Essays in Macroeconomics and Fiscal Policy

Tobias Kawalec

Nuffield College

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy in Economics

Trinity Term 2026

Supervised by Martin Ellison and Michael McMahon

Abstract

Fiscal and monetary policy do not operate in a vacuum of complete independence of each other; in fact, they are wholly *interdependent*. This thesis is concerned with a thorough analysis of the interactions between monetary and fiscal policymakers and their ramifications for macroeconomic dynamics in business-cycle frequency.

The first and second chapters centre around the role of inflation-indexed debt; that is, government-issued obligations whose face value changes itself with inflation. Despite their significant presence in the United Kingdom, where indexed debt makes up 25.8% of the total market value of outstanding UK government debt as of end-2025, such obligations have thus far not been explored in macroeconomic models of business cycle fluctuations.

The first chapter introduces the concept of inflation-indexed debt into considerations of monetary and fiscal policymakers, highlighting how inflation-indexed debt can counter-intuitively *boost* inflation, ceteris paribus, in response to deficit-financed fiscal shocks. This insight is formalized in a heterogeneous-agent New Keynesian (HANK) framework, through which I quantitatively evaluate the response of the economy to deficit-financed fiscal shocks, finding that the link between deficits and inflation is amplified through the presence of indexed debt, conditional on monetary and fiscal policy aligning so as to not nullify the the wealth effect enjoyed by households in the face of this fiscal expansion.

The second chapter moves away from the quantification of the effects of indexed debt and instead zooms directly into the link between inflation-indexed debt and the concept of fiscal dominance. Through a more tractable New Keynesian-Overlapping Generations framework, I highlight how the precise formulation of fiscal rules can limit the scope for monetary policymakers to engage in stabilizing policy. Global data on the presence of

inflation-indexed debt and monetary policy independence hints at the relevance of the uncovered channel linking indexed debt to the possibility of a fiscally-driven restraint of monetary policy in certain scenarios.

The third chapter looks in greater detail at the international transmission of fiscal shocks. I uncover how effective discount rates, which matter for the valuation of government debt, must adhere to fundamental insights known from international finance, through which foreign discount rate movements matter for the valuation of domestic government debt. Consequently, policy impulses from abroad can be imported and matter for the sustainability of the budget balance. I formalize this insight using a VAR-based decomposition exercise, as well as through a model of a small open economy rationalizing these spillovers.

While the first three chapters effectively focus on departures of the aggregate-demand side of the economy from conventional frameworks, the fourth chapter of the dissertation (co-authored with Michael McMahon) looks at a novel way of modelling the supply side, which re-introduces inventories as a relevant amplification mechanism for cyclical fluctuations. We propose a novel model of inventories on supply chains, highlighting the possibility of temporary imperfect equilibria by allowing goods markets not to clear exactly, such that pent-up demand can exist and accumulate in a net inventory/order book variable. We show how the bullwhip effect arises as a state-dependent phenomenon conditional on the source of fluctuations and the stance of monetary and fiscal policy.

Declaration

Supervision: the writing of this thesis was supervised by Prof Michael McMahon and Prof Martin Ellison.

Authorship: the first three chapters of this thesis are single-authored. The fourth chapter is co-authored with Michael McMahon.

Funding: throughout the creation of this dissertation, I obtained funding from Nuffield College, the Department of Economics of the University of Oxford, and the David Walton Memorial Fund. Their support is greatly appreciated.

Total words: this thesis contains approximately 60006 words. This was calculated by counting the words on a representative page (page 9, 411 words) and multiplying by the page count, excluding the preamble, bibliography, and appendices (146 pages).

Acknowledgements

My gratitude to all the people that accompanied me on the journey to this DPhil thesis cannot possibly be captured by mere words. First and foremost, I remain deeply indebted to my supervisors, Martin Ellison and Michael McMahon. Their guidance, rigorous feedback, and overall support shaped me to become the economist that I am today, allowing me to refine my at-times unusual ideas.

Beyond the two of them, I want to thank Francesco Zanetti for his useful comments on some of my projects. In addition, I thank my Oxford-based co-author, Chris Hyland, for the countless hours of discussions on fiscal policy, monetary policy, financial markets, and so much more. I am more than happy to not only call you a co-author, but also a great friend.

Throughout the years, the numerous conversations with the Oxford Macroeconomics Group were a further source of inspiration that I am very grateful for. I also extend my gratitude to the many conference and seminar participants that I encountered in the last 3.5 years.

Special thanks are to be given to Nuffield College for being such an amazing source of support during my DPhil; this includes all staff members whom I have seen so regularly. I am very thankful to Nuffield College, the Department of Economics, and the David Walton Memorial Fund for their financial support that made this DPhil possible.

All my friends, my extended family, my brother, and my father have been a great source of support throughout my time in Oxford. My appreciation for them cannot be done justice by such few words.

Finally, I must also thank my fiancée, Pia, for enduring years of our relationship being long-distance and for being the greatest support imaginable. You helped me to keep my head up high even in the low moments.

I wholly dedicate this thesis to my mother, who passed away from cancer shortly after I arrived in Oxford. She, alongside my brother and my father, has been the single largest source of inspiration for me, as she was the one who always pushed for my education despite economic hardship.

Dominus illuminatio mea et salus mea: quem timebo?

Contents

List of Figures

List of Tables

Chapter 1

Debt Indexation, Determinacy, and Inflation*

Abstract

Contrary to popular belief, inflation-indexed government debt can boost inflation in response to deficit shocks, conditional on a lack of sufficient future fiscal backing. We formalize this insight through a state-of-the-art calibrated HANK model with multiple asset types, showing that the annual inflationary effect of a 1% deficit-to-GDP shock increases by 0.35 percentage points when 30% of the government debt stock is indexed to inflation, as is the case in the United Kingdom, relative to a baseline case calibrated to the United States. Inflation-indexed debt makes the price level partially backward-looking through the government debt valuation equation, thereby causing additional inflationary pressure. Empirical evidence from a large, narratively identified fiscal deficit shock supports this finding, which has additional implications for the distinction between 'fiscally-led' mechanisms and 'HANK-type' mechanisms surpassing Ricardian equivalence.

*I thank Margaret Jacobson (discussant), Nico Caramp, Jesús Fernández-Villaverde, Carlo Galli, Marco Garofalo, Alex Haas, Joe Hazell, Chris Hyland, Rustam Jamilov, Bartosz Maćkowiak, Zach Mazlish, Ben Moll, Jordan Roulleau-Pasdeloup, Dejanir Silva, Francesco Zanetti, and participants of the AMSE Doctoral Workshop on Quantitative Dynamic Economics 2024, the Oxford Macroeconomics Working Group Seminar, the CEPR Paris Symposium 2024, the RGS PhD Workshop 2025, the 32nd Annual Symposium of the Society for Nonlinear Dynamics and Econometrics, the ifo Dresden Workshop on Macroeconomics and International Finance 2025, the Midwest Macro Meeting 2025, the GRAPE Summer Workshop on Macro & Finance 2025, the CEBRA 2025 Annual Meeting, the European Winter Meeting of the Econometric Society 2025, and ASSA 2026 for comments.

1.1 Introduction

Rival theories, such as labour-market tightness, supply-chain disruptions, or fiscal deficits and the associated monetary response, contribute to our understanding of the 2021–2023 inflationary episode. This chapter adds a novel, yet overlooked aspect to this ongoing debate: the *composition* of government debt, with a particular distinction between purely nominal debt and *inflation-indexed debt*, which differs from the commonly modelled government obligations as its face value changes with the gross rate of inflation. As the government budget balance plays a crucial role for aggregate demand, and thereby co-determines the price level, an interesting feedback loop arises when the face value of a part of government debt itself changes with gross inflation. In this chapter, we show the qualitative and quantitative importance of this mechanism, which is powerful conditional on fiscal deficits not being completely backed by future surpluses.

Indexed debt is not a mere theoretical curiosity. Figure 1.1 shows the share of inflation-indexed debt as part of the overall sovereign debt stock over time in a number of countries. While there is considerable heterogeneity across countries, indexed bonds are present across the board, and have been so for the past three decades. This chapter mostly focuses on the United Kingdom (U.K.) and the United States (U.S.), since these two indexed debt markets are among the largest ones in both absolute and relative terms.

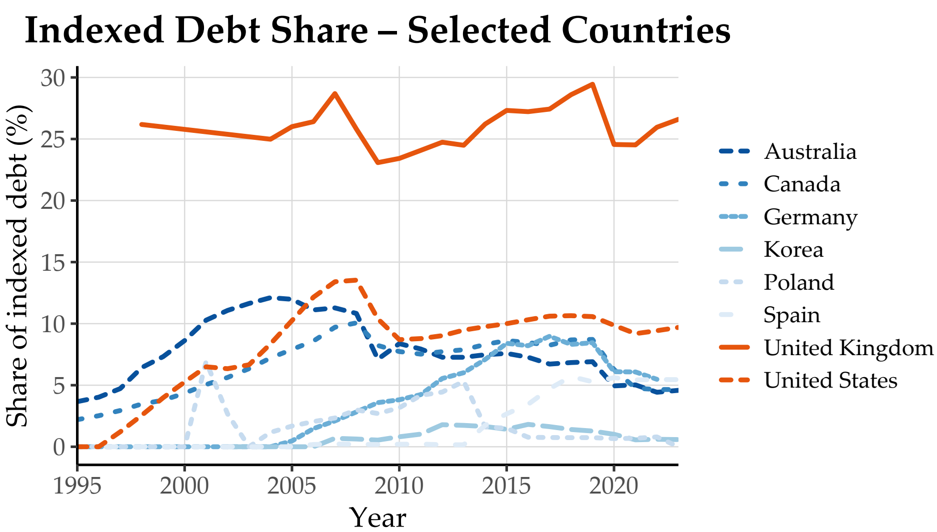

Figure 1.1: The market value-weighted share of inflation-indexed debt in sovereign debt portfolios over time. Data source: BIS (2024).

Analysing the role of the government debt structure for inflation requires a delicate treatment of the interactions between fiscal and monetary policy. While a *fiscally-led policy mix* (that is, a fiscal authority committing to policy *not* adhering to its intertemporal budget constraint), is not a prerequisite for the analysis of the role of deficits for inflation, it enhances the role of fiscal policy as drivers of inflationary dynamics in macroeconomic general equilibrium models (Leeper, 1991; Sims, 2011; Ascari et al., 2023; Rachel and Ravn, 2025). We therefore motivate this chapter further with a real-world example of fiscal policy committing to unfunded deficits and plausibly forcing the hand of the monetary policy authority (and thus informing a possibly fiscally-led policy mix): the 'U.K. mini-budget' in September 2022.[1]

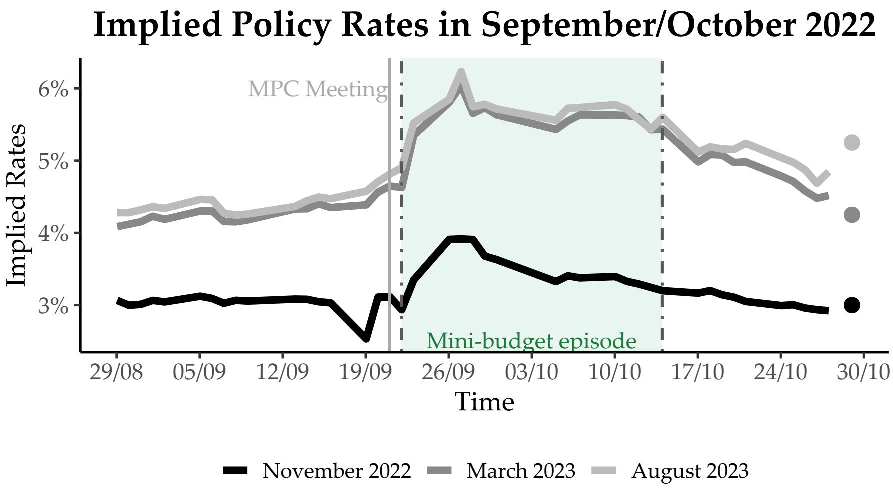

Figure 1.2: Expectations of Bank rates in the United Kingdom for three scheduled MPC meetings after the 'mini-budget' announcement in September 2022, defined as the break-even rates under which Overnight Indexed Swaps are fairly priced at the prevailing forward interest rate curve (Bloomberg data series GB0BFR). The dots at the end depict the factual policy rates after each of the meetings.

Figure 1.2 plots average market-implied expectations of the 'Bank rate', the major monetary policy rate set by the Bank of England, in the window around the 'mini-budget' announcement and its subsequent cancellation. The first solid line depicts the date of a Bank of England Monetary Policy Committee (MPC) meeting, which occurred ahead of the announcement of the 'mini-budget' fiscal policy measure, with the MPC minutes being released on the 22nd of September 2022, one day ahead of the fiscal policy announcement. This timing is useful for the argument insofar as the meeting likely communicated the Bank of England's stance on future rate changes clearly, taking all available information up to that point into account (Braun et al., 2024). Nonetheless, implied policy rates rose sharply in the subsequent days, just after the announcement of the mini-budget (denoted by the first dotted line), with the shift amounting to a 120bps peak increase in expected

[1]For a more detailed argument related to this fiscal shock, see Leeper (2023), NIESR (2022), and section 1.2 of this chapter.

policy rates for the upcoming twelve months ahead. After the scraping of the mini-budget (second dotted line), expected policy rates swiftly returned to their 'pre-shock' levels.[2]

This event resonates well with the possible idea of (at least partially) fiscally-led policy mixes: financial market participants clearly expected changes to the monetary policy stance beyond the very short term, plausibly in response to an announced fiscal policy measure, although risk or liquidity premiums likely also contributed to the observed movements.

We motivate the relevance of indexed debt for inflation further in the chapter, using data to narratively pin down the true shock component induced by the 'U.K. mini-budget', tracking it through market expectations on sovereign deficits and government bond price revaluations. We estimate a sizeable *deficit-inflation multiplier*, which grows with the share of inflation-indexed debt. Our second empirical finding, based on an exercise with local projections using exogenously identified shocks, shows that inflation-indexed debt boosts observed consumer price inflation in response to deficit shocks.

Next, we introduce inflation-indexed debt in a one-equation model to pin down the main mechanism by which indexed debt matters for the price level. It does so through the government debt valuation equation, in which the price level becomes a state variable. Thus, previous price levels matter for the determination of the current price level, even without further sources of stickiness in the economy. For such an economy with inflation-indexed debt, we prove the uniqueness of the corresponding dynamic stationary general equilibrium.

We then analyse the effects of fiscal deficit shocks in a setting featuring (i) non-Ricardian fiscal policy through imperfect risk-sharing among households and (ii) the presence of inflation-indexed debt. This framework is a heterogeneous-agent New Keynesian model á la Kaplan et al. (2018), which we solve making use of the methods pioneered in Auclert et al. (2021) to solve heterogeneous-agent models up to first-order in aggregate variables, while preserving heterogeneity with respect to the individual agents in the economy. We additionally pay attention to the different insurance properties of inflation-indexed debt in models with incomplete markets, accounting for the possible windfall gains that can be earned by governments, following Brunnermeier et al. (2024). Inflation-indexed debt matters quantitatively by increasing the volatility of inflation by 0.65 percentage points for each percentage point increase of the share of indexed debt in the government debt portfolio. In terms of the level impact, for an economy with a close to 30% share of inflation-indexed debt in the sovereign debt portfolio (as observed in the U.K.), the annual inflation rate increases by 0.35 percentage points in response to a 1% deficit-to-GDP shock relative to a baseline case calibrated to the United States (boasting a 8% share of

[2]The expected monetary policy response was partially driven by a concurrent funding mismatch in liability-driven investment strategies of defined-benefit pension funds that were closely tied to movements in yields of sovereign bonds. See Pinter (2023) for a detailed exposition of this point.

inflation-indexed debt), which is quantitatively highly relevant. We furthermore establish that the classic notions of 'active/passive monetary/fiscal policy', as derived by Leeper (1991), do not directly translate into the world with inflation-indexed debt, even though similarities in the determination of saddle path-stable equilibria prevail.

Literature Review

This chapter contributes to the burgeoning literature on fiscal-monetary interactions, pioneered in Sargent and Wallace (1981) and formalized through Leeper (1991). Initial contributions focusing on the possibility of a fiscally-led policy mix include Sims (1994) and Woodford (1995).[3] More succinct summaries of the literature are provided by Leeper and Leith (2016) and Cochrane (2023). Leeper et al. (2017), Bassetto and Cui (2018), Liemen and Posch (2022), Ascari et al. (2023), Bianchi et al. (2023), Miao and Su (2024), and Smets and Wouters (2024) provide advances in analysing fiscally-led policy mixes in OLG and New Keynesian models, while empirical support for the possibility of fiscally driven inflation has been developed in Jacobson et al. (2023), Cochrane (2022a), Cochrane (2022b), Chen et al. (2022), Cloyne et al. (2023), Ascari et al. (2024), and Barro and Bianchi (2025).[4] A narrative example of a specific fiscal shock informing inflation rates is provided by Hazell and Hobler (2024), who focus on the 2021 Georgia Senate election runoff.

The focus of models emphasizing the link between fiscal-monetary policy interactions and the price level recently shifted towards models with explicitly *non-Ricardian* fiscal policy, thereby making real interest rates endogenous to fiscal frameworks. This endogeneity is important insofar as fiscal-monetary interactions ultimately constitute criteria that constrain the *transversality condition* on government debt to hold for only *one* candidate price level, but the transversality condition itself depends on the real interest rate.[5][6]

[3] To avoid confusion, we do not explicitly define the term 'Fiscal Theory of the Price Level' (FTPL) in this chapter, which has initially been coined by Woodford (1995). Instead, we follow Brunnermeier et al. (2022), who define 'core aspects', or themes, commonplace in the FTPL literature, which broadly apply in this chapter as well. These aspects are: (a) governments issue some liabilities in their own currency, (b) the government faces a debt valuation equation over which it has some agency, (c) the price level plays a key role in ensuring debt sustainability in equilibrium, (d) additional assumptions, e.g., on fiscal-monetary policy mixes, complement equilibrium uniqueness, and (e) the framework could deliver price level uniqueness even without nominal frictions.

[4] The effects of monetary policy shocks, in turn, are also constrained and informed by the underlying fiscal reaction function, as shown by Bigio et al. (2024) and Caramp and Silva (2023, 2026).

[5] Brunnermeier et al. (2022) and Kaplan et al. (2024) provide conditions under which such models admit (unique) forward-looking equilibria expressed through the price level. Their notions of uniqueness are challenged by Hagedorn (2021, 2024), who argues that the endogeneity of the real interest rate in incomplete-markets models 'breaks' determinacy and allows a continuum of initial price levels to exist.

[6] The contribution of Niepelt (2004) highlights the possible pitfalls in specifying a model with *real* obligations and the fiscal authority ultimately pinning down the price level. Following Daniel (2007) and the distinction between real obligations and obligations indexed to the value of monetary-issued paper that we employ in this chapter, we hope to sideline these concerns.

Many such models fall under the category of heterogeneous-agent New Keynesian (HANK) frameworks. Brunnermeier et al. (2022), Kaplan et al. (2024), Campos et al. (2024), and Kwicklis (2025) have all applied fiscal price level determination to rich heterogeneous-agent frameworks. A second type of such models considers New Keynesian models with mortality frictions, exemplified by Angeletos et al. (2024b), Dupraz and Picco (2025), Nakamura et al. (2025), and Rachel and Ravn (2025). Angeletos et al. (2024b) negate the need to consider fiscally-led policy mixes when determining the price level, finding quantitatively identical responses of inflation to expansionary fiscal shocks in such NK-OLG models under monetary-led policy mixes. An overview over the state-of-the-art is provided by Kaplan (2025b), and a complementing succinct mathematical description of the link between non-Ricardian properties of households and fiscal-monetary interactions can be found in Kaplan (2025a).

Our contribution is to introduce a second type of assets (inflation-indexed debt) with a feedback loop between asset holdings and the price level, quantifying the importance that such indexed debt has for inflation dynamics in a calibrated state-of-the-art macroeconomic model. In doing so, we specifically pay attention to the different insurance properties borne by the two types of debt, thereby mitigating concerns related to misspecifications of aggregate transversality conditions (Brunnermeier et al., 2024). The specific insurance properties of inflation-related financial market instruments is also highlighted by Bahaj et al. (2023).

We also contribute explicitly to the literature on inflation-linked government bonds. Such bonds were introduced in economic and financial research long ago, especially in relation to the introduction of TIPS in the U.S. in 1997. One of the earliest contributions in this field is Fischer (1975), who derives household demand for such assets in a multi-asset framework. The special insurance properties of such inflation-linked debt are extensively discussed in Campbell and Shiller (1996), Barr and Campbell (1997), Garcia and van Rixtel (2007), Gürkaynak et al. (2010) and Andreasen et al. (2021), while Barro (2003) characterizes the optimal indexed-debt issuance strategy for a fiscal authority operating under a tax-smoothing objective, emphasizing the relevance of indexed consol bonds. Notably, Sims (2013) briefly mentions the possible detrimental consequences of indexed debt in fiscally-led policy frameworks. This chapter builds on his remarks, providing a rigorous framework nesting his intuitions. Schmid et al. (2026) provide a systematic analysis of inflation-indexed debt as a policy tool, emphasizing its role as an ex-ante commitment device against inflation. In the contribution, we leverage the unique properties of inflation-indexed debt, which express themselves mostly through the induction of a backward-looking component in the government budget equilibrium condition and through the insurance premia they bear. This chapter's focus thus effectively rests on the 'ex-post' effects that inflation-indexed debt can have in the face of expansionary government spending shocks.

In the later sections of the chapter, we rely on modern computational methods to efficiently solve and estimate heterogeneous-agent models, as in Kaplan et al. (2018) and Bayer and Luetticke (2020). In particular, we leverage the efficient computation algorithms pioneered in Auclert et al. (2021) and some of the refinements of Auclert et al. (2024b) to solve a model with heterogeneous households, two types of assets, and fiscal-monetary interactions.[7]

—

The rest of the chapter is structured as follows. Section 1.2 exposes the relevance of indexed debt for materialized and expected inflation in the face of fiscal shocks, after which we introduce inflation-indexed debt in a one-equation framework in section 1.3. We specify the main quantitative model in section 1.4. Section 1.5 discusses the calibration and the estimation of the model, and we present the quantitative findings in section 1.6. Section 1.7 concludes.

1.2 Empirical Evidence: Indexed Debt and Inflationary Dynamics

To motivate the relevance of indexed debt as a possible driver of the net present value of government debt and, therefore, of price level dynamics through the government debt valuation equation, we provide two pieces of evidence. First, we build on Hazell and Hobler (2024) and provide a narrative analysis of a large fiscal shock, but in an environment with high levels of indexed debt. We find decisively larger inflation multipliers in response to deficit shocks compared to theirs. Second, we employ a long-running series of exogenous fiscal policy shocks in a local projection to pin down the effects that inflation-indexed debt has on inflation itself when unexpected deficit-increasing policy measures occur.

Narrative evidence on the effect of deficits under indexed debt: the 2022 U.K. 'mini-budget'

We now provide narrative evidence on the effects of indexed debt through a cleanly identified fiscal policy shock: the September 2022 U.K. fiscal policy announcement, commonly dubbed the 'mini-budget'. We focus on this specific shock for two reasons: first, the event was largely unexpected in terms of its magnitude, allowing a clear identification of the effects of fiscal shortfalls on inflation. Second, this exercise is a complement to Hazell and Hobler (2024), who exploit probabilistic variation on Democrat Senate control around the 2021

[7] To motivate the relevance of household heterogeneity applied to holdings of sovereign debt, figure A.1 in the appendix provides evidence on the skew of household holdings of such debt, sorted by their respective income decile. This skew is even more pronounced for inflation-indexed debt.

Georgia Senate run-off election to infer the expected effects of expansionary fiscal policy on the price level. This chapter provides a similar exercise in an environment with high levels of inflation-indexed debt, complementing their estimates.

The institutional setup of U.K. fiscal policy serves as an excellent device for identifying the 2022 'mini-budget' episode as a clear fiscal shock. Fiscal policy in the U.K. is shaped by regular fiscal announcements, which set up the broad guidelines for expected sovereign income and spending in a given fiscal year. From 1980 to 2016, the larger 'budget announcement' usually occurred in early spring (coinciding with the beginning of a new fiscal year), supplemented by shorter budget statements in the fall of the same year. Between 2017 and 2019, the regular budget announcement was moved to fall, with the spring season being used usually for supplementary statements. Beginning in 2020, the main budget announcement started to take place in early spring again.

In spring 2022, then-Chancellor Rishi Sunak provided a budget statement, scheduled to be followed up by a full budget announcement in November 2022. In-between, and therefore outside of the usual bi-annual statement/announcement cycle, then-Chancellor Kwasi Kwarteng (who had since been appointed) presented a Ministerial Statement dubbed "The Growth Plan", with fiscal policy measures amounting to 150 Billion GBP, or approximately 5% of the GDP of the United Kingdom (NIESR, 2022). This statement did *not* constitute a budget announcement in the usual sense, being placed outside of the bi-annual statement cycle. The release of all budget statements made by the British government is usually supplemented with a concurrently released report by the Office for Budget Responsibility (OBR), an independent auditor supervising budgetary questions in the United Kingdom. In the case of the 'mini-budget', no such independent forecast of the budgetary consequences of the statement was publicly released, as the ruling government denied the release of the forecast created by the OBR.[8]

The episode of early fall 2022 is characterized by this fiscal policy announcement and its expected effects. In particular, the effects of the fiscal policy announcement are (in the very short-term) plausibly shielded from unrelated monetary policy news (both in terms of the interest rate level and in terms of the signalling of the state of the economy), since the preceding Bank of England Monetary Policy Committee decision was released one day *before* the announcement of the fiscal policy measure, on September 22, 2022.

[8] The forecast made by the OBR at that time has since been released, although it is only of limited relevance with respect to the eventual policy measures announced as the report was made 18 days ahead of the budget announcement, thus not capturing the full extent of the fiscal policy proposals. We therefore sideline this report for our analysis. The report can be found under: https://obr.uk/docs/dlm_uploads/FOI-Information-on-preparatory-work-for-the-mini-budget.pdf.

The size of the shock

The most important question is the *size* of the fiscal policy shock, which is *not* equal to the overall size of the fiscal package, as the policy announcement had been expected ahead of the budget statement. Ignoring this would bias the estimated effects of the policy announcement downwards by assuming a larger fiscal 'shock' than what has factually been observed. Additionally, the probability of the fiscal policy measures being implemented upon announcement need not equal 100%, which might also contribute to a downwards bias of the estimates.

To address these points, we follow the lead of Hazell and Hobler (2024), albeit with some limitations related to data availability. First, we establish the expected degree of debt-financing of the announced fiscal measures through their impact on the budget balance. This serves as the factual upper bound of the size of the shock component. Second, we calculate the share of debt-financing that is unexpected. For the first element, the overall increase in government debt issuance, two estimates are plausible:

- The first is based on a direct reading of the corresponding budget statement.[9] A reading of the implied policy measures yields an increase of the Debt Management Office's Net Financing Requirement from GBP 161.7 billion to GBP 234.1 billion in 2022-23, such that the corresponding upper bound of the shock (through the increase in borrowing requirements) would be GBP 72.4 billion.
- The second is an analysis by the *Institute of Fiscal Studies*, which predicted a GBP ∼60 billion funding shortfall.[10]

The *shock* impacting the expected path of debt through the fiscal announcement is not equivalent to the sum of additional deficits, since the policy package had been expected, but its full extent was simply not known. To isolate the shock component, we exploit private sector forecasts on *Public Sector Net Borrowing*, which are aggregated on a monthly basis and released by the U.K. Treasury.[11] These are forecasts of the factual borrowing requirement of the U.K. government in each fiscal year, provided both by financial market participants as well as other independent forecasters. We collect data on the forecasts provided in the period between September 1, 2022, and September 22, 2022 (i.e., until the day before the shock) and compare these forecasts with the ones collected between October 1, 2022, and October 10, 2022. Unfortunately, data is not collected at narrower

[9] The report is available under: https://www.gov.uk/government/publications/the-growth-plan-2022-documents.

[10] The report is available under: https://ifs.org.uk/articles/mini-budget-response.

[11] The forecast summaries are available under: https://www.gov.uk/government/collections/data-forecasts.

time intervals around the 'mini-budget' announcement. This limitation causes - if anything - a downwards bias of the estimated deficit-inflation multiplier.[12]

For October 2022 (i.e., after the announcement of the fiscal package), forecasts were provided by Barclays Capital, Goldman Sachs, JP Morgan, Beacon Economic Forecasting, CEBR, Heteronomics, ICAEW, Kern Consulting, and Oxford Economics. The total mean forecast revision of Public Sector Net Borrowing for the 2022-23 and 2023-24 Fiscal Years lies at GBP 47.4 billion, vastly exceeding all other non-crisis forecast revisions.[13] This confirms the initial intuition that the 'mini-budget' shock was indeed economically significant and to a large degree unexpected. Given that this forecast revision is also below the upper bound of the shock size, the following analysis works with this estimate of a GBP 47.4 billion funding shortfall, equivalent to 1.27% of annual GDP in 2022 (GBP 47.4 billion / GBP 3.732 trillion). Relative to Hazell and Hobler (2024), the 'mini-budget' shock component equals 60% of the size of their shock after normalizing by local GDP.

Linking the deficit shock to expected inflation

We now introduce data capturing changes to expected inflation through a high-frequency identification strategy. The analysis thus follows Hazell and Hobler (2024), postulating that around the 'mini-budget announcement' the dynamics of asset prices y_t can be summarized by the process:

$$y_t = \begin{cases} \varepsilon_t & \text{if } t < T, \\ \varepsilon_t + \alpha_t & \text{if } t \geq T, \end{cases} \tag{1.1}$$

where T denotes the time period at which fiscal stimulus occurred. ε_t is an arbitrary process describing asset price movements, and α_t is the effect induced by the fiscal package for all $t \geq T$. We set the shock period T to September 23, 2022, 09.30am, coinciding with the beginning of the budget statement in parliament. Denoting by j a counter of periods after the event, $\hat{\alpha}_{T+j} = y_{T+j} - \mathbb{E}_T\left[y_{T+j} | \alpha_{T+j} = 0\right]$ is the estimate of the causal effect of the shock in the narrow time window around the announcement.

[12] The deficit-inflation multiplier, measuring the effect of a change in sovereign borrowing on expected inflation, is larger for a given change in expected inflation when the borrowing shock is *smaller*. It is then easier to over-estimate the size of the shock with the available data, since the October data has been collected two weeks after the fiscal announcement, and forecasters might have by then already expected the package to be unwinding. See appendix A.1 for further details on the narrative around the 'mini-budget'.

[13] The only periods with larger absolute adjustments in the expected two-year budget deficit forecast were April 2020 (GBP 147.4 billion), May 2020 (GBP 114.9 billion), and May 2009 (GBP 50 billion). Outside of the GFC and Covid periods, the largest absolute month-on-month average forecast revision was GBP 20.8 billion in October 2019, less than half of the size of the forecast change in October 2022.

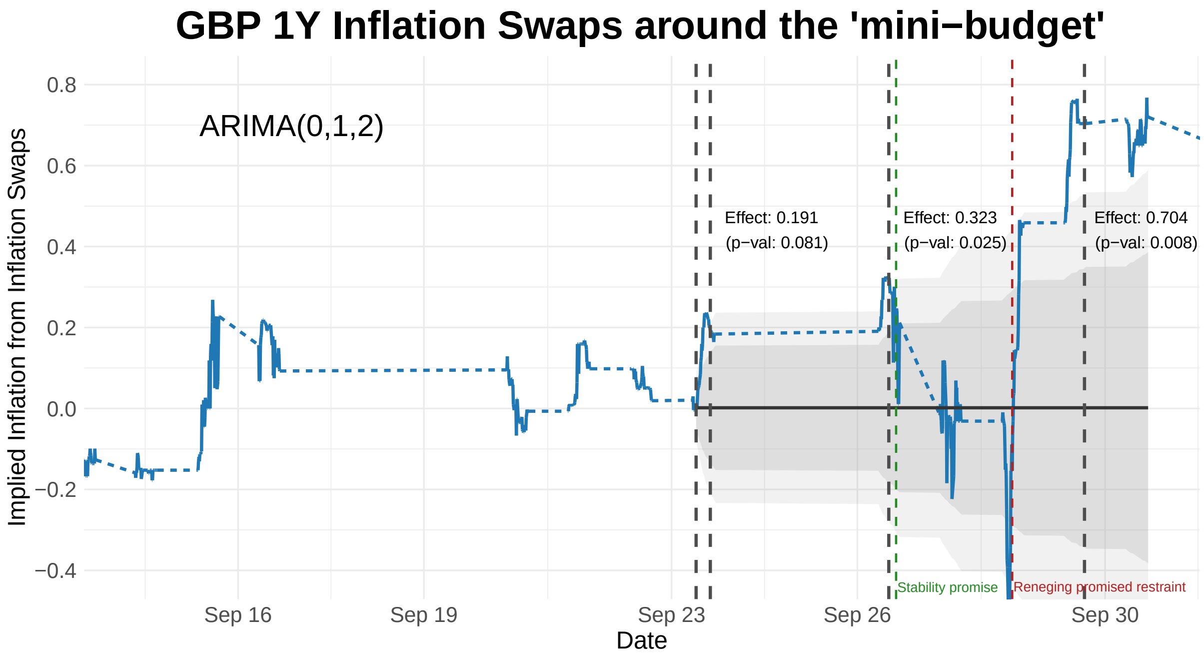

Figure 1.3: Implied inflation expectations from one-year GBP Inflation swaps in the period around the 'mini-budget' shock, with data normalized to 0 for September 23, 2022, 09:30am. The gray fan-chart depicts 68% and 95% confidence intervals for implied inflation based on a forecast of the swap price from the moment of the shock onward, with the model being chosen optimally in accordance with the Bayesian Information Criterion.

