

Who Bears the Duration?

Monetary Policy Implementation and the Allocation of Interest Rate Risk between the Private and Public Sectors

Philip Kroos*

August 2026

Abstract

That quantitative easing was in effect a fixed-for-floating swap, and that it would cost the Eurosystem when rates rose, was shown by Belhocine et al. [2023]. This paper does not claim the mechanism. It measures the composition that produced it, for all 21 national central banks monthly from June 2016 to June 2026, and reports the outturn. Four findings follow. The share of monetary policy assets held as secured, floating-rate credit fell from 37.2 per cent in March 2021 to 0.80 per cent in June 2026, widening the repricing gap, defined as reserves less secured credit, from EUR 626 billion in 2016 to EUR 3,146 billion in 2024. Capital subsequently fell for three consecutive years, by EUR 76 billion in total, extending beyond the period in which their baseline anticipated roughly two years of losses, and leaving four national central banks in negative equity. A counterfactual bounded by demonstrated collateral capacity puts the fiscal value of the composition margin between nothing and two fifths of the realised funding cost. And matching acquisition-month sovereign yields for Germany, France, Italy and Spain against the realised deposit facility rate, on windows that have actually been realised, gives a stock-weighted carry of -16 basis points at a four-year horizon, where every vintage is complete: roughly break-even, and decomposing against a random-walk benchmark into $+97$ basis points of ex-ante yield compensation less a $+113$ basis point rate surprise. The paper revises the TARGET debate accordingly: Sinn and Wollmershäuser [2012] was right that the asset side determines what the Bundesbank's claim is worth, but what materialised was interest rate risk, not the credit risk they named, and these data cannot cleanly attribute national losses to national balance sheets. A digital euro bears on the same margin, and its income effect depends jointly on remuneration and on the source of take-up: it narrows the repricing gap when the remuneration share plus the cash-substitution share is below one, is neutral at equality, and widens it above.

Keywords: central bank capital, duration risk, monetary policy implementation, TARGET, digital euro. **JEL:** E42, E52, E58, F33, G21

*philip.kroos@web.de. Replication archive: github.com/Philip-Kroos/collateralised-share. All data are published by the ECB; no proprietary source is used. All errors are my own.

1 Introduction

When a central bank expands its balance sheet, it chooses an instrument. It can lend against collateral or it can buy securities outright. Both create reserves and both are recorded the same way in any account of the policy stance. They differ in one respect that no such account records: lending leaves duration risk where it was, while buying transfers it from private portfolios to the public sector. The choice of instrument is therefore an allocation of interest rate risk between the private and public sectors, made continuously and, so far as the record shows, rarely discussed as one.

This paper asks how large that allocation was, what it cost, and whether it had to be made the way it was. The euro area is the laboratory because it is where the consequences have been realised first and most visibly, and because its disaggregated balance sheet data allow the allocation to be measured rather than inferred. But the question is not specific to it: the Federal Reserve, the Bank of England and the Swiss National Bank made the same choice with the same mechanical consequence and differ only in the arrangements for absorbing the result, which Section 9 compares.

The euro area outcome. Between 2022 and 2026 the Eurosystem lost three quarters of its capital. Aggregate national central bank capital and reserves fell from EUR 120 billion to EUR 39 billion, and four national central banks now report negative equity: Austria, Belgium, Germany and Slovakia. The Bundesbank, at minus EUR 25 billion, is the largest creditor in the system.

Not one debtor defaulted.

That the losses would come is not news. Belhocine et al. [2023] set it out in July 2023: quantitative easing removed duration risk from private balance sheets and financed the purchases with remunerated reserves, so that the Eurosystem had, in their words, executed a fixed-for-floating rate swap and would lose money when rates rose. They projected net income for the five largest national central banks over a ten-year horizon and concluded that the losses would be temporary and recoupable, with about two years of loss-making.

This paper does not claim that mechanism. What it does instead is threefold. It *measures* the composition that produced the exposure rather than projecting its consequences: the collateralised share and the repricing gap are constructed series for all 21 national central banks, monthly, from June 2016 to June 2026. It reports the *outturn* against the projection, which only the passage of time could supply. And it asks whether the composition was a *choice*, which Belhocine et al. [2023] do not, treating the swap as a feature of balance sheet policy rather than as a margin that could have been set differently.

That combination is the puzzle this paper addresses. For more than a decade Sinn and Wollmershäuser [2012] has argued that the Bundesbank's claim on the Eurosystem is a credit position whose value depends on the solvency of the debtors, and that no member state is obliged to recapitalise its national central bank. Sauer [2023] argues more generally that the value of central bank money depends on the quality of the assets standing behind it. Both are arguments about the asset side, and the losses of the last four years

look like their vindication. But the losses did not arrive through the channel either identified. There was no default. There was a change in the composition of the asset side, and that change had a price.

A central bank issues money by acquiring assets. The money is a liability; the assets stand against it. That much is accounting. What is contested is whether the composition of those assets matters, and if so for what.

Two arguments say it does, and they are usually treated as separate. Sinn and Wollmershäuser [2012] and Sinn [2014] argue that the Bundesbank's TARGET claim on the Eurosystem is a credit position whose value depends on the solvency of the debtors, and that no member state is obliged to recapitalise its national central bank. Sauer [2023] argues, from the interwar record, that the value of central bank money depends on the quality of the assets backing it, and constructs a solvency measure to make the claim operational. Both arguments are, at bottom, about the asset side. Both are usually conducted at the level of the consolidated Eurosystem, which is precisely the perimeter in which the relevant positions net out.

This paper takes their shared empirical premise and tests it with the data that can. The ECB publishes national central bank contributions to the Eurosystem balance sheet, by country and month. I use them to ask what actually backs central bank money in the euro area, how that has changed, and what a digital euro would do to it.

The distinction that organises everything. Eurosystem monetary policy assets come in two forms with entirely different risk architecture. Credit operations are secured: a counterparty posts collateral, the collateral is valued daily, a haircut is applied against adverse price movement, and a fall in value triggers a margin call. Securities held outright have none of this. There is no counterparty, no haircut, no daily valuation against a posted pool, and no call. Article 18.1 of the ESCB Statute requires lending to be against adequate collateral; it says nothing of the sort about outright purchases. Nyborg [2017] sets out how the Eurosystem's collateral framework prices that protection, and Drechsler et al. [2016] and Acharya et al. [2019] document who uses it and to what effect.

The ratio of the first to the sum of both is not an accounting curiosity. It is the share of the asset side that carries a loss-absorbing buffer negotiated in advance with an identifiable counterparty. I call it the collateralised share.

That distinction has a second reading, and it is a different property of the same instruments rather than the same property named twice. A haircut and a margin call protect against the counterparty failing and the collateral falling in value; a floating coupon protects against the asset repricing more slowly than the liability. Credit protection and repricing protection are distinct, and secured lending happens to supply both. The paper's central object is the second. Collateral enters only as the reason the two coincide in the euro area data and as a separate margin in Section 10.

The two measures are also distinct algebraically. The collateralised share is $L/(L+S)$ and the repricing gap is $R-L$; they move together only when $R \approx L+S$, which holds approximately in this sample because other balance sheet items are small relative to monetary policy assets, and would not hold in general.

The second reading is what gives the composition shift economic content. Credit operations are *floating rate*: they are rolled at the policy rate, so the asset reprices the moment the rate changes. Securities bought outright are *fixed rate*, acquired at the yields prevailing when they were bought and held to maturity. Reserves, on the liability side, reprice immediately and in full.

A shift from secured lending to outright holdings is therefore not only a shift from collateralised to uncollateralised. It is a shift from a rate-matched balance sheet to a duration-mismatched one. The exposure is the volume of liabilities that reprices without a matching asset,

$$\text{repricing gap} = \text{reserves} - \text{secured credit}, \quad (1)$$

and a rate rise of one percentage point costs that gap, per year, at once. This is the paper's central object and Section 7 is about it.

Three facts. Section 6 documents them.

The collateralised share fell from 32.8 per cent in June 2016 to a peak of 37.2 per cent in March 2021, at the height of TLTRO III, and then to 0.80 per cent in June 2026. The euro area's central bank money is now backed by EUR 3,142 billion of securities held outright against EUR 25 billion of secured credit. The collapse is not a slow drift: it happened between mid-2023 and mid-2024, as the targeted operations were repaid and the purchase portfolios ran off more slowly.

Capital and reserves across national central banks fell from 4.97 to 1.00 per cent of monetary liabilities over the same decade. Four national central banks report negative capital in June 2026, including the Bundesbank at minus EUR 25 billion, having been at plus EUR 5.7 billion as recently as December 2022.

The Bundesbank's gross intra-Eurosystem claim is EUR 1,068 billion, close to the TARGET figure usually quoted. It is offset by a EUR 605 billion intra-Eurosystem liability, arising largely because Germany puts more banknotes into circulation than its capital key share and owes the system the difference. The net position is EUR 463 billion, 43 per cent of the gross. What stands behind it on the debtor side, however, is 65 per cent securities held outright and 0.8 per cent secured credit.

The model. Section 4 asks what follows if the composition is a decision rather than an accident. Under the assumption that the two instruments differ only in repricing and duration, net interest income depends on the purchase leg alone and not at all on the size of the expansion (Proposition 1). The purchase programme is then a carry trade: the central bank is long the term premium and short the rate surprise, and it breaks even if and only if the realised short-rate surprise does not exceed the ex-ante term premium (Proposition 4). That reframes the losses. They are not an error of composition but the payout on insurance the public sector sold, and the empirical question is how the purchase leg performed ex post and what produced that outcome, which Section 8.5 evaluates.

Two further results speak directly to the data. Composition is a policy margin invisible to any description of the stance in terms of reserves: two splits with identical reserve creation differ in income variance by a factor of nine at the model's calibration (Proposition 3). And when a three-year lending operation and an eight-year portfolio are allowed to mature side by side, the balance sheet shrinks while the exposure per euro of it rises towards one, with no decision taken (Proposition 5). That is exactly the pattern the euro area data show after 2023.

The mechanism. Section 7 connects the composition to the losses. The repricing gap widened from EUR 626 billion in 2016 to EUR 3,146 billion in 2024, a fivefold increase, driven first by the growth of reserves and then by the disappearance of the secured leg that had offset them. Funding that gap at the deposit facility rate implies a carry cost of EUR 108 billion in 2023 and EUR 115 billion in 2024.

The implied carry leads the realised change in capital by one year. The correlation is 0.55 contemporaneously and 0.95 at a one-year lag, which is what provisioning implies: carry costs reach the profit and loss account first and are absorbed by risk buffers before they touch capital. Reading the contemporaneous correlation alone would understate the relationship by almost half.

This reframes rather than refutes the argument the paper begins with. Sinn and Wollmershäuser [2012] were right that the asset side determines what the Bundesbank's claim is worth, and right that a risk was accumulating that the debate was not pricing. The risk that materialised was interest rate risk on an outright portfolio, not credit risk on a collateralised loan, and Section 8.5 shows that the compensation available at acquisition was of similar size to the tightening that followed. On his own criterion the position is weaker than when they wrote, because what now stands behind the claim has neither a haircut nor a margin call nor a floating coupon.

What a digital euro would change. Section 10 turns to the question that motivated the exercise. The intuition that a digital euro creates unbacked money does not survive the second leg of the transaction. Conversion is a swap: the household gives up a deposit and receives a claim on the Eurosystem, and the bank that lost the deposit replaces the funding by borrowing. That borrowing is secured, haircut and margined. Broad money is unchanged; what changes is the composition of the central bank's assets, and it changes in the direction of more collateral, not less.

The magnitude follows from the holding limit with one behavioural parameter, the share β of the outflow banks meet from their own liquid assets rather than by borrowing. At the EUR 3,000 limit under discussion and $\beta = 0.3$, a fully taken up digital euro would raise the collateralised share from 0.80 to 16.6 per cent. Restoring the March 2021 level would require a limit of EUR 9,185 per holder, which is not on any agenda; reaching ten per cent would require EUR 1,623.

The sharper point concerns duration rather than collateral, and it is conditional. An unremunerated digital euro taken up out of deposits closes the repricing gap of (1) by the whole amount converted: at a EUR 3,000 limit, 36.5 per cent of the gap, worth EUR 17.1 billion of annual carry at a two per cent policy rate.

More generally, remuneration and cash substitution each attenuate the gain; it vanishes when their shares sum to one and reverses above that threshold.

The qualification, and where backing does fail. The rebuild is mechanical only if the collateral is there. Alexopoulou et al. [2025] document that the Additional Credit Claims frameworks, the emergency loosening of collateral standards that allowed banks to pledge loan portfolios, peaked near EUR 500 billion of post-haircut value during the pandemic. That temporary framework is being phased out while part of its credit-claim capacity migrates into the permanent one, so what matters is how much can be mobilised under the post-transition framework, and how quickly. Section 11 is explicit that this, not money creation, is where the concern belongs.

