Essays on Macrofinance and Inflation

Inauguraldissertation

zur Erlangung des Grades eines
Doktors der Wirtschaftswissenschaften

durch

die Rechts- und Staatswissenschaftliche Fakultät der
Rheinischen Friedrich-Wilhelms-Universität Bonn

vorgelegt von

Lorenzo Ranaldi

aus Sora (Italien)

2026

Dekan:	Prof. Dr. Martin Böse
Erstreferent:	Prof. Dr. Moritz Schularick
Zweitreferent:	Prof. Dr. Farzad Saidi
Tag der mündlichen Prüfung:	22. Juni 2026

Acknowledgments

This dissertation is the product of years of intellectual and personal support from mentors, coauthors, colleagues, friends, and family. The brief lines that follow cannot do justice to all of them.

I am especially grateful to my supervisors, Moritz Schularick, Farzad Saidi, and Donghai Zhang. Moritz' breadth of vision and his sense for the questions that matter have shaped this dissertation and my approach to research more generally. Farzad's sharp eye for empirical detail, his readiness to engage with ideas at every stage, and his generosity with his time have been an essential complement throughout my doctorate. Donghai's incisive feedback on my work, and his hard-earned advice on how to survive in academia, have been invaluable in seeing this dissertation through. I could not have hoped for a better team of mentors.

I owe a great deal to my coauthors. Chi Hyun Kim and Moritz Schularick made the work behind Chapter 1 a genuine collaboration: I learned from them at every step, and the chapter is better in every dimension for their input. Frederik Bennhoff, Moritz Kuhn, and Florian Scheuer welcomed me into the project that became Chapter 3 with patience and openness, and shared their insights and enthusiasm with me throughout; our conversations have left a lasting mark on how I think about research.

I have benefited enormously from the research environment of the Bonn Graduate School of Economics and the Department of Economics at the University of Bonn. I gratefully acknowledge financial support from the DAAD and the BGSE during my doctoral program, and from CRC TR 224 for funding my visiting period at the University of Chicago.

For comments, conversations, and shared coffees that have improved this work and made the past years a pleasure, I thank Timothy Meyer, Philipp Schirmer, Sebastian Hildebrand, Jörn Frerking, Martin Kornejew, Francisco Amaral, Andrea Amelio, Ricardo Duque Gabriel, Max Grimm, and Marco Forti, as well as the whole team at the Berlin office of the Kiel Institute, who supported me during the job-market phase, and the participants of seminars and conferences where these chapters were presented.

Finally, I am profoundly grateful to my family for their constant love and encouragement, and to Giada, who has been with me through all of it. This work is dedicated to her.

Contents

List of Figures

List of Tables

Introduction

The 2021 inflation surge reopened a set of questions about how inflation propagates through the financial system to asset returns, credit, and real activity — questions the macrofinance literature had largely set aside for decades. Two of the three chapters of this dissertation engage them directly: the first studies how unexpected inflation translates into the real returns on stocks, housing, and bonds, and the second how it propagates through bank balance sheets to the supply of corporate credit. The third turns to households and asks how paper capital gains are converted into realized income — which positions are liquidated, by whom, in response to what, and where the proceeds go. The common thread is the financial balance sheet — of the investors who hold the assets, of the banks that intermediate credit, and of the households that decide when to realize their gains — as the channel through which macroeconomic shocks and household decisions become real outcomes.

Chapter 1, joint with Chi Hyun Kim and Moritz Schularick, examines the response of real asset returns to unexpected inflation in a panel of 18 advanced economies from 1870 to 2023. We construct inflation surprises by digitizing OECD Economic Outlook forecasts back to 1965 and combining them with model-based forecasts that extend the series to 1870. Real returns on stocks, housing, and long-term government bonds decline persistently after an inflation surprise, with the magnitude rising sharply after the 1970s. Using exchange-rate regimes as a source of variation in monetary-policy autonomy, we identify the central-bank response — not the demand- or supply-driven origin of inflation — as the dominant channel: in pegged economies with open capital accounts, where central banks cannot raise rates, equity returns barely move after a surprise; in floating regimes, returns are persistently depressed once policy tightens. The dynamic inflation betas estimated in the chapter capture the long and variable lags of monetary transmission and recover what static measures cannot.

Chapter 2 traces a separate transmission of the same shock, this time through banks. Banks "lend long and borrow short": their assets are dominated by long-dated, fixed-rate loans, while their liabilities reprice more quickly. An unexpected rise in the price level therefore erodes the real value of bank assets relative to liabilities and tightens the constraint on lending. Using long-run macro data for 18 advanced economies and granular bank–firm administrative data from Portugal between 2011 and 2023, I show that a one-percentage-point inflation surprise reduces bank capital and real loans by roughly three percent within five years, and that this effect persists even when monetary policy does not respond. At the bank level, lenders with a one-standard-deviation larger maturity gap cut credit by about two percent; their corporate borrowers, in turn, reduce investment by 0.24 percent of assets within a year. The credit channel of inflation runs counter to the classical debt-inflation channel: inflation may relieve borrowers' real debt burdens, but it tightens the constraint on the financial intermediaries that supply their credit, with measurable consequences for aggregate investment.

Chapter 3, joint with Frederik Bennhoff, Moritz Kuhn, and Florian Scheuer, turns to the household side of the balance sheet. Using three decades of the Dutch Household Survey matched to security-level prices — in a setting without capital-gains taxation, so that selection on winners and the post-sale allocation of proceeds are not confounded by tax planning — we document four facts about how households convert paper gains into realized income. The realization channel is housing-dominated by value and stock-dominated by frequency, and its cross-sectional incidence inherits the asset-holding distribution rather than a behavioral skew. Conditional on selling, households liquidate winners over losers, a pattern that survives a position-level instrumental-variable design and inverts only under shocks sharp enough to force a joint liquidation of the household portfolio — divorce being the cleanest case. The marginal allocation of realized proceeds depends sharply on housing tenure: owners consume stock realizations and channel housing realizations into trade-ups and additional debt, while renters reinvest stock realizations back into the financial portfolio. Realization, not accrual, is the operative margin at which paper wealth becomes realized income.

Read together, the three chapters argue that the structure of balance sheets shapes the magnitude, the heterogeneity, and the persistence of the response to a shock or a decision. The first chapter shows that the monetary-policy regime conditions how unexpected inflation passes through to asset returns. The second shows that the maturity structure of bank balance sheets determines how inflation passes through to credit, independently of any contemporaneous monetary tightening. The third shows that household balance sheets — and in particular housing tenure and life-event shocks — determine which paper gains become realized income, and how those proceeds are spent. Inflation, credit, and capital income do not move in a vacuum: they move through balance sheets, and the structure of those balance sheets is what the three chapters of this dissertation set out to characterize.

1. Inflation Surprises and Asset Returns: A Macrohistory Perspective

Joint with Chi Hyun Kim and Moritz Schularick

1.1. Introduction

The global inflation surge of 2021 was followed by large negative responses in real returns of major asset classes, as shown in Figure 1.1 – a result that traces back to the influential finding of Fama and Schwert (1977). This has revived an old but important question in economics and finance: why do we observe this negative relationship between CPI inflation and asset returns? Despite its importance, the literature has yet to converge on the key driver of this nexus. The main reason for this disagreement lies in the usage of different sample periods and countries, which differ in institutional framework and economic conditions. Moreover, such heterogeneity results in mixed findings in terms of sign, magnitude, and persistence of the relationship between inflation and asset returns, further complicating the search for a comprehensive understanding of this nexus.

Motivated by this heterogeneity, we take a macrohistorical perspective to study the effects of inflation surprises on asset returns. We construct a new dataset of inflation surprises for 18 advanced economies from 1870 to 2023, where we define inflation surprises as the unanticipated difference between realized and expected inflation. We use this dataset to analyze how unexpected inflation affects real returns of three core asset classes: stocks, housing, and long-term government bonds. We digitize archival inflation forecasts from the OECD Economic Outlook for 18 countries, dating back to 1965, covering around 60 years of data. In addition, we estimate inflation expectations based on macroeconomic and financial variables of the Jordà-Schularick-Taylor Macrohistory Database during 1870-2023 using a forecasting model that combines various time series methods. This allows us to increase the sample window considerably. We show that inflation forecasts from these two sources are consistent with the major narrative evidence and mostly overlap in the recent sample.

Using this new dataset, we first estimate the so-called *inflation beta*, which measures the contemporaneous effect of inflation surprises on asset returns (Bekaert and Wang, 2010). A key contribution of our analysis, however, is to move beyond this static measure

Figure 1.1.: Asset price developments during the 2021 inflation surge

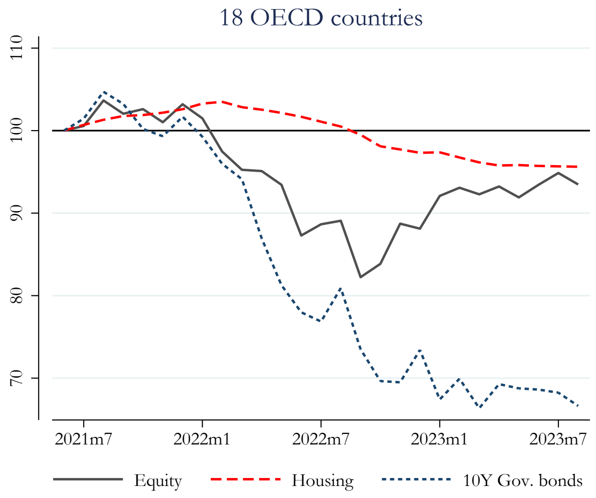

Notes: Real indices for equity, housing and 10-year government bonds from June 2021 to August 2023. The lines present the average of 18 OECD countries, i.e. Australia, Belgium, Canada, Switzerland, Germany, Denmark, Spain, Finland, France, UK, Ireland, Italy, Japan, Netherlands, Norway, Portugal, Sweden, and the USA. All indices have been adjusted for inflation using CPI and normalized to 100 in June 2021. Equity: National stock market total return indices. Housing: National house price indices (OECD database). Bonds: Total return indices, 10Y government bonds. Source: Datastream.

and study the persistence of this effect. By applying panel local projection methods, we estimate impulse response functions of cumulative asset returns to inflation surprises, which we interpret as *dynamic inflation betas*. This dynamic perspective reveals that real stock returns respond negatively not only on impact but also over the medium run, with only partial recovery in the long run. Government bond returns display a similar pattern, implying that the real stock–bond correlation becomes more positive after inflation surprises.[1] Housing offers a better but still incomplete hedge: while real returns are resilient on impact, they decline persistently in the medium to long run.

How can we explain this negative persistent relationship between inflation and real asset returns? Our long-run analysis indicates that the negative response of real assets intensified after the 1970s, a period during which short-term interest rates exhibited a stronger reaction to inflation surprises. This pattern points to monetary policy as a key driver of the inflation-asset returns nexus, in line with the literature that highlights the general importance of monetary policy in asset price fluctuations (Bernanke and Gertler, 2000; Rigobon and Sack, 2004; Gurkaynak, Sack, and Swanson, 2005). To assess this channel, we exploit the long-run variation in exchange rate regimes in our sample. Specifically, we draw on the trilemma of international finance (Obstfeld and Taylor, 1997; Jordà, Schularick, and Taylor, 2020) and compare the effect of inflation on asset returns in (i) pegged economies with open capital accounts (*open pegs*), where central banks are constrained and cannot freely

[1] This result is at the core of the debate about the changing co-movement of these two asset classes over time (see e.g., Campbell, Pflueger, and Viceira (2020) and David and Veronesi (2013)).

adjust interest rates, and (ii) floating exchange rate regimes, where monetary authorities retain policy autonomy. Our results confirm the dominant role of monetary policy: in open peg economies, cumulative equity returns barely move after an inflation surprise, while returns are *persistently* affected once the central bank raises interest rates. The absent central bank response prevents the economy from falling into a recession, leading to a more favorable outlook for future income flows such as corporate profits, dividends, and rents than in floating economies. This finding supports the proxy effect hypothesis of Fama (1981): returns on real physical assets decline in response to inflation because higher inflation proxies for weaker *future* economic activity. We show that the conduct of monetary policy—aimed at reducing inflation—is the driving force behind this proxy effect. Our empirical results thus support the growing theoretical literature showing how monetary policy generates persistent effects on asset prices through regime changes in policy rules combined with adaptive expectations (Bianchi, Lettau, and Ludvigson, 2022) or forward-looking stabilization that ties asset prices to expected future macroeconomic shocks (Caballero and Simsek, 2024).

An alternative explanation attributes the inflation–asset return relationship to the source of inflation: demand-driven surprises raise returns through stronger economic activity and cash flows, while supply-driven surprises depress them via higher costs and weaker output (Cieslak and Pflueger, 2023). Shifts in this relationship may thus reflect changes in the mix of demand- versus supply-driven shocks rather than monetary policy alone. We test this channel by exploiting the joint correlation between inflation surprises and GDP surprises in pegged countries, where short-term interest rates do not respond to inflation. Similar to our inflation surprise measure, we construct GDP surprises for 18 OECD countries over 1870-2023 using model- and OECD-based forecasts. We find no significant long-run differences in how real physical assets respond to demand- versus supply-driven surprises. On impact, returns do react more positively to demand shocks, but in both cases the overall effect is negative. This can be explained by two factors: first, cashflows adjust only gradually to a change in GDP (Leary and Michaely, 2011). Second, after a supply shock, GDP falls initially but recovers from $t + 1$ onward, while after a demand shock, GDP booms on impact but slows down thereafter. The initial benefit of demand shocks is therefore offset by the slow cashflow adjustment coupled with weaker *future* GDP growth. These factors can rationalize both the differential impact response and the similar cumulative response of asset returns after a demand- vs. supply-driven inflation surprise. This result challenges the conclusion of prior studies claiming the source of inflation being the key driver of the inflation-asset returns relationship (Hess and Lee, 1999; Fang, Liu, and Roussanov, 2022; Cieslak and Pflueger, 2023).

Our paper contributes to the literature in three different dimensions. The first contribution stems from our new long-run dataset of inflation expectations and inflation surprises, which serves as a new source for studying inflation dynamics in the long run for various developed economies. As commonly known, inflation expectations data start only from the late 20th century.[2] Previous literature estimated inflation expectations for historical samples using models or financial markets data for limited time- or country-samples (Hamilton, 1992; Cieslak and Povala, 2015; Ellison, Lee, and O'Rourke, 2024). We digitized archival inflation expectations from the OECD Economic Outlooks, and complement them with forecasts from an auxiliary model that replicates those estimates. Our final inflation surprise series is available for researchers[3] and opens the door for future research on inflation expectations proxies in both modern and historical contexts, which is proved to be essential in studying inflation dynamics (Schmitt-Grohé and Uribe, 2024).

Second, our macro-historical approach clarifies key macroeconomic channels driving the inflation-asset returns relationship by exploiting various regime shifts in policy and economic conditions. So far the literature has focused on two distinct macroeconomic channels.[4] First, Fama (1981) and Geske and Roll (1983) provide evidence on how inflation serves as a proxy for deteriorating economic conditions through monetary aggregates. Alternatively, Hess and Lee (1999) and Cieslak and Pflueger (2023) emphasize the role of the underlying demand and supply disturbances of inflation. Our results show that the primary macroeconomic channel driving the relationship between inflation and asset returns is monetary policy. Our study is the first to supply empirical evidence on this channel despite its relevance, as previously available macroeconomic datasets have precluded a formal test. Specifically, we are the first to explicitly model cumulative responses over horizons of up to five years – long enough to capture the "long and variable lags" of monetary policy (Friedman, 1961) that are essential for pinning down this channel empirically. Our results thus have important implications for central banks, highlighting how monetary policy reactions to inflation spikes can significantly and persistently disrupt financial markets.

Finally, our long-run cross-country dataset allows us to comprehensively assess not only the sign and magnitude but also the *persistence* of the effect of inflation on asset returns. While Fama and Schwert (1977) were the first to document a negative contemporaneous relationship between unexpected inflation and stock returns, later studies such as Boudoukh and Richardson (1993) emphasized that the correlation turns positive at longer horizons. Other contributions highlight high-/low-inflation regime dependence (Barnes, Boyd, and Smith, 1999), cross-country heterogeneity (Engsted and Tanggaard, 2000; Bekaert and Wang, 2010), or asset-specific differences, such as housing as a partial hedge (Fehrle, 2023). More recent work shows that the sign of the response turned positive in the post-2000s (Chaudhary and Marrow, 2022), that energy and core inflation matter differently (Fang,

[2]The earliest comes from the Livingston survey that provides inflation forecast data for the USA from 1951 onward. For other countries, similar data start only in the 1990s (e.g., Consensus economics, central bank forecasts) or they are only in a qualitative format (e.g., EU Business and Consumer survey).

[3]See e.g., Federle, Mohr, and Schularick (2024) for the use of our series in a political economy context.

[4]There exists also behavioral explanations of the negative relationship between inflation and asset returns linked to investors' discount rates. The so-called money illusion (Modigliani and Cohn, 1979) suggests that investors incorrectly discount real cash flows using nominal rates during inflation episodes. Katz, Lustig, and Nielsen (2017) similarly argue for investors' slow adjustments to changing inflation expectations. We do not directly address these channels given the aggregate nature of our data, but we provide suggestive evidence that discount rate channels play only a minor role in the overall negative response.

Liu, and Roussanov, 2022), and that markets interpret inflation as a cost shock for firms (Knox and Timmer, 2025). Our contribution is to go beyond static betas by estimating dynamic long-run inflation betas, which reveal that the effect of inflation surprises on asset returns is not only negative on impact but remains persistently so across time and countries, with the magnitude of this effect increasing substantially since the 1970s.

The structure of the paper is as follows. In Section 1.2, we present the construction of our data on inflation surprises and asset returns. In Section 1.3 we investigate the effects of inflation on asset returns. Section 1.4.1 presents results on the pivotal role of monetary policy. Section 1.4.2 analyzes the real effects of demand and supply sources, while Section 1.5 concludes.

1.2. Data

1.2.1. Construction of long-run surprises

In order to analyze the effect of inflation on asset prices, we need to identify unanticipated changes in inflation rates, namely "inflation surprises". In the spirit of Auerbach and Gorodnichenko (2012b) and Auerbach and Gorodnichenko (2013), we identify this as the difference between actual inflation and inflation expectations – the component of inflation that economic agents were not able to predict given the information on economic fundamentals. Given our focus on the long-run and across-countries, our main challenge is thus to find a consistent measure of inflation expectations that goes back 150 years in history for 18 OECD countries. Survey data are unsuitable for this purpose since, apart from the notable exception of the US, no developed country has forecasts available before the 1990s. Even in the US, both consumer and professional forecasts are unavailable before the 1950s.[5] Also, implicit inflation expectations based on financial instruments traded in financial markets, like futures or TIPS, face the problem that their prices or rates include more than implied inflation expectations, e.g. default risk or liquidity risk.[6]

OECD Economic Outlooks

We digitize historical data from archived OECD Economic Outlook reports, extracting semi-annual forecasts (June and December) of 1-year-ahead inflation covering the extensive period from 1965 to 2023. This newly digitized dataset allows us to investigate the relationship between inflation and asset returns for a rich set of countries for more than 50

[5] In the US, the Survey of Professional Forecasters (SPF) starts in 1968. The Michigan consumer survey asks about prices in the 1960s. The Livingston survey begins after WWII.

[6] For example, Treasury inflation-protected securities (TIPS) suffered, especially in their early days, from scarce liquidity in the market, implying that their rates incorporated a significant liquidity premium which is hard to estimate (e.g. Gürkaynak, Sack, and Wright (2010)).

years.[7] Apart from the long time span, one big advantage of using these series is that all forecasts are prepared using a unified methodology such that the series are comparable across countries. Furthermore, the OECD leverages its local presence to engage with government experts and policymakers, resulting in forecasts that incorporate significant local knowledge and insights into future policy changes. Reviews of the OECD's forecasts by Vogel (2007) report that they are comparable to private sector forecasts and have a number of desirable properties.

Specifically, we digitize 1-year-ahead inflation forecasts using the Consumer Price Index (CPI) inflation or the Gross National Product (GNP)/Gross Domestic Product (GDP) deflator as primary indicators. We always use the first in our estimation, and include the latter in the rare cases when CPI forecasts are not available. The dataset encompasses inflation forecasts for the following 18 countries: Australia, Belgium, Canada, Switzerland, Spain, Finland, France, Denmark, Germany, Ireland, Italy, Japan, Netherlands, Norway, Sweden, the United Kingdom, and the United States.[8] We present the forecasts in Appendix Figures A.1 and A.2. In addition, we also digitize 1-year-ahead GDP/GNP growth forecasts from the same OECD Economic Outlook documents with same country coverage, which we later use in Section 1.4.2 to disentangle the effect of demand- and supply-driven inflation.

A model of inflation expectations

As an auxiliary model to extend our sample into the late 19th century, we apply a time series approach to forecast future inflation rates. Although forecasts may not be equivalent to expectations, time series approaches use the former as a proxy for the latter, mimicking what a rational investor would expect. Various studies have used models to estimate inflation expectations in historical settings: Ellison, Lee, and O'Rourke (2024) estimate inflation expectations for many countries during the Great Depression by training a dynamic factor model on historical data; Cieslak and Povala (2015) estimate inflation expectations as a simple dynamic moving average (DMA) of past core inflation to reflect how sluggishly individuals update their expectations over time.[9]

[7]So far, only part of these forecasts were made public through the OECD's Statistics and Projections database (e.g. Auerbach and Gorodnichenko (2012a) use this dataset to exploit GDP and government spending forecasts starting in 1985). Data on inflation forecasts that are digitized and publicly available on the OECD database at `https://data.oecd.org/price/inflation-forecast.htm` overlap almost entirely with ex-post realized inflation, such that they should not be used as a reliable real-time source. Indeed, the forecast series that we digitized from the Economic Outlook significantly diverge from these values.

[8]Country coverage differs by "core" and remaining countries. 1965-2022: Canada, France, Germany, Italy, Japan, United Kingdom, United States. 1974-2022: Australia, Belgium, Netherlands, Denmark, Ireland, Finland, Norway, Sweden, Switzerland, Spain.

[9]Even though investors in the early 20th century may not have explicitly used time series methods to forecast inflation before these methods were invented, they certainly used rules of thumb that in principle resemble the results of such methods. In its *Theory of Interest*, Fisher (1930) hints at such implicit mechanism in the common business man, as he writes "*Few business men have any clear ideas about it. [...] Yet it may be true that they do take account, to some extent at least, even if unconsciously, of a change in the buying power of money [...] If inflation is going on, they will scent rising prices ahead [...] And today especially, foresight is clearer and more prevalent than ever before. The business man*

If the rationality imposed in forecasting inflation may seem too strict for a common economic agent in the 20th century economy, it has instead more solid ground in view of the task at hand. When considering asset classes traded in financial markets, such as stocks or bonds, only few investors were able to participate in those markets and are likely to process macroeconomic information in a close-to-rational manner when forming their future expectations. However, there might be several possible expectation measures among investors in a particular time, such that the resulting market-wide inflation expectations are a weighted average of this distribution.

To address the latter point, we combine different forecasting techniques in order to significantly reduce model misspecification and bias (Timmermann, 2006; Elliott and Timmermann, 2016) instead of relying on one single forecast method. Such procedure is mimicking what is normally carried out in the US Survey of Professional Forecasters, where different forecasters use their own (simple or more complicated) model and each estimate is then aggregated into a single measure of market expectations. Hendry and Clements (2004) show that the pooling of forecasts dominates other estimates even in presence of structural breaks, which have taken place throughout history for inflation expectations (Bataa et al., 2013).

Our model can be written as

$$\mathbb{E}_t[\pi_{i,t+1}] = \boldsymbol{\omega'} \quad \boldsymbol{\hat{\Pi}_{t+1|t}} = \boldsymbol{\omega'} \begin{pmatrix} \hat{\pi}^1_{t+1|t} \\ \vdots \\ \hat{\pi}^j_{t+1|t} \\ \vdots \\ \hat{\pi}^n_{t+1|t} \end{pmatrix} \tag{1.1}$$

$$\hat{\pi}^j_{t+1|t} = \boldsymbol{A'} \quad \boldsymbol{X_{t:t-k}} + \varepsilon^j_t. \tag{1.2}$$

For any model $j = 1, \ldots, n$ that we use in the pooling exercise, where $\boldsymbol{\omega}$ is the vector of weights assigned to each single model, $\boldsymbol{\hat{\Pi}_{t+1|t}}$ is the vector containing all inflation forecasts $\hat{\pi}^j_{t+1|t}$ for each model j. $\mathbb{E}_t[\pi_{i,t+1}]$ is then the resulting forecast obtained by combining those estimates. We estimate a variety of univariate and multivariate time-series and panel data models, the latter to exploit possible spillover effects for inflation between countries, e.g. imported inflation. We present the complete list of models in the Appendix Table A.1 and Table A.3. These are compactly represented in the matrix $\boldsymbol{X_{t:t-k}}$, which takes information from present and past values of various macroeconomic variables.[10]

makes a definite effort to look ahead not only as to his own particular business but as to general business conditions, including the trend of prices". For more on the relation between inflation expectations and realized inflation in historical terms, see for example Binder and Kamdar (2022).

[10]Given the emphasis in this literature on equally-weighted forecasts as best combination, we use both

All data to estimate inflation expectations and surprises are taken from the Jordà-Schularick-Taylor (JST) Macrohistory database by Jordà, Schularick, and Taylor (2017) and Jordà, Knoll, et al. (2019). This dataset contains yearly data of main aggregate macroeconomic variables such us GDP, inflation, interest rates and aggregate asset returns for 18 advanced economies for a period that spans from 1870 until 2020.[11] Since we want to include the recent inflation surge during the COVID-19 pandemic, we extend the JST database for the years 2021-2023. We present the resulting inflation expectations for the full historical sample in Appendix Figures A.3 and A.4. Appendix Figure A.5 presents all twenty inflation forecasts estimated with the single-models that we use to construct our benchmark inflation forecasts for the USA. Finally, we compare the OECD forecasts and our model forecasts for all countries in Appendix Figures D.3 and D.4, finding a correlation of 0.95+ between the two measures in the overlapping post-1965 sample.

Inflation surprises

Our main variable of interest – the unexpected component of inflation – is the residual between realized inflation and inflation expectations one year before, at time $t-1y$, for each country i. We define it as *inflation surprise* – a measure of inflation that could not be predicted in the previous period:

$$\varepsilon^{\pi}_{it} = \pi_{it} - \mathbb{E}_{t-1y}[\pi_{it}] \tag{1.3}$$

Figure 1.2 presents our post-1965 inflation surprises of all countries derived from the OECD outlooks (blue dots), together with the average across the 18 countries (red dots). In summary, only one period over the past six decades was comparable to the recent inflation spike in terms of significant inflation surprises: the Great Inflation of the 1970s and early 1980s. Since then, inflation surprises have been relatively minor and more consistent across different countries. Appendix Figure A.6 and Figure A.7 present the inflation surprise series of all 18 countries, together with realized inflation rates. Consistent with Blanco, Ottonello, and Ranošová (2025), we observe how large surges in inflation/deflation are initially unexpected, followed by a gradual catch-up. Since the Great Moderation, surprises in inflation in all countries seem to have experienced a significant reduction in volatility overall and across countries.

In Figure 1.3 we present the full sample model-based inflation surprises during 1870-2023 (left panel). Our long-run view shows how inflation surprises have been overall larger during the pre-1965 period, with large surprises being mostly concentrated around world

mean and median forecast, finding comparable results. We prefer the latter as to avoid outliers driving results. In Appendix 1.D.1 we visualize the difference between the mean and median forecast measure, producing very close estimates.

[11] The 18 advanced economies are the USA, Japan, Germany, France, United Kingdom, Ireland, Italy, Canada, Netherlands, Belgium, Sweden, Australia, Spain, Portugal, Denmark, Switzerland, Finland, Norway. See Appendix 1.B for a detailed description of the dataset.

Figure 1.2.: OECD-based inflation surprises: 1966-2023

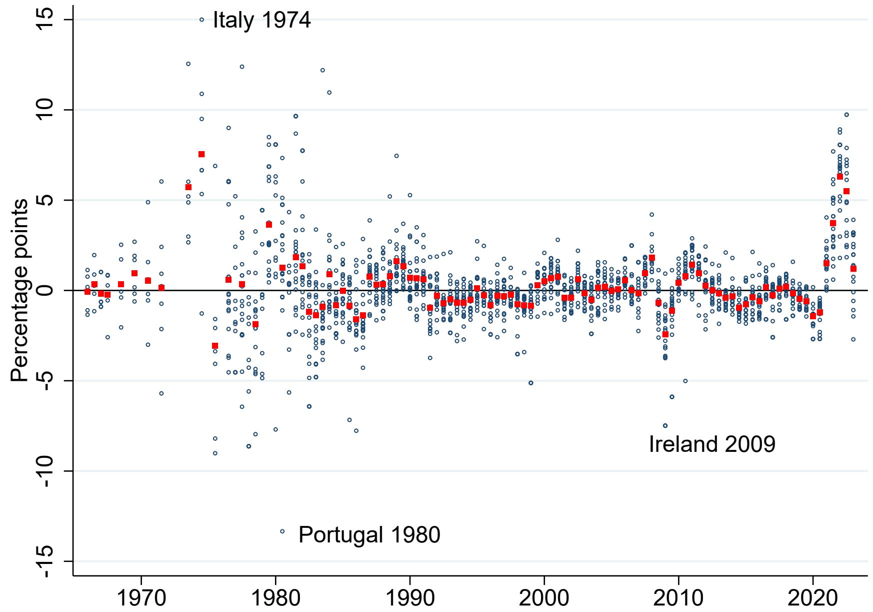

Notes: Surprises from OECD 1y-ahead forecasts, half-yearly frequency, 1966-2023. Inflation surprises are computed as the difference between realized inflation and expectations from the previous period, i.e. $\varepsilon^{\pi}_{it} = \pi_{it} - \mathbb{E}_{t-1y}[\pi_{it}]$. Inflation is measured in %, surprises in percentage points. Red dots represent the average surprise across the 18 OECD countries. Inflation surprise series of each of the countries are presented in Figure A.6 and Figure A.7 in the appendix.

wars, with spikes higher than 20 percentage points. Moreover, the volatility of inflation surprises seems to have decreased significantly in the post-WW2 period. To test for the validity of our model-based inflation surprise measures, we compare them with the inflation surprise measures derived from inflation forecasts of the OECD for the post-1965 period (right panel of Figure 1.3). Reassuringly, we observe a strong correlation betwen the two inflation surprise measures (correlation ≈ 0.75).

Narrative validation of our historical model-based inflation surprise measures Our model-based measures need external validation, especially for the pre-1965 period where OECD data are not available. We do so by zooming in on the Great Depression period during the 1930s for Germany, the UK, and the US. This historical period is particularly interesting in terms of inflation expectations and their discrepancy from realized inflation, as the interwar period was characterized by dramatic regime shifts – most notably the abandonment of the gold standard — that fundamentally altered the inflation outlook in ways that agents struggled to anticipate. As Ellison, Lee, and O'Rourke (2024) document using dynamic factor models to estimate monthly inflation expectations for 27 countries, leaving the gold standard triggered sharp increases in inflationary expectations and declines in real interest rates, with the shift in expectations driven primarily by rising inflation forecasts rather than falling nominal rates. Crucially, the timing and nature of these

Figure 1.3.: Model-based inflation surprises: 1870-2023

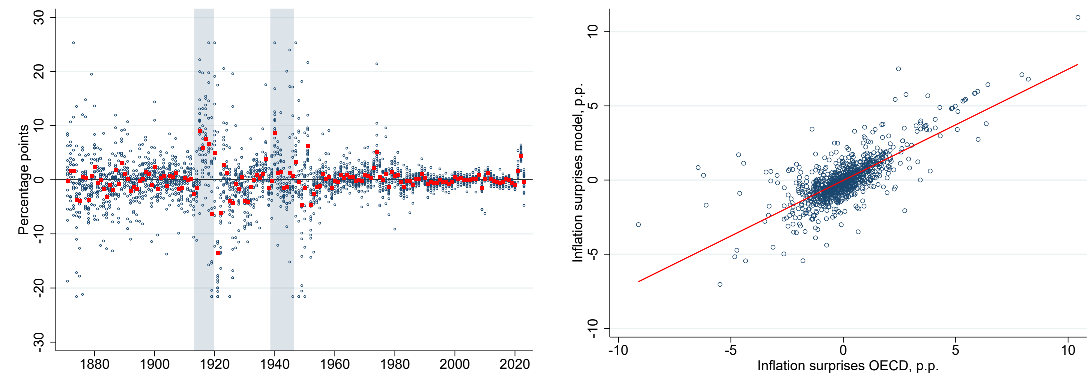

Notes: Left panel: Surprises from model 1y-ahead forecasts, yearly frequency, 1870-2023. Right panel: scatterplot of model surprises and OECD surprises (end-of-the-year), yearly frequency, 1965-2023. Inflation surprises are computed as the difference between realized inflation and expectations from the previous period, i.e. $\varepsilon_{it}^{\pi} = \pi_{it} - \mathbb{E}_{t-1}[\pi_{it}]$. Inflation is measured in %, surprises in percentage points. Red dots represent the average surprise across the 18 OECD countries. Inflation surprise series of the other countries are presented in Figure A.6 and Figure A.7 in the appendix.

expectation shifts varied across countries: in the UK and US, departure from gold produced relatively rapid reversals in inflation expectations, while in Germany, which imposed exchange controls but never formally abandoned the gold parity, the turning point was driven instead by political events, namely the suspension of reparations at the Lausanne Conference and Hitler's accession to power in January 1933. These cross-country differences in the size and timing of the gap between expected and realized inflation make the 1930s an ideal testing ground for our model-based inflation surprise measures.

Figure 1.4.: Model-based inflation surprises during the 1930s: Germany, UK, and USA

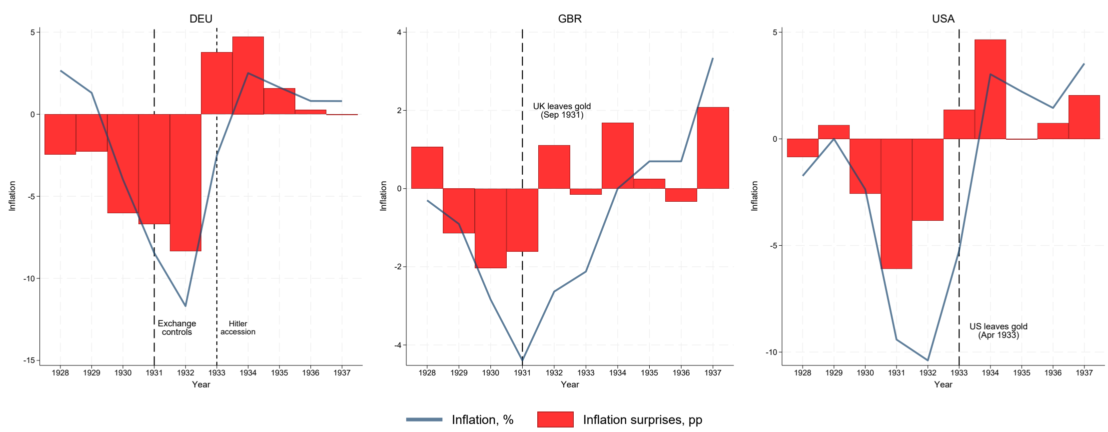

Notes: Inflation surprises (red bars) and realized inflation (blue line) for Germany, UK, and USA.

Figure 1.4 presents realized inflation and our model-based inflation surprises for Germany, the UK, and the US during 1928–1937. The dynamics are broadly consistent with the monthly estimates of Ellison, Lee, and O'Rourke (2024). In all three countries, inflation surprises are deeply negative during the early Depression years, indicating that the severity of deflation consistently exceeded what agents anticipated. The trough in surprises coincides with the trough in realized deflation around 1931–1932, after which a sharp reversal occurs. However, the nature and timing of this reversal differ across countries, reflecting distinct policy regimes. In the UK, surprises turn modestly positive in 1932, the year after departure from the gold standard, consistent with Ellison, Lee, and O'Rourke (2024)'s finding that the monetary regime shift provided a clear signal allowing agents to partially adjust expectations. The US displays a more dramatic pattern: surprises reach roughly -6 percentage points by 1931 before swinging sharply positive after departure from gold in April 1933, suggesting that the inflationary impact of leaving gold was larger than anticipated. Germany presents the starkest contrast. Despite imposing exchange controls in July 1931, inflation surprises continued to deteriorate, consistent with the finding that exchange controls without formal gold abandonment failed to shift expectations. The positive surprise in 1933 instead coincides with Hitler's accession to power and the suspension of reparations, i.e., political rather than monetary regime changes. Overall, the cross-country ordering of surprise magnitudes during the downturn mirrors the relative severity of the Depression, providing reassuring external validation for our model-based measures in the pre-OECD sample.

1.2.2. Data on asset returns

We focus on three asset classes: equity, housing and long-term government bonds. We gather data on their total returns at both semi-annual and annual frequency. We use annual data as our benchmark and semi-annual data as robustness exercises in Appendix, as OECD-based inflation surprises post-1965 are available at semi-annual frequency, while the model-based inflation surprises during 1870-2023 have annual frequency.[12] We use annual data from the Jordà-Schularick-Taylor Macrohistory Database (Jordà, Knoll, et al. (2019)).[13] We employ the total returns for housing and stocks, which include both capital gains and yields (dividends and rents, respectively), while for government bonds we target the total returns of long-term bonds (mostly at the 10 year maturity). We compute real returns by deflating the series with CPI. When analyzing returns at annual frequency, we only use end-of-the-year inflation surprises from the OECD outlooks. Finally, we add data on corporate profits from the World Inequality Database (Alvaredo et al., 2021).[14]

[12] For semi-annual data, we utilize the stock market index of Monnet, Puy, et al. (2016) and Monnet and Puy (2021) (MP), who digitize the quarterly measure from the IFS statistics, which we convert into semi-annual frequency to be matched with our OECD surprise series. The house price index is downloaded from the OECD online database.

[13] While Jordà, Knoll, et al. (2019) provides only data until 2015, the Jordà-Schularick-Taylor (JST) Macrohistory database by Jordà, Schularick, and Taylor (2017) has extended series until 2020. We further extend this series until 2023 using data for indices of house prices, rents, stock market and 10-year government bonds from IMF, OECD and Refinitiv.

[14] Specifically, we use as profits the corporate gross operating surplus. For most countries, the data are available starting in 1900.

Table 1.1.: Real asset returns, mean (and standard deviation), in percent, 1870-2023

	Overall	Pre-WWI	1919-1938	World Wars I-II	Post-WWII	Post-1982
Equity	4.27	5.64	5.28	-3.87	4.99	7.20
	(21.76)	(11.04)	(21.96)	(29.13)	(23.65)	(22.75)
Housing	6.18	6.98	7.19	0.70	6.51	6.08
	(9.09)	(8.03)	(11.03)	(14.44)	(7.67)	(6.16)
Bond	1.28	3.85	6.07	-12.03	1.05	3.59
	(13.10)	(6.83)	(14.30)	(24.72)	(11.05)	(10.81)

Notes: Average real annual returns, in percent, with standard deviation in parenthesis, for each subsample. Sample: 18 countries from the JST database, 1870-2023, yearly frequency. We exclude outlier years, i.e. hyperinflation periods.

Table 1.1 reports the average returns and its standard deviations for different historical subsamples. Historically, equity and housing returns are not very different in real terms, even though equity appears to be more profitable starting from the Great Moderation. Government bonds returns are significantly lower compared to equity and housing. Finally, government bonds and housing returns have significantly lower volatility than equity.

1.3. Inflation surprises and asset returns

In this section, we investigate the effect of inflation surprises on real asset returns. We start by focusing on the post-1965 period, where we are able to derive inflation surprises using inflation forecasts of the OECD. In Section 1.3.1, we estimate the so-called *inflation beta*, which captures the contemporaneous effect of inflation on asset returns. In Section 1.3.2, panel local projection methods are employed to investigate the dynamic effect on cumulative returns, i.e. a *dynamic* inflation beta, highlighting the persistent consequences of inflation on future asset returns. Finally, in Section 1.3.3 we take a long-run perspective and complement our OECD-based inflation surprises with model-based inflation surprises to analyze the full 1870-2023 sample. By doing so, we investigate the time-varying role of inflation on asset returns during the last 153 years.

1.3.1. Contemporaneous effect of inflation surprise on asset returns

We begin by conducting a static analysis, where we regress real asset returns on inflation surprises to examine the immediate impact of an unexpected increase in inflation on asset returns – measuring the so-called *inflation beta.* This approach has been the most commonly adopted framework in the literature, from Fama and Schwert (1977) to more recent studies such as Bekaert and Wang (2010) and Fang, Liu, and Roussanov (2022). The model we estimate is the following:

$$r_{it}^k = \beta^k \varepsilon_{it}^\pi + \gamma^k \mathbf{X}_{it} + \mu_i + \epsilon_{it}^k, \tag{1.4}$$

where r_{it}^k is the real annual return of asset k in country i at time t, with k being either equity, housing, or long-term government bonds, and ε_{it}^π is the inflation surprise of the same country at time t. Note that most of the inflation betas estimated in the previous literature regress on realized inflation instead of inflation surprises, therefore conducting an *ex-post* analysis. See for examples Bekaert and Wang (2010).[15] $\mathbf{X}$ contains a set of controls to isolate the effects of business cycles and economic conditions, such as 2 lags of ε^π, dependent variable, short-term interest rates, GDP growth, real stock price, inflation, world GDP growth & world equity price.[16] We also include country fixed-effects μ_i. The parameter of interest is β^k, which is defined as the inflation beta for asset k. $\beta^k = 0$ would imply that asset k is a perfect inflation hedge, meaning that real returns of asset k fully offset movements in the inflation rate.

Table 1.2.: Inflation beta on real returns, 1965-2023

	Equity		Housing		10Y Gov Bond	
	(1)	(2)	(3)	(4)	(5)	(6)
OECD-based	-4.34***	-2.93**	0.01	-0.11	-2.75***	-2.36**
	(0.90)	(1.28)	(0.25)	(0.25)	(0.89)	(0.65)
Model-based	-4.60***	-3.81***	-0.28	-0.00	-2.88***	-2.48***
	(1.04)	(1.16)	(0.22)	(0.23)	(0.90)	(0.50)
Country FE	No	Yes	No	Yes	No	Yes
Controls	No	Yes	No	Yes	No	Yes

Notes: The first row presents the estimated inflation betas for the period 1965-2023 using OECD-based inflation surprises. Estimates are based on returns and data at yearly frequency, using end-of-the-year OECD surprises. In the second row, we present the inflation betas estimated with model-based inflation surprises during 1965-2023. Column (1), (3), and (5) are inflation betas estimating Equation 1.4 without country fixed-effects and controls. Column (2), (4), and (6) present the inflation betas by estimating the full model of Equation 1.4. Controls include: 2-years lags of ε^π, dependent variable, short-term interest rates, GDP growth, real stock price, inflation, world GDP growth & world equity price. Driscoll and Kraay (1998) standard errors in parentheses. *, **, and *** present the significance levels at 10%, 5%, and 1%, respectively.

In Table 1.2 we present the estimated inflation betas for the post-1965 period utilizing both our OECD-based and model-based inflation surprises. Columns (1), (3), and (5) present the results where we only regress our return measures on inflation surprises without country fixed-effects and controls, while columns (2), (4), and (6) show the results with the full set of controls. Overall, inflation betas show a significant negative relationship

[15] We run a similar correlation analysis in Appendix 1.C.

[16] The latter is to conservatively control for global business cycle conditions without estimating time fixed effects, which are unfeasible in our setting with a rich set of controls. This is a standard approach in studies using this long T-short N dataset, see e.g. Jordà, Schularick, and Taylor (2020).

between inflation and returns for both equities and government bonds, while the housing betas are only insignificantly negative and close to zero, i.e. a near-perfect inflation hedge. The inflation beta for equities declines significantly more than one-to-one, with a beta around -4, whereas for government bonds the estimate falls below -2, indicating a strong inverse relationship. In contrast, the inflation beta for housing is close to 0. Reassuringly, these results remain consistent when we estimate the inflation betas using our model-based inflation surprises for the post-1965 sample (see row 2 of Table 1.2).

Our estimated inflation betas are qualitatively consistent with others from the literature: Fama and Schwert (1977) first estimated strongly negative unexpected inflation beta for stocks, around -2. Bekaert and Wang (2010) estimate an inflation beta close to zero for *nominal* returns for bonds and stocks, which are in line with the negative inflation betas in *real* terms in Table 1.2. Fang, Liu, and Roussanov (2022) estimate a similar specification to Equation (1.4) using a VAR model for the US to get inflation expectations and estimate inflation beta for *excess* returns. They find strongly negative betas for stocks (-1.33) and bonds (-2.53), while housing has a slightly positive but not significant response (0.31), roughly in line with the estimates of Table 1.2.

1.3.2. Dynamic inflation beta: panel-LPs on cumulative returns

To highlight the importance of persistence in the negative effects of inflation on asset returns, we now depart from the static framework widely employed in the literature and implement a panel local projections model à la Jordà (2005) to estimate the dynamic response of real asset returns to our inflation surprise measures, i.e. a *dynamic* inflation beta. The panel local projections (panel-LPs) estimate the average impact of a 5 percentage point surprise increase in inflation[17] on the (log) return of each asset k:

$$r^k_{i,t+h} - r^k_{i,t-1} = \beta^k_h \varepsilon^\pi_{it} + \Psi^k_h(L)\mathbf{X}_{i,t-1} + \delta^k_{i,h} + \epsilon^k_{i,t+h} \qquad \text{for} \quad h = 0, 1, 2, \dots, \tag{1.5}$$

where $r^k_{i,t+h} - r^k_{i,t-1}$ represents the *cumulative* real return over h periods for asset k and ε^π_{it} is the inflation surprise for country i at time t.[18] $\mathbf{X}_{i,t-1}$ is a vector of controls, which includes 2-year lags of ε^π, dependent variable, short-term nominal interest rate, real GDP growth, real stock price, CPI inflation, world GDP growth and world equity price (to parsimoniously control for world business cycles). $\delta_{i,h}$ are horizon-specific country fixed effects. Lag augmented shocks and controls are included to account for any serial correlation and to improve inference, as prescribed by Montiel Olea and Plagborg-Møller

[17] We scale the surprise to induce a 5 p.p. increase in inflation to simulate a scenario akin to the notable inflation surge experienced in 2022. Moreover, estimates in Blanco, Ottonello, and Ranošová (2025) produce a 4.9 p.p. increase in inflation after a "large" inflation surge.

[18] To make LP estimations on returns more meaningful, we transform the total annual returns of each asset into a *total return index*. In this way, the h-long difference Δ_h of the index is directly interpretable as cumulative (log-)return over h periods.

(2021). The inflation level is included to account for higher volatility of inflation and its unexpected component when inflation is high. Stock prices are included to capture possible high frequency information implicitly incorporated in stock markets and not yet contained in inflation forecasts. Other macroeconomic variables are included to control for country and world business cycle dynamics.[19]

Figure 1.5.: Dynamic inflation betas, 1965-2023

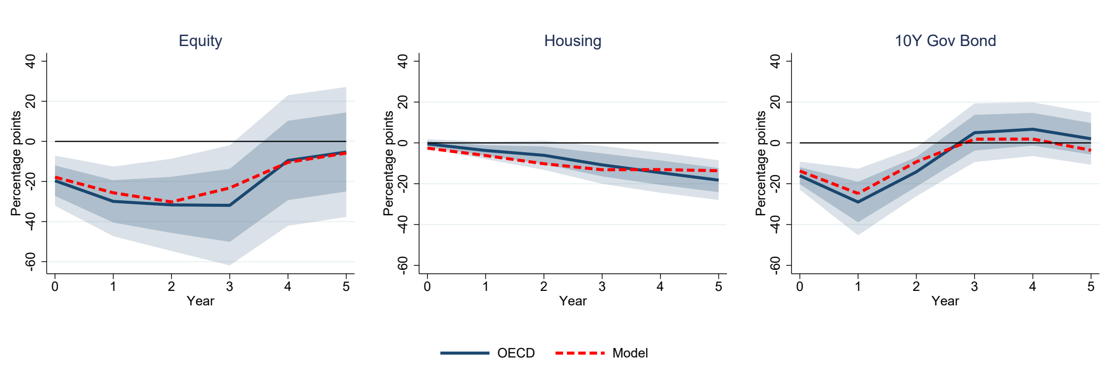

Notes: Impulse responses for cumulative real asset returns estimated with panel-LP across two sources of surprises: OECD surprises and model surprises during 1965-2023. All results are at yearly frequency, using end-of-the-year OECD- or model-surprises. Effects scaled to a 5 p.p. increase in inflation. Driscoll and Kraay (1998) standard errors. Bands are at the 68% and 90% confidence level.

Figure 1.5 reports the dynamic inflation betas of stocks, housing and bonds for the period 1965-2023. Given that in the post-1965 sample both surprise series are available, we compare the dynamic inflation betas of asset returns coming from model-based inflation surprises (red dashed line) with the OECD-based inflation surprises (blue solid line), with comparable results. In line with our estimated inflation betas, we observe a contemporaneous decrease in real returns of equity and long-term government bonds after an inflation surprise. They continue to lose value over time, but equity recovers in value after 4-5 years, while government bonds fully recover after 2-3 years. Compared to this effect, housing real returns are not strongly affected by inflation on impact, but experience persistent losses with a lag, where the returns accumulate losses up to almost -20% after 5 years.

In summary, a 5 percentage points surprise increase in inflation seems to have a significant negative effect on both real physical assets, such as equity and housing, as well as on assets with nominally-fixed cash flows, such as government bonds. Equity is the worst inflation hedge among the three asset classes: during the post-1965 period, the cumulative real response after 3 years is around -30%. In real dollar terms, this means that a 1000$ stock portfolio would lose 300$ in real value three years after a five percentage point inflation surprise compared to a baseline scenario with no surprises in inflation. Housing

[19] In robustness exercises, we further extend the set of controls. First, we use interchangeably 1-year or 3-year lags for the vector $\mathbf{X}_{i,t-1}$. Furthermore, we use corporate debt or public debt/GDP to control for leverage and financial market size. Finally, we include time fixed effects (but exclude the rest of the macro controls, due to limited statistical power). Results are robust to these different specifications.

also proves to be an ineffective inflation hedge: during the modern period, its returns initially withstand the surprise in inflation in real terms at time 0. However, it gradually declines over 1-2 years, dropping by as much as 20% after five years. Finally, government bonds perform also poorly after an inflation surprise, especially up to two years after the inflation surprise. Nevertheless, for the post-1965 period, its value converges to steady state in the medium run.

Additional exercises In Appendix Figure A.9, we investigate the effect of inflation surprises on prices and yields separately for equity and housing assets. Dividends and equity prices exhibit a similar pattern, initially reacting strongly negatively and then gradually converging back to zero within 3-4 years. In contrast, the response of house prices and rents diverge. First, real rents drop almost one-to-one with inflation on impact and recover slowly afterwards. This highlights the role of nominal rigidities in rental markets, where contracts are renewed gradually over time. Compared to this, house prices do not change in the immediate aftermath of an inflation surprise, but gradually decline over time, losing 20% over a five-year horizon. Finally, Appendix Figure A.10 reports the dynamic inflation beta for total returns using OECD surprise at semi-annual frequency, with similar magnitudes.

1.3.3. A long-run view: inflation and asset returns

So far we analyzed the effect of inflation on asset returns during the post-1965 period. In this section, we broaden our scope by taking a historical perspective, asking whether the negative relationship between inflation and asset returns is a general pattern across time or a distinctive characteristic of the post-1965 period. Figure 1.6 highlights the importance of investigating the historical period to analyze the effects of inflation on asset returns. Here, we present the frequency of large inflation surprises, defined as surprises larger than five percentage points during 1870-2023.[20]

In the full historical sample, large inflation surprises were much more frequent than in the post-1965 period. While such large surprises have been rare in recent decades – particularly during the Great Moderation – many countries often experienced inflation surprises exceeding 5 percentage points between 1870 and 1950, with a high concentration around the two world wars. This underscores the importance of studying inflation in a historical context, particularly using pre-World War II samples, as also emphasized by Schmitt-Grohé and Uribe (2024).

We start by estimating the same panel-LP model as in Equation 1.5 for the full 1870-2023 sample.[21] Figure 1.7 presents the dynamic inflation betas of stocks, housing, and bonds

[20] If not explicitly stated otherwise, our long-run inflation surprise series is a combination of model-based inflation surprises until 1965 and OECD-based inflation surprises afterwards. We choose 5 percentage points as threshold value throughout the paper to be comparable with the recent inflation surge.

[21] We include also periods of World Wars, and only exclude hyperinflation years. We report in Appendix A

Figure 1.6.: Frequency of large inflation surprises: 1870-2023

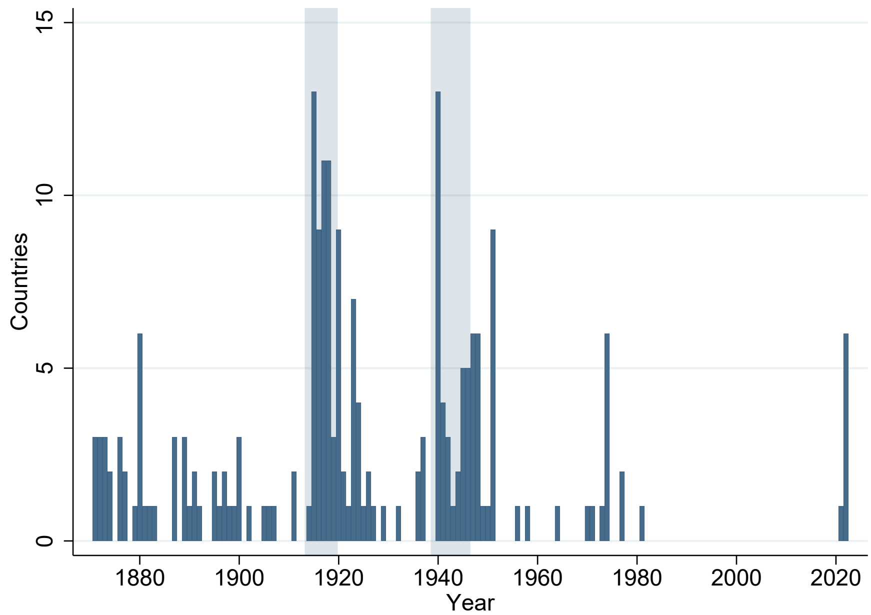

Notes: Timeline with the frequency of positive inflation surprises above 5 p.p. across the 18 OECD countries, excluding period of hyperinflation. World wars are presented as blue shaded areas. Yearly frequency, 1870-2023.

following a 5 percentage point surprise increase in inflation. Similar to the post-1965 sample, returns of equity and government bonds decrease on impact, although with smaller magnitude (less than 10 percentage points for both assets during 1870-2023, while in the post-1965 period, both equity and government bonds decreased by 20 percentage points).[22] In addition, we observe persistent effects of inflation on equity and government bonds during the full historical period, while its effects were rather transitory in the modern sample. Finally, different from the post-1965 period, housing returns drop on impact by around 5 percentage points, which further decreases by 7 percentage points 5 years after the shock. Nevertheless, like equity and government bonds, the magnitude of housing returns' response is also significantly weaker than in the post-1965 period.[23]

To explain such differences between the full historical sample and the post-1965 period, we now estimate the same impulse response functions in a *dynamic* fashion, using 40-years rolling windows over the full sample 1870-2020.[24] By doing so, we investigate whether the relationship between inflation and asset returns has experienced any structural changes

the same figures excluding World wars, with similar results.

[22]Note that the wide confidence intervals for government bonds are driven by the World War episodes, which were characterized by high volatility in the bond market, see Table 1.1. In Appendix Figure E.2 we present the IRFs excluding the World War episodes.

[23]In a robustness exercise, we include also time fixed effects in the full-sample local projections. Results align to what shown in Figure 1.7.

[24]In this exercise, we excluded the last 3 years of sample, as the 2021 inflation spike drives the whole variation in the 40-years rolling window – especially given the fact that the other inflation surprises were minor during the last 40 years.

Figure 1.7.: Dynamic inflation betas, 1870-2023

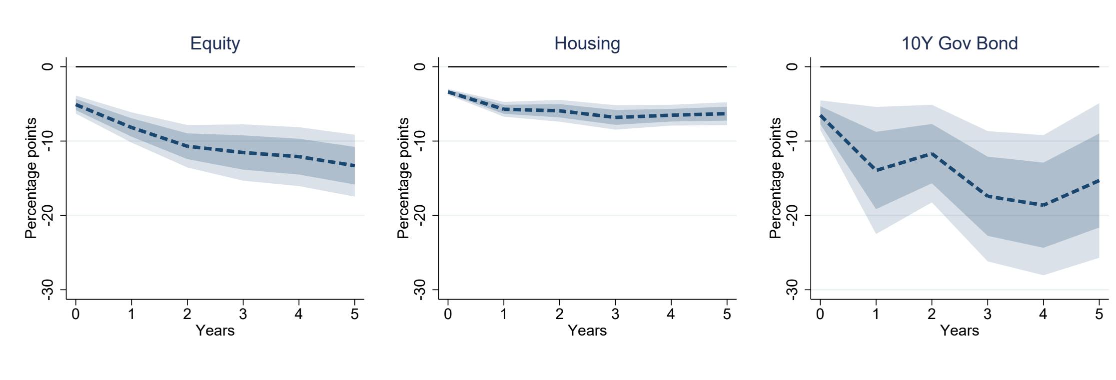

Notes: Impulse responses for cumulative real asset returns estimated with panel-LP using OECD surprises for the period 1965-2020 and model surprises during 1870-1965. Effects scaled to a 5 p.p. increase in inflation. Yearly frequency, OECD surprises are end-of-the-year. Bands are at the 68% and 90% confidence level. World wars are included in the estimation sample.

during the last 150 years. In Figure 1.8 we report in panel (a) the contemporaneous response over time (equivalent to inflation betas), and in panel (b) the response after 4 years to assess how the persistence of responses changed over time.

It is striking how the contemporaneous response of equity dropped over time from $\approx -5\%$ before 1970 and up to -40% in recent times. On the contrary, sensitivity of housing returns has remained relatively stable over time. Regarding persistence, inflation shocks had increasingly lasting effects on equity and housing returns, whereas government bonds recovered substantially more quickly in recent periods.[25] How can we explain this change in the relationship between inflation and asset returns in the mid-20th century? In Figure 1.9, we present the contemporaneous time-varying response of short-term interest rates and GDP following an inflation surprise, while in Appendix Figure A.8 we report the average response over the modern and the full sample.

[25] Results are qualitatively robust when employing 1 (rolling) standard deviation surprises within each 40-years rolling window.

Figure 1.8.: Time-varying inflation beta

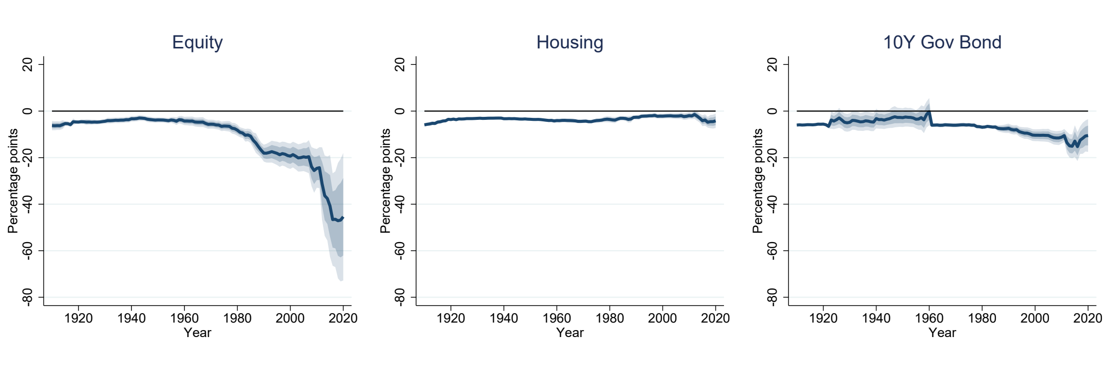

(a) Contemporaneous inflation beta ($h = 0$)

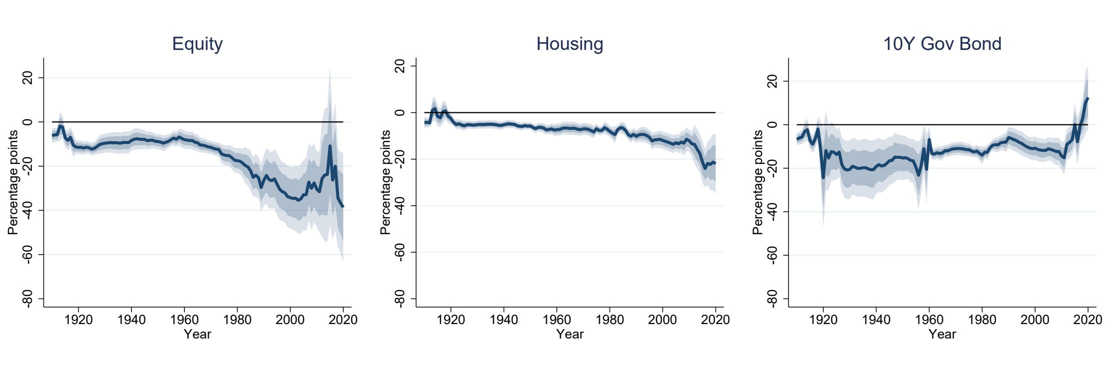

(b) Long-run inflation beta ($h = 4$)

Notes: Estimates of the inflation beta β_h at $h = 0$ (panel a) and $h = 4$ (panel b) over time, using 40-years rolling windows estimation, 1870-2020. Inflation surprises ε^{π} are derived from the model for the period 1870-1965 and from OECD for the post-1965 period and scaled as 5 p.p. increases in inflation. x-axis in the figure present the end-of-the-period year of the rolling window. Bands are at the 68% and 90% confidence level.

Interestingly, the negative response of GDP to inflation surprises reveal no common patterns with the response of asset returns, while we observe a clear structural break in how the short-term interest rate responds to an inflation surprise. Beginning in the 1970s, interest rates began responding much more significantly to inflation, suggesting that the stronger monetary policy response could be a key driver of the time-varying patterns in asset return responses to inflation.[26]

Robustness exercises In Appendix 1.E we provide a range of robustness exercises to test the validity of the results presented in Section 1.3. First, in Appendix Figure E.1, we show how baseline results of Figure 1.5 are robust to excluding the Great Inflation period, starting the sample in 1982, and the recent inflation surge. In Appendix Figure

[26] Jordà, Schularick, and Taylor (2020) demonstrate that the responses of house and stock prices to monetary policy remained consistent between the pre- and post-WW2 periods. This suggests that the greater negative impact of inflation on returns in the later period may be due to the increased magnitude of the monetary policy response, rather than to a greater sensitivity to interest rates.

Figure 1.9.: Time-varying inflation beta: macroeconomic variables

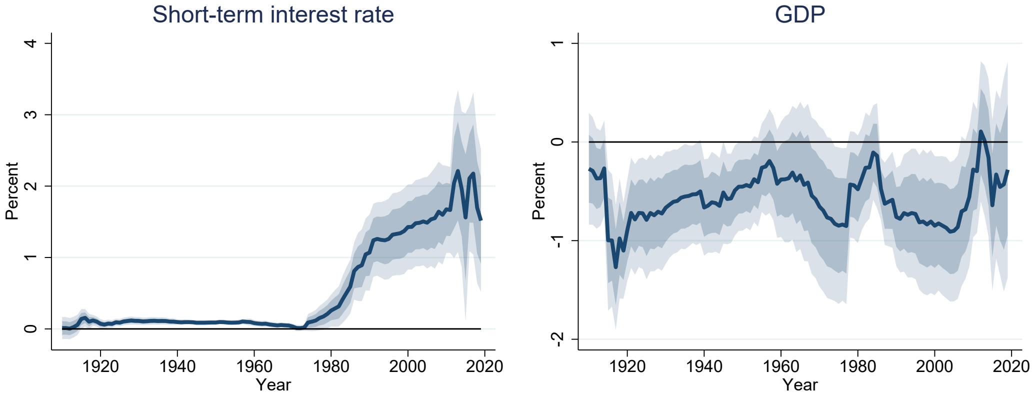

Notes: Impulse responses for short-term interest rate and real GDP through panel-LP with a 40-year rolling-window, using OECD surprises and model surprises, 1870-2020. Effects scaled to a 5 p.p. increase in inflation. Yearly frequency, OECD surprises are end-of-the-year. Bands are at the 68% and 90% confidence level.

E.2 we show how the historical sample results are mostly unchanged when excluding the two World wars. In Appendix 1.E.2 we use only the largest inflation surprises (i.e. above 5 percentage points) to run an event study analysis. All asset returns largely respond negatively and barely recover in the long run. To prove robustness of our findings using model-based inflation surprises, in Appendix 1.E.3 we present the IRFs from all inflation surprise series derived from the single models that are used in our main pooled model (see Appendix Figure E.4). Reassuringly, most of the impulse responses gravitate around our benchmark result. Nevertheless, some models fail to grasp the persistent negative effect on asset returns in the short term, which further justifies our selection of a pooled model rather than using a single model. In Appendix 1.E.4 we use the inflation surprises as instruments for overall inflation in a LP-IV framework. The estimates are closely aligned with the direct regression used in the main text. In Appendix 1.E.5 we investigate potential asymmetric effects of inflation shocks. Resassuringly, all exercises converge towards the same conclusion as our benchmark results.

1.4. Uncovering the long-run macroeconomic drivers

1.4.1. The role of monetary policy: pegged vs. float countries

Motivated by our stylized fact that the sensitivity of both asset returns and the short-term interest rate to inflation dramatically increased starting in the 1970s, we now investigate the role of monetary policy in shaping the impact of inflation surprises on asset returns. Indeed, a recent strand in the theoretical literature emphasizes the relevance of monetary policy conduct on the relationship between asset returns and inflation through sticky inflation expectations (Bianchi, Lettau, and Ludvigson, 2022) or beliefs about future demand and supply shocks (Caballero and Simsek, 2024). However, empirical evidence on this channel is scarce given the limitations of modern macroeconomic samples.[27]

Using our long-run data, we are able to study historical cases where monetary policy reacted differently to inflation due to reasons that are exogenous to business cycle conditions. Specifically, we exploit the long history of fixed exchange rate regimes and empirically investigate the effect of inflation surprises on asset returns in (i) pegged countries with open capital accounts compared to (ii) float countries. Countries peg their currencies for various reasons, such as stabilizing their exchange rates or facilitating trade with major trading partners. To support the pegged currency, the central bank and the government can intervene in several ways (e.g., announcements, buying/selling foreign exchange reserves) but the most sustainable mechanisms rely on capital controls and monetary policy alignment with the base country. However, a combination of these two is excluded by the trilemma of international finance – countries that want to maintain a fixed exchange rate cannot simultaneously have free capital flows and sovereign monetary policy (see e.g., Obstfeld and Taylor (1997) and Obstfeld, Shambaugh, and Taylor (2005)). Consequently, a central bank in a country with a pegged exchange rate system faces limited flexibility in increasing interest rates in response to an inflation shock, due to institutional constraints. In contrast, central banks in countries with floating exchange rate systems are not bound by such limitations. This contrast in monetary policy conduct is limited to short term domestic fluctuations, as the commitment of hard pegs to stable currencies avoids explosive inflation dynamics in the medium-long run.[28]

We now introduce a state-dependent panel local projections framework in the spirit of Ramey and Zubairy (2018), where the states are (i) pegged countries with open capital accounts and (ii) float countries:

$$\begin{aligned} r_{i,t+h} - r_{i,t-1} = &\mathcal{I}_{i,t-1}\Big\{\beta_{A,h}\varepsilon_{it}^{\pi} + \Psi_{A,h}(L)\mathbf{X}_{i,t-1} + \delta_{A,i,h}\Big\} \\ &+(1-\mathcal{I}_{i,t-1})\Big\{\beta_{B,h}\varepsilon_{it}^{\pi} + \Psi_{B,h}(L)\mathbf{X}_{i,t-1} + \delta_{B,i,h}\Big\} + \epsilon_{i,t+h} \quad \text{for} \quad h = 0,1,2,\ldots \end{aligned} \tag{1.6}$$

We suppress the superscript k on estimated coefficients, which refers to the different assets, for readability. $\mathcal{I}_{i,t-1}$ is a dummy variable that indicates whether a country is pegged in t and $t-1$ and in both times it has open capital accounts. The capital account opennes variable $q_{i,t-1}$ is taken from Quinn, Schindler, and Toyoda (2011), where $q_{i,t-1} = 1$

[27] A notable exception, using a high-frequency empirical strategy, is Gil de Rubio Cruz et al. (2023).