The main quantity of interest is one-year ahead expected inflation, as implied through GBP-indexed inflation swaps traded at the London Stock Exchange.[14] Figure 1.3 summarizes the movements of expected one-year ahead inflation, derived through inflation swaps, around the 'mini-budget' shock on September 23, 2022.

The gray dashed vertical lines are points at which meaningful estimates of changes to one-year ahead inflation expectations are recovered. The first vertical line depicts the beginning of the shock, as implied by the beginning of the budget speech announcing the 'mini-budget' measures in detail. The second line measures one-year ahead inflation expectations on the same day at 3:00pm, 5.5 hours after the budget speech commenced. Even though markets can credibly be expected to take a couple of days to incorporate fiscal news into forming inflation expectations (Bahaj et al., 2023), there is a significant response of implied inflation on the day of the 'mini-budget' announcement. Moving forward to September 26, i.e., the next trading day, the effect magnifies further, yielding an implied year-on-year inflationary response of 0.323% to the narratively identified shock.

Between the third and the fourth vertical line implied inflation drops sharply. This drop is consistent with the expectation that the fiscal spending announcement might end

[14]Since inflation swaps operate with a two-month indexation lag in the context of the U.K., we adjust the prices of the swaps to reflect this lag, as done by Hazell and Hobler (2024).

up being unravelled, indicated by the green dashed line labelled ’Stability Promise’. We provide a narrative description of the events in this period in appendix A.1, including a brief description of the role played by the troubles on LDI markets.

On September 28 & 29, expected inflation increased significantly again. While by then other events might contaminate the evolution of inflation swap prices, the observed sharp appreciation perfectly coincides with statements of the Treasury that *despite* the market turmoil, the proposed fiscal package will be followed through, superseding previous statements of a release of a stabilizing medium-term fiscal plan. This event is depicted by the red dashed line. The elevated levels of expected inflation then continued to persist well into October, when an eventual unravelling occurred in parallel to an overhaul of the ruling government that enacted the fiscal package in the first place.

To remain conservative in terms of the implied size of the expected inflation adjustment, yet consistent with the literature, we postulate that the response of inflation swaps until September 26 can be considered the baseline change in one-year ahead inflation expectations.

The Inflation Multiplier: the resulting baseline estimate of the one-year ahead inflation multiplier, which captures the response of year-on-year inflation to a 1% deficit-to-GDP shock, is therefore $0.323/1.27 \approx 0.254\%$. This estimate exceeds the *two-year* inflation multiplier found by Hazell and Hobler (2024) of 0.19% by 33%, despite the assumptions ensuring that the inflation multiplier estimate errs on the conservative side. Taking the point estimate of 0.704 (which aligns closest with the forecast change to the budget deficit introduced in the last subsection), the inflation multiplier would amount to $0.704/1.27 \approx 0.554\%$, more than double the estimate of Hazell and Hobler (2024) and vastly above existing estimates for other countries. The effects of the deficit shock were expected to be persistent, as implied by inflation swaps for longer horizons depicted in figure 1.4.

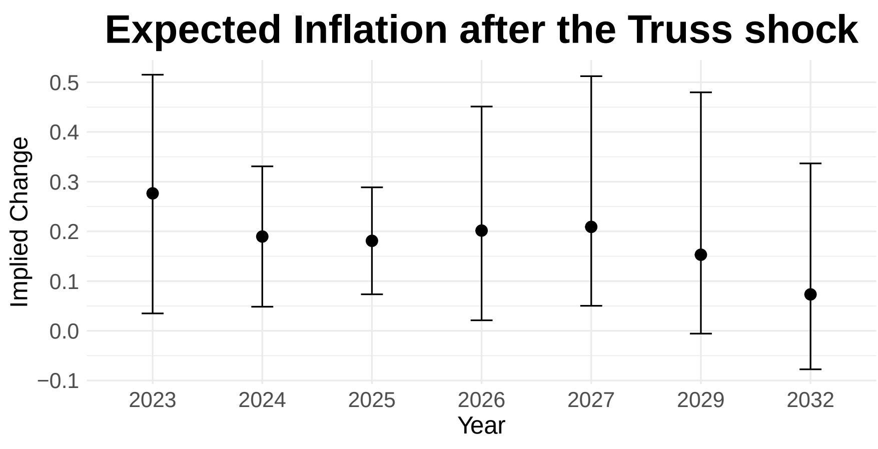

Figure 1.4: Implied change in inflation expectations for various forecast horizons, as implied by GBP inflation swaps on September 26, 2022, 12:00pm relative to the pre-shock period. 95% confidence bands are shown through bins.

Appendix section A.1 provides further evidence on the nature of the 'mini-budget shock episode', which confirms an element of surprise in relation to the size of the unveiled fiscal package that contributed to the turmoil on financial markets reflected in the pricing of inflation swaps.

Evidence on the ex-post inflationary effect of inflation-indexed debt

A limitation of the previous narrative analysis is that the inflation measure is one of *expected* inflation recovered through the pricing of inflation swaps. We now address this, providing evidence on the effect of inflation-indexed debt on *realized* inflation. To do so, we leverage the time series of narratively identified exogenous fiscal policy surprises ε_t^F provided by Mierzwa (2024), and combine it with a novel time series of the market value of inflation-indexed debt as a share of the overall market value of U.K. sovereign debt, which we label ω_t. We take this to be the main indicator for the prevalence of inflation-indexed debt. Equipped with these time series, we estimate the following local projection to measure the dynamic impact of inflation-indexed debt on the price level:

$$\log P_{t+h} - \log P_{t-1} = \alpha_h + \beta_h \Delta\omega_t \varepsilon_t^F + \delta_{1h}\Delta\omega_t + \delta_{2h}\varepsilon_t^F + \Gamma_h Z_{t-1} + e_{t+h}, \qquad (1.2)$$

where $h \geq 0$ indexes the forecast horizon considered and Z_{t-1} is a vector of control variables specified below. Of particular interest is the coefficient β_h, which captures the cross-effect of the identified fiscal shock ε_t^F and the growth in the share of inflation-indexed debt $\Delta\omega_t$ present in the economy at time t.[15]

Figure 1.5 depicts the impulse-responses estimated through the local projection (1.2). The crucial observation is that the interaction effect between the share of inflation-indexed debt and the fiscal policy shock is significantly positive in the ten quarters after the shock. In economic terms, the coefficients imply that a 1% increase in the combined measure of the change of the share of inflation-indexed debt and the narratively identified fiscal shock (measured as a percentage of GDP) itself leads to an increase of the price level of almost 1% in the two years after the shock. This evidence clearly links the share of inflation-indexed debt to inflation observed in response to expansionary fiscal shocks. Further details related to this analysis as well as an application to U.S. data are provided in appendix A.1.

[15] We work with the growth rate of the share of inflation-indexed debt in the total debt portfolio as the share of inflation-indexed debt ω_t is trending, as shown in appendix figure A.6. Using the level of inflation-indexed debt can then induce a spurious state-dependence. More generally, since ω_t is endogenous to fiscal policy decisions, using the change in the share of indexed debt $\Delta\omega_t$ can reduce the magnitude of the bias when the level of ω_t is strongly correlated with other low-frequency components of fiscal policy that affect the outcome variable (Goncalves et al., 2024). Effectively, using the change $\Delta\omega_t$ here captures the local derivative of inflation with respect to ε_t^F when the policy variable ω_t itself also changes contemporaneously. Using instead the past level of indexed debt, ω_{t-1}, does not significantly alter the results.

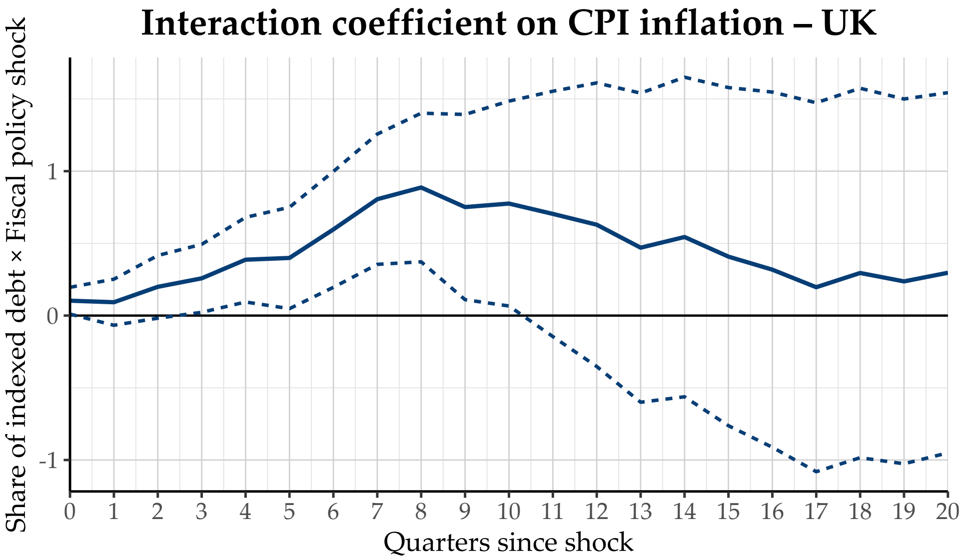

Figure 1.5: IRF implied by the local projection (1.2) through the coefficients β_h. The control vector Z contains the first four lags of the real GDP growth rate, the Bank of England Official Bank Rate, the change in the weighted real exchange rate, a same-period recession indicator, year-fixed effects, and the first lag of the price level difference. Standard errors are robust to heteroskedasticity and autocorrelation (Newey-West correction). Confidence intervals are provided at the 90% level. Sample length: 1970 Q1 - 2019 Q2.

1.3 Intuition from a One-Equation Price Level Determination Model

We now introduce inflation-indexed debt in the government budget constraint and the resulting debt valuation equation. The analysis herein can be thought of as a partial equilibrium analysis that isolates the effect of indexed debt through the government budget constraint. We derive the novel result that the price level itself becomes a *state variable* in the intertemporal government budget equilibrium, i.e., today's price level becomes a function of the past price level. This observation holds true despite the lack of other inertia, and it gives rise to a double role of the price level as a state variable and a market-clearing jump variable.

The per-period government budget constraint in a world with indexed debt is given by

$$B_{t-1} + \frac{P_t}{P_{t-1}} b_{t-1} = P_t s_t + Q_t B_t + q_t b_t,$$

where B_t denotes the face value of non-indexed government debt issued at time t at price Q_t, b_t denotes the issuance value of indexed-government debt issued at time t at price q_t, lowercase letters correspond to the values for inflation-indexed debt, s_t are primary real surpluses raised (i.e., the inverse of deficits), and P_t denotes the price level. The cost of maturing inflation-indexed debt b_{t-1} is scaled by the gross inflation rate, P_t/P_{t-1}.[16]

To close this model as simply as possible, let $Q_t = \frac{1}{1+i_t}$ and $q_t = \frac{1}{1+r_t}$, i.e., the prices of each type of bonds equal the inverse of the respective relevant gross interest rate.[17] By letting the price of indexed debt be equal to the inverse of the real interest rate, expectations of the face value adjustment of indexed debt through inflation are taken care of. Iterating this equation forward, dividing both sides by P_t, and making use of the Fisher equation gives rise to the following relationship:

$$\frac{B_{t-1}}{P_t} + \frac{b_{t-1}}{P_{t-1}} = \sum_{j=0}^{\infty} \prod_{l=1}^{j} \frac{1}{1+r_{t+l}} s_{t+j}. \tag{1.3}$$

This equation is the baseline government debt valuation equation without accounting for the differences in the insurance properties borne by the two types of debt, which made it possible to use the simplified bond pricing kernels Q_t and q_t as defined in the last paragraph. Clearly, *the price level itself becomes a state variable*: the real value of maturing inflation-indexed bonds depends on the past price level, not on today's price level. Intuitively, the real value of inflation-indexed bonds depends on the past price level because the face value payment of that bond is unity *at yesterday's prices*. The term in orange is the novel addition relative to canonical models of fiscal inflation and is the centrepiece of this chapter.

The dynamics of the government budget with indexed debt

We now explore the relationship between indexed debt and the price level within the government debt valuation equation through changes to fiscal surpluses. The goal is to explore how indexed debt changes the mechanisms inducing fiscal inflation in the clearest possible way.

The model is set up (in terms of outstanding bonds and expected surpluses), such that $P_{-1} = 1$. Therefore, the initial present discounted value (PDV) of surpluses is equal to the real value of the stock of debt. The economy has a finite horizon of 11 periods $t \in \{-1, 0, 1, ..., 9\}$, such that all debt has to be repaid by the government in period 9 using appropriate surpluses. All these assumptions jointly ensure a price level of $P_t = 1 \; \forall t$ in the absence of shocks. The impulse to the system is a one-period decrease of surpluses by 10% in period 0, announced at the same time. After period 0, the PDV of surpluses returns to its pre-shock value.

[16] See Hall and Sargent (2011) for a verification that this is the correct specification for indexed debt that is in line with how its face value adjusts empirically, absent indexation lags that we sideline here for simplicity.

[17] This equals the fair price of either type of debt for a household maximizing its cumulative discounted utility with an exogenous endowment and perfect foresight.

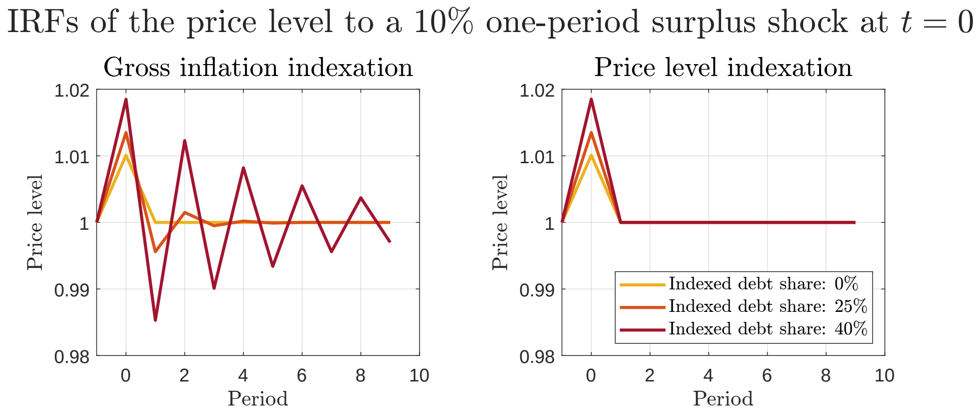

Figure 1.6: IRFs to a 10% decrease in the surplus in $t = 0$ for various levels of indexed debt.

Figure 1.6 shows the reaction of the price level in response to a one-period decrease in surpluses (which is reversed in $t = 1$), announced in the period 0. The right-side panel illustrates the response induced by the government debt valuation equation in an alternative world with *price level indexation*, which is not exactly equal and would be considered colloquially *real debt*. In period 0, the decrease in real surpluses induces a temporary upwards adjustment of the price level proportional to the decrease in surpluses, which returns back to its initial state subsequently, since the PDV of surpluses is equal to the pre-shock value after period 1.

However, when the share of inflation-indexed debt is strictly positive, the impact response is exacerbated: given that the initial price level P_{-1} is fixed in the moment of the shock at time 0, it is not possible to devalue the stock of inflation-indexed debt when the shock occurs. Therefore, the devaluation of the remaining (non-indexed) stock of bonds must be *larger* relative to the case without inflation-indexed debt: the price level must go up by a larger amount in the shock period because of the existence of inflation-indexed debt.

Once the shock vanishes, the price level oscillates when indexed debt is present instead of returning directly to steady-state. How can this be? Since the PDV of surpluses returns to its pre-shock level in $t = 1$, the stock of debt is suddenly worth *too little*: inflation-indexed debt is not worth much in $t = 1$ due to the high price level at $t = 0$, which is the factor by which b_t is normalized to 'real' terms in period 1. Since the funding shortfall caused by the deficit shock is now gone, the real value of non-indexed debt (B_1/P_1) must actually *increase* to make up the 'under-valuation' of indexed debt: therefore, P_1 must *decrease* (increasing the real value of non-indexed debt) to let the debt valuation equation hold. In the subsequent period, the price level from the previous period is now *too low*, increasing the value of indexed debt and pushing down the real value of non-indexed debt through a higher price level. This mechanism repeats itself until convergence to the initial equilibrium.[18]

[18] Cochrane (2001) explores a similar result in figure 4 of his paper, driven by a non-geometric maturity structure and the presence of long-term debt.

1.4 A Heterogeneous-Agent General Equilibrium Model with Indexed Debt

Having studied the relevance of indexed debt in a simplified framework that isolates its effect on the debt valuation equation, we now introduce inflation-indexed debt in a rich state-of-the-art macroeconomic model. Given that inflation-indexed debt delivers desirable insurance features to households by providing an income smoothing source that yields certain real returns, the model must necessarily bear relevance to imperfect consumption smoothing, borrowing constraints, and market imperfections precluding perfect risk-sharing across households. Otherwise, households would be indifferent between the two types of debt up to first-order. We work with a heterogeneous-agent model in the spirit of Kaplan et al. (2018), utilizing the efficient algorithms for solving such models provided by Auclert et al. (2021) and paying close attention to the peculiarities of determinacy in incomplete-market models exposed by Brunnermeier et al. (2024), which are intimately related to the government debt valuation equation.

Households: Heterogeneous households are indexed by i. Such households choose consumption, c_{it}, labour supply, N_{it}, and asset holdings B_{it} and b_{it} to maximize their cumulative discounted utility

$$\mathbb{E}_0 \left[\sum_{t=0}^{\infty} \beta^t \left(u(c_{it}) - v(N_{it}) \right) \right],$$

subject to two budget constraints - one for the aggregate household budget, and one for the semantically separate evolution of indexed debt in the household savings portfolio:

$$P_t c_{it} + Q_t B_{it} = \frac{e_{it}^{1-\theta}}{\int e_{it}^{1-\theta} di} (1 - \tau_{it}) W_t N_{it} + B_{i,t-1} - d_{it} 1_{\{adj_{it}=1\}},$$

$$q_t b_{it} = \Pi_t b_{i,t-1} + d_{it} 1_{\{adj_{it}=1\}},$$

where Q_t and q_t are the nominal prices for non-indexed debt B_{it} and indexed debt b_{it}, respectively. $W_{it} \equiv w_{it} P_t$ denotes the nominal wage level, adjusted by hours worked N_{it} and scaled by the idiosyncratic productivity disturbance $\frac{e_{it}^{1-\theta}}{\int e_{it}^{1-\theta} di}$ and the distortionary income tax rate τ_{it}. Through the two separate budget constraints, we posit that consumption is only possible from the non-indexed savings portfolio. In effect, indexed debt cannot be transformed to consumption as easily as non-indexed debt. This assumption reflects the significantly smaller liquidity of inflation-indexed bond markets, even relative to their market size (Andreasen et al., 2021; Fleming and Krishnan, 2012) and is necessary to ensure the existence of an ex-ante real yield differential. Without any adjustment friction, expected yields would equalize and there would be no incentive to hold both types of debt through a no-arbitrage argument.[19]

d_{it} captures household-specific transfers from the 'main' budget constraint to the relatively less accessible portfolio of indexed bond holdings, which are only allowed to happen when the exogenous portfolio rebalancing variable adj_{it} is equal to 1, which happens with probability ν. This is a Calvo-type friction, applied to portfolio holdings of the household.[20] Finally, households are subject to borrowing constraints: $B_{it} \geq -\underline{\mathrm{B}}$, $b_{it} \geq -\underline{\mathrm{b}}$.

Let $\varepsilon_i \equiv \frac{e_i^{1-\theta}}{\int e_i^{1-\theta} di}$ be a simplified descriptor of the Markov chain pinning down idiosyncratic productivity. To solve the household block, it is necessary to distinguish whether a household is able to adjust its holdings of indexed debt in a given period ($adj_{it} = 1$) or not ($adj_{it} = 0$). We now define corresponding value functions for households, noting that the state variables are therefore the household-specific past asset holdings (B_-, b_-), the Markov chain realization ε_i, and the adjustment state adj_i. The subscript i is dropped in the following for notational simplicity, yielding the following value functions:

- For households that can adjust their indexed debt holdings, $adj = 1$:

$$V_t(1, \varepsilon; B_-, b_-) = \max_{c,B,b,N} u(c) - v(N) + \beta \mathbb{E}\left[V_{t+1}(adj', \varepsilon', B, b)|\varepsilon\right] \tag{1.4}$$

  subject to a unified budget constraint (created by replacing d_{it} in the former budget constraint) and the borrowing constraints:

$$Pc + QB + qb = \varepsilon(1-\tau)WN + B_- + \Pi b_-,$$
$$B \geq -\underline{B}; \quad b \geq -\underline{b},$$

  where adj' is i.i.d., with probability $\mathbb{P}(adj' = 1) = \nu$.

- For households that cannot adjust their indexed debt holdings, $adj = 0$: b does not enter the decision set of these households and is taken to be an unchangeable state, with the next-period value of each households' indexed debt holdings being determined by the prevalent inflation rate Π and the current price of indexed debt q.

$$V_t(0, \varepsilon, B_-, b_-) = \max_{c,B,N} u(c) - v(N) + \beta \mathbb{E}\left[V_{t+1}\left(adj', \varepsilon', B, \frac{\Pi}{q} b_- |\varepsilon\right)\right] \tag{1.5}$$

  subject to the budget and borrowing constraints:

$$Pc + QB = \varepsilon(1-\tau)WN + B_-,$$
$$B \geq -\underline{B}.$$

[19] Complementary evidence on the use of inflation-indexed government bonds by households for inflation hedging within the context of the U.S. is provided by Nagel and Yan (2022).

[20] Such Calvo-type sticky portfolio arrangements have been present in macroeconomic models since at least Graham and Wright (2007) and have prominently been used in heterogeneous-agent models by Auclert et al. (2024b) and Bayer et al. (2024).

The goal is to recover policy functions $c(\cdot), B(\cdot), b(\cdot)$, and $N(\cdot)$ that solve the household problem in both instances. The above problem generally yields first-order conditions that depend on the adjustment possibilities that a household enjoys in a given period. Denote by λ_{it}, μ_{it}^B, and μ_{it}^b the respective state-dependent constraint multipliers. The relevant first-order conditions of the households that can adjust their indexed debt holdings actively ($adj_{it} = 1$) are given by

$$\begin{aligned}
\{c\}: & \qquad u'(c) = P\lambda_{it} \\
\{N\}: & \qquad v'(N) = \lambda_{it}\varepsilon(1-\tau)wP \\
\{B\}: & \qquad Q\lambda_{it} = \beta\mathbb{E}\left[V_{B,i,t+1}\right] + \mu_{it}^B \\
\{b\}: & \qquad q\lambda_{it} = \beta\mathbb{E}\left[V_{b,i,t+1}\right] + \mu_{it}^b,
\end{aligned}$$

while the envelope conditions, using $\lambda_{it} = \frac{u'(c)}{P}$ from the first-order condition on c, are

$$\begin{aligned}
V_{B,i,t} &= \frac{u'(c)}{P}, \\
V_{b,i,t} &= \begin{cases} \frac{u'(c)}{P}\Pi = \frac{u'(c)}{P_-} & \text{if } adj_{it} = 1 \\ \beta\frac{\Pi}{q}\mathbb{E}\left[V_{b,i,t+1}\right] & \text{if } adj_{it} = 0. \end{cases}
\end{aligned}$$

The conditions for equilibrium jointly imply the following Euler equations:

$$\begin{aligned}
\frac{Q}{P}u'(c) &\geq \beta\mathbb{E}\left[V_{B,i,t+1}\right], \\
\frac{q}{P}u'(c) &\geq \beta\mathbb{E}\left[V_{b,i,t+1}\right], \\
v'(N) &= u'(c)\varepsilon(1-\tau)w,
\end{aligned}$$

where the inequalities are strict if the respective asset holdings are at their lower bound.

This household block defines pricing kernels for the bonds that are on offer by the government, conditional on the households pricing the bonds being unconstrained. For non-indexed debt, the first-order conditions for households on the Euler equation imply that

$$Q_t = \beta\mathbb{E}_t\left[\frac{u'(c_{i,t+1})}{u'(c_{it})}\frac{P_t}{P_{t+1}}\right] := \mathbb{E}_t\left[\mathcal{M}_{i,t,t+1}\right], \tag{1.6}$$

where $\mathcal{M}$ denotes the household-specific stochastic discount factor (SDF). For indexed bonds, applying the definition of the SDF,

$$q_t = \beta\mathbb{E}_t\left[\frac{u'(c_{i,t+1})}{u'(c_{it})}\right] := \mathbb{E}_t\left[\mathcal{M}_{i,t,t+1}\Pi_{t+1}\right]. \tag{1.7}$$

Firms and production: To focus on the effects of indexed debt and its interaction with households facing uninsurable idiosyncratic income risk, we model the production block in a parsimonious way, following Auclert et al. (2024). The model features continuum of monopolistically competitive firms k that produce goods of variety k, with each being produced in accordance with a linear production function $Y_{kt} = A_{kt}N_{kt}$. A_{kt} evolves according to an AR(1) process in logs,

$$\log A_{kt} = \rho_a \log A_{k,t-1} + \sigma_\epsilon \epsilon_{kt},$$

where $0 \leq \rho_a \leq 1$. The firm profit function is defined as

$$D_{kt} = \frac{P_{kt}}{P_t}Y_{kt} - \frac{W_t}{P_t}N_{kt} = \left(\frac{P_{kt}}{P_t} - \frac{W_t}{P_t}\frac{1}{A_{kt}}\right) A_{kt}^{1-\zeta}\left(\frac{P_{kt}}{P_t}\right)^{-\zeta} Y_t.$$

Following Auclert et al. (2024), a log-linearized approximation to the solution of the profit-maximization problem of monopolistically competitive firms yields a Phillips Curve of the form:

$$\hat{\pi}_t = (\varphi + \sigma)\kappa \sum_{l=0}^{\infty} \beta^l \hat{y}_{t+l} \tag{1.8}$$

where $(\varphi + \sigma)$ is the sum of the Frisch elasticity of labour supply and the inverse of the elasticity of intertemporal substitution, as in standard New Keynesian models, and $\kappa = \frac{\lambda(1-\beta(1-\lambda))}{1-\lambda}$ is the slope of the Phillips curve, where $1 - \lambda$ is the probability with which monopolistically competitive firms can adjust their prices.[21]

Fiscal policy: Fiscal policy is characterized by two elements. The first one is the government debt valuation equation. Characterizing this equation is only possible by invoking a suitable transversality condition on government debt.

As pointed out by Brunnermeier et al. (2024), individual transversality conditions on household asset holdings do *not* imply directly that a similar transversality condition holds for aggregate debt stocks under incomplete markets. With incomplete markets and endogenous real interest rates, the government debt valuation equation may ultimately fail to deliver a unique price level based off a simple aggregate transversality condition on government debt, since there is no guarantee that such a condition holds when markets are incomplete. Intuitively, the government can earn a 'safe asset premium' on one type of its debt when the span of the two assets on offer is not the same. Appendix section A.2 illustrates this point in greater detail.

Instead, we here follow the approach introduced by Brunnermeier et al. (2024), which is dubbed the *dynamic trading perspective.* By that approach, we start from household

[21] This production sector requires that the aggregate effects of idiosyncratic household productivity risk are 'small' for the production firms relative to the aggregate effects of aggregate risks. See proposition 4 of Auclert et al. (2024) for a detailed exposition of this point. In addition, we assume implicitly that labour supply is intermediated by labour unions, which allows us to express the Phillips Curve in aggregate form without considering distributional wealth effects on labour supply, as outlined by McKay and Wolf (2022).

unit-level budget constraints. Aggregating them up under suitable bond pricing kernels, we obtain a dynamic aggregate constraint on sovereign debt, which constitutes the government debt valuation equation. By doing so, we account for the benefits of the two debt products in partially overcoming the effects of market incompleteness borne by households, which allows us to leverage household-level transversality conditions. This final expression of the government debt valuation equation still equates the real value of today's debt holdings to a *suitably discounted* fiscal surplus term:

Proposition 1 *In a model with both non-indexed and inflation-indexed debt and incomplete markets, the government debt valuation equation can be expressed as:*

$$\frac{B_{t-1}}{P_t} + \frac{b_{t-1}}{P_{t-1}} = \mathbb{E}_t \left[\sum_{k=0}^{\infty} \tilde{\mathcal{M}}_{t,t+k} \bar{A}_{t+k} \right], \tag{1.9}$$

where

$$\tilde{\mathcal{M}}_{t,t+k} \equiv \sum_i \mathcal{M}_{i,t,t+k} \Pi_{t+1,t+k} \frac{A_{i,t+k}}{\bar{A}_{t+k}}$$

is the weighted average stochastic discount factor across all households i, adjusted for inflation, with weights being proportionate to $A_{i,t+k}$. Furthermore,

$$\bar{A}_t \equiv \frac{1}{N_i} \sum_i A_{it}$$

is the average of the term A_{it}, which captures the surpluses raised by the government from each household i:

$$A_{it} \equiv c_{it} - \varepsilon_{it}(1 - \tau_{it}) w_t N_t.$$

Proof. See appendix A.2. ■

A_{it} captures the full portfolio return earned by household i for each additional unit of net worth, consisting of the net utility gain from saving. Depending on the precise nature of market incompleteness and the prevailing fiscal and monetary rules, equation (1.9) (co-)determines the price level at time t, given some previous price level P_{t-1}.

