 0$, the level N^B_{ss} is the reference net worth throughout the chapter. The stationary distribution is the fixed point of the law of motion given in Section 3.3 evaluated at constant policies,

$$\mu_{ss}(A, Z) = \int \mathbf{1}_{\{a'_{ss}(a,z)\in A\}}\, P(z' \in Z \mid z)\, d\mu_{ss}(a, z). \tag{3.15}$$

At this stage, the discount factor β and steady-state output Y_{ss} are solved jointly, chosen so that the goods market clears and the aggregate liquid-asset stock matches the calibrated bond-supply target B,

$$A_{ss} = B, \qquad \mathcal{G}_{ss} = 0. \tag{3.16}$$

The resulting stationary distribution determines both the borrower mass and the mass of households at the borrowing constraint. Section 4.3 calibrates the former against SHIW bank debt participation, while the latter remains an untargeted model outcome.

Event-Driven Transition Equilibrium

Conditional on the measured event, the bank block generates the paths $\{N^B_t, \phi_t\}$. These paths are taken as predetermined inputs by the downstream heterogeneous-agent equilibrium. Under the adopted small-member-state closure, the endogenous aggregate sequences are output and inflation, $\{Y_t, \pi_t\}$. Aggregate liquid assets A_t are recovered from household behaviour and never imposed as an additional transition target. The baseline is therefore a reduced member-state equilibrium, not a closed-economy transition with period-by-period asset-market clearing. The rest of the monetary union is the implicit external counterpart to a change in the household sector's desired liquid-asset position. The deviation from the stationary asset-supply target is reported as a diagnostic in Appendix A.

Sequence-Space Solution

The stationary household problem is solved using the endogenous-gridpoint method of Carroll (2006). Given next-period marginal values, the Euler equation is evaluated on a fixed grid of end of period asset positions to obtain current consumption and the associated endogenous cash on hand points, after which the policy functions are interpolated back onto the fixed state grid.

Around the stationary equilibrium, the sequence-space method of Auclert et al. (2021) computes, for each aggregate household outcome, the matrix of derivatives with respect to small perturbations in the sequences of prices and the wedge it faces. For aggregate consumption,

$$d\mathbf{C} = J^{C,w}\, d\mathbf{w} + J^{C,r}\, d\mathbf{r} + J^{C,\phi}\, d\boldsymbol{\phi}, \tag{3.17}$$

with no corresponding transfer Jacobian, since transfers are not an input to the household problem. $J^{C,r}$ captures the response to the common real rate through intertemporal substitution and households' existing liquid-asset positions. $J^{C,\phi}$ isolates the borrower-specific response to the bank credit wedge. $J^{C,w}$ is included for completeness even though $d\mathbf{w}$ is identically zero along the baseline transition.

Conditional on the event, the bank block is solved upstream to produce $d\boldsymbol{\phi}$. The household Jacobians are then composed with the analytically differentiated New Keynesian closure to form the aggregate equilibrium system $H(U, Z) = 0$, where $U = \{Y_t, \pi_t\}$ and Z contains the borrower-wedge sequence. Linearising around the steady state, $H_U\, dU + H_Z\, dZ = 0$. Provided H_U is nonsingular,

$$dU = -H_U^{-1} H_Z\, dZ. \tag{3.18}$$

Once dU is recovered, the household Jacobians deliver the corresponding first-order responses of consumption, liquid assets, borrower status, and the distributional objects used in Chapter 5.

3.6 Benchmarks and Extensions

The headline specification is the one-asset HANK model because the wedge makes household balance-sheet heterogeneity an active state variable. Two simpler benchmarks illustrate this. The removal of the coexistence of borrowers and savers by a representative agent means that the wedge has no distributional content unless debt is imposed by assumption. A two-agent benchmark restores a fixed borrower–saver split and demonstrates that borrower losses do not necessarily need to be offset by saver consumption. However, it fixes the size of the borrower group. Only the heterogeneous-household version allows households to move into and out of debt, meaning that both the proportion of households exposed to the wedge and the extent to which they are affected by it respond to the shock. Therefore, RANK and TANK are interpretive benchmarks, not competing specifications (Kaplan et al., 2018).

The bank side admits a parallel comparison. Chapter 5 reports a 25-bank aggregation alongside the representative-bank mapping, holding the household block fixed to isolate how much of the headline magnitude is an aggregation choice. The household side admits a second comparison. Appendix D reports a two-asset extension with portfolio rebalancing that has been solved and passes its own numerical validation checks, but has not been adopted. Its portfolio price is a maintained proxy. Its channel decomposition was withdrawn pending re-evaluation, and its French counterpart remains unsolved. It remains to be seen whether the incidence ranking below survives when households are given something to rebalance, and it settles no magnitude. Section 6.4 returns to this issue. Appendix A documents the numerical scope of the baseline.

CHAPTER 4

Political Episodes and Calibration

This chapter builds the sovereign-repricing input the model conditions on, and calibrates the map from that input to the borrower wedge. It produces four measured movements. Italy 2011 moves 350 basis points, Italy 2018 moves 166.1, France 2024 moves 34.1, and Italy 2022 moves 15.5. Each also carries the exposure, duration, and leverage that turn it into a bank loss, along with the pass-through scale $\phi_0 = 0.003812$ inverted from the Italy-2011 lending evidence.

Political events determine the timeframe in which market data measures movement. As observed spreads also reflect fiscal news, redenomination and liquidity risk, common rate expectations and movements in the German benchmark, these phenomena are referred to as 'event-associated sovereign-spread movements' and never as 'structural political shocks'.

4.1 Measuring Sovereign Repricing

The bank-loss mapping of Chapter 3 requires the size of the sovereign repricing in yield terms, and no political-event indicator or uncertainty index will do, because the valuation loss on banks' existing government-bond portfolios is proportional to the change in yields. Political chronology is therefore used first to identify economically meaningful episodes, following the narrative-event approach of Handler and Jankowitsch (2025) and Hazell and Hobler (2025), and market data are then used to measure the associated repricing.

For country c, sovereign stress is measured relative to the German benchmark as

$$s_{c,t} = y_{c,t}^{10y} - y_{DE,t}^{10y}, \tag{4.1}$$

where $y_{c,t}^{10y}$ and $y_{DE,t}^{10y}$ denote the national and German ten-year government bond yields, respectively. The ten-year maturity is the standard long-term sovereign benchmark and Germany provides the euro area reference yield. This national versus Bund convention is widely used to represent relative sovereign financing conditions in the euro area (Afonso et al., 2015). The analysis distinguishes between two movements in the spread because the market validation exercise and the structural model require different timeframes.

The first is the *acute market movement*,

$$\Delta s_e^{\text{acute}} = s_{c,t_e^1} - s_{c,t_e^0}, \tag{4.2}$$

measured over a narrow window surrounding the most clearly dated repricing. This is the relevant object for the bank-equity and CDS evidence in Chapter 5, where shorter windows reduce contamination from subsequent news.

The second is the *model input*,

$$\Delta s_e^{\text{model}} = \max_{t \in \mathcal{H}_e} \left(s_{c,t} - \bar{s}_{c,\mathcal{B}_e} \right), \tag{4.3}$$

defined as the peak spread deviation from the pre-specified pre-event baseline $\mathcal{B}_e$ over the broader episode horizon $\mathcal{H}_e$. This is the stress measure passed to the bank-loss mapping in Chapter 3. The peak operator is intentional. The exercise asks how the downstream mechanism responds to the maximum sovereign deterioration associated with the episode, not what average causal effect can be attributed to a single political announcement. For Italy 2018 and Italy 2022/TPI, $\Delta s_e^{\text{acute}}$ and $\Delta s_e^{\text{model}}$ are measured from different sources. The acute figure uses daily market data, while the model input uses a monthly cross-check series with fuller pre-event coverage, because the daily Italian series does not extend far enough back to support that baseline. France 2024 uses a single daily series for both objects. Appendix B documents the exact source and vintage underlying each number.

A national spread can widen because the domestic yield rises, because the Bund yield falls, or both, and broader windows may contain additional fiscal or monetary-policy information. Chronology disciplines the window. It does not partition what happens inside it.

Figure 4.1 makes the two objects concrete for Italy 2018. The acute window and the broader episode horizon are different slices of the same spread series, measured against the same pre-specified baseline.

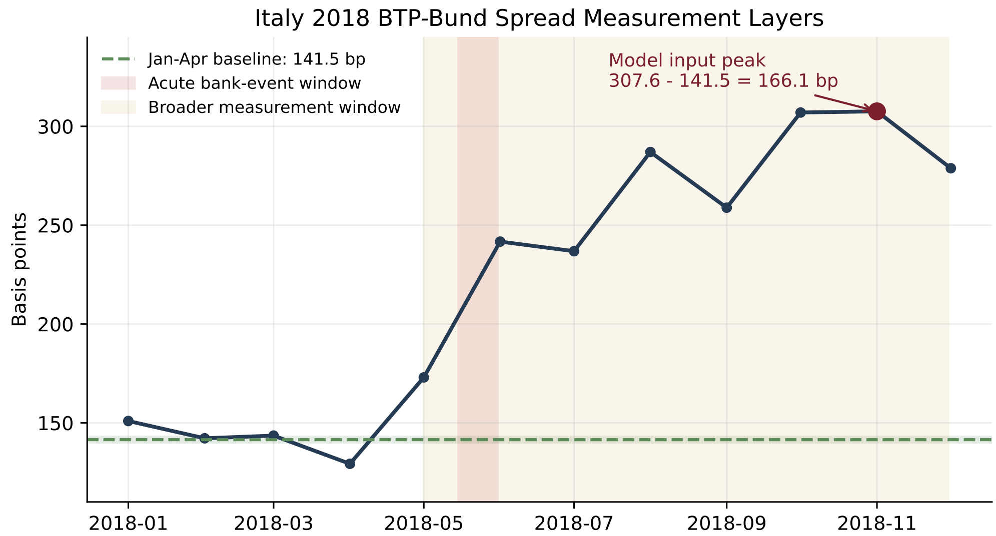

Figure 4.1. Italy 2018 BTP–Bund spread, monthly cross-check series. The dashed green line and band mark the January–April 2018 baseline. The red band marks the acute bank-event window (15–31 May). The gold band marks the broader episode horizon used to register the model input. The marked peak is the maximum deviation from baseline over that horizon, i.e. $\Delta s_{2018}^{\text{model}}$.

4.2 Political Episodes and Window Selection

The four episodes are selected for different purposes, and none of them is treated as a repeated observation of a common political shock. The design deliberately separates the episode used to calibrate the bank-to-credit mapping from the episode used as the headline application, while the remaining cases test whether the same mechanism scales with the size of sovereign repricing and domestic-bank exposure or changes under a different policy environment. Episodes are retained when the underlying political or fiscal information can be dated with reasonable precision, is plausibly relevant for sovereign valuation, produces an observable market response, and can be matched to sufficiently close balance-sheet data.

Four episodes anchor the exercise.

- **Italy 2011 supplies the external calibration.** The acute phase of the Italian sovereign-debt crisis in **autumn 2011** provides the external credit-supply episode used to calibrate the bank-to-credit pass-through. Rather than reconstructing a new event series, the thesis registers a literature-triangulated spread movement of approximately **350 basis points** over the crisis environment underlying Bofondi et al. (2018). The episode is used only to discipline ϕ_0. It is not a held-out application, and the 350-basis-point value is not presented as a newly identified political shock.

- **Italy 2018 is the flagship held-out application.** Following the inconclusive March general election, sovereign stress intensified as a Five Star Movement–League government became increasingly likely and its fiscal programme became clearer. The most abrupt repricing occurred in late May during the confrontation surrounding government formation. The acute validation window is therefore **15–31 May 2018**, over which the BTP–Bund spread widens by **114.4 basis points**. This is the window used for the bank-equity and CDS evidence. Sabatini (2022) independently identifies the sudden May-2018 rise in Italian yields and shows that banks with greater sovereign exposure subsequently contracted credit more, primarily through the capitalisation channel. The government formed at the end of May, but sovereign stress remained elevated during the subsequent budget confrontation. For the structural model, the pre-event benchmark is the **January–April 2018** average and the peak is reached in **November**, yielding $\Delta s_{2018}^{\text{model}} =$**166.1 basis points**.

- **France 2024 is the low-exposure comparison.** The dissolution of the National Assembly, announced on **9 June 2024**, and the ensuing two-round legislative election provide a comparison in which both sovereign repricing and domestic-bank exposure are much smaller than in Italy 2018. The market-validation window runs from **7 June to 8 July 2024** and produces an acute OAT–Bund widening of **14.9 basis points**. Relative to the pre-9-June baseline, the broader peak used by the model is $\Delta s_{2024}^{\text{model}} =$**34.1 basis points**. The episode is therefore used to test whether the fixed transmission mapping appropriately scales down in a lower-stress, lower-exposure banking system. Movements in the German safe-haven leg are not separately identified.

- **Italy 2022/TPI is a policy contrast.** The summer of 2022 combines domestic political instability with a change in the euro-area policy regime. The broad market window runs from **20 July to 26 September 2022** and shows a **31.4-basis-point** spread increase. For the structural comparison, the baseline is instead fixed at the **June 2022** spread level, after the initial move toward monetary-policy normalisation, and the subsequent October peak implies $\Delta s_{2022}^{\text{model}}$ =**15.5 basis points**, smaller than the acute figure above because the baseline here is fixed after the initial spread widening instead of before it, unlike Italy 2018 and France 2024. This episode is explicitly noncausal because the Governing Council approved the TPI on the same day that it raised the three key policy rates by 50 basis points (European Central Bank, 2022), and the Italian political transition unfolded over the same period. The episode therefore cannot isolate the causal effect of the backstop.

Figure 4.2 plots the underlying national and German ten-year yield series for all four episodes. Italy 2011 is the literature-triangulated external calibration episode and not a directly measured input. The other three movements are measured continuously from the same source construction.

The thesis is based on six data blocks, each documented in its own section of Appendix B. These cover the sample architecture and processing conventions (Appendix B.1), the sovereign-yield series (Appendix B.3), the country-level bank balance-sheet inputs (Appendix B.4), the separate bank-level validation panel (Appendix B.5), the household microdata (Appendix B.6), and the macro-financial control series (Appendix B.7). Section 4.3 uses these objects to discipline the bank-to-credit mapping, while Chapter 5 reports the resulting market evidence and household effects.

4.3 Calibration

This section uses the sovereign-spread and balance-sheet inputs constructed above to map sovereign repricing onto bank net-worth impairment and ultimately onto the borrower credit wedge. The bank-to-credit pass-through is calibrated externally on Italy 2011 while excluding Italy 2018. The ability of the resulting mechanism to leave the expected cross-sectional signature in market prices is tested separately, alongside the model's quantitative output, in Section 5.3. Robustness exercises and alternative calibration anchors are reported in Appendix C.

4.3.1 Calibration Architecture and Parameter Status

The quantitative exercise combines parameters that play different conceptual roles, so they should not be interpreted as having the same empirical status. Some objects are calibrated to stationary household moments. One key parameter is externally disciplined by credit supply evidence. The remaining parameters maintain structural or numerical values. The most important of these categories are those relating to the bank-to-credit channel, where the household steady state is calibrated without reference to political episodes. In this channel, the scale of the borrower wedge is anchored to independent Italy-2011 lending evidence and is then held fixed for the Italy-2018 application.

The household calibration is deliberately parsimonious. The borrowing limit $\underline{a}$ disciplines the steady-state borrower share, while the discount factor β and steady-state output are solved jointly so that the liquid-asset target and goods-market condition are satisfied. By contrast, the pass-through parameter ϕ_0 is not chosen to match a household moment or an Italy-2018 outcome. It is obtained from the external calibration developed in Section 4.3.3. Parameters governing preferences, the income process, nominal adjustment, monetary-policy feedback, and bank-net-

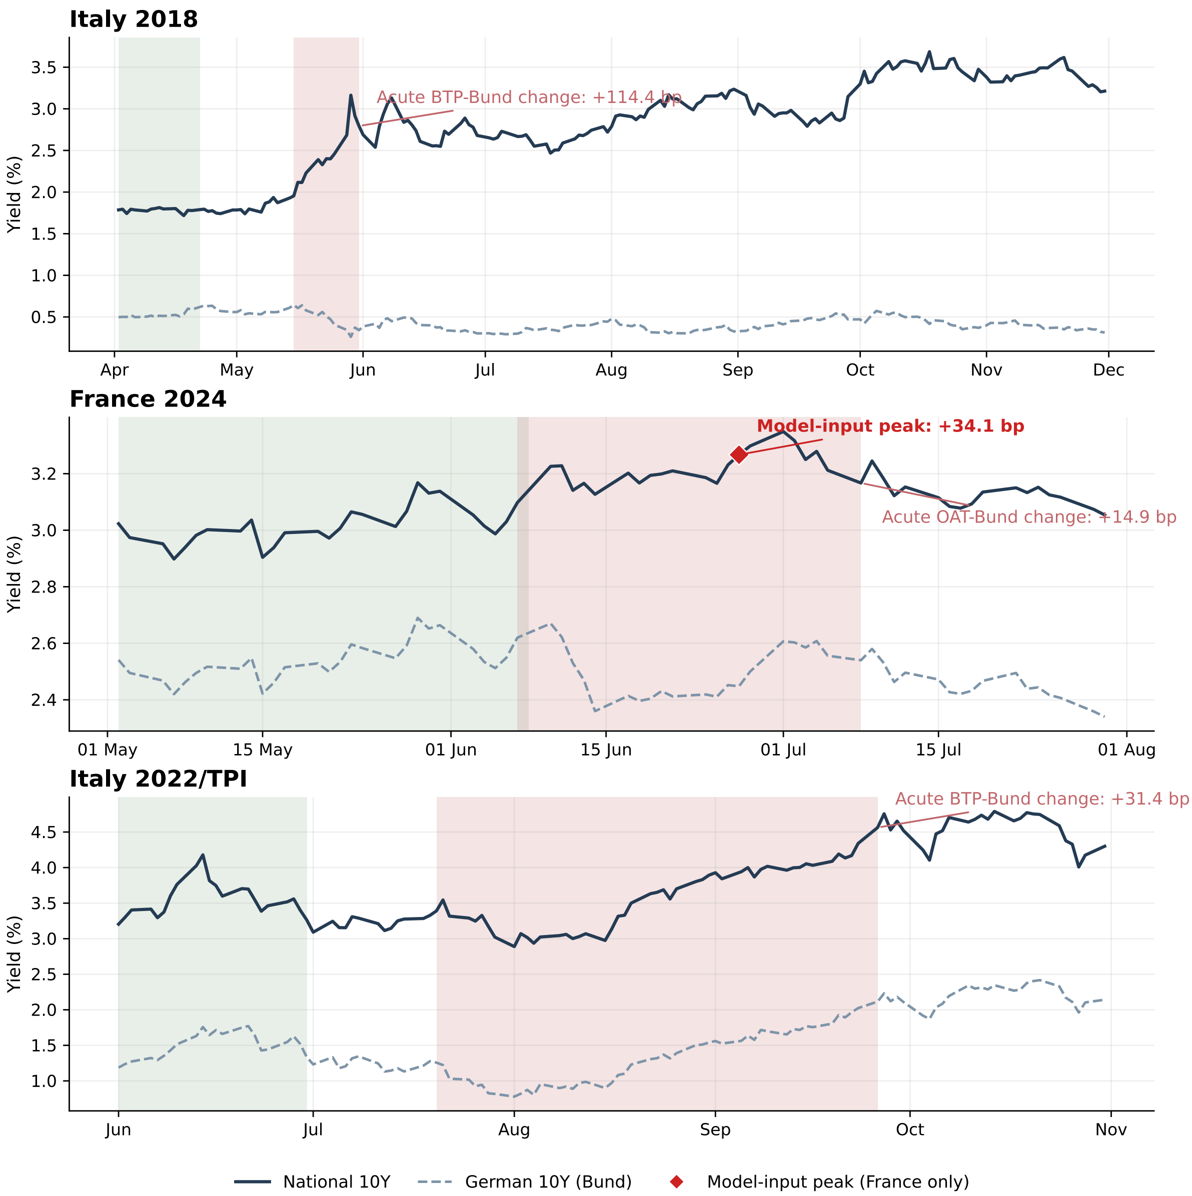

Figure 4.2. National and German ten-year sovereign yields around each episode, daily. Green shading marks the available daily pre-event reference interval and red shading the acute window. Red labels report the BTP–Bund or OAT–Bund change between the acute-window endpoints. The filled dark-red diamond marks the broader-window France model-input peak. Italy's model-input peaks are computed from separate monthly baseline series with fuller pre-event coverage than the daily series shown here, so they are not read from the plotted daily endpoints.