[28] There is ample empirical evidence that hard pegs are effective in maintaining a low and stable inflation in the medium run, see e.g., Ghosh, Qureshi, and Tsangarides (2014) and Bleaney and Francisco (2005).

when the country completely allows free capital movements.[29] We condition on being an *open peg* in both time periods t and $t-1$ to avoid cases where the decision of being peg is endogenously determined by economic conditions.[30] The main distinction between the two states of the world lies in the monetary policy conduct: whenever countries have their currency pegged while maintaining their capital accounts open, they are not able to freely react to inflation by raising interest rates. Using this specification, we can directly test the effects of inflation surprises on peg countries with open capital accounts against float countries, i.e. $\beta_{A,h} = \beta_{B,h}$. The rest of the specification is the same as in Equation 1.5.

Figure 1.10 presents the results. In the peg country scenario, cumulative real returns of equity and housing (i.e. what we define as real physical assets) respond significantly more positively than those of non-pegged countries. While the results for float countries resemble those of our benchmark estimates in Section 1.3 – responding negatively in a persistent way – equity and housing returns of pegged countries do react negatively on impact, but recover their real value afterwards in a significantly more positive manner.[31] Most of the negative response of stocks in *open pegs* realizes on impact, while from year 1 onward equity returns have positive dynamics and quickly offset all losses. We argue below that the initial negative impact on real returns is due to the stickiness of dividends, whose policies are slow moving and cannot react to sudden macro shocks within few months in real terms, but do adjust in the medium run (Leary and Michaely, 2011). In float countries, the negative response materializes not only on impact, but unfolds over several years. Compared to these results, we do not observe significantly heterogeneous responses of government bonds: in both cases, government bond returns decrease persistently and the different responses in the two states are essentially one-to-one with the response of the nominal interest rate.

To obtain a better understanding of the mechanisms, we present in Figure 1.11 the response of macroeconomic variables such as the nominal short-term interest rate, the real interest rate (computed as nominal rate minus future expected inflation, i.e., $\approx i_t - \mathbb{E}_t[\pi_{t+1}]$) and GDP. First, in line with the theory supporting the trilemma, pegged countries do not experience any change in the nominal short-term interest rate, while central banks of float countries increase their policy rate in response to higher inflation. Correspondingly, the development of GDP differs greatly across pegged and float countries. In both scenarios, GDP does not respond significantly on impact to an inflation surprise. However, in float countries GDP follows a recessionary path, while pegged countries experience a mild expansion. Such different paths of GDP have an effect on the future stream of cashflows of real physical assets, see Figure 1.12: real dividends and rents of pegged countries respond

[29] This index is available until 2012, at most. We merge this with the capital account opennes index from Chinn and Ito (2006), up to 2021. Note that the two variables are very close in the recent overlapping period.

[30] As noted in Jordà, Schularick, and Taylor (2020), restrictions on capital mobility have not been used as a high-frequency policy tool by pegging economies since the index is very slow moving in advanced economies.

[31] Appendix Figure A.11 provides the response for equity and house prices, i.e., excluding dividends and rents, respectively. The response of house prices is now aligned to equity, highlighting the paramount role of nominal rigidities in the rental markets as main source of deviation for housing.

Figure 1.10.: Asset returns responses in pegged vs non-pegged countries

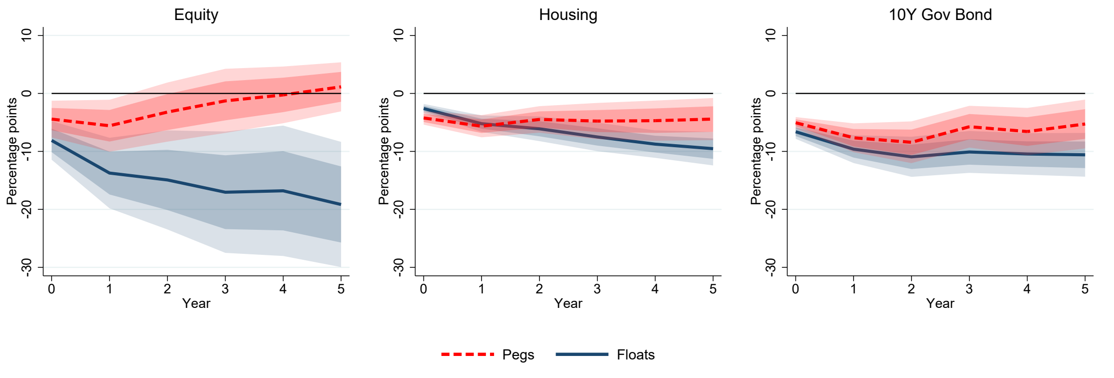

Notes: Impulse responses for cumulative real asset returns on inflation surprises in pegged countries (red dash) and in non-pegged countries (navy line). Model & OECD surprises, scaled to a 5 p.p. increase in inflation. Sample of 18 countries from the JST, 1870-2023, yearly frequency. The two lines represent the estimate of $\beta_{A,h}$ and $\beta_{B,h}$ in the state-dependent LP. Shaded areas are the 90% and 68% confidence intervals.

more positively to an inflation surprise compared to float countries, although the effect is larger for equity. Furthermore, we also present the response of real corporate profits to confirm that the dividend effects reflect income generations rather than changes in payout policy – under pegs, corporate profits rise in line with dividends, while both remain depressed under floats. This finding underscores the relevance of the proxy effect proposed by Fama (1981) in explaining the link between inflation and asset returns. According to this hypothesis, higher inflation serves as a proxy for weaker *future* economic activity, which translates into lower cash flows for real physical assets. Our analysis shows that the monetary policy reaction to inflation is the key driver of the subsequent deterioration of economic fundamentals.

Figure 1.11.: Macroeconomic responses in peg vs float

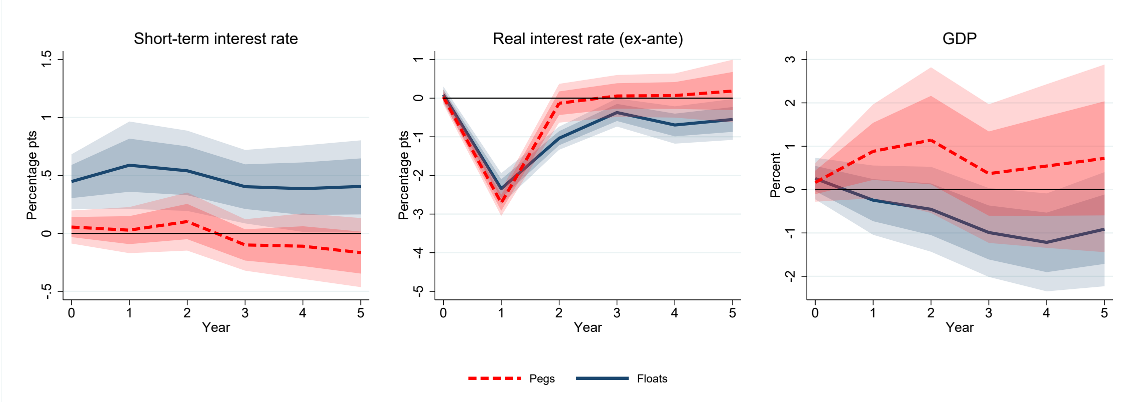

Notes: Impulse responses for nominal interest rate, (ex ante) real interest rate and real GDP on inflation surprises in pegged countries (red dash) and in float countries (navy line). Model & OECD surprises, scaled to a 5 p.p. increase in inflation. Sample of 18 countries from the JST, 1870-2023, yearly frequency. The two lines represent the estimate of $\beta_{A,h}$ and $\beta_{B,h}$ in the state-dependent LP. Shaded areas are the 90% and 68% confidence intervals.

Figure 1.12.: Cash flow and profit responses in peg vs float

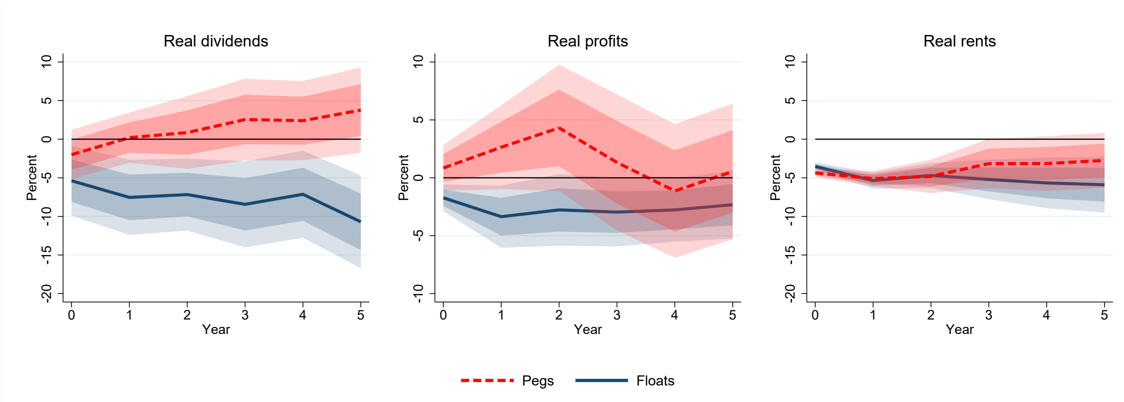

Notes: Impulse responses for real dividends, profits and rents on inflation surprises in pegged countries (red dash) and in float countries (navy line). Model & OECD surprises, scaled to a 5 p.p. increase in inflation. Sample of 18 countries from the JST, 1870-2023, yearly frequency. The two lines represent the estimate of $\beta_{A,h}$ and $\beta_{B,h}$ in the state-dependent LP. Shaded areas are the 90% and 68% confidence intervals.

Besides the cashflow channel, monetary policy can also directly influence asset returns through the change in discount rates. We investigate this by presenting the real interest rate responses to a surprise increase in inflation in Figure 1.11. In both monetary regimes, the real interest rate response mostly overlaps, with only slight differences in later years when the divergent response has already materialized. This equal response can be explained through "sticky" reaction of inflation expectations, in line with empirical evidence by Blanco, Ottonello, and Ranošová (2025) and theoretical contribution of Bianchi, Lettau, and Ludvigson (2022).[32]

[32] In a Campbell-Shiller decomposition à la Bernanke and Kuttner (2005) we can confirm that future cash flows, rather than discount rate news, drive most of the overall negative response to inflation, similarly

We conclude this section by highlighting the importance of studying the *persistent* effects of inflation on asset returns, rather than solely concentrating on contemporaneous responses, in pinning down the monetary policy channel. As shown in previous literature, the "long and variable lag of monetary policy" (Friedman, 1961) causes the full impact of inflation on asset returns to materialize with a lag. Our paper is the first to explicitly model the *cumulative* responses up to 5 years, which allows us to empirically identify the importance of this channel.[33]

Robustness exercises The robustness checks in Appendix Figures 1.F.1 and 1.F.2 confirm that excluding Euro-area countries or restricting the sample to the pre-1972 period leaves our main result unchanged: in economies with open pegs, inflation surprises more strongly boost real asset returns due to milder monetary policy reactions. Appendix Figure F.3 further shows that the gold standard regime itself does not drive our main result. Rather, only countries with genuinely constrained monetary policy — whether or not they were on the gold standard — exhibit muted return responses to inflation.

A detailed note on housing returns Similar to dividends, real rents decrease on impact, irrelevant of the exchange rate regime. Afterwards, in pegged countries, real rents begin to recover some of their value after 2-3 years, but the difference is not strongly significant compared to developments in floating countries. This suggests that heterogeneous monetary policy conduct does not fully explain the short-term negative response of real housing returns to a surprise increase in inflation, unlike the case for stocks.

A key difference between housing and other assets such as equity and bonds lies in trading frequency and rental contract renegotiation. Rental contracts are adjusted infrequently, whereas equity and bonds are traded constantly in the market. This allows the latter to quickly absorb an inflation surprise, while rents and housing returns may experience more long-lasting effects. This distinction is evident in Appendix Figure A.11, where we present the IRFs of equity and house prices, excluding dividends and rents. Now, the response of house prices aligns much more with the mechanism obseved for equity. In addition, the effect of monetary policy on housing returns may be limited, as their returns depend on long-term interest rates – attached to mortgage contracts – rather than on short-term interest rates.

All these arguments indicate that, for housing, factors beyond the monetary policy conduct – such as money illusion (Brunnermeier and Julliard, 2008) – may play a more substantial role in driving the negative reaction of returns to inflation surprises as compared to other asset classes. These observations underscore the need for future research in developing a more comprehensive framework that can capture these unique characteristics

to previous findings by Boons et al. (2020).

[33] Other studies that look at the dynamic responses of asset returns to inflation concentrate on high-frequency and short term dynamics, e.g., Chaudhary and Marrow (2022) and Gil de Rubio Cruz et al. (2023)

of the housing market.

1.4.2. Demand vs. supply driven inflation surprises

Recent literature that seeks to understand the economic consequences of inflation often emphasizes the importance of the underlying source of unexpected price increases. Cieslak and Pflueger (2023) and related studies[34] provide theoretical foundations suggesting that real physical assets should respond differently depending on whether inflation is driven by demand or supply shocks. For example, assets such as stocks are directly influenced by economic activity and the output gap and are therefore expected to react positively to demand-driven inflation. In contrast, a supply-driven inflationary shock is likely to depress returns.

So far, our analysis over the past 150 years shows that inflation surprises have often led to a decrease in GDP, which may indicate that these surprises were mostly supply-driven (see Figure 1.9). This could explain why real physical assets, such as equities and housing, have consistently reacted negatively to inflation surprises. Notably, this pattern may offer an alternative explanation for the increasingly negative relationship between inflation and asset returns that we documented in previous sections. If supply shocks became more prominent starting in the mid-20th century, they could have played a key role in depressing asset returns in response to inflation.

We test this prediction empirically by disentangling demand- and supply-driven inflation surprises. Normally, such exercise would need to be executed by identifying demand and supply shocks that cause inflation to rise. This approach in practice imposes strong assumptions or relies on sign restrictions directly from a VAR that have often severe drawbacks.[35] We therefore introduce a new identification strategy, where we identify GDP growth surprises and exploit their joint correlation with inflation surprises to disentangle demand- and supply-driven inflation surprises. The intuition is that when a positive inflation surprise is coupled with a positive GDP growth surprise, the inflation surprise is demand-driven, while supply-driven when it is coupled with a negative GDP growth surprise. This approach is similar to identification by sign restrictions, although it does not require to impose explicit assumptions on the underlying shock and parametric estimation of reduced form residuals. In particular, our approach does not identify directly the underlying shock, but rather its influence on the inflationary process. This approach is also close in spirit to Shapiro (2022), who decomposes Personal Consumption Expenditures (PCE) inflation into demand- and supply-driven contributions by looking at the sign of changes in quantities and prices at the single category level underlying the PCE basket.

Similar to how we estimate our inflation surprises, we first digitize 1-year ahead forecasts of real GDP growth from the OECD Economic Outlooks during 1965-2023 for the same

[34] See, for example, Hess and Lee (1999), Piazzesi et al. (2006), Campbell, Sunderam, and Viceira (2009), Cieslak and Pang (2021), Pflueger (2023).

[35] See Wolf (2022) for a discussion.

18 OECD countries and subtract them from realized GDP growth. In addition, we also estimate 1-year ahead GDP forecasts using the same model as in Equation 1.1 for the full 1870-2023 sample. We plot GDP surprises and inflation surprises in Appendix Figure A.12.[36] In addition, in Appendix Figure A.13 we present the fraction of demand-driven and supply-driven inflation surprises in our sample of 18 countries during 1870-2023. Interestingly, the proportion of demand-driven and supply-driven inflation shocks is nearly equal, with 51% being demand-driven and 49% supply-driven.[37] Nevertheless, especially in the modern sample, standard deviation of surprises in supply-driven cases is higher (2.03 vs 1.53). All together, these statistics might explain why we observe more often a decrease of real GDP after an inflation surprise, see e.g. Figure 1.9 and Appendix Figure A.8.

Given our GDP surprise measure ϵ^y_{it}, we then estimate the following panel-LP specification to identify the *additional* effect of demand-driven inflation surprises on cumulative real returns of asset k:

$$\begin{aligned} r^k_{i,t+h} - r^k_{i,t-1} = \beta^k_h \varepsilon^\pi_{it} + \gamma^k_h \times \mathcal{I}\{\varepsilon^y_{it} \times \varepsilon^\pi_{it} > 0\} \times \varepsilon^\pi_{it} \\ + \Psi^k_h(L)\mathbf{X}_{i,t-1} + \delta^k_{i,h} + v^k_{i,t+h} \quad \text{for} \quad h = 0, 1, 2, \ldots \end{aligned} \tag{1.7}$$

where ε^π is the inflation surprise and $\mathcal{I}\{\varepsilon^y_{it} \times \varepsilon^\pi_{it} > 0\}$ represents a dummy for when GDP growth and inflation surprises have the same sign, indicating a demand-driven response. Controls include 2 lags of inflation surprises, GDP surprises and changes in the dependent variable, short-term interest rate, real GDP growth, inflation, equity price, measures to control for global trends as world GDP growth and world equity price growth. Finally, we always include horizon-specific country fixed effects. The coefficient γ_h visualizes the *additional* effect of demand-driven inflation, compared to the response identified in β_h, which presents the response in the supply-driven inflation scenario.[38] In this context, ε^π is normalized as usual to induce a 5 p.p. increase in inflation.

[36] In our sample, the two series are mostly uncorrelated (≈ -0.11). When regressed against each other and only controlling for lags of each surprise, the contemporaneous coefficient is insignificantly different from zero.

[37] When using only the modern sample and OECD surprises, starting in 1965, this proportion shifts slightly towards supply-driven inflation (48% vs 52%), but it remains quite evenly split.

[38] Therefore, $\beta_h + \gamma_h$ shows the overall response after a demand-driven inflation surprise.

Figure 1.13.: Asset returns responses to inflation surprises: demand- vs. supply-driven

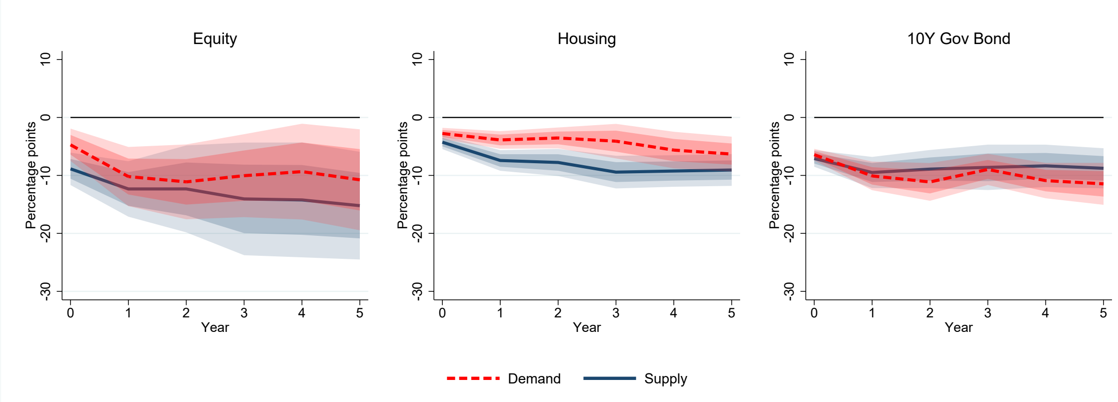

Notes: Impulse responses of cumulative real asset returns on inflation surprises due to demand or supply factors. Inflation surprise is scaled to induce an increase inflation by 5 p.p. The blue line presents estimates of β_h in Equation (1.7), while the red dashed line presents estimates of $\beta_h + \gamma_h$. Bands are at the 68% and 90% confidence level, based on Driscoll and Kraay (1998) standard errors. Yearly frequency, 1870-2023, using Model and OECD inflation and GDP growth surprises.

Figure 1.14.: Macro responses to inflation surprises: demand- vs. supply-driven

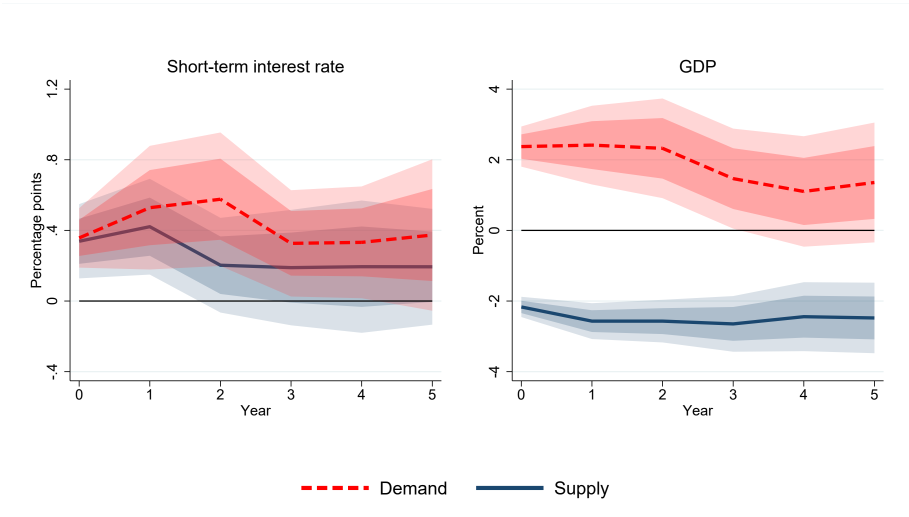

Notes: Impulse responses of short-term interest rates and real GDP on inflation surprises due to demand or supply factors. Inflation surprise is scaled to induce an increase inflation by 5 p.p. The blue line presents estimates of β_h in Equation (1.7), while the red dashed line presents estimates of $\beta_h + \gamma_h$. Bands are at the 68% and 90% confidence level, based on Driscoll and Kraay (1998) standard errors. Yearly frequency, 1870-2023, using Model and OECD inflation and GDP growth surprises.

Figure 1.13 presents the cumulative real asset returns in response to demand- and supply-driven inflation during 1870-2023.[39] In Figure 1.14 we present the response of macroeconomic variables such as GDP and short-term interest rates. In line with theoretical insights, equity and housing returns react more positively on impact when inflation is due to demand factors (equity: -4.71 p.p. vs. -8.86 p.p.; housing: -2.76 p.p. vs. -4.26

[39]In Appendix Figure G.1 we present the results using only the post-1965 sample with OECD-based inflation and GDP surprises. Results are robust.

p.p.). Nevertheless, the response is overall negative over the five years horizon regardless of the source of inflation. This is true even though real GDP significantly increases on impact after a demand-driven inflation surprise. This result is partially at odds with the prediction of Cieslak and Pflueger (2023) that a demand-driven inflation surprise leads to an increase in asset returns due to expansionary economic activity.

However, note that the short-term interest rate also increases after an inflation surprise, regardless of the source. In fact, monetary policy responds more strongly to demand-driven inflation surprises.[40] Therefore, one could argue that the results in Figure 1.14 are not capturing the pure effect of demand vs. supply-driven inflation surprises as it also includes the influence of monetary policy. We therefore purge the responses in demand and supply cases from the role of monetary policy conduct. To do so, we introduce a triple interaction between inflation surprises, GDP surprises, and the monetary policy regime dummy of Equation 1.6 and analyze how asset returns in *open peg* countries react differently to inflation in the case of demand- vs supply driven surprises. Since the response of the short-term interest is constrained due to the trilemma of international finance, as argued in Section 1.4.1, we interpret their responses as purely driven by the source of inflation.

Figure 1.15.: Asset returns responses to inflation surprises in *open pegs*: demand vs. supply

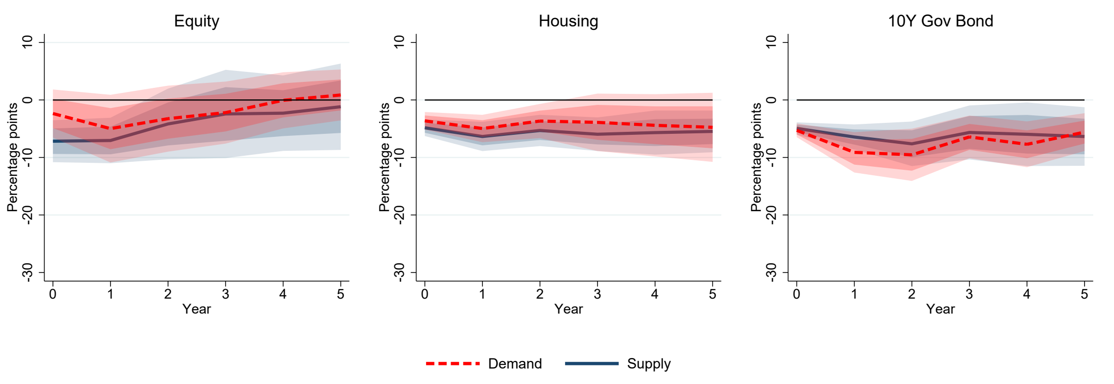

Notes: Impulse responses of cumulative real asset returns on inflation surprises due to demand or supply factors in *open pegs*. Inflation surprise is scaled to induce an increase inflation by 5 p.p. . Bands are at the 68% and 90% confidence level, based on Driscoll and Kraay (1998) standard errors. Yearly frequency, 1870-2023, using Model and OECD inflation and GDP growth surprises and open capital accounts measure by Quinn, Schindler, and Toyoda (2011).

Results are presented in Figure 1.15. Even when monetary policy is unable to intervene, the source of inflation has little impact on cumulative asset returns. The only notable difference is a more positive (and significant) immediate return for demand-driven surprises in the case of equity (-2.35 p.p. vs. -7.16 p.p.). Response for the longer horizons, however, essentially overlaps.

This can be explained by two factors: first, as shown in Appendix Figure G.6, after a

[40] Appendix Figure G.2 shows that interest rates responses are even stronger in the modern period, with short-term rates increasing by as much 4 percentage points in demand-driven case and only 1 p.p. in supply-driven case.

supply shock, GDP falls initially but recovers from $t + 1$ onward, while after a demand shock, GDP booms on impact but slows down thereafter. Second, dividends adjust only gradually to a change in GDP (Leary and Michaely, 2011) – see also Appendix Figure G.7. The initial benefit of demand shocks is therefore offset by the slow cashflow adjustment coupled with weaker *future* GDP growth.

For other asset classes the distinction between demand and supply disturbances is largely irrelevant: housing returns drop on impact and show minimal recovery due to infrequent trades and rental contract renegotiations, as discussed in Section 1.4.1; long-term government bonds, in line with theory, are mechanically affected by inflation in any case, which erodes the value of nominal coupons, yet showing a slightly higher response for supply-driven inflation surprises.[41] All in all, this evidence supports that monetary policy, rather than the source of inflation, is the key factor explaining the shift in the inflation-asset returns relationship.

Robustness exercises In Appendix Figure G.1 we replicate the results using only OECD-based surprises at the semi-annual frequency. In Appendix Figure G.3 we show and highlight that the *additional* effect of demand-driven surprises is more positive compared to supply-driven surprises. In Appendix Figure G.4 we repeat the demand vs supply decomposition using a continuous interaction between surprises rather than a dummy specification. In Appendix Figures G.5 anf G.6 we show the full set of responses in open pegs and float countries, for both demand- and supply-driven inflation surprises. Results confirm that no matter the source of the surprise, asset returns react negatively to inflation when monetary policy is free to respond to inflation and stabilize the price level.

1.5. Conclusion

In this paper, we dissect the complex relationship between inflation and asset returns by exploiting 150 years of data for 18 advanced economies. This unique setup with long-run cross-country data helps us address the long-standing debate within the inflation hedge literature. During 1870-2023, our findings show on average a negative relationship between inflation surprises and asset returns, irrelevant of whether the asset has nominally-fixed cash flows or are linked to real physical values. In addition, such negative effects seem to persist, especially for stocks and housing. This suggests that the effectiveness of real physical assets as inflation hedges is much more limited than traditionally assumed. Importantly, we show that the negative asset returns-inflation relationship has increasingly worsened starting in the in 1970s, when explicit inflation targeting policy were adopted by central banks.

Our detailed investigation on the macroeconomic drivers of this relationship reveals the importance of how monetary policy responds to inflation surprises. Regardless of the

[41] This result points towards a flight-to-safety mechanism and supports the literature that connects increased convenience yields with supply-driven inflation (Cieslak, Li, and Pflueger, 2024).

underlying source of the inflation surprise, which plays only a limited role, the increased prominence of inflation control in central banking practices translates into worsened conditions for financial markets and asset prices. Assets that are actively traded in financial markets, such as equity, adhere to this mechanism, such that when constraining monetary policy actions, inflation has essentially a zero impact on real asset returns. On the other hand, housing deserves further explorations, given the imperfections and rigidities in the renewal of rental contracts, low trading frequencies and dependence on longer-term mortgage rates, rather than the short-term interest rates directly affected by central banks.

In conclusion, our analysis emphasizes the critical role of monetary policy in mediating the impact of inflation on asset returns. As a relevant implication, this mechanism calls for central banks to assess the broader outcomes of their actions on financial markets. Recognizing this amplified effect of policy decisions on asset returns can lead to a strategic shift, where central banks could consider accepting marginally higher inflation levels to mitigate severe disruptions in asset markets. This result underlines the critical balance required between controlling inflation and ensuring financial market stability.

References

Alvaredo, Facundo, Anthony B Atkinson, Thomas Blanchet, Lucas Chancel, Luis Estevez Bauluz, Matthew Fisher-Post, Ignacio Flores, Bertrand Garbinti, Jonathan Goupille-Lebret, Clara Martínez-Toledano, et al. (2021). "Distributional national accounts guidelines methods and concepts used in the world inequality database". PhD thesis. World Inequality Lab (cit. on p. 13).

Auerbach, Alan J and Yuriy Gorodnichenko (2012a). "Fiscal multipliers in recession and expansion". In: *Fiscal policy after the financial crisis*. University of Chicago Press, pp. 63–98 (cit. on p. 8).

– (2012b). "Measuring the output responses to fiscal policy". In: *American Economic Journal: Economic Policy* 4.2, pp. 1–27 (cit. on p. 7).

– (2013). "Output spillovers from fiscal policy". In: *American Economic Review* 103.3, pp. 141–46 (cit. on p. 7).

Barnes, Michelle, John H Boyd, and Bruce D Smith (1999). "Inflation and asset returns". In: *European Economic Review*. DOI: `10.1016/S0014-2921(98)00090-7` (cit. on p. 6).

Bataa, Erdenebat, Denise R Osborn, Marianne Sensier, and Dick Van Dijk (2013). "Structural breaks in the international dynamics of inflation". In: *Review of Economics and Statistics* 95.2, pp. 646–659 (cit. on p. 9).

Bekaert, Geert and Xiaozheng Wang (2010). "Inflation risk and the inflation risk premium". In: *Economic Policy*. DOI: `10.1111/J.1468-0327.2010.00253.X` (cit. on pp. 3, 6, 14–16, 54, 55).

Bernanke, Ben S and Mark Gertler (2000). *Monetary policy and asset price volatility* (cit. on p. 4).

Bernanke, Ben S and Kenneth N Kuttner (2005). "What explains the stock market's reaction to Federal Reserve policy?" In: *The Journal of finance* 60.3, pp. 1221–1257 (cit. on p. 26).

Bianchi, Francesco, Martin Lettau, and Sydney C Ludvigson (2022). "Monetary policy and asset valuation". In: *The Journal of Finance* 77.2, pp. 967–1017 (cit. on pp. 5, 23, 26).

Binder, Carola and Rupal Kamdar (2022). "Expected and Realized Inflation in Historical Perspective". In: *Journal of Economic Perspectives* 36.3, pp. 131–56 (cit. on p. 9).

Blanco, Andrés, Pablo Ottonello, and Tereza Ranošová (2025). "The dynamics of large inflation surges". In: *Review of Economics and Statistics*, pp. 1–31 (cit. on pp. 10, 16, 26).

Bleaney, Michael and Manuela Francisco (2005). "Exchange rate regimes and inflation: only hard pegs make a difference". In: *Canadian Journal of Economics/Revue canadienne d'économique* 38.4, pp. 1453–1471 (cit. on p. 23).

Boons, Martijn, Fernando Duarte, Frans De Roon, and Marta Szymanowska (2020). "Time-varying inflation risk and stock returns". In: *Journal of Financial Economics* 136.2, pp. 444–470 (cit. on p. 27).

Boudoukh, Jacob and Matthew Richardson (1993). “Stock returns and inflation: A long-horizon perspective”. In: *The American economic review* 83.5, pp. 1346–1355 (cit. on pp. 6, 56).

Brunnermeier, Markus K and Christian Julliard (2008). “Money illusion and housing frenzies”. In: *The Review of Financial Studies* 21.1, pp. 135–180 (cit. on p. 27).

Caballero, Ricardo J and Alp Simsek (2024). “Central Banks, Stock Markets, and the Real Economy”. In: *Annual Review of Financial Economics* 16 (cit. on pp. 5, 23).

Campbell, John Y, Carolin Pflueger, and Luis M Viceira (2020). “Macroeconomic drivers of bond and equity risks”. In: *Journal of Political Economy* 128.8, pp. 3148–3185 (cit. on p. 4).

Campbell, John Y, Adi Sunderam, and Luis M Viceira (2009). *Inflation bets or deflation hedges? The changing risks of nominal bonds.* Tech. rep. National Bureau of Economic Research (cit. on p. 28).

Chaudhary, Manav and Benjamin Marrow (2022). “Revisiting the Stock-Inflation Puzzle”. In: *Available at SSRN 4154564* (cit. on pp. 6, 27).

Chinn, Menzie D and Hiro Ito (2006). “What matters for financial development? Capital controls, institutions, and interactions”. In: *Journal of development economics* 81.1, pp. 163–192 (cit. on p. 24).

Cieslak, Anna, Wenhao Li, and Carolin Pflueger (2024). *Inflation and treasury convenience.* Tech. rep. National Bureau of Economic Research (cit. on p. 32).

Cieslak, Anna and Hao Pang (2021). “Common shocks in stocks and bonds”. In: *Journal of Financial Economics* 142.2, pp. 880–904 (cit. on p. 28).

Cieslak, Anna and Carolin Pflueger (2023). “Inflation and asset returns”. In: *Annual Review of Financial Economics* 15.1, pp. 433–448 (cit. on pp. 5, 6, 28, 31).

Cieslak, Anna and Pavol Povala (2015). “Expected returns in Treasury bonds”. In: *The Review of Financial Studies* 28.10, pp. 2859–2901 (cit. on pp. 6, 8).

David, Alexander and Pietro Veronesi (2013). “What ties return volatilities to price valuations and fundamentals?” In: *Journal of Political Economy* 121.4, pp. 682–746 (cit. on p. 4).

Driscoll, John C and Aart C Kraay (1998). “Consistent covariance matrix estimation with spatially dependent panel data”. In: *Review of economics and statistics* 80.4, pp. 549–560 (cit. on pp. 15, 17, 30, 31, 54–56, 63, 71–73, 75–77).

Elliott, Graham and Allan Timmermann (2016). “Forecasting in economics and finance”. In: *Annual Review of Economics* 8, pp. 81–110 (cit. on p. 9).

Ellison, Martin, Sang Seok Lee, and Kevin Hjortshøj O’Rourke (2024). “The ends of 27 big depressions”. In: *American Economic Review* 114.1, pp. 134–168 (cit. on pp. 6, 8, 11–13).

Engsted, Tom and Carsten Tanggaard (2000). “The Relation Between Asset Returns and Inflation at Short and Long Horizons.” In: DOI: `10.2139/SSRN.221669` (cit. on p. 6).

Fama, Eugene F (1981). “Stock returns, real activity, inflation, and money”. In: *The American economic review* 71.4, pp. 545–565 (cit. on pp. 5, 6, 25).

Fama, Eugene F. and G.William Schwert (1977). “Asset returns and inflation”. In: *Journal of Financial Economics*. DOI: 10.1016/0304-405X(77)90014-9 (cit. on pp. 3, 6, 14, 16).

Fang, Xiang, Yang Liu, and Nikolai Roussanov (2022). *Getting to the core: Inflation risks within and across asset classes*. Tech. rep. National Bureau of Economic Research (cit. on pp. 5, 6, 14, 16, 55).

Federle, Jonathan, Cathrin Mohr, and Moritz Schularick (2024). “Inflation surprises and election outcomes”. In: *Available at SSRN 5032871* (cit. on p. 6).

Fehrle, Daniel (2023). “Hedging against inflation: housing versus equity”. In: *Empirical Economics* 65.6, pp. 2583–2626 (cit. on p. 6).

Fisher, Irving (1930). *Theory of interest: as determined by impatience to spend income and opportunity to invest it*. Augustusm Kelly Publishers, Clifton (cit. on p. 8).

Friedman, Milton (1961). “The lag in effect of monetary policy”. In: *Journal of Political Economy* 69.5, pp. 447–466 (cit. on pp. 6, 27).

Gabriel, Ricardo Duque (2021). “Monetary Policy and the Wage Inflation-Unemployment Tradeoff”. In: *Available at SSRN 3689791* (cit. on p. 51).

Geske, Robert and Richard Roll (1983). “The fiscal and monetary linkage between stock returns and inflation”. In: *The journal of Finance* 38.1, pp. 1–33 (cit. on p. 6).

Ghosh, Atish R, Mahvash S Qureshi, and Charalambos G Tsangarides (2014). “On the value of words: Inflation and fixed exchange rate regimes”. In: *IMF Economic Review* 62.2, pp. 288–322 (cit. on p. 23).

Gil de Rubio Cruz, Antonio, Emilio Osambela, Berardino Palazzo, Francisco Palomino, and Gustavo Suarez (2023). “Inflation Surprises and Equity Returns”. In: *Available at SSRN 4280699* (cit. on pp. 23, 27).

Grimm, Maximilian, Òscar Jordà, Moritz Schularick, and Alan M Taylor (2023). *Loose monetary policy and financial instability*. Tech. rep. National Bureau of Economic Research (cit. on p. 39).

Gurkaynak, Refet S, Brian Sack, and Eric T Swanson (2005). “Do actions speak louder than words? The response of asset prices to monetary policy actions and statements”. In: (cit. on p. 4).

Gürkaynak, Refet S, Brian Sack, and Jonathan H Wright (2010). “The TIPS yield curve and inflation compensation”. In: *American Economic Journal: Macroeconomics* 2.1, pp. 70–92 (cit. on p. 7).

Hamilton, James D (1992). “Was the deflation during the Great Depression anticipated? Evidence from the commodity futures market”. In: *The American Economic Review*, pp. 157–178 (cit. on p. 6).

Hendry, David F and Michael P Clements (2004). “Pooling of forecasts”. In: *The Econometrics Journal* 7.1, pp. 1–31 (cit. on p. 9).

Hess, Patrick J and Bong-Soo Lee (1999). “Stock returns and inflation with supply and demand disturbances”. In: *The Review of Financial Studies* 12.5, pp. 1203–1218 (cit. on pp. 5, 6, 28).

Jordà, Òscar (2005). “Estimation and inference of impulse responses by local projections”. In: *American economic review* 95.1, pp. 161–182 (cit. on p. 16).

Jordà, Òscar, Katharina Knoll, Dmitry Kuvshinov, Moritz Schularick, and Alan M Taylor (2019). “The rate of return on everything, 1870–2015”. In: *The Quarterly Journal of Economics* 134.3, pp. 1225–1298 (cit. on pp. 10, 13, 51).

Jordà, Òscar, Moritz Schularick, and Alan M Taylor (2017). “Macrofinancial history and the new business cycle facts”. In: *NBER macroeconomics annual* 31.1, pp. 213–263 (cit. on pp. 10, 13, 51).

– (2020). “The effects of quasi-random monetary experiments”. In: *Journal of Monetary Economics* 112, pp. 22–40 (cit. on pp. 4, 15, 21, 24).

Katz, Michael, Hanno Lustig, and Lars Nielsen (2017). “Are stocks real assets? Sticky discount rates in stock markets”. In: *The Review of Financial Studies* 30.2, pp. 539–587 (cit. on p. 6).

Knox, Benjamin and Yannick Timmer (2025). “Stagflationary stock returns”. In: (cit. on p. 7).

Leary, Mark T and Roni Michaely (2011). “Determinants of dividend smoothing: Empirical evidence”. In: *The Review of Financial Studies* 24.10, pp. 3197–3249 (cit. on pp. 5, 24, 32).

Modigliani, Franco and Richard A Cohn (1979). “Inflation, rational valuation and the market”. In: *Financial Analysts Journal* 35.2, pp. 24–44 (cit. on p. 6).

Monnet, Eric, Damien Puy, et al. (2016). *Has Globalization Really Increased Business Cycle Synchronization?* Tech. rep. 54. International Monetary Fund (cit. on pp. 13, 41, 42).

Monnet, Eric and Damien Puy (2021). “One Ring to Rule Them All? New Evidence on World Cycles”. In: *CEPR Discussion Paper* 15958 (cit. on p. 13).

Montiel Olea, José Luis and Mikkel Plagborg-Møller (2021). “Local projection inference is simpler and more robust than you think”. In: *Econometrica* 89.4, pp. 1789–1823 (cit. on p. 16).

Obstfeld, Maurice, Jay C Shambaugh, and Alan M Taylor (2005). “The trilemma in history: tradeoffs among exchange rates, monetary policies, and capital mobility”. In: *Review of economics and statistics* 87.3, pp. 423–438 (cit. on p. 23).

Obstfeld, Maurice and Alan M Taylor (1997). “The great depression as a watershed: international capital mobility over the long run”. In: (cit. on pp. 4, 23).

Pflueger, Carolin (2023). *Back to the 1980s or not? The drivers of inflation and real risks in Treasury bonds*. Tech. rep. National Bureau of Economic Research (cit. on p. 28).

Piazzesi, Monika, Martin Schneider, Pierpaolo Benigno, and John Y Campbell (2006). “Equilibrium yield curves [with comments and discussion]”. In: *NBER macroeconomics Annual* 21, pp. 389–472 (cit. on p. 28).

Quinn, Dennis, Martin Schindler, and A Maria Toyoda (2011). “Assessing measures of financial openness and integration”. In: *IMF Economic Review* 59.3, pp. 488–522 (cit. on pp. 23, 31, 75–77).

Ramey, Valerie A and Sarah Zubairy (2018). “Government spending multipliers in good times and in bad: evidence from US historical data”. In: *Journal of political economy* 126.2, pp. 850–901 (cit. on pp. 23, 66).

Rigobon, Roberto and Brian Sack (2004). “The impact of monetary policy on asset prices”. In: *Journal of monetary economics* 51.8, pp. 1553–1575 (cit. on p. 4).

Schmitt-Grohé, Stephanie and Martín Uribe (2024). “What Do Long Data Tell Us about the Permanent Component of Inflation?” In: *AEA Papers and Proceedings*. Vol. 114. American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203, pp. 101–105 (cit. on pp. 6, 18).

Shapiro, Adam Hale (2022). “Decomposing supply and demand driven inflation”. In: Federal Reserve Bank of San Francisco San Francisco (cit. on p. 28).

Stock, James H and Mark W Watson (2007). “Why has US inflation become harder to forecast?” In: *Journal of Money, Credit and banking* 39, pp. 3–33 (cit. on p. 39).

Taylor, Alan M (2002). “A century of purchasing-power parity”. In: *Review of economics and statistics* 84.1, pp. 139–150 (cit. on p. 51).

Timmermann, Allan (2006). “Forecast combinations”. In: *Handbook of economic forecasting* 1, pp. 135–196 (cit. on p. 9).

Vogel, Lukas (2007). “How do the OECD growth projections for the G7 economies perform? A post-mortem”. In: (cit. on p. 8).

Wolf, Christian K (2022). “What Can We Learn from Sign-Restricted VARs?” In: *AEA Papers and Proceedings*. Vol. 112. American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203, pp. 471–475 (cit. on p. 28).

1.A. Additional figures and tables

Table A.1.: List of models used, part one

Model Name	Description
Univariate Models	
ARMA(1,1)	$\pi_t = c + \phi\pi_{t-1} + \epsilon_t + \theta\epsilon_{t-1}$. 30-year rolling window (RW hereafter).
ARMA(1,1) on transitory π	As above, with 30-year RW, but using $\hat{\pi}_t = \bar{\pi}_t - \pi_t$, where $\bar{\pi}_t$ is trend inflation and follows a random walk, i.e. $\bar{\pi}_t = \bar{\pi}_{t-1} + \epsilon_t$
Backward-looking Phillips Curve	$\hat{\pi}_t = \phi_1\hat{\pi}_{t-1} + \phi_2\hat{y}_{t-1} + \epsilon_t$, where $\hat{\pi}$ and $\hat{y}$ are transitory inflation and output gap, respectively. 30-year RW.
ARMA(1,1) w/o RW	as ARMA(1,1) above but no RW
ARMA(1,1) w/o RW on transitory π	as ARMA(1,1) above but no RW, using $\hat{\pi}$
Random Walk	$\pi_t = \pi_{t-1} + \epsilon_t$
Unobserved component trend-cycle	$\pi_t = \tau_t + \eta_t$, where $\tau_t = \tau_{t-1} + \epsilon_t$ is trend inflation, as in Stock and Watson (2007).
Multivariate Models	
Benchmark VAR(2)	$Y_t = A_1 Y_{t-1} + A_2 Y_{t-2} + \nu_t$ where $Y_t =$ [inflation, output gap, int. rate]
VAR	As above, but using transitory inflation $\hat{\pi}$
VAR	As the benchmark VAR(2), but using the r^* gap measure of Grimm et al. (2023)
VAR	As above, but on transitory inflation π
VAR on asset prices	$Y_t =$ [inflation, stock price, gov. bond price]
VAR on asset prices & GDP	$Y_t =$ [inflation, stock price, house price, GDP growth]
VAR on asset & commodity prices	$Y_t =$ [inflation, stock price, house price, commodity price]
VAR on GDP components	$Y_t =$ [inflation, consumption, investment, net exports]
VAR on commodities	$Y_t =$ [inflation, grains price, metals price, softs price]

Notes: Variables are in first-differences to ensure stationarity, if needed. Lag length for VARs is equal to 2 and is chosen using information criteria.

Table A.3.: List of models used, part two

Model Name	**Description**
Panel Data Models	
Arellano-Bond (A-B), no covariates	$\pi_{it} = \alpha_1 \pi_{i,t-1} + \alpha_2 \pi_{i,t-2} + u_{it}$, 10-year rolling window
A-B on transitory π	As above, no covariates but using $\hat{\pi}$, while trend inflation follows a random walk, i.e. $\bar{\pi}_{i,t} = \bar{\pi}_{i,t-1} + \nu_{i,t}$. Two lags, 10-year rolling window
A-B with output gap and int. rate	$\pi_{it} = \sum_{j=1}^{2} \alpha_j \pi_{i,t-j} + \phi \mathbf{X_{i,t-1}} + u_{it}$, where $\mathbf{X}$ is output gap and interest rate. 10-year rolling window.
A-B with $\mathbf{X}$ on transitory π	As above, where $\mathbf{X}$ is output gap and interest rate, but using transitory inflation $\hat{\pi}$

Notes: The 10-year rolling window is kept intentionally short as the Arellano-Bond estimator suffers particularly panels with long T dimension.

Figure A.1.: OECD inflation expectations, part one

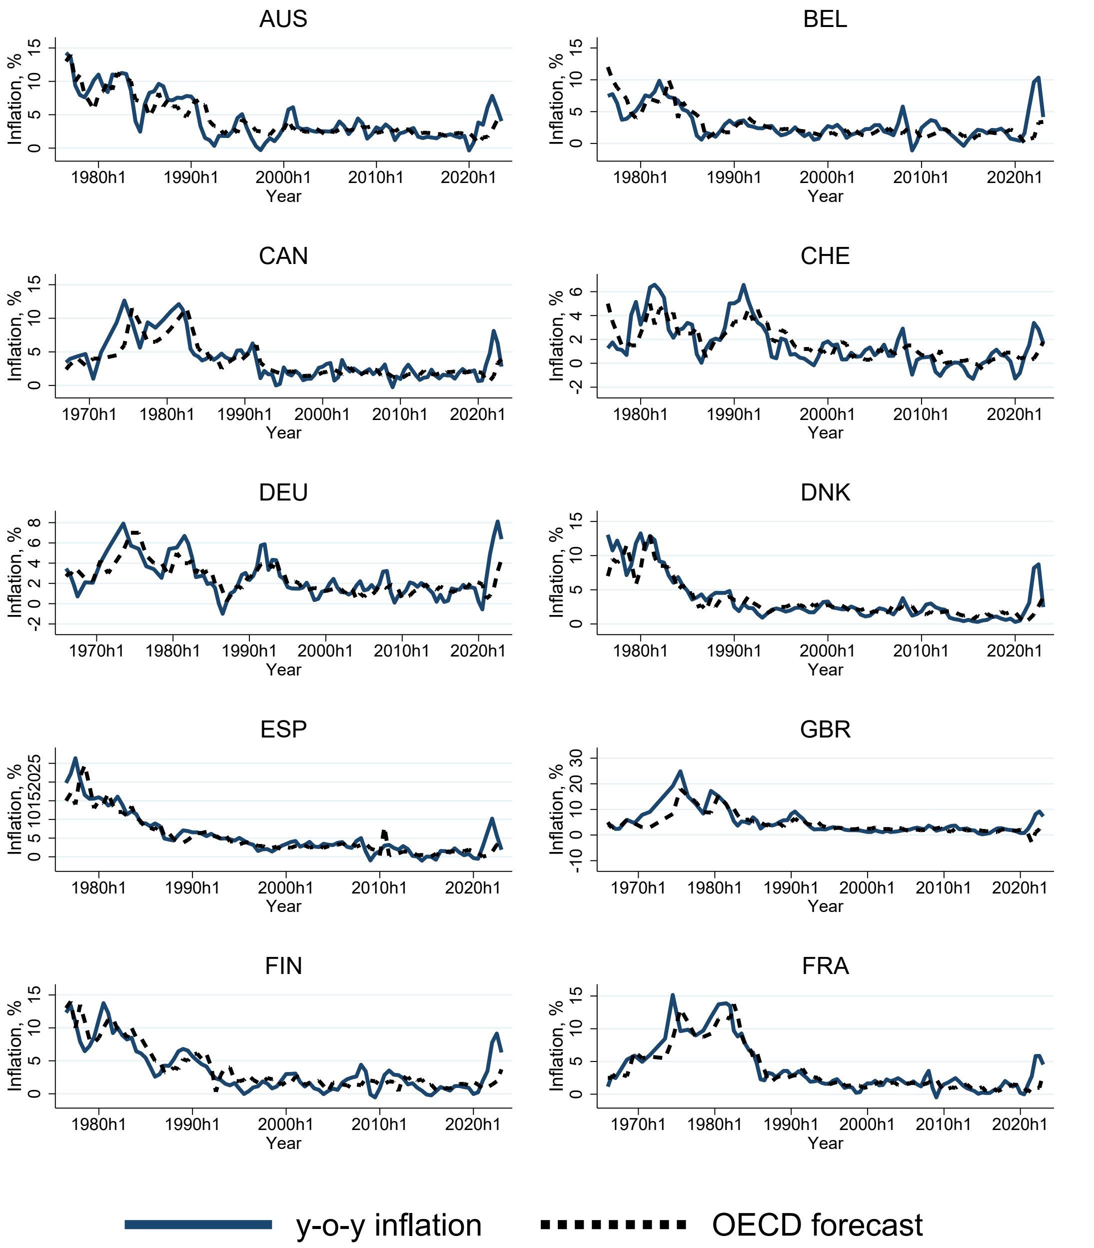

Notes: Black dashed line presents 1Y-ahead inflation expectations derived from the OECD Economic Outlooks as described in Section 1.2. Blue solid line presents the actual realized inflation rate from Monnet, Puy, et al. (2016). Half-yearly frequency, 1965-2023. Series for Australia, Belgium, Canada, Switzerland, Germany, Denmark, Spain, Great Britain, Finland, and France.

Figure A.2.: OECD inflation expectations, part two

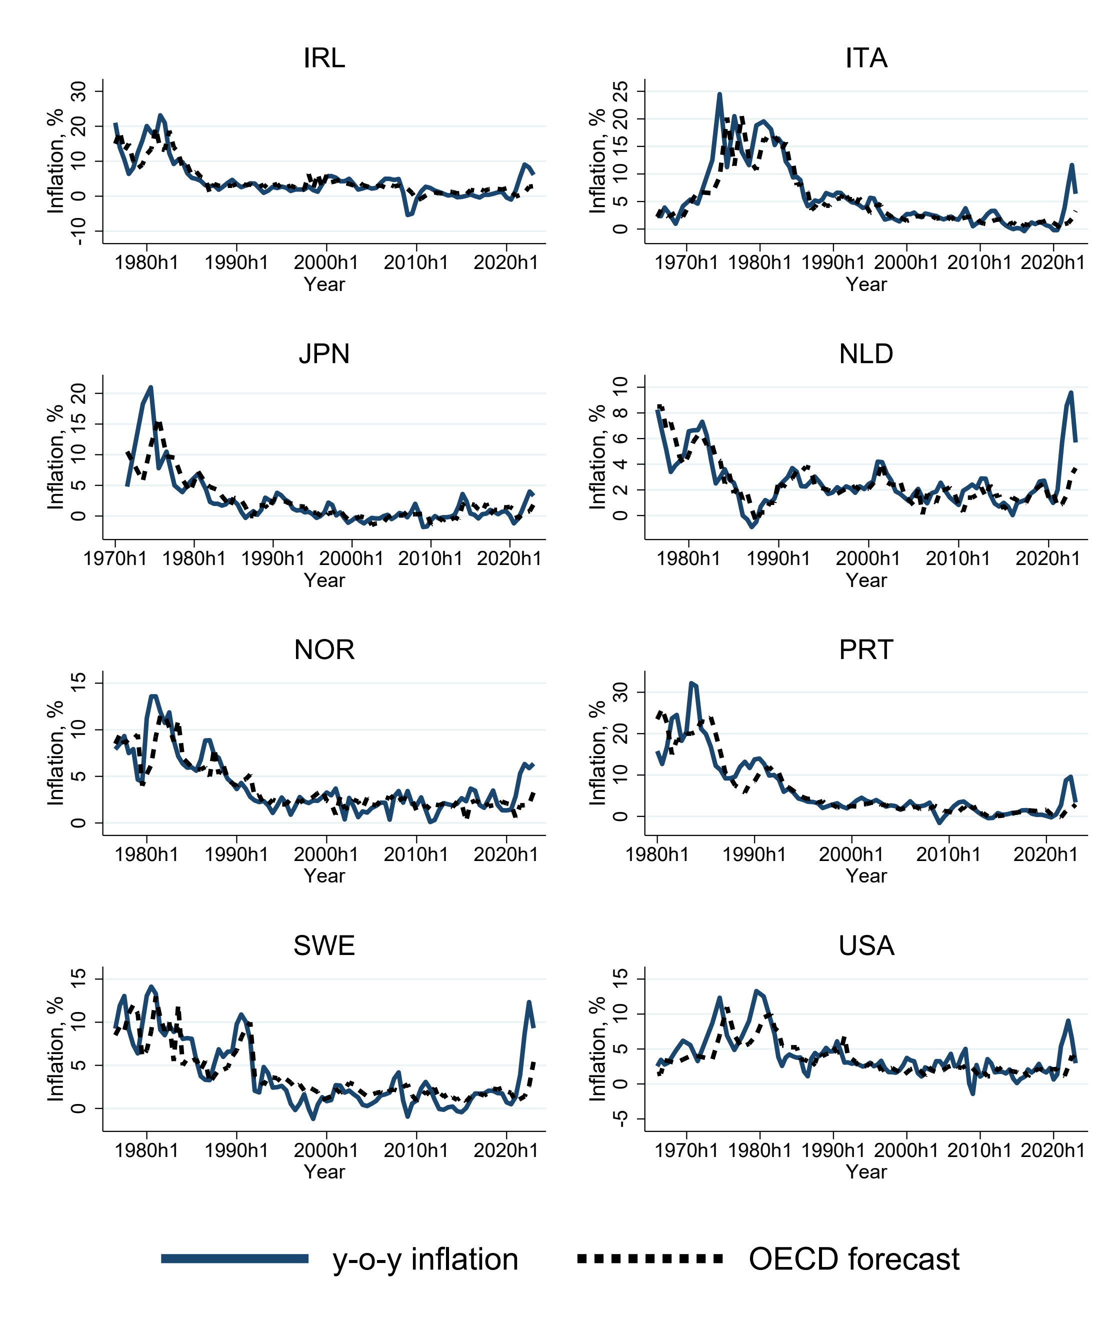

Notes: Black dashed line presents 1Y-ahead inflation expectations derived from the OECD Economic Outlooks as described in Section 1.2. Blue solid line presents the actual realized inflation rate from Monnet, Puy, et al. (2016). Half-yearly frequency, 1965-2023. Series for Ireland, Italy, Japan, Netherlands, Norway, Portugal, Sweden, and USA.

Figure A.3.: Model-based inflation expectations, part one

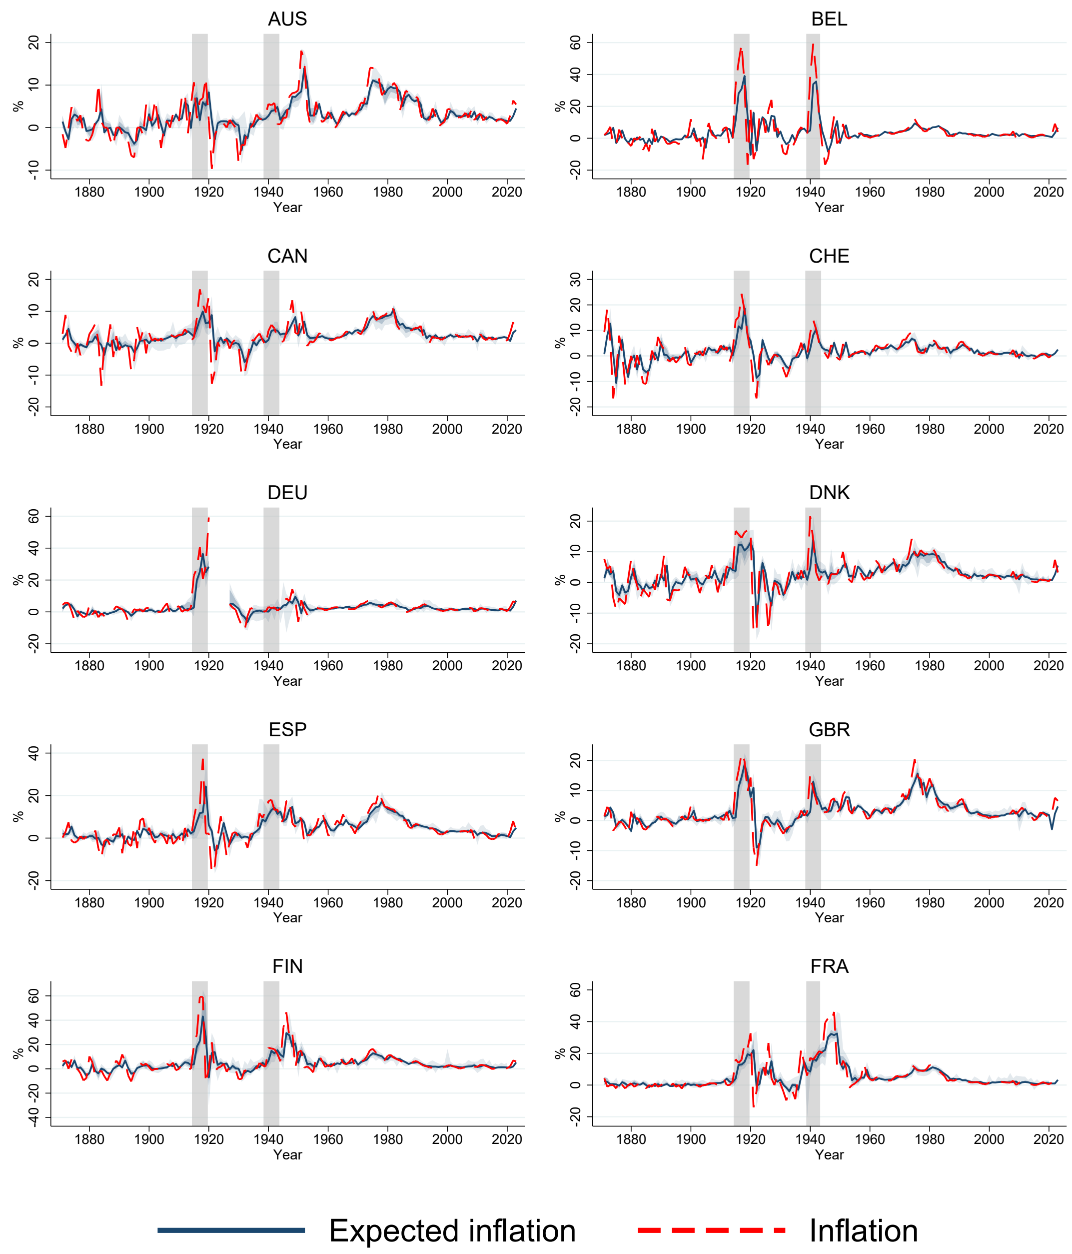

Notes: Blue solid line presents 1Y-ahead inflation expectations derived from the model as in Equation (1.1), together with bands at the interquartile range and 5th and 95th percentile of the distribution of model estimates. Red dashed line presents the actual realized inflation rate. The gray shaded areas present WWI and WWII. We exclude years around the German hyperinflation. We present the series for Australia, Belgium, Canada, Switzerland, Germany, Denmark, Spain, Great Britain, Finland, and France during 1870-2023. *Data sources*: Authors' own series and the JST database.

Figure A.4.: Model-based inflation expectations, part two

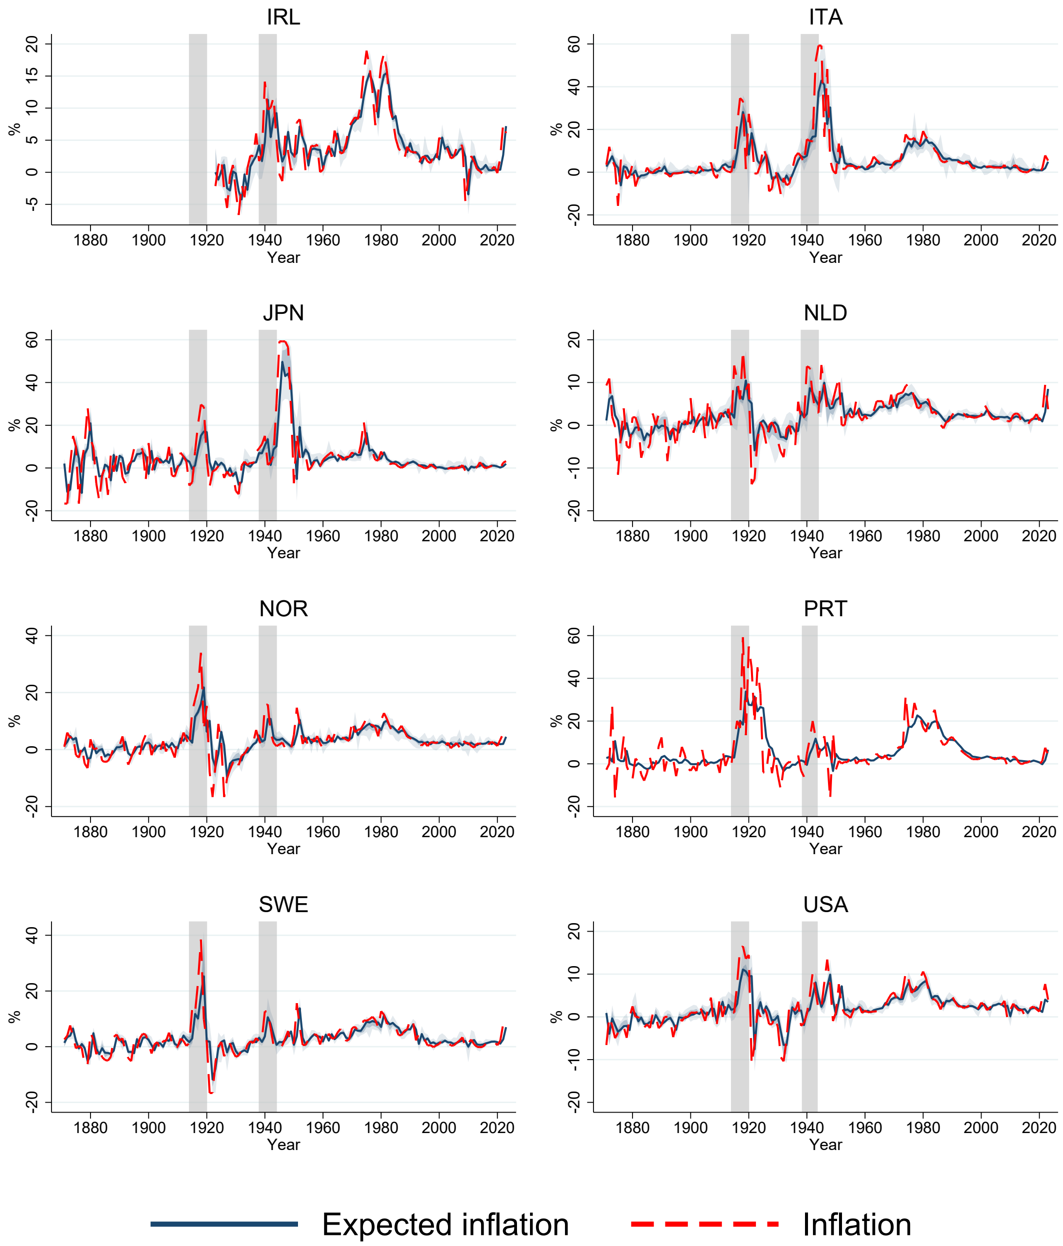

Notes: Blue solid line presents 1Y-ahead inflation expectations derived from the model as in Equation (1.1), together with bands at the interquartile range and 5th and 95th percentile of the distribution of model estimates. Red dashed line presents the actual realized inflation rate. The gray shaded areas present WWI and WWII. We present the series for Ireland, Italy, Japan, Netherlands, Norway, Portugal, Sweden, and USA during 1870-2023. *Data sources*: Authors' own series and the JST database.

Figure A.5.: Model-based inflation expectations for USA, 1870-2023

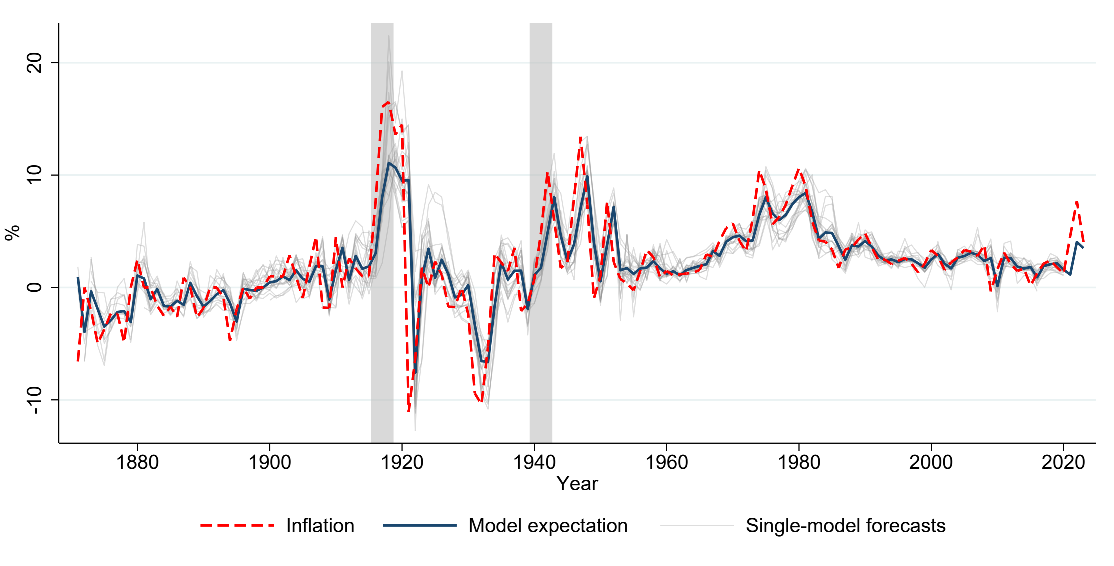

Notes: Blue solid line presents 1Y-ahead inflation expectations derived from our benchmark model as in Equation (1.1). As comparison, we also plot inflation expectations based on single models as gray lines (see Table A.1 & A.3 for a description of the models). Red dashed line presents the actual realized inflation rate. The gray shaded areas present WWI and WWII. USA, 1870-2023. *Data sources*: Authors' own series and the JST database.