We close the government block by assuming the taxation rule

$$\frac{\tau_t}{\tau} = \left(\frac{s_{B,t-1}}{s_B} \right)^{\gamma_B} \left(\frac{s_{b,t-1}}{s_b} \right)^{\gamma_b} e^{\zeta_t}, \tag{1.10}$$

where $\tau_t \equiv \frac{T_t}{Y_t}$ are surpluses raised by the government as a fraction of output, and $s_{B,t} \equiv \frac{Q_t B_t}{P_t Y_t}$, $s_{b,t} \equiv \frac{q_t b_t}{P_t Y_t}$ are the real market values of the two existing types of debt. ζ_t is an AR(1) shock to the quantity of lump-sum taxes raised, and the policy reaction coefficients to deviations of the market values of both types of debt from their steady-state values are given by γ_B and γ_b. Steady-state values are denoted without time subscripts. In log-linearized terms, this relationship becomes:

$$\hat{\tau}_t = \gamma_B \hat{s}_{B,t-1} + \gamma_b \hat{s}_{b,t-1} + \zeta_t. \tag{1.11}$$

Monetary policy: Monetary policy follows an inertial Taylor rule with weights on both inflation and output deviations from steady-state:

$$\left(\frac{R_t^n}{R^n}\right) = \left(\frac{R_{t-1}^n}{R^n}\right)^{\rho_M} \left[\left(\frac{\Pi_t}{\Pi}\right)^{\phi_\pi} \left(\frac{Y_t}{Y}\right)^{\phi_y}\right]^{1-\rho_M} e^{\nu_t} \tag{1.12}$$

where ν_t is an AR(1) shock to the conduct of monetary policy. In exact log-linearized terms,

$$\hat{r}_t^n = \rho_M \hat{r}_{t-1}^n + (1-\rho_M)\left[\phi_\pi \hat{\pi}_t + \phi_y \hat{y}_t\right] + \nu_t. \tag{1.13}$$

Market clearing: Market clearing on the three markets in this economy is defined as follows:

- Goods market: on the goods market, aggregate consumption and production are equalized, taking into account the loss from price adjustment costs on the producer's behalf:

$$C_t + G_t + \frac{\mu/(\mu-1)}{2\kappa}\left(log(1+\pi_t)\right)^2 Y_t = Y_t. \tag{1.14}$$

- labour market: labour supply and demand must be equal:

$$\sum_i N_{it} = \sum_k N_{kt}. \tag{1.15}$$

- Asset market: for each class of assets, the supply by the government must be equal to cumulative household demand:

$$B_t = \sum_i B_{it} \tag{1.16a}$$

$$b_t = \sum_i b_{it}. \tag{1.16b}$$

Equilibrium: We now define the competitive equilibrium in this economy:

Definition 1 (Competitive Equilibrium) *A* competitive equilibrium *is an allocation* $\{c_{it}, B_{it}, b_{it}, Y_{kt}, N_{it}, d_{it}, \tau_t\}_{t=0}^{\infty}$, *together with prices* $\{P_t, P_{kt}, w_t, \pi_t, Q_t, q_t, R_t^n\}_{t=0}^{\infty}$ *and exogenous variables* $\{\nu_t, A_{kt}, G_t, \zeta_t\}_{t=0}^{\infty}$, *such that:*

- *all households maximize their utility with suitable policy functions on* $c(\cdot), N(\cdot), B(\cdot)$, *and* $b(\dot{})$, *solving the type-dependent value functions (1.4) or (1.5),*

- *all firms maximize their PDV of profits,*
- *the government does not violate its per-period budget constraint, levies taxes in accordance with its fiscal rule, and the price level is determined through equation (1.9),*
- *the central bank follows its policy rule (1.12),*
- *all markets clear ((1.15), (1.16a), (1.16b), equation (1.14) follows from Walras' Law), and*
- *the distribution of household wealth and productivity* $\Gamma_t(B, b, z)$ *evolves by its law of motion and is determined in the long-run by the fixed point of its evolution:*

$$\Gamma_{t+1}(\mathcal{B}, \lfloor, z') = \int_{\{(B,b,z):g_t(B,b,z)\in(\mathcal{B},\lfloor)\}} Pr(z'|z) d\Gamma_t(\mathcal{B}, \lfloor, z).$$

We close the model by defining the utility function of consumption for each household i as $u(c) = \frac{c^{1-\sigma}-1}{1-\sigma}$, and the disutility function of labour supply as $v(N) = \frac{N^{1+\frac{1}{\varphi}}}{1+\frac{1}{\varphi}}$.

Steady-state: in the following we will consider a log-linearized approximation around the deterministic steady-state with respect to aggregate variables. That steady-state is characterized by a zero inflation rate, $\Pi = 1$, such that bond prices are equal to the household discount rate, $Q = \beta$ and $q = \beta$ in the absence of uncertainty. We furthermore normalize steady-state output to 1.

Steady-state determinacy under simplified real interest rate determination

In the following, we denote variables in steady-state with the subscript 'ss'. The analysis of the dynamic implications of indexed debt is best supported by the existence of a *unique* steady-state. This is not trivial under heterogeneous agents and non-Ricardian fiscal policy, since the real interest rate is an endogenous object and multiple values of the initial price level might be consistent with equilibrium in dependence on the prevalent real interest rate (Hagedorn, 2021). Here, we propose that the presented economy with inflation-indexed debt can yield a unique price level in a stationary equilibrium if the real interest rate is determined *outside* the government debt valuation equation. This takes off the 'double burden' of determining both the initial price level P_0 and the steady-state real interest rate r_{ss} from one equation (the government's debt valuation equation), although some additional restrictions must be made. This statement is formalized in the following:

Proposition 2 (Stationary equilibrium determinacy) *Under incomplete markets, with non-negative steady-state inflation, and abstracting from aggregate uncertainty, the intertemporal government debt valuation equation can determine a unique initial price level in stationary equilibrium even in the presence of inflation-indexed debt for non-negative steady-state inflation rates if* $\frac{b_{ss}}{b_{ss}+B_{ss}} < 1$, $r_{ss} > 0$*, the bond-issuance schedule satisfies the BGP consistency condition* $g_B = g_b$*, and if a steady-state asset demand function* $\mathcal{S}(r_{ss})$ *exists and is invertible.*

Proof. See appendix A.2. ■

The model therefore yields a unique initial price level with positive steady-state levels of inflation-indexed debt, provided that the real interest rate is pinned down outside of the government debt valuation equation. This result is reflective of Hagedorn (2021, 2024) with the added complication of inflation-indexed debt.

The intuition behind the proof is the following: the intertemporal government budget equilibrium without inflation-indexed debt relates the price level to the real interest rate, which is determined on the asset market. With inflation-indexed debt, steady-state inflation itself becomes another element of the intertemporal government budget equilibrium. That inflation rate, which is posited to be pinned down by fiscal policy in the stationary equilibrium, is directly related to the real interest rate through the Fisher equation. Then, with the real interest rate (and, thus, implicitly inflation) being pinned down by asset market equilibrium, there is only one plausible real interest rate that manages to uniquely pin down the price level from the government budget constraint.

1.5 Calibration and Computational Approach

Table 1.1 gives the overall model parametrization, while table 1.2 shows the endogenously calibrated parameters. We follow overall the approach of Auclert et al. (2021), as we apply a conceptually similar algorithm. In the preferred calibration, we vary government spending G and the household discount factor β to ensure that the goods and labour markets debt clear. The market for inflation-indexed debt is targeted with the help of ν, the probability of being able to access the portfolio of indexed debt actively. The market for non-indexed debt is not targeted, but clears with a tolerance of 10^{-5}, while targeted market clearing conditions clear with close to machine precision (10^{-15}). To compare various policy combinations, we here consider baseline active/passive policy coefficients (determining whether a given policy mix is fiscally-led or monetary-led) as given by Bianchi et al. (2023). The policy coefficients $\{\phi_\pi, \phi_y, \gamma_B, \gamma_b\}$ should thus be taken as indicative and related to suitable active/passive policy combinations in the sense of Leeper (1991), but not as calibrated feedback rules.[22] When deviating from the baseline parameterizations summarized in tables 1.1 and 1.2, we will explicitly introduce novel parameters as suitable.

Parameter	Description	Value	Source/Target
Firms			
Y	Steady-state output	1	Normalization
ε	Elasticity of substitution between product varieties	9	Firm mark-up of 11% (Auclert et al., 2024a)
κ	Slope of price Phillips curve	0.055	Hazell et al. (2022), Gagliardone et al. (2023), Benigno and Eggertsson (2023)
Households			
σ	Inverse of intertemporal elasticity of substitution	1	Simplification for simulation
φ	Inverse of Frisch elasticity of labour supply	1	Simplification for simulation
$\underline{\mathrm{B}}$	Lower bound of non-indexed debt holdings	0	
$\underline{\mathrm{b}}$	Lower bound of indexed debt holdings	0	
ρ_z	Persistence of AR(1) shocks to household productivity	0.966	Auclert et al. (2021)
σ_z	Standard deviation of AR(1) shocks to household productivity	0.92	Auclert et al. (2021)
Government			
T/G	Steady-state surplus, measured by the tax-to-government spending ratio	1.03	See explanation below
r^*	Natural rate of interest	0.0125	Benigno et al. (2024)
ρ_M	Inertia in Taylor-type interest rate rule	0	Simplification
ϕ_π	Monetary policy reaction to inflation deviations from steady-state	$\{0.5, 1.5\}$	For fiscally-led/monetary-led policy mix (Bianchi et al., 2023)
ϕ_y	Monetary policy reaction to output deviations from steady-state	0.3	
γ_B	Fiscal policy reaction to deviations of market value of non-indexed debt from steady-state	$\{0.3, 1.5\}$	For fiscally-led/monetary-led policy mix (Bianchi et al., 2023)
γ_b	Fiscal policy reaction to deviations of market value of indexed debt from steady-state	0.6	
Computational parameters			
n_z	Number of points in asset grid for household productivity shock	11	
n_b	Number of points in asset grid for indexed debt	50	
n_B	Number of points in asset grid for non-indexed debt	50	
$\bar{B}$	Maximum holdings of non-indexed debt in asset grid	5000	
$\bar{b}$	Maximum holdings of indexed debt in asset grid	5000	Approximation to Auclert et al. (2024)
T	Number of periods used in simulations of Jacobians	300	Auclert et al. (2021)

Table 1.1: Baseline parametrization for the quantitative estimation

The calibration delivers overall reasonable estimates of the endogenous parameters that are in line with the parametrization of Auclert et al. (2021). The level of government spending is not targeted to its empirical counterpart, yet the estimated government spending share of GDP is only slightly below the share of government spending in GDP in the U.K. in 2024 (44.4%).

[22]Macroeconomists tended to focus on calibrations in which ϕ_π, the parameter capturing the reaction of monetary policy to deviations of the inflation rate from steady-state, is larger than one, commonly known as the 'Taylor Principle'. As Nakamura et al. (2025) show, this notion is rejected in the data. Instead, the Taylor principle can be reinterpreted as *prescriptive*, that is, as a policy prescription according to which the central bank ought to act, but not as *descriptive*. This supports the idea of calibrating models with $\phi_\pi < 1$. Kaplan (2025a) explains further why policy rules not adhering to the Taylor principle deserve consideration in monetary macroeconomics.

Finally, to pin down both the price level and the tax rate in steady-state, we exogenously fix the tax rate to be 3% higher than government spending in GDP, such that surpluses are equal to one percent of the government spending-to-GDP ratio. This assumption runs counter to currently observed budget deficits, but solving the model under steady-state deficits is hardly feasible under positive steady-state real interest rates.[23] However, the assumption of positive surpluses in steady-state remains qualitatively and quantitatively in line with recent long-run forecasts of the current budget deficit for the United Kingdom, provided by the OBR (2024) in their historical official forecasts database (table CB).[24]

By imposing zero steady-state inflation, we nullify the possibility of distortionary inflation that could induce a wedge to the long-run natural rate. In terms of economic aggregates, the steady-state is thus well-described by the above calibration. Through the normalization of output to unity and the calibrated share of government spending of 0.4164, consumption in steady-state is implied to be equal to 0.5836 by market clearing, while taxation is equal to 0.4289.

Debt/GDP shares	HH discount factor	$\mathbb{P}$(adjustment)	Govt. spending
	Main calibration: U.K. debt portfolio		
B: 0.7154, b: 0.2646	$\beta = 0.9570$	$\nu = 0.2856$	$G = 0.4164$
	Counterfactual: U.S. debt shares		
B: 0.89, b: 0.09	$\beta = 0.9570$	$\nu = 0.1950$	$G = 0.4163$
	Counterfactual: no indexed debt		
B: 0.98, b: 0.	$\beta = 0.9570$	$\nu = 0.0064$	$G = 0.4165$

Table 1.2: Calibrated parameters across different debt scenarios

In terms of government debt, we mainly compare three different steady-state calibrations: one which follows the observed modal split of sovereign debt into non-indexed and indexed debt in the United Kingdom (which is the G7 country with the highest share of indexed debt, such that $B = 0.8176$ and $b = 0.3024$), and two counterfactual calibrations with either a split between indexed and non-indexed debt in accordance to the U.S. sovereign debt portfolio (i.e., $B = 1.0171$ and $b = 0.1029$), or the complete absence of any indexed debt (i.e., $B = 1.12$ and $b = 0$). We therefore exogenously postulate the same steady-state gross bond supply across the calibrations, given that bond supply as a share of GDP is a relatively low-frequency variable, and vary the shares of the two types of bonds. Many of the exercises will resolve around the differences between these calibrations, as we will mainly focus on the effect that indexed government debt has on economic aggregates.

Computational details - using the Sequence-Space Jacobian

[23] Kaplan et al. (2024) solve a model with negative surpluses and a negative steady-state real interest rate.

[24] Conditional on a $\sim 40\%$ share of government spending in GDP, the projected 1% budget surplus in the long-run as a share of GDP is equivalent to a ratio of sovereign income to spending of 1.025.

The simulation of the model is derived using the Sequence-Space Jacobian method developed in Auclert et al. (2021), which itself constitutes an evolution of the methods pioneered by Reiter (2009). The computational method we employ generates perfect-foresight solutions in aggregates in response to time-zero perturbations of exogenous disturbances, but it maintains the non-linearity underlying the responses of heterogeneous households.

First, the heterogeneous household block is solved, taking aggregate prices as given, for both the policy functions (through backwards iteration) and the distribution of asset holdings (through forwards iteration). Both solve with a numerical tolerance of up to 10^{-14}, and are subsequently used to inform other blocks of the model (such as firm optimality, government policies, and market clearing) and to generate updates of aggregates where necessary. The two components (heterogeneous households and aggregate) interact and iterate until convergence, which is reached with a numerical threshold of 10^{-9}, which is reasonable given the high degree of complexity underlying household behaviour in the presence of two types of assets.

1.6 Quantitative Insights: The Joint Role of Indexed Debt and Policy Rules

With the computational algorithm at hand, we solve and estimate the model's aggregate impulse-responses for a number of shocks, using the parametrization from table 1.1, but varying the calibration of the debt shares in line with table 1.2. Here, we will mostly focus on the effects of unanticipated disturbances to *government spending* G_t, which directly influence the surpluses raised by the government in any given period.[25]

We first look at the role that inflation-indexed debt plays for the amplification of shocks as evidenced through simulated moments, in line with the principal focus of the chapter. To get a detailed grasp of the effect of inflation-indexed debt for aggregate variables, we compare the simulated volatility of a number of macroeconomic aggregates across all calibrations (U.K. calibration, counterfactual U.S. distribution of debt across the two types, and issuance of non-indexed debt only) and across both 'common' policy combinations (passive monetary/active fiscal (fiscally-led) and active monetary/passive fiscal (monetary-led)). The results of this exercise are presented in table 1.3.

The three left-hand columns yield one of the major quantitative insights of the chapter: the volatility of consumption and inflation strictly increases in the presence of inflation-indexed debt, conditional on being in the fiscally-led policy case. Of particular interest in that regard is the fourth row of table 1.3, which captures the volatility of inflation in response to government spending shocks. Conditional on being in the fiscally-led policy

[25] Appendix A.3 provides an overview of the dynamic responses to expansionary monetary policy shocks.

scenario, the unweighted volatility of inflation is around 23% of the volatility induced by the government spending increase, compared to only 3.6% in the counterfactual without any inflation-indexed debt. With a calibrated share of indexed sovereign debt of about 30%, on average, a one percentage point increase in the share of inflation-indexed debt more or less then corresponds to a 0.65 percentage point increase in the volatility of inflation in response to uncovered government spending shocks.

	Normalized standard deviations across policy scenarios					
	PM/AF-U.K.	PM/AF-U.S. split	PM/AF-NoIndex	AM/PF-U.K.	AM/PF-U.S. split	AM/PF-NoIndex
G	1.000	1.000	1.000	1.000	1.000	1.000
Y	0.855	0.761	0.883	1.066	1.055	0.863
C	0.400	0.395	0.331	0.951	0.961	0.346
π	0.232	0.104	0.036	0.191	0.183	0.096
r	0.437	0.374	0.725	0.260	0.248	0.269
N	0.855	0.761	0.883	1.066	1.055	0.863

Table 1.3: Normalized standard deviations of aggregate variables in response to fiscal shocks with $\rho_G = 0.5$

The effect of indexed debt increases in the share of inflation-indexed debt, in line with the prescription from section 1.3. To the best of our knowledge, this chapter is among the first to quantitatively evaluate the impact that inflation-indexed debt can have on the volatility of inflation, and how such changes in volatility are directly related to the monetary-fiscal policy nexus. The inflation volatility increase is evidently much smaller under the monetary-led policy scenario, amounting to a difference of only 9.5% of the volatility of the government spending shock.

Impulse-responses to expansionary fiscal spending shocks

We will now look in more detail at the impulse-responses to government spending shocks and the role borne by inflation-indexed debt when an unexpected spending increase occurs. As the persistence of any shock is relevant, the model will be simulated under different possible autocorrelations of the fiscal shock to highlight the role of persistence and the forward-looking nature of the intertemporal government budget constraint. The initial focus rests on the case of a *fiscally-led policy mix* in line with the first parametrization introduced in table 1.2, i.e., the baseline calibration to the U.K. economy. Figure 1.7 plots IRFs of aggregate variables in response to a 100bp expansionary fiscal shock that increases the need for fiscal spending when the shock is highly persistent, i.e., $\rho_G = 0.8$.

A number of observations is worth highlighting: the responses of consumption and tax rates are in line with canonical macroeconomic models and the expected reactions in response to the fiscal expansion: in response to the fiscal expansion, there is an instantaneous increase in output that persists alongside the expenditure increase. The increase in government spending here does not spill over one-for-one to output, and there is no multiplying effect induced by the government spending increase. Some private consumption is crowded out, leading to a muted reaction of output, in line with modern estimates of the impact of expansionary fiscal policy (Auerbach and Gorodnichenko, 2012; Ramey and Zubairy, 2018; Ramey, 2019). The initial output increase is largest for high levels of debt indexation, but this comes at the cost of a crowding-out of private activity in later periods that is sufficiently large to temporarily decrease output marginally.

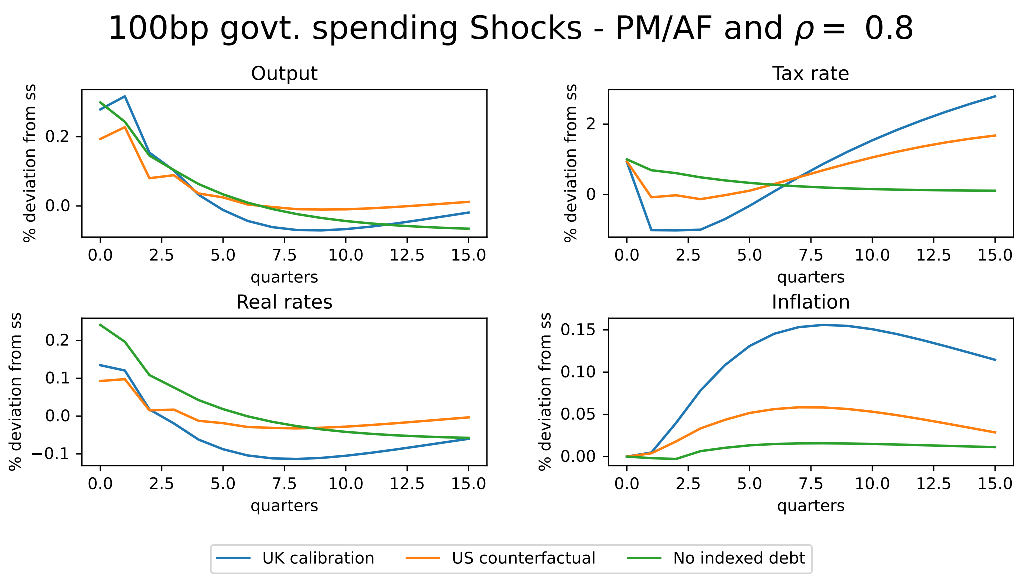

Figure 1.7: IRFs to the government spending shock under a fiscally-led policy mix.

As for the change to the marginal tax rate, its increase on impact does not fully cover additional government expenditures, in line with the specification of active fiscal policy. Without indexed debt, the tax rate monotonically converges back to zero over time. With inflation-indexed debt, however, the tax rate can increase in the medium-run due to the higher cost of serving outstanding additional payments on inflation-indexed debt, which come into play under positive changes to the rate of inflation. Once inflation-indexed debt is present, the equilibrium tax rate therefore admits a 'V-shape': after an initial increase in the tax rate on impact, it briefly decreases on the expectation of the shock being only temporary.[26] Over time, however, the tax rate subsequently increases to cover the additional expenses arising from the cost of servicing indexed debt.

The evolution of the real interest rate in response to the fiscal impulse is again tightly linked to the share of inflation-indexed debt in each calibration. In the calibration without indexed debt (green line), the real interest rate briefly appreciates on impact of the shock due to expected deflationary pressure before returning to the vicinity of its steady-state level. With positive levels of inflation-indexed debt, the impact change of the real interest rate is slightly muted, since households are better insured against the possible relocation of wealth induced by the fiscal shock due to their assets spanning more possible states. After the impact period, real interest rates depreciate in the view of expected inflationary pressure coming from the cost of serving inflation-indexed debt. This cost is increasing in the share of such debt in the economy, leading to the depression of real interest rates in the U.K. calibration below zero.

Finally, the panel on the bottom right quantifies the focal point of this chapter - the change in the rate of inflation in response to the fiscal expansion. In the present model, the price pressure arising from a fiscal expansion is only minimal without inflation-indexed debt, peaking at about 0.02% quarter-on-quarter inflation. Inflation-indexed debt, however, proves to magnify inflationary pressure quite significantly: with positive levels of inflation-indexed debt, annualized rates of inflation peak at 0.57% in the U.K. calibration (quarterly rates peak at 0.17%) and 0.22% in the counterfactual with debt shares as in the U.S., respectively. This implied deficit-inflation multiplier therefore aligns well with the empirical evidence presented in section 1.2, with the U.K. deficit-inflation multiplier being at the upper end of the range of admissible estimates arising from the large 'mini-budget shock'.

The multiplier for the U.S. calibration furthermore fits well with the evidence presented in Hazell and Hobler (2024), who find a deficit-inflation multiplier in the U.S. of 0.19%. The model therefore attributes a significant share of the differences in the deficit-inflation dynamics between the U.S. and the U.K. to the differences in the share of inflation-indexed debt, confirming the intuition laid out by the model.

Summarizing, we therefore find that turning off the debt indexation channel of government debt (i.e., setting inflation-indexed debt to zero) nullifies all dynamics beyond the first-order dynamics of the spending shock, nesting the expected reaction to a fiscal expansion under a fiscally-led policy mix with non-Ricardian households: output and inflation co-move in general, but no higher-order dynamics are observed. Once inflation-indexed debt is present, however, inflationary pressure becomes more pronounced and

[26] *Negative* deviations of the tax rate from equilibrium, as temporarily observed in the calibration to U.K. debt shares, are possible as the real value of government bonds decreases below steady-state (which is related to the large negative real return shock that lowers the prices of bonds).

persistent, accompanied by tax changes that reflect the need of the fiscal authority to cover the additional expenses arising from serving the cost of maturing inflation-indexed debt.[27]

—

Finally, we highlight what changes in the simulations when a *monetary-led policy mix* is considered instead, corresponding to fiscal policy turning 'passive' in the language of Leeper (1991). The calibration of the policy parameters in this case follows from table 1.1, with $\phi_\pi = 1.5$ and $\gamma_B = 1.5$. Figure 1.8 summarizes the results from this exercise for highly persistent fiscal shocks, $\rho_G = 0.8$.

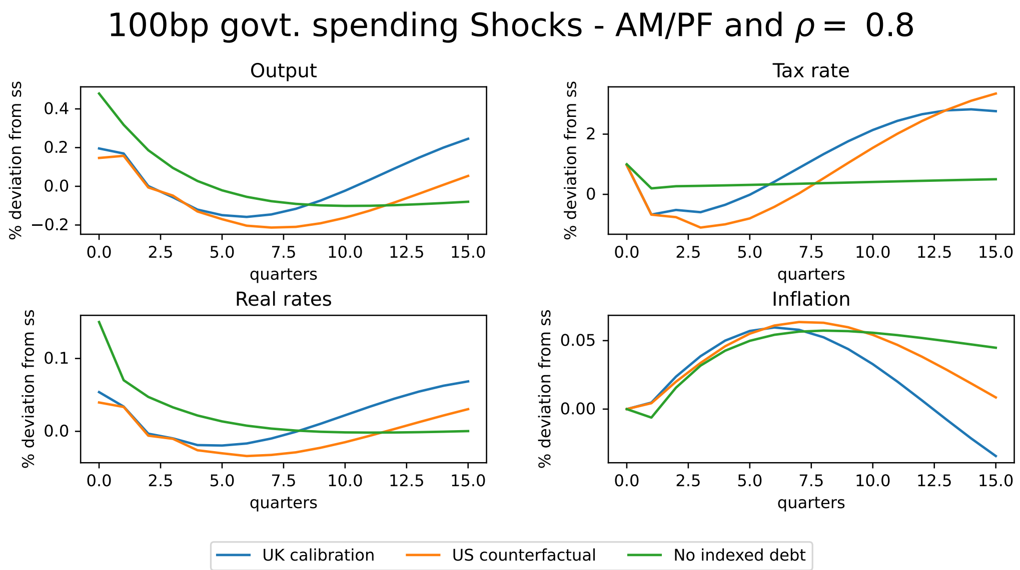

Figure 1.8: IRFs to the government spending shock under a monetary-led policy mix.

The response of output turns out to be slightly larger on impact relative to the fiscally-led policy mix for the calibration without indexed debt, indicating less of a crowding out of output in the case without indexed debt. For both calibrations with positive indexed debt levels, the output reaction to the government spending impulse turns negative in the short-to-medium run (after two quarters), reflecting that under monetary-led policy mixes fiscal policy generates smaller wealth effects on behalf of the households, such that the increased spending will be financed in part by a later reduction in available resources.

Across the board, there is little quantitative difference between the U.K. and the U.S. debt share calibrations. Since this specification follows a conventional monetary-led policy mix, fiscal policy, as exemplified through the tax rate, passively adjusts to ensure that the government budget constraint holds. It does so by increasing tax rates by a consistently higher margin relative to the fiscally-led policy case. Because the tax

[27] Complementary impulse-response functions of bond prices and quantities, as well as of the price level itself, are provided in appendix A.3.

rule shifts correspondingly, the real value of government debt remains unchanged, which is reflected in the absence of large movements of real rates and, correspondingly, of materialized inflation rates.

The model without inflation-indexed debt behaves differently in the monetary-led policy mix, generally featuring a larger response of output coincident with less volatile tax rates and a temporary increase in real rates above their equilibrium level, followed by a depreciation of real rates in the medium-run. Inflation rates depict slight positive pressure across the board, but with little difference between the various debt share calibrations. Most importantly, under the monetary-led policy mix, the reaction of inflation is muted, with deficit-inflation multipliers that are a magnitude smaller than under the fiscally-led policy mix. Nonetheless, some inflationary pressure exists, and it is somewhat larger in the medium-run without any indexed debt being present. The reason for that lies in the absence of risk-sharing among households and the imperfect insurance properties borne by normal (non-indexed) bonds: in the light of such incomplete markets, the government spending measure exhibits a greater degree of Ricardian dis-equivalence, impacting households through a wealth effect. This wealth effect contributes to a reduction in household demand, but slight inflationary pressure as the government spending shock is not fully crowded out.

Appendix A.3 presents further omitted simulation results, in particular related to the IRFs of bond prices and interest rates, household policy functions and monetary policy shocks. In particular the revaluation of the bonds as expressed through their prices are of interest, as they confirm the above arguments that the revaluation of the intertemporal government debt valuation equation belongs to the main determinants of the inflationary response.

The interaction between tax rules and the share of indexed debt

We now zoom into the joint role borne by inflation-indexed debt and the tax rule coefficients in equation (1.10). To that goal, we fix the monetary policy coefficients at the levels summarized by table 1.1 for the fiscally-led policy mix and vary the share of inflation-indexed debt in the government debt portfolio, $\omega_t = \frac{b_t}{B_t+b_t}$, between $[0, 0.25)$,[28] while also varying the fiscal policy reaction coefficient to deviations of non-indexed debt from steady-state, γ_B, between $[0, 1]$. Under these coefficients, fiscal policy is conventionally considered "active". The reaction of the price level in the first year after the fiscal impulse across the tax policy combinations and various shares of inflation-indexed debt in the sovereign debt portfolio is depicted in figure 1.9.

[28]This interval broadly captures the level of indexed debt issuance across the globe.

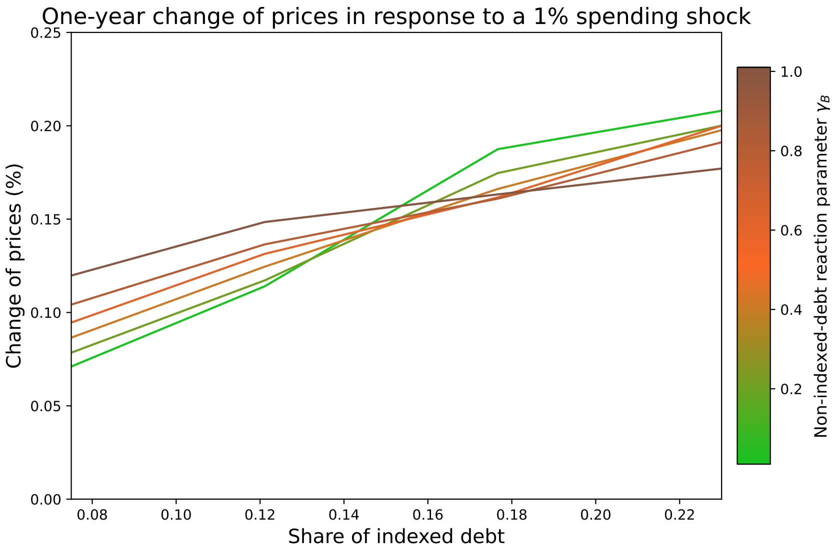

Figure 1.9: Cumulative one-year reaction of prices in response to fiscal spending shocks under a fiscally-led policy mix.

On the x-axis, we vary the share of indexed debt in the total debt portfolio (while maintaining a constant overall relation between the gross stock of debt and GDP), while the colours indicate the chosen fiscal reaction coefficient γ_B. Thus, orange and especially brown colours reflect 'less active' fiscal policy in the conventional sense (as more of the shock is covered by corresponding tax raises), while greener colours reflect 'more active' fiscal policy.

Generally, the inflationary pressure is increasing in the share of inflation-indexed debt, although the effect is non-linear. The magnitude of increasing indexed debt, however, is striking: increasing the share of indexed debt by 5 percentage points can increase the one-year deficit-inflation multiplier by 0.05 percentage points. This holds true across all fiscal reaction parameters under which fiscal policy is conventionally considered active. The marginal effect of an increase in the share of inflation-indexed debt is smaller for more restrictive fiscal policy, whereas it is larger for more expansionary fiscal policy, as indicated by the overall steeper slope of the green lines.

Fiscal-monetary policy combinations, inflation, and determinacy

The monetary reaction rule has been kept constant in the previous analysis. However, a monetary policy authority might vary its policy prescriptions to counter inflationary pressures induced by expansionary fiscal policy which does not raise taxes sufficiently to cover additional fiscal expenses. We now focus more directly on the link between the fiscal and monetary reaction rules, keeping inflation-indexed debt at constant and elevated levels corresponding to the share of inflation-indexed debt present in the U.K.

To keep the results simple, we maintain the split into policy areas that are considered conventionally fiscally-led and monetary-led. The fiscally-led policy mix is conventionally characterized by $\gamma_B \in [0, 1]$ and $\phi_\pi \in [0, 1]$, while the monetary-led policy mix is conventionally characterized by $\gamma_B > 1$ and $\phi_\pi > 1$, leaving aside in either case the possibility of negative policy parameters. In this exercise, we compare one-year inflation in response to a 100bp expansionary fiscal shock across various values of the fiscal reaction parameter γ_B and the monetary reaction parameter ϕ_π. The results for the fiscally-led policy mix are described by figure 1.10 and the results for the monetary-led policy mix are described by figure 1.11.

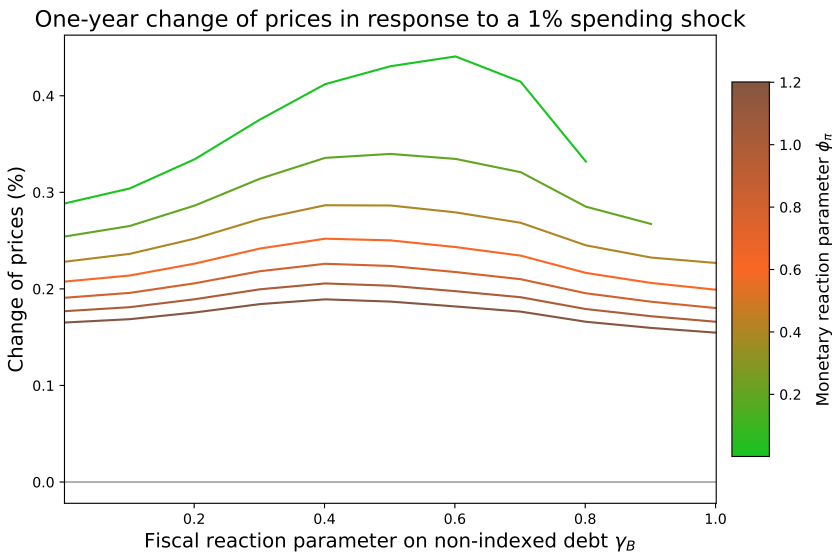

Figure 1.10: The one-year deficit-inflation multiplier under fiscally-led policy mixes.