Contribution. Section 2 sets out the TARGET debate and the case on both sides at length. In brief, the paper accepts three of the four steps in Sinn and Wollmershäuser [2012] and revises the fourth: the exposure is real and the absence of a recapitalisation obligation matters, but the risk that materialised was not the credit risk they named. To the literature on central bank capital and its fiscal backing [Buiter, 2008, Bassetto and Messer, 2013, Reis, 2013, Del Negro and Sims, 2015, Hall and Reis, 2015] it supplies the euro area quantities at national level. To the CBDC literature [Bidder et al., 2025, Burlon et al., 2024, Abad et al., 2025] it adds the asset-composition consequence, which that literature does not address because it treats the central bank as a single consolidated balance sheet.

The paper's own contribution is the link between the three. The collateralised share and the repricing gap are two complementary readings of the same shift in instrument composition, and the second converts a descriptive fact into an accounting identity with a price attached. That is what allows the losses of 2023 to 2026 to be attributed without invoking a default that did not occur.

The paper reports facts, one arithmetic counterfactual, and an empirical evaluation of the counterpart to Proposition 4 in Section 8.5. It does not estimate a causal effect from the balance-sheet regressions and does not claim one.

Section 2 sets out the TARGET debate and the literatures the paper joins. Section 3 gives the accounting framework and Section 4 the model of instrument choice. Section 5 describes the data, Section 6 the three balance-sheet facts, Section 7 the repricing-gap mechanism, and Section 8 the estimation and the identification design. Section 9 turns to fiscal backing, including the premium test in Section 8.5 and the cross-country comparison of loss-absorption arrangements. Section 10 takes up the digital euro and Section 11 the limitations.

2 The TARGET debate, and where this paper sits

The paper claims to reframe an argument that has run for more than a decade. It owes the reader that argument in enough detail to judge whether the reframing is fair, and the case against it.

2.1 Sinn's argument, in four steps

Sinn and Wollmershäuser [2012] and Sinn [2014] make a claim with four parts. Separating them matters, because this paper accepts three and revises the fourth.

First, TARGET balances are not a technical artefact. When a payment moves from a bank in one member state to a bank in another, the sending national central bank incurs a liability to the Eurosystem and the receiving one acquires a claim. If the flow persists, so do the balances. In the crisis they persisted because peripheral banks financed deposit outflows with central bank credit, and the counterpart accumulated in Frankfurt.

Second, the resulting claim is credit. The Bundesbank holds a position against the Eurosystem that bears the policy rate, has no maturity and cannot be called. The characterisation of this as an open-ended credit line has been contested on the language but not successfully on the substance: it is an asset whose realisation depends on the continued functioning of the system that owes it.

Third, and this step carries the weight, there is no recapitalisation obligation. No member state is required to recapitalise its national central bank. A creditor's claim is therefore backed in the first instance by what the debtor central banks' own assets can support, rather than by an enforceable claim on their sovereigns. The legal premise is correct; whether it exhausts the sources of backing is a separate question, and Section 9 returns to future seigniorage and to monetary income pooling as the two that remain. Section 6 makes the balance sheet side of it quantitative.

Fourth, the risk is credit risk: the danger is that a debtor cannot pay, which in practice means exit or a default impairing its central bank's assets. The remedies that follow are collateral standards, haircuts, limits and explicit risk sharing, and that is what the policy debate has argued about since.

This paper accepts the first three and revises the fourth. Between 2022 and 2026 no debtor defaulted, no member state left, and the Bundesbank's capital nonetheless fell from plus EUR 5.7 billion to minus EUR 25.3 billion. Section 7 identifies the channel, and none of the four steps names it.

2.2 The case against

The counterargument has been made carefully and is not dismissed here.

The Bundesbank's own position in 2011 was that TARGET balances constitute no independent risk to an individual national central bank, since the risks of monetary policy operations are shared under Article 32.4 of the ESCB Statute and the balances are a by-product of a common policy rather than a separate credit decision. Whelan [2014] develops the point: the balances are the accounting counterpart of decisions the Governing Council took collectively, and the exposure is not additional to what those decisions already implied.

Hellwig and Schnabel [2019] press further, arguing that the bailout framing confuses a stock arising mechanically from payments with a discretionary extension of credit. Perotti [2023] takes a third position,

treating the Eurosystem's accumulated exposures as financial repression whose incidence falls on savers rather than on any national balance sheet.

2.3 What the evidence here contributes to the dispute

Section 8 speaks to this in a way that favours neither side wholly.

It supports the Bundesbank's 2011 position on *attribution*. Regressing a national central bank's capital change on its own exposure and on its key share of the system's aggregate exposure gives two variables correlated at 0.611, each significant alone and neither identified against the other. Article 32 pooling is visible in the data, though the exemption of most purchase-programme holdings from risk sharing means the pooling is partial and the null is therefore weaker evidence for the Bundesbank's position than it first appears.

It supports Sinn on *magnitude and backing*. The losses were large, they reduced capital directly because no other arrangement exists, and they left four national central banks in negative equity. Section 9 shows seigniorage covers them, but seigniorage is a future flow, and the position today is what he described: a claim on a system with thin capital and no recapitalisation obligation.

And it revises both on the *source*. Neither anticipated that the exposure would be realised through the repricing of reserves against a fixed-coupon portfolio. That is this paper's contribution to a literature that has spent a decade on the credit dimension of a balance sheet whose interest rate dimension turned out to bind first.

2.4 The closest paper, and what is left

Belhocine et al. [2023] is the closest antecedent and it is closer than any other work cited here. Its statement of the mechanism is the one this paper uses: purchases removed duration risk from private balance sheets, reserves financed them, and the combination is a fixed-for-floating swap that loses money when the short rate rises.

Four things are nonetheless left, and they are what this paper does.

Measurement rather than projection. Belhocine et al. [2023] project net income forward using assumed rate paths and assumed portfolio yields for the five largest national central banks. The object here is different: the composition that generates the exposure, measured for all 21 national central banks at monthly frequency. The collateralised share of Fact 3 and the repricing gap of Fact 6 are constructed series, not forecasts, and they are what an observer would have needed in 2021 to see the exposure building.

The outturn. Their baseline anticipated about two years of loss-making and concluded that the losses would be temporary and recoupable. Capital fell in 2024, 2025 and 2026, and the third year is not obviously the last; the cumulative fall is EUR 76 billion. Capital and net income are not the same variable, since provisioning and revaluation accounts sit between them, so the two cannot be counted against each

other. What the comparison establishes is that the decline extended beyond the period their baseline anticipated. Section 9 supports their conclusion on recoupability and Fact 14 supports it on temporariness. Reporting the outturn at all is possible only because three more years have passed, and it is the least contestable thing this paper adds.

Composition as a margin. Belhocine et al. [2023] describe the swap as “not an unexpected outcome, but rather a feature of balance sheet policies shaped by the inflation objective”. That is a statement about intent, and it is almost certainly right. It is not a statement about whether the split between lending and buying was itself a decision variable. Proposition 3 says it was: two splits delivering identical reserve creation differ in income variance by a factor of nine at the model’s calibration. Whether that margin was considered is a question for the record; that it existed is a result.

Attribution and the digital euro. Neither the Article 32 question of Section 8 nor the digital euro of Section 10 appears in their paper.

Remark 1 (What this means for how the paper should be read). The contribution here is measurement, outturn and the composition margin, not the identification of a risk channel. A reader who wants the mechanism should read Belhocine et al. [2023]. A reader who wants to know how large the exposure was, how it was distributed, what it actually cost and whether it had to be taken should read on.

Remark 1 states how the paper should be read in consequence.

2.5 Central bank capital and its fiscal backing

The second literature asks when a central bank’s losses matter. Buiter [2008] sets out the accounting; Bassetto and Messer [2013] and Reis [2013] develop the fiscal implications of interest on reserves; Bonetti et al. [2024] examine what a central bank’s capital position means for its relationship with its shareholder; Del Negro and Sims [2015] and Hall and Reis [2015] give the condition used in Section 9. The contribution here is measurement: the euro area quantities the condition requires, at national level, which have not to my knowledge been assembled for this purpose.

2.6 Central bank digital currency

The third literature studies what a digital euro would do to banks and to the operational framework. Bidder et al. [2025] model disintermediation with endogenous runs and estimate a holding limit; Abad et al. [2025] trace the consequences for the operational regime; Burlon et al. [2024] compute an optimal quantity in a model where banks post collateral; Keister and Sanches [2023] ask whether the instrument should be issued at all. Each treats the central bank as one consolidated balance sheet, which is right for their questions and wrong for this one. Proposition 6 and Section 10 do not appear in that literature.

3 A simple framework

The accounting is elementary and the point is what it makes unavoidable, so it is worth setting out before the data.

3.1 Two ways to add reserves

A central bank that wants to expand its balance sheet by one euro has two instruments. The distinction between them is standard in the literature on monetary policy implementation [Bindseil, 2004]; what is not standard is reading it as a risk allocation. It can *lend* the euro to a bank against collateral, or it can *buy* a security outright. Both create one euro of reserves. They are usually treated as interchangeable ways of adding liquidity, and for the quantity of reserves they are.

They are not interchangeable for the income statement. Write r_t for the policy rate, L_t for secured credit outstanding, S_t for securities held outright with average coupon $\bar{c}$, and R_t for reserves. Reserves are remunerated at r_t . Secured credit is rolled at r_t . The coupon on the outright portfolio was fixed when the securities were bought. Net interest income is

$$\Pi_t = \underbrace{r_t L_t}_{\text{floats}} + \underbrace{\bar{c} S_t}_{\text{fixed}} - \underbrace{r_t R_t}_{\text{floats}}, \quad (2)$$

Equation (2) splits income into a part that floats and a part that does not. Differentiating with respect to the policy rate,

$$\frac{\partial \Pi_t}{\partial r_t} = -(R_t - L_t) \equiv -G_t. \quad (3)$$

Equation (3) is the whole mechanism. The sensitivity of the central bank's income to its own policy rate is the repricing gap G_t , and the gap is the volume of reserves that is *not* matched by an asset that reprices with them. A central bank that expands by lending has L_t rising with R_t and leaves G_t unchanged. A central bank that expands by buying leaves L_t where it was and raises G_t one for one.

Remark 2 (The choice of instrument is a risk choice). Nothing in (3) concerns credit quality, collateral eligibility or counterparty solvency. Two balance sheets identical in size, in credit risk and in the identity of the ultimate obligor can differ arbitrarily in $\partial \Pi / \partial r$ according to whether the exposure was taken as a loan or as a purchase. The operating framework is therefore an implicit decision about how much interest rate risk the central bank, and behind it the fiscal authority, will carry. It is rarely discussed in those terms.

3.2 When the loss matters

A negative Π_t reduces capital and, if capital is exhausted, must be met from somewhere. The condition under which it need not be met at all is the one Buiter [2008], Del Negro and Sims [2015] and Hall and Reis [2015] formalise, and it is the central bank's counterpart to the government budget constraint that

Leeper [1991], Sims [1994] and Cochrane [2023] place at the centre of price-level determination; Jeanne and Svensson [2007] show that a central bank’s concern for its own capital can itself be a commitment device: losses are absorbable if the present value of future seigniorage covers them. Writing B_t for banknotes, which pay no interest and are therefore costless funding, seigniorage is approximately $r_t B_t$ and the condition for a shortfall ΔK to be self-financing over T years at a constant rate is

$$\sum_{s=1}^T \frac{r B}{(1+\rho)^s} \geq \Delta K. \quad (4)$$

Section 9 evaluates (4) for the Eurosystem. The answer turns out to be reassuring, and reporting it is part of being honest about what the earlier sections do and do not establish.

3.3 What a digital euro does in this framework

A digital euro converts a deposit into a central bank liability. The bank that loses the deposit replaces the funding by borrowing, and the borrowing is secured, hence floating. In the notation of (3), conversion of X out of deposits extinguishes βX of reserves and raises L_t by $(1-\beta)X$. If the digital euro is unremunerated it does not join the remunerated liability base, so

$$\Delta G = \underbrace{-\beta X}_{\text{reserves fall}} - \underbrace{(1-\beta)X}_{\text{floating asset rises}} = -X, \quad (5)$$

and the instrument closes the gap by the whole amount converted. The buffer share and the borrowed share are two halves of one funding replacement, so counting them separately would count βX twice. Proposition 6 generalises (5) to a remunerated instrument and to take-up that displaces cash, and Section 10 quantifies it.

4 A model of instrument choice

Section 3 showed that the sensitivity of central bank income to its own policy rate is the repricing gap. That is an identity. This section asks the question the identity poses: given a target for monetary stimulus, how should a central bank divide it between lending and buying, and what does the division cost?

The answer reorganises the empirical sections. The losses documented in Sections 6 and 7 are not a mistake. They are the payout on insurance the central bank sold, and the question is whether the premium was adequate.

4.1 Environment

A central bank expands its balance sheet by Q , creating Q of reserves. It has two instruments:

- *lending* L against collateral, rolled at the policy rate, therefore floating, extracting no duration from the private sector;
- *purchases* S outright at coupon $\bar{c}$, fixed, extracting δ years of duration per unit.

The balance sheet identity is $Q = L + S$. Collateral capacity bounds the first instrument, $L \leq C$.