Figure A.6.: Inflation surprises, part one

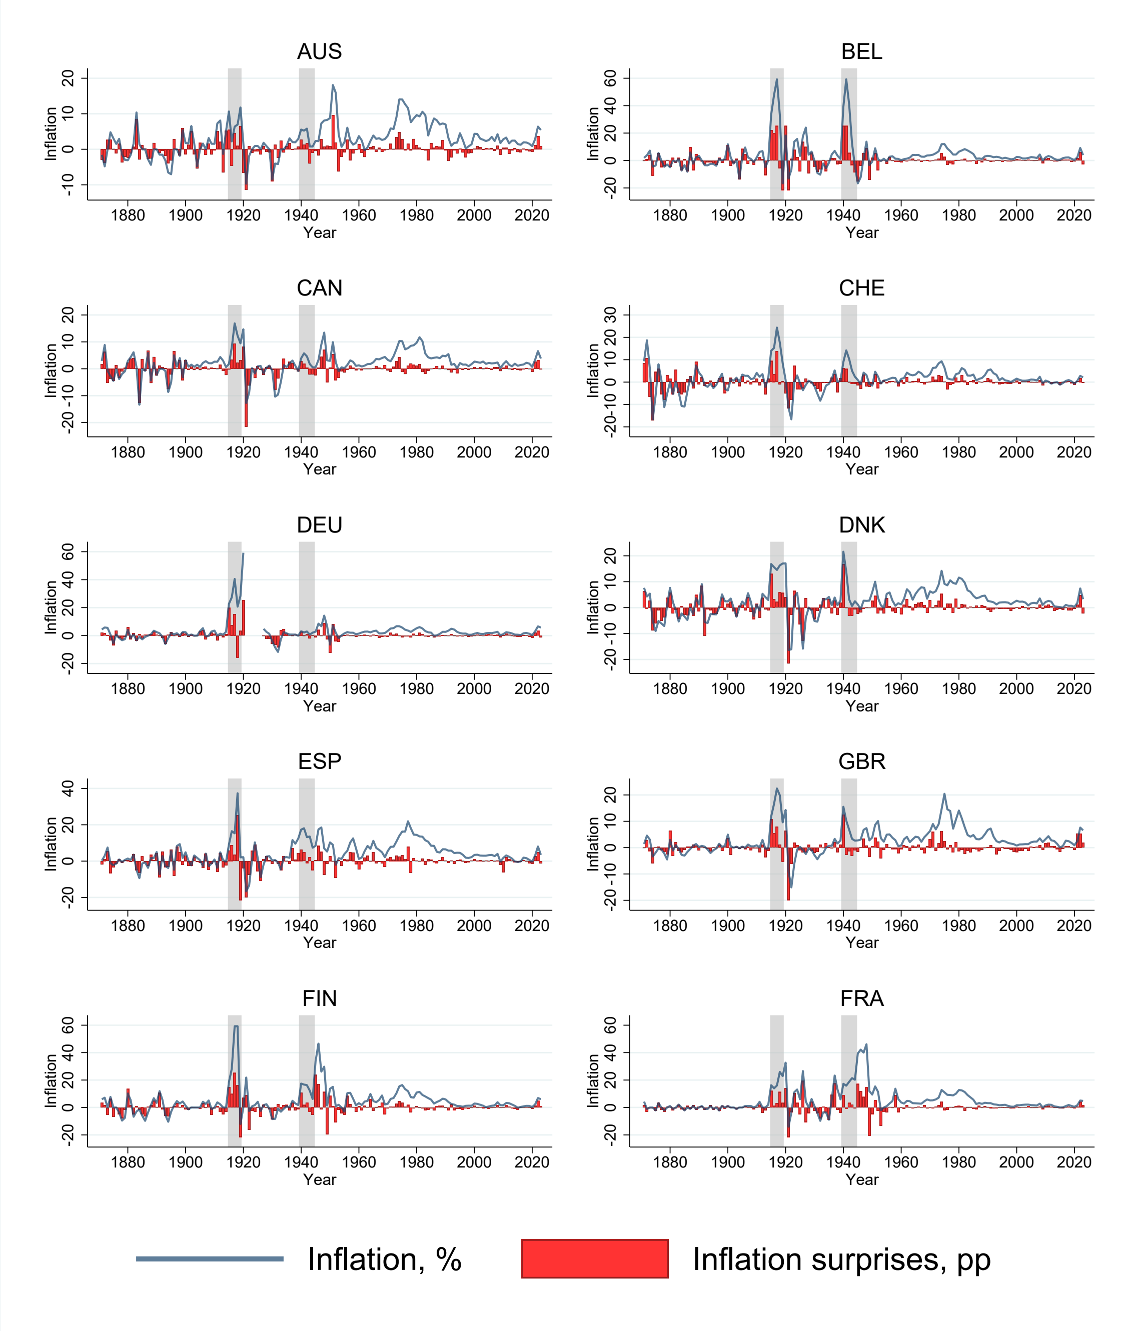

Notes: Inflation surprises (red bars) and realized inflation for Australia, Belgium, Canada, Switzerland, Germany, Denmark, Spain, Great Britain, Finland, and France during 1870-2023. Inflation surprises are computed as the difference between realized inflation and expectations from the previous period, i.e. $\varepsilon_{it}^{\pi} = \pi_{it} - \mathbb{E}_{t-1}[\pi_{it}]$, using model expectations for 1870-1965 and OECD expectations for 1965-2023. Inflation is measured in %, surprises in percentage points. Vertical gray shaded areas present world wars. We exclude the hyperinflation period for Germany.

Figure A.7.: Inflation surprises, part two

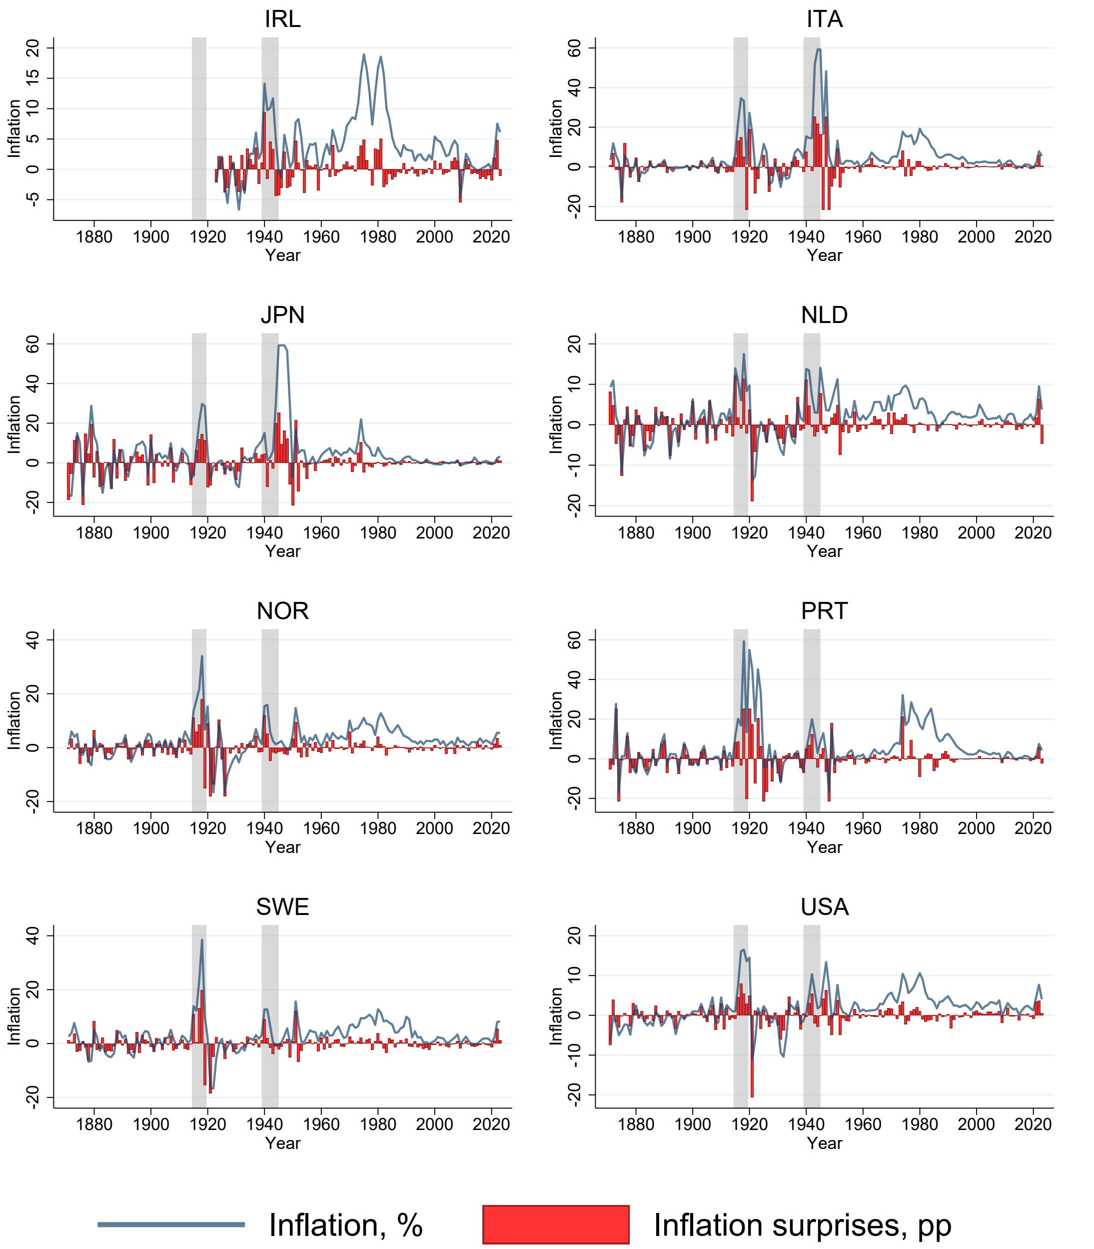

Notes: Inflation surprises (red bars) and realized inflation for Ireland, Italy, Japan, Netherlands, Norway, Portugal, Sweden, and USA during 1870-2023. Inflation surprises are computed as the difference between realized inflation and expectations from the previous period, i.e. $\varepsilon_{it}^{\pi} = \pi_{it} - \mathbb{E}_{t-1}[\pi_{it}]$, using model expectations for 1870-1965 and OECD expectations for 1965-2023. Inflation is measured in %, surprises in percentage points. Vertical gray shaded areas present world wars.

Figure A.8.: The effect of inflation surprises on the macroeconomy

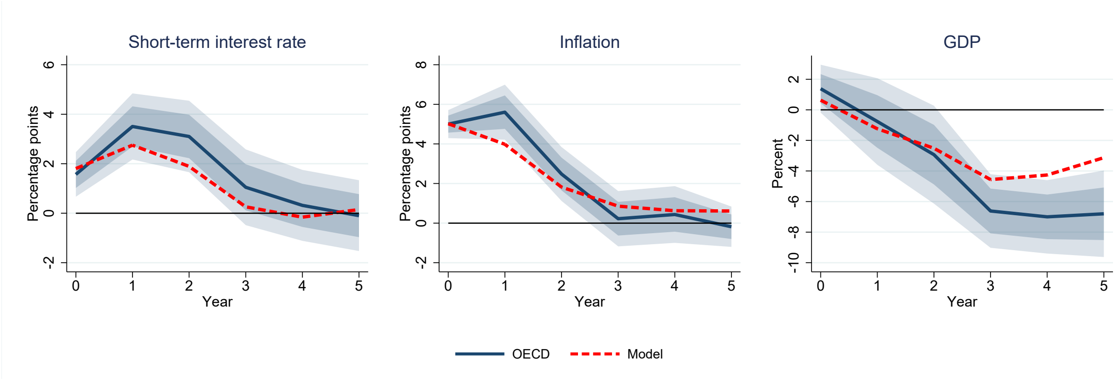

(a) Post-1965

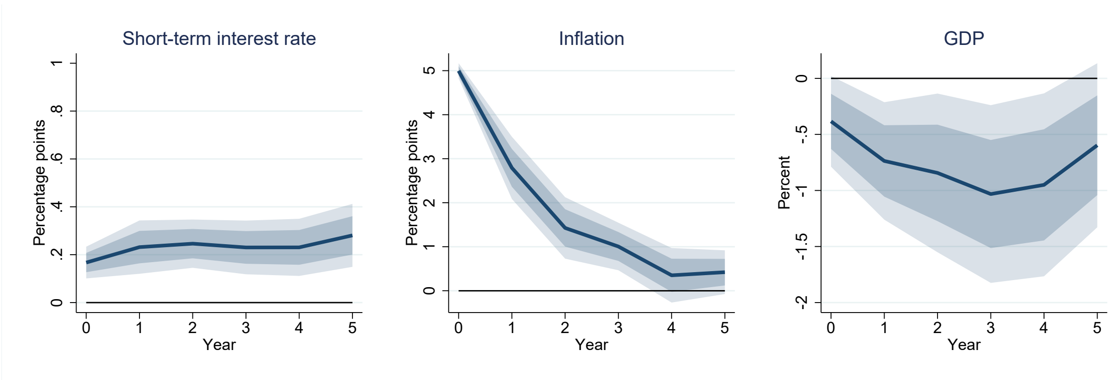

(b) Full-sample

Notes: Impulse responses for nominal short-term interest rate, inflation, and real GDP. Panel (a) presents the IRFs during the modern sample 1965-2023, where we have both set of inflation surprises (OECD-based and model-based, represented as blue solid and red dashed lines, respectively). In Panel (b) we present the IRFs for the full sample 1870-2023, using OECD surprises for the post-1965 sample and model surprises when OECD surprises are not available, i.e. 1870-1965. All effects are scaled to a 5 p.p. increase in inflation. Yearly frequency. Bands are at the 68% and 90% confidence level.

Figure A.9.: The effect of inflation surprises on asset prices and yields

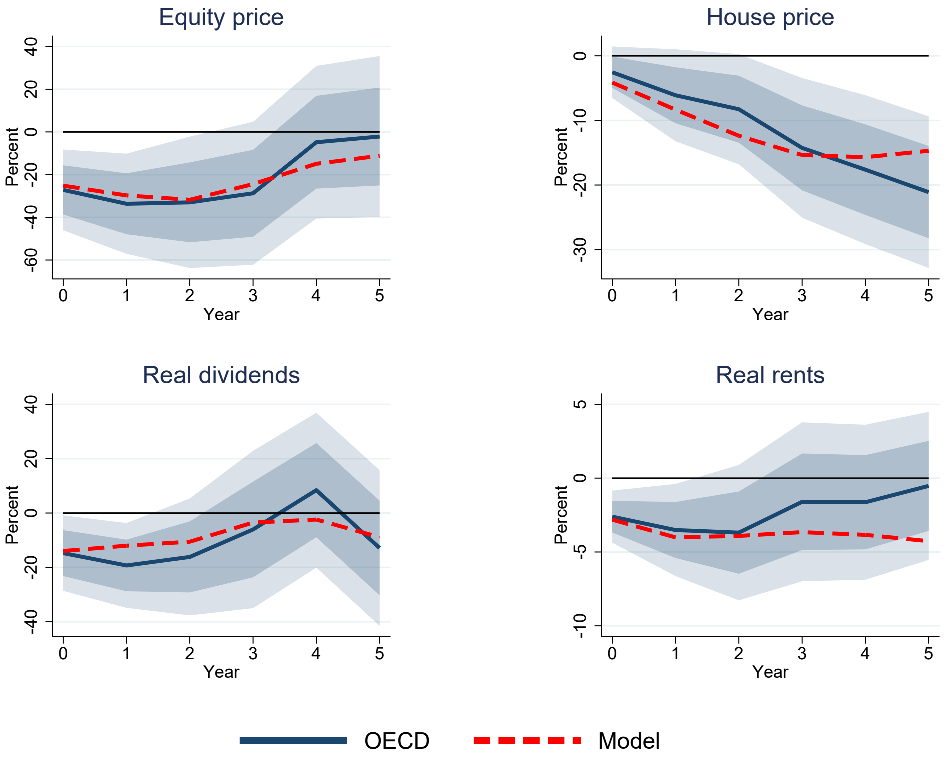

Notes: Impulse responses for real equity price, real house price, real dividends and real rents through panel-LP using OECD (end-of-the-year) surprises and model surprises, 1965-2023, yearly frequency. Effects are scaled to a 5 p.p. increase in inflation. Bands are at the 68% and 90% confidence level.

Figure A.10.: Inflation surprises on asset returns, half-yearly data during 1965-2023

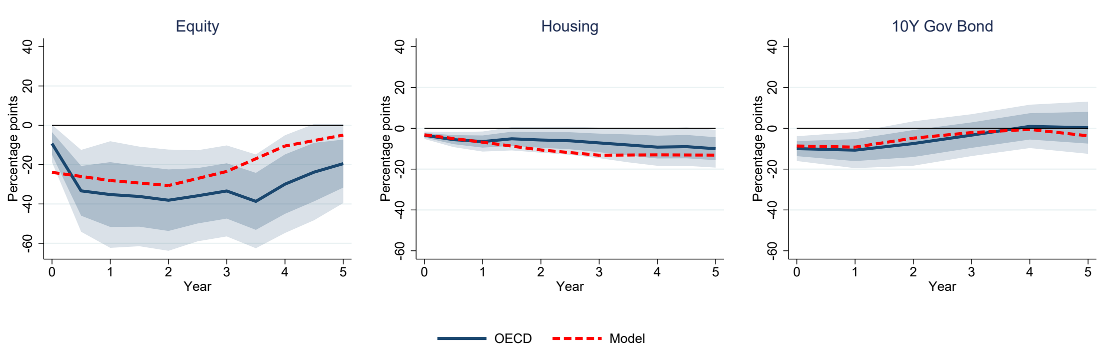

Notes: Impulse responses of cumulative real asset returns following a 5 p.p. surprise increase in inflation during 1965-2023. Results using OECD-surprises are at semi-annual frequency (blue solid line), while the results with model-surprises are at annual frequency (red dashed line). Bands are at the 68% and 90% confidence level.

Figure A.11.: Equity and house price responses in peg vs non-peg countries

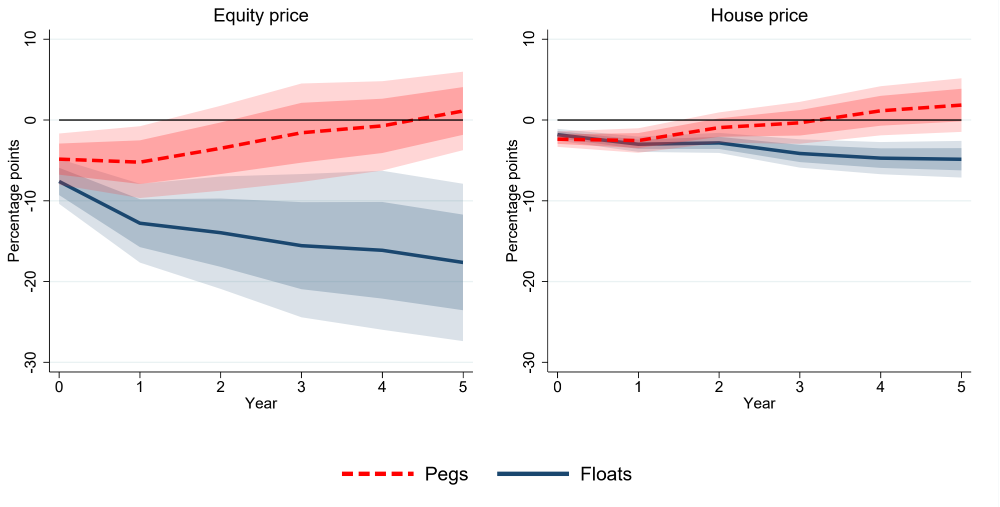

Notes: Impulse responses of real equity and house prices following a 5 p.p. inflation surprise in open-peg countries (red dashed line) and in float countries (blue solid line). Model & OECD surprises, scaled to a 5 p.p. increase in inflation. Sample of 18 countries during 1870-2023, yearly frequency. The two lines represent the estimate of $\beta_{A,h} + \beta_{A',h}$ and $\beta_{B,h}$ in the state-dependent LP. Shaded areas are the 90% and 68% confidence intervals.

Figure A.12.: Scatter of GDP surprises and inflation surprises

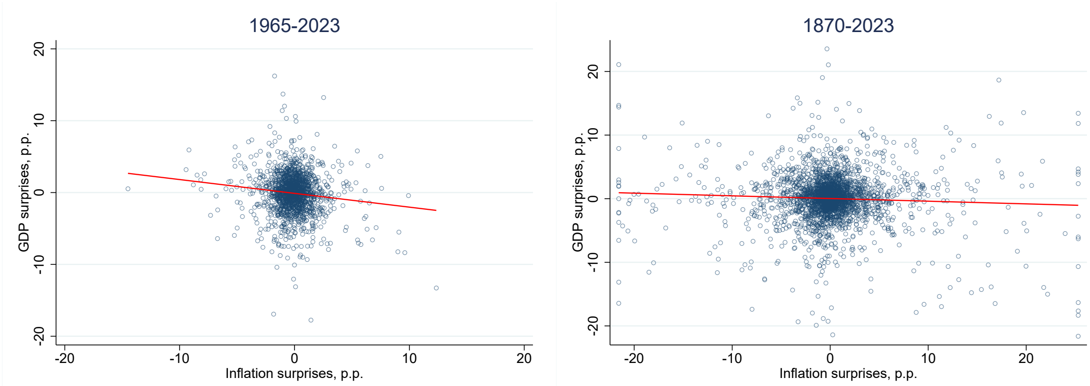

Notes: Scatterplot of GDP surprises and inflation surprises. *Left panel:* Half-yearly frequency, 1965-2023, using OECD-based inflation and GDP growth surprises. Correlation ≈ -0.11. *Right panel:* Yearly frequency, 1870-2023, using both OECD- and model-based inflation and GDP growth surprises (we used model-based measures for 1870-1965 and OECD-based measures after 1965). Correlation ≈ -0.06.

Figure A.13.: Fraction of demand- and supply-driven inflation surprises

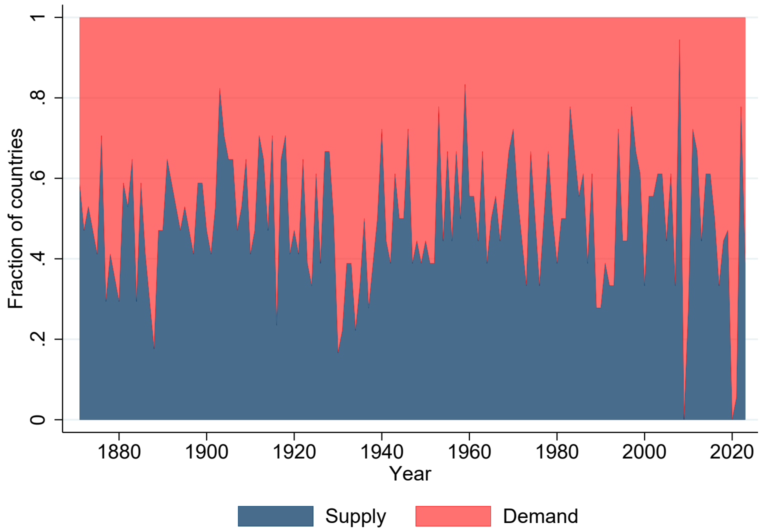

Notes: Timeline of the relative frequency of demand-driven (red) and supply-driven (blue) inflation surprises over 18 OECD countries during 1870-2020. Demand- and supply-driven inflation surprises are identified as described in Section 1.4.2.

1.B. Data description

In this appendix, we provide a more detailed description of the data we use for our analysis. The first dataset we exploit is the Jordà-Schularick-Taylor (JST) Macrohistory database by Jordà, Schularick, and Taylor (2017) and Jordà, Knoll, et al. (2019). This provides yearly data on key macroeconomic variables as well as returns on equity, housing, and bonds for 18 advanced economies for a period that spans from 1870 until 2020. The 18 advanced economies are the U.S., Japan, Germany, France, U.K., Ireland, Italy, Canada, Netherlands, Belgium, Sweden, Australia, Spain, Portugal, Denmark, Switzerland, Finland, Norway. In order to cover the recent inflation surge episode, we merge this dataset with macroeconomic and financial variables from the IMF and OECD databases, for the years 2021 to 2023.

1.B.1. Inflation rate

The first key variable that we use from the JST database is the inflation rate. Inflation is derived from the (log) changes in consumer prices, excluding observations where annual inflation exceeded 50% to remove potential outliers driving the results (Gabriel, 2021). Note that the consumer price data for the post-WWII period is the consumer price index (CPI) provided by the IMF, while for the earlier historical period the CPI as we know it today did not exist. Therefore, the JST database includes price indices of Taylor (2002), which draws on different price data such as GDP deflators and consumption deflators.

Table B.1.: Inflation rate (%) by country, mean (and standard deviation), 1870-2023

	Overall	Pre-WWI	1919-1938	World Wars	Post-WWII	Post-1982
Australia	2.91 (4.44)	0.36 (3.71)	-0.87 (3.53)	4.73 (3.32)	4.87 (3.91)	3.49 (2.36)
Belgium	3.88 (10.90)	-0.02 (4.83)	3.68 (10.86)	18.73 (27.26)	3.22 (3.03)	2.47 (1.81)
Canada	2.27 (4.45)	0.49 (3.98)	-1.38 (6.28)	4.65 (5.03)	3.66 (3.07)	2.57 (1.50)
Denmark	2.73 (5.75)	-0.03 (4.41)	-1.41 (8.47)	8.63 (7.90)	4.09 (3.26)	2.37 (1.71)
Finland	4.87 (10.06)	0.60 (5.20)	0.61 (6.38)	20.47 (21.35)	5.22 (5.63)	2.64 (2.14)
France	5.19 (9.13)	0.17 (1.46)	4.62 (12.82)	19.41 (10.59)	5.36 (7.35)	2.26 (1.86)
Germany	4.33 (10.86)	0.62 (2.70)	14.54 (25.81)	11.06 (13.17)	2.71 (2.57)	2.01 (1.46)
Ireland	4.07 (5.04)		-0.67 (3.43)	5.93 (5.62)	4.77 (4.75)	2.71 (2.62)
Italy	5.42 (10.71)	0.41 (3.97)	2.27 (9.71)	22.78 (20.32)	5.58 (6.89)	3.53 (2.92)
Japan	4.51 (11.84)	2.11 (9.33)	-1.23 (5.22)	17.72 (20.69)	4.61 (10.08)	0.71 (1.20)
Netherlands	2.24 (4.86)	0.06 (4.47)	-2.51 (5.61)	8.15 (5.62)	3.43 (2.64)	2.13 (1.63)
Norway	2.92 (5.88)	0.46 (3.24)	-2.35 (8.19)	9.61 (9.75)	4.22 (3.19)	3.11 (1.98)
Portugal	5.71 (11.21)	0.80 (6.75)	10.20 (20.23)	12.25 (15.26)	6.12 (8.04)	5.23 (5.97)
Spain	4.46 (6.76)	0.44 (3.93)	-0.40 (6.99)	13.17 (8.67)	6.14 (5.06)	3.71 (2.84)
Sweden	2.84 (5.78)	0.49 (3.43)	-2.45 (5.68)	9.05 (10.73)	4.19 (3.56)	2.92 (2.93)
Switzerland	1.58 (5.32)	-0.06 (6.03)	-2.50 (5.30)	8.22 (7.65)	2.16 (2.24)	1.31 (1.55)
UK	3.04 (5.15)	0.10 (1.90)	-1.56 (5.78)	8.59 (7.27)	4.68 (3.88)	2.91 (1.83)
USA	2.00 (4.30)	-0.64 (2.39)	-1.02 (6.08)	5.74 (6.13)	3.46 (2.74)	2.71 (1.39)

Notes: Sample: 18 countries from the JST database, 1870-2023. We exclude outlier periods, i.e. hyperinflation periods.

Figure B.1.: Historical inflation of 18 OECD countries, 1870-2023

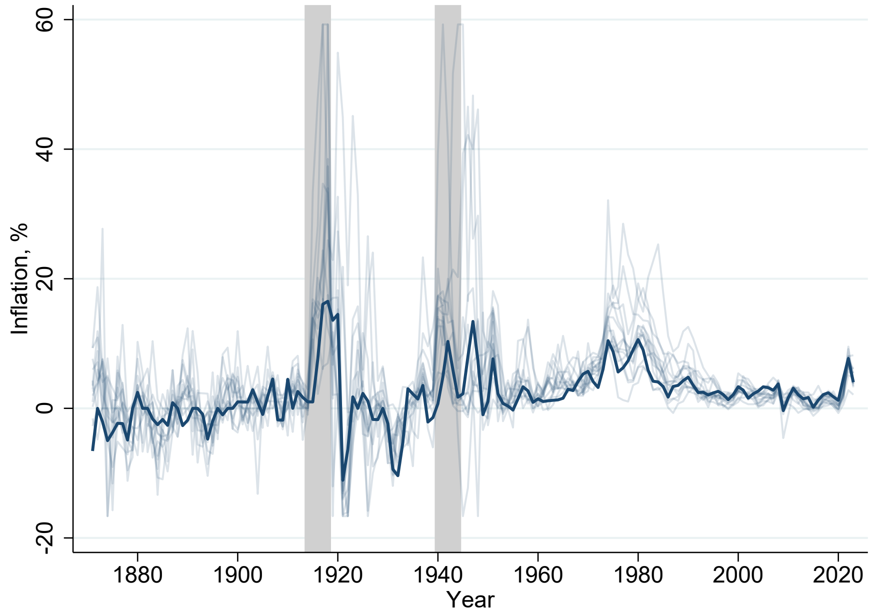

Notes: Thick blue line visualizes the inflation rate in the USA during 1870-2023. Thin blue lines represent inflation rates of the remaining 17 countries of our sample (excluding outlier years, e.g. German hyperinflation), 1870-2023. The vertical shaded areas represent the two World Wars. *Data sources*: JST database & IMF database.

In Figure B.1 we present the inflation rate of the US , together with the rest of countries in the JST sample (light blue). In addition, in Table B.1 we present the mean and standard deviation of inflation rates by country for different sub-periods. Overall, inflation rates were much lower during the pre-WWI period, which was the era of the gold standard. Between the two world wars, inflation was erratic with several countries moving even into deflation territory, especially during the Great Depression. In the aftermath of WWII, inflation rises remarkably, a period also titled as the Great Inflation period (1965-1982). Since 1982, inflation level has dropped for all countries (Great Moderation).

1.B.2. Asset returns

The JST database also provides asset returns data for equity, housing and government bonds. We extend returns data for the years 2021-2023 using the same underlying source of the JST. In Figure B.2 we present asset returns of USA as the benchmark (thick solid blue line), together with the results of all remaining countries as thin light blue lines. While housing and bond returns have low level of heterogeneity across countries, equity returns are very volatile and heterogeneous, especially after WWI. Nevertheless, country-specific heterogeneity dropped significantly during the last 20 years.

Figure B.2.: Historical asset returns

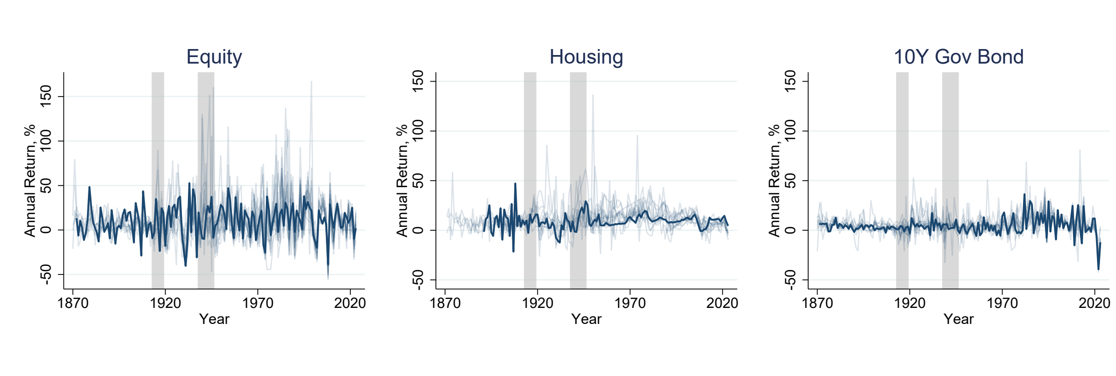

Notes: Equity, housing and 10Y government bond total nominal annual returns in USA (thick blue line). The thin blue lines represent the correspondent asset returns in the remaining 17 countries of our sample. Sample period is 1870-2023. Vertical shaded areas represent world wars. *Data sources*: JST database & IMF database

1.C. How did asset returns co-move historically with inflation?

In this section, a static *ex-post* analysis is carried out, regressing asset returns on inflation to understand how they co-moved throughout history at different horizons – measuring the so-called *inflation beta*. Even though this relation is not structural, as several confounding factors may drive some of the results, it is true that investors may not recognize the expected and unexpected component of inflation when considering inflation hedging, instead focusing more on the correlation between asset returns and past realized inflation levels (Bekaert and Wang, 2010). A more structural and dynamic analysis, looking at inflation *surprises* is carried out in the main text after introducing a measure of inflation expectations. Here, instead, the real return of asset k in country i at time t, r_{it}^k, is regressed on realized inflation in country i at time t, π_{it}:

$$r_{it}^k = \beta^k \pi_{it} + \gamma^k \mathbf{X}_{it} + \delta_i + \delta_t + \varepsilon_{it}^k, \tag{1.8}$$

where $\mathbf{X}$ is a set of controls, i.e. (2 lags of) real GDP growth, short-term interest rate, real stock price, CPI inflation, world GDP growth and world equity price, while δ_i and δ_t are country and time fixed effects, respectively. The parameter of interest is β, also called *inflation beta* (Bekaert and Wang, 2010), which represents how much asset k's return changes when realized inflation varies. A perfect inflation hedge would therefore result in $\beta = 0$. When $\beta < 0$, asset k is said to be an imperfect inflation hedge. Given the *long panel* structure of the data, i.e. with $T > 20$, the estimation is run through fixed effects regression with Driscoll and Kraay (1998) standard errors. Table C.1 shows results for the estimation of the panel regression (1.8).

For each asset, the first column presents the *unconditional* estimation of the inflation beta.

Table C.1.: Inflation beta

	Equity		Housing		10Y Gov Bonds	
	(1)	(2)	(3)	(4)	(5)	(6)
Inflation	-0.97***	-0.78***	-0.54***	-0.63***	-2.06***	-1.68***
	(0.18)	(0.18)	(0.07)	(0.08)	(0.76)	(0.38)
Country FE	No	Yes	No	Yes	No	Yes
Time FE	No	Yes	No	Yes	No	Yes
Controls	No	Yes	No	Yes	No	Yes
Observations	2298	2072	1925	1759	2326	2036

Notes: The estimates represent the β in the panel regression (1.8) for equity, housing and bond annual real returns on (contemporaneous) inflation. 18 countries from the JST panel, 1870-2023, excluding hyperinflation periods. Estimation using Driscoll and Kraay (1998) standard errors (in parentheses). Control variables are (2 lags of) real GDP growth, short-term interest rate, real stock price, CPI inflation, world GDP growth & world equity price. *,** , and *** denote significance at the 10%, 5%, and 1% level.

In the second column, the estimation is cleaned from other basic macroeconomic factors. Overall, the estimated inflation betas do not change significantly between the unconditional and conditional estimation. Stocks and bonds have a strongly negative inflation beta, very close to -1 and -2, respectively.[42] Housing seems to be the best inflation hedge among these asset classes on impact, even though not perfect, with a coefficient significantly negative around -0.5. This result is similar to Fang, Liu, and Roussanov (2022), who consider Real Estate Investment Trust (REIT) only in the US (with quarterly data), finding a (nominal) inflation beta close to zero and non significant, i.e. considering housing as a bad inflation hedge.

Table C.2.: Long-run inflation beta

	Equity		Housing		10Y Gov Bonds	
	(1)	(2)	(3)	(4)	(5)	(6)
10Y Inflation	-0.10***	-0.07***	-0.26***	-0.20**	-0.99***	-0.99***
	(0.03)	(0.02)	(0.04)	(0.10)	(0.00)	(0.00)
Country FE	No	Yes	No	Yes	No	Yes
Controls	No	Yes	No	Yes	No	Yes
Observations	2101	1933	1694	1552	2118	1900

Notes: The estimates represent the β in the panel regression (1.8) for equity, housing and bond 10-year cumulative real (log) returns on 10-year inflation. 18 countries from the JST panel, 1870-2023, excluding hyperinflation periods. Estimation using Driscoll and Kraay (1998) standard errors (in parentheses). Control variables are (2 lags of) real GDP growth, short-term interest rate, real stock price, CPI inflation, world GDP growth & world equity price. *,** , and *** denote significance at the 10%, 5%, and 1% level.

[42] Bekaert and Wang (2010) and Fang, Liu, and Roussanov (2022) find that responses of stocks to inflation are ambiguous, lying around zero, due to high standard errors. Nevertheless, note that they measure nominal or excess returns, instead of real returns.

As Boudoukh and Richardson (1993) show, extending the analysis to a longer horizon might help reconcile the theory with counter-intuitive empirical findings. Table C.2 presents results considering the 10-year inflation and 10-year asset returns. Results show that all real assets (i.e. equity and housing) tend to converge to an inflation beta of 0 in the long run. However, point estimates indicate a smaller yet significant negative relation, for all asset classes, and especially for 10Y government bonds, also in the long run.

One interesting feature of our data is the possibility to decompose the total return of housing and equity into capital gains and dividends/rents. Table C.3 shows how the different parts of asset total returns co-moved historically with inflation. Most of the fluctuations of total returns, and the whole ability to hedge against inflation, seems to be driven by capital gains for both equity and housing. This is evident when comparing the capital gains inflation beta with those on the total returns in the previous table. The intuition is, as in the main text, that dividends and rents tend to be slower in adjusting to macroeconomic fluctuations, hence suffering more from volatile inflation.

Table C.3.: Inflation beta on capital gains and dividends/rents

	Equity		Housing	
	Cap. gains	Dividends	Cap. gains	Rents
Inflation	-0.78***	-1.42***	-0.85***	-0.87***
	(0.18)	(0.31)	(0.20)	(0.08)
Country FE	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes
Observations	2074	2014	1878	1719

Notes: The estimates represent the β in the panel regression (1.8) for real equity capital gains, real dividends, real housing capital gains and real rents on (contemporaneous) inflation. 18 countries from the JST panel, 1870-2023, excluding hyperinflation periods. Estimation using Driscoll and Kraay (1998) standard errors (in parentheses). Control variables are (2 lags of) real GDP growth, short-term interest rate, real stock price, CPI inflation, world GDP growth & world equity. *,**, and *** denote significance at the 10%, 5%, and 1% level.

1.D. Quality checks for historical inflation expectation

In this appendix we provide a range of robustness checks to test the validity of our inflation forecast measures.

1.D.1. Mean vs. Median

We test the consistency of all forecast models that we use to compute a single measure of inflation expectations by taking the mean instead of the median. In terms of our model, this means that $\omega' = \frac{1}{n}$ in

$$\mathbb{E}_t[\pi_{i,t+1}] = \quad \boldsymbol{\omega'} \quad \boldsymbol{\hat{\Pi}_{t+1|t}} = \boldsymbol{\omega'} \begin{pmatrix} \hat{\pi}^1_{t+1|t} \\ \vdots \\ \hat{\pi}^j_{t+1|t} \\ \vdots \\ \hat{\pi}^n_{t+1|t} \end{pmatrix},$$

$$\hat{\pi}^j_{t+1|t} = \quad \boldsymbol{A'} \quad \boldsymbol{X_{t:t-k}} + \varepsilon^j_t.$$

Figure D.1 and Figure D.2 present the results. Overall, the mean forecast measure aligns very closely to the median for all countries.

Figure D.1.: Model forecast: median vs. mean weights, part one

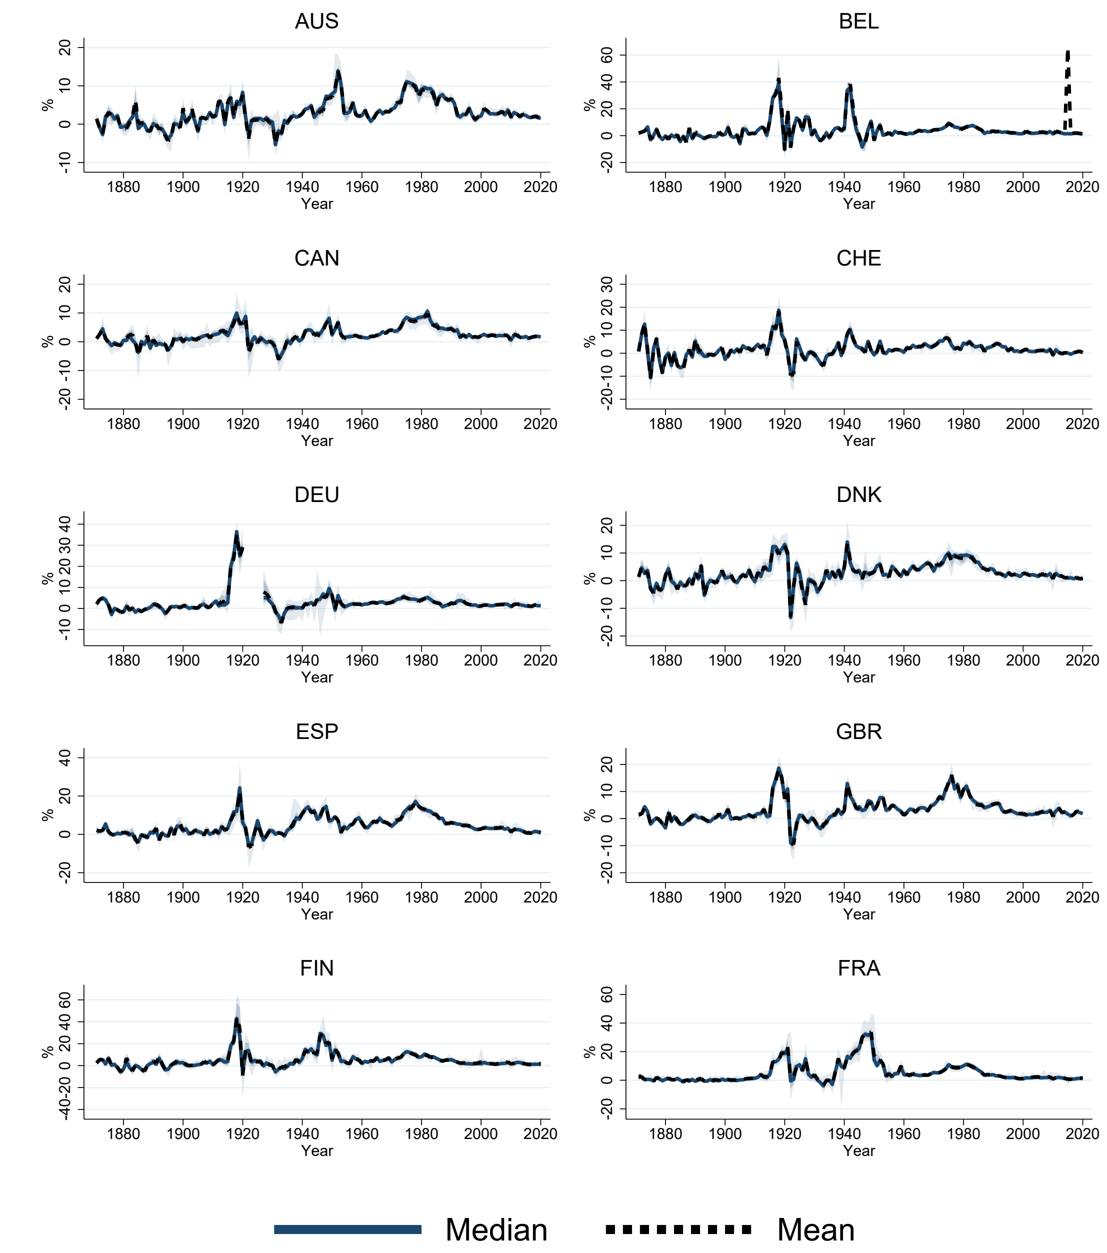

Notes: Model estimates of inflation expectations using the median (blue solid line) or the mean (dashed black line) as weights ω for the different models in Equation (1.1), together with bands at the interquartile range and 5th and 95th percentile of the distribution of model estimates. The ordering is: Australia, Belgium, Canada, Switzerland, Germany, Denmark, Spain, UK, Finland, and France. Sample period is 1870-2023. Data source: Authors' series.

Figure D.2.: Model forecast: median vs. mean weights, part two

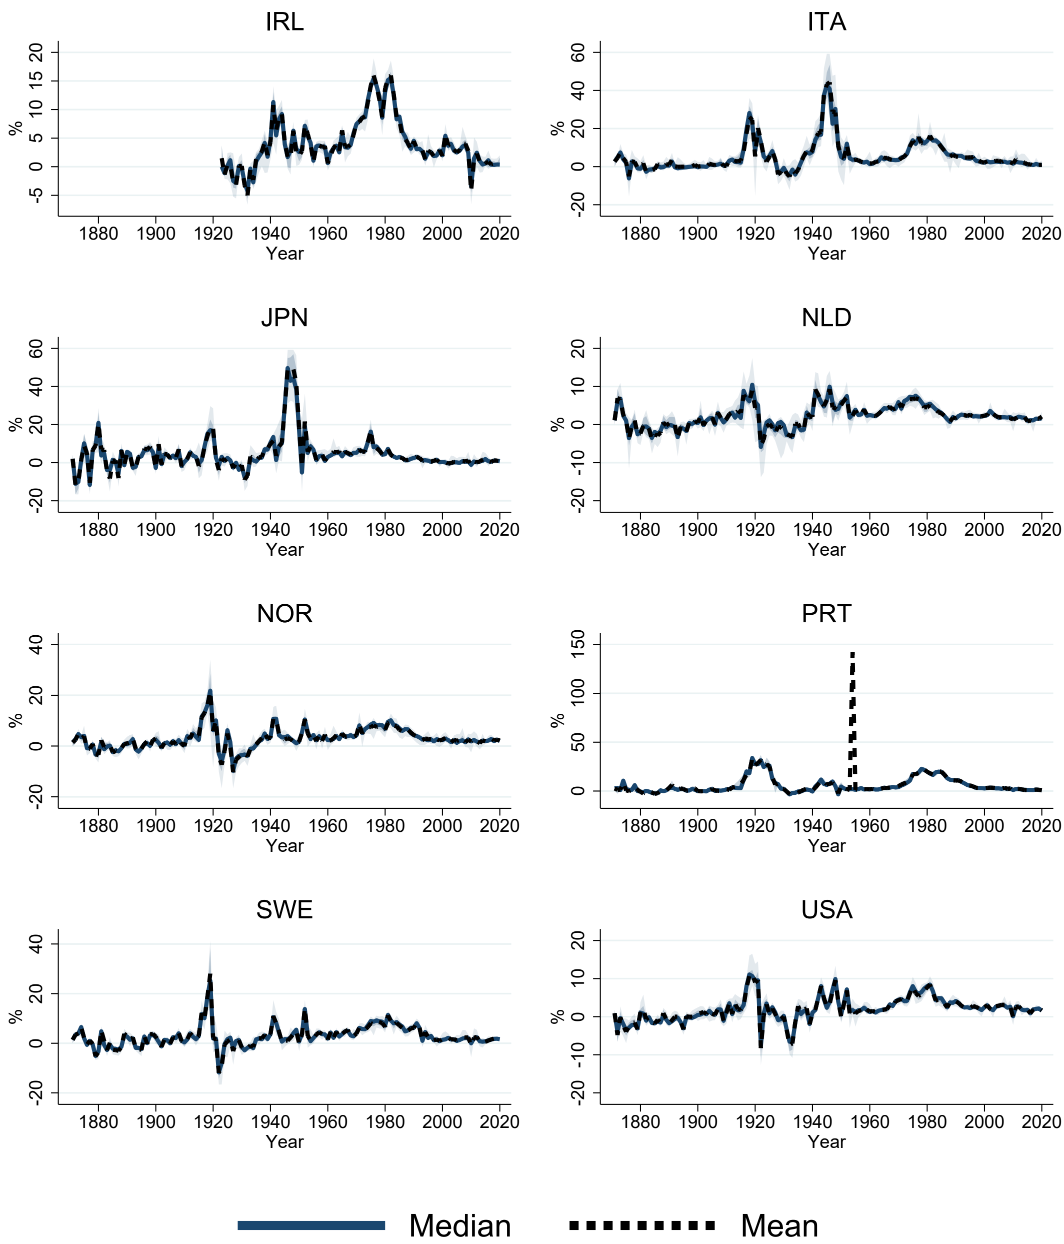

Notes: Model estimates of inflation expectations using the median (blue solid line) or the mean (dashed black line) as weights ω for the different models in Equation (1.1), together with bands at the interquartile range and 5th and 95th percentile of the distribution of model estimates. The ordering is: Ireland, Italy, Japan, Netherlands, Norway, Portugal, Sweden, and USA. Sample period is 1870-2023. Data source: Authors' series.

1.D.2. The OECD Outlook: 1965-2023

We use data from the OECD Outlooks that provide "real-time" inflation expectations. Starting from 1965, the forecast division of the OECD started to provide twice-yearly projections of various economic variables, such as inflation (the first issue published mostly in June/July and the second issue in December). For the first ten years, the OECD provided these forecasts for only Canada, France, Germany, Italy, Japan, UK and USA. After 1975, they started to provide forecasts also for the other countries (Australia, Belgium, Netherlands, Denmark, Ireland, Finland, Norway, Portugal, Sweden, Switzerland, Spain).[43]

To be consistent with our benchmark estimates of inflation expectations, we utilize the one-year-ahead inflation forecasts in the OECD Outlooks. Figure D.3 and Figure D.4 compare the OECD forecasts with our model estimates (together with bands at the interquartile range and 5th and 95th percentile of the distribution of model estimates). Reassuringly, the OECD forecasts are always within the confidence bands and often very close to our model central estimates.

[43]The forecasts include also other countries that we exclude to match the entire sample with the rest of the data in the JST database, i.e. Austria, Greece, Iceland, Luxembourg, New Zealand and Turkey.

Figure D.3.: Model-based inflation expectations vs OECD forecasts, part one

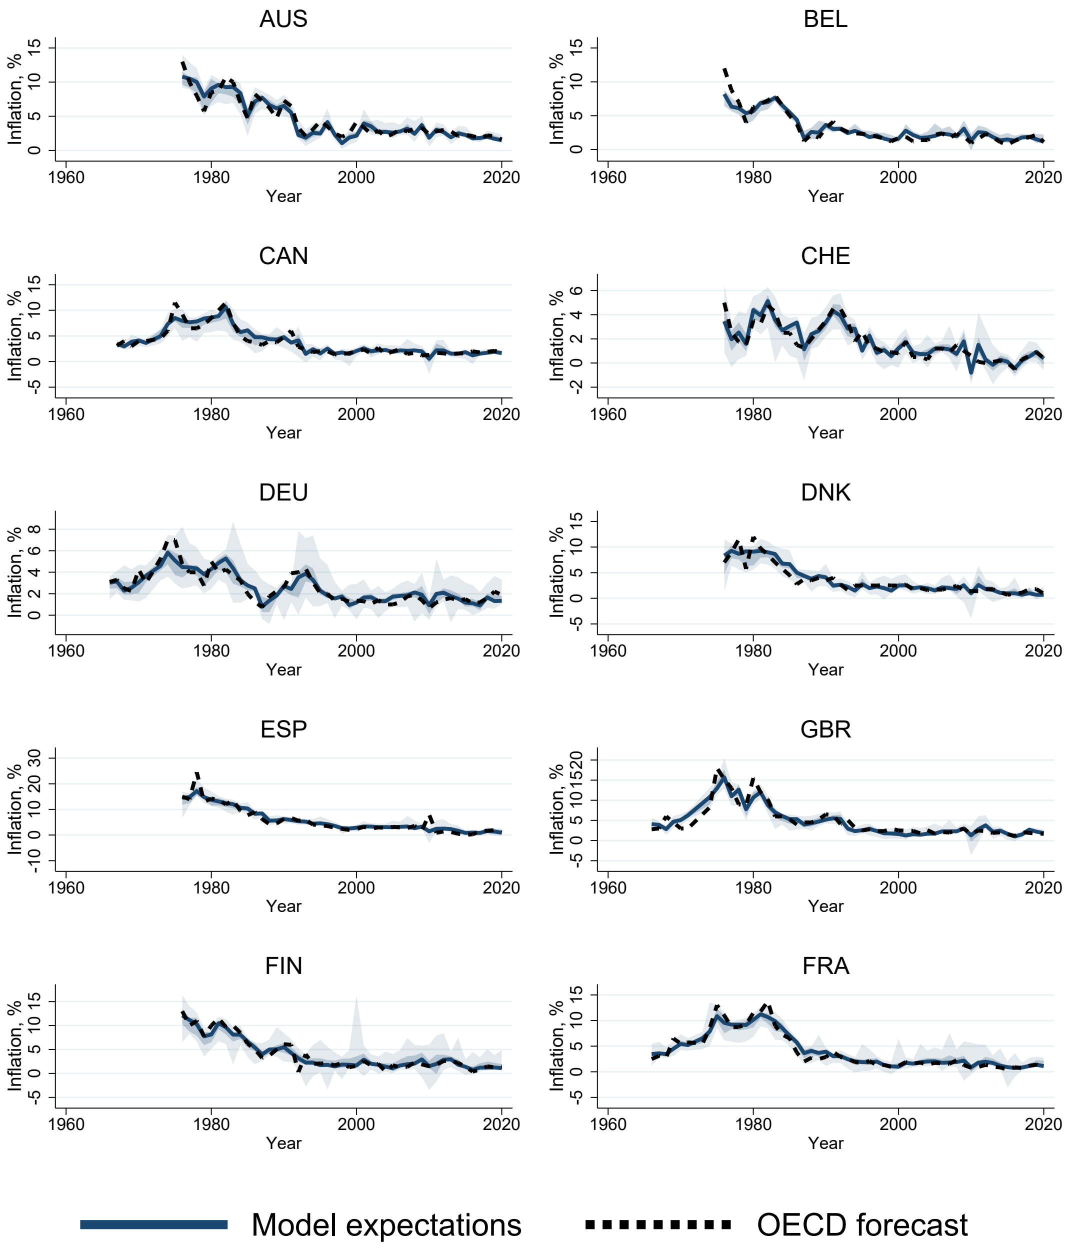

Notes: Inflation expectations (blue solid line), derived from the model in (1.1), together with bands at the interquartile range and 5th and 95th percentile of the distribution of model estimates, compared to OECD historical forecasts from Economic Outlooks (black dashed line). Australia, Belgium, Canada, Switzerland, Germany, Denmark, Spain, UK, Finland, France, 1965-2023.

Figure D.4.: Model-based inflation expectations vs OECD forecasts, part two

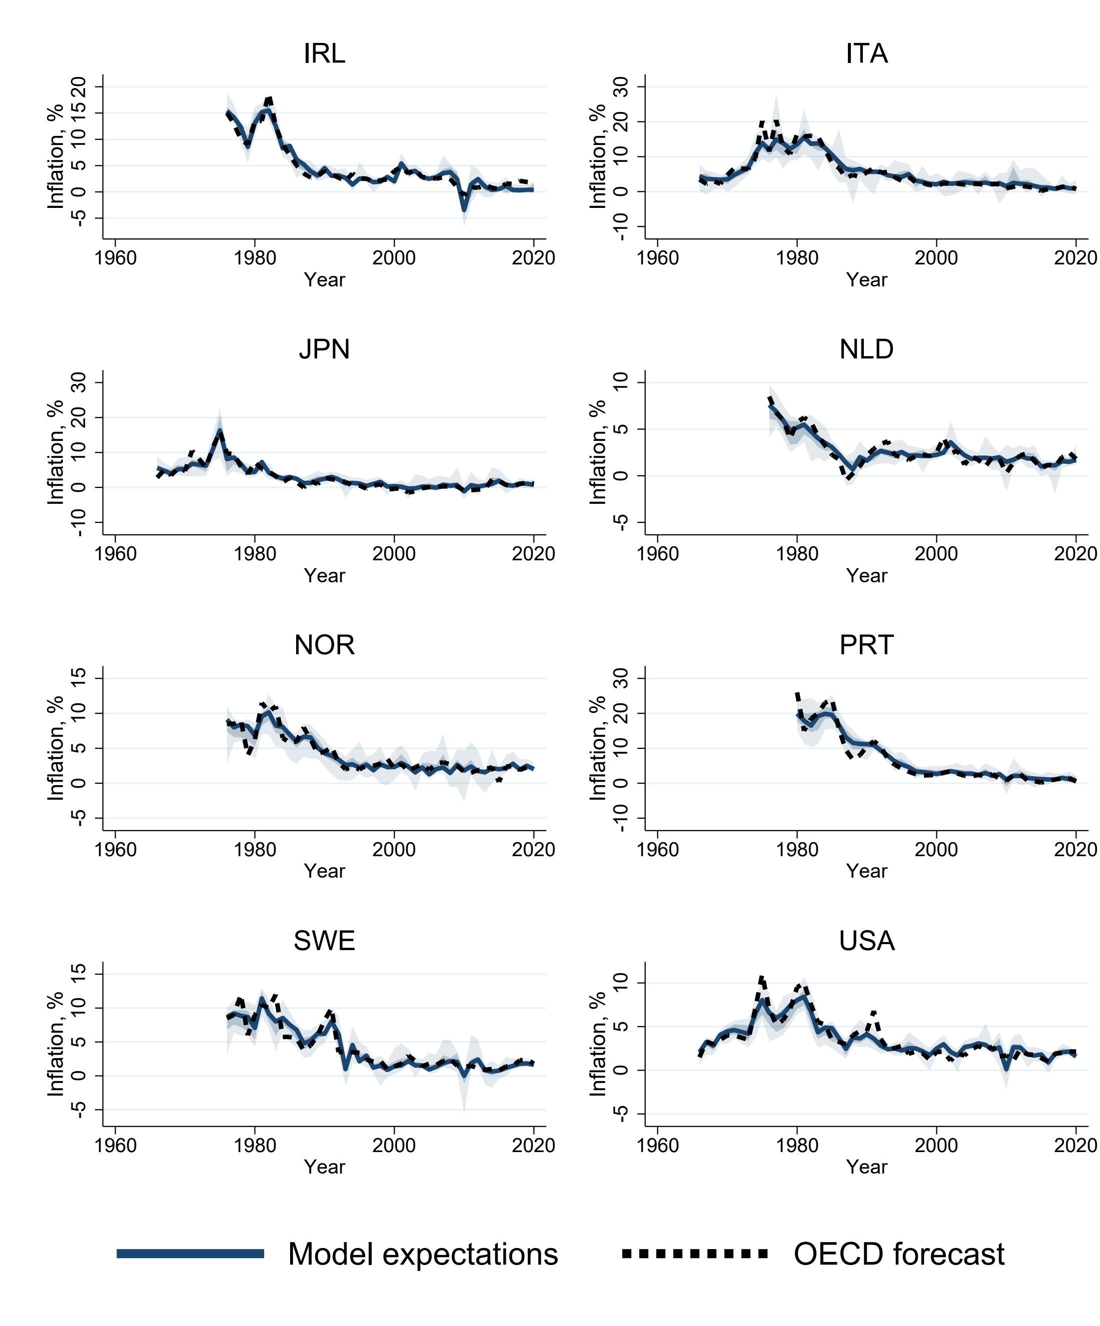

Notes: Inflation expectations (blue solid line), derived from the model in (1.1), together with bands at the interquartile range and 5th and 95th percentile of the distribution of model estimates, compared to OECD historical forecasts from Economic Outlooks (black dashed line). Ireland, Italy, Japan, Netherlands, Norway, Portugal, Sweden, USA, 1965-2023.

1.E. Robustness exercises: benchmark analysis

In this appendix, we provide a range of robustness checks of our benchmark analysis in Section 1.3.

1.E.1. Impulse responses excluding large global inflation episodes

Figure E.1.: Asset responses excluding the Great Inflation and the COVID surge

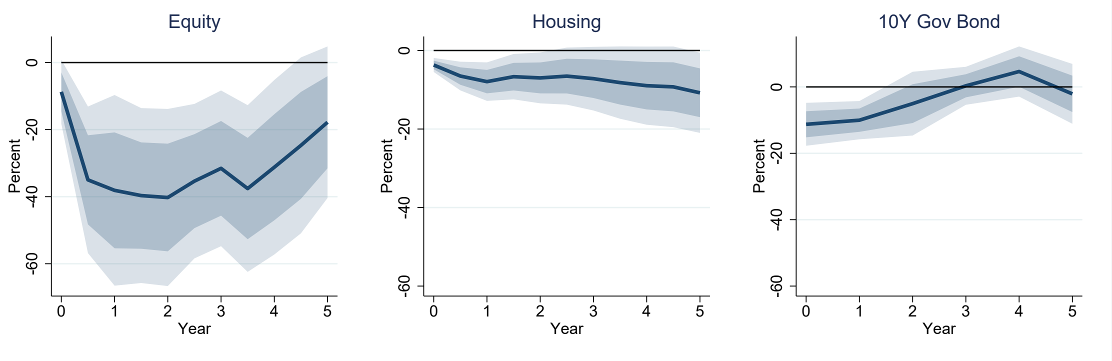

Notes: Impulse responses of the cumulative real returns of stocks, housing and 10Y government bonds to inflation surprises, using our OECD surprises that are normalized to induce a 5p.p. increase in inflation. We exclude the Great inflation period and the inflation surge during the pandemic. Sample period: 1982h1-2019h2, halfyearly frequency. Shaded areas are the 90% and 68% confidence intervals using Driscoll and Kraay (1998) standard errors.

Figure E.2.: Asset responses excluding World wars

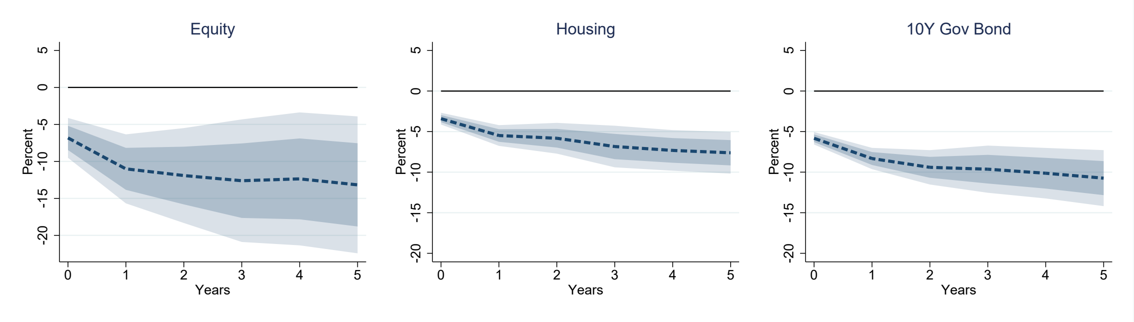

Notes: Impulse responses of the cumulative real returns of stocks, housing and 10Y government bonds to inflation surprises. Both model-based and OECD-based surprises that are normalized to induce a 5 p.p. increase in inflation. We exclude World war periods, i.e. 1913-1919 and 1939-1946. Sample period: 1870-2023, yearly frequency. Shaded areas are the 90% and 68% confidence intervals using Driscoll and Kraay (1998) standard errors.

1.E.2. Event study: large inflation surprises

In this appendix, we exploit an event study framework to investigate the contribution of large inflation surprises to asset return dynamics. The event study we conduct is as follows:

$$r_{it} = \alpha + \sum_{j=2}^{J} \beta_j (\text{lag}_j)_{it} + \sum_{k=1}^{K} \gamma_k (\text{lead}_k)_{it} + \mu_i + \lambda_t + \varepsilon_{it}, \tag{1.9}$$

where r_{it} is the (log) annual real return for each asset i, with i being either stocks, housing, or long-term government bonds. Notice that the left-hand side of (1.9) is not the cumulative return, but rather the 1-year real returns, hence results are not directly comparable to local projections in the main text. These returns are regressed with the lead and lag operators around inflation surprises larger than 5 percentage points. Finally, we also control for country- and time fixed effects (μ_i and λ_t, respectively). Figure E.3 presents the results. Signs and magnitudes in year 0 and 1 are comparable to our benchmark results in the main text. Large surprises seem to hit asset returns mostly in the very short run. At the same time, returns hardly recover the past real losses in the long run, consistent with our benchmark results. This result highlights the importance of considering milder fluctuations in inflation and *cumulative* returns for understanding overall responses of asset returns.

Figure E.3.: Event study: large inflation surprises and asset returns

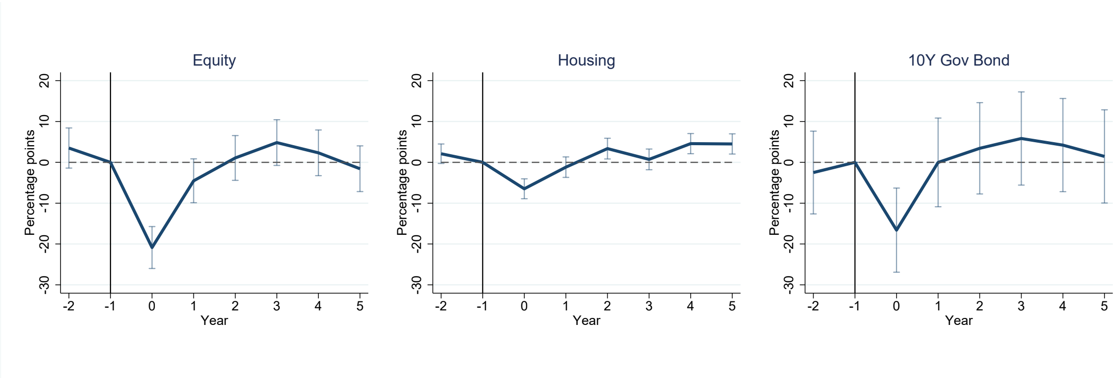

Notes: Yearly asset returns responses after a positive inflation surprise greater than 5 p.p.. Inflation surprises hit at time 0. Yearly frequency, 1870-2023, using model surprises before 1965 and OECD surprises for 1965-2023. World wars are included in the estimation sample while hyperinflation periods are excluded. Bands are at 90% confidence level.

1.E.3. Responses to single-model inflation surprises

Our benchmark model to estimate inflation expectations exploits a technique called *forecasts combination*, where we estimate a range of different models to forecast inflation. Afterwards, we pool all estimates together and take the median as our final measure for inflation expectations $E_{t-1}[\pi_{it}]$, and use them to derive inflation surprises $\varepsilon^{\pi}_{it} = \pi_{it} - \mathbb{E}_{t-1}[\pi_{it}]$. But in principle, we could have derived ε^{π}_{it} using any of the inflation forecasts from only one model of choice.

In this appendix, we derive 20 inflation surprise measures using all forecast models listed in Appendix Table A.1 and Table A.3. Afterwards, we run panel local projections using each of these series. Figure E.4 presents the impulse responses of asset returns using inflation surprises based on single models (light gray lines), compared to our benchmark response (thick blue line). Overall, all impulse responses are mostly concentrated around the main response line and within its confidence bands. This confirms the robustness of our results. However, some models consistently underestimate the response of returns, even though complying with the same sign. This result suggests that using single model estimates may not entirely grasp the true response of asset returns to an inflation surprise, thus techniques taking into account structural breaks and possible misspecification are needed. Our benchmark model, which combines a range of forecast models, is much more robust to such deficiencies. The same exercise with similar conclusions is carried out on macroeconomic variables in Figure E.5.

