Essays in Financial Economics

Jasper Knyphausen

The Queen's College

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy

April 2026

This thesis studies three sources of risk in financial markets that share a common feature: each arises from policies designed to alleviate one problem but, as this thesis shows, carrying adverse externalities for financial markets further down the line. Chapter 1 studies the effect that fiscal expansion in a safe-haven country has on the government bond universe of other countries. Using a proxy-SVAR identified with German Bund futures, it finds that fiscal shocks raise the risk-free benchmark, widen peripheral spreads, and contract lending. Crucially, these materialise before actual fiscal spending occurs. Chapter 2 studies how post-crisis regulation designed to make financial markets more resilient is shaping haircuts in repo markets in ways not previously documented. It shows that binding leverage constraints distort bilateral haircuts, with consequences for financial stability and monetary policy pass-through. Lastly, Chapter 3 investigates how climate transition risk could potentially spill into liquidity risk for commercial banks. By studying these channels, the thesis contributes to the ongoing debate about how regulation and fiscal stimulus should be designed to maximise intended outcomes while minimising unintended consequences.

Supervisors This thesis was supervised by Dimitrios Tsomocos and Mungo Wilson.

Total words This thesis contains approximately 33,000 words.

Statement of authorship The research contained in this thesis is exclusively carried out by me as a DPhil student at the University of Oxford between October 2022 and April 2026. No part of this thesis has been submitted for a degree at the University of Oxford or any other university. Chapter 2 is co-authored work with Benedikt Ballensiefen and Benoît Nguyen. Chapter 3 is co-authored work with Margherita Giuzio and Bige Kahraman Alper. Chapter 1 is entirely my own work.

Acknowledgment of funding I gratefully acknowledge funding from Saïd Business School Foundation.

Essays in Financial Economics

Jasper Knyphausen
The Queen's College
University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy
April 2026

To my goddaughter, Vivica

Acknowledgements

I would like to express my sincere gratitude towards my primary supervisor Dimitrios Tsomocos. Dimitri, without your continuous, rigorous and selfless support I would not be where I am today. I am grateful for all the debates we had in your office, for your push-back as well as for your encouragement. It does not happen often in academia to find someone who thinks so deeply about the economic questions, the theory behind the questions, and the human being behind the theory.

I would also like to thank the economics and finance faculties at Oxford. I am particularly grateful for the help of Alexandre Kohlhas, Jiri Knesl, Seung Joo Lee, Silvia Miranda-Agrippino, Tom Noe, Joel Shapiro and Ansgar Walther.

I am indebted to Benoît Nguyen, Julian Schumacher and Margherita Giuzio for helping me to get access to confidential ECB data which was extensively used in this thesis.

The time as a PhD student would have never been the same without the close PhD friends in the office, most notably Aditya Khemka and Anthony Limburg. You not only helped me grow academically, you have also been there for me when times got tough, and I will never forget that.

Lastly, I am forever grateful to my parents, siblings and friends. Your unwavering support and encouragement have made the past four challenging DPhil years an unforgettable experience.

Abstract

This thesis studies three sources of risk in financial markets that share a common feature: each arises from policies designed to alleviate one problem but, as this thesis shows, carrying adverse externalities for financial markets further down the line. Chapter 1 studies the effect that fiscal expansion in a safe-haven country has on the government bond universe of other countries. Using a proxy-SVAR identified with German Bund futures, it finds that fiscal shocks raise the risk-free benchmark, widen peripheral spreads, and contract lending. Crucially, these materialise before actual fiscal spending occurs. Chapter 2 studies how post-crisis regulation designed to make financial markets more resilient is shaping haircuts in repo markets in ways not previously documented. It shows that binding leverage constraints distort bilateral haircuts, with consequences for financial stability and monetary policy pass-through. Lastly, Chapter 3 investigates how climate transition risk could potentially spill into liquidity risk for commercial banks. By studying these channels, the thesis contributes to the ongoing debate about how regulation and fiscal stimulus should be designed to maximise intended outcomes while minimising unintended consequences.

Contents

List of Figures

Appendices

List of Tables

Appendices

Introduction

The efficient functioning of financial markets is a prerequisite for macroeconomic stability. When risks are accurately priced, capital flows to its most productive uses, monetary policy transmits smoothly to the real economy, and the financial system absorbs shocks without amplifying them. Yet financial history repeatedly demonstrates that markets can fail on all three counts: risks go unpriced or mispriced, frictions impede the transmission of policy, and shocks propagate through channels that regulators did not anticipate. Understanding the sources of financial risk and how institutions respond to them is therefore central to the design of policies that safeguard aggregate welfare.

This thesis investigates how diverse sources of risk interact with the structure of European financial markets to shape financial stability and monetary policy transmission. The origins of volatility in financial markets that I study in this thesis are very different (fiscal policy, financial regulation, and climate change) but converge on a common set of questions. How do risks propagate through the European financial system? What role do market frictions and institutional design play in amplifying them? And what are the consequences for the effectiveness of monetary policy? The euro area, which combines a single monetary union with nationally segmented sovereign bond markets, provides a particularly rich setting for these questions. The first chapter exploits this structure by tracing how fiscal news in Germany propagate through bond markets and bank balance sheets across the euro area, using high-frequency financial market data. The second and third chapters zoom in on the repurchase agreement (repo) market, the primary venue through which European banks manage short-term funding and collateral needs and the main channel through which the ECB implements its policy rate (Duffie and Krishnamurthy, 2016a; European Central Bank, 2011). Both draw on granular, transaction-level data from the ECB's Money Market Statistical Reporting (MMSR) dataset to study how market frictions and emerging risks manifest at the level of individual transactions.

A unifying theme is that risks which appear moderate in isolation can become systemically significant once they interact with the institutional features of European financial markets. A fiscal announcement that merely reprices the risk-free benchmark can activate a sovereign-bank doom loop that contracts credit in the periphery. Regulatory constraints designed to make banks safer can distort the setting of haircuts, shifting risk onto prices and quantities with adverse consequences for monetary policy pass-through. And the build-up of climate transition risk on bank balance sheets can segment the interbank lending market, amplifying financial stress when it materialises.

Chapter 1 identifies the financial spillovers of fiscal expansion in a safe-asset sovereign. The theoretical literature identifies two regimes for the effects of safe-asset supply on financial markets. In one, additional issuance supports market functioning by expanding the stock of low-risk collateral and easing funding conditions. In the other, the increase in supply raises roll-over fears and destabilises the safe-asset (He et al., 2019). Which regime prevails at any given moment is an empirical question—and one with high stakes given the rapidly rising debt levels in major safe-asset countries. I construct a novel high-frequency instrument for German fiscal policy shocks, measuring changes in Bund futures prices within narrow windows around 172 fiscal announcements between 2016 and 2025, and use this series as an external instrument in a proxy vector autoregression (Stock and Watson, 2012; Mertens and Ravn, 2013). The results reveal a striking paradox: a fiscal shock normalised to a 25 basis point increase in the 10-year Bund yield triggers an immediate widening of peripheral sovereign spreads and a surge in systemic stress. The mechanism operates through bank balance sheets, activating the sovereign-bank nexus modelled by Farhi and Tirole (2018): mark-to-market losses on sovereign holdings erode bank equity, triggering credit contraction on both the extensive and intensive margins. The resulting credit squeeze reduces quarter-on-quarter GDP growth by 11 basis points in Germany, 25 in Spain, and 29 in Italy. Importantly, this penalty materialises before any fiscal stimulus is disbursed, effectively rendering the short-run fiscal multiplier negligible or negative for the periphery.

Chapter 2 turns to the microstructure of the repo market, examining how frictions distort the setting of haircuts which is one of the primary risk management tools in collateralised lending. Using transaction-level data from the MMSR dataset covering 2019–2024, we document that haircuts in the euro area bilateral repo market are systematically suboptimal: zero haircuts are pervasive, counter-intuitive haircuts are common, and the cross-sectional pattern across collateral types contradicts standard models. We identify three frictions that explain these distortions. First, dealers' balance sheet constraints put upward pressure on haircuts in collateral-driven repos. Second, collateral heterogeneity interacts with these constraints unexpectedly: high-quality assets such as German Bunds receive *higher* haircuts because the dealers most active in these segments are especially constrained. Third, market power matters: many clients trade with only one or two dealers, limiting their ability to negotiate (Eisenschmidt et al., 2024). A Bertrand model with capacity constraints rationalises these findings. The consequences extend beyond the haircut: when haircuts fail to absorb risk, repo rates and volumes become more volatile, and more constrained dealers pass through ECB rate changes more slowly, weakening monetary policy transmission. These findings speak directly to the ongoing debate over minimum haircut requirements (Financial Stability Board, 2025) and illustrate how regulatory reforms can inadvertently impede policy effectiveness.

Chapter 3 investigates how banks' exposure to climate transition risk affects the pricing of short-term interbank lending. While a growing literature documents climate risk premia in equities (Kacperczyk and Bolton, 2021; Bolton and Kacperczyk, 2024), corporate bonds (Bua et al., 2024; Painter, 2020), and real estate (Giglio et al., 2021; Bernstein et al., 2019), virtually no evidence exists for the repo market. Using MMSR data for 2019–2022 combined with data on banks' financed greenhouse gas emissions, we show that a one standard deviation increase in transition risk exposure is associated with a 7–12 bps increase in repo rates. The evidence points to a combination of a genuine risk premium and what we term an inconvenience premium, reflecting the reluctance of climate-committed dealers to extend cheap funding to carbon-exposed counterparties. The premium amplifies three-fold during periods of financial stress, and transition risk alters monetary policy transmission: repo rates for high-emission banks adjust more rapidly to ECB rate hikes, indicating that climate-induced segmentation creates uneven pass-through across the banking sector.

Taken together, the three chapters demonstrate that fiscal shocks, regulatory distortions, and climate transition risk share a common feature: they generate frictions that segment financial markets and, to some extent, distort the pricing of risk. The findings underscore that financial stability and monetary policy effectiveness cannot be analysed in isolation, and that the institutional design of European financial markets has profound consequences for how risks propagate through the system.

Chapter 1

Contractionary stimulus: The financial channel of fiscal shocks

Abstract. This paper identifies the financial spillovers of fiscal expansion in a safe sovereign. Using high-frequency movements in German bond futures around fiscal news, I show that an expansionary shock normalised to a 25 bps increase in the 10-year Bund yield triggers an immediate widening of peripheral sovereign spreads and a surge in systemic stress. The repricing of the risk-free benchmark erodes bank capital and activates a sovereign-bank nexus that disproportionately penalises high-debt jurisdictions. Banks respond by contracting credit supply on both the extensive and intensive margins, passing on the interest rate risk to the private sector through floating-rate contracts. Critically, these contractionary forces materialise as a pre-spending penalty: the financial tightening occurs upon the news of future debt issuance, before any fiscal stimulus is actually disbursed. The resulting credit squeeze reduces real GDP growth by 11 to 29 basis points in the periphery, effectively rendering the short-run fiscal multiplier negligible or negative. These findings suggest that in a currency union, core fiscal expansions can create significant negative externalities by destabilising the financial conditions of other countries.

JEL classification: E44, E62, G21, H63

Keywords: Fiscal policy, Financial fragmentation, SVAR, External instruments, Bank lending

For helpful comments and suggestions, I thank Pedro Bordalo, Martin Ellison, Niels Johannesen, Aditya Khemka, Alexandre Kohlhas, Seung Joo Lee, Anthony Limburg, Michael McMahon, Silvia Miranda-Agrippino, Ken Okamura, Martin Schmalz, and Ansgar Walther.

1.1 Introduction

The role of fiscal policy in financial stability has taken on renewed urgency as major economies have drastically elevated their debt levels in response to unprecedented geopolitical events. The fiscal expansion of countries whose government bond is considered a safe asset, such as the U.S. or Germany, is particularly important: an increase in the supply of safe assets could stabilise financial markets by increasing the supply of low-risk collateral and easing funding conditions, provided that markets can absorb the additional issuance without repricing the sovereign's fiscal trajectory. However, in an environment where concerns about debt sustainability are elevated, additional issuance may erode the perceived safeness of the asset, raising term premia and potentially destabilising the financial system. Given the rapidly increasing debt-to-GDP level in major safe-asset countries, it is *a priori* unclear whether the increased supply of government bonds is still good news for the financial system. What are the systemic implications of expansionary fiscal policy? When does fiscal expansion stabilise or destabilise financial markets?

To answer these questions, I propose a novel approach to identify the causal effects of budget shocks in Germany on financial markets and economic output, exploiting high-frequency movements in German government bond futures around fiscal policy news. I construct a fiscal policy surprise series by measuring changes in Bund futures prices within narrow windows surrounding 187 fiscal announcements between 2016 and 2025. Reverse causality can be plausibly ruled out as macroeconomic conditions are known and priced in before the fiscal news, and they are unlikely to change within the tight window considered. Using the resulting surprise series as an external instrument in a proxy vector autoregression allows me to isolate the structural fiscal shock component. The VAR's lag structure absorbs the predictable dynamics of the system, so that the instrument identifies the fiscal shock within the unpredictable innovations, orthogonal to other contemporaneous structural disturbances. I then take the resulting structural shock to estimate the dynamic causal effects of German fiscal shocks on financial market conditions and real economic activity across Europe.

The focus on Germany as the key safe-asset anchor in the euro area is motivated by two considerations. First, the unified currency but segmented bond market structure in the euro area adds a layer of complexity that has not been studied before in this context. Because German yields act as the effective risk-free benchmark against which all euro area assets are priced, any disruption to them does not just impact Germany's government bonds but ripples through the entire euro area economy. Second, the euro area offers publicly available, disaggregated data on bank lending at the country level that does not exist for the U.S. to the best of my knowledge, enabling a direct examination of how fiscal shocks transmit through bank balance sheets to the real economy.

Even though German debt levels remain moderate by historical standards, I show that expectations about expansionary German fiscal policy have significant destabilising effects on the European financial system. A fiscal shock normalised to a 25 bps increase in the German 10-year benchmark yield triggers an immediate and persistent widening of peripheral sovereign spreads. Italian spreads increase by approximately 10–15 bps at peak and remain persistently elevated, with Spanish and French spreads following a similar trajectory at smaller magnitudes. Virtually every real-time market indicator of systemic financial stress spikes in response.

The mechanism through which these financial effects translate into real economic outcomes operates via bank balance sheets. The repricing of the German Bund, which is the *de facto* risk-free benchmark for the euro area, generates a systemic valuation shock that propagates across asset classes and jurisdictions. Because commercial banks maintain large exposures to sovereign debt, driven by both regulatory requirements and domestic home bias, any shock to the benchmark yield directly impairs the asset side of bank balance sheets. Under the theoretical framework of Farhi and Tirole (2018), this deterioration in sovereign asset quality can trigger a sovereign-bank nexus where concerns about bank solvency and sovereign creditworthiness become self-reinforcing. I document this channel empirically: the interbank funding rates spike upon impact, banking sector CDS spreads are increasing by 7% relative to their baseline levels, and banking sector equity valuations experience a sharp contraction accompanied by a surge in realised volatility.

Banks respond to this balance sheet stress by contracting credit supply on both the extensive and intensive margins. Aggregate loan supply contracts by nearly 1% below baseline, with short-maturity lending to non-financial corporations declining most sharply. Crucially, banks also pass on interest rate risk to the private sector by increasing the share of floating-rate contracts by 1 percentage point, accelerating the pass-through of elevated funding costs to borrowers. The cost of borrowing for both corporations and households rises by approximately 18 bps, reflecting both the pass-through of the fiscally induced monetary policy tightening and a widening of credit intermediation premia.

These contractionary forces materialise as a *pre-spending penalty*: the financial tightening occurs upon the news of future debt issuance, before any fiscal stimulus is actually disbursed. The resulting credit squeeze reduces Y-o-Y GDP growth by 11 bps in Germany, 25 bps in Spain, and 29 bps in Italy. The output contractions occur within the first quarter following the shock and well before significant fiscal outlays are realised. This timing reflects the nature of the instrument: German Bund futures respond to announcements and expectations of future fiscal policy, while actual government spending operates with substantial implementation lags. The financial channel therefore operates at higher frequency than the expenditure channel, effectively rendering the short-run fiscal multiplier negligible or negative for the periphery.

To rationalise these findings, I develop a minimal model that combines the moral hazard framework of Farhi and Tirole (2018) with a two-country monetary union structure. The model delivers a closed-form expression for the doom loop amplification factor that governs the pass-through of core yield shocks to periphery yields. The key insight is that the zero risk-weight privilege of core sovereign debt under the Basel framework encourages large Bund holdings on bank balance sheets, creating substantial exposure to benchmark repricing. When a fiscal shock hits, mark-to-market losses erode bank equity, the leverage constraint tightens credit supply, and the resulting growth slowdown widens sovereign spreads, feeding back into further balance sheet deterioration. The model, calibrated with independently sourced data and parameters from the literature, matches the empirical estimates for all three countries. Approximately 80% of the periphery effect operates through the direct balance sheet channel, with the remaining 20% attributable to doom loop amplification. The model also delivers a policy implication: the introduction of Eurobonds would reshape the transmission mechanism by diluting the supply shock across the European issuer base and breaking the sovereign-bank home bias channel. The periphery benefits substantially as the doom loop is eliminated, but the core absorbs a larger share of the repricing, creating a distributional trade-off whose net effect depends on the relative size of member economies.

These results are not driven by a particular sample period, identification assumption, or modelling choice. The main findings are robust to alternative event selection criteria, estimation techniques (local projections versus proxy VAR), model specifications, and controlling for confounding factors including monetary policy surprises and global risk shocks.

Because the baseline VAR conditions on the overnight index swap rate, the estimated spread widening reflects the fiscal channel operating over and above any induced monetary policy response. The distinction is intuitive: a monetary tightening shifts the term structure roughly in parallel, whereas a fiscal supply shock reprices a single sovereign's bonds, generating the asymmetric yield responses documented in this paper. To the best of my knowledge, this is the first causal evidence on the systemic financial implications of fiscal expansion in a safe-asset anchor.

Related literature and contribution. This paper contributes to several literature strands. The first is the literature on identifying fiscal policy shocks and their effects on financial markets

(Blanchard and Perotti, 2002; D'Amico and King, 2013; Greenwood and Vayanos, 2010; Ramey, 2011; Romer and Romer, 2010). A central challenge is that fiscal policy is endogenous to economic conditions. Mountford and Uhlig (2009) identify government revenue and spending shocks using sign restrictions, while Blanchard and Perotti (2002) rely on institutional features of tax and transfer systems to achieve structural identification. However, Ramey (2011) shows that standard VAR identification strategies miss the timing of fiscal shocks because markets react when they learn about policy changes, not when spending occurs. In a related but distinct strand, Romer and Romer (2010) apply a narrative approach to identify exogenous tax policy changes, filtering legislated tax measures for those not driven by contemporaneous economic conditions. Recent work has made further progress using narrative approaches that exploit the timing of fiscal news. Wiegand (2025) estimates the effects of fiscal news shocks on US asset prices, finding that they reduce convenience yields and increase Treasury riskiness. A related strategy uses high-frequency variation around announcements. Phillot (2025) applies this to US Treasury auctions, showing that supply shocks raise yields and term premia. My paper bridges these approaches but differs critically. First, I construct a news-based instrument capturing shocks to *expectations* about future fiscal positions at the budget announcement stage, not the auction stage. This timing is crucial for Germany: by the time auction details are announced, markets have incorporated this information from prior budget debates. Second, I focus on the effects of fiscal news, but I also extend the analysis beyond convenience yields to cross-border sovereign transmission in Europe's segmented bond markets.

The second strand to which this paper contributes is the literature on safe assets and the determination of the risk-free rate. A large body of work has established that government bonds of fiscally credible sovereigns carry a convenience yield that depresses their yield below what credit fundamentals alone would imply (Krishnamurthy and Vissing-Jorgensen, 2012a; Nagel, 2016; Du et al., 2018; Jiang et al., 2021). He et al. (2019) show that safety is fundamentally relative: whether an asset retains its safe status depends on its fundamentals relative to available alternatives. At the macro level, Caballero et al. (2008); Caballero and Farhi (2018); Caballero et al. (2016) model how safe asset shortages depress the natural rate of interest and can push economies into liquidity traps with destabilising consequences for aggregate demand. A parallel theoretical literature examines the microstructure of safety. Dang et al. (2015) and Gorton and Ordonez (2014) argue that safe debt is optimally designed to be information-insensitive in normal times, but sufficiently large shocks push assets beyond a safety frontier, triggering adverse selection and liquidity freezes. Brunnermeier (2023) and Barro (2006) characterise safe assets through their service flows and hedging properties against rare disasters. I contribute to this literature by providing the first systematic causal evidence of how fiscal shocks erode the safe-asset status of the benchmark sovereign, pushing Bund collateral toward information-sensitivity and forcing investors into alternative safe havens.

This paper also contributes to the literature on financial fragmentation in monetary unions. A defining feature of the euro area is the coexistence of a single monetary policy with nationally segmented sovereign bond markets, creating a unique setting where shocks to the benchmark sovereign transmit asymmetrically across member states. De Santis (2019) documents that currency union membership does not eliminate sovereign risk but transforms it into redenomination risk. Longstaff et al. (2011) show that sovereign credit risk is significantly driven by global factors, suggesting that even within a monetary union, spreads cannot be understood purely through domestic fundamentals. Beber et al. (2009) show that in times of stress, investors shift from flight-to-quality toward flight-to-liquidity, concentrating in the most liquid instruments regardless of credit quality. Lane (2012) and Shambaugh (2012) diagnose the euro area's structural vulnerability: the absence of a centrally-issued common safe asset means that fiscal stress becomes a financial stability problem for the entire union, as the sovereign bond market serves simultaneously as store of value, pricing benchmark, and source of bank capital. While German Bunds fill this role as the de facto benchmark, they remain a national instrument without the crisis-resolution capacity of a truly supranational safe asset. De Grauwe and

Ji (2013) formalises how this structure enables self-fulfilling crises even for fundamentally solvent governments. My paper extends this work by documenting a transmission channel that operates even absent periphery stress: the core country's own fiscal expansion destabilises the union's financial architecture precisely because it serves as the risk-free benchmark.

Finally, this paper contributes to the literature on the sovereign-bank nexus. Farhi and Tirole (2018) and Brunnermeier et al. (2016) model the doom loop: when banks hold domestic sovereign debt, declining bond prices erode bank equity, tightening credit and further worsening the fiscal outlook in a self-reinforcing spiral. Acharya et al. (2018), Altavilla et al. (2017), and Bocola (2016) provide empirical evidence from the European debt crisis showing this channel operates through bank balance sheet impairment and reduced lending. My contribution is twofold. First, I show the doom loop can be activated not only by deterioration in the sovereign's own fiscal position, but also by fiscal expansion in the anchor country whose bonds sit on balance sheets across the entire union. Second, I develop a tractable two-country model nesting the Farhi-Tirole framework that delivers a closed-form amplification factor, showing approximately 80% of periphery output losses operate through the direct balance sheet channel and 20% through doom loop feedback.

Methodologically, I apply proxy-SVAR techniques to financial market microstructure pioneered by Gertler and Karadi (2015) for monetary policy and extended by Känzig (2021, 2023) to oil and carbon shocks. Stock and Watson (2012) and Mertens and Ravn (2013) develop theoretical foundations for external instrument identification. Nakamura and Steinsson (2018) advance high-frequency identification in macroeconomics. Miranda-Agrippino and Ricco (2021) show how to handle weak instruments in proxy-SVARs. This is the first application of these macro techniques to study specific financial market functioning rather than economy-wide dynamics.

Outline. The paper proceeds as follows. In Section 1.2, I explain the construction of the high-frequency instrument and the theoretical background needed for the external instrument approach. In Section 1.3, I present the results of the baseline VAR in response to a fiscal shock as well as the impact on euro area bond markets. Section 1.4 traces the source of the sovereign spread response by examining the reactions of bank lending and output. In Section 1.5, I analyse the transmission from a theoretical standpoint and show that even a highly stylised model captures the underlying dynamics quite well. Section 1.6 concludes.

1.2 Identification

1.2.1 Theoretical background

We follow the external-instrument approach as developed in (Mertens and Ravn, 2013; Stock and Watson, 2012) and as applied in Känzig (2023, 2021) or Gertler and Karadi (2015), where a proxy-variable is used to identify one particular structural shock. The theoretic background of the identification method works as follows.[1]

For a given VAR of reduced form with N variables,

$$\mathbf{y}_t = \mathbf{B}_1\mathbf{y}_{t-1} + \cdots + \mathbf{B}_p\mathbf{y}_{t-p} + \mathbf{u}_t \;\forall\; t \in \{1, \ldots, T\} \tag{1.1}$$

let $\mathbf{y}_t \in \mathbb{R}^N$ be the vector of variables, $\mathbf{B}_1, \cdots, \mathbf{B}_p \in \mathbb{R}^{N\times N}$ the vector of coefficients and p the lag order of the VAR. In this representation, let $\mathbf{u}_t$ the reduced from residual and $\mathbf{\Sigma}_u \equiv \text{Var}(\mathbf{u}_t) \in \mathbb{R}^{N\times N}$ the covariance matrix of the reduced form errors. A structural shock is an economically interpretable, exogenous, and pairwise independent source of variation that drives fluctuations in a VAR system. The observed residuals are assumed to be linked to the structural shocks via the matrix $\mathbf{S} \in \mathbb{R}^{N\times N}$ as in $\mathbf{u}_t = \mathbf{S}\varepsilon_t$. it follows that the observed covariance matrix $\mathbf{\Sigma}_u \in \mathbb{R}^{N\times N}$

[1]Throughout this section, I use bold capital letters (e.g., $\mathbf{A}$, $\mathbf{B}$) for matrices in $\mathbb{R}^{n\times m}$ where $n, m > 1$, bold lowercase letters (e.g., $\mathbf{x}$, $\mathbf{y}$) for vectors in $\mathbb{R}^{n\times 1}$ or $\mathbb{R}^{1\times n}$ where $n > 1$, and non-bold lowercase letters (e.g., a, b) for scalars.

is linked to the covariance matrix of the structural shocks $\mathbf{\Sigma}_\varepsilon \in \mathbb{R}^{N \times N}$ by

$$\mathbf{\Sigma}_u = \mathbf{S}\mathbf{\Sigma}_\varepsilon\mathbf{S}'. \tag{1.2}$$

Note that the matrix $\mathbf{\Sigma}_\varepsilon$ is diagonal by definition of the structural shocks. Assuming that our VAR is ordered such that the variable whose residual we want to identify with the structural shock is in the first position, let $\varepsilon_{1,t}$ be the diagonal element belonging to the the structural shock of interest. The goal is to recover the structural shock component, i.e., the first column of **S**, by using an instrument.

In order for z_t to be a valid instrument, it has to satisfy two conditions: It has to satisfy the *relevance condition*, meaning it has to be correlated with the structural shock, i.e.,

$$\text{Cov}(z_t, \varepsilon_{1,t}) \neq 0 \tag{1.3}$$

and it has to satisfy the *exclusion restriction*, meaning it has to be uncorrelated with all other structural shocks

$$\text{Cov}(z_t, \varepsilon_{i,t}) = 0 \;\forall\, i \in \{2, \dots, N\}. \tag{1.4}$$

These two conditions together determine the first column **S** belonging to the structural shock up to scale[2]

$$\mathbf{s}_{\cdot,1} \propto \begin{pmatrix} \mathbb{E}[z_t u_{1,t}] \\ \mathbb{E}[z_t \mathbf{u}_{2:n,t}] \end{pmatrix} \in \mathbb{R}^N. \tag{1.5}$$

The vector on the right hand side can be estimated via 2SLS regression as follows. Since $u_{1,t}$ contains noise from various structural shocks, the first stage serves to isolate the variation in

[2]For a derivation of this step, please see Appendix 1.B.1.

$u_{1,t}$ that is due to the structural fiscal shock

$$u_{1,t} = \gamma z_t + v_t. \tag{1.6}$$

The second stage serves to estimate the impact on all other residuals, given the isolated effect of the structural shock on $u_{1,t}$,

$$\mathbf{u}_{2:n,t} = \beta \hat{u}_{1,t} + \xi_t. \tag{1.7}$$

The just-identified instrumental variables estimand is the ratio of reduced-form covariances with the instrument. Since $\mathbb{E}[\mathbf{u}_t] = \mathbf{0}$, covariances reduce to moments, yielding

$$\hat{\beta}_{IV} = \frac{\text{Cov}(z_t, \mathbf{u}_{2:n,t})}{\text{Cov}(z_t, u_{1,t})} = \frac{\mathbb{E}[z_t \mathbf{u}_{2:n,t}]}{\mathbb{E}[z_t u_{1,t}]} \propto \mathbf{s}_{2:n,1}. \tag{1.8}$$

This identifies the *relative* impact vector $\mathbf{s}_{2:n,1}/s_{1,1}$, i.e., the response of all other variables per unit response of the first variable. To recover the full first column $\mathbf{s}_{\cdot,1}$ an additional step is required, which I derive in Appendix 1.B.2.

To pin down the precise value of the impact vector $\mathbf{s}_{\cdot,1}$, one usually faces a choice between two modelling assumptions. The first is to set the variance of the structural shock to unity $\mathbf{\Sigma}_\varepsilon = \mathbb{I}_n$, which implies that $s_{1,1}$ is scaled such that the impulse responses represent the impact of a one-standard-deviation shock. The scalar $s_{1,1}$ can be backed out explicitly from the equations above.[3] The second option is to scale $s_{1,1}$ so that the total impulse fits the economic interpretation best.

I scale $s_{1,1}$ to a magnitude of 25 bps. While this constitutes a substantial shock, it reflects the magnitude of large but realistic fiscal surprises observed in European sovereign debt markets. For context, Germany's announcement on March 5, 2025, of a €500 bn infrastructure fund and

[3]For a derivation of this step, please see Appendix 1.B.2.

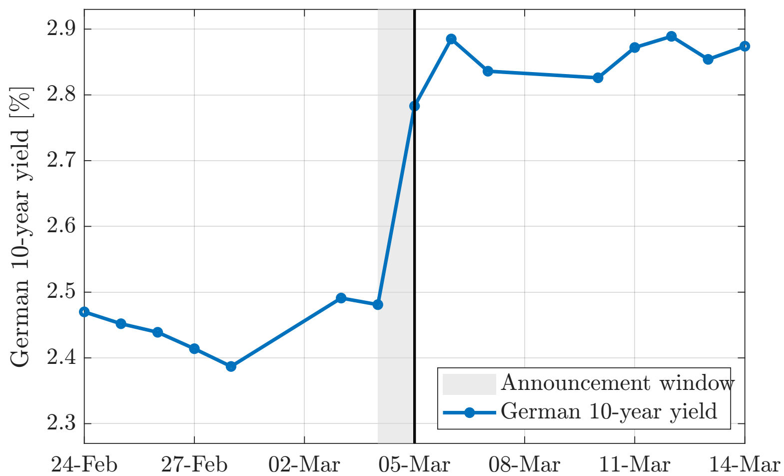

Figure 1.1. Announcement of Germany's infrastructure fund

Notes: Shaded area indicates March 5, 2025, when Germany announced a €500 billion infrastructure fund and debt brake reforms, triggering a 30 bps increase in 10-year yields.

associated borrowing rule reforms triggered an immediate 30 bps increase in German 10-year yields (see Figure 1.1). This episode demonstrates that fiscal shocks of 25 bps or larger are not merely theoretical constructs but represent the scale of policy surprises that European bond markets have experienced and are likely to encounter again.[4]

1.2.2 Instrument construction

The instrument is designed to capture shocks to expectations of future government debt, as reflected in reactions of Bund futures prices within a narrow window around news releases. The primary challenge in constructing the instrument is endogeneity: fiscal policy changes are not random but typically respond to economic shocks, such as wars, adverse macroeconomic data releases, or financial instability in other European countries. For the instrument to be valid, it must isolate the market's reaction to the fiscal policy innovation itself, not to the underlying

[4]Of course, the translation of announced spending into yield impact is likely non-linear and also depends on other factors. However, just to give some context from the Treasury market, a $1 bn net increase in supply of US Treasury securities has been found to shift yields upwards by 0.5 to 1 bps (Phillot, 2025).

economic shock that triggered the policy response. The key insight is that markets react to economic shocks and policy responses with different timing. An economic shock occurring at time t typically triggers a fiscal policy response only with a lag, at $t+1$. The total change in any fiscally-sensitive asset can thus be decomposed as:

$$\Delta \text{Asset}_{t \to t+2} = \underbrace{\Delta \text{Asset}_{t \to t+1}}_{\text{response to macro-shock}} + \underbrace{\Delta \text{Asset}_{t+1 \to t+2}}_{\text{response to fiscal policy reaction}} \tag{1.9}$$

The high-frequency identification strategy exploits this temporal separation. By measuring asset price changes in a sufficiently narrow window around the policy announcement at $t+1$, after the initial macro shock has been absorbed but while the policy surprise is being revealed, we can isolate the fiscal policy component. Setting the instrument to zero outside these event windows ensures we capture only the policy innovation, not the contaminating effects of the underlying economic shock.

I collect all news regarding the German budget between January 2016 and October 2025 from Reuters, chosen because their reporting typically leads other news sources and their coverage is widely disseminated through Bloomberg terminals to the financial sector. Since repo market data is only available from late 2016 onwards, I focus on the period thereafter, yielding 187 news releases.[5] To illustrate with one example of a news release and its market impact: on 5 March 2025, the news headline was

> ***Germany's Merz races to win support for major financial package.*** *BERLIN, March 5 (Reuters) - The parties hoping to form Germany's next government have agreed to create a 500 billion euro infrastructure fund and overhaul borrowing rules, [...]. After three years of attacking what he called the outgoing government's profligacy from opposition, Germany's likely next leader Friedrich Merz is now trying to use the narrow window before new legis-*

[5]Further information on how the news were accessed as well as the complete list of headlines can be found in Appendix 1.A.

lators officially take up their seats to pass the massive borrowing package. [. . .] – Reuters (2025)

The bond market reaction to this news was strong and began even before the official release (see Figure 1.1). It is also important to note that markets reacted well before the bill was formally signed by the German president on 22 March 2025, and long before any actual increase in bond supply materialised.

A crucial consideration is the length of the event window. There is a trade-off between capturing the full effect of the news shock and introducing excessive noise unrelated to the fiscal shock into the instrument. Potential choices range from 30 minutes to multiple days. In this context, I use a 1-day window, from the close of the previous day to the close of the current day, for several reasons. First, although Reuters is among the most prominent news sources for the financial industry, it may not always be first to release information to the public; some information could already be known to market participants by the time Reuters publishes, which is why I include the hours immediately preceding the press release in my instrument. Second, unlike scheduled announcements such as central bank interest rate decisions or debt management office issuance calendars, these news releases are unscheduled. Consequently, market participants may need additional time to become aware of the news and assess its relevance for the affected assets. For this reason, I also include several hours immediately following the announcement.

To construct a valid instrument, I need a measure of fiscal policy shocks that captures expectation changes while being orthogonal to contemporaneous risk premium variation. A wide body of literature has shown that risk-premia move slower than expectations.[6] Hence, in a short 1-day time-window, I expect most of the variation in futures prices to come from the expectations component rather than from the risk premium. This property makes futures prices

[6]See for instance Hanson and Stein (2015); Gürkaynak et al. (2007); Cochrane and Piazzesi (2005); Adrian et al. (2013)

particularly well-suited for high-frequency identification. Let $F_{h,t}$ be the futures price at time t to buy an asset at time $t+h$, S_{t+h} the spot price and $\mathbb{P}$ and $\mathbb{Q}$ the actual and the risk-neutral probability laws, respectively. With risk-neutral pricing and under the simplifying assumption that risk-premia are constant over the daily windows, we have

$$\Delta F_t = \mathbb{E}_t[S_{t+h}] - \mathbb{E}_{t-1}[S_{t+h}]. \tag{1.10}$$

Building on this insight, I construct my instrument from high-frequency movements in German Bund futures around fiscal news announcements. Let $\mathcal{N}$ be the set of days with at least one news announcement and F_t be the futures price. I then define the fiscal shock series as

$$\text{FiscalShock}_t = \begin{cases} F_t - F_{t-1} & \text{if } t \in \mathcal{N} \\ 0 & \text{otherwise.} \end{cases} \tag{1.11}$$

This daily news series is then aggregated to a weekly series by summing over all days in one week. The resulting time series is shown in Figure 1.2. Reassuringly, the fluctuations of the instrument realistically reflect the fiscal uncertainty at various points in time. Four dates show particularly high amplitudes: On 13 March 2020, Germany announced its first major COVID response, the "protective shield for employees and companies" with measures including expanded short-time work benefits and KfW loan guarantees. At 10 March 2023 Finance Minister Christian Lindner delayed presenting the 2024 budget citing unresolved spending disputes within the traffic light coalition. The delay reflected growing tensions between the fiscally conservative FDP (which insisted on reinstating the debt brake) and the SPD/Greens (which wanted continued higher spending on climate initiatives), with ministries' spending requests far exceeding available funds. By June 2024, the government was considering an additional supplementary budget of up to €11 billion to cover new shortfalls, likely related to flood relief

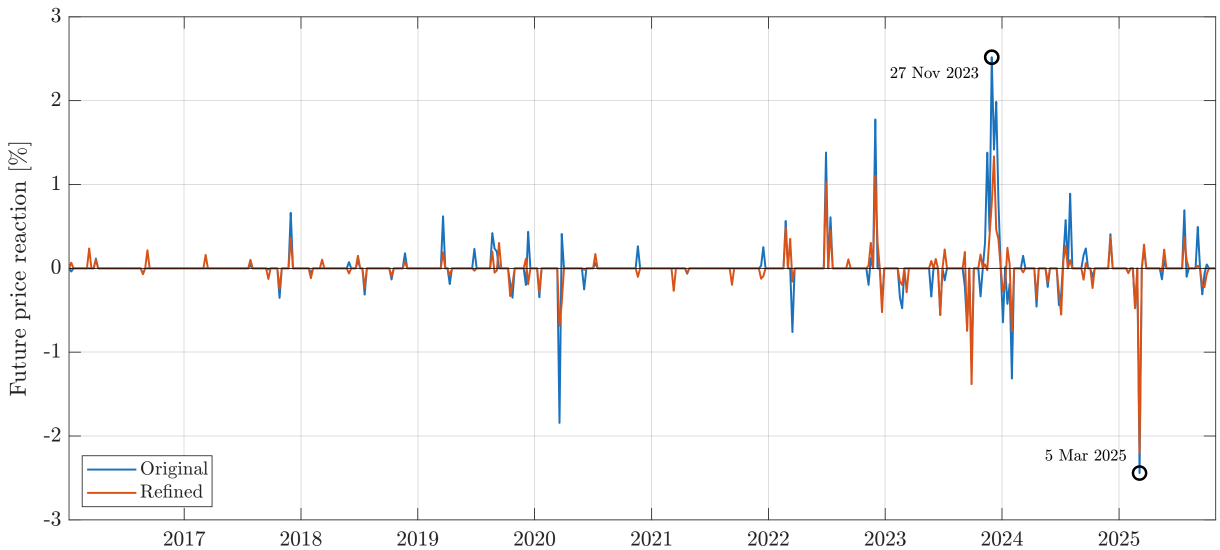

Figure 1.2. Fiscal shock series

Notes: This graph plots the weekly shock fiscal shock series as constructed by equation (1.11) and subsequent weekly aggregation.

for the devastating May-June 2024 floods in southern Germany and other unforeseen expenses that emerged during the fiscal year. In late June 2024, the coalition was still in intense negotiations over the 2025 budget amid deep divisions between Finance Minister Lindner and Chancellor Scholz and Economy Minister Habeck. Investors weighted the risks an eventual compromise requiring additional borrowing against the prospect of adhering to the restrictive debt brake.

1.2.3 Instrument validity

Relevance. As the instrument is *constructed* from the Bund reaction on event days where material news regarding the German budget was released, one important test is whether the realised variance on news dates is statistically different from placebo days that are randomly picked from the weeks where no news was released. Reassuringly, Bund futures display a sta-

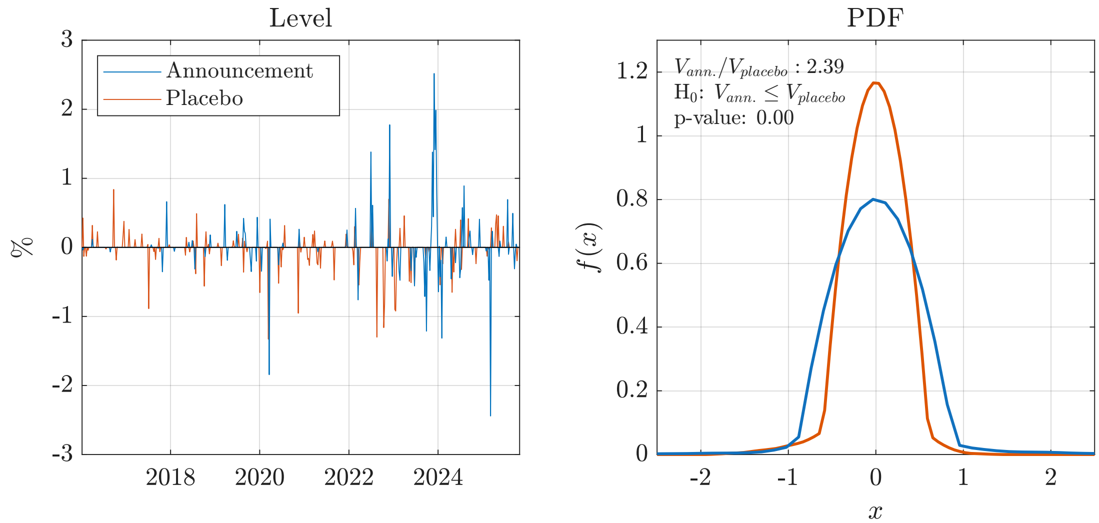

Figure 1.3. News days and placebo days in comparison.

Notes: The left hand side shows the weekly fiscal shock index constructed from real events and a weekly shock index constructed from placebo days. The right hand side chart compares the distribution of percentage changes of the daily Bund-futures series on news and placebo dates. The placebo series is constructed by drawing N days uniformly at random from all non-announcement trading days, where N equals the number of actual fiscal news dates. I repeat this procedure 100 times across independent random seeds; the placebo series shown and the associated variance ratio correspond to the draw closest to the median of the resulting distribution.

tistically significant and by 139% greater variance on days were news regarding the German budget were released. These fat tails are visible both directions (see Figure 1.3).

The relevance of the instrument can be assessed with the weak instrument test by Montiel Olea and Pflueger (2013). The robust F-statistic is 11.36 which far exceeds the critical value of 10. The adjusted R^2 is 4.55%, suggesting that the fiscal shocks explain roughly 5% of the noise in Bund yields. Thus, the null hypothesis of a weak instrument can be rejected.

Another concern could be that the instrument looks way more volatile post 2020 than pre 2020. After all, sovereign debt issuance has spiked after 2020 and asset-markets have been rattled by pandemics, trade wars and wars ever since. To alleviate concerns about my instrument picking up too much macro noise, I am comparing the impulse responses of my baseline VAR for these two historical periods in Figure 1.C.1. Reassuringly, the impulse responses for most variables are almost identical, specifically for all the fixed in come variables that I am focusing on in this paper.

Exclusion restriction. For the exclusion restriction to hold, the instrument must be orthogonal to all other structural shocks in the economy. The challenge is that structural shocks are not directly observable but must be proxied. Since I use the instrument to identify structural fiscal shocks to the German Bund yield, I must consider which other structural shocks are most likely to affect Bund yields over the sample period. The most influential shocks fall into four categories: monetary policy shocks, sovereign credit shocks spilling over from other European countries (e.g., during the euro crisis), uncertainty shocks from sources such as trade tensions, and global factors including geopolitical events or financial stress driving safe-haven flows into German Bunds. Table 1.1 tests whether these proxied shocks—either individually or jointly—predict my constructed instrument. The results show that virtually all coefficients are close to zero and statistically insignificant. In the category-specific regressions (columns 1–4), none of the 12 coefficients reaches significance at the 10% level. In the joint specification (column 5), three of the 12 coefficients reach marginal significance ($p < 0.10$). Given the number of coefficients tested, this pattern is consistent with a Type I error rate rather than genuine predictive power.

As robustness checks, I also test whether *changes* in these variables (first differences) or *lagged* values predict the fiscal shock instrument (see Appendix 1.C.3). While the contemporaneous levels specification most directly tests the exclusion restriction, these alternative specifications address concerns about non-stationarity and predictability patterns. The results remain consistent across all specifications.

To address residual concerns about confounding drivers of my instrument, I follow Känzig (2023) and purge my instrument of systematic variation attributable to global rate spillovers

Table 1.1. Predictability of fiscal shocks

Notes: OIS rates refer to the rate response around monetary policy events taken from Altavilla et al. (2019) to capture monetary policy shocks. Sovereign bond spreads refer to the difference in 10-year benchmark yields. The European sovereign CDS refers to the 5-year sovereign CDS index. Volatility is captured as the EURO STOXX 50 volatility index. German policy uncertainty is the German EPU index from Baker et al. (2016). Robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	Fiscal shock				
	(1) Monetary policy	(2) Sovereign credit	(3) Uncertainty	(4) Global factors	(5) All variables
OIS 1M	0.02				0.03
	(0.019)				(0.020)
OIS 1Y	-0.05				-0.06
	(0.040)				(0.038)
OIS 5Y	0.05				0.06*
	(0.035)				(0.034)
OIS 10Y	0.02				0.02
	(0.031)				(0.030)
Italy-Germany spread		0.01			0.01
		(0.012)			(0.016)
Spain-Germany spread		0.02			0.08
		(0.055)			(0.090)
Europe sov. CDS		-0.00			0.00
		(0.001)			(0.002)
Volatility			-0.00		-0.01*
			(0.003)		(0.005)
German policy uncertainty			0.00		0.00
			(0.000)		(0.000)
US 10Y				0.00	-0.05*
				(0.016)	(0.024)
Oil (Brent)				0.00	0.00
				(0.001)	(0.001)
Gold				-0.00	0.00
				(0.000)	(0.000)
N	514	514	514	514	514
R^2	0.01	0.00	0.01	0.00	0.03

(proxied by daily changes in the US 10-year Treasury yield), economic uncertainty (proxied by daily changes in the VSTOXX), and monetary policy surprises (proxied by the 5-year OIS response in the EA-MPD of Altavilla et al. (2019)). Specifically, I run the regression

$$s_t = \alpha + \beta_1, \text{OIS5Y}_t^{\text{EA-MPD}} + \beta_2, \Delta\text{US10Y}_t + \beta_3, \Delta\text{VSTOXX}_t + \varepsilon_t \tag{1.12}$$

on the subset of fiscal news days, and define the refined instrument as the residual $\hat{\varepsilon}_t$. Figure 1.2 compares the refined and the original series.

One desirable property of a good instrument is that it is uncorrelated with itself as this would violate the independence assumption. If the instrument is predictable from its own past, it may be endogenous which can lead to biased standard errors. By definition, the instrument should represent unexpected variation, not systematic patterns like persistent trends. Indeed, almost all autocorrelation coefficients of my instrument are statistically indistinguishable from zero (see Figure 1.D.7).

Modelling assumptions. The use of an external instrument in a proxy-VAR presumes that the system is correctly specified and captures all major structural shocks affecting government bond yields. A natural question arises: why not simply use the constructed instrument directly in a local projections framework à la Jordà (2005), as in

$$y_{i,t+h} = \beta_0^i + \psi_h^i \text{Shock}_t + \sum_{l=1}^{L} \gamma_l y_{i,t-l} + \boldsymbol{\delta}_t' \mathbf{X}_t + \xi_{i,t,h}. \tag{1.13}$$

There are two key advantages to the proxy-VAR baseline specification. First, the high-frequency instrument may still contain noise unrelated to the fiscal shock of interest. By using it as an *instrument* for the structural shock rather than treating it as the shock itself, the proxy-VAR helps filter out this measurement error. Second, the proxy-VAR framework explicitly accounts for other structural shocks occurring simultaneously, allowing for cleaner isolation of the fiscal component. Nevertheless, for robustness, I compare both approaches directly in Appendix 1.C.5. Reassuringly, the impulse responses are qualitatively very similar across methods. As expected, the LP-IV estimates exhibit somewhat greater volatility than those from the proxy-VAR, consistent with the reduced smoothing inherent in the local projections approach.

Inference. Confidence intervals are constructed using a moving block bootstrap (MBB) following Jentsch and Lunsford (2022). The MBB jointly resamples overlapping blocks of reduced-form residuals and proxy values with replacement, preserving their serial and cross-sectional dependence. Block length is set to $\ell = \lceil 1.75 \cdot T^{1/3} \rceil$ following the optimal rate in Künsch (1989), yielding $\ell = 14$ weeks (approximately one quarter). On each of $B = 10{,}000$ pseudo-samples, we re-identify the structural shock and compute IRFs. We report bias-corrected confidence bands at the 68% and 90% levels following Hall (1992), centring bootstrap quantiles on the point estimate by subtracting the bootstrap median. The MBB is preferred over wild or residual-based bootstraps because it does not require correct specification of the error distribution and accommodates the serial correlation arising from overlapping proxy windows at weekly frequency, as shown by Jentsch and Lunsford (2022). This inference approach follows Gertler and Karadi (2015), Känzig (2021, 2023).

1.3 Main results

1.3.1 Baseline VAR

The goal of my VAR specification is to account for other structural shocks that could affect the German Bund yield over time. Known drivers of yields include inflation expectations, monetary policy decisions, and EA-wide sovereign default stress. My baseline specification therefore consists of six variables: the German 10-year Bund yield, the sovereign spread[7], the corporate bond spread[8], the DAX index, the USD/EUR exchange rate[9], and the OIS 3-month rate. The OIS rate is my main control for inflation expectations and monetary policy decisions.

[7]Measured as the difference between the Italian and the German 10-year benchmark government bond yield.

[8]Measured as the difference in yields between investment-grade Euro area corporate bonds and the German benchmark bond.

[9]Just to be very clear here, this is the value of one Euro expressed in USD. If the exchange rate goes down, the Euro depreciates.

Equity and corporate bond indices are included to control for potential spillovers of risk from the corporate sector to the government. The Euro-dollar exchange rate accounts for policy risks spilling over from the US, given that the US is the Euro area's largest trading partner. Finally, I also include the sovereign spread to account for intra-Euro area shocks spilling over from other countries to Germany. I calibrate the response of all my variables to a 25 bps shock of the German bund which given its weekly standard deviation of 16 bps[10] would be a large but not unreasonable magnitude. The impulse responses to a fiscal shock in Germany on the VAR system are displayed in Figure 1.4.

Standard macroeconomic theory typically views fiscal expansions through the lens of aggregate demand and productivity. An increase in government spending might lead to higher real interest rates as a result of crowding-out or because the central bank anticipates inflationary pressures and tightens policy. Alternatively, if the spending is viewed as an investment in infrastructure or R&D, it might eventually lower real rates by shifting the supply curve and increasing long-run productivity. However, the interpretation of the shock as an expectations shock rather than an immediate spending outlay diverges from these classic channels in a fundamental way.

The baseline results suggest that a German fiscal shock operates as a systemic disruption to the euro area's financial architecture. The divergent timing of the Bund yield and the OIS rate is informative about the nature of the shock. The Bund yield spikes immediately while the OIS rate, which captures the market's expected path of ECB policy rates, rises only gradually over subsequent weeks. Since the Bund yield equals the OIS rate plus a term premium, this wedge implies that the impact effect operates almost entirely through the term premium rather than through revised expectations of monetary policy tightening. At impact, markets are repricing the compensation required for holding long-duration euro area risk, not revising their forecast

[10]The standard deviation of the 10-year German Bund yield minus the ECB deposit facility rate over the period 2020–2025.

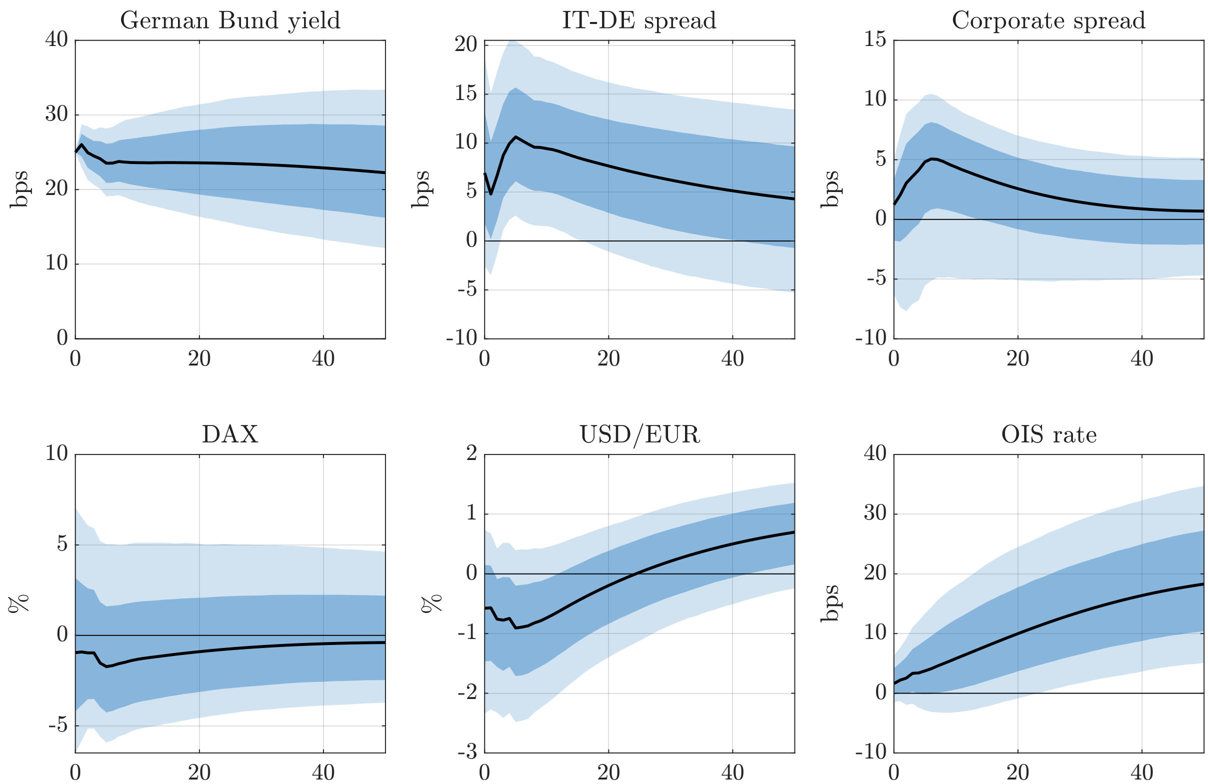

First stage regression: F: 25.38, robust F: 11.36, R^2: 4.73%, Adjusted R^2: 4.55%

Figure 1.4. Baseline response to a German fiscal shock

Notes: Figures show the response of the six baseline variables to a fiscal shock in Germany calibrated to 25 bps. Confidence intervals are shown at the 68% and the 90% level and are constructed via moving block bootstrap (10,000 iterations, 4-week blocks), robust to heteroskedasticity and autocorrelation.

of the ECB's policy rate.

The gradual OIS response is nonetheless significant. Over subsequent weeks, the 3-month OIS rate rises by approximately 15 bps, suggesting markets eventually anticipate an endogenous monetary tightening in response to the fiscal expansion, consistent with an active monetary policy regime (Leeper, 1991). Crucially, this compounds rather than offsets the initial financial tightening: by the time the ECB is expected to react, the term premium shock has already propagated through sovereign spreads and bank balance sheets.

This timing structure is central to the paper's contribution. Because the instrument captures expectations about future budgetary shifts before any stimulus enters the circular flow of income, the immediate yield increase is a pure repricing of the benchmark bond that raises the discount rate for all euro area cash flows. The financial channel therefore operates at higher frequency than either the expenditure channel or the monetary policy response. While standard theory would predict a fiscal expansion boosts equity valuations or strengthens the currency, the point estimates show a muted or negative DAX response and an ambiguous currency reaction, consistent with recent findings in the literature (Wiegand, 2025; Phillot, 2025) and supporting the interpretation that the front-loaded financial drag dominates the traditional fiscal stimulus.

Distinguishing the fiscal from the monetary channel. A natural concern is whether the documented spread widening merely reflects an induced monetary policy tightening rather than a distinct fiscal transmission mechanism. Both shocks raise government bond yields, and the literature has established that monetary policy surprises can themselves widen sovereign spreads through differential fiscal sustainability effects (Altavilla et al., 2017).

Three features of the identification address this concern. First, the baseline VAR includes the 3-month OIS rate, which serves as the primary control for monetary policy expectations and inflation risk. The impulse responses of sovereign spreads and financial stress indicators are therefore estimated conditional on the monetary policy channel, capturing the fiscal effect over and above any induced ECB response. Second, the refined instrument explicitly purges the 5-year OIS response from the EA-MPD of Altavilla et al. (2019), removing systematic covariation between the fiscal news series and contemporaneous monetary policy surprises. Third, the economic mechanisms are conceptually distinct. A monetary policy tightening operates as a roughly common factor across the euro area term structure: because it shifts the inter-temporal rate of substitution uniformly, its first-order effect is a parallel increase in sovereign yields, with spread effects arising only as a second-order consequence through differential fiscal sustainability. The German fiscal shock, by contrast, is a supply shock to a single sovereign's bonds that directly reprices the risk-free benchmark while leaving other countries' fiscal positions unchanged.

1.3.2 Sovereign spillovers

These baseline results provide an initial indication of how financial markets absorb budget surprises, but they do not reveal the underlying mechanism. One striking finding in Figure 1.4 is that the Italy-Germany spread *widens* upon impact. If the shock would just act on the risk-free rate, the change in yields should be roughly uniform across countries. However, the baseline results suggest that a more complex dynamic is at work.

A more differentiated look at all big four European countries separately, reveals striking heterogeneity across eurozone countries. Italian spreads exhibit the largest response, widening by approximately 10-15 bps at their peak, and remaining persistently elevated at 5-10 bps. Spanish and French spreads follow a similar trajectory but with slightly smaller magnitude. Importantly, this finding does not depend on the choice of the estimation method: the external instrument approach as well as the local projections show very similar results, alleviating concerns about potential misspecification driving this core result. The pattern reveals that the fiscal shock causes substantial and persistent spread widening in countries with higher debt burdens and elevated debt-to-GDP ratios.

1.3.3 Financial stress

The repricing of the Bund, which serves as the *de facto* risk-free benchmark for the euro area, generates a systemic valuation shock to sovereign bond portfolios. Given that commercial banks maintain large exposures to sovereign debt the question naturally arises whether this

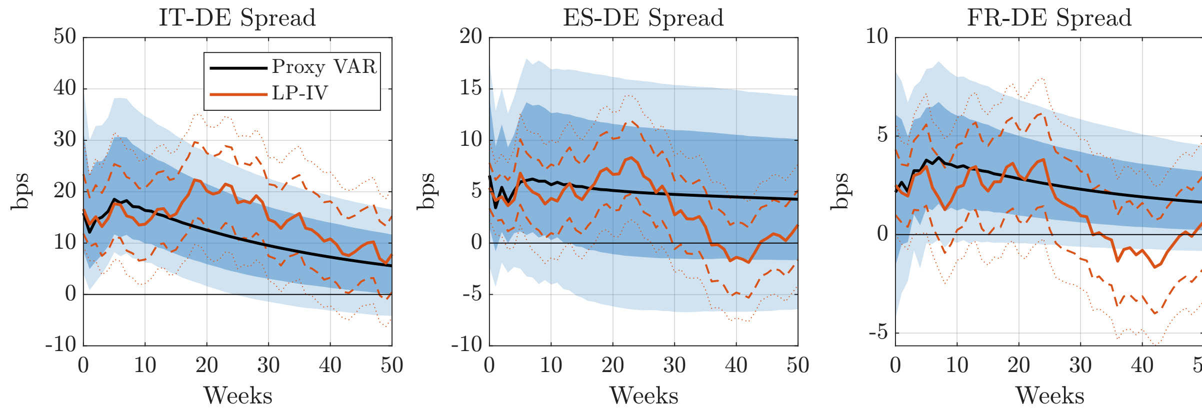

Figure 1.5. Sovereign spillovers

Notes: The charts show the impulse response of 10-year sovereign benchmark spreads with respect to the German Bund following a shock to the German 10-year yield calibrated to 25 bps. The proxy VAR is a single 8-variable system containing all three bilateral spreads simultaneously: [DE10Y, IT-DE, ES-DE, FR-DE, corporate spread, DAX, EUR/USD, OIS]. IRFs for each spread are extracted from the corresponding column. Local projections estimate each spread separately, using 6 lags of the outcome variable and the non-spread baseline controls (corporate spread, DAX, EUR/USD, OIS rate) as regressors; the other bilateral spreads are not included as LP controls. Blue shaded areas: proxy VAR 68%/90% confidence intervals. Orange dashed/dotted lines: LP 68%/90% confidence intervals. Both use moving block bootstrap (10,000 draws).

sovereign stress transmits to the banking sector. The theoretical framework of Farhi and Tirole (2018) suggests that deterioration in sovereign asset quality can trigger a sovereign-bank nexus where concerns about bank solvency and sovereign creditworthiness become self-reinforcing. To examine whether this channel operates empirically, I first turn to direct measures of banking sector stress: Euribor spreads, bank CDS premia, bank equity prices, and equity volatility.