The short rate r is random ex ante with mean $\mathbb{E}[r] = \bar{r}$ and variance σ^2 . The coupon on a long asset is the expected average short rate over its life plus a term premium,

$$\bar{c} = \bar{r} + \varphi, \quad (6)$$

so φ is compensation for bearing duration and not the yield itself. Normalising $\bar{r} = 0$ makes a realisation r a surprise and $\bar{c} = \varphi$ a normalisation rather than a definition. The distinction matters: an asset pricing reader would object, correctly, to equating a coupon with a term premium, and (6) is what makes the substitution legitimate.

Assumption 1 (Instruments differ only in repricing and duration). Lending and purchases are identical in credit quality and in the identity of the ultimate obligor. They differ in two respects: purchases carry a fixed coupon while lending floats, and purchases remove duration from private portfolios while lending does not.

Assumption 1 is stronger than reality and deliberately so. The euro area's purchases were largely of sovereign paper and its lending was against a broader collateral pool, so the credit profiles are not identical. Making them identical isolates the margin the paper is about: if the results below hold when credit risk is held constant across instruments, then credit risk is not what generates them.

4.2 The income identity

Lending earns r and reserves cost r . Purchases earn $\bar{c}$ and are funded at r . Net interest income is therefore

$$\Pi = rL + \bar{c}S - rQ = S(\bar{c} - r) = S(\varphi - r). \quad (7)$$

Proposition 1 (The profit and loss is the purchase leg alone). *Under Assumption 1, net interest income depends on S and not on Q or L . Expanding by lending is income-neutral at any policy rate; expanding by buying is a carry trade in which the central bank is long the term premium and short the rate surprise.*

Proof. Substitute $L = Q - S$ into $\Pi = rL + \bar{c}S - rQ$. The terms in Q cancel, leaving $S(\bar{c} - r)$. □

Proposition 1 is the model's organising result, and its simplicity is what makes it useful. A great deal of commentary on central bank losses treats them as a consequence of balance sheet *size*. Size does not appear in (7). Two central banks with the same balance sheet, the same credit exposure and the same

ultimate obligor can have arbitrarily different income sensitivity according to which instrument they used, and one of them can have none at all.

4.3 The optimal split

That purchases work by removing duration from private portfolios, and that the resulting price effects operate through the term premium rather than through expected short rates, is the finding of Gagnon et al. [2011], Krishnamurthy and Vissing-Jorgensen [2011] and, for the euro area, Eser and Schwaab [2016] and Altavilla et al. [2021]; the theoretical mechanism is the preferred-habitat structure of Vayanos and Vila [2021] and the bond-supply channel of Greenwood and Vayanos [2014], with Christensen and Krogstrup [2019] tracing it through the central bank's own balance sheet. The central bank values duration extraction through a concave benefit $B(\delta S) = b_1 \delta S - \frac{1}{2} b_2 (\delta S)^2$, earns expected income $S\varphi$, and dislikes the variance of its income because a loss must eventually be met by a fiscal transfer whose deadweight cost is λ per unit of variance. Its objective is

$$W(S) = B(\delta S) + S\varphi - \frac{\lambda}{2} S^2 \sigma^2. \quad (8)$$

Lending does not appear. By Proposition 1 it contributes neither income nor risk; its only role is the constraint $S \geq Q - C$, which binds when collateral capacity falls short of the expansion.

Proposition 2 (Optimal purchases). *The unconstrained optimum is*

$$S^* = \frac{b_1 \delta + \varphi}{b_2 \delta^2 + \lambda \sigma^2}, \quad (9)$$

and the chosen level is $\hat{S} = \min\{\max\{S^, Q - C\}, Q\}$. Then $\partial S^* / \partial \varphi > 0$, $\partial S^* / \partial \sigma < 0$, $\partial S^* / \partial \lambda < 0$ and $\partial S^* / \partial b_2 < 0$. The sign of $\partial S^* / \partial \delta$ is ambiguous.*

Proof. W is strictly concave in S for $b_2, \lambda > 0$, so the first-order condition $\delta(b_1 - b_2 \delta S) + \varphi - \lambda \sigma^2 S = 0$ is sufficient and rearranges to (9). Each of φ , σ , λ and b_2 appears once and with a fixed sign, which gives the four signed derivatives. In δ the derivative has numerator $b_1 \lambda \sigma^2 - b_1 b_2 \delta^2 - 2b_2 \varphi \delta$, which can take either sign over the admissible parameter region; hence the effect of duration is ambiguous, and Table 13 reports the baseline case. $\square$

Figure 4 shows the objective, the two signed comparative statics that matter most, and the run-off drift of Proposition 5. Section 4.5 asks what exercising that margin would have been worth. The four signed comparative statics have a plain reading. A central bank should buy more when the market pays more for duration, and less when rates are more uncertain, when a fiscal transfer is more costly, and when the macro benefit runs into diminishing returns sooner. None of that is controversial. What is unusual is that the euro area's purchase programmes were scaled to a stimulus target, and the parameters in (9) that concern risk, σ and λ , were not part of the calibration.

4.4 Composition is a policy margin

Proposition 3 (Reserve-quantity invariance). *For any two splits (L_a, S_a) and (L_b, S_b) with $L_a + S_a = L_b + S_b = Q$, reserve creation is identical while duration extraction differs by $\delta(S_b - S_a)$ and income variance differs by $\sigma^2(S_b^2 - S_a^2)$.*

Proposition 3 says that a description of monetary policy in terms of the quantity of reserves is incomplete in a way that matters for the fiscal position. In the calibration of Table 12, moving from $S = 1,000$ to $S = 3,000$ at an unchanged $Q = 4,000$ leaves reserves untouched, triples duration extraction, and multiplies the variance of income by nine. Nothing in a conventional account of the stance records the difference.

Remark 3 (What Proposition 3 does not say). It does not say that the same stimulus could have been delivered with less fiscal exposure. In this model duration extraction is itself a channel of stimulus, $B(\delta S)$, so the two splits above deliver different macroeconomic outcomes as well as different risk: at the calibration B rises from 108.5 to 178.5 as S goes from 1,000 to 3,000. The proposition establishes something narrower and still useful. Reserve quantities do not summarise the fiscal position, so a policy described only by how many reserves it created has left out a margin that later cost EUR 76 billion. Whether the additional duration extraction was worth the additional exposure is Proposition 2, and answering it empirically requires the premium data of Remark 6.

4.5 What the composition margin was worth

Proposition 3 says composition is a policy margin. A referee is entitled to ask what exercising it would have been worth. The data answer, but the answer is a range rather than a number, and the width of the range is the more interesting result.

4.5.1 The feasibility constraint

Substituting lending for purchases requires collateral, and eligible collateral is not in this data set. What is in it is the largest volume of secured lending the system ever actually mobilised: EUR 2,217 billion in June 2021. That is a lower bound on capacity and it is the ceiling used here.

How binding the ceiling is depends on when the substituted lending would have had to coexist with actual lending, and there are two defensible readings.

Permissive. Substituted lending is resized each period, so it uses whatever headroom exists between actual lending and the ceiling at that date. The ceiling is never violated: counterfactual lending equals EUR 2,217 billion at its maximum and exceeds it in none of the 121 months. But the substituted stock must fall from EUR 1,600 billion in December 2019 to zero by June 2021 for TLTRO III to reach its peak, which requires that pledged collateral can be released on demand.

Strict. Substituted lending must be sustainable alongside the largest actual lending stock at every date. Since actual lending itself reached the ceiling, the sustainable addition is zero.

Fact 1 (The margin was worth between nothing and two fifths). Under the permissive reading, substitution would have left the repricing gap at EUR 857 billion at the end of 2024 rather than EUR 3,040 billion, and would have avoided EUR 129 billion of the EUR 307 billion cumulative funding cost, or 42 per cent. Under the strict reading it would have avoided nothing. The range is determined entirely by whether collateral pledged in 2016 to 2019 could have been released when TLTRO III peaked.

Remark 4 explains why the range and not a point estimate is the result, and Remark 5 what the exercise does not measure. Fact 2 adds the timing.

Fact 2 (The margin was large early and gone by the peak). Under the permissive reading the substitutable volume averaged EUR 1,458 billion through 2019 and was EUR 15 billion at the end of 2021. Whatever the correct reading, the composition choice was open while the programme was being built and closed by the time it peaked.

Remark 4 (Why the range is the result). A point estimate here would be false precision. The permissive reading requires frictionless release and re-pledging of collateral across a five-year interval; the strict reading requires that no substituted lending could ever have been retired early, which is equally artificial given that the actual programme was recalibrated mid-course precisely to allow early repayment, as Section 8.4 documents. The truth lies between, and locating it needs the Eurosystem collateral data listed in the replication note rather than the demonstrated-maximum proxy used here.

What the range does establish is that within the demonstrated-capacity counterfactual the fiscal value of the composition margin runs from zero to roughly two fifths of the realised funding cost. Since EUR 2,217 billion is a lower bound on capacity rather than a measurement of it, a wider true capacity would permit more substitution, not less. That is enough to make Proposition 3 economically relevant and not enough to support a claim that the cost was avoidable.

Remark 5 (This is a fiscal counterfactual, not a welfare counterfactual). The exercise computes what the fiscal exposure would have been under a different composition. It does not hold the macroeconomic stimulus fixed and cannot, because in this model duration extraction is itself a channel of stimulus: substituting lending for purchases reduces δS and therefore reduces $B(\delta S)$. Fact 1 is a statement about cost with the benefit side left out, and Remark 3 explains why the two cannot be separated within the model. Whether the foregone stimulus was worth more or less than the avoided exposure is Proposition 2, and it is not answered here.

4.6 Was the premium adequate?

Proposition 1 makes the purchase programme a carry trade, and a carry trade can be evaluated after the fact.

Remark 3 states what the proposition does not license, which is the stronger claim that the same macroeconomic outcome was available at lower risk.

Proposition 4 (Ex-post break-even). *Write the long yield at acquisition as the expected average short rate over the holding period plus a term premium, $y_t = \mathbb{E}_t[\bar{r}_{t,t+h}] + \varphi_t$. The realised carry on a purchase made at t and held to $t + h$ is*

$$\pi_t = y_t - \bar{r}_{t,t+h}^{real} = \underbrace{\varphi_t}_{\text{ex-ante compensation}} - \underbrace{(\bar{r}_{t,t+h}^{real} - \mathbb{E}_t[\bar{r}_{t,t+h}])}_{\text{rate surprise}}, \quad (10)$$

so the purchase breaks even ex post if and only if the realised short-rate surprise does not exceed the ex-ante term premium.

Equation (10) is an identity for any expectation $\mathbb{E}_t$, and it separates two things that a realised return alone conflates: what the market paid the central bank for bearing duration, and what subsequently happened to rates. Section 8.5 measures the left-hand side directly and decomposes the right-hand side under a stated benchmark.

Table 14 evaluates Proposition 4 at the model's calibration. At a term premium of 60 basis points, a realised average rate of 1.5 per cent gives a loss of EUR 23 billion a year on the purchase leg and one of 3 per cent gives EUR 60 billion.

Remark 6 (Closed in Section 8.5). Proposition 4 is the paper's sharpest testable statement. Section 8.5 evaluates its empirical counterpart, using acquisition-month sovereign yields for the four largest purchase programmes against the realised deposit facility rate. On fully realised windows the purchase leg roughly broke even, and the realised carry decomposes into ex-ante yield compensation and a rate surprise of similar magnitude; Remark 17 states the scope, which covers four of twenty-one national central banks and uses sovereign yields as a proxy for the whole securities portfolio.

4.7 Run-off without a decision

The two instruments have different maturities. Lending in the euro area was concentrated in operations with fixed three-year terms; purchases have an average residual maturity several times longer. When neither leg is renewed,

$$L(t) = L_0 e^{-t/m_L}, \quad S(t) = S_0 e^{-t/m_S}, \quad m_S > m_L. \quad (11)$$

Proposition 5 (Passive drift). *Under (11) with $m_S > m_L$, the balance sheet $L(t) + S(t)$ is strictly decreasing while the gap share $S(t)/(L(t) + S(t))$ is strictly increasing and tends to one. The central bank becomes more exposed per euro of balance sheet while shrinking it, and no decision is taken to bring this about.*

Proof. The total is a sum of two decreasing exponentials, hence decreasing. The share is $1/(1 + (L_0/S_0)e^{-t(1/m_L - 1/m_S)})$, whose denominator decreases in t because $1/m_L > 1/m_S$. $\square$

Proposition 5 matches Figure 1 exactly. From mid-2023 the Eurosystem's balance sheet fell while the repricing gap widened, and Fact 6 records the magnitude. That pattern has been read as evidence of a policy choice to hold securities. It requires no such choice: it is what (11) produces when a three-year operation and an eight-year portfolio are allowed to mature side by side.

4.8 A digital euro in the model

Conversion of X into digital euro extinguishes βX of reserves as banks meet part of the outflow from their own liquidity and induces $(1 - \beta)X$ of secured, floating lending. By (5) the gap falls by the whole X when the instrument is unremunerated and the take-up comes out of deposits.

Neither qualification is innocuous, and together they can reverse the sign.