Parameter	Interpretation	Value	Discipline
Household calibration			
$\underline{a}$	Borrowing limit	-0.5	SHIW borrower target
β	Discount factor	0.9800	Joint steady-state solve[a]
r_{ss}	Steady-state real rate (monthly)	0.50%	Fixed calibration input[d]
B	Liquid-asset supply target	4.00	Maintained target
Bank-to-credit transmission			
ϕ_0	Externally calibrated reduced-form bank-stress-to-borrower-wedge scale	0.00381	Italy-2011 inversion[b]
Maintained structural parameters			
γ	Risk aversion	2	Maintained
ρ_e	Income persistence	0.966	Income process
σ_e	Income innovation scale	0.5	Income process
ϕ_1	Wedge curvature	1	Linear benchmark
κ	Phillips-curve slope	0.1	Maintained
ϕ_π	Inflation response	1.5	Taylor principle
μ	Price markup	1.2	Maintained
Bank recovery			
$\delta_B + \iota$	Net-worth recovery	0.1049	Gerali et al. (2010)[c]
θ_B	Retained-income loading	0.5	Maintained

[a] β and steady-state output are solved jointly so that $A_{ss} = B$ and the goods market clears.
[b] ϕ_0 is an externally calibrated reduced-form bank-stress-to-borrower-wedge scale. It inverts the Italy-2011 17.5-basis-point same-firm lending-rate target together with the measured bank-loss inputs. It is not a parameter directly estimated by Bofondi et al. (2018), whose evidence concerns firm lending, not observed household contracts.
[c] The value 0.1049 is numerically anchored to the bank-capital adjustment calibration in Gerali et al. (2010). It is not estimated there as the recovery rate appearing in equation (3.4). Sensitivity is reported in Chapter 5.
[d] $r_{ss} = 0.005$ monthly (6.17% annualised, compounding exactly) is a fixed calibration input, not a value solved jointly with β. It is the same steady-state rate that anchors the union-wide monetary rule r_t^* of Section 3.4. It is higher than typical real-rate calibrations because it is fixed to support the household liquid-asset and borrowing-participation targets, with no attempt to match an observed risk-free rate.

Table 4.1. Baseline calibration and parameter discipline.

worth recovery complete the model without adding further event-specific targets. Table 4.1 summarises these different parameter disciplines.

The resulting household steady state reproduces the targeted borrower margin closely. 24.24% of households hold negative liquid assets, compared with 24.91% reporting bank debt in SHIW 2022. Aggregate liquid assets equal $B = 4$ by construction. The 12.49% mass at the borrowing constraint is not targeted and therefore provides an additional model outcome, alongside untargeted liquid-wealth concentration and debt-conditional-on-borrowing moments and a check that the household composition is not an artefact of the baseline real-rate calibration (Appendix A.2). The non-harmonised cross-country comparisons are reported in Appendix B.

4.3.2 Economic Losses on Bank Sovereign Portfolios

For each episode e, the country-level economic loss on domestic banks' sovereign portfolio is computed using the duration approximation introduced in Chapter 3. Let Δs_e^{bp} denote the measured spread movement in basis points, ω_e^{dom} the domestic-sovereign holdings of the banking sector as a share of total assets, D_e the sovereign-portfolio duration in years, and L_e leverage, defined as assets over equity. The equity-side economic loss is

$$\ell_e^B = \omega_e^{dom} D_e L_e \frac{\Delta s_e^{bp}}{10{,}000}. \tag{4.4}$$

The exposure × duration × leverage mapping transforms the observed yield movement into an economic loss relative to equity.[1] The impact bank-net-worth ratio entering the structural model is therefore $N_{e,0}^B / N_{ss}^B = 1 - \ell_e^B$.

The three inputs to equation (4.4) are not measured from the same institutional population. The exposure share ω_e^{dom} is drawn from ECB Balance Sheet Items, a sector-wide monetary-financial-institution aggregate, because the EBA panel of large, internationally diversified banks understates domestic sovereign concentration relative to the banking sector as a whole and inverts the France–Italy exposure ranking central to the thesis's cross-event comparison (Appendix B). Duration D_e and leverage L_e, by contrast, are EBA Transparency Exercise aggregates, because sector-wide duration and leverage are not published at comparable vintages. The resulting country-level loss is therefore a composite calibration across two different institutional coverages, not a measurement on a single observed representative bank. It is retained because no single source spans exposure, duration, and leverage at the sector level, and using EBA exposure throughout would distort the cross-country comparison the thesis relies on.

The pre-event bank-market check uses a related but bank-level predicted loss. For bank i, let $H_{i,e}^{dom}$ denote domestic-sovereign exposure, D_e the event-level duration, Δs_e^{bp} the event spread movement, and $CET1_{i,e}$ pre-event Common Equity Tier 1 (CET1) capital. The predicted economic loss relative to CET1 is

$$\widehat{\ell}_{i,e}^{CET1} = \frac{H_{i,e}^{dom} D_e \Delta s_e^{bp} / 10{,}000}{CET1_{i,e}}. \tag{4.5}$$

[1] For the large Italy-2011 movement, first-order duration may overstate the loss because convexity is omitted. Appendix C.4 reports the corresponding sensitivity.

The country measure enters the structural model, while the bank measure enters the validation regression. Both are economic-value measures. Alternative accounting-recognition assumptions are reported in Appendix C.4.

4.3.3 External Calibration of the Bank-to-Credit Pass-Through

The bank-to-credit mapping is disciplined using evidence from the Italian sovereign debt crisis of 2011, as in Bofondi et al. (2018). Their same-firm design compares Italian and foreign banks lending to the same firms, controlling for fixed effects by firm and period. The rate differential is approximately 15–20 basis points, with a midpoint of 17.5 basis points serving as the target. This is a discrete domestic-versus-foreign treatment effect, not a direct estimate of ϕ_0.

With the benchmark curvature $\phi_1 = 1$, the borrower wedge on impact is

$$\phi_e = \phi_0 \left(\frac{1}{1 - \ell_e^B} - 1 \right), \tag{4.6}$$

where ϕ_e is expressed in decimal units. Calibrating this pass-through requires combining the targeted rate response with the specific balance-sheet inputs from the 2011 crisis. During that episode, the Italian banking sector held a domestic-sovereign share of $\omega_{2011}^{dom} = 0.0811$, with an average portfolio duration of $D_{2011} = 4.397$ years and a leverage ratio of $L_{2011} = 25.21$. When exposed to the $\Delta s_{2011} = 350$ basis-point spread increase, these balance-sheet conditions translate into a representative net-worth loss of $\ell_{2011}^B = 0.3146$ via equation (4.4). To rationalise the targeted 17.5-basis-point lending-rate increase against this specific capital impairment, equation (4.6) is inverted, which yields

$$\phi_0 = \frac{17.5/10{,}000}{(1/(1 - 0.3146)) - 1} \simeq 0.00381. \tag{4.7}$$

The leverage input $L_{2011} = 25.21$ divides by Core Tier 1 capital, the concept reported in the EBA-2011 disclosure. The application episodes instead divide by Tier 1 capital under later EBA templates. Harmonising the 2011 capital denominator to Tier 1 (EUR 97.2 billion against Core Tier 1's EUR 80.2 billion) raises ϕ_0 by roughly 31%, but the corresponding numerator cannot be harmonised because Basel III leverage exposure was not reported in 2011 (Appendix C.4). The baseline retains the as-calibrated pairing and treats this as an unresolved calibration limitation, which is not evidence that the headline is conservatively or aggressively biased.

The structural model is solved monthly, but the calibration target and reported wedges are annualised basis points, matching the units in which lending-rate evidence is usually stated.[2]

Italy 2018 is held out from the inversion. With the benchmark ϕ_0 and $\phi_1 = 1$, the June-2018 episode-vintage proxy produces 6.56 basis points (Chapter 5). It is a reproducible comparison point, not a pre-event estimate. The dated Italy-2018 timing table instead reports partial-pre-event bridges using December-2017 exposure and leverage. No strictly pre-event numerical bridge is claimed because the required December-2017 duration is unavailable. For any curvature

[2]Wedges are reported in annualised basis points. The monthly household rate is obtained by exact compounding, $\phi_t^m = (1 + \phi_t^{ann})^{1/12} - 1$.

sensitivity, the scale is re-inverted rather than held fixed:

$$\phi_0(\phi_1) = \frac{17.5/10{,}000}{\left(1/(1-\ell^B_{2011})\right)^{\phi_1} - 1}. \tag{4.8}$$

Appendix C uses this expression to re-fit the same 17.5-basis-point target in every curvature row, so that the exercise changes the functional shape of pass-through rather than silently changing the external calibration.

Since the 2018 banking system was less fragile than in 2011, the capital-channel evidence in Sabatini (2022) motivates an illustrative 2018 fragility attenuation of 0.94–1.31 basis points. This is neither a confidence interval nor an estimate from Sabatini. The one-fifth to one-seventh attenuation factor is a thesis construction. Alternative rate-side anchors from Altavilla et al. (2017) and De Marco (2019) are reported in Appendix C.4 because they estimate different quantities and cannot be combined with Bofondi et al. (2018). The 300–400 basis-point spread variation and convexity sensitivities are reported in Appendix C.4.

CHAPTER 5

Results

This chapter reports what the sovereign–bank bridge of Chapter 3, calibrated in Section 4.3, implies for Italian households. It opens with the distributional result, because that is what the exercise is for, and turns afterwards to the wedge that produces it, the evidence on the bank stage, and the aggregate response.

Everything below is conditional on the transported Italy-2011 credit calibration. France 2024 is a low-exposure scaling check and Italy 2022 a non-causal policy contrast.

5.1 Who Bears the Loss

Two household groups are exposed to a sovereign repricing, and they are not the same people. This section takes the model-generated borrower-wedge path as given, using the 6.56-basis-point June-2018 *peak* as the reference summary and its twelve-month mean for the first-year allocation, and asks who ends up paying. Section 5.2 then constructs that wedge and reports how far it moves with the vintage of the balance-sheet inputs behind it.

Government securities and other financial assets are concentrated among wealthier households, so they bear the largest direct losses from repricing. The borrower wedge instead loads on households with bank debt, whose participation rises from the bottom quintile to the middle and upper ones. Figure 5.1 shows that raw balance-sheet pattern, which is why the two channels are kept as separate layers below.

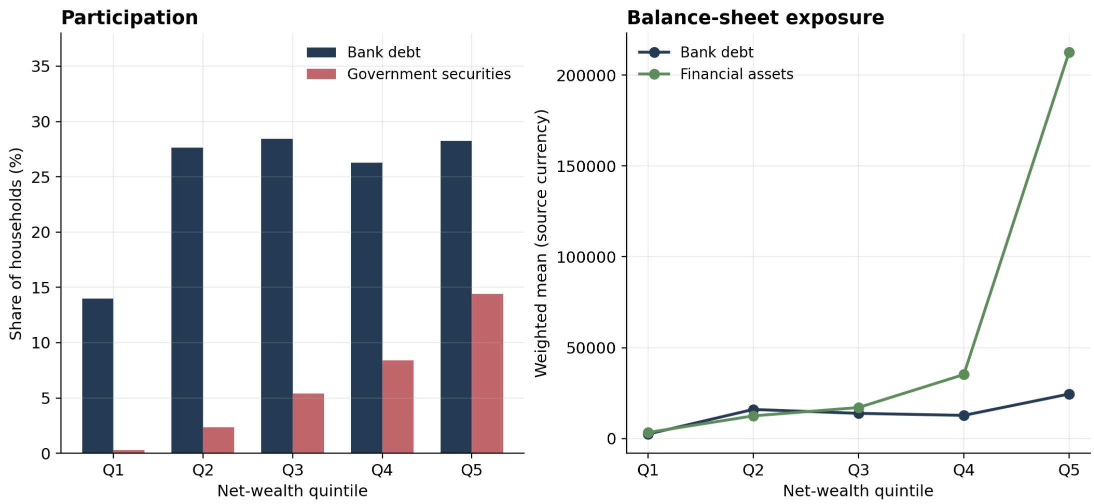

Figure 5.1. Observed household exposure by net-wealth quintile in SHIW 2022. The left panel reports participation in bank debt and government securities. The right panel reports weighted mean bank debt and financial assets. The figure shows the raw balance-sheet pattern that motivates the distinction between borrower-side incidence and asset-holding incidence.

The model supplies the bank-net-worth and borrower-wedge paths, and the steady-state borrower mass calibrated to SHIW bank-debt participation. It does not supply the household distribution used below. Its single liquid asset cannot reproduce the observed joint distribution of Italian net wealth, debt, and securities. The quintile step is therefore an accounting exercise on fixed SHIW-2022 vectors, into which the model-generated wedge and the four allocation constants of Appendix C enter as inputs.

Define $\mathcal{L}_e^{\mathrm{S1}}$ as the SHIW-weighted spread-only allocation and $\mathcal{L}_e^{\mathrm{S2}}$ as the otherwise identical allocation after the model-generated borrower wedge is added. The corresponding allocation amplification is

$$\mathcal{A}_e^{\mathrm{wedge/SHIW}} = \frac{\mathcal{L}_e^{\mathrm{S2}}}{\mathcal{L}_e^{\mathrm{S1}}}. \tag{5.1}$$

The two allocations have the same population, wealth-quintile definition, consumption weights, direct sovereign-asset loss, and denominator. The only changed object is the model-generated borrower wedge. Table 5.1 makes this counterfactual auditable. Loss entries are basis points of each quintile's consumption, while the final two columns give that quintile's share of total allocated loss before and after the wedge is activated.

The spread-only allocation is reported as its two components, because they carry different evidential weight. The bank-equity term routes the sovereign price loss to households through their indirect claims on bank equity and on fund and insurance holdings. The securities term routes it through directly held government bonds. Bank equity is the larger of the two in every quintile, contributing 2.81 of Q1's 3.09 basis points and 70.28 of Q5's 104.47. It is also the component that rests on the weakest constant, the 0.60 bank-equity exposure ratio inherited from an earlier French calibration, which Appendix C varies over 0.30–1.20. Showing the split lets a reader see exactly how much of the reference allocation depends on that constant, and rebuild the total from published inputs.

Figure 5.2 visualises the same audited comparison. The table remains the record of reference. It fixes the household set, population weights, quintile cut-offs, direct-loss vector, and denominator, while the chart makes the resulting reallocation easy to inspect.

For this matched allocation, $\mathcal{A}_{2018}^{\mathrm{wedge/SHIW}} = 1.073$. The added wedge raises total allocated loss by roughly seven percent and the bottom-two share from 4.80% to 6.28%. A separate richer SHIW accounting calculation, using the same allocation-level exposure ratios documented in Appendix C but a different, reduced-form bank-block computation, gives 1.43-times amplification and an 11.36% bottom-two share. It is useful evidence on sensitivity to the bank-block specification, but neither a second HANK estimate nor a confidence bound on 6.28%. Both figures move with the three allocation-level constants documented in Appendix C. Sweeping the stock-loss annuitisation rate and the two Italy-specific exposure ratios across their declared, sourced alternatives spans 6.28–10.54% for the baseline wedge/SHIW allocation, a wider range than the dated-bridge sensitivity in Table 5.2 alone. The 2022 vectors post-date the event, which Appendix C.5 tests by rebuilding the allocation on the 2016 and 2020 waves. Earlier vintages give a *higher* bottom-two share, so the post-event survey is the conservative choice.

Quintile	SHIW weight	Spread-only loss (S1) Bank equity	Securities	Total	Added wedge	Total loss	Share of total loss Before	After
Q1	12.8%	2.81	0.29	3.09	0.90	4.00	0.9%	1.1%
Q2	15.3%	8.85	2.86	11.70	5.00	16.70	3.9%	5.2%
Q3	17.7%	10.46	6.66	17.12	3.74	20.86	6.6%	7.5%
Q4	21.4%	17.77	10.25	28.02	2.83	30.85	13.2%	13.5%
Q5	32.8%	70.28	34.19	104.47	3.56	108.03	75.3%	72.6%
Q1+Q2	–	–	–	–	–	–	4.8%	6.3%

Notes: Losses are basis points of each quintile's own consumption. The spread-only allocation S1 splits into a bank-equity component, which routes the sovereign price loss through indirect household claims on bank equity and on fund and insurance holdings, and a securities component, which routes it through directly held government bonds. "Added wedge" is the model-generated borrower-credit cost allocated through fixed SHIW-2022 bank-debt and consumption vectors. S1 and S2 use unchanged SHIW population weights, wealth-quintile boundaries, direct losses, and denominators. The table is a microdata-assisted allocation of a structural wedge, not a model-only estimate of the household wealth distribution.

Table 5.1. Italy 2018 wedge/SHIW-assisted incidence audit.

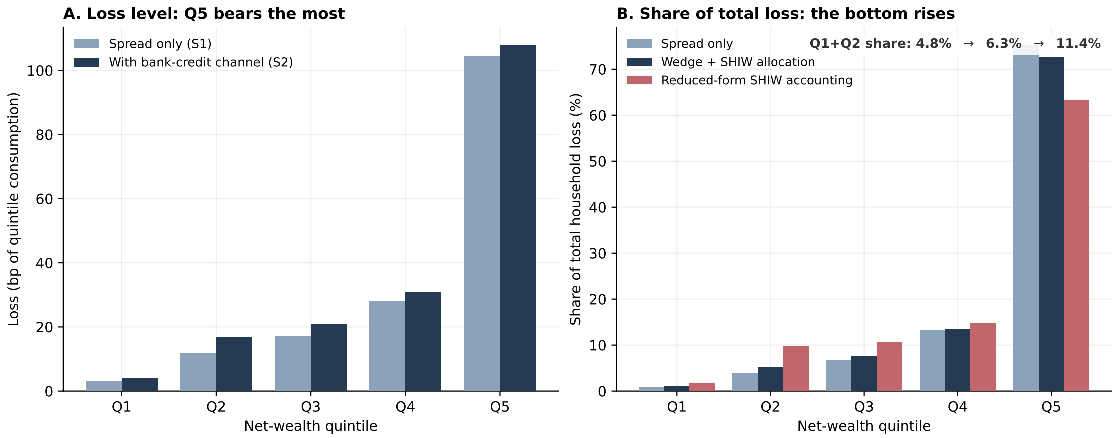

Figure 5.2. Italy 2018 wedge/SHIW-assisted incidence. The left panel plots the fixed spread-only allocation against the same allocation with the model-generated borrower-credit loss added, in basis points of each quintile's own consumption. It is the loss-level view of Table 5.1 and carries its values, so Q5 dominates the level even though the bottom-two *share* rises. The right panel reports shares of total allocated loss and adds a third series. Alongside the spread-only and preferred wedge/SHIW allocations, it shows the separate reduced-form SHIW accounting variant, which reaches an 11.4% bottom-two share. That third series is a different bank-block specification discussed in the text, not a layer of the audit table, and it appears here only so the three allocations can be compared on one axis.

The condition behind the shift

The direction of this shift is not a calibration outcome, and it is worth deriving, so that the sensitivity exercises below are read for what they actually test. Let $\bar{\phi}_{12}^{\text{ann}} \equiv \frac{1}{12}\sum_{t=0}^{11}\phi_t^{\text{ann}}$ be the annualised wedge averaged over the first twelve months. This is the allocation rate used in the incidence calculation: it places the credit-cost flow on the same first-year horizon as the annuitised direct stock loss. Write the resulting consumption-weighted loss allocated to quintile group g as

$$L_g(\bar{\phi}_{12}^{\text{ann}}) = A_g + \bar{\phi}_{12}^{\text{ann}} D_g, \tag{5.2}$$

where A_g collects the spread-only components of Table 5.1, D_g is the group's fixed consumption-weighted bank-debt exposure per unit of annualised wedge, and $A \equiv \sum_g A_g$, $D \equiv \sum_g D_g$. Thus, $\bar{\phi}_{12}^{\text{ann}} D_g$ is the added-wedge column. Any group $\mathcal{B}$ then carries a share

$$s_{\mathcal{B}}(\bar{\phi}_{12}^{\text{ann}}) = \frac{A_{\mathcal{B}} + \bar{\phi}_{12}^{\text{ann}} D_{\mathcal{B}}}{A + \bar{\phi}_{12}^{\text{ann}} D} = (1-\theta)\,\frac{A_{\mathcal{B}}}{A} + \theta\,\frac{D_{\mathcal{B}}}{D}, \qquad \theta \equiv \frac{\bar{\phi}_{12}^{\text{ann}} D}{A + \bar{\phi}_{12}^{\text{ann}} D}, \tag{5.3}$$

which is a weighted average of the share the group bears in each channel separately, with the weight equal to the credit channel's share of the total loss. Differentiating gives

$$\frac{\partial s_{\mathcal{B}}}{\partial \bar{\phi}_{12}^{\text{ann}}} = \frac{A D_{\mathcal{B}} - A_{\mathcal{B}} D}{\left(A + \bar{\phi}_{12}^{\text{ann}} D\right)^2}, \tag{5.4}$$

and, for $A, D > 0$ and a non-negative wedge,

$$\frac{\partial s_{\mathcal{B}}}{\partial \bar{\phi}_{12}^{\text{ann}}} > 0 \qquad \Longleftrightarrow \qquad \frac{D_{\mathcal{B}}}{D} \;>\; \frac{A_{\mathcal{B}}}{A}. \tag{5.5}$$

A bank-credit channel makes sovereign stress less top-heavy if and only if the borrowers it reaches are less concentrated at the top than the holders of the assets that reprice. The condition contains no parameter of the model. It compares two exposure shares that SHIW measures directly, and it can be evaluated in any country with a household balance-sheet survey and an estimate of which households borrow from banks.

