Drivers of Convenience Yields

Felix Corell
VU Amsterdam

Fédéric Holm-Hadulla
European Central Bank

Matteo Leombroni
Boston College

Lira Mota
MIT

Melina Papoutsi
European Central Bank

March 27, 2026

Abstract

Credit spreads cannot be explained by default risk alone; a significant non-default component reflects the "convenience" of safe sovereign bonds. Using comprehensive eurozone bond and holdings data, we estimate convenience premia for three services—liquidity, collateral value, and regulatory capital relief—and decompose AAA sovereign bond yields accordingly. Regulatory capital requirements have been the dominant driver of convenience yields over the past decade, reducing governments' borrowing costs by 18 bps—one-fifth of average sovereign yields. Exploiting policy-induced shocks, we show these services shape asset prices and that monetary and regulatory policies significantly affect bond markets through a "convenience yield channel."

*We thank Lorenzo Bretscher (discussant), Pierre Collin-Dufresne, Jens Dick-Nielsen (discussant), Marcio Garcia (discussant), Fotis Grigoris (discussant), Victoria Ivashina, Kristy Jansen (discussant), Ralph Koijen, Christian Kubitza, Atif Mian, Jonathan Parker, Yushi Peng (discussant), Gabor Pinter (discussant), Mirela Sandulescu (discussant), Shrihari Santosh (discussant), Ana Paula Serra, Kerry Siani, Larry Schmidt, Motohiro Yogo, conference participants at AEA, BIS-CEPR-Gerzensee-SFI Conference on Financial Intermediation, BSE Summer Forum Safety, Liquidity, and the Macroeconomy, CEPR Asset Pricing Paris Winter Meeting, Chicago Booth Asset Pricing Conference, IMF Fiscal Policy and Sovereign Debt Summit, Junior Finance Conference on Valuations, MFA, MIT Sloan Junior Finance Faculty Conference, Princeton Macro Finance Conference, Safety vs. Liquidity Workshop, SGF Conference 2025, University of Toronto Financial Economics Conference, and seminar participants at the European Central Bank, Goethe University, MIT Sloan, Princeton University, and Porto University. We thank Maui Brennan, Ziqi Li, Leon Loidl, Tomas Reis da Nobrega, and Nick von Turkovich for providing excellent research assistance. We acknowledge the contribution of the Eurosystem business areas that originate the data used, specifically the Eurosystem market operations departments as well as the CSDB and SHS teams. The views expressed in this paper are the authors' and do not necessarily reflect those of the European Central Bank or the Eurosystem. All errors and omissions are our own.

1 Introduction

The modern asset pricing literature recognizes that investors value financial assets not only for their promised cash flows, but also for the additional services they provide, such as liquidity, collateral pledgeability, and the ability to satisfy regulatory requirements.[1] A common approach to estimating the value of these services is to calculate "convenience yields"—the yield differential between assets that generate comparable cash flows, but different service flows. Although convenience yields explain a significant share of the time series and cross-sectional variation in corporate and sovereign bond valuations, they are typically measured as a residual term.[2] Consequently, convenience yield, albeit key to understanding the cost of capital for government and firm financing, remains an umbrella term that is silent about the precise forces driving the price gap.

In this paper, we estimate the premia associated with three key services that investors value in financial assets—liquidity, collateral value, and regulatory capital relief—and examine how different investors in the economy value these services. To do so, we exploit a comprehensive dataset combining bond prices, bond characteristics, and confidential granular sectoral portfolio holdings in the euro area (EA). Our empirical analysis is guided by a stylized model with heterogeneous investors in which assets provide services beyond their cash flows and investors differ in their valuation of these services. This framework allows us to quantify each service's contribution to convenience yields of AAA-rated sovereign bonds, and to provide causal evidence that policy-driven changes in these services affect investors' portfolio allocation and asset prices.

We show that over the past decade, the dominant driver of AAA-rated sovereign bond convenience yields in the euro area has been regulatory capital relief, i.e., the extent to which holding a bond affects investors' minimum capital requirements for regulated investors, like insurance companies. More specifically, the premium that sovereigns earn for providing regulatory capital relief accounts for 40–60% of their convenience yield throughout most of our

[1] See, e.g., Longstaff et al. (2005) and Bao et al. (2011) on liquidity premia, Chen et al. (2023) and Pelizzon et al. (2024) on collateral eligibility, and Jansen (2025) and Kubitza (2025) on regulatory constraints, among others.

[2] See, e.g., Krishnamurthy and Vissing-Jorgensen (2012a), Sunderam (2015), Gorton (2017), Mota (2023), Acharya and Laarits (2025), and Jiang et al. (2025).

sample period. This translates into AAA sovereigns financing themselves at yields roughly 35 basis points lower than other entities on average, with about 18 basis points attributable to the beneficial treatment of sovereign bonds under risk-based capital regulation.[3] Given the stock of euro area sovereign debt outstanding (approximately €13.2 trillion), an 18 basis point funding advantage corresponds to annual financing savings of roughly €24 billion per year. If sustained over the 27-year life of the euro, this would imply cumulative savings of approximately €650 billion, or about 4–5% of euro area current GDP. This institutional feature lowers governments' borrowing costs relative to other issuers and is consistent with mechanisms emphasized in the financial repression literature (McKinnon, 1973; Shaw, 1973).

These convenience yields are not borne equally by all investors. We leverage our comprehensive holdings data to show that intermediaries earn different returns because they value asset services differently. Banks and insurance companies and pension funds (ICPF) tilt their portfolios toward sovereign bonds and, within the corporate bond universe, toward bonds with high regulatory capital value. To quantify the cost of these portfolio tilts, we construct hedged returns that strip out compensation for credit and interest rate risk, thereby isolating the component of returns attributable to convenience yields. A key result is that, for corporate bond portfolios, banks and ICPF earn 3.5 percentage points lower cumulative hedged returns than mutual funds over our sample period—corresponding to roughly €50 billion in foregone financial returns. These return differences arise from the convenience yields embedded in the bonds that regulated intermediaries choose to hold, providing direct evidence that they accept lower financial returns in exchange for the regulatory services their portfolios provide.

Finally, we provide *causal* evidence that changes in each of our three services affect both portfolio allocation and convenience yields. We exploit policy changes that directly alter the value of each service. Specifically, we study the introduction of Solvency II, which harmonized risk-based capital requirements for insurance companies across the EU; a reform of the ECB's collateral framework, which reduced the pledgeability of a large class of bank bonds; and the ECB's Corporate Sector Purchase Programme (CSPP), which improved the

[3] In the euro area, sovereign bonds—regardless of rating—are exempt from credit risk capital charges, which mechanically gives them high regulatory value for regulated intermediaries such as banks and insurance companies.

liquidity of eligible corporate bonds. In each case, we show that the policy-induced shift in service flows leads to economically significant changes in both convenience yields and the portfolios of the investors who value the affected service most. These results validate our choice of service components and highlight an important channel through which regulatory and unconventional monetary policies transmit to asset prices and redistribute value across investors.

Our paper is organized into three parts. First, we build on Pástor et al. (2021) and present a model in which investors value not only cash flows but also derive utility from the service flows provided by financial assets. Investors differ in their valuation of these services. We define convenience yields as the equilibrium return differential between a benchmark economy in which investors value only cash flows and an economy in which assets also provide service flows. A key implication is that investors tilt their portfolios toward assets that provide the services they value most. Because these tilts concentrate holdings in high-service-flow assets, different investors earn different returns: the convenience yield embedded in each investor's portfolio reflects the return they are willing to forgo in exchange for the services their holdings provide.

Second, we operationalize our decomposition by constructing a measure of *relative convenience yields*. Guided by the model, we use the CDS-bond basis—defined as the difference between the yield of a risk-free AAA sovereign bond and that of a duration-matched risky bond whose default risk is fully hedged using a maturity-matched CDS—as our key empirical object. Because the CDS-hedged bond and the corresponding sovereign bond generate nearly identical cash flows when held to maturity, the CDS-bond basis captures the return investors are willing to forgo to hold the safe sovereign bond rather than the hedged bond. While the hedge is not perfect due to frictions in CDS markets, we argue these imperfections are small and unlikely to drive our results. The CDS-bond basis therefore measures the convenience yield of each bond relative to the corresponding AAA sovereign bond.

We then map the cross-sectional variation in CDS-bond bases to cross-sectional variation in bond-specific services to estimate the premium associated with each service. Liquidity is proxied using bond-level bid-ask spreads, collateral value using haircuts from the ECB collateral framework, and regulatory capital value using bond-specific capital charges implied

by risk-based regulation for regulated intermediaries.[4] Because the CDS-bond basis measures the convenience yield of a bond relative to AAA sovereign bonds, we run cross-sectional regressions of CDS-bond bases on the difference in each service between the bond and its duration-matched AAA sovereign counterpart. Following Fama and MacBeth (1973), this specification provides an estimate of the time-varying premium associated with each service.

We use the estimated premia to decompose the convenience yield of AAA euro area sovereign bonds into its underlying service components. While regulatory capital value is the dominant driver on average, the relative importance of the three services varies substantially over time. The regulatory capital premium spiked during the onset of COVID-19, while the liquidity and collateral premia rose markedly after 2022 as the ECB tightened monetary policy and began shrinking its balance sheet. Importantly, our three services explain over 90% of the level of the AAA sovereign CDS-bond basis, providing strong validation of the decomposition.

The second prediction of the model is that, because investors differ in their valuation of asset services, they are willing to forego returns to hold assets that provide these services. To test this prediction empirically, we link our measure of convenience yield to the portfolio holdings of different investor types. In particular, we focus on euro area banks, insurance companies and pension funds (ICPFs), and mutual funds. We find large differences in portfolio composition across these investor groups. Banks and ICPFs hold substantially larger shares of sovereign bonds relative to corporate bonds, consistent with the high regulatory capital value of sovereign debt. In contrast, mutual funds tilt their portfolios more strongly toward corporate bonds, which have wider CDS-bond bases but provide fewer regulatory services. As a result, the average portfolio basis differs markedly across investor types.

These differences translate into economically large variation in hedged portfolio returns. For corporate bonds, we compute investor-level returns that hedge both credit risk and interest rate risk, isolating the return component attributable to service flows. We find

[4]Specifically, we use the capital charges from the spread-risk module of the Solvency II regulation for insurance companies in the European Union. Under this framework, each bond is assigned a stress factor that depends on its credit rating and duration and determines the amount of regulatory capital required to hold the asset. Sovereign bonds issued by euro area governments are assigned a zero stress factor and therefore carry no spread-risk capital charges. This treatment is consistent with the regulatory framework applied to banks.

that the corporate bond portfolios held by banks and ICPFs earn approximately 20 basis points lower annual hedged returns than those held by mutual funds. While imperfections in the hedging procedure may introduce some noise, we argue these frictions are likely small. The return differences are therefore largely driven by the convenience yield embedded in the bonds each investor type holds, providing direct evidence that regulated intermediaries accept lower financial returns in exchange for the regulatory services provided by their assets.

The pattern is different for sovereign bonds. Within the sovereign bond universe, mutual funds tilt their portfolios toward more liquid bonds and are willing to forego returns to hold these assets. This highlights that different investor types value different services: regulated intermediaries primarily demand regulatory capital relief, while mutual funds place greater value on liquidity.

In the third part of the paper, we study policy shocks to the services, to provide causal evidence on how changes in services affect investors' portfolios and convenience yields. This serves two purposes: (a) it validates our choice of services for the decomposition of convenience yield, and (b) it illustrates how regulation, collateral policies, and central bank asset purchases are transmitted to asset prices through the convenience yield component.

First, we study how bond-level capital charges for European insurance companies affect convenience yield. We do so by analyzing how the CDS-bond basis and portfolio holdings behave in response to the EU-wide application of the Solvency II regulatory framework on January 1, 2016. Our approach is inspired by the finding in Jansen (2025) that Solvency II—in particular, the specific regulatory discount curve for insurers' liabilities—incentivized European insurers to hold bonds with maturities around 20 years. We show that after Solvency II became binding, the difference in convenience yields between bonds with low vs. high risk-based capital requirements (in the 20-year duration segment) increases by 20 basis points, and ICPF rebalance towards these high-regulatory value bonds. We interpret these results as evidence that regulatory constraints on the balance sheets of ICPF have a large impact on convenience yield.

Second, we focus on collateral value, which is particularly relevant for banks given their unique ability to borrow directly from the ECB against eligible securities. We exploit a 2018 change in the ECB's collateral framework that reduced the collateral value of a large

class of bank bonds. Following the reform, the average CDS-bond basis of affected bonds fell significantly (by 4–7 basis points, equivalent to 10% for the average bond) relative to otherwise similar bank bonds. Moreover, banks rebalanced their portfolios toward securities with lower haircuts, consistent with a particular preference for bonds with high collateral value.

Third and finally, we study the impact of the ECB's corporate bond quantitative easing (QE) program, which we interpret as a positive shock to the liquidity of eligible bonds. In our companion paper (Corell et al., 2025), we show that corporate QE increases the liquidity of eligible relative to ineligible bonds—a result consistent with earlier findings in Todorov (2020). By acting as a standing buyer, the central bank also reduces the daily liquidity and liquidity risk of eligible bonds (Pastor and Stambaugh (2003); Acharya and Pedersen (2005); Haddad et al. (2025)).[5] We focus on the ECB's Corporate Sector Purchase Program (CSPP), announced in March 2016, which initiated large-scale acquisitions of investment-grade, non-bank corporate bonds from the euro area. We show that the QE announcement caused the CDS-bond basis of eligible bonds to increase by 13.1 basis points relative to ineligible bonds, which accounts for 60% of the total effect on spreads. Consistent with our interpretation of QE enhancing eligible bonds' liquidity, we show that mutual funds—who are subject to redemption risk and thus prefer liquid assets—rebalance their portfolio towards QE-eligible bonds, despite their relative price increase. This finding demonstrates that corporate QE affects not only eligible bonds' default risk, but also their service flows.

Our work offers important insights into corporate bond pricing and how it is shaped by policy decisions. For example, demand-based asset pricing uses characteristics known to predict returns to estimate investors' heterogeneous demand curves and their impact on asset prices (Koijen and Yogo, 2019; Bretscher et al., 2025). However, our paper shows that characteristics driving risk and returns alone are not sufficient to fully explain demand for assets. Preferences for specific service flows might be one of the reasons why investors'

[5] Although for most of the paper our liquidity measure relies on bid-ask spreads, we recognize that liquidity extends beyond traditional measures of transaction costs. Broader notions of liquidity risk are inherently difficult to quantify, and we do not attempt to isolate them in our decompositions of convenience yields. We argue, however, that this broader notion is the right lens through which to interpret the impact of QE, as we discuss in our companion paper (Corell et al., 2025).

demand for bonds appears so inelastic to price variation. Furthermore, our empirical tests show that asset-specific service flows vary over time, particularly in response to regulatory and monetary policy changes. Therefore, models that assume time-invariant demand for assets such as Treasuries, due to their non-pecuniary benefits, cannot accurately estimate the impact of supply channels on bond yields. Fundamentally, such models will miscalculate the ability of governments and firms to issue bonds without affecting prices. This demand elasticity is key to understanding governments' fiscal space (Mian et al., 2022).

In addition, the direct impact of regulation and monetary policy on asset characteristics, such as liquidity or collateral value, highlights an important consideration for policymakers regarding the effects of policies on asset prices and capital allocation. In a world with high public debt, particularly among governments of developed countries, actions that force regulated institutions to hold certain assets have a large impact on asset prices and wealth allocation. A large literature on "financial repression" has studied the important role of banks holding sovereign debt (some recent references are Becker and Ivashina (2018), Chari et al. (2020), Jeanne (2025), Payne and Szoke (2025)). Our work shows that the focus should not only be on banks, but also on long-term investors like insurance companies and pension funds that are incentivized to hold certain assets due to the regulatory framework and duration exposure. Furthermore, even if indirectly, policies such as the collateral framework imposed by the ECB affect investors who value collateral, and therefore prices and portfolio allocation. In summary, our paper highlights that policy decisions have a large and heterogeneous impact on the prices of corporate and sovereign bonds due to a"convenience yield channel".

1.1 Related Literature

This paper connects several important strands of literature. First, we build on the growing research on safe assets and convenience yields, which highlights the large impact of convenience yields on bond valuation and therefore on the cost of capital for countries and firms. Following Krishnamurthy and Vissing-Jorgensen (2012b) and Mota (2023), we exploit the fact that corporate bonds are imperfect substitutes for traditional safe assets like Treasuries and command their own convenience yields. Key to our analysis is that both asset classes carry distinct service flows, implying that variation in the value of these services acts as

a common factor linking convenience yields across different fixed income assets. Our contribution is to exploit the rich cross-section of assets to identify the value of each service component, which is considerably harder to achieve when studying a single asset class in isolation.

We are not the first to study the nature of convenience yields. Acharya and Laarits (2025) explore the time-varying nature of Treasury convenience yields and their connection to broader market conditions. Jiang et al. (2025) study the link between fiscal surpluses and convenience yields, showing that convenience yields act as an important fiscal shock absorber. Fontana and Scheicher (2016) study the CDS-bond basis of euro area sovereign bonds, documenting that convenience yields explain a large share of the variation in sovereign yields. Our contribution relative to this literature is to analyze the full cross-section of euro area corporate and sovereign bond convenience yields, decompose them into their key service components, map these services to the portfolios of different investor types, and provide causal evidence on how regulation and unconventional monetary policy affect investor portfolio allocations and convenience yields. A key conclusion is that no single force can explain convenience yields in isolation; rather, they are shaped by multiple, distinct service flows whose relative importance varies across assets and investors.

Second, we contribute to the extensive literature on corporate bond pricing. Numerous studies establish that default risk alone cannot explain corporate bond valuations (Elton et al., 2001; Huang and Huang, 2012) and that liquidity significantly affects prices (see Amihud et al. (2006) and Vayanos and Wang (2013) for surveys). Pelizzon et al. (2024) document an "eligibility premium" in European corporate bonds when they become eligible as collateral in ECB repo operations, while Chen et al. (2023) show analogous effects in China. Bai and Collin-Dufresne (2019) show that multiple measures of arbitrage limits, including collateral value and trading liquidity, affect the US corporate CDS-bond basis. We build on the idea that corporate bond valuation reflects various factors beyond cash flows. Our innovation is to map cross-sectional variation in the basis to differences in service flows between the benchmark sovereign and the risky bond, which allows us to apply traditional asset pricing techniques to the cross-section of corporate and sovereign bonds and attribute yield spread premiums to distinct service components.

Third, our paper advances the literature on clientele and preferred habitat effects on asset prices (Vayanos and Vila, 2021). Motivated by models where heterogeneous liquidity needs shape equilibrium asset allocation (Amihud and Mendelson, 1986; Chen et al., 2020), we extend the framework to incorporate pledgeability (Chen et al., 2023) and regulatory capital relief (Becker and Ivashina, 2015; Koijen and Yogo, 2022a; Sen, 2023). Our model captures demand for multiple services in reduced form, a simplifying choice that enables a direct mapping between model primitives and the granular holdings data we observe, allowing us to document heterogeneity in both portfolio composition and the valuation of different services across investor types. By leveraging pricing and holdings data jointly, we show how these heterogeneous preferences for service flows drive both portfolio decisions and bond valuation, contributing to the emerging literature on demand-based asset pricing (Koijen and Yogo, 2019; Gabaix and Koijen, 2020).

Fourth, our work also contributes with the recent literature on how regulatory and business-model-related constraints affect portfolio decisions for insurance companies and pension funds (ICPF) (see, e.g., Koijen and Yogo (2015), Koijen and Yogo (2022b)), and how these in turn affect asset prices (Kubitza, 2025; Li, 2025). Specifically, due to their negative duration gap, ICPF have an incentive to hedge against falling interest rates to maintain solvency and—in the case of insurers—comply with Solvency II regulation (Jansen, 2025). Our paper contributes to this literature by showing how these constraints have a large and direct impact on governments' cost of funding, as regulatory capital alone accounts for up to 60% of the AAA sovereign convenience yield. Moreover, we confirm (and extend) Jansen (2025) by showing that in response to a reform in the context of Solvency II, European ICPFs shifted their portfolios not only towards bonds with maturities around 20 years, but also toward high regulatory value bonds, affecting both the term structure of government bond yields and credit spreads.

In summary, our paper makes four main contributions to these literature strands. First, we take the idea seriously that convenience yields are multidimensional and cannot be explained by a single force alone. We provide a methodology to decompose convenience yields into their underlying service components, and we apply it to liquidity, collateral value, and regulatory capital relief. Regulatory capital emerges as the dominant driver of AAA sovereign

convenience yields in the euro area. Second, we map cross-sectional variation in the CDS-bond basis to differences in service flows between benchmark sovereigns and risky bonds, allowing us to apply traditional asset pricing techniques to attribute yield spread premiums to distinct service components. Third, we connect convenience yields to investor heterogeneity by leveraging confidential granular holdings data to show that different investor types value different services, leading to economically large differences in hedged portfolio returns: banks and ICPFs earn approximately 3.5 percentage points lower cumulative hedged returns than mutual funds on their corporate bond portfolios. Fourth, we provide causal evidence—exploiting the introduction of Solvency II, a reform of the ECB's collateral framework, and the launch of the CSPP—that policy-induced changes in each service flow significantly affect both convenience yields and investors' portfolio allocations, establishing a "convenience yield channel" through which regulatory and monetary policies transmit to asset prices.

2 Model

In this section, we develop an asset pricing model in which investors value financial assets not only for their cash flows but also for the service attributes they provide. Our model builds on Pástor et al. (2021), generalizing their framework from a single non-pecuniary benefit (greenness) to accommodate multiple service dimensions simultaneously. The model characterizes how convenience yields arise endogenously from heterogeneous investor preferences over these services and how they relate to equilibrium portfolio allocations. The model delivers a set of clear insights that guide our empirical analysis.

There are N assets indexed by $n \in \{1, 2, \ldots, N\}$ and a continuum of heterogeneous investors indexed by i, each with distinct preferences over the various service attributes. There exists a risk-free asset with return r_t^f in zero net supply. Let $\boldsymbol{r}_{t+1}$ be the $(N \times 1)$ vector of returns in excess of the risk-free rate $\boldsymbol{r}_{t+1} = \boldsymbol{\mu}_t + \boldsymbol{\epsilon}_{t+1}$, where $\boldsymbol{\mu}_t$ contains the expected equilibrium excess returns and $\boldsymbol{\epsilon}_{t+1} \sim \mathcal{N}(\mathbf{0}, \Sigma_t)$.

In addition to returns, assets provide a set of service flows. The services are collected in $\boldsymbol{Z}_t$, a $(N \times K)$ matrix of asset characteristics that determines the quantity of service k provided by asset n. Agents have heterogeneous valuations of services, collected in the loading vector $\tilde{\boldsymbol{\lambda}}_{it}$, a $(K \times 1)$ vector that represents how much investor i values each characteristic

k. The service flows therefore have an asset-specific and an agent-specific component:

$$\tilde{\boldsymbol{s}}_{it} = \boldsymbol{Z}_t \tilde{\boldsymbol{\lambda}}_{it}. \tag{1}$$

Investor i's wealth at time $t+1$ is given by $w_{i,t+1} = w_{i0}\left(1 + r_t^f + \boldsymbol{x}_{it}'\boldsymbol{r}_{t+1}\right)$ where w_{i0} denotes initial wealth and $\boldsymbol{x}_{it}$ is the vector of portfolio weights. Each investor has exponential utility over next period wealth and services:

$$U_i(w_{i,t+1}, \boldsymbol{x}_{it}) = -\exp(-A_i w_{i,t+1} - \boldsymbol{x}_{it}'\tilde{\boldsymbol{s}}_{it}),$$

where A_i is the agent's absolute risk aversion and $\boldsymbol{x}_{it}'\tilde{\boldsymbol{s}}_{it}$ is the total value of service flows the investor derives from her portfolio.