Figure 1.6 provides empirical evidence for this intermediary distress channel. The EURIBOR-OIS spread[11] spikes by approximately 3 bps upon impact. This suggests an immediate tighten-

[11] The EURIBOR-OIS spread measures the difference between the unsecured interbank lending rate (EURIBOR) and the risk-free overnight index swap rate (OIS). Because OIS reflects only expected central bank policy rates while EURIBOR includes credit and liquidity risk premia, the spread isolates the risk premium banks charge each other for unsecured lending. A wider spread indicates higher perceived counterparty risk and tighter interbank funding conditions, making it a commonly used real-time proxy for financial stress in the banking system.

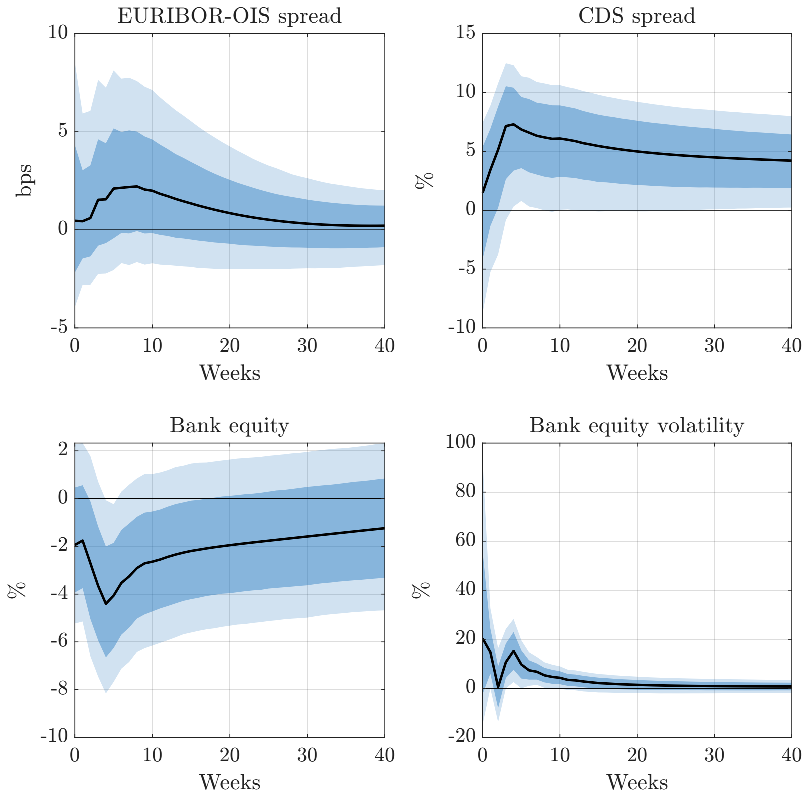

Figure 1.6. Financial stress

Notes: The charts show the impulse responses of the EURIBOR-OIS spread, the European senior financials CDS index, the European banking sector sub-index and its weekly realised return volatility. Confidence intervals are shown at the 68% and the 90% level and are constructed via moving block bootstrap (10,000 iterations, 4-week blocks), robust to heteroskedasticity and autocorrelation.

ing of interbank funding conditions as counterparty risk concerns intensify.

Even more pronounced is the reaction of the European senior financials CDS index, which rises significantly by 7% (relative to its baseline level) and remains persistently elevated. This surge in the price of default protection confirms that the market perceives a fundamental increase in the fragility of the banking sector's capital position. These concerns are mirrored in the equity markets: banking sector equity valuations experience a sharp contraction, accompanied by a surge in realised equity volatility. The co-movement of these indicators (a spike in

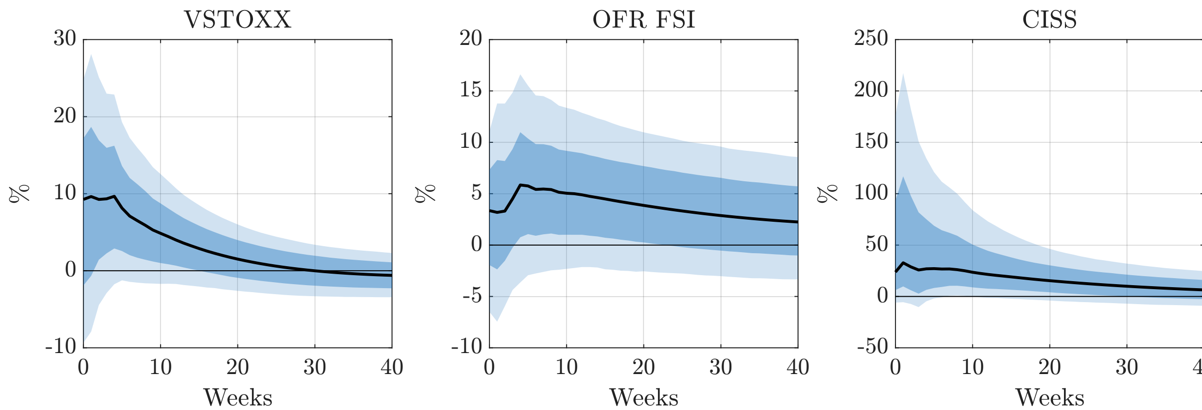

Figure 1.7. Financial stress indicators

Notes: Confidence intervals are shown at the 68% and the 90% level and are constructed via moving block bootstrap (10,000 iterations, 4-week blocks), robust to heteroskedasticity and autocorrelation.

interbank funding costs, rising default risk premia, and collapsing equity values) constitutes a classic signature of financial stress.

The finding that fiscal shocks elevate financial stress proves robust across alternative market-based stress indicators. To strengthen this result, I compare impulse response functions across several widely-used, real-time measures of financial stress in the Euro area: the EA VIX index, the OFR Financial Stress Index, and the European Central Bank's Composite Indicator of Systemic Stress (CISS). While these indicators differ in construction and scale, making quantitative comparisons challenging, they exhibit a consistent qualitative pattern that corroborates the earlier findings.

Safe-asset status and the convenience yield. A central question in the safe-asset literature is whether an increase in the supply of government debt erodes the safety premium or convenience yield (He et al., 2019), the non-pecuniary benefit investors derive from the liquidity and safety of an asset. I estimate the cumulative historical contribution of these shocks to a proxy for the convenience yield: the spread between the German Bund and an index of AAA-rated

Euro area corporate bonds (as also done in Krishnamurthy and Vissing-Jorgensen (2012a) for instance).

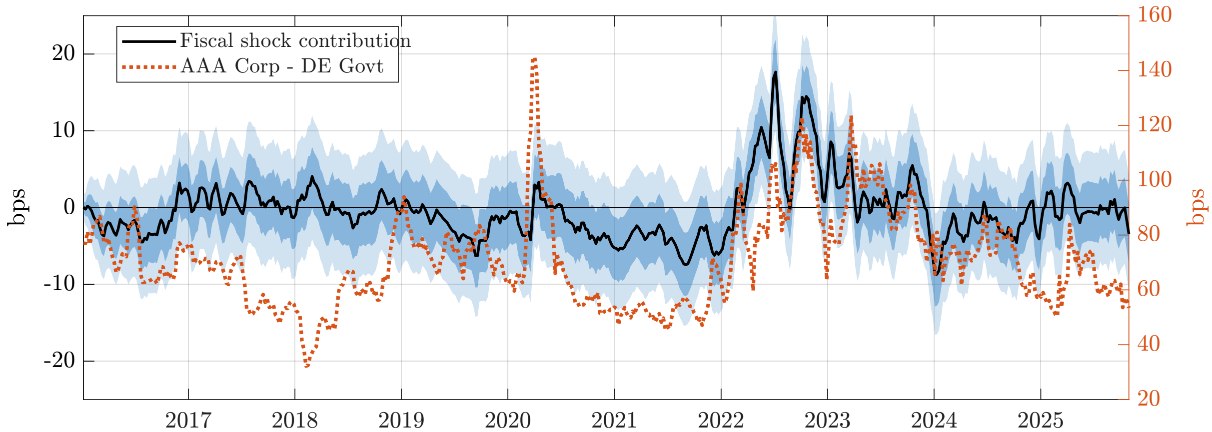

Figure 1.8. Contribution to convenience yield

Notes: The chart shows the cumulative historical contribution of German budget news shocks to the German convenience yield (measured as the spread between the German Bund and the yield to maturity of an index comprised of AAA-rated corporate Euro area bonds) and the 68 and 90 percent confidence bands together with the actual convenience yield (in percent deviations from mean).

As shown in Figure 1.8, the average convenience yield in the sample (2016–2025) is 67 bps. This is remarkably consistent with the 73 bps found in the U.S. Treasury market by Krishnamurthy and Vissing-Jorgensen (2012a), validating the Bund's role as the euro area's primary safe haven.

However, the contribution of fiscal shocks to this premium is regime-dependent, providing empirical support for the theoretical framework of He et al. (2019). Fiscal shocks compressed the convenience yield by up to 10 bps during periods of elevated issuance (2020–2022, 2024), but contributed positively to the premium in late 2022. As shown below, these periods of convenience yield compression coincide closely with positive fiscal shock contributions to the price of gold (Figure 1.9), suggesting that investors substitute toward alternative safe assets at the margin when the Bund's safety premium is under pressure.

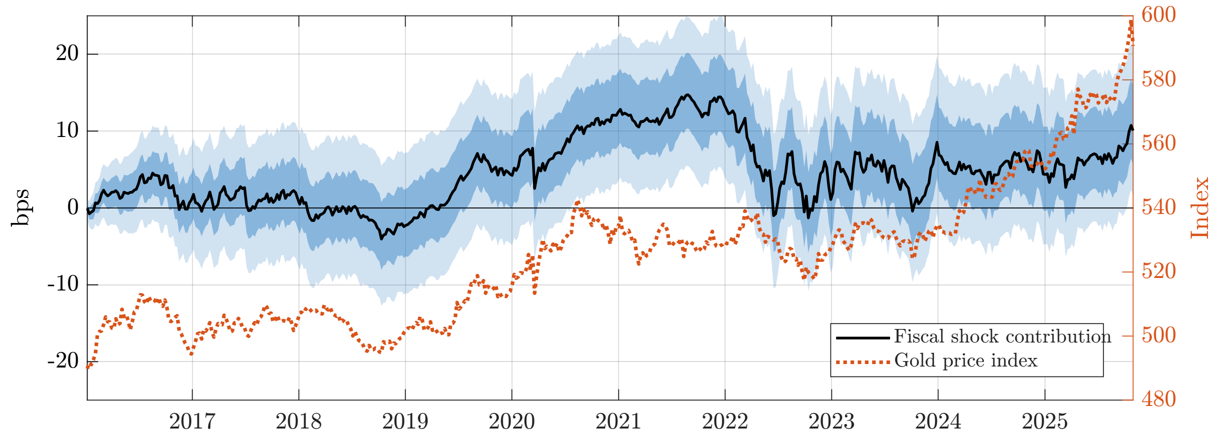

Figure 1.9. Contribution to the price of gold

Notes: The chart shows the cumulative historical contribution of German budget news shocks to the gold price index and the 68 and 90 percent confidence bands together with the actual gold price index (in percent deviations from mean).

Substitution into alternative safe havens: the case of gold. If the German Bund's safety premium is periodically squeezed by fiscal shocks, do investors seek refuge in alternative assets? To test this flight-to-safety hypothesis without the confounding effects of other sovereigns' fiscal policies[12], I examine the Gold Price Index as an outside safe asset.

Figure 1.9 reveals a striking symmetry: the periods in which fiscal shocks made their most negative contributions to the Bund's convenience yield (2020–2022) align closely with the periods in which those same shocks drove significant price increases in gold.

While the point estimates are subject to the wide confidence intervals typical of historical decompositions in volatile periods, the co-movement is theoretically telling. It suggests that at the margin, investors look for other safety assets when there is doubt about the German fiscal balance sheet. When the risk-free status of the Bund is even slightly diluted by fiscal expectations, capital migrates toward other alternatives such as gold.

[12] A natural choice for another safe-asset would be U.S. treasuries. However, they were rocked by U.S. fiscal shocks in the same time period, raising concerns about confounding factors.

1.4 Economic consequences

1.4.1 Bank lending contraction

In the previous section, I have analysed the reactions of various financial variables to a German budget news shock. The main working hypothesis is that this shock transmits directly to bank's balance sheets by shifting the risk-free rate in the fixed income market and effectively acts as a shock of financial stress. In this section, I aim to go one step further and ask: what are the economic consequences of this? The first thing to look at would be the response of bank's lending to businesses and households.

The fiscal-financial shock triggers a substantial and heterogeneous contraction in bank lending. Figure 1.10 shows that aggregate total assets of monetary financial institutions decline by approximately 0.35% at their trough, while total loans contract more severely, falling by nearly 1% below baseline. This differential response of loans contracting more sharply than total assets suggests that banks are not merely undergoing a mechanical balance sheet compression but are actively rebalancing their portfolios away from credit provision in response to heightened risk.

The primary mechanism for this contraction operates through the management of duration risk. When the fiscal shock widens sovereign spreads, banks holding these securities face mark-to-market losses that erode equity, increasing their duration exposure relative to capital. Rather than absorbing this heightened exposure, banks actively hedge by restructuring the terms of their private-sector lending. Figure 1.10 provides the clearest evidence of this risk-transfer channel: the share of floating-rate loans increases by 1 percentage point. By shifting borrowers onto floating-rate contracts, banks effectively transfer interest rate risk from their own balance sheets to households and firms, allowing them to shorten their effective portfolio duration while protecting their capital positions.

This credit contraction exhibits sharp cross-country and maturity heterogeneity that mirrors

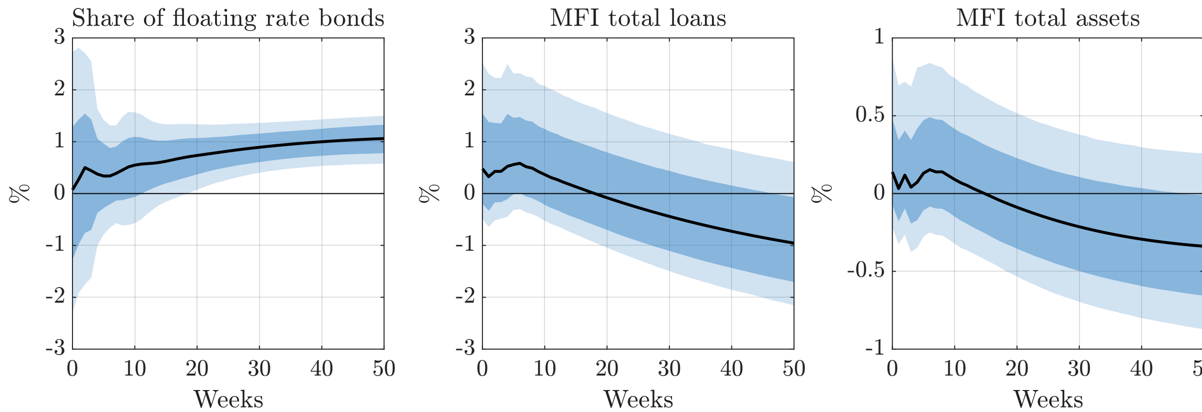

Figure 1.10. Aggregate banking sector reaction

Notes: The charts show the impulse response of the share of floating-rate loans extended to euro area non-financial corporations, the stock of loans by monetary financial institutions (the ECB's terminology for commercial banks) and the stock of assets of monetary financial institutions. Confidence intervals are shown at the 68% and the 90% level and are constructed via moving block bootstrap (10,000 iterations, 4-week blocks), robust to heteroskedasticity and autocorrelation.

the sovereign spread responses. As shown in Figure 1.11, short-maturity loans (1–5 years) contract substantially more than longer-maturity credit across all big four economies. In Germany, 1–5 year lending falls by approximately 0.6%, while longer-maturity credit declines by only 0.4%. The divergence is even more pronounced in Spain and Italy, where short-maturity lending drops by nearly 0.9%. One explanation for why short maturities react more strongly than longer ones is of technical nature. Short-term loans reach their maturity date more frequently, offering banks regular opportunities to adjust their credit supply at the extensive margin. Following a negative capital shock from the repricing of sovereign holdings, banks can passively deleverage by choosing not to renew or extend short-term credit lines as they expire, without requiring explicit contract renegotiation or default. In contrast, longer-term loans are contractually locked in until maturity, offering fewer natural adjustment points. This mechanical difference in adjustment frequencies may channel the credit supply contraction disproportionately toward firms dependent on short-term bank financing.

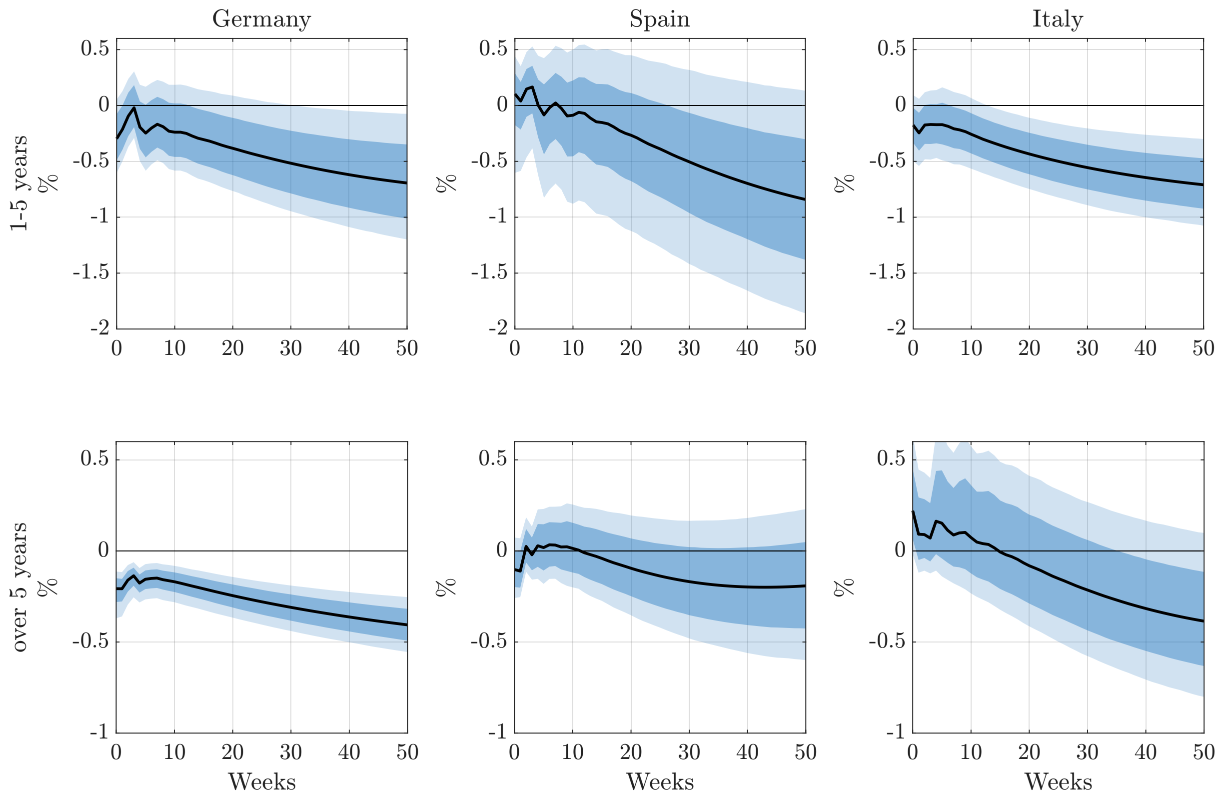

Figure 1.11. Lending to non-financial corporations

Notes: Confidence intervals are shown at the 68% and the 90% level and are constructed via moving block bootstrap (10,000 iterations, 4-week blocks), robust to heteroskedasticity and autocorrelation.

These supply-side frictions manifest as a direct transmission of bank funding stress to the real economy. Figure 1.12 illustrates that the cost of borrowing for non-financial corporations and households rises by about 18 bps. This increase likely reflects both the pass-through of the anticipated monetary policy tightening induced by the fiscal shock and an additional credit intermediation premium as interbank and CDS spreads widen. While the borrowing cost response alone is consistent with standard interest rate transmission, the concurrent lending contraction and shift toward floating-rate contracts documented above are not, pointing to a balance sheet stress channel operating on top of the induced monetary policy pass-through.

The shift toward floating-rate contracts documented earlier has two implications for the

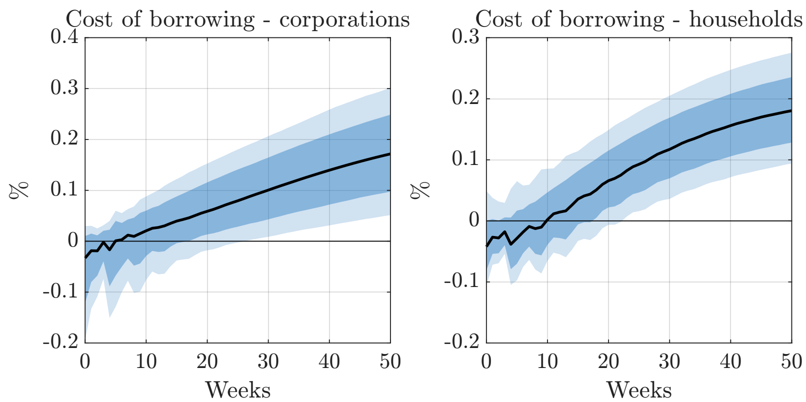

Figure 1.12. Cost of borrowing

Notes: Confidence intervals are shown at the 68% and the 90% level and are constructed via moving block bootstrap (10,000 iterations, 4-week blocks), robust to heteroskedasticity and autocorrelation.

transmission of fiscal shocks. First, it insulates bank balance sheets from interest rate risk by transferring duration exposure to borrowers. Second, it accelerates the pass-through of elevated interbank rates to the real economy through contractual loan re-indexing. This repricing mechanism ensures that increases in funding costs are reflected in borrower-facing rates. The persistence of these rate increases indicates a sustained rather than transitory adjustment in credit pricing. This sustained increase in borrowing costs likely dampens investment and consumption, particularly in economies with high levels of private debt where exposure to variable-rate contracts is most prevalent.

1.4.2 The fiscal multiplier

The preceding analysis establishes that a German fiscal shock generates financial stress that propagates heterogeneously across the euro area. This impacts bank's balance sheets and leads to a contraction in lending activity. I now examine the dynamic response of real GDP growth to quantify these real-sector consequences.

Figure 1.13 shows that the Y-o-Y GDP growth rate contracts in all three countries follow-

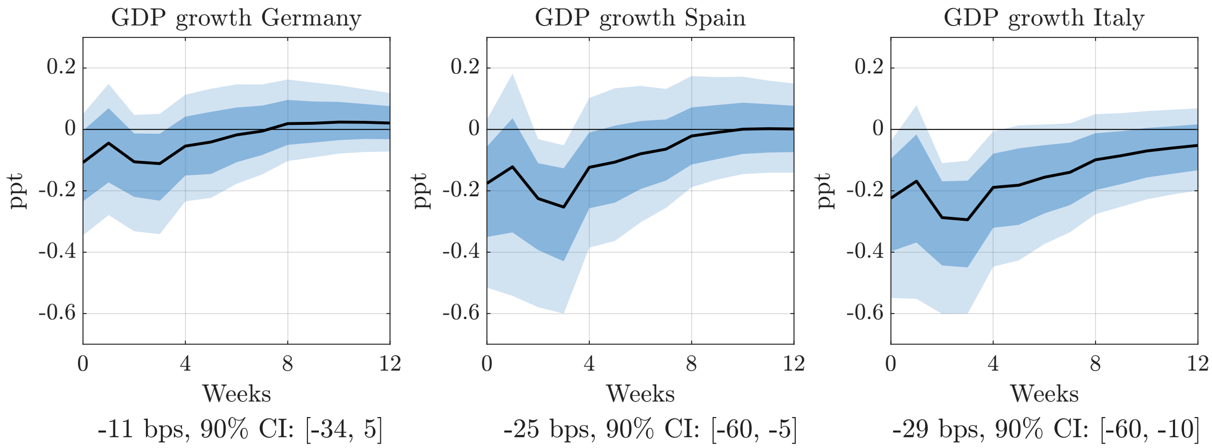

Figure 1.13. GDP growth

Notes: Confidence intervals are shown at the 68% and the 90% level and are constructed via moving block bootstrap (10,000 iterations, 4-week blocks), robust to heteroskedasticity and autocorrelation.

ing the shock, with substantial cross-country heterogeneity: German GDP declines by 11 basis points at its trough, while Spanish and Italian GDP fall by 25 and 29 basis points respectively.

The output contractions occur within the first quarter following the shock—well before significant fiscal outlays are realised. This timing reflects the nature of the instrument: German Bund futures respond to announcements and expectations of future fiscal policy, while actual government spending operates with substantial implementation lags. The financial channel therefore operates at higher frequency than the expenditure channel. Sovereign spread widening and credit supply adjustments occur within weeks of the announcement, while the stimulative effects of government purchases materialise gradually over subsequent quarters. Said differently, the fiscal multiplier is decreased by 10 to 20 bps before the spending even took place, just through the implication to financial markets.

Why do Euro area spreads widen so much? The striking divergence in output responses, where Italy and Spain contract nearly three times as much as Germany, can be structurally explained by the impact of the shock on sovereign debt sustainability: as documented in Subsection 1.3.2, the German fiscal shock drives a systemic repricing of the risk-free benchmark, raising the marginal cost of borrowing r across the entire union. Second, the banking sector deleveraging and credit supply tightening documented in Section 1.4.1 drive an immediate contraction in economic activity, thereby lowering the expected growth path g. The widening of peripheral spreads is the direct market reaction to this deteriorating $r - g$ outlook. In countries where the $r - g$ margin is already thin due to high initial debt levels, even a moderate increase in the risk-free rate combined with a credit-led slowdown can shift the perceived fiscal trajectory toward an unsustainable path. This increases the sovereign risk premium, which in turn feeds back into higher borrowing costs for domestic banks, reinforcing the sovereign-bank doom loop. To test this hypothesis, I estimate a threshold proxy-VAR in which the identification is conditioned on the lagged interest rate–growth differential ($r - g$), splitting the sample into periods of relatively high and low fiscal vulnerability over 2016–2025.[13] Figure 1.C.5 reports the regime-dependent impulse responses. The results reveal a clear asymmetry: sovereign spreads widen significantly more in response to German fiscal shocks when $r - g$ is above the 25th percentile, consistent with amplified transmission under fiscal stress. For France and Spain, the spread response is muted or even reverses sign in the low-$r - g$ regime, suggesting that fiscal spillovers to these countries operate primarily when debt sustainability concerns are elevated.

This mechanism explains the persistent widening of Italian and Spanish spreads observed throughout the 50-week horizon. While Germany's fiscal stimulus may eventually support aggregate demand, the financial chilling effect is instantaneous and concentrated in jurisdictions with high debt-to-GDP ratios. In these economies, the financial tightening dominates the traditional fiscal multiplier, as the increase in r and the decline in g occur long before any positive spillovers from German spending can materialize. Consequently, a core fiscal expansion could paradoxically become a source of recessionary pressure for the periphery by destabilising the financial conditions necessary for debt sustainability.

[13] Details on the threshold proxy-VAR methodology are provided in Appendix 1.C.4.

1.4.2.1 Does the benchmark role matter?

The preceding results show that a German fiscal shock generates large spillovers to peripheral sovereign yields and real activity. A natural question is whether these spillovers simply reflect the cross-border transmission of any large fiscal shock in the euro area, or whether they are specific to shocks originating in the country whose debt serves as the risk-free benchmark.

To distinguish between these hypotheses, I estimate a symmetry test. I construct an analogous fiscal surprise instrument for Italy and estimate a proxy-VAR with identical specification.[14] The test compares two ratios. The first, DE→IT, measures the Italian yield response per unit of German yield movement following a German fiscal shock: $(\Delta y^{\mathrm{BTP}})/(\Delta y^{\mathrm{Bund}})$. The second, IT→DE, measures the German yield response per unit of Italian yield movement following an Italian fiscal shock: $(\Delta y^{\mathrm{Bund}})/(\Delta y^{\mathrm{BTP}})$. If the transmission were symmetric—driven, for instance, by trade linkages or common monetary policy expectations—these two ratios should be statistically indistinguishable.

Table 1.2 reports the results. The asymmetry is stark. A German fiscal shock that raises Bund yields by 1 basis point raises BTP yields by approximately 1.45 basis points at horizons of 4 to 12 weeks. By contrast, an Italian fiscal shock that raises BTP yields by 1 basis point raises Bund yields by only 0.3 basis points, and this estimate is not statistically different from zero at any horizon. The difference between the two ratios is significant at the 5% level at the 4-week horizon and at the 10% level at all other horizons.

This pattern is consistent with the benchmark channel. When Germany, the issuer of the euro area's de facto risk-free asset, expands fiscally, the repricing of the Bund mechanically raises the floor under all euro-denominated sovereign yields and triggers the amplification mechanisms documented above. An Italian fiscal expansion, while relevant for Italian spreads,

[14] Since BTP futures are insufficiently liquid for high-frequency identification, the Italian instrument uses the daily change in the 10-year BTP benchmark yield on Italian fiscal announcement days. The same orthogonalisation controls (EA-MPD monetary policy surprises, US 10-year yield changes, and VSTOXX changes) are applied. Details are provided in Appendix 1.G.1.

Table 1.2. Asymmetry in fiscal shock transmission

Notes: DE→IT reports the ratio of the Italian BTP yield response to the German Bund yield response following a German fiscal shock, $(\Delta y^{\text{BTP}})/(\Delta y^{\text{Bund}})$. IT→DE reports the ratio of the Bund yield response to the BTP yield response following an Italian fiscal shock. Difference is DE→IT minus IT→DE. 90% confidence intervals constructed from 10000 moving block bootstrap replications, recentered around point estimates. *, **, *** denote significance of the difference at the 10%, 5%, and 1% level, respectively.

	DE → IT		IT → DE		Difference	
Horizon	Estimate	90% CI	Estimate	90% CI	Estimate	90% CI
1	1.25	[0.86, 1.71]	0.26	[-0.52, 0.78]	0.99*	[0.26, 1.93]
4	1.46	[1.07, 2.02]	0.35	[-0.24, 0.79]	1.11**	[0.48, 2.01]
8	1.47	[1.06, 1.98]	0.28	[-0.54, 0.83]	1.19*	[0.45, 2.22]
12	1.45	[1.03, 1.97]	0.29	[-0.65, 1.03]	1.17*	[0.29, 2.30]

does not reprice the risk-free benchmark and therefore does not generate comparable union-wide spillovers. The asymmetry confirms that what distinguishes the German fiscal shock is not its size, but the special role of the Bund in the euro area's financial architecture.

1.5 Model

This section develops a minimal model that rationalises the empirical findings. The model combines the moral hazard framework of Farhi and Tirole (2018) with a two-country monetary union structure to trace how fiscal news in the core country transmit through bank balance sheets to credit and output across the union.

1.5.1 Model set-up

Environment. Consider a monetary union consisting of two regions: core (C) and periphery (P), with GDP weights n and $1 - n$. The model has two dates, $t = 0$ and $t = 1$, and two equilibria are compared: a baseline in which core fiscal policy is unchanged, and a counterfactual in which the core government announces a fiscal expansion $\Delta G^C > 0$ at $t = 0$. The key timing assumption is that forward-looking bond markets reprice immediately upon announcement, before any fiscal outlays occur. Each region is populated by a representative bank, a fiscal authority, and a production sector that requires bank credit to produce output. There are no households. Banks fund themselves from a perfectly elastic pool of wholesale funding at exogenous rate $\bar{r}$, set by the common central bank or international capital markets. This follows the convention in Farhi and Tirole (2018): depositors are passive, and all the action is on bank balance sheets. Core sovereign bonds serve as the risk-free benchmark for the monetary union, capturing the role of the German Bund as the de facto safe asset in the Euro area. All other yields are priced as spreads over the core yield.

Production. Output in country $i \in \{C, P\}$ is a reduced-form function of bank credit

$$y_i = A_i \, \ell_i^{\alpha}, \quad \alpha \in (0, 1) \tag{1.14}$$

There is no labour market. Credit is the sole bottleneck for production. This is sufficient to close the real side of the model: anything that restricts ℓ_i contracts y_i.

Banks and timing. The representative bank in country i is managed by a risk-neutral banker. The model proceeds in two stages.

Stage 1 (portfolio formation, $t = 0^-$). Before any fiscal news arrives, the bank holds a legacy portfolio of sovereign bonds and loans, funded by wholesale deposits and equity:

Assets	**Liabilities**
Core bonds b_i^C at price q^C	Wholesale funding d_i at rate $\bar{r}$
Periphery bonds b_i^P at price q^P	Equity e_i
Loans ℓ_i	

The balance sheet identity is

$$q^C b_i^C + q^P b_i^P + \ell_i = e_i + d_i \tag{1.15}$$

where bond prices are $q^j = (1+r^j)^{-1}$. The bond portfolio (b_i^C, b_i^P) is predetermined at this stage—it reflects past investment decisions, regulatory requirements, and domestic home bias. These holdings are not adjusted in response to the fiscal news shock. This is the empirically relevant assumption: sovereign bond portfolios adjust slowly over quarters, whereas the fiscal news shock reprices assets within days.[15]

Stage 2 (repricing and adjustment, $t = 0^+$). The core government announces a fiscal expansion $\Delta G^C > 0$. Forward-looking bond markets reprice: yields adjust, bond prices change from q^j to $\hat{q}^j$, and mark-to-market gains or losses hit bank equity on the *fixed* bond portfolio:

$$\hat{e}_i = e_i + (\hat{q}^C - q^C)\, b_i^C + (\hat{q}^P - q^P)\, b_i^P \tag{1.16}$$

Given the new equity $\hat{e}_i$, the bank adjusts its lending ℓ_i and deposits d_i to maximise expected terminal dividends, subject to balance sheet feasibility and the incentive compatibility constraint defined below. The bond portfolio remains fixed at (b_i^C, b_i^P).

The bank's objective at Stage 2 is to maximise the $t = 1$ terminal value of its portfolio:

$$\max_{\ell_i, d_i} \quad \text{div}_i = (1 + r_i^\ell)\, \ell_i + (1 - \hat{\pi}^C) b_i^C + (1 - \hat{\pi}^P) b_i^P - (1 + \bar{r})\, d_i \tag{1.17}$$

where $\hat{\pi}^j$ denotes the (post-announcement) default probability of country j's bonds, with $\hat{\pi}^C = 0$ for the core.[16] Subject to (EQ) evaluated at post-shock prices $(\hat{q}^C, \hat{q}^P)$ and $\text{div}_i \geq 0$. Since bond holdings are fixed, the only margin of adjustment is ℓ_i (with d_i determined residually by the balance sheet identity).

[15] The EBA Transparency Exercise data show that the cross-sectional distribution of sovereign exposures across euro area banks is highly persistent at semi-annual frequency.

[16] Since $\hat{q}^j = (1 - \hat{\pi}^j)/(1 + \bar{r})$ under risk-neutral pricing, the objective can equivalently be written in $t = 0^+$ market-value terms as $(1 + \bar{r})^{-1}\text{div}_i = \ell_i/(1 + \bar{r}) \cdot (r_i^\ell - \bar{r}) + \hat{e}_i$. The banker thus maximises a weighted combination of lending spread income and equity value.

Credit risk. Assets differ in their exposure to credit and default risk. Let ω^j capture the asset-specific risk parameter:

$$\omega^C = 0, \qquad \omega^P \geq 0, \qquad \omega^\ell > 0$$

The assumption $\omega^C = 0$ reflects the zero risk-weight treatment of core sovereign debt under the Basel framework. These bonds are considered risk-free and face no default risk. Periphery sovereign bonds carry some credit risk ($\omega^P > 0$), while bank loans carry the highest credit risk (ω^ℓ highest).

Incentive compatibility. Depositors are rational and atomistic. They observe the banker's portfolio and supply funding d_i only if the banker has sufficient "skin in the game" to deter asset diversion. Following Farhi and Tirole (2018), I model this as a reduced-form capital constraint: the banker can divert a fraction $\phi\omega^j$ of each asset class j, and depositors require that the banker's equity stake—which equals the continuation value under the normalisation that the banker is risk-neutral and discounts at rate $\bar{r}$—weakly exceeds the diversion payoff. The payoff from diversion is $\phi(\omega^C \cdot \hat{q}^C b_i^C + \omega^P \cdot \hat{q}^P b_i^P + \omega^\ell \cdot \ell_i)$. Using $\omega^C = 0$, the constraint requires

$$\hat{e}_i \geq \phi \left(\omega^P \cdot \hat{q}^P b_i^P + \omega^\ell \cdot \ell_i \right) \tag{1.18}$$

Core bonds do not appear in the constraint. This is crucial: it means banks accumulated large core bond portfolios in Stage 1 without tightening their constraint, and are now simultaneously exposed to large mark-to-market losses when the benchmark reprices in Stage 2.

Lending earns a spread over the funding cost ($r_i^\ell > \bar{r}$), since firms' marginal product of credit exceeds the risk-free rate. The banker therefore always wishes to expand lending, and the only limit is the depositors' willingness to fund. In Stage 2, where equity has been depleted by mark-to-market losses, the constraint (1.18) binds, and the maximum loan supply is

$$\ell_i^s = \frac{\hat{e}_i - \phi\,\omega^P\,\hat{q}^P b_i^P}{\phi\,\omega^\ell} \tag{1.19}$$

This is the key equation. Credit supply is pinned by post-shock equity, with leverage multiplier $(\phi\omega^\ell)^{-1}$.

Yield determination. The core government announces at $t = 0$ a future fiscal expansion $\Delta G^C > 0$. Forward-looking markets require a higher yield to absorb the anticipated additional supply. In reduced form

$$r^C = \bar{r}^C + \mu \cdot \Delta G^C, \quad \mu > 0 \tag{1.20}$$

where μ captures the supply elasticity of the core bond market. This is the exogenous impulse. The periphery yield[17] is the core yield plus a spread

$$r^P = r^C + \sigma^P. \tag{1.21}$$

The spread depends on the $r - g$ differential

$$\sigma^P = \bar{\sigma} + \delta\left(r^C - g^P\right), \quad \delta > 0, \tag{1.22}$$

where g^P is Periphery growth, which depends on credit

$$g^P = \psi \cdot \ell^P, \quad \psi > 0. \tag{1.23}$$

[17]The model takes bond yields as reduced-form objects rather than deriving them from investor optimisation and market clearing. This is without loss for the transmission results. The parameter μ is disciplined by the observed yield response to fiscal news, and δ is taken from independent estimates in De Grauwe and Ji (2013). Alternatively, both could be micro-founded via standard CARA demand for core bonds and risk-neutral pricing of periphery bonds without altering the equilibrium equations or the comparative statics of Lemma 1 and Proposition 1.

This is the minimal specification that generates the doom loop. Substituting the credit supply and growth equations into the spread equation yields a single equation in r^P. Since $\lambda < 1$ implies that this mapping is a contraction, the Banach fixed-point theorem guarantees that a unique equilibrium exists (see Appendix 1.F.3 for details). The model has six equations and six endogenous variables. Table 1.3 summarises the mapping.

Table 1.3. Model components

Equation	Endogenous variable
$r^C = \bar{r}^C + \mu\Delta G^C$	Core yield
$\hat{e}_i = e_i + \Delta q^C b_i^C + \Delta q^P b_i^P$	Bank equity (post-shock)
$\ell_i = (\hat{e}_i - \phi\omega^P \hat{q}^P b_i^P)/\phi\omega^\ell$	Credit supply
$r^P = r^C + \bar{\sigma} + \delta(r^C - \psi\ell^P)$	Periphery yield
$y_i = A_i \ell_i^\alpha$	Output
$q^j = (1 + r^j)^{-1}$	Bond prices

1.5.2 Transmission

We now trace the impact of the fiscal news shock $\Delta G^C > 0$ announced at $t = 0$. From the yield determination equation, the announcement raises r^C by $\mu \cdot \Delta G^C$. The bond price change is obtained by a first-order Taylor approximation, $\Delta q^C \approx \frac{dq^C}{dr^C}\big|_{r^C} \cdot \Delta r^C$,

$$\Delta q^C = -\frac{\mu\Delta G^C}{(1 + r^C)^2} < 0 \tag{1.24}$$

From (EQ), mark-to-market losses hit both banks

$$\Delta\hat{e}_i = \Delta q^C \cdot b_i^C + \Delta q^P \cdot b_i^P < 0 \tag{1.25}$$

The loss is proportional to sovereign bond holdings. Because $\omega^C = 0$ encouraged large core bond portfolios, the exposure is substantial. From (1.19), with the (1.18) constraint binding:

$$\Delta \ell_i = \frac{\Delta \hat{e}_i}{\phi \omega^\ell} - \frac{\omega^P}{\omega^\ell} \Delta(\hat{q}^P b_i^P) \tag{1.26}$$

The leverage multiplier $(\phi\omega^\ell)^{-1}$ amplifies equity losses into credit cuts. From the production function, we get $\Delta y_i = \alpha A_i \ell_i^{\alpha-1} \cdot \Delta \ell_i < 0$. Output contracts before any fiscal spending occurs. For notational ease, I define the following objects.

Definition 1. Let the following objects be defined:

1. The effective periphery bond exposure: $\tilde{b}_P^P \equiv (1 - \phi\omega^P) b_P^P$.

2. The price sensitivity of periphery bonds: $\mathcal{D}^P \equiv (1 + r^P)^{-2}$.[18]

3. The leverage multiplier: $\Lambda \equiv (\phi\omega^\ell)^{-1}$.

4. The doom loop eigenvalue:

$$\lambda \equiv \delta \cdot \psi \cdot \Lambda \cdot \mathcal{D}^P \cdot \tilde{b}_P^P \tag{1.27}$$

   which is the product of the spread sensitivity to r–g (δ), the credit-growth elasticity (ψ), the leverage multiplier (Λ), and the duration-weighted effective sovereign exposure ($\mathcal{D}^P \tilde{b}_P^P$).[19]

The amplification multiplier that acts on the sovereign spreads following a fiscal shock in a core country is then pinned down by

[18] This is the absolute price sensitivity $|dq^P/dr^P|$ of a one-period zero-coupon bond, not the modified duration in the standard fixed-income sense (which would be $(1 + r^P)^{-1}$). In the calibration of Section 1.5.3, $(1 + r^P)^{-2}$ is replaced by the portfolio-weighted modified duration $\bar{\mathcal{D}}^P$ of actual sovereign bond holdings, which captures the maturity structure of bank portfolios.

[19] In the calibration of Section 1.5.3, the one-period price sensitivity $(1 + r^P)^{-2}$ is replaced by the portfolio-weighted modified duration $\bar{\mathcal{D}}^P$ of actual sovereign bond holdings, and the bond exposure $\tilde{b}_P^P$ is scaled by a constraint-relevant share θ_P to reflect accounting classification and capital buffer attenuation. See Section 1.5.3 for details.

Lemma 1 (Doom loop amplification). *Suppose the incentive compatibility constraint (1.18) binds for the periphery bank and $\lambda < 1$. Then the pass-through of core yields to periphery yields is*

$$\frac{\Delta r^P}{\Delta r^C} \geq (1+\delta) \cdot \mathcal{A} \tag{1.28}$$

where

$$\mathcal{A} = \frac{1}{1-\lambda} > 1 \tag{1.29}$$

is the doom loop amplification factor. If $\lambda \geq 1$, no stable equilibrium exists.

Proof. See Appendix 1.F. □

The bound is tight when periphery banks' core bond holdings are small relative to their domestic sovereign exposure, which is the empirically relevant case given the strong home bias documented in the data. This provides a tractable equation for the amplification factor. Specialising to the benchmark case $\omega^P = 0$ (zero risk weight for euro area sovereigns under the Basel standardised approach), so that $\tilde{b}_P^P = b_P^P$:

$$\mathcal{A} = \frac{1}{1 - \underbrace{\delta \cdot \psi}_{\text{spread-credit-growth}} \cdot \underbrace{\frac{1}{\phi\omega^\ell}}_{\text{leverage}} \cdot \underbrace{\frac{b_P^P}{(1+r^P)^2}}_{\text{bond exposure}}} \tag{1.30}$$

The doom loop is stable ($\mathcal{A}$ finite) when the product in the denominator is less than one, and explosive (crisis) otherwise. The amplification is increasing in δ (the sensitivity of spread to $r-g$), ψ (the credit-growth elasticity), $(\phi\omega^\ell)^{-1}$ (the bank leverage), and b_P^P (the periphery banks' domestic sovereign exposure).

The announcement of a core fiscal expansion triggers immediate credit contraction and output decline in both countries. However, the periphery suffers disproportionately. The core economy absorbs only the direct benchmark repricing effect, while the periphery bears a double burden: the same benchmark repricing plus the amplified impact transmitted through the doom loop channel captured by $\mathcal{A}$. This asymmetric transmission explains why announcement effects are systematically larger in countries with stronger sovereign-bank linkages.

Definition 2. Let the core financial drag be defined by

$$F_C \equiv \alpha A_C \ell_C^{\alpha-1} \cdot \Lambda \cdot \mathcal{D}^C \cdot b_C^C \cdot \mu, \tag{1.31}$$

and the periphery financial drag be defined by

$$F_P \equiv \alpha A_P \ell_P^{\alpha-1} \cdot \Lambda \cdot \left[\mathcal{D}^C b_P^C + \mathcal{D}^P \tilde{b}_P^P \cdot \Gamma \right] \cdot \mu, \tag{1.32}$$

where $\Gamma = (1+\delta) \cdot \mathcal{A}$ is the total periphery yield pass-through from Lemma 1.

The term in brackets has two components: the direct core bond exposure of periphery banks $\mathcal{D}^C b_P^C$ and the doom-loop-amplified periphery bond exposure $\mathcal{D}^P \tilde{b}_P^P \cdot \Gamma$. Let $\mathcal{M}_C > 0$ be the textbook fiscal multiplier for the core country capturing the standard Keynesian demand channel of domestic fiscal spending on domestic output, absent financial frictions.[20]

Proposition 1 (Fiscal multiplier erosion)**.** *Suppose the incentive compatibility constraint* (1.18) *binds for banks in both countries and the doom loop is stable* ($\lambda < 1$)*. Then the effective union-wide fiscal multiplier from a core fiscal expansion* $\Delta G^C > 0$ *announced at* $t = 0$ *is*

$$\hat{\mathcal{M}}^{MU} = n \cdot \mathcal{M}_C - n \cdot F_C - (1-n) \cdot F_P. \tag{1.33}$$

Proof. See Appendix 1.F. □

Some immediate conclusions are worth highlighting. First, the erosion is strictly positive ($\hat{\mathcal{M}}^{MU} < n \cdot \mathcal{M}_C$) for any $\mu > 0$ and non-zero bond holdings. Second, the periphery drag

[20] To be precise: $\mathcal{M}_C \equiv \frac{\Delta y}{\Delta G^C}$.

exceeds the core drag ($F_P > F_C$) whenever $\Gamma > 1$ and periphery banks' domestic sovereign exposure is sufficiently large relative to core banks' Bund holdings, a condition that is amply satisfied empirically given the home bias documented in the EBA data.[21] And lastly, the erosion is increasing in the doom loop eigenvalue λ. As $\lambda \to 1$, $F_P \to \infty$, approaching the crisis threshold which causes unbounded multiplier destruction.

1.5.3 Calibration

The closed-form expressions of Proposition 1 provide the structural link between observable yield movements and output losses. This section feeds independently sourced data into each link of the transmission chain and tests whether the model-implied output losses match the empirical estimates from Section 1.3.

The calibration strategy is partial-equilibrium: I take the yield movements as given from the empirical proxy-VAR and ask whether the model's balance sheet chain—from yield shocks through bank equity losses, leverage constraints, and credit contraction to output—produces GDP responses consistent with those estimated independently. This approach follows the calibration methodology of Farhi and Tirole (2018) and Bocola (2016), who derive structural amplification formulae analytically and evaluate them using supervisory data, rather than solving for bond market equilibrium numerically. The advantage is transparency: every step from yields to output is observable and independently verifiable, without relying on equilibrium selection in a nonlinear system.

I evaluate the model for three countries—Germany, Spain, and Italy—spanning the full range of estimated output responses from Section 1.3. All results are conditional on a 25 bps increase in the Bund yield, matching the normalisation in the empirical identification.

From Proposition 1, the model-implied output loss in country j can be written in two equiv-

[21] Formally, $F_P > F_C$ requires $A_P \ell_P^{\alpha-1}[\mathcal{D}^C b_P^C + \mathcal{D}^P \tilde{b}_P^P \Gamma] > A_C \ell_C^{\alpha-1} \mathcal{D}^C b_C^C$. Given that $\tilde{b}_P^P \gg b_P^C$ and $\Gamma > 1$, this holds for the empirically relevant parameter range.

alent forms. In levels:

$$\Delta y_j \;\approx\; \underbrace{\alpha \frac{y_j}{\ell_j}}_{\substack{\text{marginal product}\\ \text{of credit}}} \times \Delta \ell_j \tag{1.34}$$

where $\alpha(y_j/\ell_j)$ is the marginal product of credit evaluated at the pre-shock allocation. Equivalently, in percentage terms:

$$\frac{\Delta y_j}{y_j} \;\approx\; \alpha \;\times\; \frac{\Delta \ell_j}{\ell_j} \tag{1.35}$$

where α is the credit-output elasticity. I use the levels form (1.34) in the calibration (Table 1.5), since country-specific marginal products matter for cross-country comparisons.

The lending contraction $\Delta \ell_j$ is itself determined by the mark-to-market loss on the bank's sovereign portfolio, scaled by effective leverage:

$$\frac{\Delta \ell_j}{\ell_j} \;=\; -\underbrace{\Lambda_j \cdot \theta_j}_{\substack{\text{effective}\\ \text{leverage}}} \;\times\; \underbrace{\left[\bar{\mathcal{D}}^{C} \frac{b_j^{C}}{\ell_j}\, dr^{C} \;+\; \bar{\mathcal{D}}^{P_j} \frac{\tilde{b}_j^{P_j}}{\ell_j}\, dr^{P_j} \right]}_{\text{duration-weighted mark-to-market loss per unit of credit}} \tag{1.36}$$

Each component on the right-hand side is either directly observable or well-identified in the literature. Table 1.4 reports all parameter values, organised by source.

Crucially, the empirical analysis provides an intermediate target in addition to the final output loss. Section 1.3 estimates that the aggregate lending contraction falls in the range of -0.4% to -0.8% of baseline credit. This pins down equation (1.36) independently of the output elasticity, permitting validation at two separate points in the transmission chain.

Duration convention. The model prices bonds as one-period zero-coupon instruments for tractability, implying a price sensitivity of $(1 + r^P)^{-2} < 1$, which is far below the interest rate exposure of actual bank portfolios (where typical remaining maturities are between 5 and 10 years). To bridge this gap without complicating the analytical framework, I replace $(1 + r^j)^{-2}$

Table 1.4. Baseline parameter values

Notes: Credit data are outstanding loans to euro area non-financial corporations from ECB Balance Sheet Item statistics, sample average 2016–2024. Sovereign exposures are from EBA Transparency Exercise, 2023 vintage. GDP is Eurostat annual nominal GDP, sample average. Modified durations are portfolio-weighted for the 10-year benchmark of each sovereign. The fair-value share denotes the approximate fraction of the bank's domestic sovereign bond portfolio that is marked to market.

Parameter	Symbol	Germany	Spain	Italy	Source
Panel A: From the empirical analysis					
Bund yield shock	dr^C		25 bps		Normalisation
Spread reaction	$dr^{P_j} - dr^C$	—	+8 bps	+11 bps	Section 1.3
Lending contraction	$\Delta\ell_j/\ell_j$	−0.4%	−0.7%	−0.8%	Section 1.3
Panel B: From ECB and supervisory data					
GDP	y_j (€bn)	4,100	1,400	2,100	Eurostat
Bank credit to NFCs	ℓ_j (€bn)	1,500	500	700	ECB BSI
Credit-to-GDP ratio	ℓ_j/y_j	0.37	0.36	0.33	Computed
Sovereign exposure	$b_j^{P_j}$ (€bn)	—	≈ 170	≈ 400	EBA Transparency
Fair-value share		—	$\approx 50\%$	$\approx 29\%$	ECB supervisory
Eff. exposed	$\tilde{b}_j^{P_j}$ (€bn)	—	≈ 85	≈ 180	See text
Core bond holdings	b_j^C (€bn)	≈ 500	≈ 40	≈ 60	ECB SHS-S
Modified duration	$\bar{\mathcal{D}}^j$	7.5	6.5	7.0	Bloomberg
CET1 ratio		16.0%	13.5%	15.5%	ECB CBD
Leverage multiplier	Λ_j	8	8	8	See text
Panel C: From the literature					
Credit-output elasticity	α		0.30		Schularick and Taylor (2012, Table 4)
Spread sensitivity to $r-g$	δ		0.10		De Grauwe and Ji (2013)

with the portfolio-weighted modified duration $\bar{\mathcal{D}}^j$ in all calibration equations. This preserves the closed-form structure of Lemma 1 and Proposition 1 while matching the empirically relevant maturity profile.[22]

Sovereign exposures and accounting treatment. The single largest source of measurement uncertainty is the effective sovereign exposure of bank balance sheets. Under IFRS 9, bonds

[22] Equivalently, one can re-derive the model with multi-period bonds, replacing $q^j = (1 + r^j)^{-1}$ with $q^j = (1 + r^j)^{-\tau_j}$. The first-order approximation $\Delta q^j \approx -\bar{\mathcal{D}}^j \cdot \Delta r^j$ then holds.

classified as “held to collect” (amortised cost) do not generate income-statement losses from yield changes. For Italian banks, roughly 71% of BTP holdings are classified at amortised cost.[23] However, the full portfolio remains relevant for the economic value of equity, for repo collateral valuation, and for supervisory stress tests. The baseline uses an “effectively exposed” share of approximately 40–50% of the total domestic sovereign portfolio for Italy and Spain, intermediate between the fair-value-only and full-portfolio assumptions, yielding $\tilde{b}^P_{\text{IT}} \approx 180$ €bn and $\tilde{b}^P_{\text{ES}} \approx 85$ €bn.

Leverage multiplier. The leverage multiplier $\Lambda_j = (\phi\omega^\ell)^{-1}$ maps equity losses into credit contractions. Banks face two capital requirements that imply different values for Λ: the risk-weighted CET1 ratio (13–16%), which excludes zero-risk-weight sovereign bonds and implies $\Lambda \approx 6$–7, and the unweighted leverage ratio (5–6%), which includes them and implies $\Lambda \approx 17$–18. Since neither requirement consistently dominates, the baseline sets $\Lambda = 8$, an intermediate value.

Constraint-relevant equity loss. The leverage multiplier $\Lambda = 8$ maps equity losses into credit contractions when the entire mark-to-market loss tightens the binding constraint. In practice, three frictions attenuate this mapping. First, sovereign bonds classified as “held to collect” under IFRS 9 do not generate regulatory capital losses from yield changes; the fair-value share is approximately 29% for Italian banks and 40–50% for Spanish and German banks. Second, banks carry CET1 buffers of 3–5 percentage points above the minimum requirement, so only the fraction of the equity loss that exhausts the buffer tightens lending. Third, the model’s comparative static maps to a gradual empirical adjustment unfolding over 10–20 weeks rather than a discrete one-period jump. I capture these attenuations through a constraint-relevant

[23]European Commission In-Depth Review, 2024.

share $\theta_j \in (0, 1)$, so that the effective credit contraction becomes

$$\Delta \ell_j \;=\; \Lambda \;\cdot\; \theta_j \;\cdot\; \Delta e_j. \tag{1.37}$$

The parameter θ_j is calibrated to match the empirical intermediate target for lending contractions from Panel A of Table 1.4. The resulting values are $\theta_{\text{DE}} = 0.08$, $\theta_{\text{ES}} = 0.17$, and $\theta_{\text{IT}} = 0.12$. Spanish banks have the highest fair-value share of their domestic sovereign portfolio ($\approx 50\%$) combined with the thinnest CET1 buffer above the minimum (13.5%), making the largest fraction of mark-to-market losses constraint-relevant. Italian banks have a lower fair-value share ($\approx 29\%$) but intermediate buffers (15.5%). German banks operate with the largest capital cushion (16.0%), implying that only a small fraction of Bund revaluation losses tightens the lending constraint.

Spread sensitivity and doom loop amplification. The parameter δ governs the sensitivity of periphery spreads to the interest rate–growth differential $r - g$. From Lemma 1, the total pass-through from core to periphery yields is

$$\frac{dr^P}{dr^C} \;=\; (1 + \delta) \cdot \mathcal{A}. \tag{1.38}$$

For Italy, the empirical pass-through is $36/25 = 1.44$; for Spain, $33/25 = 1.32$. Since δ and $\mathcal{A}$ enter multiplicatively, they cannot both be pinned by this single moment. I therefore calibrate δ independently and let $\mathcal{A}$ follow from the model's structural equations.

The baseline sets $\delta = 0.10$, at the lower end of the estimates of sovereign spread sensitivity to fiscal fundamentals in De Grauwe and Ji (2013), who report elasticities of peripheral spreads to debt-to-GDP ratios in the range of 0.10–0.25.[24] Given $\delta = 0.10$, the counterfactual periphery

[24] Similar magnitudes are implied by Bocola (2016) and by Altavilla et al. (2017). The parameter δ captures only the $r{-}g$ component of spread sensitivity; the lower end of the range is appropriate because the total empirical spread response already includes non-$r - g$ channels such as flight-to-quality and

yield absent any doom loop feedback ($\mathcal{A} = 1$) is $dr^{P,\text{cf}} = (1 + \delta) \times 25 = 27.5$ bps, corresponding to a mechanical spread of 2.5 bps. The observed spreads of 8 bps (Spain) and 11 bps (Italy) substantially exceed this counterfactual, and the gap pins the amplification factors from Lemma 1:

$$\mathcal{A}_{\text{ES}} = \frac{1.32}{1.10} = 1.20, \qquad \mathcal{A}_{\text{IT}} = \frac{1.44}{1.10} = 1.31. \tag{1.39}$$

Both factors exceed unity, confirming that the doom loop is active, but remain well below the explosive threshold ($\lambda < 1$). The implied eigenvalues are $\lambda_{\text{ES}} = 0.17$ and $\lambda_{\text{IT}} = 0.24$.

Credit-output elasticity. The parameter α governs the sensitivity of output to credit supply. The baseline value of 0.30 draws on the long-run estimates of Schularick and Taylor (2012) and is consistent with the range reported in the credit channel literature.[25] Because the model uses the elasticity at observed credit-to-GDP ratios, the effective marginal product of credit $\alpha \cdot y_j/\ell_j$ varies across countries even with a common α. The ratio y_j/ℓ_j ranges from 2.7 (Germany) to 3.0 (Italy), generating a 10% differential in the output sensitivity to a given percentage credit contraction.

Results. Table 1.5 reports the model-implied output loss for each country, decomposed into each link of the transmission chain. The table uses the observed periphery yields throughout, which embed the doom loop amplification from Lemma 1. The doom loop decomposition is presented separately below.

For Germany, the model produces an output contraction of -12 bps against an empirical target of -11 bps. Because Germany experiences no domestic spread reaction, the doom loop is inoperative by construction ($\mathcal{A}_{\text{DE}} = 1$), and the entire effect operates through mark-to-market losses on the Bund portfolio.

liquidity premia.

[25] See Mian et al. (2017) and the survey in Brunnermeier et al. (2016). Values in the literature range from 0.15 to 0.40 depending on identification strategy and time horizon.

Table 1.5. Model-implied output loss: three-block decomposition

Notes: All figures conditional on a 25 bps Bund yield shock. The table uses observed periphery yields from Panel A, which embed the doom loop amplification $\mathcal{A}_j$. Sovereign exposures use the effectively-exposed assumption. The constraint-relevant share θ_j is calibrated to match the empirical lending contraction (see text for the interpretation of this calibration choice). Block 2 follows from Block 1 via $\Delta\ell_j = \Lambda \cdot \theta_j \cdot \Delta e_j$. Output loss is computed in levels as $\Delta y_j = \alpha(y_j/\ell_j) \cdot \Delta\ell_j$ per equation (1.34), then expressed as a share of GDP.

	Germany	**Spain**	**Italy**
Block 1: Yield shock to equity loss			
Core yield shock (dr^C)	25 bps	25 bps	25 bps
Periphery yield (dr^{P_j})	—	33 bps	36 bps
Equity loss: core bonds (€bn)	-9.38	-0.75	-1.13
Equity loss: domestic bonds (€bn)	—	-1.82	-4.54
Total equity loss (Δe_j, €bn)	-9.38	-2.57	-5.66
Block 2: Equity loss to credit contraction			
Leverage multiplier (Λ)	8	8	8
Constraint-relevant share (θ_j)	0.08	0.17	0.12
Credit contraction ($\Lambda \cdot \theta_j \cdot \Delta e_j$, €bn)	-6.0	-3.5	-5.4
Credit contraction ($\Delta\ell_j/\ell_j$)	-0.40%	-0.70%	-0.78%
Block 3: Credit contraction to output loss			
Marginal product of credit ($\alpha \cdot y_j/\ell_j$)	0.82	0.84	0.90
Model output loss ($\Delta y_j/y_j$)	**-12 bps**	**-21 bps**	**-24 bps**
Empirical estimate	-11 bps	-27 bps	-29 bps

For Spain and Italy, the model produces -21 and -24 bps, capturing approximately 78% and 83% of the respective empirical estimates. The model undershoots because it isolates the bank balance sheet channel and omits several complementary mechanisms: non-bank financial intermediation, confidence effects on investment demand beyond the credit margin, and second-round demand propagation as the initial credit contraction ripples through the economy. Since the output loss is a direct function of the calibrated lending contraction (via equation (1.34)), the more informative comparison is at the intermediate step: Block 1 sovereign exposures and durations, combined with $\Lambda = 8$, imply constraint-relevant shares of $\theta_{\text{DE}} = 0.08$, $\theta_{\text{ES}} = 0.17$, and $\theta_{\text{IT}} = 0.12$. These values are independently consistent with the product of observed fair-value classification shares and estimated buffer-utilisation fractions, as discussed above. This consistency across three countries with different banking structures and sovereign exposures is the model's core quantitative validation.