Under the fiscally-led policy mix (figure 1.10), the deficit-inflation multipliers are principally larger than under the monetary-led policy mix (figure 1.11). This holds across all particular calibrations. Relative monetary passivity (embedded by $\phi_\pi \to 0$, i.e., by the green lines in figure 1.10) induces larger inflation multipliers, whereas the effect of the fiscal reaction parameter γ_B is ambiguous. Generally, fiscal reaction parameters of around 0.5 appear to induce the largest inflationary pressure, whereas larger or smaller values curb some of the inflationary pressure in the ballpark of a couple basis points. In general, the distance between the lines is larger than the slope of each line; therefore, under such fiscally-led policy mixes, the deficit-inflation multiplier is more sensitive to the monetary reaction parameter.

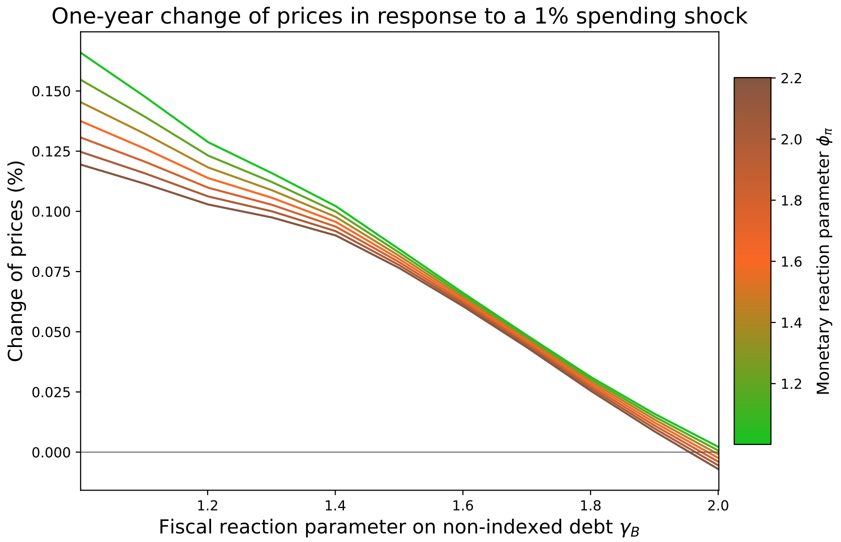

Figure 1.11: The one-year deficit-inflation multiplier under monetary-led policy mixes.

Within the confines of the monetary-led policy mix explored in figure 1.11, more restrictive monetary policy (increasing ϕ_π) always decreases the size of the deficit-inflation multiplier. A more restrictive fiscal policy (increasing γ_B) similarly always decreases the inflationary pressure in response to the fiscal impulse, and the magnitude of the effect of changing the fiscal policy parameter far outweighs the effects induced by changes to monetary policy. Under the monetary-led policy mix, therefore, the fiscal reaction parameter is particularly informative about the deficit-inflation multiplier, and very high values of γ_B can even induce *dis*inflationary pressure in response to a fiscal expenditure shock, which happens through a very restrictive tax increase.

In sum, the policy authority that is usually considered 'passive' has the greater influence on the exact size of the deficit-inflation multiplier within its constraint set. This result is surprising and amplified by the presence of inflation-indexed debt.

—

As a final exercise, we consider explicitly for which values of the fiscal and monetary policy parameters the linearized system implies a unique, saddle path-stable equilibrium. In doing so, we exploit the 'winding number criterion' developed in Auclert et al. (2023), which is consistent with the use of the associated sequence-space Jacobian methodology to solve the full dynamic model.[29]

[29] A detailed exposition of the 'winding number criterion' can be found in Auclert et al. (2023). Intuitively, it is related to the Blanchard and Kahn (1980)-conditions, which are cast in state-space. The winding number criterion provides a generalizable mapping of the Blanchard-Kahn conditions onto the sequence-space, i.e., allowing for infinitely many 'quasi-roots' of the linearized system. The prerequisites to apply the winding number criterion, such as the quasi-Toeplitz property of the generalized Jacobian, are not violated (the corresponding results are available upon request).

Figure 1.12 summarizes the determinacy properties of the model on an equispaced grid of the fiscal policy reaction parameter γ_B and the main monetary policy reaction parameter ϕ_π, with the remainder of the calibration being unchanged.

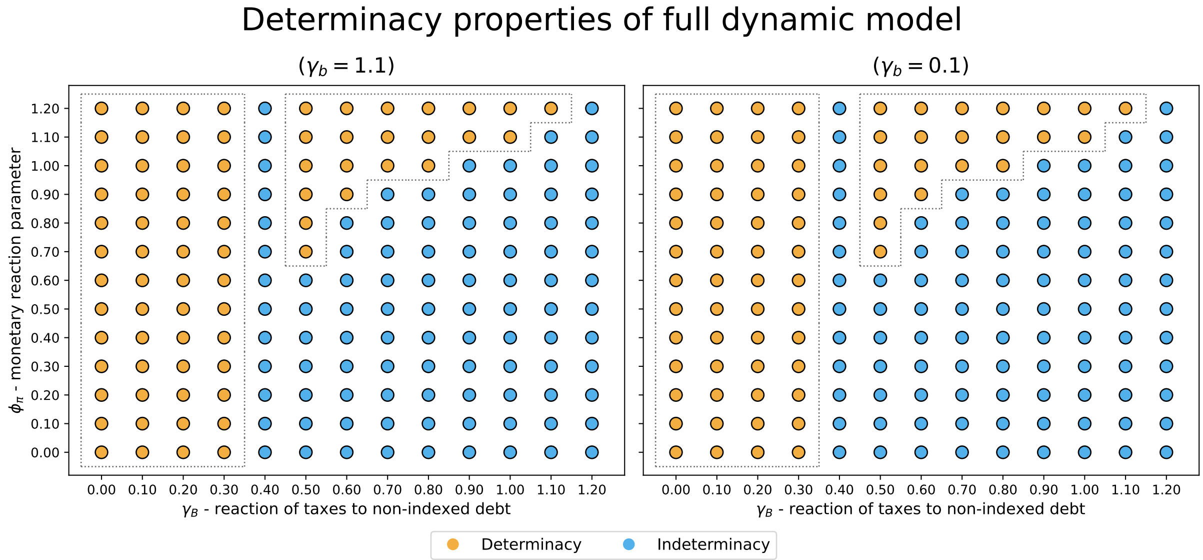

Figure 1.12: Determinacy of the generalized Jacobian in relation to choices for the fiscal and monetary policy reaction coefficients under two different values of γ_b.

In line with recent evidence on determinacy properties in non-Ricardian models (Kaplan, 2025a; Rachel and Ravn, 2025), the standard notions of determinate policy spaces do not apply one-for-one when fiscal policy is non-Ricardian. Adding the dimension of indexed debt, as we have done in this exercise, adds some further interesting insights. Under conventionally very active fiscal policy, indicated by low values of γ_B, the derived equilibrium is clearly unique and saddle path-stable for *all* plotted values of the monetary reaction parameter to deviations of inflation from steady-state. While this is not surprising for monetary policy conventionally considered passive (the bottom-left area of each panel), it is a novel result for the policy space under which monetary policy is conventionally *also* considered active (top-left area). Intuitively, the presence of indexed debt acts as an automatic stabilizer in the government debt valuation equation, allowing the admittance of policy areas under which monetary policy follows the Taylor principle and the fiscal authority does not commit to repaying any novel debt in equivalent nominal terms. The top-right area in each panel under which the economy admits unique saddle path-stable equilibria is the conventionally known active-monetary/passive-fiscal area. The monetary authority here acts in a restrictive way, and fiscal policy is not expansionary in the sense that the taxation rule allows the fiscal authority to cover additional expenses with sufficient surpluses.

The fiscal reaction parameter capturing responses to deviations of the equilibrium value of indexed debt from steady-state (γ_b) does not play a role for the model's determinacy properties, as evidenced by the fact that both panels admit the same determinacy space even though the values of γ_b are vastly different. This does not come as a surprise: except for any valuation differences in the impact period, the market price of indexed debt should reflect expectations of additional face value repayments in the presence of inflationary pressure. Thus, the precise variation of the tax rule in dependence on the market value of indexed debt does not have an effect on the model's determinacy properties.

1.7 Conclusion

This chapter introduced inflation-indexed debt into non-Ricardian general equilibrium models. We first provided empirical evidence on the role of inflation-indexed debt as a major determinant of inflationary dynamics with the help of local projections applied to the U.K. and the U.S., as well as with a specific large fiscal shock in the U.K. in September 2022. Next, we established in a simplified model that such debt itself suffices to make the price level a backward-looking state variable: the previous price level therefore matters directly for the determination of today's price level. We then introduced inflation-indexed debt in a state-of-the-art macroeconomic model with imperfect markets and household heterogeneity, ensuring the existence of a unique steady-state. Then, we provided model-driven evidence that inflation-indexed debt can indeed exacerbate the inflationary response to government spending shocks, in particular when fiscal policy is considered conventionally 'active' in the sense of Leeper (1991).

Both the empirical and theoretical results derived in this chapter thus tarnish the classic notion that inflation-indexed bonds always limit inflation in a given country by offering governments a commitment device to 'not inflate the debt away', as argued by Campbell and Shiller (1996). While this notion can remain true absent incompletely funded government deficit shocks, the results point out that once the government budget is ex-post (after debt issuance) in disarray, the inflationary consequences of funding shortfalls can increase in the share of inflation-indexed debt. Issuance of indexed debt can therefore backfire despite its great ability to serve as an ex ante commitment device following Schmid et al. (2026).

Despite these conclusions, more can be done to emphasize the interaction between inflation-indexed debt and inflation. A complete estimation of the model based off long-running samples of U.K. and U.S. data with a particular focus to the fiscal and monetary rules would further strengthen the conclusions of this chapter. Another refinement can be the inclusion of long-term government debt: as Cochrane (2001) and Barro and Bianchi

(2025) show, the maturity structure of government debt matters for the trade-off between front-loaded and delayed inflation responses to deficit shocks.[30]

Finally, inflation-indexed debt can inform recent policy debates on the possible regressivity or progressivity of inflation as implicit taxation. As evidenced by figure A.1 in the appendix, inflation-indexed debt, which serves as an insurance device against unexpected inflation, seems to be particularly skewed in household portfolios towards the highest decile of the income distribution. A more thorough analysis of the welfare effects of unexpected inflation to households at varying income levels should therefore be considered as a further policy-relevant application.

[30] Appendix A.4 briefly characterizes how to model long-term indexed debt in the government debt valuation equation.

Chapter 2

Inflation-Indexed Debt and the Risks of Fiscal Dominance

Abstract

The origins of the post-Covid inflationary episode are the subject of much debate. One argument has its roots in an unfunded expansion of debt-driven government spending, in what has been labelled *fiscal dominance* (Leeper, 1991) or a *fiscally-led policy mix* (Bianchi et al., 2023). We show that the risks of such fiscal dominance depend on the degree to which government debt is indexed to inflation. Inflation-indexation has a nonlinear effect on the existence of saddle-path equilibria, and amplifies the inflationary effects of deficit shocks when policy is fiscally led. Empirical evidence links inflation-indexed debt to low central bank independence, a high probability of suspending fiscal rules, and a larger reaction of inflation to fiscal shocks.

2.1 Introduction

Recent high inflation has renewed interest in the interactions between fiscal and monetary policy. The inflation experienced in the U.S. in 2021-23 coincided with a large debt-financed fiscal expansion that, if not accompanied by corresponding promises to repay the debt in equivalent real terms in the future, must lead to a depreciation in the real value of government debt through an inflation-driven debt devaluation process. The possibility of fiscal policy dominating in this way has contributed to concerns about the sustainability of government finances, especially given record absolute debt levels and debt-to-GDP ratios previously only seen in the aftermath of wars.[1]

This chapter shows that the risks associated with fiscal dominance are heightened when the face value of at least some of a government's outstanding debt is automatically uprated by the gross rate of inflation, as with Treasury Inflation-Protected Securities (TIPS) in the U.S. and Index-Linked Gilts in the U.K. We find that inflation-indexed debt has a nonlinear impact on the existence of saddle path-stable equilibria, and that it can induce more volatility in inflation if there are deficit-driven government transfer shocks and policy is fiscally-led. In short, fiscal dominance is more consequential when a higher proportion of government debt is indexed to inflation.

That inflation-indexed debt affects the risks of fiscal dominance has so far not been recognised. Although a number of theoretical (Angeletos et al., 2024b; Ascari et al., 2023; Bianchi et al., 2023; Cochrane, 2022b) and empirical (Barro and Bianchi, 2025; Hazell and Hobler, 2024; Hilscher et al., 2022) papers explore the recent inflationary episode using models that emphasise fiscal-monetary interactions and the extent to which debt-driven fiscal expansions can fuel inflation, they all abstract from inflation-indexed debt and so overlook the role it plays in determining the risks and consequences of fiscal dominance and a fiscally-led policy mix.

We begin with a simple Fisherian environment where households save using government debt that is partially indexed to inflation. Following the analytical approach in Leeper (1991), we show that inflation-indexed debt restricts the admissible space of monetary policies that supports a fiscally-led policy mix. This occurs because inflation-indexed debt imposes an additional constraint on fiscal policy, namely that it must fund any uprating of face values required in response to inflation. We then progress to a non-Ricardian general equilibrium model in the spirit of Blanchard (1985) and Angeletos et al. (2024b), where we find that inflation-indexed debt tends to amplify the reaction of inflation to debt-financed transfer shocks. The relationship is nonlinear since it flattens out and even inverts at high levels of inflation-indexed debt.

[1] The increased burden of servicing debt is one area of concern. For example, interest payments on federal debt in the U.S. are expected to reach 7% of GDP by 2050.

It is challenging to use empirical data to establish whether the policy mix in a particular period is monetary-led or fiscally-led (Cochrane, 2011; Neumeyer and Nicolini, 2025). We offer three distinct pieces of evidence to suggest that inflation-indexed debt plays a role in inflation dynamics through the risks of fiscal dominance. First, using a long sample of global data, we see that the share of inflation-indexed debt correlates positively with inflation and negatively with the Central Bank Independence Index of Romelli (2024). Second, we use cross-country data from Davoodi et al. (2022) to show that fiscal rules are more likely to be suspended in countries that have a high proportion of government debt indexed to inflation. Third, we combine the narratively-identified U.S. tax shocks of Mierzwa (2024) with the classification of policy in Chen et al. (2022) to demonstrate that amplification of fiscal shocks is only meaningfully amplified by indexed debt when policy is fiscally-led.

The chapter is organised as follows. Section 2.2 introduces inflation-indexed debt in a simple Fisherian model, deriving determinacy conditions and characterising the behaviour of inflation when the policy mix is fiscally led. Section 2.3 deepens the analysis of inflation-indexed debt in a general equilibrium setting. The empirical evidence is in Section 2.5, before Section 2.6 concludes.

Literature Review

This chapter relates squarely to the literature on fiscal-monetary interactions, as surveyed in chapter 1.

Closest related to us in terms of our modelling framework are Angeletos et al. (2024b) and Rachel and Ravn (2025), who themselves work with a model building on the seminal contribution of Blanchard (1985), facilitating a tractable way of introducing non-Ricardian fiscal policies using simple transfer shocks, which do not directly alter the resource constraint present in the economy.[2] Angeletos et al. (2024b) and Rachel and Ravn (2025) prove the conditions for existence and uniqueness of dynamic competitive equilibria in a NK-OLG model, delivering boundary conditions under which equilibria can be characterized. Rachel and Ravn (2025) deliver two particularly insightful results related to our work; first, indeterminacy becomes a less pervasive issue once households are non-Ricardian, i.e., adherence to the Taylor principle is not strictly necessary to deliver unique saddle-path stable equilibria even when fiscal policy is conventionally 'passive'; and second, the classic distinction between 'active' and 'passive' policies should not be taken literally once households are non-Ricardian. Relative to this body of work, our additional consideration of inflation-indexed debt innovates on a number of dimensions. First, inflation-indexed

[2] Working with government spending shocks, for instance, would alter the resource constraint in the economy. Without further assumptions on the household utility borne by such government spending, however, any such government spending would be 'wasteful'.

debt overcomes the singularity result of Angeletos et al. (2024b) that RANK-FTPL and HANK models converge to the same limit point; second, inflation-indexed debt generally increases the response of inflation in response to transfer shocks; and finally, inflation-indexed debt alters one of the conclusions of Rachel and Ravn (2025), namely that the explosive region becomes more problematic under non-Ricardianness, providing an additional anchor favouring determinate equilibria.

Banerjee et al. (2022) evaluates the prevalence of fiscally-driven policy mixes empirically with a focus on a wider set of advanced economies. A narrative example of a recent fiscal shock informing inflation rates is provided by Hazell and Hobler (2024), who focus on the 2021 Georgia Senate election run-off.

Albeit indirectly, we also contribute to the literature on inflation-linked government bonds. The plausibly earliest contribution in this field is Fischer (1975), who derives household demand for indexed bonds in a multi-asset optimal allocation framework. The special insurance properties of such inflation-linked debt are extensively discussed in Barr and Campbell (1997), Garcia and van Rixtel (2007), Gürkaynak et al. (2010) and Andreasen et al. (2021). Notably, Sims (2013) briefly mentions the possible detrimental consequences of indexed debt in frameworks that extensively feature fiscally-driven policy mixes. This chapter expands on his remarks using simple, analytical modelling frameworks. Schmid et al. (2026) provide a systematic analysis of inflation-indexed debt as a policy tool, emphasizing its role as an ex-ante commitment device against inflation. With our contribution, we do not invalidate the conclusions of Schmid et al. (2026) about the desire for indexed bonds as plausibly overcoming a certain degree of market incompleteness, but we instead add to their results by qualifying the possibility of risks associated with the issuance of inflation-indexed debt that have henceforth been hidden from the literature.

2.2 Example from a Fisherian model

We illustrate how inflation-indexed debt matters in a canonical Fisherian model with a representative household saving in a government bond that is partially indexed to inflation. We do so to analyse how the presence of indexed debt can influence the equilibrium properties in dynamic economies, and provide a brief characterization of possible price level dynamics.

A representative household receives a constant stream of goods Y, maximizing its expected present value of utility from consumption, discounted at the time-invariant rate β:

$$\max_{\{c_t,b_t\}_{t=0}^{\infty}} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t u(c_t),$$

subject to the flow budget constraint

$$P_t c_t + q_t b_t = P_t(Y - T_t) + \Pi_t^{\theta}\, b_{t-1}, \tag{2.1}$$

where b_t is an asset that is a *partially*-indexed government-issued security. The asset is indexed to the gross rate of inflation by θ, in a manner that is reflective of how the face value of indexed debt is formed empirically (Hall and Sargent, 2011). The government-issued asset is bought at the market price q_t. T_t denotes real lump-sum taxes raised by the government.

The household optimality conditions are standard:

$$\begin{aligned} \{c_t\}: &\qquad u'(c_t) = \lambda_t P_t \\ \{b_t\}: &\qquad \lambda_t q_t = \beta\, \mathbb{E}_t\Big[\lambda_{t+1}\, \Pi_{t+1}^{\theta}\Big] \end{aligned}$$

The asset pricing equation is:

$$q_t = \beta\, \mathbb{E}_t\left[\frac{u'(c_{t+1})}{u'(c_t)}\, \Pi_{t+1}^{\theta-1}\right]. \tag{2.2}$$

The flow budget constraint of the government is:

$$\Pi_t^{\theta} b_{t-1} = P_t T_t + q_t b_t, \tag{2.3}$$

with the bond pricing kernel being given by the household first-order conditions and T_t denoting real surpluses raised by the government, which are equivalent to lump-sum taxes raised in our simple environment.

The fiscal policy rule reacts to the debt-to-GDP ratio by adjusting the tax rate τ_t:

$$\frac{\tau_t}{\tau} = \left(\frac{s_{t-1}}{s}\right)^{\gamma} e^{\varphi_t}; \qquad s_t \equiv \frac{q_t b_t}{P_t Y}, \quad \tau_t \equiv \frac{T_t}{Y},$$

where $\tau_t \equiv \frac{T_t}{Y}$ are surpluses raised by the government as a fraction of output, and $s_t \equiv \frac{q_t b_t}{P_t Y}$ is the market value of the government-issued partially-indexed bond relative to GDP. φ_t is a standard AR(1) shock to the present lump-sum tax rate (the exogenous fiscal policy disturbance in the model), and the fiscal policy reaction coefficient is γ.

The monetary policy rule targets the return $R_t = 1/q_t$ on the partially-indexed bond:

$$R_t = \frac{1}{\beta}\, \Pi_t^{\phi}. \tag{2.4}$$

For completeness, we explore the model properties under a standard Taylor Rule in appendix B.1.

Equilibrium

Combining the household budget constraint (2.1) with the government budget constraint (2.3), consumption must be equal to the endowment:

$$c_t = Y \ \forall t,$$

and the government bond pricing kernel simplifies to:

$$q_t = \beta \, \mathbb{E}_t \left[\Pi_{t+1}^{\theta-1} \right],$$

reflecting the partial indexation of the government-issued asset and the resulting dependence of its price on expected inflation. We now linearize the equilibrium conditions around a deterministic zero-inflation steady-state ($\Pi = 1$), consistent with household optimality and the fiscal and monetary rules. We denote variables in their log-deviations from steady-state with hats. Log-linearization around the zero-inflation steady-state gives the following system of equations:

1. Nominal bond prices:

$$-\hat{R}_t = (\theta - 1) \, \mathbb{E}_t \hat{\pi}_{t+1}. \tag{2.5}$$

2. Monetary rule:

$$\hat{R}_t = \phi \hat{\pi}_t. \tag{2.6}$$

3. Law of motion of debt:

$$\frac{\pi_t^{\theta} b_{t-1}}{P_t Y} = \frac{P_t T_t}{P_t Y} + \frac{q_t b_t}{P_t Y} \quad \Longrightarrow \quad \pi_t^{\theta-1} \, s_{t-1} \, R_{t-1} \; = \; \tau_t + s_t.$$

   Log-linearized:

$$(\theta - 1) \, \hat{\pi}_t + \hat{s}_{t-1} + \hat{R}_{t-1} = \frac{\tau}{\tau + s} \, \hat{\tau}_t + \frac{s}{\tau + s} \, \hat{s}_t.$$

$$\tau + s = sR\pi^{\theta-1}, \qquad \frac{1}{R} = \beta \pi^{\theta-1} \quad \Longrightarrow \quad \tau + s = \frac{s}{\beta}, \qquad \frac{\tau}{s} = \frac{1-\beta}{\beta}.$$

$$(\theta - 1) \, \hat{\pi}_t + \hat{s}_{t-1} + \hat{R}_{t-1} = (1 - \beta) \, \hat{\tau}_t + \beta \, \hat{s}_t. \tag{2.7}$$

4. Fiscal rule:

$$\hat{\tau}_t = \gamma \, \hat{s}_{t-1} + \varphi_t. \tag{2.8}$$

Combining (2.5) and (2.6):

$$\phi\,\hat{\pi}_t = (1-\theta)\,\mathbb{E}_t\hat{\pi}_{t+1}, \tag{2.9}$$

Forward (2.7) one period, take expectations at t and use $-\hat{R}_t = (\theta-1)\mathbb{E}_t\hat{\pi}_{t+1}$:

$$\hat{s}_t = (1-\beta)\,\mathbb{E}_t\hat{\tau}_{t+1} + \mathbb{E}_t\beta\hat{s}_{t+1}, \tag{2.10}$$

in which inflation does not directly appear. Substitute in for $\hat{\tau}_{t+1} = \gamma\hat{s}_t + \varphi_{t+1}$ to obtain:

$$\hat{s}_t = (1-\beta)\gamma\hat{s}_t + (1-\beta)\,\mathbb{E}_t\varphi_{t+1} + \mathbb{E}_t\beta\hat{s}_{t+1}. \tag{2.11}$$

The system can be written as:

$$\underbrace{\begin{pmatrix} 1-\theta & 0 \\ 0 & \beta \end{pmatrix}}_{\equiv A_0} \mathbb{E}_t \begin{pmatrix} \hat{\pi}_{t+1} \\ \hat{s}_{t+1} \end{pmatrix} = \underbrace{\begin{pmatrix} \phi & 0 \\ 0 & 1-\gamma(1-\beta) \end{pmatrix}}_{\equiv A_1} \begin{pmatrix} \hat{\pi}_t \\ \hat{s}_t \end{pmatrix} + \underbrace{\begin{pmatrix} 0 \\ -(1-\beta) \end{pmatrix}}_{\equiv C} \mathbb{E}_t\varphi_{t+1}. \tag{2.12}$$

The matrix relevant for the determinacy properties is

$$Z = A_0^{-1}A_1 = \begin{pmatrix} \dfrac{\phi}{1-\theta} & 0 \\ 0 & \dfrac{1-\gamma(1-\beta)}{\beta} \end{pmatrix}. \tag{2.13}$$

Its eigenvalues are $\dfrac{\phi}{1-\theta}$ and $\dfrac{1-\gamma(1-\beta)}{\beta}$. A *monetary-led* equilibrium requires $\phi > 1-\theta$ and $\gamma > 1$. A *fiscally-led* equilibrium requires $\phi < 1-\theta$ and $\gamma < 1$. The introduction of inflation-indexed debt hence matters only for the determinacy boundaries on the monetary side of the model, but not on the fiscal side.

Figure 2.1 visualizes the determinacy properties of the model. In a nutshell, the presence of inflation-indexed debt shifts how *monetary* policy impacts the existence of a unique, saddle path-stable equilibrium. Once the share of inflation-indexed debt turns positive (right panel), the space under which a unique monetary-led equilibrium exists increases; that is, central banks can follow "more passive" monetary policies without running into the indeterminate (yellow) space. At the same time, a fiscal authority wishing to engage in activist fiscal policy without negative repercussions on the stability properties of the system must be aware that this requires an even more 'passive' monetary authority relative to conventional wisdom.

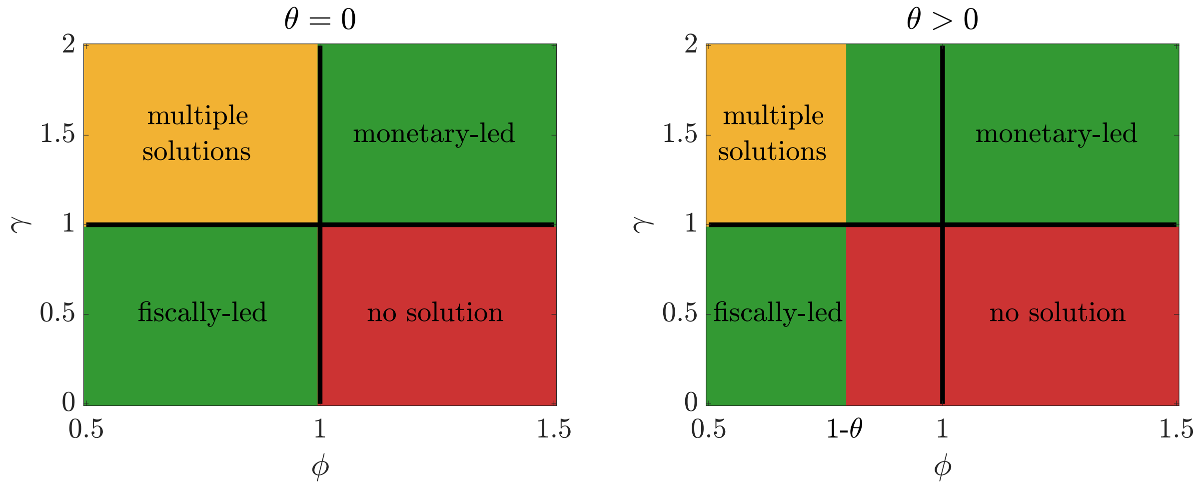

Figure 2.1: Determinacy properties of the Fisherian model.

The intuition behind this phenomenon is that indexed debt works as an 'automatic stabilizer' in terms of the determinacy properties of the model: as the central bank targets the return rate of the partially-indexed bond, any change it induces on its net return rate has *more* stabilizing effects relative to a canonical Taylor Rule, as a presumed increase in the net policy rate directly decreases the expected future inflationary pressure, thereby stabilizing the system beyond the simple first-order effect of the nominal policy rate adjustment.

To derive debt-inflation dynamics under the fiscally-led policy mix, the block-recursivity of the model means we need to solve the debt equation forward.[3] The debt equation is:

$$\Big(1-(1-\beta)\gamma\Big)\,\hat{s}_t = (1-\beta)\,\mathbb{E}_t\varphi_{t+1} + \mathbb{E}_t\beta\,\hat{s}_{t+1}.$$

so solving forward implies:

$$\hat{s}_t = \sum_{i=0}^{\infty}\left(\frac{\beta}{1-(1-\beta)\gamma}\right)^i\left(\frac{1-\beta}{1-(1-\beta)\gamma}\right)\mathbb{E}_t\varphi_{t+i+1}.$$

When φ_t is an AR(1) process with persistence ρ then $\mathbb{E}_t\varphi_{t+i+1} = \rho^{i+1}\varphi_t)$ and:

$$\begin{aligned}
\hat{s}_t &= \sum_{i=0}^{\infty}\left(\frac{\beta\rho}{1-(1-\beta)\gamma}\right)^i\left(\frac{1-\beta}{1-(1-\beta)\gamma}\right)\varphi_t \\
&= \frac{1}{1-\dfrac{\beta\rho}{1-(1-\beta)\gamma}}\left(\frac{1-\beta}{1-(1-\beta)\gamma}\right)\varphi_t \\
&= \left(\frac{1-(1-\beta)\gamma}{1-(1-\beta)\gamma-\beta\rho}\right)\left(\frac{1-\beta}{1-(1-\beta)\gamma}\right)\varphi_t,
\end{aligned}$$

[3]The solution for the monetary-led policy mix trivially features $\pi_t = 0\,\forall t$, as otherwise inflation would be explosive. The value of debt under a monetary-led policy mix evolves as $\hat{s}_t = \frac{1-(1-\beta)\gamma}{\beta}\hat{s}_{t-1} - \frac{1-\beta}{\beta}\varphi_t$.

which is independent of θ.

We solve for inflation dynamics in equilibrium by taking the linearized budget constraint and substituting in for the monetary policy rule to obtain:

$$\hat{\pi}_t = \frac{\phi}{1-\theta}\hat{\pi}_{t-1} - \frac{\beta}{1-\theta}\hat{s}_t + \frac{1-(1-\beta)\gamma}{1-\theta}\hat{s}_{t-1} - \frac{1-\beta}{1-\theta}\varphi_t. \tag{2.14}$$

Inflation-indexed debt here has the expected effect of scaling *surprise* inflation: When $\mathbb{E}_{t-1}\hat{\pi}_t \neq \hat{\pi}_t$, the fact that the equilibrium price for indexed debt had been *too low* in period $t-1$ induces additional surprise inflation in period t.

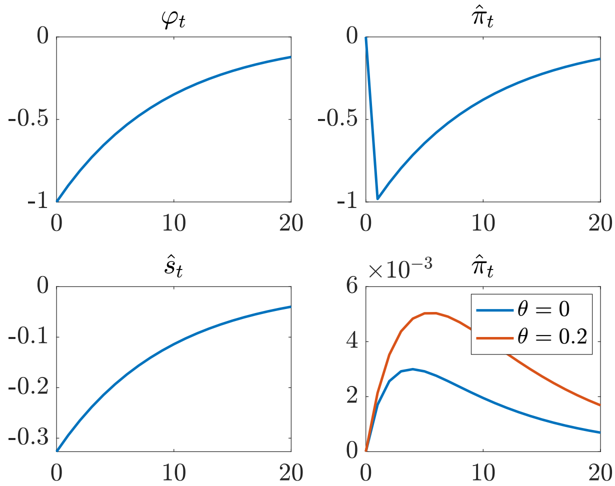

Figure 2.2: Dynamics of fiscally-led equilibrium in Ricardian model.