Italy satisfies it by a wide margin. The bottom two quintiles carry 4.80% of the spread-only loss and 26.51% of the borrower-wedge loss, a factor of 5.5. That gap fixes the sign of the reallocation for any positive wedge, which matters because ϕ_t is one-sided by construction, since Equation (3.5) admits no borrower subsidy. What the calibration determines is θ, the weight the credit channel receives, which is 6.81% at the reference wedge. The move from 4.80% to 6.28% is that arithmetic and nothing more. Every sensitivity exercise below is therefore a statement about θ, and none of them can reach the sign. The pattern is also economically consistent with evidence that consumer-credit spreads bear differently on households according to debt and liquidity (Faccini et al., 2026).

5.2 The Wedge Behind the Allocation

The allocation above applies a wedge that has to be built from a balance sheet, and the balance sheet has a date. The Italy-2018 episode pairs a 166.1-basis-point sovereign repricing with the largest domestic-bank exposure of the four cases. Table 5.2 reports what the bridge produces. Under the preferred partial-pre-event bridge, the borrower wedge is 4.03–7.84 basis points, output falls by 0.0018% to 0.0034% under the fixed-wage closure, and the bottom two wealth quintiles' share of the total allocated loss rises from 4.80% to 5.99–6.08%. Using the June-2018 balance sheet throughout gives 6.56 basis points, −0.0029%, and 6.28%.

The rows differ only in the date of the balance-sheet inputs, and that difference is the point of the table. Exposure and leverage are observable for December 2017, but the December-2017 EBA disclosure carries no residual-maturity measure, so no strictly pre-event duration exists. The partial-pre-event rows therefore pair dated exposure and leverage with a later duration input and report its range instead of settling on a single value. The all-June mapping remains exactly reproducible, but June follows the May event, which makes it an *episode-vintage proxy* and not a pre-event estimate.

The direct allocation is invariant at 4.80% across every row, because the price-loss factor cancels from a bottom-two *share* even though it does not cancel from the level. The rows therefore vary θ alone.

This is the *dated-input range*. It varies the vintage of the balance-sheet inputs and holds the firm-lending calibration fixed. Appendix C reports a *household-transport range* of roughly 3.0–6.6 basis points that does the opposite, and the two measure different uncertainties, so neither contains the other. Within the dated-input range, the 5.99–6.08% figure is the preferred result, because only its duration input post-dates the event. The quintile audit in Table 5.1, the cross-event comparison in Table 5.4, the aggregation sensitivity in Table 5.5, and the appendix allocation grids all carry the 6.28% June figure, because those exercises vary one object at a time against a single coherent balance sheet, while the partial-pre-event rows deliberately mix two vintages. Rescaling the wedge moves between them. The wedge-to-share mapping is close to linear over this interval. No conclusion below depends on which row serves as the reference.

The same conclusion survives the separately recalibrated curvature sweep in Appendix C. Each value of ϕ_1 is paired with the ϕ_0 that exactly re-fits the Italy-2011 17.5-bp target. This avoids treating a changed curvature parameter as a silent change in the external calibration. The twenty-five-bank aggregation result (1.161 bp) is shown next to the representative-bank June proxy (6.563 bp) as a different weighting object. It is a 25-bank loan-book-proxy aggregation sensitivity, not a replacement estimate of the representative Italian banking sector.

The table separates a dated representative-sector bridge from a heterogeneous-bank aggregation sensitivity. Neither is an estimate of observed household contracts. A sovereign repricing that does not generate a large aggregate contraction can nevertheless shift the allocation of losses through an added borrower-cost channel. That conditional allocation shift is the thesis's central quantitative result, and the June 6.563-bp proxy is not. Appendix C.5 reports further allocation-level diagnostic scenarios. They are not an Italian confidence interval.

The extended-window rows above feed a peak spread deviation, reached months after the May event, into an impact-loss formula that treats the balance sheet as struck at that same instant. This is a deliberate modelling choice (Section 4.1), not an oversight, but its consequence is worth making visible. The acute-window row does exactly that. Holding the June-2018 balance sheet fixed and substituting the 114.4-bp acute spread for the 166.1-bp extended peak lowers the

Bridge / weighting object	Date status	ϕ_1	ϕ_0	Peak wedge	Mean wedge (12m)	Fixed Y	Wage-moving Y	Direct S1 share	Total S2 share
Partial pre-event, low duration	Dec.-2017 inputs / June duration	1.0	0.003812	4.03 bp	3.78 bp	−0.0018%	−0.0003%	4.80%	5.99%
Partial pre-event, central duration	Dec.-2017 inputs / June duration	1.0	0.003812	5.33 bp	4.99 bp	−0.0023%	−0.0003%	4.80%	6.02%
Partial pre-event, high duration	Dec.-2017 inputs / June duration	1.0	0.003812	7.84 bp	7.30 bp	−0.0034%	−0.0005%	4.80%	6.08%
June-2018 representative-bank mapping	Episode-vintage proxy	1.0	0.003812	6.56 bp	6.13 bp	−0.0029%	−0.0004%	4.80%	6.28%
June-2018 mapping, acute window	Timing-consistent (acute)	1.0	0.003812	4.29 bp	4.02 bp	−0.0019%	−0.0003%	4.80%	6.21%
25-bank aggregation sensitivity	Loan-book-proxy weighting	1.0	0.003812	1.161 bp	1.09 bp	−0.0005%	n/a	4.80%	5.10%

Notes: The first four rows are freshly re-solved from the dated-bridge timing table. The full acute- and extended-window timing table and validation checks are in the replication package. **Two wedge columns appear because they do different jobs.** The peak wedge is the maximum of the borrower-wedge path and is the summary statistic comparable across episodes. The incidence allocation applies the twelve-month mean instead, because the direct component is a stock revaluation annuitised into a first-year flow, and both channels must be measured over the same horizon. An instantaneous peak rate would overstate a year's interest cost. The S1 and S2 columns are therefore generated by the mean-wedge column. "Direct S1 share" is the bottom-two share of the unchanged SHIW spread-only allocation. "Total S2 share" adds the model-generated borrower wedge using the same SHIW population weights, wealth-quintile boundaries, direct losses, and denominator. The fifth row re-prices the same June-2018 balance sheet with the 114.4-bp acute window instead of the 166.1-bp extended-episode peak, so that the spread input and the balance-sheet date are struck at the same moment. It is the timing-consistent counterpart to the row above it, not a preferred estimate. The final row is the twenty-five-bank aggregation, which holds the household block fixed but aggregates 25 named bank paths with wider-panel loan-book proxy weights. The richer SHIW accounting allocation (11.36%) is reported separately in Section 5.1.

Table 5.2. Evidence-first Italy-2018 calibration and incidence results, extended 166.1-bp window.

wedge from 6.56 to 4.29 basis points and the bottom-two share from 6.28% to 6.21%, a materially smaller change than moving between the dated-bridge or allocation-constant variants elsewhere in this chapter. The direction and rough magnitude of the headline result are therefore not an artefact of using the extended window in place of the acute one.

5.3 Market Evidence on the Balance-Sheet Mechanism

The bank stage in the transmission bridge has an observable market implication. Institutions with larger predicted economic losses relative to CET1 should experience more negative equity returns over the event window. This is a test of the upstream sovereign-exposure loss scaling, not of the bank-to-household-credit elasticity. Italy 2018 provides a useful acute check because the 15–31 May event window records a 114.4-bp movement, whereas the distinct 166.1-bp January–November input is used for the structural stress scenario.

An earlier draft of this section used a strict pre-event specification based on the December-2017 EBA disclosure. That sovereign template does not provide a reconstructible Italian-bank exposure regressor, so the draft specification is withdrawn and carries no result here. Appendix B and the provenance note document the finding. The replacement below is the only pre-event market-validation result used here.

Exposure gradient, June-2017 vintage. The replacement specification uses the EBA's *2017* EU-wide Transparency Exercise (European Banking Authority, 2017), whose sovereign template is fully populated at 30 June 2017, eleven months before the event window and with no look-ahead.

$$CAR_i = \alpha + \beta \widehat{\ell}^{CET1}_{i,2018} + X_i'\gamma + \varepsilon_i. \tag{5.6}$$

The outcome, controls, spread movement, and duration proxy follow the pre-specified design. Sample coverage is 13 banks because June-2017 EBA and pre-event control coverage must be complete for every regression field (Appendix B reports the sample-construction table).

Specification	$\hat{\beta}$	HC3 s.e.	p (t_7)	p (perm.)	N
Continuous predicted CET1 loss (main spec)	-1.072	1.213	0.406	0.222	13
Binary Italy indicator (diagnostic)	-14.715	7.817	0.102	0.013	13

Notes: Sovereign debt securities and CET1 are both from the 30 June 2017 EBA Transparency Exercise, eleven months before the May 2018 event. The loss uses the 114.4-bp acute spread movement and a common 4.95-year duration proxy. Both rows use the same 13-bank sample, the same four pre-event controls, and the same HC3 sandwich estimator. The t_7 column is the two-sided t reference with seven residual degrees of freedom. The final column is a permutation reference distribution for the same HC3 t statistic. The regressor is permuted across banks while each bank keeps its own return and controls. The continuous regressor uses 200,000 permutations. The Italy indicator has only $\binom{13}{4} = 715$ distinct assignments, so its distribution is enumerated exactly. The second row is a diagnostic, not a replacement specification. It isolates how much of the main coefficient's identifying variation comes from country grouping and how much from a continuous exposure measure.

Table 5.3. Strictly pre-event sovereign-loss scaling and bank equity returns, Italy 2018 (June-2017 EBA vintage).

The coefficient is negative, -1.07, with an HC3 standard error of 1.21 and a small-sample $p = 0.406$. The relationship cannot be distinguished from zero at conventional levels. A leave-one-bank-out audit shows the sign is stable across all 13 specifications (range -1.90 to -0.30), none reaching conventional significance. With thirteen banks and seven residual degrees of freedom the HC3 sandwich is itself an asymptotic object, so the p-value is also computed against a permutation reference distribution, permuting the exposure regressor across banks while each bank keeps its own return and controls. That gives $p = 0.222$ over 200,000 permutations. The permutation test is roughly twice as favourable as the t_7 reference and still far from conventional levels. The sample is very imprecise, which is different from evidence that the exposure gradient is absent. The placebo is retained as a robustness check, but it is not used as evidence for the main result, because it cannot change how that evidence is classified.

Source of the identifying variation. Only four of the thirteen banks are Italian (BPER, Intesa Sanpaolo, Mediobanca, UniCredit), and the predicted-loss regressor is close to zero for the remaining nine. The continuous coefficient is therefore identified partly by a comparison between Italian and non-Italian institutions, and only partly by smooth within-country variation in exposure. The second row of Table 5.3 isolates that comparison. Replacing the continuous regressor with a binary Italy indicator, holding the controls and estimator fixed, gives -14.72 with an HC3 standard error of 7.82. Here the choice of reference distribution decides the verdict. The t_7 reference returns $p = 0.102$. Enumerating all $\binom{13}{4} = 715$ ways of labelling four of these banks "Italian" returns $p = 0.013$, because only nine of the 715 possible groupings produce a t statistic as large in absolute value as the true one. The unconditional gap behind that number is stark, with the four Italian banks averaging -8.3% abnormal return against $+1.5\%$ for the remaining nine.

Two readings follow, and they point in different directions. The country grouping is a strong pattern that the HC3 sandwich with seven degrees of freedom was hiding. But nationality is not randomly assigned, so the exact p-value measures how unusual the Italian group looks against all equally sized groupings of these thirteen banks, not the effect of an exposure treatment. Italian banks in May 2018 also differed from their peers in domestic macroeconomic exposure, funding base, and political news flow, and four controls do not exhaust that. The surviving evidence is asymmetric. Something distinctive happened to Italian bank equity in this window, at a magnitude the sovereign-loss mapping predicts. A sample of four cannot say how much of it scales with measured exposure and how much with nationality.

A supplementary diagnostic. The June-2018 EBA vintage, which post-dates the event and is therefore not treated as pre-event evidence, gives $\hat{\beta} = -1.326$, HC3 s.e. 0.335, $p < 0.001$, $N = 19$, reproduced exactly from the same underlying input files. The sign agrees with the valid June-2017 estimate, while conventional significance appears only in the later vintage. That pattern is consistent with a look-ahead concern, although the later vintage also has richer coverage and a closer balance-sheet measurement. It is therefore a diagnostic, not a definitive test of timing bias. Aggregate household-credit quantities and rates are discussed separately in Appendix C.2. Their formal local-projection and synthetic-control tests do not identify a credit deviation from counterfactual paths.

5.4 Aggregate Transition Dynamics

The loss of bank net worth raises the borrower wedge and reduces the resources of households with negative liquid assets. Aggregate consumption and output consequently decline together (Figure 5.3). In the fixed-wage baseline, both the wage and technology terms in the Phillips curve are held at their steady-state values. Inflation and the real interest rate therefore remain exactly at their steady-state values. The response is a direct borrower-cost result, not a price- or monetary-policy propagation result.

The closure is intentionally narrow, and the alternative changes the aggregate profile more than its size. A parameter-free sensitivity in which the real wage moves one-for-one with output gives an impact response of −0.0004%, about one seventh of the fixed-wage −0.0029%. Falling wages are disinflationary here, and the lower real rate offsets more consumption at impact than the direct labour-income loss removes. The trough, however, moves the other way. Output bottoms out at −0.0045% under the moving wage against −0.0029% under the fixed one, because the labour-income loss accumulates while the rate offset fades with the wedge. The fixed-wage path is therefore not a lower bound on the aggregate effect. It is a front-loaded one. Two considerations argue for keeping it as the baseline nonetheless. It isolates the borrower-cost channel the thesis measures, with no Phillips-curve or policy-rule propagation mixed in, and the alternative's own magnitude depends on a wage rule that no Italian evidence here disciplines. Slacalek et al. (2020) find that indirect labour-income effects dominate euro-area consumption responses, which is a reason to read the fixed-wage aggregate as conservative about the medium-run response, and not as the full picture. The SHIW-assisted incidence calculation is unaffected either way, since its inputs are the model-generated wedge and the fixed observed vectors, never the aggregate wage path.

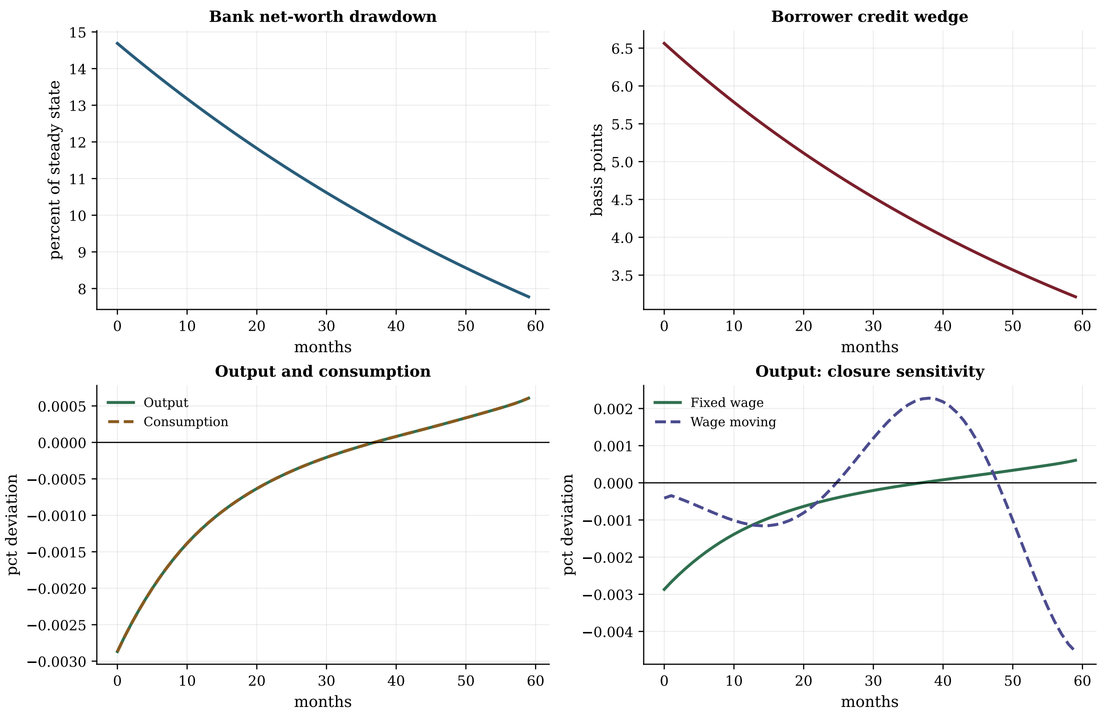

Figure 5.3. Italy 2018 transition paths under the one-asset HANK model, monthly. The top row reports the bank-net-worth and borrower-wedge margins. The lower-left panel reports output and consumption under the fixed-wage baseline. The lower-right panel makes the pre-declared wage-moving closure sensitivity visible. Inflation and the real rate are omitted because they are identically flat under the fixed-wage closure.

This result is also consistent in direction with models in which sovereign risk constrains intermediation through bank funding and precautionary deleveraging (Bocola, 2016). The channel is quantitatively visible but not crisis-sized at the aggregate level.

5.5 Cross-Event Scaling and the Monetary-Union Comparison

France 2024 is a low-exposure scaling check for whether the same sovereign-bank-household mapping generates proportionately weaker transmission in a lower-stress, lower-exposure banking system. It is not an external validation of Italian household pass-through. For comparability, Table 5.4 displays the same June-2018 Italy proxy that is explicitly dated in Table 5.2.

Object	Italy 2018	France 2024
Extended-episode spread input (bp)	166.1	34.1
Domestic sovereign assets/total assets	10.15%	2.11%
Portfolio duration (years)	4.95	7.28
Assets/equity	17.60	20.23
Economic-value loss / representative equity	14.69%	1.06%
Peak borrower wedge (bp)	6.563	0.408
Peak output/consumption deviation	$\approx -0.0029\%$	$\approx -0.0002\%$

Notes: The comparison covers only objects measured identically in both countries. Spread inputs come from the same source construction, and exposure, duration and leverage from the same ECB and EBA templates. The loss row is a synthetic sector construct, not an observed institution. The domestic-exposure share is an ECB sector aggregate for the whole national banking system, while duration and leverage come from the EBA's large-bank sample, because no single source supplies all three. Appendix B.4 shows the choice matters. EBA exposure understates domestic concentration enough to reverse the France–Italy ranking. The 14.69% is a calibration input, not a measured loss at any Italian bank. Household-incidence shares are deliberately excluded. The French allocation would need a debt-to-consumption ratio and a government-securities-to-consumption ratio with no French counterpart among the sourced inputs, and the published HFCS tables do not separate government securities from total financial assets. The available French constants are maintained values inherited from an earlier calibration, not sourced counterparts to the Italian SHIW figures, so a French bottom-two share would not be comparable with the Italian one.

Table 5.4. Italy June-2018 proxy versus France 2024: scaling comparison.

Although the representative French portfolio has longer duration and slightly higher leverage, the 34.1-basis-point spread input and 2.11% domestic-sovereign asset share generate a bank-net-worth loss of only 1.06% and a borrower wedge of 0.408 basis points. Output and consumption fall by approximately 0.0002%. The Italy-to-France ratio in the bank-net-worth loss is roughly fourteen to one, and the wedge ratio roughly sixteen to one, despite a spread ratio of only about five to one. This asymmetry arises because domestic exposure enters multiplicatively with the spread in the bank-loss formula. France's much smaller asset share compounds the smaller shock, not just one or the other. The France valuation is also the most sensitive to the choice of benchmark. Roughly half of the acute-window OAT–Bund widening reflects a rally in German yields and not a rise in French ones, so pricing the French sovereign portfolio off its own yield alone gives a materially smaller acute-window impairment (0.21% of representative equity, against 0.46% under the spread-based benchmark used throughout). Appendix C reports the full national-yield sensitivity and why the spread-based measure remains preferred.

The comparison is a bank-side result, and it stops at the bank. Exposure and the size of the repricing jointly determine how much of a sovereign movement reaches private credit, and

France 2024 is an order of magnitude smaller than Italy 2018 on both. That is exactly the pattern a monetary union produces. A common policy regime coexists with highly unequal national exposure (Bayer et al., 2024). The household leg stays out because the French allocation constants are inherited values with no sourced French counterparts, so a French bottom-two share would differ from the Italian one partly through the provenance of its own inputs.

A backstop-regime contrast: Italy 2022

Using the June-2022 baseline, the Italy-2022 episode produces a 15.5-basis-point spread input, a 1.49% bank-net-worth loss, and a 0.576-basis-point borrower wedge. The period combines the TPI announcement, ECB rate tightening, and the collapse of the Draghi government. It is therefore only a non-causal policy-regime contrast, not an estimate of the TPI effect.

5.6 Bank-Aggregation Sensitivity

The representative-bank calibration aggregates the domestic banking sector into one exposure path. This sensitivity instead aggregates 25 named bank paths with loan-book proxy weights, leaving the household block, aggregate closure, population weights, and SHIW allocation unchanged.