We model the demand for services in a reduced-form framework which, in many applications, is isomorphic to modeling mean–variance agents subject to regulatory, liquidity, or collateral constraints (e.g. Frazzini and Pedersen (2014)). Moreover, even if investors are myopic in our model, dynamic hedging demand à la ICAPM (Merton (1973)) is also nested in our framework, as hedging properties manifest as an additional service flow. The key advantage of our approach is its ability to incorporate multiple frictions within a unified framework.

The investor's problem admits a closed-form solution. As shown in Appendix A, the investor's optimal allocation is given by

$$\boldsymbol{x}_{it} = \frac{1}{a_i}\Sigma_t^{-1}\left(\boldsymbol{\mu}_t + \boldsymbol{Z}_t\tilde{\boldsymbol{\lambda}}_{it}\right), \tag{2}$$

where $a_i = A_i w_{i0}$ is agent i's relative risk aversion.

For tractability, we assume that $a_i = a$ for all agents and define $\boldsymbol{\lambda}_{it} = \tilde{\boldsymbol{\lambda}}_{it}/a$ the scaled valuation vector of service flows.

Expected Excess Returns We can use market clearing to compute expected returns. Let $\tilde{w}_{i0} := w_{i0}/w_0$, where w_0 denotes total wealth in the economy. Because the risk-free asset is in zero net supply, market clearing implies that the weights of the market portfolio $\boldsymbol{x}_{mt}$ satisfy

$$\begin{aligned} \boldsymbol{x}_{mt} &= \int_i \tilde{w}_{i0}\boldsymbol{x}_{it}\, di \\ &= \frac{1}{a}\Sigma_t^{-1}\left(\boldsymbol{\mu}_t + \boldsymbol{Z}_t\bar{\boldsymbol{\lambda}}_t\right), \end{aligned} \tag{3}$$

where $\bar{\boldsymbol{\lambda}}_t = \int_i \tilde{w}_{i0} \boldsymbol{\lambda}_{it} \, di$ denotes the wealth-weighted valuation of service flows across investors.

For simplicity, we further assume that the market portfolio has a zero market-weighted service vector, i.e., $\boldsymbol{x}'_{mt} \boldsymbol{Z}_t = \mathbf{0}$. Equivalently, we can view $\boldsymbol{Z}_t$ as being defined so that agents derive utility (disutility) from holdings that provide more (or less) services than the market. As we show in Appendix A, expected excess returns are given by

$$\boldsymbol{\mu}_t = \boldsymbol{\beta}_t \mu_{mt} - \boldsymbol{Z}_t \bar{\boldsymbol{\lambda}}_t, \tag{4}$$

where $\boldsymbol{\beta}_t = \frac{\Sigma_t \boldsymbol{x}_{mt}}{\sigma^2_{mt}}$ denotes the vector of market betas and μ_{mt} is the expected excess return on the market portfolio. Equation 4 shows that the CAPM does not hold. In equilibrium, two assets with the same exposure to market risk, $\beta_t(n)$, can have different expected returns because they provide different service flows $\boldsymbol{Z}_t$. We capture this wedge by defining the convenience value of asset n as

$$cy_t(n) = \boldsymbol{Z}_t(n) \, \bar{\boldsymbol{\lambda}}_t, \tag{5}$$

which represents the returns that investors are willing to forgo because of the additional benefits the asset provides.

Investors' Heterogeneous Allocation Using equation (2) together with market clearing, investor i's allocation can be written as

$$\boldsymbol{x}_{it} = \boldsymbol{x}_{mt} + \frac{1}{a} \Sigma_t^{-1} \boldsymbol{Z}_t \left(\boldsymbol{\lambda}_{it} - \bar{\boldsymbol{\lambda}}_t \right). \tag{6}$$

This equation clearly shows that, compared to a benchmark model without service flows, investors deviate from the standard tangency portfolio x_m and tilt their portfolios toward assets that provide services $\boldsymbol{Z}_t$. This represents a deviation from the two-fund separation theorem. Private valuations for services cause investors to tilt their portfolios away from the market portfolio, and the magnitude of this tilt increases with investors' preferences for services, $\boldsymbol{\lambda}_{it}$, relative to the wealth-weighted average $\bar{\boldsymbol{\lambda}}_t$.

2.1 Model Predictions

We derive three empirical predictions, which we test in the remainder of the paper. The proofs are in Appendix A.

Prediction 1. *The cross-sectional variation in convenience yields is driven by cross-sectional variation in bond-specific service flows,* $\boldsymbol{Z}_t$.

As shown in equation (5), the convenience yield of asset n is given by $cy_t(n) = \boldsymbol{Z}_t(n)\,\bar{\boldsymbol{\lambda}}_t$. Thus, differences in expected returns across bonds with identical exposure to systematic risk arise solely from differences in the services they provide and from the equilibrium price of these services, $\bar{\boldsymbol{\lambda}}_t$.

There is a large literature exploring the characteristics that drive expected returns and their relationship to exposure to systematic risk factors (Fama and French, 1993; Daniel and Titman, 1997; Daniel et al., 2020). In our setting, not only characteristics related to cash flows, but also characteristics related to service flows explain returns. Thus, omitting characteristics related to service flows would lead to misspecified asset pricing models. The private valuation of these specific services causes bonds to have different prices, even when bonds have identical expected cash flows.

Prediction 2. *Investors hold portfolios with different convenience yields. In particular, investors whose service valuations are high relative to those of the market hold portfolios with higher convenience yields and, conditional on identical exposure to systematic risk, earn lower expected excess returns. Specifically, the expected return differential between investors* i *and* j *is*

$$\boldsymbol{x}'_{it}\boldsymbol{\mu}_t - \boldsymbol{x}'_{jt}\boldsymbol{\mu}_t = -(cy^p_{it} - cy^p_{jt}), \tag{7}$$

where $cy^p_{it} = \boldsymbol{x}'_{it}\boldsymbol{Z}_t\bar{\boldsymbol{\lambda}}_t$ *denotes the portfolio-level convenience yield of investor* i.

By equation (6), investor i tilts her portfolio away from the market toward assets whose service flows match her preferences. Using the normalization $\boldsymbol{x}'_{mt}\boldsymbol{Z}_t = \boldsymbol{0}$, the portfolio convenience yield can be expressed as

$$cy^p_{it} = \frac{1}{a}(\boldsymbol{\lambda}_{it} - \bar{\boldsymbol{\lambda}}_t)'\boldsymbol{Z}'_t\Sigma_t^{-1}\boldsymbol{Z}_t\,\bar{\boldsymbol{\lambda}}_t, \tag{8}$$

which is increasing in the alignment between investor i's service valuations, $\boldsymbol{\lambda}_{it}$, and the aggregate valuation, $\bar{\boldsymbol{\lambda}}_t$.

To fix ideas, consider the case $K = 1$. The expression simplifies to

$$cy^p_{it} = \frac{\bar{\lambda}_t}{a}(\lambda_{it} - \bar{\lambda}_t)\,\boldsymbol{z}'_t\Sigma_t^{-1}\boldsymbol{z}_t. \tag{9}$$

Since $\boldsymbol{z}_t' \Sigma_t^{-1} \boldsymbol{z}_t > 0$ and $\bar{\lambda}_t > 0$, the portfolio convenience yield is monotonically increasing in λ_{it}, and investor i underperforms investor j if and only if $\lambda_{it} > \lambda_{jt}$.

Prediction 3. *Holding all else constant, when investor i's preference for service k increases, the convenience yield rises more for assets with higher Z_t^k. Conditional on risk exposure, the investor increases her allocation toward assets with higher Z_t^k relative to assets with lower Z_t^k.*[6]

An increase in $\lambda_{it}(k)$ raises the wealth-weighted valuation of service k, $\bar{\lambda}_t(k)$. Through equation (5), this increases the convenience yield of assets with higher Z_t^k and, by equation (4), lowers their expected excess returns.

Moreover, equation (6) implies that investors hold heterogeneous portfolios depending on their valuations of services, $\boldsymbol{\lambda}_{it}$. An increase in investor i's preference for service k induces a portfolio tilt toward assets that provide more of that service, conditional on risk exposure, creating a direct link between preference heterogeneity and portfolio heterogeneity across investors.

Prediction 4. *Holding all else constant, an increase in the service flow k of asset n, $Z_t^k(n)$, induces investors with a higher valuation for service k, to rebalance more strongly toward asset n than investors with lower valuations. As a result, the convenience yield of asset n increases.*

This proposition characterizes the effects of a change in asset characteristics, rather than a change in investor preferences. Equation (6) implies that, for a given distribution of preferences, investors with a higher valuation for service k respond more strongly to an increase in $Z_t^k(n)$ by increasing their holdings of asset n. The resulting increase in demand raises the convenience yield of asset n.

Together, Prediction 3 and 4 highlight two distinct but complementary channels through which service flows affect asset prices: changes in investor demand driven by preference heterogeneity, and changes in asset characteristics that alter the service flow. Both channels

[6] By *risk exposure*, we mean that assets share the same covariance structure of returns (equivalently, the relevant risk adjustment through Σ_t^{-1} is the same within the comparison set).

generate portfolio rebalancing across investors and lead to systematic differences in expected returns across assets.

3 Data and Measurement of Service Flows

Our objective is to exploit variation in convenience yields across assets and across investors' portfolios to infer the nature of the service flows valued by investors. This exercise requires a rich cross-section of securities with similar cash-flow characteristics but that provide different service flows and returns, together with substantial heterogeneity in holdings across investor types. Using confidential ECB data, we construct a panel of euro area corporate and sovereign bonds containing information on yields, CDS spreads, and a broad set of bond characteristics. We then link these bond-level data to the portfolio holdings of the main euro area sectors, including banks, mutual funds, insurance companies, and pension funds.

3.1 Bond Price Data

The first dataset comes from the Centralised Securities Database (CSDB), which is maintained by the Eurosystem and consists of monthly security-by-security information on instruments, issuers, and prices for debt securities traded in the euro area. For the purpose of our analysis, we restrict our sample to senior unsecured bonds with fixed or zero coupon that are issued by euro area entities (corporate or sovereign), euro-denominated, and that have a residual maturity of at least one year. Moreover, we exclude bonds with an outstanding amount of less than €10 million.

For each outstanding bond, we observe its issuer, issuance size, maturity date, price and yield information, as well as other bond characteristics. Secondary market prices are recorded in a monthly panel. We enrich the dataset with long-term credit ratings from the four main rating agencies: Standard and Poor's, Fitch, Moody's, and DBRS. When available, we use the best instrument-level rating; otherwise, we rely on the best issuer-level rating across agencies. In addition, we incorporate bond-level measures of monthly turnover and average bid-ask spreads from the TraX dataset provided by MarketAxess. Finally, for each bond in the sample, we collect its coupon payment schedule and compute modified duration.[7]

[7]For a subset of bonds, we obtain duration estimates from the iBoxx index provider. We validate our calculations

For risk-free rates, we use the ECB benchmark interest rate provided under the euro area yield curves estimation.[8] These are estimated daily using AAA-rated euro area central government bonds. Figure C.1 in the appendix presents the time series of the implied benchmark interest rate for different maturities. We merge the benchmark interest rate with the CSDB data based on the month-year and on the bond's duration to compute the credit spread for each bond. We thus obtain credit spreads for a sample of 29,869 bonds (27,726 of which are corporate) by 2,335 distinct issuers across the sample.

To hedge bond credit risk, we use credit default swaps (CDS). Specifically, we merge the CSDB bond data with single-name CDS data from Markit. The merge relies on the Markit reference obligation mapping, LEI codes, name string matching, and manual verification. CDS spreads are quoted in basis points per annum for a notional amount of $10 million and correspond to the standard ISDA contract with physical settlement. The raw dataset provides daily CDS spreads under several restructuring documentation clauses. Following standard practice, we focus on spreads quoted under the "modified restructuring" (MR) clause, which is the benchmark contract type in Europe. Markit reports a term structure of CDS spreads across multiple maturities.[9] We linearly interpolate these spreads to obtain maturity-matched CDS spreads for each bond. Markit also provides a measure of quote reliability, and we restrict our sample to CDS spreads with an average quote rating above BBB.[10] Finally, to align the seniority and guarantee features of the underlying bonds, we only consider senior unsecured euro-denominated CDS contracts.

The combined dataset allows us to calculate the *CDS-bond basis* as the difference between a bond's CDS spread and its credit spread. We discard all observations below the 1st and above the 99th monthly percentile of the CDS-bond basis. Our final sample consists of 17,406 distinct bonds (issued by 486 distinct entities) and 574,805 bond-month observations between January 2014 and August 2025. In terms of outstanding amounts, this CDS-matched sample covers approximately 80% of the initial CDSB sample. Summary statistics regarding all variables used in our analysis can be found in Table 1.

by achieving a 99% correlation between our estimated durations and those provided by iBoxx.

[8] These benchmark rates are published by the ECB on this website.

[9] Specifically, the maturities included in the data set are 1, 2, 3, 4, 5, 7, 10, 15, 20, and 30 years.

[10] This filter mitigates the concern that the CDS market is most liquid for contracts with five-year maturity.

Table 1: Summary Statistics. This table contains descriptive statistics about key variables used during our analysis, separately for corporate and sovereign bonds. *Modified duration* measures the percentage bond price change in response to a change in yields. *Credit spreads* are the difference between a bond's yield to maturity and the duration-matched AAA sovereign bond yield. The *CDS-bond basis* is the difference between a bond's CDS spread and credit spread. *Collateral value* is calculated as $1 - haircut$, *regulatory capital value* as $1 - stress$, where the stress parameter is an increasing function of duration and credit risk, defined in Art. 176 of Delegated Regulation (EU) 2015/35. *Bid-ask spreads* are calculated as $\frac{ask-bid}{bid} \times 100$. The sample period is January 2014 – August 2025.

(a) Corporate bonds

	Mean	SD	p10	Median	p90	N
Amount outstanding (bln EUR)	0.33	0.63	0.01	0.05	1.00	468,795
Residual maturity (years)	6.08	5.22	1.62	4.66	11.68	468,795
Modified duration (%)	5.44	4.15	1.58	4.39	10.22	468,795
Yield to maturity (%)	1.71	1.53	0.03	1.21	3.92	468,795
Credit spread (%)	1.10	0.68	0.44	0.97	1.90	468,795
CDS-bond basis (%)	-0.40	0.55	-1.08	-0.39	0.20	468,795
Collateral value	0.60	0.41	0.00	0.85	0.95	468,795
Regulatory capital value	0.93	0.05	0.87	0.94	0.98	468,795
Bid-ask spread (%)	0.49	0.47	0.14	0.38	0.93	153,833

(b) Sovereign bonds

	Mean	SD	p10	Median	p90	N
Amount outstanding (bln EUR)	7.35	10.16	0.04	1.43	22.37	105,988
Residual maturity (years)	11.84	10.29	2.16	8.32	26.83	106,010
Modified duration (%)	10.27	8.46	2.09	7.64	22.02	106,010
Yield to maturity (%)	1.51	1.50	-0.30	1.25	3.57	106,010
Credit spread (%)	0.75	0.80	0.03	0.57	1.65	106,010
CDS-bond basis (%)	-0.29	0.46	-0.88	-0.22	0.16	106,010
Collateral value	0.91	0.20	0.90	0.95	0.99	106,010
Regulatory capital value	1.00	0.00	1.00	1.00	1.00	106,010
Bid-ask spread (%)	0.22	0.41	0.02	0.08	0.60	68,953

Figure 1: Average credit and CDS spread for euro area bonds by rating. This plot presents average credit and CDS spreads (weighted by outstanding amounts) for each rating category, separately for corporate and sovereign bonds. *AAA – A* includes all bonds with a rating of at least *A-*; *BBB* includes bonds with ratings between *BBB+* and *BBB-*; and *HY* includes all rated bonds with a rating of *BB+* or below. For this plot, we excluded bonds of three Italian banks (Banco Popolare, Banca Monte dei Paschi di Siena, and UBI Banca) that exhibited excessively volatile CDS spreads in 2012–2013, possibly related to the ECB's OMT announcement. The time series for high-yield sovereign debt stop in September 2023 when the rating agency DBRS upgraded Greek sovereign debt to the investment-grade category.

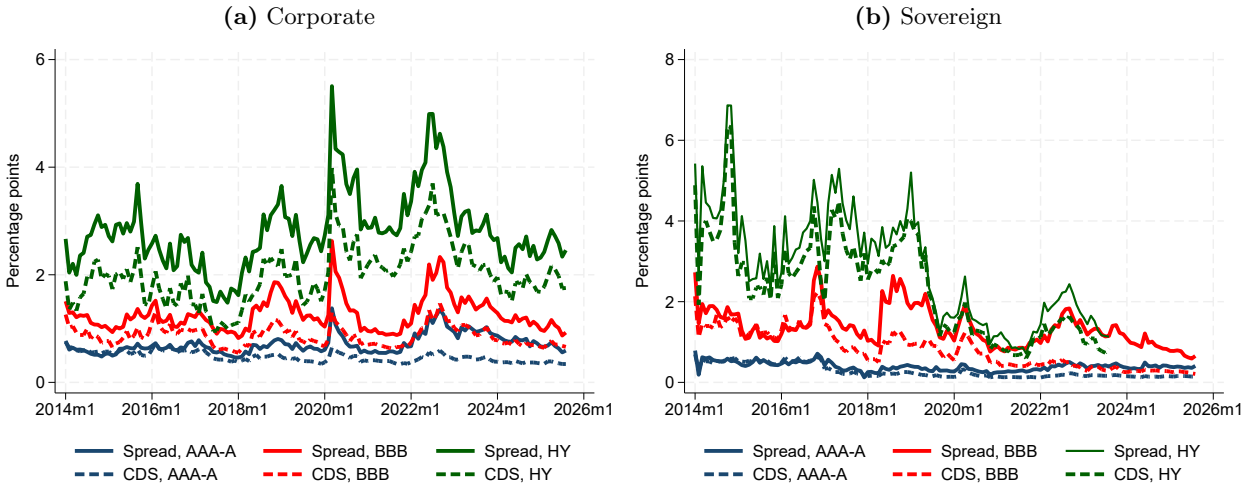
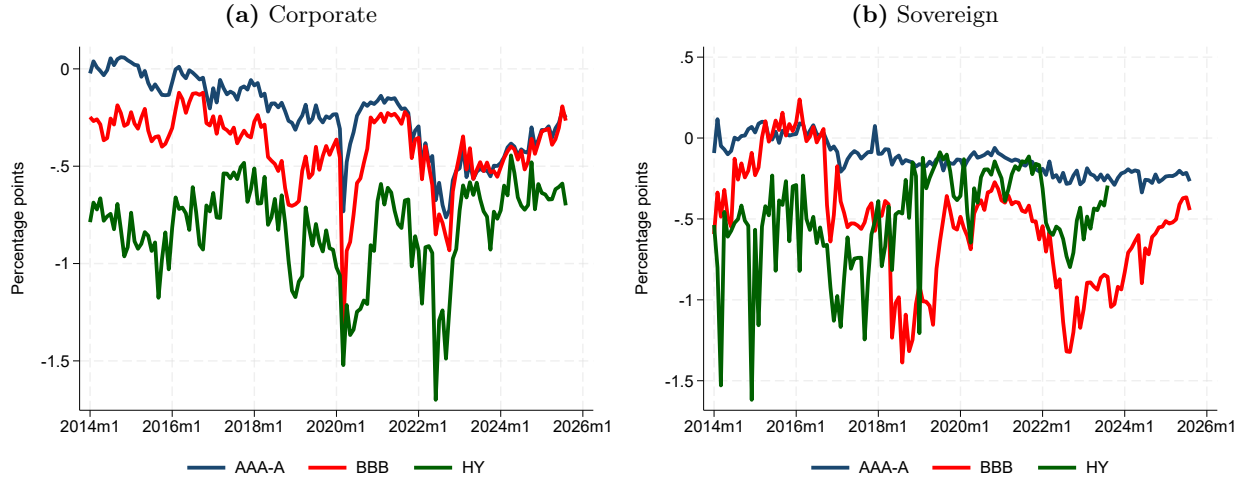


Figure 1 shows the average credit spreads and the CDS spreads separately for corporate and sovereign bonds for different rating categories. This is a period of significant fluctuations in credit spreads across all ratings, with a significant spike in March 2020 that was caused by the COVID-19 outbreak. As expected, the credit spread and the CDS spread are higher for lower rated categories with a significant difference between high-yield (HY) and investment-grade (IG) bonds.

Figure 2 shows the average CDS-bond basis across different rating categories. Throughout most of our sample period, the average CDS-bond basis remains significantly negative for all rating categories, indicating that investors are willing to forgo yields to hold AAA sovereign bonds rather than other bonds, even after adjusting for credit risk through CDS hedging. Moreover, the CDS-bond basis exhibits a monotonic relationship across ratings, suggesting a consistent preference for higher-rated bonds, even after controlling for credit risk.

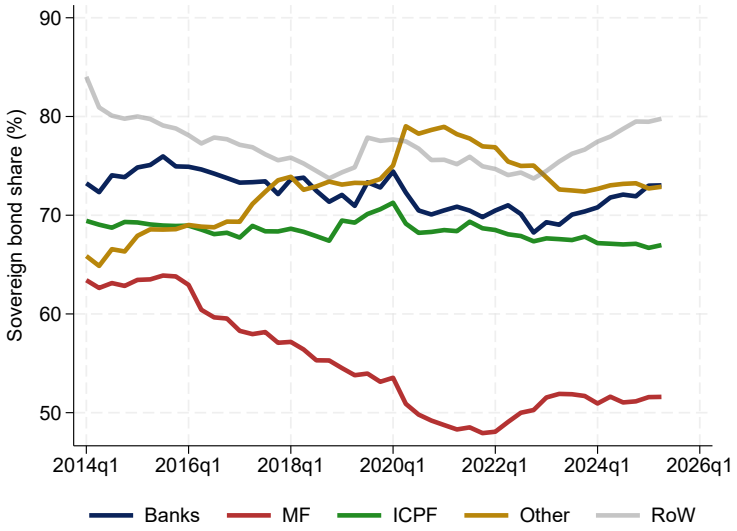
In the absence of frictions, the CDS-bond basis should be zero: a long-short portfolio that buys a AAA sovereign bond and sells short a duration-matched corporate bond whose credit risk is fully hedged through a CDS contract generates identical cash flows when held to maturity, so both positions should command the same price. In practice, however, the basis

Figure 2: Average CDS-bond basis for euro area bonds by rating. This plot presents the average CDS-bond basis (weighted by outstanding amounts) for each rating category, separately for corporate and sovereign bonds. *AAA – A* includes all bonds with a rating of at least *A*-; *BBB* includes bonds with ratings between *BBB+* and *BBB*-; and *HY* includes all rated bonds with a rating of *BB+* or below. For this plot, we excluded bonds of three Italian banks (Banco Popolare, Banca Monte dei Paschi di Siena, and UBI Banca) that exhibited excessively volatile CDS spreads in 2012–2013, possibly related to the ECB's OMT announcement. The time series for high-yield sovereign debt stop in September 2023 when the rating agency DBRS upgraded Greek sovereign debt to the investment-grade category.



is persistently negative—a well-documented regularity in the corporate bond literature (e.g., Longstaff et al., 2005; Oehmke and Zawadowski, 2017; Bai and Collin-Dufresne, 2019)—and we interpret its persistent cross-sectional variation as reflecting the premium investors attach to services associated with holding high-quality bonds, such as liquidity, collateral value, and regulatory relief, all consistent with the model in Section 2. However, the basis may also reflect frictions in the CDS market, to the extent that CDS contracts do not perfectly hedge default risk. As Mota (2023) discusses in detail, CDS market frictions are unlikely to drive cross-sectional variation in the basis: CDS are derivatives with highly elastic supply and limited counterparty risk due to overcollateralization, so these frictions are plausibly small relative to those in the cash-bond market and, as long as they are not strongly correlated with the service flows we aim to estimate, they add noise without meaningfully affecting our results.