Figure E.4.: Asset responses: main model vs single models

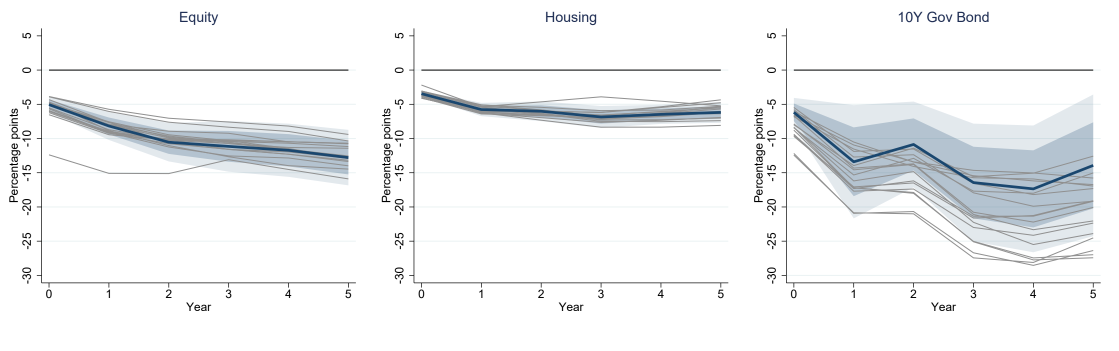

Notes: Impulse responses for cumulative real returns of stocks, housing and 10Y government bonds on inflation surprises, using our main model (blue thick line) or the single models (grey thin line) included in Equation (1.1). Surprises are always normalized to induce a 5p.p. increase in inflation. 1870-2023. Countries: USA, Canada, Germany, France, UK, Italy, Japan, Australia, Belgium, Switzerland, Denmark, Spain, Finland, Ireland, Netherlands, Norway, Portugal and Sweden. Shaded areas are the 90% and 68% confidence intervals.

1.E.4. Impulses responses using LP-IV

Another popular method to draw impulses responses in the macroeconomic literature is to use estimated shocks as proxies for their fundamental counterpart. Using local projections, this method entails using the estimated shocks as instrumental variable for the proxied variable. We adopt such LP-IV methodology as robustness exercise, using our inflation surprise series as an IV for contemporaneous inflation. In practice, we use contemporaneous and (up to 2) lags of the inflation surprise to instrument contemporaneous inflation. The

Figure E.5.: Macroeconomic responses: main model vs single models

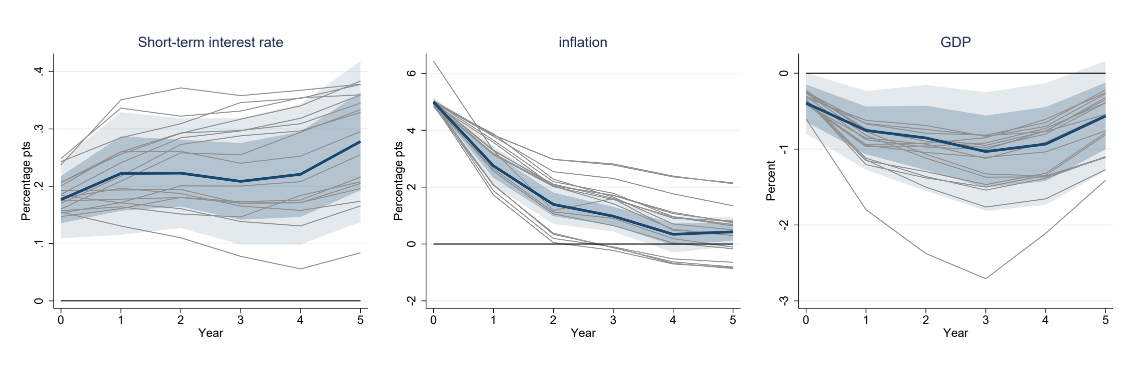

Notes: Impulse responses for interest rate, inflation and real GDP on inflation surprises, using our main model (blue thick line) or the single models (grey thin line) included in Equation (1.1). Surprises are always normalized to induce a 5p.p. increase in inflation. 1870-2023. Countries: USA, Canada, Germany, France, UK, Italy, Japan, Australia, Belgium, Switzerland, Denmark, Spain, Finland, Ireland, Netherlands, Norway, Portugal and Sweden. Shaded areas are the 90% and 68% confidence intervals.

rest of the specification is equal to the one in Equation 1.5, with the same set of controls. Figure E.6 reports the impulse responses using LP-IV. Estimates align closely to the direct OLS regression, reported also in the main text in Figure 1.5 – highlighting the robustness of the negative and persistent response of cumulative real asset returns.

Figure E.6.: Inflation surprises as instruments

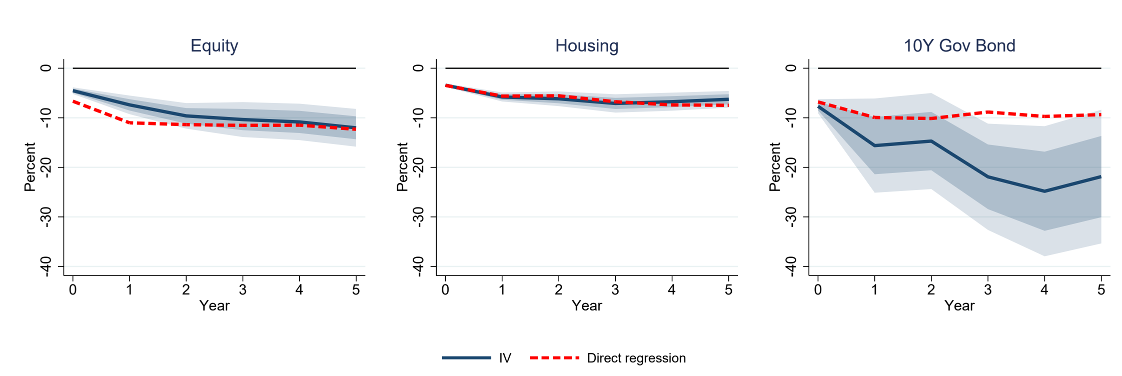

Notes: Impulse responses for equity, housing and 10Y government bond cumulative real returns on inflation, using (contemporaneous and 2 lags of) inflation surprise as IV. The surprise series is the OECD surprises during the post-1965 period and the model-based surprises for the historical sample 1870-1965. Sample of 18 countries from the JST. Shaded areas are the 90% and 68% confidence intervals.

1.E.5. Do asset returns react symmetrically to inflation?

Asymmetric effects of macroeconomic shocks are widely present in the literature, e.g. Ramey and Zubairy (2018). In this appendix, we investigate whether a surprise increase in inflation has a different impact on asset returns compared to a surprise decrease in inflation. Figure E.7 reports the impulse response functions when we estimate our benchmark panel

LP model using only positive or negative inflation surprises (blue solid line and red dashed lines, respectively). Our results show that all three asset classes react in a similar way to inflation surprises of both directions.

Figure E.7.: Asset returns and asymmetric responses to inflation

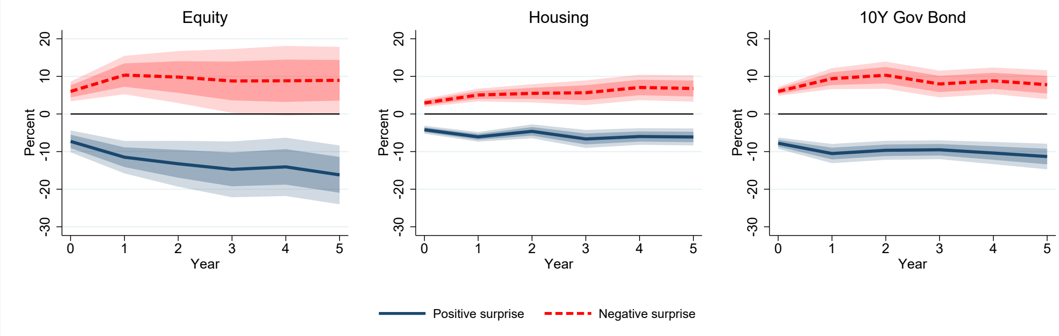

Notes: Impulse responses for equity, housing, bond on positive (blue solid line) and negative (red dash line) inflation surprises. Sample of 18 countries from the JST, 1870-2023, excluding world wars. The two lines represent the estimate of $\beta_{A,h}$ (response to a positive surprise) and $-\beta_{B,h}$ (response to a negative surprise) that we estimate using Equation (1.6), at various horizons $h = 0, 1, 2, \ldots$. Shaded areas are the 90% and 68% confidence intervals.

1.F. Robustness exercises: peg vs. float countries

In this appendix we provide the following robustness exercises for Section 1.4.1. In Appendix Section 1.F.1 we exclude the euro areas from the peg country sample. In Appendix Section 1.F.2 we focus on the pre-Bretton Woods periods and investigate the effect of inflation surprises on asset returns for pegged and float countries. Finally, in Appendix Section 1.F.3 we compare the impulse responses of real asset returns to an inflation surprise in Gold-standard countries vs. non-Gold standard countries.

1.F.1. Open pegs vs floats: excluding Euro countries

Figure F.1.: Asset returns and macroeconomic responses in peg vs non-peg, excluding euro

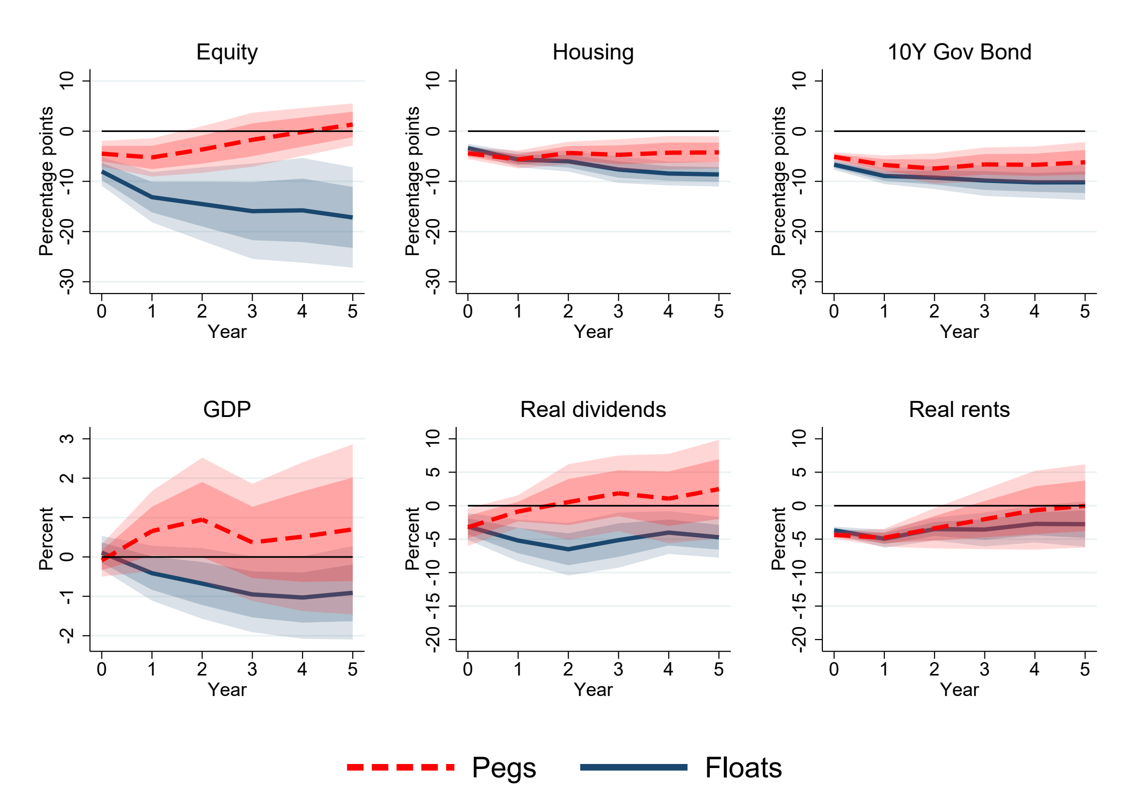

Notes: Impulse responses for real asset returns, real GDP, and real dividends to inflation surprises in open-peg countries (red dash line) and in float countries (blue solid line). We exclude Belgium, Germany, Spain, Finland, France, Ireland, Italy, Netherlands and Portugal as pegs starting from introduction of the euro (i.e. year 2000). Model & OECD surprises, scaled to a 5 p.p. increase in inflation. Sample of 18 countries from the JST, 1870-2023, yearly frequency. The two lines represent the estimate of $\beta_{A,h}$ and $\beta_{B,h}$ of the state-dependent LP model. Shaded areas are the 90% and 68% confidence intervals.

1.F.2. Open pegs vs floats: pre-Bretton Woods

Figure F.2.: Asset returns and macroeconomic responses in peg vs non-peg, pre-1973

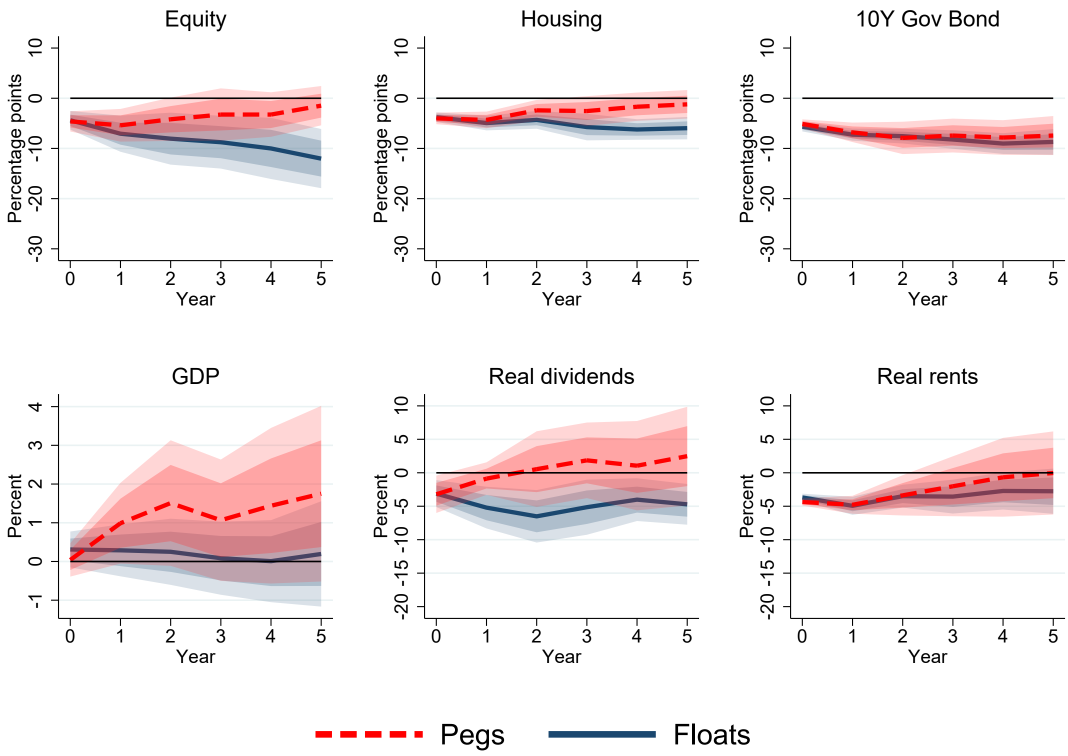

Notes: Impulse responses for real asset returns, real GDP, and real dividends to inflation surprises in open-peg countries (red dash line) and in float countries (blue solid line) for the period pre-1973. Model & OECD surprises, scaled to a 5 p.p. increase in inflation. Sample of 18 countries from the JST, 1870-1972, yearly frequency. The two lines represent the estimate of $\beta_{A,h}$ and $\beta_{B,h}$ of the state-dependent LP model. Shaded areas are the 90% and 68% confidence intervals.

1.F.3. The (non-)influence of Gold-standard

Figure F.3.: Asset returns responses with or without the Gold Standard

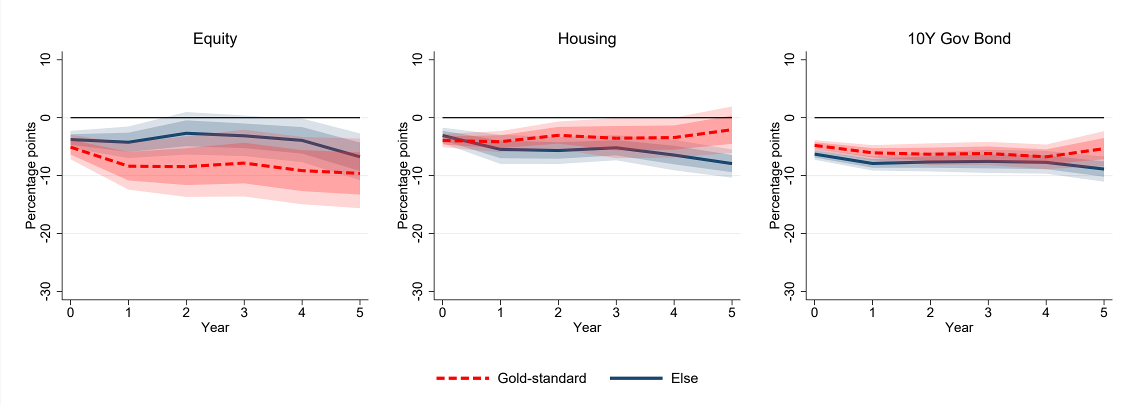

Notes: Impulse responses of real asset returns to inflation surprises in Gold-standard countries (red dash line) and in non-Gold standard countries (blue solid line) during the pre-1973 period. Model & OECD surprises, scaled to a 5 p.p. increase in inflation. Sample of 18 countries from the JST, 1870-1972, yearly frequency. The two lines represent the estimate of $\beta_{A,h}$ and $\beta_{B,h}$ using the state-dependent LP model. Shaded areas are the 90% and 68% confidence intervals.

1.G. Robustness exercises: Demand- vs. supply-driven inflation surprises

In this appendix we provide the following robustness exercises for Section 1.4.2. In Appendix Section 1.G.1 we exclude the historical period and focus only on post-1965 sample. In Appendix Section 1.G.2 we modify the dummy specification of Equation 1.7 and introduce a continuous interaction between inflation and GDP surprises. Finally, in Appendix Section 1.G.3, we combine the demand-supply exercise of Section 1.4.2 with the open peg and float exercise of Section 1.4.1.

1.G.1. Demand- and supply-driven surprises: Post-1965 sample

Figure G.1.: Asset returns responses to inflation surprises, 1965-2023: demand vs. supply

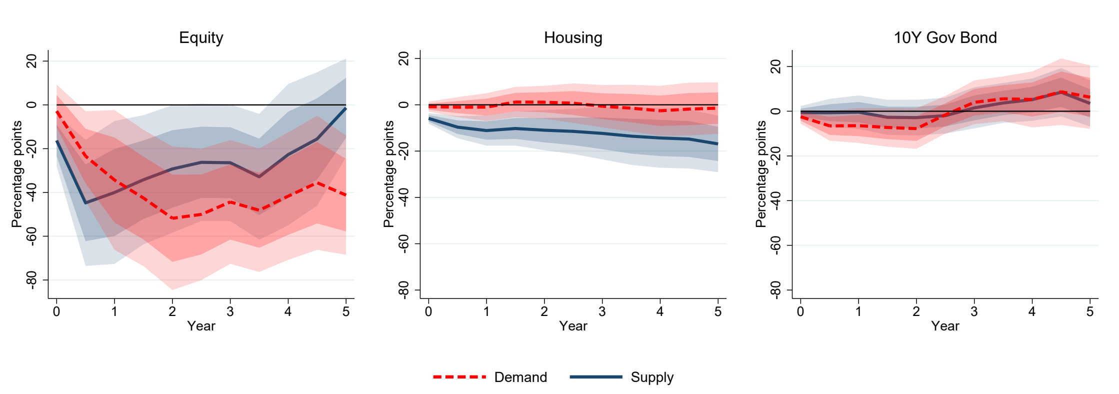

Notes: Impulse responses of cumulative real asset returns to demand-driven vs. supply-driven inflation surprises (red dashed lines and blue solid lines, respectively). Inflation surprise is scaled to increase inflation by 5 percentage points. Bands are at the 68% and 90% confidence level, based on Driscoll and Kraay (1998) standard errors. Half-yearly frequency, 1965-2020, using OECD inflation and GDP growth surprises.

Figure G.2.: Macroeconomic responses to inflation surprises, 1965-2023: demand vs. supply

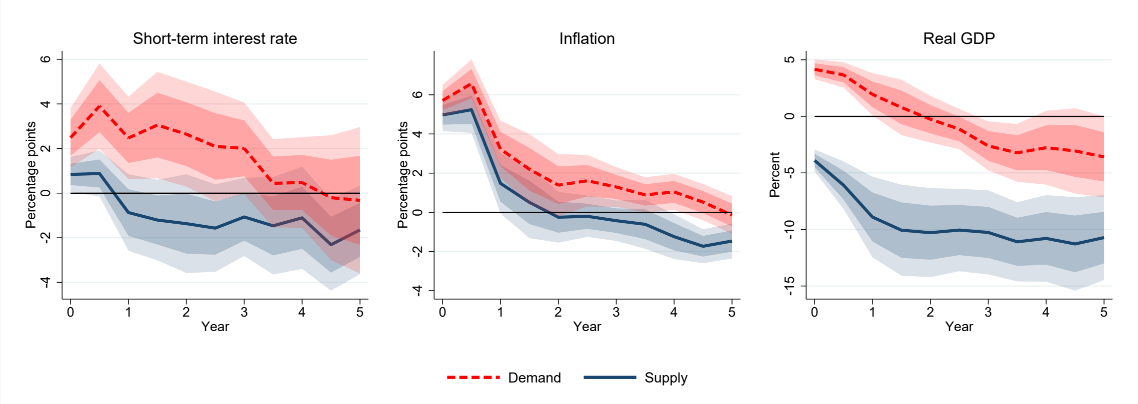

Notes: Impulse responses of nominal interest rate, inflation, and real GDP to demand-driven vs. supply-driven inflation surprises (red dashed lines and blue solid lines, respectively). Inflation surprise is scaled to increase inflation by 5 percentage points. Bands are at the 68% and 90% confidence level. Half-yearly frequency, 1965-2020, using OECD inflation and GDP growth surprises.

Figure G.3.: The additional effect of demand-driven inflation surprises on asset returns

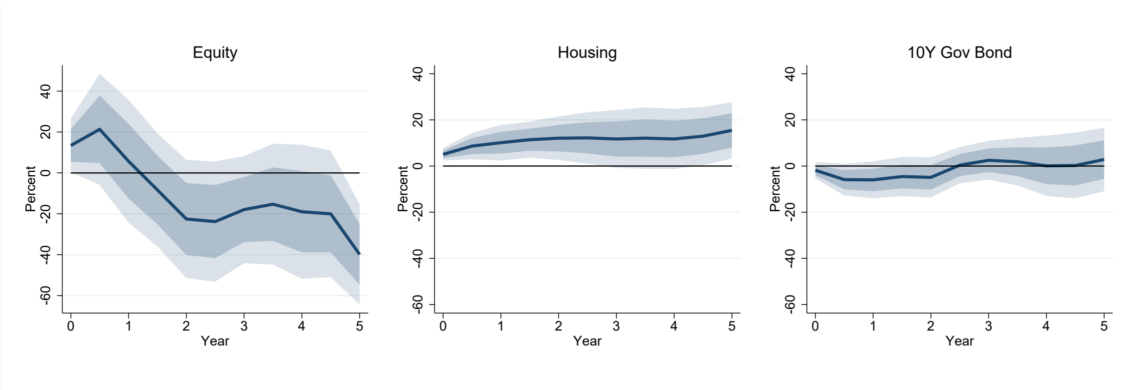

Notes: The blue line presents estimates of γ_h in Equation (1.7), i.e. the *additional* effect of a demand-driven inflation surprise compared to a supply-driven baseline case. Bands are at the 68% and 90% confidence level, based on Driscoll and Kraay (1998) standard errors. Half-yearly frequency, 1965-2020, using OECD inflation and GDP growth surprises.

1.G.2. Demand- and supply-driven surprises: continuous interaction

In this appendix we run a similar LP specification as in Equation (1.7), but now we directly exploit the joint correlation of the GDP growth and inflation surprises as a continuous interaction. The specification is as follows:

$$\begin{aligned} r^k_{i,t+h} - r^k_{i,t-1} = \beta^k_h \varepsilon^\pi_{it} + \gamma^k_h \quad \varepsilon^y_{it} \times \varepsilon^\pi_{it} \\ + \Psi^k_h(L)\mathbf{X}_{i,t-1} + \delta^k_{i,h} + v^k_{i,t+h} \quad \text{for} \quad h = 0, 1, 2, \ldots \end{aligned} \tag{1.10}$$

We normalize ε^π to induce a 5 p.p. increase in inflation and ε^y to induce a 1 p.p. increase in GDP growth. In this setting, β_h represent the average response to an inflation surprise, $\beta_h + \gamma_h$ the response to a demand-driven surprise, and $\beta_h - \gamma_h$ the response to a supply-driven surprise. The rest of the specification is unchanged with respect to Equation 1.7.

Figure G.4.: Asset returns responses to inflation surprises using continuous interaction

Notes: Impulse responses of cumulative real asset returns on inflation surprises due to demand or supply factors. Inflation surprise is scaled to increase inflation by 5 percentage points. The blue line presents estimates of $\beta_h - \gamma_h$ in Equation 1.7, while the red dashed line presents estimates of $\beta_h + \gamma_h$. Bands are at the 68% and 90% confidence level, based on Driscoll and Kraay (1998) standard errors. Half-yearly frequency, 1965-2020, using OECD inflation and GDP growth surprises.

1.G.3. Demand- vs. supply-driven inflation surprises in peg and float countries

To jointly test for the source of inflation surprises and isolate the response of monetary policy, we apply the following LP model:

$$\begin{aligned} r^k_{i,t+h} - r^k_{i,t-1} = \quad & \mathcal{I}_{i,t-1}\Big\{\beta^k_{A,h}\varepsilon^\pi_{it} + \gamma^k_{A,h} \times \mathcal{J}\{\varepsilon^y_{it} \times \varepsilon^\pi_{it} > 0\} \times \varepsilon^\pi_{it} + \Psi^k_h(L)\mathbf{X}_{i,t-1}\Big\} \\ & + (1-\mathcal{I}_{i,t-1})\Big\{\beta^k_{B,h}\varepsilon^\pi_{it} + \gamma^k_{B,h} \times \mathcal{J}\{\varepsilon^y_{it} \times \varepsilon^\pi_{it} > 0\} \times \varepsilon^\pi_{it} + \Psi^k_h(L)\mathbf{X}_{i,t-1}\Big\} \\ & + \delta^k_{i,h} + v^k_{i,t+h} \quad \text{for} \quad h = 0, 1, 2, \ldots, \end{aligned} \tag{1.11}$$

where $\mathcal{I}$ is the open peg dummy as in Section 1.4.1 while $\mathcal{J}\{\varepsilon^y_{it} \times \varepsilon^\pi_{it} > 0\}$ is the dummy for the additional effect of demand-driven inflation surprises. All in all, $\beta_{A,h}$ will show the response of asset returns in an open peg country to a supply-driven inflation surprise, while $\beta_{A,h} + \gamma_{A,h}$ will show the response of asset returns in open peg countries to a demand-driven inflation surprise. On the other hand, $\beta_{B,h}$ will show the response of asset returns in an float country to a supply-driven inflation surprise, while $\beta_{B,h} + \gamma_{B,h}$ will show the response of asset returns in float countries to a demand-driven inflation surprise.

Figure G.5.: Asset returns responses in *open pegs* and float: demand vs. supply

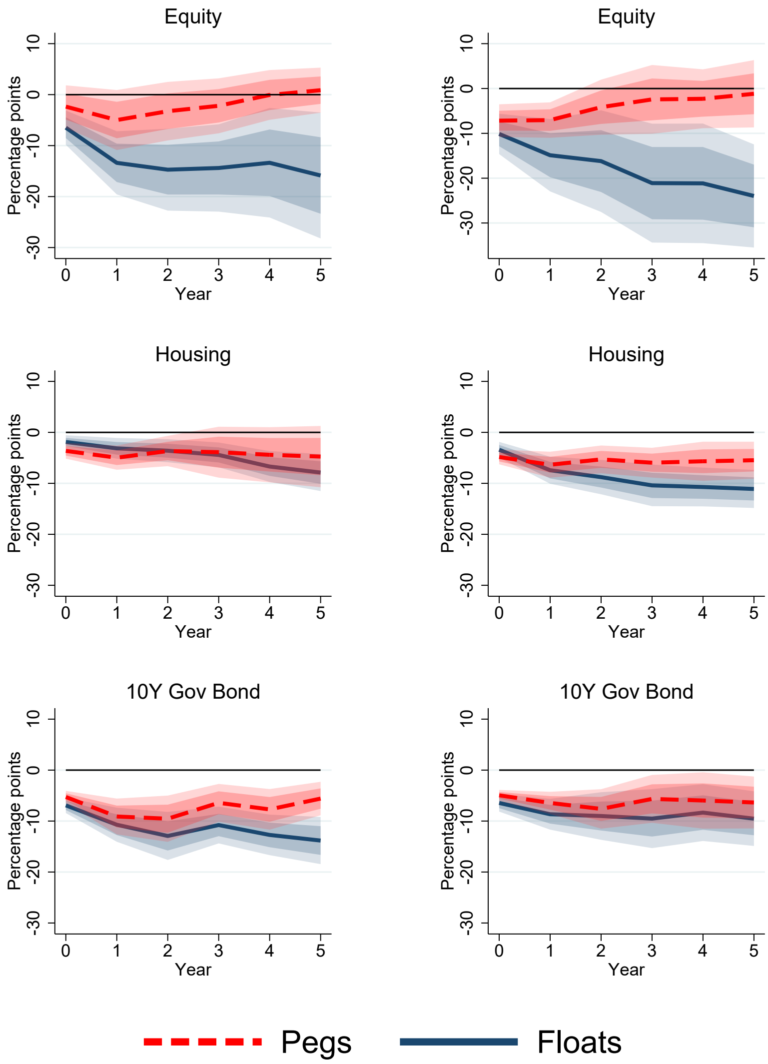

Notes: Impulse responses of cumulative real asset returns to inflation surprises in *open pegs* (red dashed line) and float countries (blue solid line) due to demand (left column) or supply factors (right column). Inflation surprise is scaled to induce an increase inflation by 5 percentage points. Bands are at the 68% and 90% confidence level, based on Driscoll and Kraay (1998) standard errors. Yearly frequency, 1870-2023, using Model and OECD inflation and GDP growth surprises and open capital accounts measure by Quinn, Schindler, and Toyoda (2011).

Figure G.6.: Macroeconomic responses to inflation surprises in *open pegs*: demand vs. supply

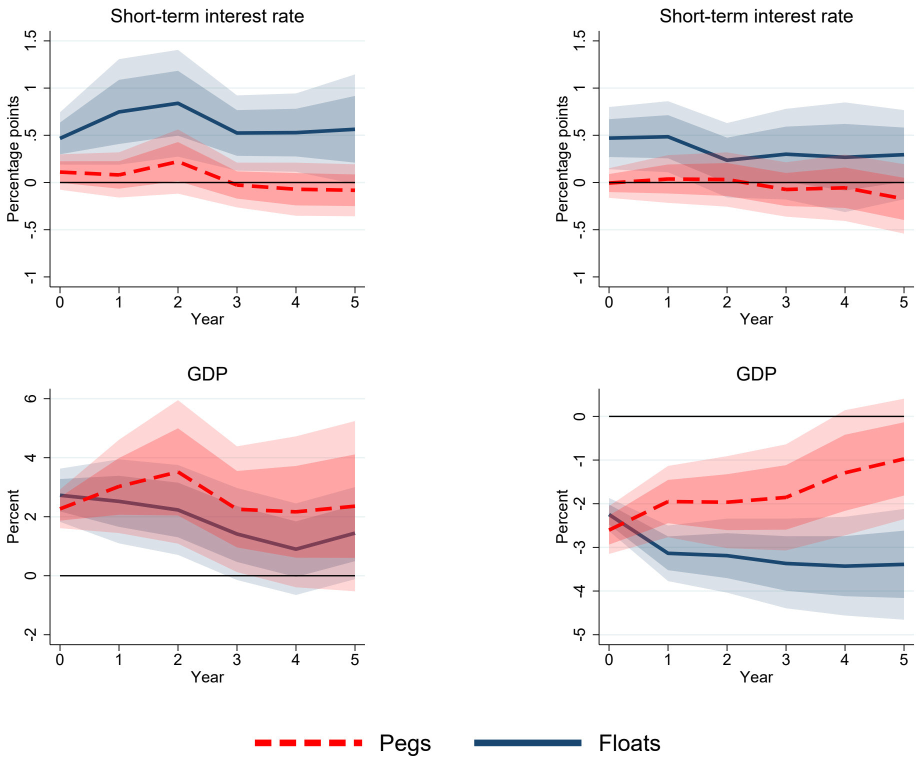

Notes: Impulse responses of short-term interest rates and real GDP to inflation surprises in *open pegs* (red dashed line) and float countries (blue solid line) due to demand (left column) or supply factors (right column). Inflation surprise is scaled to induce an increase inflation by 5 percentage points. Bands are at the 68% and 90% confidence level, based on Driscoll and Kraay (1998) standard errors. Yearly frequency, 1870-2023, using Model and OECD inflation and GDP growth surprises and open capital accounts measure by Quinn, Schindler, and Toyoda (2011).

Figure G.7.: Price and cash flow responses to inflation in *open pegs*: demand vs. supply

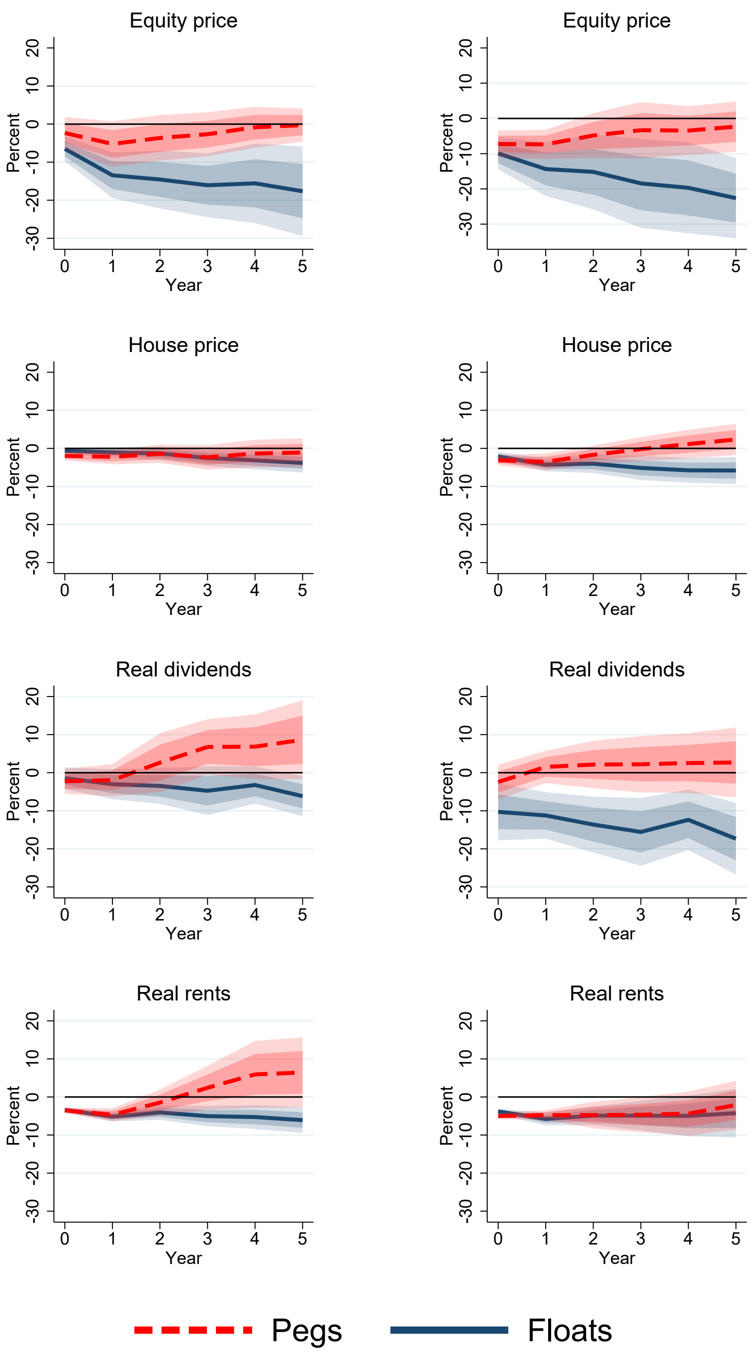

Notes: Impulse responses of real equity prices, real house prices, real dividends and real rents to inflation surprises in *open pegs* (red dashed line) and float countries (blue solid line) due to demand (left column) or supply factors (right column). Inflation surprise is scaled to induce an increase inflation by 5 percentage points. Bands are at the 68% and 90% confidence level, based on Driscoll and Kraay (1998) standard errors. Yearly frequency, 1870-2023, using Model and OECD inflation and GDP growth surprises and open capital accounts measure by Quinn, Schindler, and Toyoda (2011).

2. The Credit Channel of Inflation

2.1. Introduction

A century ago, Keynes (1923) and Fisher (1933) argued that a burst of inflation shifts wealth from creditors to debtors, easing the real burden of debt and thereby stimulating spending.[1] That intuition has returned to center stage after the sharp inflation spike of 2021. Standard models, and recent empirical work, embrace the idea that nominal frictions on debt contracts can have expansionary effects: when prices rise unexpectedly, the real value of nominal liabilities falls, giving leveraged firms a windfall that can boost real activity (Gomes, Jermann, and Schmid, 2016; Brunnermeier et al., 2025). This so-called *debt-inflation* channel, however, focuses on settings where firms are directly linked to debt markets, such as bond financing, abstracting from the balance-sheet frictions that may arise within banks. In reality, a large share of corporate debt is bank-intermediated,[2] and banks are naturally exposed to inflation via a maturity mismatch on their balance sheet. Banks "lend long and borrow short": their assets are dominated by long-dated, fixed-rate loans, whereas their liabilities are rather short-lived and reprice more quickly. An unexpected rise in the price level therefore erodes the real value of the asset side while leaving the liability side relatively less affected, possibly compressing their net worth and their capacity to supply new credit. Understanding this bank-side mechanism is key to assessing the overall impact of inflation on credit and real activity.

This paper shows that inflation can contract credit supply by weakening bank balance sheets, thereby constraining firm investment and economic activity. Using novel long-run aggregate data, I show that a one-percentage-point inflation surprise reduces bank capital and real loans in the economy by roughly three percent within five years. This result holds even in the absence of a typical contractionary monetary policy response to inflation, showing that inflation exerts real effects beyond those transmitted through monetary policy. At the micro level, using rich administrative bank–firm data, I find that inflation hits maturity-exposed banks, which cut credit supply to firms that subsequently reduce investment. Taken together, the evidence establishes a *credit channel of inflation*: by weakening bank balance sheets, inflation tightens credit and drags on real activity, offsetting – and possibly outweighing – the classical debt-inflation gains enjoyed by borrowers.

[1] The original intuition was formulated in the symmetric case due to the deflationary episode of the Great Depression, i.e., a *debt-deflation* that increases the real debt burden for debtors.

[2] In developed economies, the majority of corporate debt, including drawn credit lines, is bank-intermediated: around 80% in Europe (Holm-Hadulla et al., 2022) and 50% in the US (Colla, Ippolito, and Li, 2020)

The analysis proceeds in three main steps that build on two complementary data sets and a model. First, in Section 2.2, I assemble a new dataset of inflation surprises for 18 advanced economies over the period 1946–2020 from my earlier work (Kim, Ranaldi, and Schularick, 2025), constructed from digitized inflation forecasts published in the OECD's long-standing Economic Outlooks. I combine these surprises with macro-historical data from various sources (e.g., Jordà, Richter, et al., 2021; Baron, Verner, and Xiong, 2021). The dataset includes banking aggregates (loans, capital, deposits) and bank equity returns, together with key macroeconomic variables. Using local projections, I show that a one-percentage-point inflation surprise trims bank assets and loans by 3 percent after five years, a reduction comparable to roughly 60 percent of the contraction observed in a typical recession.[3] As a consequence, bank equity returns fall persistently and significantly more than the overall stock market.

These effects, however, may reflect the contractionary response of monetary policy rather than inflation itself. Central banks adjust policy rates to contain inflation, which indirectly curtails bank lending (e.g., Bernanke and Blinder, 1992; Jiménez, Ongena, et al., 2012). To isolate the effect of inflation without the accompanying monetary tightening, I exploit the long history of exchange-rate regimes in the sample. Due to the trilemma of international finance (Obstfeld, 1997; Obstfeld, Shambaugh, and Taylor, 2005), countries that maintain a hard currency peg while allowing free capital flows forgo monetary independence and cannot raise policy rates in response to domestic inflation.[4] Comparing these episodes to those under floating exchange rates allows me to estimate the effect of inflation surprises on bank aggregates absent the confounding influence of monetary policy. This *direct* effect explains nearly all of the credit response on impact and about half of the effect after three years, underscoring the importance of nominal frictions beyond the conventional monetary channel.

Time-series aggregates, however, cannot separate falling credit supply from contemporaneous dynamics of credit demand, nor can they track the direct consequences for firms. To fill that gap, Section 2.3 turns to administrative data from Portugal, which is particularly well suited for the empirical design for three reasons. First, Portugal is a small open economy – about two percent of euro-area GDP – and its business cycle correlates only weakly with that of the euro area (Oman, 2019). As a result, ECB policy is largely unresponsive to Portuguese-specific conditions, making the country an ideal setting to study the effects of inflation while holding monetary policy constant.[5] Second, the

[3]The consequences of inflation may depend on its underlying source. These baseline estimates should therefore be interpreted as *average* responses across episodes. In Kim, Ranaldi, and Schularick (2025), we show that demand- and supply-driven inflation surprises occurred with roughly equal frequency in the post-war period.

[4]I test this hypothesis and confirm that *open-peg* countries do not increase interest rates after an inflation surprise, unlike *float* countries.

[5]To strengthen this assumption, I extract the component of Portuguese inflation orthogonal to euro-area inflation and use it as a plausibly exogenous regressor with respect to euro-area monetary policy. In the aggregate data, I confirm that this *idiosyncratic* inflation component does not respond to euro-area monetary policy shocks and does not trigger ECB rate adjustments.

detailed supervisory balance sheet data allow me to compute a measure of each bank's *maturity gap* (English, Van den Heuvel, and Zakrajšek, 2018). This measure captures banks' heterogeneous exposure to inflation through the maturity mismatch between their long-term assets and short-term liabilities. Third, firms in Portugal typically borrow from multiple banks at the same time, a feature that enables a powerful identification design (Khwaja and Mian, 2008). By exploiting within-firm variation across banks, I can include firm–time fixed effects to control for credit *demand* and isolate changes in credit *supply*. Merging the credit registry and balance sheet datasets, I can observe the near universe of banks and firms in Portugal between 2011 and 2023, and I use this framework to estimate how inflation affects bank lending behavior through the maturity-mismatch channel.

Banks with greater maturity gaps respond more sharply to inflation. After a one-percentage-point increase in local inflation, a bank with a one-standard-deviation larger maturity gap cuts lending by roughly two percent. The contraction is stronger for highly leveraged banks and for those banks that rely less on deposits as funding source, consistent with a mechanism where tighter balance-sheet constraints and less stable funding amplify the effect.[6] Most of the adjustment occurs along the intensive margin: exposed banks scale back the size of existing credit relationships rather than terminating them, although the probability of ending a relationship increases slightly. Finally, I exploit the recent 2021 inflation surge – where inflation was mostly unexpected, together with no imminent increase in policy rates – to run an event study. I show that maturity-exposed banks reduced credit supply by 5 percentage points within one year, even before the ECB was expected to raise interest rates. Overall, this evidence highlights a contractionary credit channel of inflation operating through banks' maturity mismatches and balance-sheet constraints.

I then study whether the reduction in corporate credit by banks has real effects on investment for its corporate borrowers. I compare firms borrowing from more maturity-exposed banks to firms borrowing from non-exposed banks.[7] I only compare firms in the same municipality×sector×time cell. They are therefore subject to the same common shocks, but differ in their exposure to inflation because they borrow from different sets of banks.[8] The identifying assumption is that, conditional on controls and within the same municipality×sector×time, there are no shocks to firm fundamentals correlated with bank affiliation. I perform several checks and find support for this assumption.

At the firm level, inflation surprises weakly reduce total credit available to the firm. This masks sharp heterogeneity across borrowers: smaller and younger firms – and, generally, more bank-dependent firms – experience significant contractions in borrowing, while leveraged firms and those reliant on short-term credit face more favorable conditions,

[6] This observation is also in line with the argument that the deposit franchise value (partially) insulates banks' cash flows (Drechsler, Savov, and Schnabl, 2021). Banks with lower deposit reliance and more wholesale funding policies cut credit supply more significantly.

[7] More precisely, I define firm-level exposure to the *credit-inflation* channel as the credit-share weighted average of its banks' maturity gap.

[8] I also control for an estimate of firm-level demand shocks obtained from the loan-level regression (Cingano, Manaresi, and Sette, 2016; Jiménez, Mian, et al., 2020; Pinardon-Touati, 2024).

consistent with a debt–inflation mechanism and with banks' incentives to shorten loan maturity (Agarwal and Baron, 2024). These patterns show that firms differ in their ability to substitute credit across relationships and that the reduced credit provision by banks cannot be entirely substituted away, hence firms can face credit rationing after an inflation spike. As a consequence, firms scale back investment by 0.24% of assets within a year. The magnitude is substantial, as inflation reduces the investment rate through the credit channel by $\approx 3\%$, comparable to a 50 basis point monetary policy shock (Ottonello and Winberry, 2020).

To rationalize these findings, I develop a stylized model in which inflation surprises generate opposing forces in the credit market. The firm block builds on Kang and Pflueger (2015): here, inflation surprises lower the real burden of firms' nominal long-term debt, relaxing their balance sheets and raising credit demand. At the same time, banks – financing long-term fixed rate loans with short-term liabilities – face tighter capital constraints after an inflation surprise, because they need to rollover their short-term debt at higher price, hence reducing credit supply. The model underscores the importance of controlling for credit demand in the empirical analysis of Sections 2.3 and 2.4, and formalizes the intuition that, even with flexible prices, inflation can affect credit supply and real activity when banks are exposed to a maturity mismatch in their balance sheets.

Taken together, these patterns trace a clear chain from inflation, through bank exposure, to tighter credit and weaker real activity that counterbalance the classic debt-inflation channel. Although some firms – i.e., the more leveraged ones – may initially benefit from lower real debt burdens, the generalized outcome may revolve around the deteriorating effects of credit rationing. As inflation erodes banks' balance sheets, their capacity to lend across the system diminishes, which in turn raises lending constraints and curbs corporate investment. While offering short-term gains to debt-heavy firms, inflation can ultimately suppress aggregate activity when financial intermediaries face balance sheet pressures. The *credit channel of inflation* highlights the central role of financial intermediary health in determining the macroeconomic consequences of inflation.

Other related literature The paper bridges the literature on inflation with macro-financial research on how banks and firms respond to macroeconomic shocks.

First, this study connects to the debt-inflation literature, whose roots trace back to Fisher (1933), for the deflationary episode during the Great Depression. This strand of work typically concludes that moderate inflation benefits highly leveraged firms by eroding the real value of their nominal liabilities and freeing capacity for new investment. Gomes, Jermann, and Schmid (2016) show theoretically that the nominal rigidity embedded in long-term debt can transmit real effects through the debt-inflation channel. Kang and Pflueger (2015) find that inflation risk is priced in corporate bond spreads. Brunnermeier et al. (2025) show empirically that highly leveraged firms benefited from the German hyperinflation. Finally, D'Andrea et al. (2025) provide empirical evidence that highly leveraged firms enjoyed higher stock returns than non-leveraged firms during the 2021 inflation surge. Both these studies focus on large, publicly listed firms, whereas I observe the full distribution of firms in the economy, including many small firms that are naturally more bank-dependent. My paper complements these findings by showing that inflation simultaneously weakens banks' balance sheets and contracts credit supply, offsetting the gains for firms – especially the ones with lower leverage and with high bank dependence.

Second, the paper is related to the recent literature that links banks and inflation.[9] Drechsler, Savov, and Schnabl (2022) link credit crunches to the onset of the Great Inflation. Banks can be naturally exposed to inflation through their maturity mismatch or, as shown by Bahaj et al. (2023), by selling inflation protection in the swaps market. Agarwal and Baron (2024) shows that inflation depressed bank balance sheets during the 1977 U.S. inflationary spike. Corhay and Tong (2025) develop a New-Keynesian model with banks that supports a credit channel also in sticky-price environments. These papers do not empirically distinguish between credit supply and demand and capture responses that can also reflect monetary policy effects.[10] In contrast, I provide empirical evidence that inflation itself has a significant negative effect on banks' credit supply, holding credit demand constant and beyond the contractionary effects of monetary policy.

Finally, the paper relates to the large literature on the economic costs of inflation. Early work (Fischer, Hall, and Taylor, 1981; Cooley and Hansen, 1989; Lucas, 1994) quantified the welfare loss from reduced money holdings, generally finding modest costs at moderate inflation rates. Subsequent research highlighted broader channels: inflation raises uncertainty and relative price variability (Ball, Cecchetti, and Gordon, 1990), distorts saving and investment decisions through unindexed tax systems (Feldstein, 1999), and generates misallocation under nominal rigidities (Dotsey and Ireland, 1996). More recent work (Nakamura et al., 2018) shows that the welfare costs of inflation from price dispersion may be lower than previously thought. My paper contributes to this literature by providing direct evidence that inflation reduces bank credit supply and firm investment, highlighting a concrete channel, beyond price distortions, through which inflation negatively influences real activity.

The remainder of the paper is organised as follows. Section 2.2 details the macro-historical data and presents the local projection evidence on aggregate credit. Section 2.3 describes the bank-firm data and shows the effect of inflation on banks' credit supply, while Section 2.4 tracks the consequences on borrower firms. Section 2.5 introduces a simple model to rationalize these findings. Section 2.6 concludes.

---

[9] There is also a literature on inflation and the broader financial sector: Choi, Smith, and Boyd (1996) and Huybens and Smith (1999) argue theoretically that sustained inflation introduces credit market frictions that impair financial intermediation and slow capital formation. Boyd, Levine, and Smith (2001) find that higher long-run inflation is associated with less active banking sectors and stock markets, while English (1999) reports a positive cross-sectional link between inflation and financial sector size.

[10] A recent exception is Bergant et al. (2025), who provide suggestive evidence that bank *profitability* is on average insulated from inflation, though with significant heterogeneity. In contrast, I focus on bank *balance sheet position* and *credit supply*, and show that banks are significantly affected in their value.

2.2. Inflation and credit in the long run

In this section, I provide long-run evidence that inflation surprises are followed by a contraction of aggregate credit, bank capital, and bank stock returns. The effect is economically large and not only the consequence of tightening monetary policy. These facts highlight that inflation can directly affect banks' balance sheets, motivating the micro evidence that disentangle credit demand from credit supply.

2.2.1. Data

I combine aggregate data from post-WWII and across 18 OECD countries on bank balance sheets and macro-financial aggregates with a novel measure of inflation surprises. This unique coverage allows me to provide the first long-run evidence on how unexpected inflation affects bank balance sheets and credit provision.

The baseline period is between 1946 and 2020. In Section 2.2.3 I will extend the available data series to 1870 to exploit variation in monetary policy regimes.[11] The complete set of countries includes: Australia, Belgium, Canada, Switzerland, Spain, Finland, France, Denmark, Germany, Ireland, Italy, Japan, Netherlands, Norway, Portugal, Sweden, the United Kingdom, and the United States.

Bank aggregates and macroeconomic controls Data for aggregate loans, deposits, capital and total assets come from Jordà, Richter, et al. (2021), available from 1870 to 2020. These are book values from banking sector balance sheets, where bank capital corresponds to the Basel III definition of Common Equity Tier 1 capital,[12] loans includes mortgages and business loans, while deposits correspond to both term and sight deposits, and both checking and savings accounts. I deflate each series with the national CPI and expressed in logarithms, and merge them with aggregate bank equity total-return indices, available since 1870 and provided by Baron, Verner, and Xiong (2021).
Finally, I complement the dataset with macroeconomic aggregates from the Macrohistory Database (Jordà, Schularick, and Taylor, 2017; Jordà, Knoll, et al., 2019), including GDP, short-term interest rates, CPI, and the aggregate stock market return.

Inflation surprises The core regressor is the one-year-ahead inflation surprise:

$$\varepsilon^{\pi}_{i,t} \;=\; \pi_{i,t} - \mathbb{E}_{t-1}[\pi_{i,t}]$$

[11]The baseline sample (1946–2020) is chosen to ensure a relatively homogeneous monetary and financial environment across countries, avoiding regime breaks associated with the interwar period. In Section 2.2.3, extending the sample back to 1870 yields similar qualitative results. In that case, years around the world wars (1914–19 and 1939–46) are always excluded.

[12]This normally includes common stock (paid-up capital), reserves, and retained earnings.

where $\pi_{i,t}$ is CPI inflation realized in year t in country i and $\mathbb{E}_{t-1}[\pi_{i,t}]$ is the forecast for inflation made at the end of year $t-1$. I construct the inflation surprises in my earlier work (Kim, Ranaldi, and Schularick, 2025, KRS henceforth), where we digitize semi-annual forecasts from the OECD Economic Outlook for these 18 countries, starting in 1965. All forecasts are prepared using a unified methodology so that series are comparable across countries. Furthermore, the OECD leverages its local presence to engage with government experts and policymakers, resulting in forecasts that incorporate significant local knowledge and insights into future policy changes.[13] This measure provides a novel source for studying inflation dynamics in the long run for various developed economies. As commonly known, inflation expectations data begin to be available only from the late 20th century.[14] Previous literature estimated inflation expectations for historical samples using models or financial markets data for limited time- or country-samples (Hamilton, 1992; Cieslak and Povala, 2015; Ellison, Lee, and O'Rourke, 2024).

Figure 2.1.: Inflation surprises: 1946-2023

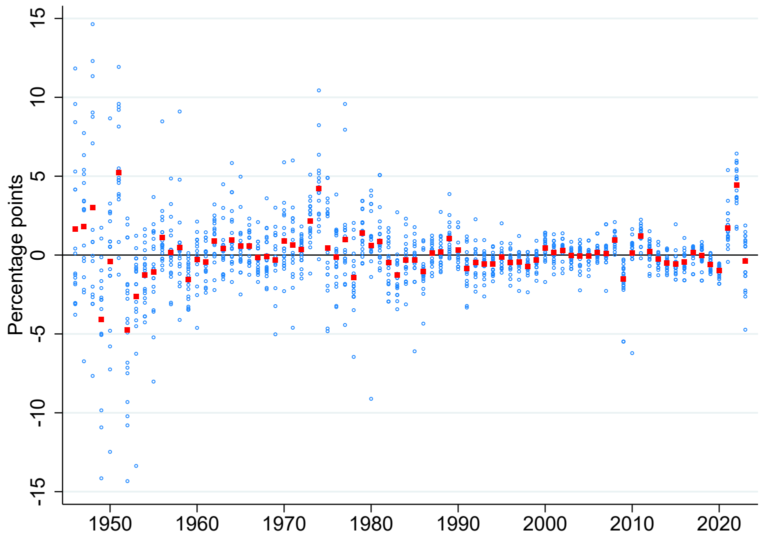

Notes: Inflation surprises from Kim, Ranaldi, and Schularick (2025), from 1-year-ahead forecasts, annual frequency, 1946-2020. Inflation surprises are computed as the difference between realized inflation and expectations from the end of the previous period, i.e. $\varepsilon^{\pi}_{it} = \pi_{it} - \mathbb{E}_{t-1y}[\pi_{it}]$. Inflation is measured in %, surprises in percentage points. Red dots represent the average surprise across the 18 OECD countries.

Figure 2.1 presents the inflation surprises for all countries (blue dots) in the period 1946-2023, together with the average across the 18 countries (red dots). In summary, over the past 80 years, only the Great Inflation episode in the 1970s was comparable to the recent inflation spike in terms of significant inflation surprises, while also the post-war decade shows large volatility. During the Great Moderation, inflation surprises have been

[13]Reviews of the OECD's forecasts by Vogel (2007) report that they are comparable to private sector forecasts and have a number of desirable properties.

[14]The earliest comes from the Livingston survey that provides inflation forecast data for the USA from 1951 onward. For other countries, similar data start only in the 1990s (e.g., Consensus economics, central bank forecasts) or they are only in a qualitative format (e.g., EU Business and Consumer survey).

relatively minor and more convergent across different countries. This is consistent with Blanco, Ottonello, and Ranošová (2025), who also observe that large surges in inflation or deflation are initially unexpected.

2.2.2. Benchmark results: inflation hurts the aggregate banking sector

I use this long-run dataset to estimate impulse response of aggregate bank balance sheet variables on inflation surprises via local projections. The results provide evidence that inflation significantly shrinks bank balance sheets and, therefore, can hurt their ability to supply credit. These findings motivate Section 2.3, where I turn to bank-firm microdata to examine the underlying mechanism.

Empirical framework To trace the dynamic response of aggregate bank balance sheets to inflation surprises, I estimate panel local projections à la Jordà (2005):

$$\Delta^h y_{i,t+h} = \beta_h \, \varepsilon^{\pi}_{i,t} + \Psi_h(L)\,\mathbf{X}_{i,t-1} + \delta_{i,h} + \delta_{t,h} + u_{i,t+h}, \qquad h = 0, 1, 2, \ldots \tag{2.1}$$

where $\Delta^h y_{i,t+h} \equiv y_{i,t+h} - y_{i,t-1}$ is the h-year log-difference in the outcome variable y in country i.[15] $\varepsilon^{\pi}_{i,t}$ is the inflation surprise defined in Section 2.2.1. $\mathbf{X}_{i,t-1}$ collects two lags of $\varepsilon^{\pi}_{i,t}$ and changes in the dependent variable, short-term interest rate, real GDP, the stock market index, and inflation. $\delta_{i,h}$ and $\delta_{t,h}$ denote horizon-specific country and year fixed effects and control for time-invariant country characteristics and common shocks, respectively. Lagged shocks and controls are included to account for any serial correlation and to improve inference, as prescribed by Montiel Olea and Plagborg-Møller (2021). I include the inflation level because both inflation and its surprises are more volatile in high-inflation states. The stock-market index proxies for high-frequency information that may not yet be reflected in inflation forecasts. Other macroeconomic variables are included to control for business cycle conditions. I use in every specification the Driscoll and Kraay (1998) heteroskedasticity- and autocorrelation-consistent standard errors. Equation (2.1) yields a sequence $\{\beta_h\}_{h=0}^{H}$ that maps a one-percentage-point inflation surprise into percentage responses of each balance-sheet item.

Results Figure 2.2 plots the impulse responses of real loans, deposits, bank capital and *cumulative* equity returns to a 1 p.p. inflation surprise. Unexpected inflation cuts real loans by 1% on impact and by more than 2% after five years, in line with cross-sectional estimates in Agarwal and Baron (2024), while bank capital declines more steadily by up to 3% in five years. Bank equity returns fall sharply after an inflation surprise, with long-lasting effects that hurt banks more than the rest of the corporate non-financial sector

[15] In the case of equity returns, $y_{i,t}$ is the real total return index, so that $y_{i,t+h} - y_{i,t-1}$ can be interpreted as the h-period *cumulative* real return.

Figure 2.2.: Response of bank balance sheet items to inflation surprises, 1946-2020

Notes: Impulse responses for aggregate loans, deposits, bank capital and cumulative bank equity returns on inflation surprises, estimated with panel-LP as in eq. (2.1). 1946-2020, annual frequency, 18 countries. All variables are expressed in real terms. Driscoll and Kraay (1998) standard errors. Bands are at the 68% and 90% confidence level.

(see Appendix Figure A.2).[16] This fact is in line with the idea that banks are particularly exposed to inflation given the maturity structure of their balance sheet, with long-term nominal assets matched with short-term liabilities.

In the aggregate, I also find that real deposits decline after an inflation surprise. While deposits account for about 60% of total bank liabilities, the remainder consists mainly of wholesale funding (*non-core funding*) which is contractually and de facto short term. Indeed, banks' reliance on non-core funding increases after an inflation surprise (see Appendix Figure A.3), as wholesale funding is of shorter maturity and its value is therefore less impaired. The contraction in real deposits is therefore consistent with banks' overall balance-sheet compression.[17]

[16]There is a rich literature that finds extensive evidence that inflation hurts stock market returns, see e.g., Fama and Schwert (1977), Cieslak and Pflueger (2023), and Kim, Ranaldi, and Schularick (2025). The evidence in Appendix Figure A.2 shows that bank returns are significantly more affected than the average market. This is in line with results in Corhay and Tong (2025) who find the same results on disaggregated data and controlling for other firm-specific factors, e.g., leverage.

[17]There is a long debate around the nature of deposits (e.g., Drechsler, Savov, and Schnabl (2021), Greenwald, Schulhofer-Wohl, and Younger (2023), and DeMarzo, Krishnamurthy, and Nagel (2024)), either as short-term liabilities, per their contractual maturity, or as de facto longer-run liabilities, because

All in all, much of the effect on the aggregate bank balance sheet materializes not only on impact but also steadily over the five-year horizon. The contraction is not a price-level artifact but reflects a genuine reduction in aggregate credit, also in nominal terms (see Appendix Figure A.1). Overall, the whole aggregate bank balance sheet tends to shrink more than one-to-one after an inflation surprise. The magnitude of these responses (-3%) is economically large. For example, during the average recession in this sample, real loan growth slows by roughly 5% – a figure nearly matched by a two-percentage-point inflation surprise.

2.2.3. It's not all about monetary policy

An inflation surprise typically triggers central banks to raise interest rates, which in turn reduces bank lending (Bernanke and Blinder, 1992; Jiménez, Ongena, et al., 2012). This makes it difficult to disentangle the direct effect of inflation from the consequences of monetary policy. To address this empirical challenge, I exploit variation in exchange rate regimes to separate the two mechanisms. I show that much of the immediate contraction in bank balance sheets is due to inflation itself, confirming that inflation surprises can hurt banks even in the absence of contractionary monetary policy responses.

Exploiting historical data permits comparison of episodes in which central banks faced different, plausibly exogenous, degrees of policy autonomy. The key source of variation is the exchange rate regime. Governments adopt pegs to anchor prices or ease trade with a major partner, but doing so tightens their monetary policy margin. By the trilemma of international finance, a country cannot simultaneously enjoy a fixed exchange rate, open capital markets, and an independent monetary policy (Obstfeld, 1997; Obstfeld, Shambaugh, and Taylor, 2005).[18] Therefore, a peg with open capital accounts forces the central bank to align with the base country's rates and limits its ability to raise short-term rates after a *local* inflation surprise. On the other hand, countries with float exchange rate regimes may tighten after an inflation surprise as they see fit.[19] Therefore, I interpret the effect of an inflation surprise in an *open peg* country as the direct effect of inflation controlling for endogenous interest rates response. On the other hand, the response in the *float* country shows the total effect of inflation, including the consequences of tighter

deposit rates are sticky and depositors are slow to withdraw from low-yielding accounts. Drechsler, Savov, and Schnabl (2021) argue that this *deposit franchise* helps banks insulate their *cash flows* from interest rate risk. However, in line with DeMarzo, Krishnamurthy, and Nagel (2024), banks' *value* can still be affected by interest rates. My results extend this point to inflation: even if cash flows can be partially insulated (Bergant et al., 2025), the real value of bank balance sheets compresses after an inflation surprise.

[18] See e.g., Bazot, Monnet, and Morys (2024) for an empirical confirmation of the trilemma over history and across countries. Other empirical studies use this setting to control for the endogenous response of monetary policy, see e.g., Andersen et al. (2023). Also, I verify in Appendix Figure A.5 that *open peg* countries do not raise their short-term interest rate in response to an inflation surprise.

[19] The constraint is short-run: hard pegs still deliver low and stable inflation, see e.g., Ghosh, Qureshi, and Tsangarides (2014) and Bleaney and Francisco (2005) for evidence that hard pegs anchor inflation in the medium run.

monetary policy.

Empirical framework I modify equation (2.1) using the extended 1870–2020 sample separately for pegs and floats.[20] The estimating equation is now a state-dependent panel local projections framework in the spirit of Ramey and Zubairy (2018), where the states are (i) pegged countries with open capital accounts and (ii) float countries:

$$\begin{aligned}\Delta^h y_{i,t+h} = &\mathcal{I}_{i,t-1}\Big\{\beta_{A,h}\varepsilon^{\pi}_{it} + \Psi_{A,h}(L)\mathbf{X}_{i,t-1} + \delta_{A,i,h} + \delta_{A,t,h}\Big\} \\ &+(1-\mathcal{I}_{i,t-1})\Big\{\beta_{B,h}\varepsilon^{\pi}_{it} + \Psi_{B,h}(L)\mathbf{X}_{i,t-1} + \delta_{B,i,h} + \delta_{B,t,h}\Big\} + \epsilon_{i,t+h} \quad \text{for} \quad h = 0,1,2,\ldots\end{aligned} \tag{2.2}$$

where $\mathcal{I}_{i,t-1}$ is a dummy variable that indicates whether a country is pegged in t and $t-1$ and in both times it has open capital accounts. The capital-account openness variable $q_{i,t-1}$ is taken from Quinn, Schindler, and Toyoda (2011), where $q_{i,t-1} = 1$ when the country completely allows free capital movements.[21] I condition on being an *open peg* in both time periods t and $t-1$ to avoid cases where the decision to peg is endogenously determined by economic conditions.[22] The main distinction between the two states lies in the monetary policy conduct: whenever countries have their currency pegged while maintaining their capital accounts open, they are not able to freely react to an inflation surprise by raising interest rates. The rest of the specification and estimation details mirror those in equation (2.1).[23]

Results Figure 2.3 shows the responses for both states and compares the effect of an inflation surprise with and without a reaction of the central bank. In pegs (red line), a 1 p.p. inflation surprise cuts real loans and the other bank balance sheet items by 1 % on impact, but slowly converges back to steady state within 5 years. In floats (blue dashed line), the contraction is approximately the same on impact but losses are persistent within the five-year horizon.[24] A similar picture applies to banks' cumulative equity returns,

[20]The need for exploiting the full-sample in this exercise comes from the fact that the balance of pegs and floats in the post-1946 period is dramatically skewed towards float countries, especially after the collapse of the Bretton-Woods system in 1973.

[21]This index is available until 2012, at most. I merge this with the capital account openness index from Chinn and Ito (2006), up to 2020. Note that the two variables are very close in the recent overlapping period.

[22]As noted in Jordà, Schularick, and Taylor (2020), restrictions on capital mobility have not been used as a high-frequency policy tool by pegging economies since the index is very slow-moving in advanced economies.

[23]I also add the inflation rate in the *base* country as additional control in $\mathbf{X}_{i,t-1}$, to exclude the cases where the peg country raises interest rates to an inflation surprise in the base country. Results are robust to a variety of alternative specifications.

[24]The total response is more muted compared to Figure 2.2, which uses the post-1946 sample only. The reason lies in the different volatility of surprises in the pre-1946 sample, where surprises were historically stronger (SD: 5.1 in pre-1946 vs. 2.78 in post-1946). This fact mechanically leads to a milder overall response to an *average* 1 p.p. surprise. Rather than the overall magnitude, this section focuses on the

falling on impact with the same magnitude across pegs and float countries but quickly offsetting losses in the former, while in floats real returns continue to drop steadily.