Disaggregating the representative bank into its 25 named constituent institutions, each institution $i \in \{1, \dots, N\}$ in the panel ($N = 25$) has its own impact equity loss $\ell^B_{i,0}$. Using its domestic-sovereign exposure $H^{dom}_{i,e}$, the common event duration D_e, and pre-event CET1 capital gives

$$\ell^B_{i,0} = \frac{H^{dom}_{i,e} D_e \left(\Delta s^{bp}_e / 10{,}000\right)}{CET1_{i,e}}, \qquad N^B_{i,0} = (1 - \ell^B_{i,0}) N^B_{ss}. \tag{5.7}$$

This is algebraically equivalent to the former leverage-exposure representation after substituting the bank's reported leverage ratio. Each bank's net-worth recovery path $N^B_{i,t}$ is solved independently through the dynamic law of Section 3.2, yielding a bank-specific annualised borrower wedge $\phi^{\text{ann}}_{i,t} = \phi_0 \max\left[\left(\frac{N^B_{ss}}{N^B_{i,t}}\right) - 1, 0\right]$, evaluated at the baseline curvature $\phi_1 = 1$. The economy-wide borrower wedge reaching households is then the loan-book-proxy-weighted aggregate.

$$\phi^{\text{agg,ann}}_t = \sum_{i=1}^{N} w_i \, \phi^{\text{ann}}_{i,t}, \qquad \text{where } w_i = \frac{\text{Loan Book}_i}{\sum_{j=1}^{N} \text{Loan Book}_j}, \tag{5.8}$$

with $\text{Loan Book}_i = \text{Total Assets}_i \times (1 - \text{Liquid Assets Share}_i)$. The annualised aggregate path is converted to its monthly equivalent before it enters the household block. Its peak is 1.161 basis points, below the representative-bank proxy value of 6.563 basis points, because large lower-stress banks receive substantial proxy weight. The corresponding output response is -0.0005% and the bottom-two loss share is 5.1%, against 6.3% under the representative-bank calibration (Table 5.5). The same direct-loss allocation and SHIW quintile definition are retained in both rows. Only the bank path changes.

The gap between 1.161 and 6.563 basis points is largely a statement about which banks are in the panel. The aggregation weights twenty-five *euro-area* institutions by loan book, and Italian banks carry only 15.5% of that weight. The remaining 84.5% belongs to French, Spanish, German, Dutch, Austrian and Irish banks holding almost no Italian sovereign debt, whose weighted mean

Bank-path specification	Peak wedge (bp)	Impact output (%)	Bottom-two share: S1	Bottom-two share: S2	Amplification
Representative-bank calibration	6.563	-0.0029	4.8%	6.3%	1.073×
25-bank wider-panel proxy aggregation	1.161	-0.0005	4.8%	5.1%	1.013×

Notes: Both specifications use the accepted household and aggregate blocks and the fixed SHIW-2022 allocation. The aggregation row changes only the borrower-wedge path, using 25 wider euro-area bank paths weighted by a loan-book-size proxy. These are not observed Italian household-credit market shares: households are not matched to lenders.

Table 5.5. Italy 2018 bank-aggregation sensitivity.

wedge is 0.415 basis points. Table 5.6 re-weights the same bank paths over Italian banks alone, which is the population that actually lends to the households in the incidence exercise. That gives 5.33 basis points against the representative bank's 6.563, a difference of nineteen per cent, not a factor of five and a half. Heterogeneous-bank aggregation performed on the relevant population therefore comes close to reproducing the representative-bank calibration, and the headline magnitude is not an artefact of aggregating away bank heterogeneity.

Weighting panel	Share of loan-book weight	Aggregate peak wedge (bp)	
	(%)	As shipped	Mediobanca repaired
All 25 euro-area banks	100.0	1.161	1.175
Italian banks only	15.5	5.238	5.329
Non-Italian banks only	84.5	0.415	0.415
Representative-bank calibration	–	6.563	6.563

Notes: All rows use the same bank-level wedge paths and the same loan-book proxy weights. Only the set of banks entering the average changes. The final column repairs one exposure. Mediobanca is recorded with zero Italian sovereign exposure in the 2018 Transparency Exercise, flagged as a reported zero, while the same bank reports EUR 4,650 million in the June-2017 exercise and EUR 5,827 million in 2022. That vintage carries eight such zeros against one in 2022 and none in 2024. The exposure is imputed as the mean of those two bracketing observations, using only the bank's own reported history, which implies an exposure ratio of 0.068 against UniCredit's 0.068 and a 2.60-basis-point wedge at Mediobanca's own leverage. The repair moves the aggregate by 0.014 basis points and is reported for completeness, since it changes nothing.

Table 5.6. Bank coverage in the 25-bank aggregation, Italy 2018.

A second reason the two rows are not a like-for-like substitution concerns the calibration. Section 4.3.3 inverted the pass-through scale ϕ_0 from a *representative-sector* Italy-2011 loss of 31.46% of equity. Holding ϕ_0 fixed while replacing that representative bank with a loan-book-weighted average of 25 individual paths changes the bank aggregation without re-fitting the external target that disciplines the mapping. The curvature sweep in Appendix C deliberately does the opposite, pairing every ϕ_1 with the ϕ_0 that exactly re-fits the 17.5-basis-point target. A consistent 25-bank specification would meet the same standard, aggregating the 2011 episode the same way before re-inverting ϕ_0 against it. That is not attempted here, and the reason is comparability, not missing data. A 2011 aggregation would have to be built from the EBA 2011 stress-test disclosure, which reports bank-level sovereign exposures by maturity bucket but is dated 31 December 2011, after the autumn episode the calibration targets, and covers five Italian banks against the twenty-five paths aggregated for 2018. Such a re-calibration would differ from the 2018 aggregation in composition and in date at once, which makes it a separate exercise and no extension of this one.

The direction of the missing correction is nevertheless clear. The same weighting that attenuates the 2018 wedge would attenuate the 2011 loss that ϕ_0 is fitted to, so re-calibrating would raise ϕ_0 and move the 25-bank wedge back toward the representative-bank value. How much of the gap would close cannot be determined from the available data, which is precisely why the row is reported as an aggregation diagnostic and never as a smaller estimate of the same object.

Bank-level dispersion is economically visible even though the proxy-weighted average is smaller. The 90th-percentile institution faces 6.23 basis points, the maximum is 12.30 basis points, and the top five banks account for 77.2% of the aggregate peak wedge. A borrower's cost therefore depends heavily on which bank lends to them, which the representative-bank calibration averages away. Recovering that margin needs household–lender links, and public Italian data do not supply them.

5.7 Interpreting the Magnitudes

Two magnitudes carry the chapter, and they are uncertain in different ways. The aggregate response is small under the fixed-wage closure and would grow under a moving wage, as Section 5.4 showed. The conditional allocation is larger and turns on two objects, the transported bank-to-credit scale and the dated balance-sheet bridge whose partial-pre-event row is the primary quantitative reading. The June proxy, attenuation, curvature, and allocation-constant exercises trace how those two objects move the answer.

Those two objects have so far been reported apart, as a dated-input range in Section 5.2 and a household-transport range in Appendix C, which leaves the reader to combine them. Equation (5.3) removes the need to, because the share depends on the wedge only through θ. Table 5.7 therefore crosses the two. Rows vary the vintage of the balance-sheet inputs behind the wedge, columns vary the anchor that scales it, and every cell holds the SHIW vectors, quintile boundaries, direct-loss allocation, and denominator fixed.

	Bank-to-credit anchor	
Balance-sheet bridge	Italy-2011 firm target	Household-adjusted
Partial pre-event, low duration	5.99%	5.37%
Partial pre-event, central duration	6.02%	5.38%
Partial pre-event, high duration	6.08%	5.41%
June-2018 representative proxy	6.28%	5.51%
Spread-only allocation, no credit channel	4.80%	

Notes: Each cell is the bottom two quintiles' share of the total allocated loss. The left column reproduces the corresponding rows of Table 5.2, in which the wedge is anchored on the Italy-2011 same-firm lending differential. The right column applies the Neri (2013) household-to-firm attenuation of Appendix C to the same wedge path, evaluated through Equation (5.3) with the direct loss and the household vectors held fixed. The June row is the check on that construction. It returns 5.5106%, against the 5.51% obtained independently in Appendix C by re-solving the model at the re-fitted ϕ_0, so the closed form and the re-solved model agree to four significant figures. This is a deterministic envelope over declared assumptions. It varies the balance-sheet vintage and the bank-to-household anchor while holding the allocation constants and the bank-block specification fixed. It is not a confidence interval and not an identified set, because its edges are assigned inputs, not estimated ones.

Table 5.7. Joint sensitivity envelope for the bottom-two loss share, Italy 2018.

Across the whole grid the credit channel raises the bottom-two share from 4.80% to between 5.37% and 6.08% under the preferred partial-pre-event bridge, and to at most 6.28% under the June proxy. The lower edge takes the least favourable reading of both objects at once, an early balance sheet and a household-adjusted anchor, and the reallocation survives there at roughly twelve per cent of the spread-only share. The allocation-constant grid of Appendix C varies a third object and reaches higher, which is why it is reported separately instead of folded into this table.

Figure 5.4 makes the first of them visible. Holding the direct loss and the SHIW vectors fixed, a larger maintained borrower wedge raises the bottom-two share along a smooth and near-linear path, which is why the transported anchor from Chapter 4 matters as much as any object in the thesis. The frontier maps calibration choices into incidence. It carries no information about which point on it the data select.

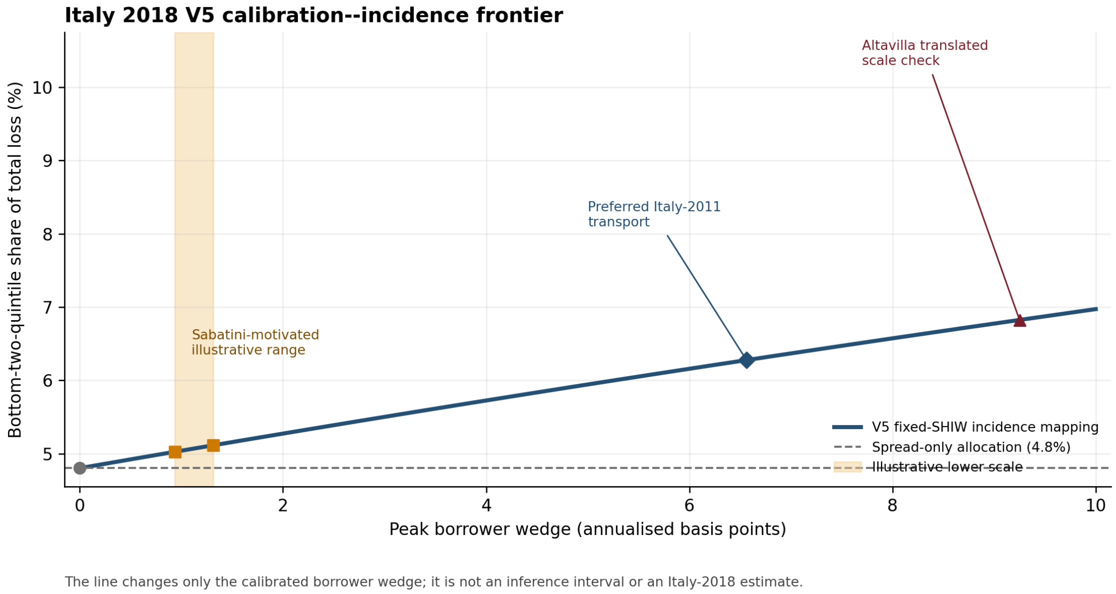

Figure 5.4. Italy-2018 calibration–incidence frontier. The line recomputes the fixed-SHIW allocation as the borrower-wedge path is scaled uniformly. The axis reports the resulting peak wedge, while the allocation applies the corresponding twelve-month mean, a fixed 0.93 of the peak under uniform scaling (see the notes to Table 5.2). The sovereign repricing, household vectors, quintile boundaries, and direct loss are fixed. Markers identify the preferred Italy-2011 transport, the two illustrative lower scales, and the thesis-translated Altavilla scale check. The frontier is a deterministic sensitivity mapping.

The composition shift is robust. It survives every sensitivity because condition (5.5) holds in the data by a wide margin and the wedge is one-sided, neither of which a calibration can overturn. Its size is conditional on constants the appendices document one by one, spanning 5.37% to 6.28% across the joint envelope of Table 5.7 and reaching 10.54% once the allocation constants are varied as well, against a spread-only share of 4.80%. The aggregate effect stands on weaker ground. It is small under this closure and larger under a moving wage. Appendix A records the numerical checks behind these paths, and Appendix C the calibration and allocation diagnostics.

CHAPTER 6

Discussion and Conclusion

6.1 Answer to the Research Question

Chapter 1 asked how domestic banks' holdings of sovereign debt amplify the effects of political risk on households, and who ultimately bears the losses. The general lesson answers the first half, and it does not depend on this thesis's calibration. The incidence of sovereign stress depends on the relative strength of the asset-price and bank-credit channels, because the two reach different households. Government securities and financial assets are held at the top of the wealth distribution, so a shock carried by bond prices is borne by the wealthy. Bank debt is spread far more evenly, so the same shock carried by lending conditions reaches households who own little and owe something. Any sovereign episode sits somewhere between those poles, and where it sits is an empirical question about the banking system, not a property of sovereign risk.

In the second half, losses are specific to one episode, so the answer is conditional rather than general. A wider sovereign spread reduces the economic value of domestic banks' bond portfolios. Under the Italy-2011 calibrated mapping, weaker bank net worth generates a borrower wedge. Fixed SHIW balance-sheet vectors then allocate that wedge across wealth quintiles.

Italy 2018 is the central application. The preferred dated bridge combines December-2017 exposure and leverage with a transparent later-duration sensitivity, producing a 4.03–7.84-basis-point wedge and a 5.99–6.08% bottom-two loss share, against a fixed 4.80% direct-loss share. This is the *dated-input range*, which holds the firm-lending calibration fixed and varies only the vintage of the balance-sheet inputs, and it measures something different from the household-transport range of Section 6.4. The June-2018 6.56-basis-point result is reproducible but remains an episode-vintage proxy. The June-2017 bank-market regression is negative and very imprecise ($\hat{\beta} = -1.07$, HC3 s.e. 1.21, small-sample $p = 0.406$, $N = 13$). It is supplementary evidence on the upstream exposure-loss relationship, not a confirmation of household pass-through.

6.2 What the Composition Result Means

The level and the composition of the loss point in opposite directions, and only one of them is the thesis's claim.

In levels, the Italy-2018 episode is borne overwhelmingly by wealthy households. The top quintile loses 104 basis points of its own consumption against 3 for the bottom quintile, and holds three-quarters of the total allocated loss even after the borrower channel is switched on. Sovereign repricing is, first and foremost, a shock to the holders of sovereign and financial assets, and in Italy those are concentrated at the top of the wealth distribution. Nothing in this thesis contradicts that, and any reading of the results as showing that a sovereign shock falls mainly on the poor would be wrong.

The claim is compositional. Adding the bank-credit channel moves the bottom two quintiles from 4.80% to 6.28% of the total, a proportional reallocation of roughly thirty per cent, because the two channels load on structurally different parts of the distribution. Government securities and financial assets are concentrated at the top, where the top quintile holds 74% and 76% respectively, while bank debt is nearly flat from the second quintile upward, with the second quintile alone holding 23%. That asymmetry is measured in SHIW, not modelled, which is why the composition result does not depend on this thesis's calibration. What the calibration fixes is where Italy 2018 sits between the two poles of Section 6.1, not which way the credit channel pulls.

The result's sign is decided by condition (5.5) rather than by anything estimated here, so no calibration exercise, including this one, can confirm or reverse it. The interesting question is whether other countries and other episodes also satisfy the condition, which their own balance-sheet surveys can answer without this model. Its magnitude, meanwhile, is a share of a modelled loss, not a euro amount, a mortgage payment, or a welfare equivalent.

6.3 Policy Implications

- **Exposure, duration, and leverage cannot be monitored separately.** The bank-loss mapping is multiplicative in the three. France 2024 illustrates this sharply. The French representative portfolio has *longer* duration and *higher* leverage than the Italian one, yet its mapped net-worth loss is fourteen times smaller, because its domestic sovereign share is roughly one fifth. A supervisory dashboard that reads the three inputs as independent risk indicators would rank France as the more vulnerable system on two of three metrics and reach the wrong conclusion. This is the most robust finding here, because it rests only on identically constructed inputs and an accounting identity.
- **A sovereign backstop has a distributional rationale, not only a financial-stability one.** A backstop that limits disorderly sovereign repricing compresses every downstream term in the maintained mapping, including the borrower component that reaches indebted households. It does not merely prevent a bank capital event. It prevents a transmission channel whose clientele differs from that of the asset losses. The Italy-2022 comparison illustrates the logic but identifies nothing, since the TPI announcement coincided with a 50-basis-point tightening and the collapse of the Draghi government. Whether such a backstop improves welfare cannot be answered here, because who ultimately holds the displaced risk is outside the model (Rojas and Thaler, 2024).
- **An economy-wide average can hide who actually faces the risk.** Replacing the representative bank with 25 named paths weighted by loan book lowers the peak wedge from 6.56 to 1.16 basis points, but Table 5.6 shows most of that gap is composition, not heterogeneity. Italian banks carry only 15.5% of the panel's weight, and re-weighting over the institutions that actually lend to these households gives 5.33 basis points, within nineteen per cent of the representative value. Aggregating over the wrong population moves the answer more than aggregating at all does. The dispersion underneath it survives either way, and

for supervisory purposes it is the point. The ninetieth-percentile institution faces 6.23 basis points and the most exposed 12.30, so a modest sector average is consistent with a handful of banks facing an order of magnitude more, and it is those banks whose borrowers would actually experience the channel described here.

6.4 Limitations and Future Research

Limitations

- **Transportability.** The wedge scale comes from firm lending during the Italy-2011 crisis and is applied to household marginal borrowing in 2018. Neri (2013) reports both legs of the same counterfactual for the same episode, 130 basis points on new corporate loans and 60 on residential mortgages in the peripheral countries, and the ratio of roughly 0.46 is the household-versus-firm adjustment the design requires. Applying it gives an 8.1-basis-point household target and a 3.03-basis-point wedge, against 6.56 under the untransported firm anchor, and lowers the bottom-two share to 5.51%. The defensible statement is therefore a *household-transport range* of roughly 3.0–6.6 basis points with the firm anchor at its upper end.
- **Identification and data.** Public aggregates do not identify how much outstanding household debt repriced at the modelled frequency. The model's negative liquid-asset position is a stylised rolling borrowing position, not a claim that mortgage debt repriced. The allocation also depends on four maintained constants, of which the bank-equity exposure ratio is the weakest sourced and is swept explicitly in Appendix C. Aggregate-credit diagnostics are non-confirmatory at every horizon, and the valid pre-event market test is underpowered at thirteen observations.
- **Scope.** The bank block is in reduced form. There are no lender-specific household contracts or reverse sovereign–bank feedback. A small-member-state closure holds the real wage fixed, thereby shutting down the indirect labour-income channel that much of the heterogeneous-agent literature deems dominant.

Future Research

- **The two-asset extension is closest to resolving the one-asset restriction.** Appendix D reports how far it has got. The restriction matters because the incidence exercise allocates a revaluation loss to households who, with a second asset, could rebalance in response to it instead of absorbing it passively. That works against the direct-loss channel and in favour of households liquid enough to buy depressed assets, which is the mechanism Mendicino et al. (2025) identify. The extension confirms that the incidence ranking of the baseline survives a materially richer household problem, so that ranking is not an artefact of holding portfolios fixed. It withholds any claim about magnitude, because the portfolio price is a maintained proxy with no estimated relation behind it, and the channel decomposition has not been re-run.
- **Matching households to lenders would make the mechanism falsifiable.** Instrumenting lender stress in the Italian Credit Register with pre-event domestic sovereign exposure

would allow estimation of the differential change in rates on new and refinanced household loans at more- versus less-exposed banks after May 2018, the household analogue of the same-firm design that anchors this thesis. That estimate would replace the transported calibration with a measured one. The maintained map predicts a rate response that increases with lender sovereign exposure and a loss composition that shifts toward indebted lower-wealth households.

6.5 Conclusion

A wider sovereign spread lowers the value of government debt held by domestic banks. Weaker banks charge more to lend, and it is the households carrying bank debt who pay that. This thesis measures that channel. In Italy in 2018, it runs from a repricing of 166.1 basis points, through 14.7% of the representative banking sector's economic value, to a borrower wedge of between 4.03 and 7.84 basis points. When allocated across SHIW wealth quintiles, the loss remains concentrated at the top, with the top fifth giving up 104 basis points of their consumption, compared to 3 basis points for the bottom fifth. However, its composition tilts the other way. The bottom two-fifths carry 5.99–6.08% of the total, compared to 4.80% from asset prices alone, roughly a quarter more of the burden, while output moves by 0.003%. The episode redistributes far more than it destroys.

The condition behind that shift extends beyond Italy. Bank credit spreads a sovereign loss further down the wealth distribution, reaching borrowers who are less concentrated at the top than the holders of the repriced assets. Italian households bear 4.8% of the asset-price loss and 26.5% of the added borrowing cost, which is a wide margin. This comparison is based on two exposure shares measured directly by a household survey, so any country with one can use it to determine the impact of sovereign stress on its own banking system. France shows what the same arithmetic reveals at the other extreme. In 2024, it had a higher exposure to longer-term debt and higher leverage than Italy, yet it incurred a loss that was only one-fourteenth of Italy's, because it held only about one-fifth of its own government's debt. This is what makes exposure a product. A common currency means that sovereign holdings and household balance sheets remain national, so the same shock has a very different impact on either side of a border. When a government's borrowing costs rise, the bill travels through the banks and reaches households who own very little.