Figure 3: Sovereign bond share by investor sector. This plot shows how the proportion of sovereign bonds within each investor sector's bond portfolio evolves over time, expressed as the share of total bond holdings (face value). *MF* stands for mutual funds, *ICPF* for insurance corporations and pension funds, *Other* includes non-financial corporations, households, and government agencies, and *RoW* (rest of the world) is calculated as a residual to match total outstanding amounts. The sample period is 2011Q4–2025Q2.



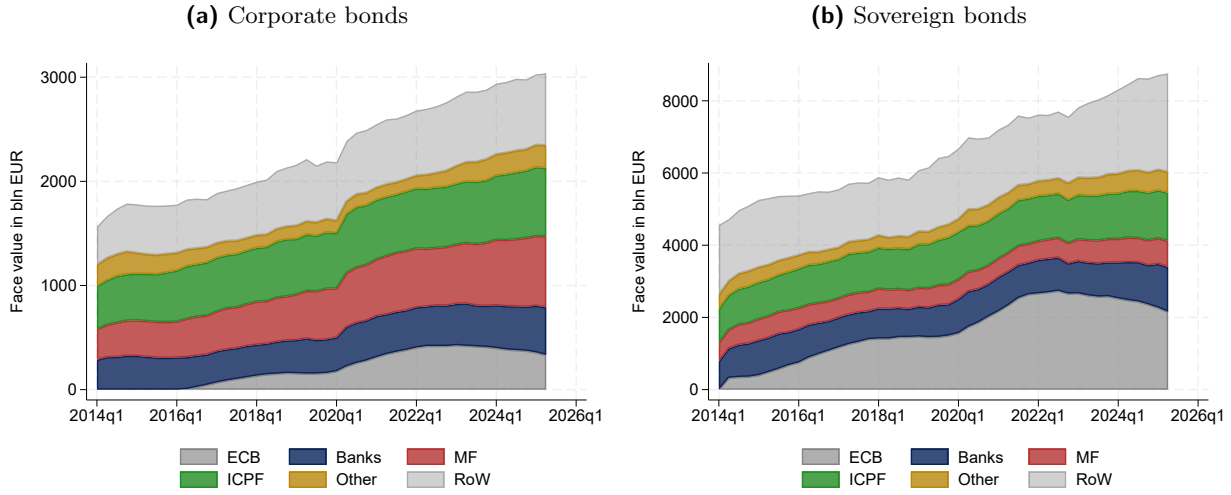
3.2 Holdings Data

We also study how investors differ in the value they assign to the services embedded in bonds. To analyze this investor heterogeneity, we use data from the Securities Holdings Statistics by Sector (SHSS), which are collected by the Eurosystem. The SHSS dataset provides information at the security level on the amount held by euro area investors. The data is aggregated at the selected investor type level. For our analysis, we group the investor types into the following categories: Banks (defined as deposit-taking corporations), insurance companies and pension funds (ICPF), mutual funds (MF), central banks, and others (which include non-financial corporations, households, general government, and other financial institutions such as hedge funds).

We infer the holdings by the rest of the world (RoW) as the residual—i.e., the difference between the total amount outstanding and the amount held by the sectors mentioned above. The frequency of these data is quarterly. As a final step, we match the price dataset at the ISIN-quarter level with the data on investor holdings.

The holdings data shows large heterogeneity in bond portfolio holdings across investors. For example, as illustrated in Figure 3, sovereign bonds account for 75–85% of the foreign

Figure 4: Sectoral holdings of euro area corporate and sovereign bonds. This plot presents the total outstanding face value (in bln EUR) of euro area corporate (left-hand side) and sovereign (right-hand side) bonds over time, split by investor sectors. *MF* stands for mutual funds, *ICPF* for insurance corporations and pension funds, *Other* includes non-financial corporations, households, and government agencies, and *RoW* (rest of the world) is calculated as a residual to match total outstanding amounts. The sample period is 2011Q2–2025Q2.



sector's bond portfolio in euro area bonds, while the corresponding share for euro area mutual funds is only between 50–65%.

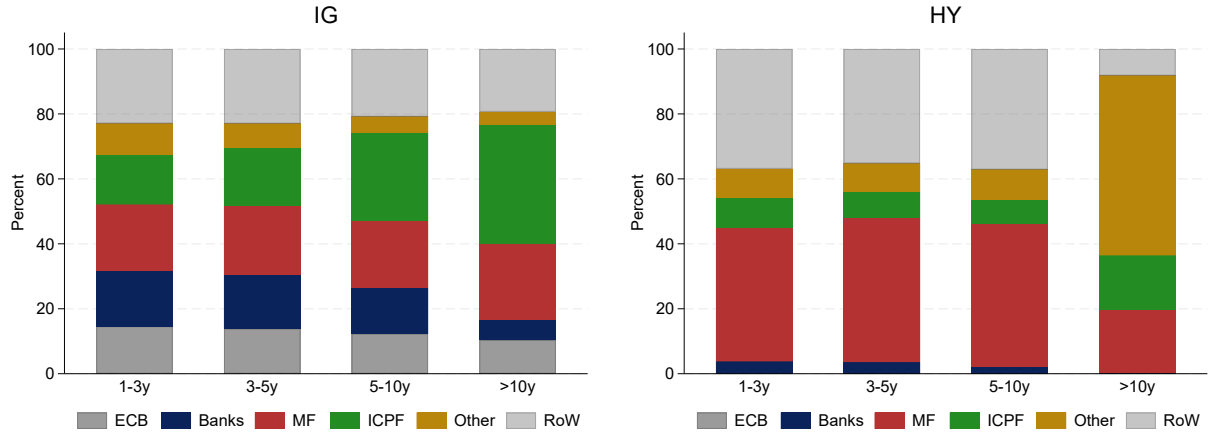
Figure 4 presents the investor composition for euro area corporate and sovereign bonds during the period 2011Q2–2025Q2. In aggregate, the largest investors in corporate bonds are insurance companies and pension funds (ICPF) and mutual funds (MF), which play a similarly important role. For sovereign bonds, the largest investors are the foreign sector and the ECB, whose direct holdings have grown significantly since the initiation of its asset purchase programs. As of June 2024, the ECB held approximately 13.1% of the total outstanding euro area corporate bond market, banks 13.9%, mutual funds 21.1%, ICPF 22.0%, other sectors 7.0%, and the rest of the world 22.8%.

The investor composition differs significantly in the cross-section of bonds, as also documented by Darmouni and Papoutsi (2023).[11] To illustrate the heterogeneity, Figure 5 presents the investor composition of EA corporate and sovereign bonds by rating and maturity at the end of 2024. Insurance companies and pension funds hold a large share of investment grade corporate bonds, and their holding share increases with maturity, reach-

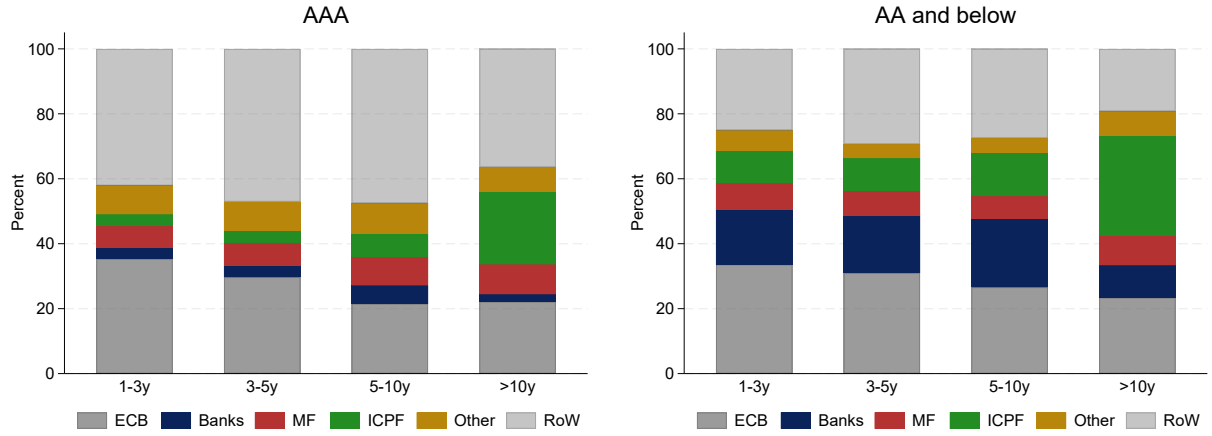
[11] Darmouni and Papoutsi (2023) focus on bonds issued by EA non-financial corporations, while we also include bonds issued by financial corporations.

Figure 5: Investor composition of euro area bonds by rating and maturity. This plot illustrates the sectoral holder composition of euro-denominated, EA-issued bonds in 2024Q4, for different residual maturity brackets (x-axis) and rating categories. The top panels distinguish between investment grade (IG) and high-yield (HY) corporate bonds, where IG bonds have a credit rating of at least *BBB-*. The bottom panel splits the sample into sovereign bonds with a *AAA* rating vs. all other sovereign bonds.

(a) Corporate



(b) Sovereign



ing almost 40% of the amount outstanding for IG bonds with remaining maturity above 10 years. Mutual funds, on the contrary, hold a significantly larger share of HY bonds across all maturity categories. The ECB holds only investment-grade bonds, and its holding share decreases with maturity.

In the sovereign bond segment, the rest of the world and the ECB are the two largest investors, with particularly high foreign shares among AAA-rated sovereign bonds. Banks and ICPF hold large amounts of sub-AAA sovereign bonds, with opposite maturity gradients, as expected. The documented segmentation in the euro area bond market demonstrates that

investors value specific bond characteristics differently, and for that reason, we expect this heterogeneity in preferences to have implications for bond prices.

3.3 Measuring Service flows

The paper aims to better understand the sources of convenience yields. To this end, we construct proxies for the service flows that drive these premia. We identify three measurable services: liquidity, collateral value, and regulatory value.

Liquidity: Liquidity captures the ability to liquidate bonds at little or no discount. We proxy for liquidity using $1 -$ bid-ask spread, where bid-ask spreads are measured as $\frac{\text{ask}-\text{bid}}{\text{bid}}$. We acknowledge that bid-ask spreads capture the *level* of transaction costs rather than liquidity *risk*—that is, the extent to which transaction costs may deteriorate precisely when investors need to trade the most, such as during periods of market stress. Liquidity risk is inherently difficult to measure, but to the extent that bonds with wider spreads also tend to experience larger price dislocations during crises, our proxy is informative about both dimensions.[12]

Collateral value: Bonds provide collateral services by allowing investors to raise secured funding. In particular, banks can borrow directly from the ECB by pledging eligible bonds. The central bank publishes daily lists of the ISINs of securities that it considers eligible collateral—the so-called *Eligible Assets Database* (EADB). Among other things, these lists report the haircut applied to the market value of a given bond. For example, if a bond carries a 5% haircut and a bank pledges 10 million euros of that bond, it can obtain only 9.5 million euros of cash in return.[13] We measure the collateral value of bond n at time t as $1 - \text{haircut}_t(n) \in [0, 1]$. Our collateral value measure is intended to capture pledgeability

[12] We also consider turnover as an alternative liquidity proxy. However, in the cross-section of the CDS-bond basis, we find no premium associated with high-turnover bonds, and the estimated turnover premium remains negligible over time. We therefore conclude that turnover is not a reliable measure of the liquidity services valued by investors in the euro area. Furthermore, we lack transaction-level data analogous to TRACE for US bonds, which would allow for more sophisticated liquidity measures.

[13] These haircuts depend on the bond's issuer sector, residual maturity, credit rating, and coupon structure. For more details, see https://www.ecb.europa.eu/mopo/coll/standards/html/index.en.html or Appendix F.

more broadly across all investors, under the assumption that ECB haircuts provide a good proxy for haircuts applied in private secured funding markets.

Regulatory value: Regulated financial intermediaries such as banks (under Basel III) and insurance companies (under Solvency II) must hold capital in proportion to their risk-weighted assets. These capital requirements typically apply along at least two dimensions: credit risk and interest rate risk. We construct a measure of regulatory value ($RegValue$) that captures both.

Risk-based capital (RBC): Capital requirements based on credit risk give rise to what is commonly known as risk-based capital charges. We measure these charges using the *spread risk* formula under Solvency II, the regulatory framework for insurance companies in the European Union. Solvency II follows a risk-based capital approach in which the overall Solvency Capital Requirement (SCR) is determined across several modules, including underwriting risk, market risk, and counterparty risk. Within the market risk module, the spread risk submodule captures potential losses from a widening of corporate bond spreads. The regulation specifies a stress factor, $stress_t(n)$, that assigns each bond n a capital charge increasing in duration and varying with credit rating. We define the RBC value of a bond as $1 - stress_t(n) \in [0, 1]$, so that bonds with lower capital charges have higher regulatory value. Importantly, sovereign bonds issued by euro area governments are fully exempt from spread risk charges under both Solvency II and euro area bank capital regulation.

Interest rate risk: The capital charge for interest rate risk depends on each institution's overall balance sheet exposure to rate changes, which is institution-specific and not directly observable. We therefore adopt an indirect approach: we interact modified duration with the RBC measure, reflecting that—conditional on credit risk—duration may carry an additional regulatory benefit through the interest rate risk module.

We define $RegValue$ as the sum of the credit risk and interest rate risk components. For further details on the construction of regulatory capital and collateral value measures, see Sections 6.1 and 6.2.

The summary statistics for the three service proxies are reported in Table 1. Liquidity, measured by $1-$bid-ask spread, is substantially lower for corporate bonds than for sovereign

bonds. The average bid-ask spread is 0.49% for corporates (with most observations ranging from 0.14% at the 10th percentile to 0.93% at the 90th percentile), compared to only 0.22% for sovereigns (0.02% to 0.60%).

Collateral value, measured as $1 - \text{haircut}$, also differs sharply across asset classes: corporate bonds exhibit considerable dispersion, with an average of 0.60 and values spanning from zero at the 10th percentile to 0.95 at the 90th, whereas sovereign bonds are consistently high-collateral assets, with an average of 0.91 and a much tighter distribution (0.90 to 0.99).

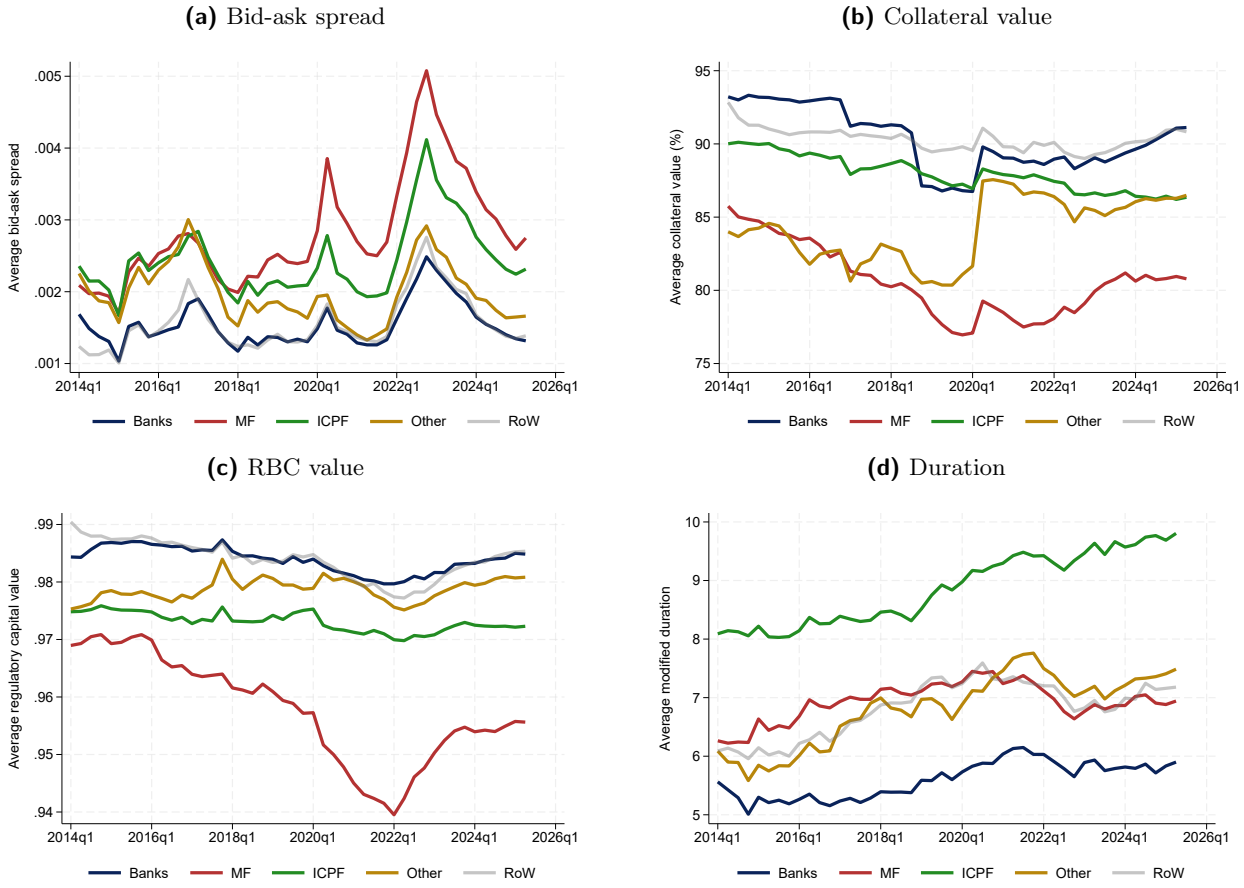
Regulatory capital value displays an even starker contrast. Corporate bonds have an average regulatory value of 0.93, with meaningful cross-sectional variation (0.87 at the 10th percentile and 0.98 at the 90th). Sovereign bonds, by contrast, are treated as uniformly safe from a regulatory perspective: their regulatory value is equal to one, reflecting their privileged treatment under capital requirements.

3.4 Service Flows Across Investors' Portfolios

With proxies for service flows in hand, we turn to the holdings data to analyze the service content of each investor sector's portfolio. The model in Section 2 predicts that heterogeneous investors hold different portfolios, reflecting differences in their valuations of service flows. To characterize these differences empirically, we compute, for each sector and quarter, value-weighted averages of the three service proxies discussed above, using sector-level portfolio shares measured by face value at the end of each quarter as weights. Figure 6 presents the results. The main takeaway is that service content differs substantially across investor sectors, consistent with heterogeneous valuations of these services.

Figure 6 also highlights several clear patterns. First, the ICPF portfolio has by far the longest duration, while banks lie at the opposite end of the spectrum. ICPFs tend to have long-duration liabilities and therefore hold long-duration bonds to reduce asset-liability duration mismatch and the associated regulatory capital for interest rate risk. The opposite is true for banks. Second, banks hold portfolios with the highest average collateral value for most of the sample period. High-collateral bonds are particularly valuable for banks because they can be pledged in secured funding markets. Third, mutual funds hold portfolios with relatively low regulatory capital value. Banks and insurance companies are subject to risk-

Figure 6: Average service content by investor sector over time. This plot shows the average service content of each investor sectors' bond portfolio over time for four services—liquidity value, collateral value, risk-based capital value (RBC), and modified duration. Panel (a) shows the average bid-ask spread, measured as $\frac{ask-bid}{bid}$; collateral value in panel (b) is measured as $1 - haircut \in [0, 1]$; RBC value in panel (c) is measured as $1 - stress \in [0, 1]$, where the *stress* parameter is defined in Appendix E. Panel (d) shows the average modified duration of each sector's bond portfolio. To calculate the sectoral averages, bond-level observations have been weighted by the face value held. The sample period is 2011Q4–2025Q2.



based capital regulation and therefore tilt their portfolios toward assets with high regulatory capital value. Mutual funds, in contrast, are not subject to the same regulatory framework, making them the natural holders of low regulatory value bonds. These bonds are mostly risky corporate bonds and tend to exhibit wider bid-ask spreads. Finally, the high average liquidity, collateral, and RBC content of the rest-of-the-world (RoW) portfolio is driven by its large exposure to sovereign bonds.[14]

[14] Figures C.2 and C.3 in the appendix plot the average service content by investor sector separately for sovereign and corporate bonds.

4 Decomposing Convenience Yields

Prediction 1 from the model posits that cross-sectional variation in bond convenience yields is driven by cross-sectional variation in bond-specific service flows. In this section, we use the service flow measures described in Section 3.3 to quantify the importance of each service in determining the convenience yields of AAA sovereign bonds.

We use the CDS-bond basis—defined as the difference between the yields of a risk-free AAA sovereign bond and the duration-matched risky bond whose default risk is fully hedged using a maturity-matched CDS—as a measure of convenience yields of the safe sovereign relative to that of the hedged bond. To see how our model-implied measure of convenience yield maps to the empirically measured CDS-bond basis, consider the following: The CDS-bond basis represents the yearly expected return of a portfolio that keeps a long position in the sovereign risk-free bond and a short position in the hedged risky bond, holding both positions to maturity. Based on equation (4) in the model section, under the assumptions that (i) the CDS fully hedges the bond's credit risk, (ii) AAA sovereign bonds are default-free, and (iii) holding CDS contracts provides no additional services, we can express the CDS-bond basis of bond n as follows:

$$b_t(n) = y_{T,t} - (y_{i,t} - CDS_{i,t}) = [\boldsymbol{Z}_t(n) - \boldsymbol{Z}_t(T)] \times \bar{\boldsymbol{\lambda}}_t \tag{10}$$

The CDS-bond basis thus directly maps to a bond's convenience yield relative to that of safe sovereigns. Once we make specific assumptions about the number and type of service flows in the matrix $\boldsymbol{Z}_t$, we can estimate the average service value $\bar{\lambda}_t(k)$ for each of the K services using cross-sectional regressions of the basis on the spread in services between bond n and the duration matched sovereign T.

We follow this insight in our empirical exercise. We assume that the matrix $\boldsymbol{Z}_t$ is formed by three service flows: liquidity, collateral, and regulatory capital value, measured as described in Section 3.3.

Because the observed CDS-bond basis is subject to measurement error and idiosyncratic shocks, we take Equation 10 to the data using the following cross-sectional regression for

each month t:

$$b_t(n) = \lambda_{lt} Liquidity_t^s(n) + \lambda_{ct} Collateral_t^s(n) + \lambda_{rt} RegValue_t^s(n) + \varepsilon_t(n) \tag{11}$$

where $b_t(n)$ is the CDS-bond basis of each bond n, the superscript s stands for the spread in service between each bond n and the duration-matched AAA sovereign.[15] We interpret the estimated $\boldsymbol{\lambda}_t$ as a measure of the value of the three services described above—specifically, the yield investors are willing to forgo to hold a portfolio that, all else constant, has one standard deviation higher service flow.

The estimated premium for each service flow enables us to decompose the convenience yield of any bond or bond portfolio for which we can measure the specific service flows. One portfolio of particular interest is the AAA sovereign bond portfolio, which serves as the safe asset in the euro area. We decompose its convenience yield as:

$$cy_t^{AAAsovereign} = \sum_k \hat{\lambda}_{kt} \times \bar{z}_{kt} \tag{12}$$

where $\bar{z}_{kt}$ represents the face-value weighted average of service k across all outstanding AAA sovereign bonds.

We now walk through the steps and identifying assumptions in estimating regression (11). For each bond in the sample, we construct service spreads by measuring the difference in service between bond n and its duration-matched safe sovereign T. We directly observe collateral values, risk-based capital requirements, and duration for both corporate bonds and sovereigns across all maturities. For liquidity, when an exact maturity match with a sovereign bond is unavailable, we use linear interpolation between the closest-duration sovereigns. We assume that hedging a bond's credit risk with a CDS does not affect its service flows, and therefore use the unhedged bond's service values directly.[16]

To ensure comparability across services and time periods, we standardize each service measure by subtracting its cross-sectional mean and dividing by its standard deviation before calculating spreads. This normalization allows our estimated coefficients λ_k to be interpreted

[15] Recall that we define $RegValue_t(n) = RBC_t(n) + Duration_t(n) \times RBC_t(n)$.