Figure 2.2 plots the equilibrium paths of the tax rate, the debt-to-GDP ratio, and the rate of inflation in response to a deficit shock φ_t for two levels of inflation-indexed debt: $\theta = 0$ and $\theta = 0.2$. As expected, the reaction of the tax rate and the debt-to-GDP ratio are independent of the share of inflation-indexed debt. The key difference lies in the observed rate of inflation: inflation-indexed debt causes additional loading on inflation under a fiscally-led policy mix. Higher equilibrium inflation is needed to ensure sufficient devaluation of the government debt stock, since only the *non-indexed* share of the debt stock can be devalued. The inflation that follows a deficit shock is increasing in the share of inflation-indexed debt when the fiscal authority does *not* promise a complete repayment of borrowing.

In sum, inflation-indexed debt in the simple model makes the policy space qualifying as a 'fiscally-led policy mix' smaller, but the consequences of being within that space become more detrimental: the inflationary pressure arising from an expansionary deficit shock is increasing in the share of outstanding inflation-indexed debt.

2.3 A non-Ricardian General Equilibrium Model with Indexed Debt

We proceed with a richer, albeit tractable, general equilibrium framework based on Angeletos et al. (2024b), who develop an OLG-NK model in the spirit of Blanchard (1985).[4] They introduce a mortality friction that breaks Ricardian equivalence, which can be considered a general proxy for the liquidity risk in canonical heterogeneous-agent models. It allows for the analysis of purely distributive *transfer* shocks with relevant real and nominal effects, facilitating a clean analysis of the effects of inflation-indexed debt.

Model framework

As in the previous section, government bonds D_t are partially inflation-indexed, and can heuristically be decomposed into constant proportions of non-indexed and indexed debt:

$$D_t \equiv (1-\theta)B_t + \theta\tilde{B}_t, \tag{2.15}$$

where θ is the time-invariant share of inflation-indexed debt. Inflation is zero in steady-state, with debt d_t, taxes t_t and household assets a_t measured in absolute deviations to allow for the possibility of zero-debt steady-states.

Households

The probability of the household surviving (i.e., the inverse of the mortality risk) from one period to another is $\omega \in (0,1]$. Households are replaced whenever they die, keeping overall population levels constant. Household i maximizes the present discounted value of its utility:

$$\mathbb{E}_t \sum_{t=0}^{\infty} (\beta\omega)^k \left(\frac{C_{i,t+k}^{1-\frac{1}{\sigma}} - 1}{1-\frac{1}{\sigma}} - \nu \frac{L_{i,t+k}^{1+\frac{1}{\varphi}}}{1+\frac{1}{\varphi}} \right), \tag{2.16}$$

where β is a common discount factor, ν scales the disutility of labour supply L_{it}, σ is the elasticity of intertemporal substitution and φ is the Frisch elasticity of labour supply.

Households trade the partially-indexed government bond, D_t, earning a nominal rate of return R_t^D/ω if the household survives. In effect, there is a risk-free return of R_t^n/ω on the fraction $1-\theta$ of debt that is not indexed and a return of $\frac{R_t^n}{\omega}\left(\frac{P_{t+1}}{P_t}\right)$ on the fraction θ that is indexed:

$$R_t^D = (1-\theta)R_t^n + \theta R_t^n \left(\frac{P_{t+1}}{P_t}\right) = R_t^n \left(1 + \theta\left(\frac{P_{t+1}}{P_t} - 1\right)\right). \tag{2.17}$$

[4] Rachel and Ravn (2025) work in a similar framework.

The budget constraint follows Angeletos et al. (2024b), in that all households receive labour income $W_t L_{it}$ and dividends X_{it}, and pay lump-sum taxes T_{it}. Existing households make a contribution S_{it} to a social fund whose proceeds are distributed to newborn households, eliminating wealth effects from mortality risk.[5] The household-specific budget constraint is:

$$P_{t+1}A_{i,t+1} = \frac{R_t^D}{\omega}P_t\left(A_{i,t} + \underbrace{Y_{i,t}}_{\equiv W_t L_{i,t}+X_{i,t}} - C_{i,t} - T_{i,t} + S_{i,t}\right). \tag{2.18}$$

We retain all other household-side assumptions from Angeletos et al. (2024b). Dividends are identical across households i, labour supply is intermediated by unions to obtain $L_{it} = L_t$, and income and taxes faced by households are equalized. Taking expectations and defining the *ex ante* real interest rate through the Fisher equation as $R_t = R_t^n \mathbb{E}_t P_t / P_{t+1}$, the first-order conditions are:

$$\{C_t\}: \qquad C_t^{-\frac{1}{\sigma}} - \lambda_t \frac{R_t}{\omega}(1 + \theta\mathbb{E}_t(\Pi_{t+1} - 1)) = 0, \tag{2.19a}$$

$$\{L_t\}: \qquad -\nu L_t^{\frac{1}{\varphi}} + \lambda_t \frac{R_t}{\omega}(1 + \mathbb{E}_t\theta(\Pi_{t+1} - 1))\frac{\partial Y_t}{\partial L_t} = 0, \tag{2.19b}$$

$$\{A_{t+1}\}: \qquad -\lambda_t + \beta\omega\mathbb{E}_t\left(\lambda_{t+1}\frac{R_{t+1}}{\omega}(1 + \theta\mathbb{E}_{t+1}(\Pi_{t+2} - 1))\right) = 0, \tag{2.19c}$$

which implies the Euler equation for consumption:

$$\frac{C_t^{-\frac{1}{\sigma}}}{R_t(1 + \mathbb{E}_t\theta(\Pi_{t+1} - 1))} = \beta\mathbb{E}_t C_{t+1}^{-\frac{1}{\sigma}}. \tag{2.20}$$

The log-linearized Euler equation is:

$$c_t = -\sigma(r_t + \theta\mathbb{E}_t\pi_{t+1}) + \mathbb{E}_t c_{t+1}. \tag{2.21}$$

The present value budget constraint is obtained from the Euler equation and per-period budget constraint. As in Angeletos et al. (2024b), it is a generalization of the Permanent Income Hypothesis:

$$\begin{aligned} c_t =& (1 - \beta\omega)\left(a_t + \mathbb{E}_t\sum_{s=0}^{\infty}(\beta\omega)^s(y_{t+s} - t_{t+s})\right) \\ & - \beta\left(\sigma\omega - (1 - \beta\omega)\frac{A^{SS}}{Y^{SS}}\right)\mathbb{E}_t\left[\sum_{s=0}^{\infty}(\beta\omega)^s(r_{t+s} + \theta\pi_{t+1+s})\right]. \end{aligned} \tag{2.22}$$

a_t are net asset holdings of households and, crucially, there is an additional inflation adjustment in the last term in proportion to the share of indexed debt. This is the novelty relative to Angeletos et al. (2024b), capturing income from inflation for holders of inflation-indexed debt.

[5] The transfers are $S_{it}^{new} = D^{SS} \geq 0$ and $S_{it}^{old} = -\frac{1-\omega}{\omega}D^{SS} \leq 0$, such that $(1-\omega)S^{new} + \omega S^{old} = 0$.

Aggregate supply

The supply side of the model is standard and follows the canonical New Keynesian structure. The New Keynesian Phillips Curve is due to pricing frictions in the firm's problem:

$$\pi_t = \kappa y_t + \beta \mathbb{E}_t \pi_{t+1}, \tag{2.23}$$

which, iterated forward, gives inflation as a function of current and future output gaps:

$$\pi_t = \kappa \sum_{k=0}^{\infty} \beta^k \mathbb{E}_t y_{t+k}. \tag{2.24}$$

Government

The government budget constraint is:

$$D_{t+1} = R_t^D (D_t - P_t T_t),$$

where D_t is the stock of government debt and $R_t^D = R_t^n(1 + \theta(\Pi_{t+1} - 1))$ is the *ex post* return the government has to pay. The return is defined in terms of the *realised* rate of inflation, from t to $t+1$.

Dividing both sides by P_t and denoting the real value of government debt by $D_t^R = D_t/P_t$,[6] a linear approximation around a zero-inflation steady-state yields:

$$(D_{t+1}^R - D^{R,SS}) + D^{R,SS}(\Pi_{t+1} - 1) = R^{D,SS}[(D_t^{R,SS} - D^{R,SS}) - (T_t - T^{SS})] + (D^{R,SS} - T^{SS})(R_t^D - R^{D,SS}).$$

Dividing both sides by steady-state output Y^{SS} and defining $d_t = (D_t^R - D^{R,SS})/Y^{SS}$ for debt and $t_t = (T_t^R - T^{SS})/Y^{SS}$ for taxes, the linearized budget constraint becomes:

$$d_{t+1} + \frac{D^{SS}}{Y^{SS}} \pi_{t+1} = \frac{1}{\beta}(d_t - t_t) + \frac{D^{SS}}{Y^{SS}} r_t^D,$$

where we simplify using the steady-state nominal return $R^{D,SS} = 1/\beta$ and the steady-state government budget constraint $D^{SS} = (D^{SS} - T^{SS})/\beta$.

The *ex post* return r_t^D is related to the risk-free nominal return r_t^n and *realised* inflation π_{t+1} through the inflation-indexation of the government bond, with the risk-free nominal return r_t^n itself related to the *ex ante* real return r_t and *expected* inflation $\mathbb{E}_t \pi_{t+1}$ through the Fisher equation. As approximate log-deviations from steady state, the returns satisfy:

$$\begin{aligned} r_t^D &= r_t^n + \theta \pi_{t+1}, \\ r_t &= r_t^n - \mathbb{E}_t \pi_{t+1}, \\ r_t^D &= r_t + \mathbb{E}_t \pi_{t+1} + \theta \pi_{t+1}. \end{aligned}$$

[6] The resulting government budget constraint is $D_{t+1}^R \Pi_{t+1} = R_t^D (D_t^R - T_t)$.

The final expression for the government budget constraint is:

$$d_{t+1} = \frac{1}{\beta}(d_t - t_t) + \frac{D^{SS}}{Y^{SS}} r_t - \frac{D^{SS}}{Y^{SS}} \left((1-\theta)\pi_{t+1} - \mathbb{E}_t \pi_{t+1}\right). \tag{2.25}$$

The novelty here is the multiplication of *realized* future inflation by $(1-\theta)$, which is the share of non-indexed debt. Intuitively, inflation-indexed debt cannot be devalued through surprise inflation. The ability of governments to devalue their debt in real terms is therefore constrained by indexed debt, with implications for the evolution of the debt stock in the presence of inflation.[7]

The steady-state of the model requires $x_{-1} = 0 \; \forall x \in \{d, t, r, y, \pi\}$. Starting from the steady state, equation (2.25) defines the initial change in the debt stock as a function of surprise inflation:

$$d_0 = -\frac{D^{SS}}{Y^{SS}}(1-\theta)\pi_0,$$

and the presence of inflation-indexed debt generally reduces the initial reaction of the value of the debt stock to inflationary pressure.

We close the model with the monetary and fiscal and policy rules in Angeletos et al. (2024b):

$$r_t = \phi y_t, \tag{2.26}$$

$$t_t = -\varepsilon_t + \tau_d(d_t + \varepsilon_t) + \tau_y y_t. \tag{2.27}$$

Aggregate demand

Aggregate demand is derived by imposing market clearing conditions $a_t = d_t$ and $y_t = c_t$ and substituting the monetary policy rule (2.26) into the intertemporal budget constraint (2.22):

$$y_t = (1-\beta\omega)\left(d_t + \mathbb{E}_t \sum_{s=0}^{\infty} (\beta\omega)^s \left[\left(1 - \frac{\Gamma\phi}{1-\beta\omega}\right) y_{t+s} - t_{t+s} - \frac{\Gamma\theta}{1-\beta\omega}\pi_{t+1+s}\right]\right), \tag{2.28}$$

where $\Gamma \equiv \beta\left(\sigma\omega - (1-\beta\omega)\frac{D^{SS}}{Y^{SS}}\right)$.

In recursive form, the aggregate demand relationship simplifies to:

$$\mathbb{E}_t y_{t+1} - \frac{\Gamma\theta}{\beta\omega}\mathbb{E}_t\pi_{t+1} - (1-\beta\omega)\mathbb{E}_t d_{t+1} = \frac{\beta\omega(1-\tau_y) + \Gamma\phi + \tau_y}{\beta\omega} y_t - \frac{(1-\beta\omega)(1-\tau_d)}{\beta\omega}(d_t + \varepsilon_t), \tag{2.29}$$

where taxes have been substituted out using the fiscal policy rule (2.27).

[7] We furthermore impose the no-Ponzi condition following Angeletos et al. (2024b), i.e., $\mathbb{E}_t \underset{T\to\infty}{lim} \beta^T d_{t+T} = 0$, which arises from rational optimizing behaviour on behalf of the households.

Equation (2.29) generalizes the aggregate demand equation to account for inflation-indexed debt. Through θ, inflation indexation weakens the relationship between current output and future inflation. Together with the New Keynesian Phillips Curve (2.23) and the government debt valuation equation (2.25), it gives a full characterization of the model's dynamics.

Equilibrium

Definition 2 *A competitive equilibrium is a series* $\{c_t, y_t, \pi_t, a_t, d_t, t_t, r_t\}_{t=0}^{\infty}$ *that satisfies the aggregate demand function, the New Keynesian Phillips Curve, market clearing, the government's flow budget constraint, the no-Ponzi condition, and the monetary and fiscal policy rules.*

Without mortality risk

We begin by abstaining from mortality risk and setting $\omega = 1$. This representative agent New Keynesian (RANK) environment retains the fundamental block recursivity of the system, whereby the dynamics of output and inflation are independent of the dynamics of debt. To see this, substitute the government budget constraint (2.25) with the policy rules (2.26) and (2.27) into the aggregate demand equation (2.29) under $\omega = 1$:

$$\mathbb{E}_t y_{t+1} - \theta\sigma\mathbb{E}_t\pi_{t+1} = (1 + \sigma\phi)y_t, \tag{2.30}$$

The output-inflation block in the absence of mortality risk is therefore independent of the evolution of government debt, even with indexed debt. This mirrors the canonical formulations of RANK models (Kaplan, 2025a; Rachel and Ravn, 2025), except for the adjustment for indexed debt. The above equation combines with the New Keynesian Phillips Curve (2.23) and the government budget constraint (2.25) to define the first-order system:

$$\mathbb{E}_t \begin{bmatrix} y_{t+1} \\ \pi_{t+1} \\ d_{t+1} \end{bmatrix} = \begin{bmatrix} 1 + \phi\sigma - \frac{\kappa\sigma\theta}{\beta} & \frac{\sigma\theta}{\beta} & 0 \\ -\frac{\kappa}{\beta} & \frac{1}{\beta} & 0 \\ \frac{D^{SS}}{Y^{SS}}\phi - \frac{\tau_y}{\beta} - \frac{D^{SS}}{Y^{SS}}\frac{\kappa\theta}{\beta} & \frac{D^{SS}}{Y^{SS}}\frac{\theta}{\beta} & \frac{1}{\beta}(1 - \tau_d) \end{bmatrix} \begin{bmatrix} y_t \\ \pi_t \\ d_t + \varepsilon_t \end{bmatrix}. \tag{2.31}$$

The system is block recursive for all $\theta \geq 0$, so inflation indexation of debt in itself does not overcome the result that government debt is block-exogenous with respect to output and inflation in RANK models (Leeper, 1991; Angeletos et al., 2024b). The policy space that supports a unique saddle-path stable equilibrium will therefore still split into areas conventionally labelled 'active" and 'passive' with respect to the actions of each policymaker.

The system has one predetermined variable d_t and two forward looking variables (y_t, π_t), so determinacy requires one eigenvalue inside and two eigenvalues outside the unit circle. In the fiscally-led equilibrium, the eigenvalue $\beta^{-1}(1-\tau_d)$ associated with the government budget constraint must be outside the unit circle, which requires $\tau_d < 1-\beta$. The conditions under which the other two eigenvalues straddle the unit circles in the fiscally-led equilibrium are as follows:

Proposition 3 *To ensure a unique saddle-path equilibrium when $\omega = 1$ under a fiscally-led policy mix ($\tau_d < 1-\beta$), the restriction on ϕ, the monetary policy parameter, tightens with the degree of inflation-indexed debt. The formal restriction for a fiscally-led equilibrium in the absence of mortality risk is:*

$$\max\left\{-\frac{1}{\sigma}, -\frac{1}{\sigma}\left[2-\frac{\kappa\sigma\theta}{1+\beta}\right]\right\} < \phi < -\frac{\kappa\theta}{1-\beta}. \tag{2.32}$$

Proof. See appendix B.2. ■

The conditions on the monetary policy parameter ϕ are tighter for determinacy in the fiscally-led regime with a positive share of inflation-indexed debt. A larger share θ of indexed debt, or a steeper Phillips curve (higher κ) all push the right-hand bound in (2.32) leftward, so increased monetary passivity is needed to sustain fiscal dominance. Otherwise, there are multiple equilibria with self-fulfilling fluctuations. The bounds on ϕ for monetary dominance can also tighten when the share of indexed debt is sufficiently high.[8] Figure 2.3 depicts Proposition 3, without indexed debt in the left panel and with indexed debt ($\theta = 0.2$) in the right panel. The affected combinations for a unique monetary-led equilibrium are not restricted by the presence of inflation-indexed debt (the top-right green area stays the same in the left and right panels), but any fiscally-led equilibrium can only be sustained by increased degrees of monetary passivity (the bottom-left green area gets smaller).

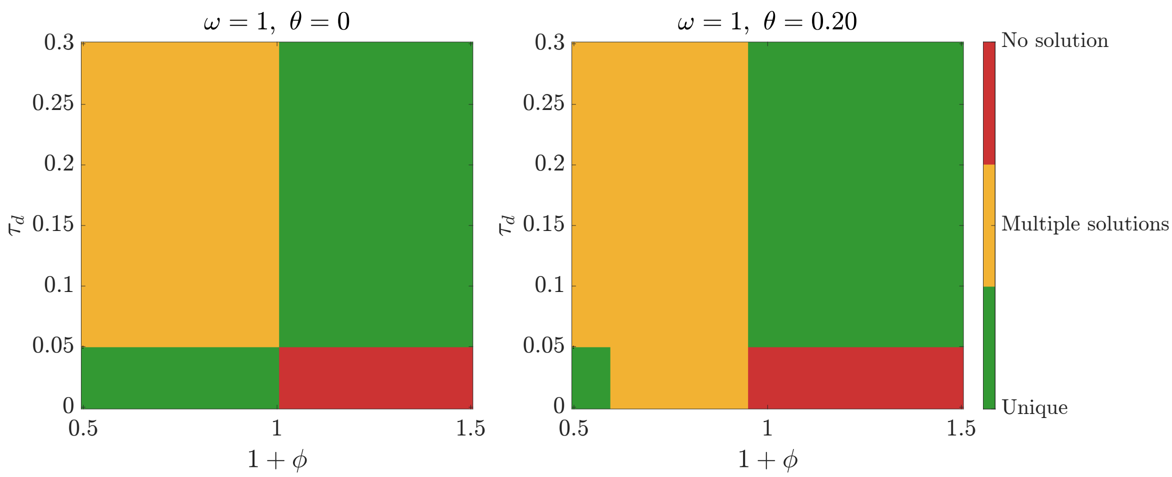

Figure 2.3: The determinacy properties of the RANK model.

[8] See Appendix B.3 for technical details.

With mortality risk

The first-order system for the joint dynamics of $\{y_t, \pi_t, d_t\}$ with $\omega \neq 1$ is:

$$\begin{pmatrix} 1 & -\frac{\Gamma\theta}{\beta\omega} & -(1-\beta\omega) \\ 0 & 1 & 0 \\ 0 & -\theta\frac{D^{SS}}{Y^{SS}} & 1 \end{pmatrix} \mathbb{E}_t \begin{pmatrix} y_{t+1} \\ \pi_{t+1} \\ d_{t+1} \end{pmatrix} = \begin{pmatrix} \frac{\beta\omega(1-\tau_y)+\Gamma\phi+\tau_y}{\beta\omega} & 0 & -\frac{(1-\beta\omega)(1-\tau_d)}{\beta\omega} \\ -\frac{\kappa}{\beta} & \frac{1}{\beta} & 0 \\ \frac{D^{SS}}{Y^{SS}}\phi - \frac{\tau_y}{\beta} & 0 & \frac{1}{\beta}(1-\tau_d) \end{pmatrix} \begin{pmatrix} y_t \\ \pi_t \\ d_t + \varepsilon_t \end{pmatrix}. \tag{2.33}$$

Multiplying by the inverse of the left-hand matrix defines a system in the usual form:

$$\mathbb{E}_t \begin{pmatrix} y_{t+1} \\ \pi_{t+1} \\ d_{t+1} \end{pmatrix} = A \begin{pmatrix} y_t \\ \pi_t \\ d_t + \varepsilon_t \end{pmatrix}. \tag{2.34}$$

This characteristic polynomial of A is relatively complex, precluding an *exact* parametric characterization of the conditions for saddle-path stable equilibria. Nonetheless, it is possible to provide an approximate characterization of the policy space. The details are in Appendix B.4.

Proposition 4 *The feasible region for a unique saddle path-stable equilibrium is constrained by a band $\phi^-(\tau_d;\theta) < \phi < \phi^+(\tau_d;\theta)$ when $\tau_d^0 < \tau_d < \tau_d^{\#}$. The band shifts with θ, with higher levels of θ making it less likely that a unique saddle-path equilibrium exists for low values of τ_d.*

Proof. See appendix B.4. ∎

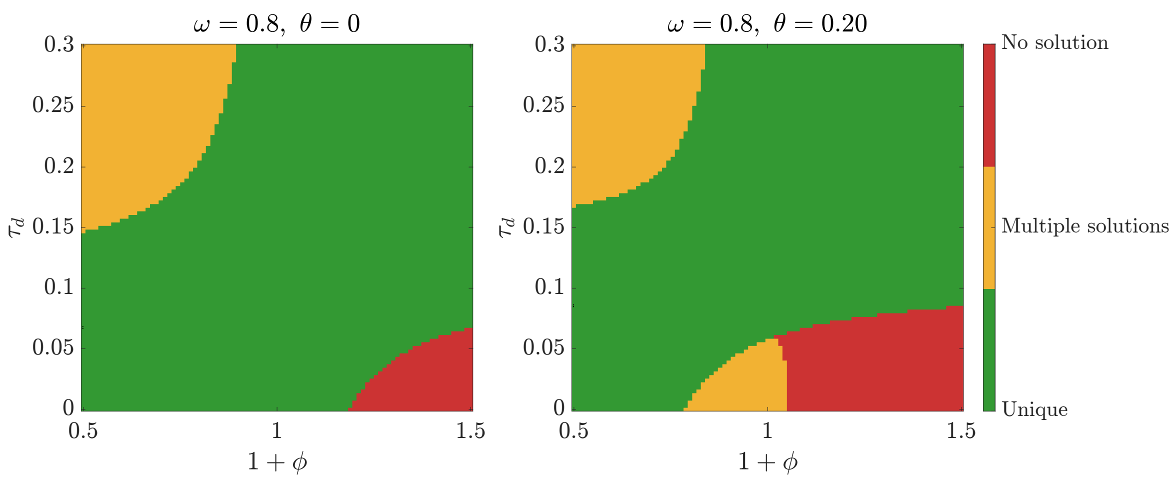

Figure 2.4: Visualization of Proposition 4.

Figure 2.4 visualizes Proposition 4, showing parametrizations of fiscal policy τ_d and monetary policy $1+\phi$ for which the first-order system with mortality risk has a unique saddle-path solution. The left panel depicts the policy space without inflation-indexed debt, replicating the result in Angeletos et al. (2024b) and Rachel and Ravn (2025) that mortality risk supports a unique equilibrium for a continuous policy space with no clear distinction between fiscally and monetary led equilibria. The right panel of Figure 2.4 depicts the same system when inflation-indexed debt is 20% of the overall government debt stock.[9] In line with Proposition 4, the green area shrinks at low values of τ_d and a lower monetary policy parameter $1 + \phi$ is needed to support a unique saddle-path equilibrium. In contrast, the green area expands slightly for high values of τ_d and the range of monetary policy parameters that supports equilibrium is larger with inflation-indexed debt.

What is the intuition behind this result? Inflation-indexed debt can be thought of as a type of *automatic stabilizer* inherent to the conduct of fiscal policy. For a given expansionary, deficit-financed fiscal measure, for instance, the presence of inflation-indexed debt should stabilize the value of that part of the government debt stock, as such debt retains its real value from the perspective of the household. However, when a fiscal policy authority commits to being *very* expansionary (i.e., committing to $\tau_d \searrow 0$), we run into a scenario where self-fulfilling fluctuations are possible if the *monetary* authority in turn engages in interest rate policies by which real interest rates remain approximately stable.

To see this, start from considering the fiscal policy authority. When it reacts insufficiently by appropriate taxation to changes of the real value of debt from equilibrium, it is possible that the raised additional taxes are *insufficient* to serve the additional face value of maturing inflation-indexed debt in response to an expansionary fiscal shock. When the monetary authority intends to keep real rates at least approximately stable, this problem compounds for the fiscal authority as the overall debt service cost also does not decrease. Through the combination of the presence of indexed debt, very expansionary fiscal policy, and a monetary authority that intends to 'fight back' against inflationary pressure, it is then possible for self-fulfilling fluctuations to materialize.

—

Under such a rich system, characterizing inflation analytically becomes difficult. We are nonetheless interested in the effect of inflation-indexed debt, and therefore perform the following experiment: for the entirety of the (relevant) policy space, we calculate the one-year change in the price level in response to a 1% deficit-to-GDP shock under two scenarios: once setting the share of inflation-indexed debt to 20% ($\theta = 0.2$), and once without any inflation-indexed debt ($\theta = 0$). Then, we take the difference of the two observed inflation rates, allowing us to characterize the impact of inflation-indexed debt on price level dynamics in the model. Figure 2.5 shows the results of this exercise, varying the core fiscal reaction parameter τ_d on the y-axis and the monetary policy parameter ϕ on the x-axis.

[9] This is above the share of TIPS in the stock of US Federal Debt, which is about 7.5%, but below the share of inflation-linked Gilts in the stock of UK Government debt, which is around 25%.

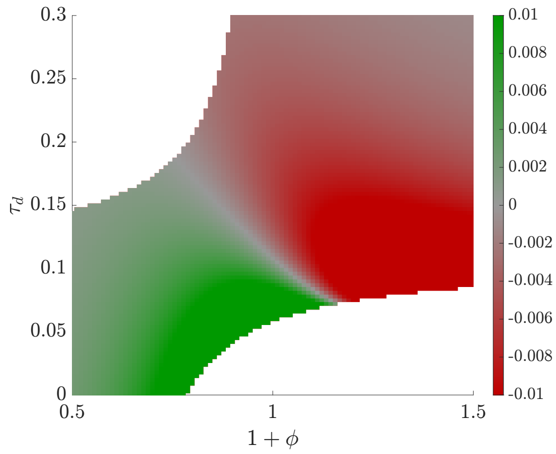

Figure 2.5: Difference in one-year inflation rates due to indexed debt in response to a 1% deficit-to-GDP shock in dependence on the policy parameters.

While the clear-cut distinction between fiscally-led and monetary-led policy mixes has become seemingly difficult under the presence of mortality risk, as clear from figure 2.4 and the interpretation in Rachel and Ravn (2025), figure 2.5 shows a very interesting novel result that arises due to the presence of inflation-indexed debt: we have one clear area in which inflation-indexed debt has a net *positive* impact on inflation, and one clear area in which it has a net *negative* impact on inflation, i.e., an area where inflation-indexed debt can be *dis*inflationary.[10] Low values of τ_d and ϕ are increasing the propensity of indexed debt to be inflationary, which can therefore be seen as a tool in helping measure approximate fiscal dominance, even when no clearly distinct policy areas persist.

How is it possible for inflation-indexed debt to induce slight *dis*inflationary pressure in response to a deficit shock? This holds particularly true for high values of ϕ, i.e., for monetary reaction parameters that induce a real rate increase in equilibrium. In such circumstances, the real value of inflation-indexed debt is *increasing* once a deficit shock materializes through the promise of the central bank to crush additional inflationary pressure arising from the deficit shock, yielding a surprise *positive* revaluation of the outstanding indexed debt stock. Then, the overall value of the stock of government debt is increasing through the presence of inflation-indexed debt, yielding some disinflationary pressure in equilibrium. The presented response of equilibrium inflation in dependence on the share of inflation-indexed debt also rationalizes the result in figure 2.4 to some extent, as the overall volatility of the price level is clearly most pronounced through the presence of indexed debt for low values of τ_d, irrespective of the monetary reaction parameter, which mostly influences the *direction* of the observed price level volatility.

[10]This was not the case in the RANK model, where inflation-indexed debt was either inflationary (under a fiscally-led policy mix), or had no impact on the unique solution $\pi_t = 0$ (in the monetary-led policy mix).

Relative to proposition 4, figure 2.6 sheds additional light on the effect of risk aversion. In particular, the relatively even-sided effect on determinacy properties caused by the presence of indexed debt (shifting the determinate space generally upwards) becomes increasingly *one-sided* as risk aversion goes up: for higher levels of risk aversion, inflation-indexed debt still reduces the feasible policy space for activist fiscal policy (restricting the number of monetary reaction parameters for which any given value of τ_d yields a unique saddle path-stable equilibrium); however, it does *not* increase the policy space under relatively passive fiscal policy (in the plot, the ratio of orange dots is virtually zero). Therefore, for high levels of risk aversion, inflation-indexed debt is generally only reducing the feasible policy space by eliminating some fiscal-monetary policy combinations.

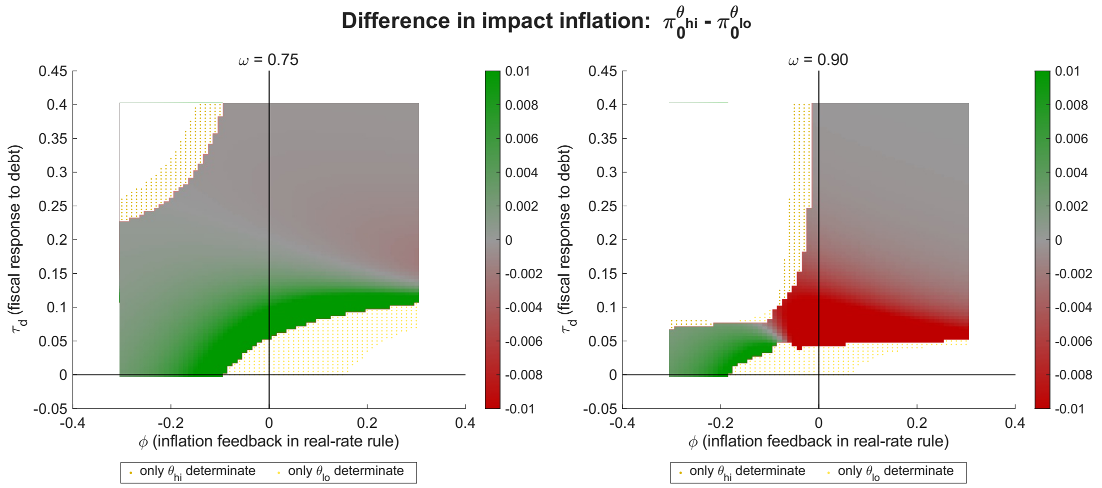

(a) $\sigma = 1$

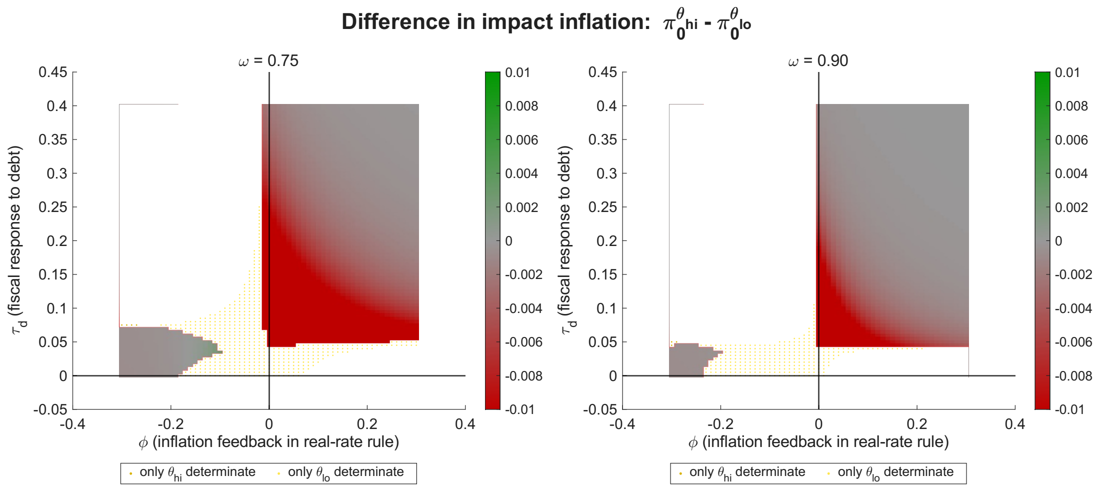

(b) $\sigma = 6$

Figure 2.6: The effect of risk aversion on the model, contrasting the cases of $\sigma = 1$ and $\sigma = 6$.