AI-use statement: During the preparation of this thesis, the author used generative-AI writing and coding tools (Claude [Anthropic] and Codex [OpenAI]) to support prose editing, code drafting, repository diagnostics, and checks of numerical outputs against their underlying data and code, under the author's continuous direction and review. The author takes full responsibility for the content, analysis, and conclusions of this thesis.

Bibliography

Acharya, Viral V., Tim Eisert, Christian Eufinger, and Christian W. Hirsch (2018). “Real Effects of the Sovereign Debt Crisis in Europe: Evidence from Syndicated Loans”. In: *Review of Financial Studies* 31.8, pp. 2855–2896.

Acharya, Viral V. and Sascha Steffen (2015). “The "greatest" carry trade ever? Understanding eurozone bank risks”. In: *Journal of Financial Economics* 115.2, pp. 215–236.

Afonso, António, Michael G. Arghyrou, and Alexandros Kontonikas (2015). *The Determinants of Sovereign Bond Yield Spreads in the EMU*. Working Paper Series 1781. European Central Bank.

Albertazzi, Ugo, Tiziano Ropele, Gabriele Sene, and Federico M. Signoretti (2014). “The impact of the sovereign debt crisis on the activity of Italian banks”. In: *Journal of Banking & Finance* 46, pp. 387–402. DOI: 10.1016/j.jbankfin.2014.05.005.

Altavilla, Carlo, Marco Pagano, and Saverio Simonelli (2017). “Bank Exposures and Sovereign Stress Transmission”. In: *Review of Finance* 21.6. Earlier circulated as ECB Working Paper No. 1969 (2016) and ESRB Working Paper No. 11 (2016), pp. 2103–2139.

Auclert, Adrien (2019). “Monetary Policy and the Redistribution Channel”. In: *American Economic Review* 109.6, pp. 2333–2367. DOI: 10.1257/aer.20160137.

Auclert, Adrien, Bence Bardóczy, Matthew Rognlie, and Ludwig Straub (2021). “Using the Sequence-Space Jacobian to Solve and Estimate Heterogeneous-Agent Models”. In: *Econometrica* 89.5, pp. 2375–2408.

Battistini, Niccolò, Marco Pagano, and Saverio Simonelli (2014). “Systemic Risk, Sovereign Yields and Bank Exposures in the Euro Crisis”. In: *Economic Policy* 29.78, pp. 203–251.

Bayer, Christian, Alexander Kriwoluzky, Gernot J. Müller, and Fabian Seyrich (2024). “A HANK^2 Model of Monetary Unions”. In: *Journal of Monetary Economics* 147. Supplement, p. 103579. DOI: 10.1016/j.jmoneco.2024.103579.

Bocola, Luigi (2016). “The Pass-Through of Sovereign Risk”. In: *Journal of Political Economy* 124.4, pp. 879–926.

Bofondi, Marcello, Luisa Carpinelli, and Enrico Sette (2018). “Credit Supply During a Sovereign Debt Crisis”. In: *Journal of the European Economic Association* 16.3, pp. 696–729. DOI: 10.1093/jeea/jvx020.

Cameron, A. Colin, Jonah B. Gelbach, and Douglas L. Miller (2008). “Bootstrap-Based Improvements for Inference with Clustered Errors”. In: *The Review of Economics and Statistics* 90.3, pp. 414–427. DOI: 10.1162/rest.90.3.414.

Carroll, Christopher D. (2006). “The Method of Endogenous Gridpoints for Solving Dynamic Stochastic Optimization Problems”. In: *Economics Letters* 91.3, pp. 312–320.

Corradin, Stefano and Bernd Schwaab (2023). “Euro Area Sovereign Bond Risk Premia Before and During the COVID-19 Pandemic”. In: *European Economic Review* 153, p. 104402. DOI: 10.1016/j.euroecorev.2023.104402.

De Marco, Filippo (2019). “Bank Lending and the European Sovereign Debt Crisis”. In: *Journal of Financial and Quantitative Analysis* 54.1, pp. 155–182. DOI: 10.1017/S0022109018000510.

Equiza-Goñi, Juan, Elisa Faraglia, and Rigas Oikonomou (2023). “Union debt management”. In: *Journal of International Money and Finance* 130, p. 102747. DOI: 10.1016/j.jimonfin.2022.102747.

European Banking Authority (2017). *2017 EU-Wide Transparency Exercise.* Bank-by-bank credit-risk disclosure, reference dates 31 December 2016 and 30 June 2017, 132 banks. European Banking Authority, Paris.

European Central Bank (2022). *The Transmission Protection Instrument.* ECB press release, 21 July 2022. https://www.ecb.europa.eu/press/pr/date/2022/html/ecb.pr220721~973e6e7273.en.html.

Faccini, Renato, Seungcheol Lee, Ralph Luetticke, Morten O. Ravn, and Tobias Renkin (2026). “Financial Frictions: Micro versus Macro Volatility”. In: *American Economic Review* 116.2, pp. 464–501. DOI: 10.1257/aer.20211219.

Farhi, Emmanuel and Jean Tirole (2018). “Deadly Embrace: Sovereign and Financial Balance Sheets Doom Loops”. In: *Review of Economic Studies* 85.3, pp. 1781–1823.

Gennaioli, Nicola, Alberto Martin, and Stefano Rossi (2014). “Sovereign Default, Domestic Banks, and Financial Institutions”. In: *Journal of Finance* 69.2, pp. 819–866.

Gerali, Andrea, Stefano Neri, Luca Sessa, and Federico M. Signoretti (2010). “Credit and Banking in a DSGE Model of the Euro Area”. In: *Journal of Money, Credit and Banking* 42.s1, pp. 107–141.

Giannetti, Mariassunta, Martina Jasova, Caterina Mendicino, and Dominik Supera (2026). *Securities Losses and the Bank Collateral Channel of Monetary Transmission.* Working Paper Series 3209. European Central Bank.

Grandi, Pietro (2019). “Sovereign Stress and Heterogeneous Monetary Transmission to Bank Lending in the Euro Area”. In: *European Economic Review* 119, pp. 251–273. DOI: 10.1016/j.euroecorev.2019.07.011.

Handler, Lukas and Rainer Jankowitsch (2025). “Political Uncertainty and Sovereign Bond Markets”. In: *Financial Markets and Portfolio Management* 39, pp. 47–97.

Hazell, Jonathon and Stephan Hobler (2025). *Do Deficits Cause Inflation? A High Frequency Narrative Approach.* CEPR Discussion Paper No. 19928. Centre for Economic Policy Research.

Kaplan, Greg, Benjamin Moll, and Giovanni L. Violante (2018). “Monetary Policy According to HANK”. In: *American Economic Review* 108.3, pp. 697–743. DOI: 10.1257/aer.20160042.

Kaplan, Greg and Giovanni L. Violante (2014). “A Model of the Consumption Response to Fiscal Stimulus Payments”. In: *Econometrica* 82.4, pp. 1199–1239.

Kopecky, Karen A. and Richard M. H. Suen (2010). “Finite State Markov-Chain Approximations to Highly Persistent Processes”. In: *Review of Economic Dynamics* 13.3, pp. 701–714. DOI: 10.1016/j.red.2010.02.002.

Krishnamurthy, Arvind, Stefan Nagel, and Annette Vissing-Jorgensen (2018). “ECB Policies Involving Government Bond Purchases: Impact and Channels”. In: *Review of Finance* 22.1, pp. 1–44.

Mendicino, Caterina, Lukas Nord, and Marcel Peruffo (2025). *Distributive Effects of Banking Sector Losses*. Working paper. Revised April 2025. SSRN. DOI: 10.2139/ssrn.5012454.

Neri, Stefano (2013). *The Impact of the Sovereign Debt Crisis on Bank Lending Rates in the Euro Area*. Questioni di Economia e Finanza (Occasional Papers) 170. Banca d’Italia.

Pástor, Ľuboš and Pietro Veronesi (2013). “Political Uncertainty and Risk Premia”. In: *Journal of Financial Economics* 110.3, pp. 520–545.

Rojas, Luis E. and Dominik Thaler (2024). “The Bright Side of the Doom Loop: Banks’ Sovereign Exposure and Default Incentives”. In: *European Economic Review* 170, p. 104876. DOI: 10.1016/j.euroecorev.2024.104876.

Roldán, Francisco (2025). “The Aggregate-Demand Doom Loop: Precautionary Motives and the Welfare Costs of Sovereign Risk”. In: *American Economic Journal: Macroeconomics* 17.3, pp. 160–204. DOI: 10.1257/mac.20220248.

Sabatini, Fabiana (2022). *The Impact of Sovereign Tensions on Bank Lending: Identifying the Channels at Work*. Temi di Discussione 1397. Banca d’Italia.

Slacalek, Jiri, Oreste Tristani, and Giovanni L. Violante (2020). “Household Balance Sheet Channels of Monetary Policy: A Back of the Envelope Calculation for the Euro Area”. In: *Journal of Economic Dynamics and Control* 115, p. 103879. DOI: 10.1016/j.jedc.2020.103879.

Theobald, Thomas and Silke Tober (2020). “Euro Area Sovereign Yield Spreads as Determinants of Private Sector Borrowing Costs”. In: *Economic Modelling* 84, pp. 27–37. DOI: 10.1016/j.econmod.2019.03.004.

Uribarri Sánchez, Jauder (2026). *KS_Models: Reusable Krusell–Smith Numerical Methods and Solver-Validation Tools*. Version 0.1.1. Python package (MIT licence, distributed on PyPI as `ks-models`) and companion numerical-methods repository.

CHAPTER A

Numerical Validation of the Baseline Model

This appendix documents the numerical evidence for the one-asset heterogeneous-household benchmark used in the thesis. It does not document the development history of the project or validate the exploratory two-asset and heterogeneous-bank extensions retained in the private replication archive. The baseline stationary household problem is solved by endogenous-grid methods (EGM), and its local aggregate transition is solved with sequence-space Jacobians (Carroll, 2006; Auclert et al., 2021). All stochastic routines use seed 33.

A.1 Implementation

The stationary problem uses a 100-point liquid-asset grid on $[-0.5, 150]$ and a three-state Rouwenhorst income process. The discount factor and steady-state output jointly match the aggregate liquid-asset target and borrower-participation moment. EGM iterates on marginal value to a sup-norm tolerance of 10^{-6}; the stationary distribution is iterated to the same tolerance. Around that stationary economy, the fake-news algorithm constructs household Jacobians for the borrower-wedge path, and the linear equilibrium system delivers the transition. The baseline fixes the real wage, treats liquid assets as a diagnostic and not a period-by-period clearing target, and is a first-order local solution, not a nonlinear perfect-foresight model.

The numerical state and the economic scope should not be conflated. The asset grid permits a borrowing region and a saving region, while the income process supplies idiosyncratic risk; the stationary distribution determines the mass at each occupied state. The transition then evaluates how a small, time-varying borrower wedge changes policies and aggregates around that stationary distribution. It does not solve a nonlinear path with occasionally binding aggregate constraints, nor does it make the asset-market diagnostic an additional market-clearing equation. These choices are part of the small-member-state closure described in Chapter 3, and the checks below assess accuracy for precisely that implementation.

Reproducibility and validation criteria

A baseline replication run fixes the calibration, grid, horizon, seed, and input hashes, and writes the numerical diagnostics before any manuscript-facing table or figure is accepted. Exact checks and tolerance checks are kept apart. Accounting identities, no-shock invariance, and seed-dependent outputs must agree to machine precision. Derivative, grid, horizon, and cross-solver comparisons are judged against declared tolerances, because they compare different approximations to the same object. The gates run from elementary identities upward. The sovereign-loss identity is checked before the loss enters the recovery equation, the stationary policy is checked for feasibility and Euler optimality, and only then is the sequence-space derivative compared with a finite-difference re-solve and EGM with the independent VFI code. That order localises failures, since an accounting mismatch cannot be repaired by a finer grid.

Two boundaries govern what these checks establish. First, the reported output response comes from the production EGM/sequence-space configuration, never from the VFI cross-check or a sensitivity grid; the finite-difference calculation is not averaged with the Jacobian but used to test it. Second, reproducibility is not economic validation. Re-running identical code with the same seed shows the published object is recoverable from documented inputs. It says nothing about whether the one-asset closure, the Italy-2011 calibration, or the partial-pre-event bridge are right. Some quantities are also deliberately *not* validation targets. The model does not target the distribution of individual Italian loan contracts, a realised credit impulse response, or clearing in an omitted foreign asset market. The non-targeted liquid-asset diagnostic is reported as such so it cannot be mistaken for an equilibrium condition.

A.2 Validation

Numerical validation addresses implementation accuracy conditional on the maintained model; it cannot resolve the economic calibration or identification limitations discussed in the main text. Table A.1 reports feasibility, optimality, and no-shock diagnostics. Budget residuals are at machine precision, Euler residuals are small over occupied states, and the no-shock transition is exactly zero.

Check	Metric	Interpretation
No-shock invariance	max IRF $= 0.00$	No artificial transition
Sovereign-loss accounting	max relative error $= 1.30 \times 10^{-16}$	$\omega DL\Delta s$ identity
Steady-state calibration	borrower error $= 0.675$ pp; asset residual $= 1.6 \times 10^{-14}$	Targets met
Budget residuals	max $= 1.1 \times 10^{-16}$	Feasible occupied states
Euler residuals	median $= 7.5 \times 10^{-7}$; p99 $= 4.9 \times 10^{-6}$	Household optimality
Constraint diagnostics	0 violations; 219 occupied states	Complementarity and feasibility

Notes: Euler residuals are relative deviations between the two sides of the consumption Euler equation, evaluated at the solved policy on occupied states.

Table A.1. Baseline-model household and equilibrium diagnostics.

The diagnostics are deliberately heterogeneous. A no-shock check establishes that the transition machinery does not create motion without an input. The accounting check verifies the deterministic sovereign-loss identity before it enters the household block. Budget and constraint checks establish feasibility at occupied household states, whereas Euler residuals assess local optimality of the consumption policy. None of these checks validates the Italy-2011 calibration or the interpretation of a borrower wedge as a household contract; those are economic evidence questions addressed separately in the main text and robustness appendix.

Sequence-space derivative check

The central local-solution test compares a column of the sequence-space household Jacobian with a direct finite-difference re-solve under a perturbation to the annualised borrower wedge. Here e_s is the unit vector that perturbs only date s, and $\varepsilon > 0$ is the scalar perturbation size.

$$J^{Y,x}_{\cdot s} \quad \text{against} \quad \frac{Y(x+\varepsilon e_s) - Y(x)}{\varepsilon}. \tag{A.1}$$

For economically informative perturbations of 10^{-4} and 10^{-3}, the maximum relative discrepancy is 0.90% for aggregate consumption and output, 0.46% for borrower consumption, and 0.09% for assets. Figure A.1 shows the active-path comparison. This supports the derivative used in the linear equilibrium inversion; it is not a claim that the model is globally linear.

The finite-difference calculation has a specific diagnostic purpose. It perturbs one element of the wedge sequence, re-solves the household block, and compares the resulting path with the corresponding Jacobian column. A perturbation that is too small is dominated by numerical round-off; one that is too large tests nonlinear curvature instead of the derivative. The reported range is therefore chosen to be economically small while still producing a stable re-solved response. Agreement over the active horizon is the relevant criterion because the baseline uses the Jacobian only inside the local equilibrium system.

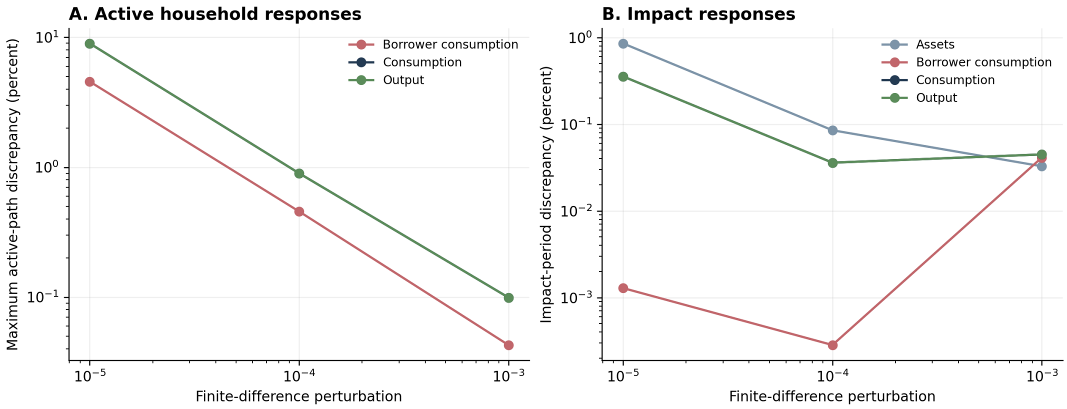

Figure A.1. Finite-difference validation of sequence-space household Jacobians. The comparison uses perturbations in the economically informative range 10^{-4} to 10^{-3}.

Resolution, horizon, and independent-solver checks

Changing the asset grid from 70 to 150 points moves the Italy-2018 impact output response only from −0.00285% to −0.00297%. Extending the transition horizon from 60 to 120 periods changes the peak by 4.57% and the cumulative response by 1.45%. The non-targeted asset-market diagnostic has a maximum absolute deviation of 0.00305, normalised by steady-state output; this is a consequence of the stated small-member-state closure, and no market-clearing condition is claimed.

The grid comparison is useful because it varies the resolution of the borrowing region where the wedge matters most; the preserved sign and close impact responses indicate that the reported aggregate result is not an artefact of the production grid. The horizon comparison tests whether the finite transition truncates the bank-recovery path too early. The modest movement in peak and cumulative responses is recorded as a numerical sensitivity, not averaged into a preferred result. Incidence shares are computed from fixed observed vectors after the model transition and thus are not an additional claim about grid convergence of unobserved household contracts.

An independently coded grid-search VFI solver provides a separate stationary-policy check. It lives outside this thesis's codebase, in `ks-models`, a public MIT-licensed Python library that packages the Krusell–Smith grid, Markov, and solver-validation utilities developed in a companion methods project (Uribarri Sánchez, 2026). An optional adapter in that library calls the production model once, extracts its exact stationary household environment, and hands it to the library's own value-function solver, so Table A.2 is generated by code that shares no solution routine with the model it checks. The companion note, *Three Numerical Implementations of the Krusell–Smith Model: KS-VFI, KS-HARK, and KS-NN*, documents those methods on the frictionless Krusell–Smith economy; it is a methodological companion and validates none of the political-risk results reported here.

As its asset grid is refined, the VFI policy and distribution converge toward the production EGM allocation. Refinement sharply reduces both the aggregate-asset discrepancy and the weighted policy gap, and the finest 2,000-point grid leaves a 1.14% aggregate-asset difference with borrower and constrained masses close to the EGM reference. The residual gap is consistent with the VFI's on-grid choice restriction near the borrowing constraint. Convergence from two independent solution routines is stronger evidence of a shared stationary solution than a second run of the same solver would be. It remains a stationary cross-check, and does not independently solve the sequence-space transition.

Grid n_a	A^{VFI}	A^{EGM}	Relative A error	Borrower mass	Weighted policy gap
100	1.719	4.000	57.03%	27.98%	0.10656
500	3.774	4.000	5.64%	25.16%	0.01262
2000	3.954	4.000	1.14%	24.88%	0.00354

Notes: The EGM reference borrower mass is approximately 24.24%. The VFI restricts next-period assets to grid nodes, whereas EGM uses an interpolated continuous policy.

Table A.2. Independent VFI convergence toward the production EGM stationary allocation.

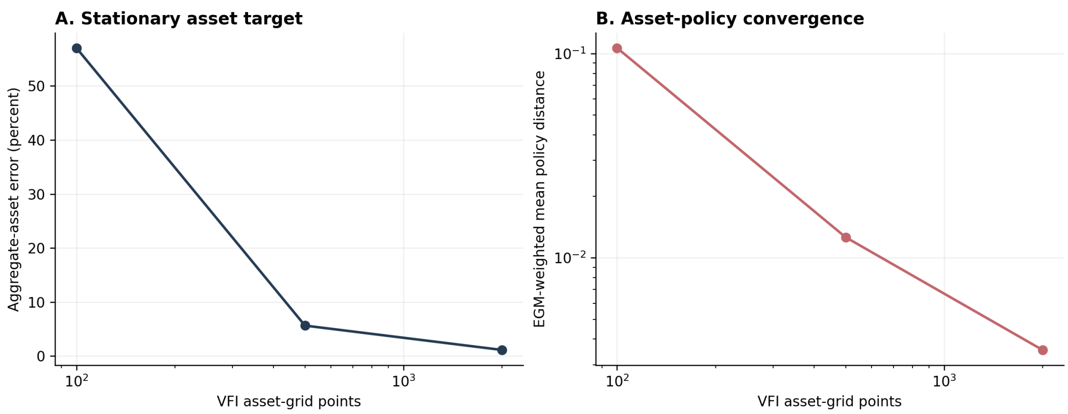

Figure A.2. Convergence of the independent grid-search VFI solution toward the production EGM allocation under grid refinement.