[16] We view this as a reasonable assumption. The hedged bond retains the same haircut in the collateral and regulatory framework as the underlying bond. Our liquidity measure assumes that CDS contracts are much more liquid than the cash bond such that their own liquidity does not affect the basis.

as the basis points earned by a zero-investment portfolio with one standard deviation of exposure to service k and zero exposure to all other services.[17]

The advantage of using the CDS-bond basis rather than (hedged or unhedged) bond yields is that it allows us to sidestep the difficulties of estimating a credit-risk model or term-premium model, making our decomposition model-free. The disadvantage is that we cannot capture the value of services along dimensions where the basis exhibits no variation. For instance, if high-duration bonds provide service flows to long-term investors relative to short-term bonds, we cannot identify this value because the duration spread in the basis is zero by construction for all duration-matched bond pairs. Consequently, our approach can only capture the value of services that vary in the cross-section of our hedged bonds. This limitation means we potentially identify a lower bound on total convenience yields, isolating only those service flows that generate cross-sectional dispersion beyond what is captured by duration matching and credit hedging.

The estimation of regression (11) excludes a constant term. In principle, including one would capture the component of AAA sovereign bond value *not* attributable to the measured services, allowing for a level misspecification of the model. In practice, however, this is not feasible: because risk-based capital requirements (RBC) take a constant value of one for all sovereign bonds, a constant would mechanically absorb part of the RBC coefficient, biasing the estimated value of RBC and of the total $RegValue$.[18]

Furthermore, the reader should view this exercise as a projection of the CDS-bond basis onto these three services. If there are additional services not included in our analysis that also drive convenience yields, our results could be affected in two ways. First, a large portion of the convenience yield would remain unexplained by our chosen services. Second, if omitted services are correlated with any of our included services, omitted variable bias would arise, with the correlation determining the sign of the bias. One can always extend our framework

[17] To see this, let Z_{st} be the service spread. Notice that the OLS coefficient implies that $\hat{\boldsymbol{\lambda}}_t = (Z_{st}^\top Z_{st})^{-1} Z_{st}^\top \boldsymbol{b}_t$. Define $W_t^\top = (Z_{st}^\top Z_{st})^{-1} Z_{st}^\top$, the $K \times N$ matrix of portfolio weights. The service flows of these portfolios are $W_t^\top Z_{st} = I$, where I is the identity matrix.

[18] In unreported results, we reestimate the service premium including a constant. The constant has little effect on the estimated collateral and liquidity premiums but, as expected, renders the RBC estimate noisy and difficult to interpret, given the high correlation between the two.

Table 2: Average value of services. This table shows the results of the Fama-Macbeth regression of the CDS-bond basis on the proxy of four service flows in consideration.

	CDS-bond basis (%)				
	(1)	(2)	(3)	(4)	(5)
Liquidity	0.202***				0.110***
	(0.008)				(0.005)
Collateral		0.245***			0.122***
		(0.009)			(0.004)
RBC			0.197***		0.161***
			(0.008)		(0.006)
Duration:RBC				−0.082***	0.075***
				(0.004)	(0.002)
Avg. R^2	0.273	0.290	0.269	0.132	0.403
Observations	201,898	201,898	201,898	201,898	201,898
Note:					*p<0.1; **p<0.05; ***p<0.01

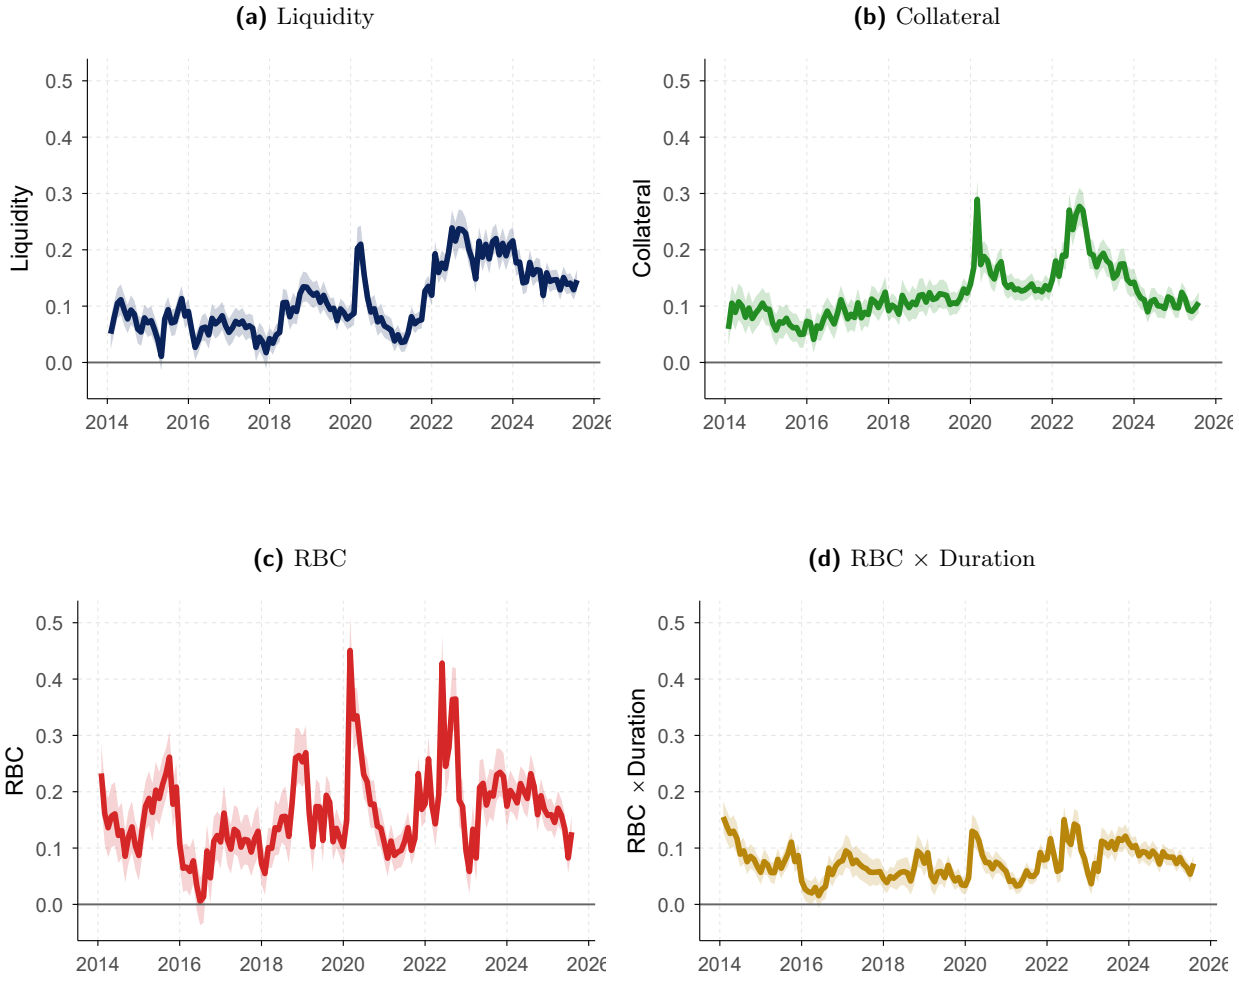
to incorporate additional services. We believe that this three-service specification strikes an optimal balance between comprehensiveness and tractability: it captures the economically important drivers of convenience yield for which we have reliable measures, and it delivers strong empirical performance.

Table 2 presents the estimated average service values for the full sample.[19] Column (5) demonstrates that premiums for all services are positive on average. Bonds with one standard deviation higher liquidity have, on average, an 11 bps higher (less negative) CDS-bond basis, or equivalently, lower credit spreads. The effect is 12 bps for collateral and 16 bps for RBC. Note that conditional on RBC, an additional one standard deviation in duration (around 5 years) increases the basis by another 7 bps. These effects are economically significant relative to the average CDS-bond basis of 34 bps. Furthermore, the spread in services explain on average 40% of the cross-section variation of the CDS-bond basis.

The premium for each service is not constant over time. In Figure 7, we report the time series of each λ_k estimate with 5% confidence intervals. Several interesting patterns emerge. First, all services exhibit a pronounced spike in March 2020, though the increase is particularly strong for the RBC component, consistent with investors' rapid need to satisfy

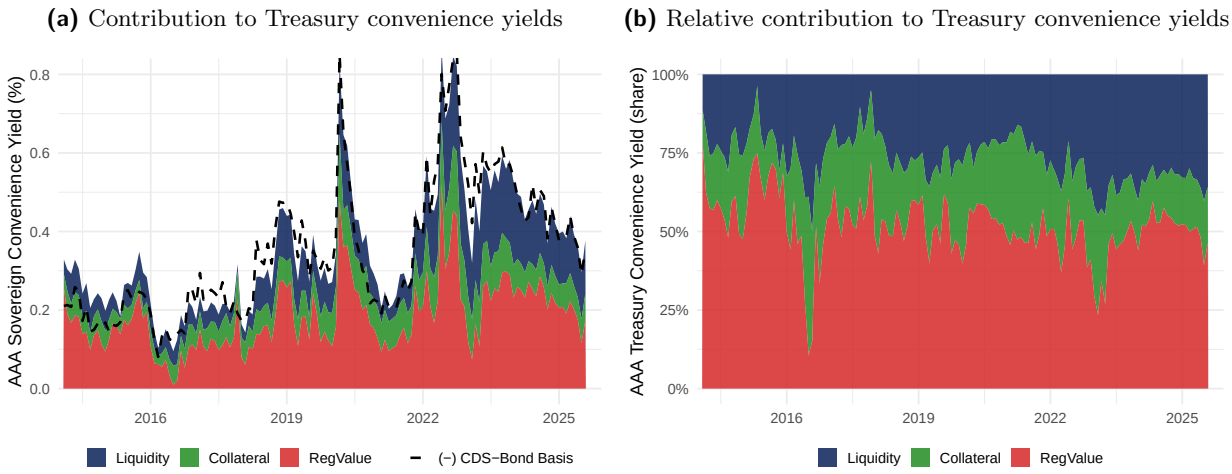
[19]The coefficients are the average coefficients from the cross-sectional regressions.

Figure 7: Service value estimates This plot shows the monthly service value estimated from equation 11 with 5% confidence interval.



regulatory requirements and seeking safety during the onset of the COVID-19 shock. Second, beginning in 2022—when the ECB initiated the interest-rate tightening cycle, soon followed by a decline in its balance sheet—we observe a clear rise in the value of both collateral and liquidity services relative to regulatory value. Notably, the collateral premium subsequently declined again once the ECB started normalizing policy rates in 2024, mirroring the reduced scarcity value of high-quality collateral in a lower-rate environment. Third, the combined RBC-duration service, which captures the additional value associated with longer-maturity securities, stands out as the most stable component over the full sample, showing comparatively limited fluctuations despite major shifts in monetary and market conditions.

Figure 8: Contributions of service flows to AAA sovereign bonds convenience yield. Panel (a) shows the contribution of each service, which is the value estimated in (11) times the average service, $Z(T)$, of all AAA sovereign bonds outstanding. Panel (b) shows the relative contribution of each service.



We now use Equation (12) to study the convenience yields of AAA sovereign bonds. Panel (a) of Figure 8 plots the premium in sovereign bonds for each service component, together with the negative of the CDS-bond basis.[20] Panel (b) shows each service's share relative to the total. For most of our sample period, the explained convenience yield is mainly driven by regulatory capital value. This service flow accounts for, on average, 50% of the convenience yields of AAA sovereigns in our sample, highlighting its dominant role across periods. While the relative importance of liquidity has increased over the last four years—reflecting tighter monetary policy, a declining supply of central bank reserves as the ECB balance sheet started shrinking and greater demand for liquid assets—its contribution remains below that of regulatory capital value. Notably, even during periods of elevated liquidity demand, the regulatory component continues to be the primary driver of the convenience yield.

How do we validate that our services do a good job in explaining convenience yields? Beyond the high R^2 values of our cross-section regressions, we can examine how well the chosen services explain the level of the CDS-bond basis. The regression does not match the average CDS-bond basis by construction. The fact that our estimated services nonetheless

[20] Recall that the CDS-bond basis is negative most of the time, as investors are willing to forgo returns to hold AAA sovereigns rather than synthetic risk-free bonds.

explain over 90% of the negative basis therefore validates the modeling approach rather than reflecting a mechanical relationship.

5 Investor-Specific Valuation of Service Flows and Expected Returns

The equilibrium service premia estimated in the previous section are silent about which investors value which services. Yet understanding whose demand drives these premia is important: if premia arise through regulatory or financial constraints, identifying who values what can inform policy design and clarify transmission mechanisms.

At the same time, the nature of the services we measure implies substantial heterogeneity in how different investors value each one. For instance, both a mutual fund and an insurance company may be willing to pay a premium to hold a AAA sovereign over a corporate bond, beyond what is explained by credit risk alone—but for different reasons. The mutual fund may do so for the extra liquidity, while the insurance company values the regulatory relief. Since demand and supply must equilibrate, the premia we estimate above reflect the aggregation of these diverse valuations. In this section, we use detailed investor-level data to shed light on whose excess demand drives convenience yields. In the next section, we use event studies to disentangle which service is valued most by each type of investor.

Prediction 2 posits that investors differ in their valuation of bond service flows and therefore in their expected returns. Investors with higher demand for services are willing to accept lower expected returns relative to the market benchmark, as the convenience yield of their portfolios increases. Conversely, sectors with lower overall demand for service flows should earn higher returns. By revealed preferences, then, if a sector systematically underperforms another—even after we control for traditional sources of risk in fixed income, such as duration and credit risk—we posit that this investor has high demand for service flows, which not only pushes up the premium for the services they value, but also lowers the overall performance of their portfolio.

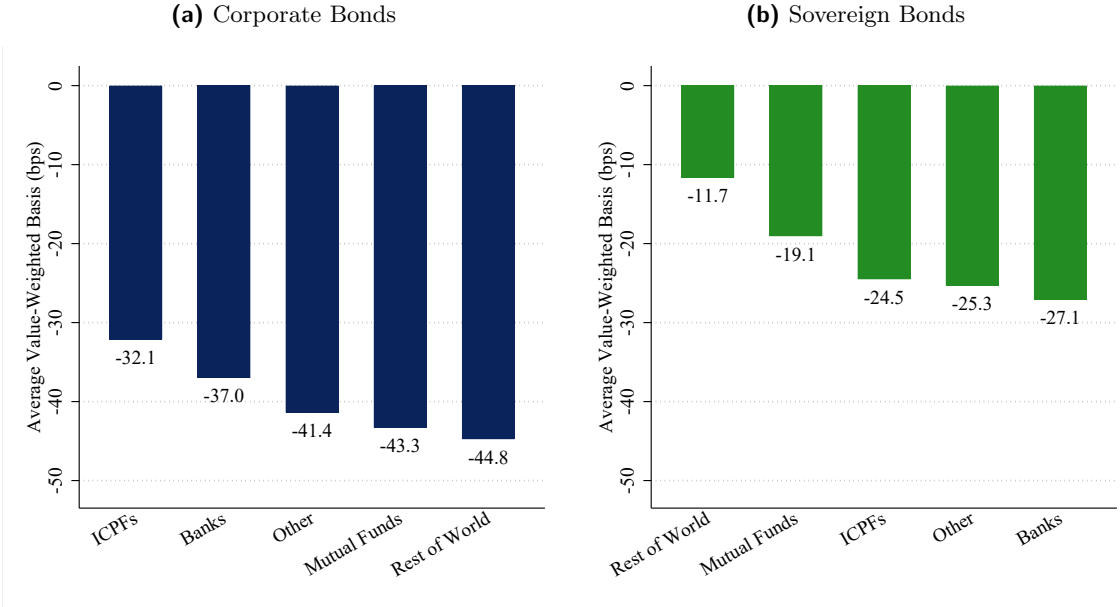
We proceed in two steps. First, we link convenience yields to holdings data and compute the average yield borne by each investor sector, quantifying the income that each is willing

to forgo to hold bonds with valuable service flows. Second, we compute sector-level portfolio returns, showing how differences in service flow demand translate into heterogeneous realized performance.

Our performance analysis relies on returns sourced from iBoxx, so our sample is the intersection of iBoxx and the holdings data. In the appendix, we present corresponding results using the broader universe of bonds in the CSDB database; the main findings are consistent across the two samples.

Portfolio-level CDS-bond basis and convenience yield. In Figure 9, we report the average value-weighted CDS-bond basis for each investor sector. Panel (a) shows results for corporate bonds, while Panel (b) shows results for sovereign bonds. We compute weights using each sector's portfolio shares, based on face value, within the sovereign and corporate bond universes, respectively.

Figure 9: Average CDS-Bond Basis by Sector (iBoxx). The figure reports the average value-weighted CDS–bond basis for each investor sector. Panel (a) shows results for corporate bonds, while Panel (b) shows results for sovereign bonds. We compute weights using each sector's portfolio shares, based on face value, within the sovereign and corporate bond universes, respectively.



Focusing on corporate bonds (Panel (a)), we observe substantial heterogeneity across investor sectors. Mutual funds hold bonds with a wider CDS-bond basis, averaging −43.3

basis points. Similarly, other investors—a category that includes hedge funds—also hold bonds with relatively wide CDS-bond bases, averaging −41.4 basis points. In contrast, banks and insurance companies and pension funds (ICPFs) hold corporate bonds with higher convenience yields, reflected in a less negative basis. The average basis is −37.0 basis points for banks and −32.1 basis points for ICPFs.

This pattern is consistent with the importance of regulatory and collateral service flows. Lower-rated corporate bonds provide low regulatory value for insurance companies, since they require substantial capital charges under Solvency II. As a result, insurers tilt their portfolios toward higher-rated corporate bonds, which carry higher regulatory value and higher convenience yields. Overall, these findings suggest that banks, insurers, and pension funds are willing to forego expected returns in order to hold bonds that deliver regulatory capital benefits. In equilibrium, market clearing implies that other investors—such as mutual funds—must absorb corporate bonds with lower regulatory value and are compensated through higher expected returns, reflected in a wider CDS-bond basis. In Figure D.2 Panel (a) in the Appendix, we repeat the exercise using the broader universe of bonds from the CSDB database. The results are consistent: banks and ICPFs hold portfolios with higher convenience yields, whereas mutual funds hold bonds with lower convenience yields. The only exception concerns the rest of the world sector: when considering the full CSDB universe, we find that RoW holds corporate bond portfolios with higher convenience yields.

Interestingly, when we turn to sovereign bonds (Panel (b)), the ranking across sectors is essentially reversed. Banks and ICPFs hold sovereign bonds with the lowest convenience yields, with an average basis of −27.1 and −24.5, respectively. In contrast, mutual funds hold sovereign bonds with a substantially narrower basis, averaging −19.1. In this case, mutual funds and the rest of the world appear willing to forgo expected returns in order to hold sovereign bonds with higher convenience yields. As previously discussed, for insurance companies and banks, the regulatory value of all euro area sovereign bonds is broadly the same. In fact, banking and insurance regulations do not assign a capital charge related to the spread risk on euro area government bonds. Moreover, investors like insurers likely do not particularly value liquidity, as they tend to be long-term investors. As a consequence, they are the natural holders of illiquid sovereign bonds, such as those issued by "vulnerable"

countries.

A natural interpretation is that mutual funds benefit from the high liquidity of safer sovereign bonds. They may tilt their portfolios toward the most liquid sovereign securities, thereby offsetting the lower liquidity of the corporate bonds they hold. Ma et al. (2022) shows that, when meeting redemptions, mutual funds follow a pecking order by first selling their most liquid assets. In this sense, as mutual funds absorb relatively illiquid corporate bonds—given that banks and insurance companies refrain from holding these securities—they may build a liquidity buffer through holdings of liquid sovereign bonds. Such a buffer reduces liquidity risk and facilitates redemption management.

The rest-of-the-world sector, in contrast, comprises a diverse set of investors, ranging from foreign hedge funds to insurance companies and central banks. The motives for holding euro area bonds are therefore highly heterogeneous. Importantly, unlike domestic banks and insurance companies, foreign banks and insurers may face capital charges on euro area government bonds that depend on the bonds' credit ratings. This regulatory asymmetry implies that foreign investors tilt their sovereign bond portfolios toward safer and more liquid euro area sovereign bonds, which exhibit higher convenience yields.

Beyond regulation, foreign investors may value euro area sovereign bonds for the services they provide in international financial markets. In particular, investors seeking exposure to the euro, funding in euros, or access to euro-denominated repo markets have incentives to hold bonds with high collateral value and secondary-market liquidity. Highly rated sovereign bonds are especially attractive along these dimensions, as they are more easily pledged in secured borrowing and serve as benchmark safe assets. As a result, foreign investors concentrate their sovereign holdings in securities with higher convenience yields and therefore a narrower CDS-bond basis.

That being said, because the rest-of-the-world sector pools heterogeneous investor groups, we refrain from drawing strong conclusions about the precise mechanisms underlying this pattern.

Finally, comparing the results for corporate and sovereign bonds in Panels (a) and (b) of Figure 9, the level difference in the average CDS-bond basis between the two asset classes stands out clearly. The figure shows that the corporate bond basis is roughly twice as

large as the sovereign bond basis. This difference remains significant even when controlling for credit ratings, suggesting that investors are willing to forego higher returns in order to hold sovereign bonds. A natural explanation is that sovereign bonds receive preferential regulatory treatment, which enhances their value as a source of convenience services.

Hedged bond returns and convenience yield. We now compute sector-level portfolio returns attributable to differences in convenience yields, thereby measuring how much return investors forgo in order to hold bonds that provide valuable service flows. Our objective is to isolate the component of bond returns that is not compensation for standard sources of risk, such as interest rate and default risk. This addresses two concerns that arise from examining the portfolio basis alone: First, the basis measures a hedged yield, but this translates into a realized return only if the bond is held to maturity. In practice, sectors differ in their trading behavior and timing, so realized hedged returns capture whether investors actually earn the spread implied by the basis. Second, the basis has important dynamic properties: bonds with a more negative basis tend to earn higher returns on average, but not in periods when the basis widens. Looking at realized returns incorporates these dynamics, rather than treating the basis as a static measure of the convenience yield cost borne by each sector.

For each bond n, we form a portfolio that is long the physical bond plus a maturity-matched CDS contract written on the same reference entity. The CDS removes default risk, leaving a return that reflects the bond's non-default component. We then remove interest rate exposure by subtracting the return on a government bond matched to bond n's modified duration, constructed from the ECB zero-coupon yield curve. The resulting basis-trade return is given by

$$r_t^{\text{ls}}(n) = r_t^{\text{hedged}}(n) - r_t^{\text{gov}}. \tag{13}$$

This long-short return captures the residual compensation associated with holding the physical bond relative to its synthetic equivalent. Appendix B provides full implementation details, including the treatment of standardized CDS contract pricing and upfront fees. Under a first-order approximation, it is closely linked to the CDS-bond basis:

$$r_t^{\text{ls}}(n) \approx -b_{t-1}(n) + D_t(n)\,\Delta b_t(n), \tag{14}$$

where $b_t(n)$ denotes the CDS-bond basis and $D_t(n)$ is modified duration. The first term represents the carry from holding a bond with a negative basis, while the second term captures mark-to-market variation as the basis changes over time.