2.4 Alternative Policy Rules

In the previous section, we restricted ourselves to relatively simple policy rules that were broadly in line with Angeletos et al. (2024b) and Rachel and Ravn (2025). However, it is logical for the presence of inflation-indexed debt to alter the conduct of policymakers: both monetary and fiscal policy might anticipate some of the effects of indexed debt explored in section 2.2, internalizing them in their own policymaking. We here explore such alternative policy rules, showing how the shift in the determinate space can be unwound, or even inverted, with appropriate policy rules.

In section 2.2, we specified the monetary rule as equation (2.26) and the fiscal rule as equation (2.27). However, the presence of inflation-indexed debt can plausibly be accounted for by either policy authority, especially in the light of the possibly detrimental consequences in terms of equilibrium (in)existence or (dis)inflationary pressure.

It is possible to create a myriad of such policy specifications. Here, we focus on one particular instance, comprised of a monetary authority that aims to mitigate the impact of inflation-indexed debt on the aggregate demand equation, paired with a fiscal authority that is aware of the additional cost of serving such indexed debt upon the realization of a positive inflationary shock.

The alternative monetary rule: following equation (2.22), inflation-indexed debt manifests in the household policy function by inducing a direct wedge on equilibrium real interest rates. A monetary authority tasked with providing broad economic stabilization might then prefer to nullify this effect to the degree possible within its mandate. In this case, the central bank can simply follow a real rate rule that internalizes the additional pressure on the portfolio-adjusted real rate earned by the household; formally:

$$r_t = \phi y_t - \theta \mathbb{E}_t \pi_{t+1} \tag{2.35}$$

after log-linearization. The novel term is given by $-\theta \mathbb{E}_t \pi_{t+1}$. Under such a monetary rule, the policy rate set is effectively a *portfolio-adjusted real rate*, which acts as a stabilizing tool on the aggregate demand side in the presence of unexpected inflationary shocks.

Fiscal policy, in turn, can try to internalize anticipated expenditures arising from the presence of such debt. In that case, fiscal policy can simply be specified to follow:

$$t_t = -\varepsilon_t + \tau_d(d_t + \varepsilon_t) + \tau_y y_t + \frac{\beta\omega}{1-\beta\omega}\frac{D^{SS}}{Y^{SS}}\theta\tau_\pi \mathbb{E}_t \pi_{t+1} \tag{2.36}$$

after log-linearization. Here, the novel term in the rule is $\frac{\beta\omega}{1-\beta\omega}\frac{D^{SS}}{Y^{SS}}\theta\tau_\pi \mathbb{E}_t \pi_{t+1}$, with the novel policy parameter τ_π. The terms $\frac{\beta\omega}{1-\beta\omega}\frac{D^{SS}}{Y^{SS}}\theta$ weight this novel reaction parameter in line with intertemporal discounting and the overall outstanding indexed debt stock.

Factually, such a fiscal rule ensures that a fiscal policy authority expecting increased outlays for inflation-indexed debt attempts to raise additional funds through taxation to cover such expenses. Conditional on the return rates for either assets being equivalent in real terms, *expected* inflationary pressure would then even yield a windfall gain on behalf of the fiscal authority.

With these two rules, the aggregate demand equation (2.29) changes to the following specification:

$$\mathbb{E}_t y_{t+1} - \tau_\pi \frac{D^{SS}}{Y^{SS}} \theta \mathbb{E}_t \pi_{t+1} - (1-\beta\omega)\mathbb{E}_t d_{t+1} = \frac{\beta\omega(1-\tau_y) + \Gamma\phi + \tau_y}{\beta\omega} y_t - \frac{(1-\beta\omega)(1-\tau_d)}{\beta\omega}(d_t + \varepsilon_t). \tag{2.37}$$

Future inflation continues to matter for household aggregate demand, but now only through the distortionary taxation induced by the fiscal authority in proportion to its anticipated additional cost of serving maturing indexed debt. The initial 'base' effect through a wealth channel induced by a mismatch between the real interest rate set by the central bank and the factual portfolio return has been nullified.

The third equilibrium condition, the intertemporal debt valuation equation (initially given by equation (2.25)), similarly changes to the following expression after taking expectations and inserting our novel tax rule:

$$\mathbb{E}_t d_{t+1} + \tau_\pi \theta \frac{D^{SS}}{Y^{SS}} \mathbb{E}_t \pi_{t+1} = \left(\frac{D^{SS}}{Y^{SS}} \phi - \frac{\tau_y}{\beta} \right) y_t + \frac{1-\tau_d}{\beta}(d_t + \varepsilon_t). \tag{2.38}$$

The wedge term arising from the additional cost of serving indexed debt has vanished due to the adjustment of the monetary rule. Additionally, expected inflationary pressure can have a net positive impact on fiscal balance sheet through the taxation adjustment.

Summarizing equations (2.37) and (2.38), together with the NKPC (2.23), gives us the following 3×3 first-difference system:

$$\mathbb{E}_t \begin{bmatrix} y_{t+1} \\ \pi_{t+1} \\ d_{t+1} \end{bmatrix} = \begin{bmatrix} a_{11} & \frac{D^{SS}}{Y^{SS}} \tau_\pi \theta\omega & \frac{1-\beta\omega}{\beta}(1-\tau_d)\left(1-\frac{1}{\omega}\right) \\ -\frac{\kappa}{\beta} & \frac{1}{\beta} & 0 \\ \frac{D^{SS}}{Y^{SS}}\left(\phi + \frac{\kappa\,\tau_\pi\,\theta}{\beta}\right) - \frac{\tau_y}{\beta} & -\frac{D^{SS}}{Y^{SS}} \frac{\tau_\pi\,\theta}{\beta} & \frac{1}{\beta}(1-\tau_d) \end{bmatrix} \begin{bmatrix} y_t \\ \pi_t \\ d_t + \varepsilon_t \end{bmatrix}, \tag{2.39}$$

where $a_{11} \equiv 1 + \phi\sigma + \left(1-\beta\omega\right)\left(1-\frac{1}{\omega}\right)\left(\frac{D^{SS}}{Y^{SS}}\phi - \frac{\tau_y}{\beta}\right) - \frac{D^{SS}}{Y^{SS}}\,\kappa\,\tau_\pi\,\theta\omega.$

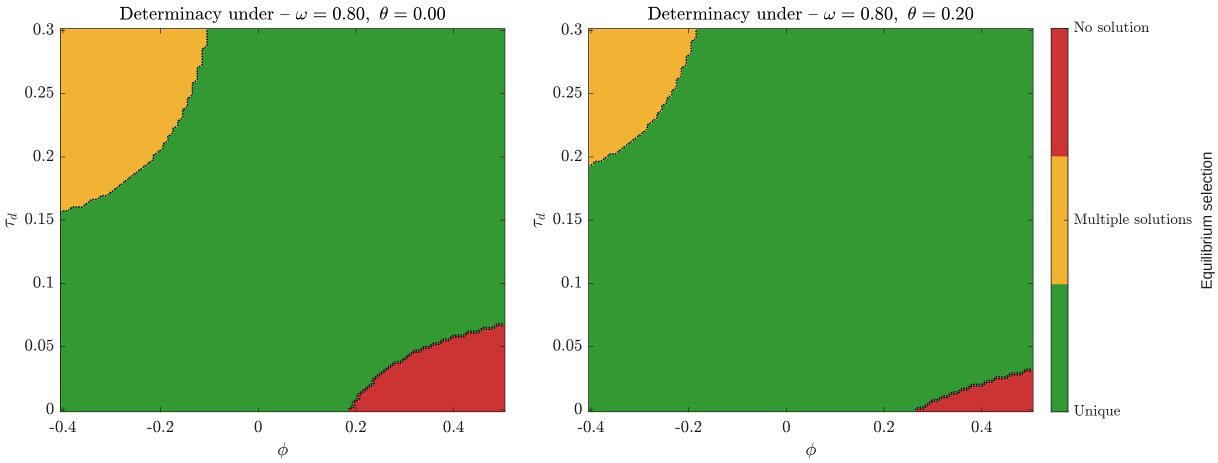

Figure 2.7: The determinacy properties of the model with mortality risk in dependence on the presence of indexed debt for the novel monetary and fiscal rules (2.35) and (2.36). The novel policy parameter τ_π has been set to 1.0.

Figure 2.7 depicts the impact of these novel rules on the determinacy properties of the system. In short, changing the fiscal and monetary rules to reflect a partial internalization of the anticipated effects of indexed debt *widens* the space under which saddle path-stable equilibria are possible unilaterally. In particular, even for relatively low values of τ_d, which indicate fiscal authorities committing to very expansionary policy measures, determinate equilibria are possible even for central bank policies that are tightening the economy in the presence of high inflation rates. In this scenario, the presence of indexed debt allows a foray into the area conventionally labelled as "active/active", which has thus far not been possible in other models, despite its possible real-world relevance in the light of the post-Covid inflationary episode.

Intuitively, the presence of the novel term in the fiscal rule facilitates more 'active' fiscal rules, since the presence of indexed debt does not weigh negatively on the fiscal budget constraint, since additional outlays for inflation-indexed debt are covered by the novel 'inflation stabilizer' in the fiscal rule (2.36).

What is the impact of the novel rules on inflationary pressure in the face of deficit shocks? Figure 2.8 sheds some light on this answer.

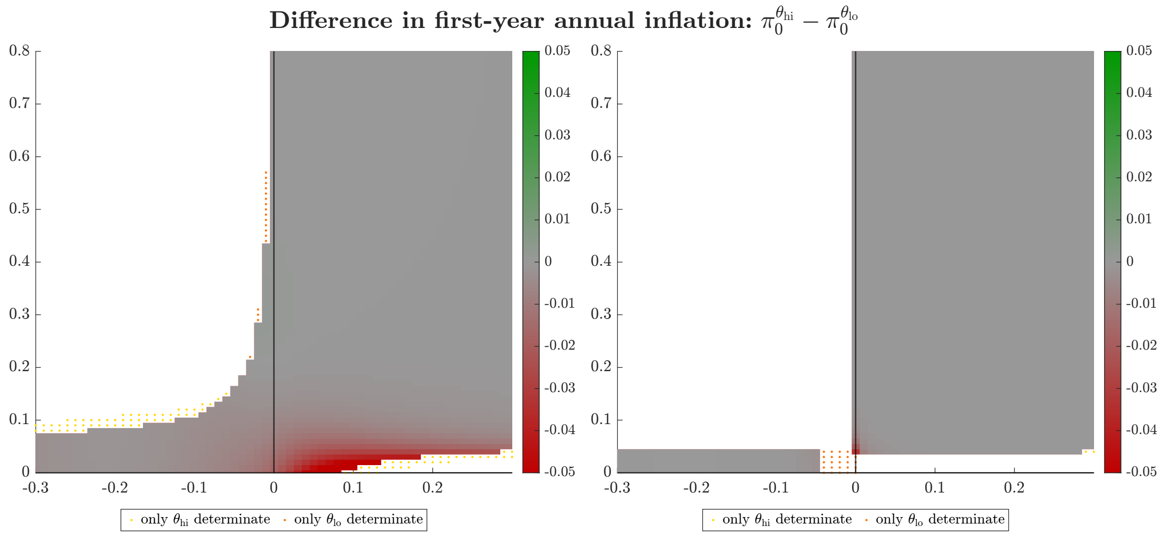

Figure 2.8: The difference in one-year inflation in the model with mortality risk between the cases $\theta = 0.2$ and $\theta = 0$ with fiscal rules (2.35) and (2.36). The novel policy parameter τ_π has been set to 1.0. In the left panel, $\omega = 0.8$; in the right panel, $\omega = 0.99$.

When the monetary rule absorbs the household-level windfall gain arising from their holdings of indexed debt in the presence of surprise inflation, the presence of indexed debt does *not* induce any additional inflationary pressure; by that same logic, the government's intertemporal budget constraint does not deteriorate through the presence of additional indexed debt. If anything, it is possible for inflation-indexed debt to depress inflationary pressure in response to deficit shocks when the government receives an additional windfall gain through its taxation in the wake of expected inflation. Note that, of course, the inflationary pressure in response to the deficit shock remains positive and particularly elevated for low levels of τ_d, but the change in the fiscal and monetary rules simply allows for a reduction in the magnitude of that pressure when indexed debt is present. Simply put, the *additional* inflationary pressure coming through indexed debt can be nullified.

2.5 The Empirics of Fiscally-led Policy Mixes and Inflation

We now supplement the model-based evaluations with an empirical analysis evaluating the possibility of fiscally-led policy regimes. Given that empirically distinguishing between policy regimes is a futile task (Neumeyer and Nicolini, 2025), we mostly provide correlational evidence that hints at the relevance of our indexed debt mechanism in periods that are more likely to be associated with fiscally-led policy mixes. At the end of this section, we utilize exogenously identified fiscal shocks and an exogenously identified regime classification to provide quasi-causal evidence in favour of our mechanism for the specific context of the US.

Correlations and non-causal regressions: the coincidence of probable fiscally-led policy mixes and inflation-indexed debt

At first, we intend to provide certain aspects motivating our analysis of inflation-indexed debt as a potential driver of inflationary episodes, as well as of the broad relevance of our indexed debt mechanism.

To do so, we first examine in a simple descriptive statistic whether inflation-indexed debt levels are descriptively related to materialized inflation rates. This task is conducted in figure 2.9, which plots the share of inflation-indexed debt against observed gross annual CPI inflation rates in the sample of BIS Member Countries issuing inflation-indexed debt since 1990.

While only correlational, the evidence presented here hinders the traditional message of inflation-indexed debt serving as a device ensuring that issuing governments will successfully curb inflationary pressure (Campbell and Shiller, 1996): there seems to be a clear positive correlation between observed shares of indexed debt and annual inflation rates in our long-running sample.

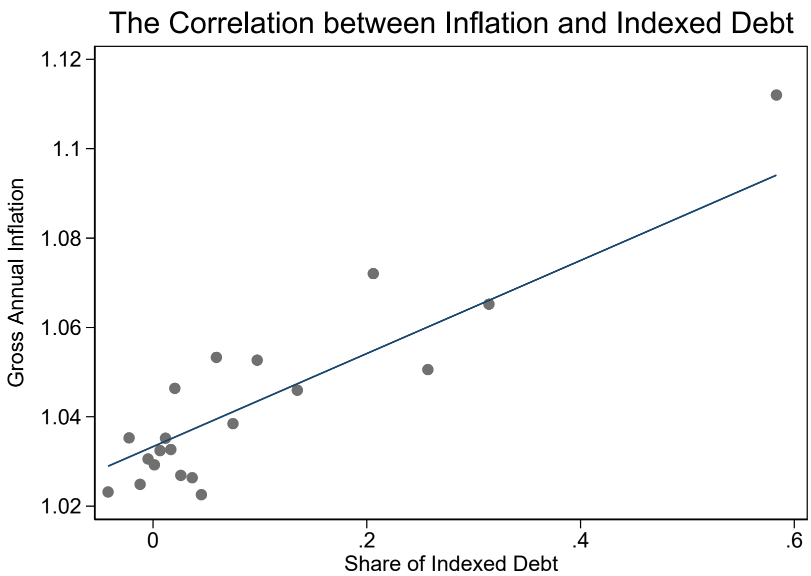

Figure 2.9: Binscatter of inflation-indexed debt against materialized inflation rates, controlling for the presence of fiscal rules, the degree of central bank independence, and a time-fixed effect in an annual sample of all BIS Member Countries issuing inflation-indexed debt.

Inflation-indexed debt therefore does not seem to eliminate inflationary pressures altogether. But how much of this is informed through the link between inflation-indexed debt and the monetary-fiscal policy mix? As alluded to before, and as established rigorously by Cochrane (2011) and Neumeyer and Nicolini (2025), estimating policy rules with empirical data in the hopes of distinguishing between fiscally-led and monetary-led policy mixes is a futile task. We therefore explicitly do not do so, and now mostly present non-causal evidence that suggestively hints at an effect of inflation-indexed debt on measures that can, to a certain extent, be proxies for the possibilities of monetary-led and fiscally-led policy mixes.

A natural proxy evaluating the monetary side is the degree of central bank independence, here measured through the Central Bank Independence Index of Romelli (2024). Intuitively, the index captures the degree to which a central bank is not constrained through other factors (e.g., partisan considerations) in its decision-making. A less independent central bank might hint at a reduced probability of a monetary-led policy mix. Figure 2.10 relates this independence index to the share of inflation-indexed debt present at a given country and at a given point in time. In this logic, we therefore follow Banerjee et al. (2022).

While again at best correlational, the evidence is again striking: this time, we observe a clearly negative correlational relationship between the share of inflation-indexed debt in the overall debt stock and the measured degree of central bank independence, in particular across countries since the results are less pronounced when focussing on within-country variation. While not making any causal statements, the observed data hints at a diminished probability of monetary-led policy mixes occurring when large swaths of the government debt stock are indexed to gross inflation.

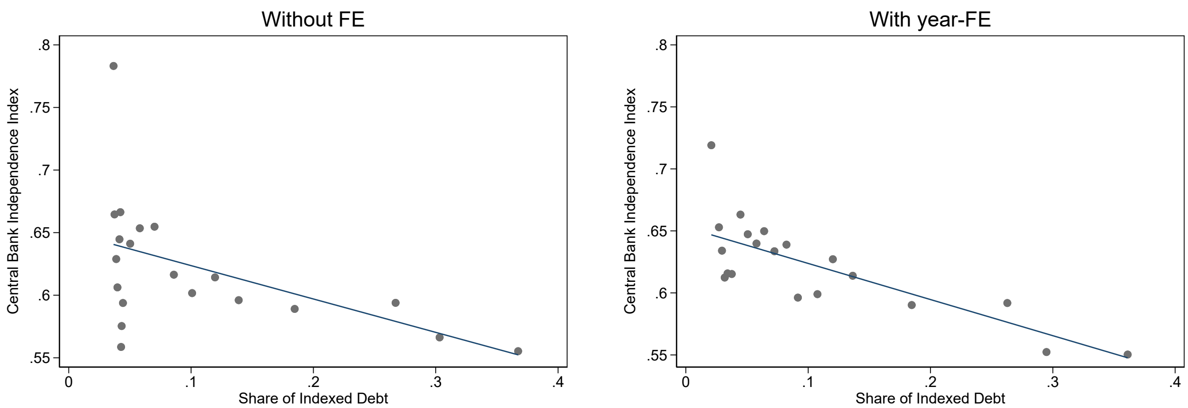

(a) Estimates without fixed effects   **(b)** Estimates with year fixed effects

Figure 2.10: Binscatter of the relationship between the share of inflation-indexed debt and the measured degree of central bank independence following Romelli (2024).

The following table 2.1 provides the numbers corresponding to figure 2.10, providing evidence on the correlation between measured central bank independence and the share of inflation-indexed debt in the overall government debt stock.[11]

[11] In appendix B.5, we provide some further robustness exercises.

Dep. var.	Central Bank Independence Index (Romelli, 2024)							
CB board index	0.687***	0.680***	0.505***	0.483***	0.686***	0.615***	0.500***	0.485***
	(0.0194)	(0.0201)	(0.0152)	(0.0169)	(0.0192)	(0.0187)	(0.0149)	(0.0164)
Fiscal Rule Intensity	0.0102***	0.0130***	0.0118***	0.00798***	0.0114***	0.0447***	0.00994***	0.00621***
	(0.00361)	(0.00397)	(0.00183)	(0.00183)	(0.00370)	(0.00445)	(0.00191)	(0.00181)
Indexed debt share	-0.182***	-0.198***	0.258***	0.105***	-0.313***	0.410***	0.253***	0.0666**
	(0.0413)	(0.0424)	(0.0336)	(0.0287)	(0.0490)	(0.0763)	(0.0414)	(0.0322)
FisRules × IndexDebt			-0.0893***			-0.490***	-0.0955***	
			(0.0153)			(0.0411)	(0.0161)	
Inflation					-0.00206*	-0.00125	-0.00155***	-0.00222***
					(0.00113)	(0.00108)	(0.000376)	(0.000377)
IndexDebt × Inflation					0.0169***	0.000288	0.00147	0.00382***
					(0.00387)	(0.00383)	(0.00132)	(0.00123)
Constant	0.251***	0.252***	0.326***	0.344***	0.260***	0.253***	0.337***	0.365***
	(0.0115)	(0.0120)	(0.00885)	(0.0112)	(0.0121)	(0.0112)	(0.00927)	(0.0115)
Obs.	605	605	605	605	605	605	605	605
R^2	0.707	0.713	0.978	0.980	0.717	0.778	0.979	0.981
F	482.5	451.3	296.8	41.32	304.1	317.9	210.2	42.46
R^2_{adj}	0.705	0.697	0.977	0.978	0.715	0.765	0.978	0.979
RMSE	0.106	0.107	0.030	0.029	0.104	0.095	0.029	0.028
Year-FE	No	Yes	No	Yes	No	Yes	No	Yes
Country-FE	No	No	Yes	Yes	No	No	Yes	Yes

Robust standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table 2.1: Regressions on possible predictors of central bank independence following Romelli (2024)

While establishing a proxy for the monetary side of the policy mix has already been ambitious, guidance for the fiscal side is even more difficult to come by due to the clearly simultaneous nature of fiscal policy setting with overall macroeconomic dynamics (Ramey and Zubairy, 2018). To nonetheless get a semblance of the idea, we exploit data from Davoodi et al. (2022) on the implementation of *fiscal rules* around the globe. While the nature of such rules and their implementation are relatively arbitrary, we zoom into *periods of rule suspension*; that is, periods in which previously established fiscal rules were suspended. In doing so, we focus on the question whether any rule had been suspended, as well as whether specific fiscal rules (expenditure rules, budget balancing rules, and deficit rules) were suspended. While again an imperfect proxy, we interpret here periods of fiscal rule suspension as periods of governments being willing to run fiscal policies that are normally considered unsustainable, therefore contributing to a possible fiscally-led policy mix.

The results of the logit regression predicting factors influencing the probability of a given fiscal rule suspension (for each of the four previously mentioned sets of rule suspensions) are presented in table 2.2.[12]

[12] Appendix B.5 shows additional results in relation to past levels of inflation-indexed debt.

Dep. var.	AnyR	ER	BR	DR	AnyR	ER	BR	DR
CB independence	-2.2267	-3.5126	-2.9157	0.4831	0.6227	13.5366	0.0185	0.9860
	(1.8226)	(3.9415)	(2.0313)	(4.9218)	(2.1084)	(10.9178)	(1.9724)	(4.8700)
CB board index	1.6759	4.8314	1.7574	-0.2362	-1.2981	-6.5182	-1.5719	-1.2533
	(1.5110)	(3.7678)	(1.7029)	(3.0294)	(1.4926)	(6.1250)	(1.4332)	(2.7346)
Indexed debt share	3.0572***	4.7854***	3.1249***	0.3678	6.3852***	13.6669***	6.4125***	2.0403
	(0.6578)	(1.0479)	(0.6751)	(0.9102)	(1.1965)	(4.0720)	(1.2509)	(1.3204)
Inflation	-0.0162	-0.0255	-0.0767*	-0.1554*	-0.0204	-0.1393**	-0.1335**	-0.2160*
	(0.0154)	(0.0219)	(0.0379)	(0.0674)	(0.0775)	(0.0432)	(0.0458)	(0.0958)
Constant	-3.1932***	-5.6121***	-2.7544***	-4.1307**	-1.1790	-8.8706*	-0.4267	-1.0932
	(0.5008)	(0.5327)	(0.4616)	(1.5229)	(0.9511)	(4.1165)	(0.8817)	(1.8025)
Obs.	652	652	652	652	285	261	285	95
ll	-103.8148	-51.7716	-91.2659	-42.2145	-59.0467	-18.4910	-56.0453	-24.6104
χ^2	27.9986	66.0330	27.5217	16.0084	70.0490	69.0798	70.0949	8.8445
p	0.0000	0.0000	0.0000	0.0030	0.0000	0.0000	0.0000	0.2640
R^2	0.1014	0.2750	0.1122	0.0218	0.3550	0.6779	0.3194	0.1034
Year-FE	No	No	No	No	Yes	Yes	Yes	Yes

Robust standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 2.2: Rule Suspension Regression. The outcome variables are defined as: $AnyR$ = "Any Fiscal Rule Suspended", ER = "Expenditure Rule Suspended", BR = "Budget Balancing Rule Suspended", DR = "Deficit Rule Suspended".

We generally observe a clear positive correlation between the share of inflation-indexed debt being present and the probability of rule suspension, except for periods of deficit rule suspension.[13] Taking our most comprehensive estimation in column five, we can observe that each percentage point increase in the share of inflation-indexed debt is correlated with a 0.634% increase in the probability of the suspension of any fiscal rule that had previously been put in place.

To visualize the results of table 2.2, figure 2.11 predicts the probability of the suspension of any fiscal rule as a function of the share of inflation-indexed debt. To maximize insights, figure 2.11 performs this analysis at both higher inflation levels of 5% annualized CPI inflation, as well as for lower levels of 1% annualized CPI inflation.

Irrespective of the level of inflation, figure 2.11 shows a general clear positive link between the share of inflation-indexed debt and the propensity of the suspension of any fiscal rule. That being said, the effect seems to be quite limited for most advanced economies, which boast inflation-indexed debt shares of below 30%.[14] There is, however, a significant difference between low-inflationary and high-inflationary environments in the propensities of a rule suspension, with the probability being higher in *low* inflation environments. While this might seem surprising, it can easily be rationalized as high-inflation environments

[13]The cross-country variation on the suspension of debt rules is quite low, however, impeding the statistical power of the estimation.

[14]Note, however that up to 80% of debt being indexed to factors other than the price level is possible. For instance, many emerging market economies issue large amounts of debt denominated to foreign currencies, which work tantamount to inflation-indexed debt in simple models of the Fiscal Theory of the Price Level (Cochrane, 2023).

are usually characterized by a lack of *surprise* inflationary pressure, as higher inflation levels usually imply that corresponding market expectations (and therefore market prices of inflation-indexed debt) reflect expectations of pronounced inflationary episodes. But inflation-indexed debt becomes especially costly during surprise inflation episodes, which must arise from an initial point of relative price stability.

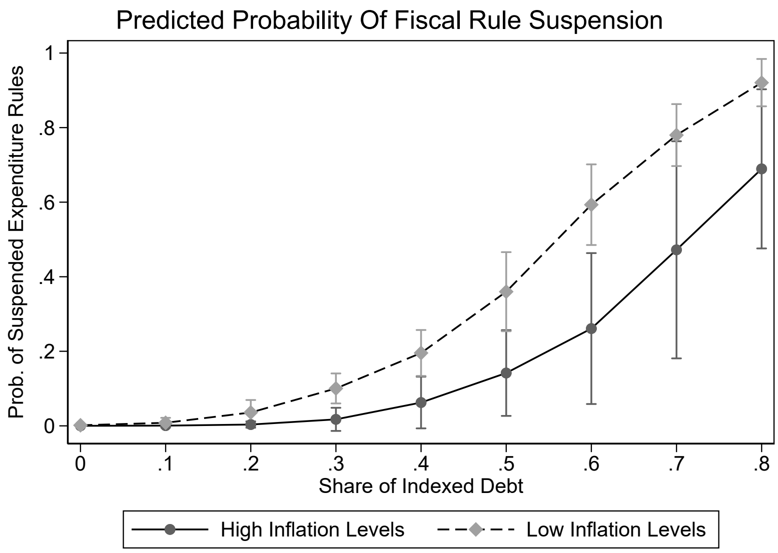

Figure 2.11: The relationship between the share of inflation-indexed debt and the predicted probability of a suspension of fiscal rules in high- and low-inflationary environments, following the estimation in table 2.2.

Even though the non-causal evidence presented here is useful to set the ideas, we now intend to move on to the degree possible and provide a semblance of causal evidence on the role that inflation-indexed debt might play in shaping inflationary dynamics across policy regimes. While identifying policy regimes is generally hardly possible (meaning that evaluating the effect of indexed debt on the propensity of either policy regime in a causal sense is prohibitively difficult), we can attempt to find evidence on the inflationary consequence of inflation-indexed debt in either policy regime, conditional on supplying the prevalence of each regime exogenously.

Causal evidence on the link between fiscally-led policy mixes and materialized inflation in the light of high shares of indexed debt

To round up our empirical exercise, this subsection presents direct evidence on the effect that inflation-indexed debt can have on inflation, making use of the series of narratively identified tax shocks in the US provided by Mierzwa (2024), similarly to chapter 1, but making explicit the *state-dependence* of the prevalent monetary/fiscal policy mix.

We leverage his time series of exogenous fiscal policy surprises, and combine it with a corresponding series of inflation-indexed debt, taking the share of inflation-indexed debt in the overall sovereign debt portfolio as our main indicator for the intensity of the prevalence of inflation-indexed debt. Equipped with these time series, we estimate the following local projection (Jordà, 2005) to measure the dynamic impact of inflation-indexed debt on changes in the rate of inflation:

$$\log P_{t+h} - \log P_{t-1} = \alpha_h + \beta_h \omega_{t-1} \varepsilon_t^F + \delta_{1h} \omega_{t-1} + \delta_{2h} \varepsilon_t^F + \Gamma_h Z_{t-1} + e_{t+h}, \tag{2.40}$$

where $h \geq 0$ indexes the forecast horizon considered and Z_{t-1} is a vector of control variables specified below. Of particular interest to us is the coefficient β_h, which captures the cross-effect of the identified fiscal shock ε_t^F and the past level of the share of inflation-indexed debt ω_{t-1}.[15] We estimate this equation separately for two different periods, creating effectively a *state-dependent local projection*, as described by Jordà and Taylor (2025). We estimate the specification (2.40) separately for periods of *fiscally-led policy mixes* and *monetary-led policy mixes*, as identified by Chen et al. (2022) for the United States.[16] We label the regimes here 'active fiscal policy' and 'passive fiscal policy' in line with Chen et al. (2022).

Figure 2.12 summarizes this exercise, with the left-hand panel showcasing the fiscally-led policy mix and the right-hand panel showing the monetary-led policy mix. Under the monetary-led policy mix, the response of the price level on the interaction between the share of inflation-indexed debt and the identified fiscal policy shock is relatively imprecisely estimated, but - if anything - there is at least a temporary negative effect on the price level when more debt is indexed. In the case of a fiscally-led policy mix, however, medium-term price pressures arising from the identified fiscal policy disturbance are significantly increasing in the share of inflation-indexed debt. In the period 4-5 years after the initial shock, a 1% deficit-to-GDP shock accompanied by a corresponding level difference of 1% in the share of inflation-indexed debt boosts observed changes to the price level by up to 1.5% relative to a case where the share of inflation-indexed debt is lower. This is an economically meaningful effect that would remain hidden without consideration of inflation-indexed debt for inflationary dynamics.

[15] We use last period's level of inflation-indexed debt to avoid running into a simultaneity bias.

[16] While the results in their paper only extend for the period until the Great Financial Crisis, we utilize their characterization to further pin down the prevalent policy mix until 2020, inclusive. Our results are qualitatively robust to characterizing the prevalent fiscal-monetary regime as Bianchi and Melosi (2022) do.