Figure 2.3.: Direct effect of inflation surprises, 1870-2020

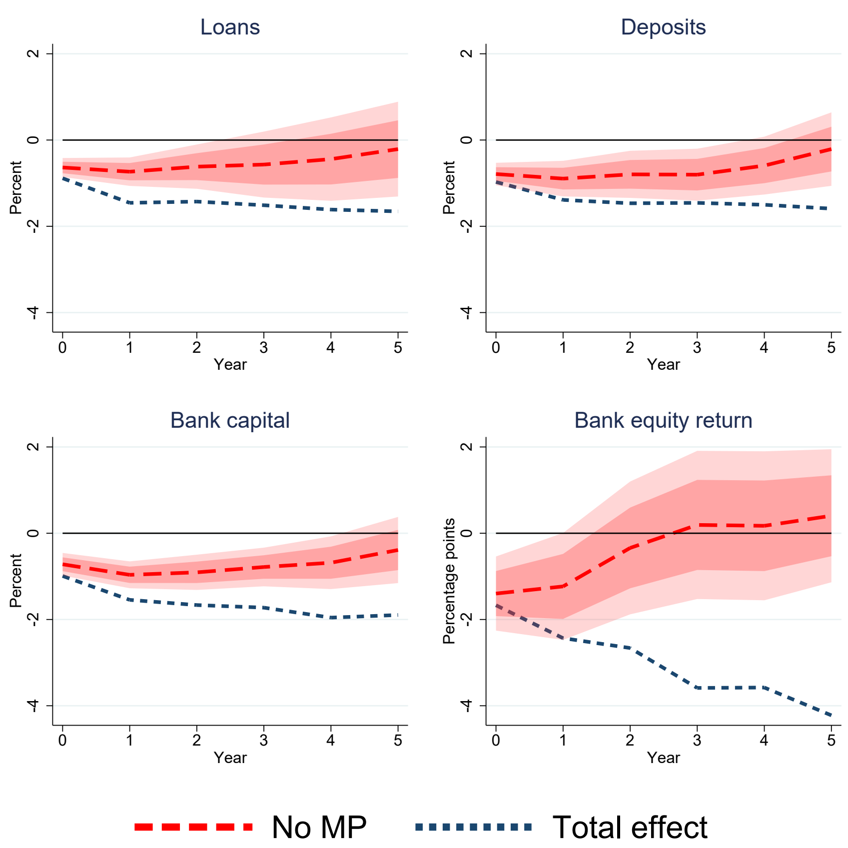

Notes: Impulse responses for aggregate loans, deposits, bank capital, cumulative bank equity returns on inflation surprises, estimated with panel-LP as in eq. (2.2). The red solid line represents the "direct" effect of inflation, i.e., the response of outcome y_t in open peg countries ($\beta_{A,h}$), where monetary policy cannot respond to inflation. The blue dashed line shows the "total" effect of inflation, i.e., the response of outcome y_t in float countries ($\beta_{B,h}$). 1870-2020, annual frequency, 18 countries. Driscoll and Kraay (1998) standard errors. Bands are at the 68% and 90% confidence level.

Overall, these dynamics suggest that inflation has a large direct impact on the banking sector even controlling for the indirect response of monetary policy. Roughly 90% of the impact response can be attributed to the direct effect of inflation. However, much of the persistence of these effects is due to the contractionary response of monetary policy: at $h = 3$, approximately 50% of the total response of bank balance sheet items is due to inflation directly; at $h = 5$, this direct effect of inflation has vanished, while the overall response remains significantly negative, drawn down by the contractionary rise in interest rates. This result is also in line with the literature on the long-run non-neutrality of monetary policy (Mankiw, 2001; Jordà, Singh, and Taylor, 2024).

fraction of the total effect explained by inflation per se, without contractionary policy responses.

Inflation can directly, substantially, and persistently affect banks. However, this section showed that failing to control for the countercyclical reaction of monetary policy can bias the results. When turning to the micro data, in Section 2.3, I exploit the case of Portugal as a laboratory to study a small country in a monetary union, to keep fixed the fluctuations in the monetary reaction.

Additional exercises In Appendix 2.C, I provide suggestive evidence that the debt–inflation channel can, in the aggregate, be offset by the credit channel of inflation. Specifically, I show that when corporate indebtedness is high, i.e. when the debt-inflation channel should have its higher positive effects on firms' real debt burden, inflation surprises trigger sharper declines in bank fundamentals, and are also followed by lower aggregate investment.

2.3. Micro-data: the effect of inflation on credit supply

In this section, I turn to micro data to establish whether the contraction of credit after inflation surprises reflects supply (from banks) rather than demand (from firms). To do so, I exploit administrative data on the universe of bank–firm credit relationships in Portugal, where monetary policy is fixed at the euro-area level. Banks that are more exposed to inflation through the maturity gap in their balance sheets cut lending more to the same firm. This provides evidence that the contraction documented in the aggregate data can operate through a bank credit-supply channel.

2.3.1. Institutional background

Portugal provides an ideal setting to study the credit channel of inflation, since local inflation surprises do not trigger monetary policy responses by the ECB. This ensures that the micro evidence identifies the *direct* effect of inflation on banks and credit supply, cleaned of any indirect outcome due to the response of the central bank.

Portugal accounts for only 2% of euro-area GDP, and its business cycle correlates weakly with the aggregate euro-area cycle (Giannone, Lenza, and Reichlin, 2008; Oman, 2019). Because the ECB targets euro-area inflation as a whole, Portuguese-specific inflation surprises do not elicit a country-specific monetary policy response.[25] This identifying strategy is similar to Jiménez, Ongena, et al. (2012), who apply it to Spain.[26] To further support this assumption, the analysis uses idiosyncratic Portuguese inflation, net of euro-area inflation, as the main regressor, which strengthens the credibility of treating monetary policy as exogenous in this environment.

[25] See Article 127(1) TFEU: "*implementing the monetary policy of the Union.*" The 2021 ECB strategy review makes this scope explicit: "The ECB's inflation target makes clear that the focus of the ECB's monetary policy is on the euro area as a whole. Therefore, the 2% inflation target is assessed on the basis of inflation developments in the euro-area economy."

[26] In Jiménez, Ongena, et al. (2012), Spain accounts for approximately 10% of euro-area GDP, with higher business-cycle convergence compared to Portugal.

Specifically, to capture the component of Portuguese inflation that is unrelated to area-wide monetary policy, I construct an orthogonalized series. Aggregate inflation across euro-area countries π_{it} is decomposed into a common component, π_{it}^{Common}, and a residual country-specific component, $\pi_{it}^{\perp}$:

$$\pi_{it} = \pi_{it}^{Common} + \pi_{it}^{\perp} = \Gamma_i F_t + \pi_{it}^{\perp} \tag{2.3}$$

where π_{it} is the inflation rate in country i at time t, F_t is the common factor, and Γ_i are country loadings.[27] The common component π_{it}^{Common} captures euro-area inflation that guides ECB policy, while the orthogonalized component $\pi_{it}^{\perp}$ reflects local country-specific inflation fluctuations that do not trigger monetary tightening. Appendix Figure A.6 plots the decomposition for Portugal.

I further test the credibility of this assumption with two separate exercises. First, I estimate impulse response functions of the orthogonalized inflation $\pi^{\perp}$ for Portugal on the Jarociński and Karadi (2020) euro-area monetary policy shocks. $\pi_{it}^{\perp}$ does not respond to monetary policy shocks within the first three years, unlike euro-area and overall Portuguese inflation, see Appendix Figure A.7.[28] Second, I estimate impulse response functions of the short-term interest on $\pi^{\perp}$ in a similar fashion and find that it does not trigger any significant movement within three years, see Appendix Figure A.8. This alleviates concerns that my measure might still reflect euro-area policy reactions, showing that it can be used in my main regressions as proxy for the *direct* effect of inflation.

These institutional features make Portugal a perfect laboratory to identify the direct effect of inflation on banks and credit supply. The next subsections introduce the data (Section 2.3.2), outline the empirical strategy (Section 2.3.3), and present the results (Section 2.3.4). Finally, Section 2.3.5 exploits the 2021 inflation surge as a clean inflationary episode for an event study.

2.3.2. Data

The analysis relies on exceptionally rich administrative data from Banco de Portugal, covering the entire universe of banks and firms in the country. This unique feature allows me to observe every bank–firm credit relationship together with detailed balance sheet information on both sides.

The study combines four datasets. The *Harmonized Central Credit Responsibility Database* (HCRC, 2009–2023, monthly) records all bank–firm credit exposures above €50,

[27] Specifically, I use year-over-year CPI inflation at monthly frequency for all euro-area countries, 1999m1-2024m12. F_t and Γ_i are estimated using the first principal component, which captures ≈ 90% of overall variation in π_{it}.

[28] Specifically, using the Jarociński and Karadi (2020) monetary policy shocks as instruments for changes in the three-month EONIA rate, I estimate a local projection instrumental variable framework (LP-IV) over the 1999m1–2024m12 period. The results in Appendix Figure A.7 show that both euro-area π_t^{Common} and Portuguese inflation π_t fall significantly after a contractionary monetary policy shock, as theory predicts, but the orthogonalized Portuguese inflation $\pi_t^{\perp}$ does not respond within the first three years.

including credit lines, overdue or written-off loans, and whether the exposures have short original or residual maturity (below one year). The *Monetary Financial Institutions Balance Sheet Database* (BBS, 1997–2023, monthly) reports the assets and liabilities of all monetary financial institutions at the item–counterparty level. Finally, the *Central Balance-Sheet Database* (CB, 2006–2022, annual) covers the universe of non-financial firms with mandatory accounting and tax information.[29]

I apply conservative sample restrictions to avoid outliers. Overall, I restrict the sample to start in 2011, as in 2010 the CB database undergoes a change in reporting standards for firms, reflecting a significant structural change. On the firm side, I exclude liquidated or inactive firms, financial and public administration firms, those with missing or zero assets, firms with no employees, firms located outside mainland Portugal, firms that are directly controlled by foreign groups, and firms without any bank relationship. On the bank side, I retain only monetary financial institutions present in the BBS and exclude credit exposures within two years of major corporate events such as mergers, resolutions, or sales.

The resulting dataset includes 385,255 unique firms and 72 banks, linked by 900,138 unique bank–firm relationships. I winsorize all variables in the sample at the 1st and 99th percentile. On average, firms borrow from 3 different banks (median = 2), and 60% maintain multiple bank relationships—a share substantially higher than in other European countries.[30] This dense network of multibank lending enables the use of firm–time fixed effects while maintaining a wide representativeness. Table 2.1 reports loan-level characteristics (Panel A) and firm-level characteristics (Panel B). Notice that the distributions of the variables at the loan level are similar between multibank relationships and the unconditional case. Moreover, the firm data covers a significant share of small firms, with the majority of companies having less than 10 employees. These facts confirm the broad representativeness of the data and capture the broad structure of the Portuguese corporate sector.

2.3.3. Empirical strategy: banks are exposed to inflation

The goal is to identify how different banks adjust lending to the same firm after an increase in inflation. The empirical strategy exploits multibank firms to hold credit demand fixed and to identify credit-supply effects of inflation through banks' balance-sheet exposures.

I estimate the following specification:

$$\begin{aligned} L_{bft} = & \beta \left[\pi_t^{\perp} \times \sigma_{b,t-1}\right] + \alpha_{ft} + \alpha_{bf} \\ & + \gamma_1' \mathbf{X}_{bf,t-1} \times \pi_t^{\perp} + \gamma_2' \mathbf{X}_{bf,t-1} + \gamma_3' \sigma_{b,t-1} + u_{bft} \end{aligned} \tag{2.4}$$

[29]Banco de Portugal Microdata Research Laboratory (BPLIM), 2023; Banco de Portugal Microdata Research Laboratory (BPLIM), 2024a; Banco de Portugal Microdata Research Laboratory (BPLIM), 2024b.

[30]In France, only about 30% of firms have multiple banks, while in Germany and Belgium the share is below 50%. Appendix Figure A.9 shows the distribution of bank relationships per firm in the sample.

Table 2.1.: Summary statistics

Panel A: Firm×bank-level variables

	All					Multibank				
	Mean	SD	p10	p50	p90	Mean	SD	p10	p50	p90
Credit growth $\Delta L_{bft}(\%)$	-2.18	18.20	-11.78	-1.62	1.49	-2.11	18.62	-12.15	-1.59	2.58
Credit outstanding L_{bft} (€)	395,544	4,051,223	7,506	50,560	568,991	473,177	4,630,339	8,446	62,340	711,212
Length of relationship$_{bft}$ (months)	40.69	33.92	5.00	31.00	93.00	41.40	34.24	6.00	32.00	95.00
Maturity gap σ_{bt} (years)	4.05	0.82	3.17	4.12	4.88	4.01	0.86	2.98	4.06	4.89
Assets$_{bt}$ (€Mn)	10.17	1.39	7.84	10.73	11.34	10.04	1.47	7.63	10.67	11.34
Overdue ratio$_{bt}$	0.03	0.03	0.00	0.02	0.07	0.03	0.03	0.00	0.02	0.07
Liquidity ratio$_{bt}$	0.04	0.03	0.00	0.04	0.08	0.04	0.03	0.00	0.04	0.08
Deposit ratio$_{bt}$	0.20	0.12	0.02	0.20	0.33	0.19	0.12	0.00	0.19	0.33
Loan ratio$_{bt}$	0.65	0.11	0.54	0.63	0.85	0.66	0.12	0.54	0.63	0.88
Leverage$_{bt}$	0.11	0.05	0.06	0.11	0.16	0.11	0.05	0.06	0.11	0.16
Observations	40,488,265					26,488,990				

Panel B: Firm-level variables

	Mean	SD	p10	p50	p90
Inflation surprise, p.p.	0.59	1.88	-0.76	0.14	1.46
Firm-banks avg maturity exposure (θ_{ft}, in years)	7.61	6.04	3.18	4.99	15.43
Δ Investment (%)	0.92	21.11	-14.22	0.00	14.00
Δ Revenues (%)	28.18	143.94	-46.41	2.18	99.42
Δ Profits (%)	4.87	49.61	-23.42	0.14	31.38
Δ Interest expenses (%)	-0.00	0.97	-0.74	-0.00	0.76
Δ Number of employees (%)	1.80	25.43	-28.77	0.00	33.65
Δ Hours worked (%)	4.03	51.40	-37.76	0.00	48.03
Δ Credit (%)	2.52	38.15	-24.80	-0.87	33.69
Assets (€)	1,629,126	58,759,630	13,426	122,259	1,331,360
Number of employees	9.14	100.42	1.00	3.00	13.00
Firm age (years)	13.15	12.64	1.00	10.00	30.00
Loan-Assets ratio	0.54	0.67	0.01	0.32	1.32
Ratio of short-term loans	0.49	0.37	0.00	0.45	1.00
Return on Assets (ROA)	-0.13	0.95	-0.44	0.05	0.31
Capital per worker (€)	304.79	923.25	0.00	45.61	593.32
Exporting firm (dummy)	0.19	0.39	0.00	0.00	1.00
Observations	3,461,350				

Notes: This table reports summary statistics of the variables used in the micro analysis. Panel A presents firm–bank level variables from the HCRC database. Panel B reports firm-level characteristics from the CB database. Statistics are reported as mean, standard deviation (SD) and percentiles. All variable definitions are presented in Appendix Table B.1.

where b indexes banks, f indexes firms, and t indexes time in months. L_{bft} is the log-level of credit between bank b and firm f.[31] I focus on long-term credit, defined as loans with original maturity above one year, following the theoretical insights of Gomes, Jermann, and Schmid (2016), who show that inflation has real effects only when debt has long maturity.[32] $\pi_t^{\perp}$ is the local inflation in Portugal, as described in Section 2.3.1.[33] α_{ft} and α_{bf} are firm-time and bank-firm fixed effects, respectively. $\mathbf{X}_{t-1}$ is a rich vector of controls. It includes the lag of the dependent variable to control for autocorrelation, the share of loans

[31]This specification exploits the same within-bank–firm variation as a log-difference formulation in two-period designs, but does so more efficiently once the bank–firm fixed effect α_{bf} is included. Differencing introduces additional noise and mean reversion, and mechanically excludes newly established relationships when $L_{t-1} = 0$. For these reasons, I work with log-levels as the baseline and report results using log-differences as a robustness check.

[32]In robustness, I show that results are qualitatively similar when extending L_{bft} to measure overall credit.

[33]In the baseline specification, this is the orthogonalized series from year-over-year CPI inflation, similar to other studies in the literature (Core et al., 2025). Results are robust to using higher-frequency measures, although these introduce additional volatility and seasonality that can bias estimates.

in a bank with payments past due ($\geq$ 90 days) to measure loan portfolio risk, and the length of the credit relationship. I further control for the deposit ratio (deposits over assets), the loan ratio (loans over assets), and the short-term loan ratio (share of loans with maturity $\leq$ 1 year) to capture funding structure and maturity composition. Finally, I include other determinants emphasized in the bank-lending literature: size (Kashyap and Stein, 1995), leverage (Kishan and Opiela, 2000), and the share of liquid assets (Kashyap and Stein, 2000). Following Gomez et al. (2021) and English, Van den Heuvel, and Zakrajšek (2018), these controls enter both directly and interacted with current inflation $\pi^{\perp}$.[34] Results are robust to excluding these measures one by one.

I measure a bank's exposure to inflation through its *maturity gap*, $\sigma_{b,t-1}$. This variable captures the extent to which a bank borrows short and lends long, and has been widely used in the banking literature as a core characteristic of maturity transformation. A larger gap implies greater sensitivity of the balance sheet to nominal shocks, since liabilities reprice faster than assets. Prior work shows that the maturity gap is a strong predictor of banks' responses to monetary policy shocks (English, Van den Heuvel, and Zakrajšek, 2018; Drechsler, Savov, and Schnabl, 2021), which makes it well-suited to study the effect of inflation. I construct the measure using the BBS dataset, which reports the maturity class of most balance-sheet items (e.g., loans with maturity between one and five years). I assign each item the midpoint of its maturity interval and compute the average maturity of assets and liabilities. The maturity gap $\sigma_{b,t}$ is then defined as the difference between the weighted average maturity of assets and that of liabilities. The complete classification of maturity categories is reported in Appendix Table A.1, following the methodology in English, Van den Heuvel, and Zakrajšek (2018). $\sigma_{b,t-1}$ enters equation (2.4) as predetermined with respect to $\pi_t^{\perp}$, to avoid endogenous adjustments of the bank's maturity position caused by higher inflation: in the baseline, this is the average maturity gap of bank b in the last 12 months.[35] Figure 2.4 depicts the median maturity gap in the sample, together with its cross-sectional dispersion. The median is stable over time and fluctuating between 3 and 4 years, similar to what English, Van den Heuvel, and Zakrajšek (2018) report for the US.

Identifying assumptions The goal is to identify effects of inflation to credit supply through banks' maturity gap, i.e. the coefficient β in (2.4). My empirical design will be valid if the standard orthogonality condition is satisfied:[36]

[34]To clarify why the inclusion of the interaction term is crucial, suppose we exclude the interaction term $\mathbf{X}_{t-1} \times \pi_t^{\perp}$ from the specification, i.e., $L_{bft} = \beta\left[\pi_t^{\perp} \times \sigma_{b,t-1}\right] + \gamma_2 \mathbf{X}_{b,t-1} + \gamma_3 \sigma_{b,t-1} + \alpha_b + \alpha_{ft} + \underbrace{\gamma_1 \mathbf{X}_{b,t-1} \times \pi_t^{\perp} + u_{bft}}_{v_{bft}}$. If v_{bft} is correlated with $\pi_t^{\perp} \times \sigma_{b,t-1}$ — e.g., if $\mathrm{Cov}(\text{Deposit ratio}, \sigma_b) < 0$ — the OLS estimator becomes $\hat{\beta} = \beta + \frac{\mathrm{Cov}(\pi_t^{\perp} \times \sigma_{b,t-1}, v_{bft})}{\mathrm{Var}(\pi_t^{\perp} \times \sigma_{b,t-1})}$ and is biased towards zero.

[35]This is robust to both using time $t-1$ or time $t-12$.

[36]More formally, let $Z_{bft} \equiv \pi_t^{\perp} \times \sigma_{b,t-1}$ and $W_{bft} \equiv \{X_{t-1}, X_{t-1} \times \pi_t^{\perp}, \sigma_{b,t-1}, \alpha_{ft}, \alpha_{bf}\}$. The key assumption is the conditional mean independence $\mathbb{E}[u_{bft} \mid Z_{bft}, W_{bft}] = 0$ which implies the conditional moment (2.5).

Figure 2.4.: Maturity gap

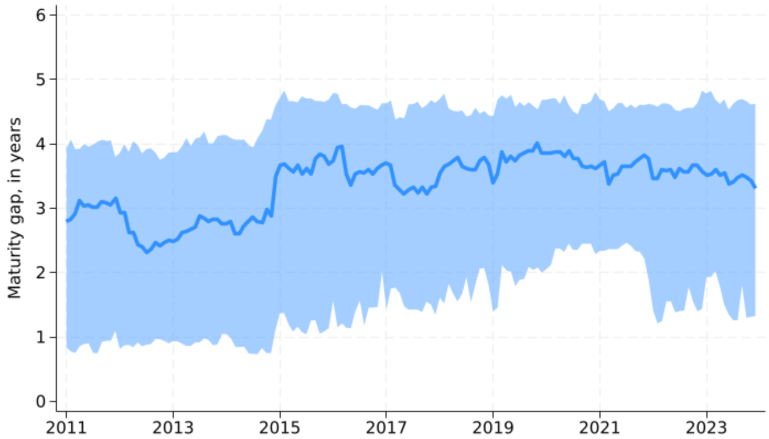

Notes: Sample period: 2011m1 - 2023m12. The solid blue line depicts the median maturity gap for the sample of 72 banks; the shaded blue band depicts the corresponding inter-quartile range. The maturity gap is defined as the weighted average maturity time of assets less the weighted average maturity time of liabilities. For details on each balance sheet item reported maturity, see Table A.1

$$\mathbb{E}\left[(\pi_t^{\perp} \times \sigma_{b,t-1})\ u_{bft}\ \mid\ \mathbf{X}_{bf,t-1}, \mathbf{X}_{bf,t-1} \times \pi_t^{\perp}, \sigma_{b,t-1}, \alpha_{ft}, \alpha_{bf}\right] = 0 \tag{2.5}$$

Intuitively, conditional on firm–time and bank–firm fixed effects and rich bank controls, the remaining variation in the interacted exposure $\pi_t^{\perp} \times \sigma_{b,t-1}$ is uncorrelated with the residual lending shock. This design is meant to address three main threats to identification that can arise in this framework.

First, credit demand shocks could confound the estimates if banks reduce lending simply because firms lower their borrowing needs. I address this identification problem by focusing on firms with multiple lending relationships and adding firm-time fixed effects α_{ft}. Any firm-level demand shock that is symmetric across lenders will be absorbed by the fixed effects (Khwaja and Mian, 2008). This assumption entails little costs in the Portuguese setting: the majority of firms maintain multibank relationships, providing high representativeness.

Second, an increase in Portuguese inflation might be correlated with a contemporaneous reaction of monetary policy. I alleviate this concern by exploiting both firm-time fixed effects and the fact that movements in Portuguese (idiosyncratic) inflation $\pi_t^{\perp}$ do not trigger a response from the ECB, which sets monetary policy at the Euro area level, as in Jiménez, Ongena, et al. (2012) (see Appendix Figure A.7 and A.8 for an empirical validation).[37]

The final identification threat is anticipation: banks might anticipate higher inflation and

[37] Firm–time fixed effects only absorb shocks that affect *homogeneously* all banks lending to the same firm. If the policy rate R_t affects banks differentially through the maturity gap $\sigma_{b,t-1}$, the term $R_t\sigma_{b,t-1}$ survives the within transformation. English, Van den Heuvel, and Zakrajšek (2018) show empirically that the coefficient on $R_t\sigma_{b,t-1}$ is indeed positive ($\delta > 0$), implying that OLS estimates of β would be biased towards zero when $\text{Cov}(\pi_t, R_t) > 0$. Therefore, other maintained assumptions are essential to disentangle the direct impact of inflation from monetary policy.

adjust their maturity gaps σ_b before the shock occurs. However, such behavior could bias the estimate only if banks react heterogeneously. If all banks uniformly shorten asset duration in anticipation, the *relative* maturity gap across banks remains intact and identification still follows. On the other hand, if for example only more exposed banks adjust, then the relative maturity gap used in the interaction $\pi_t^{\perp} \times \sigma_{b,t-1}$ becomes endogenous, biasing $\hat{\beta}$.[38] It can be shown that in case of anticipation, $\hat{\beta}$ represents an attenuated response, i.e., a lower bound (in absolute value) of the effect. One way to deal with this threat would be to use inflation surprises, which are unfortunately not available for Portugal at the monthly or quarterly frequency.[39] To mitigate this concern, I use the predetermined average $\sigma_{b,t-1}$ in the last 12 months, and use the 12-month lag in robustness specifications, with comparable results.

2.3.4. Results: inflation shrinks credit supply

Table 2.2 presents the estimation results of equation (2.4). A 1p.p. increase in inflation reduces the supply of long-term credit by a 1SD more maturity-exposed bank already after 1 month, by 0.22%.[40] The results are robust to the inclusion of various fixed effects and set of controls. Importantly, in column (2)-(5), I exploit *within firm across banks* variation through firm-time fixed effect: the effect controls for any demand variation, as it compares different banks lending to the same firm in the same month. The effect is persistent and grows with time: after 12 months, inflation reduces the supply of long-term credit by 1.52%.

There is no direct comparison in the literature to contrast this results. The closest instance is Agarwal and Baron (2024), who estimate a reduction in total lending during the 1977 inflationary episode (an increase in inflation from 5% to 7%) by roughly 3p.p.. However, their estimates do not identify purely credit supply effects and cannot disentangle from the effects caused by contractionary monetary policy. The results of Table 2.2 draw a stronger picture once these two elements are taken into account. Results are robust to using alternative lags of σ_b (Table C.1). Appendix Table C.2 reports alternative more granular specifications that support the main results: columns (1)-(2) further interact the firm-time fixed effect with a dummy equal to 1 if the bank is specialized in the firm's industry; columns (3)-(4) interact the firm-time fixed effect with a dummy equal to 1 if the bank is specialized in the firm's geographical area; moreover, results are strong also when restricting the set of banks to main entities (column (5)), when using a minimal set of

[38] More formally, suppose the data-generating process is $y_{bft} = \beta\pi_t^{\perp} \times \sigma_{b,t-1} + \delta\Delta\sigma_{b,t-1} + \alpha_{ft} + u_{bft}$ where $\Delta\sigma_{b,t-1}$ is an anticipated change in the maturity gap for some banks. If $\Delta\sigma_{b,t-1}$ correlates with both $\pi_t^{\perp}$ and $\sigma_{b,t-1}$, the OLS estimate mixes the true inflation-supply effect β with the anticipation effect δ, leading to attenuation bias, i.e., if $\beta < 0$, then $\hat{\beta} > \beta$.

[39] The only sources for inflation expectations in Portugal target annual inflation, even though they are updated more frequently during the year. Other methodologies (e.g., implied forecasts through the TIPS market) are not directly available in Portugal.

[40] $\beta \times \mathrm{SD}(\sigma_b) = 0.28 \times 0.77 = 0.22$

Table 2.2.: Inflation effect on banks' credit supply

	$\ln(\text{Credit}_{bf,t+h})$				
	$h=0$			$h=3$	$h=12$
	(1)	(2)	(3)	(4)	(5)
Inflation × $\bar{\sigma}_{b,t-1}$	-0.39*	-0.43**	-0.28*	-0.90**	-1.98***
	(0.22)	(0.19)	(0.14)	(0.42)	(0.74)
Controls	No	No	Yes	Yes	Yes
Time FE	Yes	No	No	No	No
Bank-firm FE	Yes	Yes	Yes	Yes	Yes
Firm-time FE	No	Yes	Yes	Yes	Yes
N	26,995,370	16,465,587	16,465,545	14,854,041	11,014,677
R^2	0.98	0.98	0.98	0.96	0.93

Notes: This table reports the results of estimating specification (2.4). The dependent variable is the log-level of long-term credit (maturity above one year) between bank b and firm f. The main explanatory variable is the interaction between orthogonalized Portuguese inflation $\pi_t^{\perp}$ and the bank's maturity gap $\sigma_{b,t-1}$. Controls include lagged credit, loan portfolio risk (share of overdue loans), relationship length, deposit ratio, loan ratio, short-term loan ratio, size, leverage, and liquidity, as well as their interactions with inflation. Firm–time and bank–firm fixed effects are included throughout. Standard errors are double-clustered at the bank and time level. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively.

controls (e.g., excluding leverage; column (6)) and when further controlling for changes in interest rates interacted with inflation – to explicitly control for interest rate risk – (column (7)). Estimation results with loan growth as dependent variables or using raw inflation as main regressor are more modest in magnitude (Table C.3 and Table C.4). In Section 2.3.5, I will provide clean evidence that during the 2021 inflation surge credit supply declined by about 5% for more exposed banks.

Extensive vs. intensive margin The previous framework can analyze the reduction in the intensive margin of the credit relationship. To investigate if the mechanism works also through the extensive margin, i.e., banks' establishment of new lending relationships or the closure of existing ones, I use the strategy proposed by Bittner et al. (2022). For each firm f that has an outstanding loan from bank b at some point in time t, I fill the respective bft panel with zeros for all other periods (indicating no credit exposure). This allows to examine the timing of new bank-firm relationships or of existing relationships that eventually terminate. The specification is similar to equation (2.4) but includes as dependent variable a dummy NewEntry_{bft} that equals one when a new relationship starts, and a dummy NewExit_{bft} that equals one when an existing relationship ends. The coefficient β can be interpreted as the additional probability of starting or ending a credit relationship with a 1SD more maturity-exposed bank, following an increase in inflation.

Table 2.3 reports the estimates for the contemporaneous response after a rise in inflation. Inflation pushes maturity-exposed banks to terminate existing credit relationships, although the magnitude is rather small. After a 1p.p. increase in inflation, a 1SD more maturity-exposed bank is only 0.03 percentage points more likely to terminate a credit

Table 2.3.: Extensive margin: Effect of inflation on bank-firm relationships

	New Entry			New Exit		
	(1)	(2)	(3)	(4)	(5)	(6)
Inflation × $\bar{\sigma}_{b,t-1}$	-0.00028	-0.00004	0.00000	0.00034*	0.00033**	0.00038***
	(0.00129)	(0.00086)	(0.00004)	(0.00019)	(0.00013)	(0.00009)
Controls	No	No	Yes	No	No	Yes
Bank FE	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	No	No	Yes	No	No
Time FE	Yes	No	No	Yes	No	No
Firm-Time FE	No	Yes	Yes	No	Yes	Yes
N	32,515,619	25,187,982	24,309,387	32,515,619	25,187,982	24,309,387
R^2	0.29	0.54	0.30	0.02	0.30	0.30

Notes: This table reports estimates of specification (2.4) for extensive-margin outcomes. The dependent variable is either NewEntry$_{bft}$, a dummy equal to one when a new bank–firm relationship is formed, or NewExit$_{bft}$, a dummy equal to one when an existing relationship ends. The main explanatory variable is the interaction between orthogonalized Portuguese inflation $\pi_t^{\perp}$ and the bank's maturity gap $\sigma_{b,t-1}$. Controls and fixed effects are as in Table 2.2. Standard errors are double-clustered at the bank and time level. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively.

relationship. By contrast, there is no significant effect on the creation of new lending relationships. These findings suggest that the credit-supply effect of inflation operates primarily through the intensive margin: banks tend to scale back lending volumes within existing relationships rather than frequently initiating or terminating ties. The extensive margin plays a secondary role, reflecting that banks usually adjust quantities before resorting to relationship termination.

Heterogeneous effects on banks The impact of inflation on credit supply need not be uniform across banks. Differences in size, capitalization, liquidity, and business models shape how strongly banks are exposed to maturity transformation and thus to inflation risk. Identifying such heterogeneity is important both to understand the transmission mechanism and to assess potential financial stability implications. In this subsection, I investigate whether the credit-supply effect of inflation varies systematically with bank characteristics beyond the maturity gap, focusing on balance-sheet strength and funding structure.

Table 2.4 presents the results at horizon $h = 12$ months. Banks with size below the median reduce particularly their credit supply. This is in line with Kashyap and Stein (1995), who rationalize the opposite sign response along the size distribution by acknowledging frictions for small banks in raising non-deposit external finance. Column (3) and (4) show that highly leveraged banks have higher balance sheet pressures when inflation rises, forcing them to reduce credit supply relatively more. The model in Section 2.5 exploits this financial constraint as part of the transmission mechanism. Finally, banks with low deposit reliance (defined as a deposit ratio below the median) drastically cut their lending

Table 2.4.: Heterogeneity: Effect of inflation on credit supply by bank characteristics

	ln(Credit$_{t+1y}$)					
	Large	Small	Lev.	Non-lev.	High Dep.	Low Dep.
	(1)	(2)	(3)	(4)	(5)	(6)
Inflation × $\bar{\sigma}_{b,t-1}$	0.91	-4.26***	-2.47**	-1.87**	-0.85	-6.59***
	(3.64)	(0.83)	(0.96)	(0.90)	(1.48)	(1.05)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Bank FE	Yes	Yes	Yes	Yes	Yes	Yes
Firm-Time FE	Yes	Yes	Yes	Yes	Yes	Yes
N	3,204,948	4,134,107	3,398,625	3,962,225	4,453,330	4,860,882
R^2	0.93	0.95	0.95	0.95	0.94	0.95

Notes: This table reports the results of estimating specification (2.4) at horizon $h = 12$ months, allowing for heterogeneity across banks. The dependent variable is the log-level of long-term credit. The main explanatory variable is the interaction between orthogonalized Portuguese inflation $\pi_t^{\perp}$, the bank's maturity gap $\sigma_{b,t-1}$, and bank characteristics. Small bank is a dummy equal to one if total assets are below the sample median. High leverage is a dummy equal to one if the capital ratio (equity/assets) is below the median. Low deposit reliance is a dummy equal to one if the deposit ratio (deposits/assets) is below the median. Controls and fixed effects are as in Table 2.2, with the exception that I exclude the control used to split the distribution, e.g., in column (1) and (2) I exclude log(Assets) from $\mathbf{X}_{t-1}$, and so on. Standard errors are double-clustered at the bank and time level. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively.

when inflation increases. This result connects to the aggregate evidence in Section 2.2.2: while deposits decline after an inflation surprise, banks with weaker deposit franchises are especially vulnerable because they rely more heavily on wholesale funding, which depreciates less relative to their asset side. This perspective helps reconcile the evidence with the argument of Drechsler, Savov, and Schnabl (2021): the stabilizing role of deposits is limited, but when this buffer is thin, inflation translates more directly into credit-supply contractions.

Appendix Table A.2 explores further dimensions of heterogeneity: banks with a higher reliance on loans in their asset portfolio and with longer average loan maturities cut lending more sharply, consistent with their greater maturity-gap exposure to inflation. Finally, the effect of inflation on bank credit supply appears broadly symmetric across positive and negative fluctuations.

2.3.5. Event study: the 2021 inflation surge

The 2021 inflation surge provides a clean setting for an event study, as the spike in CPI was mostly unexpected and policy rates did not rise until the second half of 2022. I exploit this episode to show that more maturity-exposed banks cut lending sharply even before markets anticipated any monetary tightening, confirming the relevance of the credit channel in cleaner and more parsimonious settings.

The surge of inflation in 2021 was largely unforeseen by professional forecasters, with prices increasing by more than 10% within few months and forecasts lagging behind at

3-4%.[41] Furthermore, expectations of an interest-rate hike by the ECB did not materialize until mid-2022 (see Appendix Figure A.10). This ensures that the first year of the inflation spike was not contaminated by actual or expected policy tightening.

I estimate the following event-study regression for the period 2021m1–2022m12:

$$L_{bft} = \sum_{t \neq \text{June 2021}} \beta_t \times \sigma_{b,2020} \times D_t + \gamma \mathbf{X}_{b,t-1} + \alpha_{ft} + \alpha_b + u_{bft}, \tag{2.6}$$

where $\sigma_{b,2020}$ is the pre-determined maturity gap of bank b, measured as the average between July and December 2020,[42] and D_t are month dummies. The coefficients β_t trace the differential response of more maturity-exposed banks relative to less exposed ones in each month after the shock. Identification is therefore obtained by comparing lending patterns before and after June 2021, within the same firm, across banks with different maturity gaps.

Figure 2.5.: Event study: credit supply during the 2021 inflation surge

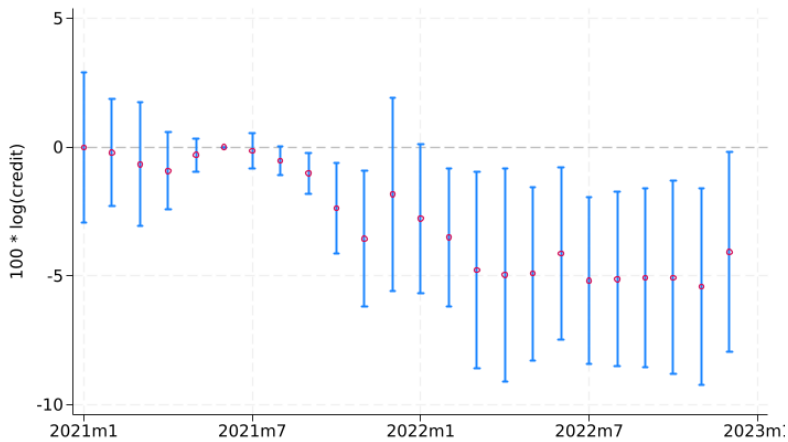

Notes: This figure reports the coefficients β_t from the event-study specification in equation (2.6). The dependent variable is the log-level of long-term credit at the bank–firm level. The sample covers 2021m1–2022m12. $\sigma_{b,2020}$ is the bank average maturity gap measured pre-sample (July–December 2020). Dotted red circles denote point estimates, vertical bars the 90% confidence intervals. Standard errors are clustered at the bank level.

Figure 2.5 plots the estimated coefficients of the event study around the spike in June 2021.[43] Banks one standard deviation more exposed to inflation reduce their credit supply by about 2 percentage points within six months of the surge, with the effect reaching 5 percentage points at the onset of the tightening cycle (July 2022). The decline is statistically

[41] These dynamics are not new to this episode, as Blanco, Ottonello, and Ranošová (2025) show that large inflation surges are usually unexpected.

[42] I exclude the previous period of 2020 to limit as much as possible the influence of the COVID shock on bank exposure. Results are robust to using alternative time periods.

[43] The results are robust to different starting periods of the inflation spike, from February 2021 to September 2021.

and economically significant, and it arises well before policy rates or expectations adjust, pointing to a direct credit-supply effect of inflation.

Overall, the event study validates the baseline evidence of Section 2.3.4: inflation erodes bank balance sheets and tightens credit supply. Crucially, this contraction is visible even in the absence of monetary tightening, underscoring that inflation itself operates as a distinct credit-supply shock. This evidence complements also the aggregate results in Section 2.2.2, showing that the contraction in credit supply following inflation surprises also emerges clearly in a clean, higher-frequency setting.

2.4. Micro-data: the consequences on borrower firms

This section shows that the contraction in credit supply after an inflation surprise spills over to borrower firms. Firms more exposed to inflation-sensitive banks cannot fully substitute credit from other lenders and cut back investment, with particularly sharp effects for smaller and younger firms. These findings confirm that the credit-supply mechanism can counteract the traditional debt-inflation channel and transmit inflation surprises to real activity.

2.4.1. Empirical strategy: firms' exposure through their lenders

The aim of this step is to assess how the contraction in bank credit supply transmits to firms. I measure each firm's exposure to inflation through the credit-weighted exposure of its lending banks:

$$\theta_{ft} = \sum_{b \in B_{ft}} w_{bft}\, \sigma_{b,t}, \tag{2.7}$$

where B_{ft} denotes the set of banks lending to firm f at time t, w_{bft} is the share of bank b in the firm's outstanding credit, and σ_{bt} is the maturity gap of bank b. Because the data are annual, both weights and exposures are averaged within the year.

Firm outcomes are modeled as

$$\frac{Y_{f,t+h} - Y_{f,t-1}}{\text{Assets}_{f,t-1}} = \beta \left[\pi_t^{\perp} \times \theta_{f,t-1}\right] + \gamma' \mathbf{X}_{f,t-1} + \hat{\alpha}_{ft} + \alpha_{g(f),t} + u_{f,t}, \tag{2.8}$$

where $Y_{f,t}$ is either total credit or investment.[44] The key regressor is the interaction between a measure of Portuguese inflation and lagged firm exposure $\theta_{f,t-1}$. In the baseline, I use the idiosyncratic inflation $\pi^{\perp}$, as in Section 2.3, while robustness exercises show the

[44] Investment is defined as the change in tangible and intangible capital plus depreciation. Credit is measured as total bank credit, matched from the credit registry. At the firm level, I express outcomes such as investment and profits as changes relative to lagged total assets, $\frac{Y_{f,t+h}-Y_{f,t-1}}{\text{Asset}_{f,t-1}}$. This scaling follows the standard approach, as in Gomez et al. (2021), and allows comparing firms of different sizes.

results using the inflation surprise ε_t^π (constructed from OECD forecasts, see Section 2.2), given the annual frequency of firms' balance sheet data. The coefficient β measures whether firms borrowing from more maturity-exposed banks perform worse after an increase in inflation.

The identification builds on the bank–firm regression in Equation (2.4). At that level, firm–time fixed effects α_{ft} absorb firm-specific credit demand shocks. Aggregating to the firm level introduces a challenge: if bank exposures $\sigma_{b,t}$ correlate with the firm–time shock α_{ft}, then the portfolio measure $\theta_{f,t}$ is also correlated with it. In addition, the firm-level specification cannot include firm-time fixed effects to absorb the firm-specific shocks. To address this, I follow Cingano, Manaresi, and Sette (2016), Jiménez, Mian, et al. (2020), and Pinardon-Touati (2024), and use as control $\hat{\alpha}_{ft}$, the estimated firm–time fixed effect from the bank–firm regression (2.4).[45] This control captures the firm-specific component of credit demand, ensuring that identification of β inherits from the bank–firm level strategy in Section 2.3.3.

$\alpha_{g(f),t}$ are firm fixed effects and sector–municipality–year fixed effects that absorb local and sectoral shocks. As a consequence, I only compare firms in the same municipality-sector-year cell. These firms are therefore subject to the same common shocks, but differ in their exposure to inflation because they borrow from different sets of banks. The vector $\mathbf{X}_{f,t-1}$ includes lagged firm-level controls: (i) log assets and firm age; (ii) leverage (debt-to-assets ratio), to account for the direct debt-inflation channel; (iii) turnover growth scaled by assets and return on assets (ROA), as proxies for demand conditions; (iv) capital intensity; (v) liquidity (cash and short-term assets relative to total assets); (vi) collateral capacity (tangible assets over total assets); and (vii) an export dummy.

A negative β implies that credit constraints at more maturity-exposed banks are transmitted to their borrowing firms, reducing firms' real activity after an inflation surprise.

2.4.2. Results: borrower firms cut back on investment

The response of firms' total credit Table 2.5 shows the estimated results of equation (2.8) on firm's total credit. Panel A reports the average effects. The results across all specifications indicate that firms borrowing from more exposed banks reduce credit after an increase in local inflation. After 1 year from a 1 p.p. increase in inflation, firms see their total credit reduced by $0.15 \times 6.04 = 0.9$ percentage points of assets, a decline of roughly 1.5%, close to the estimates in the loan-level regression. This fact highlights how firms cannot easily substitute their main lenders and are affected by shocks to banks' balance sheets.

The average response may mask significant heterogeneity. In Panel B, I explore ex-ante

[45] Because I use lagged exposure $\theta_{f,t-1}$, I also use the lagged estimated fixed effect from regression (2.4) at horizon $h = 12$. Results are robust also estimating $\hat{\alpha}_{ft}$ directly from a regression of L_{bft} on firm-time and bank-time fixed effects only, as in Pinardon-Touati (2024), and averaging it within the calendar year.

heterogeneity that reveals where the transmission mechanism of the credit channel of inflation concentrates. The negative effect on credit is sizable and precisely estimated in groups with tighter financing needs, e.g., (i) smaller firms and (ii) younger firms. These are firms with more difficulties in accessing other financing instruments, thus struggling to substitute the reduced credit provision by banks. On the other hand, highly leveraged firms and firms that rely on short-term financing face better credit conditions: leveraged firms benefit from a lower real debt burden, they appear as safer borrowers, and banks are more willing to roll over or even increase lending to them; short-term credit-reliant firms benefit from the fact that more maturity-exposed banks are incentivized to reduce loan maturity to hedge against inflation.[46] Together, these dynamics create a clear picture of where the credit reduction hits: while some firms see their credit contract sharply due to exposure of their lenders, others experience offsetting gains that dilute the mean effect. The heterogeneity thus helps reconcile the credit supply channel of inflation from banks to firm-level credit rationing.

[46]This mechanism is in line with predictions in the bank model of Agarwal and Baron (2024).

Table 2.5.: Effect of inflation on firms' total credit

Panel A: Benchmark: average effect

	$\frac{\text{Credit}_{f,t+h}-\text{Credit}_{f,t-1}}{\text{Assets}_{f,t-1}}$			
	$h=0$			$h=1$
	(1)	(2)	(3)	(4)
Inflation × $\theta_{f,t-1}$	-0.1016***	-0.0723***	-0.0762***	-0.1551***
	(0.0083)	(0.0173)	(0.0180)	(0.0270)
Firm FE	No	No	Yes	Yes
Sector-Municipality-Year FE	Yes	Yes	Yes	Yes
Firm Demand	No	Yes	Yes	Yes
Controls	No	No	Yes	Yes
N	1,766,226	358,565	338,049	278,872
R sq.	0.08	0.17	0.50	0.65

Panel B: Heterogeneity: effects by firm characteristics

	$\frac{\text{Credit}_{f,t+1}-\text{Credit}_{f,t-1}}{\text{Assets}_{f,t-1}}$							
	Large	Small	Lev.	Non-Lev.	Short	Long	Young	Old
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Inflation × $\theta_{f,t-1}$	-0.0699***	-0.2925***	0.2019***	-0.0317	0.0634**	-0.0275	-0.2429***	-0.0106
	(0.0180)	(0.0623)	(0.0266)	(0.0196)	(0.0322)	(0.0214)	(0.0697)	(0.0164)
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Sector-Municipality-Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm Demand	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
N	164,484	78,922	104,989	126,532	63,094	168,735	41,223	203,291
R sq.	0.68	0.75	0.76	0.68	0.73	0.71	0.83	0.66

Notes: This table reports the results of estimating specification (2.8) with the dependent variable equal to the annual change in firms' bank credit scaled by lagged assets. The key explanatory variable is the interaction between the idiosyncratic Portuguese inflation $\pi_t^{\perp}$ and the firm's loan-weighted exposure to banks' maturity gaps $\theta_{f,t-1}$. Panel A shows the benchmark estimates of the average effect. Panel B allows for heterogeneity by splitting firms' sample by the relevant distribution at horizon $h = 1$ year. Size (Large/Small) is measured above/below the 75th percentile of the asset distribution. Leverage (Lev./Non-Lev.) is measured above/below the 75th percentile of the Loan/Asset ratio distribution. Maturity structure of borrowing (Short/Long) is measured above/below the median of the Short-Loan/Asset distribution. Age (Young/Old) is measured below/above 8 years since the foundation – this is done to keep the sample balanced. Controls include lagged firm-level covariates (log assets, age, leverage, turnover growth, return on assets, capital intensity, liquidity, collateral ratio, export dummy), excluding one-by-one the covariate used to split the sample, and the lagged firm–time demand component $\hat{\alpha}_{ft}$. All regressions include firm fixed effects and sector–municipality–year fixed effects. Standard errors are clustered at the firm level. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

The response of firms' investment Table 2.6 reports the estimated effects on firms' investment. Panel A shows that, on average, investment declines significantly for firms borrowing from more inflation-exposed banks. Notice that the coefficient is initially positive when credit demand is not controlled for by $\hat{\alpha}_{ft}$, in line with a debt-inflation intuition. I formalize this opposing forces in the credit market in the model in Section 2.5, which highlight the importance of controlling for credit demand when assessing the effects of the credit channel of inflaton. Overall, after a 1 p.p. increase in local inflation, a firm one standard deviation more exposed through its banks reduces investment by about $0.04 \times 6.04 = 0.24$ percentage points of assets within 1 year. Relative to an average investment rate of 8% in the sample, this corresponds to a decline of about 3%, comparable to the effects of a 50 basis point monetary policy shock (Ottonello and Winberry, 2020). This results confirms that the contraction in bank credit supply translates into lower real investment at the firm level.

Panel B highlights the heterogeneity behind these results. The negative effect is strongest for firms that rely on banks to finance investment: those borrowing primarily through long-term credit—consistent with theoretical insights (Gomes, Jermann, and Schmid, 2016)—and young firms, which face limited access to financial markets.[47] For leveraged firms, the pattern is different. Although they do not reduce borrowing when attached to maturity-exposed banks – consistent with a debt-inflation mechanism – they still cut investment. This suggests that credit flows to leveraged firms are redirected toward rollover rather than new capital expenditures. The evidence underscores that the credit channel of inflation operates not only through the volume of credit but also through its allocation to productive investment.

In Appendix 2.E, I confirm the results above for credit and investment (Table D.1 and D.2) using as regressor the inflation surprise ε_t^{π}, rather then realized idiosyncratic. Then I present additional evidence that firms borrowing from maturity-exposed banks experience weaker fundamentals after an inflation surprise, facing higher interest expenses (Table D.3), lower revenues (Table D.4), and reduced profits (Table D.5). Moreover, the adverse effect on investment remains strong when focusing only on tangible investment (Table D.6).

Taken together, the firm-level evidence shows that inflation-induced credit supply shocks do not stop at the balance sheets of banks but propagate to the real economy by curtailing firms' borrowing and investment. The transmission is highly heterogeneous: while some borrowers manage to offset tighter bank lending, smaller, younger, and more bank-dependent firms face sharp investment cuts. These findings provide the motivation for the theoretical model that formalizes the credit channel of inflation in the next section.

[47] Young firms often drive innovation, so credit rationing via inflation could carry broader welfare consequences (Kim, 2025).

Table 2.6.: Effect of inflation on firms' investment

Panel A: Benchmark: average effect

	$\frac{\text{Investment}_{f,t+h} - \text{Investment}_{f,t-1}}{\text{Assets}_{f,t-1}}$			
	$h=0$			$h=1$
	(1)	(2)	(3)	(4)
Inflation × $\theta_{f,t-1}$	0.0094**	-0.0060	-0.0219**	-0.0420***
	(0.0045)	(0.0079)	(0.0091)	(0.0106)
Firm FE	No	No	Yes	Yes
Sector-Municipality-Year FE	Yes	Yes	Yes	Yes
Firm Demand	No	Yes	Yes	Yes
Controls	No	No	Yes	Yes
N	1,397,137	332,212	312,166	256,843
R sq.	0.06	0.13	0.33	0.40

Panel B: Heterogeneity: effects by firm characteristics

	$\frac{\text{Investment}_{f,t+1} - \text{Investment}_{f,t-1}}{\text{Assets}_{f,t-1}}$							
	Large	Small	Lev.	Non-Lev.	Short	Long	Young	Old
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Inflation × $\theta_{f,t-1}$	-0.0122	-0.0114	-0.0382***	-0.0286***	-0.0172	-0.0226**	-0.0686**	-0.0113
	(0.0080)	(0.0311)	(0.0122)	(0.0104)	(0.0108)	(0.0104)	(0.0349)	(0.0073)
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Sector-Municipality-Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm Demand	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
N	157,320	66,464	95,935	116,627	58,112	155,099	36,194	188,562
R sq.	0.42	0.55	0.50	0.50	0.54	0.46	0.61	0.41

Notes: This table reports the results of estimating specification (2.8) with the dependent variable equal to the annual change in firms' investment (Capital expenditure) scaled by lagged assets. The key explanatory variable is the interaction between the Portuguese inflation surprise $\pi_t^{\perp}$ and the firm's loan-weighted exposure to banks' maturity gaps $\theta_{f,t-1}$. Panel A shows the benchmark estimates of the average effect. Panel B allows for heterogeneity by splitting firms' sample by the relevant distribution. Size (Large/Small) is measured above/below the 75th percentile of the asset distribution. Leverage (Lev./Non-Lev.) is measured above/below the 75th percentile of the Loan/Asset ratio distribution. Maturity structure of borrowing (Short/Long) is measured above/below the median of the Short-Loan/Asset distribution. Age (Young/Old) is measured below/above 8 years since the foundation – this is done to keep the sample balanced. Controls include lagged firm-level covariates (log assets, age, leverage, turnover growth, return on assets, capital intensity, liquidity, collateral ratio, export dummy), excluding one-by-one the covariate used to split the sample, and the lagged firm–time demand component $\hat{\alpha}_{ft}$. All regressions include firm fixed effects and sector–municipality–year fixed effects. Standard errors are clustered at the firm level. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

2.5. A stylized model of inflation in the credit market

To rationalize the findings in Section 2.3 and 2.4, I develop a stylized model in which inflation surprises generate opposing forces in credit markets. Here, inflation surprises reduce the real burden of firms' nominal long-term debt while simultaneously eroding the real value of banks' long-term assets. As a result, firms would like to increase leverage and credit demand, but banks – financing long-term loans with short-term liabilities – face tighter capital constraints because they need to rollover their short-term debt at higher price, hence reducing credit supply. The model underscores the importance of controlling for credit demand in the empirical analysis of Sections 2.3 and 2.4, and formalizes the intuition that, even with flexible prices, inflation can affect credit supply and real activity when banks are exposed to a maturity mismatch in their balance sheets.

The framework builds on Kang and Pflueger (2015), who study firms with long-term nominal debt in an overlapping-generations environment that captures infrequent refinancing under flexible prices. In their setup, higher inflation lowers default risk and raises optimal leverage. I add a parsimonious bank block with a binding capital (leverage) constraint and costly external financing to show how inflation can contract credit supply and limit firms' effective leverage. In this literature, Agarwal and Baron (2024) develop a static partial equilibrium model to study how inflation affects banks. I extend this perspective, using a dynamic model and by emphasizing the interaction between bank balance sheets and credit demand from firms, which is essential to fully capture the impact of inflation on credit dynamics. In contrast, Corhay and Tong (2025) propose a full New Keynesian model with a banking sector, where inflation driven by monetary policy shocks constrains credit supply. My model is deliberately simpler and does not require that banks are subject to liquidity shocks. It highlights that the contractionary effect of inflation on credit supply is not necessarily tied to monetary policy and can arise even in flexible-price environments, isolating the mechanism underlying the *credit channel of inflation*.

2.5.1. Firm problem

In this subsection, I simply sketch the firm problem, describing the timing and main objects. I defer the overall derivation and discussion of the firm block in Appendix 2.F, while focusing on the bank block in the next subsection.

A. Firm timeline

Figure 2.6.: Timeline of firm i in cohort t

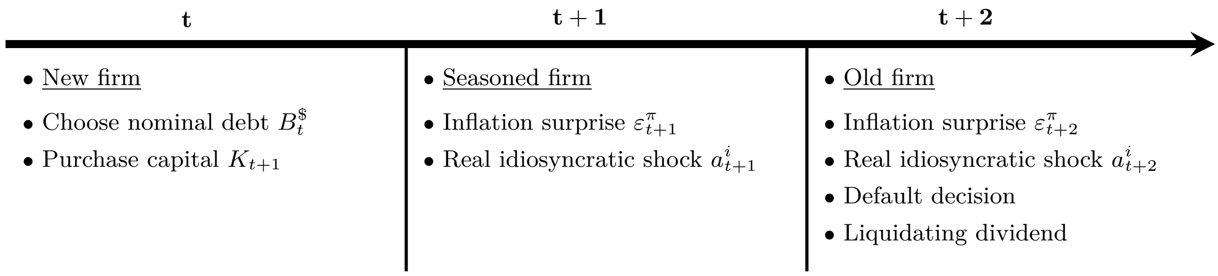

Figure 2.6 illustrates the timing for a firm that enters at the end of period t and produces for two periods. At the end of period t, the firm chooses the face value of nominal two-period debt $B_t^{\$}$ provided by the bank and purchases capital K_{t+1}, which becomes available for production at time $t+1$. The firm's corporate debt then has two periods remaining until maturity.

In period $t+1$, an aggregate inflation shock realizes, the firm experiences an idiosyncratic shock to its capital stock, and it produces. The firm cannot modify its capital structure, so leverage is sticky. The firm's seasoned corporate debt has one period remaining until maturity.

In period $t+2$, the firm again receives shocks and produces. At the end of period $t+2$, equity holders decide whether to default. Equity and debt holders then receive payments.

B. Inflation, stochastic discount factor and prices of loans

Let P_t denote the price level at time t and π_t the log-inflation from $t-1$ to t, $\pi_t = \log(P_t/P_{t-1})$. Consistent with Stock and Watson (2007), I let log inflation follow a random walk.

$$\begin{aligned} \pi_{t+1} &= \pi_t + \varepsilon_{t+1}^{\pi}, \\ \mathbb{E}_t[\pi_{t+1}] &= \pi_t. \end{aligned} \tag{2.9}$$

The two-period stochastic discount factors for pricing two-period real and nominal payoffs are, respectively,[48]

$$M_{t,t+2}\,, \quad M^{\$}_{t,t+2} = \frac{M_{t,t+2}}{\exp(\pi_{t+2}+\pi_{t+1})} = \frac{M_{t,t+2}}{\exp(2\pi_t + 2\varepsilon^{\pi}_{t+1} + \varepsilon^{\pi}_{t+2})}. \tag{2.10}$$

The price for a new corporate loan q_t^{new}, issued at time t with no partial repayment at $t+1$ and full repayment at $t+2$, is derived as the present value of the expected repayment, weighted by the probability of not defaulting in $t+2$. This is presented in Equation (2.29) in Appendix 2.F. There, I formally derive the following prediction:

Prediction 1. After a positive inflation shock, firms increase their credit demand.

A positive inflation shock ε^{π}_{t+1} affects the first-order condition for firm leverage (2.32) through the default threshold a^*_{t+2}: higher inflation lowers the probability of default, so the marginal benefit of debt increases. Firms therefore wish to raise leverage (i.e., issue more debt) after a positive inflation shock. This implication motivates the need in the empirical analysis of Section 2.3 and 2.4 to control for credit demand when assessing the impact of inflation on credit supply.

2.5.2. Bank problem

In this subsection, I show how banks with a positive maturity gap—financing long-term nominal loans through short-term liabilities—and a leverage constraint reduce credit supply after an inflation shock ε^{π}_{t+1}. This ultimately implies that firms end up with a suboptimal level of debt after the inflation shock.

A. Bank timeline and environment

A representative, competitive bank intermediates between international capital markets and firms. At each period t, it (i) grants two-period nominal loans to the new cohort with face value $B^{\$}_t$ at price $q_t^{\text{corp,new}}$, and (ii) raises short-term one-period nominal funding L_t at price $q_t^D = \mathbb{E}_t[M^{\$}_{t,t+1}]$ from international markets. Assets and liabilities are carried at face value; the bank receives cash funding $D_t = q_t^D L_t$ from its liabilities L_t. Balance-sheet items are recorded in real units by dividing by the price level P_t. The bank is price taking, distributes zero dividends in equilibrium, and is subject to a leverage constraint.

At $t+1$, the inflation shock ε^{π}_{t+1} realizes. The bank then repays L_t and rolls it over, adjusts its funding mix with costly external equity financing, and—before granting new

[48] As in Kang and Pflueger (2015), I can explicitly model the SDF using a representative consumer with expected power utility over consumption, risk aversion γ, and discount rate β,

$$U_t = \mathbb{E}_t \sum_{s=t}^{\infty} \exp(\beta(s-t)) \frac{C_s^{1-\gamma}}{1-\gamma}.$$

Then, the SDF is $M_{t,t+2} = \exp(-2\beta)(C_{t+2}/C_t)^{-\gamma}$.

two-period loans $B^{\$}_{t+1}$ to cohort $t+1$—faces again the leverage constraint.

B. Credit provision

At the beginning of period t, the bank is endowed with (real) equity E_t. It lends a nominal fixed-rate two-period loan $B^{\$}_t$ to firms, which will repay $B^{\$}_t/q^{new}_t$ in $t+2$, with no partial repayments in $t+1$,[49] and raises the remaining funds via international markets as one-period short-term debt with face value L_t—to be repaid in $t+1$—hence receiving cash $D_t = q^D_t L_t$ in period t. The balance-sheet identity, evaluated in nominal terms, is $B^{\$}_t = \underbrace{E^{nom}_t}_{=P_t E_t} + L_t$.

As in Iacoviello (2015), I assume that the bank is constrained in its ability to issue new loans by the amount of equity: $E^{nom}_t \geq \gamma B^{\$}_t$ with $\gamma \in (0,1)$. This constraint on the capital ratio can be motivated by standard limited-commitment problems or by regulatory concerns. Under competitive equilibrium, the capital ratio binds

$$\gamma B^{\$}_t = E^{nom}_t \implies L_t = (1-\gamma) B^{\$}_t \tag{2.11}$$

Notice that, in equilibrium at t, the change in credit supplied by the bank is proportional to its change in equity, i.e., $\Delta\text{Credit}_{t+1} \propto \Delta E_{t+1}$, yielding a *quantity cap* on new lending. No residual cash is carried over to $t+1$.

In $t+1$, the new price level P_{t+1} realizes, together with the inflation shock ε^{π}_{t+1}. No cash flow enters from the loan, while the one-period debt L_t is to be repaid. To do so, the bank rolls it over by raising new short-term debt L^{roll}_{t+1} at the new price $q^D_{t+1} = \mathbb{E}_t[M^{\$}_{t+1,t+2}]$, which needs to be repaid in $t+2$. By (2.11), the bank effectively borrows $L^{roll}_{t+1} = (1-\gamma)B^{\$}_t/q^D_{t+1}$ to roll over the initial short-term debt. Before any new lending and funding decision, at the beginning of $t+1$, real equity is therefore evaluated at the new price level as follows

$$E^{pre}_{t+1} = \frac{B^{\$}_t}{P_{t+1}} - \frac{L^{roll}_{t+1}}{P_{t+1}} = \frac{B^{\$}_t}{P_{t+1}}\Big(1 - \frac{1-\gamma}{q^D_{t+1}}\Big) \tag{2.12}$$

This is the core maturity-gap object: while long-term assets are nominally fixed and do not reprice in $t+1$, short-term liabilities are rolled over and embed the new price level through q^D_{t+1}.

C. Credit provision at $t+1$ and funding mix

During $t+1$, the bank lends a new two-period loan $B^{\$}_{t+1}$ to the $t+1$ cohort of new firms. It finances the loan through a mix of new external equity and short-term debt financing, since no internal financing is available from period t.[50] Define the debt share of the new

[49]This formulation without partial repayments is due to the *sticky leverage* of firms, which allows having firms defaulting only in $t+2$.

[50]This is due to the absence of partial repayment of $B^{\$}_t$ and zero profits of the competitive bank.

lending as

$$\theta \equiv \frac{q^D_{t+1} L^{new}_{t+1}}{B^\$_{t+1}} \in [0,1] \implies L^{new}_{t+1} = \frac{\theta}{q^D_{t+1}} B^\$_{t+1} \tag{2.13}$$

Resources raised by new equity financing, used for new lending, is therefore $X_{t+1} = (1-\theta)B^\$_{t+1}$. However, external equity financing is particularly costly for the bank. Let $\bar{E}$ denote *gross* external equity issued at $t+1$. Only a fraction $1-\rho$ arrives as deployable cash, i.e., $X_{t+1} = (1-\rho)\,\bar{E}$ with $\rho \in (0,1)$, where $\bar{E}$ represents a clear supply capacity, i.e. the bank cannot issue more new external equity financincg than $\bar{E}$.

After choosing the new lending and new funding, equity evolves as[51]

$$E^{post}_{t+1} = E^{pre}_{t+1} + \frac{B^\$_{t+1}}{P_{t+1}} - \frac{L^{new}_{t+1}}{P_{t+1}} \tag{2.14}$$

with the new capital constraint on total assets, binding in equilibrium:

$$E^{post}_{t+1} = \gamma\Big(\frac{B^\$_t}{P_{t+1}} + \frac{B^\$_{t+1}}{P_{t+1}}\Big). \tag{2.15}$$

Substitute (2.12) and (2.13) into (2.14), and impose (2.15). Rearranging terms by $(B^\$_t, B^\$_{t+1})$ yields

$$\Big[(1-\gamma) - \frac{\theta}{q^D_{t+1}}\Big]\, B^\$_{t+1} \;=\; (1-\gamma)\,\frac{1-q^D_{t+1}}{q^D_{t+1}}\, B^\$_t. \tag{2.16}$$

Since $B^\$_t \geq 0$ and $\frac{1-q^D_{t+1}}{q^D_{t+1}} \geq 0$ for $q^D_{t+1} \in (0,1]$, the right-hand side is nonnegative. Feasibility requires $B^\$_{t+1} \geq 0$, which in turn implies that the bracket on the left-hand side must be nonnegative, i.e., $\theta \;\leq\; q^D_{t+1}(1-\gamma)$. Therefore the *maximal feasible debt share* at $t+1$ is $\theta^{\max}(q^D_{t+1}) \equiv q^D_{t+1}(1-\gamma)$, and the corresponding *minimal equity share* per unit of new lending is

$$s_{\min}(q^D_{t+1}) \;\equiv\; 1-\theta^{\max}(q^D_{t+1}) \;=\; 1-q^D_{t+1}(1-\gamma) \;\in (0,1) \tag{2.17}$$

Intuitively, after the rollover tightens the bank's balance-sheet position, a fraction $s_{\min}$ of every euro of new loans must be financed with equity to satisfy the new capital constraint. To maximize feasible lending at $t+1$ for given (q^D_{t+1}, γ), the bank chooses the most debt-intensive feasible mix, $\theta = \theta^{\max}$, so that equity use per unit is minimized:

$$X_{t+1} \;=\; s_{\min}(q^D_{t+1})\, B^\$_{t+1} \tag{2.18}$$

Using the wedge $X_{t+1} = (1-\theta)B^\$_{t+1}$, and allowing at most $\bar{E}$ of gross issuance, we get

[51] Notice that the book equity increment $B^\$_{t+1} - L^{new}_{t+1} = \big(1 - \frac{\theta}{q^D_{t+1}}\big)B^\$_{t+1}$ is generally different from the resources raised through equity $X_{t+1} = (1-\theta)B^\$_{t+1}$. Only when $q^D_{t+1} = 1$ (par funding) and $\rho = 0$ (no wedge) do they coincide.

$s_{\min}(q^D_{t+1})\, B^{\$}_{t+1} \;\le\; (1-\rho)\,\bar{E}$, so that the new lending supply at t+1 is

$$B^{\$}_{t+1} \;=\; \frac{(1-\rho)\,\bar{E}}{s_{\min}(q^D_{t+1})} \;=\; \frac{(1-\rho)\,\bar{E}}{1-q^D_{t+1}(1-\gamma)}\;. \tag{2.19}$$

Notice that the price level P_{t+1} drops out of (2.19) and the maturity-mismatch transmission at t+1 operates entirely through q^D_{t+1} (liability repricing and rollover), not through pure price-level scaling or net nominal position.[52]

D. Inflation, maturity gap, and credit supply

In this paragraph, I show that, when a bank has a maturity gap, the period t short-term liability *rolls* at t+1 into the new price q^D_{t+1}. Hence, an inflation surprise lowers q^D_{t+1} and tightens new credit supply.

First, notice that $\frac{\partial q^D_{t+1}}{\partial \varepsilon^{\pi}_{t+1}} < 0$: an inflation shock decreases the price of short-term debt (raises short-term funding interest rates) because, intuitively, it is inversely related to the nominal discount factor, i.e., $q^D_{t+1} = \mathbb{E}_{t+1}[M^{\$}_{t+1,t+2}] = \mathbb{E}_{t+1}[\frac{M_{t+1,t+2}}{\exp(\pi_{t+2})}] = \mathbb{E}_{t+1}[\frac{M_{t+1,t+2}}{\exp(\pi_t+\varepsilon^{\pi}_{t+1}+\varepsilon^{\pi}_{t+2})}]$.

The next proposition shows that an inflation shock that decreases the price of short-term debt pushes the bank to contract credit supply.

Proposition 1 (Inflation shock constrains credit supply)**.** *After an inflation shock ε^{π}_{t+1} that raises the price level P_{t+1} and satisfies $\frac{\partial q^D_{t+1}}{\partial \varepsilon^{\pi}_{t+1}} < 0$, the change in new lending at t+1 is negative:*

$$\frac{\partial B^{\$}_{t+1}}{\partial \varepsilon^{\pi}_{t+1}} < 0$$

Proof. From (2.19), $B^{\$}_{t+1} = \frac{(1-\rho)\,\bar{E}}{1-q^D_{t+1}(1-\gamma)}$. Differentiate with respect to $q \equiv q^D_{t+1}$:

$$\frac{\partial B^{\$}_{t+1}}{\partial q} \;=\; (1-\rho)\,\bar{E}\;\cdot\;\frac{1-\gamma}{[1-q(1-\gamma)]^2} \;>\; 0.$$

Therefore, an inflation shock that *lowers* q^D_{t+1} implies $\frac{\partial B^{\$}_{t+1}}{\partial \varepsilon^{\pi}_{t+1}} < 0$. □

The mechanism is transparent: inflation forces short-term liability rollover at a worse price for the bank ($q^D_{t+1}\downarrow$), which raises the minimal equity share $s_{\min} = 1 - q^D_{t+1}(1-\gamma)$ required by the capital constraint; with costly equity, each euro of issuance yields only $1-\rho$ euros of usable cash, so lending capacity contracts.