After constructing bond-level hedged returns, we compute sector-level portfolio returns by combining daily returns with portfolio weights measured at the beginning of each quarter. Specifically, for each investor sector we form value-weighted portfolios using holdings shares observed at the quarterly rebalancing date and apply these weights to daily returns over the subsequent quarter. We then compute the cumulative performance of each sector portfolio by tracking the value of one dollar invested at the start of 2014 and held through the end of the sample period. In this exercise, we focus on the three major euro area investor sectors: mutual funds, banks, and insurance companies and pension funds (ICPFs).[21]

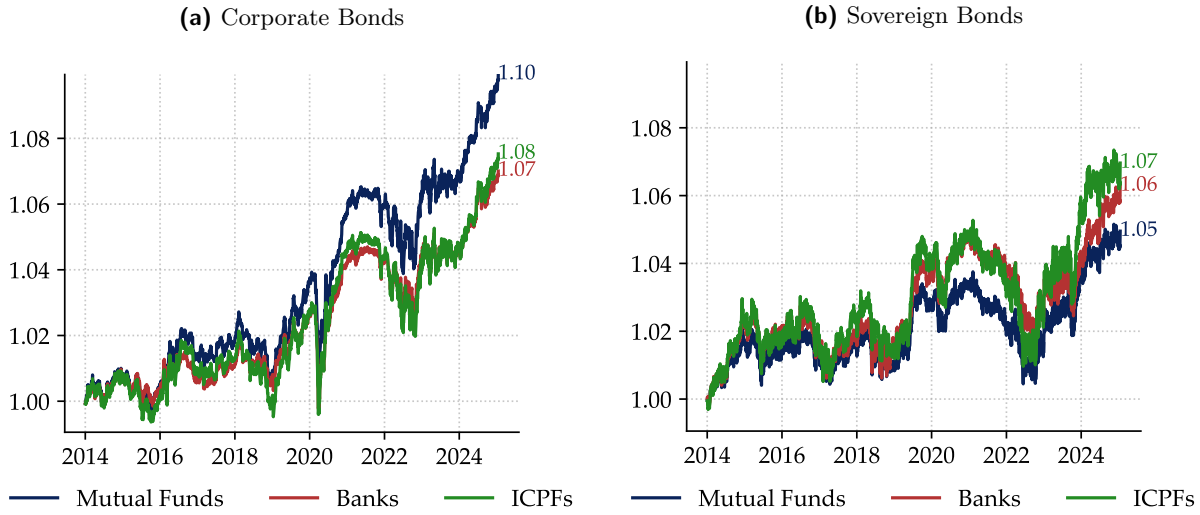
The results are reported in Figure 10. Panel (a) plots the cumulative hedged returns for the different sectors. Mutual funds earn a cumulative return of approximately 10% over our sample period. In contrast, banks and ICPFs earn cumulative returns of about 7.5%, implying a sizable return gap of roughly 3.5 percentage points. Given that these sectors collectively hold corporate bond portfolios on the order of 1.5 trillion euros, this gap corresponds to approximately 50 billion euros in cumulative foregone returns over the sample period.

Consistent with our service-flow interpretation, banks and ICPFs tilt their corporate bond portfolios toward securities that provide higher regulatory capital value and collateral services. Overall, these findings suggest that intermediaries are willing to incur substantial economic costs in order to hold assets that deliver valuable service flows, consistent with the central role of convenience yields in shaping equilibrium expected returns.

This estimate is conservative. Our return series is constructed using senior unsecured bonds that are relatively liquid and for which reliable CDS matches are available. The return differential is therefore likely even larger for lower-rated and less liquid corporate bonds—which are underrepresented in our sample—and for which the divergence in holdings between

[21] We exclude the rest of the world (RoW) because, as discussed earlier, we cannot clearly identify the type of investor in this category. Moreover, because RoW holdings are constructed as a residual, the results vary depending on whether we use the iBoxx or the CSDB sample of bonds. For the other investor sectors, the choice of bond sample does not materially affect the results.

Figure 10: Cumulative Returns by Sector. The figure reports the cumulative returns for each investor sector. Panel (a) shows results for corporate bonds, while Panel (b) shows results for sovereign bonds. We compute weights using each sector's portfolio shares, based on face value, within the sovereign and corporate bond universes, respectively.



mutual funds and the other two sectors is more pronounced (as discussed in Section 3).

Panel (b) shows that the pattern is reversed in sovereign bond markets. Banks and ICPFs earn cumulative returns of around 6.5%, compared to 5.0% for mutual funds, implying a return advantage of roughly 1.5 percentage points. For portfolios of comparable size, this corresponds to an economically significant difference of approximately 22 billion euros over the sample period.

Finally, the hedged corporate bond portfolio exhibits substantially higher cumulative returns than the hedged sovereign bond portfolio. Over the sample period, the return gap ranges from about 1 percentage point for banks, insurance companies, and pension funds to roughly 5 percentage points for mutual funds. This pattern implies that some investors are willing to forego sizeable returns in order to hold sovereign bonds.

Consistent with this interpretation, banks and ICPFs hold roughly 70% of their portfolios in sovereign bonds, as documented in Figure 3. This large exposure indicates that these investors are willing to sacrifice income to obtain the service flows provided by sovereign bonds.

Mutual funds, in contrast, which are less constrained by regulatory requirements, allocate

a larger share of their portfolios to corporate bonds (roughly 50%) and earn a sizeable additional return by doing so. Importantly, this extra return is not compensation for credit risk or interest rate risk, but instead reflects the lower convenience yields embedded in corporate bonds.

Overall, the results in this section confirm Prediction 2: investors value asset services differently, tilt their portfolios toward assets that provide the services they value most, and are willing to forgo financial returns to do so. A stark difference emerges between corporate and sovereign bonds. In the corporate bond universe, mutual funds earn significantly higher hedged returns than banks and insurance companies, who pay a premium for regulatory capital value. In the sovereign bond universe, the pattern reverses: banks and insurance companies earn higher returns by absorbing less liquid sovereigns, while mutual funds forgo returns to hold the most liquid sovereign bonds as a buffer against redemption risk.

While these results provide suggestive evidence on how different investors value different services, they do not establish causality. In the next section, we sharpen the identification by exploiting policy-driven shocks to each of the three services.

6 Policy-Driven Changes to Service Flows

So far, we have analyzed and estimated the drivers of convenience yields in the cross-section of European bonds. To do so, we had to make explicit assumptions about the services that investors value. However, without exogenous variation in services or preferences, we cannot gauge the causal impact of services on prices or portfolio allocation. In the remainder of this paper, we exploit policy shocks that provide exogenous variation in services or preferences to validate our empirical approach. We study three types of policy shocks, each targeting a different service. Our objectives are threefold: first, to provide causal evidence that the services we measure impact convenience yields and portfolio allocation; second, to demonstrate that investors' private valuations for services are heterogeneous across sectors; third, to show that policy decisions have substantial effects on both convenience yields and portfolio allocation. In this way, the empirical results that follow allow us to test the predictions 3 and 4 of the model.

6.1 Regulatory Capital Requirements in the Insurance Sector

First, we turn our attention to the insurance sector, which underwent significant regulatory changes in January 2016, when the Solvency II framework became fully applicable in the EU.[22] Solvency II harmonized the regulation of insurance companies across Europe, which had previously been subject to national rules. In particular, regulators specified common standards for the calculation of insurers' "Solvency Capital Requirement" (SCR)—capital buffers designed to withstand losses emanating from different risk sources. As a result, bonds differ in how much equity capital they implicitly require on the holder's balance sheet. To the extent that equity issuance is costly, among two otherwise identical bonds, insurers prefer the one that requires less capital. Our model then suggests that insurers may be willing to forgo yields to hold such assets, resulting in portfolios with high convenience yields due to capital constraints.

To test our model-implied Prediction 3, we interpret the implementation of Solvency II as a positive shock to insurance companies' appetite for (a) bonds with low risk-based capital charges, and (b) bonds with durations around 20 years (Jansen, 2025). We explain our reasoning in the following paragraphs.

The basic SCR, as defined in Art. 87 of Delegated Regulation (EU) 2015/35, is a function of the capital charges arising from each of the following "risk modules" and their correlations: (a) *Underwriting risk*, which refers to the risks that are inherently associated with the insurance business, e.g., the risk of claim settlements, natural disasters, or longevity/mortality risk in the case of life insurers; (b) *Market risk* that arises from fluctuations in the market value of insurers' assets and liabilities; and (c) *Counterparty risk*. Since underwriting risk depends purely on insurers' liabilities, for the purpose of our analysis, we focus on (b), and more specifically on the so-called *spread risk* and *interest rate risk* sub-modules.

First, *spread risk* is defined in Art. 176 of the above regulation, which specifies a bond-specific *stress* function that takes values between 0 and 1, decreases with the bond's credit

[22] Key legal acts included the "Solvency II Directive" 2009/138/EC and the Delegated Regulation (EU) 2015/35. As any Directive, Solvency II had to be transformed into national legislation in each Member State separately, and the implementation timelines varied by country. By January 1, 2016, all aspects of the framework—including technical details specified in the 2015 Delegated Regulation—had to be fully implemented.

quality and increases with modified duration.[23] The *stress* value reflects the hypothetical adverse-scenario loss in the bond's value that the insurer's capital needs to withstand. Hence, bonds with a higher *stress* value command more capital per unit of market value held. Importantly, any sovereign bond issued by Member States is assigned a *stress* value of zero, i.e., they carry no spread risk-related capital charges. As described in Section 3.3 above, we use $RBC_t(n) = 1 - stress_t(n)$ as a measure of the risk-based regulatory capital value of bond n at time t.[24]

Second, to capture changes in the *interest rate risk* sub-module, we closely follow Jansen (2025), who argues that Solvency II incentivized European insurers—particularly in countries that had previously relied on historical cost accounting—to hold bonds with durations around 20 years. In these countries, insurers' liabilities had historically been largely insensitive to long-term interest rate fluctuations. Solvency II, however, introduced fair value accounting and a market-consistent discounting framework in which the risk-free curve is constructed from observed swap rates only up to the *last liquid point* (LLP), located around the 20-year maturity in the euro area, and extrapolated toward the "ultimate forward rate" (UFR) beyond that horizon.

Insurance companies and pension funds typically exhibit a *negative duration gap* between their assets and liabilities (Domanski et al., 2015; Koijen and Yogo, 2022b). When long-term interest rates decline, the market value of liabilities therefore increases by more than that of assets, pushing insurers closer to their solvency constraints.[25] Under the market-consistent discounting framework, maintaining such a negative duration gap requires substantial eq-

[23] We report the exact formula in Appendix E.

[24] It is worth noting that according to Art. 83(4) of the Delegated Regulation, insurers can engage in *risk mitigation* techniques that lower their implied capital requirements if they comply with certain conditions (specified in Arts. 209–215 of the same Regulation). For example, if an insurer hedges a bond's credit risk by holding the corresponding CDS, the hypothetical adverse-scenario net loss (and thus the required capital under the spread risk module) is reduced. However, holding CDS will give rise to additional capital requirements through the *counterparty risk* module in most cases. As a result, to maintain low overall capital requirements, bonds with high *RBC* value are generally preferable. In practice, hedging bond credit risk via CDS is rarely relevant for insurers anyway. Their bond portfolios are typically investment-grade and well diversified, whereas CDS are issuer-specific and would require complex matching procedures.

[25] Although pension funds in the EU are not subject to a harmonized regulatory framework like Solvency II, their negative duration gap (and hence the risk of underfunding when interest rates are low) suggests that they face qualitatively similar constraints.

uity capital. As a result, insurers were incentivized to increase the duration of their asset portfolios in order to hedge interest rate risk and limit implied capital requirements.

Moreover, because liability valuations become particularly sensitive to interest rate movements around the LLP—the point at which the discount curve transitions from market swap rates to UFR-based extrapolation—insurers in euro area countries were incentivized to rebalance their sovereign bond portfolios toward 20-year bonds, and away from both longer and shorter maturities.[26]

Leveraging the framework established in this paper, we examine how regulatory changes differentially affect the convenience yields (or the CDS-bond basis) of bonds with different regulatory values, using their two dimensions of spread and interest rate risk. This goes beyond Jansen (2025), who focus on the effects of regulation across bonds with different durations and the impact on prices.

We start by examining whether the EU-wide application of the Solvency II framework in January 2016 influenced the convenience yield of bonds with high RBC value differently than those with lower regulatory value, especially in the 20-year duration segment. Specifically, we estimate the following model:[27]

$$b_t(n) = \beta \times HighRBC_t(n) \times 20y_t(n) \times Post_t + \alpha_{r\times m} + \alpha_t + u_t(n) \tag{15}$$

In this equation, $b_t(n)$ represents the CDS-bond basis of bond n and $HighRBC_t(n)$ is a dummy indicating whether bond n's risk-based capital relief exceeds the cross-sectional median at time t. $20y_t(n)$ is another dummy that is 1 if bond n has a modified duration between 15–25 years at time t. Finally, the dummy variable $Post_t$ is equal to 1 for all months after (including) January 2016, and zero otherwise.[28] We also include month and rating-by-time-to-maturity fixed effects,[29] and we estimate the triple-difference model in the period from

[26] We verified that these results from Jansen (2025) hold in our own sample. The results of this replication exercise are available upon request.

[27] The regression we estimate is fully saturated, but we omit some interaction effects and standalone coefficients here for brevity.

[28] As acknowledged above, the EU-wide application of Solvency II cannot be considered an exogenous shock, and anticipation effects are possible. However, in unreported results, we document that the parallel trends assumption holds for the chosen "treatment" date.

[29] Specifically, we construct twelve rating-maturity buckets where each bucket is a combination of one of three rating categories (AAA–A, BBB, or high-yield) with one of four residual maturity categories (1-3y, 3-5y, 5-10y, 10-30y). The

January 2015 – December 2016.

Prediction 3 in our model states that when demand for a specific service increases, bonds that provide more of that service will experience an increase in convenience yields. This maps directly to the hypothesis that β will be positive: following the introduction of Solvency II, bonds with high regulatory capital value should exhibit increased convenience yields relative to low-RBC bonds, conditional on having a duration of around 20 years. Conditioning on duration is important because the formula to calculate RBC assigns different weights to bonds based on their duration. Therefore, omitting the interaction with duration would imply that RBC largely proxies inversely for duration. An equivalent interpretation of our finding is that, conditional on duration, the CDS-bond basis of safer bonds (including sovereigns) increases after Solvency II implementation relative to that of riskier bonds.

The results reported in Table 3 are consistent with our predictions. Specifically, column (1) shows that after the Solvency II reform and conditional on having a duration around 20 years, the CDS-bond basis is approximately 20 basis points higher for bonds with above-median RCB value compared to those below the median. This result is highly statistically significant and persists qualitatively even if we replace the $20y$ dummy by modified duration as a continuous variable (see column (2)).

Our model further implies that if investors experience a positive change in their taste for specific services, they rebalance their portfolios towards assets providing these services. Therefore, we next analyze portfolio rebalancing behavior. Specifically, for each investor sector i we estimate[30]

$$x^i_{ct}(n) = \beta^i \times HighRBC_t(n) \times 20y_t(n) \times Post_t + \alpha_c + \alpha_{r\times m} + \alpha_t + u_{ict}(n), \qquad (16)$$

where $x^i_{ct}(n)$ is the portfolio weight of bond n in the portfolio of investor type i from country c in quarter t, and all other variable definitions remain unchanged. The regression includes investor country, time, and rating-by-time-to-maturity fixed effects, and our estimation window is 2015Q1 – 2016Q4. Following Jansen (2025), we only include investor countries that

corresponding fixed effects control for duration and credit risk. In a benchmark model with no demand for service flows, these would be the two main sources of risk driving investors' portfolio allocation and prices.

[30] Again, the underlying regression we estimate is fully saturated, but we only report the coefficients of interest here for brevity.

Table 3: CDS-bond basis and regulatory value. This table displays the results of estimating equation (15). *Post* is a dummy variable equal to one for all months after (including) January 2016, when the Solvency II regulation entered into force. *HighRBC* is a binary variable that indicates whether a bond's regulatory value (defined as $1 - stress$) exceeds the monthly sample median. $20y$ is a dummy variable equal to 1 iff the modified duration of the bond lies within 15–25 years (Jansen, 2025), and *Duration* is the modified bond duration in years. The sample period is January 2015 – December 2016, at monthly frequency. The regression we estimate is fully saturated, but we omit some interaction effects and standalone coefficients here for brevity.

	Dependent variable: CDS-bond basis	
	(1)	(2)
High RBC=1 × 20y=1 × Post=1	0.207***	
	(0.0362)	
High RBC=1 × Post=1 × Duration		0.00453***
		(0.00155)
High RBC=1 × Post=1 × Duration		
High RBC=1 × Post=1	-0.0712***	-0.0941***
	(0.00691)	(0.0111)
20y=1 × Post=1	-0.164***	
	(0.0295)	
Post=1 × Duration		-0.00624***
		(0.00123)
Time FE	✓	✓
Rating x Maturity FE	✓	✓
R^2	0.14	0.14
Observations	87,829	87,829

Standard errors in parentheses * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

had applied historical cost accounting prior to Solvency II, which should amplify the incentive to rebalance toward longer maturities, especially around the 20-year mark. Our hypothesis is that β^i is positive for ICPF, as they rebalance toward high-RBC bonds (conditional on 20-year duration) in response to the implementation of Solvency II on January 1, 2016.

Strikingly, as column (3) of Table 4 shows, our estimate is positive and significant *only* for ICPF—consistent with the implicit assumption that Solvency II and our RBC measure are specific to this sector. Numerically, the estimate means that, after Solvency II became binding, the portfolio share of high-RBC bonds increased by 1.5 basis points more than that of low-RBC bonds (within the set of 20-year bonds). Given that ISIN-level portfolio shares are quite small, this estimate corresponds to an economically meaningful increase of 44% relative to the average portfolio share of 3.5 basis points.

In conclusion, the results presented in this subsection serve two purposes: (a) they validate the argument in Jansen (2025) about how Solvency II (and, in particular, the introduc-

Table 4: Sectoral holdings and regulatory value. This table displays the results of estimating equation (16) separately for four investor sectors: Banks, mutual funds, insurance companies & pension funds, and others (incl. households and non-financial corporations). The sample includes only investor sectors from countries where insurers did not have to mark-to-market their liabilities to determine regulatory capital requirements prior to Solvency II (i.e., they relied on historical cost accounting). These countries are Austria, Belgium, Finland, France, Germany, Greece, Italy, Portugal, Slovakia, and Spain (Jansen, 2025). *Post* is a dummy variable equal to one for all quarters after (including) 2016Q1, when the Solvency II regulation entered into force. $HighRBC$ is a binary variable that indicates whether a bond's regulatory value (defined as $1 - stress$) exceeds the monthly sample median. $20y$ is a dummy variable equal to 1 iff the modified duration of the bond falls within 15–25 years (Jansen, 2025). The sample period is 2015Q1 – 2016Q4, at quarterly frequency. The dependent variable has been rescaled to be measured in basis points.

	Dependent variable: Portfolio share			
	(1) Banks	(2) MF	(3) ICPF	(4) Other
High RBC=1 × 20y=1 × Post=1	-1.062	-0.874	1.565***	-0.221
	(0.909)	(0.812)	(0.537)	(0.713)
High RBC=1 × Post=1	-0.0170	-0.494***	-0.301***	-0.115
	(0.106)	(0.110)	(0.0937)	(0.0769)
20y=1 × Post=1	-0.108	-0.272	-0.323	0.0445
	(0.841)	(0.736)	(0.455)	(0.652)
Time FE	✓	✓	✓	✓
Rating x Maturity FE	✓	✓	✓	✓
Holder country FE	✓	✓	✓	✓
R^2	0.12	0.15	0.17	0.08
Observations	74,445	87,793	98,712	102,116

Standard errors in parentheses * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

tion of the UFR in discounting insurers' liabilities) affected insurers' regulatory incentives to hold long-term, safe bonds; and (b) they validate our measure of risk-based regulatory capital value as a quantifiable service that ICPF value and that gives rise to convenience yield.

6.2 Collateral Value for the Banking Sector

We study the collateral value of bonds held by different sectors. Collateral is especially valuable for banks, as they rely on it for liquidity management and are the only institutions that can borrow directly from the ECB's lending facilities by pledging high-quality bonds—including corporate bonds—as collateral. ECB collateral eligibility is therefore a service that banks are particularly likely to value. Consistent with this hypothesis, banks clearly emerge as the sector holding bonds of the highest collateral value (i.e., the lowest haircuts) in Figure

6 above. After a sharp drop in December 2018, euro area banks replenished their collateral pools by purchasing disproportionately more collateral-eligible bonds.

The striking drop in the average collateral value of banks' bond holdings in December 2018 was caused by the loss of eligibility of a large set of German bank bonds, providing us with a useful policy experiment that can help understand the drivers of bond convenience yields. More precisely, in December 2017, the ECB announced that (statutorily) *subordinated* unsecured bank bonds (sUBB) issued before April 16, 2018 would cease to be eligible for ECB collateral after December 31, 2018. Due to differences in the national implementation of Directives across European Union (EU) Member States, this shock only affected sUBBs issued by German banks.[31]

This change in collateral eligibility offers a clean shock for studying the value of pledgeability and how it maps to the CDS-bond basis. The most useful aspect of this experiment is that the loss of eligibility occurs orthogonally to other bond characteristics, such as credit risk and residual maturity. Using the publicly available EADB data, we can identify the set of German sUBBs and interpret the outright loss of their collateral value as a reduction in the service flows that these bonds provide. To test this hypothesis formally, we compare the average CDS-bond basis of eligible German sUBBs issued before April 2018 (the "treatment group") with that of a control group (consisting of eligible bank bonds from other EA countries), before and after the loss of eligibility at the end of 2018. Specifically, we estimate the DiD equation

$$b_t(n) = \beta\, German(n) \times Post_t + \alpha_t + \alpha_n + \alpha_{r\times m} + \varepsilon_t(n), \tag{17}$$

where $b_t(n)$ denotes the CDS-bond basis of bond n in month t, $German(n)$ is a dummy variable identifying bonds issued by German banks, and $Post_t$ is a dummy variable equal to one for all months after (and including) December 2018, and zero otherwise. α_t, α_n, and $\alpha_{r\times m}$ are placeholders for time, bond, and rating $\times$ residual maturity fixed effects, respectively. Since bank bonds often exhibit small lot sizes, we weight observations by outstanding amounts to reduce the impact of small issuances. We estimate equation (17)

[31] See Appendix F for a more detailed account of the collateral framework amendment regarding sUBBs, including a timeline of events.

Table 5: The effect of ECB collateral eligibility on convenience yield. This table displays the results of estimating equation (17), where each bond-month observation is weighted by its outstanding amount. The sample consists of senior unsecured bank bonds issued before April 2018 by euro area banks (which were eligible to be pledged as collateral with the ECB). The binary treatment variable *German* identifies German bank bonds that became ineligible for use as ECB collateral after December 2018. The indicator variable *Post* is equal to one for all observations after that event, and zero before. The sample period is June 2018 – June 2019, at monthly frequency. Standard errors are clustered at the bond-level.

	Dependent variable: CDS-bond basis			
	(1)	(2)	(3)	(4)
German=1 × Post=1	-0.0364	-0.0401*	-0.0434*	-0.0699***
	(0.0226)	(0.0218)	(0.0226)	(0.0139)
German=1	0.0195	0.0228	0.0384	
	(0.0418)	(0.0414)	(0.0350)	
Post=1	-0.0594***			
	(0.0165)			
Time FE		✓	✓	✓
ISIN FE				✓
Rating x Maturity FE			✓	
R^2	0.01	0.03	0.15	0.85
Observations	16,940	16,940	16,940	16,894

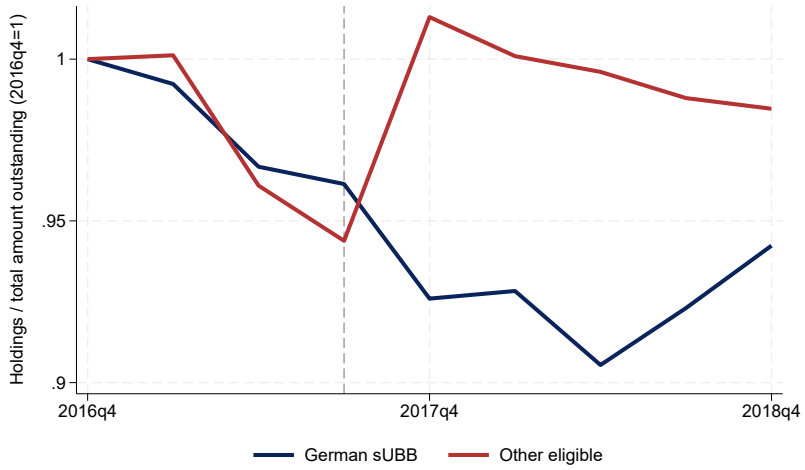
Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

during the period June 2018 – June 2019 and present the results in Table 5. In line with our hypothesis that collateral eligibility is a major driver of convenience yield, the loss of eligibility of German sUBB leads to a statistically significant drop in the CDS-bond basis of 4–7 basis points relative to bank bonds from other countries. The baseline estimate is robust to using variation within rating × maturity buckets and even within ISINs for identification. Given an average pre-shock CDS-bond basis of -40 basis points, even our smallest estimate corresponds to an economically meaningful 10% reduction in convenience yield.[32]

To complement our analysis of the impact of ECB collateral eligibility on bond prices, we now look at the quantity response, i.e., the portfolio rebalancing behavior of European banks. As Figure 11 illustrates, banks started decreasing their aggregate holdings of German sUBB relative to other eligible senior bank bonds immediately after the ECB announced in

[32]We also verified that, before the loss of eligibility, the average CDS-bond basis of both treatment and control group followed parallel trends, strengthening our identifying assumption.