Proposition 6 (The sign turns on one inequality). *Let χ be the share of take-up displacing banknotes rather than deposits and θ the digital euro's remuneration as a fraction of the policy rate. Conversion of X changes the repricing gap by*

$$\Delta G = \underbrace{-(1 - \theta)(1 - \chi)X}_{\text{out of deposits}} + \underbrace{\theta\chi X}_{\text{out of banknotes}} = (\theta + \chi - 1)X, \quad (12)$$

so the instrument improves the central bank's income position if and only if $1 - \chi > \theta$. The annual effect is $-\Delta G \cdot r$.

Proof. Take-up out of deposits, a share $1 - \chi$ of X , extinguishes $\beta(1 - \chi)X$ of reserves, induces $(1 - \beta)(1 - \chi)X$ of secured floating lending, and adds $\theta(1 - \chi)X$ to the remunerated liability base, contributing $-(1 - \theta)(1 - \chi)X$. The buffer share and the borrowed share are two halves of one funding replacement, so β cancels. Take-up out of banknotes, a share χ , removes an unremunerated liability, which is not in the base, and adds $\theta\chi X$ to it, contributing $+\theta\chi X$. Summing and expanding gives $(\theta + \chi - 1)X$. $\square$

The four corners make the content plain. An unremunerated digital euro taken up out of deposits closes the gap by the whole conversion. Remunerating it at the policy rate, or having it displace banknotes, removes the effect. Doing both *reverses* it: a remunerated digital euro that displaces banknotes replaces a costless liability with a paying one and widens the gap by X . The common reading that a digital euro is at worst neutral for the central bank's income position is therefore wrong, and (12) says exactly when. Table 15 and Figure 5 report the four corners.

Corollary 1 (Design choices are not separable). *By (12) the holding limit sets X , remuneration sets θ and the substitution margin sets χ , and only the first is scale. At EUR 855 billion and a two per cent policy rate the annual effect is +EUR 17.1 billion when the instrument is unremunerated and displaces deposits, zero when either condition fails alone, and -EUR 17.1 billion when both fail together. A design that remunerates the instrument to limit disintermediation and succeeds mainly in displacing cash is the corner that costs money.*

Corollary 1 also identifies the condition under which the reassuring result of Section 9 would fail. That section shows that seigniorage on the banknote stock covers the losses in under three years. Seigniorage is the policy rate on non-interest-bearing liabilities, and a digital euro that displaces cash substitutes one such liability for another while an unremunerated one displacing deposits adds to the stock; a remunerated one subtracts from it. The instrument can therefore strengthen, leave unchanged, or weaken the condition that makes the losses absorbable, depending jointly on remuneration and on the source of take-up.

5 Data

The source is the ECB’s national central bank contributions to the Eurosystem balance sheet, listed in Table 1: ten items per country per month, 21 countries, June 2016 to June 2026. Croatia enters in 2023 and Bulgaria in 2026; the panel is unbalanced by accession and is not back-filled.

Table 1: Balance sheet items used

Code	Short name	Item
A050000	refi	Lending to euro area credit institutions related to monetary policy operations. Secured, haircut, daily valuation, margin calls.
A070100	secs	Securities held for monetary policy purposes. Outright: no counterparty, no haircut, no margin call.
A010000	gold	Gold and gold receivables
A020000, A030000	fx	Claims in foreign currency
A090000	intra_claims	Intra-Eurosystem claims, gross
L100000	intra_liab	Intra-Eurosystem liabilities, gross
L010000	banknotes	Banknotes in circulation
L020000	reserves	Liabilities to credit institutions (MPO)
L150000	capital	Capital and reserves

Remark 7 (Gross claims are not TARGET balances). A090000 and L100000 are gross positions. For the Bundesbank in June 2026 they are EUR 1,068 and 605 billion. The first is close to the TARGET claim usually quoted in the debate; the second is largely the banknote allocation liability, which arises because Germany issues more banknotes than its capital key share. Quoting the gross claim as the exposure ignores an offset of more than half its size; treating the net position as the TARGET balance is equally wrong. Both are carried separately throughout, and the replication code asserts that they differ.

All statements in what follows are checked by `src/verify.py` and listed in Table 23; the build stops if any check fails.

Ratios are formed by summing levels across countries and dividing, not by averaging country ratios. The distinction is material: at June 2026 the level-weighted collateralised share is 0.80 per cent while the unweighted mean across national central banks is 5.74 per cent, because small central banks with residual credit operations are weighted like Germany.

6 Three facts

6.1 The collateralised share

Fact 3 (The secured leg has effectively disappeared). The share of Eurosystem monetary policy assets that is collateralised fell from 32.8 per cent in June 2016 to a peak of 37.2 per cent in March 2021 and then to 0.80 per cent in June 2026. Secured credit stands at EUR 25 billion against EUR 3,142 billion of securities held outright.

Figure 1 shows both the levels and the ratio. Panel (a) makes clear that this is not a story about the size of the balance sheet: monetary liabilities have fallen from their 2022 peak, and the outright portfolio is running off. It is a story about composition. The secured block, which reached EUR 2.2 trillion at the height of the targeted operations, is gone; the outright block has fallen by roughly a third and now constitutes essentially the whole of the monetary policy asset side.

Table 2 gives the dates. The collapse is concentrated in the twelve months from mid-2023, when TLTRO III was repaid.

What this means for the two arguments the paper takes as its starting point is direct. Sauer's measure of backing is a ratio of assets to money; it is unchanged by a swap of one asset class for another, and at June 2026 own assets are 1.25 times monetary liabilities. But the *kind* of asset has changed completely, and a ratio cannot see that. The asset that now stands against the euro is a sovereign bond bought outright, held to maturity, with no counterparty who can be called for margin. That is exactly the object Sauer's historical work is about, and it is now almost the entire asset side.

6.2 Capital

Fact 4 (Loss absorption has fallen fourfold, and is negative in four countries). Aggregate national central bank capital and reserves fell from 4.97 per cent of monetary liabilities in June 2016 to 1.00 per cent in June 2026. In June 2026 Austria, Belgium, Germany and Slovakia report negative capital. The Bundesbank is at minus EUR 25.3 billion, against plus EUR 5.7 billion in December 2022.

Figure 2 shows the aggregate path and the cross-section.

Whether negative central bank capital matters is exactly the question Buiter [2008], Del Negro and Sims [2015] and Hall and Reis [2015] address, and their answer is conditional: a central bank can operate with negative accounting equity provided the present value of future seigniorage, or fiscal support, is credible. The relevance here is not that the Bundesbank is insolvent. It is that the buffer between the asset side and the liability side, which is what Fact 3 is about, is no longer supplemented by an accounting cushion. Both layers have thinned at the same time, and the second is the one Sinn's argument turns on, because a national central bank with no capital and no recapitalisation obligation is precisely the counterparty whose claims he questions.

6.3 What stands behind the German claim

Fact 5 (Gross and net differ by a factor of two; the backing is outright). In June 2026 the Bundesbank holds EUR 1,068 billion of gross intra-Eurosystem claims against EUR 605 billion of gross liabilities, a net position of EUR 463 billion. Weighting debtor national central banks by their net intra-Eurosystem liabilities, 65.3 per cent of their own assets are securities held outright and 0.77 per cent are secured credit.

Figure 8 and Table 4 report both, and Table 3 gives the full cross-section. The decomposition is where the paper can be useful to a debate that has largely been conducted with a single number.

To those who quote the gross claim, the offset is real and large: more than half the gross position is a banknote allocation liability that would be netted in any dissolution. To those who dismiss the exposure as an accounting artefact of a payment system, the composition of the debtor side is the answer: what ultimately stands behind the claim is not a collateralised loan that can be called, but a portfolio of sovereign bonds held to maturity by the very central banks whose sovereigns issued them.

Sinn and Wollmershäuser [2012] made the argument when the counterpart was refinancing credit. The counterpart is now overwhelmingly outright holdings, which have neither haircut nor margin call. On his own criterion the position is weaker than when they wrote, not stronger.

7 The repricing gap

7.1 What the composition shift created

Facts 3 and 4 are usually read as two separate developments: the Eurosystem stopped lending and started holding, and separately it lost money. They are the same development.

Secured credit is rolled at the policy rate. When the rate rises, the asset repriced within weeks. Securities bought outright between 2015 and 2021 carry the coupons of that period and are held at amortised cost to maturity; their contractual coupons do not reprice, whatever happens to their market value. Reserves reprice immediately and in full. The exposure is therefore (1): the volume of liabilities that reprices without a matching asset.

Fact 6 (The gap widened fivefold). The repricing gap rose from EUR 626 billion in 2016 to EUR 3,146 billion in 2024. Reserves grew and then shrank; what left the gap open was the disappearance of the secured leg that had offset roughly half of them.

Figure 7(a) separates the two forces. Until 2020 the gap grew with reserves, but secured credit grew alongside, so the gap stayed near EUR 1.2 trillion. The pandemic operations then lifted secured credit to EUR 2.1 trillion, and the gap grew more slowly than reserves. From mid-2023 secured credit collapsed while reserves fell only gradually, so the gap kept widening even as the balance sheet shrank. That is the sequence Fact 6 records, and it is the mechanical consequence of Fact 3.

7.2 What it cost

Funding the gap at the deposit facility rate gives an implied annual carry, reported in Table 7. It is a first-order approximation in three identified respects, and the reader should hold the level loosely while holding the pattern firmly.

It ignores the coupon earned on the outright portfolio, so it understates net interest income in the negative-rate years and overstates the loss in the tightening years. It is the funding leg of (7), not net income.

It treats all reserves as remunerated at the deposit facility rate. Three institutional features cut against this. Under the two-tier system in force from October 2019 to September 2022, a multiple of minimum reserves was exempt from the negative rate and remunerated at zero. Since September 2023 minimum reserves have been remunerated at zero rather than at the deposit facility rate. And the third series of targeted operations was indexed to average key rates only from November 2022; before that its terms were fixed and in part below the deposit rate.¹

Each of these makes the effective funding rate lower than the deposit facility rate in the affected periods, so the implied carry overstates the cost. A proper measurement would construct an effective remuneration rate from the composition of reserve balances and the actual terms of outstanding operations, and would produce a smaller number.

Remark 8 (What survives the approximation). The magnitude can be bounded even without the reserve composition data. If a share of reserves equal to five per cent were exempt, cumulative carry since 2023 falls from EUR 307 to EUR 282 billion; at eight per cent it falls to EUR 266 billion; Table 17 reports the band. The two-tier system cuts the other side, reducing what the Eurosystem earned in the negative-rate years from EUR 25 to about EUR 22 billion. Both corrections make the net funding cost smaller than the headline figure, by something in the order of ten per cent at plausible exemption shares.

The corrections above are level effects concentrated in identifiable periods. They do not touch the sign of the mechanism, the timing of the lead-lag pattern in Fact 7, or the direction of Fact 6. They do mean that the figures of EUR 108 and 115 billion should be read as upper bounds on the funding cost rather than as measurements of it, and the same applies to the cumulative EUR 280 billion in Fact 15.

Remark 8 states what survives the approximation. What the calculation captures is the marginal cost of a rate change, which is the object of interest.

Fact 7 (The carry leads capital by a year). Implied carry was EUR 108 billion in 2023 and EUR 115 billion in 2024. Its correlation with the realised change in capital and reserves is 0.55 contemporaneously, 0.95 at a one-year lag, and 0.66 at two years.

The lag is not a nuisance; it is what provisioning implies. Carry costs reach the profit and loss account first and are absorbed by risk provisions and revaluation accounts before they touch capital. In 2023 the

¹These features enter the paper only as reasons the implied carry is an upper bound rather than a measurement; the direction of every result below is unaffected by their exact dates.

implied carry was EUR 108 billion and capital *rose* by EUR 5.7 billion; in 2024 the carry was similar and capital fell by EUR 27 billion. Figure 7(b) shows the pattern. A reader who took the contemporaneous correlation as the answer would conclude the relationship is weak; it is not, and the year is the buffer being consumed.

Remark 9 (What this is and is not). Fact 7 is a lead-lag relationship in eleven annual observations. It is not an identified causal estimate and the paper does not claim one. What it establishes is that a mechanism true by construction, in which a repriced liability against an unrepriced asset costs the difference, has the right sign, the right order of magnitude and the right timing to account for the observed capital path, and that no default is required to explain it. The alternative explanations are coupon income, revaluation and operating costs, all of which are in the published accounts and none of which moved by EUR 100 billion.

7.3 Why this reframes the argument

Sinn and Wollmershäuser [2012] argued that the Bundesbank’s position is a credit exposure whose value depends on debtors who might not pay, and that the debate was not pricing it. Between 2022 and 2026 the Bundesbank’s capital went from plus EUR 5.7 billion to minus EUR 25.3 billion. The concern was warranted. But the loss was not a credit loss, and the mechanism was not the one named: it was the repricing of reserves against a portfolio that does not reprice.

This is Remark 2 in the data: the instrument chosen, not the credit taken, determined the exposure. The distinction matters for what follows from it. If the risk is credit risk, the remedies are collateral standards, haircuts and risk sharing, which are the instruments the TARGET debate has argued over. If the risk is duration risk, those instruments do nothing, and the remedy is a shorter or floating asset side. That is precisely what secured refinancing is, and it is what the Eurosystem has spent five years replacing.