It is important to be precise about what the calibration does and does not test. Because θ_j is calibrated to match the empirical lending contraction, the output loss in Block 3 is largely a mechanical transformation of the lending target and the assumed elasticity α: the model does not independently predict the output response. What the calibration *does* test is whether the implied θ_j values—backed out from Block 1 exposures and the lending target—are plausible given observable institutional features. If the sovereign exposures or durations in Block 1 were wrong, or if the leverage multiplier Λ were miscalibrated, then θ_j would take implausible values (outside the range consistent with fair-value shares and capital buffers), and the calibration would fail this plausibility check. The substantive discipline therefore flows from Block 1 to θ_j, not from θ_j to output. The cross-country variation in θ_j provides an additional consistency test: it must line up with observable differences in accounting classification and bank capitalisation, which it does.

Figure 1.14 visualises the cross-country pattern by varying the periphery debt-to-GDP ratio B_P/Y in the model while holding all other parameters fixed. The model-implied output loss is increasing and convex in sovereign indebtedness, consistent with the nonlinear amplification structure of Lemma 1. The empirical point estimates for Germany, Spain, and Italy—identified independently in Section 1.3—fall close to the model curve, with all three lying within their respective 90% confidence intervals. The dashed line shows the core output response, which is approximately flat since the core country does not experience spread feedback.

Doom loop decomposition. To isolate the contribution of the doom loop, I compare the full model to a counterfactual in which the feedback from credit to growth to spreads is shut off. In this counterfactual, periphery yields rise by the mechanical pass-through $dr^{P,\text{cf}} = (1+\delta) \cdot dr^C = 27.5$ bps rather than the observed equilibrium values. The doom loop—the self-reinforcing cir-

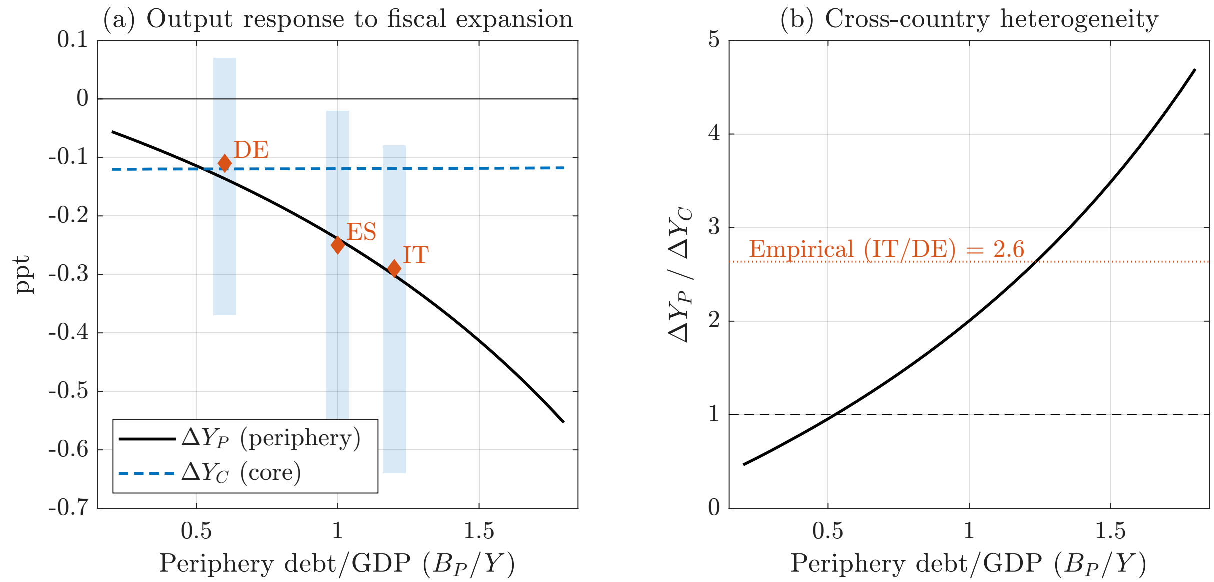

Figure 1.14. Model-implied output response to a core fiscal expansion

Notes: Panel (a) plots the model-implied output response (ppt) as a function of the periphery debt-to-GDP ratio B_P/Y, holding all other parameters at their baseline values from Table 1.4. The solid line is the periphery response ΔY_P; the dashed line is the core response ΔY_C. Diamond markers show the empirical point estimates from Section 1.3 for Germany ($B/Y = 0.60$), Spain ($B/Y = 1.00$), and Italy ($B/Y = 1.20$); shaded rectangles indicate the corresponding 90% confidence intervals. Panel (b) plots the ratio $\Delta Y_P/\Delta Y_C$, with the empirical IT/DE ratio shown as a dotted horizontal line.

cuit from sovereign repricing through bank equity erosion, credit contraction, growth slowdown, and back to wider spreads—is switched off by setting $\mathcal{A} = 1$.

Table 1.6 reports the decomposition. For Italy, the counterfactual output loss is -19 bps compared to -24 bps in the full model. The 5 bps difference—approximately 21% of the total model-implied effect—is the doom loop contribution. For Spain, the doom loop contributes approximately 2 bps (10% of the total). The cross-country pattern is consistent with Lemma 1: Italy has a larger amplification factor ($\mathcal{A} = 1.31$ versus 1.20) because its banks hold a larger domestic sovereign portfolio relative to credit, and the country's higher debt-to-GDP ratio makes the $r - g$ margin more sensitive to the initial shock.

Table 1.6. Doom loop decomposition

Notes: "Direct" uses the counterfactual periphery yield $dr^{P,\text{cf}} = (1+\delta) \cdot dr^{C} = 27.5$ bps that obtains when the doom loop is shut off ($\mathcal{A} = 1$). "Full model" uses the observed equilibrium yield. The doom loop contribution is the difference. All other parameters held fixed.

	Spain	**Italy**
Periphery yield		
Direct ($\mathcal{A} = 1$)	27.5 bps	27.5 bps
Full model (observed)	33 bps	36 bps
Doom loop adds to spread	+5.5 bps	+8.5 bps
Equity loss (€bn)		
Direct	-2.27	-4.59
Full model	-2.57	-5.66
Output loss (bps)		
Direct channel	-19	-19
Full model	-21	-24
Doom loop contribution	-2	-5
Amplification factor $\mathcal{A}_j$	1.20	1.31
Doom loop share of model output	10%	21%

1.5.4 Discussion

Two features of the calibration stand out. First, approximately 80% of the model-implied periphery output loss operates through the direct balance sheet channel (mark-to-market losses eroding equity and tightening the leverage constraint), and the remaining 20% can be attributed to the doom loop amplification. The direct channel is common to all countries. Cross-country variation arises from differences in sovereign exposures, capital buffers, and credit-to-GDP ratios. Second, the doom loop is active but stable: the eigenvalue λ lies between 0.17 and 0.24, generating finite amplification of $\mathcal{A} \approx 1.2$–1.3. This contrasts with the euro area debt crisis of 2010–12, where estimates suggest λ approached the explosive threshold.

The Eurobond counterfactual. A natural policy question is whether mutualising sovereign debt would attenuate the transmission mechanism. Figure 1.15 decomposes the model-implied output loss into three components: the direct balance sheet channel (with the doom loop shut off, $\lambda = 0$), the additional loss from doom loop amplification, and the counterfactual under unified Eurobond issuance. If Eurobonds replaced national sovereign bonds as the benchmark safe asset, three things would change simultaneously: the supply shock would be diluted across the entire European issuer base (μ falls), the sovereign-bank home bias channel would be broken as banks hold Eurobonds b_i^{EU} instead of domestic debt b_i^P, and the doom loop amplification factor would collapse toward one ($\mathcal{A} \to 1$). However, the gains are not uniformly distributed. While the periphery benefits substantially (the elimination of the spread channel reduces its output loss by approximately 17%), the core absorbs a larger share of the repricing shock, since mark-to-market losses previously concentrated in national Bund holdings are now spread across a jointly guaranteed bond stock. The net union-wide effect depends on the relative economic size of core and periphery, a distributional trade-off at the heart of the political economy of Eurobond proposals.

What explains the cross-country pattern? The model decomposes the Germany–Italy gap into three sources. First, Italian banks absorb larger mark-to-market losses (from both Bund and BTP holdings) and have a higher constraint-relevant share ($\theta_{\text{IT}} = 0.12$ versus $\theta_{\text{DE}} = 0.08$), producing a credit contraction of 0.8% compared to 0.4% in Germany. Second, Italian credit is more productive at the margin: the ratio y_j/ℓ_j is 3.0 in Italy versus 2.7 in Germany, so the same percentage credit contraction generates a proportionally larger output loss. Third, the doom loop amplifies the Italian effect by an additional 5 basis points while leaving Germany unaffected. The relative importance of these three channels is approximately 35%, 25%, and 40% of the total Germany–Italy differential.

What the model does not capture. The model accounts for approximately 80% of the empirical periphery output losses. The remaining gap likely reflects channels the model omits: non-bank financial intermediation, confidence effects on investment beyond the credit margin,

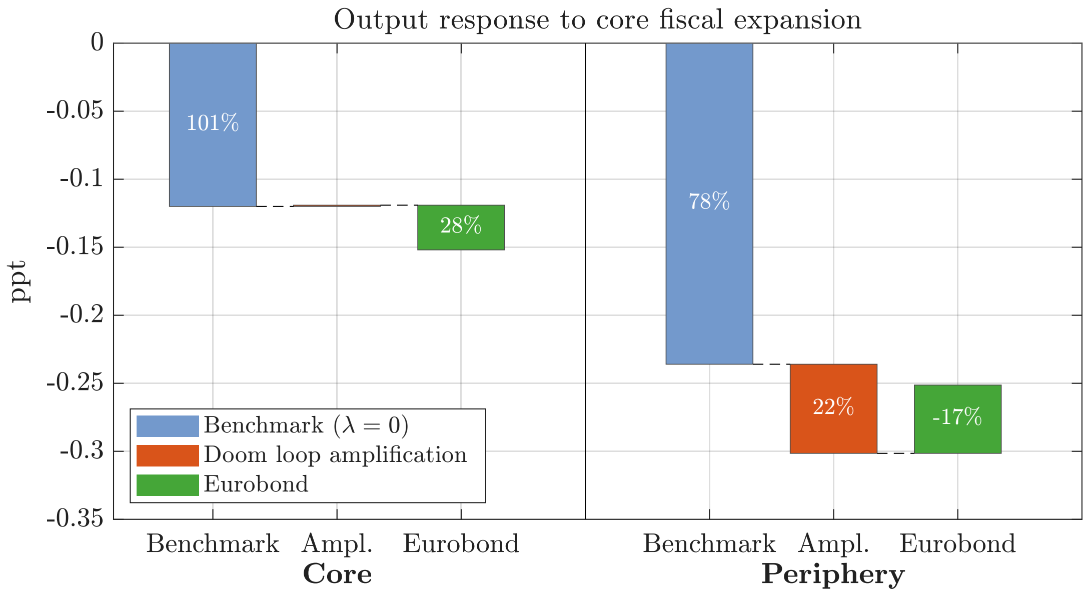

Figure 1.15. Doom loop amplification and Eurobond counterfactual

Notes: Decomposition of the model-implied output response (ppt) to a 25 bps core yield shock. "Benchmark" is the output loss when the doom loop is shut off ($\mathcal{A} = 1$). "Ampl." is the additional loss from sovereign-bank feedback. "Eurobond" is the counterfactual under unified bond issuance, in which the supply shock is diluted, home bias is eliminated, and the doom loop is attenuated. Percentage labels denote each component's share of the total national output loss.

second-round demand propagation, and direct pass-through of sovereign yields to corporate borrowing costs outside the banking sector. The calibration also abstracts from ECB intervention (OMO, TPI), portfolio rebalancing, and within-country bank heterogeneity. To the extent that markets price in the ECB backstop, the observed spread reaction already incorporates this truncation, and the model captures the net effect.

1.6 Conclusion

This paper documents a previously understudied externality of fiscal shocks. The empirical evidence demonstrates that the news of future budget expansions impacts asset markets significantly even before any fiscal outlays occur. This reaction is particularly detrimental in a monetary union where one sovereign serves as the risk-free benchmark, as illustrated here by the role of Germany within the Euro area.

The mechanism operates through a sequence of financial-sector frictions. By repricing the German government bond yield, the de facto risk-free rate in the Euro area is pushed upward. This causes a systemic revaluation of all asset classes, directly impairing the perceived riskiness of bank balance sheets. Intermediaries respond by cutting the supply of credit to the economy on both the extensive and intensive margins, which creates an immediate drag on growth.

For countries where the $r - g$ differential is relatively narrow, concerns about debt sustainability exacerbate the impact on domestic sovereign yields. This sets up a doom loop where the initial benchmark shock is amplified, ultimately hitting the yields of peripheral Euro area countries far more severely than those of Germany. The real economic impact of this channel is considerable: even before any fiscal funds are spent, the transmission through financial markets reduces future multipliers by 11 to 29 basis points depending on the country.

Ultimately, these findings suggest that in a financially integrated currency union, the fiscal space of the core is constrained by the financial stability of the periphery. Without a common safe asset to decouple the risk-free benchmark from national fiscal trajectories, core expansions may inadvertently act as a contractionary force for the union at large, as the financial chilling effect eclipses the traditional expansionary effects of government spending.

Chapter 2

Stop Believing in Haircuts

Co-authors: Benedikt Ballensiefen and Benoît Nguyen

Abstract. This paper studies how frictions impact haircut setting in repo markets. Using transaction-level supervisory data and exploiting the 2021 introduction of the binding leverage ratio requirement, we show that three frictions jointly shape haircuts in the bilateral dealer-to-client segment: dealer balance sheet constraints, collateral convenience yields, and dealer market power. In collateral-driven repos, binding leverage constraints raise haircuts, higher convenience yields push them down, and greater client outside options compress them further. In funding-driven trades the regulation does not bind, while convenience yields and market power continue to shape haircuts with opposite signs. We rationalise these patterns in a partial-equilibrium bargaining model and show that when haircuts cannot adjust freely, risk is displaced onto repo rates and volumes. Consistent with the model, constrained dealers exhibit higher rate volatility and weaker pass-through of ECB policy rate changes to their clients. Post-crisis bank regulation thus carries potential unintended consequences for repo market functioning and the transmission of monetary policy.

Keywords: Repo markets, haircuts, frictions, balance sheet constraints, convenience yield, financial stability, monetary policy transmission.

JEL Classification: E43, E44, E52, G12, G21, G28.

For helpful comments and suggestions, we thank Stefano Corradin, Andy Hill, Felix Hermes, Seung Joo Lee, Anthony Limburg, Edouard Mattille, Andrea Poinelli, Patrick Schaffner, Martin Scheicher, Andreas Schrimpf, Dimitrios Tsomocos, Jonathan Wallen, and ECB DGM colleagues as well as seminar participants at the ECB.

2.1 Introduction

The money market plays a central role in the efficient allocation of short-term funding, securities financing, and collateral. A deep and well-functioning money market is also a prerequisite for the effective transmission of monetary policy. Given its importance to the broader financial and economic system, it is essential to understand the frictions that arise in money markets and the mechanisms through which they impact market efficiency, financial stability, and the transmission of monetary policy.

The repurchase agreement (repo) is the main money market instrument. It is a short-term loan secured by different types of government securities. A unique feature of repos is that they can serve a dual role in obtaining either cash or collateral. Repo haircuts are a key risk management tool in repurchase agreements as they protect the lender in case of a counterparty default. Recently, the prevalence of zero and negative haircuts has moved to the center of the policy debate over the financial stability risks posed by their effects on the leverage of non-bank financial institutions, with one recommendation under discussion being the imposition of minimum haircuts (Financial Stability Board, 2025). The market practice of zero and negative haircuts is confirmed *inter alia* by Kahn and McCormick (2025) and Cenicola et al. (2025). This paper builds on this debate and studies (i) how different money market frictions affect haircut-setting practices in general and (ii) the implications of suboptimal haircut choices for financial stability and monetary policy.

As transactions in the inter-dealer segment tend to be centrally cleared and hence their haircuts are subject to portfolio margining that we do not observe, we focus on the bilateral segment of the repo market. Because dealers have access to a centrally cleared interdealer market—offering capital netting benefits, anonymous centralized trading, homogeneous counterparty risk, and efficient price discovery—we assume that most trades in the dealer-to-client segment are client-initiated. Accordingly, we classify transactions in which a dealer provides cash to a client as *funding*-driven and those in which a dealer provides collateral as *collateral-driven*.[1] The haircut is defined as the percentage difference between the market value of the collateral and the cash lent. The basic intuition behind collateralised lending is that haircuts should protect the cash- or collateral-lender from the counterparty's default, and that rates are a function of supply and demand for cash and scarce collateral. If a dealer lends out cash (in a *funding-driven* transaction), we would expect haircuts to be positive to protect the dealer against a borrower default; if the borrower defaults, the dealer must sell the collateral at a potentially lower price than at inception. We refer to this risk as *liquidation risk*. By contrast, if a dealer lends out collateral (in a *collateral-driven* transaction), we would expect haircuts to be negative to protect dealers against the risk of the collateral not being returned. If the collateral is not returned, the dealer must re-source it in the market at a potentially higher price plus foregone convenience yields. We refer to this risk as *replacement risk*.[2]

For the empirical analysis, we leverage a unique transaction-level dataset of euro area repo transactions in the bilateral dealer-to-client segment, drawn from confidential supervisory reporting to the ECB. Specifically, we use transaction-level information from the ECB's Money Market Statistical Reporting (MMSR) dataset. Empirically, there are several important patterns that we observe in these data. First, we document systematic deviations of observed haircuts from our prior expectations. Haircuts are often set at 0%, in which case they are not serving a risk-mitigation purpose, as noted by Financial Stability Board (2025). In addition, haircuts are frequently negative in funding-driven trades and positive in collateral-driven trades. Both ob-

[1] As there is no universally accepted definition for classifying repo transactions as funding-driven or collateral-driven, we also consider alternative classification schemes to ensure the robustness of our results. In particular, we exploit the general collateral versus special collateral designation.

[2] Repo contracts embed additional risk management tools beyond the initial haircut, most notably daily margining. Margining refers to the daily mark-to-market adjustment of the collateral value, which requires the posting of variation margin when price changes alter the exposure between counterparties. While the haircut provides an ex ante buffer against potential losses in the event of default, daily margining manages ex post exposure by ensuring that the collateral value remains aligned with the outstanding cash position over the life of the transaction. Thus, although both mechanisms mitigate counterparty and market risk, they address distinct dimensions of risk: haircuts protect against liquidation and replacement risk, whereas margining limits the accumulation of mark-to-market losses.

servations suggest that haircut setting reflects strategic considerations beyond managing conventional liquidity and replacement risk. Second, standard theory predicts that haircuts should depend on collateral characteristics such that scarcer or safer assets are expected to command lower haircuts. However, our second stylized empirical finding shows that this prediction does not strictly hold in the data. There are periods in which haircuts on collateral from less risky countries are higher than those on collateral from riskier countries. In short, we observe that many repo contracts are priced as if there is no liquidation or replacement risk, and are even priced contrary to ex ante expectations.

We document these empirical patterns both at the transaction level and at the dealer–client level. This dual approach mitigates concerns that haircuts are determined at the client portfolio level rather than at the level of individual transactions. Consistent with this interpretation, we show that the distribution of transaction-level haircuts closely mirrors the distribution implied by dealer–client–day aggregated cash and collateral positions. The tight correspondence between these measures suggests that our haircut estimates are not materially distorted by portfolio-level margining practices and is consistent with institutional features documented for the U.S. Treasury repo market. For example, Hempel et al. (2023) highlight that in bilateral repo transactions transaction-level haircuts determine margins collected for that transaction, whereas for centrally cleared repos, margins are calculated on the portfolio level.

Having established the stylized facts on repo haircuts, we turn to the empirical analysis. In the first part, we examine how three key frictions known to shape repo markets—dealer balance-sheet constraints (e.g., Ranaldo et al., 2021), collateral convenience yield and scarcity (e.g., Buraschi and Menini, 2002; Corradin and Maddaloni, 2020; Arrata et al., 2020; Ballensiefen et al., 2023a), and market power (e.g., Eisenschmidt et al., 2024)— jointly affect haircut setting.

We first show that dealer balance-sheet constraints put upward pressure on haircuts in collateral-driven repos.[3] From a dealer's perspective, the regulatory treatment of repos and re-

[3]Dealer leverage ratios are sourced from the ECB's harmonised financial reporting FINREP dataset.

verse repos is asymmetric: in collateral-driven repos, the cash leg increases total assets and enters the non-risk-weighted leverage ratio, making these trades balance-sheet intensive, whereas funding-driven repos do not tighten the leverage constraint in the same way. This is because, under standard accounting treatments, securities received as collateral remain on the balance sheet of the collateral provider, whereas the cash leg of the transaction is recognized as an asset on the bank's balance sheet. Moreover, haircuts also carry capital cost. Lower haircuts in collateral-driven repos increase the amount of cash received by the dealer and thus the balance-sheet usage, giving constrained dealers an incentive to raise haircuts when leverage capacity becomes scarce. Consistent with this mechanism, we find, in the *cross-section*, that dealers with lower leverage ratios relatively increased haircuts compared to less-constrained dealers after the introduction of the leverage ratio requirements. Interestingly, dealers in German collateral, for example, are - on average - more balance-sheet constrained than dealers in Italian securities, which provides intuition for why haircuts for collateral countries with higher convenience yields and replacement risk can still be higher. In the *time-series*, we show that this effect has only emerged since the leverage ratio requirements for dealer banks in the euro area have become binding. Importantly, the leverage ratio in the euro area becomes binding only on reporting days (quarter-end days), however, as of June 2021, euro area banks are also required to report an "averaging" version of their leverage ratio. In line with this, we show that the effect of the leverage ratio on haircuts is present not only on quarter-end days but throughout the entire quarter. As the leverage ratio regulation does not penalize reverse repos, there is no effect on funding-driven repos. Figure 2.1 illustrates these results: after the introduction of the leverage ratio requirement, the most leverage-constrained dealers increased haircuts relative to less constrained dealers, but only in collateral-driven trades. The estimates rely on within dealer–client variation and include rich fixed effects, thereby isolating within-relationship changes in haircuts around the regulatory shock.

Second, we show that collateral heterogeneity matters not only for repo rates but also for

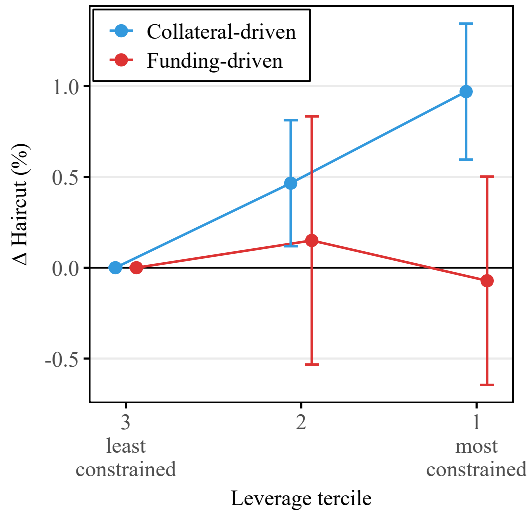

Figure 2.1. Haircut adjustments to leverage constraints

Notes: The plot shows the estimated differences in haircuts relative to the tercile with the least constrained banks (least constrained = 3, most constrained = 1) after the 3% leverage ratio requirement became binding for all EU banks on June 28, 2021. The coefficients and standard errors shown in this graph correspond to regressions results presented in columns (3) and (6) in Table 2.A.1. Coefficients are obtained from regressions that include collateral-country, dealer×client, month, and leverage-tercile fixed effects. 'Collateral-driven' and 'funding-driven' defined from the perspective of the dealer's clients.

haircuts. Assets with higher convenience yields—reflecting, for example, scarcity and specialness (Krishnamurthy and Vissing-Jorgensen, 2012b) and their value to market participants (Ballensiefen and Ranaldo, 2023)—entail greater replacement risk if they are not returned. Consistent with this mechanism, we find that convenience yields influence haircuts in the direction of the dealer's market power: they reduce haircuts in collateral-driven trades to mitigate replacement risk, but increase them in funding-driven markets to insure against liquidity risk.

Third, we show that market segmentation shapes haircut setting. The bilateral repo market is segmented in two dimensions: dealer–client relationships are concentrated, and not all dealers intermediate all types of collateral. This structure gives rise to market power within trading relationships. In funding-driven repos, clients with relationships to a larger set of dealers are able to negotiate lower haircuts. In contrast, collateral-driven repos are typically concentrated within a narrow set of dealers specialized in a given collateral class, limiting clients' outside options and allowing dealers to exert greater market power. Overall, this segmentation implies that haircut levels are largely relationship-specific and shaped by bargaining positions and market conventions, rather than adjusting fully to time-varying risk.

In the second part of the paper, we develop a theoretical framework that embeds these frictions into a partial-equilibrium model of the repo market. We model competition among dealers as a Bertrand game with both soft and strict capacity constraints on collateral and cash. This structure is motivated by the highly oligopolistic nature of the repo market (Eisenschmidt et al., 2024) and allows us to capture limited balance-sheet capacity and segmented intermediation in a tractable way. The model explicitly distinguishes between collateral-driven and funding-driven transactions and shows how balance-sheet costs, collateral convenience yield, and market power jointly determine equilibrium haircuts. For collateral-driven repos, comparative statics predict that haircuts increase in dealer balance-sheet costs but decrease in the collateral's convenience yield. These predictions align closely with our empirical findings. Importantly, the model implies that haircuts are more sensitive to changes in dealer balance-sheet costs than to changes in collateral convenience yields. This asymmetry provides a theoretical explanation for the empirical puzzle that collateral with higher convenience yields can nonetheless be associated with higher haircuts when dealers are balance-sheet constrained.

Finally, we use the model to make two predictions about the implications for market stability and monetary policy when haircuts are not adjusted optimally. We show that repo rates and transaction volumes are more volatile in trades with constrained haircuts. One potential explanation is that, when haircuts are sticky or constrained, dealers adjust rates and volumes instead, shifting risk absorption from margins to prices and quantities. We also provide suggestive evidence that dealer balance-sheet constraints weaken monetary policy transmission. If, in collateral-driven markets, more constrained dealers choose suboptimal haircuts and instead use repo rates as a risk-mitigation tool, then we expect the monetary policy transmission from dealers to clients to be weaker for more balance sheet constrained dealers. Consistent with this mechanism, repo rates offered by more constrained dealers adjust more slowly to changes in policy rates. These results suggest that the regulatory measures implemented by financial stability authorities have affected monetary policy and may have introduced even more volatility to market outcomes.

Our analysis mainly contributes to two strands of the literature. First, we add to the literature on short-term money markets. The innovations we bring are twofold: First, we document how money market frictions impact haircuts. Second, we differentiate those effects for repos that either serve a funding- or collateral-purpose. Four papers are most closely related to our work. Julliard et al. (2022) study the determinants of repo haircuts using transaction-level data from six major UK banks and highlight the role of money market frictions. Hermes et al. (2025) document that even relatively risky euro-area collateral frequently carries zero or negative haircuts. Focusing on the U.S., Bejarano et al. (2025) show that major dealer banks sometimes quote negative haircuts on Treasury repos when they have excess balance-sheet capacity, indicating that haircuts do not solely reflect collateral risk. Bittner and Jank (2025) show that competition among lenders weakens risk-based haircut pricing, pointing to market power as a further determinant. We contribute to this literature in two ways. First, we identify dealer balance sheet constraints as an understudied friction in repo pricing and examine how they interact with collateral scarcity and dealer market power. Second, we distinguish systematically between collateral-driven and funding-driven trades, showing that these frictions operate differently across the two market segments. Our paper also relates to other work focusing on the collateral demand in repo markets (e.g., Mancini et al., 2016; Arrata et al., 2020; Corradin and Maddaloni, 2020, and Coen et al., 2024).[4]

[4]For example, studies on safe asset scarcity and collateral "specialness" show that high-quality government bonds in Europe command convenience yields that can impact repo markets. In line with this, Krishnamurthy and Vissing-Jorgensen (2012b) illustrate how safe assets carry a premium in investor

Second, we add to the literature on monetary policy, financial regulation, and the monetary policy transmission process (e.g., Duffie and Krishnamurthy, 2016b; Drechsler et al., 2017; Ballensiefen et al., 2023a; Eisenschmidt et al., 2024). It has been shown that post-crisis bank regulations can have unintended effects on short-term funding markets. In particular, Ranaldo et al. (2021) find that the Basel III leverage ratio constraint induces "window dressing" behavior by European banks. We add to this by shedding light on how central bank regulation impacts euro area repo markets throughout the entire quarter, leads to suboptimally chosen haircuts, and thereby adversely impacts the monetary policy transmission. This adds to the work by, for example, Ballensiefen et al. (2023a) and Eisenschmidt et al. (2024) on the monetary policy transmission process. We show that haircuts are sub-optimally chosen so that adjustments in central bank rates translate into heightened volatility in repo rates and volumes rather than uniform interest rate pass-through. Such concerns resonate with findings by Duffie and Krishnamurthy (2016b) and Drechsler et al. (2017).

Overall, our study contributes to the existing literature by introducing several novel dimensions. First, we specifically incorporate dealers' balance-sheet constraints by focusing on regulatory ratios, allowing us to directly assess how post-crisis regulation shapes haircut setting. Second, exploiting detailed euro area data, we distinguish between collateral- and funding-driven repo transactions and show that these two segments are affected differently by money market frictions, with important implications beyond haircut formation on repo market outcomes and monetary policy transmission. And third, we add a theoretical model that rationalizes our empirical findings.

The remainder of the paper proceeds as follows. Section 2.2 describes our setting and the data. Section 2.3 documents empirically how frictions in the money market impact repo haircuts. Section 2.4 introduces our model, calibrates it to the data and makes predictions about

preferences that translates into repo pricing (Ballensiefen and Ranaldo, 2023). Focusing more on crisis periods, Gorton and Metrick (2012), for example, documented how repo haircuts spiked during the 2008 financial crisis.

market outcomes which are then tested empirically in Section 2.5. Section 2.6 concludes.

2.2 Data and Institutional Setting

This section provides an overview of the euro area repo market and describes the data used in our analysis. The focus of this paper is on the bilateral dealer-to-client segment of the euro area money market, which provides a uniquely suitable setting to study how repo haircuts are set in collateral- and funding-driven environments. This segment is characterised by bilateral contracting, relationship-specific bargaining, and limited transparency, making market frictions particularly salient. Moreover, the dealer-to-client segment represents the primary channel through which monetary policy is transmitted from regulated dealer banks to non-bank financial institutions, allowing us to study directly how regulatory constraints affect both contract design and policy pass-through.

2.2.1 The Euro Area Repo Market

A repurchase agreement, or repo, is a short-term secured financing transaction in which one counterparty sells a security (most commonly a sovereign bond) to another, with an obligation to repurchase it at a predetermined price after an agreed period. A repo contract is characterised by four key terms, each negotiated at the transaction level: (i) the volume of cash exchanged, (ii) the repo rate, (iii) the specific security posted as collateral, and (iv) the haircut applied to the collateral.[5] The collateralised nature of repos means they serve a dual function: they provide a vehicle both for obtaining short-term funding and for sourcing specific securities.

Market structure. The euro area repo market comprises two broad segments: an *inter-dealer* segment, in which dealers trade with one another, and a *dealer-to-client* segment, in which deal-

[5]The cash lender in a repo transaction becomes the legal owner of the collateral for the term of the repo, and may use it for example, to settle delivery obligations or to re-pledge in a subsequent transaction (Duffie, 1996).

ers intermediate trades with non-dealer institutions such as investment funds, hedge funds, money market funds, insurance companies, and pension funds.

The inter-dealer segment is predominantly centrally cleared. Transactions are executed on electronic platforms (such as BrokerTec, Eurex Repo, and MTS Repo) and novated to a central counterparty (CCP), which becomes the legal counterparty to both sides of each trade. Central clearing allows dealers to net their positions against a single counterparty across many trades, substantially reducing gross balance-sheet exposures. Moreover, the CCP applies standardised risk-management tools, including initial and variation margin, and facilitates efficient price formation through anonymous order books.

The dealer-to-client segment, by contrast, is an over-the-counter (OTC) bilateral market. Because netting is only possible within each individual dealer–client relationship, gross exposures remain larger and balance-sheet constraints are more binding than in centrally cleared markets. There is no centralised price formation mechanism; instead, contract terms are negotiated bilaterally and are relationship-specific. As Eisenschmidt et al. (2024) show using the same MMSR data, the median non-bank customer trades with only a single dealer, and this concentrated intermediation structure endows dealers with significant market power.[6] This paper focuses on the non-cleared, bilateral dealer-to-client segment, precisely because it is in this setting that frictions in haircut-setting are most salient and empirically identifiable.

Figure 2.2 shows the composition of dealer-to-client activity by client sector, separately for collateral-driven and funding-driven repos. In both segments, non-money-market-fund (non-MMF) investment funds are the single largest client category, accounting for 56.9% of collateral-driven volume and 41.7% of funding-driven volume. Other financial intermediaries[7] constitute the second-largest group. In funding-driven repos, deposit-taking corporations play a more

[6] See also Schaffner et al. (2019) for evidence on dealer trading patterns and home bias in the inter-dealer segment.

[7] e.g. securities and derivatives dealers, financial vehicle corporations, venture capital and development capital firms

significant role (24.0% of volume), reflecting these institutions' demand for short-term secured funding. The remaining volume is distributed across money market funds, pension funds, insurance corporations, and other sectors.

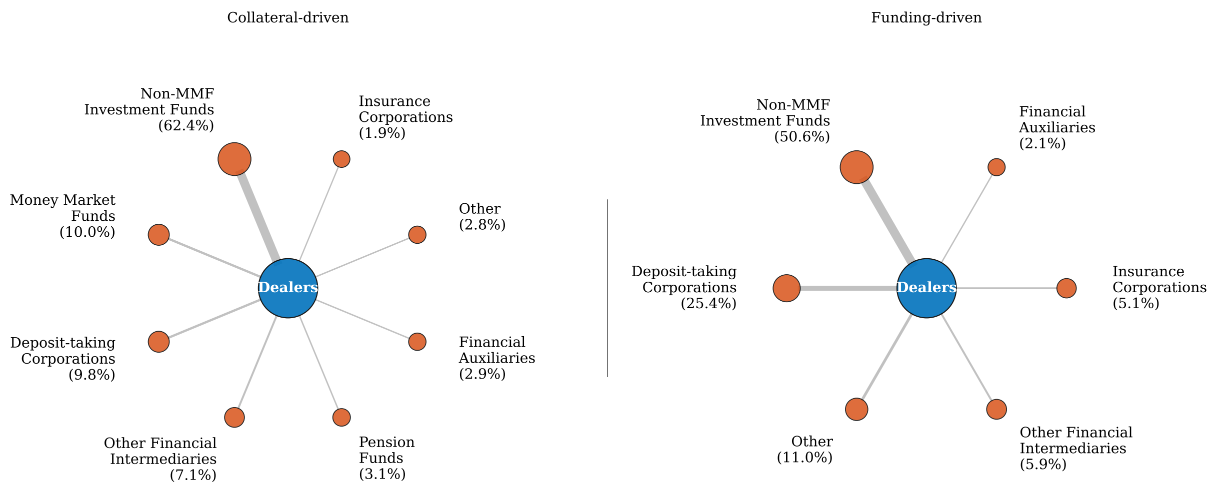

Figure 2.2. Client composition of dealer-to-client repos

Notes: This figure displays the share of dealer-to-client repo volume by client sector, separately for collateral-driven and funding-driven repos. The sample covers bilateral repos collateralised by government bonds (S13), January 2019 to June 2024.

Initial margin versus variation margin. It is important to distinguish between the two forms of margining in repo markets. The *initial margin*, commonly referred to as the *haircut*, is set at the inception of the trade and determines the extent to which the cash exchanged deviates from the market value of the collateral. It remains fixed over the life of the transaction and serves as an ex-ante buffer against counterparty default risk—specifically, the risk that a defaulted counterparty's collateral must be liquidated under potentially stressed market conditions (*liquidation risk*) or that the collateral must be repurchased at a cost (*replacement risk*). The *variation margin*, by contrast, addresses ex-post price fluctuations: it is exchanged on an ongoing basis as the market value of the collateral moves, ensuring that mark-to-market losses are continuously covered. Because haircuts are not adjusted during the life of a transaction, they reflect how dealers price risk, balance-sheet constraints, and market frictions at the point of contract initiation. This paper focuses on initial margins, as these are the terms over which dealers and clients negotiate ex ante. Our data do not contain information on variation margin calls.[8]

Formally, the MMSR reporting guidelines define the haircut as:

$$\text{Haircut} = 100 - \left(\frac{\text{Cash lent or borrowed at inception}}{\text{Collateral market value at inception including accrued interest}} \times 100 \right). \tag{2.1}$$

MMSR reporting explicitly permits negative haircuts, which arise when the cash amount at inception exceeds the market value of the collateral.[9] A positive haircut means the lender of cash receives collateral worth more than the cash disbursed (overcollateralisation), while a negative haircut means the cash amount exceeds the collateral value (undercollateralisation from the lender's perspective).

Collateral-driven versus funding-driven repos. A key distinction for this paper is between *funding-driven* and *collateral-driven* repo transactions. In a funding-driven repo, the initiating party seeks cash; in a collateral-driven repo, the initiating party seeks a specific security or, more broadly, eligible collateral. This distinction matters fundamentally for the economics of haircut-setting. In a funding-driven repo, the haircut protects the cash lender against the liquidation risk of the collateral in the event of borrower default. The cash lender therefore has an incentive to demand overcollateralisation, implying positive haircuts. In a collateral-driven repo, the haircut protects the collateral provider (the dealer) against the replacement risk of the security: should the cash lender default and fail to return the collateral, the dealer must repurchase it in

[8]For cleared transactions, it is common practice to report a zero haircut because risk management is handled through the CCP's initial and variation margin framework. Our focus on the bilateral segment avoids this confound.

[9]See Reporting Instructions for the Electronic Transmission of Money Market Statistical Reporting (MMSR), version 3.6.1 (30 August 2023), available at https://www.ecb.europa.eu/stats/money/mmss/shared/files/MMSR-Reporting_instructions.pdf.

the open market, potentially at a premium. The dealer therefore has an incentive to deliver less cash than the collateral is worth, implying negative haircuts.

We do not directly observe who initiated a given trade. However, given the netting benefits of central clearing, the superior pricing information available on centralised platforms, and the broad pool of both cash and collateral accessible through CCPs, it is considerably more likely that dealers would use the inter-dealer market to fulfil their own cash or collateral needs. In the bilateral dealer-to-client segment, trades are therefore predominantly initiated by clients.[10] We classify repo transactions accordingly: those in which the dealer lends cash and receives collateral are classified as *funding-driven* (the client sought cash), while those in which the dealer lends securities and receives cash are classified as *collateral-driven* (the client sought collateral).

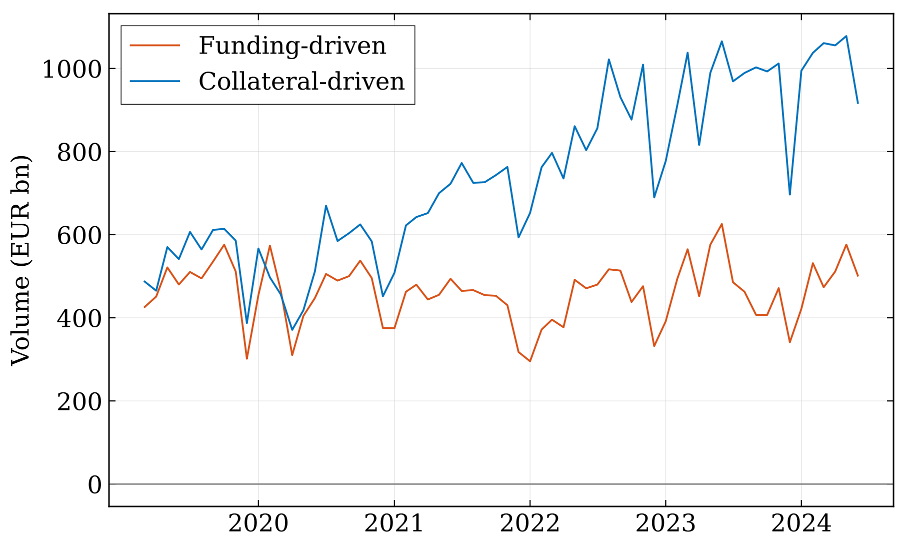

Figure 2.3. Trading volumes in funding- and collateral-driven repos

Notes: This figure shows monthly aggregated volume of one-day repos in the bilateral dealer-to-client segment, collateralised by German, French, Italian, or Spanish general government bonds (S13 sector), January 2019 to June 2024.

[10]There may be instances in which dealers access the dealer–client segment to source specific collateral. However, such cases are relatively rare, as clients—particularly hedge funds—typically have stronger demand for particular securities.

Figure 2.3 shows the evolution of monthly trading volumes in the two segments. At the beginning of 2019, collateral-driven and funding-driven volumes are roughly comparable. By 2024, however, collateral-driven volumes have nearly doubled while funding-driven volumes have remained largely flat. This divergence underscores the growing importance of the collateral motive in the euro area repo market and reinforces the relevance of understanding how haircuts are set in this segment.

An alternative classification used in the literature distinguishes between general collateral (GC) and special collateral (SC) repos. While SC repos are more likely to be collateral-driven and GC repos more likely to be funding-driven, this mapping is imperfect: a client needing any eligible collateral may enter a GC repo for collateral reasons, while a client holding scarce collateral may enter an SC repo primarily to obtain favourable funding terms. Moreover, the SC/GC flag is reported for only about 20% of transactions in the MMSR data and is concentrated among a small number of dealers, making it unsuitable as the primary classification criterion.

2.2.2 Data

The MMSR dataset. Our primary data source is the Money Market Statistical Reporting (MMSR) dataset, a confidential supervisory dataset collected by the European Central Bank. The MMSR requires 53 euro area banks to report daily transaction-level information on all their euro-denominated money market activity, covering both secured (repo) and unsecured segments.[11] The secured segment covers all repo transactions with a maturity of up to one year.[12] For each transaction, we observe the repo rate, the volume, the maturity, the haircut, the collateral security at the ISIN level, and pseudonymised identifiers for both the reporting dealer and the counterparty, including the counterparty's sector classification. Our sample covers the

[11] See the list of reporting agents at https://www.ecb.europa.eu/stats/financial_markets_and_interest_rates/money_market/html/index.en.html. The MMSR data is described in detail by Eisenschmidt et al. (2024).

[12] Defined as transactions with a maturity date of not more than 397 days after the settlement date.

period from January 2019 to June 2024.

Sample restrictions. We impose several restrictions to obtain a sample that is well suited to studying haircut determination in the bilateral dealer-to-client segment. First, we restrict the sample to bilateral (non-CCP-cleared) repo transactions in which the reporting agent is the dealer and the counterparty is a non-dealer client. This excludes inter-dealer trades and CCP-cleared transactions, where haircuts follow standardised schedules and margining is handled by the clearing house rather than negotiated bilaterally. Second, we focus on transactions collateralised by general government bonds (ESA sector S13). These collateral class accounts for the vast majority of euro area sovereign repo activity and provide sufficient cross-sectional variation in sovereign risk characteristics. Third, we exclude transactions between dealers and the central bank in our analysis as trading incentives and terms likely follow different rules than laid out in the bilateral master agreements. Finally, we winsorise repo rates at the 1st and 99th percentiles to mitigate the influence of outliers and potential reporting errors.[13] Our sample period spans from January 2019 to June 2024.[14]

Table 2.1 provides an overview of the resulting sample. The dataset comprises approximately 1.8 million collateral-driven and 1.3 million funding-driven transactions, involving 17 dealers and over 1,500 unique clients. Italian collateral accounts for the largest share of funding-driven volume (29.1%), while in the collateral-driven segment, volumes are slightly more evenly distributed across Germany (22.9%), France (19.3%), Italy (25.8%), and Spain (14.1%). The vast majority of transactions (over 90%) have an overnight maturity.

A striking feature of the data is the distribution of haircut signs. In collateral-driven repos, 71.1% of transactions have a zero haircut and only 22.1% have a strictly negative haircut—even though economic logic would suggest negative haircuts should be the norm when dealers are

[13] In the MMSR dataset, different banks have different ways of reporting haircuts or margins. We are converting all the different definitions to a uniformly defined haircut as in equation (2.1)

[14] While MMSR data is available from Oct 2016, we exclude the pre-2019 period due to significant reporting inconsistencies regarding haircuts that could bias our estimates.

exposed to replacement risk. In funding-driven repos, 61.2% of transactions have a zero haircut and only 28.2% have a strictly positive haircut. These patterns motivate the central question of this paper: why are so many haircuts set at zero or in the "wrong" direction, and what are the determinants of the remaining cross-sectional and time-series variation?

Table 2.1. Sample coverage

Notes: Bilateral dealer-to-client repos collateralised by euro area government bonds (S13), January 2019–June 2024. Under the ESA 2010 classification, S124 comprises investment funds excluding money market funds while S125 includes other financial intermediaries such as security dealers and financial vehicle corporations but excludes insurance corporations and pension funds. Rates and haircuts winsorised at 1st/99th percentiles.

	Collateral-driven	Funding-driven
Transactions		
Total number of transactions	1,786,980	1,332,839
Unique dealers	17	16
Unique clients	1,512	856
of which S124	867	381
of which S125	69	69
Unique dealer–client pairs	2,763	1,792
Avg. transactions per pair per month	52.6	62.0
Collateral country (% of volume)		
Germany (DE)	22.9	24.4
France (FR)	19.3	17.9
Italy (IT)	25.8	29.1
Spain (ES)	14.1	11.3
Haircut sign (% of transactions)		
Negative (<0)	22.1	10.6
Zero (=0)	71.7	61.2
Positive (>0)	6.2	28.2

Haircut distributions. Figure 2.4 shows the estimated kernel density of haircuts for funding-driven and collateral-driven repos. Both distributions are sharply peaked at zero, with secondary mass points visible around ±2% and ±5%, likely reflecting standardised haircut schedules used by some dealer–client pairs. The collateral-driven distribution is more tightly centred around zero than the funding-driven distribution, which exhibits greater dispersion and a heavier right tail. The prevalence of zero haircuts across both segments—and the presence of positive haircuts in collateral-driven repos and negative haircuts in funding-driven repos—are difficult to rationalise from a pure risk-management perspective. Understanding the economic forces that generate these patterns is a central objective of our analysis.

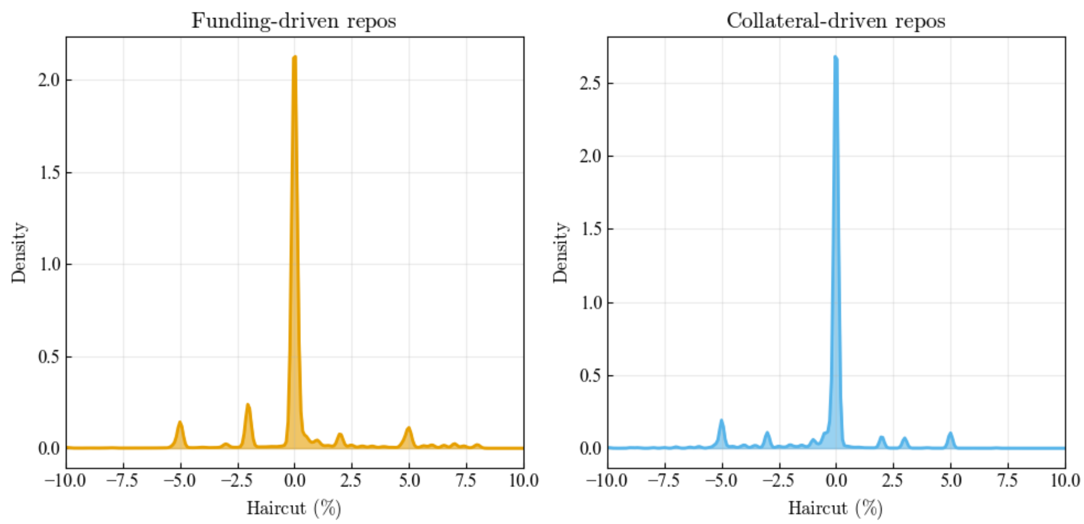

Figure 2.4. Distribution of repo haircuts

Notes: This figure shows the estimated kernel density of haircuts in funding-driven and collateral-driven bilateral dealer-to-client repos, January 2019 to June 2024. The estimation uses a Gaussian kernel for confidentiality reasons.

Haircut dynamics. Figure 2.5 decomposes monthly repo volumes by haircut sign, separately for funding-driven and collateral-driven repos. In funding-driven repos, the composition is relatively stable over time: the majority of volume has a positive haircut, with a modest share carrying zero or negative haircuts. In collateral-driven repos, a more striking pattern emerges: the rapid growth in collateral-driven volumes since 2021 is almost entirely accounted for by transactions with a zero or positive haircut, while the volume of transactions with negative haircuts—the theoretically expected sign—has remained largely flat. This raises the question of whether the growing collateral-driven segment is characterised by systematically weaker risk management, or whether other forces, such as dealer balance-sheet constraints or market power, are shaping haircut outcomes.

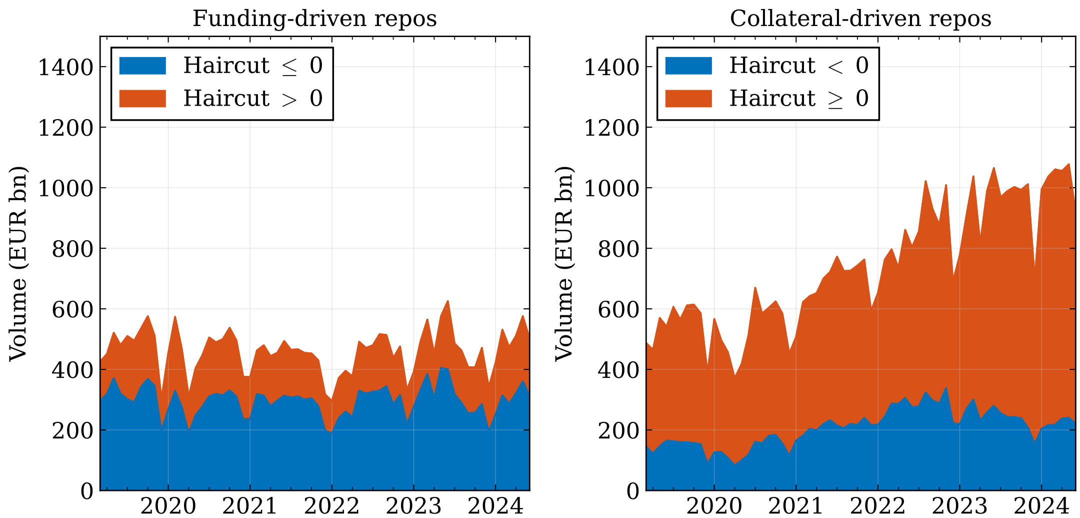

Figure 2.5. Monthly repo volume by haircut sign

Notes: This figure shows monthly aggregated repo volume in the bilateral dealer-to-client segment, collateralised by German, French, Italian, or Spanish general government bonds. Each panel is segmented by haircut sign: in the funding-driven panel, volumes are split into non-positive (≤ 0) and positive (> 0) haircuts; in the collateral-driven panel, into negative (< 0) and non-negative (≥ 0) haircuts. Time series show the trend and seasonal components for confidentiality.

Figure 2.6 adds a cross-sectional dimension by plotting the monthly volume-weighted average haircut by collateral country. In funding-driven repos, haircuts across collateral countries have converged since 2021, with slightly higher haircuts for Italian and Spanish collateral reflecting their greater sovereign credit risk. In collateral-driven repos, dispersion has increased since 2021, with haircuts for German and French collateral (which tend to be more "special" and scarce) being less negative than haircuts for Italian and Spanish collateral. This cross-country heterogeneity motivates our analysis of how collateral characteristics interact with dealer constraints to determine haircut outcomes.

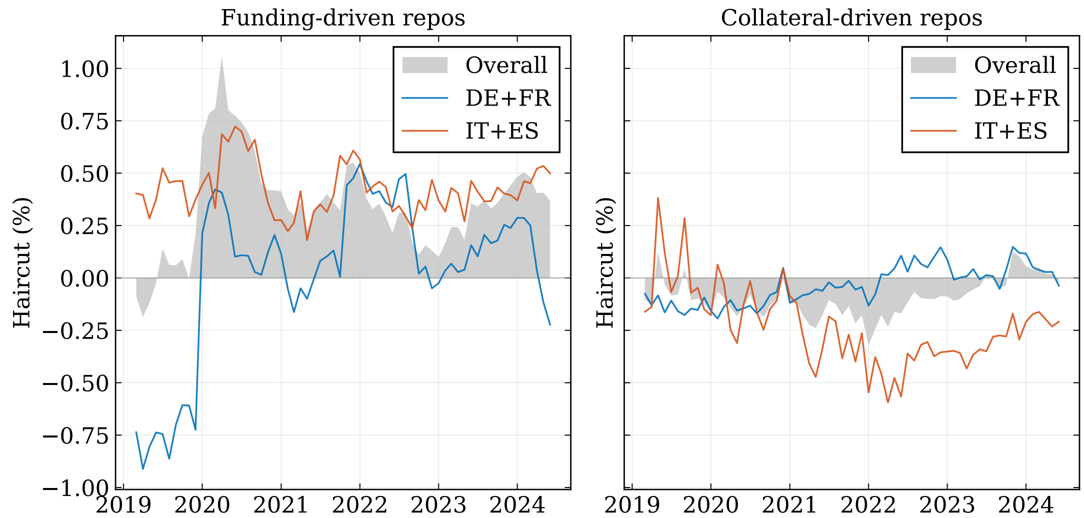

Figure 2.6. Volume-weighted average haircuts by collateral country

Notes: This figure shows the monthly volume-weighted average haircut in bilateral dealer-to-client repos, separately for funding-driven and collateral-driven transactions, January 2019 to June 2024. The shaded area represents the overall average across all countries.

Dealer and client characteristics. Table 2.2 reports dealer-level characteristics by leverage ratio tercile. The average dealer leverage ratio is 4.8%, with meaningful dispersion: dealers in the most constrained tercile have an average ratio of 3.87%, while those in the least constrained tercile average 5.82%. Leverage ratio data is sourced from the ECB's financial reporting (FINREP) dataset and are available at quarterly frequency. As Schaffner et al. (2019) document, there is a significant home bias in dealer trading, so that the average leverage ratio of dealers active in a given collateral segment varies systematically across countries: dealers trading in German and French collateral tend to be more balance-sheet constrained than those active in Italian and Spanish collateral.

Table 2.3 reports client-level characteristics. The mean client trades with only one or two dealers (measured over a rolling 12-month window), consistent with the sparse connectivity documented by Eisenschmidt et al. (2024). Average trade sizes are approximately EUR 30 million in both segments, and the vast majority of trades have overnight maturities. The mean haircut is -0.053% in collateral-driven repos and 1.169% in funding-driven repos, but with substantial dispersion within each segment.

Table 2.2. Dealer summary statistics

Notes: Dealer-level statistics by leverage ratio tercile (within quarter). Tercile 1 = most constrained (lowest leverage ratio). Rates and haircuts winsorised at 1st/99th percentiles. Volume in EUR millions.

	Tercile 1 (most constr.)	Tercile 2	Tercile 3 (least constr.)	All
Mean leverage ratio (%)	3.87	4.70	5.82	4.80
Panel A: Collateral-driven				
Mean active client relationships	194.7	103.3	86.7	172.6
Share CD volume (%)	66.3	63.9	54.7	60.6
Share DE+FR collateral (%)	42.4	51.1	16.0	42.2
Haircut (%)	-0.04	-0.15	-0.06	-0.08
Rate, DFR-adjusted (%)	-0.27	-0.25	-0.12	-0.24
Volume per transaction (EUR mn)	28.81	27.12	17.41	26.49
Haircut, pre-June 2021 (%)	-0.21	-0.16	0.03	-0.13
Haircut, post-June 2021 (%)	0.01	-0.15	-0.15	-0.05
Panel B: Funding-driven				
Mean active client relationships	135.1	67.9	39.9	112.0
Share CD volume (%)	66.3	63.9	54.7	60.6
Share DE+FR collateral (%)	39.1	51.0	30.0	42.2
Haircut (%)	0.75	0.44	-0.18	0.55
Rate, DFR-adjusted (%)	-0.23	-0.09	-0.21	-0.19
Volume per transaction (EUR mn)	21.93	27.62	12.85	22.15
Haircut, pre-June 2021 (%)	1.48	0.53	-0.14	0.91
Haircut, post-June 2021 (%)	0.40	0.34	-0.21	0.31

2.3 Empirical analysis

In our analysis, we consider three main money market frictions and their impact on repo haircuts: *(i)* dealer balance sheet constraints, *(ii)* the convenience yield embedded in the collateral value, and *(iii)* dealer market power.

Table 2.3. Client summary statistics

Notes: Client-level statistics. Q1–Q5 report the mean within each quintile of the distribution. All counterparty sectors included. Rates and haircuts winsorised at 1st/99th percentiles.

	Mean	SD	Q1	Q2	Q3	Q4	Q5
Panel A: Collateral-driven							
Number of dealers (rolling 12m)	1.5	1.1	1.1	3.2	3.2	3.4	3.6
Avg. dealer leverage ratio (%)	4.97	1.10	3.82	4.23	4.43	5.74	6.61
Rate, DFR-adjusted (%)	-0.154	0.196	-0.442	-0.198	-0.109	-0.040	0.017
Haircut (%)	-0.053	0.762	-0.823	-0.000	0.939	0.969	1.143
Volume (EUR mn)	31.31	252.72	0.00	0.05	0.29	2.20	153.72
Maturity (days)	2.2	2.1	1.0	1.7	6.3	6.9	7.3
Panel B: Funding-driven							
Number of dealers (rolling 12m)	1.6	1.3	1.1	3.7	3.9	4.2	4.3
Avg. dealer leverage ratio (%)	4.31	0.58	3.74	4.01	4.20	4.44	5.19
Rate, DFR-adjusted (%)	-0.087	0.314	-0.512	-0.107	-0.033	0.025	0.196
Haircut (%)	1.169	2.696	-0.155	0.021	0.754	5.521	6.635
Volume (EUR mn)	34.49	182.43	0.01	0.06	0.42	3.97	168.18
Maturity (days)	3.5	2.6	1.0	2.5	6.6	7.3	8.1

2.3.1 Dealer balance sheet constraints

The post-crisis leverage ratio (LR) regulation treats repo transactions asymmetrically. Lending collateral against cash in a repo extends the dealer's balance sheet, whereas lending cash against collateral in a reverse repo does not. Dealers are therefore more disincentivized to intermediate collateral-driven than funding-driven trades, as documented for repo rates by Ranaldo et al. (2021). The asymmetry extends naturally to haircuts: a negative haircut on a collateral-driven trade enters gross exposure at full value, so haircuts become a principal margin of adjustment whenever the constraint tightens.

The regulatory shock. On June 28, 2021, the 3% LR requirement under CRR II became a binding Pillar 1 obligation for all ECB-supervised banks. Article 429e permits netting of securities financing transactions only under narrow conditions, notably identical counterparties and settlement dates. Matched-book intermediation across different clients therefore expands gross exposure on both legs and directly consumes scarce leverage capacity. Although central bank reserves were temporarily excluded from the exposure measure during the pandemic, the effective date was predetermined years in advance, providing time-series exogeneity; cross-sectional dispersion in how close each dealer sits to the 3% threshold generates variation in treatment intensity. Because euro-area banks must report both a point-in-time and an averaged version of the ratio, the constraint remains relevant beyond quarter-end.

Figure 2.7 offers a first visual. Haircuts on funding-driven trades are essentially flat across reporting and non-reporting days, while collateral-driven haircuts shift clearly upward on reporting days, consistent with dealers repricing balance sheet usage when the constraint is most visible.

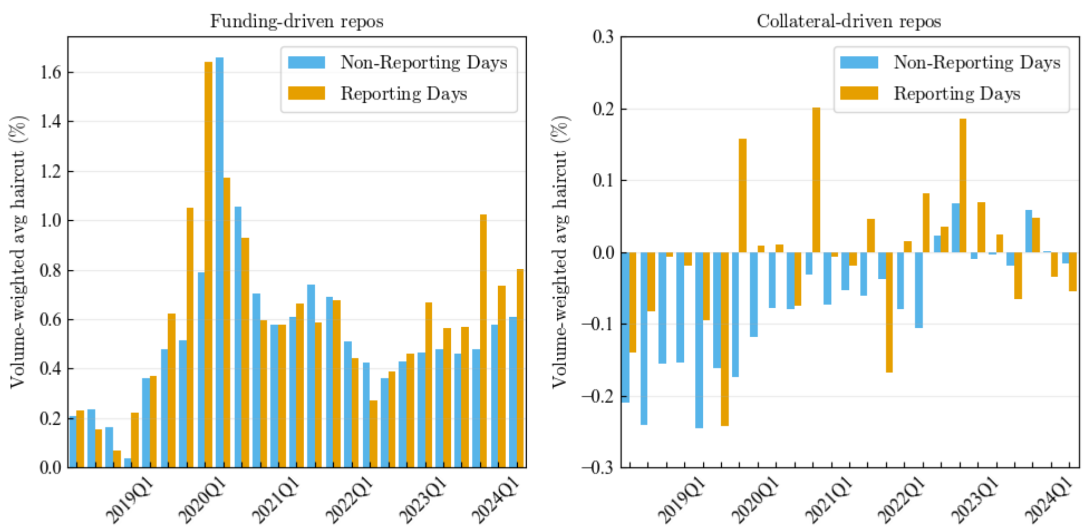

Figure 2.7. Do haircuts matter for balance sheet space?