[52] In this baseline case, there is no internal financing for the bank in t+1, so external issuance is the only source for X_{t+1}. If one allowed retained earnings or a carried buffer, this would simply add to the numerator, leaving intact the main intuition.

Now, I compare the effects of an inflation shock on a *maturity-mismatched* bank ("MM"; two-period loan; one-period funding rolled at t+1) and a *non-mismatched* bank ("NM"; one-period loan; one-period funding repaid at t+1), under the same costly external equity wedge $1-\rho$ and the same equity issuance capacity $\bar{E}$. With a gross issuance capacity $\bar{E}$ and wedge $1-\rho$, the new lending capacities are

$$B_{t+1}^{\$,\,\mathrm{MM}} \;=\; \frac{(1-\rho)\,\bar{E}}{s_{\min}(q)}, \qquad B_{t+1}^{\$,\,\mathrm{NM}} \;=\; \frac{X_{t+1}^{\mathrm{int}} + (1-\rho)\,\bar{E}}{s_{\min}(q)} \tag{2.20}$$

where $s_{\min} = 1 - q_{t+1}^{D}(1-\gamma)$ and, in the non-mismatched case, the period t-position produces internal financing at t+1,

$$X_{t+1}^{\mathrm{int}} \;=\; \Big(\frac{1}{q_t^{\mathrm{new}}} - (1-\gamma)\Big) B_t^{\$}, \tag{2.21}$$

because both the loan the debt are repaid after one period, in $t+1$. Notice that the internal financing does not depend on $q \equiv q_{t+1}^{D}$, as it is fixed at origination in t.
Let $\bar{E} = \bar{E}(E_{t+1}^{pre})$ with $\partial \bar{E}/\partial E^{pre} \geq 0$, capturing that a stronger pre-lending equity position relaxes market access for additional equity financing. From (2.12), the equity positions before lending in $t+1$ are

$$E_{t+1}^{pre,\,\mathrm{MM}} = \frac{B_t^{\$}}{P_{t+1}}\Big(1 - \frac{1-\gamma}{q}\Big), \qquad E_{t+1}^{pre,\,\mathrm{NM}} = \frac{B_t^{\$}}{P_{t+1}}\Big(\frac{1}{q_t^{\mathrm{new}}} - (1-\gamma)\Big), \tag{2.22}$$

so only the mismatched case has *direct* q-sensitivity in E^{pre}.

Proposition 2 (Milder credit contraction without maturity mismatch)**.** *Suppose $\bar{E} = \bar{E}(E_{t+1}^{pre})$ with $\partial \bar{E}/\partial E^{pre} > 0$, and let $B_{t+1}^{\$,\,\mathrm{MM}}$ and $B_{t+1}^{\$,\,\mathrm{NM}}$ be given by* (2.20)*. After an inflation shock ε_{t+1}^{π} that raises the price level P_{t+1} and satisfies $\frac{\partial q_{t+1}^{D}}{\partial \varepsilon_{t+1}^{\pi}} < 0$, the change in new lending at* t+1 *is more negative for the mismatched bank than for the non-mismatched bank:*

$$\frac{\partial B_{t+1}^{\$,\,\mathrm{MM}}}{\partial \varepsilon_{t+1}^{\pi}} < \frac{\partial B_{t+1}^{\$,\,\mathrm{NM}}}{\partial \varepsilon_{t+1}^{\pi}},$$

Proof. See Appendix 2.F.3. □

In the mismatched case, the period-t short-term liability rolls at t+1 into the new price q_{t+1}^{D}, so an inflation surprise that lowers q_{t+1}^{D} tightens two margins at once: (i) it raises the minimal equity share $s_{\min} = 1 - q(1-\gamma)$ required by the capital constraint (so every euro of new lending needs more equity financing), and simultaneously (ii) it erodes pre-new lending equity E^{pre} via rollover, which contracts external equity issuance capacity $\bar{E}$ when market access is linked to the bank's net worth. The non-mismatched bank, by contrast, has no need to rollover debt at t+1: both asset and liability positions settle and even generate internal cash, so the inflation surprise affects it only by raising $s_{\min}$; it does *not* worsen $\bar{E}$ through a contemporaneous hit to E_{t+1}^{pre} via higher q_{t+1}^{D}. Hence, the credit contraction is milder without maturity mismatch.

Overall, the model rationalizes the main intuition of the credit channel of inflation: while inflation reduces firms' real debt burden and stimulates credit demand, it simultaneously weakens banks' balance sheets, leading them to restrict lending and constrain firm investment.

2.6. Conclusion

This paper shows that inflation can operate through a *credit channel*. When prices spike, banks with a maturity gap suffer equity losses and tighter capital constraints. They curtail lending, and firms borrowing from these banks reduce investment. In macro data for advanced economies, a 1 p.p. inflation surprise is followed by a contraction in loans and bank capital of about 3%, even in the absence of monetary policy moves. In micro data from Portugal, more exposed banks supply less credit, consistent with evidence during the 2021 inflation surge. At the firm level, this has real effects, forcing firms to cut their investment rate. These findings highlight a new cost of inflation: it reallocates and compresses credit through banks, offsetting the debt–inflation relief for borrowers and dampening real activity.

This channel has several implications. First, assessments of inflation shocks should account for bank capital and maturity-transformation exposures. Macroprudential buffers and banks' funding structure shape the strength of the channel and therefore the transmission of inflation to the real economy. Moreover, the heterogeneity of the effects – stronger for younger, bank-dependent firms – points to distributional and welfare consequences beyond aggregates. Taken together, this paper shows that controlling inflation is also a financial-stability policy: the consequences of inflation for the macroeconomy can crucially depend on the health of financial intermediaries.

References

Agarwal, Isha and Matthew Baron (2024). "Inflation and disintermediation". In: *Journal of Financial Economics* 160, p. 103902 (cit. on pp. 82, 83, 86, 97, 104, 108).

Almeida, Heitor and Thomas Philippon (2007). "The risk-adjusted cost of financial distress". In: *The Journal of Finance* 62.6, pp. 2557–2586 (cit. on p. 139).

Andersen, Asger Lau, Niels Johannesen, Mia Jørgensen, and José-Luis Peydró (2023). "Monetary policy and inequality". In: *The Journal of Finance* 78.5, pp. 2945–2989 (cit. on p. 88).

Bahaj, Saleem, Robert Czech, Sitong Ding, and Ricardo Reis (2023). *The market for inflation risk*. Tech. rep. Bank of England (cit. on p. 83).

Ball, Laurence, Stephen G Cecchetti, and Robert J Gordon (1990). “Inflation and uncertainty at short and long horizons”. In: *Brookings Papers on Economic Activity* 1990.1, pp. 215–254 (cit. on p. 83).

Banco de Portugal Microdata Research Laboratory (BPLIM) (2023). *Harmonized Central Credit Responsibility Database*. Version V1, Extraction: June 2023. Banco de Portugal - BPLIM. Dataset. DOI: `10.17900/HCRC.Jun2023.V1` (cit. on p. 93).

– (2024a). *Bank Balance Sheet Monthly Data*. Version V1, Extraction: June 2024. Banco de Portugal - BPLIM. Dataset. DOI: `10.17900/BBS.Jun2024.V1` (cit. on p. 93).

– (2024b). *Central Balance Sheet Annual Data*. Version V1, Extraction: June 2024. Banco de Portugal - BPLIM. Dataset. DOI: `10.17900/CB.CBA.Jun2024.V1` (cit. on p. 93).

Baron, Matthew, Emil Verner, and Wei Xiong (2021). “Banking crises without panics”. In: *The Quarterly Journal of Economics* 136.1, pp. 51–113 (cit. on pp. 80, 84).

Bazot, Guillaume, Eric Monnet, and Matthias Morys (2024). “Central banks and the absorption of international shocks (1891-2019)”. In: (cit. on p. 88).

Bergant, Katharina, Mr Divya Kirti, Ms Mai Hakamada, and Rui Mano (2025). *Inflation and Bank Profits: Monetary Policy Trade-offs*. Tech. rep. International Monetary Fund (cit. on pp. 83, 88).

Bernanke, Ben S. and Alan S. Blinder (1992). “The Federal Funds Rate and the Channels of Monetary Transmission”. In: *The American Economic Review* 82.4, pp. 901–921 (cit. on pp. 80, 88).

Bittner, Christian, Diana Bonfim, Florian Heider, Farzad Saidi, Glenn Schepens, and Carla Soares (2022). “The augmented bank balance-sheet channel of monetary policy”. In: *Available at SSRN 4047820* (cit. on p. 98).

Blanco, Andrés, Pablo Ottonello, and Tereza Ranošová (2025). “The dynamics of large inflation surges”. In: *Review of Economics and Statistics*, pp. 1–31 (cit. on pp. 86, 101).

Bleaney, Michael and Manuela Francisco (2005). “Exchange rate regimes and inflation: only hard pegs make a difference”. In: *Canadian Journal of Economics/Revue canadienne d’économique* 38.4, pp. 1453–1471 (cit. on p. 88).

Boyd, John H, Ross Levine, and Bruce D Smith (2001). “The impact of inflation on financial sector performance”. In: *Journal of monetary Economics* 47.2, pp. 221–248 (cit. on p. 83).

Brunnermeier, Markus, Sergio Correia, Stephan Luck, Emil Verner, and Tom Zimmermann (2025). “The Debt-Inflation Channel of the German (Hyper) Inflation”. In: *American Economic Review* 115.7, pp. 2111–2150 (cit. on pp. 79, 82).

Chinn, Menzie D and Hiro Ito (2006). “What matters for financial development? Capital controls, institutions, and interactions”. In: *Journal of development economics* 81.1, pp. 163–192 (cit. on p. 89).

Choi, Sangmok, Bruce D Smith, and John H Boyd (1996). “Inflation, financial markets, and capital formation”. In: *Review-Federal Reserve Bank of Saint Louis* 78, pp. 9–35 (cit. on p. 83).

Cieslak, Anna and Carolin Pflueger (2023). "Inflation and asset returns". In: *Annual Review of Financial Economics* 15.1, pp. 433–448 (cit. on p. 87).
Cieslak, Anna and Pavol Povala (2015). "Expected returns in Treasury bonds". In: *The Review of Financial Studies* 28.10, pp. 2859–2901 (cit. on p. 85).
Cingano, Federico, Francesco Manaresi, and Enrico Sette (2016). "Does credit crunch investment down? New evidence on the real effects of the bank-lending channel". In: *The Review of Financial Studies* 29.10, pp. 2737–2773 (cit. on pp. 81, 103).
Colla, Paolo, Filippo Ippolito, and Kai Li (2020). "Debt structure". In: *Annual Review of Financial Economics* 12.1, pp. 193–215 (cit. on p. 79).
Cooley, Thomas F and Gary D Hansen (1989). "The inflation tax in a real business cycle model". In: *The American Economic Review*, pp. 733–748 (cit. on p. 83).
Core, Fabrizio, Filippo De Marco, Tim Eisert, and Glenn Schepens (2025). "Inflation and floating-rate loans: Evidence from the euro-area". In: (cit. on p. 94).
Corhay, Alexandre and Jincheng Tong (2025). "Inflation risk and the finance-growth nexus". In: *Rotman School of Management Working Paper* 3795679 (cit. on pp. 83, 87, 108).
D'Andrea, Angelo, Andrea Fabiani, Fabio Massimo Piersanti, and Anatoli Segura (2025). *Inflation, Leverage and Stock Returns*. Tech. rep. CEPR Discussion Paper (cit. on pp. 82, 129).
DeMarzo, Peter M, Arvind Krishnamurthy, and Stefan Nagel (2024). *Interest rate risk in banking*. Tech. rep. National Bureau of Economic Research (cit. on pp. 87, 88).
Dotsey, Michael and Peter Ireland (1996). "The welfare cost of inflation in general equilibrium". In: *Journal of Monetary Economics* 37.1, pp. 29–47 (cit. on p. 83).
Drechsler, Itamar, Alexi Savov, and Philipp Schnabl (2021). "Banking on deposits: Maturity transformation without interest rate risk". In: *The Journal of Finance* 76.3, pp. 1091–1143 (cit. on pp. 81, 87, 88, 95, 100).
– (2022). *Credit crunches and the great stagflation* (cit. on p. 83).
Driscoll, John C and Aart C Kraay (1998). "Consistent covariance matrix estimation with spatially dependent panel data". In: *Review of economics and statistics* 80.4, pp. 549–560 (cit. on pp. 86, 87, 90, 121–123, 130, 131).
Ellison, Martin, Sang Seok Lee, and Kevin Hjortshøj O'Rourke (2024). "The ends of 27 big depressions". In: *American Economic Review* 114.1, pp. 134–168 (cit. on p. 85).
English, William B (1999). "Inflation and financial sector size". In: *Journal of Monetary Economics* 44.3, pp. 379–400 (cit. on p. 83).
English, William B, Skander J Van den Heuvel, and Egon Zakrajšek (2018). "Interest rate risk and bank equity valuations". In: *Journal of Monetary Economics* 98, pp. 80–97 (cit. on pp. 81, 95, 96).
Fama, Eugene F and G William Schwert (1977). "Asset returns and inflation". In: *Journal of financial economics* 5.2, pp. 115–146 (cit. on p. 87).
Feldstein, Martin S (1999). "Capital income taxes and the benefit of price stability". In: *The Costs and Benefits of Price Stability*. University of Chicago Press, pp. 9–46 (cit. on p. 83).

Fischer, Stanley, Robert E Hall, and John B Taylor (1981). "Relative shocks, relative price variability, and inflation". In: *Brookings Papers on Economic Activity* 1981.2, pp. 381–441 (cit. on p. 83).

Fisher, Irving (1933). "The Debt-Deflation Theory of Great Depressions". In: *Econometrica* 1.4, pp. 337–357 (cit. on pp. 79, 82).

Ghosh, Atish R, Mahvash S Qureshi, and Charalambos G Tsangarides (2014). "On the value of words: Inflation and fixed exchange rate regimes". In: *IMF Economic Review* 62.2, pp. 288–322 (cit. on p. 88).

Giannone, Domenico, Michele Lenza, and Lucrezia Reichlin (2008). *Business cycles in the euro area.* Tech. rep. National Bureau of Economic Research (cit. on p. 91).

Gomes, Joao, Urban Jermann, and Lukas Schmid (2016). "Sticky leverage". In: *American Economic Review* 106.12, pp. 3800–3828 (cit. on pp. 79, 82, 94, 106).

Gomez, Matthieu, Augustin Landier, David Sraer, and David Thesmar (2021). "Banks' exposure to interest rate risk and the transmission of monetary policy". In: *Journal of Monetary Economics* 117, pp. 543–570 (cit. on pp. 95, 102).

Gonçalves, Sílvia, Ana María Herrera, Lutz Kilian, and Elena Pesavento (2022). "When do state-dependent local projections work?" In: (cit. on p. 129).

Gourio, Francois (2013). "Credit risk and disaster risk". In: *American Economic Journal: Macroeconomics* 5.3, pp. 1–34 (cit. on p. 139).

Graham, John R (2000). "How big are the tax benefits of debt?" In: *The journal of finance* 55.5, pp. 1901–1941 (cit. on p. 139).

Greenwald, Emily, Sam Schulhofer-Wohl, and Josh Younger (2023). "Deposit convexity, monetary policy and financial stability". In: (cit. on p. 87).

Hamilton, James D (1992). "Was the deflation during the Great Depression anticipated? Evidence from the commodity futures market". In: *The American Economic Review*, pp. 157–178 (cit. on p. 85).

Holm-Hadulla, Fédéric, Alberto Musso, Giulio Nicoletti, and Mika Tujula (2022). "Firm debt financing structures and the transmission of shocks in the euro area". In: *Economic Bulletin Articles* 4 (cit. on p. 79).

Huybens, Elisabeth and Bruce D Smith (1999). "Inflation, financial markets and long-run real activity". In: *Journal of monetary economics* 43.2, pp. 283–315 (cit. on p. 83).

Iacoviello, Matteo (2015). "Financial business cycles". In: *Review of Economic Dynamics* 18.1, pp. 140–163 (cit. on p. 111).

Jarociński, Marek and Peter Karadi (2020). "Deconstructing monetary policy surprises—the role of information shocks". In: *American Economic Journal: Macroeconomics* 12.2, pp. 1–43 (cit. on pp. 92, 124).

Jensen, Michael C (1986). "Agency costs of free cash flow, corporate finance, and takeovers". In: *The American economic review* 76.2, pp. 323–329 (cit. on p. 139).

Jiménez, Gabriel, Atif Mian, José-Luis Peydró, and Jesús Saurina (2020). "The real effects of the bank lending channel". In: *Journal of Monetary Economics* 115, pp. 162–179 (cit. on pp. 81, 103).

Jiménez, Gabriel, Steven Ongena, José-Luis Peydró, and Jesús Saurina (2012). “Credit supply and monetary policy: Identifying the bank balance-sheet channel with loan applications”. In: *American Economic Review* 102.5, pp. 2301–2326 (cit. on pp. 80, 88, 91, 96).

Jordà, Òscar (2005). “Estimation and inference of impulse responses by local projections”. In: *American economic review* 95.1, pp. 161–182 (cit. on p. 86).

Jordà, Òscar, Katharina Knoll, Dmitry Kuvshinov, Moritz Schularick, and Alan M Taylor (2019). “The rate of return on everything, 1870–2015”. In: *The Quarterly Journal of Economics* 134.3, pp. 1225–1298 (cit. on p. 84).

Jordà, Òscar, Björn Richter, Moritz Schularick, and Alan M Taylor (2021). “Bank capital redux: Solvency, liquidity, and crisis”. In: *The Review of economic studies* 88.1, pp. 260–286 (cit. on pp. 80, 84).

Jordà, Òscar, Moritz Schularick, and Alan M Taylor (2017). “Macrofinancial history and the new business cycle facts”. In: *NBER macroeconomics annual* 31.1, pp. 213–263 (cit. on p. 84).

– (2020). “The effects of quasi-random monetary experiments”. In: *Journal of Monetary Economics* 112, pp. 22–40 (cit. on p. 89).

Jordà, Òscar, Sanjay R Singh, and Alan M Taylor (2024). “The long-run effects of monetary policy”. In: *Review of Economics and Statistics*, pp. 1–49 (cit. on p. 90).

Kang, Johnny and Carolin E Pflueger (2015). “Inflation risk in corporate bonds”. In: *The journal of finance* 70.1, pp. 115–162 (cit. on pp. 82, 108, 110, 137).

Kashyap, Anil K and Jeremy C Stein (1995). “The impact of monetary policy on bank balance sheets”. In: *Carnegie-rochester conference series on public policy*. Vol. 42. Elsevier, pp. 151–195 (cit. on pp. 95, 99).

– (2000). “What do a million observations on banks say about the transmission of monetary policy?” In: *American Economic Review* 90.3, pp. 407–428 (cit. on p. 95).

Keynes, JM (1923). *A Tract on Monetary Reform*. Macmillan (cit. on p. 79).

Khwaja, Asim Ijaz and Atif Mian (2008). “Tracing the impact of bank liquidity shocks: Evidence from an emerging market”. In: *American Economic Review* 98.4, pp. 1413–1442 (cit. on pp. 81, 96).

Kim, Chi Hyun, Lorenzo Ranaldi, and Moritz Schularick (2025). “Inflation Surprises and Asset Returns: a Macrohistory Perspective”. In: *Working Paper* (cit. on pp. 80, 85, 87, 128).

Kim, Seula (2025). *Workers’ Job Prospects and Young Firm Dynamics*. US Census Bureau, Center for Economic Studies (cit. on p. 106).

Kishan, Ruby P and Timothy P Opiela (2000). “Bank size, bank capital, and the bank lending channel”. In: *Journal of Money, credit and banking*, pp. 121–141 (cit. on p. 95).

Lucas, Robert E (1994). “On the welfare cost of inflation”. In: (cit. on p. 83).

Mankiw, N Gregory (2001). “The inexorable and mysterious tradeoff between inflation and unemployment”. In: *The Economic Journal* 111.471, pp. 45–61 (cit. on p. 90).

Meckling, William H and Michael C Jensen (1976). “Theory of the Firm”. In: *Managerial behavior, agency costs and ownership structure* 3.4, pp. 305–360 (cit. on p. 139).

Montiel Olea, José Luis and Mikkel Plagborg-Møller (2021). “Local projection inference is simpler and more robust than you think”. In: *Econometrica* 89.4, pp. 1789–1823 (cit. on p. 86).

Nakamura, Emi, Jón Steinsson, Patrick Sun, and Daniel Villar (2018). “The elusive costs of inflation: Price dispersion during the US great inflation”. In: *The Quarterly Journal of Economics* 133.4, pp. 1933–1980 (cit. on p. 83).

Obstfeld, Maurice (1997). “Destabilizing effects of exchange-rate escape clauses”. In: *Journal of International Economics* 43.1-2, pp. 61–77 (cit. on pp. 80, 88).

Obstfeld, Maurice, Jay C Shambaugh, and Alan M Taylor (2005). “The trilemma in history: tradeoffs among exchange rates, monetary policies, and capital mobility”. In: *Review of economics and statistics* 87.3, pp. 423–438 (cit. on pp. 80, 88).

Oman, William (2019). “The synchronization of business cycles and financial cycles in the euro area”. In: *International Journal of Central Banking* 15.1, pp. 327–362 (cit. on pp. 80, 91).

Ottonello, Pablo and Thomas Winberry (2020). “Financial heterogeneity and the investment channel of monetary policy”. In: *Econometrica* 88.6, pp. 2473–2502 (cit. on pp. 82, 106).

Pinardon-Touati, Noémie (2024). *The crowding out effect of local government debt: Micro- and macro-estimates.* Tech. rep. (cit. on pp. 81, 103).

Quinn, Dennis, Martin Schindler, and A Maria Toyoda (2011). “Assessing measures of financial openness and integration”. In: *IMF Economic Review* 59.3, pp. 488–522 (cit. on p. 89).

Ramey, Valerie A and Sarah Zubairy (2018). “Government spending multipliers in good times and in bad: evidence from US historical data”. In: *Journal of political economy* 126.2, pp. 850–901 (cit. on p. 89).

Stock, James H and Mark W Watson (2007). “Why has US inflation become harder to forecast?” In: *Journal of Money, Credit and banking* 39, pp. 3–33 (cit. on p. 109).

Vogel, Lukas (2007). “How do the OECD growth projections for the G7 economies perform? A post-mortem”. In: (cit. on p. 85).

2.A. Additional figures and tables

2.A.1. Figures

Figure A.1.: Response of *nominal* bank balance sheet items to inflation surprises, 1946-2020

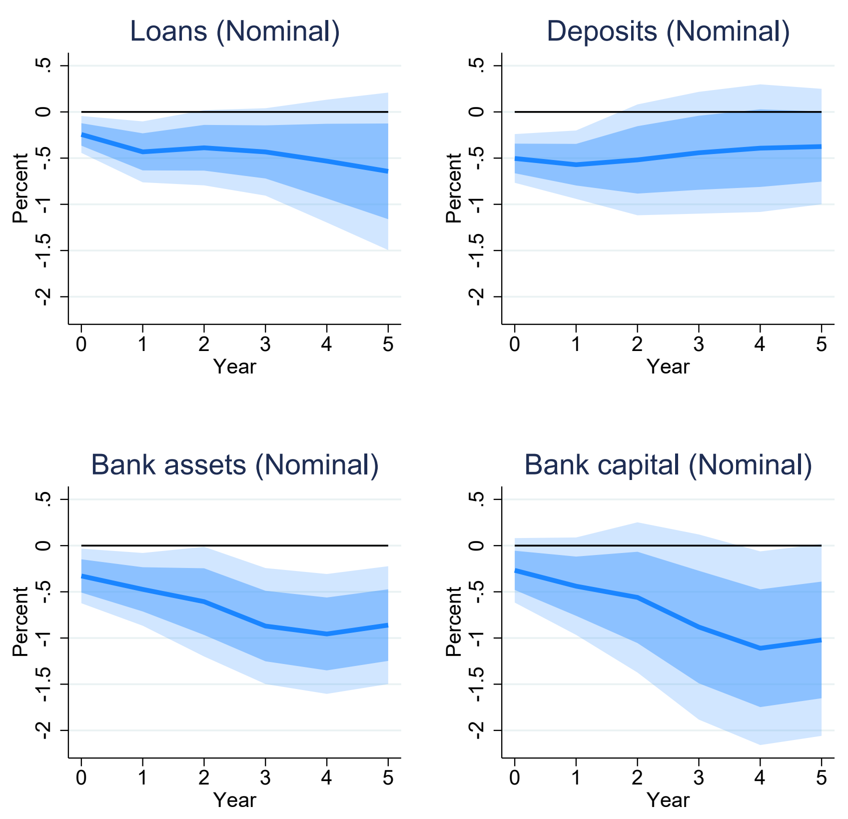

Notes: Impulse responses for aggregate loans, deposits, bank capital and cumulative equity returns on inflation surprises, estimated with panel-LP as in eq. (2.1). 1946-2020, annual frequency, 18 countries. All variables are expressed in nominal terms. Driscoll and Kraay (1998) standard errors. Bands are at the 68% and 90% confidence level.

Figure A.2.: Response of equity returns to inflation surprises, 1946-2020

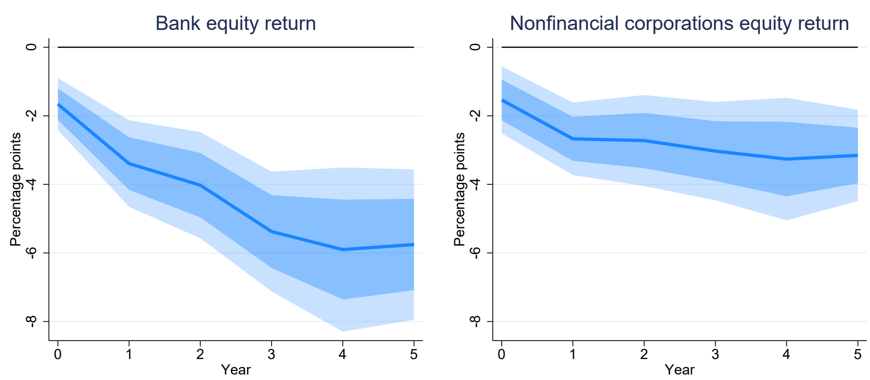

Notes: Impulse responses for banks' and non-financial corporations' cumulative real equity returns on inflation surprises, estimated with panel-LP as in eq. (2.1). 1946-2020, annual frequency, 18 countries. Driscoll and Kraay (1998) standard errors. Bands are at the 68% and 90% confidence level.

Figure A.3.: Response of non-core funding to inflation surprises, 1946-2020

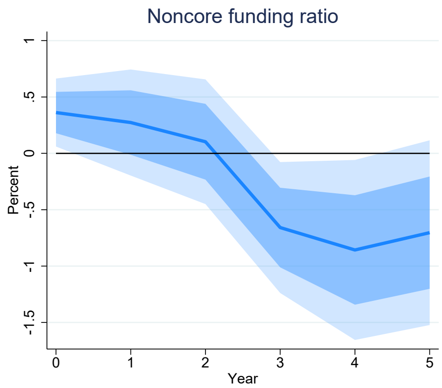

Notes: Impulse responses of banks' non-core funding ratio on inflation surprises, estimated with panel-LP as in eq. (2.1). 1946-2020, annual frequency, 18 countries. Non-core funding ratio is defined as $\frac{\text{Non-core}_{it}}{\text{Non-core}_{it}+\text{Deposits}_{it}}$, where Non-core_{it} includes wholesale funding and interbank lending. Driscoll and Kraay (1998) standard errors. Bands are at the 68% and 90% confidence level.

Figure A.4.: Effect of inflation surprises on the macroeconomy, 1946-2020

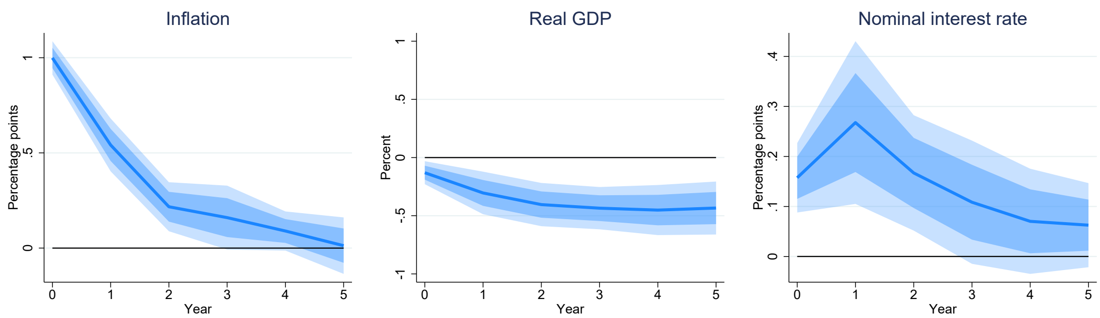

Notes: Impulse responses for inflation, real GDP and nominal short-term interest rates on inflation surprises, estimated with panel-LP as in eq. (2.1). 1946-2020, annual frequency, 18 countries. Driscoll and Kraay (1998) standard errors. Bands are at the 68% and 90% confidence level.

Figure A.5.: Direct effect of inflation surprises on the macroeconomy, 1870-2020

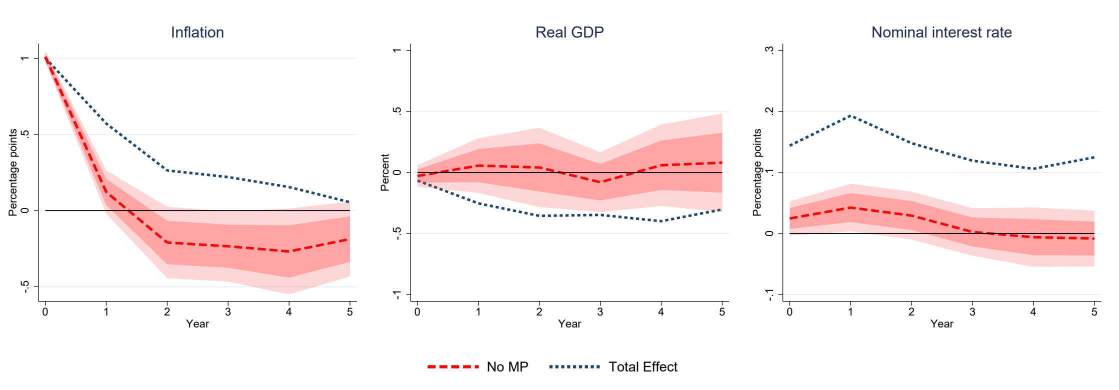

Notes: Impulse responses for inflation, real GDP and nominal short-term interest rates on inflation surprises, estimated with panel-LP as in eq. (2.2). The red solid line represent the "direct" effect of inflation, i.e., the response of outcome y_t in open peg countries ($\beta_{A,h}$). The blue dashed line shows the "total" effect of inflation, i.e., the response of outcome y_t in float countries ($\beta_{B,h}$). 1870-2020, annual frequency, 18 countries. Driscoll and Kraay (1998) standard errors. Bands are at the 68% and 90% confidence level.

Figure A.6.: Orthogonalized inflation in Portugal, $\pi_t^{\perp}$

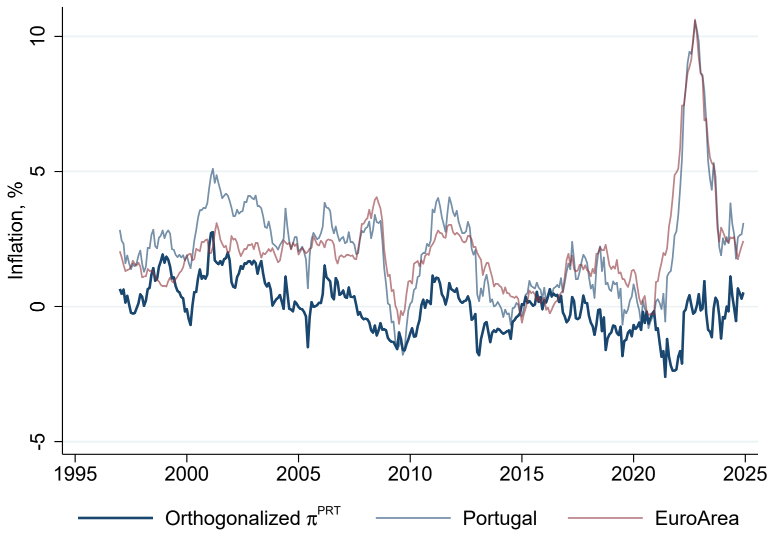

Notes: Orthogonalized Portugues inflation $\pi_t^{\perp}$, Euro-area aggregate inflation π_t^{EA} and raw Portuguese inflation π_t^{PT}. 1999m1-2024m12, year-over-year CPI inflation, monthly frequency.

Figure A.7.: Response of inflation to monetary policy shocks, 1999-2024

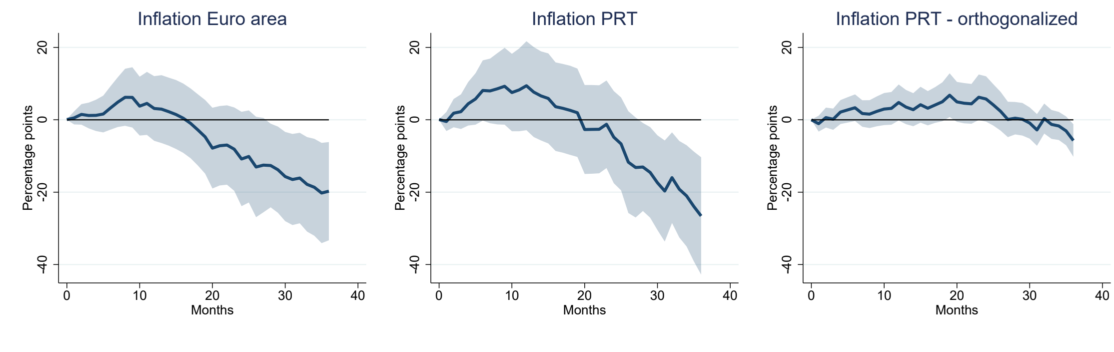

Notes: Impulse responses for euro area inflation, portuguese inflation and orthogonal portuguese inflation on Jarociński and Karadi (2020) monetary policy shocks, estimated with LP-IV, instrumenting the 3-month interest rate with the contemporaneous shock. 1999m1-2024m12, monthly frequency. Bands are at the 90% confidence level.

Figure A.8.: Response of the short-term interest rate to $\pi^{\perp}$, 1999-2024

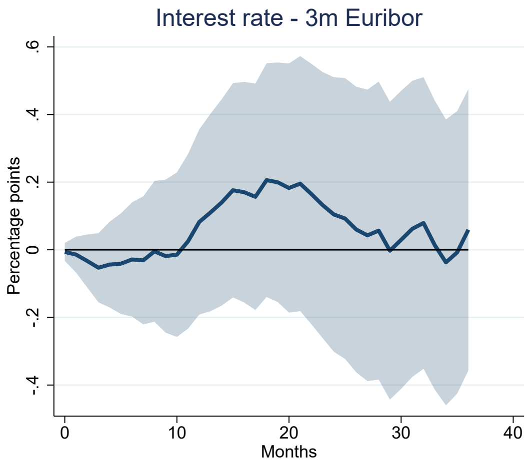

Notes: Impulse responses for the 3-months EURIBOR on the orthogonalized Portuguese inflation $\pi^{\perp}$ estimated with local projections. Controls include the 12-month lags of: $\pi^{\perp}$, euro-area inflation, 3m EURIBOR, real GDP growth of Portugal, real GDP growth of the aggregate Euro Area. 1999m1-2024m12, monthly frequency. Bands are at the 90% confidence level.

Figure A.9.: Firms with multibank relationships

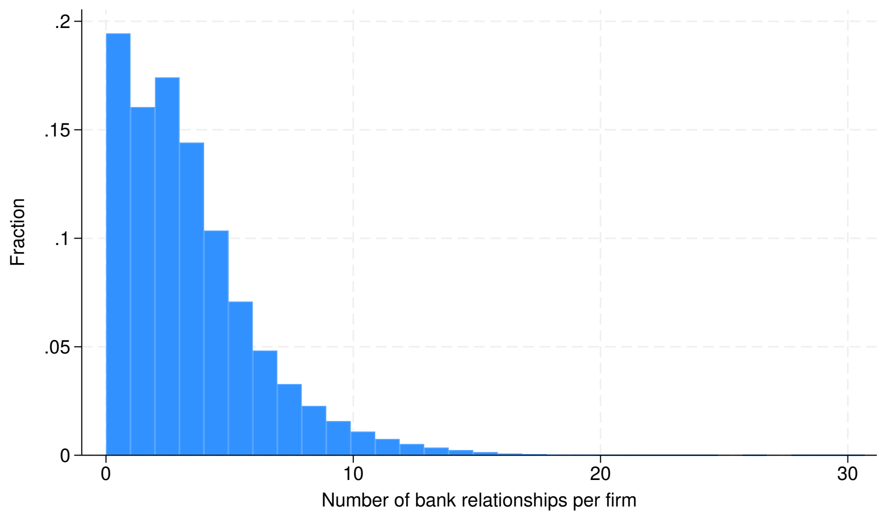

Notes: Histogram of the number of multibank relationships. HCRC dataset. Portugal, 2011-2022.

Figure A.10.: Event study: Expectations during the 2021 inflation surge

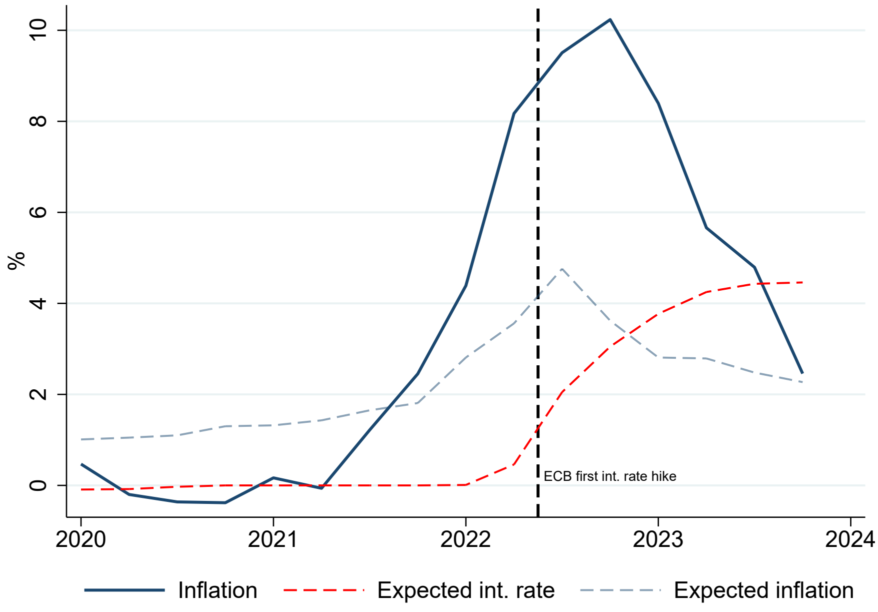

Notes: Inflation (blue solid line) shows the year-over-year inflation rate in Portugal. Expected inflation (light blue dashed line) reports the forecasted inflation for the next 12 months (data source: Survey of Professional Forecasters, ECB). Expected interest rate (red dashed line) reports the forecasted average interest rate for the next 3 months (data source: Survey of Professional Forecasters, ECB). Sample period: 2020Q1 - 2023Q4.

2.A.2. Tables

Table A.1.: Maturity of bank balance sheet items

Balance-sheet item	Assigned maturity (years)
Cash	0
Loans $\leq$ 1 y	0.5
Loans 1–5 y	3
Loans $\geq$5 y	10
Bonds $\leq$ 1y	0.5
Bonds 1-2y	1.5
Bonds $\geq$ 2y	5
Stocks	10
Properties/Equipments	10
Demand & Reedemable deposits ($\leq$ 3 m)	0.25
Deposits $\leq$ 1y	0.5
Deposits 1-2y	1.5
Deposits $\geq$ 2y	5
Trading liabilities	1

Notes: Assigned maturity for each balance sheet item, in the computation of the *maturity gap*. *Notes:* Stocks and properties have no contractual maturity but are generally treated as long-term, illiquid assets funded by equity, which itself is permanent capital. Banking practice and regulation (e.g., Basel NSFR, IRRBB) classify them as requiring stable, long-term funding and often assign them behavioral maturities of 5–10 years in asset–liability management.

Table A.2.: Heterogeneity: Effect of inflation on credit supply by bank characteristics

	High loans	Low loans	Short	Long	Positive	Negative
	(1)	(2)	(3)	(4)	(5)	(6)
Inflation × $\bar{\sigma}_{b,t-1}$	-2.69***	-1.03	1.84	-4.73***	-2.86***	-2.19***
	(0.73)	(2.11)	(1.38)	(1.45)	(1.07)	(0.81)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Bank FE	Yes	Yes	Yes	Yes	Yes	Yes
Firm-Time FE	Yes	Yes	Yes	Yes	Yes	Yes
N	3,718,703	3,843,317	3,730,395	4,146,349	5,457,807	5,532,399
R^2	0.96	0.94	0.95	0.95	0.93	0.94

Notes: This table reports the results of estimating specification (2.4) at horizon $h = 12$ months, allowing for heterogeneity across banks. The dependent variable is the log-level of long-term credit. The main explanatory variable is the interaction between orthogonalized Portuguese inflation $\pi_t^{\perp}$ and the bank's maturity gap $\bar{\sigma}_{b,t-1}$ averaged in past 12 months. Sample are split according to the following bank characteristics. *High Loans* restricts the sample to banks with the share of loans to total assets that is above the sample median. *Short* restricts the sample to banks with the share of short-term loans to total loans that is above the sample median. Positive restricts the sample to positive variations of $\pi^{\perp}$. Controls and fixed effects are as in Table 2.2, with the exception that I exclude the control used to split the distribution, e.g., in column (1) and (2) I exclude the Loan ratio from $\mathbf{X}_{t-1}$, and so on. Standard errors are double-clustered at the bank and time level. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively.

2.B. Data description

Table B.1.: Variable Definitions

Variable Name	**Definition**
Panel A: Bank–firm (loan-level) variables	
Credit outstanding L_{bft} (€)	Credit with original maturity above 1 year (euros)
Length of relationship$_{bft}$ (months)	Duration of the bank–firm relationship (b, f) up to t, in months
Maturity gap σ_{bt} (years)	Bank-level maturity gap for bank b, as midpoint of item-level maturity (Table A.1).
Assets$_{bt}$ (€Mn)	Total assets, in €millions.
Overdue ratio$_{bt}$	Share of overdue credit (more than 90 days) over total credit.
Liquidity ratio$_{bt}$	Liquid assets (cash and short-term bonds) divided by total assets
Deposit ratio$_{bt}$	Customer deposits (Demand, sight, savings) divided by total assets
Loan ratio$_{bt}$	Total loans divided by total assets
Leverage$_{bt}$	Capital ratio (i.e., equity over total assets).
Firm-level variables	
Inflation surprises	CPI inflation surprise in Portugal, $\pi_t - \mathbb{E}_{t-1}\pi_t$, from Kim, Ranaldi, and Schularick (2025)
Firm-banks avg maturity exposure (θ_{ft}, in years)	Lender-weighted average maturity gap of firm f's banks, as in 2.7
Investment	CapEX: change in fixed tangible and intangible assets + depreciation
Revenues	Turnover
Profits	Total income − Total expenses
Interest expenses	Interest expenses
Hours worked	Hours worked by paid employees
Credit	Total bank credit over total lagged assets
Assets (€)	Total assets, in euros
Number of employees	Headcount of paid employees
Firm age (years)	Years since foundation of the firm
Loan-Assets ratio	Total bank credit divided by total assets
Ratio of short-term loans	Share of outstanding bank credit with original (or residual) maturity ≤ 1 year
Return on Assets (ROA)	Profits divided by total assets
Capital per worker (€)	Fixed capital stock divided by number of employees; euros per worker
Exporting firm (dummy)	Indicator equal to 1 if the firm reports positive exports in year t, 0 otherwise.

2.C. Additional macro-level impulse responses

This section provides a back-of-the-envelope exercise to evaluate whether the debt–inflation channel can, in the aggregate, be offset by the credit channel of inflation and become contractionary. I show that when corporate indebtedness is high, inflation surprises trigger sharper declines in bank fundamentals, and are also followed by lower investment. These results suggest that the credit channel of inflation can more than offset the debt-relief benefits for firms.

The *debt–inflation* channel argues that higher inflation reduces the real debt burden of leveraged firms, freeing up additional resources to invest and produce. Especially when corporate debt is at high levels, we should expect inflation to particularly boost economic activity via higher investment. I test this argument by implementing a state-dependent local projection on the macro data of Section 2.2, using the specification (2.2), where the state $\mathcal{I}$ is defined by the level of corporate debt relative to GDP. Country–years with debt above the country median are classified as high-debt states.[53] The debt–inflation mechanism would predict that inflation surprises should be least contractionary in these states, as the real burden of nominal liabilities falls most strongly.

Figure C.1 shows the main results. In high-debt states, aggregate credit and bank capital decline significantly more after an inflation surprise than in low-debt states, and bank equity returns also register larger losses. The response of non-financial corporations' equity returns is partially in line with a debt–inflation mechanism: firms benefit initially from the reduction in real debt burdens (as in D'Andrea et al. (2025)) when debt levels are high, yet the cumulative response turns significantly negative in the long run.

What do these dynamics imply for aggregate investment? I show in Figure C.2 that aggregate investment also contracts more strongly in high-debt states. This provides suggestive evidence that, even in periods of elevated corporate leverage, the contraction in bank credit supply can dominate the debt-relief channel and can depress real activity. More broadly, it illustrates that inflation can undermine the very mechanism often thought to make it expansionary. Inflation does not simply redistribute between creditors and debtors; it destabilizes banks' balance sheets and tightens financing conditions – an insight that underscores the importance of recognizing the credit channel as a central cost of inflation.

[53] As prescribed by Gonçalves et al. (2022), I test for the state variable $\mathcal{I}$ to be independent from the inflation surprise. In a regression of the state on the surprises, the coefficient of interest is always statistically insignificant, confirming the validity of the state-dependent exercise.

Figure C.1.: Effect of inflation surprises with high corporate debt

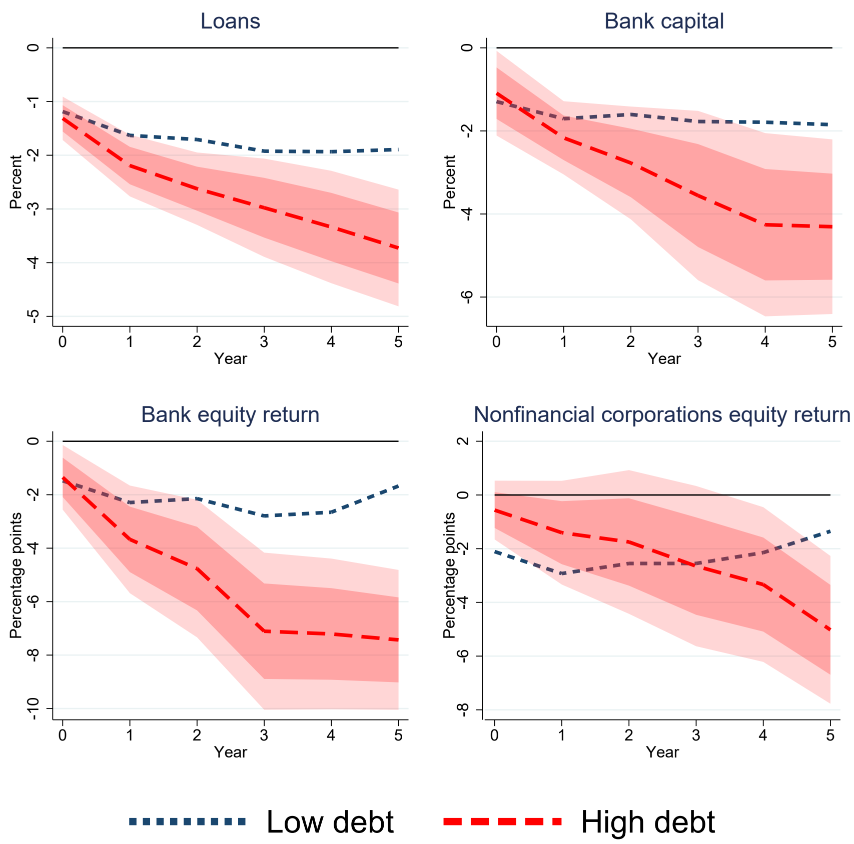

Notes: Impulse responses for aggregate loans, bank capital, cumulative bank equity returns and cumulative non-financial corporations equity returns on inflation surprises, estimated with panel-LP as in eq. (2.2). The red dashed line represents the effect of inflation when corporate debt-over-GDP is above the median. 1946-2020, annual frequency, 18 countries. Driscoll and Kraay (1998) standard errors. Bands are at the 68% and 90% confidence level.

Figure C.2.: Effect of inflation surprises with high corporate debt

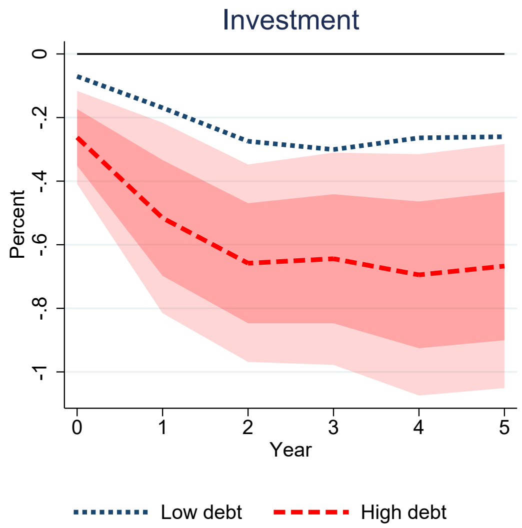

Notes: Impulse responses for aggregate investment (measured as the logarithm of real investment) on inflation surprises, estimated with panel-LP as in eq. (2.2). The red dashed line represents the effect of inflation when corporate debt-over-GDP is above the median. 1946-2020, annual frequency, 18 countries. Driscoll and Kraay (1998) standard errors. Bands are at the 68% and 90% confidence level.

2.D. Robustness: loan-level regressions

Table C.1.: Inflation effect on banks' credit supply: lag of σ_b

	$\ln(\text{Credit}_{bf,t+h})$					
	$h=0$				$h=12$	
	(1)	(2)	(3)	(4)	(5)	(6)
Inflation $\times\sigma_{b,t-1}$	-0.43**	-0.31*			-2.49***	
	(0.18)	(0.17)			(0.79)	
Inflation $\times\sigma_{b,t-12}$			-0.23	-0.07		-1.79**
			(0.15)	(0.16)		(0.73)
Controls	No	Yes	No	Yes	Yes	Yes
Time FE	No	No	No	No	No	No
Bank-firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Firm-time FE	Yes	Yes	Yes	Yes	Yes	Yes
N	18,174,350	18,174,185	12,819,261	12,819,237	11,859,962	8,280,271
R^2	0.98	0.98	0.99	0.99	0.93	0.94

Notes: This table reports the results of estimating specification (2.4). The dependent variable is the log-level of long-term credit (maturity above one year) between bank b and firm f. The main explanatory variable is the interaction between orthogonalized Portuguese inflation $\pi_t^{\perp}$ and the bank's maturity gap σ_b. The first line reports the 1-month lag, i.e., $\sigma_{b,t-1}$. The second line reports the 12-month lag, i.e., $\sigma_{b,t-12}$. Controls include lagged credit, loan portfolio risk (share of overdue loans), relationship length, deposit ratio, loan ratio, short-term loan ratio, size, leverage, and liquidity, as well as their interactions with inflation. Firm–time and bank–firm fixed effects are included throughout. Standard errors are double-clustered at the bank and time level. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively.

Table C.2.: Inflation effect on banks' credit supply: more granular specifications

	$\ln(\text{Credit}_{bf,t+12})$						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Inflation $\times\bar{\sigma}_{b,t-1}$	-3.73**	-1.88**	-4.05**	-2.00**	-1.45*	-2.90**	-1.99***
	(1.69)	(0.77)	(1.74)	(0.78)	(0.77)	(1.21)	(0.74)
Controls	No	Yes	No	Yes	Yes	Yes	Yes
Bank-Firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm-Time FE	No	No	No	No	Yes	Yes	Yes
Firm-Ind. spe.-Time FE	Yes	Yes	No	No	No	No	No
Firm-Loc. spe.-Time FE	No	No	Yes	Yes	No	No	No
N	9,660,033	9,660,033	10,104,051	10,104,051	6,386,256	11,660,269	11,014,677
R^2	0.94	0.94	0.93	0.94	0.94	0.91	0.93

Notes: This table reports the results of estimating specification (2.4) with the dependent variable equal to the log-level of long-term credit (maturity above one year) between bank b and firm f. The main explanatory variable is the interaction between orthogonalized Portuguese inflation $\pi_t^{\perp}$ and the bank's average maturity gap in the last 12 months $\bar{\sigma}_{b,t-1}$. Columns (1)–(2) interact the firm–time fixed effect with a dummy equal to one if the firm's industry belongs to the top 3 industries for the bank; columns (3)–(4) interact the firm–time fixed effect with a dummy equal to one the firm's municipality belongs to the top 10 municipalities for the bank. Column (5) restricts the sample to main banking entities of the banking group, column (6) uses a reduced set of controls (size, overdue ratio, liquidity, share of short loans, length of the bank-firm relationship), and column (7) additionally controls for changes in interest rates interacted with inflation $\bar{\sigma}_{b,t-1}$, to control explicitly for interest rate risk. Standard errors are double-clustered at the bank and time level. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

Table C.3.: Inflation effect on banks' credit supply: using ΔL_{bft}

	ΔL_{bft}			$L_{bf,t+3} - L_{bf,t-1}$	$L_{bf,t+12} - L_{bf,t-1}$
	(1)	(2)	(3)	(4)	(5)
Inflation $\times \bar{\sigma}_{b,t-1,t-12}$	-0.14	-0.16	-0.08	-0.21	-0.19
	(0.11)	(0.12)	(0.11)	(0.40)	(1.10)
Controls	No	No	Yes	Yes	Yes
Time FE	Yes	No	No	No	No
Bank FE	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	No	No	No	No
Firm-time FE	No	Yes	Yes	Yes	Yes
N	27,001,834	16,476,181	16,476,139	14,863,657	11,024,475
R^2	0.04	0.37	0.37	0.39	0.43

Notes: This table reports the results of estimating specification (2.4). The dependent variable is the log-difference of long-term credit (maturity above one year) between bank b and firm f. The main explanatory variable is the interaction between orthogonalized Portuguese inflation $\pi_t^{\perp}$ and the bank's average maturity gap in the last 12 months $\bar{\sigma}_{b,t-1}$. Controls include lagged credit, loan portfolio risk (share of overdue loans), relationship length, deposit ratio, loan ratio, short-term loan ratio, size, leverage, and liquidity, as well as their interactions with inflation. Firm–time and bank–firm fixed effects are included throughout. Standard errors are double-clustered at the bank and time level. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively.

Table C.4.: Inflation effect on banks' credit supply: using raw inflation π_t

	$h = 0$			$h = 3$	$h = 12$
	(1)	(2)	(3)	(4)	(5)
Inflation $\times \bar{\sigma}_{b,t-1,t-12}$	0.01	-0.00	-0.07	-0.25	-0.96
	(0.04)	(0.04)	(0.05)	(0.17)	(0.75)
Controls	No	No	Yes	Yes	Yes
Time FE	Yes	No	No	No	No
Bank FE	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	No	No	No	No
Firm-time FE	No	Yes	Yes	Yes	Yes
N	27,001,834	16,476,181	16,476,139	14,863,657	11,024,475
R^2	0.98	0.98	0.98	0.95	0.89

Notes: This table reports the results of estimating specification (2.4). The dependent variable is the log-level of long-term credit (maturity above one year) between bank b and firm f. The main explanatory variable is the interaction between raw Portuguese inflation π_t and the bank's average maturity gap in the last 12 months $\bar{\sigma}_{b,t-1}$. Controls include lagged credit, loan portfolio risk (share of overdue loans), relationship length, deposit ratio, loan ratio, short-term loan ratio, size, leverage, and liquidity, as well as their interactions with inflation. Firm–time and bank–firm fixed effects are included throughout. Standard errors are double-clustered at the bank and time level. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively.

2.E. Robustness: firm-level regressions

Table D.1.: Effect of inflation on firms' interest expenses

	$\frac{\text{Credit}_{f,t+h} - \text{Credit}_{f,t-1}}{\text{Assets}_{f,t-1}}$			
	$h = 0$			$h = 1$
	(1)	(2)	(3)	(4)
Inflation × $\theta_{b,t-1}$	-0.0936***	-0.0218***	0.0024	-0.0166
	(0.0026)	(0.0049)	(0.0044)	(0.0162)
Firm FE	Yes	Yes	Yes	Yes
Sector-Municipality-Year FE	Yes	Yes	Yes	Yes
Firm Demand	No	Yes	Yes	Yes
Controls	No	No	Yes	Yes
N	1,713,367	371,460	324,511	264,653
R sq.	0.28	0.44	0.55	0.67

Notes: This table reports the results of estimating specification (2.8) with the dependent variable equal to the annual change in total bank credit scaled by lagged assets. The key explanatory variable is the interaction between the Portuguese inflation surprise ε_t^{π} and the firm's loan-weighted exposure to banks' maturity gaps $\theta_{f,t-1}$. All regressions include firm fixed effects and sector–municipality–year fixed effects. Firm demand is proxied by the estimated firm–time demand component $\hat{a}_{ft}$ from the credit supply regression (2.4). Controls include lagged firm-level covariates (log assets, age, leverage, turnover growth, return on assets, capital intensity, liquidity, collateral ratio, export dummy). Standard errors are clustered at the firm level. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

Table D.2.: Effect of inflation on firms' interest expenses

	$\frac{\text{Investment}_{f,t+h} - \text{Investment}_{f,t-1}}{\text{Assets}_{f,t-1}}$			
	$h = 0$			$h = 1$
	(1)	(2)	(3)	(4)
Inflation × $\theta_{b,t-1}$	-0.0155***	-0.0253***	-0.0141***	-0.0205***
	(0.0013)	(0.0024)	(0.0024)	(0.0073)
Firm FE	Yes	Yes	Yes	Yes
Sector-Municipality-Year FE	Yes	Yes	Yes	Yes
Firm Demand	No	Yes	Yes	Yes
Controls	No	No	Yes	Yes
N	1,352,414	303,048	300,830	243,988
R sq.	0.16	0.27	0.34	0.42

Notes: This table reports the results of estimating specification (2.8) with the dependent variable equal to the annual change in firm's investment (capital expenditure) by lagged assets. The key explanatory variable is the interaction between the Portuguese inflation surprise ε_t^{π} and the firm's loan-weighted exposure to banks' maturity gaps $\theta_{f,t-1}$. All regressions include firm fixed effects and sector–municipality–year fixed effects. Firm demand is proxied by the estimated firm–time demand component $\hat{a}_{ft}$ from the credit supply regression (2.4). Controls include lagged firm-level covariates (log assets, age, leverage, turnover growth, return on assets, capital intensity, liquidity, collateral ratio, export dummy). Standard errors are clustered at the firm level. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

Table D.3.: Effect of inflation on firms' interest expenses

	$\frac{\text{Int. expenses}_{f,t+h} - \text{Int. expenses}_{f,t-1}}{\text{Assets}_{f,t-1}}$			
	$h=0$			$h=1$
	(1)	(2)	(3)	(4)
Inflation × $\theta_{f,t-1}$	0.0012***	0.0007***	0.0006***	0.0020***
	(0.0001)	(0.0001)	(0.0001)	(0.0004)
Firm FE	Yes	Yes	Yes	Yes
Sector-Municipality-Year FE	Yes	Yes	Yes	Yes
Firm Demand	No	Yes	Yes	Yes
Controls	No	No	Yes	Yes
N	1,246,436	346,785	304,042	245,928
R sq.	0.23	0.35	0.35	0.49

Notes: This table reports the results of estimating specification (2.8) with the dependent variable equal to the annual change in firms' interest expenses scaled by lagged assets. The key explanatory variable is the interaction between the Portuguese inflation surprise ε_t^π and the firm's loan-weighted exposure to banks' maturity gaps $\theta_{f,t-1}$. All regressions include firm fixed effects and sector–municipality–year fixed effects. Firm demand is proxied by the estimated firm–time demand component $\hat{a}_{ft}$ from the credit supply regression (2.4). Controls include lagged firm-level covariates (log assets, age, leverage, turnover growth, return on assets, capital intensity, liquidity, collateral ratio, export dummy). Standard errors are clustered at the firm level. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

Table D.4.: Effect of inflation on firms' revenues

	$\frac{\text{Revenues}_{f,t+h} - \text{Revenues}_{f,t-1}}{\text{Assets}_{f,t-1}}$			
	$h=0$			$h=1$
	(1)	(2)	(3)	(4)
Inflation × $\theta_{b,t-1}$	0.0606***	-0.0277***	-0.0168**	-0.0127
	(0.0055)	(0.0078)	(0.0077)	(0.0274)
Firm FE	Yes	Yes	Yes	Yes
Sector-Municipality-Year FE	Yes	Yes	Yes	Yes
Firm Demand	No	Yes	Yes	Yes
Controls	No	No	Yes	Yes
N	1,798,127	371,456	324,510	265,164
R sq.	0.34	0.44	0.47	0.61

Notes: This table reports the results of estimating specification (2.8) with the dependent variable equal to the annual change in firms' turnover scaled by lagged assets. The key explanatory variable is the interaction between the Portuguese inflation surprise ε_t^π and the firm's loan-weighted exposure to banks' maturity gaps $\theta_{f,t-1}$. All regressions include firm fixed effects and sector–municipality–year fixed effects. Firm demand is proxied by the estimated firm–time demand component $\hat{a}_{ft}$ from the credit supply regression (2.4). Controls include lagged firm-level covariates (log assets, age, leverage, turnover growth, return on assets, capital intensity, liquidity, collateral ratio, export dummy). Standard errors are clustered at the firm level. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

Table D.5.: Effect of inflation on firms' profits

	$\frac{\text{Profits}_{f,t+h} - \text{Profits}_{f,t-1}}{\text{Assets}_{f,t-1}}$			
	$h=0$			$h=1$
	(1)	(2)	(3)	(4)
Inflation × $\theta_{b,t-1}$	-0.0067***	-0.0087***	-0.0012	-0.0142***
	(0.0019)	(0.0024)	(0.0020)	(0.0055)
Firm FE	Yes	Yes	Yes	Yes
Sector-Municipality-Year FE	Yes	Yes	Yes	Yes
Firm Demand	No	Yes	Yes	Yes
Controls	No	No	Yes	Yes
N	1,798,163	371,460	324,511	265,165
R sq.	0.21	0.33	0.64	0.69

Notes: This table reports the results of estimating specification (2.8) with the dependent variable equal to the annual change in firms' profits scaled by lagged assets. The key explanatory variable is the interaction between the Portuguese inflation surprise ε_t^π and the firm's loan-weighted exposure to banks' maturity gaps $\theta_{f,t-1}$. All regressions include firm fixed effects and sector–municipality–year fixed effects. Firm demand is proxied by the estimated firm–time demand component $\hat{a}_{ft}$ from the credit supply regression (2.4). Controls include lagged firm-level covariates (log assets, age, leverage, turnover growth, return on assets, capital intensity, liquidity, collateral ratio, export dummy). Standard errors are clustered at the firm level. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

Table D.6.: Effect of inflation on firms' tangible investment

	$\frac{\text{Tang. Investment}_{f,t+h} - \text{Tang. Investment}_{f,t-1}}{\text{Assets}_{f,t-1}}$			
	$h=0$			$h=1$
	(1)	(2)	(3)	(4)
Inflation × $\theta_{b,t-1}$	-0.0153***	-0.0253***	-0.0146***	-0.0233***
	(0.0012)	(0.0023)	(0.0023)	(0.0070)
Firm FE	Yes	Yes	Yes	Yes
Sector-Municipality-Year FE	Yes	Yes	Yes	Yes
Firm Demand	No	Yes	Yes	Yes
Controls	No	No	Yes	Yes
N	1,352,380	303,048	300,830	243,988
R sq.	0.16	0.27	0.33	0.41

Notes: This table reports the results of estimating specification (2.8) with the dependent variable equal to the annual change in firms' tangible investment scaled by lagged assets. The key explanatory variable is the interaction between the Portuguese inflation surprise ε_t^π and the firm's loan-weighted exposure to banks' maturity gaps $\theta_{f,t-1}$. All regressions include firm fixed effects and sector–municipality–year fixed effects. Firm demand is proxied by the estimated firm–time demand component $\hat{a}_{ft}$ from the credit supply regression (2.4). Controls include lagged firm-level covariates (log assets, age, leverage, turnover growth, return on assets, capital intensity, liquidity, collateral ratio, export dummy). Standard errors are clustered at the firm level. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

2.F. Model derivations

2.F.1. Firm problem

This part borrows heavily from Kang and Pflueger (2015). I continue the description of the firm block of the model from Section 2.5.

C. Production

Firm i has a standard Cobb–Douglas production function with capital (K_t^i) and labor (N_t^i) inputs to produce output (Y_t^i).[54]

$$Y_t^i = (N_t^i)^{1-\alpha}(K_t^i)^{\alpha} \tag{2.23}$$

Firm i chooses labor to maximize single-period operating revenue, taking the aggregate wage as given. I assume the aggregate supply of labor is fixed at one, abstracting from unemployment. The equilibrium wage adjusts to clear the labor market. Capital depreciates at a constant rate δ, and I impose the resource constraint that total output equals aggregate consumption plus investment,[55]

$$\begin{aligned} K_{t+1} = & I_t + (1-\delta)K_t, \\ Y_t = & C_t + I_t. \end{aligned}$$

Young and old firms are heterogeneous in their capital stocks, but constant returns to scale imply that, for any firm, the gross real return on capital from time t to $t+1$ equals

$$R_{t+1}^K = \left[\alpha\left(\frac{1}{K_{t+1}}\right)^{1-\alpha} + (1-\delta)\right]. \tag{2.24}$$

D. Firm default

A firm's default decision depends on the initial level of debt, aggregate nominal shocks, and idiosyncratic real shocks.

Corporate debt promises a fixed nominal payment after two periods, when the firm pays a liquidating dividend. I denote logs by small letters throughout. All firms in cohort t are identical ex ante. Denote the initial log nominal face value of debt by $b_t^{\$}$ and initial log

[54] Differently from Kang and Pflueger (2015), who are interested in the correlation of the inflation shock with an aggregate real shock, I assume TFP is at steady state throughout, i.e., $z = 1$.

[55] I can formally define aggregate output, capital, labor, and investment at time t by integrating over all firms,

$$Y_t = \int_i Y_t^i \mathrm{d}i, \quad K_t = \int_i K_t^i \mathrm{d}i, \quad N_t = \int_i N_t^i \mathrm{d}i, \quad I_t = \int_i I_t^i \mathrm{d}i.$$

leverage adjusted for expected inflation by l_t. Then firms choose

$$l_t = b_t^{\$} - 2\pi_t - k_{t+1}. \tag{2.25}$$

Inflation persistence implies that the inflation shock in period $t+1$ enters twice into the real log liabilities of the old firm:

$$b_{t+2}^{real,old} = l_t + k_{t+1} - 2\varepsilon_{t+1}^{\pi} - \varepsilon_{t+2}^{\pi}. \tag{2.26}$$

Firm i in cohort t experiences independent idiosyncratic shocks to log capital at times $t+1$ and $t+2$, $a_{t+1}^{i,1}$ and $a_{t+2}^{i,2}$, with $a_{t+2}^{i,id} = a_{t+1}^{i,1} + a_{t+2}^{i,2}$. Using (2.24), the log real value of an old firm at the end of period $t+2$ equals

$$v_{t+2}^{i,old} = k_{t+1} + r_{t+1}^{K} + r_{t+2}^{K} + a_{t+2}^{i,id}. \tag{2.27}$$

Equity holders have the option to default on debt payments and receive a zero liquidating dividend. They optimally default if and only if the real value of the firm (2.27) is less than its real liabilities (2.26).[56] Firms with the most adverse idiosyncratic shocks default if

$$a_{t+2}^{i,id} < \underbrace{l_t - 2\varepsilon_{t+1}^{\pi} - \varepsilon_{t+2}^{\pi} - r_{t+1}^{K} - r_{t+2}^{K}}_{\text{survival threshold } a_{t+2}^{*}}. \tag{2.28}$$

Equation (2.28) formalizes the debt–inflation channel: positive inflation shocks ε_{t+1}^{π} and ε_{t+2}^{π} lower the survival threshold a_{t+2}^{*} and, in turn, defaults.