Sauer [2023] is closer to the mark in insisting that the composition of the asset side determines what money is worth, but his measure is a ratio of assets to money, and a ratio is blind to the maturity structure of its numerator. Own assets remain 1.25 times monetary liabilities at June 2026, a figure that has barely moved. What changed is what those assets are, and the change is invisible in the ratio.

8 Estimation

Section 7 reported a lead-lag pattern. This section puts it on a regression footing, explains why the natural cross-sectional design does not work, and argues that the reason it does not work is itself informative.

8.1 Frequency

Capital is an annual quantity. Across the panel, reported capital moves in 21 per cent of months, 56 per cent of all changes fall in the first quarter, and for the Bundesbank there are nine changes in 121 months,

all in February or March. Monthly regressions would fit the calendar on which accounts are adopted. All estimation is therefore annual.

8.2 The aggregate relation

Table 9 regresses the annual change in Eurosystem capital on implied carry, contemporaneously and at a one-year lag. The contemporaneous coefficient is 0.230 with $R^2 = 0.42$; the lagged coefficient is 0.333 with a t statistic of 7.16 and $R^2 = 0.88$.

The coefficient is well below one, and should be. Implied carry is the cost of funding the gap and ignores the coupon earned on the outright portfolio, operating income and revaluation. A pass-through of roughly a third from implied carry to reported capital is what remains after those offsets and after provisions absorb their share.

Remark 10 (What nine observations can and cannot do). Table 9 has nine annual observations and contains essentially one tightening episode. An R^2 of 0.88 on nine points that share a single common shock is visually compelling and statistically close to meaningless as inference: it does not convert an accounting mechanism into empirical identification, and no reader should take it as doing so. The relation's sign and timing follow from (3) by construction; the regression adds only the magnitude of pass-through and the location of the lag, and it adds those with nine degrees of freedom minus two. What the regression adds to the identity is the magnitude of pass-through and the location of the lag. A reader unwilling to draw inference from nine points can read Table 7 directly: implied carry of EUR 108 and 115 billion in 2023 and 2024 against capital that fell by EUR 27 billion in 2024 and EUR 45 billion in 2025, with no default in either year.

8.3 Why there is no cross-sectional design

The natural identification is differential exposure to a common shock. Every national central bank faces the same policy rate, set by the Governing Council, but holds a different repricing gap: the standard deviation of exposure between countries is 0.304 against 0.191 within, so the cross-sectional variation is there and is the larger part.

It cannot be used, and Table 10 shows why. Regressing a country's capital change on its own carry intensity gives -0.168 with $t = -3.13$. Regressing it instead on its capital key share of the *system's* aggregate exposure gives -0.095 with $t = -3.52$. Both are significant alone. Neither survives when both are included, at $t = -1.84$ and -1.82 , because the two regressors are correlated at 0.611. Adding year fixed effects, which absorb the common rate movement, leaves the own-exposure coefficient at -0.041 with $t = -0.59$.

Section 8.4 returns to this with a source of predetermined heterogeneity that breaks the collinearity, though not with identification. The explanation is institutional, and it is only partial. Article 32 of the ESCB

Statute pools monetary income across the Eurosystem and redistributes it by capital key, which removes much of the cross-sectional variation the design would need.

The pooling is not complete, and the paper should not overstate it. Under the public sector purchase programme roughly eighty per cent of holdings of national government paper sit on the acquiring national central bank's own books at its own risk, with only about twenty per cent subject to risk sharing; the pandemic programme follows similar principles.² So a national central bank's result is partly a function of what it holds. What the evidence in Table 10 shows is that the shared and own components cannot be separated in these data, not that the own component is absent.

Fact 8 (Own and shared exposure cannot be separated). Own exposure and the key share of aggregate exposure are correlated at 0.611 and neither is identified against the other. These data cannot attribute a national central bank's capital loss to its own asset composition, and cannot rule it out either.

Fact 8 bears directly on the argument this paper began with. Sinn and Wollmershäuser [2012] write about the Bundesbank's exposure. The Bundesbank's loss reflects both system-wide exposures allocated through pooling and programme-specific exposures retained on its own balance sheet, and these data cannot cleanly separate the two. That is a weaker statement than either side of the debate has made, and it is what the evidence supports.

8.4 A predetermined-exposure design, and its limits

Section 8 reported that own and pooled exposure cannot be separated because they are collinear. That is a statement about one set of regressors, not about the question. This subsection asks whether a source of predetermined heterogeneity exists that breaks the collinearity, finds that one does, and then sets out why it does not deliver identification. It is a diagnostic exercise and the paper's contribution does not rest on it.

8.4.1 The idea

Banks borrowed very different amounts under the third series of targeted operations. A national central bank whose banks had borrowed more saw a larger widening of its repricing gap when those operations ran off. Treatment intensity is refinancing over monetary liabilities at the end of 2021, before the tightening cycle began. It ranges from 0.05 to 0.70 across the eighteen national central banks with monetary liabilities above EUR 10 billion, with a standard deviation of 0.195.

8.4.2 What the design achieves

Fact 9 (It breaks the collinearity). Using predetermined repayment intensity in place of contemporaneous exposure reduces the correlation between own and pooled exposure from 0.611 to -0.077 . The collinearity that defeated Table 10 is a property of how exposure was measured there, not of the question.

²The split is used only to establish that pooling is partial. Nothing below depends on the exact figure.

Fact 9 matters independently of anything that follows. It says that institutional pooling does not make the separation of own from shared exposure mechanically impossible, which is a weaker and more defensible claim than the one Section 8 might have invited.

8.4.3 Why it does not identify

Remark 12 states what survives the three objections below.

Three reasons, in ascending order of seriousness.

The first stage is weak. Table 16 reports a coefficient of 0.399 with $t = 2.00$ and an F statistic of 4.01, far below the thresholds in Stock and Yogo [2005]. Conventional two-stage inference is not reliable at this strength and none is offered.

The sample is eighteen units and one episode. The reduced form is -2.23 with $t = -0.96$; in the joint least squares specification the own coefficient is -7.26 with $t = -1.54$ and the pooled coefficient -41.60 with $t = -0.95$. Nothing is distinguishable from zero or from anything else.

The repayment date was not predetermined. This is the substantive objection, Remark 11 states it in full, and the data make it plainly. On 27 October 2022 the Governing Council recalibrated the third series and created additional voluntary early repayment windows. The refinancing stock fell by EUR 1 billion in October 2022, by EUR 297 billion in November and by EUR 499 billion in December. Those two months are the largest and third largest monthly declines in the entire sample; the contractual maturity in June 2023 accounts for EUR 487 billion, the second largest. Whatever was fixed in 2019 and 2020, the timing of most of the run-off was determined by a policy decision taken in the middle of the tightening cycle and by banks' responses to it.

Remark 11 (Predetermined exposure, not predetermined repayment). Measuring intensity at end-2021 rules out reverse causality from outcomes in 2023 to 2025. It does not establish an exclusion restriction. Banks in Italy, Spain and Greece borrowed more under the programme than banks elsewhere, and they did so for reasons of funding conditions, sovereign spreads and banking system structure, any of which could plausibly bear on their national central banks' income independently of the repricing channel. The design shows timing, not exogeneity, and this paper does not defend the exclusion restriction it would need.

Remark 12 (What the exercise establishes). Only Fact 9: pooling does not make separation mechanically impossible. It does not establish that sample size is the only obstacle, because the exclusion restriction remains an additional and undefended assumption. A credible version would need several tightening episodes, or a source of variation in repayment timing that is not itself a policy response, and would have to argue exogeneity rather than assume it. What can be said with the data in hand is that the point estimates are both negative and the own-exposure coefficient is the more precisely estimated, which is weakly against the strong pooling reading and consistent with most purchase-programme holdings being retained at national risk. That is an observation, not a result.

8.5 The realised carry, and what produced it

Section 4 showed that a purchase programme is a carry trade whose ex-post outcome splits, by (10), into the compensation available at acquisition and the rate surprise that followed. This subsection measures both, using ten-year sovereign yields for Germany, France, Italy and Spain against the deposit facility rate.

8.5.1 Construction, and the constraint that governs it

For each country and acquisition month the acquisition-yield proxy is the ten-year yield prevailing in that month. It is a market yield to maturity, not the coupon of the securities actually bought, and is used as a proxy for the return on paper acquired at that date. The realised funding cost is the average deposit facility rate over the following h months. The realised carry is the difference, and break-even for it is zero.

Remark 13 (Two thresholds that are not the same quantity). Section 7.2 reports that covering the EUR 280 billion cumulative funding cost would have required roughly 80 basis points of average coupon on a EUR 3.5 trillion portfolio over a decade. That is a required *coupon*. The carry measured here is an acquisition-yield proxy *less* realised funding, for which break-even is zero. The two are different objects and are not compared.

The binding constraint is that a realised carry requires a window that has actually been realised. The deposit facility series ends in September 2026, so an eight-year window is complete only for acquisitions up to September 2018. Including later vintages would average the 2023 to 2026 tightening over a window that stops in 2026, weighting the high-rate years far more heavily than the full window would.

Remark 14 (A truncation that reverses the sign). The magnitude is not small. Computed over all vintages using whatever window exists, the eight-year figure is -93 basis points; computed over the 42 per cent whose eight-year window is complete, it is $+17$. The 107 basis point gap is truncation, not economics.

Remark 14 shows how much this matters. Completeness and coverage trade off directly against each other, and Table 19 reports both at every horizon rather than choosing one; Fact 10 states the result.

8.5.2 The result

Fact 10 (The realised carry is small and not robustly signed). On fully realised windows the stock-weighted carry is $+18$ basis points at three years, -16 at four, -30 at five, -4 at six, $+12$ at seven and $+17$ at eight. At four years every vintage in the window is complete and the carry is -16 basis points, or roughly break-even. The sign is not robust either across horizons or across weighting schemes.

The variation across horizons is mostly composition rather than economics. At longer horizons only the earlier vintages are complete, and those were bought at higher yields: ex-ante compensation rises from

+97 basis points at three and four years, where coverage is near total, to +143 at eight years, where only the 2016 to 2018 purchases survive the completeness filter. Reading the eight-year figure as evidence that longer holding helped would confuse the sample with the horizon.

Weighting is not innocuous either. At the four-year horizon, capital-key weighting gives +19 basis points and equal weighting +32, against the −16 obtained by weighting on holdings. The country weights are nearly identical across the three schemes; the difference is entirely in time. Holdings rose from about a twentieth of the window total in 2016 to more than a quarter in 2021, while the carry on those vintages fell from +120 to −182 basis points, so weighting on holdings puts most weight on the worst vintages. That is the economically relevant weighting, since the programme was largest when the compensation was thinnest, and it is the one reported throughout. A reader should nonetheless know that a flat weighting reverses the sign.

8.5.3 What produced it

Fact 11 (Compensation and surprise were of similar size). At four years, the longest horizon with complete coverage, ex-ante compensation under the random-walk benchmark was +97 basis points and the realised rate surprise +113, giving the −16 basis point carry exactly. Compensation was positive on 96 per cent of acquisitions. Table 20 reports the three terms at that horizon.

This is the substantive result and it is a modest one. The Eurosystem was compensated something close to a percentage point a year for holding fixed-rate paper against floating liabilities, and the tightening cycle that followed exceeded that compensation by roughly a sixth. The position lost money, but not by an amount that suggests it was priced without regard to the risk.

Remark 15 (What the benchmark assumes). The decomposition is an identity for any $\mathbb{E}_t[\bar{r}]$; the random walk sets $\mathbb{E}_t[\bar{r}] = r_t$. That is a benchmark, not a measurement, and it likely biases the compensation term upward relative to one that allows for expected policy-rate normalisation, since forward guidance led markets to price some eventual tightening. Establishing the level needs an overnight index swap forward curve at each acquisition date. A cleaner formulation still would replace the expectation with the swap rate directly, splitting the carry into a yield spread over hedged funding and the cost of having remained unhedged; that requires no claim about true expectations and is the natural next step.

Remark 16 (Yield compensation, not duration compensation). The quantity measured is the bond yield less the expected funding rate. For Italian or Spanish paper that spread contains a sovereign credit component, a liquidity component and a safe-asset convenience component alongside any term premium. The cross-country range in Table 21, +31 basis points for Germany against +204 for Italy, is substantially sovereign spread rather than compensation for duration alone. The paper therefore calls it ex-ante yield compensation and does not identify ϕ in the model with it.

Remark 16 is why the country detail below is labelled yield compensation throughout.

8.5.4 By country

Table 22 gives the underlying acquisition-yield proxies and funding rates, and Fact 12 what they show.

Fact 12 (The surprise was common; the compensation was not). At the four-year horizon the realised surprise lies between +111 and +114 basis points for all four countries, since all face the same policy rate. Ex-ante yield compensation ranges from +31 basis points for Germany to +204 for Italy, and the realised carry from –82 for Germany to +93 for Italy.

Figure 6 shows the underlying series. Where German and French yields sat close to the policy rate, the compensation available was thin; Italian and Spanish paper, carrying a sovereign spread, offered more and absorbed the surprise.

Remark 17 (Scope). Four of twenty-one national central banks; sovereign yields as a proxy for the whole securities portfolio, close by construction under PSPP and PEPP but not exact for the corporate and covered-bond components; and, at the eight-year horizon, only acquisitions through September 2018, which excludes the pandemic purchase surge where most of the volume sits. The four-year horizon covers the whole window and is the basis for Fact 11 for that reason.