Notes: The graph plots the volume-weighted average haircut over time of funding- and collateral-driven transactions between the client sector (non-deposit-taking institutions) and the dealer sector. It considers only transactions collateralised by general government bonds. Reporting days are identified as the last two days of each quarter, whereas non-reporting days refer to all remaining days of each quarter.

Reporting-day effects are only part of the story. Because of the obligation to report monthly averages of leverage ratios, traders internalize capital constraints in day-to-day pricing and hence, cross-sectional differences in dealer balance sheet capacity should shape haircuts throughout the quarter. To test this, we sort dealers each quarter into leverage ratio (LR) terciles and estimate the transaction-level difference-in-differences:

$$Haircut_{i,j,c,t} = \sum_{k=1}^{2} \beta_k \left(Tercile_{k,j,q} \times Post_t\right) + \gamma' \mathbf{X}_{i,j,c,t} + \mu_{i,t} + \eta_{j,c} + \varepsilon_{i,j,c,t} \tag{2.2}$$

where $Haircut_{i,j,c,t}$ is the percentage haircut on a repo for security i between dealer j and client c on day t, $Tercile_{1,j,q}$ identifies the most constrained dealers and $Tercile_{3,j,q}$ the omitted least-constrained group, and $Post_t$ equals one after June 28, 2021. Controls $\mathbf{X}_{i,j,c,t}$ comprise log transaction volume and maturity bucket dummies. ISIN×day fixed effects ($\mu_{i,t}$) absorb all daily bond-specific conditions, including collateral scarcity and specialness, so that β_k is identified purely from variation across dealers for the same bond on the same day. Our preferred specification additionally includes dealer×client fixed effects ($\eta_{j,c}$) to control for endogenous relationship sorting. Standard errors are clustered at the dealer level.

Table 2.4 reports the results of this estimation. The coefficients of interest provide clear evidence that haircuts in collateral-driven repos increase monotonically with the strictness of a dealer's leverage constraint in the post-regulatory period. In our baseline and within-relationship specifications (Columns 1 and 2), the interaction terms are positive, monotonically ordered ($Tercile\,1 > Tercile\,2$), and highly statistically significant. Column (3) confirms this relationship using a continuous measure of leverage ratio headroom, demonstrating that as balance sheet space increases, haircuts significantly decrease. Furthermore, the balanced panel estimation in Column (4) confirms that this effect is driven by active repricing within surviving relationships, rather than the entry or exit of specific counterparties.

Importantly, our results demonstrate that this regulatory friction fundamentally alters daily market pricing rather than merely inducing transient window-dressing behaviour. In Column (5), the interaction with a quarter-end dummy (QE) is small and statistically insignificant, in-

Table 2.4. Do balance sheet constraints affect repo haircuts?

Notes: Transaction-level regressions of repo haircuts (%) on dealer leverage constraints. The sample covers bilateral dealer-to-client repos collateralised by euro area government bonds, January 2019 to June 2024. Dealers are sorted into leverage ratio terciles each quarter; Tercile 1 denotes the most constrained dealers and Tercile 3 (omitted) the least. $Post = 1$ after June 28, 2021, when the binding 3% leverage ratio requirement took effect. All specifications include tercile and Post main effects. Controls: log volume, maturity bucket dummies. Standard errors clustered at the dealer level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

	Collateral-driven					Funding-driven
	(1)	(2)	(3)	(4)	(5)	(6)
Tercile 1 × Post	0.224**	0.289***		0.291***	0.224**	−0.245
	(0.079)	(0.071)		(0.068)	(0.078)	(0.298)
Tercile 2 × Post	0.183*	0.191***		0.195***	0.182*	−0.107
	(0.090)	(0.062)		(0.055)	(0.089)	(0.267)
LR headroom × Post			−0.048*			
			(0.025)			
T1 × Post × QE					0.044	
					(0.069)	
T2 × Post × QE					0.039	
					(0.069)	
ISIN × day FE	✓	✓	✓	✓	✓	✓
Client FE	✓		✓		✓	✓
Dealer × client FE		✓		✓		
Dealer FE			✓			
Balanced relationships				✓		
N	1,556,748	1,556,533	1,556,748	1,141,552	1,556,748	1,210,025
R^2	0.82	0.88	0.83	0.88	0.82	0.87

dicating that constrained dealers apply these higher haircuts persistently throughout the entire quarter. Finally, Column (6) re-estimates the baseline on funding-driven repos. Because the leverage ratio regulation does not asymmetrically penalize the balance sheet for reverse repos (cash lending), this segment serves as a natural placebo. As expected, the leverage constraint coefficients in the funding-driven market are statistically indistinguishable from zero.

To quantify this regulatory friction: on an average collateral-driven transaction of € 26.3

million, a highly constrained dealer (Tercile 1) asks the client to lend € 76,000 ($= 26.49$ mn $\times$ 0.289%) less cash than an unconstrained dealer (Tercile 3) would to secure the exact same bond. Given that a typical dealer executes roughly 2,650 collateral-driven trades per month, this pricing wedge effectively saves the constrained dealer over € 200 million in balance sheet capacity monthly. Ultimately, the strict shadow cost of the Basel III leverage ratio creates an incentive for constrained dealers to hand out zero or even positive haircuts, rather than using a negative haircut to hedge against replacement risk.

2.3.2 Convenience yields and haircut setting

In the preceding analysis we examined how dealer's balance-sheet constraints shape haircuts. Of course, client and collateral characteristics should also matter. A large literature documents that investors pay a convenience premium for the safety and liquidity services embedded in high-quality assets (Krishnamurthy and Vissing-Jorgensen, 2012b; Gorton, 2017), and that this premium is reflected in repo rates (Ballensiefen and Ranaldo, 2023; Duffie and Krishnamurthy, 2016b). Convenience yields vary across securities (driven by differences in issuer credit quality, remaining maturity, and benchmark status) and over time with the opportunity cost of holding money (Nagel, 2016) or macroeconomic uncertainty (Moreira and Savov, 2017). Yet while the link between convenience yields and repo *rates* is well established, far less is known about how collateral valuations affect the *haircut* margin.

We argue that convenience yields should influence haircuts through a replacement-risk channel. In a collateral-driven repo the dealer pledges a specific bond in exchange for cash; if the counterparty defaults and fails to return the security, the dealer must repurchase a potentially scarce asset in the open market. The higher the convenience yield on the pledged bond, the costlier this replacement and the greater the dealer's incentive to protect herself by demanding a lower haircut, i.e., demanding more cash in return for the collateral pledged.

We test this prediction by regressing transaction-level haircuts on an ISIN-level convenience yield measure. Specifically, we estimate

$$h_{i,d,c,t} = \beta \, \mathrm{CY}_{i,t} + \mathbf{X}'_{i,d,c,t} \, \delta + \alpha_{d \times c} + \alpha_{k(i) \times t} + \varepsilon_{i,d,c,t}, \tag{2.3}$$

where $h_{i,d,c,t}$ is the haircut (in percent) on transaction i between dealer d and client c on day t, $\mathrm{CY}_{i,t}$ is the ISIN-level convenience yield defined as the spread between the sovereign benchmark yield (interpolated via a Nelson–Siegel curve at the bond's remaining maturity) and the corresponding OIS rate, $\mathbf{X}_{i,d,c,t}$ collects controls for log transaction volume and maturity bucket dummies, $\alpha_{d \times c}$ denotes dealer×client fixed effects, and $\alpha_{k(i) \times t}$ denotes country×day fixed effects (replaced by separate ISIN and day effects in the most saturated specification). Standard errors are clustered at the dealer×client level. Note that we cannot use the most granular ISIN×day fixed effects specification as we have used before as our convenience yield measure varies exactly on this level and all variation would be taken away. The coefficient of interest is β, which we expect to be negative in collateral-driven repos and zero or positive in funding-driven repos.

Table 2.5 reports results separately for collateral-driven repos (columns 1–4) and funding-driven repos (columns 5–8), progressively saturating the specification with fixed effects. Column (1) includes month and collateral country fixed effects; column (2) replaces the collateral country fixed effect with dealer×client fixed effects, absorbing all time-invariant relationship characteristics; column (3) moves to country×day fixed effects alongside dealer×client effects, thereby controlling for all daily collateral country-level demand and supply shocks; Column (4) includes ISIN, day, and dealer×client fixed effects, so all remaining variation in the convenience yield is within-ISIN and within-day, i.e., net of any time-invariant ISIN characteristics and any market-wide daily factors. All columns control for the log transaction volume and maturity bucket dummies.

Across all four specifications, the convenience yield enters with a negative and statistically

Table 2.5. Convenience yields and repo haircuts

Notes: Transaction-level regressions of repo haircuts (%) on the ISIN-level convenience yield (country benchmark yield minus the OIS yield, computed via Nelson–Siegel interpolation at the bond's remaining maturity). Columns (1)–(4) use the collateral-driven sample (dealer borrows cash, pledges bonds); columns (5)–(8) the funding-driven sample (dealer lends cash, receives bonds). Controls: log volume, maturity bucket dummies. Standard errors clustered at the dealer×client level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

	Collateral-driven				Funding-driven			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Convenience yield	−0.138**	−0.127***	−0.121**	−0.057*	0.232**	0.225***	0.279***	0.272**
	(0.057)	(0.034)	(0.057)	(0.032)	(0.096)	(0.055)	(0.096)	(0.129)
Month FE	✓	✓			✓	✓		
Country FE	✓				✓			
Country × day FE			✓				✓	
ISIN + day FE				✓				✓
Dealer × client FE		✓	✓	✓		✓	✓	✓
N	1,440,780	1,440,479	1,440,463	1,440,317	963,571	963,286	963,231	963,128
R^2	0.07	0.69	0.70	0.73	0.09	0.66	0.69	0.73

significant coefficient in the collateral-driven sample. The point estimate in our preferred specification (column 3) implies that a one-standard-deviation increase in the ISIN-level convenience yield is associated with a reduction in haircuts of roughly 12 basis points. This is consistent with the replacement-risk channel: when the pledged bond carries a high convenience premium, the dealer demands more cash in return for parting with it, reflected in a lower haircut, so as to be better insured against the cost of replacing a scarce and valuable security in the event of counterparty default.

In the funding-driven sample (columns 5–8), the coefficient is positive and significant, indicating that dealers *charge* higher haircuts when the collateral they receive is more special. This sign reversal is consistent with dealer market power in funding-driven trades: when the collateral received carries a high convenience yield, the dealer has a stronger incentive to demand over-collateralisation, as a higher haircut allows them to acquire more of the valuable asset per dollar lent and capture the associated specialness premium through rehypothecation. The borrower, needing the funding, absorbs the worse terms.

2.3.3 Dealer market power

A third key friction in bilateral money market segments arises from market power. Establishing and maintaining trading relationships in OTC markets is costly and time-consuming, limiting the number of dealer links that clients can form. As a result, dealer banks may exercise market power over their OTC counterparties (Eisenschmidt et al., 2024; Bittner and Jank, 2025). We expect this effect to vary depending on the underlying motive of the trade. Collateral-driven markets are often concentrated among a small set of dealers specialized in specific collateral classes. In contrast, funding-driven repo markets are more homogeneous and typically involve a broader set of potential counterparties, making them theoretically less susceptible to dealer market power.

To analyse the impact of dealer market power on repo haircuts, we estimate transaction-level panel regressions using both collateral-driven and funding-driven samples. We proxy a client's negotiation power by calculating the number of distinct dealers a customer has traded with over the previous twelve months. We consider these to be the *active* trading relationships that clients have in place, reflecting established master agreements and trading infrastructure. We assume that clients with more active dealer connections have greater negotiation power (and dealers conversely have less market power) in adjusting repo terms. Our baseline empirical specification is estimated as follows:

$$Haircut_{i,j,c,t} = \beta MarketPower_{c,t} + \gamma' \mathbf{X}_{i,j,c,t} + \mu_{i,t} + \eta_c + \lambda_{j,q} + \varepsilon_{i,j,c,t} \tag{2.4}$$

where $Haircut_{i,j,c,t}$ is the haircut (expressed in percentage points) on a repo transaction for security i, between dealer j and client c, on day t. The primary variable of interest, $MarketPower_{c,t}$, counts the number of unique dealers the client has transacted with over a rolling 365-day look-back window. The vector of controls, $X_{i,j,c,t}$, includes the natural logarithm of the transaction volume and fixed effects for maturity buckets. To rigorously isolate the impact of client market power from collateral-specific shocks and macroeconomic conditions, the model incorporates high-dimensional fixed effects. Specifically, $\mu_{i,t}$ denotes security-by-day (ISIN$\times$day) fixed effects, fully absorbing unobserved daily variation in collateral quality and liquidity. Unobserved, time-invariant client characteristics are captured by client fixed effects, η_c, which are replaced by dealer-by-client fixed effects in our within-relationship specifications to control for endogenous relationship capital. Finally, $\lambda_{j,q}$ represents dealer leverage tercile fixed effects, absorbing time-varying balance sheet constraints at the dealer level. Standard errors are clustered at the dealer level across all specifications.

Table 2.6 reports the results of these panel regressions. Economically, clients in the collateral-driven repo market segment are incentivised to negotiate repo haircuts upwards (i.e., lending less cash for a given amount of highly sought-after securities), whereas clients in the funding-driven segment have the opposite incentive to negotiate haircuts downwards (i.e., borrowing more cash for a given amount of securities posted as collateral). The regression results confirm this dual dynamic with striking symmetry. In the funding-driven segment (Columns 4–6), the coefficient on $\#Dealers$ is negative and highly significant, indicating that clients with a broader network of active dealer relationships successfully negotiate lower haircuts. Conversely, in the collateral-driven segment (Columns 1–3), the coefficient is positive and statistically significant. This confirms that when clients hold scarce collateral and possess greater negotiation leverage, they successfully push haircuts upwards to secure more favorable terms.

Furthermore, the interaction terms between our market power proxy and the Post-June 2021 indicator ($\#D \times Post$) in Columns (3) and (6) reveal a significant attenuation of this market power premium in both market segments. We interpret this reversal through the lens of strict dealer balance sheet constraints. On June 28, 2021, the 3% minimum leverage ratio requirement

Table 2.6. Market power

Notes: Transaction-level regressions of repo haircuts (%) on the number of dealers a client transacted with over the prior 365 days. The sample covers bilateral dealer-to-client repos collateralised by euro area government bonds, January 2019 to June 2024. Columns (7)–(8) restrict to dealer–client pairs active both pre- and post-June 2021. "Post" denotes observations after June 28, 2021. Controls: log volume, maturity bucket dummies. Standard errors clustered at the dealer level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

	Collateral-driven			Funding-driven			Balanced	
	(1)	(2)	(3)	(4)	(5)	(6)	(7) CD	(8) FD
#Dealers	0.038***	0.028*	0.077***	−0.107***	−0.093***	−0.142***	0.033**	−0.096***
	(0.012)	(0.014)	(0.021)	(0.024)	(0.018)	(0.033)	(0.014)	(0.015)
#D × Post			−0.064***			0.121**		
			(0.018)			(0.048)		
ISIN × day FE	✓	✓	✓	✓	✓	✓	✓	✓
Client FE	✓		✓	✓		✓		
Dealer × client FE		✓			✓		✓	✓
Leverage tercile FE	✓	✓	✓	✓	✓	✓	✓	✓
Controls	✓	✓	✓	✓	✓	✓	✓	✓
Balanced							✓	✓
N	1,556,748	1,556,533	1,556,748	1,210,025	1,209,798	1,210,025	1,141,552	968,758
R^2	0.82	0.88	0.83	0.87	0.89	0.87	0.88	0.89

became formally binding for European banks under the Capital Requirements Regulation (CRR II). This regulatory shift exogenously increased the cost of balance sheet capacity. Because repo transactions are highly balance-sheet intensive but generate relatively low margins, the strictly binding leverage ratio severely restricted dealers' ability to flexibly accommodate the costly pricing preferences of their counterparties. Consequently, the strict shadow cost of the regulatory constraint outweighed the negotiation leverage of even highly connected clients, forcing dealers to retract favourable haircut terms and effectively compressing the market power premium across the bilateral repo market. These results mirror the findings in Bittner and Jank (2025).

To quantify the economic magnitude of this effect, we can evaluate the impact of a marginal increase in dealer relationships for an average client in the funding-driven segment. Based on our sample statistics, the average funding-driven transaction volume is € 34.49 million, subject to an average haircut of 1.169%. Our baseline regression estimate indicates that establishing one additional active dealer relationship (an increase of 1 in the $\#Dealers$ variable) reduces the applied haircut by 0.107 percentage points (approximately 11 basis points). For an average-sized transaction of € 34.49 million, this translates to a client having to post roughly € 38,000 less in sovereign collateral to secure the same amount of cash (€ 34.49 mn $\times$ 0.107%). While this per-transaction saving may appear modest in isolation, its aggregate impact is substantial: given that an active dealer–client pair executes an average of 46 funding-driven transactions per month, this enhanced negotiation leverage frees up nearly € 1.7 million in encumbered collateral per month for a typical client, significantly easing their operational liquidity constraints.

2.3.4 Robustness

Some facts may bias our results. First, the haircut data is highly centred around zero and next to the three frictions explained above, some stickiness within client-dealer pairs or contractual agreements unobserved from the data could be the reason for the zero haircuts, and hence falsely attribute the statistical effect. To address this concern we are repeating all three regression tables on the basis of transactions which had a non-zero haircut (hence considering only about 30% of the sample) in Appendix 2.A.2. The directions of the estimates as well as their significance are widely unchanged.

Second, different counterparties may have different contracting terms and trade conditions that we do not observe and which may change our results. As we are mostly interested in the NBFI-sector, we repeat the exercise only considering the sectors S124 (non-money market investment funds) and S125 (non-bank financial intermediaries) and thus mirroring the sample in Bittner and Jank (2025). Again, the significance as well as the direction of the estimates remains stable (cmp. Appendix 2.A.2).

Third, haircuts in bilateral repos may be negotiated at the portfolio level rather than individually for each transaction. If a dealer and client agree on a standing haircut schedule for a given collateral class, individual transaction-level observations reflect a single underlying contracting decision observed repeatedly. To ensure that our results are not driven by this within-relationship granularity, we aggregate the data to the dealer–client–collateral country–month level and re-estimate all main specifications on these aggregated cells. This makes the unit of observation coincide with the plausible unit of haircut-setting. The significance, direction, and economic interpretation of all estimates remain stable under aggregation (cmp. Appendix 2.A.1).

Lastly, we ask whether our results are driven by the leverage ratio specifically or by general balance-sheet scarcity. If leverage-constrained dealers are also capital-constrained, our estimates might reflect risk-weighted capital costs rather than the non-risk-weighted leverage penalty. Two observations argue against this. First, while the leverage ratio and the consolidated total capital ratio are positively correlated in our sample, the association is moderate ($corr = 0.4$), indicating that the two ratios capture meaningfully different dimensions of balance-sheet capacity. Second, we re-estimate our baseline specification (Table 2.4) replacing leverage ratio terciles with total capital ratio terciles. The results, reported in Table 2.A.8, are mostly insignificant or, if anything, carry a negative sign for collateral-driven repos. This is consistent with the regulatory asymmetry at the heart of our mechanism: under risk-weighted capital requirements, sovereign collateral attracts zero or near-zero risk weights, so collateral-driven repos are inexpensive in capital terms. The leverage ratio, by contrast, is non-risk-weighted and penalises the cash leg one-for-one regardless of collateral quality.

2.3.5 Takeaways

The implications of our empirical findings are notable. In collateral-driven repos, dealer balance-sheet constraints exert upward pressure on haircuts, while higher convenience yields, reflecting greater replacement risk, exert downward pressure. As a result, government bonds with high convenience yields can nonetheless carry relatively higher haircuts if the effect of dealer balance-sheet constraints dominates the convenience-yield effect. In funding-driven repos, by contrast, balance-sheet constraints and asset convenience yield do not play a comparable role. Haircuts, therefore, display a more monotonic relationship with collateral characteristics, with lower (higher) haircuts for bonds with lower (higher) risks. However, because the funding-driven segment is more homogeneous, dealer—and to some extent client—market power becomes a more important determinant of haircut levels.

2.4 Model

We develop a partial-equilibrium model of the bilateral repo market that embeds dealer balance-sheet constraints, collateral convenience yields, and market power. The model delivers an equilibrium in which haircuts, rates, and volumes are jointly determined through a two-stage bargaining game à la Kreps and Scheinkman (1983). We model the repo market as a Bertrand competition model with capacity pre-commitment (Kreps and Scheinkman, 1983). The key departure from their model is that we allow for varying degrees of dealer market power by modelling the second stage of the bargaining problem as a bilateral Nash-bargaining problem similar to Infante (2019). The limiting case where dealers are holding the entire market power and clients are assumed to be price takers ($\theta_j = 1$) leads us back to the classic equilibrium of Kreps and Scheinkman (1983), which, of course, is essentially a Cournot equilibrium. When clients hold positive bargaining power ($\theta_j < 1$), the surplus split departs from the competitive benchmark but crucially, the jointly efficient haircut remains determined by surplus maximisation, not by the distribution of bargaining power. This two-stage structure is motivated by two institutional features of the bilateral repo market. First, dealers face hard collateral capacity constraints determined by their bond inventories and prime brokerage agreements, making volume pre-commitment a natural first-stage choice. Second, once capacity is allocated, individual contract terms are negotiated bilaterally in an OTC setting where dealers hold significant market power Eisenschmidt et al. (2024), making Nash-in-Nash bargaining the appropriate Stage 2 model.

Subsequently, we derive three corollaries on how frictions impact haircuts, revisiting earlier empirical findings. The model also delivers two predictions about the consequences for market outcomes (i.e., financial stability and monetary policy) that we take to the data in Section 2.5.

2.4.1 Model structure

We focus on the *collateral-driven segment*; the funding-driven structure is mostly analogous and covered in the Appendix 2.C. The market consists of D dealers indexed by d, $J = 4$ collateral countries indexed by $j \in \{\mathrm{DE}, \mathrm{FR}, \mathrm{IT}, \mathrm{ES}\}$, and a continuum of clients. Each dealer is characterised by a leverage ratio ℓ_d, predetermined within the trading period, and a collateral capacity $\bar{C}_{dj}$, the maximum amount of country-j collateral the dealer can lend.

A bilateral repo trade between dealer d and a client over country-j collateral is specified by three terms: the face value of collateral posted C_{dj}, the haircut h_{dj}, and the repo rate r_{dj}. The cash exchanged at inception is $V_{dj} = (1 - h_{dj})\, C_{dj}$, analogous to the official definition in (2.1).

Timing. The game proceeds in two stages within each trading period. In Stage 1, each dealer d simultaneously commits a volume of country-j collateral $C_{dj} \leq \bar{C}_{dj}$ to lend in the bilateral market; committed volumes are publicly observed before Stage 2 begins. In Stage 2, aggregate supply $Q_j = \sum_d C_{dj}$ pins down the market-clearing effective cost ρ_j through inverse demand, and each dealer–client pair then bargains bilaterally over the haircut h_{dj} and the repo rate r_{dj}, taking all other pairs' outcomes as given (Nash-in-Nash). The client allocates demand to the dealer offering the lowest effective cost $\rho_{dj} = -r_{dj} - \lambda\, h_{dj}$.

Dealer profit. Dealer d's expected profit from a collateral-driven trade in country j is

$$\pi_{dj}^{\mathrm{CD}} \;=\; \Big[r_{\mathrm{DFR}} - (1-p)\,r_{dj} - \frac{y_j + p\,L}{1-h_{dj}}\Big](1-h_{dj})\,C_{dj} \;-\; \frac{\gamma_d\,\mu_j}{2}\,V_{dj}^2, \tag{2.5}$$

where r_{DFR} is the deposit facility rate earned on the cash received, r_{dj} is the repo rate paid at maturity conditional on the counterparty surviving with probability $1-p$, y_j is the foregone convenience yield on country-j collateral, pL is the expected replacement cost, γ_d is a dealer-specific balance-sheet cost, and μ_j is a country-specific multiplier, which helps us to incorporate the fact that due to home bias in dealer trading, the dealers active in DE/FR collateral tend to be more leverage-constrained.

The balance-sheet penalty is quadratic in the cash leg $V_{dj} = (1-h_{dj})C_{dj}$ because the leverage-ratio exposure measure expands one-for-one with the cash the dealer receives, and marginal balance-sheet costs are increasing as the dealer approaches its regulatory limit. In funding-driven repos, the balance-sheet cost is normalised to zero ($\Phi \equiv 0$). The dealer-specific cost γ_d is linked to the observed leverage ratio through the mapping

$$\gamma_d = \gamma_0 + \gamma_1 \log\!\big(1 + e^{-\eta(\ell_d - \ell_{\min})}\big), \tag{2.6}$$

centred on the regulatory minimum $\ell_{\min} = 3\%$.[15]

Client surplus. A client with valuation v for country-j collateral evaluates a contract by its effective cost

$$\rho_{dj} = -r_{dj} - \lambda\, h_{dj}, \tag{2.7}$$

[15]The softplus function $\log(1+e^x)$ is a smooth, everywhere-differentiable approximation to $\max(0,x)$. We prefer it to the inverse-distance form $\gamma_0 + \gamma_1/(\ell_d - \ell_{\min})^\eta$ because the latter diverges as $\ell_d \to \ell_{\min}$, causing numerical instability for the most constrained dealers.

where $\lambda > 0$ is the client's sensitivity to the haircut margin. The client prefers a higher repo rate (which she receives from the dealer at the second leg of the transaction) and a higher haircut (which reduces the cash she must post in return for the collateral at the first leg). The client's surplus from trading with dealer d at effective cost ρ_{dj} is

$$U_{dj} = (v - \rho_{dj})\, C_{dj}. \tag{2.8}$$

The client trades if and only if $v > \rho_{dj}$.

Bargaining. The bilateral dealer-to-client repo market is a highly oligopolistic OTC market in which a small number of dealer banks exert significant market power over their counterparties. Each dealer–client pair bargains bilaterally, taking the outcomes of all other pairs as given (Nash-in-Nash bargaining). For a given match, the dealer and client jointly choose the effective cost ρ_{dj} and the haircut h_{dj} to maximise the generalised Nash product

$$\max_{\rho_{dj},\, h_{dj}} \; N \;=\; \left(\pi_{dj}^{\mathrm{CD}}\right)^{\theta_j} \left(U_{dj}\right)^{1-\theta_j}, \tag{2.9}$$

where the dealer's bargaining share in country j is

$$\theta_j = \frac{1}{1 + \kappa(n_j - 1)}, \tag{2.10}$$

n_j is the number of dealers active in country j, and $\kappa > 0$ governs the decay of bargaining power with competition. Equation (2.9) is reminiscent of the model structure in Infante (2019).

Demand. On the client side, a continuum of potential counterparties with heterogeneous valuations for country-j collateral generates aggregate demand $D_j(\rho) = M_j(\bar{v}_j - \rho)/\bar{v}_j$, where $\bar{v}_j$ is the upper bound of the valuation distribution and M_j scales the market size. The linear form follows from a uniform distribution of client valuations on $[0, \bar{v}_j]$: a client with valuation v_k trades if and only if $v_k > \rho_j$, so the mass of active clients is $M_j(1 - \rho_j/\bar{v}_j)$.

2.4.2 Equilibrium

Dealers compete in a two-stage game: in Stage 1, each dealer d commits a volume $C_{dj} \leq \bar{C}_{dj}$; in Stage 2, each dealer–client pair bargains over (r_{dj}, h_{dj}) by solving the Nash product (2.9), and the client allocates demand to the dealer offering the lowest effective cost ρ_{dj}.

Definition 3 (Equilibrium)**.** An equilibrium is a collection of contracts $\{(r^*_{dj}, h^*_{dj}, C^*_{dj})\}$ for all active dealer–country pairs such that:

(i) each (r^*_{dj}, h^*_{dj}) solves the Nash bargaining problem (2.9) given C^*_{dj} (Stage-2 optimality),

(ii) each C^*_{dj} maximises π^{CD}_{dj} evaluated at the Stage-2 optimum, subject to $C_{dj} \leq \bar{C}_{dj}$ (Stage-1 optimality), and

(iii) markets clear: $\sum_d C^*_{dj} = D_j(\rho^*_j)$ for each country j.

We solve the game by backward induction. Without loss of generality, we set $r_{\text{DFR}} = 0$, expressing repo rates as spreads from the deposit facility rate.

Proposition 2 (Equilibrium of the two-stage game)**.** *The equilibrium of the collateral-driven repo bargaining game is characterised as follows.*

(i) *For a given committed volume C and effective cost ρ_j, the jointly efficient haircut maximises the dealer's profit and is given by*

$$h^*(C) = \frac{(1-p)\,\lambda + \gamma_d \mu_j\, C - (1-p)\,\rho_j}{2(1-p)\,\lambda + \gamma_d \mu_j\, C}. \tag{2.11}$$

(ii) *The profit-maximising volume satisfies*

$$(1-p)(\rho_j+\lambda\, h^*)(1-h^*)-(y_j+pL)-\gamma_d\mu_j(1-h^*)^2\, C+(1-p)(1-h^*)\,\frac{\partial \rho_j}{\partial C_{dj}}\, C=0, \tag{2.12}$$

where $\partial\rho_j/\partial C_{dj} = -\bar{v}_j/M_j < 0$ *is the dealer's price impact through market clearing.*

Proof. See Appendix 2.B.1. □

Note that the bargaining share θ_j does not appear in the optimal haircut: under generalised Nash bargaining with transferable utility, the haircut is chosen to maximise the total surplus, and bargaining power affects only the surplus split through the effective cost ρ_j.[16] In Appendix 2.B.2 we also show that for this game, an equilibrium always exists .

The following three corollaries revisit the empirical findings from Section 2.3 and claim that in equilibrium, the effects observed in the data are predicted by the model.

Corollary 1. *The equilibrium haircut is increasing in the dealer's balance-sheet cost*

$$\frac{\partial h^*}{\partial \gamma_d} > 0.$$

Proof. See Appendix 2.B.3. □

Corollary 2. *The equilibrium haircut is decreasing in the collateral's convenience yield*

$$\frac{dh^*}{dy_j} < 0.$$

Proof. See Appendix 2.B.4. □

[16]When $\theta_j = 1$, the Stage-2 bargaining gives the dealer the entire surplus, so the negotiated effective cost already maximises π over ρ_j. In this limiting case $\partial\pi/\partial\rho_j = 0$ at the bargaining solution, the Cournot term in (2.12) vanishes, and the volume FOC reduces to $\partial\pi/\partial C = 0$. One can then show by substitution that the equilibrium haircut collapses to $h^* = 1 - \sqrt{(y_j + pL)/[(1-p)\lambda]}$, which is independent of the dealer's balance-sheet cost γ_d. The balance-sheet channel (Corollary 1) therefore operates in general equilibrium only when $\theta_j < 1$, i.e. when the bargaining constrains the dealer's effective cost below the profit-maximising level.

Corollary 3. *An increase in dealer market power triggers two opposing forces on the equilibrium haircut: a negative direct effect via the negotiated effective cost, and a positive indirect effect via an expansion in dealer supply*

$$\frac{dh^*}{d\theta_j} = \underbrace{\frac{\partial h^*}{\partial \rho_j}\frac{d\rho_j}{d\theta_j}}_{\textit{Direct effect } (<0)} + \underbrace{\frac{\partial h^*}{\partial C}\frac{dC}{d\theta_j}}_{\textit{Indirect effect } (>0)} .$$

Proof. See Appendix 2.B.5. □

Together, the three corollaries rationalise the cross-country haircut puzzle documented in Section 2.3: German and French collateral carry higher haircuts than Italian and Spanish collateral in collateral-driven repos, despite having higher convenience yields (contrary to classic risk-management theory which would predict that the party with the majority of the market power would ask for more cash in return for scarce and valuable collateral). This is because the dealers intermediating core-country collateral tend to be more balance-sheet constrained. The balance-sheet channel (Corollary 1) dominates the convenience-yield channel (Corollary 2) when dealers are sufficiently constrained.

Model fit. To test whether our model able to produce results that are roughly in line with earlier empirical results, we are using observed MMSR averages and calibrate the μ_j-parameters to match the observed coefficients in Table 2.4.

This simple baseline calibration yields a heat map that mirrors the empirical magnitudes reasonably well. As expected, tighter dealers are incentivised to offer structurally higher haircuts in collateral-driven markets. Noticeable is also that this effect is higher in the jurisdictions Germany and France, which matches the empirical observation that haircuts tend to be higher in these jurisdiction despite their higher convenience yield. As explained in the summary statistics, dealers that are predominantly operating in the German and the French jurisdiction tend to have higher tighter leverage ratios and scarcity makes adds to the market power of some dealers which in effect are pushing haircuts upwards. In the model these channels are captured by the parameter μ_j.

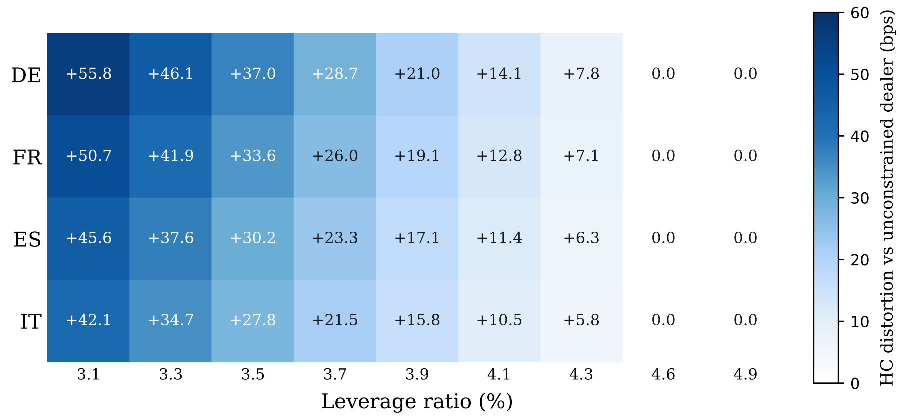

Figure 2.8. Haircut distortion and balance sheet space

Notes: Each cell reports the model-implied haircut distortion (in basis points) relative to an unconstrained dealer, evaluated at the equilibrium of the two-stage game described in Section 2.4. Rows correspond to collateral countries; columns to dealer leverage ratios. The country multipliers, μ_j are calibrated to match the regression coefficients in Table 4; all other parameters are set to economically motivated values described in Appendix 2.D.

2.4.3 Consequences of constrained haircuts

We define a haircut as *constrained* when balance-sheet constraints push it away from the level that would be set by an unconstrained dealer ($\gamma_d = \gamma_0$). When haircuts cannot absorb risk and price signals—because the balance-sheet channel stiffens the first-order condition—the resulting adjustment is displaced onto other contract terms: rates and volumes. We derive two predictions that formalise this displacement and that we test in Section 5.

Prediction 1 (Suboptimal haircuts increase rate volatility)**.** *The standard deviation of the bilateral repo rate is increasing in the dealer's balance-sheet cost: constrained dealers display higher rate volatility than unconstrained dealers, conditional on the same collateral and the same client.*

The mechanism operates through a displacement effect. Consider for instance, a shock to the convenience yield, Δy_j, reflecting a shift in the perceived safety or convenience yield of this collateral. For an unconstrained dealer, the haircut absorbs this shock: by Corollary 2, $\partial h^*/\partial y_j < 0$, so a rise in y_j lowers the haircut, insulating the bilateral rate from the collateral revaluation. For a constrained dealer, however, the balance-sheet penalty stiffens the optimal haircut, shrinking $|\partial h^*/\partial y_j|$ toward zero. Because the haircut cannot adjust, the shock that would ordinarily be absorbed by the margin instead spills over into the rate via the pricing identity $r^*_{dj} = -\rho_j - \lambda h^*_{dj}$: if h does not move, r must. Under repeated convenience yield fluctuations, this displacement generates structurally higher rate volatility for constrained dealers, precisely in the collateral-driven segment where the balance-sheet channel is operative.

Prediction 2 (Suboptimal haircuts impair monetary policy pass-through)**.** *The pass-through of a deposit facility rate change to bilateral repo rates is decreasing in the dealer's balance-sheet cost: constrained dealers transmit a smaller share of policy-rate changes than unconstrained dealers.*

The mechanism operates through two channels. First, a hike in r_{DFR} mechanically raises the return on the existing cash leg, generating a surplus windfall that both constrained and unconstrained dealers share with clients through bargaining. This direct channel transmits the bulk of the rate change to bilateral repo rates regardless of dealer leverage. Second, the higher return on cash creates an incentive to *expand* the cash leg: an unconstrained dealer optimally lowers the haircut, increasing V_{dj}, which generates additional surplus and pushes the bilateral rate further via the pricing identity ($r^*_{dj} = -\rho_j - \lambda h^*_{dj}$). A constrained dealer faces a quadratic balance-sheet penalty (γ_d) that stiffens the optimal haircut, preventing this amplification. Because the haircut cannot adjust, the rate moves only through the direct surplus-sharing channel, not through contract redesign.

Figure 2.9 summarises the two prediction in figures. Both predictions are qualitatively strongly supported by the model, but may quantitively vary significantly across different calibrations. Pinning down the precise magnitude will ultimately be the empirical exercise of

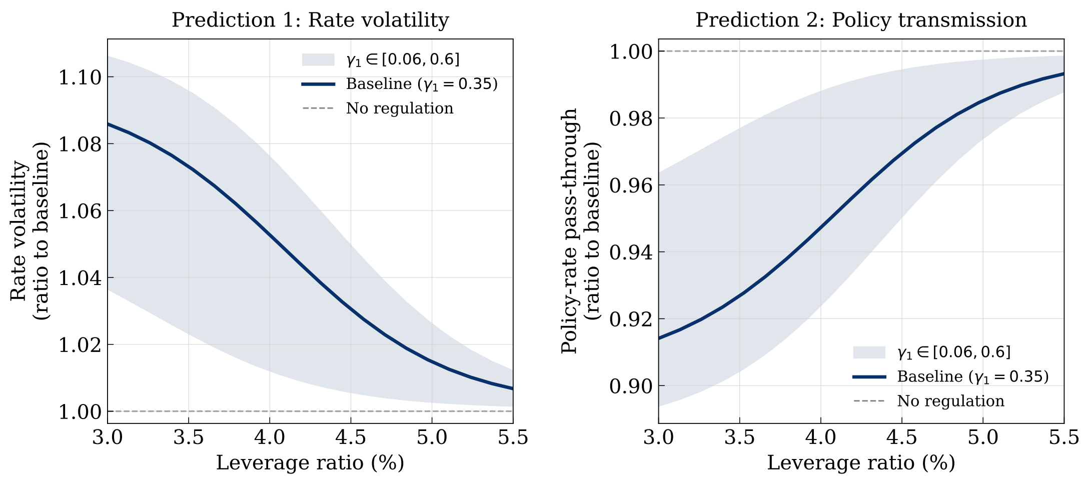

Figure 2.9. Model predictions about market outcomes

Notes: The left panel plots the standard deviation of bilateral repo rates (ratio to the unconstrained baseline) as a function of the dealer's leverage ratio; the right panel plots the pass-through of a deposit facility rate change to bilateral repo rates (ratio to full pass-through). Both panels are evaluated at the equilibrium of the collateral-driven model described in Section 2.4, averaging across collateral countries.

Section 2.5.

2.5 Implications for market stability and monetary policy

In Section 2.3, we documented how balance sheet space,convenience yields and market power are impacting haircuts. In Section 2.4 we designed a model which embedded these frictions in a partial-equilibrium, oligopolistic model which featured to testable predictions regarding the market outcome: constrained dealers should exhibit higher rate volatility (Prediction 1) and weaker monetary policy pass-through (Prediction 2). We now test these two predictions empirically.

2.5.1 Market stability

Prediction 1 states that the standard deviation of repo rates is increasing the more balance sheet constraint the dealer is. The intuition is that that when is stiffened by regulatory constraints, it can now longer react as much as needed to shocks to the underlying collateral or changes in the clients riskiness. Hence other contracting terms have to absorb the fluctuations, rendering the rate more volatile.

Reduced-form evidence. We aggregate transactions to the dealer–collateral country–month level, retaining cells with at least three transactions, and compute the standard deviation of DFR-adjusted repo rates within each cell, $\sigma_{r,d,j,t}$. Our treatment exploits the cross-sectional variation in leverage constraints documented in Sections 3 and 4. We sort dealers into terciles of the leverage ratio distribution each quarter and estimate

$$\sigma_{r,d,j,t} = \sum_{k=1}^{2} \beta_k \, \mathbb{1}_{\{\text{Tercile}_{d,t}=1\}} \times \text{Post}_t + \mathbf{X}'_{d,j,t}\,\gamma + \alpha_j + \delta_t + \omega_d + \eta_k + \varepsilon_{d,j,t}\,, \tag{2.13}$$

where Tercile_1 denotes the most constrained dealers (lowest leverage ratio headroom), Post_t equals one after June 2021 and $\mathbf{X}_{d,j,t}$ includes average log volume, average maturity, and $\log N_{\text{tx}}$ as controls. We include collateral country, month, dealer, and leverage tercile fixed effects. The omitted category is the least constrained tercile. As an alternative, we replace the tercile interactions with a continuous measure, $(\text{LR}_{d,t} - 3\%) \times \text{Post}_t$, where $\text{LR}_{d,t} - 3\%$ is the dealer's leverage ratio headroom above the regulatory minimum in percentage points.

Table 2.7 reports the results. In columns (1) and (2), estimated on collateral-driven (CD) repos, the most constrained dealers exhibit a statistically significant increase in rate volatility after the relief expires: the coefficient on $\text{Tercile}_1 \times \text{Post}$ is 0.032, significant at the 10% level when clustering at the dealer level and at 5% when clustering at the dealer–collateral country level. The continuous specification in columns (5)–(6) confirms: a one percentage point reduction in

Table 2.7. Leverage constraints and repo market volatility

Notes: Unit of observation is dealer × collateral country × month. The dependent variable is the standard deviation of DFR-adjusted rates (σ_r) within each cell. Cells with fewer than 3 transactions are dropped. Controls: average log volume, average maturity, $\log(N_{\text{tx}})$. Tercile 1 = most constrained; Tercile 3 (omitted) = least. *Post* = 1 after June 28, 2021. Odd columns cluster at the dealer level; even columns cluster at dealer × collateral country. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

	CD		FD (placebo)		CD — continuous	
	(1)	(2)	(3)	(4)	(5)	(6)
Tercile 1 × Post	0.032*	0.032**	0.032	0.032*		
	(0.016)	(0.010)	(0.023)	(0.018)		
Tercile 2 × Post	0.016	0.016*	−0.003	−0.003		
	(0.012)	(0.009)	(0.019)	(0.016)		
LR headroom × Post					−0.013**	−0.013***
					(0.006)	(0.005)
FE: Coll. country	✓	✓	✓	✓	✓	✓
FE: Month	✓	✓	✓	✓	✓	✓
FE: Dealer	✓	✓	✓	✓	✓	✓
Cluster: Dealer	✓		✓		✓	
Cluster: Dealer × Country		✓		✓		✓
N	10,683	10,683	12,829	12,829	10,683	10,683
R^2	0.30	0.30	0.25	0.25	0.30	0.30

leverage headroom raises σ_r by approximately 1.3 basis points post-treatment ($p < 0.01$ with dealer–country clustering).

Columns (3) and (4) provide a placebo test. We estimate the identical specification on funding-driven (FD) repos, where the dealer is on the lending side and the leverage constraint mechanism of Section 2.3 does not predict haircut distortion. The point estimates are economically similar but lack precision and, crucially, the coefficient on Tercile_2 × Post flips sign, consistent with the absence of a treatment effect in this segment.

IV estimates. The reduced-form estimates in Table 2.7 establish that leverage-constrained dealers generate more volatile repo pricing after the regulatory change took effect, but they do not isolate the haircut distortion channel from other behavioural changes that constrained dealers may undertake simultaneously (e.g. reduced market-making, client selection). To sharpen the channel, we instrument the mean haircut in each cell, $h_{d,j,t}$, with the leverage constraint variables from equation (2.13). The first stage regresses the mean haircut on the leverage constraint variables (which served as direct regressors in the reduced form) and the same controls and fixed effects:

$$h_{d,j,t} = \psi\, Z_{d,t} \times \text{Post}_t + \mathbf{X}'_{d,j,t}\,\phi + \alpha_j + \delta_t + \omega_d + \eta_k + u_{d,j,t}\,, \tag{2.14}$$

where $Z_{d,t}$ is either the continuous leverage headroom, $\text{LR}_{d,t} - 3\%$, or the pair of tercile indicators $(\mathbb{1}_{\{\text{Tercile}_{d,t}=1\}},\, \mathbb{1}_{\{\text{Tercile}_{d,t}=2\}})$. The second stage replaces the leverage constraint variables with the instrumented mean haircut:

$$\sigma_{r,d,j,t} = \beta\, \widehat{h}_{d,j,t} + \mathbf{X}'_{d,j,t}\,\gamma + \alpha_j + \delta_t + \omega_d + \eta_k + \varepsilon_{d,j,t}\,. \tag{2.15}$$

We are upfront about the limitations of this exercise. A clean causal interpretation of β requires that leverage constraints affect rate volatility *only* through their effect on the haircut, i.e.

$$\mathbb{E}[Z_{d,t} \times \text{Post}_t \cdot \varepsilon_{d,j,t} \mid \mathbf{X}, \text{FE}] = 0. \tag{2.16}$$

This is a strong assumption in our setting. The literature on window dressing (e.g. Ranaldo et al., 2021) documents that leverage-constrained dealers adjust repo rates directly around quarter-ends, which would be a violation of the exclusion restriction. Our specification addresses this by including dealer, collateral country, and month fixed effects, and excluding quarter-end days which absorb the average quarter-end effect. The residual identifying variation is within-dealer, within-month, across collateral countries. Whether leverage constraints also influence rates directly *throughout* the quarter (independent of the haircut channel) is an empirical question that we cannot resolve with the data at hand. Under this assumption, the exclusion restriction is plausibly satisfied and β admits a causal interpretation; to the extent that it is not, β should be read as an estimate of the haircut–volatility link, consistent with (but not proof of) the structural mechanism in Section 2.4.

Table 2.8. Haircut distortion and repo rate volatility (IV estimates)

Notes: 2SLS estimates of the effect of mean haircut distortion (h) on repo rate volatility (σ_r, standard deviation of DFR-adjusted rates). Unit of observation is dealer × collateral country × month. Panel A reports first-stage coefficients and the F-statistic; Panel B reports the 2SLS coefficient on h. Controls: average log volume, average maturity, $\log(N_{\text{tx}})$. Column (6): robustness with minimum 5 transactions per cell. Column (1) clusters at the dealer level; all other columns cluster at dealer × collateral country. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

	LR headroom		Tercile	Saturated FE		Robust.
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: First stage — dependent variable: h (mean haircut)						
LR headroom × Post	−0.164	−0.164		−0.135		−0.168
Tercile 1 × Post			0.306		0.244	
Tercile 2 × Post			0.106		0.076	
F-statistic	14.19	16.80	7.21	12.36	6.55	16.08
Panel B: Second stage (2SLS) — dependent variable: σ_r						
$\hat{h}$	0.077	0.077^{**}	0.096^{**}	0.078^{*}	0.096^{**}	0.075^{**}
	(0.047)	(0.033)	(0.039)	(0.041)	(0.044)	(0.034)
Instrument	LR cont.	LR cont.	Tercile	LR cont.	Tercile	LR cont.
Controls	✓	✓	✓	✓	✓	✓
FE: Coll. country, month, dealer	✓	✓	✓			✓
FE: Coll. country × month, dealer				✓	✓	
Cluster: Dealer	✓					
Cluster: Dealer × Country		✓	✓	✓	✓	✓
N	10,683	10,683	10,683	9,333	9,333	10,007

Table 2.8 reports the results. Panel A presents the first stage. A one percentage point increase in leverage headroom reduces the mean haircut by approximately 0.16 percentage points after the relief expires (columns 1–2), with F-statistics of 14–17, comfortably above the Stock and Yogo (2005) threshold for 10% maximal bias. The tercile instruments in column (3) yield a

weaker first stage ($F = 7.2$), reflecting the coarser variation, though the pattern of coefficients—most constrained dealers concede the largest haircuts—is consistent with the reduced-form evidence. Columns (4)–(5) saturate the model with collateral country×month fixed effects, absorbing time-varying sovereign risk; first-stage relevance is attenuated but preserved for the continuous instrument ($F = 12.4$).

Panel B reports the second-stage estimates. Across all specifications, the coefficient on h is positive: a one percentage point increase in the instrumented mean haircut—driven by tighter leverage constraints—raises rate volatility by 7.5–9.6 basis points. The estimates are statistically significant at the 5% level in all columns that cluster at the dealer–country level (columns 2–6), and are robust to saturated fixed effects (column 4) and a stricter minimum transaction threshold (column 6). Column (1), which clusters at the dealer level only and thus relies on fewer effective clusters, yields the same point estimate but wider confidence intervals ($p = 0.13$).

The economic magnitude is meaningful. Taking the preferred estimate of 0.077 (column 2), the implied effect of moving from the least to the most constrained tercile—corresponding to a mean haircut shift of roughly 0.3 percentage points in the first stage—is an increase in rate volatility of approximately 2.3 basis points, or about 7% of the unconditional standard deviation of σ_r. While this is a modest effect in absolute terms, it arises in overnight sovereign repo—a market segment where participants value pricing stability and small dislocations can propagate through collateral chains.

Taken together, the reduced-form and IV results point in the same direction: when leverage constraints prevent dealers from pricing haircuts optimally, the resulting risk is not absorbed costlessly but instead transmits into rate volatility. The evidence is most compelling for collateral-driven repos, where the theoretical mechanism is sharpest.

2.5.2 Monetary policy transmission

Our results also have *suggestive* implications for the monetary policy transmission process. The "monetary policy transmission pipeline" involves three distinct steps: (i) the central bank's monetary policy actions pass through into the money market, (ii) money market conditions propagate to debt and equity markets, and (iii) asset price changes transmit into the real sector. Focusing on the first step, "in an idealized money market, any change in the main monetary policy rate should pass through perfectly to all money market rates" (Corradin et al., 2020, p. 13). Recent work has shown that frictions in the repo market can impair this pass-through. In particular, Eisenschmidt et al. (2024) document that dealer market power in OTC repo markets reduces pass-through of the ECB's policy rate to non-dealer counterparties by 20 to 30 percentage points relative to the inter-dealer benchmark. Our analysis complements this finding by identifying a distinct, *within*-dealer-sector channel: balance sheet constraints impede the transmission of rate changes specifically in collateral-driven repos.

To measure the monetary policy pass-through, we exploit the eleven changes in the ECB's Governing Council (GC) rate—ten increases in the deposit facility rate (DFR) between July 2022 and September 2023, and the June 2024 rate cut. For each rate decision, we construct a transaction-level panel by stacking a symmetric ± 7 business day window around the announcement. The dependent variable is the *DFR-adjusted* repo rate, defined as the transaction-level repo rate minus the prevailing DFR. Define $\text{PostHike}_t = 1$ for transactions on or after the rate implementation date. The key estimating equation pools collateral-driven (CD) and funding-driven (FD) repos within dealer–client pairs:

$$\begin{aligned} r^{\text{adj}}_{d,c,i,t} = {} & \beta_1 \,\text{PostHike}_t \times LR^{\text{std}}_d \times \text{CD}_i + \beta_2 \,\text{PostHike}_t \times LR^{\text{std}}_d \\ & + \beta_3 \,\text{PostHike}_t \times \text{CD}_i + \beta_4 \,\text{PostHike}_t + \gamma \mathbf{X}_{d,c,i,t} + \alpha_{d,c} + \delta_{e,i} + \varepsilon_{d,c,i,t}, \end{aligned} \tag{2.17}$$

where LR_d^{std} is the dealer's leverage ratio standardized to mean zero and unit standard deviation, CD_i is an indicator for collateral-driven repos, $\alpha_{d,c}$ denotes dealer–client fixed effects, and $\delta_{e,i}$ denotes event–ISIN (or event–country) fixed effects. The controls $\mathbf{X}$ include log volume and maturity bucket dummies. The triple interaction β_1 identifies the *differential* effect of dealer leverage on rate adjustment in CD repos, netting out any common leverage effect in FD repos. The coefficient β_2 (on $\text{PostHike}_t \times LR_d^{\text{std}}$ without the CD interaction) serves as a placebo: it captures the leverage effect in funding-driven repos, where balance sheet constraints should be less binding because the transaction is motivated by cash needs rather than collateral demand.

Table 2.9. Passthrough of monetary policy

Notes: Transaction-level panel stacking ±7-day windows around GC rate decisions, pooling CD and FD repos. The dependent variable is the DFR-adjusted repo rate. PostHike $= 1$ on or after the rate decision. LR_{std} is the dealer's leverage ratio, standardised (mean 0, sd 1). The triple interaction PostHike $\times$ LR_{std} $\times$ CD identifies the differential rate adjustment by constrained dealers in collateral-driven repos, netting out common effects in funding-driven repos. The coefficient on PostHike $\times$ LR_{std} (without $\times CD$) captures the leverage effect in FD repos and serves as a placebo. Column (4) restricts to the 10 rate hikes, excluding the June 2024 rate cut. Controls: log volume, maturity bucket dummies. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

	All events			Hikes only
	(1)	(2)	(3)	(4)
PostHike × LR_{std} × CD	0.066***	0.066***	0.071	0.071***
	(0.022)	(0.017)	(0.045)	(0.022)
PostHike × LR_{std}	−0.004	−0.004	−0.004	−0.006
	(0.024)	(0.018)	(0.038)	(0.023)
PostHike × CD	−0.002	−0.002	−0.037	0.002
	(0.021)	(0.014)	(0.043)	(0.020)
PostHike	−0.025	−0.025	−0.001	−0.031
	(0.025)	(0.016)	(0.038)	(0.023)
Event × ISIN FE	✓	✓		✓
Event × Country FE			✓	
Dealer × Client FE	✓	✓	✓	✓
Cluster: Dealer	✓		✓	✓
Cluster: Dealer × Client		✓		
N	321,857	321,857	322,636	299,266
R^2	0.85	0.85	0.70	0.86

Table 2.9 reports the results. Across all four specifications, the triple interaction $\hat{\beta}_1$ is positive, ranging from 0.066 to 0.071, and statistically significant at the 1% level in three of four columns.[17] The estimate is stable whether standard errors are clustered at the dealer level (columns 1, 3, 4) or the dealer–client level (column 2), and whether the sample includes the June 2024 rate cut (columns 1–3) or is restricted to the ten hikes only (column 4). The placebo coefficient $\hat{\beta}_2$ is economically small and statistically insignificant across all columns, confirming that the leverage–pass-through nexus is specific to collateral-driven repos.

The magnitude is economically meaningful. Based on the preferred specification in column (1), a one-standard-deviation increase in the leverage ratio is associated with a 6.6 percentage point improvement in DFR pass-through for CD repos. This 7 percentage point pass-through differential is modest relative to the impairment documented by Eisenschmidt et al. (2024), whose estimates imply that dealer market power reduces OTC pass-through by 20 to 30 percentage points relative to the inter-dealer benchmark. The two channels are, however, complementary and operate at different margins. Market power acts *between* the inter-dealer and dealer-to-client segments, compressing pass-through for all OTC customers. Balance sheet constraints, by contrast, generate heterogeneity *within* the dealer sector: they distort rate adjustment specifically in collateral-driven trades, where—as documented in the preceding sections—constrained dealers substitute away from haircut-based risk mitigation toward rate-based adjustments. The finding that the placebo coefficient $\hat{\beta}_2$ is insignificant in FD repos further supports this interpretation: when the transaction is driven by funding needs rather than collateral demand, the dealer's leverage position is less relevant for rate setting.

[17] In column (3), which replaces event–ISIN with coarser event–country fixed effects, the point estimate is 0.071 with a standard error of 0.045, significant at the 12% level.

2.6 Conclusion

The primary reason to implement the Basel III framework in the aftermath of the global financial crisis was to make the financial world more resilient. After the implementation, it is now one of the most important tasks to understand how market participants are practically dealing with the new constraints. Balance sheet space is one of the most biting constraints for the repo market. Our paper is the first to show that the initial margin or haircut, which was initially designed to insure the lenders of cash or collateral from market volatility or counterparty default, is directly impacted by this regulation. At the face of it, the setting of initial margins in the Euro area repo markets displays strange patterns. Our model and empirical results rationalise that these are the consequence of three core frictions, namely the balance sheet space, the collateral heterogeneity (or scarcity) and the dealer market power. This has important consequences for market stability and monetary policy. The theoretical channel shows that if the haircut cannot adjust as freely as it is supposed to, more of the volatility has to be absorbed by the rate, which potentially also results in a weaker pass-through from central banks to the broader economy.

Chapter 3

Climate Change, Bank Liquidity and Systemic Risk

Co-authors: Margherita Giuzio and Bige Kahraman

Abstract This paper examines the relevance of banks' exposure to climate transition risk in the interbank lending market. Using transaction-level data on repo agreements, we first establish that banks with higher exposure to transition risk face significantly higher borrowing costs. This premium is a combination of a risk premium, compensating lenders for increased credit risk, and an inconvenience premium, reflecting the sustainability preferences of key dealer banks. We also find that the transition risk premium intensifies during periods of financial stress, indicating that climate-induced risks amplify existing vulnerabilities in financial markets. Furthermore, the rate segmentation caused by transition risk premium has implications for the transmission of monetary policy. Transition risk is an important factor in financial stability and policy design.

Keywords: climate finance, transition risk, repo markets, risk premium, financial stability

JEL Classification: Q54, G21, G32, Q58

For helpful comments and suggestions, we thank Anthony Limburg, Alba Patozi (discussant), Dimitrios Tsomocos and participants at the ESCB Climate Cluster Seminar, the University of Oxford FAME seminar, the CCC Green Seminar Series at the ECB, and the ECB Monetary Policy Strategy Division.

3.1 Introduction

Climate change presents a profound global risk with potential implications for financial markets. Financial market participants and policymakers are increasingly concerned that climate change related risks may disrupt financial stability and inflict negative economic consequences (Acharya et al., 2023). While a growing body of research has explored how climate risk affects asset prices and corporate financing, its impact on systemic risk remains less understood despite these concerns. As Acharya et al. (2023) highlight, climate change can increase market risk through multiple channels, potentially depressing property and corporate values, reducing corporate profits, and eroding household wealth. These effects, in turn, can exacerbate credit, market, and liquidity risks across the financial system.[1]

In this paper, we explore a critical, yet understudied, channel: whether and how climate transition risks —those arising from climate policies, regulations, or technological shocks—may threaten liquidity in the banking sector. We specifically investigate three key questions:

- Does increased exposure to transition risk affect a bank's refinancing needs and alter the pricing of short-term interbank loans, known as repos?
- Does exposure to transition risk amplify other exogenous shocks in financial markets, making banks with higher climate risk exposure more vulnerable during periods of financial stress?
- Does exposure to transition risk impact monetary policy transmission?

By addressing these questions, our research contributes to the academic literature on climate finance by shedding light on the interplay between **climate risk** and both **systemic risk** and **monetary policy transmission**. Our findings have implications for policymakers, as climate

[1] See Figure 3.A.1.

risk-induced frictions in the repo market could hinder the effectiveness of financial stability measures and monetary policy.

When the literature discusses climate risk, it typically considers two distinct types of risk relevant to financial markets: transition risk and physical risk. We focus on transition risk as it is more likely than physical risk to affect firms and financial institutions in the near future (Krueger et al., 2020; Acharya et al., 2023). Although physical risks, such as extreme weather events, have become more prevalent in recent years, their financial impact has been relatively contained due to ample insurance coverage and liquidity buffers. In contrast, recent research shows that transition risk is already materializing, with investors increasingly taking this risk into account when making portfolio choices (e.g., Kacperczyk and Bolton (2021); Bolton and Kacperczyk (2024)).

Our analysis employs high-frequency, trade-level data for the European repo market from 2019 to 2022. This market is the primary venue for managing funding and collateral for European banks, with daily transaction volumes ranging between EUR 500 and EUR 1,500 billion. It is also the main venue through which the European Central Bank implements its policy rate. Given this central role for financial intermediaries, any disruption in the functioning of the repo market could have serious implications for liquidity across the European financial system. Our data source is the Money Market Statistical Reporting (MMSR) dataset from the European Central Bank, which covers all transactions in the repo market by the 46 biggest European banks. For each transaction, we observe the trade date, maturity, volume, loan rate, haircut, collateral (identified by a unique ISIN), lender, borrower, and transaction type. Our analysis focuses on reverse repo transactions conducted by the 46 MMSR-reporting banks (which we refer to as "lenders") with 223 other banks in our sample (their cash-seeking trading partners, referred to as "borrowers").