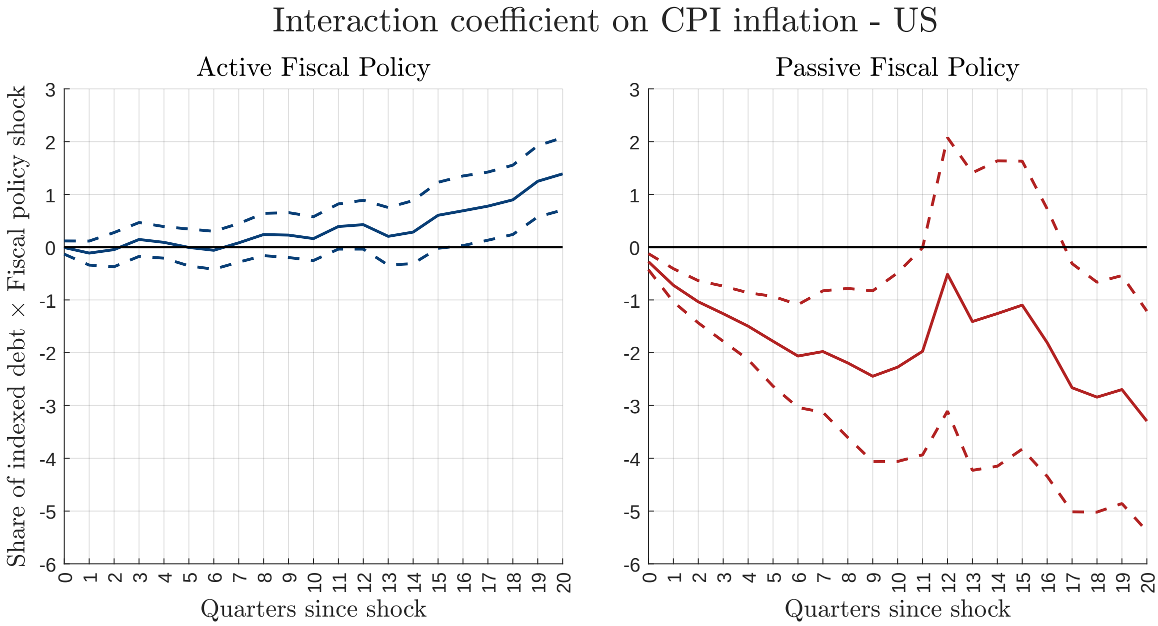

Figure 2.12: IRF implied by the local projection (2.40) through the coefficients β_h. The control vector Z contains the first four lags of the real GDP growth rate, the Fed Funds rate, the change in the weighted real exchange rate, a same-period recession indicator, year-fixed effects, and the first lag of the price level difference. Standard errors are robust to heteroscedasticity and autocorrelation. Confidence intervals at the 90% level. Sample length: 1970 Q1 - 2019 Q2.

2.6 Conclusion

This chapter has looked at fiscal-monetary interactions through the lens of a policy instrument insufficiently considered in dynamic macroeconomic models thus far: inflation-indexed sovereign debt, issued by a multitude of sovereigns.

Inflation-indexed debt highlights two aspects associated with the risks of what is conventionally known as fiscally-led policy mixes. First, recent results emphasizing that the clear distinction between fiscally- and monetary-led policy mixes are overcome once market imperfections such as liquidity risk are added appear to be qualified by the presence of indexed debt: in particular, fiscal policy conventionally considered active is strikingly related to additional inflationary pressure through the presence of inflation-indexed debt, even under such market imperfections. Intuitively, since cost pressures translate into higher debt service burdens for governments when they have issued more indexed debt, monetary and fiscal policies must react sufficiently to cover this additional cost, leading to the emergence of 'quasi fiscally-led' policy mechanisms. Second, it increases the risk of inflationary disasters when we actually are within such fiscally-led policy mixes - that is, when a fiscal deficit shock materializes, it is more inflationary when we actually have higher shares of inflation-indexed debt *and* we are simultaneously operating under policy rules considered to be conventionally fiscally-led.

Chapter 3

The Open-Economy Debt Valuation Equation and International Shock Transmission*

Abstract

This chapter is concerned with the drivers of the government debt valuation equation in the context of international macroeconomics. We first focus on government debt arithmetics and show how foreign interest rates, and therefore all factors influencing these rates directly, can transmit to the valuation of government debt in a small open economy (SOE). Using market value data for U.S. and U.K. public debt covering the period from 1975 to 2024, we show how real exchange rate and interest rate revisions contribute to unexpected changes in the value of government debt. Focusing on fiscal deficit innovations that are derived in a model-consistent way, we show how both domestic and foreign deficits contribute to debt revaluations through nominal adjustment channels. We rationalize these findings in a tractable two-period framework that isolates the international policy transmission channels as well as in a calibrated model of a SOE, in which households in the SOE favour holding rest of world (ROW) bonds through a convenience yield on these bonds.

*We are grateful to Yan Bai, Francesco Bianchi, Martin Ellison, Seung Joo Lee, Zach Mazlish, Andrew Scott, and seminar participants at the Austrian National Bank, the University of Konstanz, and the 2025 London PhD Macro Workshop for comments. Parts of this chapter are based on work with Charles Christopher Hyland, which circulates as "Our Deficit, Your Problem? Fiscal Sustainability, Exchange Rates, and Inflation". For the purposes of our thesis submissions, we apportioned the paper into our own contributions, i.e., the content of this chapter fully constitutes my own contribution with appropriate credit given to my co-author's independent contribution where required.

3.1 Introduction

As the debates on the drivers of inflation intensified in the post-Covid period, bringing about an increased focus on the interactions between monetary and fiscal policymakers, the overall stock of government debt became a focal point of attention again. A significant number of papers went beyond the standard paradigm of considering the face value of public obligations, highlighting the importance of considering the *market value* of government debt (Chen et al., 2022; Reis, 2022; Barro and Bianchi, 2025; Hall and Sargent, 2025).

This distinction between the face value and the market value is of primary relevance to ensure the proper accounting of the effects of shocks on debt sustainability, conditional on the present policy response. Scarcely explored within that perspective focussing on market value-accounting has been the *international* dimension of the valuation of government debt.

This chapter attempts to fill precisely that gap by taking the perspective of an international financial investor determining the pricing, and thereby the valuation, of government debt in a small open economy, in which we tie the cost of borrowing to the interest rate in the world market. By this channel, employing mobility of capital of international investors, we link exchange rates and discount factors to the government debt valuation equation. Households engage in international trade of goods and financial products, and the ROW bond features a convenience yield, consistent with evidence on such yields flowing from U.S. debt (Jiang et al., 2024). While we do not model the fiscal position of the ROW explicitly (in line with the implicit partial-equilibrium perspective taken by SOE models), our headline exercise nonetheless shows the possibility of international contagion of the valuation of government debt: in our calibrated model, we find that a foreign interest tightening, which, ceteris paribus, worsens the foreign fiscal position, can contribute to inflationary pressure in the SOE, consistent with a necessary revaluation of the debt issued by the SOE itself.

Empirically, we start from a present-value identity that embeds the exchange rate, nominal discount rates, and time-varying foreign-exchange risk premia directly in the government budget constraint, building on the present-value logic of Hall and Sargent (2011), Cochrane (2019), and Jiang et al. (2024). Using market-value data for U.K. (and U.S.) sovereign liabilities spanning 1975–2024, we quantify how much of local debt dynamics is driven by primary surplus news, by global real-rate shocks, and by exchange rate or risk-premium repricing. Our results establish that almost 30% of the *financial* factors driving the year-by-year revaluation of the market value of government debt in the United Kingdom, a representative G10 open economy, can be attributed to movements in the real exchange rate. While the domestic fiscal-surplus process remains the largest determinant of the value of government debt, accounting for real-exchange-rate-induced movements helps reconcile episodes that appear contradictory when considering only surpluses and interest rate movements in isolation.

To rationalise the empirical observations, we construct a simple two-period two-country model with two explicit fiscal blocks and the possibility of fiscal contagion in an economy despite no changes to the fiscal stance of that economy. In that simple model, a worsening of the fiscal position in the other economy transmits to a necessary revaluation of government obligations through exchange rate adjustments, affecting the tax revenues that must be raised to ensure the absence of a forced nominal revaluation of outstanding claims to ensure budget balance.

We finally generalize these insights using a standard SOE model in the tradition of Gali and Monacelli (2005). Our two major departures are the characterization of a rich debt maturity structure in the SOE following Cochrane (2001) and using convenience yields on foreign debt following the specification of Hyland (2026). Since we retain effectively a partial-equilibrium perspective by taking the debt supply from the ROW as exogenous and disconnected from fiscal fundamentals, we explore two specific, but model-consistent proxies that influence the attractiveness of ROW sovereign debt: first, a deterioration in convenience yields earned by owning ROW debt; second, an increase in the bond return rate on ROW debt. In analysing the responses to these two impulses, we postulate and compare two different policy regimes, one being akin to a fiscally-led policy mix, and one being akin to a monetary-led policy mix (in the language of Bianchi et al. (2023)).

The decrease in ROW convenience yields principally makes SOE debt more attractive. The Home currency appreciates, however, there is no expansion since SOE households lose out on their ROW debt holdings, from which they substitute away, pushing up the price of SOE debt. Ultimately, the change to convenience yields has similar effects on the SOE under all considered monetary-fiscal policy mixes.

The situation changes once we consider the responses to an increase in the bond return rate of ROW debt. As before, we initially have a depreciation in the SOE currency, alongside a decrease in the value of long-term bonds, which are relatively less attractive. The SOE households, which own a chunk of ROW debt, obtain windfall gains on these ROW debt holdings, leading to a real expansion. In this case, the fiscal-monetary policy mix matters quite significantly. Under a monetary-led policy mix, the corresponding real expansion is *larger*. The reason is that the fiscal authority, whose real debt value shrinks but whose expenditures are tied to goods produced in the SOE, engages in a very short-lived fiscal expansion concurrent to a decreased attractiveness of its debt, followed by a gradual tightening to ensure budget balance. Under a fiscally-led policy mix, the fiscal authority's reaction stays muted, and all observed effects are just the unfiltered consequence of the windfall gain that the SOE households enjoy. These results are consistent with mechanisms under which a worsening of the fiscal position in the ROW can transmit fiscal pressure to other economies by increasing the cost of servicing debt across countries, tantamount to internationally transmitted fiscal inflation even when the fiscal budgeting of the SOE is unchanged.

Our analysis is consistent with mechanisms in which changes in expected, rather than realised, deficits drive the equilibrium valuation of debt across borders and therefore both the exchange rate and inflation, which refutes the apparent contradiction of an unclear relationship between *realized* deficits and exchange rate movements/inflation over time. The accompanying theoretical framework shows how expected deficits in one country affect both domestic and foreign inflation by inducing exchange rate and interest rate movements that jointly determine the real value of government debt in equilibrium. The necessary deviation from a standard complete-markets benchmark that we propose is minimal: the existence of convenience yields on government bonds, which are non-pecuniary benefits associated with holding quasi-safe assets and which have been documented in relation to government debt (Valchev, 2020; Bianchi and Bigio, 2022; Krishnamurthy and Ma, 2025).

Literature Review

We primarily contribute to the literature on sovereign deficits and their ties to inflation and exchange rate dynamics. The chapter is also adjacent to work on fiscally driven price-level determination, which depends on the broader fiscal-monetary policy mix. That literature was inaugurated by Sargent and Wallace (1981), with the subsequent work of Leeper (1991) and Woodford (1995) laying much of the modern groundwork for models that assign at least a share of inflationary dynamics to fiscal aggregate demand management. Literature reviews in this field are provided by Leeper and Leith (2016) and Cochrane (2023), among others. Support for fiscally driven inflation has been developed both theoretically (Ascari et al., 2023; Bassetto and Cui, 2018; Bassetto et al., 2024; Bianchi et al., 2023; Campos et al., 2024; Caramp and Silva, 2023; Bigio et al., 2024; Corsetti and Maćkowiak, 2024; Kaplan et al., 2024; Miao and Su, 2024) and empirically (Ascari et al., 2024; Banerjee et al., 2022; Barro and Bianchi, 2025; Chen et al., 2022; Cochrane, 2022a; Reichlin et al., 2023).

While most of this work is set in closed-economy environments, a small but growing literature brings these insights to international models. Early explorations linking fiscal price-level determination and exchange rate dynamics include Bergin (2000), Dupor (2000), Sims (1997), Sims (1999), and Woodford (1998). Initial applications of such ideas include Benigno and Missale (2004) and Corsetti and Maćkowiak (2006), who focus on the ties between public debt and currency stability, especially during episodes of severe currency turmoil. International aspects of government debt valuation and associated effects on inflation and exchange rates were further explored in the run-up to the Global Financial Crisis by Giannitsarou et al. (2006) and Gourinchas and Rey (2007).

Alongside the renewed interest in deficit-inflation links, research has revisited the role of deficits in exchange rate determination. Schmitt-Grohé and Uribe (2003) provided an early analysis of the role of exchange rate adjustments relative to inflation in financing fiscal deficits within emerging markets. Valchev (2020) establishes that excess returns on currency holdings can be traced to endogenous movements in bond convenience yields, which in turn depend on the prevailing interactions between monetary and fiscal policy. Particularly related to our contribution are Alberola et al. (2021) and Jiang (2022). Alberola et al. (2021) distinguish between Ricardian and non-Ricardian regimes, tying unexpected exchange rate movements (relative to a complete-market frictionless benchmark) to non-Ricardian fiscal regimes in which government debt lacks sufficient fiscal backing. Jiang (2022) makes explicit some of the implications of applying fiscal price-level determination to exchange rate dynamics, pointing out that real exchange rate adjustments are linked to government surplus shocks through the cost of issuing government debt, among other factors.

The chapter also relates to the large macro-finance literature on nominal and real exchange rate determination, which we cannot survey fully here. A partial review is provided by Itskhoki (2021), including an overview of the many puzzles related to exchange rate dynamics that standard complete-market models fail to resolve.[1]

The chapter is organised as follows: section 3.2 lays out the present value accounting framework and derives its open economy linearisation, from which we derive a market-implied measure of surplus/deficit innovations. We then use that measure to link deficits, exchange rates, and inflation empirically. A toy model rationalizing these ideas is presented in section 3.3 followed by our main small open economy in section 3.4. Section 3.5 concludes with policy implications for debt management and for the coordination of fiscal and exchange rate policies. Further details on the data, additional model derivations, and all proofs can be found in the appendix.

3.2 Empirical Framework

The correct measurement of surplus and deficit innovations is non-trivial due to the fundamental endogeneity of fiscal policy. No single method is considered a clear-cut standard for measuring fiscal innovations (Ramey, 2016, 2019). The consideration of

[1] These include the possibility of a disconnect from macroeconomic fundamentals (Meese and Rogoff, 1983), the forward premium puzzle (Fama, 1984), the slightly negative correlation of real exchange rates and relative consumption (Backus and Smith, 1993; Corsetti et al., 2008), the Purchasing Power Parity puzzle (Rogoff, 1996), and the lower volatility of the terms of trade relative to the real exchange rate (Atkeson and Burstein, 2008), among others.

international aspects of cross-border bond trade likely amplifies existing concerns while making these questions even more important for understanding exchange rate determination and the international transmission of inflation.

Our empirical framework places the pricing of government debt on financial markets at its centre. Informed by the market-value identities of Cochrane (2019), we develop an empirical open-economy framework in which we price U.K. and U.S. sovereign debt in line with a government debt valuation equation, taking into account exchange rate adjustments. We embed exchange rates and the UIP condition in the pricing of government debt to show that foreign discount-rate news and currency movements enter fiscal constraints on equal footing with domestic surpluses and inflation. We then use the resulting identity, which holds as an accounting relation in every model, to decompose debt dynamics into domestic channels (surpluses, deficits, growth) and international channels (UIP, foreign surpluses and deficits). While our use of the *market value* of government debt (as opposed to its *face value*) is not novel (see Hall and Sargent (2011); Cochrane (2019)), it is a necessary ingredient for recovering the only feasible surplus surprises consistent with the value of government debt itself.

Data. We collect U.K. public debt market values from Ellison and Scott (2020) and update them through 2024 with the Gilt database of Cairns and Wilkie (2026). U.S. debt market values are taken from Cochrane (2019), again suitably updated. Macroeconomic aggregates for both countries are from Müller et al. (2025). The real exchange rate is recovered from the BIS Real Effective Exchange Rate index. We use market values throughout rather than face (book) values. Book-value measures omit capital gains within the portfolio, so the period-by-period budget identity requires ad-hoc valuation adjustments and any present-value decomposition is biased (Hall and Sargent, 2011). Market values embed time-varying real yields and risk premia by construction, ensuring that the decompositions below close exactly without a residual discrepancy, which allows us to recover the model-consistent surplus/deficit surprise.

The Open-Economy Debt Valuation Identity

We begin by describing the procedure that allows us to recover realised surplus and deficit innovations. We first present the baseline result of Cochrane (2019) before extending it to our open-economy setting.

Single-country baseline. We define the nominal end-of-period market value of government liabilities as

$$V_t \equiv M_t + \sum_{j=0}^{\infty} Q_t^{(t+1+j)} B_t^{(t+1+j)}, \tag{3.1}$$

where M_t is non-interest-bearing money, $B_t^{(t+j)}$ is the face value of a zero-coupon bond due at $t+j$, and $Q_t^{(t+j)}$ is its time-t price. The log debt-to-GDP ratio is $v_t \equiv \log(V_t/P_t Y_t)$ and the log nominal holding-period return on the portfolio is $r_{t+1}^n \equiv \log R_{t+1}^n$, where

$$R_{t+1}^n \equiv \frac{M_t + \sum_{j=1}^{\infty} Q_{t+1}^{(t+j)} B_t^{(t+j)}}{M_t + \sum_{j=1}^{\infty} Q_t^{(t+j)} B_t^{(t+j)}}.$$

Log-linearising the nonlinear flow identity following Campbell and Shiller (1988) and Cochrane (2019) yields the following result.

Lemma 1 (Linearised Flow Identity; Cochrane, 2019, 2022b) *The linearised government debt valuation equation is given by*

$$\rho\, v_{t+1} = v_t + r_{t+1}^n - \pi_{t+1} - g_{t+1} - s_{t+1}, \tag{3.2}$$

where π_{t+1} is CPI inflation, g_{t+1} is log real GDP growth, s_{t+1} is the real primary surplus scaled by GDP, and ρ is the discount factor.

Equation (3.2) is a pure accounting identity that must hold exactly period by period in market values. It states that the debt ratio next period is determined by today's debt ratio, augmented by the nominal return on the government bond portfolio, and eroded by inflation, growth, and the primary surplus. Crucially, the identity is silent on causation: in a Ricardian world, surpluses adjust to satisfy (3.2) for any given path of prices and interest rates; in a fiscal-theory world, prices and interest rates adjust to satisfy it for a given path of surpluses. Our empirical strategy remains agnostic between these interpretations, since the identity holds under both.

Open-economy extension. Consider a Home country, labelled H, and a Foreign country, labelled F. Applying (3.2) to each country separately yields

$$\rho_H v_{t+1}^H = v_t^H + r_{t+1}^{n,H} - \pi_{t+1}^H - g_{t+1}^H - s_{t+1}^H, \tag{3.3}$$

$$\rho_F v_{t+1}^F = v_t^F + r_{t+1}^{n,F} - \pi_{t+1}^F - g_{t+1}^F - s_{t+1}^F. \tag{3.4}$$

Let $e_t \equiv \log \mathcal{E}_t$ denote the log nominal exchange rate. For concreteness, we let the Home country be the U.S. and the Foreign country be the U.K., so that an increase in e_t represents an appreciation from the U.K. perspective. The UIP condition with a mean-zero risk-premium shock is

$$r_{t+1}^{n,H} - r_{t+1}^{n,F} = \Delta e_{t+1} + \varphi_{t+1}, \qquad \mathbb{E}_t[\varphi_{t+1}] = 0. \tag{3.5}$$

The key step in our open-economy extension is to substitute (3.5) into (3.4), replacing the Foreign nominal return with its UIP-implied counterpart. This yields the central empirical equation by which we recover surpluses:

$$\rho_F v_{t+1}^F = v_t^F + \underbrace{\left(r_{t+1}^{n,H} - \Delta e_{t+1} - \varphi_{t+1}\right)}_{\text{external discount factor}} - \pi_{t+1}^F - g_{t+1}^F - s_{t+1}^F. \tag{3.6}$$

The foreign discount rate $r_{t+1}^{n,H}$, the exchange rate appreciation Δe_{t+1}, and the UIP premium φ_{t+1} enter the U.K. budget constraint on exactly the same footing as domestic inflation and growth. In effect, we treat U.S. discount rates as an anchor for international financial flows, with U.K. discount rates following as a direct function thereof after accounting for exchange rate-related revaluations.

The economic content of equation (3.6) is best understood through a concrete example. Suppose that U.S. Treasury yields rise by 100 basis points because markets revise down expected U.S. surpluses, raising $r_{t+1}^{n,H}$ in (3.4). Under the UIP link, this increase in the Foreign discount rate feeds one-for-one into the U.K. budget constraint through the external discount factor in (3.6). The U.K. debt ratio must therefore be higher next period than it would have been absent the U.S. shock, unless at least one of the remaining terms adjusts. This observation is the core of the international transmission mechanism: the U.K. intertemporal budget constraint tightens in response to a fiscal deterioration abroad.

The adjustment can come through at least one of four channels. *Fiscal inflation* erodes the real value of outstanding Sterling debt, effectively taxing existing bondholders. Stronger *growth* raises the tax base and partially self-finances the higher debt burden (Angeletos et al., 2024a). A *Dollar depreciation* occurs as investors reallocate away from Dollar assets, mechanically increasing the Sterling value of the external discount factor. Finally, *UIP compression* arises if the convenience yield on U.S. treasuries falls, narrowing the wedge between Home and Foreign returns. Any combination of these four adjustments that satisfies (3.6) is consistent with U.K. fiscal solvency without requiring any change in U.K. surpluses.

Two features of this result merit emphasis. The first is that the transmission is symmetric in magnitude but asymmetric in direction: a U.S. deficit shock tightens the U.K. constraint, but the reverse need not hold with the same force, because the UIP substitution replaces the Foreign return with the Home return, not vice versa. This asymmetry reflects a modelling choice that assigns the anchor role to the larger economy's discount rate, consistent with the empirical dominance of U.S. Treasury yields in global fixed-income markets. The second is that the identity (3.6) places no restriction on which channel does the adjusting. In the data, all four channels operate simultaneously, and the backward and forward decompositions in Section 3.2 are designed to measure the relative contribution of each. This is the empirical counterpart of the fiscal contagion mechanism documented in our theoretical model.

Backward and Forward Decompositions

Before examining the direct effects of surpluses on exchange rates and inflation, we first analyse the properties of the market-implied surpluses we recover. Following Cochrane (2019), we decompose the drivers of the value of government debt using both backward and forward perspectives.

Backward Decomposition

Iterating (3.6) forward from a base year (setting $\rho_F \approx 1$ throughout)[2] yields the backward accounting identity:

$$v_t^F = v_0^F + \sum_{j=1}^{t}\Big[r_j^{n,H} - \Delta e_j - \varphi_j\Big] - \sum_{j=1}^{t}\pi_j^F - \sum_{j=1}^{t}g_j^F - \sum_{j=1}^{t}s_j^F. \tag{3.7}$$

This is a pure arithmetic attribution: it measures, rather than tests, the present-value relation (Cochrane, 2019). Each term cumulates one channel's contribution, and their sum equals the realised debt path at every date by construction.

Figure 3.1 applies (3.7) to U.K. data from 1975, with each series starting at zero in that year. Table 3.1 defines each series formally. Three results stand out. First, primary surpluses account for approximately half of the cumulative change in the U.K. debt ratio since 1975; the remainder is accounted for by valuation effects, split broadly between the global discount factor ($r^{n,H} - \pi^F - g^F$) and the exchange rate channels (REER and UIP). Second, every major regime break identified by Bordo et al. (2022) appears as a kink in a valuation curve rather than in the surplus path, confirming that discount-rate repricing is closely related to observed debt dynamics. Third, the open-economy channels are quantitatively material: a permanent 10 % Dollar depreciation raises the domestic-currency value of Sterling coupons by the same 10 %, while a 200 bp UIP shock adds roughly 10% of GDP to the market value of outstanding debt without any change in cash surpluses.

[2]To be precise, we set $\rho_F \simeq 0.97$ at annual frequency. Suppressing it keeps notation light and has negligible effect on the decomposition over the sample.

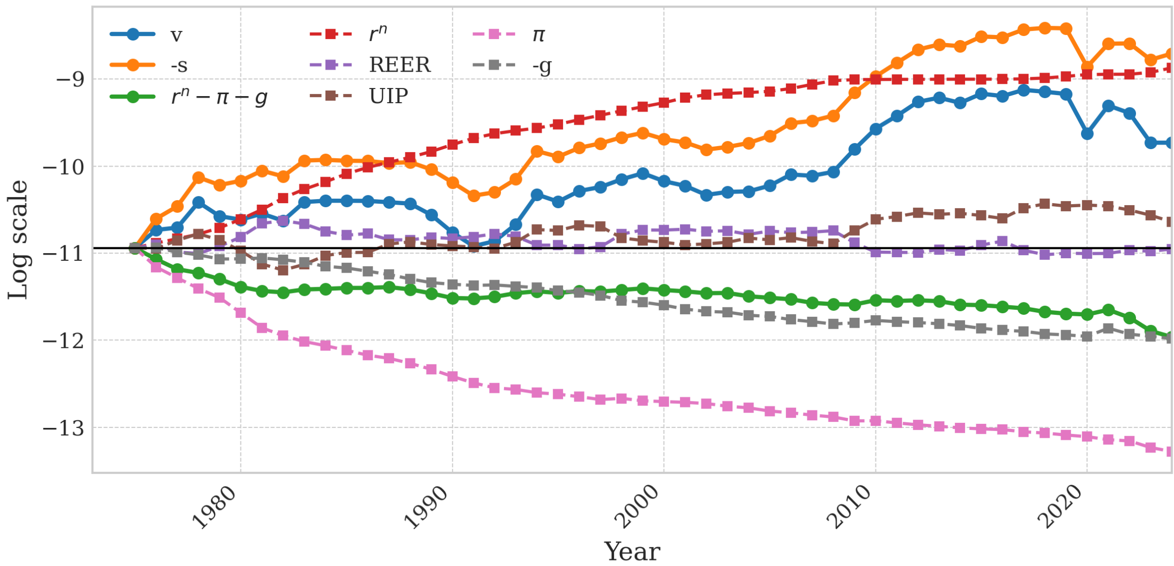

Figure 3.1: Backward decomposition of U.K government debt, 1975–2024. Each path cumulates one term of identity (3.7) from the 1975 base year. The vertical sum of all series equals the realised market-value debt-to-GDP ratio (blue line, v) at every date.

Series	Formula (base 1975)	Interpretation
v	$v_0^F + \sum_{j\leq t}(r_j^{n,H} - \Delta e_j - \varphi_j - \pi_j^F - g_j^F - s_j^F)$	Realised market-value debt/GDP ratio.
$-s$	$v_0^F - \sum_{j\leq t} s_j^F$	Counterfactual driven only by realised surpluses.
$r^n - \pi - g$	$v_0^F + \sum_{j\leq t}(r_j^{n,H} - \pi_j^F - g_j^F)$	Cumulated real discount-rate contribution.
r^n	$v_0^F + \sum_{j\leq t} r_j^{n,H}$	Cumulated U.S. nominal return.
$-\pi$	$v_0^F - \sum_{j\leq t} \pi_j^F$	Inflation erosion of the real debt value.
$-g$	$v_0^F - \sum_{j\leq t} g_j^F$	Growth contribution.
REER	$v_0^F - \sum_{j\leq t} \Delta e_j$	Cumulated Sterling appreciation.
UIP	$v_0^F - \sum_{j\leq t} \varphi_j$	Cumulated FX risk-premium shocks.

Table 3.1: Components of the backward decomposition, equation (3.7), and their mapping to Figure 3.1.

Forward Decomposition

The forward decomposition replaces realised values with conditional expectations: the present-value contribution of each channel at date t is the amount that channel alone would need to provide to balance the intertemporal budget constraint given information at t. Taking expectations of the iterated identity yields the forward balance condition:[3]

$$v_t^F = PV_t^{(s)} + PV_t^{(g)} - PV_t^{(d)}, \qquad PV_t^{(d)} \equiv PV_t^{(r)} - PV_t^{(\pi)}, \tag{3.8}$$

where each $PV_t^{(k)}$ is the expected discounted sum of component k from date t forward.

VAR implementation. We bring (3.8) to the data via a first-order annual VAR following Cochrane (2019, 2022b):

$$x_{t+1} = Ax_t + \varepsilon_{t+1}, \qquad \varepsilon_{t+1} \sim \mathcal{N}(0, \Sigma), \tag{3.9}$$

where the vector x_t contains the U.S. nominal return $r_t^{n,H}$, U.K. surplus/deficit innovations s_t^F, inflation π_t^F, growth g_t^F, and the log debt ratio v_t^F, as well as the real exchange rate change Δe_t, the ex-ante UIP premium φ_t, and short and long U.K. Treasury yields as predictors of future returns. All eigenvalues of A lie strictly inside the unit circle, ensuring finite long-run multipliers. With row-selection vectors a_k picking out component k, the present-value contributions are

$$PV_t^{(k)} \equiv a_k'(I - A)^{-1}Ax_t, \qquad k \in \{s,\, g,\, r,\, \pi,\, \Delta e,\, \varphi\}, \tag{3.10}$$

yielding an infinite-horizon present value in a single matrix multiplication in which all projections are internally consistent with the accounting identity.

Figure 3.2 and Table 3.2 present the results of the forward decomposition exercise. Deficit/surplus news and real discount-rate news are the dominant contributors to debt-ratio fluctuations across the full sample. The exchange rate and UIP channels are nevertheless economically significant: they reconcile the relative undervaluation of U.K. debt implied by the surplus process alone in the pre-2000 period and the relative overvaluation thereafter.[4]

[3] The derivation of the relevant steps is provided in Appendix C.2.

[4] A historical variance decomposition assigning shares of total variance to each channel is provided in Appendix C.2.

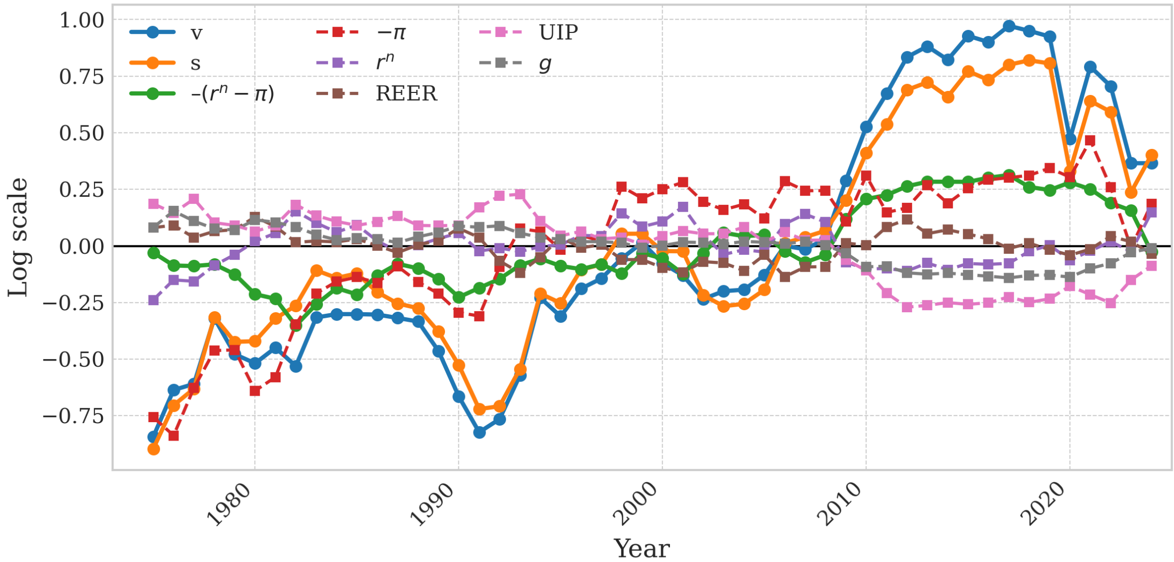

Figure 3.2: Forward decomposition of the value of U.K. government debt. The blue line is the realised debt ratio v_t^F. Every other series is the corresponding present-value component $PV_t^{(k)}$ from (3.10), formed with information available at date t.

Series	Object	Economic meaning
v	v_t^F	Market-value debt/GDP ratio.
s	$PV_t^{(s)}$	PV of expected future primary surpluses.
$-(r^n - \pi)$	$-PV_t^{(d)}$	PV of expected future real discount rates.
r^n	$PV_t^{(r)}$	PV of expected future nominal returns.
$-\pi$	$-PV_t^{(\pi)}$	PV of expected future U.K. inflation.
REER	$PV_t^{(\Delta e)}$	PV of expected future exchange rate changes.
UIP	$PV_t^{(\varphi)}$	PV of expected future UIP-premium shocks.
g	$PV_t^{(g)}$	PV of expected future real growth.

Table 3.2: Components of the forward decomposition in Figure 3.2. By construction, $v_t^F = PV_t^{(s)} + PV_t^{(g)} - PV_t^{(d)}$.