E. Corporate loan prices

Let $H(a_{t+2}^{*})$ and $\Omega(a_{t+2}^{*})$ denote the time $t+2$ default probability and the average defaulted firm value conditional on the survival threshold a_{t+2}^{*}, respectively. Let $G(a_{t+1}^{*}, a_t^{i,1})$ and $W(a_{t+1}^{*}, a_t^{i,1})$ denote the time $t+1$ default probability and the average defaulted firm value of a cohort $t-1$ firm conditional on both the survival threshold a_{t+1}^{*} and the firm's first-period idiosyncratic shock $a_t^{i,1}$.

$$H\left(a_{t+2}^{*}\right) = \mathbb{P}\left(a_{t+2}^{i,id} < a_{t+2}^{*}\right), \quad \Omega\left(a_{t+2}^{*}\right) = \mathbb{E}\left(\exp\left(a_{t+2}^{i,id}\right)\mathbb{I}\left(a_{t+2}^{i,id} < a_{t+2}^{*}\right)\right),$$
$$G\left(a_{t+1}^{*}, a_t^{i,1}\right) = \mathbb{P}\left(a_{t+1}^{i,id} < a_{t+1}^{*} \mid a_t^{i,1}\right), \quad W\left(a_{t+1}^{*}, a_t^{i,1}\right) = \mathbb{E}\left(\exp\left(a_{t+1}^{i,id}\right)\mathbb{I}\left(a_{t+1}^{i,id} < a_{t+1}^{*}\right) \mid a_t^{i,1}\right).$$

where $\mathbb{I}$ denotes the indicator function. The prices for a new corporate loan and a seasoned

[56] The firm never finds it optimal to default in its intermediate period because no debt payments are due then.

corporate loan at time t then equal

$$q_t^{\text{corp,new}} = \mathbb{E}_t\left[M_{t,t+2}^{\$}\Big(1 - \underbrace{H\left(a_{t+2}^*\right)}_{\text{Default Rate}} + \underbrace{\theta\frac{\Omega\left(a_{t+2}^*\right)}{\exp\left(a_{t+2}^*\right)}}_{\text{Expected Recovery Payment}}\Big)\right], \tag{2.29}$$

$$q_t^{i,\text{ seas}} = \mathbb{E}_t\left[M_{t,t+1}^{\$}\Big(1 - \underbrace{G\left(a_{t+1}^*, a_t^{i,1}\right)}_{\text{Cond. Def. Rate}} + \underbrace{\theta\frac{W\left(a_{t+1}^*, a_t^{i,1}\right)}{\exp\left(a_{t+1}^*\right)}}_{\text{Cond. Recovery}}\Big)\right], \tag{2.30}$$

where $\theta < 1$ is the fraction recovered by the bank when the firm is in bankruptcy.[57]

F. Capital structure choice

Firms choose leverage according to a standard trade-off view of capital structure. Following Gourio (2013), firms receive benefits $\chi > 1$ for each dollar of debt issued. Equity holders of cohort t firms choose capital K_{t+1} and nominal liabilities $B_t^{\$}$ subject to the budget constraint

$$K_{t+1} = \underbrace{S_t}_{\text{Value of new equity}} + \chi \times \underbrace{q_t^{\text{corp,new}}}_{\text{New nominal debt price}} \times B_t^{\$}. \tag{2.31}$$

Higher χ increases the incentive to raise leverage.[58] Equity holders trade off the benefits of debt with expected bankruptcy costs. I assume the bank recovers only a constant fraction $\theta < 1$ of firm value in bankruptcy. A lower recovery rate reduces the incentive to lever up. There exists an interior optimal leverage ratio if bankruptcy costs are sufficiently large relative to debt benefits. Following Gourio (2013), I formally assume $\theta\chi < 1$.

Define the marginal probability of default as $h(a_{t+2}^*) = H'(a_{t+2}^*)$. Solving the first-order condition for leverage (see Appendix 2.F for derivations) yields

$$0 = \underbrace{-\chi(1-\theta)\mathbb{E}_t\left(M_{t,t+2}^{\$} h\left(a_{t+2}^*\right)\right)}_{\text{Marginal Bankruptcy Cost (MBC)}} + \underbrace{(\chi - 1)\mathbb{E}_t\left(M_{t,t+2}^{\$}\left(1 - H\left(a_{t+2}^*\right)\right)\right)}_{\text{Marginal Benefit of Debt (MBD)}}. \tag{2.32}$$

Equation (2.32) shows that equity holders equate the marginal benefit of issuing an additional dollar of debt with the marginal increase in bankruptcy costs.

Prediction 1. After a positive inflation shock, firms increase their credit demand.

A positive inflation shock ε_{t+1}^{π} affects the first-order condition (2.32) through the default threshold a_{t+2}^*: higher inflation lowers a_{t+2}^* and the probability of default $H\left(a_{t+2}^*\right)$, so the marginal benefit of debt increases. Firms therefore wish to raise leverage (i.e., issue more debt) after a positive inflation shock. This implication motivates the need in the empirical

[57] All derivations are presented in Appendix 2.F.

[58] There is a debate over whether tax benefits are sufficiently large to explain observed leverage ratios (Graham, 2000; Almeida and Philippon, 2007). χ is interpreted to include broader benefits and costs of debt, such as constraining managers from empire-building and reducing informational asymmetries (Meckling and Jensen, 1976; Jensen, 1986).

analysis of Section 2.3 and 2.4 to control for credit demand when assessing the impact of inflation on credit supply.

2.F.2. Derivations for firm problem

Optimal choice of labor

Firm i chooses labor optimally to maximize single period operating revenue, while taking the aggregate wage W_t as given:

$$N_t^i = \arg\max_{N_t^i} \left\{ \underbrace{Y_t^i - W_t N_t^i}_{\text{Operating Revenue}} \right\}. \tag{2.33}$$

From the firm's single period optimization we obtain the first-order condition with respect to labor:

$$(1-\alpha)\left(\frac{K_t^i}{N_t^i}\right)^{\alpha} = W_t. \tag{2.34}$$

The capital to labor ratio is constant across firms and equal to K_t. Substituting back into operating revenue gives firm i's one-period equilibrium revenue as $\alpha K_t^i \left(\frac{1}{K_t}\right)^{1-\alpha}$. The expression for the equilibrium return on capital follows as:

$$R_{t+1}^K = \left[\alpha\left(\frac{1}{K_{t+1}}\right)^{1-\alpha} + (1-\delta)\right]. \tag{2.35}$$

Debt price

The time $t+2$ real cash flow of a corporate bond issued by firm i at time t is:

$$\frac{\left(1-\mathbb{I}\left\{a_{t+2}^{i,id} < a_{t+2}^*\right\}\right)}{\exp\left(2\pi_t + 2\varepsilon_{t+1}^{\pi} + \varepsilon_{t+2}^{\pi}\right)} + \theta\frac{K_{t+1}^y}{B_t^\$} R_{t+1}^K R_{t+2}^K \exp\left(a_{t+2}^{i,id}\right)\mathbb{I}\left\{a_{t+2}^{i,id} < a_{t+2}^*\right\} \tag{2.36}$$

The time t price of the bond is given by the expected stochastic discounted value of real cash flows:

$$q_t^{\text{corp,new}} = \mathbb{E}_t\left[M_{t,t+2}^\$ \left(1 - H\left(a_{t+2}^*\right)\right)\right] + \theta\mathbb{E}_t\left[M_{t,t+2}\frac{K_{t+1}^y}{B_t^\$} R_{t+1}^K R_{t+2}^K \Omega\left(a_{t+2}^*\right)\right] \tag{2.37}$$

The expression for the survival threshold then implies:

$$q_t^{\text{corp,new}} = \mathbb{E}_t\left[M^{\$}_{t,t+2}\Big(1-\underbrace{H\left(a^*_{t+2}\right)}_{\text{Default Rate}}+\underbrace{\theta\frac{\Omega\left(a^*_{t+2}\right)}{\exp\left(a^*_{t+2}\right)}}_{\text{Recovery Rate}}\Big)\right] \tag{2.38}$$

as in (2.29) in the main text.

First-order conditions

Equity holders maximize:

$$\mathbb{E}_t\left[M_{t,t+2}\max\left(V^{i,old}_{t+2}-B^{\$}_t\exp\left(-2\pi_t-2\varepsilon^{\pi}_{t+1}-\varepsilon^{\pi}_{t+2}\right),0\right)\right]-S_t \tag{2.39}$$

subject to:

$$v^{i,old}_{t+2} = k^y_{t+1}+r^K_{t+1}+r^K_{t+2}+a^{i,id}_{t+2}, \tag{2.40}$$

$$K^y_{t+1} = S_t+\chi q^{corp,new}_t B^{\$}_t. \tag{2.41}$$

Given constant returns to scale and no equity issuance costs, the net equity value (2.39) will equal zero in equilibrium, reflecting free entry. Substituting the constraints and (2.38) into (2.39) we can rewrite the firm's problem as maximizing:

$$\exp\left(2\pi_t\right)K^y_{t+1}L_t\mathbb{E}_t\left[M^{\$}_{t,t+2}\begin{pmatrix}\exp\left(-a^*_{t+2}\right)+\left(\chi-1\right)\left(1-H\left(a^*_{t+2}\right)\right)\\ +(\chi\theta-1)\Omega\left(a^*_{t+2}\right)\exp\left(-a^*_{t+2}\right)\end{pmatrix}\right]-K^y_{t+1} \tag{2.42}$$

Differentiating (2.42) with respect to leverage L_t while holding constant the level of capital K_{t+1} gives:

$$0=\left(1+\frac{\partial}{\partial a^*_{t+2}}\right)\mathbb{E}_t\left[M^{\$}_{t,t+2}\left(\exp(-a^*_{t+2})+\left(\chi-1\right)\left(1-H(a^*_{t+2})\right)+\left(\chi\theta-1\right)\Omega(a^*_{t+2})\exp(-a^*_{t+2})\right)\right] \tag{2.43}$$

Using $\frac{\partial}{\partial a^*_{t+2}}\Omega(a^*_{t+2})=\exp(a^*_{t+2})h(a^*_{t+2})$ gives the first-order condition with respect to leverage:

$$0=-\chi(1-\theta)\,\mathbb{E}_t\left(M^{\$}_{t,t+2}h(a^*_{t+2})\right)+\left(\chi-1\right)\mathbb{E}_t\left(M^{\$}_{t,t+2}\left(1-H(a^*_{t+2})\right)\right) \tag{2.44}$$

2.F.3. Derivations for the bank problem

Proof of proposition 2

By the chain rule and (2.20),

$$\frac{\partial B^{\$,\,\mathrm{MM}}_{t+1}}{\partial \varepsilon^{\pi}_{t+1}} = \underbrace{\frac{(1-\rho)}{s_{\min}(q)}\,\frac{\partial \bar{E}_{t+1}}{\partial E^{pre}}\,\frac{\partial E^{pre,\,\mathrm{MM}}_{t+1}}{\partial q}\,\frac{\partial q}{\partial \varepsilon^{\pi}_{t+1}}}_{\text{issuance-capacity channel}} + \underbrace{(1-\rho)\,\bar{E}_{t+1}\,\frac{\partial}{\partial q}\Big(\frac{1}{s_{\min}(q)}\Big)\,\frac{\partial q}{\partial \varepsilon^{\pi}_{t+1}}}_{\text{capital-share (denominator) channel}}. \tag{2.45}$$

Using $s_{\min}(q) = 1 - q(1-\gamma)$ and $E^{pre,\,\mathrm{MM}}_{t+1}$ from (2.22),

$$\frac{\partial}{\partial q}\Big(\frac{1}{s_{\min}(q)}\Big) = \frac{1-\gamma}{[1-q(1-\gamma)]^2} \;>\; 0, \qquad \frac{\partial E^{pre,\,\mathrm{MM}}_{t+1}}{\partial q} = \frac{B^{\$}_{t}}{P_{t+1}}\cdot\frac{1-\gamma}{q^2} \;>\; 0, \tag{2.46}$$

and $\frac{\partial q}{\partial \varepsilon^{\pi}_{t+1}} < 0$. Hence both terms in (2.45) are *negative*: the inflation surprise (i) worsens the equity-per-euro requirement via $s_{\min}$, and (ii) directly depresses issuance capacity by lowering $E^{\mathrm{pre},\,\mathrm{MM}}_{t+1}$, which tightens $\bar{E}_{t+1}$ when $\partial \bar{E}/\partial E^{\mathrm{pre}} > 0$.

For the non-mismatched bank,

$$\frac{\partial B^{\$,\,\mathrm{NM}}_{t+1}}{\partial \varepsilon^{\pi}_{t+1}} = \underbrace{\frac{1}{s_{\min}(q)}\,\frac{\partial\Big((1-\rho)\bar{E}\Big)}{\partial E^{pre}}\,\frac{\partial E^{pre,\,\mathrm{NM}}_{t+1}}{\partial q}\,\frac{\partial q}{\partial \varepsilon^{\pi}_{t+1}}}_{=0\ \ \text{since } E^{pre,\,\mathrm{NM}} \text{ does not depend on } q} + \underbrace{(X^{\mathrm{int}}_{t+1} + (1-\rho)\bar{E})\,\frac{\partial}{\partial q}\Big(\frac{1}{s_{\min}(q)}\Big)\,\frac{\partial q}{\partial \varepsilon^{\pi}_{t+1}}}_{<0}. \tag{2.47}$$

Thus the non-mismatched bank is hit only through the denominator channel; it *avoids* the extra issuance-capacity deterioration term that is present under mismatch. This establishes the ordering stated in Proposition 2. □

Other derivations

(i) Derivative of E^{pre} with respect to q. From (2.22),

$$\frac{\partial E^{pre,\,\mathrm{MM}}_{t+1}}{\partial q} = \frac{B^{\$}_{t}}{P_{t+1}}\cdot\frac{1-\gamma}{q^2} \;>\; 0, \qquad \frac{\partial E^{pre,\,\mathrm{NM}}_{t+1}}{\partial q} = 0.$$

(ii) Derivative of $1/s_{\min}(q)$. With $s_{\min}(q) = 1 - q(1-\gamma)$,

$$\frac{\partial}{\partial q}\Big(\frac{1}{s_{\min}(q)}\Big) = \frac{1-\gamma}{[1-q(1-\gamma)]^2} \;>\; 0.$$

(iii) Chain rule with issuance capacity. Assuming $\bar{E} = \bar{E}(E^{pre}_{t+1})$,

$$\frac{\partial \bar{E}}{\partial \varepsilon^{\pi}_{t+1}} = \frac{\partial \bar{E}}{\partial E^{pre}}\,\frac{\partial E^{pre}_{t+1}}{\partial q}\,\frac{\partial q}{\partial \varepsilon^{\pi}_{t+1}}.$$

(iv) Putting the pieces together. Plugging (i)–(iii) yields the claimed ordering in the proof of proposition 2:

$$\frac{\partial B^{\$,\mathrm{MM}}_{t+1}}{\partial \varepsilon^{\pi}_{t+1}} = \underbrace{\frac{(1-\rho)}{s_{\min}} \frac{\partial \bar{E}}{\partial E^{\mathrm{pre}}} \frac{B^{\$}_{t}}{P_{t+1}} \frac{1-\gamma}{q^2} \frac{\partial q}{\partial \varepsilon^{\pi}_{t+1}}}_{\text{extra negative term (mismatch only)}} + [\text{common negative denominator term}],$$

while

$$\frac{\partial B^{\$,\mathrm{NM}}_{t+1}}{\partial \varepsilon^{\pi}_{t+1}} = [\text{common negative denominator term}].$$

Since $\partial \bar{E}/\partial E^{\mathrm{pre}} \geq 0$ and $\partial q/\partial \varepsilon^{\pi}_{t+1} < 0$, the extra term is strictly negative, implying a larger drop in credit for the mismatched bank.

3. Who Gains from Capital Gains? The Realization Channel

Joint with Frederik Bennhoff, Moritz Kuhn, and Florian Scheuer

3.1. Introduction

Realized capital gains are a large and highly cyclical component of household income. Figure 3.1 reports the U.S. aggregate from 1997 to 2012 by source: the flow ranges between roughly $400 and $900 billion per year in the run-up to the Global Financial Crisis, collapses by more than ninety percent in 2009, and partially recovers thereafter, with stock and mutual fund sales accounting for the bulk of both the level and the variation. The series moves closely with GDP growth and at a scale comparable to other major household income aggregates. Yet the household-level mechanics of this flow remain poorly understood: which households realize accrued gains, on which positions, in response to what, and what they do with the proceeds once a sale closes. These questions matter because realized gains, not paper gains, are the part of capital income that households actually convert into spending power.

Figure 3.1.: U.S. realized capital gains and GDP growth, 1997–2012.

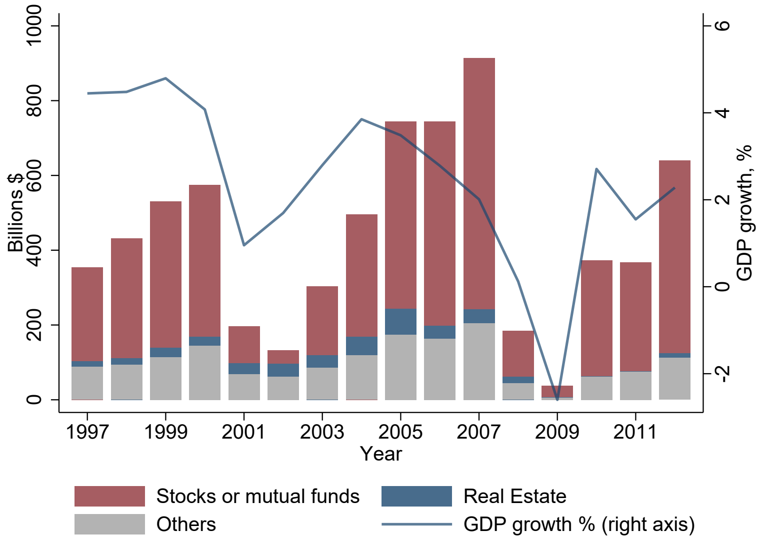

Notes: Aggregate U.S. realized capital gains in billions of dollars by asset class (stocks and mutual funds; real estate; other). Real GDP growth on the right axis. Source: IRS Statistics of Income, Sales of Capital Assets tables, and U.S. Bureau of Economic Analysis.

This paper documents the realization channel at the household level using three decades of panel data from the DNB Household Survey (DHS) matched to security-level prices. The Dutch setting is particularly clean for this question because realized capital gains are not taxed in the Netherlands: lock-in on accrued gains and year-end tax-loss harvesting are absent, so the patterns we report must originate in non-tax behavioral or liquidity channels rather than in tax planning. We organize the household-level evidence into four facts. First, the realization channel is housing-dominated by size and stock-dominated by frequency, and its cross-sectional incidence inherits the asset-holding distribution more than any behavioral skew that may change along the income distribution. Second, conditional on selling a share, households liquidate winners over losers in a way that survives a position-level instrumental variable. Third, life events matter more for realization behavior than actual market price dynamics, as the selection on winners bends sharply around divorce – and less clearly for marriage, unemployment and retirement. Fourth, the marginal use of proceeds from house or stocks selling depends strongly on housing tenure: homeowners spend stock realizations and channel housing realizations into trade-ups and additional debt, while renters cycle stock realizations back into the financial portfolio – a behavior consistent with how households' balance sheet and characteristics change dramatically spending patterns (Kaplan and Violante, 2014; Fagereng, Holm, and Natvik, 2021).

Section 3.2 describes the data and the empirical setting. The DNB Household Survey (DHS) is an annual panel of about two thousand Dutch households in which respondents report their portfolio holdings *by name*—up to ten directly held stock positions and five mutual fund positions per household-year, together with the year-end euro value of each.

We aggregate these holdings into a household–position–year panel covering 2001–2023, manually match security names to Datastream end-of-year prices, and recover implied share counts and a household-specific average purchasing price for each position; realized stock gains then follow from the standard average-cost rule whenever shares change hands. Housing is recorded at the household level (purchase year and price, current market value, and—on exit from owner-occupation—year and price of sale), from which we construct realized housing gains. Two features of the institutional setting make the Dutch case useful for this question. The Netherlands does not tax realized capital gains, so lock-in on accrued gains and year-end tax-loss harvesting are absent. And homeownership is broadly distributed (roughly sixty percent of households), so the housing realization channel is not concentrated at the top of the wealth distribution as it would be for stocks.

Section 3.3 measures the size and the cross-sectional shape of the realization channel. Stocks and housing realize on opposite frequency–size schedules: about half of stockholding households sell at least one position in a given year, in amounts averaging €8,000, while only about one homeowner in a hundred sells, in amounts averaging €255,000—more than thirty times larger. Aggregating to a national series, Dutch realized gains add up to a few percent of GDP and are housing-dominated by value. To separate what households *own* from how aggressively they *convert* paper gains into cash, we decompose the ratio of realized gains to income into a stock-of-capital term (A/Y), an embedded-gain term (G/A), and a behavioral term—the *average propensity to realize* (APR), the share of accrued gains that the household converts into cash within the year. The income gradient in realized gains comes significantly from A/Y, while the APR is more flat across the distribution. Cross-sectional differences in realized volumes therefore reflect what households own more than how aggressively they convert paper gains into cash.

Section 3.4 sharpens the question to which positions households actually liquidate, conditional on selling. We work at the position level and define a position's *relative return* as its annual return measured against the value-weighted return on the rest of the same household's portfolio—a within-household-year benchmark that nets out aggregate market movements. Two facts emerge. The return distribution of sold positions sits to the right of that of held positions across the full support: sold positions earned higher returns than retained ones, the empirical signature of disposition behavior originally identified by Odean (1998). And in a linear-probability model with household-year fixed effects a position's sale probability rises sharply with its relative return, a pattern that survives a position-level instrument in the spirit of Di Maggio, Kermani, and Majlesi (2020): a ten-log-point increase in relative return raises the sale probability by about seventeen percentage points against a baseline of forty-nine percent. Because the Netherlands taxes neither realized gains nor losses, this selection cannot reflect lock-in or tax-loss harvesting; it has to come from non-tax channels such as belief inertia, attachment to losing positions, or rule-of-thumb rebalancing.

Section 3.5 asks whether this selection margin is a stable preference of households or one that bends under pressure. We compare the return distribution of sold positions with that of held positions inside and outside windows around four life events—unemployment, divorce, marriage, and retirement. Divorce is the cleanest case and the most informative: the disposition pattern not only disappears, it inverts. When a joint household portfolio is dissolved, households end up selling positions that on average underperform what they retain—they cannot cherry-pick winners and shed their losers along with the rest. Marriage, unemployment, and retirement point in the same direction but the in-window samples are too small to read with confidence. The lesson is that selection on winners is a discretionary margin, not a hard preference: it survives ordinary liquidity needs and breaks only under shocks sharp enough to force the joint liquidation of the household's portfolio.

Section 3.6 traces where a euro of proceeds goes once a sale closes. We assemble a household budget identity in which every euro of source—labor income, dividends, rental income, stock sales, housing sales—must equal the sum of uses—stock investment, housing investment, debt service, consumption, and the change in liquid balances. Estimating the marginal allocation of sources to uses separately for owners and renters then asks the natural question: when one source rises by a euro, where does that euro appear on the use side? The same realized euro funds different things in different parts of the housing-tenure distribution. For homeowners, a marginal euro of stock realization splits roughly between contemporaneous consumption and net cash flows to creditors, with effectively zero reinvestment into financial assets; a marginal euro of housing realization is associated with additional housing investment and *additional* mortgage borrowing—the signature of a trade-up rather than a draw-down. For renters, a marginal euro of stock realization is reinvested approximately one-for-one back into the financial portfolio. The owner–renter contrast on the same source—stocks—is the central finding of the section. It is consistent with the two-asset (liquid–illiquid) interpretation of MPC heterogeneity in Kaplan and Violante (2014) and Fagereng, Holm, and Natvik (2021), in which housing acts as the illiquid asset and the financial portfolio as the liquid buffer drawn down at the margin.

Related literature. The paper sits at the intersection of four strands of research. On the distributional and aggregate side, work using U.S. tax records and harmonized national accounts has documented that realized capital gains are large, cyclical, and concentrated at the top of the distribution (Burman and Ricoy, 1997; Robbins, 2018; Piketty, Saez, and Zucman, 2018; Larrimore et al., 2021). The closest contemporary analogue is Msall and Næss (2025), who use Norwegian administrative data to document never-realized gains and an APR analogue. We differ by working in a setting without capital gains taxation, by recovering a causal slope on the selection margin, and by tracing where the proceeds go once a sale closes. On position-level realization behavior, our selection-on-winners results extend the disposition-effect literature from Odean (1998) and Odean (1999) to Meyer and Pagel (2022) by isolating the pattern in a no-capital gains taxation (CGT) environment and by instrumenting relative return at the position level following Di Maggio, Kermani, and Majlesi (2020); the role of life events as a moderator of the selection margin is, to our knowledge, new and relates to the literature on family-structure shocks and household saving (Borella et al., 2025). The flow-of-funds analysis connects to causal estimates of the consumption response to wealth in Di Maggio, Kermani, and Majlesi (2020) and Meeuwis (2022) and to the two-asset MPC interpretation of Kaplan and Violante (2014) and Fagereng, Holm, and Natvik (2021); we contribute tenure-dependent marginal allocations at the realization level rather than at the level of paper wealth changes, which sharpens the empirical content of these models. The institutional background draws on Aguiar, Moll, and Scheuer (2024), who articulate the theoretical case for realization-based capital-income taxation and motivate the empirical importance of understanding realization behavior independently of tax incentives. The Dutch panel and position-level reconstruction build on Alessie, Hochgürtel, and Soest (2004), Gaudecker (2015), and Adam and Wagner (2023). Taken together, the patterns we document point to three disciplines for quantitative lifecycle models of household saving and portfolio choice: the realization decision itself, not the accrued gain, should be the operative decision margin; housing tenure must enter the state space, because the same realized euro from the same asset is allocated very differently by owners and renters; and idiosyncratic non-financial shocks—divorce most clearly—move realizations more than returns do. These results call for embedding an explicit realization decision in a two-asset lifecycle model with housing tenure and family-structure shocks calibrated to the moments documented here.

Roadmap. Section 3.2 describes the data and the position-level construction of realized gains. Section 3.3 measures the size and the cross-sectional shape of the realization channel. Section 3.4 documents and instruments selection on winners at the position level. Section 3.5 tests whether selection bends under life-event-driven liquidity shocks. Section 3.6 estimates the marginal allocation of realized proceeds by housing tenure. Section 3.7 concludes.

3.2. Data and Measurement

The analysis combines a long Dutch household panel with a position-level reconstruction of stock and mutual fund holdings matched to Datastream price series. Section 3.2.1 describes the panel and its lifecycle representativeness; Section 3.2.2 the position-level portfolio reconstruction and cost-basis rule; Section 3.2.3 the resulting realization variables.

3.2.1. The Dutch Household Survey panel

The DNB Household Survey (DHS) is an annual household panel collected by CentERdata and launched in 1993[1]. Each wave consists of six questionnaires covering household composition, employment, accommodation and mortgages, health and income, assets

[1] Data can be accessed at www.dhsdata.nl

and liabilities, and economic and psychological concepts. We use waves covering years 2001–2023; item names for stocks and funds are recorded as free-text only from 2001 onward, while earlier waves used a closed list of Dutch issuers whose accompanying free-text supplement is no longer archived (Adam and Wagner, 2023), which fixes the start of our sample window. All monetary variables are reported in euros, and households are linked across waves by a household identifier.

About 2,000 households participate in each wave, and the median household appears in four waves (Appendix 3.D, Figure 3.13). Attrition is moderate and not concentrated by tenure, and the survey is representative of the Dutch population through the use of sampling weights. Respondents without internet access are supplied with a set-top box and reimbursed for connection costs, so panel membership is not selected on internet adoption. Table 3.1 reports the pooled distribution of household resources, portfolio composition, and leverage across the sample; the corresponding age-group and year-by-year cross-sections are in Appendix 3.D, Tables 3.20–3.21. The median household earns about €35,000 in gross income and holds €64,000 in net wealth, with mean wealth (€168,000). Deposits are near-universal at 84 %, while only about 21 % of households participate in stocks or mutual funds; stocks and housing together account for roughly two thirds of non-business wealth at the median.

The Dutch lifecycle profile of asset participation tracks the U.S. pattern. Stock and mutual fund participation in the DHS is hump-shaped over the lifecycle and stays slightly below the Survey of Consumer Finances (SCF) analogue at every age (Figure 3.2), with a more pronounced stock market exit after retirement age. Over time, both the US and the Netherlands observed a drastic decline in stock market participation for directly held stocks starting in 2002 (after the dotcom bubble) and continuing after the 2008 financial crisis. Homeownership broadly tracks the SCF profile through middle age and remains lower than in the U.S. at older ages, consistent with Dutch households divesting housing later in life (Appendix 3.D, Figure 3.14).

Table 3.1.: Summary statistics

	N	Mean	SD	p10	p50	p90
Gross income (EUR)	33,779	40,364.6	29,496.1	6,000.0	35,000.0	79,257.0
Net wealth (EUR)	42,015	167,577.6	243,400.7	0.0	64,014.0	459,453.0
Risky asset participation (%)	37,953	21.2	40.9	0.0	0.0	100.0
Direct stock participation (%)	37,953	10.6	30.8	0.0	0.0	100.0
Mutual fund participation (%)	37,953	16.0	36.7	0.0	0.0	100.0
Other instruments participation (%)	39,461	6.8	25.1	0.0	0.0	0.0
Bond participation (%)	37,953	3.5	18.4	0.0	0.0	0.0
Deposit participation (%)	37,953	84.4	36.2	0.0	100.0	100.0
Insurance product participation (%)	37,953	21.8	41.3	0.0	0.0	100.0
Any financial instrument (%)	39,461	35.3	47.8	0.0	0.0	100.0
Total debt (EUR)	42,015	63,157.2	103,659.9	0.0	875.0	210,000.0
Leverage: debt/assets (%)	39,879	159.1	3,244.9	0.0	5.9	90.1
Leverage: debt/income (%)	30,900	294.9	4,314.4	0.0	17.4	499.2
Value: stocks and housing (EUR)	42,014	182,204.3	260,724.2	0.0	123,236.1	452,041.9
Business assets (EUR)	37,953	4,592.0	57,083.2	0.0	0.0	0.0
Non-business assets (EUR)	42,015	229,090.1	277,279.4	830.0	162,704.9	568,533.9
Stocks-housing share of non-bus. wealth (%)	39,876	48.9	61.9	0.0	65.7	97.1
Employer-sponsored share (%)	37,122	1.3	8.3	0.0	0.0	1.0
Risky share of financial wealth (%)	36,938	8.6	38.2	0.0	0.0	37.0

Notes: Pooled across all sample years. Mean, standard deviation, and percentiles computed with time-averaged household weights; N is the unweighted number of household-years. Sample restricted to ages 20–85. *Gross income* and *Net wealth:* household totals, EUR, winsorized at 1/99. *Risky asset participation:* holds stocks or mutual funds. *Other instruments:* growth funds, bonds, options, warrants, etc. *Deposits:* savings accounts and certificates (excludes checking). *Insurance products:* annuity or endowment policies. *Any financial instrument:* excludes deposits. *Debt:* consumer plus mortgage debt, winsorized at 1/99. *Stocks and housing:* directly held stocks, mutual funds, and housing value. *Non-business assets:* total assets minus business holdings. *Stock–housing share:* stocks plus housing over non-business wealth. *Employer-sponsored share:* employer savings plans over non-business wealth. *Risky financial share:* stocks plus funds over financial wealth.

Figure 3.2.: Stock and mutual fund participation: Netherlands and United States.

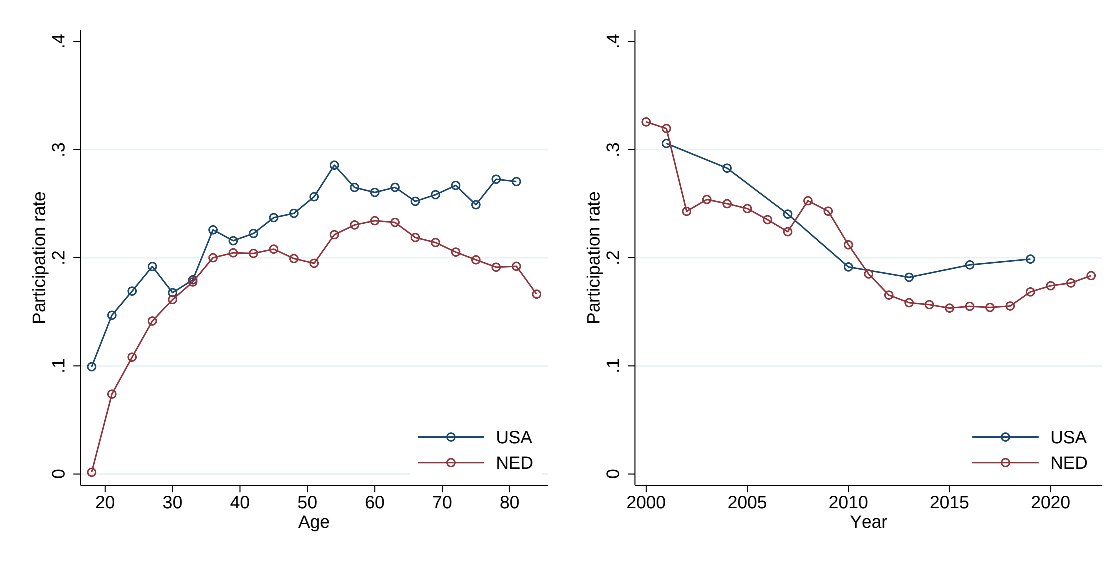

Notes: Lifecycle (left) and time-series (right) profiles of stock and mutual fund participation. Dutch profile estimated from the DHS; U.S. profile from the SCF. The corresponding homeownership profiles are reported in Appendix 3.D, Figure 3.14.

3.2.2. Reconstructing position-level portfolios

Households can report up to ten directly held stock positions and up to five mutual fund positions *by name*, with the year-end euro value of each holding. Most respondents report fewer than the allowed maximum, consistent with the underdiversification documented in retail portfolios (Grinblatt and Keloharju, 2000; Barber and Odean, 2000; Barber and Odean, 2001). We aggregate the raw holdings files into a household–position–year panel and manually match security names to Datastream end-of-year price series through a name-matching pipeline; this position-level reconstruction follows Gaudecker (2015)'s use of the DHS holdings names and parallels Calvet, Campbell, and Sodini (2007) and Calvet, Campbell, and Sodini (2009) on Swedish administrative data. Approximately 80 % of household-years that report any direct stockholding are matched to at least one priced position – the remaining are either not available stock names in Datastream or unidentifiable due to severe mistakes in the transcriptions. For each matched position j, the reported value V_{ijt} and price P_{jt} pin down the implicit number of shares, $Q_{ijt} = V_{ijt}/P_{jt}$. Appendix 3.D, Figure 3.15 reports the distribution of the number of positions per household-year. The median stockowner holds 2 (mean 2.8) distinct shares and 1 (mean 1.8) mutual funds.

Realized gains require both a price change and a change in shares held. We maintain a household-specific average purchasing price (APP) per position, $\bar{P}_{ijt}$, following the standard average-cost convention (e.g., Constantinides, 1983), updating only on net purchases:

$$\bar{P}_{ijt} = \begin{cases} [Q_{ijt-1}\,\bar{P}_{ijt-1} + (Q_{ijt} - Q_{ijt-1})\,P_{jt}]/Q_{ijt} & \text{if } Q_{ijt} > Q_{ijt-1}, \\ \bar{P}_{ijt-1} & \text{if } Q_{ijt} \leq Q_{ijt-1}. \end{cases} \tag{3.1}$$

Households are followed across waves so $\bar{P}_{ijt}$ accumulates over the panel, and non-response gaps are bridged with a carry-forward rule. All prices are Datastream series converted to euros at the corresponding-frequency exchange rate. End-of-year prices value the household's holdings Q_{ijt} and pin down the level of stock gains; for the APP law of motion, the unobserved within-year transaction date is handled by valuing net purchases at the mid-year (June) price. As a robustness variant we also compute an annual-average-price construction in the spirit of Fagereng, Guiso, et al. (2020)'s modified-Dietz measure, which combines a held-portion gain with a within-year-mean trade adjustment and reports a single annual flow rather than separate realized and unrealized components.

Housing has a different reporting structure. Each household reports tenure status, current market value, the year of purchase, the original purchase price, and—for households who exit owner-occupation—the year and price of sale. The questionnaire elicits a cost basis only for the primary residence, so housing gains are necessarily measured at the household–year level. With the APP $\bar{P}_{ijt}$ for stocks and the reported purchase price for housing, the formal definitions of realized and unrealized gains follow in §3.2.3.

A cost-basis caveat deserves emphasis. Households entering the panel with positive stock positions have an unknown initial APP; we treat such position-years as left-censored and verify robustness to imputed first-year APPs.[2]

3.2.3. Definitions: realized gains, APR, and relative return

Realized gains R are the trade-induced gains and losses booked when shares change hands; *unrealized gains* U are the paper gains on currently held positions valued at year-end prices. For stocks, both objects are constructed at the position–year level from the average purchasing price $\bar{P}_{ijt}$ defined in Equation (3.1):

$$R^S_{it} = \sum_j (P_{jt} - \bar{P}_{ijt})\,\max\{0,\, Q_{ijt-1} - Q_{ijt}\}, \tag{3.2}$$

$$U^S_{it} = \sum_j (P_{jt} - \bar{P}_{ijt})\,Q_{ijt}. \tag{3.3}$$

For housing, gains are measured at the household–year level using the reported purchase price:

$$R^H_{it} = (P^H_{it} - \bar{P}^H_{i,t-k}) \cdot \mathcal{I}_{it}, \qquad U^H_{it} = (P^H_{it} - \bar{P}^H_{i,t-k}) \cdot (1 - \mathcal{I}_{it}), \tag{3.4}$$

[2]In practice, we impute the average of the stock price in the last 2 years – this is consistent with the average household in sample holding a single position on average for 2 years, as we find in the data.

where $\mathcal{I}_{it} = 1$ indicates a sale in year t and k is the holding period implied by the reported purchase year.

The central summary statistic is the *average propensity to realize* (APR), the totals-based ratio

$$\text{APR}_g = \frac{\sum_{i \in g} R_i}{\sum_{i \in g} R_i + \sum_{i \in g} U_i}, \tag{3.5}$$

where g indexes a group of households or positions. The APR captures the share of accrued capital gains (realized plus unrealized) that a group converts into cash within the year—a single dimensionless measure of activity on the realization margin. At the position level the APR is mechanically bounded in $[0, 1]$: R and U share the factor $(P_{jt} - \bar{P}_{ijt})$, so it reduces to the ratio of shares sold over total shares. At the household–year level the APR is unbounded: aggregation across positions with different cost bases and signs can place the ratio outside $[0, 1]$. Section 3.A.2 exploits this distinction in a sign decomposition of household-level realization.

Let Π_j denote the gross return on position j over the year. The *relative return*, $\log(\Pi_j/\bar{\Pi}_{-j})$, measures Π_j against $\bar{\Pi}_{-j}$, the value-weighted gross return on the rest of the same household's portfolio. It measures how much share j performed compared to the rest of the portfolio and is used as regressor in Section 3.4 to show how households typically realize *winners* rather than *losers* in a capital gain tax-free environment.

Coverage limits and measurement caveats The DHS is recruited through an Internet panel and undercovers the very top of the wealth distribution, with consequences that vary by asset class. Stocks are highly top-concentrated, so DHS-based aggregate stock realizations are a strict lower bound on the population total. Housing is much less concentrated—Dutch homeownership stands at roughly sixty percent of households and is broad-based across the distribution—so housing realizations from the DHS are closer to representative.

Three additional measurement issues bear on interpretation. The cost basis for stock positions held at panel entry is unknown; we treat such position-years as left-censored and verify robustness to imputed first-year APPs. Self-reported income and wealth carry the standard non-response and recall biases of household surveys.[3] Housing cost bases are recorded only for the primary residence, so household-level housing realizations omit gains on second properties and rentals.

3.3. The Realization Channel

Two facts organize this section. First, the realization channel in the Netherlands is housing-dominated, with stocks and housing sitting at opposite ends of a frequency–size

[3] A matching with administrative data is ongoing and should resolve item non-response on income and net worth.

spectrum: housing realizes rarely, in large amounts, and almost always at a gain, while stock realizations are frequent, small, and centered near zero (Section 3.3.1). Second, a multiplicative decomposition of R/Y into a stock-of-capital, an embedded-gain, and a realization-margin term shows that pre-exposure levels—the composition and depth of household portfolios—account for a substantial share of the income gradient in realized gains, while the realization margin determines what fraction of those accrued gains is actually converted into cash and is the central object of the rest of the paper (Section 3.3.2). The position-level analysis of Section 3.4 then sharpens the behavioral question by asking *which* stocks households choose to liquidate.

3.3.1. Stocks and housing: a frequency–size asymmetry

We begin by contrasting realization patterns in the two biggest household asset classes, stocks and housing. Table 3.2 reports sale frequency and conditional sale size for each asset class, both overall and across age and income groups. The two assets sit on opposite ends of the frequency–size spectrum. Forty-seven percent of stockholding households sell at least one position in an average year; only one percent of homeowners do. Conditional sale sizes run in the opposite direction: a stock sale realizes about €8,000 on average, a housing sale roughly €255,000—more than thirty times larger. The stock margin sharpens with age along both dimensions, with sale frequency rising from 38 % for households under 45 to 52 % for those over 65 and average sale size doubling. Housing sale frequency stays low and roughly flat across age and income, but the conditional gain at sale rises with both, consistent with larger and longer-held properties at the top.

Table 3.2.: Sale frequency and conditional sale size: stocks vs. housing

	Stocks & mutual funds		Housing	
	Sale freq. (%)	Cond. size (€K)	Sale freq. (%)	Cond. size (€K)
Overall	46.5	8.0	1.0	255
Age <45	38.0	4.7	1.4	226
Age 45–64	46.8	8.0	0.7	241
Age 65+	52.0	9.6	0.9	315
Income p0–50	43.6	7.4	1.2	258
Income p50–90	48.1	7.6	0.8	265
Income top 10%	46.8	9.0	0.7	312

Notes: Sale frequency is the share of stockholding (resp. homeowning) household-years with a recorded sale. *Conditional sale size* is the population-weighted mean realized gain among sellers in the asset class. Sample: stockholding household-years for the stock columns, homeowning household-years for the housing columns; ages 20–85, 2001–2023. Wealth-tertile variant in Appendix Table 3.12.

The realization pattern translates into very different realized gains distributions, reported in Figure 3.3. Stock realizations cluster tightly around zero—both mean and median sit within a few thousand euros of the origin—but the tails are very fat: realizations in the 1st and 99th percentiles reach tens of thousands of euros in either direction, and excess kurtosis exceeds 60. Housing is the mirror image: nearly four out of five sales close at a gain, the median seller realizes about €65,000, and the bulk of the distribution sits above zero with a right tail reaching several hundred thousand euros. In addition, stock realized gains suffer also from a substantial part of the time sample where the Dutch stock market did not perform well: the AEX index did not recover its peak level at the onset of the dotcom bubble in 2000 until the late 2019 and was on a decreasing trend for about 10 years until the end of the 2008 financial crisis. In a post-2014 subsample, for example, the realized gains distribution for stocks is shifted to the right, but keeping its main features: small (positive) realizations and fat tails.

Figure 3.3.: Distribution of realized capital gains: stocks and housing.

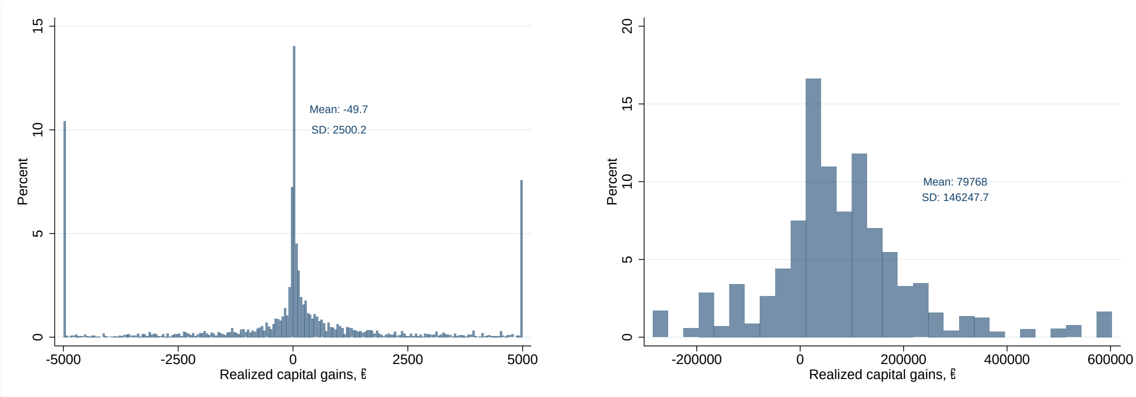

(a) Stocks (b) Housing

Notes: Population-weighted distributions of household-year realized gains in euros, conditional on a sale in the asset class. *Stocks* (panel a) computed at the household-year level by aggregating position-level realized gains; tails truncated at ±€5,000 for visibility. *Housing* (panel b) measured at the household-year level on owner-occupiers who sell.

The asymmetry shapes how the channel aggregates. Household-level realized gains are a mixture of two distributions: a high-frequency, small-magnitude stock component active most years for stockholders, and a low-frequency, large-magnitude housing component active in roughly one year out of fifty for homeowners. The cross-section of *who* realizes shifts with the asset class as well: housing realizations spread broadly across the wealth distribution because Dutch homeownership is broad, while stock realizations concentrate at the top because stock holdings do.

3.3.2. Decomposing gains: assets, accumulation, and the propensity to realize

The asset-class asymmetry motivates a clarifying framework. Stocks and housing realize on completely different frequency–size schedules, and the two assets also enter household portfolios at very different scales: housing wealth dominates the balance sheet of the median Dutch household, while equity is concentrated at the top. To interpret the cross-section of realized gains and average propensity to realize (APR), it helps to separate what households own from the rate at which they convert it into cash. This decomposition can be used in matching moments in models of realization behavior.

To make this concrete, let A_{it} denote household i's gross assets in year t, Y_{it} its gross income, and $G_{it} = R_{it} + U_{it}$ the total accrued capital gain on those assets (realized plus unrealized). Realized gains relative to income then decompose multiplicatively as

$$\frac{R_{it}}{Y_{it}} = \underbrace{\frac{A_{it}}{Y_{it}}}_{\text{stock of capital}} \times \underbrace{\frac{G_{it}}{A_{it}}}_{\text{embedded gain}} \times \underbrace{\frac{R_{it}}{G_{it}}}_{\text{realization margin}} . \tag{3.6}$$

The three terms isolate distinct primitives. The asset-to-income ratio A/Y reflects the *stock of capital* households have accumulated relative to current income; the gains-to-asset ratio G/A measures the *embedded gain* on that stock, which depends on past returns (i.e., market price dynamics not directly affected by the household) and on the asset mix; the R/G is the APR, that captures *behavior* on the realization margin—the share of embedded gains the household converts into cash within the year. A model of realization behavior need only discipline the third term; the first two are pinned down by demographics, history, and prices and can be read directly from the data.

Applying (3.6) and averaging within income and age tertile delivers Table 3.3. Two patterns stand out. First, the channel is asymmetric across assets: housing realizations average a positive R/Y of about 1.5 %, while stock realizations are slightly negative on average, reflecting the fat left tail of the realized-gain distribution. Restricting the sample to the post-2014 window flips the stock channel positive ($R/Y \approx 1.8\,\%$, Appendix Table 3.11), tracking the bull market over those years. Second, the cross-sectional gradient within either asset is driven mostly by A/Y, which falls by a factor of three to six from the bottom half to the top decile of income; G/A stays roughly flat and the APR contributes only modestly. The wealth-tertile cut sharpens the A/Y gradient but preserves the qualitative shape (Appendix Table 3.10).

Table 3.3.: Decomposition of $R/Y = (A/Y) \times (G/A) \times \text{APR}$

	Stocks & mutual funds				Housing			
	$\frac{R}{Y}$ (%)	$\frac{A}{Y}$	$\frac{R+U}{A}$ (%)	APR (%)	$\frac{R}{Y}$ (%)	$\frac{A}{Y}$	$\frac{R+U}{A}$ (%)	APR (%)
Overall	-1.42	1.11	-1.95	65.4	1.54	6.53	45.16	0.5
Age <45	-0.51	1.47	-0.53	65.8	1.06	5.68	26.99	0.7
Age 45–64	-1.38	0.62	-3.63	61.5	0.92	6.06	47.41	0.3
Age 65+	-2.11	1.58	-1.91	69.6	3.10	8.19	54.96	0.7
Income p0–50	-4.71	3.99	-1.93	61.2	5.08	13.40	47.66	0.8
Income p50–90	-1.08	0.74	-1.70	85.2	1.14	5.81	45.10	0.4
Income top 10%	-0.75	0.64	-2.49	47.5	0.43	4.21	40.85	0.3

Notes: Population-weighted means within income tertile. A denotes total household assets; Y gross household income; $G = R + U$ total gains (realized plus unrealized) on stocks and housing combined; $\text{APR} = \sum R / \sum(R + U)$ computed at the tertile level (totals basis). Identity in Equation (3.6) closes by construction up to rounding. Sample: stockholding or homeowning household-years, ages 20–85, full sample 2001–2023. Wealth-tertile and post-2014 variants reported in Tables 3.10 and 3.11.

The implication for the channel is structural rather than behavioral. Cross-sectional differences in realization volumes reflect differences in accumulated portfolios, not differences in households' propensity to convert paper gains into income. Whether the propensity itself shifts within households over the lifecycle, in response to liquidity shocks, or with market conditions, is the subject of the mechanism sections that follow.

Two other characterizations of the realization margin are further investigated in Appendix 3.A. First, the household-year APR is markedly heterogeneous: most household-years sit at zero because no sale takes place, and the remainder spread across a wide tail, with positive gains realized at noticeably higher rates than negative ones, i.e., $\text{APR}^+ \gg$ Loss realization rate (LRR). The asymmetry is sharpest in the top of the income distribution, suggesting that high-income households are more selective about which positions they liquidate. Second, household-year aggregation across positions with different cost bases can flip the sign of $R+U$ relative to R and place the APR outside $[0, 1]$. This is a measurement artifact rather than a behavioral statement, but it motivates the position-level analysis of Section 3.4, where the APR is mechanically bounded in $[0, 1]$ and *which* positions households realize, conditional on selling, becomes the right object.

3.4. Which Positions Are Liquidated? Realization of the Winners

Section 3.3 closed with a question: conditional on selling, which stocks do Dutch households actually liquidate? The empirical answer is sharp—they sell winners disproportionately. Section 3.4.1 documents this as a gap between the realized and unrealized return distributions at the position level; Section 3.4.2 estimates the sale-probability–relative-return relationship and instruments relative return with a Di Maggio-style market shock. The Dutch setting taxes neither realized gains nor losses, so the patterns we report cannot be attributed to tax-loss harvesting or lock-in—a sharpening relative to U.S. evidence, where the disposition signature is partially confounded with year-end tax planning.

3.4.1. The realized-versus-unrealized return gap

Figure 3.4 plots the position-level return distribution conditional on a sale against the corresponding distribution conditional on continued holding. Sold positions earn higher returns than held positions: the realized distribution lies to the right of the unrealized one throughout the support, with a heavier right tail, a thinner left tail, and a higher mode. On average, realized returns are 2 percentage points higher than returns on unrealized shares. This is the empirical signature of disposition behaviour first identified by Odean (1998).

Figure 3.4.: Realized vs. unrealized return distributions, position level, stocks.

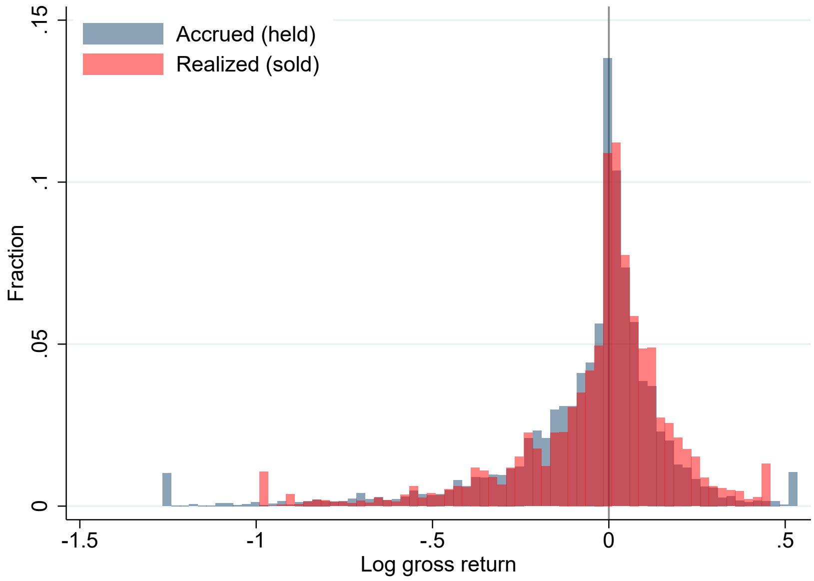

Notes: Empirical densities of position-level annualized returns. *Realized*: positions sold during the year. *Unrealized*: positions held at year-end. Both distributions are weighted by position euro-value and winsorized at the 1st/99th percentiles.

The Netherlands taxation is on imputed (rather than actual) returns to financial wealth (Box 3) but levies no capital gains tax on realization, so the standard tax-driven channels—lock-in on accrued gains, end-of-year tax-loss harvesting on losses—are absent. The right-shift of the realized distribution must therefore come from non-tax sources: behavioural attachment to losing positions, mental accounting, or rebalancing rules that respond to relative-return signals. Section 3.4.2 takes this pattern from descriptive to identified.

3.4.2. Sale probability rises sharply with relative return

We formally estimate the winners-bias at the position level. Define the relative return on position j held by household i at time t as $\log(\Pi_{ijt}/\bar{\Pi}_{-j,it})$, where $\bar{\Pi}_{-j,it}$ is the value-weighted gross return on the rest of household i's portfolio—that is, position j's return measured against its remaining holdings in the same household-year. We estimate the linear probability model

$$\Pr(\text{sale}_{ijt} = 1) \;=\; \alpha + \beta \, \log(\Pi_{ijt}/\bar{\Pi}_{-j,it}) + \gamma' X_{ijt} + \mu_j + \tau_{it} + \varepsilon_{ijt}, \tag{3.7}$$

on the full position-year panel. The stock fixed effects μ_j absorb cross-stock differences in baseline turnover propensity; the household-by-year fixed effects τ_{it} absorb everything time-varying at the household level—liquidity needs, overall trading activity, portfolio-wide taste shifts—so β is identified strictly *within* a household-year, by comparing the sale propensity of the same household's positions that earned different relative returns. What remains is a position-specific concern: a household that receives news about one of its stocks may both update its valuation and sell it, biasing β. We address this with an instrument that interacts position j's lagged portfolio weight with the change in j's market price return,

$$z_{ijt} \;=\; w_{ij,t-1} \, \Delta r^{\text{mkt}}_{j,t}, \tag{3.8}$$

following Di Maggio, Kermani, and Majlesi (2020), who apply this strategy at the portfolio level; our position-level implementation, combined with the household-by-year fixed effects, tightens the exclusion restriction. The lagged weight is predetermined relative to the time-t sale and the market price change is exogenous to the household's selection problem, so the instrument shifts relative return through portfolio composition without responding to the household's contemporaneous selling decisions.

Table 3.4 reports five specifications. Column 1 is OLS without fixed effects; columns 2–4 progressively add household, year, position, and household-by-year fixed effects; column 5 instruments relative return with (3.8) on the saturated-FE sample. The slope is positive and significant throughout, with a strong first stage (Kleibergen–Paap $F \approx 83$). The IV point estimate of 1.68 implies that a ten-log-point increase in relative return—roughly a 10 % outperformance of the rest of the household's portfolio—raises the position-year sale probability by about 17 percentage points, against a baseline sale rate of 49 %. The IV exceeds the within-household-year OLS slope of 0.26 by more than classical measurement-error attenuation alone would explain; the gap is also consistent with selection bias in OLS (households retaining positions they privately expect to outperform) and with the instrument recovering a steeper local effect for compliers—positions whose past portfolio weight happened to be high in stocks that subsequently experienced large market moves. The underlying binned scatter, reported in Appendix Figure 3.8, confirms that the relationship is monotone over the bulk of the relative-return distribution and modestly flatter at the extremes.

Table 3.4.: Sale probability and relative return

	OLS				IV
	(1)	(2)	(3)	(4)	(5)
Relative return (log)	0.052***	0.050*	0.038	0.264***	1.682***
	(0.014)	(0.026)	(0.024)	(0.048)	(0.334)
Controls	No	Yes	Yes	No	No
Household FE	Yes	Yes	Yes	No	No
Position FE	No	No	Yes	Yes	Yes
Year FE	No	No	Yes	No	No
Household × Year FE	No	No	No	Yes	Yes
Observations	11,873	5,125	5,015	6,124	5,250
R^2	0.111	0.115	0.219	0.487	-0.075

Notes: Dependent variable: indicator equal to one if the household sold (part of) the position during the year. Regressor of interest: relative return $\log(\Pi_{ijt}/\bar{\Pi}_{-j,it})$. Column 5 instruments relative return with z_{ijt} defined in Equation (3.8). Standard errors clustered at the household level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

The position-level pattern aggregates predictably to the household-year: under a level classification of sold positions ($R_j \gtrless 0$), 54 % of seller-years sell only winners against 27 % selling only losers, with the ranking preserved when winners and losers are defined relative to the rest of the portfolio (Appendix 3.B.2, Table 3.15). Trade intensity itself is heterogeneous, rising sharply with age, stock-portfolio value, and gross income at the household-year level, and concentrating on losing positions at the position level (Appendix 3.B.3, Figures 3.9–3.10).

Taken together, the position-level evidence in this section establishes that selection on winners is the dominant feature of how Dutch households realize stock gains: the realized-vs-unrealized gap, the sale-probability–relative-return slope, the household-year netting of sellers, and the concentration of partial-liquidation activity on losers all line up. Because the Netherlands does not tax realized gains or losses, this selection pattern cannot be a tax artefact and must originate in non-tax channels—belief inertia, attachment to losing positions, mental accounting, or rule-of-thumb rebalancing. The natural next question is whether the selection margin is a stable preference of households or a discretionary one that bends under pressure. Section 3.5 answers it by tracking the disposition wedge through life events—divorce, marriage, unemployment, retirement—that force liquidity onto the household.

3.5. Life Events as Stress Tests of the Selection Margin

Section 3.4 established that selection on winners is the dominant feature of how Dutch households realize stock gains. The natural follow-up question is whether this margin is a stable preference of households or one that bends under pressure—a quantitatively meaningful question, since marital transitions, job loss, and retirement are first-order drivers of household saving, labor supply, and portfolio composition over the lifecycle (e.g., Borella et al., 2025). If realization decisions respond to the same forces, the disposition pattern should narrow precisely when these events demand cash from the household. We test this by comparing the realized and accrued return distributions inside and outside windows around four life events: unemployment, divorce, marriage, and retirement.

Figure 3.5 reports the histograms. Each panel overlays two position-level distributions of annualized log gross returns: the *accrued* distribution (returns on positions still held at year-end) in navy and the *realized* distribution (returns on positions sold during the year) in red. Dashed vertical lines mark the means of the two distributions, and the gap between them—the *return wedge*—is reported in each panel as the difference (realized minus accrued) in log points. The wedge is the diagnostic object: a positive value means the typical sold position outperformed the typical held position, the empirical signature of selection on winners; a wedge near zero means sold and held positions earned similar returns, i.e., realization is no longer selective; a negative wedge means sold positions *underperformed* held ones, the opposite of disposition behaviour. Comparing the wedge across panels therefore tells us, panel by panel, whether the selection margin survives, weakens, or inverts inside each life-event window.

Outside any event window (panel a), the realized distribution sits visibly to the right of the accrued one, with a wedge of +0.027 log points—the disposition signature of §3.4.1, recovered on the full position-year sample. Inside event windows, the wedge attenuates or inverts in every case, but the magnitudes vary sharply. Unemployment leaves it positive but compressed (+0.019); marriage and retirement bring it essentially to zero (−0.019 and −0.006 respectively), with the realized and accrued distributions nearly coinciding; divorce delivers a sharp inversion (−0.132), with the realized distribution shifting to the *left* of the accrued one—households selling under divorce (and partly the other life events) are no longer cherry-picking winners and, if anything, are forced to liquidate several positions, including the worse ones in their portfolio. Given that these events are quite rare in the sample, they are severely underpowered, which motivates the use of simple comparison of means of distributions.

Figure 3.5.: Realized vs. accrued return distributions: life-events.

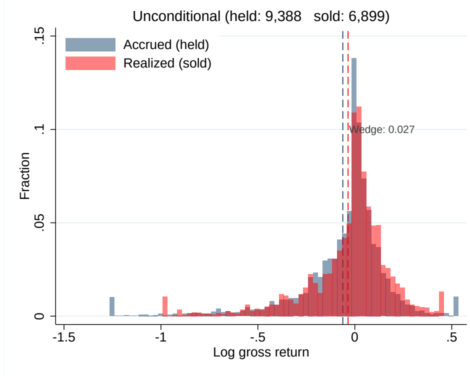

(a) Unconditional baseline

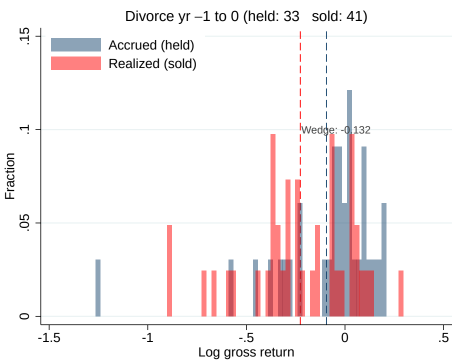

(b) Divorce window

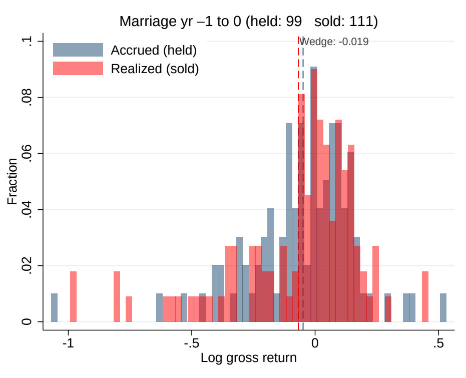

(c) Marriage window

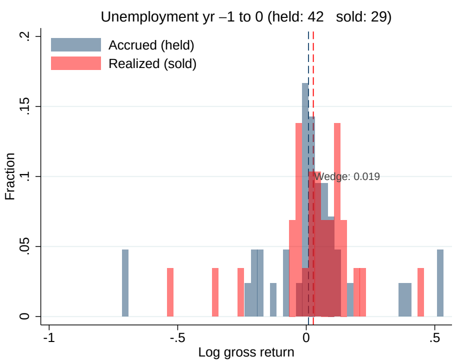

(d) Unemployment window

(e) Retirement window

Notes: Position-level distributions of annualized log gross returns. *Accrued (held)*: positions held at year-end. *Realized (sold)*: positions sold during the year. Dashed vertical lines mark the means; "Wedge" is the difference (realized − accrued) in log points. Event windows are $t - t^* \in \{-1, 0\}$, where t^* is the event year. Sample sizes (held / sold) reported in each subtitle. All distributions winsorized at the 1st/99th percentiles.

A simple interaction regression formalizes the comparison. Pool realized and accrued returns into a single position-year variable, $\log r_{ijt}$, and estimate

$$\log r_{ijt} \;=\; \beta_1 \,\text{sold}_{ijt} + \beta_2 \,\text{event}_{it} + \beta_3 \,(\text{sold} \times \text{event})_{ijt} + \varepsilon_{ijt}, \tag{3.9}$$

where sold_{ijt} is a sale indicator and event_{it} is one if year t falls in the household's event window. β_1 recovers the baseline disposition gap; β_3 measures the change in that gap during the event. Standard errors are clustered at the household level.

Table 3.5 reports β_3 for each of the four events. Divorce delivers $\beta_3 = -0.161$, the largest negative estimate in the table: the realized–accrued gap contracts by sixteen percentage points inside the divorce window, exceeding the baseline gap of 2.8 percentage points and confirming the visual inversion in panel (b). Marriage, unemployment, and retirement also deliver negative point estimates (-0.047, -0.009, -0.034), consistent with the descriptive narrowing visible in the histograms, but the in-window samples are an order of magnitude smaller and the estimates are imprecise. The divorce coefficient is not driven by sample size: its event window has the smallest total position count (503) of the four, and the largest absolute estimate.

Table 3.5.: Disposition gap inside event windows

	Unemployment	Divorce	Marriage	Retirement
β_1 (baseline disposition gap)	0.027^{***}	0.028^{***}	0.028^{***}	0.028^{***}
	(0.004)	(0.004)	(0.004)	(0.004)
β_3 (sold × event window)	-0.009	-0.161^{***}	-0.047	-0.034
	(0.045)	(0.034)	(0.057)	(0.030)
Position-years in window	638	503	1,490	2,493
Sold (with return)	29	41	111	176

Notes: Each column is a separate OLS regression of position-level log return on a sale indicator, an event-window indicator, and their interaction. Event window is $t \in \{-1, 0\}$ relative to the event year. β_3 is the coefficient of interest: a negative value indicates the disposition gap narrows inside the event window. Standard errors clustered at the household level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

The divorce result is the cleanest finding in the table. It shows that the disposition pattern is not a hard preference of households—it can be broken—and that breaking it requires a household disruption sharp enough to force the liquidation of jointly held positions, regardless of which positions households would otherwise have chosen to keep. The other three events are underpowered: marriage, unemployment, and retirement all yield negative point estimates, directionally consistent with some weakening of the selection margin under liquidity pressure, but the small in-window samples do not allow us to separate genuine narrowing from noise.

3.6. What Happens After a Sale? Owners Extract, Renters Rebalance

The previous sections established two facts about how Dutch households realize stock gains: in the aggregate the channel is housing-dominated, with stock realizations small and concentrated at the top of the income distribution (Section 3.3); at the position level, households selectively liquidate winners over losers (Section 3.4), a margin that attenuates only under sharp household disruptions like divorce (Section 3.5). We now ask what happens to a euro of proceeds once a position is sold. The answer depends sharply on tenure. For homeowners, stock realizations are absorbed by current consumption and net cash flows to creditors, with rebalancing back into financial assets essentially zero; housing realizations, in turn, finance trade-up purchases that expand rather than contract household leverage. For renters, stock proceeds are reinvested approximately one-for-one back into financial assets once cross-sectional scale is purged. The tenure asymmetry pins down the behavioural content of the realization channel: a realized stock gain functions as spendable income for homeowners and as portfolio rotation for renters—a within-asset, between-tenure heterogeneity in the marginal propensity to consume out of wealth that the consumption–savings literature on lumpy liquidity has emphasized (Kaplan and Violante, 2014; Fagereng, Holm, and Natvik, 2021).

3.6.1. Sources and uses of funds: a household budget identity

To trace where a marginal euro of proceeds goes once a position is sold, we need a framework that imposes adding-up across every candidate use—investment, debt service, consumption, and liquid balances. We obtain it from a household budget identity in which sources and uses of funds must balance period by period; this identity disciplines the marginal-allocation regressions that follow and backs out consumption as a residual, since DHS does not observe it directly.