8.5.5 What this changes

The purchase programmes were a carry trade that, on the four largest and on the horizons the data can support, roughly broke even or lost modestly. That is what the evidence supports. It bears on Section 4.5 in the same direction: substituting floating-rate lending for fixed-coupon purchases would have given up the compensation in Fact 11 as well as avoiding the surprise, so the composition counterfactual there is an upper bound on the fiscal gain and not an estimate of it.

9 Fiscal backing

Section 3 set out the condition under which central bank losses need no fiscal transfer: the present value of seigniorage must cover them. Table 11 evaluates it, and Remark 18 states plainly what such an evaluation can and cannot support.

The shortfall relative to the 2023 peak is EUR 82 billion. Banknotes in circulation are EUR 1,489 billion, and they pay no interest, so seigniorage at a policy rate of two per cent is EUR 30 billion a year.

Fact 13 (The losses are covered by seigniorage in under three years). At a two per cent policy rate, 2.7 years of seigniorage on the banknote stock would restore Eurosystem capital to its 2023 level. At one per cent it takes 5.5 years; at three per cent, 1.8 years.

This is the most important qualification in the paper and it points against the alarm the earlier sections might invite. The composition shift was real, its cost was large, and it produced the first negative capital positions in the Eurosystem's history. On this arithmetic it did not produce anything approaching a solvency problem, and the reason has a pleasing symmetry: the same rate that makes the gap expensive makes banknotes valuable.

Remark 18 (This is a back-of-the-envelope calculation, not a solvency test). Table 11 is undiscounted, holds the banknote stock and the policy rate fixed, and ignores every other item in the income statement. Discounting at two per cent moves the figure from 2.7 to 2.8 years and at four per cent to 3.0, so that particular omission is not what matters. What matters is that a present value of seigniorage is a general equilibrium object. Del Negro and Sims [2015] show that the size and duration of the balance sheet themselves bear on whether fiscal support is required, and Reis [2013] that computing the present value of seigniorage is not a routine exercise. The calculation here establishes an order of magnitude and is offered as one. It is not a test of (4) and should not be cited as one.

The finding also disciplines the inflation question. Sauer [2023] argues that deterioration in central bank asset quality is inflationary. The channel through which that could operate here is the one in (4): if seigniorage could not cover the losses, either a fiscal transfer or a higher price level would be required [Del Negro and Sims, 2015, Hall and Reis, 2015]. Fact 13 says the first condition does not bind at present rates. It would bind if banknote demand fell substantially, which is precisely what a digital euro that substitutes for cash rather than for deposits would do, and that is a genuine risk this paper does not quantify.

9.1 Is the exposure self-correcting?

A referee's first question about any exposure is whether it shrinks on its own. This one does, and the answer qualifies the paper's own alarm.

Fact 14 (The gap is closing). Over the twenty-four months to June 2026 the repricing gap fell by EUR 34.4 billion a month. On a linear extrapolation it reaches zero in 5.7 years. The implied carry at a two per cent policy rate falls from EUR 47 billion a year now to EUR 22 billion in three years and EUR 6 billion in five.

Table 18 reports the projection. It is crude: the trend is linear, the portfolio's maturity profile is not, and the policy rate is held fixed. What it establishes is an order of magnitude, and the order of magnitude is that the exposure documented in this paper is a stock in run-off rather than a permanent feature. Proposition 5 says the *share* of the balance sheet at risk rises during run-off; Fact 14 says the *level* of the exposure falls. Both are true and they are not in tension: the balance sheet shrinks faster than the gap within it.

9.2 What the funding leg cost, and what is still unknown

Section 4 showed that the purchase programme is a carry trade. This subsection reports the funding half of it; Section 8.5 adds a proxy for the coupon half and finds that on complete observable windows the combined carry is close to break-even.

Fact 15 (The funding leg). Cumulating implied carry over the sample, the Eurosystem earned EUR 34 billion on the gap across 5.6 years of negative policy rates and paid EUR 315 billion across the 4.5 years since, a net funding cost of EUR 280 billion.

Fact 15 is the funding half of (7) and omits the coupon earned on the purchase portfolio, which is the other half. It therefore bounds the question rather than answering it: the programme broke even if and only if coupon income over the same period exceeded EUR 280 billion. On a portfolio averaging roughly EUR 3.5 trillion over a decade, that requires an average coupon of about 80 basis points.

Whether the portfolio earned that is the question Section 8.5 takes up with acquisition-month yields; the answer there is that on fully observed windows it roughly did. The point of Fact 15 is that the threshold is not obviously out of reach, which is a different and more useful statement than either that the programme was ruinous or that its losses were costless.

To put the realised loss in scale: cumulative Eurosystem capital fell by EUR 76 billion between 2022 and June 2026, or about 0.51 per cent of euro area annual output.

9.3 How other central banks arranged the same exposure

The Eurosystem is not alone in having expanded through purchases and then paid for it. What differs across central banks is not the exposure but the arrangement for absorbing the loss, and the contrast is instructive because it shows that the arrangement is a choice made in advance rather than a consequence of the accounting.

Four cases are worth setting side by side. The descriptions below are stated from the published accounts and policy documents of the four institutions; readers should treat them as summaries rather than as legal characterisations.

The *Federal Reserve* suspends remittances and records the shortfall in its accrued remittances to Treasury, or deferred asset, account. A debit balance there represents the net earnings that must be realised before remittances resume; Reserve Bank capital is not reduced. The arrangement makes the fiscal backing in (4) explicit rather than implicit.

The *Bank of England* operates its purchase portfolio in a separate facility indemnified by HM Treasury. Gains were remitted to the Treasury during the low-rate years and losses are reimbursed by it now. The exposure never touches the central bank's own capital, and the fiscal transfer is contemporaneous rather than deferred.

The *Swiss National Bank* has no indemnity either. Losses reduce equity directly, and when they did so on a large scale the consequence was visible and political: distributions to the Confederation and the cantons, which several cantonal budgets had come to treat as recurring, were suspended.³ The Swiss case is the closest analogue to the Eurosystem's, and it shows what the absence of an arrangement produces once the loss is large enough to matter fiscally.

The *Eurosystem* has neither arrangement. There is no deferred asset, no indemnity, and, as Sinn and Wollmershäuser [2012] emphasises, no member state is obliged to recapitalise its national central bank. Losses reduce capital directly, which is why Fact 4 records four central banks in negative equity while the Federal Reserve, with a comparable exposure, reports none.

Remark 19 (The accounting differs; the economics does not). None of the three arrangements changes the resource cost. In each case the public sector has sold duration insurance and is paying the claim. What differs is where the payment is recorded and how visible it is: a deferred asset makes it a future reduction in remittances, an indemnity makes it a current expenditure, and the Eurosystem's arrangement makes it a reduction in central bank capital without a corresponding claim on the fiscal authority. The last is the least visible of the three and, precisely for that reason, the one most likely to be read as a solvency event rather than as the insurance payout that Proposition 1 shows it to be.

Remark 19 draws the conclusion. The comparison also bears on Fact 8. Within the Eurosystem, monetary income pooling under Article 32 and the programme-specific risk-sharing rules intermingle own and shared exposure, so a national central bank's capital reflects neither its own balance sheet alone nor the system's alone. Across the three central banks above, the loss is allocated by three different rules, none of which is the balance sheet either. In both dimensions the location of a central bank loss is an institutional convention, and reading it as information about the underlying exposure is a mistake.

10 What a digital euro would change

10.1 The mechanics

A common intuition holds that a digital euro is money creation without backing. The intuition does not survive the second leg of the transaction. Conversion is a swap. A household gives up a deposit and receives a claim on the Eurosystem. Broad money is unchanged. Central bank money rises, and the offsetting asset is a claim on the bank that lost the deposit, secured under Article 18.1.

Write X for the amount converted and β for the share of the outflow the bank meets from its own liquid assets rather than by borrowing. New secured credit is $(1 - \beta)X$, and the collateralised share becomes

$$c_1 = \frac{c_0 M + (1 - \beta)X}{M + (1 - \beta)X}, \quad M = \text{refi} + \text{secs}, \quad (13)$$

³The suspension and the underlying loss are reported in the Swiss National Bank's annual accounts for the years concerned; the point here is the arrangement, not the amount.

where c_0 is the current share. Equation (13) holds the outright portfolio fixed, which is conservative for the paper’s own argument: if the portfolio ran off while a digital euro was taken up, the share would rise by more.

10.2 Magnitudes

Table 5 and Figure 3 report the arithmetic at 285 million holders; Table 6 inverts it, giving the limit each target share would require.

Fact 16 (A digital euro rebuilds the secured leg). At a EUR 3,000 limit, fully taken up, with $\beta = 0.3$, the collateralised share rises from 0.80 to 16.6 per cent. At $\beta = 0$ it rises to 21.9 per cent. Restoring the March 2021 peak would require EUR 9,185 per holder; reaching ten per cent requires EUR 1,623.

Remark 20 (Two margins the counterfactual omits). Equation (13) assumes a bank meeting a deposit outflow either draws on its own liquidity or borrows from the central bank. It has two other options: it can replace the deposit with wholesale funding, or it can shrink its asset side. Both are modelled in the CBDC literature [Burlon et al., 2024, Bidder et al., 2025] and neither appears here. The consequence differs across the paper’s two objects. The collateralised share of Fact 3 depends on β directly, since only central bank credit adds secured assets, so allowing wholesale funding or asset shrinkage lowers the figures in Table 5 proportionately. The repricing gap of Proposition 6 does not: there β cancels, because reserves drawn down and credit taken up are two halves of the same replacement. What alternative adjustment would change in that case is the composition of the bank’s response, not the central bank’s exposure, provided the deposit itself is still converted. The counterfactual is arithmetic under a stated split, not a behavioural prediction, and Table 5 reports it across the whole range of β for that reason.

Fact 17 (A digital euro closes part of the gap). By Proposition 6, an unremunerated digital euro taken up out of deposits closes the gap one for one in the amount converted. At a EUR 3,000 limit the gap falls from EUR 2,343 to EUR 1,488 billion, a reduction of 36.5 per cent, worth EUR 17.1 billion of annual carry at a two per cent policy rate.

Table 8 reports the arithmetic across limits and rates, subject to the two omitted adjustment margins of Remark 20. The effect is proportional to the policy rate, so the instrument works as insurance. It does most for the balance sheet exactly when a tightening cycle is under way.

Two readings follow, and they point in opposite directions.

On the composition of the asset side, a digital euro is an improvement by the standard both Sinn and Sauer apply. It replaces a liability backed by an outright sovereign portfolio with one backed by a secured, haircut, daily margined claim on a commercial bank. If the concern is that the Eurosystem’s assets carry no loss-absorbing buffer negotiated in advance, then the digital euro supplies exactly that buffer, mechanically, at a magnitude the limit controls.

Facts 16 and 17 together are the case for the instrument on this margin. On the level, no plausible limit restores the position of five years ago. The EUR 3,000 limit under discussion recovers less than half the ground lost since March 2021. The instrument is not a solution to Fact 3; it is a partial and automatic offset.

10.3 Where backing does fail

The rebuild in (13) is arithmetic, and arithmetic assumes the collateral exists. Two constraints bind and neither is in the equation.

First, capacity. Alexopoulou et al. [2025] document that Additional Credit Claims frameworks, under which national central banks accept loan portfolios falling outside the general framework, rose to roughly EUR 500 billion of post-haircut collateral value during the pandemic. The temporary framework is being phased out, while part of its credit-claim capacity is being migrated into the permanent general collateral framework. The relevant constraint is therefore how much collateral can be mobilised under the post-transition framework, and at what speed.

Second, speed. Collateral mobilisation is a legal and operational process measured in days. A digital euro conversion is measured in seconds. Where the holding limit is high relative to the collateral a bank can mobilise in the time available, the mechanism in (13) does not operate smoothly; it fails discontinuously, and the bank goes to emergency liquidity assistance, which sits on the national balance sheet at national risk.

These are the places where the concern about backing is well founded. They are not about money creation, and they are not addressed by a lower holding limit alone, since a lower limit reduces both the risk and the offset.

11 Limitations

Ten years, one regime. The panel begins in June 2016, inside the asset purchase era. It cannot show what the collateralised share looked like before 2015, when refinancing was the whole of the asset side, and it therefore cannot distinguish a return to normal from a departure from it.

The counterfactual is arithmetic, not behaviour. Equation (13) takes take-up as given and the outright portfolio as fixed. It is a mapping from a policy parameter to a balance sheet composition, not a prediction. The one behavioural parameter, β , is varied over its whole range and the results are reported for each value rather than at a preferred one.

The debtor decomposition is a weighted average. Fact 5 weights debtor national central banks by their net intra-Eurosystem liabilities and reports their aggregate asset mix. It does not trace individual claims, which the published data do not permit.

The exposure is in run-off. Fact 14 projects the gap to zero in under six years on a linear trend. The paper documents a stock that is shrinking, and a reader who takes the level rather than the composition as the object of concern should weigh that.