Untargeted household-model moments

Only the borrower share is targeted, to SHIW bank-debt participation (Section 4.3). Table A.3 reports the remaining moments of the model's own stationary distribution, none of which enters the calibration.

Moment	Value	Status
Borrower share (mass with $a < 0$)	24.24%	Targeted (SHIW: 24.91%)
Mass at the borrowing constraint	12.49%	Untargeted
Mean debt conditional on borrowing, $\mathbb{E}[a \mid a < 0]$	-0.384	Untargeted
Mean liquid assets, all households	4.00	Untargeted (equals B by construction)
Median liquid assets, all households	1.36	Untargeted
Top-20% share of aggregate net liquid wealth	66.6%	Untargeted
Gini coefficient, positive-asset holders only[a]	0.544	Untargeted

Table A.3. Household steady-state moments: targeted and untargeted

[a]A standard Gini coefficient requires non-negative wealth; since 24% of the stationary distribution is in debt, the coefficient is computed only over the renormalised positive-asset sub-population (76% of households), and is never left ill-defined or silently clipped.

The model concentrates liquid wealth more than SHIW concentrates net wealth, consistent with a one-asset household block matched only to a debt-participation target and never to the joint distribution of assets and debt. As a stability check, re-solving the steady state at a more conventional annualised real rate (2.43%, against the baseline's 6.17%, with β re-solved to clear the same market-clearing targets) moves the borrower share by -0.17 percentage points and the constrained share by $+0.06$ percentage points, both an order of magnitude smaller than the differences the incidence exercise studies, so the qualitative household composition is not an artefact of the baseline real-rate calibration.

Together, these checks support the numerical implementation of the baseline at its intended local scope. They do not validate the transported Italy-2011 credit calibration, observed household contract repricing, or the exploratory extensions.

CHAPTER B

Data Sources, Sample Construction, and Measurement

This appendix documents the data, not the findings. Table B.2 does most of the work; the sections that follow are organised around its five panels and add only the measurement decisions a table cannot carry.

B.1 Data Architecture and Processing Conventions

No single source spans the transmission chain: sovereign-market data provide the upstream shock, EBA disclosures and ECB sector aggregates map it into bank net-worth losses, listed-bank market outcomes test whether the exposure gradient is visible in prices, and household microdata allocate the credit tightening across the wealth distribution. The conventions below distinguish the structural country mapping from the timing-sensitive market validation and record the remaining measurement choices.

- *Vintage and role.* The pre-event bank-market check uses strictly pre-event inputs (the 30 June 2017 EBA Transparency Exercise disclosure (European Banking Authority, 2017) and pre-event controls from ORBIS, the Bureau van Dijk bank-financial-statement database); the country-level structural mapping instead uses the closest usable episode vintage, which can post-date the acute window—most notably the June 2018 EBA disclosure for the Italy-2018 country mapping. An earlier draft used the 31 December 2017 disclosure for the bank-market test; that vintage was withdrawn (Section B.5 gives the full provenance finding).
- *Units.* Yields in percent, spreads in basis points, bank amounts in millions of euro, household amounts in thousands. Wedges are annualised though the model is solved monthly.
- *Weighting.* Household statistics use survey weights; bank statistics are unweighted, the panel being a census of matchable institutions and not a sample.
- *Missing observations.* Market series are never interpolated: a non-trading target date is matched to the nearest market date. Bank controls are complete-case, never imputed.
- *Computed versus cited.* A number is data only if computed here from a raw source. The Italy-2011 lending-rate target and 350-basis-point movement are calibration targets from published work.

B.2 Reproducibility Records and Data-Access Boundary

The private submission package contains only the inputs and derived objects needed to reproduce the documented calculations; it is not a redistribution of commercial databases, restricted survey microdata, or supervisory source files. Each documented input is assigned a source, reference date, transformation script, access classification, schema description, and SHA-256 hash in the data manifest. The corresponding output and parameter manifests record the replication seed, the generating command, and the hash of every table and figure in the manuscript. These records make the research workflow auditable without implying that every underlying source can be shared publicly.

Three distinctions govern the package. First, a cited source is a published paper or public document whose result may discipline a calibration; it need not be an input file. Second, a documented input is an object held inside the replication directory and used by a replication script. Its inclusion is conditional on the licence and on its containing no redistributable raw or respondent-level material. Third, a derived output is a reproducible result of a replication script. It can be retained for audit even when its underlying raw source is licensed, provided the output itself does not disclose protected observations or recreate the source database. The separation prevents a reference, an input, and a result from being treated as interchangeable evidence.

Class	Examples	Included in private package?	Permitted role
Published reference	Italy-2011 lending study, EBA methodology	Citation or documentation only	External calibration or measurement definition
Raw/licensed source	Bloomberg prices, ORBIS controls, survey micro-records	No	Rebuilt only in the authorised private workspace
Documented derived input	Quintile vectors; June-2017 bank-validation panel	Yes, with hash and codebook	Baseline in-release calculation
Derived audit output	Tables, figures, diagnostics, validation reports	Yes, with generating command	Verification of reported manuscript objects

Notes: The separate public GitHub export contains no data or results. It documents access requirements and provides model-only smoke tests; exact empirical reproduction remains a private-submission activity.

Table B.1. Data and output classes in the private submission package.

The hash record is not a substitute for provenance. A matching hash verifies that a particular file has not changed, whereas the source date and transformation description establish what the file represents. Both are needed for timing-sensitive objects. This is especially important for the Italy-2018 bridge, where a later, better-populated disclosure is informative but must not be labelled pre-event. The manifests consequently record both the observation date and the date on which a source was retrieved or transformed. A change in either the data definition or the documented processing code requires a new record, and never an informal replacement of the old file.

The household allocation illustrates the same discipline. The release retains a compact quintile-level representation in place of respondent-level SHIW observations. A reconstruction record re-computes the committed incidence table from that documented representation and checks byte-for-byte equality before the result is used. This validates the reproducibility of the published accounting object, not the availability of the underlying microdata to an unrestricted user. The codebook records the remaining aggregation loss: the release cannot recover household-level lender links, contract fixation, or longitudinal refinancing decisions from the released quintile vectors.

The access boundary is therefore part of the scientific claim. A reader can determine which results are reproducible from the supplied private inputs, which rely on a published calibration, and which would require an authorised source-data reconstruction. No replication script silently searches outside the replication directory or downloads a substitute when a documented input is absent. It stops with an explicit data-availability message. Stopping is preferable to a superficially complete run that uses an undocumented local path or a different vintage of the data.

Complete dataset inventory, contents, analysis samples, and role in the thesis

Dataset and provider	Contents, coverage, and frequency	Unit of observation and analysis sample	Role in the thesis	Access
Panel A. Sovereign yields and event spreads				
Generic ten-year sovereign yields (Bloomberg)	Daily closes for Germany (`GDBR10`), Italy (`GBTPGR10`) and France (`GFRN10`)	Country-day. $s_{c,t} = y^{10y}_{c,t} - y^{10y}_{DE,t}$; nearest-date matched	Acute window: 114.4 bp (Italy 2018), 31.4 bp (Italy 2022), 14.9 bp (France 2024); sets bank-market event windows	Bloomberg terminal; licensed
Long-term government-bond yields, OECD via FRED	Ten-year benchmark yields, monthly: `IRLTLT01ITM156N`, `IRLTLT01DEM156N`, `IRLTLT01ESM156N`; retrieved 8 July 2026	Country-month. Baseline mean vs. peak month	Model-input movements: 166.1 bp Italy 2018, 15.5 bp Italy 2022; Spanish series retained for a private-archive robustness exercise	Public
Episode spread series (project-assembled)	Daily OAT/Bund/BTP closes with derived spreads for France 2024, plus France 2017 and Netherlands 2023	Country-day, one file per episode	France-2024 input 34.1 bp; historical France-2017 and Netherlands-2023 diagnostics retained in the private archive	Derived
Panel B. Bank balance sheets and sovereign exposures				
EBA 2011 EU-wide stress test	Sovereign holdings by residual-maturity bucket and Core Tier 1 capital, 2011 vintage	Bank-vintage; Italian banking sector	Duration (4.397y) and leverage (25.21) for the Italy-2011 bridge disciplining ϕ_0	Public (EBA)
EBA Transparency Exercises 2017, 2018, 2022, 2024	Sovereign templates by counterparty country and residual maturity, own-funds and leverage templates	Bank-vintage; closest usable vintage per episode; 30 June 2017 for the replacement market-validation panel	Country-mapping duration/leverage (Italy 2018: $D = 4.95$, $L = 17.60$)	Public (EBA)
ECB Balance Sheet Items (BSI)	MFI holdings of domestic general-government debt and total assets, DE/ES/FR/IT/NL, monthly	Country-sector-month; derived ω^{dom}	Sector ω^{dom}_e (Italy Jun. 2018: 10.15%; France 2024: 2.11%); used because the EBA panel inverts the France–Italy ranking	Public (ECB)
ORBIS / Orbis BankFocus	Annual bank financial statements, 24 listed euro-area banks, last fiscal year before each event	Bank-year; 23 with total assets, 19 with non-performing loan (NPL) ratio	Pre-event controls; complete coverage for the June-2017 strict specification is 13 banks	Licensed (BvD)
Panel C. Bank market outcomes				
Listed-bank equity prices (Bloomberg)	Daily closes, listed euro-area banks, estimation and event windows	Bank-day; 25-bank panel (7 IT, 5 ES, 3 FR, 3 AT, 2 DE, 2 NL, 2 IE, 1 BE)	Beta-adjusted abnormal returns, dependent variable of the exposure-gradient regression; Italian banks average −7.25%	Bloomberg; licensed
Market indices (Bloomberg)	Daily closes: Euro STOXX (`SX5E`), Euro STOXX Banks (`SX7E`), FTSE MIB, CAC 40	Index-day	Benchmark for raw and beta-adjusted abnormal returns	Bloomberg; licensed
Bank CDS spreads (Bloomberg)	Five-year senior CDS mid spreads, ten banks	Bank-day; 7 with a usable change over the acute window	Corroborating outcome; thin coverage keeps it at appendix weight	Bloomberg; licensed

Continued on the next page

Table B.2. Complete dataset inventory *(continued)*

Dataset and provider	Contents, coverage, and frequency	Unit of observation and analysis sample	Role in the thesis	Access
Yahoo Finance equity/index history	Daily adjusted prices, 29 tickers and two indices, JSON	Bank-day	Independent source for Spain-2017, Italy-2022 and exposure-sorted exercises; not the headline 25/19-bank panels	Public
Panel D. Household microdata				
SHIW (Banca d'Italia)	Survey on Household Income and Wealth, 2022 wave; 2016 and 2020 waves used for the vintage robustness of Section C.5	Household records, weighted; 9,641 households	Disciplines borrower mass against the 24.91% bank-debt participation target; supplies quintile vectors	Public on request
EFF (Banco de España)	Encuesta Financiera de las Familias, 2022 wave, implicate 1	Household records, weighted; 6,385 households	Private-archive robustness input; not used in the submitted PDF	Public on request
Household Finance and Consumption Survey (HFCS, ECB)	Published statistical tables, Wave 2021 (July 2023 release)	Published wealth-group aggregates, not micro records	Household side of France inside the structural model (Chapter 5)	Public
Luxembourg Wealth Study (LWS), France 2020	Four LISSY remote-execution jobs: net worth, bond exposure, incidence, capital-to-consumption	Aggregate LISSY output only	Household moments for reduced-form French illustrations; not substituted for HFCS in the structural model	Registered access (LIS)
Panel E. Macro-financial controls and reference documents				
ECB Balance Sheet Items, credit series	Monthly outstanding MFI loans to households and NFCs	Country-month; donor pool AT, BE, DE, FI, IE, PT	Donor-pool counterfactual, synthetic control, panel local projection (Appendix C.2)	Public (ECB)
Eurostat HICP inflation	Monthly HICP rate of change (`prc_hicp_manr`)	Country-month	Inflation control, local-projection specification	Public
Eurostat consumer confidence	Monthly indicator (`ei_bsco_m`)	Country-month	Demand-sentiment control, panel local projection	Public
ECB policy-rate history	Dated deposit-facility-rate changes	Event-date list	Separates July-2022 tightening from the political component of Italy 2022	Public (ECB)
Supervisory and methodological documents	Banca d'Italia Financial Stability Reports 1/2018, 2/2018, 2/2022; EBA template guides	Published reports	Cross-checks on sector duration/exposure; EBA template field definitions	Public

Table B.2. Complete dataset inventory, contents, analysis samples, and role in the thesis

Sources are not samples: Bloomberg is not the bank-validation sample, which is the set of listed banks matchable to a pre-event EBA exposure and to ORBIS controls.

B.3 Panel A. Sovereign Yields and Event Spreads

The spread is $s_{c,t} = y_{c,t}^{10y} - y_{DE,t}^{10y}$. Two independent yield sources are used at different frequencies. Daily Bloomberg generic indices, preferred to individual benchmark bonds so the series does not jump when the on-the-run bond is replaced, measure the acute reaction and set the event-study windows. Monthly OECD yields distributed through FRED measure the broader baseline-to-peak repricing that enters the model, reporting averages and not point-in-time quotes, so a gap of ten to thirty basis points against any single-day figure is expected – deliberate, since the Italian episodes unfold over months, not days. Italy 2011 is not reconstructed as a series; it enters as a calibration bridge whose 350-basis-point movement is triangulated from the identification window of Bofondi et al. (2018).

Episode	Window	Ten-year yields	Spread movement
Italy 2018 acute bank window	15–31 May 2018	Italy: 1.954% to 2.794%; Bund: 0.645% to 0.341%	130.9 to 245.3 bp; change 114.4 bp
Italy 2022/TPI window	20 Jul.–26 Sep. 2022	Italy: 3.391% to 4.563%; Bund: 1.257% to 2.115%	213.4 to 244.8 bp; change 31.4 bp
France 2024 snap election	7 Jun.–8 Jul. 2024	France: 3.098% to 3.167%; Bund: 2.620% to 2.540%	47.8 to 62.7 bp; change 14.9 bp

Table B.3. Daily sovereign-yield windows used for market measurement

Reporting both legs separately shows the Italian and French movements are not driven by the same mechanics: in May 2018 the Italian leg rose while the Bund leg fell, whereas in June 2024 the French leg barely moved and most of the widening came from the Bund.

Episode	Role	Δs_e	ω^{dom}	D	L	ℓ_e
Italy 2011	External calibration bridge	350.0 bp	0.0811	4.397	25.21	31.46%
Italy 2018	Partial-pre-event duration sensitivity	166.1 bp	0.0900	3.86–6.88	16.58	9.57–17.06%
Italy 2018	June-2018 episode-vintage proxy	166.1 bp	0.1015	4.95	17.60	14.69%
France 2024	Low-exposure check	34.1 bp	0.0211	7.28	20.23	1.06%
Italy 2022/TPI	Policy/backstop contrast	15.5 bp	0.1027	5.07	18.45	1.49%

Table B.4. Event inputs for the bank-loss bridge

Table B.4 belongs here because $\ell_e = \omega_e^{dom} D_e L_e \Delta s_e/10{,}000$ is an accounting transformation of measured quantities, not a model output. The first Italy-2018 row is partial pre-event: exposure and leverage are dated December 2017 but duration is a June-2018 sensitivity value; the second is retained solely as the closest episode-vintage proxy. France 2024 has the longest duration of the four but the smallest loss, its exposure share being roughly one fifth of Italy's.

B.4 Panel B. Bank Balance Sheets and Sovereign Exposures

The model uses three country-level inputs, deliberately not all from the same source. Duration and leverage come from EBA disclosures. Duration, a maturity-bucket weighted average with the open-ended bucket assigned fifteen years, is computed identically for 2011 and the application episodes. Leverage is not: the application episodes use exposure over Tier 1 capital from the later EBA templates, while the 2011 vintage only reports total assets over Core Tier 1 capital, so the 2011 bridge uses that narrower pairing. Each application is paired with the closest usable EBA episode vintage: Italy 2022 uses June 2022, and France 2024's nearest disclosure predates the event by roughly two years. Italy 2018 receives separate treatment: the dated-bridge timing table uses December-2017 ECB BSI exposure and EBA leverage where possible, records the June-2018 duration sensitivity explicitly, and retains the all-June mapping only as an episode-vintage proxy. The December-2017 EBA template does not disclose residual maturity, so this release contains no fabricated or falsely labelled strictly pre-event Italy-2018 numerical bridge; the June fields are never pre-event inputs to the May-2018 market test. ORBIS supplies the pre-event controls, and its coverage sets the sample size: the NPL ratio exists for 19 of the 25 banks.

Domestic sovereign exposure is measured from ECB Balance Sheet Items as MFI holdings of domestic general-government debt over total assets of the same reporting sector, and not from the EBA panel, because the EBA panel alone inverts the France–Italy exposure ranking relative to the whole banking sector (large diversified banks understate domestic concentration to different degrees across countries), and that contrast is one of the thesis's main results. Sector-wide duration and leverage are unavailable at comparable vintages, so EBA aggregates remain the proxy for those two inputs only.

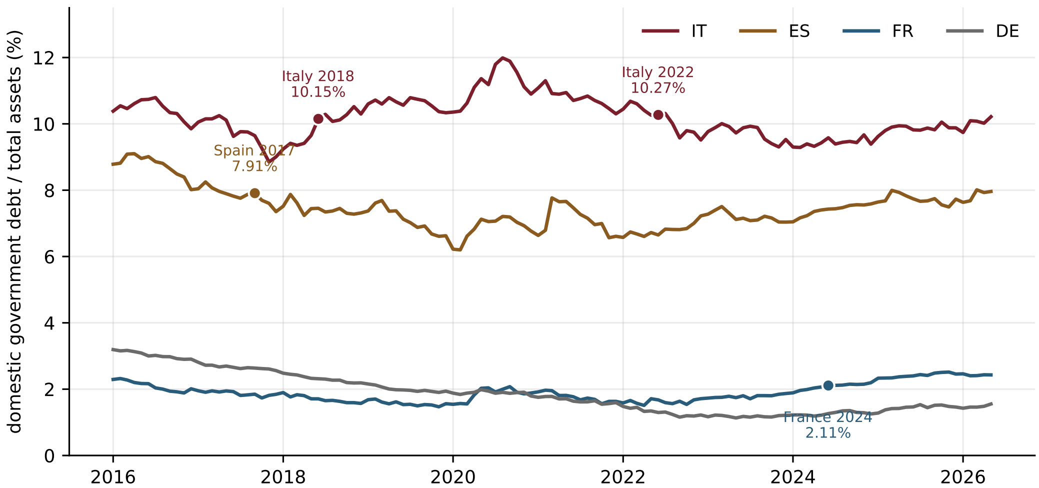

Figure B.1. Domestic sovereign exposure of the banking sector, ECB Balance Sheet Items, January 2016 to May 2026. Each line is holdings of domestic general-government debt securities as a share of total assets of the reporting sector. Markers show the episode-specific observations paired with the structural country mapping; the separate pre-event bank-market check uses strictly pre-event inputs.

Two features of Figure B.1 matter. Italian and Spanish banks hold domestic government debt at roughly five times the French share throughout the period, so the France–Italy contrast is a persistent structural difference and no vintage artefact. And the Italian series shows no discrete jump at either Italian event date, so the exposure input is not reacting to the announcement itself. It does drift over the surrounding months, and that drift is material: the December-2017 and June-2018 readings used in the dated bridge are 9.00% and 10.15% of total assets, a 12.8% relative difference over six months, alongside leverage moving from 16.58 to 17.60. This is precisely the vintage uncertainty that Table B.4 carries explicitly and that the partial-pre-event bridge in Chapter 5 is designed to expose it instead of averaging it away.

B.5 Panel C. Bank Market Outcomes

The listed-bank panel is separate from the country-level inputs and does not feed the structural model. Its function is to test whether banks with larger *pre-event* predicted sovereign losses suffered weaker market outcomes in the acute window. Returns are log differences of Bloomberg closing prices; abnormal returns are formed against the Euro STOXX Banks benchmark, raw and beta-adjusted, with beta estimated over a window ending before the event opens. The headline specification in Chapter 5 uses country-level Italian debt securities and CET1 from the EBA's 2017 EU-wide Transparency Exercise, reference date 30 June 2017, together with the 114.4-bp May spread movement and a common duration proxy. More detailed June-2018 EBA positions are retained only for supplementary diagnostics because they post-date the May event. Yahoo Finance provides a second, public price source used only outside the headline panel, so no coefficient mixes price conventions across providers.

A withdrawn December-2017 vintage. An earlier draft used the 31 December 2017 disclosure from the *2018* Transparency Exercise instead. That template's country-broken-down debt-securities item (1820806) is populated for only a minority of the 131 reporting institutions, and reads zero for every Italian bank in the 25-bank panel; the corresponding CET1 field is likewise incomplete. Because the regressor cannot be reconstructed from the December-2017 disclosure, that draft specification is unavailable and reports no result. The EBA's separate *2017* Transparency Exercise (reference dates December 2016 and June 2017) reports the same sovereign template fully populated, so the June-2017 reference date is used instead: eleven months before the event window, with no look-ahead relative to the withdrawn design's intent. The withdrawn coefficient is retained only in the internal provenance record, never as evidence.