Figure 11: Banks' holdings of bank bonds around loss of collateral eligibility. This plot displays the nominal value of banks' total holdings of German sUBB (blue) and those of other collateral-eligible bank bonds (pink), each divided by their total amount outstanding and normalized to their respective 2016Q4 holdings. In December 2017, German sUBB were announced to lose ECB collateral eligibility one year later, at the end of 2018.



December 2017 that German bonds would lose eligibility one year later. More precisely, between the start of 2017 and the actual loss of eligibility at the end of 2018, banks' holdings of German sUBB (as a share of their total outstanding amount) fell by almost 5% whereas holdings of other eligible bank bonds rose by 5%.

For a formal test, we again employ a DiD design and estimate the following equation, where we only focus on collateral-eligible senior bank bonds issued before April 2018 to keep the treatment and control group as similar as possible:

$$\frac{B_{it}(n)}{\sum_n B_{it}(n)} = \beta_i German(n) \times Post_t + \gamma_i German(n) + \alpha_c + \alpha_t + \alpha_{r\times m} + \varepsilon_{it}(n) \qquad (18)$$

To ensure a tight correspondence between the empirical specification and the theoretical model predictions in equation (2), $B_{it}(n)$ denotes the market value of investor i's holdings of bond n in quarter t, so the left-hand side expression can be interpreted as the weight of bond n in investor i's portfolio of bank bonds. As in the previous equation, $German(n)$ is a dummy variable identifying bonds issued by German banks, and $Post_t$ is a dummy variable equal to one for all quarters after (and including) December 2017, and zero otherwise. α_c, α_t, and $\alpha_{r\times m}$ denote investor country, time, and rating $\times$ residual maturity fixed effects, respectively. Controlling for unobservable time-invariant investor country characteristics via

Table 6: The effect of ECB collateral eligibility on sectoral bank bond holdings. This table displays the results of estimating equation (18) separately for five holdings sectors: Banks, mutual funds, insurance companies & pension funds, others (incl. households and non-financial corporations), and rest of the world (i.e., non-euro area holdings estimated as the residual). The sample consists of senior unsecured bank bonds issued before April 2018 by euro area banks (which were eligible to be pledged as collateral with the ECB). The binary treatment variable *German* identifies German bank bonds that became ineligible for use as ECB collateral after December 2018. The indicator variable *Post* is equal to one for all observations after the announcement in December 2017, and zero before. The sample period is 2016Q4 – 2018Q4, at quarterly frequency. The dependent variable has been rescaled to be measured in basis points. Standard errors are clustered at the bond-level.

	Dependent variable: Portfolio share				
	(1)	(2)	(3)	(4)	(5)
	Banks	MF	ICPF	Other	RoW
German=1 × Post=1	-4.532***	-0.686	-3.018	-0.661	-1.289
	(1.090)	(2.122)	(1.912)	(1.043)	(0.827)
German=1	-6.172***	-12.43***	-16.93***	-9.515***	-14.30***
	(1.201)	(2.765)	(2.090)	(1.197)	(1.497)
Time FE	✓	✓	✓	✓	✓
Holder country FE	✓	✓	✓	✓	✓
Rating x Maturity FE	✓	✓	✓	✓	✓
R^2	0.40	0.21	0.19	0.17	0.10
Observations	26,860	21,351	24,221	27,677	10,888

Standard errors in parentheses: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

α_c is particularly important here, since the loss of eligibility is country-specific and it is well known that investors' portfolios exhibit "home bias". We estimate equation (18) separately for each investor sector during the period 2016Q4 – 2018Q4.

The results are summarized in Table 6. The statistically significant estimate of interest in column (1) implies that European banks reduced the share of German bank bonds in their bank bond portfolios by 4.5 basis points on average (compared with the control group) in response to the announced loss of collateral eligibility. Since the average portfolio share at the bond level in the estimation sample amounts to 0.3%, the estimate corresponds to a significant reduction of 15% for the average bond. Crucially, we do not find the same rebalancing response for the other investor sectors (which arguably do not value ECB collateral

eligibility).

Taken together, these results confirm Prediction 4 above and suggest that the "service flow" of collateral eligibility is a major component of convenience yields in euro area bonds.

6.3 Corporate QE as a Shock to Bond Liquidity

Finally, we study the effect of corporate QE announcements on corporate bond prices and convenience yields. We show evidence that bonds eligible for the ECB purchases experienced an increase in their liquidity service—a service especially valued by mutual funds. In our companion paper, Corell et al. (2025), we develop this argument in more detail.

We study the Corporate Sector Purchase Program (CSPP), announced in March 2016 and launched in June 2016, which introduced large-scale net purchases of investment-grade non-bank corporate bonds in the euro area. In total, the ECB accumulated more than €400 billion by 2023, or approximately 15% of the total European corporate bond market.

Various criteria determined the eligibility of the bonds to be purchased under the CSPP. Specifically, the CSPP was limited to investment-grade non-bank corporate bonds issued in the euro area. Only bonds with a remaining maturity of between 6 months and 31 years and with a yield to maturity (YTM) above the ECB's deposit facility rate were CSPP-eligible.

The exogenous nature of the CSPP announcement allows us to employ a difference-in-differences (DiD) identification strategy to study the impact of the program on bond yields and their decomposition into credit spreads, CDS spreads, and the CDS-bond basis for eligible versus ineligible bonds. Our main specification is:

$$y_t(n) = \theta \, Treat_t(n) \times Post_t + \alpha_{r \times m} + \alpha_t + u_t(n) \tag{19}$$

For each bond n and month t the dependent variables $y_t(n)$ are: corporate bond yields, credit spreads, CDS-spread, or the CDS-bond basis. $Treat_t(n)$ is a dummy equal to 1 if the bond is eligible to be purchased under the CSPP and 0 otherwise;[33] $Post_t$ is a dummy equal to 0 for the period preceding the announcement date of the respective purchase program,

[33]The eligibility criteria for the CSPP can be summarized in the following: a) to be rated by accepted external credit assessment institutions and third-party rating tools and to have a minimum rating of BBB-; b) to have initial maturity of 367 days or more and the minimum remaining maturity to be six months and the maximum remaining maturity 30 years and 364 days; c) to be euro-denominated; d) to be issued in the euro area; e) the issuer to be established in the euro area and the place of settlement to be in the euro area; and f) to have either a fixed and

and 1 after it. We include month and rating × maturity fixed effects, just as in the previous subsection.

The identifying assumption is that, in the absence of the ECB purchases, corporate bond yields, credit spreads, CDS spreads, and CDS-bond basis of the treatment group and the control group would have maintained parallel trends. Regardless of the parallel trend assumption, one could still worry that eligible and ineligible bonds are fundamentally different, and that ECB decisions could potentially be concurrent with other news that affect the two groups differently. For example, the announcement of ECB purchase programs is usually a response to the deterioration of an economic scenario. Since eligible bonds are, in general, of higher credit quality than non-eligible bonds, they might be less exposed to macroeconomic risks. Therefore, the ex-post credit spread gap between the groups might just be the consequence of a higher risk premium. In contrast, analyses based on the CDS-bond basis are not subject to this critique. Any variation in the risk premium that affects default should be reflected in the CDS-spreads, since the payoff of the CDS is directly linked to the default probability. Furthermore, our empirical methodology mitigates concerns about the mismeasurement of the default component of credit spreads. Any friction that could affect CDS spreads and is not related to default risk, such as counterparty risk, would have to affect treatment and control group differently in order to affect our results. This is not likely to be the case.

To estimate the baseline equation (19), we define six months before the announcement date as the pre-period and six months after as the post-period. We differ from other studies (most importantly, Todorov (2020)) and include a longer window around the event. This longer window is crucial to connect the results on the market price responses to the portfolio rebalancing exercise that follows. Holdings are unlikely to change significantly in very short windows, especially when considering long-term investors such as insurance companies and pension funds.

In Table 7, we examine the impact of the CSPP in March 2016. We find that the CSPP

unconditional principal amount or an unconditional principal amount that is linked, on a flat basis, to only one euro area inflation index at a single point in time, containing no other complex structures. Further criteria are included with respect to acceptable markets and coupon payments, but they do not apply for our analysis as the sample used in the analysis is already restricted.

Table 7: The effect of the CSPP on corporate bond spreads. This table displays the results of estimating equation (19) for different outcome variables related to bond yields. Standard errors are clustered at the firm and month level. The sample period is September 2015 – September 2016.

	Dependent variable: CDS - Bond Basis			
	(1)	(2)	(3)	(4)
Post=1 × CSPP eligible=1	0.122**	0.123**	0.131**	0.143**
	(0.0472)	(0.0482)	(0.0481)	(0.0612)
Post=1	0.0198			
	(0.0448)			
CSPP eligible=1	-0.159**	-0.160**	-0.199***	
	(0.0555)	(0.0575)	(0.0546)	
Time FE		✓	✓	✓
ISIN FE				✓
Rating x Maturity FE			✓	
R^2	0.01	0.01	0.11	0.79
Observations	37,999	37,999	37,999	37,782

Standard errors in parentheses * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

decreased the credit spread of eligible bonds with respect to ineligible bonds by 22.1 basis points, consistent with our Prediction 4. Columns (3) and (4) show that the decrease in credit spreads is explained mainly by the non-default component of the spread, the CDS-bond basis, and less by the default component, the CDS spread.[34]

Our last step is to understand *why* ECB eligibility caused an increase in the convenience yield of eligible bonds. To unveil the mechanism through which the corporate QE program of the ECB affect corporate bond prices, we look at the response of different investors around the program announcement.

For that reason, we estimate the effect of the ECB corporate QE program on portfolio rebalancing, using micro-level holdings data. Holdings information at the ISIN-level allows us to analyze the heterogeneous investor composition and infer how different investors value the QE-eligibility feature.

Specifically, we use a DiD specification where the dependent variable is the portfolio share

[34]We verify the parallel trend assumption in our companion paper, Corell et al. (2025).

Table 8: The effect of corporate QE on sectoral corporate bond holdings. This table displays the results of estimating equation (20) separately for five holdings sectors: Banks, mutual funds, insurance companies & pension funds, others (incl. households and non-financial corporations), and rest of the world (i.e., non-euro area holdings estimated as the residual). The CSPP announcement date defining the *Post* dummies is 2016Q1. The dependent variable has been rescaled to be measured in basis points. Standard errors are clustered at the firm and quarter level. The sample period is 2015Q1 – 2016Q4.

	Dependent variable: Portfolio share				
	(1) Banks	(2) MF	(3) ICPF	(4) Other	(5) RoW
Post=1 × CSPP eligible=1	-0.335	0.271**	-0.00468	-0.444***	-0.408
	(0.189)	(0.0814)	(0.0742)	(0.0910)	(0.254)
CSPP eligible=1	-2.533**	-0.967*	0.362	-1.430**	1.715*
	(0.661)	(0.361)	(0.410)	(0.348)	(0.650)
Holder country x Time FE	✓	✓	✓	✓	✓
Rating x Maturity FE	✓	✓	✓	✓	✓
R^2	0.29	0.22	0.25	0.15	0.02
Observations	45,888	71,403	74,935	74,312	20,335

Standard errors in parentheses: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

$x^c_{it}(n)$ for each investor type i, in country c, in asset n, and quarter t.

$$x^c_{it}(n) = \psi_i Treat(n) \times Post_t + \alpha_{r\times m} + \alpha_{c\times t} + \varepsilon_{i,t}(n) \tag{20}$$

where the time-varying outside option of investors in country c is captured by $\alpha_c \times \alpha_t$.

Maturity × rating fixed effects are meant to capture duration and credit risk. Holder country × quarter fixed effects are meant to capture investor-specific demand (which is likely to be country-specific) and time-varying outside options.

The results are reported in Table 8. They show that the CSPP announcement caused an increase in holdings for eligible bonds only for mutual funds, while all other investor sectors decreased their holdings of eligible bonds relative to ineligible. This finding aligns with the hypothesis that mutual funds strategically maintain eligible bonds, even when they appear relatively expensive compared to other bonds with similar credit and duration risks.

These results are striking, in that we observe that mutual funds buy more eligible bonds as they become more expensive, even though the mutual fund sector has been shown to

be very price elastic (Koijen et al. (2021)). We explain this pattern by QE announcements affecting prices not only through supply changes, but also by altering the liquidity properties of eligible bonds.[35] Because mutual funds place particular value on liquidity, and they are the investors that rebalance most strongly toward eligible bonds, we view this as evidence that QE fundamentally changes the liquidity services of eligible bonds. We provide comprehensive additional evidence for this interpretation in our companion paper Corell et al. (2025).

7 Conclusion

This paper studies the sources of convenience yields in fixed-income markets and how they vary across assets and investors. We develop a framework in which financial assets provide services beyond their cash flows and investors differ in the value they attach to these services. Guided by this framework, we decompose bond convenience yields into three key components—liquidity, collateral value, and regulatory capital value—and estimate the premium associated with each service using a comprehensive dataset that combines bond characteristics, prices, and confidential sectoral portfolio holdings in the euro area.

Our results show that regulatory capital value is the dominant driver of the convenience yield of AAA euro area sovereign bonds. Favorable regulatory treatment under risk-based capital rules lowers sovereign borrowing costs by roughly 18 basis points on average, generating substantial financing savings for governments. Because sovereign bonds carry zero capital charges for regulated intermediaries, banks and insurance companies have strong incentives to hold them even at lower returns. Our results therefore provide micro-level evidence of a mechanism commonly associated with financial repression: regulation induces intermediaries to absorb government debt at yields below those that would prevail in the absence of these institutional incentives.

Convenience yields are not borne uniformly across investors. Regulated intermediaries tilt their portfolios toward sovereign bonds and, within the corporate bond universe, toward bonds with high regulatory value. As a result, they earn lower hedged returns than mutual funds, which allocate more strongly toward corporate bonds with lower service values.

[35] The liquidity concept here is broader than the bid-ask spreads we measured in the previous section. It encompasses the expectation that the ECB will step in to purchase these bonds, especially in times of distress.

Finally, we show that policy-driven changes in asset services have causal effects on both portfolio allocation and asset prices. Reforms to capital regulation, changes in the ECB collateral framework, and the ECB's corporate bond purchase program all generate measurable shifts in service flows that lead investors to rebalance their portfolios and affect convenience yields.

In conclusion, service flows are fundamental drivers of asset prices with significant implications for monetary policy transmission, financial stability, and market efficiency. By providing a framework for measuring the value of these service flows, we open new avenues for understanding how financial markets function and how policy interventions can achieve more targeted outcomes.

References

Acharya, Viral V, and Toomas Laarits, 2025, When do treasuries earn the convenience yield? A hedging perspective, *Journal of Finance* forthcoming.

Acharya, Viral V., and Lasse Heje Pedersen, 2005, Asset pricing with liquidity risk, *Journal of Financial Economics* 77, 375–410.

Amihud, Yakov, and Haim Mendelson, 1986, Asset pricing and the bid-ask spread, *Journal of Financial Economics* 17, 223–249.

Amihud, Yakov, Haim Mendelson, Lasse Heje Pedersen, et al., 2006, Liquidity and asset prices, *Foundations and Trends® in Finance* 1, 269–364.

Bai, Jennie, and Pierre Collin-Dufresne, 2019, The cds-bond basis, *Financial Management* 48, 417–439.

Bao, JACK, JUN Pan, and JIANG Wang, 2011, The illiquidity of corporate bonds, *The Journal of Finance* 66, 911–946.

Becker, Bo, and Victoria Ivashina, 2015, Reaching for yield in the bond market, *The Journal of Finance* 70, 1863–1902.

Becker, Bo, and Victoria Ivashina, 2018, Financial repression in the european sovereign debt crisis [sovereign debt, government myopia and the financial sector], *Review of Finance* 22, 83–115.

Bretscher, Lorenzo, Lukas Schmid, Varun Sharma, and Ishita Sen, 2025, Institutional corporate bond pricing, *Review of Financial Studies* forthcoming.

Chari, Varadarajan Venkata, Alessandro Dovis, and Patrick J Kehoe, 2020, On the optimality of financial repression, *Journal of Political Economy* 128, 710–739.

Chen, Hui, Zhuo Chen, Zhiguo He, Jinyu Liu, and Rengming Xie, 2023, Pledgeability and asset prices: Evidence from the chinese corporate bond markets, *The Journal of Finance* 78, 2563–2620.

Chen, Xuanjuan, Jing-Zhi Huang, Zhenzhen Sun, Tong Yao, and Tong Yu, 2020, Liquidity premium in the eye of the beholder: An analysis of the clientele effect in the corporate bond market, *Management Science* 66, 932–957.

Corell, Felix, Fédéric Holm-Hadulla, Matteo Leombroni, Lira Mota, and Melina Papoutsi, 2025, The Liquidity Promises of QE .

Daniel, Kent, Lira Mota, Simon Rottke, and Tano Santos, 2020, The cross-section of risk and returns, *The Review of Financial Studies* 33, 1927–1979.

Daniel, Kent D., and Sheridan Titman, 1997, Evidence on the characteristics of cross sectional variation in stock returns, *The Journal of Finance* 52, 1–33.

Darmouni, Olivier, and Melina Papoutsi, 2023, The rise of bond financing in europe: Five facts about new and small issuers .

Domanski, Dietrich, Hyun Song Shin, and Vladyslav Sushko, 2015, The Hunt for Duration: Not Waving but Drowning?, *BIS Working Papers No 519* .

Elton, Edwin J, Martin J Gruber, Deepak Agrawal, and Christopher Mann, 2001, Explaining the rate spread on corporate bonds, *the journal of finance* 56, 247–277.

Fama, Eugene F, and Kenneth R French, 1993, Common risk factors in the returns on stocks and bonds, *Journal of financial economics* 33, 3–56.

Fama, Eugene F., and James D. MacBeth, 1973, Risk, return, and equilibrium: Empirical tests, *Journal of Political Economy* 81, 607–636.

Fontana, Alessandro, and Martin Scheicher, 2016, An analysis of euro area sovereign cds and their relation with government bonds, *Journal of Banking & Finance* 62, 126–140.

Frazzini, Andrea, and Lasse Heje Pedersen, 2014, Betting against beta, *Journal of Financial Economics* 111, 1–25.

Gabaix, Xavier, and Ralph S. J. Koijen, 2020, In Search of the Origins of Financial Fluctuations: The Inelastic Markets Hypothesis, *SSRN Electronic Journal* .

Gorton, Gary, 2017, The History and Economics of Safe Assets, *Annual Review ofEconomics* 9, 547–86.

Haddad, Valentin, Alan Moreira, and Tyler Muir, 2025, Whatever it takes? The impact of conditional policy promises, *American Economic Review* 115, 295–329.

Huang, Jing-Zhi, and Ming Huang, 2012, How much of the corporate-treasury yield spread is due to credit risk?, *The Review of Asset Pricing Studies* 2, 153–202.

Jansen, Kristy AE, 2025, Long-term investors, demand shifts, and yields, *The Review of Financial Studies* 38, 114–157.

Jeanne, Olivier, 2025, From fiscal deadlock to financial repression: Anatomy of a fall, Technical report, National Bureau of Economic Research.

Jiang, Zhengyang, Hanno N Lustig, Stijn Van Nieuwerburgh, and Mindy Z Xiaolan, 2025, Bond convenience yields in the eurozone currency union, *Review of Financial Studies* forthcoming.

Koijen, Ralph SJ, François Koulischer, Benoît Nguyen, and Motohiro Yogo, 2021, Inspecting the mechanism of quantitative easing in the euro area, *Journal of Financial Economics* 140, 1–20.

Koijen, Ralph SJ, and Motohiro Yogo, 2015, The cost of financial frictions for life insurers, *American Economic Review* 105, 445–475.

Koijen, Ralph S.J., and Motohiro Yogo, 2019, A demand system approach to asset pricing, *Journal of Political Economy* 127, 1475–1515.

Koijen, Ralph SJ, and Motohiro Yogo, 2022a, The fragility of market risk insurance, *The Journal of Finance* 77, 815–862.

Koijen, Ralph SJ, and Motohiro Yogo, 2022b, Global life insurers during a low interest rate environment, in *AEA Papers and Proceedings*, volume 112, 503–508, American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203.

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

Krishnamurthy, Arvind, and Annette Vissing-Jorgensen, 2012b, Why an MBS-Treasury swap is better policy than the Treasury twist, Technical report.

Kubitza, Christian, 2025, *Investor-driven corporate finance: Evidence from insurance markets* (ECB Working Paper No. 2816).

Li, Ziang, 2025, Long rates, life insurers, and credit spreads .

Longstaff, Francis A., Sanjay Mithal, and Eric Neis, 2005, Corporate yield spreads: Default risk or liquidity? New evidence from the credit default swap market, *Journal of Finance* 60, 2213–2253.

Ma, Yiming, Kairong Xiao, and Yao Zeng, 2022, Mutual fund liquidity transformation and reverse flight to liquidity, *Review of Financial Studies* 35, 4674–4711.

McKinnon, Ronald I., 1973, *Money and Capital in Economic Development* (Brookings Institution, Washington, DC).

Merton, Robert C., 1973, An intertemporal capital asset pricing model, *Econometrica* 41, 867–887.

Mian, Atif R, Ludwig Straub, and Amir Sufi, 2022, A goldilocks theory of fiscal deficits, Technical report, National Bureau of Economic Research.

Mota, Lira, 2023, The corporate supply of (quasi) safe assets, *Available at SSRN* .

Oehmke, Martin, and Adam Zawadowski, 2017, The anatomy of the cds market, *The Review of Financial Studies* 30, 80–119.

Pastor, Lubos, and Robert F Stambaugh, 2003, Liquidity Risk and Expected Stock Returns, *Journal of Political Economy* 111, 642–685.

Payne, Jonathan, and Bálint Szoke, 2025, The fragility of government funding advantage, *Princeton University Department of Economics* .

Pelizzon, Loriana, Max Riedel, Zorka Simon, and Marti G Subrahmanyam, 2024, Collateral eligibility of corporate debt in the eurosystem, *Journal of financial economics* 153, 103777.

Pástor, Ľuboš, Robert F. Stambaugh, and Lucian A. Taylor, 2021, Sustainable investing in equilibrium, *Journal of Financial Economics* 142, 550–571.

Sen, Ishita, 2023, Regulatory limits to risk management, *The Review of Financial Studies* 36, 2175–2223.

Shaw, Edward S., 1973, *Financial Deepening in Economic Development* (Oxford University Press, New York).

Sunderam, Adi, 2015, Money creation and the shadow banking system, *Review of Financial Studies* 28, 939–977.

Todorov, Karamfil, 2020, Quantify the quantitative easing: Impact on bonds and corporate debt issuance, *Journal of Financial Economics* 135, 340–358.

Vayanos, Dimitri, and Jean-Luc Vila, 2021, A preferred-habitat model of the term structure of interest rates, *Econometrica* 89, 77–112.