The policy rate is an input, not a measurement. The deposit facility rates in Table 7 are annual averages taken from the published series rather than reconstructed from the underlying decisions. The implied carry scales linearly in them, so the level of Fact 7 depends on them; the lead-lag structure does not.

Eleven annual observations. Fact 7 is a correlation in a short series, discussed in Remark 9. It is consistent with the mechanism and inconsistent with a credit-loss account, but it does not identify a causal magnitude and no attempt is made to.

Whether composition matters is not settled here. The paper documents that the composition of Eurosystem assets has changed and quantifies by how much. Whether that has consequences for the price level is the question Sauer [2023] raises and Del Negro and Sims [2015] and Hall and Reis [2015] address formally, and it depends on fiscal backing, which is outside this data set. Readers who hold that a fiat central bank's asset composition is irrelevant will read Section 6 as a description without implications. The facts hold under either reading, and that is why they are reported.

12 Conclusion

The question this paper began with was how the implementation of monetary policy allocates duration risk between the private and public sectors. The euro area's answer over the past decade is that it allocated a great deal of it, that the allocation was made by instrument choice rather than by any explicit decision about risk, and that its realisation coincided with a decline of EUR 76 billion in central bank capital when the rate cycle turned. The paper does not identify that decline as a causal effect and Remark 10 says why.

What backs the euro has changed, and the change is larger and faster than the debate has registered. A decade ago roughly a third of Eurosystem monetary policy assets were secured claims with haircuts and margin calls. Today it is 0.80 per cent. Capital and reserves have fallen from five per cent of monetary liabilities to one, and four national central banks, including the largest creditor in the system, report negative capital. And the purchases made across the four largest programmes roughly broke even on the horizons the data can support: a realised carry of minus 16 basis points at four years, decomposing into about 97 basis points of ex-ante yield compensation and a rate surprise a sixth larger.

For the argument Sinn and Wollmershäuser [2012] has made for over a decade, this is a strengthening of the premise rather than a refutation of it. The counterpart to the Bundesbank's claim is no longer collateralised credit that can be called; it is a sovereign bond portfolio held outright. For the argument Sauer [2023] makes about the backing of money, it supplies the measurement his framework implies but has not been given at national level.

The digital euro, on this evidence, is not the problem it is often taken to be, but neither is it unambiguously a remedy. It is a swap, and whether it swaps in a helpful direction turns on two design choices rather than

on the instrument itself: by Proposition 6 it narrows the repricing gap whenever the remuneration share plus the cash-substitution share falls below one, is neutral at equality, and widens the gap above it. Even in the favourable corner it does not restore the position that has been lost, and the collateral it requires at the speed of a digital conversion is being reduced at the same time the instrument is being built. That is where the concern about backing belongs.

References

- Jorge Abad, Galo Nuño, and Carlos Thomas. CBDC and the operational framework of monetary policy. *Journal of Monetary Economics*, 151:103762, 2025.
- Viral V. Acharya, Tim Eisert, Christian Eufinger, and Christian Hirsch. Whatever it takes: The real effects of unconventional monetary policy. *Review of Financial Studies*, 32(9):3366–3411, 2019.
- Ioana Alexopoulou, Calogero Brancatelli, Adina-Elena Fudulache, Diana Gomes, Daniel Gybas, and Stephan Sauer. Collateral easing in non-standard times: a review of the role of additional credit claims in the eurosystem collateral framework. Occasional Paper Series 378, European Central Bank, 2025.
- Carlo Altavilla, Giacomo Carboni, and Roberto Motto. Asset purchase programs and financial markets: Lessons from the euro area. *International Journal of Central Banking*, 17(4):1–48, 2021.
- Marco Bassetto and Todd Messer. Fiscal consequences of paying interest on reserves. *Fiscal Studies*, 34(4):413–436, 2013.
- Nazim Belhocine, Ashok Vir Bhatia, and Jan Frie. Raising rates with a large balance sheet: The eurosystem’s net income and its fiscal implications. IMF Working Paper WP/23/145, International Monetary Fund, 2023.
- Rhys Bidder, Timothy Jackson, and Matthias Rottner. CBDC and banks: disintermediating fast and slow. BIS Working Paper 1280, Bank for International Settlements, 2025.
- Ulrich Bindseil. Monetary policy implementation: Theory, past, and present. *Oxford University Press*, 2004.
- Matteo Bonetti, Dirk Broeders, Damiaan Chen, and Daniel Dimitrov. Central bank capital and shareholder relationship. *De Nederlandsche Bank Working Paper*, (809), 2024.
- Willem H. Buiter. Can central banks go broke? *CEPR Policy Insight*, (24), 2008.
- Lorenzo Burlon, Manuel A. Muñoz, and Frank Smets. The optimal quantity of CBDC in a bank-based economy. *American Economic Journal: Macroeconomics*, 16(4):172–217, 2024.
- Jens H. E. Christensen and Signe Krogstrup. The slope of the term structure and the central bank balance sheet. *Economic Journal*, 129(617):249–272, 2019.
- John H. Cochrane. The fiscal theory of the price level. 2023.
- Marco Del Negro and Christopher A. Sims. When does a central bank’s balance sheet require fiscal support? *Journal of Monetary Economics*, 73:1–19, 2015. doi: 10.1016/j.jmoneco.2015.05.001.
- Itamar Drechsler, Thomas Drechsel, David Marques-Ibanez, and Philipp Schnabl. Who borrows from the lender of last resort? *Journal of Finance*, 71(5):1933–1974, 2016.
- Fabian Eser and Bernd Schwaab. The yield impact of the ECB’s asset purchase programme. *Journal of Financial Economics*, 119(1):147–167, 2016.
- Joseph Gagnon, Matthew Raskin, Julie Remache, and Brian Sack. The financial market effects of the federal reserve’s large-scale asset purchases. *International Journal of Central Banking*, 7(1):3–43, 2011.
- Robin Greenwood and Dimitri Vayanos. Bond supply and excess bond returns. *Review of Financial Studies*, 27(3):663–713, 2014.
- Robert E. Hall and Ricardo Reis. Maintaining central-bank financial stability under new-style central banking. *NBER Working Paper*, (21173), 2015.
- Martin Hellwig and Isabel Schnabel. TARGET-Falle oder Empörungsfalle? Zur deutschen Diskussion um die Europäische Währungsunion. *Wirtschaftsdienst*, 99(9):632–640, 2019.
- Olivier Jeanne and Lars E. O. Svensson. Credible commitment to optimal escape from a liquidity trap: The role of the balance sheet of an independent central bank. *American Economic Review*, 97(1):474–490, 2007. doi: 10.1257/aer.97.1.474.
- Todd Keister and Daniel Sanches. Should central banks issue digital currency? *Review of Economic Studies*, 90(1):404–431, 2023.
- Arvind Krishnamurthy and Annette Vissing-Jorgensen. The effects of quantitative easing on interest rates: Channels and implications for policy. *Brookings Papers on Economic Activity*, 42(2):215–287, 2011.

- Eric M. Leeper. Equilibria under active and passive monetary and fiscal policies. *Journal of Monetary Economics*, 27(1): 129–147, 1991.
- Kjell G. Nyborg. Central bank collateral frameworks. *Journal of Banking and Finance*, 76:198–214, 2017.
- Roberto Perotti. Government debt, financial repression, and the eurosystem. Working paper, 2023.
- Ricardo Reis. The mystique surrounding the central bank’s balance sheet, applied to the european crisis. *American Economic Review*, 103(3):135–140, 2013. doi: 10.1257/aer.103.3.135.
- Ingo Sauer. The lessons from 1923 for the euro area: Enlightening the dark side of (in-)solvent central banks’ balance sheets. Technical report, Goethe University Frankfurt, 2023. SSRN Working Paper 4620462.
- Christopher A. Sims. A simple model for study of the determination of the price level and the interaction of monetary and fiscal policy. *Economic Theory*, 4(3):381–399, 1994.
- Hans-Werner Sinn. *The Euro Trap: On Bursting Bubbles, Budgets and Beliefs*. Oxford University Press, Oxford, 2014.
- Hans-Werner Sinn and Timo Wollmershäuser. TARGET loans, current account balances and capital flows: the ECB’s rescue facility. *International Tax and Public Finance*, 19(4):468–508, 2012. doi: 10.1007/s10797-012-9236-x.
- James H. Stock and Motohiro Yogo. Testing for weak instruments in linear IV regression. In Donald W. K. Andrews and James H. Stock, editors, *Identification and Inference for Econometric Models: Essays in Honor of Thomas Rothenberg*, pages 80–108. Cambridge University Press, 2005.
- Dimitri Vayanos and Jean-Luc Vila. A preferred-habitat model of the term structure of interest rates. *Econometrica*, 89(1): 77–112, 2021. doi: 10.3982/ECTA17440.
- Karl Whelan. TARGET2 and central bank balance sheets. *Economic Policy*, 29(77):79–137, 2014.

Tables and Figures

Table 2: The collapse of the collateralised share, EUR billion and per cent

	Secured credit	Securities outright	Secured share (%)	Monetary liabilities	Capital	Capital (% of money)
Jun 2016	540	1,106	32.80	2,025	100.7	4.97
Dec 2018	732	2,395	23.41	3,098	104.2	3.37
Dec 2020	1,793	3,345	34.90	4,809	108.7	2.26
Mar 2021	2,107	3,565	37.15	5,469	110.0	2.01
Dec 2022	1,324	4,480	22.82	5,445	114.5	2.10
Dec 2023	410	4,269	8.77	4,951	120.2	2.43
Dec 2024	34	3,906	0.87	4,532	93.2	2.06
Jun 2026	25	3,142	0.80	3,870	38.5	1.00

Table 3: By national central bank, June 2026

	Secured credit	Securities outright	Secured share (%)	Monetary liabilities	Capital	Net intra-Eurosystem
DE	6.7	720	0.92	1,203	-25.3	462.7
FR	0.1	712	0.01	804	18.6	-128.6
ES	0.3	467	0.06	298	1.9	-233.0
IT	13.1	467	2.72	328	26.3	-283.9
BE	0.0	160	0.00	281	-2.1	115.3
NL	0.0	153	0.00	270	3.2	106.3
GR	3.9	114	3.33	53	3.5	-70.6
AT	1.0	84	1.14	101	-1.0	-7.1
PT	0.0	66	0.03	56	2.0	-21.9
FI	0.0	65	0.00	97	2.9	30.2
IE	0.0	46	0.05	104	6.5	60.7
SK	0.0	32	0.00	24	-3.8	-20.4
SI	0.0	13	0.00	12	1.3	0.9
LV	0.0	9	0.21	12	0.3	0.4
LU	0.0	9	0.00	122	0.2	112.2
LT	0.0	9	0.34	23	0.6	16.3
EE	0.0	8	0.00	12	0.6	2.4
CY	0.0	5	0.02	19	0.3	16.8
BG	0.0	0		25	1.8	26.1
HR	0.3	0	100.00	20	0.5	10.9

Central banks with monetary liabilities below EUR 10 billion are omitted. Bulgaria and Croatia hold no securities for monetary policy purposes, so their share is undefined or degenerate and is reported as published.

Table 4: German intra-Eurosystem position, June 2026

	EUR billion
Gross intra-Eurosystem claims	1,067.83
Gross intra-Eurosystem liabilities	605.12
Net position	462.71

See Remark 7. The gross claim is close to the TARGET figure usually quoted; the offset is largely the banknote allocation liability.

Table 5: Counterfactual: the collateralised share under a digital euro

Limit (EUR)	β	Conversion	New secured credit	Secured share (%)
500	0	142.50	142.50	5.07
	0.3	142.50	99.75	3.83
	0.6	142.50	57.00	2.56
1000	0	285.00	285.00	8.99
	0.3	285.00	199.50	6.68
	0.6	285.00	114.00	4.25
2000	0	570.00	570.00	15.93
	0.3	570.00	399.00	11.90
	0.6	570.00	228.00	7.46
3000	0	855.00	855.00	21.89
	0.3	855.00	598.50	16.57
	0.6	855.00	342.00	10.47
5000	0	1,425.00	1,425.00	31.58
	0.3	1,425.00	997.50	24.56
	0.6	1,425.00	570.00	15.93

285 million holders, full take-up. β is the share of the outflow banks meet from own liquid assets. Conversion and new credit in EUR billion.

Table 6: Holding limit required to restore a given share, $\beta = 0.3$

	Required limit (EUR)
March 2021 peak	9,185
December 2016	7,562
Ten per cent	1,623
Twenty per cent	3,810

Table 7: The repricing gap and the capital path

	Reserves	Secured credit	Repricing gap	Policy rate (%)	Implied carry	Capital	Change in capital
2016	1,175.0	548.6	626.4	-0.4	2.5	99.4	
2017	1,770.8	742.3	1,028.5	-0.4	4.1	102.2	2.8
2018	1,994.2	747.3	1,246.8	-0.4	5.0	104.2	2.1
2019	1,939.6	695.0	1,244.6	-0.5	5.6	107.1	2.8
2020	2,758.3	1,323.8	1,434.6	-0.5	7.2	108.7	1.6
2021	4,299.5	2,114.9	2,184.6	-0.5	10.9	109.2	0.5
2022	4,630.4	2,064.2	2,566.3	0.0	0.0	114.5	5.3
2023	3,924.3	795.9	3,128.4	3.5	-107.9	120.2	5.7
2024	3,289.9	144.0	3,145.9	3.6	-114.8	93.2	-27.0
2025	2,830.2	25.7	2,804.5	2.4	-67.3	48.1	-45.0
2026	2,477.5	26.1	2,451.4	2.0	-49.0	38.5	-9.6

EUR billion, annual averages except capital, which is end of year. The policy rate is the annual average of the ECB deposit facility rate; the series and its construction are documented in the replication note. Implied carry scales linearly in it, so the level of the column depends on it while the lead-lag structure of Fact 7 does not. Implied carry ignores coupon income on the outright portfolio.