We work with a discrete-time household balance sheet with four asset classes: financial assets S_t, housing assets H_t, liquid assets A_t, and liabilities $B_t < 0$. The laws of motion are

$$\begin{aligned} S_{t+1} &= S_t + G_t^S + N_t^S, & N_t^S &= I_t^S - R_t^S, \\ H_{t+1} &= H_t + G_t^H + N_t^H, & N_t^H &= I_t^H - R_t^H, \\ B_{t+1} &= B_t(1 + r_t) - b_t, \\ A_{t+1} &= A_t + D_t^S + D_t^H + W_t - C_t - N_t^S - N_t^H - b_t. \end{aligned}$$

Here G_t^j is unrealized capital gains on asset class $j \in \{S, H\}$, R_t^j is realizations (sales), I_t^j is gross investment, N_t^j is net flows, D_t^j is dividends or rental income, W_t is labor income, C_t is consumption, and b_t is the net cash flow to creditors (interest plus principal repayment minus new borrowing): $b_t > 0$ when the household pays down debt on net and $b_t < 0$ when it borrows on net. Expanding the net-flow terms in the liquid-asset equation and rearranging yields the budget identity

$$I_t^S + I_t^H + b_t + C_t + (A_{t+1} - A_t) \; = \; D_t^S + D_t^H + W_t + R_t^S + R_t^H. \tag{3.10}$$

The right-hand side collects the five sources of funds available to the household in period t:

financial dividends, housing rental income, labor income, stock sales, and housing sales. The left-hand side collects five uses: gross stock investment, gross housing investment, net cash to creditors, consumption, and the change in liquid balances. Since labor income, dividends, sales, gross investments, debt service, and the change in liquid balances are all observed in DHS, equation (3.10) backs out consumption as a residual,

$$C_t \;=\; D_t^S + D_t^H + W_t + R_t^S + R_t^H \;-\; I_t^S - I_t^H - b_t - (A_{t+1} - A_t). \tag{3.11}$$

3.6.2. Estimating the marginal allocation of proceeds

Let $M_t \equiv A_{t+1} - A_t$ denote liquidity accumulation. Stack the uses into $Z_t \equiv (I_t^S, I_t^H, b_t, C_t, M_t)'$ and the sources into $X_t \equiv (D_t^S, D_t^H, W_t, R_t^S, R_t^H)'$. The budget identity (3.10) can then be written compactly as $\mathbf{1}'Z_t = \mathbf{1}'X_t$. We estimate the linear projection of each use on the full vector of sources,

$$Z_t \;=\; \delta + \Gamma X_t + u_t, \qquad \mathbb{E}[u_t] = 0, \quad \mathrm{Cov}(X_t, u_t) = 0, \tag{3.12}$$

which defines δ and Γ as population OLS coefficients without any assumption on the conditional distribution of u_t. Summing equation (3.12) across uses and applying $\mathbf{1}'Z_t = \mathbf{1}'X_t$ implies $(\mathbf{1}' - \mathbf{1}'\Gamma)X_t - \mathbf{1}'\delta = \mathbf{1}'u_t$. The orthogonality $\mathrm{Cov}(X_t, \mathbf{1}'u_t) = 0$ and nonsingular $\mathrm{Var}(X_t)$ then force $\mathbf{1}'\Gamma = \mathbf{1}'$ and $\mathbf{1}'\delta = 0$. Equivalently, for each source k,

$$\gamma_{I^S,k} + \gamma_{I^H,k} + \gamma_{b,k} + \gamma_{C,k} + \gamma_{M,k} \;=\; 1. \tag{3.13}$$

The coefficients $\gamma_{j,k}$ are naturally interpreted as *marginal allocation shares*: the share of a marginal euro from source k that the household routes to use j. They are, however, *prediction coefficients*, not causal parameters; absent exogenous variation in the sources, they summarize conditional associations rather than counterfactual marginal propensities. We treat them as descriptive statistics that the realization channel needs to be made of.

We estimate (3.12) separately for homeowners and renters. The two groups face different tenure decisions and different access to mortgage credit, and the housing-related sources and uses (R^H, D^H, I^H) are absent for renters by construction. The owner system has five sources and five uses; the renter system collapses to three sources (W, R^S, D^S) and four uses (I^S, b, C, M). The main specification is OLS without fixed effects; standard errors are clustered at the household level. Robustness specifications add household and year fixed effects, replace the level scaling with shares of total sources or with lagged household wealth as denominator, and tighten winsorization to the 5th–95th percentiles; these are reported in Appendix 3.C.

Table 3.6.: Marginal allocation of sources to uses, homeowners

	(1) ΔLiquid	(2) Stock inv. (I^S)	(3) Housing inv. (I^H)	(4) Debt service (b)	(5) Consumption
Stock sales (R^S)	-0.355^{***}	0.026	-0.057^{***}	0.727^{***}	0.660^{***}
	(0.090)	(0.060)	(0.012)	(0.156)	(0.189)
Housing sales (R^H)	0.040^{**}	-0.006	0.038^{***}	-0.390^{***}	1.318^{***}
	(0.018)	(0.004)	(0.006)	(0.078)	(0.083)
Net non-capital income (W)	0.062^{***}	0.010^{**}	0.017^{***}	0.021	0.890^{***}
	(0.012)	(0.005)	(0.003)	(0.030)	(0.033)
Interest/dividends	0.228^{***}	0.265^{***}	-0.038^{***}	-0.101	0.646^{***}
	(0.078)	(0.065)	(0.014)	(0.151)	(0.172)
Real estate income (D^H)	-0.128	-0.021	0.229^{***}	2.036^{***}	-1.116
	(0.341)	(0.125)	(0.084)	(0.694)	(0.778)
N	12853	12853	12853	12853	12853
R^2	0.006	0.013	0.021	0.009	0.123

Notes: OLS, no fixed effects, SE clustered on household. Each column is a use of funds; rows give the projection coefficient on each source. By construction, row sums equal one across all uses. b_t is net cash to creditors (Convention A): $b_t > 0$ is net debt repayment, $b_t < 0$ is net borrowing. ΔLiquid is $A_{t+1} - A_t$. Sample: homeowner-years in DHS, 2001–2023. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3.7.: Marginal allocation of sources to uses, renters

	(1) ΔLiquid	(2) Stock inv. (I^S)	(3) Debt service (b)	(4) Consumption
Stock sales (R^S)	0.025	0.112	-0.006	0.869^{***}
	(0.235)	(0.156)	(0.032)	(0.312)
Net non-capital income (W)	0.038^{*}	0.009	-0.012	0.964^{***}
	(0.022)	(0.009)	(0.013)	(0.028)
Interest/dividends	0.107	0.237^{**}	0.008	0.649^{***}
	(0.140)	(0.111)	(0.037)	(0.182)
Real estate income (D^H)	1.913	1.158	0.733	-2.803
	(2.081)	(0.713)	(1.380)	(3.479)
N	4527	4527	4527	4527
R^2	0.003	0.033	0.001	0.308

Notes: OLS, no fixed effects, SE clustered on household. Renter system omits housing-related sources and uses. Conventions as in Table 3.6. Sample: renter-years in DHS, 2001–2023.

3.6.3. Marginal allocation of realization proceeds

Tables 3.6 and 3.7 report the main specification. Each column corresponds to one use; rows give the projection coefficients on each source; and by construction the coefficients in any row sum to one across columns up to rounding. We focus the discussion on the two realization sources, R^S and R^H.

Owners: stock sales. For homeowners (Table 3.6), a marginal euro of stock realization splits roughly in half between consumption and net cash flows to creditors: $\hat{\gamma}_{C,R^S} = 0.660$ and $\hat{\gamma}_{b,R^S} = 0.727$. Liquid balances absorb a negative share, $\hat{\gamma}_{M,R^S} = -0.355$, so proceeds do not pile up in cash; reinvestment into stocks is essentially zero, $\hat{\gamma}_{I^S,R^S} = 0.026$; reinvestment into housing is small and negative, $\hat{\gamma}_{I^H,R^S} = -0.057$. The five coefficients sum to 1.001, as required by (3.13). Two clarifications matter for the interpretation. First, $\hat{\gamma}_{C,R^S}$ is identified off realized euros, not gains, so it is not directly comparable to the textbook MPC out of wealth: it is the share of liquidated principal-plus-gain that contemporaneously shows up as consumption flow. Either way, owners do not treat stock sales as portfolio rotation; they treat them as spendable cash that leaves the household sector via consumption or via creditors.

Owners: housing sales. Housing realizations look very different. The contemporaneous consumption coefficient is large, $\hat{\gamma}_{C,R^H} = 1.318$, and the debt-service coefficient is sharply *negative*, $\hat{\gamma}_{b,R^H} = -0.390$. A negative $\hat{\gamma}_b$ on R^H is the signature of a trade-up: a euro of housing sale is associated with 0.39 euro of *additional* net borrowing, not with debt repayment. The remaining coefficients are small: $\hat{\gamma}_{M,R^H} = 0.040$, $\hat{\gamma}_{I^H,R^H} = 0.038$, and $\hat{\gamma}_{I^S,R^H} = -0.006$; the row sums to 1.000. The implied behavioural pattern is that DHS households who sell a home typically buy a more expensive one, finance the difference with new mortgage debt, and the residual consumption flow absorbs the timing gap between sale proceeds, transaction costs, and renovation expenditures captured by (3.11). A consumption coefficient above one is not anomalous in this design, since C_t is constructed as a residual that includes any unobserved gross outlays in the year of the move. The qualitative point is that housing sales in this sample do not finance retirement consumption out of accumulated home equity; they finance further housing accumulation and additional leverage—the realization channel for housing operates inside the housing sector.

Renters: stock sales. For renters (Table 3.7), only R^S is observed. In the level specification, a marginal euro of stock realization appears almost entirely as consumption, $\hat{\gamma}_{C,R^S} = 0.869$, with imprecise coefficients on the other three uses ($\hat{\gamma}_{M,R^S} = 0.025$, $\hat{\gamma}_{I^S,R^S} = 0.112$, $\hat{\gamma}_{b,R^S} = -0.006$). The level system is, however, dominated by cross-sectional heterogeneity in scale: renters with small and large portfolios are pooled, and the estimate confounds within- and between-household variation. We therefore view the lagged-wealth-scaled specification (Table 3.19, Appendix 3.C) as the more informative cut. There, $\hat{\gamma}_{I^S,R^S} = 1.097$: a euro of stock proceeds is reinvested approximately one-for-one back into financial assets.

The owner–renter contrast. The contrast on the same source—stock realizations—is the central finding of this section. For owners, the marginal euro flows to consumption and debt service with essentially no reinvestment; for renters, it cycles back into the financial portfolio almost in full. The two groups use stocks for different purposes: renters hold stocks as their primary risky savings vehicle and treat realizations as portfolio churn, while owners hold housing as their dominant risky asset and use stock sales as the marginal source of spendable income. The pattern is consistent with the two-asset (liquid–illiquid) interpretation of MPC heterogeneity in Fagereng, Holm, and Natvik (2021): owners hold housing as the illiquid stock and use the financial portfolio as the liquid buffer drawn down at the margin, while renters have no illiquid asset and therefore keep financial wealth recycling inside the portfolio. The same realized euro funds different things in different parts of the housing-tenure distribution, and the gap is large.

Robustness. We assess the level estimates against four alternatives, reported in Appendix 3.C: (i) tighter winsorization at the 5th–95th percentiles; (ii) scaling each source and use by total sources, which removes the level dependence of the estimates; (iii) the same scaled specification with household and year fixed effects, which absorbs persistent heterogeneity in tenure status, family structure, and labor income; and (iv) scaling by lagged household wealth, used in §3.6.3 as the preferred cut for renters. The qualitative pattern—owners route R^S to consumption and debt service, owners route R^H to additional housing and additional debt, renters reinvest R^S back into stocks—is stable across all four. Magnitudes shift with the scaling: the owner consumption coefficient on R^S ranges from 0.27 to 0.66 across specifications, the owner debt-service coefficient on R^H from -0.39 to -0.59, and the renter stock-reinvestment coefficient on R^S from 0.11 (level) to 1.10 (wealth-scaled), but the signs and the rank ordering across uses are unchanged. As a further check, Appendix 3.C also reports local projections of each use $Z_{j,t+h}$ on the full source vector X_t at horizons $h \in \{-2, \dots, +2\}$ with household and year fixed effects; the contemporaneous coefficients reproduce the static rows and the dynamics show no systematic pre-trends, so the level estimates do not appear to be picking up offsetting flows in adjacent years.

3.7. Conclusion

This paper has used three decades of Dutch household-level data to trace capital gains realizations from the size of the realized flow, through the household-level decision to sell, to the marginal use of the proceeds. The three layers line up. Realized gains are a sizeable household income flow, on the order of several percent of GDP, with two assets contributing on opposite frequency–size schedules: roughly half of stockholding households sell at least one position in a given year but in small amounts, while only about one homeowner in a hundred sells, in amounts more than thirty times larger. Cross-sectional differences in realized volumes are driven mostly by what households own rather than by how aggressively they convert paper gains into cash: the income gradient in R/Y runs through the asset-to-income ratio A/Y, not through the average propensity to realize. Sales respond to portfolio composition and to a small set of life events—divorce in particular—rather than to high-frequency variation in returns. And once the sale occurs, where the proceeds go depends sharply on housing tenure: owners spend stock realizations and channel housing realizations into trade-ups and mortgages, while renters cycle stock realizations back into the financial portfolio.

The implications run in two directions. On the distributional side, the no-capital gains taxation environment makes the wedge between realized and accrued gains transparent: households who sell receive a spendable flow at the moment of realization, while non-sellers hold gains as paper wealth. The incidence of that spendable flow inherits the underlying asset-holding distribution rather than any behavioral skew: stock realizations are concentrated at the top because stock holdings are, while housing realizations spread broadly because Dutch homeownership is broad. On the macro side, the marginal-allocation results imply that the consumption response to a realized euro is far from uniform: it depends on which asset was sold and on whether the household owns its home. Aggregating over households without these splits will mismeasure the spending response to asset-price movements.

The findings discipline the next generation of models in three concrete ways. First, the realization choice itself, not the accrued gain, is the relevant decision margin: a model in which households mark to market and consume a fixed fraction of wealth changes will miss the bunching of behavioral response at the moment of sale. Second, tenure status belongs in the state space: the same realized euro from the same asset is allocated differently by owners and renters, so a representative-agent or single-asset framework cannot reproduce the cross-section of marginal propensities. Third, life events—divorce most clearly—move realizations more than returns do, which suggests that liquidity and idiosyncratic non-financial shocks, rather than rebalancing motives alone, drive the timing of sales. A natural next step is to embed an explicit realization decision in a two-asset lifecycle model with housing tenure and family-structure shocks, calibrated to the moments documented here, and to ask how much of the cross-section of MPCs out of wealth and the macro response to asset-price movements such a model can recover.

References

Adam, Klaus and Ulrich J. Wagner (2023). “Composition vs. Intensity Timing in Household Portfolio Returns”. Working paper (cit. on pp. 149, 150).

Aguiar, Mark, Benjamin Moll, and Florian Scheuer (Sept. 2024). *Putting the “Finance” into “Public Finance”: A Theory of Capital Gains Taxation*. NBER Working Paper 32951. Cambridge, MA: National Bureau of Economic Research. DOI: 10.3386/w32951 (cit. on p. 149).

Alessie, Rob, Stefan Hochgürtel, and Arthur van Soest (2004). “Ownership of Stocks and Mutual Funds: A Panel Data Analysis”. In: *Review of Economics and Statistics* 86.3, pp. 783–796. DOI: 10.1162/0034653041811761 (cit. on p. 149).

Barber, Brad M. and Terrance Odean (2000). “Trading Is Hazardous to Your Wealth: The Common Stock Investment Performance of Individual Investors”. In: *Journal of Finance* 55.2, pp. 773–806. DOI: 10.1111/0022-1082.00226 (cit. on p. 152).

– (2001). “Boys Will Be Boys: Gender, Overconfidence, and Common Stock Investment”. In: *Quarterly Journal of Economics* 116.1, pp. 261–292. DOI: 10.1162/003355301556400 (cit. on p. 152).

Borella, Margherita, Mariacristina De Nardi, Fang Yang, and Johanna P. Torres Chain (May 2025). *Why Do Households Save and Work?* NBER Working Paper 33874. Revised January 2026. Cambridge, MA: National Bureau of Economic Research. DOI: 10.3386/w33874 (cit. on pp. 149, 162).

Burman, Leonard E. and Peter D. Ricoy (Sept. 1997). “Capital Gains and the People Who Realize Them”. In: *National Tax Journal* 50.3, pp. 427–451 (cit. on p. 148).

Calvet, Laurent E., John Y. Campbell, and Paolo Sodini (2007). “Down or Out: Assessing the Welfare Costs of Household Investment Mistakes”. In: *Journal of Political Economy* 115.5, pp. 707–747. DOI: 10.1086/524204 (cit. on p. 152).

– (2009). “Fight or Flight? Portfolio Rebalancing by Individual Investors”. In: *Quarterly Journal of Economics* 124.1, pp. 301–348. DOI: 10.1162/qjec.2009.124.1.301 (cit. on p. 152).

Constantinides, George M. (May 1983). “Capital Market Equilibrium with Personal Tax”. In: *Econometrica* 51.3, pp. 611–636. DOI: 10.2307/1912150 (cit. on p. 153).

Di Maggio, Marco, Amir Kermani, and Kaveh Majlesi (Dec. 2020). “Stock Market Returns and Consumption”. In: *Journal of Finance* 75.6, pp. 3175–3219. DOI: 10.1111/jofi.12968 (cit. on pp. 147–149, 160).

Fagereng, Andreas, Luigi Guiso, Davide Malacrino, and Luigi Pistaferri (Jan. 2020). “Heterogeneity and Persistence in Returns to Wealth”. In: *Econometrica* 88.1, pp. 115–170. DOI: 10.3982/ECTA14835 (cit. on p. 153).

Fagereng, Andreas, Martin B. Holm, and Gisle J. Natvik (Oct. 2021). “MPC Heterogeneity and Household Balance Sheets”. In: *American Economic Journal: Macroeconomics* 13.4, pp. 1–54. DOI: 10.1257/mac.20190211 (cit. on pp. 146, 148, 149, 165, 169).

Gaudecker, Hans-Martin von (Apr. 2015). "How Does Household Portfolio Diversification Vary with Financial Literacy and Financial Advice?" In: *Journal of Finance* 70.2, pp. 489–507. DOI: 10.1111/jofi.12231 (cit. on pp. 149, 152).

Grinblatt, Mark and Matti Keloharju (Jan. 2000). "The Investment Behavior and Performance of Various Investor Types: A Study of Finland's Unique Data Set". In: *Journal of Financial Economics* 55.1, pp. 43–67. DOI: 10.1016/S0304-405X(99)00044-6 (cit. on p. 152).

Kaplan, Greg and Giovanni L. Violante (July 2014). "A Model of the Consumption Response to Fiscal Stimulus Payments". In: *Econometrica* 82.4, pp. 1199–1239. DOI: 10.3982/ECTA10528 (cit. on pp. 146, 148, 149, 165).

Larrimore, Jeff, Richard V. Burkhauser, Gerald Auten, and Philip Armour (2021). "Recent Trends in U.S. Income Distributions in Tax Record Data Using More Comprehensive Measures of Income Including Real Accrued Capital Gains". In: *Journal of Political Economy* 129.5, pp. 1319–1360. DOI: 10.1086/713098 (cit. on p. 148).

Meeuwis, Maarten (Feb. 2022). "Wealth Fluctuations and Risk Preferences: Evidence from U.S. Investor Portfolios". Working paper, Washington University in St. Louis (cit. on p. 149).

Meyer, Steffen and Michaela Pagel (June 2022). "Fully Closed: Individual Responses to Realized Gains and Losses". In: *Journal of Finance* 77.3, pp. 1529–1585. DOI: 10.1111/jofi.13122 (cit. on p. 148).

Msall, Lucy and Ole-Andreas Næss (2025). "Never-Realized Capital Gains". Job market paper, University of Chicago Booth School of Business (cit. on p. 148).

Odean, Terrance (Oct. 1998). "Are Investors Reluctant to Realize Their Losses?" In: *Journal of Finance* 53.5, pp. 1775–1798. DOI: 10.1111/0022-1082.00072 (cit. on pp. 147, 148, 159).

– (Dec. 1999). "Do Investors Trade Too Much?" In: *American Economic Review* 89.5, pp. 1279–1298. DOI: 10.1257/aer.89.5.1279 (cit. on p. 148).

Piketty, Thomas, Emmanuel Saez, and Gabriel Zucman (2018). "Distributional national accounts: methods and estimates for the United States". In: *The Quarterly Journal of Economics* 133.2, pp. 553–609 (cit. on p. 148).

Robbins, Jacob A. (2018). "Capital Gains and the Distribution of Income in the United States". Job market paper, Brown University (cit. on p. 148).

3.A. Aggregate Channel and APR: Supporting Material

3.A.1. Distribution and gain–loss asymmetry of the household-year APR

The household-year APR is markedly heterogeneous and asymmetric in gains and losses. Two features of the distribution stand out. First, the household-year APR is bimodal: a large mass at zero (non-sellers) coexists with a long right tail extending beyond unity (Figure 3.6 panel a). Conditioning on sellers (panel b) removes the spike at zero but preserves the heavy tail, indicating that within-seller heterogeneity is itself substantial. Second, gains and losses are realized at different rates: the positive-gain APR (APR^+) and the loss realization rate (LRR) move differently across income and wealth tertiles (Table 3.8), with the gains–loss differential most pronounced in the top tertile.

Figure 3.6.: Household-year APR distribution.

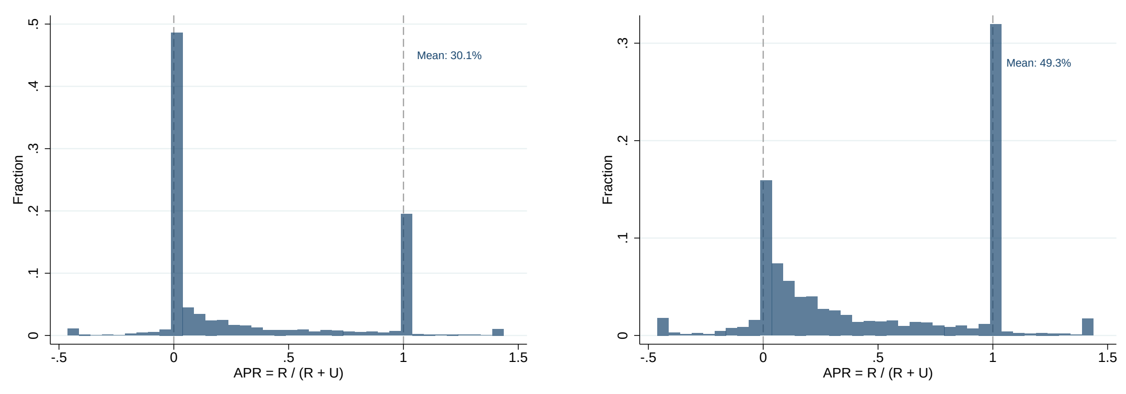

(a) All household-years (b) Sellers only

Notes: Household-year APR computed as $R/(R+U)$ on stocks aggregated within household-year. APR is unbounded at the household-year level: aggregation across positions with different cost bases can place the ratio outside $[0, 1]$. Histograms winsorized at 1st/99th percentile. Household-years with $R + U = 0$ are excluded.

Table 3.8.: Positive-gain APR (APR^+) and loss realization rate (LRR), stocks, by income tertile

	Gains ($R+U>0$)			Losses ($R+U<0$)		
	APR (%)	Avg. gain (€K)	N	LRR (%)	Avg. loss (€K)	N
Overall	14.3	2.4	5537	15.0	-3.7	4497
Age <45	11.4	1.0	913	10.3	-1.4	824
Age 45–64	13.3	2.1	2308	10.2	-3.3	1835
Age 65+	15.2	3.4	2316	18.7	-5.1	1838
Income p0–50	14.5	2.2	1737	14.7	-3.7	1459
Income p50–90	14.3	2.4	3031	16.5	-3.4	2362
Income top 10%	13.8	3.3	769	11.5	-4.6	676

Notes: $\text{APR}^+ = \sum R^+ / \sum(R^+ + U^+)$ within tertile, where R^+ and U^+ restrict to positions or household-years with positive total gains. LRR defined symmetrically on negative gains. Both statistics computed on the totals basis. Wealth-tertile and post-2014 variants in Tables 3.13 and 3.14.

The income gradient of the household-year APR is consistent with what Equation (3.6) suggests, though the lifecycle profile is hump-shaped rather than monotone.

3.A.2. Sign aggregation in the household-year APR

At the household-year level, realized and total gains can take opposite signs. A household selling a winner while continuing to hold losers has $R > 0$ even when $R + U < 0$ on the portfolio; a household holding mostly winners but selling a loss position has $R < 0$ with $R + U > 0$. Sign discordance is impossible at the position level—R and U share the factor $(P_{jt} - \bar{P}_{ijt})$—and it appears only when household-year aggregates fold across positions with different cost bases.

Table 3.9 cross-tabulates household-years by sign(R) and sign($R + U$), with within-cell totals-based APRs in parentheses. The concordant cells—top-left ($R > 0$, $R + U > 0$) and bottom-right ($R < 0$, $R + U < 0$)—account for the majority of household-years and have within-cell APR in $[0, 1]$ as expected. The discordant cells—top-right and bottom-left—contain a meaningful share of household-years and host APRs that are negative or exceed unity. The aggregate household-year APR of roughly 66 % reflects totals-based aggregation across all four cells weighted by the magnitudes of R and U. A mean-of-ratios alternative, weighting every household-year equally, lands closer to 30 %; the gap is a Jensen's-inequality artifact of cross-position aggregation rather than a behavioral statement.

Table 3.9.: Sign decomposition of realized and total gains, household-year level, stocks

	$R > 0$	$R = 0$	$R < 0$	Total
$R + U > 0$	31.7%	22.9%	1.9%	56.6%
	[1,766]	[1,268]	[103]	[3,137]
	(43.9%)	(0.0%)	(-25.2%)	(31.6%)
$R + U < 0$	3.1%	17.2%	23.1%	43.4%
	[179]	[936]	[1,286]	[2,401]
	(-11.6%)	(0.0%)	(61.9%)	(43.9%)
Total	34.9%	40.1%	25.0%	100.0%
	[1,945]	[2,204]	[1,389]	[5,538]
	(50.2%)	(0.0%)	(64.0%)	(66.0%)

Notes: Each cell reports the population-weighted share of household-years (top), unweighted count (brackets), and totals-based APR in percent (parentheses). The $R = 0$ column captures non-sellers, for whom APR is mechanically zero. Sample restricted to stockholding household-years with non-zero total gains. Weights are time-averaged household weights.

Figure 3.7.: Household-year APR distribution by sign of realized and total gains.

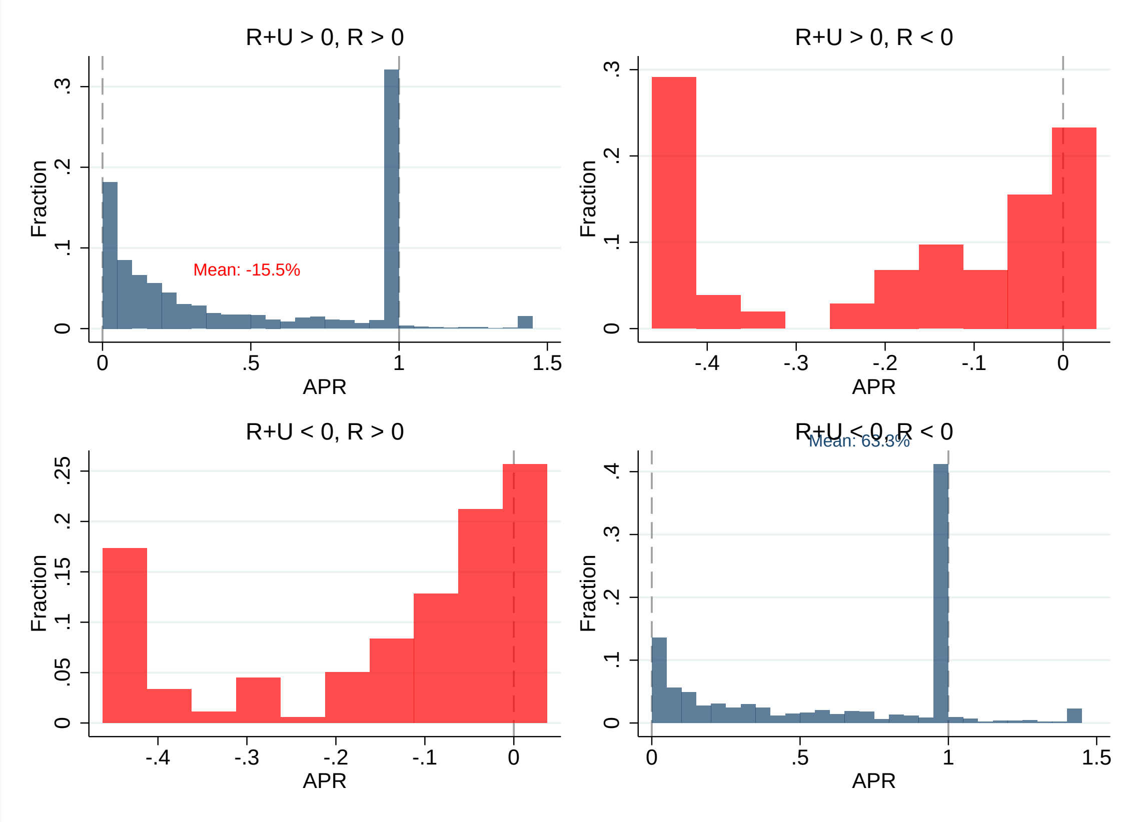

Notes: Top row: portfolios with positive total gains ($R + U > 0$). Bottom row: $R + U < 0$. Left column: $R > 0$. Right column: $R < 0$. Concordant cells (top-left, bottom-right) show APR concentrated in $[0, 1]$; discordant cells show negative APR or APR above unity, reflecting sign-switching from cross-position aggregation. APR winsorized at 1st/99th percentile.

The position-level analysis of Section 3.4 sidesteps the aggregation issue by construction. At the position level the APR is bounded in $[0, 1]$, the sign discordance vanishes, and the question of *which* positions households realize, conditional on selling, becomes well-posed.

3.A.3. Wealth-tertile and post-2014 variants

Table 3.10.: Decomposition of $R/Y = (A/Y) \times (G/A) \times \text{APR}$, by wealth tertile

	Stocks & mutual funds				Housing			
	$\frac{R}{Y}$ (%)	$\frac{A}{Y}$	$\frac{R+U}{A}$ (%)	APR (%)	$\frac{R}{Y}$ (%)	$\frac{A}{Y}$	$\frac{R+U}{A}$ (%)	APR (%)
Overall	-1.42	1.11	-1.95	65.4	1.54	6.53	45.16	0.5
Age <45	-0.51	1.47	-0.53	65.8	1.06	5.68	26.99	0.7
Age 45–64	-1.38	0.62	-3.63	61.5	0.92	6.06	47.41	0.3
Age 65+	-2.11	1.58	-1.91	69.6	3.10	8.19	54.96	0.7
Wealth p0–50	-1.40	2.53	-0.63	87.5	5.51	4.66	19.14	6.2
Wealth p50–90	-1.12	0.39	-3.71	76.5	1.01	6.19	45.29	0.4
Wealth top 10%	-2.07	1.89	-1.98	55.2	0.90	8.53	52.75	0.2

Notes: Wealth-tertile analogue of Table 3.3. See notes there for definitions.

Table 3.11.: Decomposition of R/Y, by income tertile, post-2014

	Stocks & mutual funds				Housing			
	$\frac{R}{Y}$ (%)	$\frac{A}{Y}$	$\frac{R+U}{A}$ (%)	APR (%)	$\frac{R}{Y}$ (%)	$\frac{A}{Y}$	$\frac{R+U}{A}$ (%)	APR (%)
Overall	1.77	1.64	3.63	29.6	2.11	6.68	39.53	0.8
Age <45	0.14	3.65	0.35	11.4	1.89	5.79	18.41	1.8
Age 45–64	0.95	0.66	6.95	20.6	1.61	6.08	37.77	0.7
Age 65+	3.55	1.72	5.81	35.5	2.91	8.08	52.37	0.7
Income p0–50	3.62	7.89	1.52	30.1	6.87	13.81	43.74	1.1
Income p50–90	1.75	0.87	6.79	29.6	1.50	5.89	39.48	0.6
Income top 10%	1.13	0.60	6.44	29.1	0.58	4.07	30.91	0.5

Notes: Same specification as Table 3.3, restricted to financial reference years 2014–2023.

Table 3.12.: Sale frequency, conditional sale size, and APR: stocks vs. housing, by wealth tertile

	Stocks & mutual funds			Housing		
	Sale freq. (%)	Cond. size (€K)	APR (%)	Sale freq. (%)	Cond. size (€K)	APR (%)
Overall	46.5	8.0	65.4	1.0	255	0.5
Age <45	38.0	4.7	65.8	1.4	226	0.7
Age 45–64	46.8	8.0	61.5	0.7	241	0.3
Age 65+	52.0	9.6	69.6	0.9	315	0.7
Wealth p0–50	39.6	4.4	87.5	3.8	229	6.2
Wealth p50–90	46.8	7.2	76.5	0.5	275	0.4
Wealth top 10%	47.5	11.5	55.2	0.3	383	0.2

Notes: Wealth-tertile analogue of Table 3.2.

Table 3.13.: Positive-gain APR (APR^+) and loss realization rate (LRR), stocks, by wealth tertile

	Gains ($R + U > 0$)			Losses ($R + U < 0$)		
	APR (%)	Avg. gain (€K)	N	LRR (%)	Avg. loss (€K)	N
Overall	14.3	2.4	5537	15.0	-3.7	4497
Age <45	11.4	1.0	913	10.3	-1.4	824
Age 45–64	13.3	2.1	2308	10.2	-3.3	1835
Age 65+	15.2	3.4	2316	18.7	-5.1	1838
Wealth p0–50	15.2	1.1	1080	20.4	-2.0	840
Wealth p50–90	14.2	2.0	3187	10.7	-3.0	2680
Wealth top 10%	14.2	4.6	1270	18.4	-7.2	977

Notes: Wealth-tertile analogue of Table 3.8.

Table 3.14.: Positive-gain APR (APR^+) and loss realization rate (LRR), stocks, by income tertile, post-2014

	Gains ($R + U > 0$)			Losses ($R + U < 0$)		
	APR (%)	Avg. gain (€K)	N	LRR (%)	Avg. loss (€K)	N
Overall	14.4	3.8	2337	8.0	-1.8	1166
Age <45	7.8	1.7	350	2.6	-1.2	200
Age 45–64	11.6	3.4	870	11.7	-1.3	400
Age 65+	16.7	4.8	1117	7.6	-2.3	566
Income p0–50	15.8	3.7	735	7.8	-1.7	376
Income p50–90	14.9	3.5	1290	8.6	-1.7	648
Income top 10%	10.6	5.2	312	6.5	-2.5	142

Notes: Same specification as Table 3.8, restricted to financial reference years 2014–2023.

3.B. Position Selection: Supporting Material

This appendix collects three extensions of the position-level evidence in Section 3.4: the binned scatter underlying the sale-probability regression, the household-year aggregation of the winner-bias, and the cross-sectional heterogeneity in trade intensity.

3.B.1. Binned scatter: sale probability vs. relative return

Figure 3.8.: Sale probability as a function of relative portfolio return.

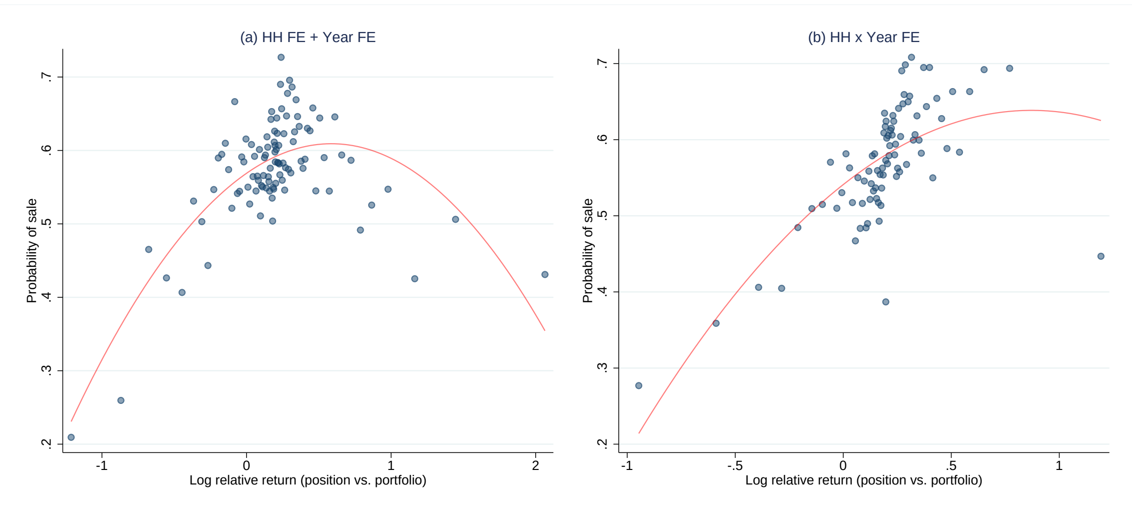

Notes: Binned scatter of $\Pr(\text{sale}_{ijt} = 1)$ on $\log(\Pi_{ijt}/\bar{\Pi}_{-j,it})$, residualized on household and year fixed effects. Bins are equal-mass.

3.B.2. Household-year seller netting

The position-level pattern aggregates predictably to the household-year. Table 3.15 classifies seller household-years by the realized-gain composition of their sold positions. Under a level cut ($R_j \gtrless 0$), 54 % of seller-years sell only winners, 27 % sell only losers, 16 % sell a mix of both, and the remaining 4 % are unclassified for lack of position-level signing. Under a relative-return cut ($\log(\Pi_j/\bar{\Pi}_{-j}) \gtrless 0$, the regressor of Equation (3.7)), the winners-only category narrows because the bar is now higher than zero, but the rank ordering is preserved.

Table 3.15.: Seller netting: winners vs. losers (household-year level, stocks)

(a) Level classification: $R_j \gtrless 0$

	N	Share (%)	Mean gain (EUR)	Mean loss (EUR)	Mean net (EUR)
Winners only	1,142	53.7	1,458	0	1,458
Losers only	568	26.6	0	-3,325	-3,325
Both	331	15.8	1,546	-2,369	-822
Unclassified	78	3.9	0	0	0
All sellers	2,119	100.0	1,028	-1,259	-231

(b) Relative-return classification: $\log(\Pi_j/\bar{\Pi}_{-j}) \gtrless 0$

	N	Share (%)	Mean gain (EUR)	Mean loss (EUR)	Mean net (EUR)
Winners only	986	46.1	-345	0	-345
Losers only	209	10.0	0	-832	-832
Both	378	18.1	979	-1,588	-609
Unclassified	546	25.8	0	0	0
All sellers	2,119	100.0	18	-370	-352

Notes: Each seller household-year is classified by whether its sold positions are winners, losers, or both. Panel (a) uses the level cut; Panel (b) uses the relative-return cut from Equation (3.7). Mean gain, loss, and net columns report level realized gains in euros even under the relative-return cut, so a relative winner can carry a negative level gain when the rest of the portfolio lost more. "Unclassified" captures household-years with zero or missing position-level realized gains; in Panel (b) it additionally absorbs single-stock households for whom no counterfactual portfolio return is defined. Shares are population-weighted; realized gains winsorized at 1st/99th percentile.

The composition of seller-years also varies with the size of the position pool. Households holding many stocks tend to be partial sellers; households holding few tend to be full liquidators. Among full liquidators, 57 % are *market exiters*—they liquidate every held position and leave the asset class entirely—while the remaining 43 % are *position cleaners*, retaining at least one position. Table 3.16 reports the position-count distributions and sub-row breakdowns by aggregate sign of $R + U$.

Table 3.16.: Positions held by seller type, with sign of $R + U$ split (household-year level, stocks)

	N	Share (%)	Mean	Median	p25	p75	Mean sells
Full liquidator	175	4.1	2.6	2	1	3	1.9
$R + U > 0$	79	1.9	2.4	1	1	3	1.8
$R + U < 0$	90	2.1	2.9	2	1	4	2.1
Market exiter	99	2.4	1.6	1	1	2	1.6
Position cleanup	76	1.7	4.0	3	2	5	2.4
Partial seller	1,944	44.8	3.2	2	1	4	1.6
$R + U > 0$	1,168	27.1	2.8	2	1	4	1.6
$R + U < 0$	745	16.9	3.8	3	2	5	1.7
Non-seller	2,227	51.1	1.9	1	1	2	0.0
$R + U > 0$	1,237	28.2	1.8	1	1	2	0.0
$R + U < 0$	872	20.1	2.0	1	1	2	0.0
All stockholders	4,346	100.0	2.5	2	1	3	0.8

Notes: Full liquidator: at least one position fully sold. *Partial seller:* sold shares but no position fully liquidated. Sub-rows split household-years by the aggregate sign of $R + U$ at the household level. For full liquidators, additional sub-rows separate market exiters ($n_{\text{full liq}} = n_{\text{held}}$) from position cleanups. Shares are population-weighted; main and sub-rows nest. Mean, median, p25, p75 refer to the number of positions held with positive year-end value or sold during the year.

3.B.3. Trade intensity across households and positions

Household-year trade intensity rises with age, stock-portfolio value, and gross income. Figure 3.9 plots the average number of trades per household-year against four characteristics. The relationship is steeply rising in stock-portfolio value (from roughly 0.7 trades per year at the bottom to 1.7 at portfolios of €45,000) and in age (from below 0.1 at age 30 to roughly 0.4 at age 80), more modestly rising in log gross income, and essentially flat in average realized return. The aggregate time series of trades per household-year, reported in Table 3.17, shows substantial year-to-year variation tracking market conditions.

Figure 3.9.: Trades per household-year by household characteristics.

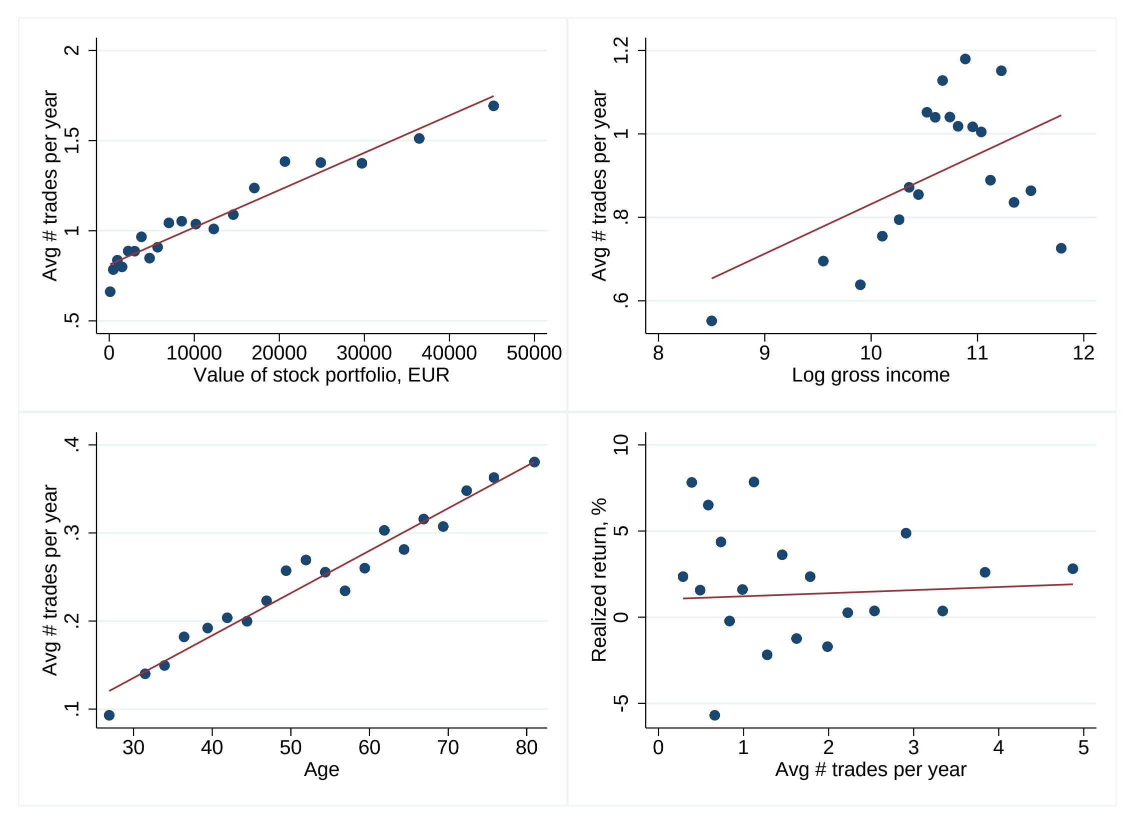

Notes: Binned scatter of the number of stock trades per household-year against the value of the stock portfolio (top left), log gross income (top right), age (bottom left), and the relationship between realized return and trade frequency (bottom right). Bins equal-mass; year fixed effects partialled out.

At the position level, trade frequency is associated with realized losses. Figure 3.10 re-plots intensity at the position-year level: the age gradient persists (trades per position rise from 0.5 at age 30 to 0.8 at age 80), the income gradient flattens, and positions with lower realized returns exhibit higher trade frequency. The negative slope between trades per position and realized return is consistent with partial-liquidation activity concentrating on losing positions, while winners tend to be liquidated in single transactions—a pattern aligned with the seller-netting evidence of Appendix 3.B.2, where full liquidators were over-represented in the winners-only category.

Figure 3.10.: Trades per position by household and position characteristics.

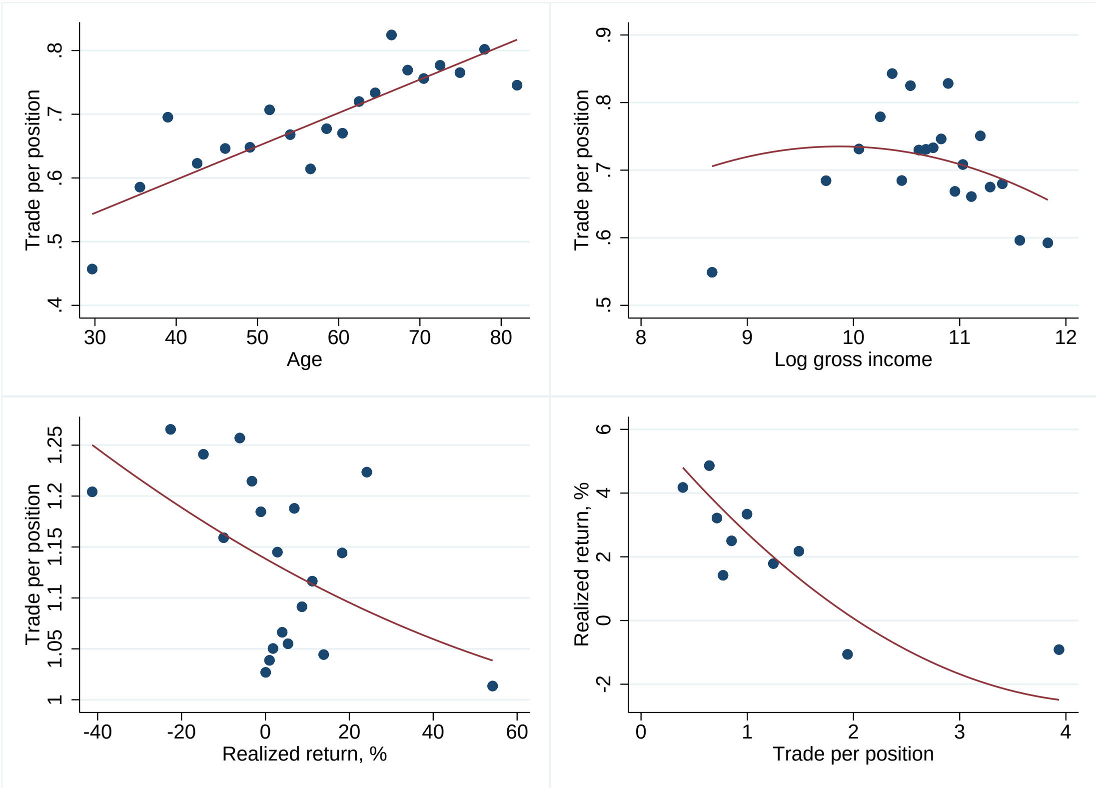

Notes: Binned scatter of the number of trades per position-year against age (top left), log gross income (top right), realized return on the position in percent (bottom left), and the converse view of trades per position against realized return (bottom right). Bins equal-mass; year fixed effects partialled out.

Table 3.17.: Trades per household-year, by year

	Mean	P10	Median	P90	SD	N
Trades per year	0.77	0.00	0.29	2.25	1.23	15911
Average trades per year, over life	0.77	0.00	0.35	2.19	1.07	15911
Trades per year, purchases	0.39	0.00	0.08	1.20	0.68	15911
Trades per year, sales	0.39	0.00	0.13	1.12	0.62	15911

Notes: Sample mean of the number of stock trades per household-year, restricted to stockholding household-years. Population-weighted.

3.C. Flow-of-Funds Robustness Specifications

This appendix collects the robustness specifications referenced in §3.6.3. Tables 3.18 and 3.19 report the marginal-allocation regressions with each source and use scaled by lagged household wealth. The wealth-scaled specification is the preferred cut for renters

(§3.6.3) because it removes the cross-sectional dispersion in portfolio scale. Adding-up holds within rounding error in every row.

Table 3.18.: Marginal allocation, homeowners — scaled by lagged wealth

	(1) ΔLiquid	(2) Stock inv. (I^S)	(3) Housing inv. (I^H)	(4) Debt service (b)	(5) Consumption
Stock sales / L.wealth	−0.280*	−0.151	−0.084	1.247	0.268
	(0.148)	(0.156)	(0.055)	(0.940)	(0.843)
Housing sales / L.wealth	0.017	0.002	0.006	−0.585***	1.559***
	(0.016)	(0.002)	(0.007)	(0.149)	(0.143)
Non-capital income / L.wealth	0.069*	−0.006	0.026***	−0.184	1.095***
	(0.035)	(0.005)	(0.009)	(0.221)	(0.213)
Interest/dividends / L.wealth	−0.638	0.071	0.667**	−0.877	1.777
	(0.412)	(0.060)	(0.313)	(5.735)	(5.929)
N	12528	12528	12528	12528	12528
R^2	0.059	0.022	0.301	0.032	0.347

Notes: As in Table 3.6, with each variable divided by lagged total household wealth. Sample: homeowner-years.

Table 3.19.: Marginal allocation, renters — scaled by lagged wealth

	(1) ΔLiquid	(2) Stock inv. (I^S)	(3) Debt service (b)	(4) Consumption
Stock sales / L.wealth	−0.244*	1.097***	−0.442***	0.590**
	(0.146)	(0.191)	(0.133)	(0.229)
Non-capital income / L.wealth	0.045***	0.002*	−0.014	0.966***
	(0.010)	(0.001)	(0.016)	(0.019)
Interest/dividends / L.wealth	0.014	−0.004	0.003	0.988***
	(0.142)	(0.021)	(0.107)	(0.197)
N	3535	3535	3535	3535
R^2	0.030	0.335	0.001	0.755

Notes: As in Table 3.7, with each variable divided by lagged total household wealth. Sample: renter-years.

Local projections: dynamics around realization

To check that the static estimates in §3.6.3 are not picking up offsetting flows in adjacent years, we estimate local projections of each use $Z_{j,t+h}$ on the full source vector X_t for horizons $h \in \{-2, -1, 0, +1, +2\}$, with household and year fixed effects and standard errors clustered at the household level. The contemporaneous coefficients ($h = 0$) reproduce the static rows of Tables 3.6 and 3.7; the surrounding horizons describe the timing of the response. Figures 3.11 and 3.12 report the full source/use grids for owners and renters respectively, including labour income W, dividends and interest D^S, and other non-capital income D^H. The dynamics show no systematic pre-trends and the post-impact responses are mean-reverting, consistent with the static specification capturing the bulk of the marginal-allocation pattern.

Figure 3.11.: Local projections of uses on sources, homeowners — full source/use grid.

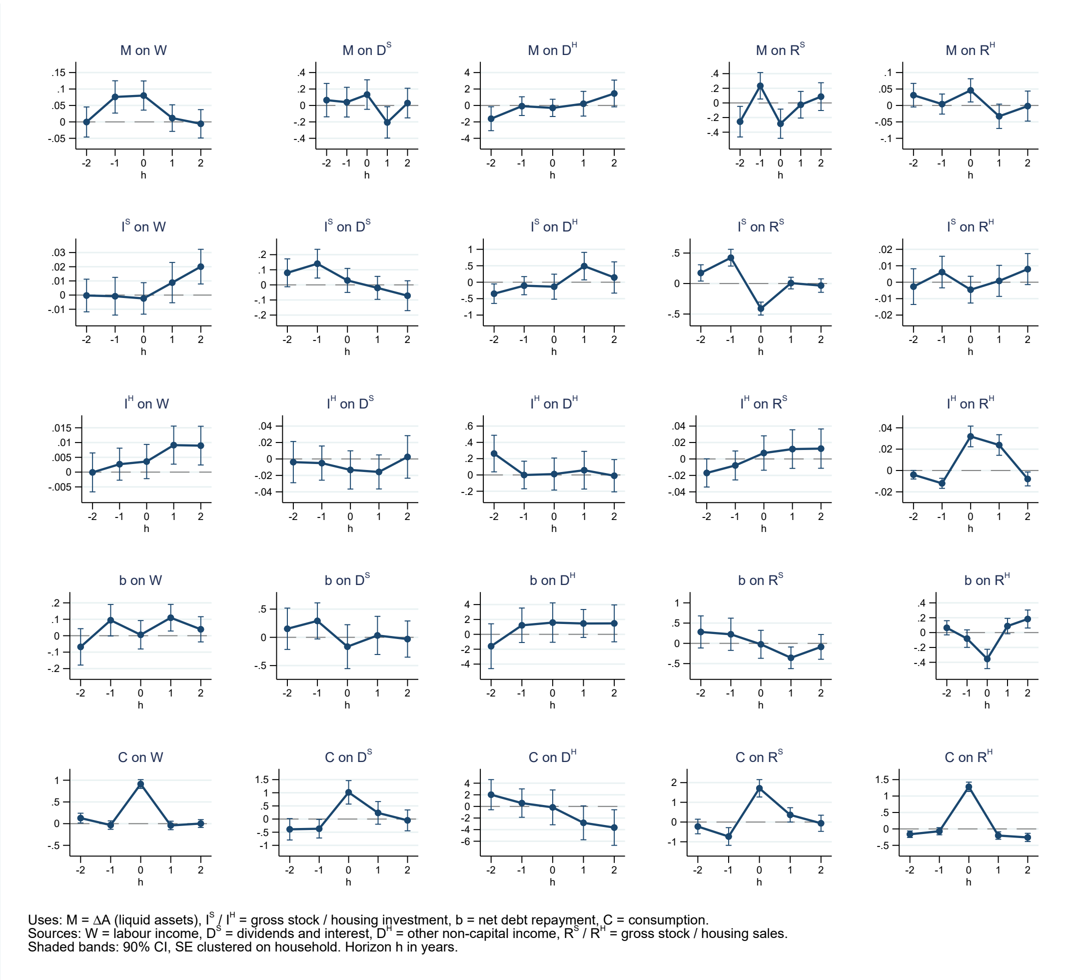

Notes: Each panel plots the coefficient on the column source at horizon h in a regression of the row use $Z_{j,t+h}$ on the full source vector X_t, with household and year fixed effects. Whiskers are 90% CIs from household-clustered SE. Sample: homeowner-years.

Figure 3.12.: Local projections of uses on sources, renters — full source/use grid.

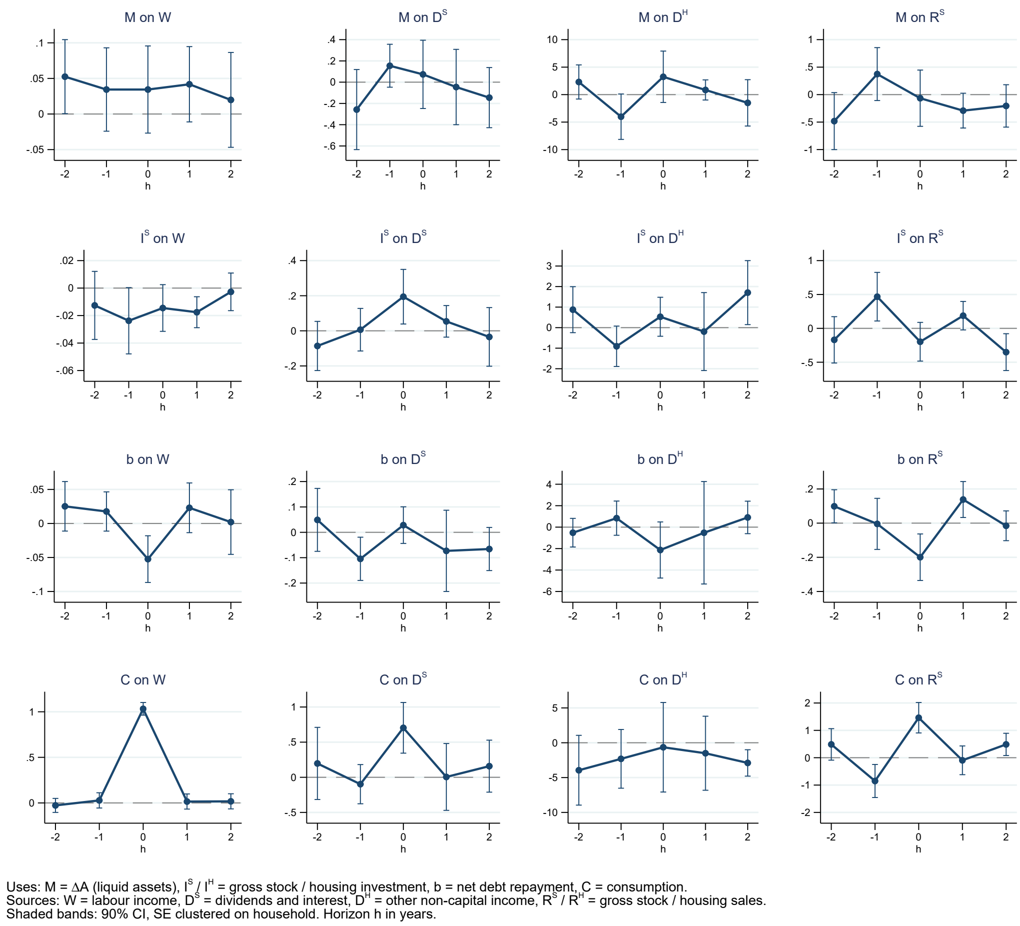

Notes: As in Figure 3.11, restricted to renter-years and to the renter source/use system.

3.D. DHS Sample, Panel, and Portfolio Diagnostics

Figure 3.13.: Panel structure of the DHS sample.

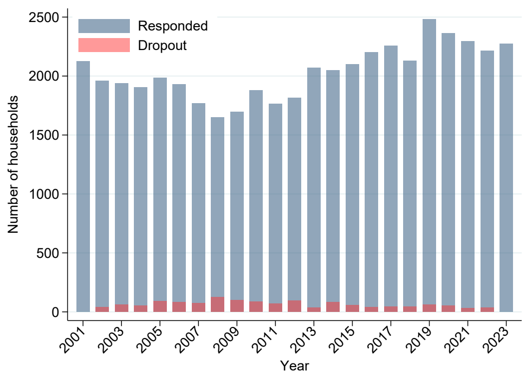

(a) Panel composition over time

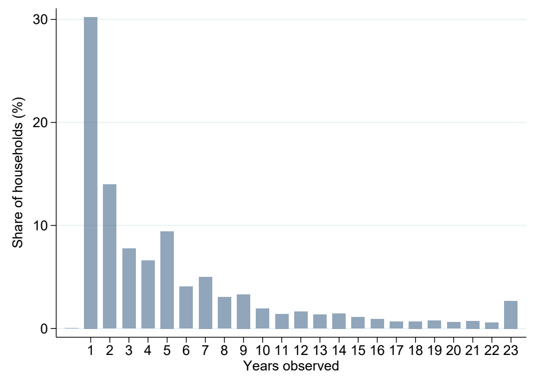

(b) Years observed per household

Notes: Left: number of households observed per wave. Right: distribution of waves per household across all DHS waves used. Sample: 2001–2023.

Figure 3.14.: Homeownership: Netherlands and United States.

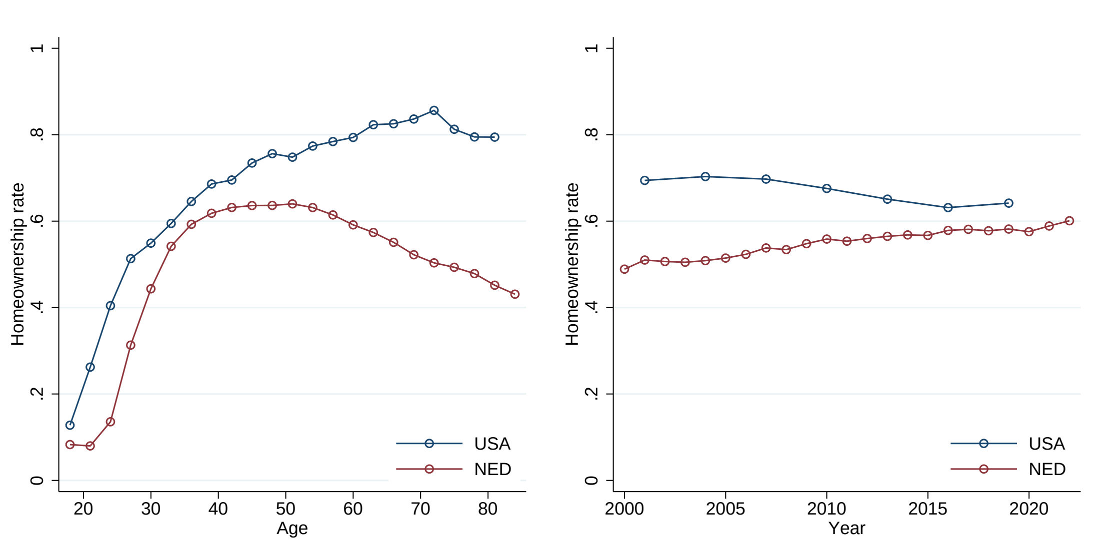

Notes: Lifecycle (left) and time-series (right) profiles of homeownership. Dutch profile estimated from the DHS; U.S. profile from the SCF. Companion to Figure 3.2 in the main text.

Figure 3.15.: Distribution of the number of positions held per household-year.

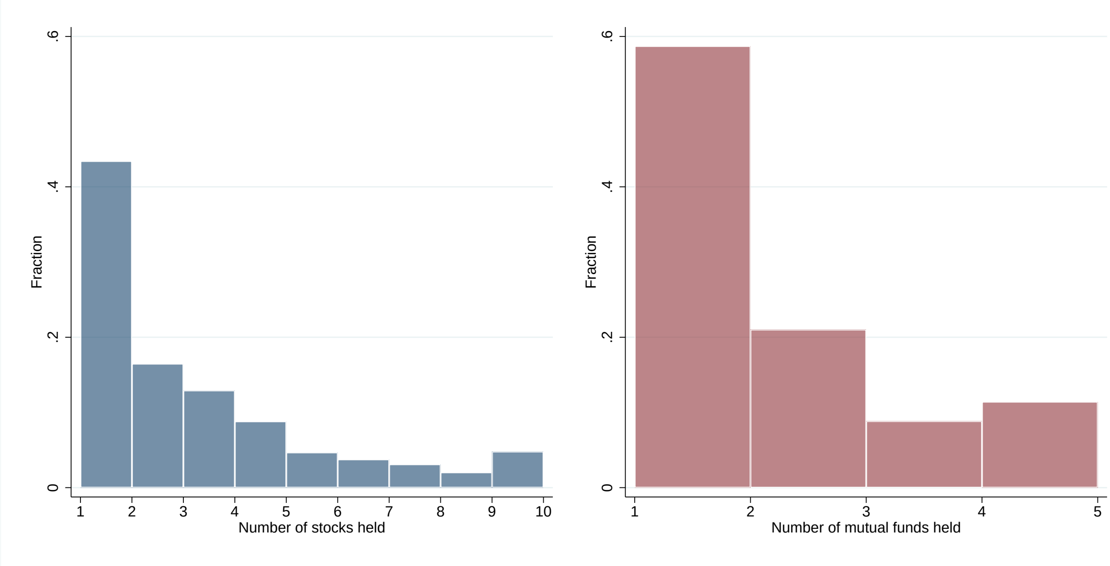

Notes: Stocks (left) and mutual funds (right). DHS records up to ten stock positions and up to five mutual fund positions per household.

Table 3.20.: Summary statistics by age group

	Age 30	Age 40	Age 50	Age 60	Age 70
Risky asset participation (%)	17.5	20.0	23.2	23.8	23.6
Direct stock participation (%)	8.8	9.9	12.1	11.4	11.6
Mutual fund participation (%)	12.4	15.3	17.2	19.1	17.7
Other instruments participation (%)	6.1	6.5	6.9	6.3	9.1
Bond participation (%)	1.6	2.7	3.4	4.0	6.4
Deposit participation (%)	87.9	83.7	81.1	83.0	85.5
Insurance product participation (%)	19.6	28.3	32.1	26.2	5.1
Any financial instrument (%)	32.5	37.5	42.3	39.3	28.7
Entry: risky assets (%)	1.2	1.2	1.3	0.9	0.9
Exit: risky assets (%)	1.2	1.7	1.8	1.9	1.4
Entry: direct stocks (%)	0.7	0.6	0.8	0.6	0.6
Exit: direct stocks (%)	0.6	0.9	0.9	1.0	0.8
Entry: mutual funds (%)	1.0	1.2	1.2	0.8	0.9
Exit: mutual funds (%)	1.1	1.7	1.6	1.7	1.4
Total debt (EUR)	86694.6	90378.2	65626.7	47444.1	34210.7
Leverage: debt/assets (%)	199.6	218.1	100.7	54.2	16.9
Leverage: debt/income (%)	345.7	528.7	303.1	182.0	127.0
Value: stocks and housing (EUR)	141978.4	189745.3	200042.2	211298.9	210615.4
Business assets (EUR)	3786.7	4216.9	8994.4	5159.4	1115.6
Non-business assets (EUR)	169864.0	232961.2	252451.4	274038.9	270017.2
Stocks-housing share of non-bus. wealth (%)	48.6	52.6	51.1	49.8	49.3
Employer-sponsored share (%)	2.2	1.6	1.8	0.9	0.1
Risky share of financial wealth (%)	5.6	7.7	8.9	10.5	10.3

Table 3.21.: Summary statistics by year

	2001	2005	2008	2015	2022
Risky asset participation (%)	27.9	22.7	25.2	17.7	21.9
Direct stock participation (%)	13.8	11.7	10.6	8.1	12.2
Mutual fund participation (%)	21.5	17.2	21.4	13.3	16.4
Other instruments participation (%)	12.5	10.6	6.9	3.3	4.4
Bond participation (%)	3.1	4.0	4.0	1.8	2.2
Deposit participation (%)	81.9	84.2	86.2	86.9	85.4
Insurance product participation (%)	32.4	29.7	27.5	17.3	15.4
Any financial instrument (%)	49.0	45.1	43.0	27.9	32.5
Entry: risky assets (%)	1.0	1.2	3.8	0.9	1.3
Exit: risky assets (%)	1.2	1.5	2.3	0.8	0.8
Entry: direct stocks (%)	0.7	0.9	1.8	0.4	0.9
Exit: direct stocks (%)	0.7	1.1	2.0	0.7	0.5
Entry: mutual funds (%)	1.0	1.1	4.4	0.8	0.8
Exit: mutual funds (%)	1.3	0.9	2.1	0.6	0.7
Total debt (EUR)	87921.9	48539.9	54489.8	70443.4	81090.6
Leverage: debt/assets (%)	211.6	81.4	85.8	134.4	81.5
Leverage: debt/income (%)	385.8	404.4	170.8	249.1	504.5
Value: stocks and housing (EUR)	286248.3	146948.8	160647.2	173537.8	274204.4
Business assets (EUR)	4363.8	1480.6	4225.8	4844.8	8061.4
Non-business assets (EUR)	323341.0	184675.1	203726.2	223491.4	334964.0
Stocks-housing share of non-bus. wealth (%)	49.0	48.5	46.1	51.5	51.3
Employer-sponsored share (%)	4.9	2.5	2.4	0.2	0.0
Risky share of financial wealth (%)	11.2	9.8	9.3	6.8	9.0

Notes: Population-weighted means (using time-averaged household weights). Year refers to the financial reference year (end-of-year balances). Sample restricted to ages 20–85. *Entry / Exit:* year-over-year $0 \to 1$ / $1 \to 0$ transitions. All other variable definitions follow Table 3.1.