Construction stage	Banks
Listed euro-area banks matched to Bloomberg equity and EBA exposure	25
With usable beta-adjusted equity return	25
With June-2017 event-sovereign debt and CET1>0	21
With log total assets from ORBIS/BankFocus	19
With NPL ratio from ORBIS/BankFocus	16
Complete strictly pre-event controlled regression sample	13
With CDS change in the Italy 2018 acute window	7

Table B.5. Italy 2018 bank-panel sample construction, June-2017 EBA vintage

Every loss in Table B.5 comes from balance-sheet or control coverage, not from a screen on the outcome. Table B.6 defines each constructed variable; several are ratios whose denominator differs by source: exposure is scaled by the leverage measure at bank level but by the sector balance sheet at country level.

Variable	Definition	Unit	Timing
Domestic sovereign exposure	Domestic general-government securities or claims divided by leverage exposure or sector balance-sheet denominator	Ratio	Episode-paired BSI vintage for country mapping
Event-sovereign debt	Country-level debt securities used with CET1 to construct predicted loss	EUR million	30 June 2017 EBA Transparency Exercise
Duration	Maturity-bucket weighted average of sovereign holdings; open-ended bucket assigned 15 years	Years	Closest usable EBA episode vintage for country mapping
Leverage	Exposure measure divided by Tier 1 capital; 2011 bridge uses total assets over Core Tier 1 capital	Multiple	Closest usable EBA episode vintage for country mapping
Economic CET1 loss	Event-sovereign exposure multiplied by duration and spread movement, scaled by capital	Basis points of CET1	Event window
Abnormal equity return	Bank return minus Euro STOXX Banks or beta-adjusted Euro STOXX Banks benchmark return	Percent	Acute market window
Bank controls	CET1 ratio, leverage ratio, liquid assets to assets, NPL ratio, and log total assets	Percent or log	Last fiscal year before event

Table B.6. Bank-variable definitions

Variable	N	Mean	SD	P25	Median	P75
Italian sovereign exposure / leverage	25	2.86%	4.52%	0.12%	0.53%	2.71%
Domestic sovereign exposure / leverage	25	5.53%	4.17%	2.03%	5.55%	8.11%
CET1 ratio	23	13.87	2.25	12.08	13.44	14.79
Leverage ratio	22	5.78	1.55	4.97	5.57	6.18
Liquid assets / assets	23	30.30	10.09	23.99	28.40	34.22
NPL ratio	19	8.09	8.46	2.95	4.80	11.12
Log total assets	23	19.79	1.19	18.85	19.95	20.73
Beta-adjusted abnormal return	25	-0.76	5.00	-2.89	0.58	2.28
CDS change	7	51.66	31.96	28.22	49.15	73.18

Notes: Each variable's N is its own ORBIS/BankFocus coverage, not a stage of the funnel in Table B.5, which tracks only the two controls, log total assets and the NPL ratio, that gate the preferred regression sample. The leverage ratio and CET1 ratio are reported for descriptive purposes and are not both required for the preferred specification.

Table B.7. Descriptive statistics for the Italy 2018 listed-bank panel

Figure B.2 shows the same 25-bank Italy-2018-acute sample as a scatter of raw domestic sovereign exposure against the abnormal-return outcome, in place of the CET1-normalised predicted-loss regressor used in the preferred regression of Chapter 5; it is a descriptive companion to Table B.7, not a reproduction of the regression figure.

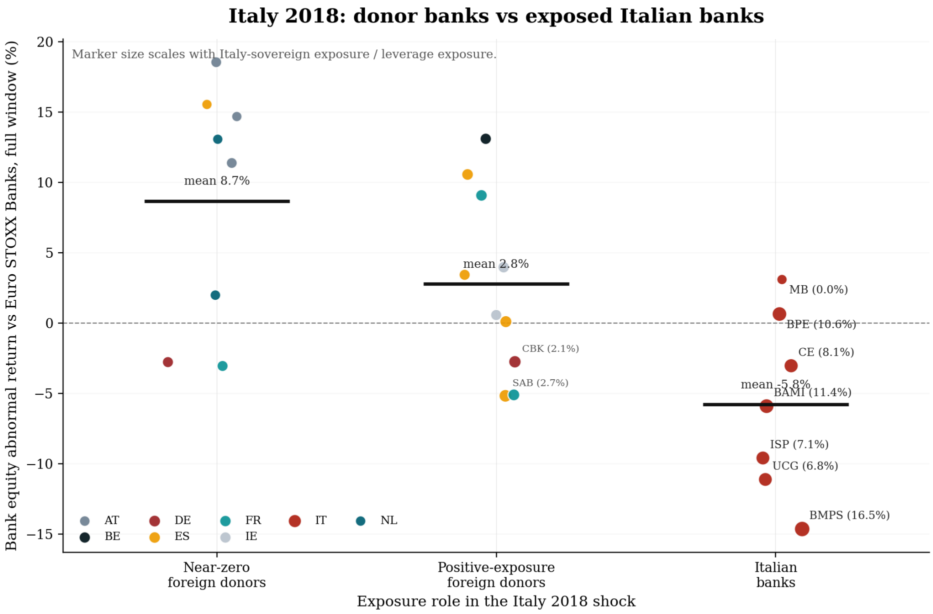

Figure B.2. Italy 2018 acute window: domestic sovereign exposure (EUR million, unnormalised by capital) against beta-adjusted abnormal equity return, 25-bank panel. Italian banks are marked separately from the remaining euro-area banks; the line is an unweighted linear fit shown for visual reference only. The Italian group mean of −5.4% shown here covers the seven Italian banks in the descriptive 25-bank panel; the −8.3% quoted in Chapter 5 covers the four Italian banks that survive into the controlled 13-bank regression sample of Table 5.3. The two figures describe different samples and are not alternative estimates of the same quantity.

The identifying variation is skewed: median Italian-sovereign exposure is 0.53% of the leverage measure against a mean of 2.86%, so a few banks carry most of the cross-sectional spread. The natural robustness check is a sample-exclusion variant dropping each bank in turn. This requires the same licensed, not-redistributable Bloomberg-and-ORBIS panel as the headline regression, so the underlying returns and controls are not part of what accompanies this release; the derived, non-licensed leave-one-bank-out coefficients are released instead (`results/acceptance/bank_validation/bank_loss_validation_jun2017_leave_one_out.csv`). The sign is negative in all 13 one-bank-dropped specifications, ranging from −1.90 to −0.30, none reaching conventional significance; Banco de Sabadell is the most influential single observation. The null result in Chapter 5 is therefore not driven by one institution.

The stricter pre-event regression is necessarily coarser than the later detailed EBA template, but it is the appropriate design for the May event. Its 13-bank controlled coefficient is negative but statistically indistinguishable from zero (Chapter 5): a very imprecise small-sample check on upstream loss scaling, neither confirming a household lending-rate pass-through nor ruling out the loss-scaling channel.

B.6 Panel D. Household Microdata and Harmonisation

Household data enter in two layers: debt participation disciplines the steady-state borrower mass in the one-asset HANK, and survey balance-sheet vectors allocate model-implied losses across net-wealth quintiles after the model is solved. Four sources serve different purposes and are never pooled. SHIW is preferred for Italy because it reports direct holdings of government securities as a distinct field, unlike comparable surveys. EFF is read using implicate 1, understating imputation uncertainty, acceptable only because it supports a robustness comparison and no headline estimate. HFCS enters through published tables and not micro records, so no household-level moment can be recomputed, but it is harmonised across euro-area countries by construction, which a cross-country incidence comparison requires. LWS returns only aggregate outputs from the LISSY remote-execution system, so its moments cannot be re-cut without resubmitting a job.

Stage	SHIW	EFF	LWS	HFCS
Country/wave	Italy 2022	Spain 2022	France 2020	Wave 2021
Unit retained in the thesis	Household micro record	Household micro record	Aggregate LISSY output	Published group table
Final analysis sample	9,641	6,385	Aggregate only	Published wealth groups
Weights used	yes	yes	LWS weights	Published weights
Government-security detail	Direct measure	Broader fixed income	Public and private debt securities	Grouped public tables

Table B.8. Household sample construction

Only SHIW and EFF are retained as micro records, so only they can be re-cut into the quintiles the allocation requires. Table B.9 documents how each concept is built in each survey; the weakest correspondence is debt service outside SHIW.

Final concept	SHIW construction	EFF construction	LWS construction	Remaining limitation
Net wealth	Survey net-wealth measure, used for quintiles	Survey net-wealth measure, implicate 1	Disposable net worth from LISSY output	Definitions and imputation rules differ across surveys
Bank debt	Bank and finance debt fields	Total household debt fields	Household liabilities aggregate	Loan categories are not identical
Debt service	Annual debt-service field	Debt-service field where available	Aggregate debt-service output if retained	Debt-service concept is weakest outside SHIW
Financial assets	Deposits plus securities and other financial assets	Financial-asset aggregate	Financial assets aggregate	Treatment of pensions and business assets differs
Government securities	Direct government-security holding	Broader fixed-income securities	Debt securities combine public and private issuers	Cross-country bond-holding comparisons are not exact
Consumption / income	SHIW income and expenditure concepts	Selected-goods expenditure proxy	Aggregate disposable-income moments	Consumption concepts are not fully harmonised

Table B.9. Cross-survey variable harmonisation

Monetary entries in Table B.10 are weighted means in thousands of euros. Means and not medians are reported because the incidence accounting allocates aggregate euro losses, which a median would not aggregate to. Participation rates appear alongside amounts because holdings are skewed within each quintile: a quintile can have a high mean government-security holding while most of its households hold none.

Quintile	Pop. share	Net wealth	Fin. assets	Govt. sec.	Bank debt	Bank debt part.	Govt. sec. part.
Q1	19.99%	3.2	3.3	0.04	2.4	13.98%	0.31%
Q2	20.00%	71.6	12.5	0.42	16.0	27.65%	2.34%
Q3	19.99%	152.7	17.1	1.12	13.9	28.44%	5.38%
Q4	19.98%	258.1	35.2	2.10	12.7	26.25%	8.38%
Q5	20.03%	986.0	212.6	10.71	24.5	28.22%	14.40%

Table B.10. Italy household balance-sheet moments by net-wealth quintile

Figure 5.1 in Chapter 5 makes the resulting trade-off visible before any model result is imposed: bank-debt participation rises from the bottom quintile to the middle and upper quintiles, while financial assets and government securities are concentrated at the top, which is why the borrower-credit channel and the asset-holding allocation are kept as separate layers. EFF vectors are retained in the private archive for a historical Spain robustness exercise, which is not part of the submitted PDF.

B.7 Panel E. Macro-Financial Control Series

The remaining sources do not enter the structural model. ECB Balance Sheet Items credit series are converted to growth rates and compared with a donor pool of Austria, Belgium, Germany, Finland, Ireland and Portugal; the Netherlands is excluded as itself treated in the Netherlands-2023 null episode. Eurostat HICP inflation and consumer confidence enter as controls, the latter addressing the obvious alternative reading of the credit results – that households simply borrowed less because political turmoil made them pessimistic. The ECB policy-rate history separates the July-2022 tightening from the Italian political component of the same window. None of these series can sort banks by sovereign exposure or separate credit supply from demand, which is why the exercises built on them are consistency evidence, not an independent test.

CHAPTER C

Robustness, Calibration, and Allocation Sensitivity

This appendix retains only checks that qualify the preferred Italy-2018 borrower-wedge and SHIW allocation exercise. Historical episode comparisons, stress grids, and exploratory model extensions remain in the private replication archive. The checks below concern aggregate-credit diagnostics, the external calibration, and the constants that convert modelled losses into the fixed-vector incidence calculation.

C.1 Dated-Bridge Timing Classification

The timing classification is itself a robustness result. For Italy 2018, the relevant political window begins in May, while the supervisory templates that contain useful maturity information have different reference dates. The release does not use the word “pre-event” as a generic label for all balance-sheet information available during 2018. Instead, every bridge is classified according to the dates of the exposure, leverage, and duration components separately. A strict pre-event bridge would require all three components to be observed no later than 14 May 2018; no such duration measure is available in the sourced data.

Classification	Exposure and leverage	Duration	Status	Permitted use
Strict pre-event	Required by 14 May 2018	Required by 14 May 2018	Unavailable	Never fabricated or reported
Partial pre-event	December 2017	June-2018 range used explicitly as sensitivity	Preferred	Main quantitative result
Episode-vintage proxy	June 2018	June 2018	Post-event proxy	Context; never principal estimate

Notes: The partial-pre-event result is preferred because its exposure and leverage precede the May event. Its duration range is not relabelled as pre-event. The all-June calculation is retained for comparability with the closest episode-vintage mapping.

Table C.1. Italy-2018 bridge timing classification and permissible use.

This classification explains why the thesis gives the 4.03–7.84-basis-point wedge a different status from the June-2018 6.56-basis-point value. The first result preserves the earlier balance-sheet margins and makes the later duration information visible as an input sensitivity. The second uses one coherent, richer June snapshot, but that coherence does not remove its post-event timing. Neither number is a statistical estimate of an unobserved household pass-through; the distinction concerns the credibility of the dated input to the same conditional mapping. Reporting both prevents completeness of a later disclosure from being mistaken for pre-event identification.

C.2 Aggregate-Credit Diagnostics

ECB country aggregates are a deliberately weak diagnostic, not a test of the household borrower wedge: they combine credit supply and demand, monetary policy, housing conditions, and macroeconomic activity, and contain neither lender exposure nor borrower-level outcomes. Italian new-business mortgage and consumer-credit rates do not show a discrete deterioration after May 2018; the outstanding stock is never multiplied by the wedge to construct an immediate payment-flow loss.

Table C.2 reports descriptive donor-pool and synthetic-control comparisons. Italy's 2018 gaps point toward weaker relative credit growth, but the formal local-projection exercise with country and calendar-month fixed effects, lagged inflation and confidence controls, and country-clustered standard errors is non-confirmatory at every horizon. With ten country clusters and overlapping Italian windows, these diagnostics cannot validate or recalibrate the structural pass-through.

Event	Relative-growth change	Synthetic gap	Pre-RMSPE
Italy 2018 acute	-0.97 pp	-1.48%	1.15
Italy 2018 full	-1.09 pp	-2.95%	1.29
Italy 2022/TPI	+0.19 pp	-1.01%	1.78
France 2024	-0.73 pp	-1.21%	0.82

Table C.2. Descriptive household-credit diagnostics for retained comparison episodes.

For completeness, Table C.3 reports selected horizons from the formal pooled local-projection diagnostic. The coefficients are small relative to their country-clustered uncertainty, and every interval includes zero. The negative pre-event coefficients are also a reason not to use this panel as confirmatory causal evidence: with only ten country clusters, common macro shocks and overlapping Italian windows cannot be separated sharply from the impulse indicator. The donor-pool and synthetic-control summaries above are retained as descriptive context, while the formal panel result sets the appropriate evidentiary boundary.

Horizon (months)	Coefficient	Clustered s.e.	95% interval	Observations
-2	-0.637	1.105	[-2.803, 1.530]	2,629
0	-0.237	1.370	[-2.924, 2.450]	2,649
3	-0.249	1.523	[-3.236, 2.739]	2,667
6	-0.330	1.525	[-3.320, 2.660]	2,671
12	-0.109	1.508	[-3.067, 2.849]	2,661

Notes: Two-way fixed-effects local projections pool seven event windows against the six-country donor pool, with one-month-lagged HICP inflation and consumer confidence controls. Standard errors are clustered by country. The table is a diagnostic for aggregate household-credit growth, not a household-credit pass-through estimate.

Table C.3. Selected horizons from the pooled local-projection credit diagnostic.

Ten clusters is well below the number at which cluster-robust standard errors are reliable, and the usual failure is over-rejection (Cameron et al., 2008). The null result above could therefore be understating the evidence against a credit effect and not overstating it, so each horizon's p-value is recomputed by a wild cluster bootstrap with Rademacher weights and the null imposed. With ten clusters there are only $2^{10} = 1024$ distinct weight vectors, so the bootstrap distribution is enumerated exhaustively and the resulting p-values carry no simulation error. They are indistinguishable from the clustered ones at every one of the nineteen horizons: 0.877 against 0.873 at impact, 0.852 against 0.842 at six months, and 0.953 against 0.947 at twelve. No horizon crosses 5% under either reference. The panel's failure to detect an aggregate credit response is a property of the data, not an artefact of the standard-error formula.

C.3 Contract-Exposure Boundary

The structural wedge is a condition on marginal borrowing and refinancing in the model: the household's negative liquid-asset position is a stylised rolling borrowing position, not a literal mortgage or loan contract, and the wedge applies to that stylised position and never to an observed debt instrument. It cannot be converted into an immediate cash-flow loss by multiplying it by the observed stock of Italian household debt. Public ECB series identify new-business mortgage and consumer-credit rates and the aggregate stock, but not the share of the outstanding stock that reprices or refinances at the modelled frequency. The information required to make that conversion is contract-level: initial rate fixation, reset dates, remaining maturity, refinancing behaviour, and lender-specific pricing. Those records are not available to this thesis.

The public aggregates are nevertheless useful for ruling out an overstatement. New-business house-purchase and consumer-credit rates do not show a discrete increase in the selected aggregate comparison, and the household loan stock is reported only as a scale reference. This does not contradict the model: new business and outstanding contracts are different objects, and aggregate prices combine supply, demand, composition, and monetary policy. It does establish that the thesis should retain the narrower marginal-credit interpretation.

Object	Pre-event	Post-event	Permitted use
New-business house-purchase rate	1.903%	1.825%	Marginal mortgage-pricing description only
New-business consumer-credit rate	6.723%	6.672%	Marginal consumer-credit description only
Italian household loan stock		EUR 634.2bn	Scale reference; never multiplied by the model wedge
Outstanding-stock repricing share	unavailable		No immediate payment-flow estimate is made

Notes: Inputs are the ECB MIR/BSI aggregate proxy file. Dates and units are documented in the release manifest. The unavailable row is a substantive limitation, not a zero value.

Table C.4. Public contract-exposure information and permitted interpretation.

C.4 Calibration and Accounting Sensitivity

The bank-loss object is an economic-value revaluation, not a claim about contemporaneous regulatory-capital recognition. Market losses can impair intermediation through capital, collateral, and funding even when their accounting recognition differs (De Marco, 2019; Giannetti et al., 2026). The Italy-2011 bridge uses first-order duration and a Core Tier 1 denominator; harmonising only the reported capital denominator to later Tier 1 definitions raises ϕ_0, but the numerator cannot be harmonised because Basel-III leverage exposure was not reported in 2011. This is an unresolved calibration limitation, not evidence that the headline is conservative.

Table C.5 distinguishes the preferred external target from quantities translated by this thesis. Only the Bofondi–Carpinelli–Sette Italy-2011 same-firm lending differential is the preferred calibration target. The other rows are diagnostic readings, not confidence intervals.

Source	Source object	Thesis-implied wedge	Status
Bofondi et al. (2018)	15–20 bp same-firm Italian-versus-foreign lending differential, 2011	6.56 bp	Preferred external target
Altavilla et al. (2017)	Euro-area bank-lending response to sovereign stress	9.25 bp	Translated scale check
Sabatini (2022)	Italy-2018 credit-quantity response through capitalisation	0.94–1.31 bp	Illustrative attenuation
Neri (2013)	130 bp firm and 60 bp household counterfactual rate increases, peripheral euro area, 2010–11	3.03 bp	Household-transport adjustment

Notes: The final three wedges are computed by this thesis from empirically distinct objects. They do not estimate Italian household-rate pass-through. The Neri row applies a ratio, not a level: see the discussion below.

Table C.5. Calibration anchors and their status.

Transporting the firm target to households. The preferred anchor measures firm lending, while the model wedge applies to household borrowing. Neri (2013) permits a direct adjustment because it reports both legs of the same counterfactual on the same sample: had sovereign spreads remained at their April-2010 levels, rates on new loans to non-financial corporations and on residential mortgages in the peripheral countries would have been lower by approximately 130 and 60 basis points respectively at the end of 2011. The levels are not comparable with the 17.5-basis-point figure used here, because Neri estimates a total spread-to-rate pass-through combining funding, capital, risk and demand channels with no control group, whereas the same-firm design differences out borrower demand and common shocks. The *ratio* is more informative: household rates moved approximately 0.46 times as much as firm rates over the same episode. It provides an illustrative same-episode attenuation factor with no claim to identification, because Neri fits separate country-level equations to two lending categories differing in contract structure, borrower composition and transmission channels. Applying it to the 17.5-basis-point same-firm differential gives an 8.1-basis-point household target, a re-calibrated ϕ_0 of 0.001764, a peak wedge of 3.03 basis points and a bottom-two total-loss share of 5.51%.

This adjustment lowers the headline magnitude and is reported for that reason. It implies that the baseline wedge is, on the only household-versus-firm evidence available for this episode, approximately twice too large when read as a household object. The defensible statement is therefore a calibration range of roughly 3.0–6.6 basis points, with the Italy-2011 firm target at its upper end, not its centre. Table 5.7 crosses this attenuation with the dated-input bridge. The direction of the allocation shift is unchanged throughout, because condition (5.5) turns on the shape of the SHIW vectors and not on the scale of the wedge.

The lower-scale re-solves reduce the fixed-wage output response to −0.0004% to −0.0006% and the matched wedge/SHIW bottom-two share to 5.03–5.11%. Re-calibrating ϕ_0 for each curvature value so that every row reproduces the Italy-2011 17.5-basis-point target gives a 5.80–6.96-basis-point wedge and a 6.11–6.36% bottom-two share (Table C.6). These are deterministic calibration sensitivities over maintained assumptions.