Vayanos, Dimitri, and Jiang Wang, 2013, Chapter 19 - market liquidity—theory and empirical evidence *, volume 2 of *Handbook of the Economics of Finance*, 1289–1361 (Elsevier).

Appendix

Appendix A Model Derivations and Proofs

Appendix A.1 Optimal Portfolio Choice

Investor i chooses portfolio weights $\boldsymbol{x}_{it}$ to maximize expected utility

$$\mathbb{E}_t\Big[-\exp\Big(-A_i w_{i,t+1} - \boldsymbol{x}_{it}'\boldsymbol{Z}_t\tilde{\boldsymbol{\lambda}}_{it}\Big)\Big], \tag{A.1}$$

where next-period wealth satisfies

$$w_{i,t+1} = w_{i0}\big(1 + r_t^f + \boldsymbol{x}_{it}'\boldsymbol{r}_{t+1}\big)$$

and excess returns are normally distributed according to

$$\boldsymbol{r}_{t+1} = \boldsymbol{\mu}_t + \boldsymbol{\epsilon}_{t+1}, \qquad \boldsymbol{\epsilon}_{t+1} \sim \mathcal{N}(\mathbf{0}, \Sigma_t).$$

Under CARA preferences and normal returns, maximizing expected utility is equivalent to maximizing a mean–variance objective augmented by service flows. Substituting for wealth and collecting all terms that depend on $\boldsymbol{x}_{it}$, investor i solves

$$\max_{\boldsymbol{x}_{it}} \left\{\boldsymbol{x}_{it}'\boldsymbol{\mu}_t + \boldsymbol{x}_{it}'\boldsymbol{Z}_t\tilde{\boldsymbol{\lambda}}_{it} - \frac{a_i}{2}\boldsymbol{x}_{it}'\Sigma_t\boldsymbol{x}_{it}\right\}, \tag{A.2}$$

where $a_i = A_i w_{i0}$ is relative risk aversion.

The first-order condition yields

$$\boldsymbol{x}_{it} = \frac{1}{a_i}\Sigma_t^{-1}\left(\boldsymbol{\mu}_t + \boldsymbol{Z}_t\tilde{\boldsymbol{\lambda}}_{it}\right), \tag{A.3}$$

which corresponds to equation (2) in the main text.

Market Clearing and Expected Excess Returns Let $\tilde{w}_{i0} = w_{i0}/w_0$ denote investor i's share of total wealth. With the risk-free asset in zero net supply, market clearing for risky assets implies that the market portfolio satisfies

$$\boldsymbol{x}_{mt} = \int_i \tilde{w}_{i0}\boldsymbol{x}_{it}\, di. \tag{A.4}$$

Imposing the common risk-aversion assumption $a_i = a$ and defining $\boldsymbol{\lambda}_{it} = \tilde{\boldsymbol{\lambda}}_{it}/a$, aggregation yields

$$\boldsymbol{x}_{mt} = \frac{1}{a}\Sigma_t^{-1}\left(\boldsymbol{\mu}_t + \boldsymbol{Z}_t\bar{\boldsymbol{\lambda}}_t\right), \tag{A.5}$$

where

$$\bar{\boldsymbol{\lambda}}_t = \int_i \tilde{w}_{i0}\boldsymbol{\lambda}_{it}\, di$$

is the wealth-weighted valuation of service flows.

Multiplying both sides by Σ_t and rearranging gives

$$\boldsymbol{\mu}_t = a\,\Sigma_t\boldsymbol{x}_{mt} - \boldsymbol{Z}_t\bar{\boldsymbol{\lambda}}_t. \tag{A.6}$$

Define the market excess return as

$$r_{m,t+1} = \boldsymbol{x}_{mt}'\boldsymbol{r}_{t+1},$$

with expected excess return

$$\mu_{mt} = \boldsymbol{x}_{mt}'\boldsymbol{\mu}_t$$

and variance

$$\sigma_{mt}^2 = \boldsymbol{x}_{mt}'\Sigma_t\boldsymbol{x}_{mt}.$$

Let $\boldsymbol{\beta}_t$ denote the vector of market betas,

$$\boldsymbol{\beta}_t = \frac{\Sigma_t\boldsymbol{x}_{mt}}{\sigma_{mt}^2}. \tag{A.7}$$

Substituting equation (A.6) into the definition of μ_{mt} yields

$$\mu_{mt} = a\,\sigma_{mt}^2 - \boldsymbol{x}_{mt}'\boldsymbol{Z}_t\bar{\boldsymbol{\lambda}}_t. \tag{A.8}$$

Under the normalization imposed in the main text,

$$\boldsymbol{x}_{mt}'\boldsymbol{Z}_t = \boldsymbol{0},$$

the market risk premium satisfies

$$\mu_{mt} = a\,\sigma_{mt}^2, \qquad a = \frac{\mu_{mt}}{\sigma_{mt}^2}.$$

Plugging this expression for a back into equation (A.6) yields

$$\boldsymbol{\mu}_t = \frac{\mu_{mt}}{\sigma_{mt}^2}\Sigma_t\boldsymbol{x}_{mt} - \boldsymbol{Z}_t\bar{\boldsymbol{\lambda}}_t \tag{A.9}$$

$$= \boldsymbol{\beta}_t\,\mu_{mt} - \boldsymbol{Z}_t\bar{\boldsymbol{\lambda}}_t, \tag{A.10}$$

which corresponds to equation (4) in the main text.

Appendix A.2 Proofs of Model Predictions

Throughout, we maintain the normalization $\boldsymbol{x}'_{mt}\boldsymbol{Z}_t = \mathbf{0}$ imposed in the main text, and we use the equilibrium expected return relation

$$\boldsymbol{\mu}_t = \boldsymbol{\beta}_t \, \mu_{mt} - \boldsymbol{Z}_t \bar{\boldsymbol{\lambda}}_t, \tag{A.11}$$

together with the definition of the convenience value (convenience yield) of asset n,

$$cy_t(n) = \boldsymbol{Z}_t(n) \, \bar{\boldsymbol{\lambda}}_t. \tag{A.12}$$

We also use the portfolio-tilt representation

$$\boldsymbol{x}_{it} = \boldsymbol{x}_{mt} + \frac{1}{a} \Sigma_t^{-1} \boldsymbol{Z}_t \left(\boldsymbol{\lambda}_{it} - \bar{\boldsymbol{\lambda}}_t \right), \tag{A.13}$$

which follows from equation (2) and market clearing.

Proof of Prediction 1. Fix t and the equilibrium price of services $\bar{\boldsymbol{\lambda}}_t$. From equation (5), the convenience yield of asset n is

$$cy_t(n) = \boldsymbol{Z}_t(n) \, \bar{\boldsymbol{\lambda}}_t.$$

For any two assets n and m, it follows that

$$cy_t(n) - cy_t(m) = \left(\boldsymbol{Z}_t(n) - \boldsymbol{Z}_t(m) \right) \bar{\boldsymbol{\lambda}}_t. \tag{A.14}$$

Thus, conditional on $\bar{\boldsymbol{\lambda}}_t$, cross-sectional differences in convenience yields are driven solely by cross-sectional differences in service flows $\boldsymbol{Z}_t(\cdot)$. This proves Prediction 1. □

Proof of Prediction 2. Consider a change in investor i's valuation of service k such that $\Delta\lambda_{it}(k) > 0$. This change raises the wealth-weighted valuation of service k,

$$\Delta\bar{\lambda}_t(k) = \tilde{w}_{i0} \, \Delta\lambda_{it}(k) > 0,$$

holding all other investors' valuations fixed. From equation (5),

$$cy_t(n) = \sum_{l=1}^{K} Z_t^l(n) \, \bar{\lambda}_t(l),$$

so, holding $\bar{\lambda}_t(l)$ fixed for all $l \neq k$,

$$\Delta cy_t(n) = Z_t^k(n)\, \Delta \bar{\lambda}_t(k).$$

Therefore, if $Z_t^k(n) > Z_t^k(m)$ and $\Delta \bar{\lambda}_t(k) > 0$, then $\Delta cy_t(n) > \Delta cy_t(m)$, proving the first statement in Prediction 2.

Next, from equation (6),

$$\boldsymbol{x}_{it} = \boldsymbol{x}_{mt} + \frac{1}{a}\Sigma_t^{-1} \boldsymbol{Z}_t \left(\boldsymbol{\lambda}_{it} - \bar{\boldsymbol{\lambda}}_t\right).$$

Holding $(\boldsymbol{\mu}_t, \Sigma_t, \boldsymbol{Z}_t)$ fixed, a change in $\lambda_{it}(k)$ affects $\boldsymbol{x}_{it}$ through the kth service flow only:

$$\Delta \boldsymbol{x}_{it} = \frac{1}{a}\Sigma_t^{-1} \boldsymbol{Z}_t^k\, \Delta \lambda_{it}(k).$$

Conditional on risk exposure (i.e., within a set of assets for which the risk adjustment through Σ_t^{-1} is the same), higher Z_t^k implies a larger increase in holdings. This proves the second statement in Prediction 2. □

Proof of Prediction 3. Fix an asset n and consider an increase in its service flow k, so that $\Delta Z_t^k(n) > 0$, holding all other characteristics fixed. Let e_n be the nth unit vector in $\mathbb{R}^N$ and e_k the kth unit vector in $\mathbb{R}^K$. Then the change in the characteristics matrix is $\Delta \boldsymbol{Z}_t = \Delta Z_t^k(n)\, e_n e_k^\top$.

From equation (6),

$$\boldsymbol{x}_{it} - \boldsymbol{x}_{mt} = \frac{1}{a}\Sigma_t^{-1} \boldsymbol{Z}_t \left(\boldsymbol{\lambda}_{it} - \bar{\boldsymbol{\lambda}}_t\right),$$

so the induced change in investor i's holdings is

$$\Delta \boldsymbol{x}_{it} = \frac{1}{a}\Sigma_t^{-1} \Delta \boldsymbol{Z}_t \left(\boldsymbol{\lambda}_{it} - \bar{\boldsymbol{\lambda}}_t\right) = \frac{\Delta Z_t^k(n)}{a}\, \Sigma_t^{-1} e_n \left(\lambda_{it}(k) - \bar{\lambda}_t(k)\right).$$

Taking the nth component gives

$$\Delta x_{it}(n) = \frac{\Delta Z_t^k(n)}{a}\, [\Sigma_t^{-1}]_{nn} \left(\lambda_{it}(k) - \bar{\lambda}_t(k)\right).$$

Because Σ_t is positive definite, $[\Sigma_t^{-1}]_{nn} > 0$, and $a > 0$, the response of $x_{it}(n)$ is increasing in $\lambda_{it}(k)$. Hence investors with higher valuation for service k rebalance more strongly toward asset n than investors with lower valuations.

Finally, from equation (5),

$$cy_t(n) = \boldsymbol{Z}_t(n)\bar{\boldsymbol{\lambda}}_t,$$

so holding $\bar{\boldsymbol{\lambda}}_t$ fixed implies

$$\Delta cy_t(n) = \Delta Z_t^k(n)\,\bar{\lambda}_t(k),$$

which is positive when $\bar{\lambda}_t(k) > 0$. This proves Prediction 3. □

Appendix B Construction of CDS-Hedged and Basis-Trade Returns

This appendix provides full implementation details for the construction of CDS-hedged bond returns and basis-trade returns used in the empirical analysis. The goal of these return measures is to isolate the component of bond returns that is not compensation for default risk or interest rate risk, thereby capturing residual returns associated with bond-specific convenience services.

Appendix B.1 Standardized CDS contract mechanics

To hedge default risk, we rely on single-name CDS contracts from Markit. Following the 2009 ISDA "Big Bang" protocol, standardized CDS contracts trade with a fixed running coupon, typically 100 basis points for investment-grade reference entities and 500 basis points for high-yield entities, paid quarterly. Because the fixed coupon generally differs from the market par spread, CDS contracts typically have a non-zero market value at inception.

This difference is settled through an upfront payment $U_{i,t}$, expressed as a fraction of the notional amount. The upfront fee equals the present value of the protection leg minus the present value of the premium leg:

$$U_{i,t} = \text{PV(Protection Leg)} - \text{PV(Premium Leg)}\,. \tag{B.1}$$

A positive upfront fee implies that the protection buyer pays the seller at inception (when the par spread exceeds the fixed coupon), whereas a negative value implies the reverse. As the credit quality of the reference entity evolves, the upfront fee changes over time, generating mark-to-market gains or losses on the CDS position. The running coupon, by contrast, remains fixed for the life of the contract and represents the predictable cost of maintaining the hedge.

Appendix B.2 Interpolation of CDS spreads and upfront fees

The Markit dataset provides a term structure of CDS par spreads and upfront fees across standard maturities (1, 2, 3, 5, 7, 10, 15, 20, and 30 years). For each bond i with remaining

time to maturity $\tau_{i,t}$, we construct maturity-matched CDS quantities by interpolating across available tenors.

Let (τ^L, τ^U) denote the closest lower and upper CDS maturities surrounding $\tau_{i,t}$, with corresponding par spreads (s_t^L, s_t^U). We compute the interpolated CDS spread as

$$s_{i,t}^{\text{CDS}} = s_t^L + \frac{\tau_{i,t} - \tau^L}{\tau^U - \tau^L} \left(s_t^U - s_t^L \right), \tag{B.2}$$

and apply the same procedure to interpolate the upfront fee $U_{i,t}$. When $\tau_{i,t}$ lies outside the available tenor range, we use the closest observed maturity.

Appendix B.3 Synthetic government bond return

To remove interest rate exposure, we construct a synthetic government bond return matched to each corporate bond's modified duration. Using the ECB zero-coupon yield curve, we interpolate the government yield corresponding to duration $D_{i,t}$.

We then compute the implied government bond price and daily return as

$$P_{i,t}^{\text{gov}} = \frac{1}{\left(1 + y_{i,t}^{\text{gov}}\right)^{D_{i,t}}}, \qquad r_{i,t}^{\text{gov}} = \frac{P_{i,t}^{\text{gov}}}{P_{i,t-1}^{\text{gov}}} - 1. \tag{B.3}$$

This construction yields a duration-matched safe return that captures the bond's exposure to interest rate movements.

Appendix B.4 CDS-hedged bond return

Consider an investor who holds one unit of bond i and hedges default risk by purchasing CDS protection on the same notional amount $N_{i,t}$. The beginning-of-day portfolio value is given by the bond's market value plus the mark-to-market value of the CDS position:

$$V_{i,t-1} + N_{i,t-1} U_{i,t-1}. \tag{B.4}$$

Over day t, the investor receives bond cash flows $C_{i,t}$, the CDS mark-to-market value changes through the upfront fee $U_{i,t}$, and the investor pays the daily accrual of the fixed running coupon. Defining the daily premium payment as

$$\pi_{i,t} = (1 + c_{i,t-1})^{1/252} - 1, \tag{B.5}$$

where $c_{i,t-1}$ is the annualized fixed coupon rate on the CDS contract, the CDS-hedged return is

$$r_{i,t}^{\text{hedged}} = \frac{V_{i,t} + C_{i,t} + N_{i,t}\,(U_{i,t} - \pi_{i,t})}{(V_{i,t-1} + N_{i,t-1}U_{i,t-1}) \cdot \frac{N_{i,t}}{N_{i,t-1}}} - 1. \tag{B.6}$$

The numerator represents the end-of-day portfolio value: the bond position (including coupon cash flows) plus the CDS mark-to-market value net of the running premium paid. The denominator is the beginning-of-day portfolio value, adjusted for any changes in the bond's outstanding notional.

When day t is not the first business day of the month, we replace coupon payments with changes in accrued interest to ensure continuity of daily returns.

Appendix B.5 Basis-trade return

Finally, we remove interest rate exposure by subtracting the duration-matched government bond return. The resulting basis-trade return is

$$r_{i,t}^{\text{ls}} = r_{i,t}^{\text{hedged}} - r_{i,t}^{\text{gov}}. \tag{B.7}$$

This long–short return captures the residual compensation from holding the physical bond relative to its synthetic equivalent. Under a first-order approximation, it is closely linked to the CDS–bond basis:

$$r_{i,t}^{\text{ls}} \approx -\frac{b_{i,t-1}}{252} + \Delta b_{i,t}\, D_{i,t}, \tag{B.8}$$

where $b_{i,t} = s_{i,t}^{\text{CDS}} - s_{i,t}^{\text{credit}}$ denotes the CDS–bond basis. The first term represents the daily carry from holding a bond with a negative basis, while the second term captures mark-to-market variation as the basis changes over time.

Appendix B.6 Sector-level aggregation

We aggregate bond-level basis-trade returns into sector portfolios using value weights based on each sector's bond holdings. The cross-section of these portfolio returns across investor sectors is informative about the distribution of convenience yields and the value of bond-specific service flows in equilibrium.

Appendix C Extra Descriptive Information

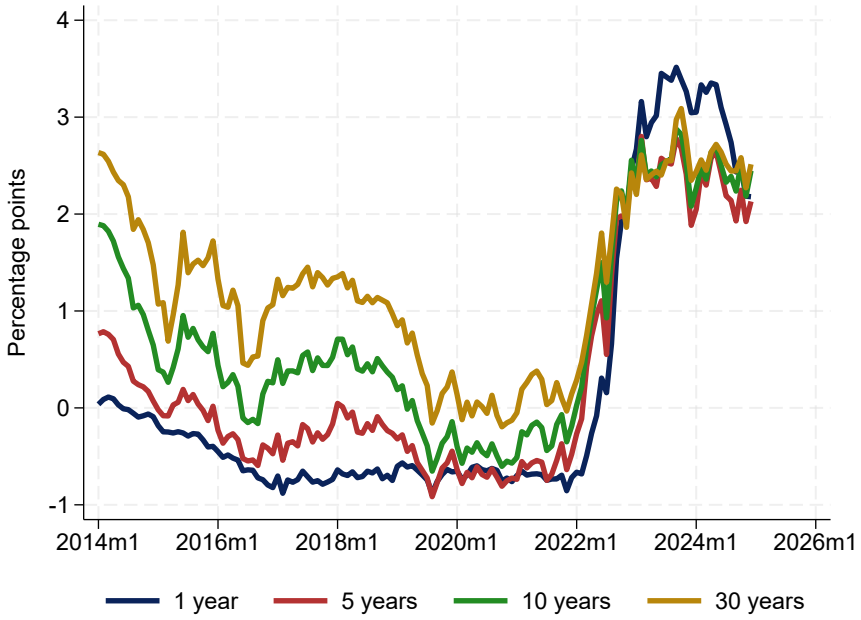


Figure C.1: ECB benchmark interest rates for different maturities. This plot presents the time series of the benchmark interest rate for 1Y, 5Y,10Y, and 30Y maturity. The ECB benchmark interest rate is provided under the euro area yield curves estimation and is used as an approximate measure of risk-free rate in the euro area. It is estimated daily by using the AAA-rated euro area central government bonds.

Figure C.2: Average service content by investor sector over time (only sovereign bonds). This plot shows the average service content of each investor sectors' sovereign bond portfolio over time for four services—liquidity value, collateral value, risk-based capital value (RBC), and modified duration. Panel (a) shows the average bid-ask spread, measured as $\frac{ask-bid}{bid}$; collateral value in panel (b) is measured as $1 - haircut \in [0, 1]$; RBC value in panel (c) is measured as $1 - stress \in [0, 1]$, where the $stress$ parameter is defined in Appendix E (note that, since EA-issued sovereign bonds all have a $stress$ parameter of zero, the average RBC value of sovereign bonds trivially collapses to 1); Panel (d) shows the average modified duration of each sector's bond portfolio. To calculate the sectoral averages, bond-level observations have been weighted by the face value held. The sample period is 2011Q4–2025Q2.

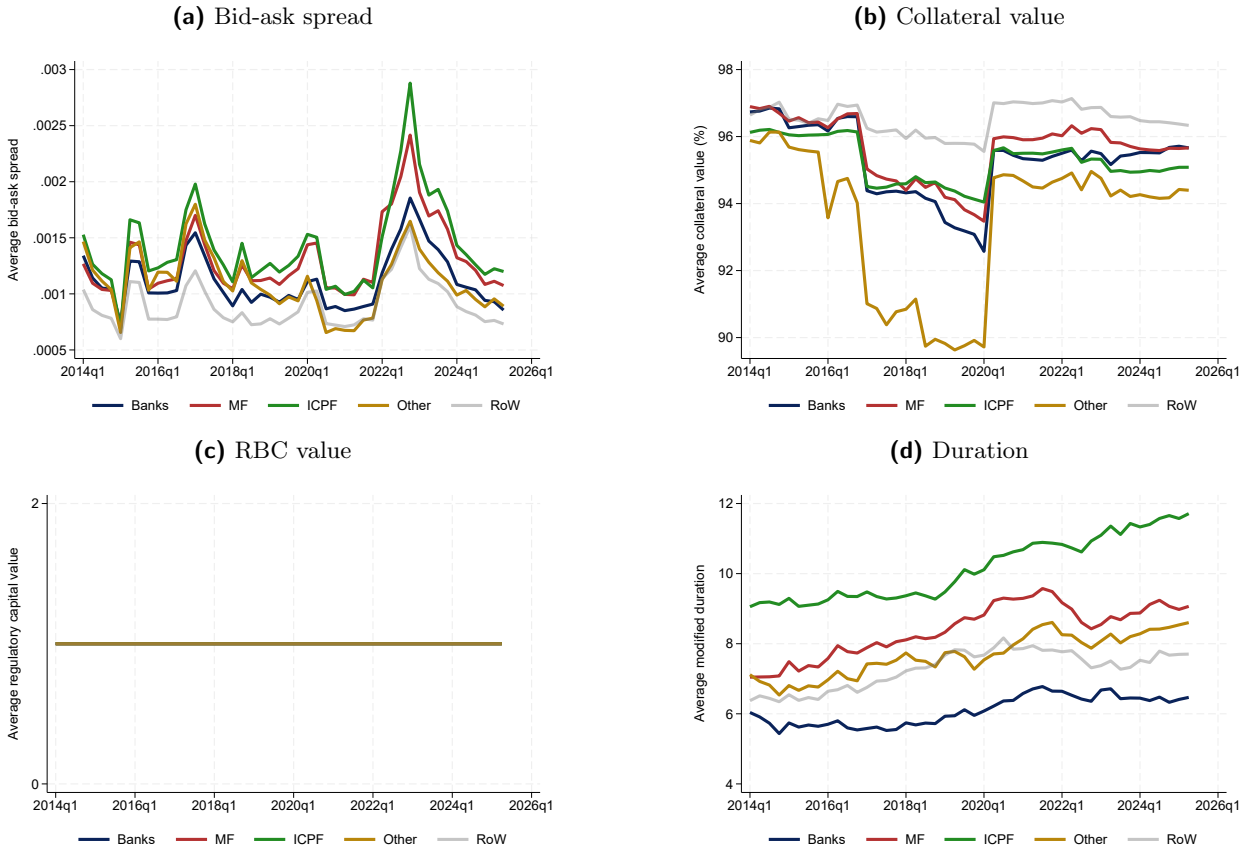
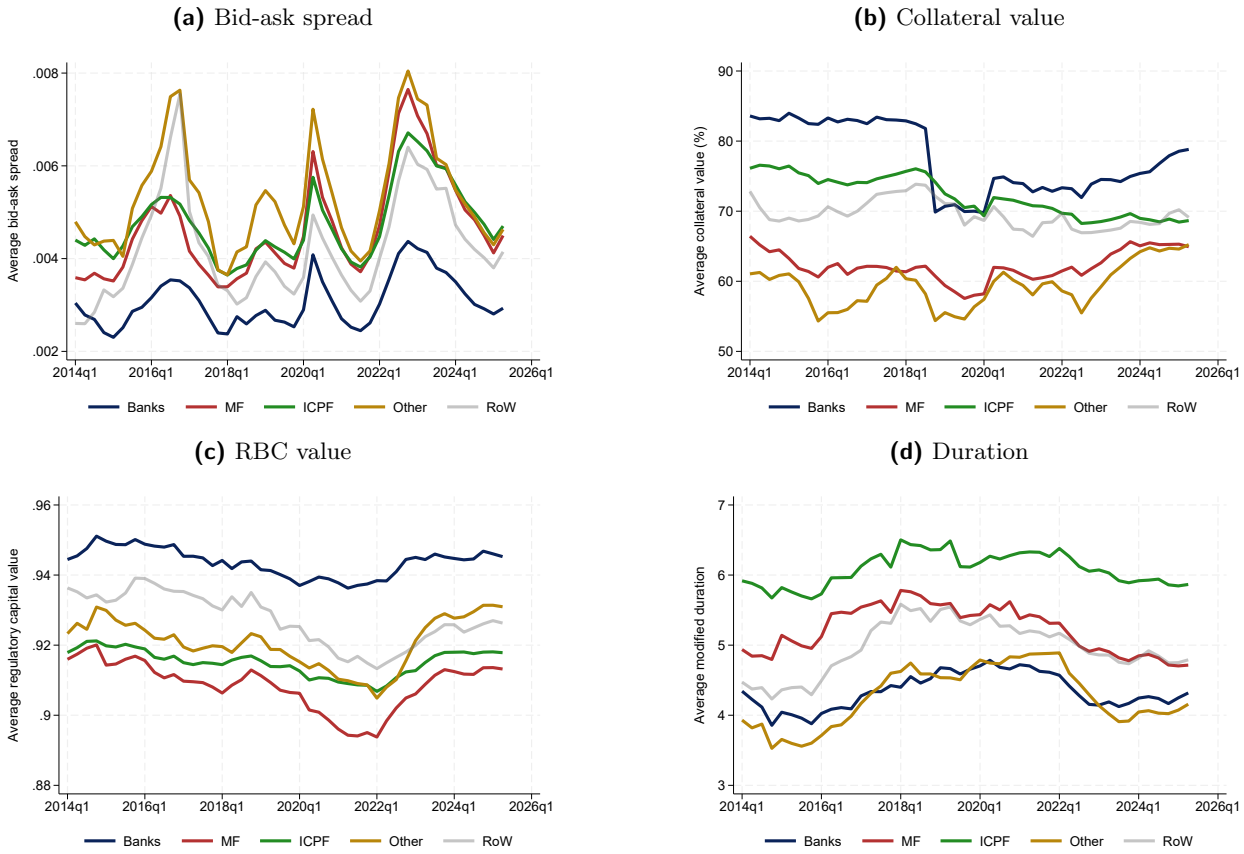


Figure C.3: Average service content by investor sector over time (only corporate bonds). This plot shows the average service content of each investor sectors' corporate bond portfolio over time for four services—liquidity value, collateral value, risk-based capital value (RBC), and modified duration. Panel (a) shows the average bid-ask spread, measured as $\frac{ask-bid}{bid}$; collateral value in panel (b) is measured as $1-haircut \in [0,1]$; RBC value in panel (c) is measured as $1-stress \in [0,1]$, where the *stress* parameter is defined in Appendix E. Panel (d) shows the average modified duration of each sector's bond portfolio. To calculate the sectoral averages, bond-level observations have been weighted by the face value held. The sample period is 2011Q4–2025Q2.