We find that exposure to higher transition risk, proxied by banks' financed GHG emissions, is significantly linked to a premium in repo rates. Specifically, our baseline regressions show that a one standard deviation increase in transition risk exposure is associated with a 7–12% increase in repo rates, on average. This finding is particularly notable because the repo market is generally considered very safe due to its short-term nature and the fact that loans are collateralized. To ensure the robustness of this result, we control for a variety of deal- and bank-specific characteristics that are typically associated with risk factors in the repo market. Our findings remain stable even in extended tests where we also account for borrower as well as borrower-lender fixed effects. To further validate our findings, we conduct a natural experiment using exogenous shocks to transition risk. Specifically, we use the staggered increases in carbon prices that occurred when the EU reduced the supply of carbon allowances traded on its Emissions Trading System (ETS) market as exogenous shocks. We find both statistically and economically significant transition risk premium after the supply shock introduced in 2022.

Next, we aim to understand the source of the risk premium, and to this end, we consider three hypotheses. First, the transition risk premium may reflect a genuine risk premium, compensating lenders for the increased risk associated with carbon-intensive exposures. Second, it could be related to the preferences of dealer banks, particularly large European institutions that have committed to various climate initiatives. In this case, higher rates charged to carbon-intensive banks may reflect a loss of convenience or a misalignment with these banks' sustainability goals. Finally, the premium might be a result of increased funding demand from carbon-intensive firms, which their banks must support by raising funds in the money market. This greater funding demand could, in equilibrium, lead to higher borrowing costs.

Taken altogether, our findings suggest that the results can be understood as a combination of a risk and an inconvenience premium that reflects the preferences of dealer banks. First, we find that the premium is more pronounced in less-standardized transactions and when the collateral is generally considered riskier, such as in the non-government sector. Risk premium is associated with longer maturities and higher haircuts, but not with higher trading volumes. Risk premium is associated with higher *lagged* financed emissions, but not with contemporaneous ones. Finally, the premium becomes particularly noticeable after dealer banks joined voluntary climate commitments.

If repo rates are associated with banks' transition risk exposure, and the repo market is the core channel for both funding and monetary policy transmission, what are the implications for financial stability and monetary policy? Focusing first on financial stability question, we examine the pricing of the transition risk in the repo market during periods of heightened financial stress. Using various different measures of market stress, we find that transition risk premium significantly increases (three-fold) during periods of heightened financial stress. This finding further supports the earlier interpretation of a risk premium and highlights that the real dangers of climate change for financial stability may not be from direct exposure, which currently seems contained, but rather from its compounding effect with broader financial vulnerabilities. This finding underscores the need for policymakers to consider these interactions when designing new regulations.

For monetary policy, the key concern is whether interest rate changes set by the central bank are passed on quickly and fully to market participants, a process which is called the transmission of monetary policy. In reality, this is not necessarily the case, partly because the repo market is highly segmented and influenced by various factors that blur the direct effect of policy changes. When the level of repo rates is affected by transition risk exposure, does this mean that transmission is also impacted? Our data shows that during half of the rate hikes in our sample, borrowing rates for "brown" banks (those with high financed emissions) adjusted approximately 7% faster to new policy rates than those for "green" banks. This suggests that the transmission of monetary policy is quicker to brown banks during rate hikes, possibly because dealer banks are less willing to offer them favorable terms after a rate increase due to risk concerns or their own preferences.

3.1.1 Related literature

This study makes a two-fold contribution to the literature on climate finance and to policy discussions. First, to the best of our knowledge, this study provides the first empirical evidence on the dynamic relationship between **climate risks** and **bank liquidity**. In doing so, the analysis offers new insights into how climate risks can amplify existing financial vulnerabilities and affect the transmission of monetary policy, filling a key gap in the literature identified by Acharya et al. (2023). Second, our findings offer guidance for policymakers and central banks. They highlight the need for a new policy paradigm that integrates climate risk into systemic risk assessments, especially given the repo market's vital role in the financial system. Policies aimed at mitigating the climate risk's impact on interbank markets, along with enhanced climate risk disclosure and the implementation of robust stress testing frameworks, would help assessing and reducing transition-related systemic risk.

Much of the existing research examines how climate risks affect individual asset classes in isolation. Equity markets, in particular, have been extensively studied, with most findings suggesting a positive risk premium for carbon-intensive or "brown" stocks. For example, Kacperczyk and Bolton (2021, 2022); Faccini et al. (2023); Ramelli et al. (2021); Hong et al. (2019); Hsu et al. (2020); Barnett (2020) document that investors demand higher expected returns for holding brown stocks, reflecting compensation for regulatory, reputational, or transition risks. Engle et al. (2020) provide insight into how news about climate change could potentially be hedged in the equity market.

Similar patterns emerge in fixed income markets. Studies on corporate bonds suggest that firms with higher carbon footprints face higher borrowing costs, consistent with the idea that climate risks influence credit spreads and bond yields (Bua et al., 2024; Painter, 2020; Huynh and Xia, 2020; Goldsmith-Pinkham et al., 2019; Baker et al., 2018). Government bond markets also exhibit climate-related pricing effects, with sovereign debt instruments reflecting exposure to climate vulnerabilities or policy shifts (Collender et al., 2023).

Climate risks extend beyond equities and fixed income. In corporate loan markets, Altavilla et al. (2024) show that transition risk affects loan pricing, with lenders adjusting credit terms based on borrowers' climate-related exposures. Short-term options markets are also influenced, as Ilhan et al. (2020) find that options on firms with high carbon exposure tend to exhibit higher implied volatility, reflecting uncertainty about climate policies and stranded asset risks. Even very long-term assets, such as real estate (Foerster et al., 2025), are impacted. Research indicates that climate risks influence property valuations, with coastal properties exposed to sea level rise experiencing price discounts, while more resilient locations see price premia (Giglio et al., 2021; Bernstein et al., 2019; Bakkensen and Barrage, 2017; Baldauf et al., 2020).

Finally, our paper also contributes to the literature on repo markets. Several studies provide insights into key determinants of repo market dynamics. For example, Julliard et al. (2022) offer an in-depth analysis of the factors driving haircuts in the UK repo market, while Ballensiefen et al. (2023b) highlight the strong segmentation between collateral-driven and funding-driven repo markets. Eisenschmidt et al. (2024) demonstrate that oligopolistic effects play a crucial role in shaping interbank borrowing rates, and Duffie and Krishnamurthy (2016a) provide a foundational explanation of how the US repo market functions, including the role of liquidity demand and collateral quality.

3.2 Institutional background and hypothesis development

3.2.1 Climate risks

The finance literature typically categorises climate-related risks into two broad types: physical risks (e.g., climate disasters, rising sea levels) and transition risks (e.g., policy changes, market shifts, technological innovations). In its essence, transition risk refers to the financial risks associated with the process of shifting towards a low-carbon economy. [2]

Transition risk could manifest through multiple channels, including through (1) Regulatory risk: Stricter carbon pricing mechanisms (e.g., carbon taxes, cap-and-trade systems) increase operational costs for high-emitting firms; (2) Market risk: Shifts in investor sentiment and consumer preferences lower the valuation of carbon-intensive assets; (3) Technological risk: Disruptive clean technologies (e.g., electric vehicles, renewable energy) erode the market share of fossil-fuel-based industries; (4) Legal risk: Climate litigation against polluting firms raises liability costs and insurance risks.

Our analysis focuses on the regulatory risk channel of transition risk. This is motivated by two factors. First, regulatory risk has a direct and quantifiable impact on financial institutions and markets. Unlike market or technological risks, which evolve gradually, policy changes such as carbon taxes, cap-and-trade systems, and disclosure mandates can immediately alter firms' cost structures, profitability, and creditworthiness. This makes regulatory risk a more tangible and measurable potential driver of financial instability, particularly in the interbank lending market, where liquidity constraints and risk assessments are highly sensitive to sudden policy shifts. Secondly, large institutional investors see transition risk as the more immediate consequence of climate change for their portfolios. While they believe that physical risk will only really start to materialise in their portfolios in the long-term, the impact of transitioning to a net-zero economy is already felt today (Krueger et al., 2020).

While greenhouse gas emissions are a widely used proxy for transition risk, physical risk exposure is more difficult to measure. In the euro area, physical risk can to some extent be proxied by country or region, but it is important to consider how much of the risk is borne by firms and households versus insurance corporations and the public sector. This makes it more

[2] The Task Force on Climate-related Financial Disclosures (TCFD) defines transition risk as the potential financial losses that firms, investors, and financial institutions may face due to changes in policies, regulations, technologies, and consumer preferences aimed at mitigating climate change (Task Force on Climate-related Financial Disclosures, 2017).

complex to estimate the share of risk ultimately held by banks.

3.2.2 Carbon markets

The European Union Emissions Trading System (EU ETS) is the world's largest and most established carbon market, designed to help the EU meet its climate targets by reducing greenhouse gas (GHG) emissions in a cost-effective manner. It was established in 2005 as a central policy instrument to meet the EU's obligations under the Kyoto Protocol and has since been a key part of the region's strategy to achieve carbon neutrality by 2050.

The EU ETS operates as a cap-and-trade system, meaning that a maximum limit (or cap) is set on the total amount of GHG emissions allowed from covered sectors. This cap is progressively reduced over time, aligning with the EU's long-term climate goals. Companies within these sectors are allocated a certain number of emission allowances (EUAs), which represent the right to emit one ton of CO2 or its equivalent in other GHGs. If a company emits more than its allocated allowances, it must buy additional permits from others, incentivising emissions reductions. Conversely, firms that reduce emissions below their allocation can sell surplus allowances, creating a financial incentive to lower emissions.

The EU ETS covers emissions from over 11,000 power plants, industrial facilities, and airlines operating within the EU, as well as a growing number of sectors under its new expansion. The system initially focused on energy-intensive industries, such as power generation, cement, steel, and chemicals, but its coverage has expanded over time, now including aviation and potentially extending to maritime shipping and other sectors in the future. Collectively, the EU ETS accounts for around 40% of the EU's total GHG emissions.

A key reform to enhance the system's efficiency is the Market Stability Reserve (MSR), introduced in 2019. The MSR was designed to address the surplus of allowances in the market and ensure that carbon prices remain at levels that incentivize genuine emissions reductions. Around 57% of the allowances in the EU ETS are auctioned, with auctions held about three

Figure 3.1. EUA futures prices and MSR event dates

Notes: The graph shows the EU carbon price as measured by the price of the first EUA futures contract over the different phases of the EU ETS together with key events regarding the market stability reserve (MSR). Dotted lines indicate the announcement of each supply reduction, and solid lines indicate the actual start of the supply reduction.

times a week, while the remainder is allocated for free to sectors at risk of carbon leakage. If the surplus of allowances in the market exceeds a threshold of 833 million, excess allowances are removed from future auctions. This mechanism automatically sets aside unused allowances in the MSR, tightening supply and supporting higher carbon prices. As a result, the market price for allowances has nearly quadrupled over the observation period from 2019 to 2022.

3.2.3 The interbank lending market

A repurchase agreement (repo) is a short-term financial transaction in which one party sells securities to another with the agreement to repurchase them at a set price and date. Essentially, it functions as a secured loan, where the securities serve as collateral. Repos are widely used for short-term funding and liquidity management in financial markets.

The repo market plays a crucial role for financial institutions in sourcing funding or specific securities through short-term, secured transactions. With an average daily trading volume of approximately EUR 700 billion and an outstanding stock of EUR 2.15 trillion (European Central Bank, 2025), it is vital to the functioning of monetary policy, as it "ensures an even distribution of central bank liquidity and a homogeneous level of short-term interest rates" (European Central Bank, 2011).

For this paper it will be particularly important to understand *why* participants trade. There are roughly three reasons: (1) *funding-driven*: cash-borrowers seek short term funding, (2) *return-driven:* cash-lenders seek to collect a return on their excess liquidity, and (3) *collateral-driven:* cash-lenders seek to source specific asset. While a significant share of the transactions are initiated because of the latter two reasons, the secured market remains the predominant source of short-term funding and in this paper, we will therefore entirely focus on the funding-driven segment of the market. Throughout this paper, we will refer to the borrower as the cash-borrower and to the lender as the cash-lender on the first leg of the transaction.

Repos can be further sub-classified into general collateral (GC) repos and special collateral

(SC) repos. In a GC repo, the collateral is selected from a predefined basket of eligible securities, meaning the specific ISIN is not known to the cash lender in advance. These transactions are primarily funding-driven, focusing on liquidity rather than a particular security. In contrast, a SC repo involves a predefined ISIN as collateral, making it the preferred option for collateral-driven transactions. These transactions are typically used when a specific security is in demand, such as for short-covering or settlement purposes.

Repos can be traded through three main channels. The first is via a central clearing counterparty (CCP), which, under the European Markets Infrastructure Regulation (EMIR), acts as an intermediary by becoming the buyer to the repo seller and the seller to the repo buyer (open order). By assuming counterparty risk, the CCP also determines the haircut, making it an exogenous factor for the original contracting parties. The second method is trading through a tri-party agent, a financial intermediary that assists with administrative tasks and settlement but does not assume counterparty risk. Lastly, repos can be traded bilaterally in the OTC segment, where transactions occur directly between financial institutions without an intermediary. In the EU, the tri-party segment is relatively small, representing only 10% of the market. Around 70% of transactions in a dataset covering 46 major European commercial banks are cleared via a CCP, with 90% of these being SC repos, reflecting the demand for specific securities (European Central Bank, 2025). However, many financial institutions lack direct access to a CCP and are therefore restricted to trading in the OTC market

3.2.4 Testable hypotheses

We now list the hypotheses on if and how climate change impacts the repo market which we aim to assess in our empirical analysis. A substantial body of research indicates that transition risk associated with climate change permeates most markets studied to date. In most cases, firms' GHG emissions have been used as a proxy for transition risk and a price premium for securities of low GHG-intensive firms has been discovered (Kacperczyk and Bolton, 2021; Bua et al., 2024; Ilhan et al., 2020). On the other hand, we are looking at mostly short-term transactions in a market that is perceived as being very liquid and secure. If transition risk does not play a role in this market, the null hypothesis should find some support:

Hypothesis 1. *A bank's exposure to transition risk, as measured by financed GHG emissions, has no impact on the pricing of their funding-driven repo transactions.*

Transition risk undoubtedly impacts the risk profile, cash demand, and potentially the profitability of firms. This effect could spill over to banks with higher financed GHG emissions, leading to increased risk premiums in their refinancing transactions, or driving greater demand for borrowing, which they may source from the interbank market. If either of these scenarios holds true, we would expect to find support for the following alternative hypothesis.

Hypothesis 2. *A bank's exposure to transition risk, as measured by financed GHG emissions, leads to higher rates on its funding-driven repo transactions.*

3.3 Data

3.3.1 Data on European repo transactions

For our analysis we employ high-frequency, trade-level data for the European repo market for the time span from 2019 to 2022. Our data source is the Money Market Statistical Reporting (MMSR) data set from the European Central Bank, which covers all transactions in the repo market made by the 46 biggest European banks.[3] For each transaction we observe the trade date, the maturity, the trade volume, the loan rate, the haircut, the collateral identified by a unique ISIN, the lender, the borrower and the transaction type. Unless stated otherwise, our analysis focuses on *reverse* repo transactions conducted by the 46 MMSR-reporting banks with 223 banks in our sample, which act as cash borrowers in these transactions. For simplicity,

[3]For a detailed list of the reporting banks see European Central Bank (2025).

we refer to the 46 MMSR-reporting banks as *lenders* and their cash-seeking trading partners as *borrowers*.

As we are interested in cash-driven transactions, i.e., transactions which were initiated because the borrower needed cash, we only consider a part of the transactions from the MMSR dataset. First, we focus on bilateral transactions and exclude those cleared by Central Counterparties (CCPs) for three key reasons: (1) To assess counterparty risk, it is essential to observe the bank-counterparty pair directly; (2) CCPs often determine rates at the portfolio level, which can introduce noise that makes it difficult to attribute specific risk to collateral alone; (3) Many CCP transactions are primarily used by banks to manage collateral demands rather than cash demands, which reduces their relevance to our analysis.[4] Second, we are excluding special collateral transactions and all transactions made with haircuts below zero as both are most likely collateral-driven. Third, we focus on transactions between deposit-taking financial institutions for the simple reason that we cannot observe loan-portfolios for other market participants like money-market funds, hedge funds or pension funds. We end up with a sample of 2.3 million transactions, with an average transaction volume of EUR 7.8 million.

All rates, volumes and haircuts are winsorised at the 1% and 99% level. This was done to clean the dataset from flawed data points, but it has no impact on the qualitative nature of our findings. Figure 3.2 shows the maturity and collateral sector distribution in our sample. While most repos are typically collateralised by a government bond, our sample features also a significant share of transactions that is collateralised with corporate sector bonds.

We also note that our classification into funding- and collateral driven transactions may not be perfect for two reasons. One is that the transaction type classifier (whether it was a generic or a specific transaction) is voluntary and hence missing in many cases. The second reason is that even GC transactions are used for collateral purposes in some rare cases. However, including

[4] About 90% of the cleared transactions are SC repos which tend to be collateral-driven (European Central Bank, 2025).

Figure 3.2. Collateral sector and maturity distribution of transactions in our sample

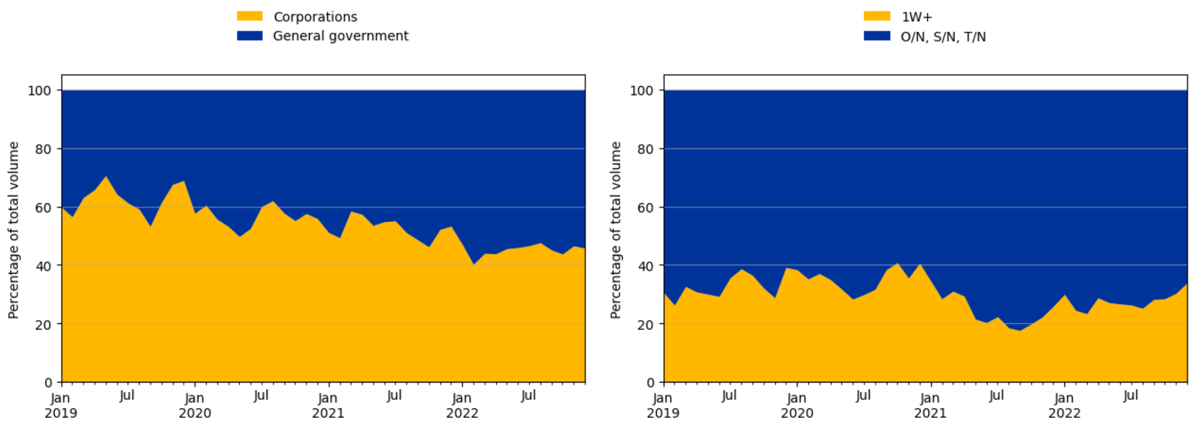

(A) Collateral sector distribution **(B)** Maturity distribution

Notes: The left graph illustrates the relative volume-weighted share of transactions collateralized by securities across the two primary collateral sectors in our sample. Meanwhile, the right graph displays the relative volume-weighted share of transactions within the two main maturity buckets in our sample.

collateral-driven transactions in our sample would only go against us as these transactions are neither reflective of the demand for cash, nor about the counterparty-risk of the cash-borrower. As a consequence, by including some of the collateral-driven transactions in our sample, we would be underestimating the true effect.

The average borrowing rate of our sample is plotted in Figure 3.3. It is very similar to the average GC rate which is a close proxy for the cost of money (Ballensiefen et al., 2023b).

3.3.2 Data on financed GHG emissions

Our primary proxy for transition risk will be the financed GHG emissions of each bank, calculated as follows. For each year and each bank, we examine all outstanding corporate loans with a value exceeding EUR 25,000.[5] These loans are paired with Scope 1 GHG intensity data

[5]These loans include credit card debt, credit lines, deposits (excluding reverse repurchase agreements), reverse repurchase agreements, trade receivables, loans, revolving credit (excluding overdrafts and credit card debt), overdrafts, and finance leases.

Figure 3.3. Average borrowing rate

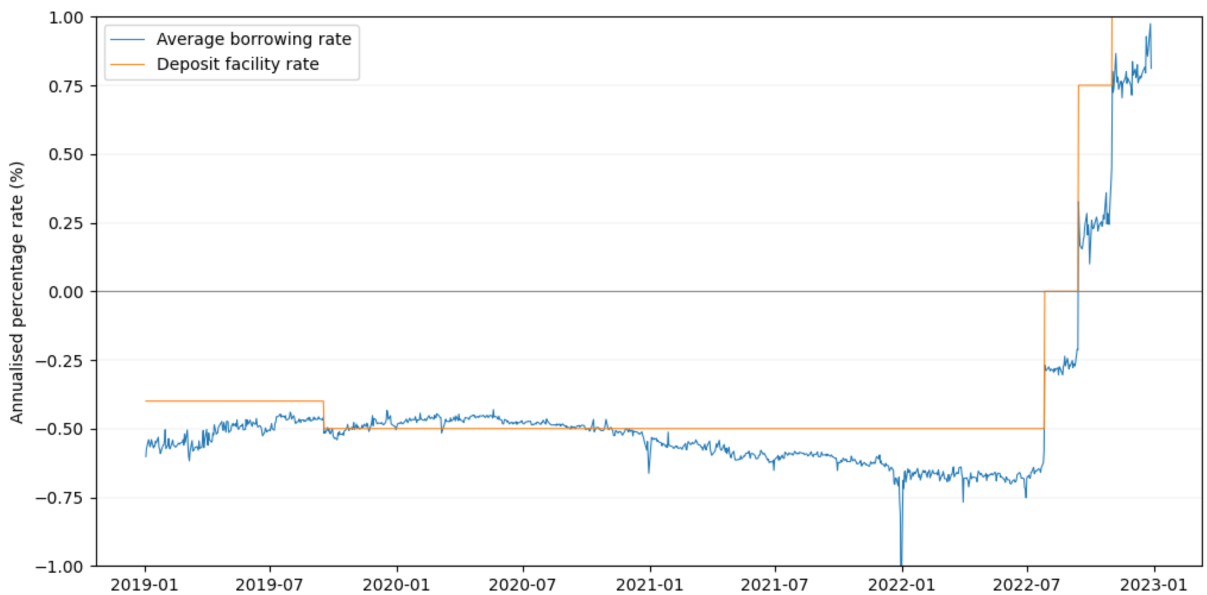

Notes: The graph shows the average daily borrowing rate and the ECB's deposit facility rate. The average daily borrowing rate is calculated as the volume weighted average rate of all borrowing transactions made by 223 counterparty banks in our sample with the 46 MMSR reporting banks.

from Urgentem. We use Scope 1 emissions as it is the most likely level subject to regulatory measures.[6] Additionally, we focus on intensity rather than absolute emissions, as firms with lower revenues relative to emissions will likely face greater regulatory burden. This approach aligns with the climate finance literature (e.g., Ilhan et al. (2020)).

Next, we calculate the volume-weighted Scope 1 intensity for each year and bank, then divide it by the bank's revenues. This method is motivated by the fact that banks with relatively low revenues are most likely to face challenges in accommodating the borrowing demand arising from transition risks.

Table 3.1. Summary statistics and correlation matrix

Notes: For all repo loan rates we show the difference of the actual loan rate minus the ECB's deposit facility rate (DFR). The collateral rating is a mapping from Moody's rating into integer numbers from 1 to 20, with 20 representing AAA-quality.

Panel A: Summary Statistics

Variable	N	Mean	Std. Dev.	Min	1%	25%	50%	75%	99%	Max
$Rate$	2,334,623	-0.53	0.83	-6.17	-6.10	-0.55	-0.35	-0.20	0.35	8.50
$FinancedEmissions$	2,334,623	0.03	0.02	0.00	0.00	0.02	0.02	0.03	0.10	0.43
$Haircut$	2,334,623	0.37	1.62	0.00	0.00	0.00	0.00	0.00	5.00	100.00
$Volume$	2,334,623	7.88	30.61	0.00	0.02	0.53	1.10	3.68	111.60	6,755.21
$Rating$	1,959,710	17.29	2.81	11.00	12.00	16.00	18.00	20.00	20.00	20.00
$Maturity$	2,334,623	2.88	16.61	1.00	1.00	1.00	1.00	1.00	30.00	365.00
$TotalAssets$	2,006,742	721.49	840.19	0.64	26.82	101.98	348.87	1,299.70	2,766.39	2,766.39
$LeverageRatio$	2,006,742	0.05	0.02	0.02	0.03	0.04	0.05	0.06	0.12	0.42
$LiquidityRatio$	2,006,742	24.74	4,787.70	1.01	1.20	1.36	1.49	1.71	9.28	999,999
$CapitalRatio$	2,006,742	0.17	0.04	0.06	0.13	0.15	0.17	0.18	0.38	1.37
ROA	2,006,742	0.01	0.01	-0.01	0.00	0.01	0.01	0.01	0.03	0.11

Panel B: Correlation Matrix

	$FinancedEmissions$	$TotalAssets$	$LeverageRatio$	$LiquidityRatio$	$CapitalRatio$	ROA
$FinancedEmissions$	1.00	-0.13	-0.02	0.00	-0.06	-0.12
$TotalAssets$	-0.13	1.00	-0.39	0.00	-0.45	0.06
$LeverageRatio$	-0.02	-0.39	1.00	0.01	0.46	0.06
$LiquidityRatio$	0.00	0.00	0.01	1.00	0.00	0.00
$CapitalRatio$	-0.06	-0.45	0.46	0.00	1.00	0.10
ROA	-0.12	0.06	0.06	0.00	0.10	1.00

Table 3.1 summarises the key variables used in the following empirical analysis.

[6] The aggregate time series is published in European Central Bank (2024).

Before turning to the regression analysis, we verify that banks with high and low financed emissions are comparable along key dimensions. Table 3.2 reports mean values for the bottom and top terciles of financed emissions. The groups are statistically indistinguishable on all risk measures (leverage, liquidity, capital ratios, and profitability) as well as on deal characteristics such as haircuts, volumes, and maturities. The only marginally significant difference is in total assets, which we control for throughout our analysis.

Table 3.2. Bank characteristics across emissions terciles

Notes: The table reports mean values of bank-level characteristics for borrowers in the bottom and top terciles of average financed GHG emissions, computed at the borrower level over the full sample. Each bank enters once, with variables averaged across its observations. The Difference column reports the top-minus-bottom mean gap. Statistical significance is based on a two-sided Welch t-test of equal means (unequal variances). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Variable	Low emissions	High emissions	Difference
Total assets (EUR bn)	50.00	108.25	58.25*
Leverage ratio (%)	7.86	7.64	−0.22
Liquidity coverage ratio	2.79	2.10	−0.69
Capital ratio (%)	23.70	20.21	−3.48
ROA (%)	1.08	1.02	−0.06
Haircut (%)	1.12	1.76	0.64
Transaction volume (EUR bn)	0.04	0.04	0.00
Deal maturity (days)	34.45	37.14	2.69

3.4 The climate premium

In this section, we establish the existence of a strong correlation between repo rates and financed GHG emissions. We then proceed to investigate *why* this carbon premium arises. A common explanation in the climate finance literature is that it reflects compensation for risk, i.e., a risk premium. We assess this hypothesis against alternative explanations based on investor preferences or demand shocks, using additional panel regressions and identification strategies.

3.4.1 Panel regressions of interbank borrowing rates on financed emissions

We begin by presenting our main findings on the pricing of emissions in the European repo market. We first report results for the full set of available collateral assets, and then examine how the carbon premium varies across different collateral classes. At a high level, Figure 3.4 illustrates the distribution of financed GHG emissions and repo rates, even before controlling for repo-specific characteristics or borrower fixed effects. Across all collateral sectors, the figure suggests that banks with lower financed GHG emissions tend to receive more favourable rates.

Figure 3.4. Financed emissions and average borrowing rates

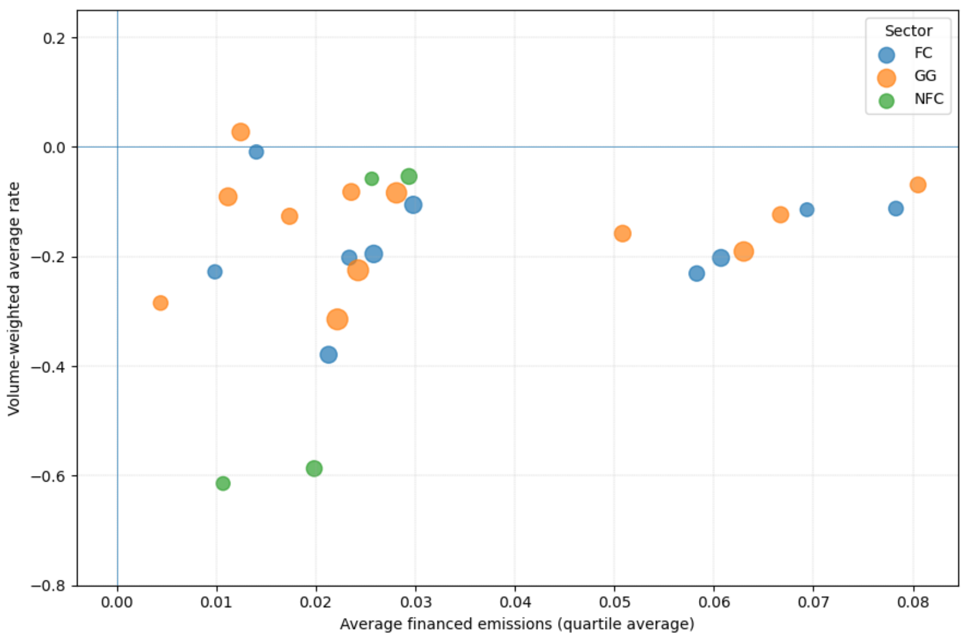

(A) Distribution across collateral types

Notes: The chart shows the volume-weighted yearly average loan rate in cash-borrowing transactions of 223 banks against their financed GHG-emissions, subdivided by collateral issuer sector. Financed GHG emissions are calculated as the volume-weighted GHG-intensities per bank-year from all commercial loans over EUR 25k, divided by the bank's revenues and lagged by one year. Only O/N, S/N and T/N transactions are included.

Our main analysis of the carbon premium relies on two fixed-effects specifications. The first controls for all time-series variation in rates, allowing us to focus on cross-sectional heterogeneity in banks' financed emissions. The second controls for bank-specific heterogeneity, thereby emphasizing time-series variation. Given that our sample spans only four years, there is substantially more variation across banks than over time. As a result, the second specification generally yields weaker, though still largely significant results.

Even though most studies focus on general firms rather than banks (e.g., Kacperczyk and Bolton (2021, 2022)), several characteristics are particularly relevant when using carbon emissions as the main dependent variable. These include total assets or bank size, risk measures such as the liquidity and leverage ratios, capital ratios, and return on assets as a proxy for profitability. We also note considerable cross-sectional heterogeneity in repo rates, which largely depends on deal-specific characteristics such as the ISIN used as collateral, maturity, volume, haircut, and dealer-specific factors.

We begin by linking banks' borrowing rates to their corresponding financed emissions intensities and other characteristics, all lagged by one year. Specifically, we estimate the following specification:

$$\begin{aligned} Rate_{l,b,t,i} - DFR_t = {} & \beta_1 \times FinancedEmissions_{b,t-365} \\ & + \beta_2 \times DealControls_i + \beta_3 \times BorrowerControls_{b,t} + \delta_t + \gamma_{l,b} + \varepsilon_{l,b,t,i}, \end{aligned} \tag{3.1}$$

where δ_t and $\gamma_{l,b}$ are time and borrower-lender fixed effects. The index i represents the transaction level on which this regression is executed which – strictly speaking – makes the indices l for lender, b for borrower and t for the specific day obsolete. As deal controls, we include the collateral ISIN, the transaction maturity, volume and haircut. As borrower controls, we include total assets to proxy for the borrower's size, the return on assets to proxy for its profitability and the leverage ratio, the liquidity ratio as well as the capital ratio to proxy for the riskiness of the borrower. In all specifications, standard errors are clustered at the lender-borrower level, as we expect errors to be correlated at this level due to the strong impact of relationships within this market (Eisenschmidt et al., 2024). The results are reported in Table 3.3.

Table 3.3. Transition risk premia

Notes: The dependent variable is the loan rate minus the ECB's deposit facility rate expressed in percentage points. Financed GHG emissions are calculated as the volume-weighted GHG-intensities per bank-year from all commercial loans over EUR 25,000, divided by the bank's revenues. All continuous variables are standardised. Standard errors are clustered at the bank-counterparty-month-level and reported in parentheses. Controls include the collateral ISIN, total assets, the leverage ratio, the liquidity coverage ratio, the capital ratio, and the return on assets. $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

	DFR-adjusted rate (%)			
	(1)	(2)	(3)	(4)
$FinancedEmissions$	0.066***	0.058***	0.031***	0.048***
	(0.006)	(0.005)	(0.004)	(0.006)
$Haircut$		−0.148***	−0.054***	0.006
		(0.017)	(0.009)	(0.007)
$Volume$		0.104***	0.075***	0.051***
		(0.005)	(0.004)	(0.003)
$Maturity$		0.003***	0.002***	0.002***
		(0.000)	(0.000)	(0.000)
Controls	✓	✓	✓	✓
FE: Date			✓	
FE: Lender			✓	
FE: Lender × Borrower				✓
N	2 006 742	2 006 742	2 006 742	2 006 742
R^2	0.16	0.33	0.43	0.50

Across all specifications, we find a positive and significant effect of financed GHG emissions on a bank's interbank borrowing rate. Most importantly, this effect cannot be explained by individual deal characteristics or lender-borrower relationships—both of which are known to be key determinants of repo rates. Isolating the time-series variation in financed GHG emissions yields larger positive coefficients than focusing on cross-sectional differences. This may reflect the growing salience of transition risk over the observation period.

The presence of a carbon premium in generally safe and short-term loans is both novel and surprising. Taken at face value, our results suggest that a one standard deviation increase in a bank's financed GHG emissions is associated with an average increase of 3–5 basis points in its borrowing rate. To put this into perspective, the overall standard deviation of adjusted rates in our sample is 83 basis points. However, this figure includes substantial heterogeneity driven by deal-specific characteristics. After controlling for such factors, the standard deviation of the residual loan rates is 0.43, implying that a one standard deviation increase in financed emissions is associated with a rate increase equivalent to approximately 7–12% of this residual variation.

As the scatter plot in Figure 3.4 indicates, there may be significant differences between the collateral sectors. Collateral management is highly relevant for banks and one of the key reasons why repos are so often used next to its funding purpose. We repeat the above equation (3.1) for each collateral class separately. The results are reported in Table 3.4.

The differences between collateral classes are economically meaningful and, depending on the specification, range from around 4 basis points for transactions collateralised by government bonds and financial corporations, to up to 21 basis points in the case of non-financial corporate collateral. Two potential factors may explain this variation. First, repos backed by general government bonds tend to be more standardised and liquid instruments compared to those involving non-financial corporate bonds. As a result, their contractual terms are less flexible, and climate-related risks may be less directly reflected in pricing. Second, transactions involving general government collateral are more commonly used for collateral management purposes, rather than strictly for funding. In such cases, the pricing of borrower-specific risks (such as financed emissions exposure) may be less relevant, which could weaken the observed relationship in this segment. We also find that the results become stronger once lender-borrower fixed effects are included. This specification isolates time variation within individual lending relationships, suggesting that dealers may have increasingly priced in climate transition risks over time. One possible explanation is that awareness and salience of climate-related financial risks grew between 2019 and 2022, leading to a gradual incorporation of such premia

Table 3.4. Transition risk premia by collateral sector

Notes: The dependent variable is the loan rate minus the ECB's deposit facility rate expressed in percentage points. Financed GHG emissions are calculated as the volume-weighted GHG-intensities per bank-year from all commercial loans over EUR 25,000, divided by the bank's revenues. Regressions are based on the subset of transactions from three different collateral sectors: Non-financial corporations (NFC), Financial corporations (FC), and General government (GG). All continuous variables are standardised. Controls include the collateral ISIN, total assets, the leverage ratio, the liquidity coverage ratio, the capital ratio, and the return on assets. Standard errors are clustered at the bank-counterparty-month level. $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

	DFR-adjusted rate (%)					
	NFC		**FC**		**GG**	
	(1)	(2)	(3)	(4)	(5)	(6)
$FinancedEmissions$	0.055***	0.209***	0.025***	0.039***	−0.004	0.043***
	(0.012)	(0.029)	(0.005)	(0.006)	(0.004)	(0.005)
$Haircut$	0.015*	0.025***	−0.005	0.003	−0.013***	−0.022***
	(0.009)	(0.009)	(0.004)	(0.002)	(0.004)	(0.004)
$Volume$	0.010	0.021***	0.091***	0.080***	0.029***	0.025***
	(0.006)	(0.005)	(0.007)	(0.006)	(0.003)	(0.003)
$Maturity$	0.005***	0.000	0.003***	0.002***	0.001***	0.002***
	(0.001)	(0.001)	(0.000)	(0.000)	(0.000)	(0.000)
Controls	✓	✓	✓	✓	✓	✓
FE: Lender	✓		✓		✓	
FE: Date	✓		✓		✓	
FE: Lender × Borrower		✓		✓		✓
N	466 410	466 410	1 095 577	1 095 577	434 832	434 832
R^2	0.50	0.54	0.31	0.45	0.49	0.54

into repo rates.

3.4.2 Mechanism analysis

The panel regression in the previous section provides suggestive evidence that transition risk leads to higher borrowing rates for affected banks. But we have not identified the transition risk yet and cannot say anything yet about the origin of this premium. There are three main hypotheses as to where this premium comes from.

The first one is that it reflects *credit risk*. If firms with higher emissions get hit by climate policies, for instance by rising carbon prices, then these firms face higher thresholds for profitability and thus become riskier for their lenders. The risk that these lenders carry if they are highly exposed to the brown sector, could be internalised into the lending rates by other peer banks in the interbank lending market.

The second hypothesis is that the observed effect reflects an *equilibrium adjustment* in rates driven by increased demand for cash. When carbon-intensive (or "brown") firms are impacted by regulation, they face higher production costs due to new climate policies. To cover these costs, they may draw more heavily on their credit lines. Banks that lend to these firms therefore experience a surge in credit demand, prompting them to seek additional liquidity in the interbank repo market. As a result, banks with greater exposure to carbon-intensive borrowers face increased refinancing needs, which raises their demand for cash in the interbank market. This heightened demand, in turn, drives up the price of liquidity, reflected in higher repo rates.

A third hypothesis relates to *institutional preferences*. Most dealer banks in the euro area have joined at least one major sustainability initiative, such as the Net-Zero Banking Alliance (NZBA) or the Science Based Targets initiative (SBTi).[7] These frameworks typically involve commitments to reducing financed emissions or excluding carbon-intensive sectors from lending portfolios. Joining such initiatives not only reflects a strategic orientation towards climate-aligned finance but also signals the bank's broader preferences, whether driven by regulation, stakeholder pressure, or internal values. As a result, it is plausible that dealer banks with these commitments may have been less willing to extend funding to counterparties with high exposure to carbon-intensive sectors.

In the following sections we will look at various different variables and events to differentiate between these three hypotheses and eventually answer the question where this premium

[7] The NZBA and SBTi are voluntary initiatives through which financial institutions commit to aligning their lending and investment portfolios with net-zero greenhouse gas emissions by 2050 (NZBA), or to setting emission reduction targets consistent with climate science and the goals of the Paris Agreement (SBTi).

really comes from.

3.4.2.1 Maturities

Even in this ultra short-term market, maturities matter and are priced in, as evidenced by the consistently positive and significant coefficients on maturity across all panel regressions in subsection 3.4.1. Longer maturities amplify credit risk: the longer a bank is exposed to its counterparty through a loan contract, the more relevant the assessment of counterparty risk becomes. This finding is particularly relevant for our second hypothesis. If the observed rate premium were driven primarily by a demand-side explanation—for example, a surge in repo demand due to climate transition pressures—then maturity should not systematically affect pricing. In such a case, the resulting rate adjustment would affect transactions uniformly, regardless of their term. A reasonable null hypothesis would therefore be to expect no amplification effect. We are testing this by estimating a modified version of equation (3.1):

$$
\begin{aligned}
Rate_{l,b,t,i} - DFR_t = {} & \beta_1 \times FinancedEmissions_{b,t-365} \\
& + \beta_2 \times LongMaturity + \beta_3 \times LongMaturity \times FinancedEmissions_{b,t-365} \\
& + \beta_4 \times DealControls_i + \beta_5 \times BorrowerControls_{b,t} + \delta_t + \gamma_{l,b} + \varepsilon_{l,b,t,i},
\end{aligned}
\tag{3.2}
$$

where δ_t and $\gamma_{l,b}$ are time and borrower-lender fixed effects. The variable $LongMaturity$ is a dummy variable assuming the value 1 whenever the maturity of the repo was longer than one month. As deal controls, we include the collateral ISIN, volume and haircut. As borrower controls, we include total assets to proxy for the borrower's size, the return on assets to proxy for its profitability and the leverage ratio, the liquidity ratio as well as the capital ratio to proxy for the riskiness of the borrower. In all specifications, standard errors are clustered at the lender-borrower level. The results are reported in Table 3.5.

Table 3.5. Transition risk premia and longer maturities

Notes: The dependent variable is the loan rate minus the ECB's deposit facility rate expressed in percentage points. The variable $LongMaturity$ is a dummy indicating if the repo has a maturity greater than one month. Financed GHG emissions are calculated as the volume-weighted GHG-intensities per bank-year from all commercial loans over EUR 25,000, divided by the bank's revenues. All continuous variables are standardised. Standard errors are clustered at the bank-counterparty-month-level and reported in parentheses. Controls include the collateral ISIN, total assets, the leverage ratio, the liquidity coverage ratio, the capital ratio, and the return on assets. $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

	DFR-adjusted Rate (%)			
	(1)	(2)	(3)	(4)
$FinancedEmissions$	0.064***	0.057***	0.030***	0.048***
	(0.006)	(0.005)	(0.004)	(0.006)
$LongMaturity$	0.526***	0.352***	0.343***	0.156***
	(0.033)	(0.039)	(0.030)	(0.022)
$FinancedEmissions \times LongMaturity$	0.061**	0.054**	0.065***	−0.010
	(0.025)	(0.026)	(0.022)	(0.014)
$Haircut$		−0.146***	−0.052***	0.008
		(0.017)	(0.009)	(0.007)
$Volume$		0.107***	0.078***	0.055***
		(0.005)	(0.004)	(0.003)
Controls	✓	✓	✓	✓
FE: Date			✓	
FE: Lender			✓	
FE: Lender × Borrower				✓
N	2 006 742	2 006 742	2 006 742	2 006 742
R^2	0.16	0.32	0.43	0.50

While the coefficient on $FinancedEmissions$ remains largely unchanged relative to the baseline reported in Table 3.3, and the coefficient on $LongMaturity$ is positive and significant as expected, the interaction term between the two depends on the specification. In the cross-sectional analysis, there is a strong additional premium of approximately 7 basis points for longer-term loans issued to borrowers with high financed emissions. However, in the time-series specification the interaction term is insignificant and close to zero. This suggests that within stable lending relationships, longer maturities do not amplify the pricing of transition risk, whereas such interactions do matter across borrowers. One possible explanation is that within lender-borrower pairs, repo contracts tend to become more stable over time, partly due to operational convenience, and partly because lenders develop a better understanding of their counterparties. Nevertheless, the fact that longer maturities are associated with higher rates in the cross-section, and that they interact positively with financed emissions, is consistent with a credit risk channel. It may also indicate that dealers exhibit a preference for lower-emission counterparties, particularly when their exposure spans longer durations.

3.4.2.2 Haircuts

When investigating credit risk, the haircut is another natural variable to look at. Julliard et al. (2022) show that counterparties matter for haircut determination: riskier clients such as hedge funds are charged significantly higher haircuts, whereas larger borrowers with higher ratings receive lower haircuts. One important caveat is that haircuts in repo markets change much less frequently than rates and are set to zero in more than 80% of the transactions in our sample. Nevertheless, we still see that haircuts spike around important macroeconomic events such as the onset of the Covid crisis in 2020, indicating that it still captures some degree of credit risk. That there is a premium for additional credit risk originating from transition risk in these ultra-short term and secured loans is surprising so a reasonable null hypothesis would be for the financed emissions to not have any effect on the haircut.

We are testing this hypothesis in two ways. First, we use the haircut itself as the dependent variable in the following regression:

$$\begin{aligned} Haircut_{l,b,t,i} = \beta_1 \times FinancedEmissions_{b,t-365} \\ + \beta_1 \times DealControls_i + \beta_5 \times BorrowerControls_{b,t} + \delta_t + \gamma_{l,b} + \varepsilon_{l,b,t,i}. \end{aligned} \tag{3.3}$$

Second, because of the limited variation in our dataset, we are using a dummy indicating whether the haircut has been strictly positive with the regression

$$D_{l,b,t,i}^{Haircut>0} = \beta_1 \times FinancedEmissions_{b,t-365} \tag{3.4}$$
$$+ \beta_1 \times DealControls_i + \beta_5 \times BorrowerControls_{b,t} + \delta_t + \gamma_{l,b} + \varepsilon_{l,b,t,i}.$$

In both cases δ_t and $\gamma_{l,b}$ are used as time and borrower-lender fixed effects. As deal controls, we include the collateral ISIN, volume and the rate. As borrower controls, we include total assets to proxy for the borrower's size, the return on assets to proxy for its profitability and the leverage ratio, the liquidity ratio as well as the capital ratio to proxy for the riskiness of the borrower. In all specifications, standard errors are clustered at the lender-borrower level. The results are reported in Table 3.6.

In the first specification, we observe a pattern consistent with earlier regressions. While the effect is present and significant in the cross-sectional analysis, the coefficient diminishes once we focus on within lender–borrower pair variation – likely for the same reasons discussed above.

The coefficients for $FinancedEmissions$ in the second specification are positive and statistically significant. They suggest that a one standard deviation increase in a borrower's financed emissions raises the probability of receiving a positive haircut by approximately 1%. While this may seem modest, it should be viewed in the context of haircuts being generally rigid and often zero.

3.4.2.3 Identification via ETS-prices

In this section, we aim to identify transition risk through shocks to carbon prices. The core idea is that an unexpected increase in the price of emissions allowances represents a direct realization of transition risk. For carbon-intensive ("brown") firms, such shocks raise production costs. Given nominal rigidities, these cost increases are likely to compress short-term profits,

Table 3.6. Transition risk impact on haircuts

Notes: The dependent variable in the left two regressions is the transaction haircut, and a dummy indicator in the two right hand side regressions indicating whether or not the transaction has had a positive haircut. Financed GHG emissions are calculated as the volume-weighted GHG-intensities per bank-year from all commercial loans over EUR 25,000, divided by the bank's revenues. All continuous variables are standardised. Standard errors are clustered at the bank-counterparty-month-level and reported in parentheses. Controls include the collateral ISIN, total assets, the leverage ratio, the liquidity coverage ratio, the capital ratio, and the return on assets. $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

	Haircut (%)		$D^{Haircut>0}$	
	(1)	(2)	(3)	(4)
$FinancedEmissions$	0.043***	−0.012	0.008***	0.006***
	(0.011)	(0.008)	(0.001)	(0.002)
$Rate$	−0.087***	0.006	−0.020***	0.004***
	(0.014)	(0.006)	(0.002)	(0.001)
$Volume$	0.043***	0.046***	0.011***	0.012***
	(0.016)	(0.014)	(0.002)	(0.002)
$Maturity$	0.005***	0.005***	0.001***	0.001***
	(0.001)	(0.001)	(0.000)	(0.000)
Controls	✓	✓	✓	✓
FE: Date	✓		✓	
FE: Lender	✓		✓	
FE: Lender × Borrower		✓		✓
N	2 006 742	2 006 742	2 006 742	2 006 742
R^2	0.41	0.70	0.74	0.84

which can in turn increase funding needs and, in more severe cases, contribute to financial distress if sufficient external financing is not secured.

If we can use carbon price shocks as an instrument and demonstrate that they are associated with rising repo rates, this would provide compelling evidence that credit risk or demand-side channels are active in this market.

One potential concern, however, is that MSR-related supply reductions (i.e., cuts in the Market Stability Reserve) are announced well in advance, typically in May or June before being implemented on September 1st. As a result, these events may not constitute truly "unexpected" shocks. Yet for our purposes, what matters is not the surprise element of the supply announcement, but rather the actual shock to carbon prices. While the reduction in allowance supply is anticipated, the market's reaction in the form of price changes is an equilibrium outcome that unfolds around the implementation date. Due to the ultra-short-term nature of the repo market, much of this price effect is revealed only after September 1st. That said, firms anticipating tighter supply may warehouse allowances in advance, which could mute the price adjustment on the implementation date and potentially weaken the predictive power of our identification strategy.

Another potential concern with this approach is serial correlation. If we would be using transaction level data, the number of observations per treated unit are much more granular than the level on which the intervention takes place. As a consequence, results will be highly significant (and indeed positive) but no longer trustworthy (Angrist and Pischke, 2004). Hence, we are doing two things to reduce this problem: The first is that we don't do this estimation on a transaction level, but rather average transactions by counterparty-week. As there is important heterogeneity in repo contracts that we need to control for, we are first stripping repo rates off fundamental factors that are known to influence repo rates such as the collateral ISIN, the maturity, and the haircut. To fix ideas, we are computing residual repo rates as in

$$Rate_{l,b,i,t} - DFR_t = \beta_1 Haircut_{i,t} + \beta_2 Maturity_{i,t} + \beta_3 CollISIN_{i,t} + \delta_t + \varepsilon_{l,b,i,t}, \tag{3.5}$$

where δ_t are month-year fixed effects. We average these residuals as in

$$Rate^{resid}_{b,w} \equiv \frac{\sum_{i \in \mathcal{T}_{b,w}} LoanVolume_{i,t} \times \hat{\varepsilon}_{l,b,i,t}}{\sum_{i \in \mathcal{T}_{b,w}} LoanVolume_{i,t}}, \tag{3.6}$$

where $\mathcal{T}_{b,w}$ denotes all borrowing transactions made by counterparty b in week w.

The second important step is to cluster the standard errors at the counterparty level to account for errors being correlated across time for each borrower. For each year in our sample for which we have MSR-related supply reductions, we are estimating the following specification

$$Rate^{resid}_{b,w} = \beta^{IV} D^{IV}_{w} \times FinancedEmissions_{w-52,b} + X_{b,w} + \delta_w + \gamma_b + \varepsilon_{b,w}, \tag{3.7}$$

where D^{IV}_w is a dummy taking the value 1 for all weeks after the supply reduction in a given year. The results are reported in Table 3.7.

Table 3.7. Identification via ETS-supply shocks

Notes: The dependent variable is the loan rate residual as defined in equation 3.6. Financed GHG emissions are calculated as the volume-weighted GHG-intensities per bank-year from all commercial loans over EUR 25,000, divided by the bank's revenues. All continuous variables are standardised. Standard errors are clustered at the borrower-level and reported in parentheses. Controls include total assets, the leverage ratio, the liquidity coverage ratio, the capital ratio, and the return on assets. $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

	$Rate^{resid}$ (%)			
	2019	2020	2021	2022
$Post \times FinancedEmissions$	0.022	0.000	0.000	0.070**
	(0.018)	(0.010)	(0.007)	(0.030)
Controls	✓	✓	✓	✓
FE: Counterparty	✓	✓	✓	✓
FE: Week	✓	✓	✓	✓
Num.Obs.	1343	1139	1321	1439
R2	0.605	0.677	0.713	0.788

In three out of four cases, we reject the claim that MSR supply reductions directly increased repo rates for banks more heavily exposed to the brown sector. The year 2022 stands out: the MSR supply reduction can be causally linked to a 7 basis point increase in repo rates over 12 weeks. Why was this intervention different from the others? There are several potential reasons. First, climate awareness has grown over the years, so while earlier transactions did not carry a climate premium, later interventions may have reflected stronger premia. Second, energy prices soared in the second half of 2022 due to the war in Ukraine, meaning that high emitting companies already facing elevated input costs were hit even harder by higher carbon prices. Third, carbon prices themselves increased substantially over the years and were gen-

Figure 3.5. Estimated coefficients of the 2022 MSR treatment indicator

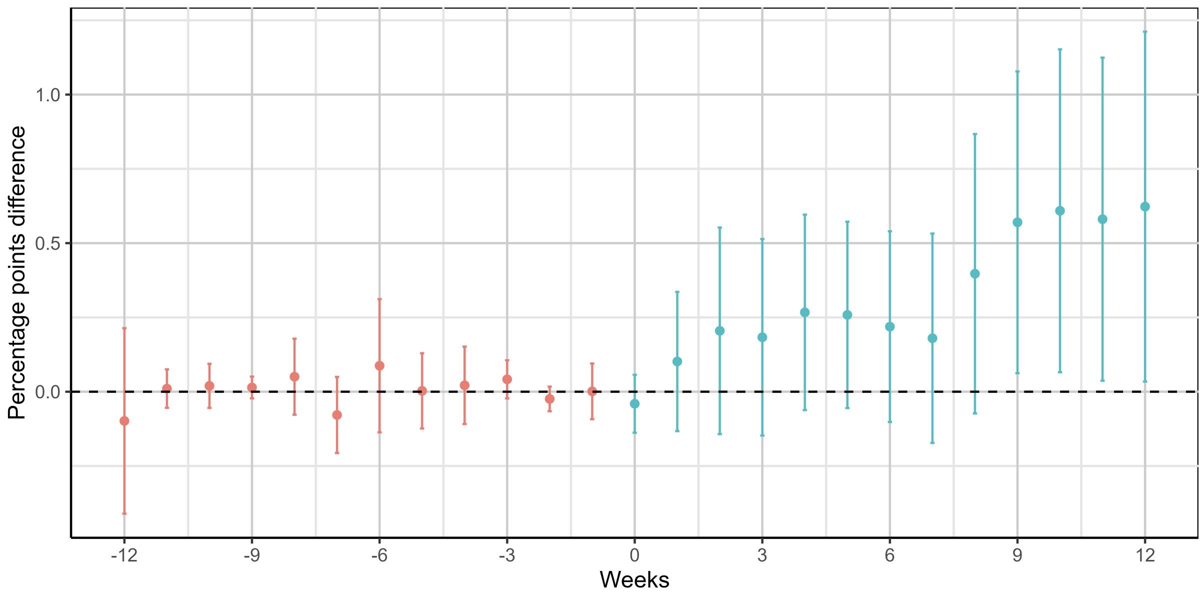

Notes: The graph plots the estimated coefficients on interactions of the treatment indicator variable with year-week fixed effects. We drop the interaction for the last week of August in 2022 (week 35, which started on Monday, August 29, 2022, and ended on Sunday, September 4, 2022) and thus the effect is normalized to zero for that week. We control for deal maturity, haircut, collateral ISIN, and bank-counterparty relationships. Standard errors are clustered at the borrower-level.

erally much higher in 2022 than before. Firms that might have coped with 20 Euros per ton of carbon previously faced prices up to 90 Euros per ton in 2022, which likely posed a different challenge. Fourth, successive interest rate hikes in the second half of 2022 driven by rising prices may confound our estimates. This is why we used adjusted rates and included month-year fixed effects when calculating residuals. The DiD-plot in Figure 3.5 for this year shows almost no pre-trend and a sharp incline right after the introduction which alleviates this concern. Lastly, warehousing carbon allowances or the expectation of supply reductions may have erased our effects in the first three years.

3.4.2.4 Voluntary climate commitments

A third explanation could be that the premium is due to preferences, i.e., banks could have a preference for borrowers with lower financed GHG-emissions, and hence they charge lower rates for green banks, or higher rates to brown banks to compensate themselves for the loss in convenience.

Often introduced due to stakeholder pressure, voluntary climate commitments offer insight into each bank's preferences. To the extent that banks genuinely follow through on these commitments — and there is ongoing debate about whether they do[8] — we can examine whether banks' preferences influence pricing in repo markets.

There are two main types of climate commitments we could use in our analysis: SBTi and NZBA. The SBTi distinguishes between the date of commitment, when a bank formally requests a decarbonization timeline, and the date when concrete targets are actually set and implementation can begin. While some dealer banks in our sample do make SBTi commitments, most had not established specific targets by the end of the sample period. Therefore, we should not expect observable changes in behavior based solely on initial SBTi commitments.

In contrast, NZBA includes a signing date that is directly linked to specific decarbonization targets. Most dealer banks in our sample committed to NZBA during the sample period,[9] which is the key reason we rely on NZBA commitments rather than SBTi to conduct our empirical tests.

[8]Kacperczyk and Peydró (2022) and Ye (2023), for instance, find that banks reduce exposure to high emitters or coal-related sectors in syndicated loans following voluntary climate commitments. Additionally, Altavilla et al. (2024), Delis et al. (2018), and Degryse et al. (2023) find that banks charge higher interest rates to polluting firms, suggesting a risk-based pricing response aligned with climate goals. On the other hand, Giannetti et al. (2025) and Sastry et al. (2024) show that banks highlighting the sustainability of their lending practices in public disclosures do not necessarily reduce their environmental impact and in fact, may extend more credit to brown borrowers.

[9]A striking number of banks have since exited the NZBA, but this occurred only after the end of our sample period.

We specify our test in the following form:

$$Rate_{l,b,t,i} - DFR_t = \beta_1 \times FinancedEmissions_{b,t-365} + \beta_2 \times VCC_{b,t} + \beta_3 \times FinancedEmissions_{b,t-365} \times VCC_{b,t} + \beta_4 \times DealControls_i + \beta_5 \times BorrowerControls_{b,t} + \delta_t + \gamma_{l,b} + \varepsilon_{l,b,t,i}, \quad (3.8)$$

where δ_t and $\gamma_{l,b}$ are time and borrower-lender fixed effects. The variable VCC is a dummy variable assuming the value 1 whenever the transaction day of the specific repo is on or after the day on which the bank committed to their Specific NZBA targets. As deal controls, we include the collateral ISIN, volume and haircut. As borrower controls, we include total assets to proxy for the borrower's size, the return on assets to proxy for its profitability and the leverage ratio, the liquidity ratio as well as the capital ratio to proxy for the riskiness of the borrower. In all specifications, standard errors are clustered at the lender-borrower level. The results are reported in Table 3.8.

The coefficients of interest, β_1 and β_3, are both significantly positive across all specifications. Relative to the baseline results in Table 3.3, the isolated effect of higher financed emissions captured by β_1 is somewhat smaller but remains statistically significant. This suggests that voluntary climate commitments account for a substantial part of the observed carbon premium, though they do not fully explain it.

3.4.2.5 Volumes

In theory, the climate premium that we are observing could be due to an equilibrium effect and not really a risk premium. For instance, brown firms may have more demand for funding which means that banks lending to these firms have to source the money from somewhere. The most popular way to source cash for commercial banks is the repo market. Consistent with this hypothesis would be if there is a positive correlation coefficient of *contemporaneous* financed

Table 3.8. Transition risk premia and preferences

Notes: The dependent variable is the loan rate minus the ECB's deposit facility rate expressed in percentage points. VCC stands for voluntary climate commitments and assumes the value one whenever the respective bank has committed to follow the NZBA guidelines. Financed GHG emissions are calculated as the volume-weighted GHG-intensities per bank-year from all commercial loans over EUR 25,000, divided by the bank's revenues. All continuous variables are standardised. Standard errors are clustered at the bank-counterparty-month-level and reported in parentheses. Controls include the collateral ISIN, total assets, the leverage ratio, the liquidity coverage ratio, the capital ratio, and the return on assets. $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

	DFR-adjusted Rate (%)			
	(1)	(2)	(3)	(4)
$FinancedEmissions$	0.048***	0.033***	0.015***	0.023***
	(0.007)	(0.005)	(0.005)	(0.007)
VCC	−0.039**	−0.046***	0.004	−0.040**
	(0.018)	(0.013)	(0.021)	(0.017)
$FinancedEmissions \times VCC$	0.036***	0.055***	0.039***	0.088***
	(0.012)	(0.010)	(0.009)	(0.016)
$Haircut$		−0.151***	−0.055***	0.006
		(0.017)	(0.009)	(0.007)
$Volume$		0.104***	0.076***	0.052***
		(0.005)	(0.004)	(0.003)
$Maturity$		0.003***	0.002***	0.002***
		(0.000)	(0.000)	(0.000)
Controls	✓	✓	✓	✓
FE: Date			✓	
FE: Lender			✓	
FE: Lender × Borrower				✓
N	2 006 742	2 006 742	2 006 742	2 006 742
R^2	0.16	0.33	0.43	0.50

emissions on volumes after controlling for counterparty specific characteristics.

To test this idea we are estimating the following regression:

$$Volume_{l,b,t,i} = \beta_1 \times FinancedEmissions_{b,t-365} + \beta_2 \times DealControls_i + \beta_3 \times BorrowerControls_{b,t} + \delta_t + \gamma_{l,b} + \varepsilon_{l,b,t,i}, \tag{3.9}$$

where δ_t and $\gamma_{l,b}$ are time and borrower-lender fixed effects. As deal controls, we include the collateral ISIN, the transaction maturity, adjusted rate and haircut. As borrower controls, we include total assets to proxy for the borrower's size, the return on assets to proxy for its profitability and the leverage ratio, the liquidity ratio as well as the capital ratio to proxy for the riskiness of the borrower. The results are reported in Table 3.9.

Table 3.9. Transition risk premia and volumes

Notes: The dependent variable is the transaction volume. Financed GHG emissions are calculated as the volume-weighted GHG-intensities per bank-year from all commercial loans over EUR 25,000, divided by the bank's revenues. All continuous variables are standardised. Standard errors are clustered at the bank-counterparty-month-level and reported in parentheses. Controls include the collateral ISIN, total assets, the leverage ratio, the liquidity coverage ratio, the capital ratio, and the return on assets. $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

	Volume (EUR mln)			
	(1)	(2)	(3)	(4)
$FinancedEmissions$	0.916	0.781	−0.799***	−1.076***
	(0.849)	(0.879)	(0.166)	(0.274)
$Haircut$	0.549	0.863	1.264	1.255
	(1.645)	(1.261)	(1.336)	(1.355)
$Maturity$	0.437***	0.357***	0.315***	0.314***
	(0.122)	(0.094)	(0.084)	(0.084)
Controls	✓	✓	✓	✓
FE: Date		✓		✓
FE: Lender		✓		
FE: Lender × Borrower			✓	✓
N	2 006 742	2 006 742	2 006 742	2 006 742
R^2	0.28	0.39	0.64	0.64

While coefficients for the first two specifications are not significant, the coefficients for regressions where we control for borrower fixed effects are actually negative and significant, which is inconsistent with the demand side channel.