The two exercises answer different concerns. The lower-scale calculation asks what happens when a distinct Italy-2018 capital-channel quantity result is translated into a smaller borrower wedge; it is not a re-estimation of the preferred price target. The curvature exercise instead preserves the external 2011 target in every row, changing only the functional shape that maps bank stress into the wedge. Recalibrating ϕ_0 is essential: varying curvature while holding its scale fixed would confound a functional-form sensitivity with an unannounced change in the external target.

Across the curvature rows the sign of the borrower component and the ordering relative to the direct allocation are unchanged. This should not be overstated as empirical confirmation. Given the positive model wedge and fixed SHIW vectors, the direction of the allocation shift follows mechanically; the useful robustness information is how much the reported magnitude moves when the calibrated mapping changes. The timing bridge in the main results remains the preferred way to express the additional vintage uncertainty.

ϕ_1	Calibrated ϕ_0	2011 fit error	Peak wedge	Fixed-wage Y	Wage-moving Y	Bottom-two S2
0.5	0.008416	0.000 bp	6.96 bp	−0.0031%	−0.0004%	6.36%
1.0	0.003812	0.000 bp	6.56 bp	−0.0029%	−0.0004%	6.28%
1.5	0.002295	0.000 bp	6.18 bp	−0.0027%	−0.0004%	6.19%
2.0	0.001550	0.000 bp	5.80 bp	−0.0025%	−0.0004%	6.11%

Table C.6. Curvature sensitivity with the Italy-2011 target re-calibrated in every row.

C.5 Allocation-Level Sensitivity

The fixed SHIW allocation uses four constants: the stock-loss annuitisation rate, debt relative to consumption, government-securities exposure relative to consumption, and the bank-equity exposure ratio that scales the indirect direct-loss channel. The preferred combination uses the baseline model's steady-state rate and Italy-specific SHIW or macro sources. The alternatives include a legacy quarterly rate, a French-stylised debt ratio, and a survey-only securities ratio; they are disclosed diagnostic scenarios, not an Italian statistical interval. Across the eight combinations, the bottom-two total-loss share ranges from 6.28% to 10.54%. A separate levered bank-equity-price basis produces 6.06%, because enlarging the top-loaded direct-loss channel dilutes the relative borrower component.

The fourth constant requires a separate statement because its provenance is weaker than the other three. The bank-equity exposure ratio proxies household indirect exposure to revalued bank equity and to fund and insurance holdings of sovereign bonds. Its value of 0.60 is inherited from an earlier French-calibrated proof of concept and is not Italy-specific; no Italian survey aggregate among the sourced inputs separates indirect bank-equity exposure from the financial-asset total, so it is carried as a maintained level, not a sourced one. It matters because it scales the larger of the two direct-loss components. The final two rows of Table C.7 therefore vary it over 0.30–1.20 around the maintained 0.60, holding everything else at baseline. The direct-only share moves from 4.39% to 5.13% and the total-loss share from 6.57% to 6.03%: the reference point against which the borrower channel is measured is materially sensitive to this constant, while the headline total is not. Both endpoints preserve the sign of the allocation shift.

The table separates two distinct sources of variation. Changing the annuitisation rate rescales a stock loss relative to consumption and therefore changes the balance between direct and borrower components. Changing the debt-to-consumption or direct-bond constants changes the cross-quintile allocation itself. The Italian SHIW debt ratio and the MEF–Eurostat direct-bond mapping are the preferred inputs because they match the country and the stated accounting object. The alternatives remain visible to show that allocation-side uncertainty can be larger than solver-side uncertainty; they are not candidates for selective headline choice.

Annuit. rate	Debt/cons.	Bonds/cons.	A1 price basis	S1 (%)	S2 (%)	Amplification
6.17%	0.541 (SHIW IT)	0.297 (MEF)	Unlevered	4.802	6.277	1.073
6.17%	0.541 (SHIW IT)	0.112 (SHIW)	Unlevered	5.234	7.023	1.092
6.17%	1.30 (FR stylised)	0.297 (MEF)	Unlevered	4.802	8.037	1.175
6.17%	1.30 (FR stylised)	0.112 (SHIW)	Unlevered	5.234	9.080	1.221
4.10%	0.541 (SHIW IT)	0.297 (MEF)	Unlevered	4.802	6.947	1.110
4.10%	0.541 (SHIW IT)	0.112 (SHIW)	Unlevered	5.234	7.816	1.138
4.10%	1.30 (FR stylised)	0.297 (MEF)	Unlevered	4.802	9.329	1.263
4.10%	1.30 (FR stylised)	0.112 (SHIW)	Unlevered	5.234	10.536	1.332
6.17%	0.541 (SHIW IT)	0.297 (MEF)	Levered	5.086	6.062	1.048
6.17%	0.541 (SHIW IT)	0.297 (MEF)	Unlevered, ratio 0.30	4.388	6.572	1.110
6.17%	0.541 (SHIW IT)	0.297 (MEF)	Unlevered, ratio 1.20	5.132	6.026	1.044

Notes: The first row is the baseline. Annuitisation rates are this model's own steady-state rate (6.17% annualised) and the superseded quarterly proof of concept's rate (4.10%). Debt/consumption is either the SHIW 2022 Italian weighted mean (0.541) or the earlier French-stylised default (1.30); government securities per unit of consumption is either the MEF retail share applied to the Eurostat debt/GDP ratio (0.297) or the SHIW directly reported holding (0.112). The A1 price basis is the unlevered duration price loss, the levered loss the bank itself absorbs at impact, or the unlevered loss with the bank-equity exposure ratio moved from its maintained 0.60 to 0.30 or 1.20. The bank path, population weights, quintile boundaries, and household steady state are held fixed throughout. The table is a deterministic accounting sensitivity and does not supply a confidence interval.

Table C.7. Household-incidence allocation sensitivity, Italy 2018.

Illustrative new-loan attenuation scenario

Public data do not identify the share of outstanding Italian household debt that reprices at the modelled frequency. EMF Hypostat reports an approximately 33% variable-rate share for *new* Italian mortgage lending in 2023; it is a manually transcribed chart value and not evidence on the 2018 outstanding stock. Table C.8 therefore scales only the credit-exposure vector by 0.33 as an illustrative attenuation scenario. It is neither a bound nor an estimate of contractual pass-through.

The calculation moves the bottom-two total-loss share from 6.277% under full model pass-through to 5.312% under the illustrative 33% scaling, while leaving the 4.802% direct share unchanged. It is useful precisely because it demonstrates the sensitivity of the accounting allocation to a contract assumption that cannot be recovered from the public stock data. Calling it a lower bound would be incorrect: new-loan variable-rate shares, 2018 contract composition, and the refinancing margin are different objects.

Scenario	S1 (%)	S2 (%)	Amplification
Full pass-through (baseline)	4.802	6.277	1.073
Floating-rate bound (33% of credit exposure)	4.802	5.312	1.024

Notes: The scenario leaves direct losses, the bank path, household weights, and the steady state unchanged. It uses a 2023 new-lending share solely to illustrate the dependence of the accounting result on the full-pass-through assumption.

Table C.8. Illustrative new-loan attenuation of the borrower-credit allocation, Italy 2018.

Pre-event vintage robustness

The event is May 2018 and the preferred vectors come from the 2022 wave, which post-dates it. That gap is worth testing instead of asserting away, because Italian household balance sheets moved over the period. The indebted share alone rose from roughly 21% in 2016 to 26% in 2022. Table C.9 therefore holds the model completely fixed, using the same bank path, borrower wedge, price loss, and all four allocation constants, and swaps only the household vectors for the 2016 and 2020 waves from the historical SHIW archive. Banca d'Italia supplies that archive with its own harmonised population weight, which is what makes waves either side of the 2020 sampling redesign comparable at all. The historical archive carries total household debt without the bank-financed detail, so every historical row uses total debt and the 2022 historical row is the like-for-like comparator.

The definitional step is small. Moving from bank-financed to total debt within the 2022 wave changes the bottom-two share only from 6.28% to 6.23%. The vintage effect runs the other way and is larger. Earlier waves give a *higher* bottom-two share, 7.49% in 2016 and 7.28% in 2020 against 6.23% in 2022, because household debt was less concentrated at the top of the wealth distribution before 2020 than after it. Using post-event vectors is therefore the conservative choice, and the composition shift survives in every wave. The aggregate amplification is identical at 1.073 throughout, which is not a coincidence. The exposure vectors are shares that sum to one in each wave, so redistributing them across quintiles changes who bears the loss without changing its total.

SHIW wave	Weighting	Credit vector	S1 (%)	S2 (%)	Amplification
2016	Historical	Total debt	6.08	7.49	1.073
2020	Historical	Total debt	5.50	7.28	1.073
2022	Historical	Total debt	4.84	6.23	1.073
2022	2022 cross-section	Bank debt (preferred)	4.80	6.28	1.073

Notes: S1 and S2 are the bottom-two quintiles' shares of total allocated loss before and after the model-generated borrower wedge is switched on. Only the household vectors change across rows; the bank path, wedge, price loss, and the four allocation constants of this appendix are held at their accepted Italy-2018 values. "Historical" weighting is Banca d'Italia's harmonised population weight for cross-wave comparison; the final row is the preferred 2022 cross-section allocation reported in Chapter 5.

Table C.9. Incidence allocation under pre-event SHIW waves, Italy 2018.

C.6 Reading the Sensitivity Evidence

The exercises in this appendix answer different questions and should not be pooled into one range. The timing bridge varies the balance-sheet inputs behind the wedge and is the preferred way to show vintage uncertainty. The lower-scale, curvature, and Neri transport rows vary the calibrated map from bank stress to the wedge. The allocation and new-loan rows leave the bank path untouched and vary how the borrower component reaches the household vectors.

Three things stay fixed throughout, which is what makes the rows comparable. The direct spread-only component, the population weights and quintile boundaries, and the model closure and household steady state are the same in every row. A change in a reported share therefore reflects the named assumption, and never a redefinition of the population or the loss measure. A preserved positive borrower component does not test its sign, which is mechanical given a positive wedge and the observed debt vector, and a span of curvature values is not a confidence interval, because its endpoints are assigned functional forms, not estimated ones. The appendix gives priority to the partial-pre-event bridge, reports the June mapping as a dated proxy, and reports the Neri transport even though it lowers the headline magnitude.

CHAPTER D

The Two-Asset Portfolio-Rebalancing Extension

The baseline gives households one liquid asset, so a sovereign repricing reaches them only through the borrower wedge and through their fixed holdings of revalued assets. Nobody can respond by changing what they hold. That is a substantive restriction for an incidence question, because a household with a second, illiquid asset can rebalance while its price is temporarily low, and the direction of that response works against the direct-loss channel the incidence exercise measures.

This appendix reports a two-asset extension that relaxes exactly that restriction. It is an extension, not a rival specification, and Section D.3 states plainly why it stays out of the headline. It is included because leaving it out would present the one-asset incidence result as though the rebalancing margin had never been examined.

D.1 Household Problem and Portfolio Price

Each household holds liquid assets a, illiquid capital k, and an exogenous income state $z \in \{z_L, z_H, z_C\}$. Liquid assets may be negative, so households with $a < 0$ face the same borrower wedge ϕ_t the sovereign–bank block generates in the baseline. Illiquid capital is non-negative and trades at the portfolio price q_t. Following Kaplan and Violante (2014), households draw a stochastic adjustment opportunity that permits a new illiquid position; the capitalist state z_C additionally receives aggregate bank dividends d_t. The probability of drawing that opportunity responds to the value of taking it,

$$p_{\text{adj}} = 1 - \exp\left[-\frac{\max(V^{\text{adjust}} - V^{\text{no adjust}}, 0)}{\bar{\eta}}\right], \tag{D.1}$$

with the cost scale $\bar{\eta} = 12$ set so that the capitalist state's own adjustment frequency in the top net-wealth quintile falls inside the 15–20% annual range used in the household-finance literature (Kaplan and Violante, 2014). At that value the implied probability for that state is 1.56% monthly, or 17.20% annualised; it is not a single rate applied uniformly to all households.

With an adjustment opportunity the household chooses consumption and both asset positions subject to

$$c_t + a_{t+1} + q_t k_{t+1} = (1 + r_t^a)a_t + \left[r_t^k + (1-\delta)q_t\right]k_t + w_t z_t + \mathbf{1}_{z_t = z_C} d_t, \tag{D.2}$$

liquidating the surviving value of the existing holding and buying a new one. Without an opportunity the illiquid holding stays at k_t and the budget is

$$c_t + a_{t+1} = (1 + r_t^a)a_t + (r_t^k - \delta q_t)k_t + w_t z_t + \mathbf{1}_{z_t = z_C} d_t, \tag{D.3}$$

where $r_t^k - \delta q_t$ is the rental return net of the maintenance investment needed to hold the position in place. Liquid saving and borrowing remain adjustable in both branches.

Total capital combines the household portfolio with bank-intermediated capital, $K_t^{\text{tot}} = K_t^H + L^B N_t^B$, where L^B is the maintained bank-leverage scale. Defining investment as $I_t = K_t^{\text{tot}} - (1-\delta)K_{t-1}^{\text{tot}}$, with the same expression evaluated on the no-event path for I_t^{base}, the portfolio price follows a baseline-relative proxy,

$$q_t = 1 + \psi_q \frac{I_t - I_t^{\text{base}}}{K_{t-1}^{\text{base}}}, \qquad \psi_q = 0.10, \tag{D.4}$$

iterated with the household solution until the capital path and the price are mutually consistent. Equation (D.4) is a maintained proxy that introduces a temporary rebalancing return around the event path. It is not a structural asset-pricing equation, and it is the reason this extension is reported as a mechanism and not as a second estimate.

D.2 Archival Value-Ratio Diagnostic for Italy 2018

The current archival V6 run reports a normalised value-ratio diagnostic, comparing the event path with the no-event transition at the initial household state:

$$VR_i = 100 \left(\frac{V_i^{\text{base}}}{V_i^{\text{shock}}} - 1 \right). \tag{D.5}$$

It is a lifetime diagnostic, not an observed consumption change in any month or a promoted consumption-equivalent welfare measure. Reported group figures are population-weighted means of household-level values over roughly 9,640 SHIW-2022 households.

Every quintile records a negative average value-ratio diagnostic, from -2.483% in Q1 to -0.496% in Q5, with a population-weighted aggregate of -1.616%. The portfolio price troughs at 0.9868, and 11.64% of households nonetheless record a positive diagnostic. In Q5 the representative-point evaluation turns positive while the full-distribution mean stays negative, which is why the full-distribution values are the reported record: averaging a representative point over a quintile whose portfolio holdings are widely dispersed gives the wrong sign for that quintile. The run carries the corrected pre-event capital initialisation and records zero boundary clipping on the illiquid-asset grid.

Figure D.2 sorts households by their initial liquid and illiquid positions instead of by net wealth, and shows that composition matters on its own. Two groups do comparatively well. Non-adjusting households gain mechanically while the price is low, because their net return $r_t^k - \delta q_t$ in Equation (D.3) rises as q_t cuts the replacement cost of maintaining the holding. Adjusting households with enough liquidity buy capital cheaply and hold it into the recovery of q_t. Both channels favour the liquid, which is the mechanism Mendicino et al. (2025) identify, and both are unavailable to the borrowers who carry the wedge in Chapter 5.

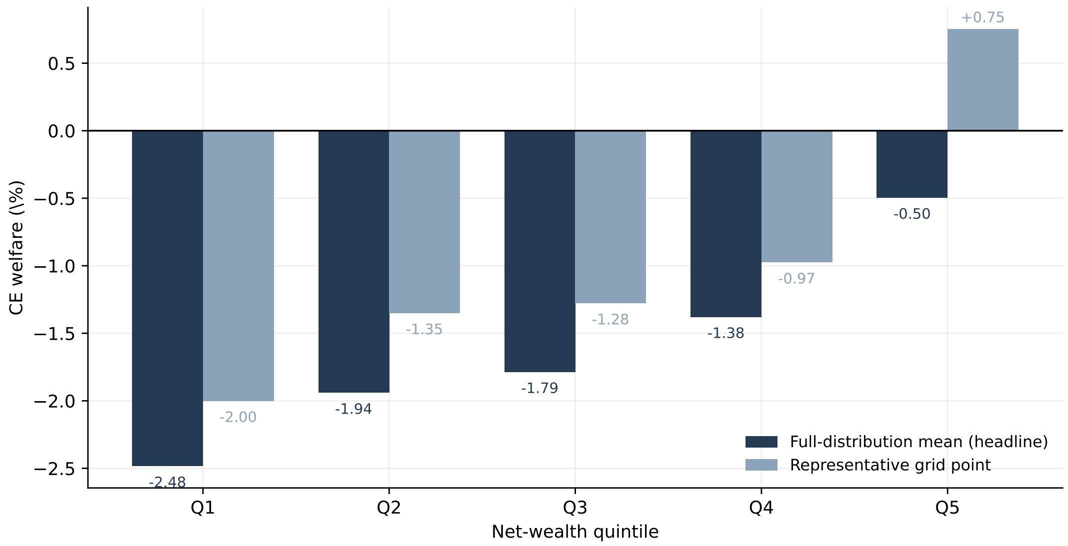

Figure D.1. Archival V6 value-ratio diagnostic by net-wealth quintile, Italy 2018. Full-distribution averages differ from representative-point evaluations most clearly in the top quintile, where portfolio holdings are most dispersed.

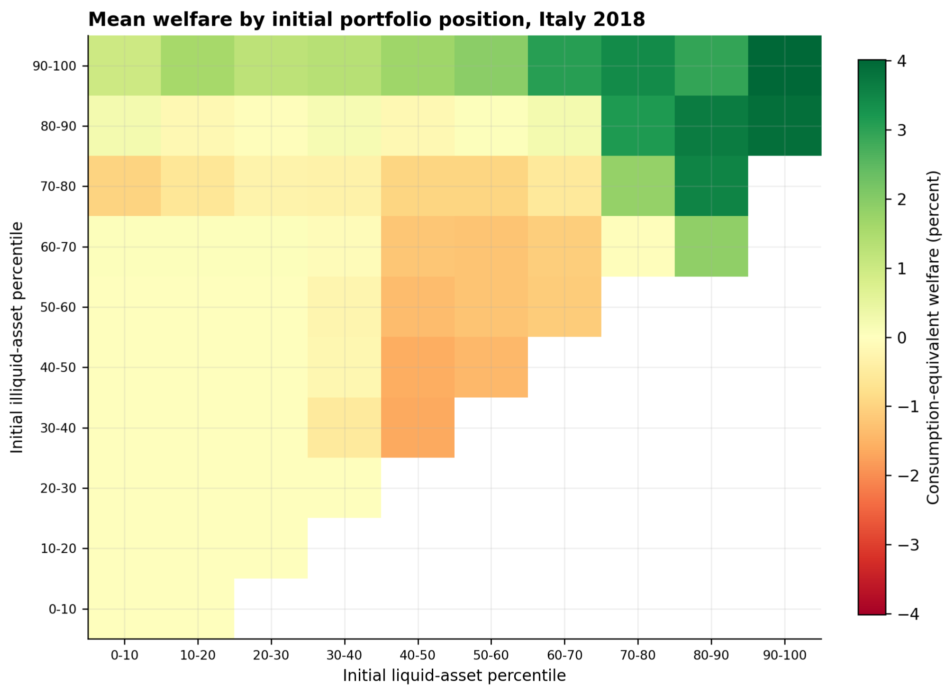

Figure D.2. Portfolio-position map for Italy 2018. Cell colours give the mean archival value-ratio diagnostic across survey-weighted percentile bins of initial liquid and illiquid holdings.

D.3 Status of This Extension

The extension is an exploratory diagnostic, not a robustness test that validates the baseline ranking. Its Italy-only output is directionally consistent with that ranking: the bottom quintile loses most, the top least, and the households able to benefit from the price trough are the liquid ones. That comparison is informative about the rebalancing margin, but it cannot establish that the baseline ranking survives a materially richer household problem.

Three limits keep it in an appendix. Equation (D.4) is a maintained price proxy, not an estimated asset-pricing relation, so the size of the rebalancing return is assumed and never identified. The archived V6 calculation also uses a value-ratio normalisation and retains a legacy annual depreciation setting in a monthly code path; a frequency-clean consumption-equivalent calculation requires a corrected implementation and a fresh re-solve, so the figures above are not promoted as welfare evidence. The channel decomposition that would attribute the diagnostic to the wedge, the price path and dividends separately was withdrawn after the capital-initialisation correction and has not been re-run, so no decomposition is reported here. And the France-2024 counterpart has not been re-solved since that correction, so this appendix reports Italy 2018 only and no cross-country ratio should be taken from it.

The consequence for the thesis is narrow and worth stating exactly. The composition result of Chapter 5 is unchanged, because it is an accounting allocation over observed vectors and does not depend on this extension. What this appendix adds is an exploratory illustration of the rebalancing margin; it neither confirms nor refutes that the baseline ranking is robust to the one-asset restriction. Any claim about that ranking or its magnitude would require a structural price relation, a re-run channel decomposition, and comparable cross-event solutions.