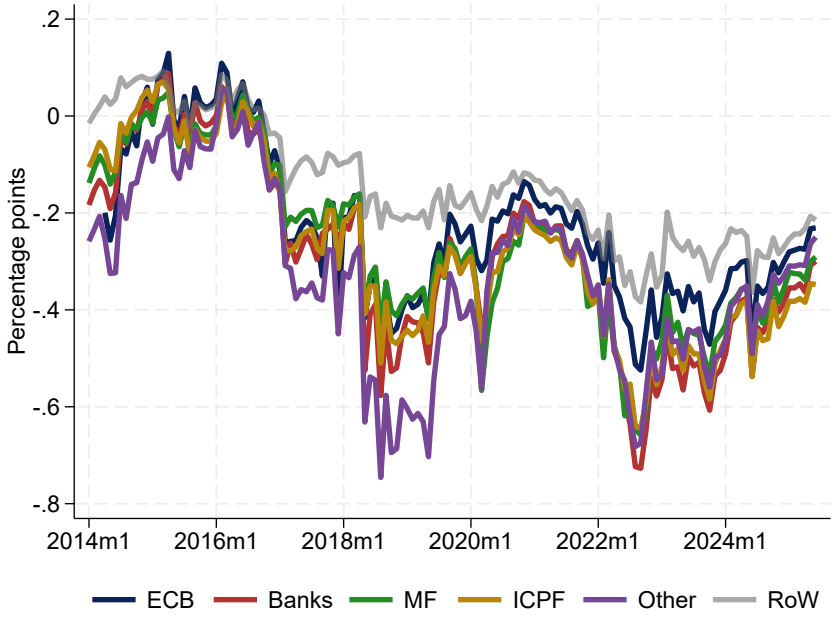


Appendix D Heterogeneity in convenience yields across investors

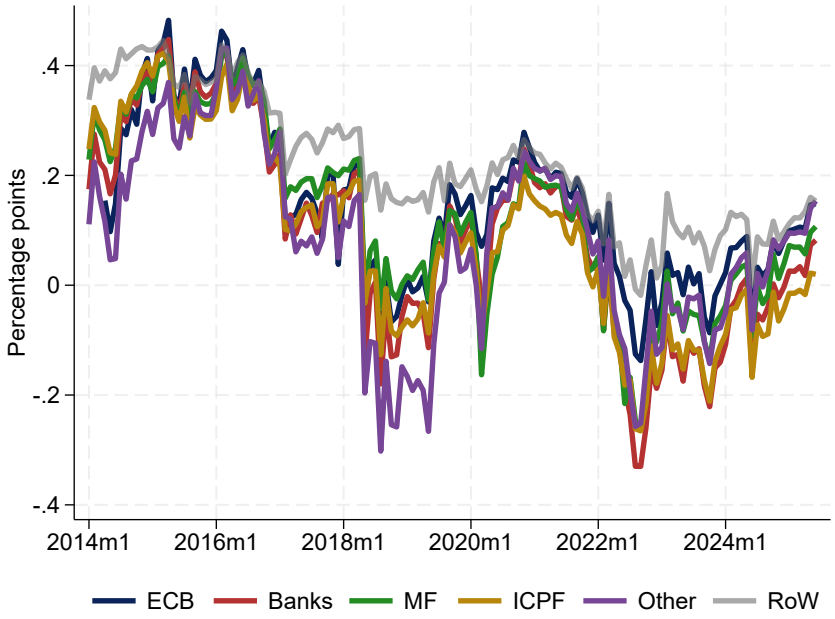
Table D.1 and Figure D.1 contain the average CDS-bond basis for each investor sector's bond portfolio.

Table D.1: CDS-bond basis by investor sector This table contains distributional summary statistics of the CDS-bond basis (calculated by subtracting a bond's credit spread from its CDS spread), separately for each investor type's portfolio. All moments are obtained by weighting observations by nominal holdings.

	Mean	SD	Min	p10	p25	p50	p75	p90	Max
ECB	-0.28	0.36	-2.89	-0.75	-0.45	-0.21	-0.03	0.10	2.68
Banks	-0.29	0.39	-3.17	-0.76	-0.47	-0.24	-0.05	0.11	2.74
ICPF	-0.29	0.39	-3.07	-0.77	-0.51	-0.24	-0.03	0.14	2.74
MF	-0.27	0.40	-3.14	-0.76	-0.48	-0.23	0.00	0.15	2.52
Other	-0.32	0.45	-3.17	-0.95	-0.58	-0.23	0.02	0.14	2.74
RoW	-0.15	0.33	-3.17	-0.57	-0.29	-0.07	0.06	0.17	2.74



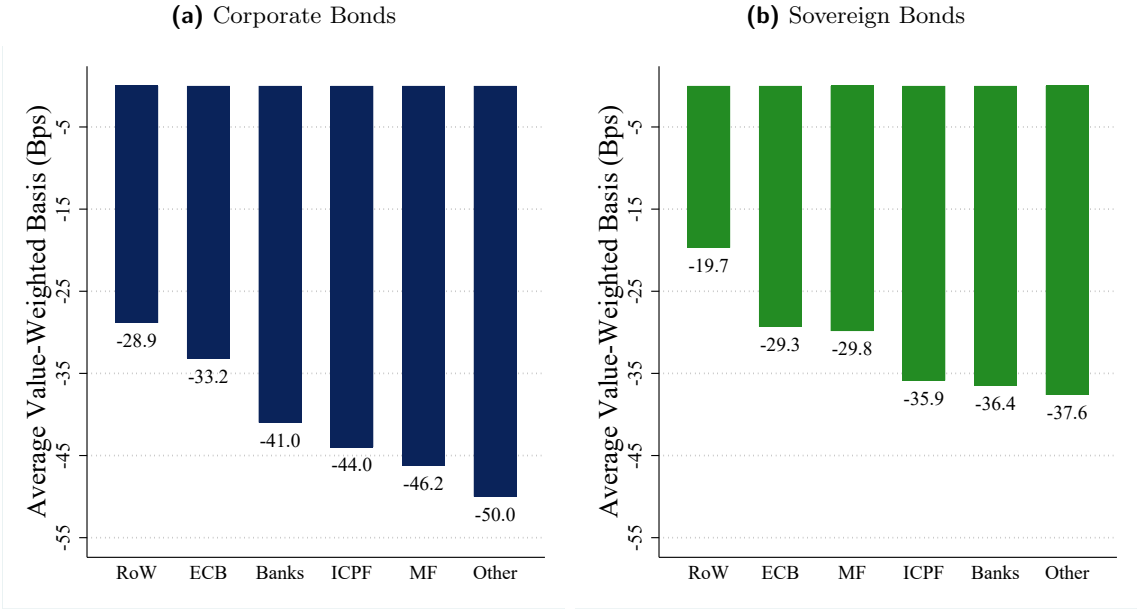
(a) CDS-bond basis of portfolios of different investor classes



(b) CDS-bond basis of portfolios of different investor classes after absorbing Rating x Maturity FE

Figure D.1: Value-weighted average CDS-bond basis of each investor sector's bond portfolio Panel (a) of this plot presents the average CDS-bond basis (weighted by nominal holdings) of each investor sector's bond portfolio. Panel (b) plots the average residual from a regression of the CDS-bond basis on rating x maturity dummies for each investor sector, again weighted by nominal holdings.

Figure D.2: Average CDS-Bond Basis by Sector (CSDB). The figure reports the average value-weighted CDS-bond basis for each investor sector. Panel (a) shows results for corporate bonds, while Panel (b) shows results for sovereign bonds. We compute weights using each sector's portfolio shares, based on face value, within the sovereign and corporate bond universes, respectively.



Appendix E Solvency II Details

The table below shows the stress parameter calculation according to Article 176 of the Solvency II Directive 2009/138/EC.

Furthermore, Article 180 paragraph 2 determines specific stress calculations for certain bonds. In particular, exposures in the form of bonds and loans to the following shall be assigned a risk factor $stress_i$ of 0 %:

(a) the European Central Bank;

(b) Member States' central government and central banks denominated and funded in the domestic currency of that central government and the central bank;

(c) multilateral development banks referred to in paragraph 2 of Article 117 of Regulation (EU) No 575/2013;

(d) international organisations referred to in Article 118 of Regulation (EU) No 575/2013.

Table E.1: Credit Quality Step Stress Parameters by Duration

Duration (dur_i)	$stress_i$	Credit quality step											
		0		1		2		3		4		5 and 6	
		a_i	b_i	a_i	b_i	a_i	b_i	a_i	b_i	a_i	b_i	a_i	b_i
up to 5	$b_i \cdot dur_i$	—	0.9%	—	1.1%	—	1.4%	—	2.5%	—	4.5%	—	7.5%
More than 5 and up to 10	$a_i + b_i \cdot (dur_i - 5)$	4.5%	0.5%	5.5%	0.6%	7.0%	0.7%	12.5%	1.5%	22.5%	2.5%	37.5%	4.2%
More than 10 and up to 15	$a_i + b_i \cdot (dur_i - 10)$	7.0%	0.5%	8.4%	0.5%	10.5%	0.5%	20.0%	1.0%	35.0%	1.8%	58.5%	0.5%
More than 15 and up to 20	$a_i + b_i \cdot (dur_i - 15)$	9.5%	0.5%	10.9%	0.5%	13.0%	0.5%	25.0%	1.0%	44.0%	0.5%	61.0%	0.5%
More than 20	$\min[a_i + b_i \cdot (dur_i - 20); 1]$	12.0%	0.5%	13.4%	0.5%	15.5%	0.5%	30.0%	0.5%	46.5%	0.5%	63.5%	0.5%

Furthermore, exposures in the form of bonds and loans that are fully, unconditionally, and irrevocably guaranteed by one of the counterparties mentioned in points (a) to (d), where the guarantee meets the requirements set out in Article 215, shall also be assigned a risk factor stress of 0%.

Appendix F ECB Collateral Eligibility

Appendix F.1 A snapshot of ECB haircuts as 06/29/2023

Table F.1: ECB haircuts for different bond categories. This table contains the haircuts that the ECB applies to the market value of securities pledged as collateral. The maximum amount a bank can borrow against a given amount of collateral is (100-haircut)% times the market value of the pledged security. The collateral framework is periodically reviewed and changes to the haircut schedule are not unusual. Per Guideline ECB/2022/49 the schedule below is effective as of 29 June 2023.

Rating	Residual Maturity	Central government		Local gov. + covered bonds		NFC		FC	
		Fix/Float	Zero	Fix/Float	Zero	Fix/Float	Zero	Fix/Float	Zero
AAA-A	0-1y	0.5	0.5	1.0	1.0	1.0	1.0	7.5	7.5
	1-3y	1.0	2.0	1.5	2.5	2.0	3.0	10.0	11.5
	3-5y	1.5	2.5	2.5	3.5	3.0	4.5	12.0	13.0
	5-7y	2.0	3.0	3.5	4.5	4.5	6.0	14.0	15.0
	7-10y	3.0	4.0	4.5	6.5	6.0	8.0	16.0	17.5
	10-15y	4.0	5.0	6.5	8.5	7.5	10.0	18.0	22.5
	15-30y	5.0	6.0	8.0	11.5	9.0	13.0	21.0	25.0
	>30y	6.0	9.0	10.0	13.0	11.0	16.0	24.0	31.5
BBB	0-1y	5.0	5.0	5.5	5.5	6.5	6.5	11.5	11.5
	1-3y	6.0	7.0	7.5	10.5	9.5	12.0	18.5	20.0
	3-5y	8.5	10.0	11.0	16.0	13.0	18.0	23.0	27.0
	5-7y	10.0	11.5	12.5	17.0	15.0	21.5	25.5	29.5
	7-10y	11.5	13.0	14.0	21.0	17.0	23.5	26.5	31.5
	10-15y	12.5	14.0	17.0	25.5	19.5	28.0	28.5	35.0
	15-30y	13.5	15.0	20.0	28.5	22.0	31.0	31.5	39.0
	>30y	14.0	17.0	22.0	32.5	25.0	35.5	34.5	43.0

Appendix F.2 The Change of Eligibility of Subordinated Unsecured Bank Bonds

The ultimate cause of the change in ECB collateral eligibility of senior unsecured bank bonds in December 2018 is the Directive of the European Union 2014/59/EU, also called the Bank Recovery and Resolution Directive (BRRD). Among other things, the directive required banks to increase their total loss-absorbing capital (TLAC) and comply with the new minimum requirement for own funds and eligible liabilities (MREL), so that a significant share of liabilities could be written down before public funds be used in the recovery of a

failing bank. In other words, banks had to issue more "bail-in-able" securities that are subordinated to deposits, covered liabilities, and other senior liabilities such as wages and taxes.

Importantly, the directive allowed three ways of implementing the new rules, leaving it open to member states to decide which approach they preferred:

- *Contractual subordination*, where the subordination is explicitly stated in the bond's terms. This effectively amounts to the creation of an entirely new asset class, sometimes called *senior non-preferred* liabilities. France and Spain were among the early adopters of this option.

- *Structural subordination*, which applies only to banks with a holding structure: Debt issued by a holding company is automatically structurally subordinated to debt issued by the operating company. This applied mostly to banks from the UK, Belgium, Switzerland, and the Netherlands.

- *Statutory subordination*, where all senior bank bonds would be subordinated by law, through a modification of national insolvency laws. Germany was the only member state to take this route.

Regardless of which approach was taken, subordinated uncovered bank bonds (sUBB), even if called "senior", are more risky than plain-vanilla senior bank debt by construction. Any exposure to potential bail-in losses during a bank's resolution could jeopardize the ECB's authority and even create conflicts of interest. This is what motivated the ECB to reconsider its collateral framework and, in particular, which bank bonds it should accept as collateral going forward. The most immediate solution would have been to simply stop accepting all subordinated UBB. However, this would have put German banks at a significant disadvantage. While, for example, French banks could have issued both senior preferred and senior non-preferred UBB (the former of which would have been eligible for collateral), German banks could not have issued any ECB-eligible UBB anymore because of statutory subordination.

That is why, on October 5$^{\text{th}}$, 2016 the ECB announced to *maintain* the eligibility of statutorily or structurally subordinated UBB for the time being until member states would

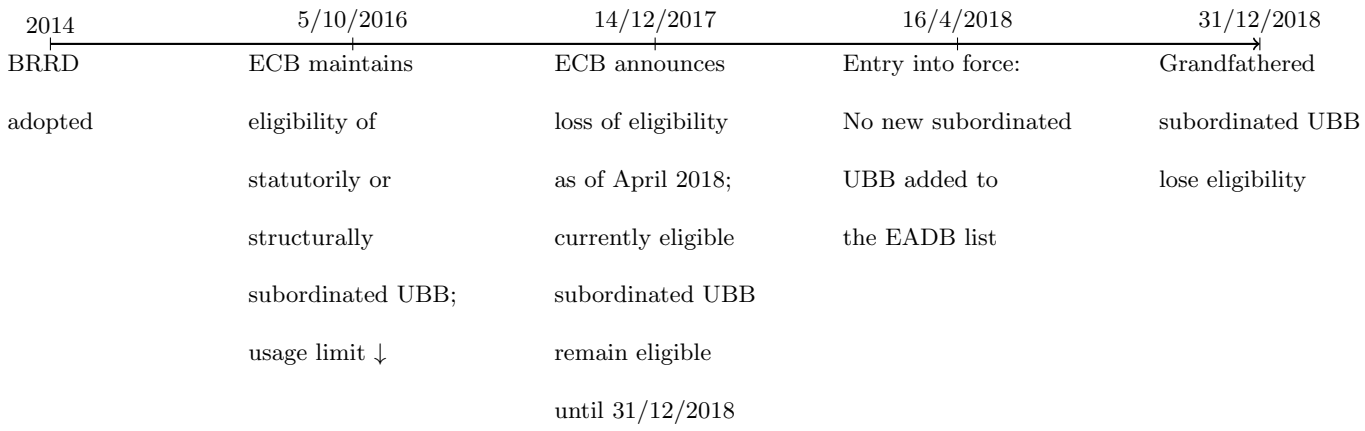


Figure F.1: Timeline of events

converge to a harmonized approach. Without this temporary waiver, the sUBB in question would have become ineligible as of January 1st, 2017. Obviously, the main beneficiaries of this decision were German banks and bank holding companies which could now continue to issue ECB-eligible sUBB which, at the same time, counted towards the issuer's MREL requirement. In other words, for a while, German sUBB enjoyed the best of both worlds.[36]

However, on December 14th, 2017, the ECB announced the end of this exception: As of April 2018, no new subordinated UBB would be eligible for collateral anymore (with the exception of government-guaranteed bonds). To ease the transition, the ECB adopted a grandfathering approach for those subordinated UBB that were still eligible at the time of the announcement. Those securities would maintain their eligibility until December 31st, 2018. We summarize the most important events in Figure F.1.

Eventually, to level the playing field, the German Banking Act (Kreditwesengesetz) was amended to adopt the contractual subordination option as well.[37] This allowed German banks to issue both senior preferred and senior non-preferred bonds as of July 21st, 2018.

As further evidence of our interpretation of ECB collateral eligibility as a service that banks value, Figure F.2 plots the average CDS-bond basis of senior preferred vs. senior non-preferred bonds. Note that the series only starts in 2018, after senior non-preferred

[36] However, the ECB lowered the usage limit of sUBB from 5% to 2.5%, i.e. the value of sUBB pledged as collateral by a bank could not exceed 2.5% of its total pledged collateral.

[37] See section 46f paragraphs 6, 7, and 9 KWG.

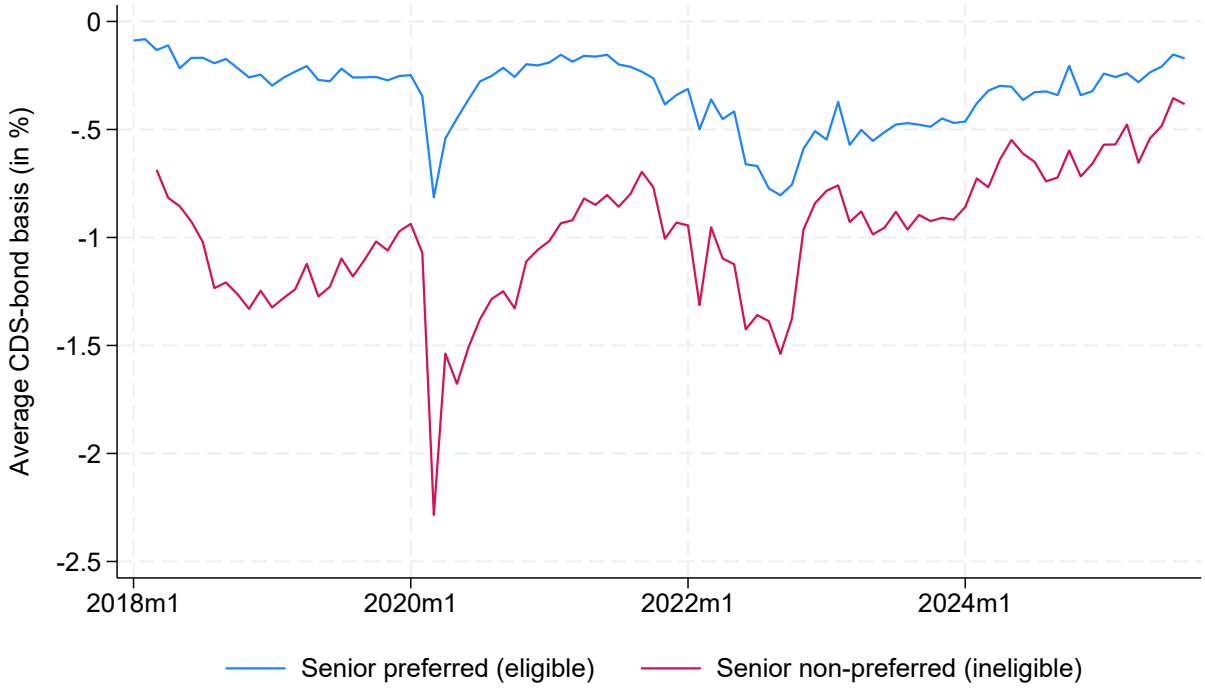


Figure F.2: Average CDS-bond basis of bank bonds of different seniority. This plot displays the average CDS-bond basis of senior preferred bank bonds (blue) vs. that of senior non-preferred bank bonds (pink). The former are generally collateral-eligible, the latter are not. To obtain the average basis, observations are weighted by outstanding amounts. Note that because the two debt classes have different seniority, the CDS spreads used to calculate the CDS-bond basis are also different.

bonds and, even more importantly, their corresponding CDS contracts, had been widely introduced. The plot clearly shows that senior preferred bank bonds (which are generally eligible for collateral at the ECB) have a higher CDS-bond basis than (ineligible) senior non-preferred bank bonds, even after we account for the different default risk of each asset class (by using the seniority-appropriate CDS spread when calculating the CDS-bond basis).