3.5 Consequences to market stability and monetary policy

In the main body of this paper, we established a strong relationship between a bank's financed greenhouse gas (GHG) emissions and the repo rates it faces. As we argued, this relationship can be partly attributed to a risk premium and partly to investor preferences. With financed emissions likely to become an increasingly important focus of future policy, this pricing premium is poised to become even more significant. This raises an important question: what are the implications for the repo market when a factor seemingly exogenous to it is associated with such pricing differences? We explore this question along two key dimensions. First, we examine whether this relationship poses concerns for market stability or systemic risk. Second, we consider its implications for the transmission of monetary policy.

3.5.1 Amplification of liquidity risks

In this part of the paper, we want to examine whether this rate-sensitivity to transition risk has implications for market stability. More precisely we are asking: *Are banks that have higher financed GHG emissions disproportionally affected by market stress?* Our null hypothesis would be that transition risk, to the extent that it can be measured by financed GHG emissions, is independent of other exogenous liquidity risks in the funding-driven repo market. The alternative hypothesis would mean that the rate premium is exacerbated or even driven by stressed times, meaning that banks with higher emissions tend to get into funding stress quicker than less-heavily exposed banks.

To measure financial stress we first use the OFR Financial Stress Index (Office of Financial Research, 2024), which measures systemic financial stress using 33 market variables across five categories: credit, equity valuation, funding, safe assets, and volatility. It applies a dynamic factor model, standardizing each indicator and weighting them based on co-movement to capture overall market stress. Positive values indicate above-average stress, while negative values suggest calmer conditions. This way, the index provides a real-time gauge of financial stress, and helps to identify periods of market strain. Figure 3.6 shows the index values over our observation period. For the sake of interpretation, we interact the financed GHG emissions with a dummy variable indicating whether the Financial Stress index was in the top tercile of

Figure 3.6. OFR Financial Stress Index

Notes: This graph shows the OFR financial stress index over time. Shaded areas indicate days where the index has been in its top tercile regarding all values contained in the time span from 2019 to 2022.

observations over the observation window from 2019 to 2022. Specifically, the regression will be

$$
\begin{aligned}
Rate_{l,b,t,i} - DFR_t = {} & \beta_1 \times FinancedEmissions_{b,t-365} \\
& + \beta_2 \times FinRisk_t + \beta_3 \times FinRisk_t \times FinancedEmissions_{b,t-365} \\
& + \beta_4 \times DealControls_i + \beta_5 \times BorrowerControls_{b,t} + \delta_t + \gamma_l + \varepsilon_{l,b,t,i},
\end{aligned}
\tag{3.10}
$$

where δ_t and γ_l are date and lender fixed effects. The results are reported in Table 3.10.

Table 3.10. The amplification effect of financed GHG-emissions.

Notes: The dependent variable is the loan rate minus the ECB's deposit facility rate expressed in percentage points. Individual regressions are based on the sub-samples of transactions collateralised by bonds issued in the sectors shown. Financed GHG emissions are calculated as the volume-weighted GHG-intensities per bank-year from all commercial loans over EUR 25,000, divided by the bank's revenues. The variable $FinRisk$ is a dummy taking the value one whenever the OFR Financial Stress Index is in its top tercile. All continuous variables are standardised. Standard errors are clustered at the bank-counterparty-level and reported in parentheses. ${}^{*}p < 0.1$; ${}^{**}p < 0.05$; ${}^{***}p < 0.01$.

	DFR-adjusted rate (%)			
	All collateral	**NFC**	**FC**	**GG**
$FinancedEmissions$	0.020***	0.019*	0.014***	−0.007*
	(0.004)	(0.011)	(0.005)	(0.004)
$FinancedEmissions \times FinRisk$	0.052***	0.159***	0.049***	0.016*
	(0.010)	(0.029)	(0.010)	(0.010)
$Haircut$	−0.083***	0.011	−0.037***	−0.031***
	(0.009)	(0.017)	(0.012)	(0.006)
$Volume$	0.079***	0.010*	0.095***	0.030***
	(0.004)	(0.006)	(0.008)	(0.003)
$Maturity$	0.002***	0.004***	0.003***	0.001***
	(0.000)	(0.001)	(0.000)	(0.000)
Controls	✓	✓	✓	✓
FE: Date	✓	✓	✓	✓
FE: Lender	✓	✓	✓	✓
N	1 968 692	457 400	1 074 985	426 563
R^2	0.44	0.51	0.32	0.49

We first note that all coefficients for the interaction term as well as for the financed emissions are positive and significant. This is inconsistent with the previous hypothesis of complete independence. Moreover, the coefficients for the financed emissions term is slightly lower than those of our baseline regression in Table 3.3 suggesting that the coefficient for financed emissions is to some extent driven by the observation from stressed financial conditions. Yet they are still mostly positive and significant which indicates that if there is an amplification effect, a one standard deviation difference in finance emissions adds another 2 to 16 bps to rates in times of financial stress to an already existing but low premium of 0 to 2 bps. Again, the results for transactions collateralised by government bonds are weaker. Several factors could explain why we do not observe an effect in this segment. First, repos backed by government bonds are more standardised and conventional than those in other segments, meaning market standards may play a larger role. Second, government bond repos are more likely to be used for collateral purposes, which could obscure the true effect. Lastly, in times of financial stress, there may be stronger incentives to use high-quality government bonds specifically to avoid higher rates.

As a robustness test we are also conducting the same regressions with the (European) VIX index as a measure for financial stress and alternative fixed effects specifications which can be found in the appendix.

3.5.2 Implication for monetary policy

The repo market serves as the primary channel through which the ECB's monetary policy rates are transmitted. We have shown that banks with higher financed GHG emissions face a rate premium, paying more for repo funding. A natural follow-up question is whether this carbon exposure also affects the transmission of policy rate changes. Specifically, when the ECB raises rates, do high emission banks adjust to the new borrowing rate more slowly or more quickly than their peers?

There are four rate hikes in our sample that we could look at.[10] In a first specification we

[10] In 2022, the ECB raised its key policy rates four times - by 50 basis points in July, 75 in both September

are testing the rate response of brown firms following each of these policy events in a tight 6 weeks window around the decision. Specifically, we are testing

$$Rate^{resid}_{b,w} = \beta^{IV} D^{IV}_{w} \times FinancedEmissions_{w-52,b} + X_{b,w} + \delta_w + \gamma_b + \varepsilon_{b,w}, \quad (3.11)$$

where δ_w and γ_b are bank and week fixed effects. D^{IV}_w is an indicator which is 1 after the monetary policy announcement. We cluster at the borrower level and by using only three pre/post weekly aggregated loan rate residuals per bank we are mitigating concerns about serial correlation. The results are reported in Table 3.11.

Table 3.11. Monetary policy shocks

Notes: The dependent variable is the loan rate residual as defined in equation 3.6. Financed GHG emissions are calculated as the volume-weighted GHG-intensities per bank-year from all commercial loans over EUR 25,000, divided by the bank's revenues. All continuous variables are standardised. Standard errors are clustered at the borrower-level and reported in parentheses. Controls include the collateral ISIN, total assets, the leverage ratio, the liquidity coverage ratio, the capital ratio, and the return on assets. $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

	$Rate^{resid}$ (%)			
	(1)	(2)	(3)	(4)
$Jul27 \times FinancedEmissions$	0.049**			
	(0.022)			
$Sep14 \times FinancedEmissions$		0.027		
		(0.024)		
$Nov2 \times FinancedEmissions$			0.056***	
			(0.021)	
$Dec21 \times FinancedEmissions$				0.052
				(0.044)
FE: Counterparty	✓	✓	✓	✓
FE: Week	✓	✓	✓	✓
N	338	340	351	216
R2	0.862	0.898	0.922	0.938
R2 Adj.	0.814	0.858	0.890	0.900

The coefficients for all policy events are positive, ranging between 3 and 6 basis points,

and October, and another 50 in December - in response to rising inflation.

though only two are statistically significant. While it is difficult to pinpoint exactly what differentiates the significant events from the others, the results do not allow us to reject the null hypothesis of no effect on policy transmission. That said, the estimates suggest that brown banks, on average, adjust more quickly to higher rates following a hike, by around 5 basis points.

We complement this regression with an alternative approach measuring the passthrough directly as in Eisenschmidt et al. (2024). First, we consider loan rate residuals as in 3.6 to be able to average rates across dates while respecting the transaction heterogeneity. We then measure the passthrough as

$$Passthrough_{b,t}^{DFR\_OTC} = \frac{Rate_{b,t}^{resid,\ post} - Rate_{b,t}^{resid,\ pre}}{DFR_t^{post} - DFR_t^{pre}}, \tag{3.12}$$

where $Rate_{b,t}^{resid,\ pre}$ denotes the average pre-announcement rate per counterparty banks from the week before the policy decision, and $Rate_{b,t}^{resid,\ post}$ denotes the average post-implementation rate from the week after the new rate was implemented. The distribution for banks divided into above and below median financed emissions groups is given in Table 3.12. The table shows that brown banks adjust *quicker* to the new higher rate than green banks. This observation is also

Table 3.12. Passthrough by financed emissions group

Emission Group	Mean	P10	P25	P50	P75	P90
High emissions	64.01	16.58	41.95	72.15	91.67	96.68
Low emissions	54.75	2.71	22.15	60.00	86.33	94.44

confirmed again with the following regression

$$Passthrough_{b,t}^{DFR\_OTC} = \beta_1 \times FinancedEmissions_{t-52,b} + X_b + \delta_w + \varepsilon_{b,w}, \tag{3.13}$$

where X_b are borrower-level controls and δ_w are time fixed effects. As all rate hikes happened in the same year 2022 for our sample, we cannot use borrower fixed-effects, as this would make financed emissions obsolete which also vary only on the year level. The coefficient on pass-

Table 3.13. Monetary policy shocks

Notes: The dependent variable is the Passthrough as defined in equation 3.12. Financed GHG emissions are calculated as the volume-weighted GHG-intensities per bank-year from all commercial loans over EUR 25,000, divided by the bank's revenues. The variable $FinancedEmissions$ is standardised. Standard errors are clustered at the counterparty-level and reported in parentheses. $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

	$Passthrough^{DFR\_OTC}$ (%)		
	(1)	(2)	(3)
$FinancedEmissions$	4.625**	7.196**	7.435**
	(2.175)	(3.163)	(3.132)
$Assets$		0.000	0.000
		(0.000)	(0.000)
ROA		595.256	736.065
		(978.246)	(983.308)
$Leverage$			137.709
			(258.111)
$Liquidity$			−2.617
			(2.658)
$Capital$			65.801
			(144.769)
FE: Month	✓	✓	✓
N	278	195	195
R^2	0.01	0.02	0.03

through is positive across all specifications, indicating that a one standard deviation increase in financed GHG emissions is associated with a roughly 7% faster adjustment to the new policy rate in the first week after implementation. This suggests that brown banks experience quicker rate pass-through following a hike. Whether this reflects dealers' preferences for greener counterparties or the pricing in of a brown risk premium, the evidence points to a consistent pattern: dealers appear to be less accommodating toward brown banks, while green banks tend to face a more gradual adjustment to higher rates.

3.6 Conclusion

How is climate change affecting financial stability? This is a fundamental question not only for the emerging field of climate change and finance but also for policymakers who are striving to design stress tests that adequately capture all potential factors impacting the financial system.

To address this question, we conduct a detailed cross-sectional and time-series analysis of interbank lending rates, haircuts, and volumes, using financed carbon emissions as a key bank characteristic. We find robust evidence that financed carbon emissions significantly and positively affect borrowing rates in the funding-driven repo market which can be rationalised in part as a result of banks' preferences for greener counterparties and in part also as a carbon risk premium.

This provides clear evidence of a link between transition risk, financial stability, and monetary policy. Transition risk raises firms' production costs and puts downward pressure on profits, increasing the risk of default. As a result, banks lending to these firms face higher credit risk. In the repo market, brown banks are penalised with higher rates, and during periods of rate hikes, this also leads to a faster transmission of new policy rates. The repo market is a key source of daily liquidity for both banks and non-banks, and its importance is likely to grow as excess liquidity declines.

Additionally, our findings suggest that transition risk could amplify other exogenous financial risks. While transition risk alone might not induce immediate stress in the repo market, its mere presence implies that banks are subject to differential treatment in terms of their cost of funding. This dynamic suggests that, in the presence of significant market shocks, such amplification could indeed exacerbate broader financial instability.

Conclusion

This thesis provides new evidence on how fiscal policy, financial regulation, and climate transition risk interact with the institutional architecture of European financial markets. While the three chapters differ in their specific focus, they share a common finding: risks that appear contained when viewed in isolation can become systemically significant once they encounter the frictions and structural features of the euro area's financial system.

Several policy implications emerge. First, the design of fiscal policy in a monetary union cannot abstract from its financial stability consequences. Chapter 1 shows that expansionary fiscal policy by the anchor sovereign generates contractionary financial spillovers that materialise before any stimulus is disbursed. This pre-spending penalty suggests that assessments of fiscal multipliers in currency unions must account for the financial channel. The case for a common safe asset, such as jointly issued Eurobonds, is strengthened by the finding that the entire transmission mechanism depends on the national segmentation of sovereign debt markets.

Second, financial regulation entails trade-offs that extend beyond its immediate objectives. Chapter 2 demonstrates that the Basel III leverage ratio, while designed to limit excessive bank leverage, distorts the setting of haircuts in the repo market and weakens the pass-through of monetary policy. Policymakers debating the introduction of minimum haircut floors should recognise that the current dysfunction in haircut-setting is not primarily a failure of risk management but a consequence of binding regulatory constraints interacting with market structure.

Third, climate transition risk is no longer confined to long-horizon asset classes. Chapter 3 establishes that it is already priced in the repo market and that it amplifies during periods of financial stress. This finding reinforces the case for integrating climate risk into systemic risk assessments and central bank stress-testing frameworks, as the compounding of transition risk with broader financial vulnerabilities may pose greater dangers than direct climate exposures alone.

More broadly, the thesis underscores that financial stability and monetary policy transmission are deeply intertwined, and that both are shaped by institutional design choices that are themselves subject to change. As the euro area confronts a period of fiscal expansion, evolving financial regulation, and accelerating climate transition, understanding how these forces interact within its financial architecture is not merely an academic exercise but a prerequisite for effective policy.

Bibliography

Acharya, Viral V., Richard Berner, Robert Engle, Hyeyoon Jung, Johannes Stroebel, Xuran Zeng, and Yihao Zhao, "Climate Stress Testing," *Annual Review of Financial Economics*, 2023, *15*, 291–326.

_, Tim Eisert, Christian Eufinger, and Christian Hirsch, "Real Effects of the Sovereign Debt Crisis in Europe: Evidence from Syndicated Loans," *The Review of Financial Studies*, August 2018, *31* (8), 2855–2896.

Adrian, Tobias, Richard K. Crump, and Emanuel Moench, "Pricing the term structure with linear regressions," *Journal of Financial Economics*, 2013, *110* (1), 110–138.

Altavilla, Carlo, Luca Brugnolini, Refet S. Gürkaynak, Roberto Motto, and Giuseppe Ragusa, "Measuring euro area monetary policy," *Journal of Monetary Economics*, 2019, *108*, 162–179.

_, Marco Pagano, and Saverio Simonelli, "Bank Exposures and Sovereign Stress Transmission," *Review of Finance*, 2017, *21* (6), 2103–2139.

_, Miguel Boucinha, Marco Pagano, and Andrea Polo, "Climate risk, bank lending and monetary policy," *ECB Working Paper No 2969*, 2024.

Angrist, Joshua D. and Jörn-Steffen Pischke, "How Much Should We Trust Differences-In-Differences Estimates?," *The Quarterly Journal of Economics*, 2004, *119* (1), 249–275.

Arrata, William, Benoit Nguyen, Imene Rahmouni-Rousseau, and Miklos Vari, "The scarcity effect of quantitative easing on repo rates: Evidence from the Euro area," *Journal of Financial Economics*, 2020, *137* (3), 837–856.

Baker, Malcolm, Daniel Bergstresser, George Serafeim, and Jeffrey Wurgler, "Financing the Response to Climate Change: The Pricing and Ownership of US Green Bonds," Technical Report, National Bureau of Economic Research 2018.

Baker, S. R., N. Bloom, and S. J. Davis, "Measuring Economic Policy Uncertainty," *The Quarterly Journal of Economics*, 2016, *131* (4), 1593–1636.

Bakkensen, Laura A. and Leah Barrage, "Flood Risk Belief Heterogeneity and Coastal Home Price Dynamics: Going Under Water?," Technical Report, National Bureau of Economic Research 2017.

Baldauf, Matthias, Luca Garlappi, and Constantine Yannelis, "Does Climate Change Affect Real Estate Prices? Only If You Believe in It," *The Review of Financial Studies*, 2020, *33*, 1256–1295.

Ballensiefen, Benedikt and Angelo Ranaldo, "Safe asset carry trade," *Review of Asset Pricing Studies*, 2023, *13* (2), 223–265.

_ , _ , and Hannah Winterberg, "Money market disconnect," *Review of Financial Studies*, 2023, *36* (10), 4158–4189.

_ , _ , and _ , "Money Market Disconnect," *The Review of Financial Studies*, 04 2023, *36* (10), 4158–4189.

Barnett, Michael, "Estimated Impacts of Climate Policy Risk," Working Paper 2020.

Barro, Robert J., "Rare Disasters and Asset Markets in the Twentieth Century," *The Quarterly Journal of Economics*, 2006, *121* (3), 823–866.

Beber, Alessandro, Michael W. Brandt, and Kenneth A. Kavajecz, "Flight-to-Quality or Flight-to-Liquidity? Evidence from the Euro-Area Bond Market," *The Review of Financial Studies*, March 2009, *22* (3), 925–957.

Bejarano, Jeremy, Lina Lu, and Jonathan Wallen, "Negative Treasury Haircuts," *Harvard Business School Working Paper, No. 26-034*, 2025.

Bernstein, Asaf, Matthew T. Gustafson, and Ryan Lewis, "Disaster on the Horizon: The Price Effect of Sea Level Rise," *Journal of Financial Economics*, 2019, *134*, 253–272.

Bittner, Christian and Stephan Jank, "Lending to Hedge Funds: Does Competition Erode Bank Risk Management?," SSRN Working Paper 5151327, Deutsche Bundesbank February 2025.

Blanchard, Olivier and Roberto Perotti, "An Empirical Characterization of the Dynamic Effects of Changes in Government Spending and Taxes on Output," *The Quarterly Journal of Economics*, 2002, *117* (4), 1329–1368.

Bocola, Luigi, "The Pass-Through of Sovereign Risk," *Journal of Political Economy*, 2016, *124* (4), 879–926.

Bolton, Patrick and Marcin Kacperczyk, "Are Carbon Emissions Associated with Stock Returns? Comment?," *Review of Finance*, 2024, *28(1)*, 107–109.

Brunnermeier, Markus K., "Safe Assets," *Annual Review of Economics*, 2023, *15*, 153–178.

— , **Luis Garicano, Philip R. Lane, Marco Pagano, Ricardo Reis, Tano Santos, David Thesmar, Stijn Van Nieuwerburgh, and Dimitri Vayanos**, "The Sovereign-Bank Diabolic Loop and ESBies," *American Economic Review: Papers & Proceedings*, 2016, *106* (5), 508–512.

Bua, Giovanna, Daniel Kapp, Federico Ramella, and Lavinia Rognone, "Transition versus physical climate risk pricing in European financial markets: a text-based approach," *The European Journal of Finance*, 2024, *30* (17), 2076–2110.

Buraschi, Andrea and Davide Menini, "Liquidity risk and specialness," *Journal of Financial Economics*, 2002, *64* (2), 243–284.

Caballero, Ricardo J. and Emmanuel Farhi, "The Safety Trap," *Review of Economic Studies*, January 2018, *85* (1), 223–274.

— , — , **and Pierre-Olivier Gourinchas**, "An Equilibrium Model of "Global Imbalances" and Low Interest Rates," *American Economic Review*, March 2008, *98* (1), 358–393.

Caballero, Ricardo J, Emmanuel Farhi, and Pierre-Olivier Gourinchas, "Safe asset scarcity and aggregate demand," *American Economic Review*, 2016, *106* (5), 513–518.

Cenicola, Ashlyn, Robert Mann, and Mark Paddrik, "Are Zero-Haircut Repos as Common as Advertised?," *The OFR Blog*, 8 2025.

Cochrane, John H. and Monika Piazzesi, "Bond risk premia," *American Economic Review*, 2005, *95* (1), 138–160.

Coen, Jamie, Patrick Coen, and Anne-Caroline Hüser, "Collateral demand in wholesale funding markets," *Bank of England Staff Working Paper No. 1,082*, 2024.

Collender, Sierra, Baoqing Gan, Christina S. Nikitopoulos, Kylie-Anne Richards, and Laura Ryan, "Climate transition risk in sovereign bond markets," *Global Finance Journal*, 2023, *57*, 100868.

Corradin, Stefano and Angela Maddaloni, "The importance of being special: Repo markets during the crisis," *Journal of Financial Economics*, 2020, *137* (2), 392–429.

— , **Jens Eisenschmidt, Marie Hoerova, Tobias Linzert, Glenn Schepens, and Jean-David Sigaux**, "Money markets, central bank balance sheet and regulation," *ECB Working Paper*, 2020, *No. 2483.*

D'Amico, Stefania and Thomas B. King, "Flow and stock effects of large-scale treasury purchases: Evidence on the importance of local supply," *Journal of Financial Economics*, May 2013, *108* (2), 425–448.

Dang, Tri Vi, Gary Gorton, and Bengt Holmström, "Ignorance, debt and financial crises," 2015. Working paper, Yale School of Management.

Degryse, Hans, Roman Goncharenko, Carola Theunisz, and Tamas Vadasz, "When Green Meets Green: Do Environmentally Conscious Banks Reward Environmentally Conscious Firms?," *Journal of Corporate Finance*, 2023, *78*, 102355.

Delis, Manthos D., Kathrin de Greiff, Maria Iosifidi, and Steven R. G. Ongena, "Being Stranded with Fossil Fuel Reserves? Climate Policy Risk and the Pricing of Bank Loans," Technical Report Research Paper No. 18-10, Swiss Finance Institute (also EBRD and CEPR Working Paper) 2018.

Drechsler, Itamar, Alexi Savov, and Philipp Schnabl, "The deposits channel of monetary policy," *Quarterly Journal of Economics*, 2017, *132* (4), 1819–1876.

Du, Wenxin, Joanne Im, and Jesse Schreger, "The US Treasury premium," *Journal of International Economics*, May 2018, *112*, 167–181.

Duffie, D., "Special repo rates," *Journal of Finance*, 1996, *51* (2), 493–526.

Duffie, Darrell and Arvind Krishnamurthy, "Pass-Through Efficiency in the Fed's New Monetary Policy Setting," in "Designing Resilient Monetary Policy Frameworks for the Future" Federal Reserve Bank of Kansas City Jackson Hole 2016, pp. 21–102.

_ **and** _ , "Pass-through efficiency in the Fed's new monetary policy setting," *Designing resilient monetary policy frameworks for the future, a symposium sponsored by the Federal Reserve Bank of Kansas City, Jackson Hole*, 2016, pp. 21–102.

Eisenschmidt, Jens, Yiming Ma, and Anthony Lee Zhang, "Monetary policy transmission in segmented markets," *Journal of Financial Economics*, 2024, *151*, 103738.

Engle, Robert F, Stefano Giglio, Bryan Kelly, Heebum Lee, and Johannes Stroebel, "Hedging Climate Change News," *The Review of Financial Studies*, 02 2020, *33* (3), 1184–1216.

European Central Bank, *The Monetary Policy of the ECB*, European Central Bank, 2011. Accessed: 2025-03-29.

_ , "Analytical indicators on carbon emissions," https://www.ecb.europa.eu/stats/all-key-statistics/horizontal-indicators/sustainability-indicators/data/html/ecb.climate_indicators_carbon_emissions.en.html 2024. European Central Bank website — updated April 2024; accessed Jul 2025.

_ , "Money Market Statistics," 2025. Available at: https://www.ecb.europa.eu/stats/financial_markets_and_interest_rates/money_market/html/index.en.html (Accessed: 2025-03-09).

Faccini, Renato, Rastin Matin, and George Skiadopoulos, "Dissecting climate risks: Are they reflected in stock prices?," *Journal of Banking and Finance*, 2023, *155*, 106948.

Farhi, Emmanuel and Jean Tirole, "Deadly Embrace: Sovereign and Financial Balance Sheets Doom Loops," *The Review of Economic Studies*, 2018, *85* (3), 1781–1823.

Financial Stability Board, "Leverage in Nonbank Financial Intermediation," *FSB Final Report*, 2025.

Foerster, Kai, Ellen Ryan, and Benedikt Alois Scheid, "Pricing or Panicking? Commercial Real Estate Markets and Climate Change," ECB Working Paper 2025/3059, European Central Bank May 2025. Available at SSRN: https://ssrn.com/abstract=5265614.

Gertler, Mark and Peter Karadi, "Monetary Policy Surprises, Credit Costs, and Economic Activity," *American Economic Journal: Macroeconomics*, January 2015, *7* (1), 44–76.

Giannetti, Mariassunta, Martina Jasova, Maria Loumioti, and Caterina Mendicino, ""Glossy Green" Banks: The Disconnect Between Environmental Disclosures and Lending Activities," Technical Report 5357407, SSRN Electronic Journal 2025.

Giglio, Stefano, Matteo Maggiori, Krishna Rao, Johannes Stroebel, and Andreas Weber, "Climate Change and Long-Run Discount Rates: Evidence from Real Estate," *The Review of Financial Studies*, 03 2021, *34* (8), 3527–3571.

Goldsmith-Pinkham, Paul S., Matthew Gustafson, Ryan Lewis, and Michael Schwert, "Sea Level Rise and Municipal Bond Yields," *SSRN Electronic Journal*, 2019. Available at SSRN.

Gorton, Gary, "The history and economics of safe assets," *Annual Review of Economics*, 2017, *9*, 547–586.

_ **and Andrew Metrick**, "Securitized banking and the run on repo," *Journal of Financial Economics*, 2012, *104* (3), 425–451.

_ **and Guillermo Ordonez**, "Collateral crises," *American Economic Review*, 2014, *104* (2), 343–378.

Grauwe, Paul De and Yuemei Ji, "Self-Fulfilling Crises in the Eurozone: An Empirical Test," *Journal of International Money and Finance*, 2013, *34*, 15–36.

Greenwood, Robin and Dimitri Vayanos, "Price Pressure in the Government Bond Market," *American Economic Review*, May 2010, *100* (2), 585–590.

Gürkaynak, Refet S., Brian Sack, and Jonathan H. Wright, "The US treasury yield curve: 1961 to the present," *Journal of Monetary Economics*, 2007, *54* (8), 2291–2304.

Hall, Peter, *The Bootstrap and Edgeworth Expansion* Springer Series in Statistics, New York: Springer-Verlag, 1992.

Hanson, Samuel G. and Jeremy C. Stein, "Monetary policy and long-term real rates," *Journal of Financial Economics*, 2015, *115* (3), 429–448.

He, Zhiguo, Arvind Krishnamurthy, and Konstantin Milbradt, "A Model of Safe Asset Determination," *American Economic Review*, April 2019, *109* (4), 1230–1262.

Hempel, Samuel J., R. Jay Kahn, Robert Mann, and Mark E. Paddrik, "Why Is So Much Repo Not Centrally Cleared? Lessons from a Pilot Survey of Non-centrally Cleared Repo Data," *OFR Brief 23-01 | May 12, 2023*, 5 2023.

Hermes, Felix, Maik Schmeling, and Andreas Schrimpf, "The International Dimension of Repo: Five New Facts," *ECB Working Paper No. 2025/3065*, 2025.

Hong, Harrison, Frank W. Li, and Jiang Xu, "Climate Risks and Market Efficiency," *Journal of Econometrics*, 2019, *208*, 265–281.

Hsu, Po-Hsuan, Kai Li, and Chia-Ying Tsou, "The Pollution Premium," *SSRN Electronic Journal*, 2020. Available at SSRN 3578215.

Huynh, Toan and Yuhang Xia, "Climate Change News Risk and Corporate Bond Returns," *Journal of Financial and Quantitative Analysis*, 2020. Forthcoming.

Ilhan, Emirhan, Zacharias Sautner, and Grigory Vilkov, "Carbon Tail Risk," *The Review of Financial Studies*, 06 2020, *34* (3), 1540–1571.

Infante, Sebastian, "Liquidity windfalls: The consequences of repo rehypothecation," *Journal of Financial Economics*, 2019, *133* (1), 42–63.

Jentsch, Carsten and Kurt G. Lunsford, "Asymptotically Valid Bootstrap Inference for Proxy SVARs," *Journal of Business & Economic Statistics*, 2022, *40* (4), 1876–1891.

Jiang, Zhengyang, Arvind Krishnamurthy, and Hanno Lustig, "Foreign Safe Asset Demand and the Dollar Exchange Rate," *The Journal of Finance*, June 2021, *76* (3), 1049–1089.

Jordà, Òscar, "Estimation and Inference of Impulse Responses by Local Projections," *American Economic Review*, March 2005, *95* (1), 161–182.

Julliard, Christian, Gabor Pinter, Karamfil Todorov, Jean-Charles Wijnandts, and Kathy Yuan, "What drives repo haircuts? Evidence from the UK market," *BIS Working Papers*, 07 2022, *1027*.

Kacperczyk, Marcin and Patrick Bolton, "Do Investors Care about Carbon Risk?," *Journal of Financial Economics*, 2021, *142(2)*, 517–549.

_ and _ , "Global Pricing of Carbon-Transition Risk," *Journal of Finance (forthcoming)*, 2022.

Kacperczyk, Marcin T. and José-Luis Peydró, "Carbon Emissions and the Bank–Lending Channel," Technical Report 3915486, SSRN Electronic Journal 2022. European Corporate Governance Institute – Finance Working Paper No. 991/2024; CEPR Discussion Paper No. DP16778.

Kahn, R. Jay and Matthew McCormick, "Proportionate margining for repo transactions," *FEDS Notes*, 2 2025.

Kreps, David M. and Jose A. Scheinkman, "Quantity Precommitment and Bertrand Competition Yield Cournot Outcomes," *The Bell Journal of Economics*, 1983, *14* (2), 326–337.

Krishnamurthy, Arvind and Annette Vissing-Jorgensen, "The Aggregate Demand for Treasury Debt," *Journal of Political Economy*, April 2012, *120* (2), 233–267.

_ **and** _ , "The aggregate demand for Treasury debt," *Journal of Political Economy*, 2012, *120* (2), 233–267.

Krueger, Philipp, Zacharias Sautner, and Laura T Starks, "The Importance of Climate Risks for Institutional Investors," *The Review of Financial Studies*, 02 2020, *33* (3), 1067–1111.

Künsch, Hans R., "The Jackknife and the Bootstrap for General Stationary Observations," *The Annals of Statistics*, 1989, *17* (3), 1217–1241.

Känzig, Diego R, "The Macroeconomic Effects of Oil Supply News: Evidence from OPEC Announcements," *American Economic Review*, April 2021, *111* (4), 1092–1125.

_ , "The Unequal Economic Consequences of Carbon Pricing," Working Paper 31221, National Bureau of Economic Research May 2023.

Lane, Philip R., "The European Sovereign Debt Crisis," *Journal of Economic Perspectives*, 2012, *26* (3), 49–68.

Leeper, Eric M., "Equilibria under 'active' and 'passive' monetary and fiscal policies," *Journal of Monetary Economics*, 1991, *27* (1), 129–147.

Longstaff, Francis A., Jun Pan, Lasse H. Pedersen, and Kenneth J. Singleton, "How Sovereign Is Sovereign Credit Risk?," *American Economic Journal: Macroeconomics*, April 2011, *3* (2), 75–103.

Mancini, Loriano, Angelo Ranaldo, and Jan Wrampelmeyer, "The Euro interbank repo market," *Review of Financial Studies*, 2016, *29* (7), 1747–1779.

Mertens, Karel and Morten O. Ravn, "The Dynamic Effects of Personal and Corporate Income Tax Changes in the United States," *American Economic Review*, June 2013, *103* (4), 1212–47.

Mian, Atif, Amir Sufi, and Emil Verner, "Household Debt and Business Cycles Worldwide," *The Quarterly Journal of Economics*, 2017, *132* (4), 1755–1817.

Miranda-Agrippino, Silvia and Giovanni Ricco, "The Transmission of Monetary Policy Shocks," *American Economic Journal: Macroeconomics*, July 2021, *13* (3), 74–107.

Montiel Olea, José Luis and Carolin Pflueger, "A Robust Test for Weak Instruments," *Journal* *of Business & Economic Statistics*, 2013, *31* (3), 358–369.

Moreira, Alan and Alexi Savov, "The macroeconomics of shadow banking," *Journal of Finance*, 2017, *72* (6), 2381–2432.

Mountford, Andrew and Harald Uhlig, "What are the effects of fiscal policy shocks?," *Journal of Applied Econometrics*, 2009, *24* (6), 960–992.

Nagel, Stefan, "The liquidity premium of near-money assets," *Quarterly Journal of Economics*, 2016, *131* (4), 1927–1971.

Nakamura, Emi and Jón Steinsson, "Identification in Macroeconomics," *Journal of Economic Perspectives*, Summer 2018, *32* (3), 59–86.

Office of Financial Research, "OFR Financial Stress Index," 2024. Available at: https://www.financialresearch.gov/financial-stress-index/ (Accessed: 2024-03-24).

Painter, Marcus, "An Inconvenient Cost: The Effects of Climate Change on Municipal Bonds," *Journal of Financial Economics*, 2020, *135*, 468–482.

Phillot, Maxime, "US Treasury Auctions: A High-Frequency Identification of Supply Shocks," *American Economic Journal: Macroeconomics*, January 2025, *17* (1), 245–73.

Ramelli, Stefano, Elisa Ossola, and Michela Rancan, "Stock price effects of climate activism: Evidence from the first Global Climate Strike," *Journal of Corporate Finance*, 2021, *69*, 102018.

Ramey, Valerie A, "Identifying Government Spending Shocks: It's all in the Timing," *The Quarterly Journal of Economics*, 2011, *126* (1), 1–50.

Ranaldo, Angelo, Patrick Schaffner, and Michalis Vasios, "Regulatory effects on short-term interest rates," *Journal of Financial Economics*, 2021, *141* (2), 750–770.

Reuters, "Germany's Merz races to win support for major financial package," *Reuters*, 2025. Accessed 25 Nov 2025.

Romer, Christina D and David H Romer, "The Macroeconomic Effects of Tax Changes: Estimates Based on a New Measure of Fiscal Shocks," *American Economic Review*, 2010, *100* (3), 763–801.

Santis, Roberto A. De, "Redenomination Risk," *Journal of Money, Credit and Banking*, 2019, *51* (8), 2173–2206.

Sastry, Parinitha R., Emil Verner, and David Marques Ibanez, "Business as Usual: Bank Net Zero Commitments, Lending, and Engagement," Technical Report NBER Working Paper No. w32402, National Bureau of Economic Research May 2024.

Schaffner, Patrick, Angelo Ranaldo, and Kostas Tsatsaronis, "Euro repo market functioning:

collateral is king," *BIS Quarterly Review, December*, 2019.

Schularick, Moritz and Alan M. Taylor, "Credit Booms Gone Bust: Monetary Policy, Leverage Cycles, and Financial Crises, 1870–2008," *American Economic Review*, 2012, *102* (2), 1029–61.

Shambaugh, Jay C., "The Euro's Three Crises," *Brookings Papers on Economic Activity*, 2012, *2012* (1), 157–231.

Stock, James and Motohiro Yogo, "Testing for Weak Instruments in Linear IV Regression," in "Identification and Inference for Econometric Models: Essays in Honor of Thomas Rothenberg," New York: Cambridge University Press, 2005, pp. 80–108.

Stock, James H. and Mark W. Watson, "Disentangling the Channels of the 2007-2009 Recession," NBER Working Paper 18094, National Bureau of Economic Research May 2012.

Task Force on Climate-related Financial Disclosures, "Final Report: Recommendations of the Task Force on Climate-related Financial Disclosures," Technical Report, Financial Stability Board 2017.

Wiegand, Courtney, "The Effect of Fiscal Policy Shocks on Asset Prices," Working Paper, NYU Stern School of Business February 2025.

Ye, Zhen, "Bank Divestment and Green Innovation," Technical Report 4324996, SSRN Electronic Journal 2023.

Appendix to Chapter 1

1.A News selection

All news were taken from Reuters. Only news containing direct information regarding the German budget published were chosen. The data was accessed on 25 November 2025.

Table 1.A.1. German Fiscal Policy News Headlines

Date	Headline
02/10/2025	Germany using landmark infrastructure fund to ease budget pressures
23/09/2025	What is included in Germany's 2026 draft budget
18/09/2025	Germany approves 2025 budget ushering in new era of spending
05/09/2025	German budget committee approves investment-heavy 2025 budget
30/07/2025	German cabinet approves 2026 budget
30/07/2025	German cabinet approves 2026 budget tripling borrowing
21/07/2025	German government to approve 2026 budget with record investment sources say
21/07/2025	Aviation tax cut not a priority in 2026 budget German government source says

Continued on next page

Table 1.A.1 continued

Date	Headline
19/05/2025	Germany pushes ahead with budget preparations savings encouraged document shows
19/05/2025	Germany pushes ahead with budget preparations while seeking savings document shows
14/05/2025	German cabinet due to agree 2026 draft budget in July says finance minister
14/05/2025	Germany sees limited financial leeway on next EU budget policy paper shows
21/03/2025	German budget committee gives green light for 3 bln euros in Ukraine aid
16/03/2025	Budget committee approves Germany's massive borrowing plans
05/03/2025	Germany's Merz races to win support for major financial package
17/02/2025	German cities struggle with budget amid rising costs weak economy survey shows
29/01/2025	Germany's budget committee wants to approve 3 billion euros for Ukraine sources say
18/01/2025	German budget committee backs 4.7 billion euro deal to buy ThyssenKrupp submarines
02/12/2024	French bond risk premium rises on budget woes German yields fall
11/10/2024	Germany's Lindner to decide on 2025 budget debt after tax estimate
11/10/2024	French 10-year yields rise after belt-tightening budget German Bund yields climb
16/09/2024	Germany's coalition argues over Intel subsidies in new budget dispute

Continued on next page

Table 1.A.1 continued

Date	Headline
10/09/2024	Germany's finance minister defends 2025 budget despite 12 bln euro shortfall
01/08/2024	German finance ministry advisory board doubtful about efforts to reduce budget gap
29/07/2024	German defence minister told not to expect additional budget boost
17/07/2024	German coalition set to pass 2025 budget mid-term financial plan
17/07/2024	German coalition passes 2025 budget after months of wrangling
16/07/2024	German 2025 budget draft has a 17 billion euro hole – sources
15/07/2024	German coalition to agree on Wednesday on draft budget – sources
15/07/2024	France will have to abide by EU budget rules say Germany and Commission
08/07/2024	German defence budget for 2025 falls signicantly short of request minister says
08/07/2024	German budget deal is basis for economic confidence in H2 says econ min
08/07/2024	German defence budget for 2025 significantly less than sought minister says
08/07/2024	Germany plans 50.3 bln euros in new debt with supplementary 2024 budget report says
04/07/2024	Work to do but we will get a budget deal German Fin Min says
02/07/2024	Germany's Scholz expects deal on 2025 budget this week sources say
28/06/2024	What next for Germany's contentious 2025 budget
24/06/2024	Germany's Scholz says he will keep social agenda in budget deal
22/05/2024	German defence minister not ready to compromise on 2025 budget demand
16/04/2024	German budget needs radical consolidation says audit institute

Continued on next page

Table 1.A.1 continued

Date	Headline
07/03/2024	Preparation of Germany's 2025 budget will be challenging – ministry sources say
02/02/2024	Germany's lower house of parliament approves 2024 budget
30/01/2024	German advisory council floats debt brake reform after budget mess
26/01/2024	German government faces double-digit billion gap in 2025 budget
24/01/2024	Germany's Ifo institute trims 2024 GDP forecast again due to budget
18/01/2024	No need to suspend German debt brake in 2024 budget committee member says
16/01/2024	Germany has more funds than expected for 2024 budget
08/01/2024	German cabinet approves savings plans for 2024 budget
04/01/2024	German budget savings shrink as farm subsidy cuts delayed
20/12/2023	German budget measures to drive up inflation at start of 2024 economist says
19/12/2023	German government details where savings will be made in 2024 budget
15/12/2023	Germany approves 2023 budget with debt brake suspension
14/12/2023	DIW institute lowers German GDP forecasts after budget cuts
14/12/2023	Germany lifts spending freeze after budget deal clinched document shows
13/12/2023	Agreement on German 2024 budget reached – government sources
13/12/2023	German GDP to contract 0.5% in 2024 due to budget crisis – IW
13/12/2023	Germany to plug budget hole through savings and subsidy cuts – Scholz
13/12/2023	Germany aims to return to debt brake as part of 2024 budget deal
13/12/2023	German budget fix to further crimp growth in 2024
13/12/2023	Main elements of Germany's budget for 2024

Continued on next page

Table 1.A.1 continued

Date	Headline
13/12/2023	Germany clinches last-minute 2024 budget deal keeps debt brake
13/12/2023	German 2024 budget deal could spur EU agreement on budget revision
12/12/2023	What is happening with Germany's 2024 budget
10/12/2023	German coalition leaders to resume budget talks on Sunday – sources
09/12/2023	German 2024 budget deal possible this week – SPD
09/12/2023	Germany's Scholz confident that budget crisis can be overcome
08/12/2023	Germany's Lindner talks tough on EU deficits as budget crisis drags on at home
07/12/2023	German budget crisis seen dragging into next year as coalition haggles
06/12/2023	German coalition parties at loggerheads as budget crisis grows
05/12/2023	Germany's Scholz expects coalition to swiftly agree deal on 2024 budget
04/12/2023	No exact timeline for German budget talks – govt spokesperson
04/12/2023	German budget crisis stings already weak economy – ZEW president
04/12/2023	SCENARIOS – German budget crisis talks enter crunch week
02/12/2023	German economy minister cancels Middle East trip for budget talks
01/12/2023	Germany's climate commitments uncertain amid budget crisis – COP envoy
01/12/2023	German finance minister names target areas for cuts amid budget crisis – Funke media
29/11/2023	Germany faces high double-digit billion gap in 2024 budget – SPD politician
29/11/2023	Germany faces 17 billion euro gap in 2024 budget – finance minister
28/11/2023	German coalition leaders to hold budget talks on Weds evening – sources

Continued on next page

Table 1.A.1 continued

Date	Headline
28/11/2023	Germany's Scholz vows to modernise economy back Ukraine despite budget woes
27/11/2023	German budget row prises open debt brake debate
27/11/2023	German government unveils budget fixes as way out of crisis
27/11/2023	German cabinet agrees 2023 supplementary budget – finance ministry
21/11/2023	Germany freezes new spending commitments as budget woes deepen
16/11/2023	German coalition faces difficult choices after budget hammerblow
14/11/2023	German budget sees more cash for Ukraine green buildings industry
08/11/2023	German budget to spell out ban on funding terrorist activities
25/10/2023	Ukraine aid clear priority for Berlin in EU budget talks – official
28/09/2023	Italy bond yields surge after budget German 10-year hits 12-year high
15/09/2023	AfD's regional budget win erodes German firewall against far right
12/09/2023	Germany's 2024 budget reverses fiscal trend since 2019
05/09/2023	German finance minister 2024 budget just a first step in fiscal normalization
05/07/2023	Germany approves first draft of 2024 budget focused on fiscal realities
05/07/2023	German cabinet approves first draft of 2024 budget
05/07/2023	Germany eyes slashing parental leave allowance for higher earners as budget tightens
19/06/2023	German budget draft will be submitted in early July says finance minister
09/06/2023	Germany's Lindner plans new debt of almost 17 bln euros for 2024 budget – Spiegel
25/05/2023	No date in sight for end of Germany's budget dispute

Continued on next page

Table 1.A.1 continued

Date	Headline
11/03/2023	German finance minister says defence budget needs to rise
09/03/2023	Germany's budget for 2024 isn't in danger despite delays coalition says
24/02/2023	German minister delays 2024 budget plans as coalition squabbles – media
20/02/2023	Budget demands of German coalition partners unrealistic says finance minister
15/02/2023	Top German ministers spar over budget in sign of coalition tensions
22/12/2022	Germany faces 12 bln euro budget financing gap – Handelsblatt
09/12/2022	German government deficit to hit 4.5% in 2023
01/12/2022	German FinMin Commission proposals aren't end of debate on EU budget rules
29/11/2022	Germany can step in with OSCE funding if Russia blocks budget – foreign minister
18/11/2022	German finance minister must talk about fiscal stimulus for 2024 budget
07/11/2022	German budget experts push for dividend ban tied to energy subsidies
09/09/2022	Germany will get state support done for VNG – economy minister
09/09/2022	Euro zone to coordinate fiscal, monetary policy to fight inflation
15/07/2022	Germany plans 40 million euros in budget aid to Moldova
01/07/2022	Possible Uniper bailout would not directly affect German budget says finance minister
11/06/2022	No reserves in Germany's federal budget says finance minister
14/03/2022	Germany plans supplementary budget to cope with Ukraine impact – sources

Continued on next page

Table 1.A.1 continued

Date	Headline
09/03/2022	German 2022 budget plans 100 billion in new borrowing – source
24/02/2022	German finance minister says aims to stop downward trend of defence budget
13/12/2021	German cabinet passes climate fund booster with 60 bln euro extra budget
10/12/2021	Germany to pass extra budget for more climate funds on Monday – FinMin
09/12/2021	Germany to pass extra budget on Monday for more climate funds – sources
09/09/2021	Merkel says her conservatives face hard battle after 16 years in power
23/04/2021	German parliament to vote on extra budget to fund COVID-19 response
23/04/2021	German lawmakers approve debt-fuelled extra budget to fight COVID-19
09/03/2021	Germany suspends COVID-19 aid for firms over suspected fraud
19/11/2020	Germany EU presidency looking for solution to EU budget row – minister
19/11/2020	German bond yields edge lower EU budget spat in focus
17/11/2020	German foreign minister upbeat that EU budget issue can be solved
10/07/2020	Germany's Scholz EU will reach compromise on budget recovery fund
04/06/2020	Germany set to pay 42% more into EU budget in coming years – report
04/06/2020	German stimulus package requires extra budget of some 25 bln euros – Scholz
09/05/2020	Germany needs another extra budget to cushion coronavirus impact – Merkel ally
23/03/2020	Late German virus budget lacks neighbourly gesture
18/03/2020	Merkel to address nation on coronavirus, Germany mulls new economic steps

Continued on next page

Table 1.A.1 continued

Date	Headline
13/01/2020	German coalition quarrels over how to spend record budget surplus
13/12/2019	German balanced budget policy should not be a fetish – Bundesbank chief
11/12/2019	Germany still wants future EU budget contribution cap of 1% of GDP
04/12/2019	Germany to end 2019 with big budget surplus that then shrinks – paper
21/10/2019	EU revises up German 2018 budget surplus next year's easing seen small
16/10/2019	Finland says both EU Commission and German caps on EU budget unrealistic
09/09/2019	Germany considers shadow budget to circumvent national debt rules – sources
03/09/2019	German budget lawmakers expect Scholz to stick to balanced budget goal
27/08/2019	Weaker exports hit German economy but budget surplus still high
20/08/2019	Germany could soon bury its black zero budget rule
24/06/2019	German budget to rise 1% next year no new debt until 2023
11/04/2019	German budget C/A surpluses damaging to euro zone – Moscovici
21/03/2019	German finmin No plans for Deutsche Bank/Commerzbank merger in my budget
20/03/2019	German finmin No precautions for Deutsche/Commerzbank merger taken in 2020 budget
18/03/2019	German 2020 budget plan calls for 1.7 pct boost in spending
20/11/2018	EU welcomes Franco-German budget idea more work needed
10/10/2018	Germany says it will only assess Italy's budget once Rome presents it to EU
16/07/2018	France Germany call for stable EU farm budget post-Brexit

Continued on next page

Table 1.A.1 continued

Date	Headline
28/06/2018	German budget committee approves balanced 2018 spending plan
31/05/2018	France confident of progress towards euro zone budget with Germany
31/05/2018	Germany's Scholz keeps distance from Macron's euro zone budget
09/03/2018	Euro debt supply – Five euro zone countries to issue bonds next week
31/01/2018	German far-right AfD lawmaker becomes chair of key budget committee
01/12/2017	German conservative says Macron plan for EU budget unrealistic – Focus
29/10/2017	Germany to confirm size of expected budget surplus next month
28/10/2017	German federal budget surplus could reach 14 bln euros-magazine
24/10/2017	German parties in coalition talks agree to stick to balanced budget
22/09/2017	German liberals draw red line for Merkel on euro zone budget
24/07/2017	Euro zone budget savings could complicate ECB rate hikes Bundesbank
09/03/2017	Schaeuble says Germany can avoid new borrowing in next legislative period
15/10/2016	German finmin wants ESM to be budget watchdog for euro zone states – newspaper
10/09/2016	Germany's EU budget bill may rise by 4.5 bln euros after Brexit – report
09/09/2016	Germany's Schaeuble wants quick action on tax relief
09/09/2016	German EU insiders urge fiscal discipline ahead of "Club Med" Summit
24/08/2016	Record German budget surplus fuels investment debate
01/04/2016	Concerns raised about German spending despite balanced budget
09/03/2016	Greece lenders resume bailout talks of fiscal gap, reforms
06/03/2016	Germany rejects calls to give Greece more time for budget goals
14/01/2016	Germany aims for balanced budget in 2016 but not certain – Merkel

1.B Identification methodology

1.B.1 Determinacy up to scale by IV conditions

To see this, let

$$\mathbf{S} = \begin{pmatrix} s_{1,1} & \mathbf{s}_{1,2:n} \\ \mathbf{s}_{2:n,1} & \mathbf{S}_{22} \end{pmatrix}, \ \mathbf{u}_t = \begin{pmatrix} u_{1,t} \\ \mathbf{u}_{2:n,t} \end{pmatrix}, \text{ and } \boldsymbol{\varepsilon}_t = \begin{pmatrix} \varepsilon_{1,t} \\ \boldsymbol{\varepsilon}_{2:n,t} \end{pmatrix}. \tag{1.14}$$

Given the structural relationship between reduced-form innovations $\mathbf{u}_t$ and structural shocks $\boldsymbol{\varepsilon}_t$, we have

$$\begin{aligned} u_{1,t} &= s_{1,1}\varepsilon_{1,t} + \mathbf{s}_{1,2:n}\boldsymbol{\varepsilon}_{2:n,t} \\ \mathbf{u}_{2:n,t} &= \mathbf{s}_{2:n,1}\varepsilon_{1,t} + \mathbf{S}_{22}\boldsymbol{\varepsilon}_{2:n,t} \end{aligned} \tag{1.15}$$

Using the instrument z_t with the properties $\mathbb{E}[z_t\varepsilon_{1,t}] = \alpha$ and $\mathbb{E}[z_t\boldsymbol{\varepsilon}_{2:n,t}] = \mathbf{0}$, we take the expectations

$$\begin{aligned} \mathbb{E}[z_t u_{1,t}] &= \mathbb{E}[z_t(s_{1,1}\varepsilon_{1,t} + \mathbf{s}_{1,2:n}\boldsymbol{\varepsilon}_{2:n,t})] \\ &= s_{1,1}\mathbb{E}[z_t\varepsilon_{1,t}] + \mathbf{s}_{1,2:n}\mathbb{E}[z_t\boldsymbol{\varepsilon}_{2:n,t}] \\ &= s_{1,1}\alpha \end{aligned} \tag{1.16}$$

and

$$\begin{aligned} \mathbb{E}[z_t\mathbf{u}_{2:n,t}] &= \mathbb{E}[z_t(\mathbf{s}_{2:n,1}\varepsilon_{1,t} + \mathbf{S}_{22}\boldsymbol{\varepsilon}_{2:n,t})] \\ &= \mathbf{s}_{2:n,1}\mathbb{E}[z_t\varepsilon_{1,t}] + \mathbf{S}_{22}\mathbb{E}[z_t\boldsymbol{\varepsilon}_{2:n,t}] \\ &= \mathbf{s}_{2:n,1}\alpha \end{aligned} \tag{1.17}$$

By dividing the two expectations, the unobserved covariance α cancels out, identifying the relative impact vector up to sign and scale

$$\frac{\mathbf{s}_{2:n,1}}{s_{1,1}} = \frac{\mathbb{E}[z_t \mathbf{u}_{2:n,t}]}{\mathbb{E}[z_t u_{1,t}]}. \tag{1.18}$$

1.B.2 Derivation of the scale under unit variance

The relationship between the reduced-form variance-covariance matrix Σ_u and the structural impact matrix $\mathbf{S}$ is given by:

$$\Sigma_u = \mathbf{S}\Sigma_\varepsilon \mathbf{S}' \tag{1.19}$$

Under the normalisation that the structural shocks have unit variance, we set $\Sigma_\varepsilon = \mathbb{I}$, implying $\Sigma_u = \mathbf{SS}'$. We partition the matrices into the variable of interest (indexed by 1) and the remaining variables.

$$\begin{pmatrix} \Sigma_{11} & \Sigma_{12} \\ \Sigma_{21} & \Sigma_{22} \end{pmatrix} = \begin{pmatrix} s_{11} & \mathbf{s}_{12} \\ \mathbf{s}_{21} & \mathbf{S}_{22} \end{pmatrix} \begin{pmatrix} s_{11} & \mathbf{s}'_{21} \\ \mathbf{s}'_{12} & \mathbf{S}'_{22} \end{pmatrix} \tag{1.20}$$

All elements on the left hand side are known from the VAR estimation. The elements on the right hand side are a priori unknown, but s_{11} and $\mathbf{s}_{21}$ can be derived from the known variables. The system (1.20) yields the following set of equations:

$$\Sigma_{11} = s_{11}^2 + \mathbf{s}_{12}\mathbf{s}'_{12} \tag{1.21}$$

$$\Sigma_{21} = \mathbf{s}_{21}s_{11} + \mathbf{S}_{22}\mathbf{s}'_{12} \tag{1.22}$$

$$\Sigma_{12} = s_{11}\mathbf{s}'_{21} + \mathbf{s}_{12}\mathbf{S}'_{22} \tag{1.23}$$

$$\Sigma_{22} = \mathbf{s}_{21}\mathbf{s}'_{21} + \mathbf{S}_{22}\mathbf{S}'_{22} \tag{1.24}$$

Note that Equation (1.21) immediately pins down s_{11} provided that we know $\mathbf{s}_{12}\mathbf{s}'_{12}$. So the only thing left to show is how to infer $\mathbf{s}_{12}\mathbf{s}'_{12}$ from observables. From the 2SLS regression of

$\mathbf{u}_{2:n,t}$ on $\hat{u}_{1,t}$, we obtain a consistent estimate of the relative impact vector[11]

$$\mathbf{b} = \frac{\mathbf{s}_{21}}{s_{11}} \Leftrightarrow \mathbf{s}_{21} = \mathbf{b}s_{11} \tag{1.25}$$

Using this estimate, equation (1.22) becomes

$$\Sigma_{21} = s_{11}^2\mathbf{b} + \mathbf{S}_{22}\mathbf{s}_{12}' \tag{1.26}$$

Now, subtracting equation (1.21) left-multiplied by $\mathbf{b}$ from (1.26) yields

$$\begin{aligned}\Sigma_{21} - \mathbf{b}\Sigma_{11} &= s_{11}^2\mathbf{b} + \mathbf{S}_{22}\mathbf{s}_{12}' - \mathbf{b}(s_{11}^2 + \mathbf{s}_{12}\mathbf{s}_{12}') \\ &= \mathbf{S}_{22}\mathbf{s}_{12}' - \mathbf{b}\mathbf{s}_{12}\mathbf{s}_{12}' \\ &= (\mathbf{S}_{22} - \mathbf{b}\mathbf{s}_{12})\mathbf{s}_{12}', \end{aligned} \tag{1.27}$$

which implies $\mathbf{s}_{12}' = (\mathbf{S}_{22} - \mathbf{b}\mathbf{s}_{12})^{-1}(\Sigma_{21} - \mathbf{b}\Sigma_{11})$. Denoting $\mathbf{Q} \equiv (\mathbf{S}_{22} - \mathbf{b}\mathbf{s}_{12})(\mathbf{S}_{22} - \mathbf{b}\mathbf{s}_{12})'$, we can rewrite $\mathbf{s}_{12}\mathbf{s}_{12}'$ as

$$\mathbf{s}_{12}\mathbf{s}_{12}' = (\Sigma_{21} - \mathbf{b}\Sigma_{11})'\mathbf{Q}^{-1}(\Sigma_{21} - \mathbf{b}\Sigma_{11}). \tag{1.28}$$

As $\Sigma_{21} - \mathbf{b}\Sigma_{11}$ is observable, it remains to show that $\mathbf{Q}$ can be computed from observables. Expanding $\mathbf{Q}$ yields

$$\mathbf{Q} = \mathbf{S}_{22}\mathbf{S}_{22}' - \mathbf{b}\mathbf{s}_{12}\mathbf{S}_{22}' - \mathbf{S}_{22}\mathbf{s}_{12}'\mathbf{b}' + \mathbf{b}(\mathbf{s}_{12}\mathbf{s}_{12}')\mathbf{b}'. \tag{1.29}$$

[11]This impact vector $\mathbf{b}$ is exactly what I have denoted by β_{IV} earlier in this paper. For the sake of readability, I am just now going to call it $\mathbf{b}$.

Substituting equation (1.25) into (1.23), (1.22), and (1.24), give

$$\Sigma_{21} = \mathbf{b}s_{11}^2 + \mathbf{S}_{22}\mathbf{s}_{12}' \tag{1.30}$$

$$\Sigma_{12} = s_{11}^2\mathbf{b}' + \mathbf{s}_{12}\mathbf{S}_{22}' \tag{1.31}$$

$$\Sigma_{22} = s_{11}^2\mathbf{b}\mathbf{b}' + \mathbf{S}_{22}\mathbf{S}_{22}'. \tag{1.32}$$

Plugging equations (1.21), (1.31), (1.30), and (1.32) into (1.29) gives

$$\mathbf{Q} = \Sigma_{22} - \mathbf{b}\Sigma_{12} - \Sigma_{21}\mathbf{b}' + \Sigma_{11}\,\mathbf{b}\mathbf{b}', \tag{1.33}$$

which is completely characterised by observables.[12] To conclude, $\mathbf{s}_{12}\mathbf{s}_{12}'$ can be computed by

$$\mathbf{s}_{12}\mathbf{s}_{12}' = (\Sigma_{21} - \mathbf{b}\Sigma_{11})'[\Sigma_{22} - \mathbf{b}\Sigma_{12} - \Sigma_{21}\mathbf{b}' + \Sigma_{11}\,\mathbf{b}\mathbf{b}']^{-1}(\Sigma_{21} - \mathbf{b}\Sigma_{11}), \tag{1.34}$$

and s_{11} follows immediately by computing

$$s_{1,1} = \sqrt{\Sigma_{11} - \mathbf{s}_{12}\mathbf{s}_{12}'}. \tag{1.35}$$

[12] To verify dimensional conformability: $\Sigma_{22} \in \mathbb{R}^{(N-1)\times(N-1)}$, $\mathbf{b}\Sigma_{12} \in \mathbb{R}^{(N-1)\times 1} \cdot \mathbb{R}^{1\times(N-1)} = \mathbb{R}^{(N-1)\times(N-1)}$, $\Sigma_{21}\mathbf{b}' \in \mathbb{R}^{(N-1)\times 1} \cdot \mathbb{R}^{1\times(N-1)} = \mathbb{R}^{(N-1)\times(N-1)}$, and $\Sigma_{11}\mathbf{b}\mathbf{b}' \in \mathbb{R} \cdot \mathbb{R}^{(N-1)\times(N-1)}$. All terms are $(N-1)\times(N-1)$.

1.C Robustness tests

1.C.1 Sample restrictions

A questions that frequently came up id if the results depend qualitatively on the post-2020 period. While it is true that the futures have reacted a lot stronger from 2022 onwards than earlier, in line with worries over the fiscal burden of Germany which are only a relatively recent phenomenon, the results are quantitively similar in the pre-2020 period.