Table 8: How much of the gap a digital euro closes

Limit (EUR)	Policy rate (%)	New floating assets	Gap closed (%)	Carry saved
1000	2	199.5	12.2	5.7
	4	199.5	12.2	11.4
3000	2	598.5	36.5	17.1
	4	598.5	36.5	34.2
5000	2	997.5	60.8	28.5
	4	997.5	60.8	57.0

285 million holders, full take-up, $\beta = 0.3$. Carry saved in EUR billion per year at the stated policy rate.

Table 9: Change in Eurosystem capital on implied carry, annual

	β	s.e.	t	p	R^2	N
Contemporaneous	0.230	0.095	2.413	0.042	0.421	10.000
Lagged one year	0.333	0.047	7.157	0.000	0.880	9.000

Dependent variable: annual change in aggregate capital and reserves, EUR billion. See Remark 10 on what nine observations support.

Table 10: Own exposure against the key share of aggregate exposure

	Own exposure	Key share	R^2	N
Own exposure only	-0.168 (-3.13)		0.311	182.0
Key share of aggregate only		-0.095 (-3.52)	0.298	182.0
Both	-0.114 (-1.84)	-0.048 (-1.82)	0.323	182.0
Own, with year fixed effects	-0.041 (-0.59)		0.375	182.0

Country-year panel, dependent variable the change in capital as a per cent of lagged monetary liabilities. Country fixed effects throughout; standard errors clustered by country; t statistics in parentheses. The two regressors are correlated at 0.611.

Table 11: Seigniorage and the time to rebuild capital

	Annual seigniorage	Years to rebuild
1 per cent	14.9	5.5
2 per cent	29.8	2.7
3 per cent	44.7	1.8

Seigniorage is the policy rate applied to banknotes in circulation, EUR 1,489 billion. Target is the 2023 capital peak of EUR 120 billion, a shortfall of EUR 82 billion. Undiscounted.

Table 12: Model calibration and the optimal split

Symbol	Meaning	Value
Q	Size of the expansion, EUR bn	4,000
C	Collateral capacity, EUR bn	2,500
δ	Duration extracted per unit purchased, years	7.0
φ	Ex ante term premium, per cent	0.60
σ	Standard deviation of the short rate, per cent	1.80
λ	Deadweight cost of a fiscal transfer	0.020
b_1	Marginal benefit of duration extraction	0.019
b_2	Curvature of that benefit	1.0×10^{-6}
m_L	Average maturity of lending, years	3.0
m_S	Average maturity of purchases, years	8.0
S^*	Optimal purchases, EUR bn	2,505
	Expected income at S^* , EUR bn per year	15.0
	Standard deviation of that income, EUR bn	45.1

Illustrative, chosen to sit in the region the euro area occupied between 2015 and 2021. The collateral constraint does not bind at these values, so the optimum in Proposition 2 is interior; a grid search over $[0, Q]$ reproduces it to within one billion.

Table 13: Comparative statics of the optimal purchase volume

	Low	Baseline	High	Sign
Term premium φ	2,415	2,505	2,614	increasing
Rate volatility σ	2,765	2,505	1,891	decreasing
Transfer cost λ	2,800	2,505	1,855	decreasing
Benefit curvature b_2	8,538	2,505	553	decreasing
Duration per unit δ	4,070	2,505	1,343	decreasing

S^* in EUR billion at low, baseline and high values of each parameter. The first four signs are unambiguous, as each parameter enters (9) once. The sign in δ is specific to this calibration: at $b_2 = 10^{-7}$ the relationship is non-monotone.

Table 14: Break-even on the purchase leg

	Expected income	Realised income	Premium adequate
0.000000	15.0	15.0	yes
0.600000	15.0	0.0	yes
1.500000	15.0	-22.5	no
3.000000	15.0	-60.1	no

EUR billion per year at $S^* = 2,505$ billion and a term premium of 60 basis points, following Proposition 4. See Remark 6 on what an empirical counterpart would require.

Table 15: The two legs of a digital euro's income effect

Remunerated	From banknotes (%)	Gap change	Annual effect at 2%
no	0	-855.0	17.1
	50	-427.5	8.6
	100	0.0	-0.0
yes	0	0.0	-0.0
	50	427.5	-8.6
	100	855.0	-17.1

Take-up of EUR 855 billion at a two per cent policy rate, following Proposition 6. Gap change in EUR billion; the annual effect is that times the rate. The gain vanishes if full remuneration is introduced while take-up comes out of deposits, or if banknotes are displaced while the instrument is unremunerated. If both occur together the sign reverses.

Table 16: Predetermined repayment intensity as a source of variation

	Coefficient	<i>t</i>	<i>R</i> ²
First stage: Δ gap on intensity	+0.399	+2.00	0.172
Reduced form: Δ capital on intensity	-2.227	-0.96	0.015
Own exposure only	-6.993	-1.52	0.133
Key share of aggregate only	-34.905	-0.91	0.023
Both: own	-7.261	-1.54	0.166
Both: key share	-41.600	-0.95	

Eighteen national central banks with monetary liabilities above EUR 10 billion. Intensity is refinancing over monetary liabilities at end-2021; the dependent variable is the change from end-2022 to end-2025, scaled by end-2021 monetary liabilities. Heteroskedasticity-robust standard errors. All six rows are least squares; the last two are a horse race between own and pooled exposure, not a second stage. The first-stage *F* of 4.01 is far below the thresholds in Stock and Yogo [2005], so no two-stage estimates are computed.

Table 17: Implied carry under exempt reserves, EUR billion

	Carry since 2023	Carry before 2023
0%	-307	25
3%	-292	23
5%	-282	22
8%	-266	21

Share of reserves assumed exempt from the deposit facility rate, applied from September 2023 and at three times that share during the two-tier period. The share is not observed in these data; the table is a band, not an estimate.

Table 18: The gap in run-off

	Projected gap	Carry at 2 per cent
In 1 year	1,930.8	38.6
In 2 years	1,518.4	30.4
In 3 years	1,105.9	22.1
In 5 years	280.9	5.6

EUR billion, linear extrapolation of the trend over the twenty-four months to June 2026, at an unchanged two per cent policy rate. Crude by construction; the order of magnitude is the content.

Table 19: Realised carry by horizon, complete windows only

Horizon (years)	<i>N</i>	Coverage (%)	Last complete vintage	Ex-ante (bp)	Surprise (bp)	Realised carry (bp)
3	268	100	2021-12	96.6	78.4	18.2
4	268	100	2021-12	96.6	112.8	-16.2
5	256	96	2021-09	98.6	129.0	-30.4
6	208	78	2020-09	115.9	120.0	-4.1
7	160	60	2019-09	134.1	122.2	11.9
8	112	42	2018-09	142.5	125.7	16.8

Stock-weighted realised carry, the acquisition-yield proxy less the realised average deposit facility rate, June 2016 to December 2021, Germany, France, Italy and Spain. Break-even for this quantity is zero; see Remark 13 on why the 80 basis point figure in Section 7.2 is a different object.

Table 20: Decomposing the realised carry

	Basis points
Ex-ante compensation, random-walk benchmark	96.6
less realised rate surprise	-112.8
equals realised carry	-16.2

Stock-weighted, June 2016 to December 2021, four countries, four-year horizon, the longest with complete coverage. Ex-ante compensation uses the random-walk benchmark $\mathbb{E}_t[\tilde{r}] = r_t$; Remark 15 states the direction of its bias. The three rows are an identity and sum exactly.

Table 21: The decomposition by country

	Ex-ante (bp)	Surprise (bp)	Realised carry (bp)
DE	31.3	113.1	-81.7
ES	124.9	112.2	12.7
FR	70.1	114.0	-43.8
IT	204.3	111.5	92.8

Basis points, stock-weighted within country. The surprise is common because all four face the same policy rate; the variation is on the compensation side.

Table 22: Realised carry by country, complete four-year windows

	N	Avg. acquisition yield (%)	Realised avg. rate (%)	Realised carry (bp)	Share positive (%)
DE	67	-0.1	0.4	-81.7	52.2
FR	67	0.3	0.4	-43.8	53.7
IT	67	1.7	0.4	92.8	73.1
ES	67	0.9	0.4	12.7	56.7

Stock-weighted within country at the four-year horizon, where every vintage in the acquisition window is complete. N is the number of country-months with a weight from the balance sheet panel, out of 67 possible.




Figure 1: **What backs central bank money in the euro area.** Panel (a): monetary policy assets split into secured credit operations and securities held outright, against monetary liabilities. Panel (b): the secured share, from 37.2 per cent in March 2021 to 0.80 per cent in June 2026.

Table 23: Executable claim audit

	Result
Panel loads with the expected items, countries and frequency	Pass
Units are EUR billion, cross-checked against known magnitudes	Pass
Gross intra-Eurosystem claim and net position are distinct	Pass
Panel is unbalanced only through euro area accession	Pass
The collateralised share falls by more than an order of magnitude	Pass
Ratios are finite where defined and the share lies in the unit interval	Pass
Aggregate ratios are computed from summed levels, not averaged	Pass
The counterfactual share is increasing in the holding limit	Pass
Inverting the counterfactual reproduces the target share	Pass
The behavioural parameter brackets the effect as it should	Pass
Capital per euro of monetary liability has fallen and some NCBs are negative	Pass
The debtor side of the creditor claim is almost entirely uncollateralised	Pass
The repricing gap widened as secured credit disappeared	Pass
Implied carry leads the capital change by one year	Pass
A digital euro closes part of the repricing gap	Pass
Net income depends on purchases alone, not on the expansion size	Pass
Closed-form optimum matches a grid search and is interior	Pass
Four comparative statics are signed; the one in duration is not	Pass
Composition changes risk at an unchanged quantity of reserves	Pass
Run-off shrinks the balance sheet and raises the gap share	Pass
The digital euro effect is signed by whether $1 - \chi$ exceeds θ	Pass
The repricing gap is shrinking on a linear trend	Pass
The funding leg earned at negative rates and paid since 2022	Pass
The predetermined-exposure design breaks collinearity and is weak	Pass
Most of the refinancing run-off followed the 2022 recalibration	Pass
The composition counterfactual is a bounded range, not a number	Pass
Exempting reserves lowers the implied carry monotonically	Pass
Sovereign yields and the deposit facility rate load and overlap	Pass
Stock weighting gives a negative carry; flat weightings do not	Pass
Ex-ante compensation less rate surprise equals realised carry	Pass
Ex-ante compensation was positive on almost every acquisition	Pass
Truncated windows are excluded and materially change the answer	Pass
The realised carry changes sign across fully observed horizons	Pass

Generated by `src/verify.py`; the build stops if any check fails.




Figure 2: **Loss absorption.** Panel (a): aggregate national central bank capital and reserves as a share of monetary liabilities. Panel (b): the cross-section in June 2026; four central banks are negative.




Figure 3: **A digital euro rebuilds the secured leg.** Panel (a): the collateralised share as a function of the holding limit, for three values of the share met from own liquidity. Panel (b): the limit each target share would require.




Figure 4: **The model.** Panel (a): the objective in (8) and the interior optimum, with the collateral floor marked. Panel (b): the optimum rises in the term premium and falls in rate volatility, as Proposition 2 states. Panel (c): Proposition 5. With a three-year lending leg and an eight-year portfolio, the balance sheet shrinks while the exposure per euro of it rises.




Figure 5: **The design margin.** Panel (a): the annual income effect of an unremunerated digital euro against the share of take-up that displaces banknotes rather than deposits, at three policy rates. Panel (b): the four corners of (12) at EUR 855 billion and a two per cent rate. Either condition failing alone removes the effect; both failing together reverses it.




Figure 6: **The empirical counterpart to Proposition 4.** Panel (a): ten-year sovereign yields for the four largest purchase programmes against the deposit facility rate; the shaded band is the June 2016 to December 2021 acquisition window. Panel (b): the realised carry at each acquisition date, the acquisition-yield proxy less the realised average funding rate over the following four years, against a break-even of zero. At this horizon every vintage in the acquisition window is complete, so all points are realised carries.




Figure 7: **The mechanism.** Panel (a): reserves, secured credit and the gap between them. Until 2023 secured credit offset roughly half of reserves; after its collapse the gap widened even as the balance sheet shrank. Panel (b): implied carry on the gap against the realised change in capital, which follows a year later as provisions are consumed.




Figure 8: **Intra-Eurosystem positions and what stands behind them.** Panel (a): gross claims, gross liabilities and net positions. Panel (b): the asset mix of debtor central banks, weighted by their net liabilities.