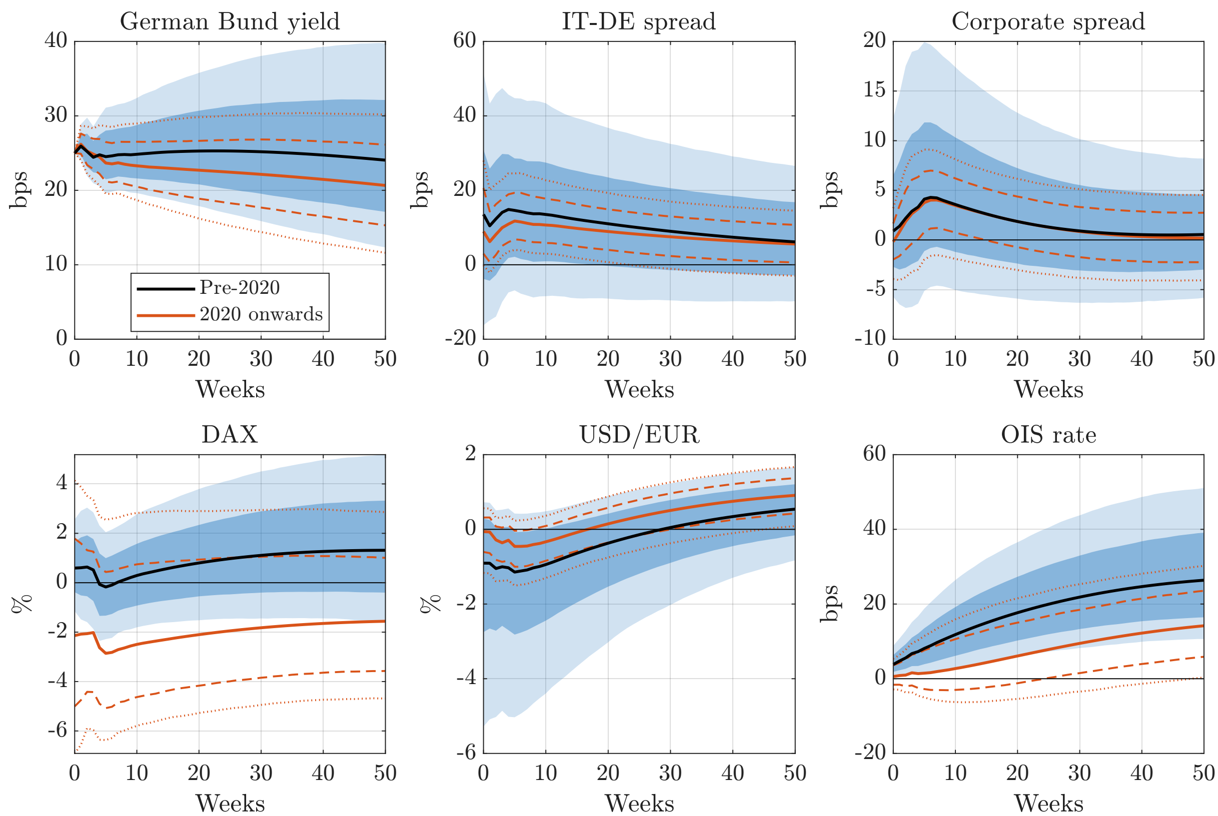

Figure 1.C.1. Comparison between pre- and post-2020 IRFs

Notes: Figures show the response of the six baseline variables to a fiscal shock in Germany calibrated to 25 bps. Confidence intervals are shown at the 68% and the 90% level and are constructed via moving block bootstrap (10,000 iterations, 4-week blocks), robust to heteroskedasticity and autocorrelation.

Another concern is that we are picking up shocks during the Covid phase that are not primarily attributable to the fiscal situation of Germany. Even though I already control for uncertainty in my refined instrument, I also run the baseline VAR excluding the year 2020 and the years 2020-2021. The instrument remains strong and the IRFs are almost unchanged.

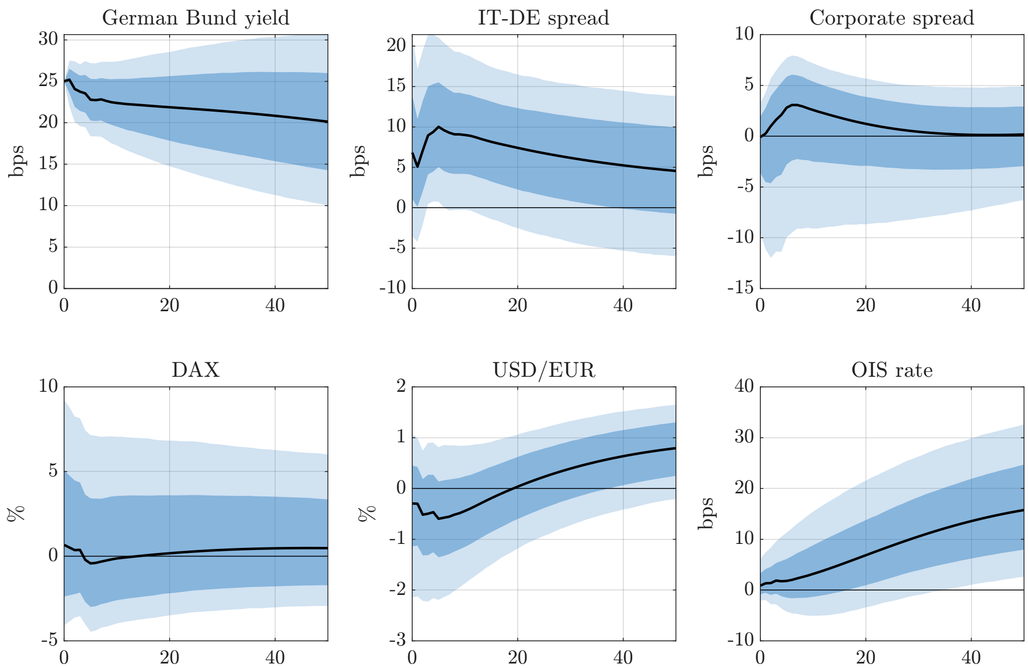

First stage regression: F: 19.89, robust F: 11.29, R^2: 3.75%, Adjusted R^2: 3.56%

Figure 1.C.2. Baseline – excluding 2020

Notes: Figure shows the response of the six baseline variables to a fiscal shock in Germany calibrated to 25 bps. Confidence intervals are shown at the 68% and the 90% level and are constructed via moving block bootstrap (10,000 iterations, 4-week blocks), robust to heteroskedasticity and autocorrelation.

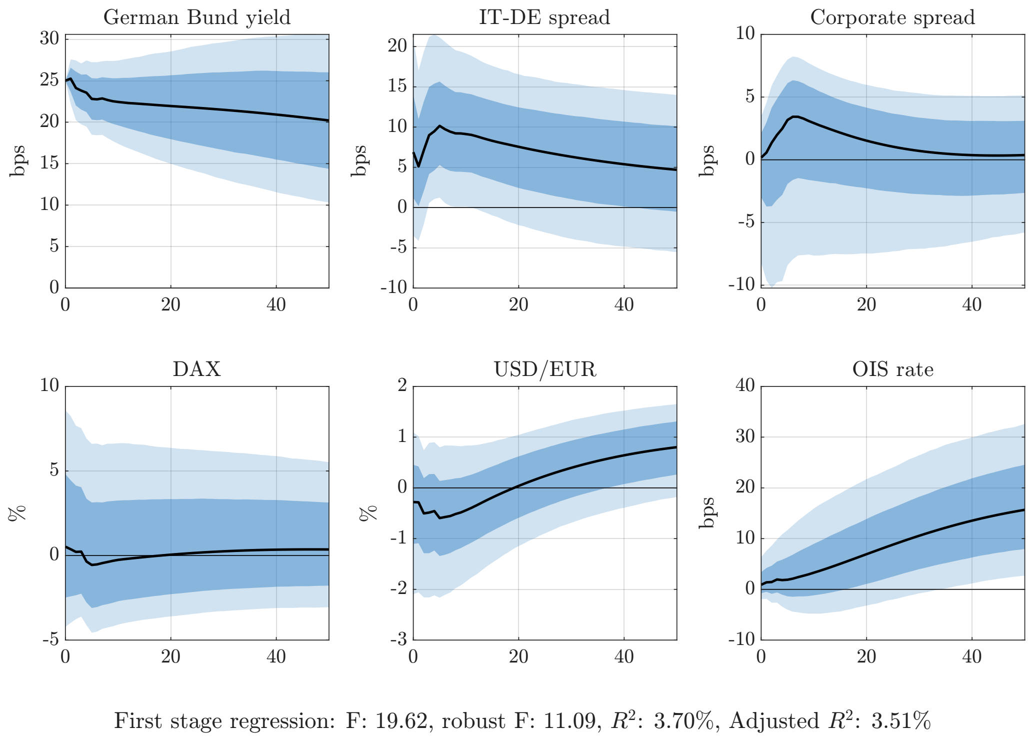

Figure 1.C.3. Baseline – excluding 2020–2021

Notes: Figure shows the response of the six baseline variables to a fiscal shock in Germany calibrated to 25 bps. Confidence intervals are shown at the 68% and the 90% level and are constructed via moving block bootstrap (10,000 iterations, 4-week blocks), robust to heteroskedasticity and autocorrelation.

1.C.2 Relevance of the instrument

To make sure that the results are not driven mechanically by the design of the instrument I am repeating the baseline regression here with the set of placebo days that we have created earlier. The IRfs look markedly different from the IRFs that we see in the baseline results. Virtually all confidence intervals are large and centred around zero.

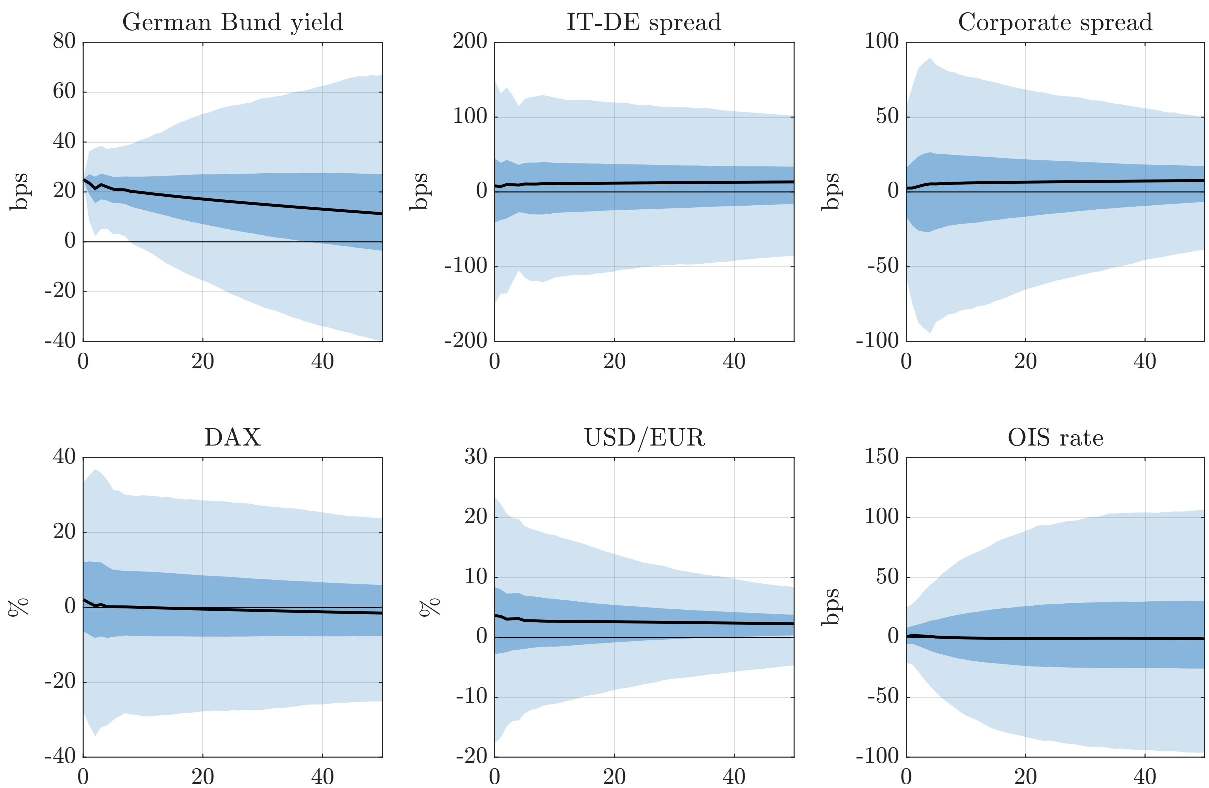

First stage regression: F: 2.16, robust F: 0.71, R^2: 0.42%, Adjusted R^2: 0.23%

Figure 1.C.4. Baseline with Placebo days

Notes: Figures show the response of the six baseline variables to a fiscal shock in Germany calibrated to 25 bps. Confidence intervals are shown at the 68% and the 90% level and are constructed via moving block bootstrap (10,000 iterations, 4-week blocks), robust to heteroskedasticity and autocorrelation.

1.C.3 Predictability of the instrument

Table 1.C.2. Predictability of fiscal shocks: First differences

Notes: OIS rates refer to the rate response around monetary policy events taken from Altavilla et al. (2019) to capture monetary policy shocks. Sovereign bond spreads refer to the difference in 10-year benchmark yields. The European sovereign CDS refers to the 5-year sovereign CDS index. Volatility is captured as the EURO STOXX 50 volatility index. German policy uncertainty is the German EPU index from Baker et al. (2016). Robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	Fiscal shock				
	(1) Monetary policy	(2) Sovereign credit	(3) Uncertainty	(4) Global factors	(5) All variables
Δ OIS 1M	0.02				0.02
	(0.017)				(0.016)
Δ OIS 1Y	-0.03				-0.03
	(0.017)				(0.018)
Δ OIS 5Y	0.03*				0.03*
	(0.018)				(0.018)
Δ OIS 10Y	-0.01				-0.01
	(0.017)				(0.018)
Δ Italy-Germany spread		-0.10			-0.02
		(0.091)			(0.097)
Δ Spain-Germany spread		0.22			0.12
		(0.178)			(0.162)
Δ Europe sov. CDS		-0.00			-0.00
		(0.004)			(0.004)
Δ Volatility			0.00		0.00
			(0.003)		(0.003)
Δ German policy uncertainty			-0.00		-0.00
			(0.000)		(0.000)
Δ US 10Y				-0.26	-0.25
				(0.166)	(0.177)
Δ Oil (Brent)				-0.00	-0.00
				(0.005)	(0.005)
Δ Gold				-0.00	-0.00
				(0.005)	(0.004)
N	513	513	513	513	513
R^2	0.01	0.00	0.00	0.01	0.03

Table 1.C.3. Predictability of fiscal shocks: Lagged variables

Notes: OIS rates refer to the rate response around monetary policy events taken from Altavilla et al. (2019) to capture monetary policy shocks. Sovereign bond spreads refer to the difference in 10-year benchmark yields. The European sovereign CDS refers to the 5-year sovereign CDS index. Volatility is captured as the EURO STOXX 50 volatility index. German policy uncertainty is the German EPU index from Baker et al. (2016). Robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	Fiscal shock				
	(1) Monetary policy	(2) Sovereign credit	(3) Uncertainty	(4) Global factors	(5) All variables
OIS 1M$_{t-1}$	-0.03				-0.02
	(0.029)				(0.024)
OIS 1Y$_{t-1}$	-0.00				-0.01
	(0.018)				(0.017)
OIS 5Y$_{t-1}$	-0.01				-0.01
	(0.023)				(0.021)
OIS 10Y$_{t-1}$	0.04				0.04
	(0.028)				(0.028)
Italy-Germany spread$_{t-1}$		0.02			0.01
		(0.013)			(0.014)
Spain-Germany spread$_{t-1}$		0.01			0.04
		(0.059)			(0.079)
Europe sov. CDS$_{t-1}$		-0.00			0.00
		(0.001)			(0.002)
Volatility$_{t-1}$			-0.00		-0.01
			(0.003)		(0.004)
German policy uncertainty$_{t-1}$			0.00		0.00
			(0.000)		(0.000)
US 10Y$_{t-1}$				0.00	-0.04
				(0.017)	(0.026)
Oil (Brent)$_{t-1}$				0.00	0.00
				(0.001)	(0.001)
Gold$_{t-1}$				-0.00	0.00
				(0.000)	(0.001)
N	513	513	513	513	513
R^2	0.02	0.00	0.01	0.00	0.03

1.C.4 Regime change test

To examine whether the transmission of German fiscal shocks to sovereign spreads depends on the fiscal vulnerability of the recipient country, I estimate a threshold proxy SVAR that allows the structural impact vector to vary across regimes defined by the interest rate–growth differential ($r - g$).

The reduced-form VAR is estimated on the full sample:

$$Y_t = c + A_1 Y_{t-1} + \cdots + A_p Y_{t-p} + u_t, \qquad u_t \sim (0, \Sigma), \tag{1.36}$$

where Y_t contains the German 10-year Bund yield, the sovereign spread (XX–DE), corporate spreads, the DAX, EUR/USD, and the OIS rate. I estimate a separate six-variable VAR for each country (France, Spain, Italy), substituting the relevant bilateral spread into the second position.

Identification proceeds via the external instrument approach of Mertens and Ravn (2013) and Stock and Watson (2012), but applied separately within each regime. I define the regime indicator using the *lagged* country-specific $r - g$ differential:

$$s_t = \mathbf{1}(r_{t-1} - g_{t-1} > \bar{\gamma}), \tag{1.37}$$

where r_t is the 10-year sovereign yield and g_t is the annualised quarterly real GDP growth rate. The threshold $\bar{\gamma}$ is set to the 25th percentile of $r - g$ over the proxy sample (2016–2025) in the baseline specification. The use of the lagged indicator avoids simultaneity between the regime classification and the contemporaneous shock.

Within each regime $s \in \{0, 1\}$, I partition the proxy-sample residuals $\hat{u}_t$ and the fiscal shock instrument m_t into the corresponding subsamples. The two-stage least squares identification—first-stage regression of $\hat{u}_{1,t}$ on m_t, second-stage projection of $\hat{u}_{-1,t}$ on $\hat{u}_{1,t}$—and the Mertens and Ravn (2013) covariance decomposition to recover the structural impact vector $b_1^{(s)}$ are performed separately for each regime. The impulse response functions are then computed using the common reduced-form VAR lag polynomial combined with the regime-specific impact vector $b_1^{(s)}$:

$$\text{IRF}_h^{(s)} = \Phi_h \, b_1^{(s)}, \qquad h = 0, 1, \ldots, H, \tag{1.38}$$

where Φ_h denotes the h-step moving average coefficient from the reduced-form VAR. This approach preserves degrees of freedom relative to estimating fully separate VARs while allowing the contemporaneous transmission of fiscal shocks to differ across regimes.

Confidence bands are constructed using the moving block bootstrap of Jentsch and Lunsford (2022), adapted to the threshold setting. In each bootstrap replication, the reduced-form VAR is re-estimated on the simulated data, and the regime-specific identification is performed within each subsample. Bootstrap draws for which identification is infeasible in either regime (e.g., due to insufficient non-zero proxy observations or singular covariance matrices) are discarded. The 68% and 90% confidence intervals are bias-corrected by re-centring the bootstrap quantiles around the point estimates.

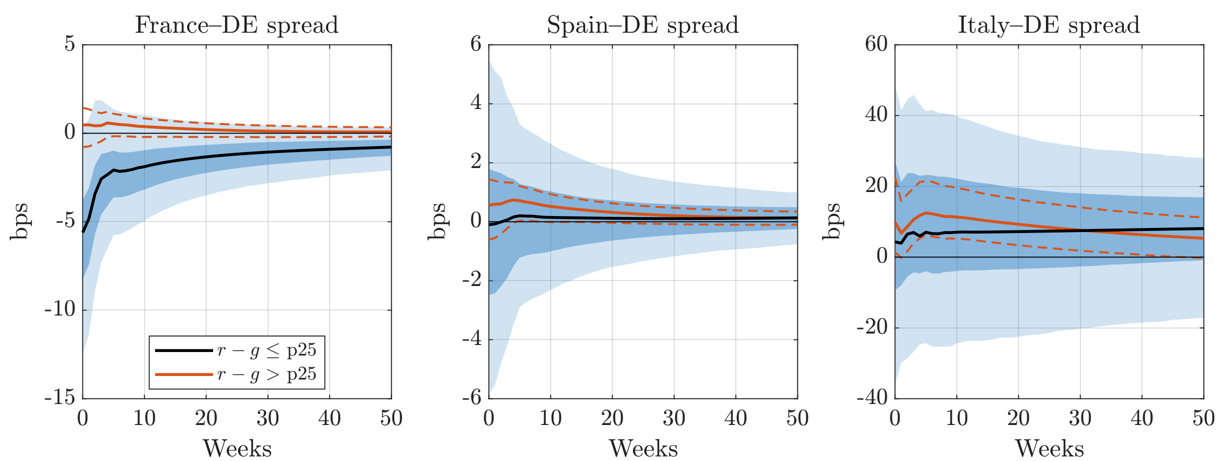

Figure 1.C.5. Amplification effects depending on debt sustainability levels
Notes: Confidence intervals are shown at the 68% and the 90% level and are constructed via moving block bootstrap (10,000 iterations, 4-week blocks), robust to heteroskedasticity and autocorrelation.

1.C.5 Comparison between LP-IV and external instrument approach

In this subsection I am comparing the impulse response functions of my baseline VAR with the more direct approach of using the constructed fiscal policy shock directly in local projections. The selected specification includes three lags of the fiscal shock and three lags of all variables in the VAR as controls (including the outcome variable and the other five variables).

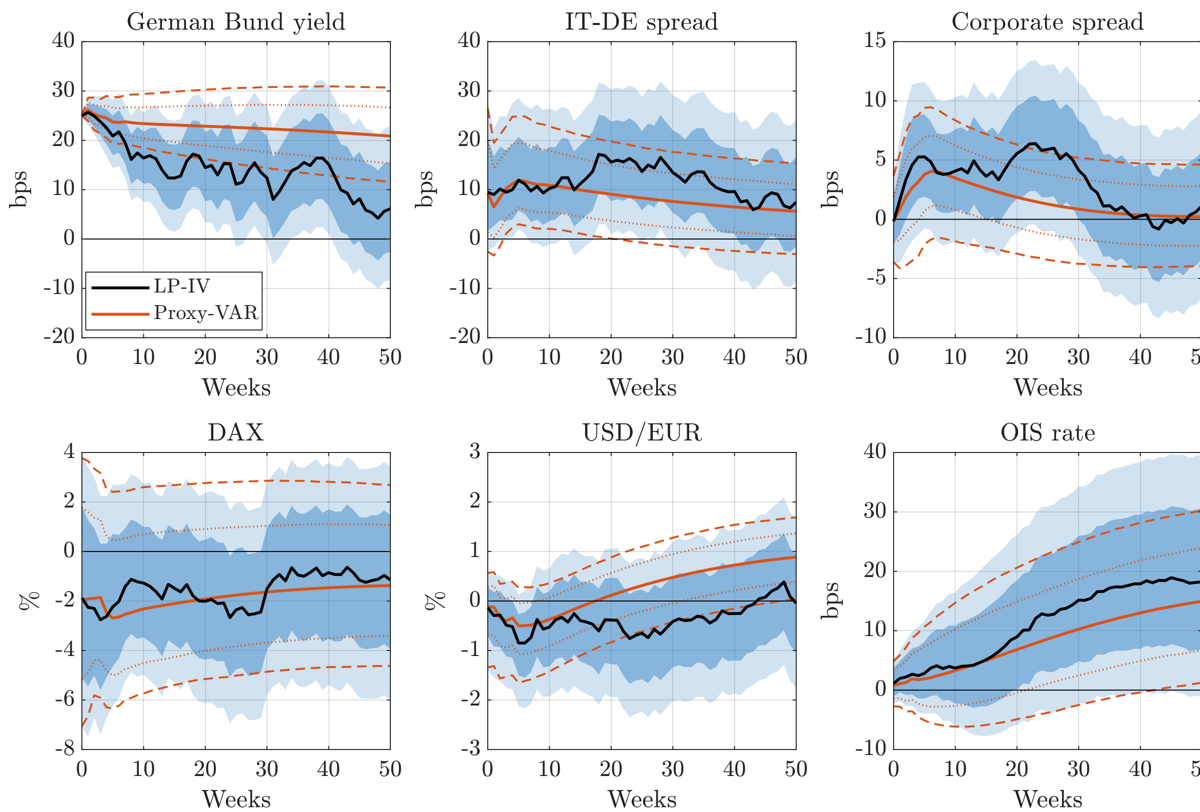

Figure 1.C.6. Local Projections and proxy-SVAR impulse responses

1.D Additional figures

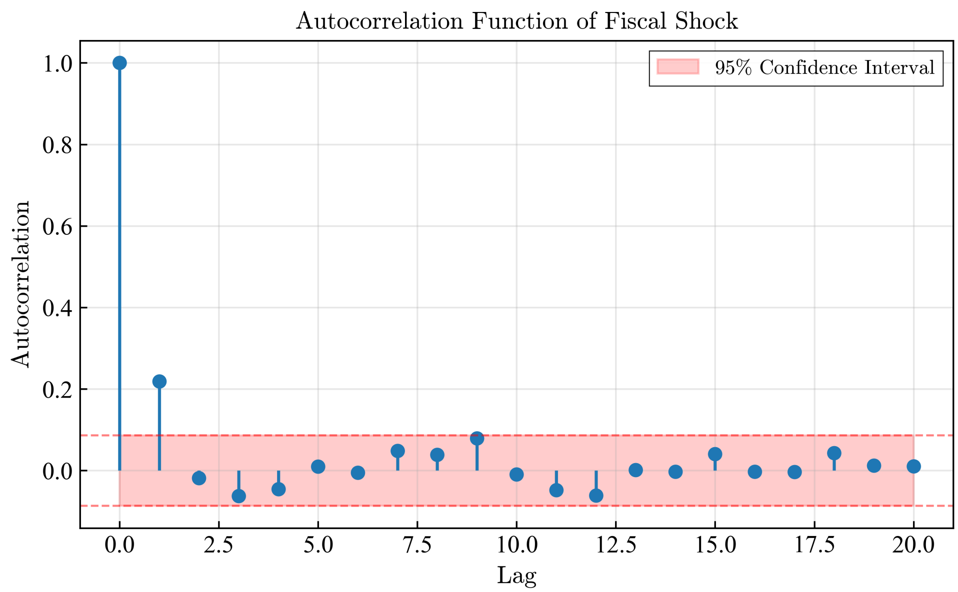

Figure 1.D.7. Autocorrelation function plot

Notes: This graph shows the autocorrelation coefficient of the fiscal news shock series at various lags. The confidence band tests whether one autocorrelation coefficient is statistically different from zero. If it falls within the red stripe, it is *not* statistically different form zero.

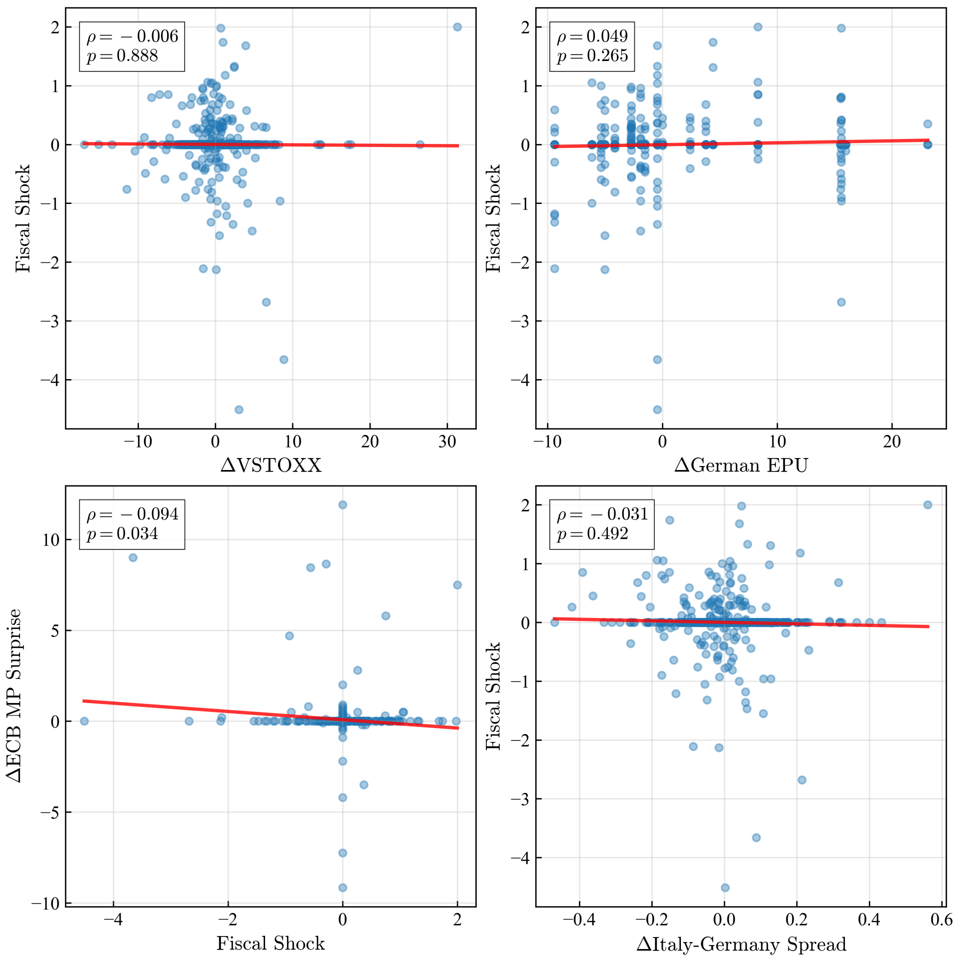

Figure 1.D.8. Correlation with other uncertainty measures

Notes: The graphs shows the correlation of the fiscal shock instrument with the European VIX index, the German economic policy index by Baker et al. (2016), the European monetary policy shock series by Altavilla et al. (2019), and the Italy-Germany spread in 10-year government bonds.

1.E Data sources

Table 1.E.4. Variables and data sources

Variable	Description	Source	Freq.
Panel A: Sovereign bond yields			
DE10YT=RR (BID_YIELD)	German 10-year benchmark yield	LSEG Refinitiv	D
FR10YT=RR (BID_YIELD)	French 10-year benchmark yield	LSEG Refinitiv	D
IT10YT=RR (BID_YIELD)	Italian 10-year benchmark yield	LSEG Refinitiv	D
ES10YT=RR (BID_YIELD)	Spanish 10-year benchmark yield	LSEG Refinitiv	D
US10YT=RR (BID_YIELD)	US 10-year benchmark yield	LSEG Refinitiv	D
FGBLc1 (TRDPRC_1)	Futures on German 10-year Bund	LSEG Refinitiv	D
Panel B: Interest rates & swaps			
EUREST3M=	OIS 3-month rate	LSEG Refinitiv	D
EUREST10Y=	OIS 10-year rate	LSEG Refinitiv	D
FM.M.U2.EUR.RT.MM.EURIBOR3MD_.HSTA	EURIBOR 3-month rate	ECB SDW	M
FM.B.U2.EUR.4F.KR.DFR.LEV	ECB deposit facility rate	ECB SDW	D
OIS_SW (Monetary Event Window)	OIS change in MP event window (bps)	CEPR	D
Panel C: Equity indices & volatility			
.GDAXI (TRDPRC_1)	Deutsche Boerse DAX Index	LSEG Refinitiv	D
.SX7E (TRDPRC_1)	EURO STOXX Banks Index	LSEG Refinitiv	D
.IBEXIB (TRDPRC_1)	IBEX 35 Banks Index (Spain)	LSEG Refinitiv	D

Continued on next page

Table 1.E.4 continued

Variable	Description	Source	Freq.
.FTIT3010E (TRDPRC_1)	FTSE Italia Banks Index	LSEG Refinitiv	D
.V2TX (TRDPRC_1)	VSTOXX (Euro STOXX 50 volatility)	LSEG Refinitiv	D
Panel D: Credit spreads & CDS			
.IBBEU003D	iBoxx EUR Corporates index	LSEG Refinitiv	D
.IBBEU0057	iBoxx EUR Corporates AAA index	LSEG Refinitiv	D
ITEFA5Y=R	iTraxx Europe Senior Financials CDS	LSEG Refinitiv	D
ITEFS5Y=R	iTraxx Europe Sub Financials CDS	LSEG Refinitiv	D
ITEUR5Y=R	iTraxx Europe Senior Financials CDS	LSEG Refinitiv	D
Panel E: Commodities & FX			
EUR=	USD/EUR spot exchange rate	LSEG Refinitiv	D
.BCOMGC	Bloomberg Gold price index	LSEG Refinitiv	D
LCOc1	ICE Europe Brent Crude Electronic Energy Future	LSEG Refinitiv	D
Panel F: Financial stress indicators			
OFR FSI	EU OFR Financial Stress Index	OFR	D
CISS.D.U2.Z0Z.4F.EC.SS_CI.IDX	Composite Indicator of Systemic Stress	ECB SDW	W
DEEPUINDXM	Economic Policy Uncertainty, Germany	FRED	M

Continued on next page

Table 1.E.4 continued

Variable	Description	Source	Freq.
Panel G: Banking sector data			
BSI.M.ES.N.A.A20.I.1.U2.2240.EUR.E	Loans 1–5 yrs to NFCs, Spain	ECB SDW	M
BSI.M.ES.N.A.A20.J.1.U2.2240.Z01.E	Loans >5 yrs to NFCs, Spain	ECB SDW	M
BSI.M.IT.N.A.A20.I.1.U2.2240.Z01.E	Loans 1–5 yrs to NFCs, Italy	ECB SDW	M
BSI.M.IT.N.A.A20.J.1.U2.2240.Z01.E	Loans >5 yrs to NFCs, Italy	ECB SDW	M
BSI.M.DE.N.A.A20.I.1.U2.2240.Z01.E	Loans 1–5 yrs to NFCs, Germany	ECB SDW	M
BSI.M.DE.N.A.A20.J.1.U2.2240.Z01.E	Loans >5 yrs to NFCs, Germany	ECB SDW	M
BSI.M.U2.N.A.A20.A.1.U2.2240.Z01.E	Total loans to EA NFCs by MFIs, stocks	ECB SDW	M
BSI.M.U2.N.A.T00.A.1.Z5.0000.Z01.E	Total assets of MFIs, stocks	ECB SDW	M
RAI.M.U2.SVLHHNFC.EUR.MIR.Z	Share of variable-rate loans to HH & NFCs	ECB SDW	M
MIR.M.U2.B.A2I.AM.R.A.2240.EUR.N	Cost of borrowing, NFCs	ECB SDW	M
MIR.M.U2.B.A2C.AM.R.A.2250.EUR.N	Cost of borrowing, households	ECB SDW	M
Panel H: Macroeconomic data			
NAMQ_10_GDP (DE, CLV_PCH_SM)	Real GDP, Y-o-Y growth, Germany	Eurostat	Q
NAMQ_10_GDP (ES, CLV_PCH_SM)	Real GDP, Y-o-Y growth, Spain	Eurostat	Q

Continued on next page

Table 1.E.4 continued

Variable	Description	Source	Freq.
NAMQ_10_GDP (IT, CLV_PCH_SM)	Real GDP, Y-o-Y growth, Italy	Eurostat	Q
Panel I: Calibration data (Table 1.4)			
EBA Transparency Exercise 2023	Domestic sovereign holdings of banks	EBA	H1 2023
CBD2.Q.{CC}.W0.11._Z._Z.A.A.I4008	CET1 ratio (CC = DE, ES, IT)	ECB CBD	Q
NAMA_10_GDP (CP_MEUR)	Annual nominal GDP	Eurostat	A
BSI.M.{CC}.N.A.A20.A.1.U2.2240.Z01.E	NFC loans, avg. 2016–2024 (CC = DE, ES, IT)	ECB SDW	M

Notes: Freq. column: D = daily, W = weekly, M = monthly, Q = quarterly, A = annual. Daily series converted to weekly (Friday) frequency; monthly and quarterly series forward-filled.

1.F Proofs

1.F.1 Proof of Lemma 1

Proof. The proof proceeds in four stages. First, we combine the Periphery equations into a single reduced-form expression. Second, we derive the sensitivity of Periphery credit to the Periphery yield through the bank balance sheet. Third, we totally differentiate the reduced-form expression and solve the fixed-point problem. Fourth, we establish the stability condition.

Begin with the three Periphery equations

$$r^P = r^C + \sigma^P \tag{1.39}$$

$$\sigma^P = \bar{\sigma} + \delta\left(r^C - g^P\right) \tag{1.40}$$

$$g^P = \psi \cdot \ell_P \tag{1.41}$$

Substitute (1.41) into (1.40)

$$\sigma^P = \bar{\sigma} + \delta\left(r^C - \psi\ell_P\right) \tag{1.42}$$

Substitute (1.42) into (1.39)

$$r^P = r^C + \bar{\sigma} + \delta\left(r^C - \psi\ell_P\right) = (1+\delta)\, r^C + \bar{\sigma} - \delta\psi\,\ell_P \tag{1.43}$$

This is the single equation governing the Periphery yield. The complication is that ℓ_P on the right-hand side depends on r^P through the bank balance sheet, creating a fixed-point problem. We now derive $\partial\ell_P/\partial r^P$ by tracing the chain of dependence $r^P \to q^P \to e_P \to \ell_P$.

Link 1: Periphery yield to bond price ($r^P \to q^P$). From the bond pricing equation $q^P = (1+r^P)^{-1}$, differentiate with respect to r^P:

$$\frac{dq^P}{dr^P} = -\frac{1}{(1+r^P)^2} \tag{1.44}$$

Link 2: Bond price to bank equity ($q^P \to e_P$). From the equity equation (1.15), with Periphery bond holdings b_P^P predetermined (legacy portfolio):

$$e_P = q^C b_P^C + q^P b_P^P + \ell_P - d_P$$

Partially differentiating with respect to q^P, holding all other terms constant:

$$\frac{\partial e_P}{\partial q^P} = b_P^P \tag{1.45}$$

Link 3: Bank equity to credit supply ($e_P \to \ell_P$). From the credit supply equation (1.19) when the IC constraint binds:

$$\ell_P = \frac{e_P - \phi\omega^P q^P b_P^P}{\phi\omega^\ell}$$

Partially differentiating with respect to e_P

$$\frac{\partial \ell_P}{\partial e_P} = \frac{1}{\phi\omega^\ell} \tag{1.46}$$

Note on the risk-weight offset. Credit supply also depends on q^P directly through the term $\phi\omega^P q^P b_P^P$ in the numerator of (1.19). The total partial derivative of ℓ_P with respect to q^P is therefore:

$$\frac{\partial \ell_P}{\partial q^P} = \frac{1}{\phi\omega^\ell} \cdot b_P^P - \frac{\omega^P}{\omega^\ell} \cdot b_P^P = \frac{b_P^P}{\phi\omega^\ell}\left(1 - \phi\omega^P\right) = \frac{\tilde{b}_P^P}{\phi\omega^\ell} \tag{1.47}$$

where we define the *effective Periphery bond exposure*:

$$\tilde{b}_P^P \equiv (1 - \phi\omega^P)\, b_P^P \tag{1.48}$$

This accounts for the fact that when q^P falls, equity drops by b_P^P but the IC constraint also loosens by $\phi\omega^P b_P^P$ (since the risk-weighted value of Periphery bond holdings decreases). The net tightening is proportional to $\tilde{b}_P^P$.

For the baseline case where Periphery sovereign bonds also carry zero regulatory risk weight ($\omega^P = 0$, as under Basel standardised approach for Euro area sovereigns), we have $\tilde{b}_P^P = b_P^P$ and the expressions simplify. We proceed with the general case.

Combining the links. Apply the chain rule:

$$\frac{\partial \ell_P}{\partial r^P} = \frac{\partial \ell_P}{\partial q^P} \cdot \frac{dq^P}{dr^P} = \frac{\tilde{b}_P^P}{\phi \omega^\ell} \cdot \left(-\frac{1}{(1+r^P)^2} \right) = -\frac{\tilde{b}_P^P}{\phi \omega^\ell (1+r^P)^2} \tag{1.49}$$

This is negative: a higher Periphery yield reduces bond prices, erodes bank equity (net of the IC offset), and tightens the leverage constraint, cutting credit supply.

Totally differentiate the reduced-form equation (1.43) with respect to the exogenous impulse dr^C:

$$dr^P = (1+\delta)\, dr^C - \delta\psi\, d\ell_P \tag{1.50}$$

The credit change $d\ell_P$ has two sources: a direct effect from the Core yield change (through Core bond losses in the Periphery bank's portfolio) and a feedback effect from the Periphery yield change (through the doom loop). Decompose:

$$d\ell_P = \underbrace{\frac{\partial \ell_P}{\partial r^C}\, dr^C}_{\text{direct effect}} + \underbrace{\frac{\partial \ell_P}{\partial r^P}\, dr^P}_{\text{feedback effect}} \tag{1.51}$$

Substitute (1.51) into (1.50):

$$dr^P = (1+\delta)\, dr^C - \delta\psi \left[\frac{\partial \ell_P}{\partial r^C}\, dr^C + \frac{\partial \ell_P}{\partial r^P}\, dr^P \right] \tag{1.52}$$

Collect all terms in dr^P on the left-hand side:

$$dr^P + \delta\psi \frac{\partial \ell_P}{\partial r^P}\, dr^P = (1+\delta)\, dr^C - \delta\psi \frac{\partial \ell_P}{\partial r^C}\, dr^C \tag{1.53}$$

$$dr^P \left[1 + \delta\psi \frac{\partial \ell_P}{\partial r^P}\right] = \left[(1+\delta) - \delta\psi \frac{\partial \ell_P}{\partial r^C}\right] dr^C \tag{1.54}$$

Substitute the expression for $\partial \ell_P / \partial r^P$ from (1.49):

$$dr^P \left[1 - \delta\psi \cdot \frac{\tilde{b}_P^P}{\phi\omega^\ell (1+r^P)^2}\right] = \left[(1+\delta) - \delta\psi \frac{\partial \ell_P}{\partial r^C}\right] dr^C \tag{1.55}$$

The expression in brackets on the left-hand side is $1 - \lambda$, where:

$$\lambda = \delta \cdot \psi \cdot \frac{1}{\phi\omega^\ell} \cdot \frac{\tilde{b}_P^P}{(1+r^P)^2} \tag{1.56}$$

which is the product of four terms: the spread sensitivity to $r - g$ (δ), the credit-growth elasticity (ψ), the leverage multiplier $(\phi\omega^\ell)^{-1}$, and the duration-weighted effective bond exposure $\tilde{b}_P^P(1 + r^P)^{-2}$.

Solving for the pass-through:

$$\frac{dr^P}{dr^C} = \frac{(1+\delta) - \delta\psi \frac{\partial \ell_P}{\partial r^C}}{1-\lambda} \tag{1.57}$$

The numerator captures the direct pass-through: the $(1+\delta)$ term is the mechanical effect of r^C on r^P through the spread function (a 1 bp increase in r^C raises r^P by $(1+\delta)$ bps directly), and the $\frac{\partial \ell_P}{\partial r^C}$ term captures the direct credit contraction from Core bond losses in the Periphery bank's portfolio. When the latter is small relative to $(1+\delta)$ the numerator simplifies to $(1+\delta)$, yielding:

$$\frac{dr^P}{dr^C} = \frac{1+\delta}{1-\lambda} = (1+\delta) \cdot \underbrace{\frac{1}{1-\lambda}}_{\mathcal{A}} \tag{1.58}$$

where:

$$\mathcal{A} = \frac{1}{1 - \underbrace{\delta \cdot \psi}_{\text{spread-credit -growth}} \cdot \underbrace{\frac{1}{\phi\omega^\ell}}_{\text{leverage}} \cdot \underbrace{\frac{\tilde{b}_P^P}{(1+r^P)^2}}_{\text{bond exposure}}} \tag{1.59}$$

The amplification factor $\mathcal{A}$ is well-defined and finite if and only if $\lambda \neq 1$. Since all constituent terms of λ are non-negative, we have $\lambda \geq 0$. There are three cases:

1. **Stable amplification** ($0 \leq \lambda < 1$): $\mathcal{A} = (1-\lambda)^{-1} > 1$. The feedback loop amplifies the initial shock but converges. The amplification factor admits the geometric series representation:

$$\mathcal{A} = \sum_{k=0}^{\infty} \lambda^k = 1 + \lambda + \lambda^2 + \lambda^3 + \cdots \tag{1.60}$$

   Each term λ^k corresponds to the k-th round of the feedback loop: r^P rises $\to$ q^P falls $\to$ e_P falls $\to$ ℓ_P falls $\to$ g^P falls $\to$ σ^P rises $\to$ r^P rises again. The parameter λ is the fraction of the shock that survives each round trip through the loop. Since $\lambda < 1$, each successive round is smaller than the last, and the infinite sum converges.

2. **Crisis threshold** ($\lambda = 1$): $\mathcal{A} \to \infty$. The feedback loop is explosive—each round of the loop reproduces the full initial shock, leading to unbounded amplification. This corresponds to a self-fulfilling sovereign debt crisis in which an arbitrarily small perturbation in r^C causes a discrete jump to a bad equilibrium.

3. **Unstable region** ($\lambda > 1$): No stable equilibrium exists. The system is beyond the crisis threshold.

Since all parameters are positive, $\lambda > 0$ and therefore $\mathcal{A} > 1$: the doom loop always amplifies. The model nests both the stable amplification region (normal times) and the crisis region, with $\lambda = 1$ as the boundary. □

1.F.2 Proof of Proposition 1

Proof. The proof proceeds in three stages. First, we derive the Core output response. Second, we derive the Periphery output response. Third, we aggregate and establish the properties. All comparative statics are first-order approximations around the pre-shock equilibrium; since the production function is concave in credit ($\alpha < 1$), the linearisation understates the true output loss for discrete credit contractions.

The Core output response to a fiscal expansion $\Delta G^C > 0$ has two components: the textbook demand channel and the financial drag. We derive the financial drag by tracing the transmission chain.

From the yield determination equation (1.20), the fiscal news shock raises the Core yield:

$$\Delta r^C = \mu \cdot \Delta G^C \tag{1.61}$$

By a first-order Taylor approximation of the bond pricing equation $q^C = (1 + r^C)^{-1}$, the Core bond price falls:

$$\Delta q^C \approx -\frac{1}{(1 + r^C)^2} \cdot \Delta r^C = -\mathcal{D}^C \cdot \Delta r^C \tag{1.62}$$

To derive the Core equity loss, I assume that Core banks hold only Core sovereign bonds ($b_C^P = 0$), reflecting the strong home bias documented in the EBA Transparency Exercise data. Core banks' periphery holdings are empirically negligible relative to their Bund portfolios.[13] Under this assumption, equation (1.15) gives the Core bank's equity loss as:

$$\Delta \hat{e}_C = b_C^C \cdot \Delta q^C = -\mathcal{D}^C \cdot b_C^C \cdot \Delta r^C \tag{1.63}$$

From the loan supply equation (1.19), with the incentive compatibility constraint (1.18) binding

[13] This assumption is conservative: if $b_C^P > 0$, the Core bank would absorb additional mark-to-market losses from periphery bond repricing, increasing F_C and strengthening the result.

and $b_C^P = 0$, the credit contraction is:

$$\Delta \ell_C = \Lambda \cdot \Delta \hat{e}_C = -\Lambda \cdot \mathcal{D}^C \cdot b_C^C \cdot \Delta r^C \tag{1.64}$$

By a first-order Taylor approximation of the production function (1.14) around the pre-shock credit level ℓ_C, the output loss through the financial channel is:

$$\Delta y_C^{\text{fin}} = \alpha A_C \ell_C^{\alpha-1} \cdot \Delta \ell_C = -\alpha A_C \ell_C^{\alpha-1} \cdot \Lambda \cdot \mathcal{D}^C \cdot b_C^C \cdot \Delta r^C \tag{1.65}$$

Substituting (1.61) and recognising F_C from Definition 2:

$$\Delta y_C^{\text{fin}} = -F_C \cdot \Delta G^C \tag{1.66}$$

Combining the textbook demand channel ($\mathcal{M}_C$) with the financial drag:

$$\frac{\Delta y_C}{\Delta G^C} = \mathcal{M}_C - F_C \tag{1.67}$$

The Periphery receives no direct fiscal stimulus—the spending occurs in the Core. Therefore there is no textbook multiplier term, and the output response operates entirely through the financial channel. We derive it by tracing the credit contraction.

From the equity equation (1.15), the Periphery bank's equity loss reflects mark-to-market losses on both Core and Periphery bond holdings:

$$\Delta \hat{e}_P = b_P^C \cdot \Delta q^C + b_P^P \cdot \Delta q^P = -\mathcal{D}^C \cdot b_P^C \cdot \Delta r^C - \mathcal{D}^P \cdot b_P^P \cdot \Delta r^P \tag{1.68}$$

From the loan supply equation (1.19), totally differentiating with b_P^P predetermined so that

$\Delta(q^P b^P_P) = b^P_P \, \Delta q^P = -\mathcal{D}^P b^P_P \, \Delta r^P$:

$$\Delta \ell_P = \Lambda \left[\Delta \hat{e}_P - \phi \omega^P \, \Delta(q^P b^P_P) \right] = \Lambda \left[\Delta \hat{e}_P + \phi \omega^P \mathcal{D}^P b^P_P \, \Delta r^P \right] \tag{1.69}$$

Substituting (1.68) and collecting terms:

$$\Delta \ell_P = -\Lambda \left[\mathcal{D}^C b^C_P \cdot \Delta r^C + (1 - \phi \omega^P) \mathcal{D}^P b^P_P \cdot \Delta r^P \right] = -\Lambda \left[\mathcal{D}^C b^C_P \cdot \Delta r^C + \mathcal{D}^P \tilde{b}^P_P \cdot \Delta r^P \right] \tag{1.70}$$

where $\tilde{b}^P_P = (1 - \phi \omega^P) b^P_P$ is the effective Periphery bond exposure from Definition 1. Now apply the result from Lemma 1: $\Delta r^P = \Gamma \cdot \Delta r^C$ where $\Gamma = (1 + \delta) \cdot \mathcal{A}$. Substituting:

$$\Delta \ell_P = -\Lambda \left[\mathcal{D}^C b^C_P + \mathcal{D}^P \tilde{b}^P_P \cdot \Gamma \right] \Delta r^C \tag{1.71}$$

By a first-order Taylor approximation of the production function (1.14) around the pre-shock credit level ℓ_P:

$$\Delta y_P = \alpha A_P \ell_P^{\alpha-1} \cdot \Delta \ell_P = -\alpha A_P \ell_P^{\alpha-1} \cdot \Lambda \left[\mathcal{D}^C b^C_P + \mathcal{D}^P \tilde{b}^P_P \cdot \Gamma \right] \Delta r^C \tag{1.72}$$

Substituting $\Delta r^C = \mu \cdot \Delta G^C$ and recognising F_P from Definition 2:

$$\frac{\Delta y_P}{\Delta G^C} = -F_P \tag{1.73}$$

Note the sign: the Periphery output response to a Core fiscal expansion is unambiguously negative. The Periphery absorbs the full financial transmission (both the direct Core bond losses and the doom-loop-amplified Periphery bond losses) without receiving any fiscal stimulus.

The effective union-wide fiscal multiplier is the GDP-weighted sum of country-specific output responses:

$$\hat{\mathcal{M}}^{MU} = n \cdot \frac{\Delta y_C}{\Delta G^C} + (1-n) \cdot \frac{\Delta y_P}{\Delta G^C} \tag{1.74}$$

Substituting (1.67) and (1.73):

$$\hat{\mathcal{M}}^{MU} = n \cdot (\mathcal{M}_C - F_C) + (1-n) \cdot (-F_P) = n \cdot \mathcal{M}_C - n \cdot F_C - (1-n) \cdot F_P \tag{1.75}$$

which establishes (1.33). □

1.F.3 Existence and uniqueness of equilibrium

Proposition 3. *If $\lambda < 1$, the model has a unique equilibrium.*

Proof. Since $r_i^\ell > \bar{r}$, each bank's objective is strictly increasing in ℓ_i, so the incentive compatibility constraint (1.18) binds and credit supply is given by (1.19). Substituting (1.19) and the growth equation into the spread equation yields a single fixed-point equation in r^P

$$r^P = \Phi(r^P) \equiv (1+\delta)\, r^C + \bar{\sigma} - \delta\psi\, \Lambda \left[\hat{e}_P(r^P) - \phi\omega^P \hat{q}^P(r^P)\, b_P^P \right], \tag{1.76}$$

where $\hat{e}_P$ and $\hat{q}^P$ depend on r^P through the bond pricing and equity equations. Differentiating yields

$$\Phi'(r^P) = \delta\psi \cdot \Lambda \cdot \mathcal{D}^P \cdot \tilde{b}_P^P = \lambda. \tag{1.77}$$

Since $\lambda < 1$ by assumption, $|\Phi'| < 1$ everywhere, so Φ is a contraction on $\mathbb{R}_+$. By the Banach fixed-point theorem, a unique fixed point exists. All remaining variables are then uniquely determined. When $\lambda \geq 1$, the mapping is no longer a contraction and no stable equilibrium exists, corresponding to the crisis threshold of Lemma 1. □

1.G Italian fiscal events

Table 1.G.5. Italian Fiscal Policy News Headlines

Date	Headline
28/12/2025	Italian Parliament approves 2026 Budget Law targeting 2.8% deficit
15/10/2025	Italy submits 2026 Draft Budgetary Plan: deficit target 2.6% for 2026
30/09/2025	Italy NADEF 2025: deficit target confirmed at 3.3% for 2025, 2.8% for 2026
09/04/2025	Italy DEF 2025: government confirms consolidation path, deficit target 3.0% for 2025
11/03/2025	European Commission puts Italy's EDP on hold, citing positive fiscal trajectory
28/02/2025	ISTAT reports 2024 deficit at 3.4% of GDP, better than 3.8% target
21/01/2025	EU Council adopts EDP recommendation: Italy must end excessive deficit by 2026
28/12/2024	Italian Parliament approves 2025 Budget Law
22/11/2024	European Commission opinion: Italy's 2025 budget broadly in line with EDP recommendation
15/10/2024	Italy submits 2025 Draft Budgetary Plan: deficit below 3% by 2026
27/09/2024	Italy fiscal-structural plan: targets deficit 3.8% in 2024, 3.3% in 2025, 2.8% in 2026
20/09/2024	Italy submits medium-term fiscal-structural plan to European Commission
26/07/2024	EU Council formally opens Excessive Deficit Procedure against Italy

Continued on next page

Table 1.G.5 continued

Date	Headline
19/06/2024	European Commission recommends EDP for Italy based on 2023 deficit of 7.4%
09/04/2024	Italy DEF 2024: skeletal update, defers detailed fiscal path to autumn amid new EU rules
03/04/2024	Economy Minister Giorgetti confirms EU will open Excessive Deficit Procedure against Italy
19/03/2024	Eurostat validates Italy's 2023 deficit at 7.4% of GDP following further Superbonus revisions
01/03/2024	ISTAT reports 2023 deficit at 7.2% of GDP (vs 5.3% target), largely due to Superbonus
29/12/2023	Italian Parliament approves 2024 Budget Law: IRPEF reform, tax wedge cut made permanent
22/11/2023	European Commission opinion: Italy's 2024 budget at risk of non-compliance
16/10/2023	Italy submits 2024 Draft Budgetary Plan: spending restraint with limited new measures
27/09/2023	Italy NADEF 2023: deficit target revised sharply upward to 5.3% for 2023; 4.3% for 2024
22/04/2023	Eurostat final data: 2022 Italy deficit revised to 8.6% of GDP
11/04/2023	Italy DEF 2023: deficit target 4.5% for 2023, acknowledges Superbonus fiscal impact

Continued on next page

Table 1.G.5 continued

Date	Headline
23/03/2023	Eurostat reclassifies Superbonus tax credits as payable, further increasing Italy's recorded deficit
01/03/2023	ISTAT revises 2022 deficit from 5.6% to 8.0% due to Superbonus accounting reclassification
29/12/2022	Italian Parliament approves 2023 Budget Law
22/11/2022	European Commission opinion: Italy's 2023 budget not fully in line with Council recommendation
17/11/2022	Italy submits 2023 Draft Budgetary Plan with 4.5% deficit target
01/11/2022	Meloni government presents NADEF Update: deficit target 4.5% for 2023
06/04/2022	Italy DEF 2022: deficit target 5.6% for 2022, debt/GDP forecast 147%
01/03/2022	ISTAT reports 2021 deficit at 7.2% of GDP (revised from 9.4% target)
30/12/2021	Italian Parliament approves 2022 Budget Law: IRPEF reform, Superbonus extension
24/11/2021	European Commission opinion: Italy's 2022 budget broadly in line with recommendation
15/10/2021	Italy submits 2022 Draft Budgetary Plan: deficit target 5.6% for 2022
29/09/2021	Italy NADEF 2021: deficit forecast improved to 9.4% of GDP, growth raised to 6.0%
01/03/2021	ISTAT reports 2020 deficit at 9.5% of GDP, debt at 155.6%
27/12/2019	Italian Parliament approves 2020 Budget Law
20/11/2019	European Commission opinion: Italy's 2020 budget at risk of non-compliance with SGP

Continued on next page

Table 1.G.5 continued

Date	Headline
15/10/2019	Italy submits 2020 Draft Budgetary Plan to European Commission
30/09/2019	Italy NADEF 2019: new government targets 2.2% deficit for 2020, VAT hike deactivated
02/07/2019	Italy avoids EDP: EU Council accepts Italy's corrective measures and fiscal commitments
05/06/2019	European Commission report says Italy's debt growth justifies opening an EDP
08/04/2019	Italy DEF 2019: growth forecast slashed to 0.2%, deficit target raised to 2.4%
01/04/2019	ISTAT reports 2018 deficit at 2.2% of GDP, within revised target
29/12/2018	Italian Parliament approves 2019 Budget Law with revised deficit target
19/12/2018	European Commission accepts Italy's revised deficit target of 2.04%, averts EDP
12/12/2018	Italy offers budget compromise: deficit reduced from 2.4% to 2.04% for 2019
19/11/2018	European Commission concludes Italy warrants Excessive Deficit Procedure
13/11/2018	Italy resubmits unchanged budget to European Commission, defying revision request
23/10/2018	European Commission unprecedented rejection of Italy's 2019 Draft Budgetary Plan
18/10/2018	European Commission sends letter expressing serious concern about Italy's budget deviation
15/10/2018	Italy submits 2019 Draft Budgetary Plan to European Commission with 2.4% deficit

Continued on next page

Table 1.G.5 continued

Date	Headline
04/10/2018	Italian government publishes detailed 2019 budget: citizens' income, Quota 100, flat tax plans
27/09/2018	Italy NADEF 2018: government announces 2.4% deficit target for 2019, triple previous commitment
07/08/2018	Italian government announces plan for flat tax, citizens' income, and pension reform in autumn budget
26/04/2018	Italy DEF 2018: outgoing government presents cautious fiscal path, deficit 1.6% of GDP
02/02/2018	ISTAT reports 2017 deficit at 2.4% of GDP
23/12/2017	Italian Parliament approves 2018 Budget Law
22/11/2017	European Commission opinion: Italy's 2018 budget at risk of non-compliance with SGP
16/10/2017	Italy submits 2018 Draft Budgetary Plan with 1.6% deficit target to Brussels
23/09/2017	Italy NADEF 2017: deficit target for 2018 set at 1.6%, revised from 1.2%
11/04/2017	Italy DEF 2017: deficit forecast 2.1% for 2017, 1.2% for 2018
03/04/2017	ISTAT confirms 2016 deficit at 2.4% of GDP, below 3% threshold
17/01/2017	EU Commission sends letter to Italy demanding €3.4bn additional budget correction
07/12/2016	Italian Parliament approves 2017 Budget Law; Renzi formally resigns after budget passage
16/11/2016	European Commission warns Italy's 2017 budget at risk of non-compliance with SGP

Continued on next page

Table 1.G.5 continued

Date	Headline
25/10/2016	European Commission requests clarification on Italy's 2017 Draft Budgetary Plan
15/10/2016	Italy submits 2017 Draft Budgetary Plan to European Commission with 2.3% deficit target
27/09/2016	Italy NADEF 2016: deficit target revised to 2.4% of GDP for 2016
17/05/2016	EU Commission grants Italy unprecedented budget flexibility for structural reforms and refugee crisis
08/04/2016	Italy DEF 2016 presented: GDP growth forecast 1.2%, deficit target 2.3% of GDP

1.G.1 Italian fiscal shock instrument

The Italian fiscal shock instrument is constructed to mirror the German instrument described in Section 1.2, with one key modification. BTP futures lack the liquidity required for reliable high-frequency identification: bid-ask spreads are wide and trading volumes are thin relative to Bund futures, so that daily price changes on BTP futures would reflect microstructure noise rather than fundamental repricing. I therefore use the daily change in the 10-year BTP benchmark yield (IT10YT=RR from LSEG Refinitiv) as the base series.

The instrument is constructed as follows. Let $\mathcal{N}^{IT}$ denote the set of days on which Italian fiscal news was released (Table 1.G.5). The raw daily Italian fiscal shock is

$$s_t^{IT} = \begin{cases} \Delta y_t^{\mathrm{BTP}} & \text{if } t \in \mathcal{N}^{IT} \\ 0 & \text{otherwise,} \end{cases} \tag{1.78}$$

where $\Delta y_t^{\mathrm{BTP}}$ is the one-day change in the 10-year BTP yield.

Using the yield directly rather than a futures price introduces a potential concern: daily yield changes contain both an expectations component and a risk-premium component. However, the same logic that motivates the use of Bund futures applies here. Over a one-day window, risk premia move slowly relative to expectations, so the dominant source of variation in $\Delta y_t^{\mathrm{BTP}}$ on announcement days is the revision in expected fiscal fundamentals rather than a shift in the term premium.

The daily series is then aggregated to weekly frequency by summing within Friday-labelled weeks, exactly as for the German instrument. The resulting weekly series is used as the external instrument in a proxy-VAR with the same specification, where the BTP yield replaces the Bund yield in the first position and the Bund yield enters in the second position. The pass-through ratio IT→DE reported in Table 1.2 is then $(\Delta y^{\mathrm{Bund}})/(\Delta y^{\mathrm{BTP}})$, the Bund response per unit of BTP response, which is directly comparable to the DE→IT ratio from the baseline German VAR.

Because the pass-through ratios are nonlinear functions of the impulse responses, I construct confidence intervals and test statistics using a moving block bootstrap. For each of the two proxy-VARs (German and Italian), I draw 10,000 bootstrap samples using the same block bootstrap procedure as in the baseline estimation (4-week blocks). Each draw yields a full set of VAR coefficients and an impact vector, from which I compute the impulse response functions and the corresponding pass-through ratio at every horizon. Draws in which the denominator of the ratio is close to zero produce explosive values; I discard draws where either ratio exceeds 50 in absolute value.

The bootstrap distributions are recentered around the point estimates to correct for finite-sample bias: for each horizon h, I shift the bootstrap distribution so that its median coincides with the point estimate. The 90% confidence intervals reported in Table 1.2 are then the 5th and 95th percentiles of the recentered distribution. To test $H_0 : \mathrm{DE} \to \mathrm{IT} = \mathrm{IT} \to \mathrm{DE}$, I compute the bootstrap distribution of the difference between the two recentered ratio distributions. The two-sided p-value is $2 \cdot \min\big(\hat{P}(d \geq 0),\ \hat{P}(d \leq 0)\big)$, where d denotes the bootstrapped difference. Since the two VARs are estimated on separate samples with independent instruments, the bootstrap draws are independent across the two systems.

Appendix to Chapter 2

2.A Robustness tests

2.A.1 Regressions based on aggregated values

Table 2.A.1. Balance sheet space and haircuts (aggregated regressions)

Notes: The dependent variable is the volume-weighted average haircut (basis points), aggregated at the dealer–client–collateral country–month level. The sample covers bilateral repos between identified dealers and non-dealer counterparties, collateralized by euro-area government bonds (sector S13) from all issuer countries and across all maturities, from January 2019 to June 2024. Columns (1)–(3) use collateral-driven repos; columns (4)–(6) use funding-driven repos. Dealers are assigned to leverage terciles within each quarter using prudential reporting data; Tercile 1 denotes the most constrained dealers. *Post* equals one after June 2021, when the leverage ratio exemption for central bank reserves expired. Standard errors clustered at the dealer level in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

	Haircut (%)					
	Collateral driven			Funding driven		
	(1)	(2)	(3)	(4)	(5)	(6)
$Tercile1 \times Post$	0.977***	1.026***	0.970***	−0.313	0.124	−0.072
	(0.176)	(0.204)	(0.191)	(0.536)	(0.440)	(0.293)
$Tercile2 \times Post$	0.582***	0.653***	0.465**	0.052	0.178	0.150
	(0.110)	(0.163)	(0.177)	(0.484)	(0.447)	(0.349)
N	75 142	75 142	75 142	49 339	49 339	49 339
R^2	0.22	0.53	0.62	0.25	0.55	0.62
FE: Coll. country	✓	✓	✓	✓	✓	✓
FE: Client		✓			✓	
FE: Dealer × Client			✓			✓
FE: Month	✓	✓	✓	✓	✓	✓
FE: LeverageTercile	✓	✓	✓	✓	✓	✓

Table 2.A.2. Convenience yields and haircuts (aggregated regressions)

Notes: The dependent variable is the volume-weighted average haircut (basis points), aggregated at the dealer–client–collateral country–month level. *ConvenienceYield* is the volume-weighted average DFR-adjusted repo rate for each collateral country on each day, then averaged to the dealer–client–country–month level. Controls include average transaction volume and volume-weighted maturity (in days). The sample and haircut cleaning follow Table 3. Columns (1)–(3) use collateral-driven repos; columns (4)–(6) use funding-driven repos. Standard errors clustered at the dealer–client level in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

	Haircut (%)					
	Collateral-driven			Funding-driven		
	(1)	(2)	(3)	(4)	(5)	(6)
$ConvenienceYield$	−0.391***	−0.418***	−0.426***	0.037	0.154	0.091
	(0.092)	(0.073)	(0.064)	(0.221)	(0.226)	(0.178)
$Volume$	−0.281	−0.180	−0.001	0.492	−0.489*	−0.539**
	(0.593)	(0.236)	(0.164)	(0.499)	(0.271)	(0.253)
$MaturityBucket$	−0.036***	0.006	0.006	0.033**	0.000	0.000
	(0.008)	(0.005)	(0.005)	(0.016)	(0.005)	(0.004)
ES	−0.172***	−0.040	−0.093***	0.192***	0.150**	0.234***
	(0.043)	(0.035)	(0.027)	(0.073)	(0.061)	(0.049)
FR	0.073***	0.041***	0.000	−0.056	−0.005	−0.012
	(0.027)	(0.015)	(0.013)	(0.046)	(0.032)	(0.030)
IT	−0.270***	−0.089***	−0.117***	0.326***	0.183***	0.200***
	(0.044)	(0.030)	(0.027)	(0.073)	(0.060)	(0.052)
N	45 378	45 378	45 378	25 803	25 803	25 803
R^2	0.05	0.45	0.61	0.03	0.41	0.65
FE: Month	✓	✓	✓	✓	✓	✓
FE: Client		✓			✓	
FE: Client x Dealer			✓			✓

Table 2.A.3. Market power and haircuts (aggregated regressions)

Notes: The dependent variable is the transaction-level haircut (basis points). $\#Dealers$ is the number of distinct dealers each client transacts with over the full sample period. Controls include transaction volume and maturity (in days). The sample and haircut cleaning follow Table 3. Columns (1)–(3) use collateral-driven repos; columns (4)–(6) use funding-driven repos. Standard errors clustered at the dealer–client level in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

	Collateral-driven			Funding-driven		
	(1)	(2)	(3)	(4)	(5)	(6)
$\#Dealers$	−0.014	−0.014	0.005	−0.198***	−0.199***	−0.209***
	(0.017)	(0.017)	(0.017)	(0.050)	(0.049)	(0.047)
$Volume$		−0.076	−0.114		2.484*	0.650
		(0.402)	(0.387)		(1.490)	(1.236)
$Maturity$		0.016	−0.002		0.024	−0.017
		(0.011)	(0.011)		(0.036)	(0.028)
N	1 803 869	1 803 869	1 803 869	1 359 609	1 359 609	1 359 609
R^2	0.15	0.15	0.21	0.36	0.36	0.41
FE: Month	✓	✓	✓	✓	✓	✓
FE: Coll. country	✓	✓	✓	✓	✓	✓
FE: Coll. sector	✓	✓	✓	✓	✓	✓
FE: Dealer			✓			✓

Table 2.A.4. Haircuts and volatility in rates and volumes (aggregated regressions)

Notes: The dependent variables are the within-group standard deviation of repo rates (columns 1–5) and of transaction volumes (columns 6–10), computed at the dealer–collateral country–month level, separately for transactions with and without a haircut. *ZeroHaircut* equals one for transactions where the haircut equals zero (collateral-driven panel) or is non-negative (funding-driven panel). Panel A uses collateral-driven repos; Panel B uses funding-driven repos. The sample and haircut cleaning follow Table 3. Standard errors clustered at the dealer–month level in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

	Panel A: Collateral-driven							
	Rate (std. dev.)				**Volume (std.dev.)**			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$Haircut \geq 0$	0.017***	0.034***	0.040***	0.038***	0.003***	0.005***	0.005***	0.004***
	(0.004)	(0.004)	(0.004)	(0.004)	(0.001)	(0.001)	(0.001)	(0.001)
$Volume$				0.001***				0.001***
				(0.000)				(0.000)
N	11 838	11 838	11 838	11 838	11 838	11 838	11 838	11 838
R^2	0.08	0.12	0.22	0.22	0.33	0.38	0.39	0.42
FE: Coll. country	✓	✓	✓	✓	✓	✓	✓	✓
FE: Dealer		✓	✓	✓		✓	✓	✓
FE: Month			✓	✓			✓	✓

	Panel B: Funding-driven							
	Rate (std. dev.)				**Volume (std.dev.)**			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$Haircut \leq 0$	0.062***	0.078***	0.079***	0.078***	0.000	0.002***	0.002***	0.002***
	(0.004)	(0.004)	(0.004)	(0.004)	(0.000)	(0.000)	(0.000)	(0.000)
$Volume$				0.001***				0.001***
				(0.000)				(0.000)
N	13 678	13 678	13 678	13 678	13 678	13 678	13 678	13 678
R^2	0.06	0.09	0.14	0.14	0.33	0.37	0.38	0.41
FE: Coll. country	✓	✓	✓	✓	✓	✓	✓	✓
FE: Dealer		✓	✓	✓		✓	✓	✓
FE: Month			✓	✓			✓	✓

2.A.2 Robustness of Subsection 2.3.2 results

Table 2.A.5. Balance sheet space and haircuts (various specifications)

Notes: Robustness of the leverage constraint regressions (Table 4, specification (1)) to two sample restrictions. Columns (1)–(2) exclude transactions with a zero haircut. Columns (3)–(4) restrict the counterparty to sector S124 (investment funds). Within each panel, odd columns report the baseline tercile specification; even columns add the quarter-end interaction (QE $= 1$ on the last two business days of each quarter). All specifications include tercile and Post main effects, with the latter absorbed by the ISIN×day fixed effects. Controls: log volume, maturity bucket dummies. Standard errors clustered at the dealer×client level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

	Collateral-driven			
	Non-zero haircuts		S124/S125 counterparties	
	(1)	(2)	(3)	(4)
Tercile 1 × Post	0.623***	0.617***	0.462*	0.462*
	(0.215)	(0.217)	(0.252)	(0.252)
Tercile 2 × Post	0.878***	0.870***	0.445*	0.445*
	(0.222)	(0.224)	(0.252)	(0.252)
T1 × Post × QE		0.204		−0.006
		(0.234)		(0.062)
T2 × Post × QE		0.283		−0.004
		(0.225)		(0.059)
ISIN × day FE	✓	✓	✓	✓
Client FE	✓	✓	✓	✓
N	387,392	387,392	787,353	787,353
R^2	0.93	0.93	0.87	0.87

Table 2.A.6. Convenience yields and haircuts (various specifications)

Notes: Robustness of the convenience yield regressions (Table 5) to two sample restrictions. Columns (1)–(4) exclude transactions with a zero haircut. Columns (5)–(8) restrict the counterparty to sectors S124 (investment funds) and S125 (other financial intermediaries). Within each panel, odd columns use dealer×client + month + country FE; even columns use dealer×client + country×month FE. Controls: log volume, maturity bucket dummies. Standard errors clustered at the dealer×client level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

	Non-zero haircuts				S124/S125 counterparties			
	Collateral-driven		Funding-driven		Collateral-driven		Funding-driven	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Convenience yield	-0.291^{***}	-0.314^{***}	0.690^{***}	0.712^{***}	-0.156^{**}	-0.164^{**}	0.396^{***}	0.420^{***}
	(0.102)	(0.108)	(0.121)	(0.122)	(0.073)	(0.078)	(0.125)	(0.135)
Month FE	✓		✓		✓		✓	
Country FE	✓		✓		✓		✓	
Country × month FE		✓		✓		✓		✓
Dealer × client FE	✓	✓	✓	✓	✓	✓	✓	✓
N	407,128	407,121	375,178	375,177	959,363	959,363	400,648	400,640
R^2	0.87	0.87	0.74	0.77	0.69	0.70	0.64	0.65

Table 2.A.7. Dealer market power and haircuts (various specifications)

Notes: Robustness of the market power regressions (Table 6) to two sample restrictions. Columns (1)–(4) exclude transactions with a zero haircut. Columns (5)–(8) restrict the counterparty to sectors S124 (investment funds) and S125 (other financial intermediaries). Within each panel, odd columns report the baseline #Dealers specification; even columns add the #Dealers×Post interaction. "Post" denotes observations after June 28, 2021. Controls: log volume, maturity bucket dummies. Standard errors clustered at the dealer level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

	Non-zero haircuts				S124/S125 counterparties			
	Collateral-driven		Funding-driven		Collateral-driven		Funding-driven	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
#Dealers	0.014	0.005	−0.155	-0.154^{*}	0.052^{***}	0.144^{***}	-0.080^{**}	-0.131^{**}
	(0.010)	(0.008)	(0.094)	(0.073)	(0.015)	(0.029)	(0.032)	(0.046)
#D × Post		0.018		−0.004		-0.113^{***}		0.113^{**}
		(0.015)		(0.104)		(0.026)		(0.044)
ISIN × day FE	✓	✓	✓	✓	✓	✓	✓	✓
Client FE	✓	✓	✓	✓	✓	✓	✓	✓
Leverage tercile FE	✓	✓	✓	✓	✓	✓	✓	✓
Controls	✓	✓	✓	✓	✓	✓	✓	✓
N	387,392	387,392	413,879	413,879	1,002,234	1,002,234	477,994	477,994
R^2	0.93	0.93	0.93	0.93	0.85	0.85	0.92	0.92

2.A.3 Capital ratio and risk-weighted-assets (RWA)

Table 2.A.8. Capital ratio and haircuts (various specifications)

Notes: Robustness of Table 4 using total capital ratios instead of leverage ratios. Transaction-level regressions of repo haircuts (%) on dealer capital constraints. The sample covers bilateral dealer-to-client repos collateralised by euro area government bonds, January 2019 to June 2024. Dealers are sorted into total capital ratio terciles each quarter; Tercile 1 denotes the most constrained dealers and Tercile 3 (omitted) the least. *Post* = 1 after June 28, 2021. All specifications include tercile and Post main effects; the latter is absorbed by the ISIN×day fixed effects. (1) Baseline with client FE. (2) Dealer×client FE absorbs time-invariant relationship characteristics. (3) Continuous: CR headroom ≡ capital ratio − 8% (pp); dealer FE absorbs the level so identification is from within-dealer quarterly variation. (4) Restricted to dealer–client pairs active both pre- and post-treatment. (5) QE = 1 on the last two business days of each quarter. (6) Funding-driven placebo. Controls: log volume, maturity bucket dummies. Standard errors clustered at the dealer level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

	Collateral-driven				
	(1)	(2)	(3)	(4)	(5)
Tercile 1 × Post	−0.026	−0.015		−0.002	−0.027
	(0.023)	(0.057)		(0.059)	(0.023)
Tercile 2 × Post	−0.140**	−0.081*		−0.058	−0.140**
	(0.063)	(0.039)		(0.035)	(0.063)
CR headroom × Post			0.007		
			(0.012)		
T1 × Post × QE					0.052**
					(0.022)
T2 × Post × QE					0.016
					(0.027)
ISIN × day FE	✓	✓	✓	✓	✓
Client FE	✓		✓		✓
Dealer × client FE		✓		✓	
Dealer FE			✓		
Balanced relationships				✓	
N	1,556,748	1,556,533	1,556,748	1,141,552	1,556,748
R^2	0.83	0.88	0.83	0.88	0.83

2.B Proofs

2.B.1 Proof of Proposition 2

Proof. (i): In Stage 2, the dealer and client jointly choose the effective cost ρ_j and the haircut h to maximise the generalised Nash product

$$\max_{\rho_j, h} N = (\pi)^{\theta_j} (U)^{1-\theta_j}, \tag{2.79}$$

where the dealer's profit (from equation (2.5) with $r = -\rho_j - \lambda h$ and $V = (1-h)C$) is

$$\pi = (1-p)(\rho_j + \lambda h)(1-h)\,C - (y_j + pL)\,C - \tfrac{\gamma_d \mu_j}{2}(1-h)^2 C^2, \tag{2.80}$$

and the client's surplus is

$$U = (v - \rho_j)\,C, \tag{2.81}$$

with v the client's valuation of the collateral. Equivalently to maximising (2.9), we can also maximise $\ln N = \theta_j \ln \pi + (1-\theta_j)\ln U$. The first-order condition with respect to h is

$$\frac{\partial \ln N}{\partial h} = \frac{\theta_j}{\pi}\frac{\partial \pi}{\partial h} + \frac{1-\theta_j}{U}\frac{\partial U}{\partial h} = 0. \tag{2.82}$$

The client's surplus $U = (v - \rho_j)C$ depends on the effective cost ρ_j but not directly on the haircut h (at fixed ρ_j, the client is indifferent over how the dealer decomposes ρ_j into a rate and a haircut). Therefore $\partial U/\partial h = 0$, and (2.82) reduces to

$$\frac{\theta_j}{\pi}\frac{\partial \pi}{\partial h} = 0 \qquad \Longrightarrow \qquad \frac{\partial \pi}{\partial h} = 0. \tag{2.83}$$

The jointly efficient haircut unconditionally maximises the dealer's profit for any given ρ_j. The bargaining share θ_j drops out: it governs the split of the surplus (through ρ_j) but not the efficient contract design. From (2.80), expand the cash-leg revenue $(1-p)(\rho_j+\lambda h)(1-h)C$:

$$(\rho_j+\lambda h)(1-h)=\rho_j-\rho_j h+\lambda h-\lambda h^2. \tag{2.84}$$

The collateral-side cost simplifies as before: $\frac{y_j+pL}{1-h}\times(1-h)C=(y_j+pL)C$, which is constant in h. The full profit is

$$\pi=(1-p)\big[\rho_j-\rho_j h+\lambda h-\lambda h^2\big]C-(y_j+pL)C-\tfrac{\gamma_d\mu_j}{2}(1-h)^2C^2. \tag{2.85}$$

From (2.84):

$$\frac{\partial}{\partial h}\big[(1-p)(\rho_j-\rho_j h+\lambda h-\lambda h^2)C\big]=(1-p)(-\rho_j+\lambda-2\lambda h)\,C. \tag{2.86}$$

The collateral-side cost $(y_j+pL)C$ is constant in h, hence

$$\frac{\partial}{\partial h}\big[-(y_j+pL)C\big]=0. \tag{2.87}$$

Apply the chain rule to the balance-sheet penalty $-\frac{\gamma_d\mu_j}{2}(1-h)^2C^2$:

$$\frac{\partial}{\partial h}\big[-\tfrac{\gamma_d\mu_j}{2}(1-h)^2C^2\big]=\gamma_d\mu_j(1-h)\,C^2. \tag{2.88}$$

Setting $\partial\pi/\partial h=(2.86)+(2.87)+(2.88)=0$:

$$(1-p)(-\rho_j+\lambda-2\lambda h)\,C+\gamma_d\mu_j(1-h)\,C^2=0. \tag{2.89}$$

Divide by $C > 0$ and expand $\gamma_d \mu_j C\,(1-h) = \gamma_d \mu_j C - \gamma_d \mu_j C\,h$. Collecting constant terms and terms in h:

$$\big[(1-p)\lambda + \gamma_d \mu_j C - (1-p)\rho_j\big] - \big[2(1-p)\lambda + \gamma_d \mu_j C\big]\,h = 0. \tag{2.90}$$

Solve for h:

$$h^*(C) = \frac{(1-p)\lambda + \gamma_d \mu_j C - (1-p)\rho_j}{2(1-p)\lambda + \gamma_d \mu_j C}. \tag{2.91}$$

The second-order condition $\partial^2 \pi / \partial h^2 = -2(1-p)\lambda\,C - \gamma_d \mu_j\,C^2 < 0$ for $C > 0$, confirms that (2.91) is a maximum.

(ii): By Step 2, the Stage-2 haircut satisfies $\partial \pi / \partial h = 0$ unconditionally. The envelope theorem therefore applies without qualification: when differentiating the dealer's profit with respect to C, the indirect effect through $h^*(C)$ vanishes. However, in the Stage-1 Cournot game, the dealer internalises that expanding volume depresses the market-clearing effective cost through $\partial \rho_j / \partial C_{dj} = -\bar{v}_j / M_j < 0$. The total derivative of the dealer's profit with respect to own volume is therefore

$$\frac{d\pi}{dC} = \frac{\partial \pi}{\partial C} + \underbrace{\frac{\partial \pi}{\partial h}}_{=0} \frac{\partial h^*}{\partial C} + \frac{\partial \pi}{\partial \rho_j} \frac{\partial \rho_j}{\partial C_{dj}} = \frac{\partial \pi}{\partial C} + \frac{\partial \pi}{\partial \rho_j} \frac{\partial \rho_j}{\partial C_{dj}}. \tag{2.92}$$

We need the direct partial derivatives of (2.85) with respect to C and ρ_j. The cash-leg revenue $(1-p)(\rho_j + \lambda h)(1-h) \times C$ is linear in C

$$\frac{\partial\,(\mathrm{I})}{\partial C} = (1-p)(\rho_j + \lambda h)(1-h). \tag{2.93}$$

The collateral-side cost $(y_j + pL)C$ is linear in C

$$\frac{\partial\,(\mathrm{II})}{\partial C} = (y_j + pL). \tag{2.94}$$

The balance-sheet penalty $(\gamma_d \mu_j/2)(1-h)^2 C^2$ is quadratic in C

$$\frac{\partial\,\text{(III)}}{\partial C} = \gamma_d \mu_j (1-h)^2 C. \tag{2.95}$$

For the strategic term, differentiating (2.85) with respect to ρ_j at fixed h and C yields

$$\frac{\partial \pi}{\partial \rho_j} = (1-p)(1-h^*)\, C. \tag{2.96}$$

Collecting terms and setting $d\pi/dC = 0$:

$$(1-p)(\rho_j + \lambda h^*)(1-h^*) - (y_j + pL) - \gamma_d \mu_j (1-h^*)^2\, C + (1-p)(1-h^*)\,\frac{\partial \rho_j}{\partial C_{dj}}\, C = 0. \tag{2.97}$$

□

2.B.2 Proof of existence

Proposition 4 (Existence of Equilibrium)**.** *There exists a pure-strategy subgame perfect Nash equilibrium in the two-stage collateral-driven repo market game.*

Proof. We prove the existence of a pure-strategy Nash equilibrium in the Stage-1 volume game using the Debreu-Glicksberg-Fan Theorem. Let the aggregate collateral volume be $Q_j = C_{dj} + C_{-dj}$. The inverse demand function $\rho_j(Q_j)$ is linear and strictly downward sloping, meaning $\frac{\partial \rho_j}{\partial C_{dj}} < 0$ and $\frac{\partial^2 \rho_j}{\partial C_{dj}^2} = 0$.

To apply the theorem, the Stage-1 game must satisfy three conditions:

(i) Compact and Convex Strategy Space: Each dealer d chooses a volume subject to a collateral capacity constraint $C_{dj} \in [0, \overline{C}_{dj}]$. This interval is a non-empty, compact, and convex subset of $\mathbb{R}$.

(ii) Continuous Payoffs: The dealer's expected profit function $\Pi_d(C_{dj}, C_{-dj})$ is composed of continuous algebraic operations on the linear inverse demand and the quadratic balance-sheet penalty. Therefore, Π_d is continuous in the strategy profile of all players.

(iii) Strictly Concave Payoffs: We must show that a dealer's profit is strictly concave in their own chosen volume ($\frac{\partial^2 \Pi_d}{\partial C_{dj}^2} < 0$). By the Envelope Theorem, the indirect effect of C_{dj} on the optimal Stage-2 haircut h^*_{dj} vanishes because the haircut maximizes joint surplus ($\frac{\partial \Pi_d}{\partial h^*_{dj}} = 0$). Differentiating the profit function with respect to C_{dj} twice yields:

$$\frac{\partial^2 \Pi_d}{\partial C_{dj}^2} = \underbrace{2(1-p)(1-h^*_{dj})\frac{\partial \rho_j}{\partial C_{dj}}}_{\text{Negative (Downward demand)}} - \underbrace{\gamma_d \mu_j (1-h^*_{dj})^2}_{\text{Positive (Convex cost)}}$$

Because the inverse demand is strictly downward sloping and the balance-sheet cost parameter ($\gamma_d \mu_j$) is strictly positive, the subtraction of the convex cost from the negative demand effect ensures the entire second derivative is strictly negative. Thus, Π_d is strictly concave in C_{dj}.

Because the strategy space is compact and convex, and the payoff function is continuous and strictly concave in the dealer's own strategy, each dealer's best-response correspondence is continuous and single-valued. By the Debreu-Glicksberg-Fan Theorem, at least one pure-strategy fixed point (Nash equilibrium) exists. $\square$

2.B.3 Proof of Corollary 1

Proof. From Proposition 2, the optimal bargained haircut that maximizes joint surplus is given by

$$h^* = \frac{(1-p)(\lambda - \rho_j) + \gamma_d \mu_j C}{2(1-p)\lambda + \gamma_d \mu_j C} \tag{2.98}$$

We want to find the partial derivative of the equilibrium haircut with respect to the dealer's balance-sheet cost, $\frac{\partial h^*}{\partial \gamma_d}$, and show that it is strictly positive. Let the numerator be N and the denominator be D

$$N = (1-p)(\lambda - \rho_j) + \gamma_d \mu_j C$$
$$D = 2(1-p)\lambda + \gamma_d \mu_j C$$

Taking the partial derivative of N and D with respect to γ_d gives

$$\frac{\partial N}{\partial \gamma_d} = \mu_j C$$
$$\frac{\partial D}{\partial \gamma_d} = \mu_j C$$

Using the quotient rule $\frac{\partial h^*}{\partial \gamma_d} = \frac{\frac{\partial N}{\partial \gamma_d} D - N \frac{\partial D}{\partial \gamma_d}}{D^2}$, we get:

$$\frac{\partial h^*}{\partial \gamma_d} = \frac{(\mu_j C)\left[2(1-p)\lambda + \gamma_d \mu_j C\right] - \left[(1-p)(\lambda - \rho_j) + \gamma_d \mu_j C\right](\mu_j C)}{\left[2(1-p)\lambda + \gamma_d \mu_j C\right]^2} \tag{2.99}$$

Factoring out the common term $(\mu_j C)$ in the numerator yields

$$\frac{\partial h^*}{\partial \gamma_d} = \frac{\mu_j C\Big(\left[2(1-p)\lambda + \gamma_d \mu_j C\right] - \left[(1-p)(\lambda - \rho_j) + \gamma_d \mu_j C\right]\Big)}{\left[2(1-p)\lambda + \gamma_d \mu_j C\right]^2} \tag{2.100}$$

Notice that the $\gamma_d \mu_j C$ terms cancel out

$$\frac{\partial h^*}{\partial \gamma_d} = \frac{\mu_j C\Big(2(1-p)\lambda - (1-p)(\lambda - \rho_j)\Big)}{\big[2(1-p)\lambda + \gamma_d \mu_j C\big]^2} \tag{2.101}$$

Factor out $(1-p)$ from the remaining terms inside the bracket

$$\frac{\partial h^*}{\partial \gamma_d} = \frac{\mu_j C(1-p)\big[2\lambda - (\lambda - \rho_j)\big]}{\big[2(1-p)\lambda + \gamma_d \mu_j C\big]^2}, \tag{2.102}$$

which simplifies to the final derivative

$$\frac{\partial h^*}{\partial \gamma_d} = \frac{\mu_j C(1-p)(\lambda + \rho_j)}{\big[2(1-p)\lambda + \gamma_d \mu_j C\big]^2}. \tag{2.103}$$

To prove that $\frac{\partial h^*}{\partial \gamma_d} > 0$, we must verify that every component of this expression is strictly positive. The denominator $\big[2(1-p)\lambda + \gamma_d \mu_j C\big]^2 > 0$ since the square of any non-zero real number is positive. The parameters $\mu_j C > 0$ (volume and the country multiplier are strictly positive) and $(1-p) > 0$ (survival probability $p \in (0,1)$). The economic restriction requires that the equilibrium haircut cannot exceed 100% ($h^* < 1$). We can express $1 - h^*$ as:

$$1 - h^* = 1 - \frac{(1-p)(\lambda - \rho_j) + \gamma_d \mu_j C}{2(1-p)\lambda + \gamma_d \mu_j C} = \frac{(1-p)(\lambda + \rho_j)}{2(1-p)\lambda + \gamma_d \mu_j C} \tag{2.104}$$

For the haircut to be less than 1, the expression $1 - h^*$ must be strictly positive. Since the denominator is positive and $(1-p)$ is positive, the term $(\lambda + \rho_j)$ must be strictly positive for the model to yield a valid economic contract. □

2.B.4 Proof of Corollary 2

Proof. From Proposition 2, the optimal bargained haircut is given by

$$h^*(C, \rho_j) = \frac{(1-p)(\lambda - \rho_j) + \gamma_d \mu_j C}{2(1-p)\lambda + \gamma_d \mu_j C}. \tag{2.105}$$

The convenience yield y_j does not enter the h^* equation directly, as it represents a constant marginal opportunity cost that drops out of the Nash Bargaining first-order condition. Therefore, its effect on the haircut operates entirely through the equilibrium volume C and the market-clearing effective cost ρ_j.

We evaluate the total derivative of h^* with respect to y_j using the chain rule:

$$\frac{dh^*}{dy_j} = \frac{\partial h^*}{\partial \rho_j}\frac{d\rho_j}{dy_j} + \frac{\partial h^*}{\partial C}\frac{dC}{dy_j} \tag{2.106}$$

First, we differentiate h^* with respect to the effective cost ρ_j

$$\frac{\partial h^*}{\partial \rho_j} = \frac{-(1-p)}{2(1-p)\lambda + \gamma_d \mu_j C} < 0 \tag{2.107}$$

Since $(1-p) > 0$ and the denominator is positive, this derivative is strictly negative.

Second, we differentiate h^* with respect to volume C. Applying the quotient rule and simplifying the algebraic terms yields

$$\frac{\partial h^*}{\partial C} = \frac{\gamma_d \mu_j (1-p)(\lambda + \rho_j)}{\left[2(1-p)\lambda + \gamma_d \mu_j C\right]^2}. \tag{2.108}$$

Because the economic equilibrium restriction dictates that the haircut cannot exceed 100% ($h^* < 1$), the term $(\lambda + \rho_j)$ must be positive. Since all other parameters are strictly positive, this partial derivative is strictly positive ($\frac{\partial h^*}{\partial C} > 0$).

In the Stage-1 volume choice, an increase in the convenience yield y_j acts as a direct increase in the marginal cost of supplying collateral-driven repo. By the standard properties of Cournot competition with downward-sloping demand, an increase in marginal cost strictly decreases the equilibrium quantity supplied by the dealer

$$\frac{dC}{dy_j} < 0. \tag{2.109}$$

Because aggregate supply decreases, the market-clearing effective cost (price) ρ_j must strictly increase along the inverse demand curve

$$\frac{d\rho_j}{dy_j} > 0. \tag{2.110}$$

Substituting the signs of these components back into the total derivative equation:

$$\frac{dh^*}{dy_j} = \underbrace{\left(\frac{\partial h^*}{\partial \rho_j}\right)}_{(-)} \underbrace{\left(\frac{d\rho_j}{dy_j}\right)}_{(+)} + \underbrace{\left(\frac{\partial h^*}{\partial C}\right)}_{(+)} \underbrace{\left(\frac{dC}{dy_j}\right)}_{(-)} \tag{2.111}$$

The first term represents the price effect: a higher convenience yield drives up the effective cost ρ_j, which allows the dealer to concede a lower haircut. The second term represents the volume effect: the higher cost shrinks the dealer's supplied volume, relaxing their balance-sheet constraint, which also puts downward pressure on the haircut. □

2.B.5 Proof of Corollary 3

Proof. The equilibrium haircut h^* maximizes joint surplus and is given by:

$$h^*(C, \rho_j) = \frac{(1-p)(\lambda - \rho_j) + \gamma_d \mu_j C}{2(1-p)\lambda + \gamma_d \mu_j C} \tag{2.112}$$

Because the bargaining parameter θ_j does not appear directly in the joint-efficiency condition, a change in dealer market power affects the haircut entirely through the negotiated effective cost ρ_j and the equilibrium volume C. Applying the chain rule yields the total derivative:

$$\frac{dh^*}{d\theta_j} = \frac{\partial h^*}{\partial \rho_j}\frac{d\rho_j}{d\theta_j} + \frac{\partial h^*}{\partial C}\frac{dC}{d\theta_j} \tag{2.113}$$

Differentiating h^* with respect to the effective cost yields $\frac{\partial h^*}{\partial \rho_j} = \frac{-(1-p)}{2(1-p)\lambda + \gamma_d \mu_j C} < 0$. By the mechanics of Generalized Nash Bargaining, an increase in dealer power θ_j allows the dealer to extract a larger share of the surplus, securing a strictly higher effective cost from the client ($\frac{d\rho_j}{d\theta_j} > 0$). Therefore, the first term of the total derivative is strictly negative:

$$\underbrace{\left(\frac{\partial h^*}{\partial \rho_j}\right)}_{(-)}\underbrace{\left(\frac{d\rho_j}{d\theta_j}\right)}_{(+)} < 0 \tag{2.114}$$

Economically, as the dealer secures a higher profit margin (ρ_j), the trade becomes more lucrative per unit of cash exchanged. To maximize total returns on this high margin, the dealer optimally expands the cash leg of the trade by pushing the haircut lower, accepting a slightly higher balance-sheet penalty in exchange for greater cash-leg revenue.

Differentiating h^* with respect to volume C via the quotient rule and simplifying yields

$$\frac{\partial h^*}{\partial C} = \frac{\gamma_d \mu_j (1-p)(\lambda + \rho_j)}{\left[2(1-p)\lambda + \gamma_d \mu_j C\right]^2}. \tag{2.115}$$

Because the economic equilibrium restriction dictates that the haircut cannot exceed 100% ($h^* < 1$), the term $(\lambda + \rho_j)$ must be strictly positive, making this partial derivative strictly positive ($\frac{\partial h^*}{\partial C} > 0$).

Simultaneously, higher dealer market power θ_j increases expected profit margins, inducing the dealer to expand the total volume of collateral supplied in the first stage ($\frac{dC}{d\theta_j} > 0$). Therefore, the second term of the total derivative is strictly positive:

$$\underbrace{\left(\frac{\partial h^*}{\partial C}\right)}_{(+)}\underbrace{\left(\frac{dC}{d\theta_j}\right)}_{(+)} > 0 \tag{2.116}$$

Economically, the expanded volume tightly binds the dealer's quadratic balance-sheet constraint, putting upward pressure on the haircut. □

2.C Funding-driven market model

In the funding-driven segment, the roles are reversed: the dealer lends cash $V_{dj} = (1 - h_{dj})C_{dj}$ to a client who posts country-j collateral C_{dj} as security. The dealer forgoes the deposit facility rate on the cash lent but earns the repo rate at maturity and receives the convenience yield on the collateral held during the term of the trade.

Dealer profit. Dealer d's expected profit from a funding-driven trade in country j is

$$\pi_{dj}^{FD} = \left[(1-p)\, r_{dj} - p\, L + \frac{y_j}{1-h_{dj}}\right] V_{dj}, \tag{2.117}$$

where r_{dj} is the repo rate received at maturity conditional on the counterparty surviving with probability $1-p$, y_j is the convenience yield earned on the collateral received, $p\,L$ is the expected loss given default, and $V_{dj} = (1 - h_{dj})\, C_{dj}$ is the cash lent. Crucially, the balance-sheet penalty is normalised to zero ($\Phi \equiv 0$) in the funding-driven segment: a reverse repo from the dealer's perspective acts as an asset swap—cash is exchanged for collateral—and does not generate the same leverage-ratio exposure as the cash absorption in collateral-driven trades.

Client surplus. In the funding-driven segment, the client borrows cash and posts collateral. A higher repo rate and a higher haircut both increase the client's cost. The effective borrowing cost is

$$\rho_{dj}^{FD} = r_{dj} + \lambda\, h_{dj}, \tag{2.118}$$

where $\lambda > 0$ is the client's sensitivity to the haircut margin (the shadow cost of collateral encumbrance). The client's surplus from trading with dealer d at effective cost ρ_{dj} is

$$U_{dj} = (v - \rho_{dj})\, V_{dj}, \tag{2.119}$$

where $v > 0$ is the client's reservation value for cash (i.e. the maximum effective cost at which the trade remains individually rational).

Bargaining. As in the collateral-driven segment, each dealer–client pair bargains bilaterally under Nash-in-Nash. The dealer and client jointly choose the effective cost ρ_{dj} and the haircut h_{dj} to maximise the generalised Nash product

$$\max_{\rho_{dj},\, h_{dj}} N = \left(\pi_{dj}^{FD}\right)^{\theta_j} (U_{dj})^{1-\theta_j}\,. \tag{2.120}$$

Demand. On the client side, aggregate demand for cash in country-j collateral is $D_j(\rho) = M_j\,(\bar{v}_j - \rho)/\bar{v}_j$, with the same linear structure as in the collateral-driven segment.

Equilibrium. Dealers compete in a two-stage game: in Stage 1, each dealer d commits a cash capacity $V_{dj} \leq \bar{V}_{dj}$; in Stage 2, each dealer–client pair bargains over (r_{dj}, h_{dj}) by solving the Nash product (2.120), and the client allocates demand to the dealer offering the lowest effective cost ρ_{dj}.

Proposition 5 (Funding-driven equilibrium)**.** *The equilibrium of the funding-driven repo bargaining game is characterised as follows.*

(i) *The dealer's profit in* (2.117) *is* convex *in h_{dj}, so the equilibrium haircut is a corner solution. Comparing profits at $h = 0$ and $h = \bar{h}$, the equilibrium haircut is*

$$h_{FD}^{*} = \begin{cases} \bar{h} & \text{if } y_j > (1-p)\,\lambda\,(1-\bar{h}), \\ 0 & \text{otherwise.} \end{cases} \tag{2.121}$$

(ii) *The dealer's profit in* (2.117) *is* linear *in V_{dj}, so the equilibrium volume is also a corner solution: each dealer deploys full capacity, $V_{dj}^{*} = \bar{V}_{dj}$, provided per-unit profit is non-negative. Dealer* *competition then bids down the repo rate until per-unit profit is zero, pinning down the equilibrium repo rate r^*_{dj} via*

$$(1-p)\, r^*_{dj} - p\, L + \frac{y_j}{1 - h^*_{FD}} = 0. \tag{2.122}$$

Proof. Part (i). Substituting $r_{dj} = \rho_{dj} - \lambda h_{dj}$ from (2.118) into (2.117), the dealer's profit at fixed ρ_{dj} and V_{dj} is

$$\pi^{FD}_{dj}(h) = \left[-(1-p)\lambda\, h + \frac{y_j}{1-h} + \text{const} \right] V_{dj}.$$

The second derivative with respect to h is $2y_j/(1-h)^3 > 0$, so π^{FD}_{dj} is strictly convex in h on $[0, \bar{h}]$. The unique interior critical point is therefore a minimum, and the maximum is attained at a corner. Evaluating $\pi^{FD}_{dj}(\bar{h}) - \pi^{FD}_{dj}(0) = \bar{h}\, V_{dj} \left[y_j/(1-\bar{h}) - (1-p)\lambda \right]$, which is positive if and only if $y_j > (1-p)\lambda(1-\bar{h})$, yielding (2.121).

Part (ii). Since (2.117) is linear in V_{dj}, the sign of per-unit profit—the bracket in (2.117) evaluated at h^*_{FD}—determines the optimal volume: if positive, the dealer sets $V^*_{dj} = \bar{V}_{dj}$; if negative, the dealer exits. In the Stage-1 capacity game, dealers compete by deploying full capacity until the marginal dealer earns zero per-unit profit, which gives equation (2.122) as the condition determining r^*_{dj}. □

The funding-driven equilibrium haircut (2.121) depends only on the collateral's convenience yield y_j, the default probability p, the client's haircut sensitivity λ, and the feasibility bound $\bar{h}$. It does *not* depend on the dealer's balance-sheet cost γ_d, the effective cost ρ_j, or the bargaining share θ_j. This stands in sharp contrast to the collateral-driven haircut (equation (2.11) in the main text), where both γ_d and ρ_j enter directly. The economic intuition is that, absent balance-sheet costs, the haircut in funding-driven repos serves purely as a risk-management tool: the dealer receives the convenience yield y_j on collateral held and therefore prefers maximum over-collateralisation when the convenience yield is sufficiently large relative to the client's collateral encumbrance cost.

We now derive the funding-driven analogues of Corollaries 1–3 from the main text.

Corollary 4. *The equilibrium haircut in the funding-driven segment does not depend on the dealer's balance-sheet cost*

$$\frac{\partial h^*_{FD}}{\partial \gamma_d} = 0.$$

Proof. Obvious. □

This is the model's strict null prediction, corresponding to the empirically insignificant coefficients in columns (4)–(6) of Table 2.4 and the funding-driven placebo tests throughout Section 5. It confirms that the haircut distortions documented in the collateral-driven segment are driven by the regulatory asymmetry in balance-sheet treatment, not by some unobserved common factor affecting both segments.

Corollary 5. *The equilibrium haircut in the funding-driven segment is non-decreasing in the collateral's convenience yield*

$$\frac{\partial h^*_{FD}}{\partial y_j} \geq 0,$$

with a discrete upward jump at the threshold $y_j = (1-p)\,\lambda\,(1-\bar{h})$.

Proof. From (2.121), h^*_{FD} switches from 0 to $\bar{h}$ as y_j crosses $(1-p)\lambda(1-\bar{h})$ from below, and is constant on either side of the threshold. Hence the comparative static is everywhere non-negative. □

The dealer receives the convenience yield y_j on collateral held throughout the term of the trade (see (2.117)). A higher convenience yield increases the marginal value of receiving more collateral per unit of cash lent, i.e. a higher haircut, so the dealer pushes toward maximum over-collateralisation once the convenience yield exceeds the client's encumbrance cost $(1-p)\lambda(1-\bar{h})$.

Corollary 6. *An increase in dealer market power has no effect on the equilibrium haircut in the funding-driven segment:*

$$\frac{dh^*_{FD}}{d\theta_j} = 0.$$

Market power affects only the surplus split through the negotiated effective cost ρ_j*, not the efficient contract design.*

Proof. Obvious. □

Corollary 6 delivers a sharp model prediction: market power should not affect FD haircuts. Empirically, Table 3 shows that clients with more dealer relationships do negotiate lower haircuts in the funding-driven segment ($\hat{\beta} < 0$). This suggests that the strict invariance result relies on the maintained assumption that $\Phi \equiv 0$ in FD repos. In practice, funding-driven repos may carry non-zero (though smaller) balance-sheet costs through counterparty credit risk add-ons and imperfect netting, which would reintroduce a channel through which bargaining power could influence haircuts. The model's null prediction should therefore be interpreted as a limiting case that captures the *direction* of the regulatory asymmetry: balance-sheet frictions are first-order in collateral-driven repos and second-order in funding-driven repos.

2.D Additional figures and tables

2.D.1 Additional figures

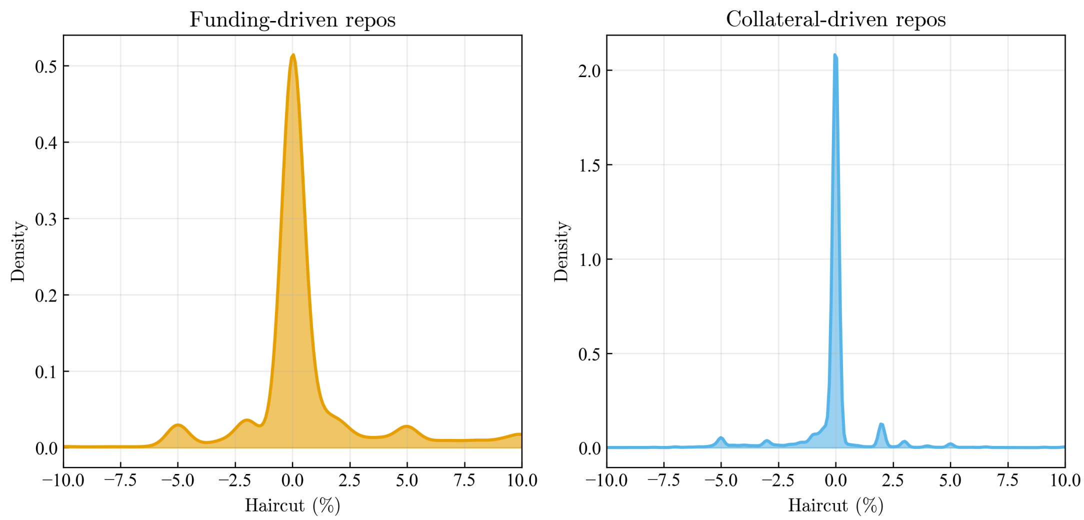

Figure 2.D.1. What is the distribution of haircuts?

Notes: This histogram displays the estimated distribution of haircuts in funding-driven and collateral-driven repo transactions between dealer banks and non-dealer clients from January 2019 to June 2024 on a *dealer-client-day* level. The estimation used a Gaussian kernel.

2.D.2 Additional tables

Table 2.D.9. Calibration

Parameter	Description	Value	Rationale
Panel A: Balance-sheet cost			
γ_0	Base balance-sheet cost (unconstrained)	0.0005	Slack dealers bear essentially no penalty
γ_1	Softplus slope (heatmap / counterfactuals)	0.00415	Matches HC distortion
η	Softplus steepness (heatmap / counterfactuals)	80	Smooth gradient / sharp kink at $\ell_{\min}$
$\ell_{\min}$	Leverage-ratio kink	3.00%	Basel-III leverage-ratio floor
Panel B: Client and risk parameters			
λ	Client haircut sensitivity	50 bp / pp	1pp more haircut ⇒ 50bp rate concession
p	Counterparty default probability	10 bp	∼0.1% annualised, IG dealer range
L	Loss given default	20%	Standard secured-exposure LGD
Panel C: Demand			
$\bar{\rho}_j$	Choke repo rate (demand-curve intercept)	200 bp	Max spread to DFR before demand goes to 0
M_j	Market size per country	0.365 bn EUR	Average daily GC volume per country
μ_{DE}	Country collateral multiplier — DE	1.10	Core-country balance-sheet loading
μ_{FR}	Country collateral multiplier — FR	1.04	Core-country balance-sheet loading
μ_{ES}	Country collateral multiplier — ES	0.96	Periphery baseline
μ_{IT}	Country collateral multiplier — IT	0.90	Periphery baseline
Panel D: Nash-in-Nash bargaining			
κ	Decay of bargaining power with competition	0.25	Intermediate market power
n_j	Active dealers per country	9	Matches MMSR data
θ_j	Dealer bargaining share (derived)	0.333	$= 1/(1+\kappa(n_j-1))$
Panel E: Collateral convenience yields			

Table 2.D.9. Calibration (continued)

Parameter	Description	Value	Rationale
y_{DE}	Convenience yield — Germany	50 bp	Core: scarcity / safety premium
y_{FR}	Convenience yield — France	40 bp	Core: scarcity / safety premium
y_{ES}	Convenience yield — Spain	25 bp	Periphery baseline
y_{IT}	Convenience yield — Italy	15 bp	Periphery baseline
Panel F: Policy environment			
r_{DFR}	ECB deposit facility rate (reference)	0 bp	Rates modelled as spreads to DFR

Notes: Calibrated parameters of the structural model (Section 1.5). Rates and yields are expressed as spreads to the ECB Deposit Facility Rate. Volumes are in EUR bn. Country-level parameters $\bar{\rho}_j$ and M_j take a common baseline across DE, FR, IT, and ES; country collateral multipliers μ_j and convenience yields y_j vary. The Nash bargaining share $\theta_j = 1/(1 + \kappa(n_j - 1))$ is derived from κ and the number of active dealers per country n_j.

Appendix to Chapter 3

3.A Tables and figures

Figure 3.A.1. Climate risk channels

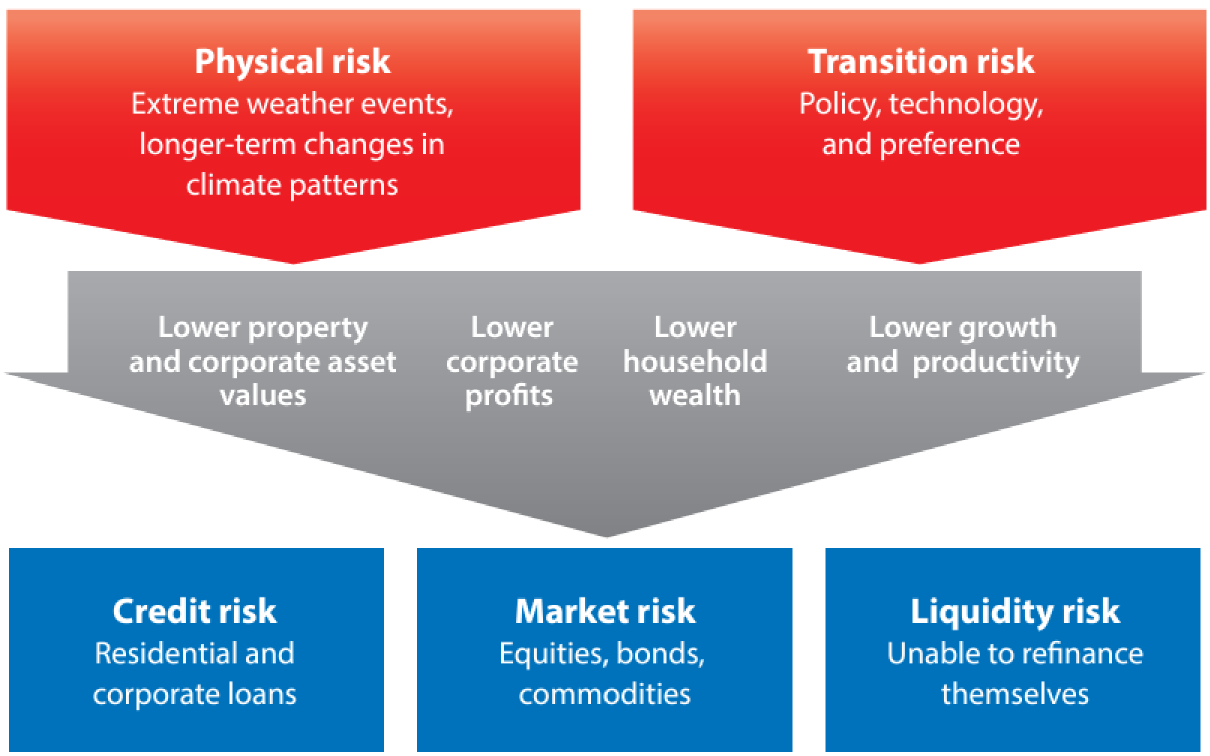

Notes: Source: Acharya et al. (2023).

Table 3.A.1. Transition risk premia dependent on collateral sector.

Notes: The dependent variable is the loan rate minus the ECB's deposit facility rate expressed in percentage points. Individual regressions are based on the sub-samples of transactions collateralised by bonds issued in the sectors shown. Financed GHG emissions are calculated as the volume-weighted GHG-intensities per bank-year from all commercial loans over EUR 25,000, divided by the bank's revenues. All continuous variables are standardised. Standard errors are clustered at the bank-counterparty-month-level and reported in parentheses. Controls include total assets, the leverage ratio, the liquidity coverage ratio, the capital ratio, and the return on assets. $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

	(1)	(2)	(3)	(4)	(5)
Panel A: Non-financial corporations					
$FinancedEmissions$	0.066**	0.062***	0.055***	0.187***	0.209***
	(0.030)	(0.013)	(0.012)	(0.030)	(0.029)
$Haircut$			0.015*		0.025***
			(0.009)		(0.009)
$Volume$			0.010		0.021***
			(0.006)		(0.005)
$Maturity$			0.005***		0.000
			(0.001)		(0.001)
N	466 410	466 410	466 410	466 410	466 410
R^2	0.20	0.45	0.50	0.51	0.54
Panel B: Financial corporations					
$FinancedEmissions$	0.064***	0.029***	0.025***	0.039***	0.039***
	(0.005)	(0.005)	(0.005)	(0.006)	(0.006)
$Haircut$			−0.005		0.003
			(0.004)		(0.002)
$Volume$			0.091***		0.080***
			(0.007)		(0.006)
$Maturity$			0.003***		0.002***
			(0.000)		(0.000)
N	1 095 577	1 095 577	1 095 577	1 095 577	1 095 577
R^2	0.07	0.22	0.31	0.40	0.45
Panel C: General government					
$FinancedEmissions$	0.020***	−0.003	−0.004	0.035***	0.043***
	(0.007)	(0.007)	(0.004)	(0.005)	(0.005)
$Haircut$			−0.013***		−0.022***
			(0.004)		(0.004)
$Volume$			0.029***		0.025***
			(0.003)		(0.003)
$Maturity$			0.001***		0.002***
			(0.000)		(0.000)
N	434 832	434 832	434 832	434 832	434 832
R^2	0.02	0.16	0.49	0.36	0.54
Controls	✓	✓	✓	✓	✓
FE: Collateral ISIN			✓		✓
FE: Lender		✓	✓		
FE: Date		✓	✓		
FE: Lender × Borrower				✓	✓

Table 3.A.2. Size and leverage effects of banks with high financed GHG emissions.

Notes: The dependent variable is the loan rate minus the ECB's deposit facility rate expressed in percentage points. Financed GHG emissions are calculated as the volume-weighted GHG-intensities per bank-year from all commercial loans over EUR 25,000, divided by the bank's revenues. All continuous variables are standardised. Standard errors are clustered at the bank-counterparty-month-level and reported in parentheses. Controls include the leverage ratio, the liquidity coverage ratio, and the capital ratio for the size regressions, and the returns on assets as well as the size for the capital ratio and the liquidity coverage ratio regressions. $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

	(1)	(2)	(3)	(4)	(5)
$FinancedEmissions$	0.067***	0.044***	0.073***	0.064***	0.079***
	(0.007)	(0.006)	(0.006)	(0.007)	(0.007)
$D^{LargeSize}$	0.040***	−0.033***			−0.003
	(0.011)	(0.011)			(0.011)
$D^{LargeSize} \times FinancedEmissions$	−0.050***	−0.027***			−0.021**
	(0.010)	(0.009)			(0.010)
$D^{LowCapital}$			0.088***	0.136***	0.089***
			(0.012)	(0.013)	(0.012)
$D^{LowCapital} \times FinancedEmissions$			−0.066***	−0.071***	−0.056***
			(0.009)	(0.009)	(0.009)
$Volume$	0.073***	0.076***	0.073***	0.073***	0.073***
	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)
$Maturity$	0.002***	0.002***	0.002***	0.002***	0.002***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
N	2 006 742	2 006 742	2 006 742	2 006 742	2 006 742
R^2	0.41	0.43	0.41	0.41	0.41
Controls		✓		✓	
FE: Lender	✓	✓	✓	✓	✓
FE: Date	✓	✓	✓	✓	✓
FE: Collateral ISIN	✓	✓	✓	✓	✓

Table 3.A.3. The amplification effect of financed GHG-emissions.

The dependent variable is the loan rate minus the ECB's deposit facility rate expressed in percentage points. Individual regressions are based on the sub-samples of transactions collateralised by bonds issued in the sectors shown. Financed GHG emissions are calculated as the volume-weighted GHG-intensities per bank-year from all commercial loans over EUR 25,000, divided by the bank's revenues. The variable $FinRisk$ is a dummy taking the value one whenever the OFR Financial Stress Index is in its top quartile. All continuous variables are standardised. Standard errors are clustered at the bank-counterparty-level and reported in parentheses. $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

	All collateral	All collateral	Non-financial corporations	Non-financial corporations	Financial corporations	Financial corporations	General government	General government
$FinancedEmissions$	0.020***	0.036***	0.019*	0.187***	0.014***	0.030***	−0.007*	0.034***
	(0.004)	(0.006)	(0.011)	(0.030)	(0.005)	(0.006)	(0.004)	(0.005)
$FinancedEmissions \times FinRisk$	0.052***	0.017**	0.159***	0.043**	0.049***	0.013*	0.016*	0.008
	(0.010)	(0.007)	(0.029)	(0.022)	(0.010)	(0.007)	(0.010)	(0.007)
$Haircut$	−0.083***	0.007	0.011	0.035*	−0.037***	0.007*	−0.031***	−0.032***
	(0.009)	(0.006)	(0.017)	(0.018)	(0.012)	(0.004)	(0.006)	(0.005)
$Volume$	0.079***	0.052***	0.010*	0.022***	0.095***	0.082***	0.030***	0.026***
	(0.004)	(0.003)	(0.006)	(0.005)	(0.008)	(0.005)	(0.003)	(0.004)
$Maturity$	0.002***	0.001***	0.004***	0.000	0.003***	0.001***	0.001***	0.002***
	(0.000)	(0.000)	(0.001)	(0.001)	(0.000)	(0.000)	(0.000)	(0.000)
N	1 968 692	1 968 692	457 400	457 400	1 074 985	1 074 985	426 563	426 563
R^2	0.44	0.50	0.51	0.54	0.32	0.45	0.49	0.54
Controls	✓	✓	✓	✓	✓	✓	✓	✓
FE: Collateral ISIN	✓	✓	✓	✓	✓	✓	✓	✓
FE: Date	✓		✓		✓		✓	
FE: Lender	✓		✓		✓		✓	
FE: Lender × Borrower		✓		✓		✓		✓

Table 3.A.4. The amplification effect of financed GHG-emissions - Robustness test using VIX instead of OFR index

The dependent variable is the loan rate minus the ECB's deposit facility rate expressed in percentage points. Individual regressions are based on the sub-samples of transactions collateralised by bonds issued in the sectors shown. Financed GHG emissions are calculated as the volume-weighted GHG-intensities per bank-year from all commercial loans over EUR 25,000, divided by the bank's revenues. The variable VIX is a dummy taking the value one whenever the VIX Index is in its top tercile. All continuous variables are standardised. Standard errors are clustered at the bank-counterparty-level and reported in parentheses. $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

	All collateral	All collateral	Non-financial corporations	Non-financial corporations	Financial corporations	Financial corporations	General government	General government
$FinancedEmissions$	0.011**	0.038***	0.023**	0.209***	0.006	0.031***	−0.012***	0.030***
	(0.005)	(0.006)	(0.011)	(0.029)	(0.006)	(0.006)	(0.004)	(0.005)
$FinancedEmissions \times VIX$	0.039***	0.015***	0.066***	−0.020	0.039***	0.013***	0.018***	0.022***
	(0.006)	(0.004)	(0.016)	(0.014)	(0.007)	(0.004)	(0.005)	(0.004)
$Haircut$	−0.083***	0.009	0.010	0.039**	−0.037***	0.008*	−0.030***	−0.032***
	(0.009)	(0.006)	(0.017)	(0.018)	(0.012)	(0.004)	(0.006)	(0.005)
$Volume$	0.079***	0.052***	0.011*	0.020***	0.095***	0.083***	0.030***	0.026***
	(0.004)	(0.003)	(0.006)	(0.005)	(0.008)	(0.006)	(0.003)	(0.004)
$Maturity$	0.002***	0.001***	0.004***	0.000	0.003***	0.001***	0.001***	0.002***
	(0.000)	(0.000)	(0.001)	(0.001)	(0.000)	(0.000)	(0.000)	(0.000)
N	2 006 740	2 006 740	466 409	466 409	1 095 577	1 095 577	434 831	434 831
R^2	0.44	0.50	0.51	0.54	0.32	0.45	0.49	0.54
Controls	✓	✓	✓	✓	✓	✓	✓	✓
FE: Collateral ISIN	✓	✓	✓	✓	✓	✓	✓	✓
FE: Date	✓		✓		✓		✓	
FE: Lender	✓		✓		✓		✓	
FE: Lender × Borrower		✓		✓		✓		✓

Table 3.A.5. Robustness test: Transition risk premia including specific collateral transactions.

Notes: The dependent variable is the loan rate minus the ECB's deposit facility rate expressed in percentage points. Financed GHG emissions are calculated as the volume-weighted GHG-intensities per bank-year from all commercial loans over EUR 25,000, divided by the bank's revenues. All continuous variables are standardised. Standard errors are clustered at the bank-counterparty-level and reported in parentheses. $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

	(1)	(2)	(3)	(4)	(5)
$FinancedEmissions$	0.090***	0.060***	0.043***	0.070***	0.008
	(0.023)	(0.013)	(0.015)	(0.021)	(0.020)
$Haircut$		−0.237***	−0.187***	−0.046	−0.046
		(0.087)	(0.069)	(0.042)	(0.044)
$Volume$		0.098***	0.080***	0.052***	0.046***
		(0.015)	(0.011)	(0.010)	(0.009)
$Maturity$		0.003***	0.003***	0.002***	0.002***
		(0.001)	(0.001)	(0.000)	(0.000)
N	2 523 574	2 523 574	2 523 574	2 523 574	2 523 574
R^2	0.01	0.24	0.31	0.45	0.47
FE: Collateral ISIN		✓	✓	✓	✓
FE: Date			✓		✓
FE: Lender			✓		
FE: Lender × Borrower				✓	✓

Table 3.A.6. Robustness test: Lagged vs. Non-lagged GHG-exposure

Notes: The dependent variable is the loan rate minus the ECB's deposit facility rate expressed in percentage points. Financed GHG emissions are calculated as the volume-weighted GHG-intensities per bank-year from all commercial loans over EUR 25,000, divided by the bank's revenues. All continuous variables are standardised. Standard errors are clustered at the bank-counterparty-month-level and reported in parentheses. Controls include total assets, the leverage ratio, the liquidity coverage ratio, the capital ratio, and the return on assets. $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

	(1)	(2)
$FinancedEmissions(lag)$	0.031***	
	(0.004)	
$FinancedEmissions$		−0.006
		(0.004)
$Haircut$	−0.050***	−0.089***
	(0.009)	(0.010)
$Volume$	0.073***	0.073***
	(0.004)	(0.004)
$Maturity$	0.002***	0.003***
	(0.000)	(0.000)
N	2 006 742	1 571 265
R^2	0.43	0.57
Controls	✓	✓
FE: Lender	✓	✓
FE: Date	✓	✓
FE: Collateral ISIN	✓	✓