Exploiting Deficit Shocks to Analyse Exchange Rate Movements

In our derivation of market value-consistent surpluses, we have already shown how UIP deviations and cross-border interest rate spillovers matter significantly for the equilibrium valuation of government debt. We now leverage the recovered surplus and *deficit* processes for the U.K. and the U.S.[5] to determine how much each contributes to medium-run movements in U.K. real exchange rates. We view this exercise as informative about the plausible secular determinants of U.K. exchange rates, quantifying how much of the

[5] We do not detail the recovery of the U.S. surplus/deficit process in detail here. The process of recovering this series mirrors our exposition of the U.K. case above.

observed variance can be attributed to movements in domestic versus foreign deficits, where foreign deficits matter through foreign interest rates and the associated exchange rate and UIP deviation channels described above.

Table 3.3 summarises the results of this exercise, linking the deficit innovations implied by the forward decomposition of the VAR in Section 3.2 to observed year-on-year movements in U.K. *real* exchange rates, that is, after accounting for differences in inflation.[6]

The results are striking. Deficit expectation revisions relate to observed movements in U.K. real exchange rates significantly and in the direction implied by the government debt valuation equation: a deficit innovation in the U.K. is associated with a *depreciation* of the U.K. real exchange rate, with a 1% deficit-to-GDP innovation corresponding to a 37 basis point relative *depreciation*. U.S. deficit dynamics matter an order of magnitude less, as expected and as prescribed by the framework. Quantitatively, a 1% deficit-to-GDP innovation in the U.S. is associated with a 1.8 basis point relative *appreciation* of U.K. real exchange rates. The second column confirms that the *differential* between U.K. and U.S. surpluses relates most materially to annual exchange rate movements. The third column introduces an interaction with domestic recessions, revealing that in such periods deficit innovations may plausibly matter through unconventional channels not otherwise captured.[7]

The observed links between deficits and the exchange rate are fairly robust over time, as we demonstrate using the following local projection on our exogenous deficit measure:

$$\Delta e_{t+h} = \alpha_h + \beta_h^{UK}\,\mathrm{Def}_t^{UK} + \beta_h^{US}\,\mathrm{Def}_t^{US} + \gamma_h\, e_{t-1} + u_{h,t}, \qquad h = 0, 1, \ldots, H,$$

where Δe_{t+h} is the change in the log real exchange rate index and the right-hand side variables are the U.K. and U.S. deficit measures, respectively. Figure 3.3 plots the local projection coefficients, showing that the exchange rate response to expected deficit innovations is relatively persistent across horizons.

Does the same relationship hold true for materialized inflation rates, in that both domestic and foreign deficits can influence inflation? We re-estimate the previous local projection, but on observed CPI inflation rates:

$$\Delta P_{t+h}^{CPIUK} = \alpha_h + \beta_h^{UK}\,\mathrm{Def}_t^{UK} + \beta_h^{US}\,\mathrm{Def}_t^{US} + \gamma_h\, P_{t-1}^{CPIUK} + u_{h,t}, \qquad h = 0, 1, \ldots, H,$$

[6] We weight the U.S. deficit innovations by relative U.S./U.K. GDP to reflect the relative size of both countries. Not weighting the deficit does not affect the qualitative results.

[7] In Appendix C.3 we provide additional results based on the Engel and West (2005) reverse-regression approach, showing that exchange rates predict materialised deficits in the U.K. The link between exchange rates and deficits follows an "overshooting" pattern.

	Dependent variable: Δ *U.K. REER*		
	(1)	(2)	(3)
$\Delta\mathbb{E}(\text{Deficit})_t^{U.K.}$	0.370***	0.352***	0.358***
	(0.077)	(0.075)	(0.107)
$\Delta\mathbb{E}(\text{Deficit})_t^{U.K.} \times \mathbf{1}^{g_Y<0}$			−2.370***
			(0.615)
$\Delta\mathbb{E}(\text{Deficit})_t^{U.S.}$	−0.018***		−0.010**
	(0.005)		(0.004)
$\Delta\mathbb{E}(\text{Deficit})_t^{U.S.} \times \mathbf{1}^{g_Y<0}$			0.210***
			(0.061)
$\Delta\mathbb{E}(\text{Deficit})_t^{U.K.-U.S.}$		0.016***	
		(0.006)	
Constant	-0.006	-0.022	-0.001
	(0.010)	(0.014)	(0.009)
Observations	25	25	25
Adjusted R^2	0.469	0.320	0.523

Note: HAC-robust standard errors. *p<0.1; **p<0.05; ***p<0.01

Table 3.3: Long-run correlates of U.K. exchange rates.

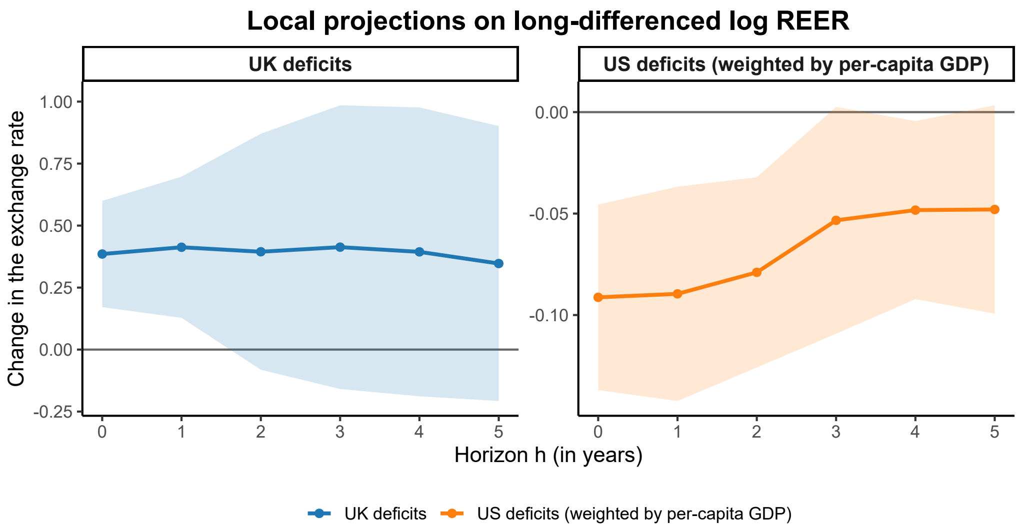

Figure 3.3: The effect of U.K.+U.S. deficits on the U.K. REER.

where ΔP_{t+h}^{CPIUK} is the change in the log consumer price index (i.e., cumulative inflation) and the right-hand side variables are the U.K. and U.S. deficit measures, respectively. Figure 3.4 plots the local projection coefficients. While U.K. deficits matter quantitatively more (even though the effect is econometrically not significant), we can also clearly observe that the U.S. deficit surprises matter both quantitatively and economically for movements of the consumer price index in the U.K.

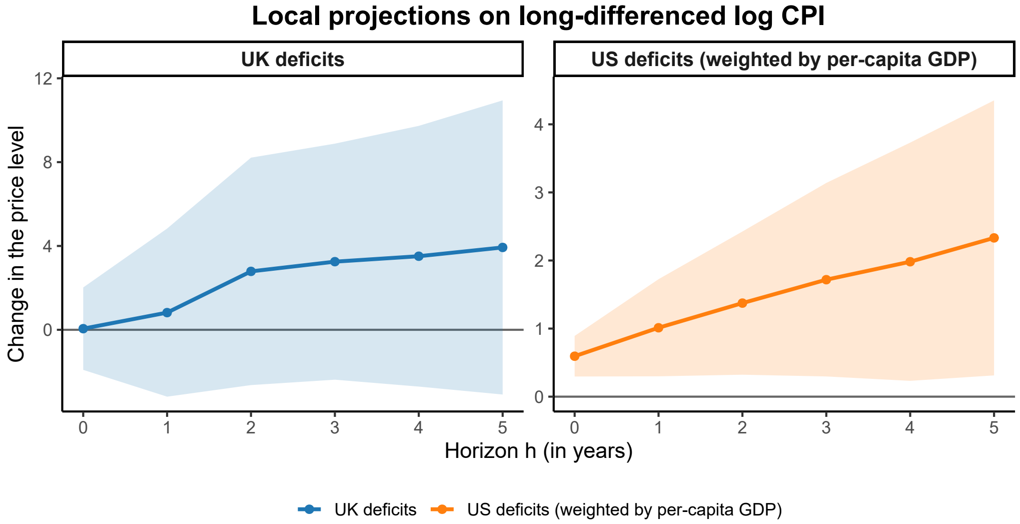

Figure 3.4: The effect of U.K.+U.S. deficits on the price level in the U.K.

In sum, our long-run historical analysis, centred on the dynamics underpinning the valuation of government debt, underscores that the domestic debt valuation equation is necessarily linked to the UIP condition and to fluctuations in the real exchange rate. We have further shown how measured surplus/deficit innovations contribute to observed movements in U.K. real exchange and inflation rates, which holds true not only for U.K. surplus/deficit innovations, but also for those of the U.S.

3.3 A Simple Model of International Fiscal Spillovers

We now present a simple model emphasizing a fiscally-led inflation mechanism in a standard international macro context, with plausible relevance of fiscal deficits for nominal exchange rate movements and international spillovers through interest rate adjustments. What we show here is a minimal mechanism of how fiscal sustainability concerns in one country can spill over to the government debt valuation of another country. The interesting messages of the model are kept fully nominal, with no relevance for international goods flows, which we nullify here for expositional simplicity.

Environment

The primitives of the model consist of a simplified iteration of the framework of Bassetto and Miller (2025), but with a diminished focus on the information structure. Instead, we add an international dimension facilitating the transmission of fiscal shocks between countries. Denote the rest of world country by ROW, the small open economy country by SOE, and let stars (*) denote variables in the SOE country. Heuristically, we will think of the ROW country as an international hegemon determining international spillovers, and of the SOE country as a small open economy subject to the financial influence of the hegemon. There are two periods:

1. t=1: there exists an initial steady-state inherited from t=0. In period one, the price levels P_1, P_1^*, as well as the nominal exchange rate $\mathcal{E}_1$, are all determined through the respective government budget constraints and cross-border debt holdings inherited from the initial period. There is a shock to ROW government spending in period 2, G_2, whose plausible size is revealed. Additionally, there is a prior about the policy regime in place at ROW. Two policy regimes are possible: a 'fiscally-led' regime, in which taxes are a fixed quantity and the spending need will be financed through inflation; and a 'monetary-led' regime in which taxes adjust to ensure that a target price level at ROW is achieved. The SOE fiscal authority levies fixed taxes $\bar{T}^*$ in all periods, while SOE interest rates on debt issuance i_1^* are determined in equilibrium through the UIP condition.

2. t=2: The tax and monetary regimes from ROW unfold, yielding the terminal ROW price level P_2. Note that the SOE price level P_2^* will already be perfectly revealed in period 1, being a function of equilibrium interest rates on SOE debt from period 1 to period 2.

ROW households Relative to standard macroeconomic models, we introduce two reduced-form preference terms capturing frictions in cross-border bond holdings. On the one hand, SOE households have a positive convenience yield from holding ROW bonds. On the other hand, ROW households face a cost (or inconvenience) when holding SOE bonds, reflecting information asymmetries, institutional frictions, or a generic home bias.

ROW households maximize the utility borne from consumption of the domestic good, adjusted for the disutility from supplying labour and from holding SOE bonds:[8]

$$max\ U = max\ \mathbb{E}_t \sum_{t=1}^{2} \beta^t \left[u(c_t) - L_t - \phi_{ROW} \frac{B_{SOE,t}\mathcal{E}_t}{P_t} \right],$$

[8]Note that there is no goods trade in equilibrium, but only trade in financial products at this stage.

where c_t is consumption of the representative household, L_t is labour supplied, and the production function of the household is given by $y_t = L_t$. Disutility from labour is kept linear on purpose to simplify determination of the real equilibrium in the presence of fiscal shocks. The term $-\phi_{ROW}\frac{B_{SOE,t}\mathcal{E}_t}{P_t}$ captures the disutility from holding SOE bonds, where $\phi_{ROW} > 0$ is an exogenous parameter and $\frac{B_{SOE,t}\mathcal{E}_t}{P_t}$ is the real value of SOE bonds held by the ROW household, converted to ROW currency units. This term ensures that the ROW household's Euler equation for SOE bonds carries a wedge, which is necessary for the equilibrium interest rate on SOE bonds to be uniquely determined.

Each household takes as given government policy and the price level. Define $B_{ROW,t-1}$ as household holdings of ROW bonds at the beginning of period t, having taken them over from period $t-1$. The same applies to holdings of SOE bonds by the ROW household, defined as $B_{SOE,t-1}$. Then, the domestic household budget constraint is given by:

$$B_{ROW,t-1}(1+\bar{i}) + B_{SOE,t-1}\mathcal{E}_t(1+i^*_{t-1}) + P_tL_t = B_{ROW,t} + B_{SOE,t}\mathcal{E}_t + P_t(c_t+T_t). \quad (3.11)$$

Note that no debt is left outstanding in the final period for either household in equilibrium; that is, $B_{ROW,2} = B_{SOE,2} = B^*_{ROW,2} = B^*_{SOE,2} = 0$.

ROW Fiscal and Monetary Authority Next, consider the ROW government. The government faces an inherited stock of debt B_0, held by both ROW and SOE households in some proportion. It can raise lump-sum taxes T_t in each period. We follow Bassetto and Miller (2025) in assuming that $G_1 = 0$ without loss of generality. In period 1, lump-sum taxes will be set to cover interest payments, such that $T_t = \frac{B_{t-1}(1+\bar{i})-B_t}{P_t}$ for $t = 1$. In period $t = 2$, ROW government spending and taxes are uncertain and plausibly disjoint, allowing for altered government debt dynamics.

In line with Bassetto and Miller (2025), two monetary-fiscal policy regimes are possible:

- ML (Monetary-led): in this regime, $T_2 = G_2 + \frac{B_1}{P_2^T}$, where P_2^T is an externally supplied targeted price level in period 2. Then, any spending requirement G_2 is fully soaked up by corresponding taxes, while the primary surplus $\frac{B_1}{P_2^T}$ soaks up the cost of outstanding bond holdings.

- FL (Fiscally-led): in regime FL, taxes do not respond to the spending requirement: $T_2 = \bar{T}$, where $\bar{T}$ is an exogenously supplied known constant.

Fundamentally, the assumption of introducing tax changes in period 2, but their announcement taking place in period 1, is meant to distinguish the effects of announced deficits from the effects once the deficits actually materialize. It therefore allows us to capture the effects of fiscal spending announcements on exchange rates through their forward-looking nature.

Monetary policy in the ROW economy will be captured by a strict nominal interest rate peg as already in use in the ROW budget constraint, $i_t = \bar{i}\,\forall t$.

SOE block The *SOE* household is described by a similar problem as the domestic household, but with all variables being starred and the exchange rate adjustment in the budget constraint reflecting that the SOE household considers the SOE bonds as domestic, and, crucially, an additional idiosyncratic preference term capturing the utility from holding ROW bonds. Formally, the representative SOE household maximizes:

$$max\ U^* = max \sum_{t=1}^{2} \beta^t \mathbb{E}_t \left[u(c_t^*) - L_t^* + \phi_{SOE} \frac{B^*_{ROW,t}}{P_t^*} \right].$$

The additional term at the end, $+\phi_{SOE}\frac{B^*_{ROW,t}}{P_t^*}$, captures the utility from holding ROW bonds.[9] ϕ_{SOE} is an exogenous parameter capturing the weight of the preference term, and $\frac{B^*_{ROW,t}}{P_t^*}$ is the real quantity of ROW bonds held by SOE households. Together with the disutility parameter ϕ_{ROW} on the ROW side, the pair (ϕ_{SOE}, ϕ_{ROW}) governs the asymmetric frictions in cross-border bond markets: SOE households enjoy a convenience yield from ROW bonds, while ROW households face a cost of holding SOE bonds.

This maximization is performed subject to the per-period budget constraint of the SOE household:

$$B^*_{SOE,t-1}(1 + i^*_{t-1}) + \frac{B^*_{ROW,t-1}}{\mathcal{E}_t}(1 + \bar{i}) + P_t^* L_t^* = B^*_{SOE,t} + \frac{B^*_{ROW,t}}{\mathcal{E}_t} + P_t^*(c_t^* + T_t^*). \quad (3.12)$$

Now, we consider the government of SOE. The SOE government differs from the ROW government in a number of important aspects:

- There is no uncertainty with respect to SOE government spending: $G_t^* = 0\ \forall t$.

[9] Note that we do not normalize this term by the inverse of the exchange rate $\mathcal{E}_t$ to translate this convenience flow into units of the local currency. The economic idea here is that the ROW bonds are also valued because of their intrinsic value in what might be a relatively *safer* currency; the mechanical reason is that such an adjustment would cause this simple model to be indeterminate, since the model's UIP conditions would then collapse to a parametric restriction under the convenience yield parameters, and not a clear relationship pinning down $1 + i_1^*$.

- In terms of the policy in use, we postulate that taxes levied in the terminal period by the SOE government are fixed; that is, $T_2^* = \bar{T}^*$. The terminal period price level in the SOE economy will be determined by the terminal government budget constraint, with the price level adjusting to ensure that the fiscal budget constraint holds.

- SOE interest rates are allowed to be time-varying: i_t^* is not constant. It adjusts, however, not based on a policy rule, but instead in a way that allows the UIP to hold while taking into account the convenience wedges.

Information structure: We introduce an explicit information structure that allows us to simplify the expectations terms. In $t = 1$, both ROW and SOE households have a (shared) prior about the policy regime and spending G_2 of ROW. To keep it simple, we denote by ψ_1 the prior that the ROW government will be in regime FL, and by $1 - \psi_1$ the prior that the ROW government will be in regime ML. These priors will be the same in both ROW and SOE. The government spending variable can follow an arbitrary distribution $\mathcal{G}$.

Optimality Conditions: We provide detailled derivations in appendix C.1. The ROW household's equilibrium conditions are characterised by:

$$u'(c_t) = 1, \qquad t = 1, 2; \tag{3.13}$$

$$1 = \beta(1 + \bar{i})\, \mathbb{E}_t\left[\frac{P_t}{P_{t+1}}\right], \qquad t = 1; \tag{3.14}$$

$$(1 + \phi_{ROW}) = \beta(1 + i_t^*)\, \mathbb{E}_t\left[\frac{P_t}{P_{t+1}}\frac{\mathcal{E}_{t+1}}{\mathcal{E}_t}\right], \qquad t = 1. \tag{3.15}$$

Note that the Euler equation on SOE bonds (3.15) carries the wedge $(1+\phi_{ROW})$ on the left-hand side, reflecting the additional cost borne by the ROW household for holding foreign bonds. For the SOE household, their equilibrium optimality conditions are characterised by:

$$u'(c_t^*) = 1, \qquad t = 1, 2; \tag{3.16}$$

$$1 = \beta(1 + \bar{i})\, \mathbb{E}_t\left[\frac{P_t^*}{P_{t+1}^*}\frac{\mathcal{E}_t}{\mathcal{E}_{t+1}}\right] + \phi_{SOE}\mathcal{E}_t, \qquad t = 1; \tag{3.17}$$

$$1 = \beta(1 + i_t^*)\, \mathbb{E}_t\left[\frac{P_t^*}{P_{t+1}^*}\right], \qquad t = 1. \tag{3.18}$$

Note that, in general, no Euler equations apply in period 2 as all bond holdings in that period are set to zero in that period by household optimality.

Solving for equilibrium

Appendix C.1 contains the definition of the competitive equilibrium we solve for, as well as details on how the effect of the behaviour of the fiscal and the monetary authority in the ROW on the SOE can succinctly be summarized by the SOE interest rate i_1^*. Our discussion of the effects of fiscal spillovers will thus centre around i_1^*:

Proposition 5 *The equilibrium gross interest rate of the SOE in period 1 is given by:*

$$(1+i_1^*) = \frac{\phi_{SOE}\,\bar{T}^*\,B^*_{ROW,1}}{\phi_{ROW}\,B_1^*\,B_{SOE,1}\left[\psi_1 \frac{\bar{T}-\mu_1}{B_1(1+i)} + (1-\psi_1)\frac{1}{P_2^T}\right]}. \tag{3.19}$$

This solution for the equilibrium interest rate delivers a clear message: if expectations on fiscally-led policy mixes at ROW grow in period 1, the equilibrium interest rate on SOE debt issued in period 1 *increases*. Formally:

Corollary 2 *If the terminal price level at ROW is higher under the fiscally-led policy mix ($P_2|_{FL} > P_2|_{ML}$), and both cross-border preference parameters are positive ($\phi_{SOE} > 0$, $\phi_{ROW} > 0$), the co-movement between SOE interest rates in period 1 and the probability of a fiscally-led policy mix at ROW is positive; that is:*

$$\frac{\partial(1+i_1^*)}{\partial\psi_1} > 0. \tag{3.20}$$

A deterioration of ROW fiscal conditions correlates with a deterioration in SOE fiscal conditions, even if SOE fiscal policy is unchanged. This result is squarely linked to the specification of the terminal exchange rate above, which induced the linkage in fiscal sustainability between ROW and SOE. Thus, an alternative fiscal-monetary policy specification limiting the factual hegemony of ROW in the determination of the exchange rate would alter our results. Our intuitions are therefore related to Ding and Jiang (2024).

Economically, the equilibrium interest rate on SOE bonds is governed by the ratio ϕ_{SOE}/ϕ_{ROW}: the convenience yield that SOE households derive from ROW bonds relative to the inconvenience cost that ROW households bear for holding SOE bonds. This ratio, scaled by the relative size of cross-border bond positions $B^*_{ROW,1}/B_{SOE,1}$ and the SOE fiscal capacity $\bar{T}^*/B_1^*$, determines the level of i_1^*. The fiscal expectations at ROW, captured by the term in square brackets, enter through the expected inverse price level and transmit the ROW fiscal stance to the SOE interest rate.

At constant taxes and no changes to the aggregate bond issuance in SOE, this higher cost of borrowing in period 1 directly translates to a higher price level in period 2 (and, thus, also higher inflation from period 1 to period 2).

Corollary 3 *The terminal price level in SOE co-moves positively with the probability of a fiscally-led policy mix at ROW:*

$$\frac{\partial P_2^*}{\partial \psi_1} > 0. \tag{3.21}$$

Proof. Follows immediately from equations (3.20) and (C.6). ∎

A higher probability of a Fiscally-led regime in ROW ($\psi_1 \uparrow$) increases the domestic price level generally by equation (C.5). Additionally, with our specification of international financial markets, this increase in the ROW price level must yield an increase in the interest rates charged on SOE bonds for the UIP to hold. This, in turn yields an *increase* in the *SOE terminal price level* P_2^* through the government budget constraint of SOE, given by $P_2^* = \frac{(1+i_1^*)B_1^*}{T^*}$. Therefore, in our simple model, we identify a minimum viable mechanism by which a deterioration in Fiscal conditions at ROW spills over to higher inflation from $t = 1$ to $t = 2$ in SOE, manifesting through a deterioration of the fiscal balance of SOE.

Figure 3.5 presents these comparative statics graphically. Both the expected price level at ROW, here denoted by P_2^e, as well as the materialized price level in SOE P_2^* (which is fully determined in period one through the SOE government budget constraint) increase in the SOE interest rate i_1^*, which in turn increases in the probability of being in the fiscally-led regime. The sensitivity of the SOE price level to the ROW fiscal shock is of a similar order of magnitude in level terms to the ROW price level sensitivity, which is a consequence of the fiscal shock being comparatively 'large' from the perspective of SOE due to their lower bond issuance. This resonates the mechanism of Barro and Bianchi (2025), by which a given fiscal deterioration has larger effects on the price level when the stock of sovereign debt is *smaller*, as a relatively larger change in the price level is needed to ensure that the real deterioration of the fiscal balance is sufficiently large.

As for the exchange rate, note that there are two plausibly opposing effects, in line with equation (C.1): there is a possible direct effect on the terminal exchange rate if the factual policy regime is actually FL through the deterioration of ROW fiscal conditions, and an opposing effect from the spillovers on the fiscal budget of the SOE. In our calibration, the latter effect dominates, which is evidenced by the relatively larger deterioration of the real value of debt at SOE relative to ROW. Here, therefore, a higher probability of the fiscally-led regime in ROW leads to an *appreciation* of the ROW currency in the terminal period. Effectively, this is the international extension of the mechanism presented in Barro and Bianchi (2025): because the relative adjustment of the SOE fiscal position is larger in response to the deterioration of the ROW fiscal condition, the bilateral exchange rate can appreciate from the perspective of ROW.

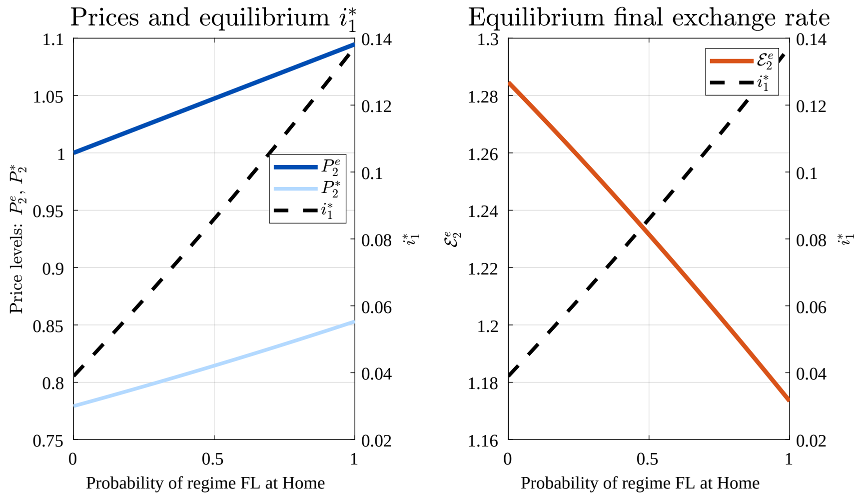

Figure 3.5: Comparative statics with respect to ψ_1, the probability of ROW being in regime FL, for prices and equilibrium exchange rates. Calibration: $B_1 = 1$, $B_1^* = 0.15$, $B^*_{ROW,1} = B_{SOE,1} = 0.1$, $\bar{i} = 0.04$, $\beta = 0.9$, $\bar{T} = 1$, $\bar{T}^* = 0.2$, $\mu_1 = 0.05$, $P_2^T = 1$, $\phi_{SOE} = 0.03$, $\phi_{ROW} = 0.01$.

3.4 A Small Open Economy Model

To rationalize the empirics and the international fiscal inflation mechanism from the previous subsection, we now introduce a discrete-time SOE New Keynesian model. In our model, the government of the SOE issues two different nominal liabilities: a simple one-period bond that carries a convenience yield, and a long-term consol-style bond with a geometric maturity structure. The SOE trades goods and financial assets with the rest of the world, which we here take to be an exogenous entity.

Households

Time is discrete, $t \in \{0, 1, 2, ...\}$. The representative household maximizes the cumulative discounted sum of a standard utility function of the CRRA class, but amended by *convenience yields.*

The representative household in this SOE maximises

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left\{ \log C_t - \frac{N_t^{1+\phi}}{1+\phi} + \nu\Big(b_t^F,\, C_t;\, \theta_t\Big) \right\}, \tag{3.22}$$

where C_t is the consumption aggregate, N_t is labour supplied by the representative household, $\beta \in (0,1)$ is the discount factor, $\phi > 0$ is the inverse Frisch elasticity, and $\nu(\cdot)$ is a convenience flow earned on *foreign* bond holdings - effectively, we postulate that the foreign country's debt features some particular benefit in facilitating payments, or it requires less stringent oversight over its usefulness as a savings asset due to its perception of being relatively safe, which is a feature that has henceforth been associated particularly with U.S. treasury debt (Krishnamurthy and Vissing-Jorgensen, 2012; Greenwood et al., 2015; Nagel, 2016; Van Binsbergen et al., 2022; Di Tella et al., 2023; Krishnamurthy and Ma, 2025). We postulate that the flow on convenience yields must be proportional to the real value of foreign bond holdings in Home consumption units, such that we must adjust the stock of foreign bond holdings B_t^F in the convenience flow function by the nominal exchange rate $\mathcal{E}_t$ and the domestic price level P_t:

$$b_t^F \equiv \frac{\mathcal{E}_t \, B_t^F}{P_t} \tag{3.23}$$

The consumption index is a CES aggregate of Home ($C_{H,t}$) and Foreign ($C_{F,t}$) goods:

$$C_t = \left[(1-\alpha)^{\frac{1}{\eta}} \, C_{H,t}^{\frac{\eta-1}{\eta}} + \alpha^{\frac{1}{\eta}} \, C_{F,t}^{\frac{\eta-1}{\eta}}\right]^{\frac{\eta}{\eta-1}}, \tag{3.24}$$

where $\alpha \in (0,1)$ is the effective home bias in consumption and $\eta > 0$ is the elasticity of substitution between Home and Foreign goods. Cost minimisation over the CES bundle yields the standard demand system and the consumer price index (CPI)

$$P_t = \left[(1-\alpha) \, P_{H,t}^{1-\eta} + \alpha \, P_{F,t}^{1-\eta}\right]^{\frac{1}{1-\eta}}, \tag{3.25}$$

with CPI inflation $\Pi_t \equiv P_t/P_{t-1}$, the terms of trade $\mathcal{S}_t \equiv P_{F,t}/P_{H,t}$, and the nominal exchange rate $\mathcal{E}_t$.

As for the convenience yield, we assume that it has an isoelastic in deviations from steady-state bond holdings, with our specification following Hyland (2026):

$$\nu\Big(b_t^F, \, C_t; \, \theta_t\Big) = \frac{\bar{\Psi} \, \theta_t \, \bar{b}^F}{(1-\zeta) \, C_t} \left[\left(\frac{b_t^F}{\bar{b}^F}\right)^{1-\zeta} - 1\right], \qquad \zeta > 0, \tag{3.26}$$

where $\bar{\Psi} > 0$ is the steady-state convenience-yield parameter, $\bar{b}^F$ is the steady-state real value of foreign bond holdings, denoted in units of the Home currency, and θ_t is an exogenous process capturing time-varying perceived liquidity or safety of foreign government bonds. Scaling by $1/C_t$ makes the convenience flow non-separable from consumption, so that $\partial\nu/\partial C_t = -\nu_t/C_t$, and it allows us to assess convenience flows in direct proportion to foregone utility from consumption when embarking in savings in that convenience-bearing asset. For the following, it is worthwhile to define the marginal convenience value as:

$$\Psi_t \equiv \frac{\partial \nu}{\partial b_t^F} \, C_t = \bar{\Psi} \, \theta_t \left(\frac{b_t^F}{\bar{b}^F}\right)^{-\zeta}, \tag{3.27}$$

and the effective convenience yield as:

$$\Phi_t \equiv \frac{\Psi_t}{1-\nu_t}. \tag{3.28}$$

The household engages in maximization of (3.22) subject to a standard flow budget constraint:

$$P_t C_t + \frac{B_t^S}{1+i_t} + Q_t^L\, B_t^L + \frac{\mathcal{E}_t\, B_t^F}{1+i_t^*} = B_{t-1}^S + \left(1 + \delta\, Q_t^L\right) B_{t-1}^L + \mathcal{E}_t\, B_{t-1}^F + W_t N_t + P_t T_t, \quad (3.29)$$

where B_t^S is the face value of one-period Home bonds purchased at t, which are issued at price $1/(1+i_t)$, B_t^L is the number of units of long-term Home bonds held at the end of period t, which are each priced at Q_t^L and are subject to a geometric decay δ (as in Cochrane (2001)), B_t^F is face value of one-period Foreign bonds purchased at t in terms of the foreign currency, issued at price $1/(1+i_t^*)$, W_t is the nominal wage, and T_t are real lump-sum transfers to or from the government. The term $(1+\delta Q_t^L)B_{t-1}^L$ reflects the fact that each long-term bond outstanding at the start of period t pays a coupon of 1 and retains a residual value of δQ_t^L.

Denoting by λ_t the Lagrange multiplier on the nominal budget constraint (3.29), we obtain a set of fairly standard optimality conditions for the household.

The first-order condition for consumption is

$$\frac{1-\nu_t}{C_t} = \lambda_t P_t, \quad (3.30)$$

through which we can also define a *modified* real stochastic discount factor that takes into account the convenience flows:

$$\tilde{m}_{t,t+1} \equiv \beta\, \frac{(1-\nu_{t+1})\, C_t}{(1-\nu_t)\, C_{t+1}}, \quad (3.31)$$

which reduces to the standard $\beta C_t/C_{t+1}$ when $\nu_t = 0\, \forall t$.

Labour supply is pinned down by a standard intratemporal optimality condition:

$$C_t\, N_t^{\phi} = (1-\nu_t)\, \frac{W_t}{P_t}. \quad (3.32)$$

For the savings assets, we obtain one condition per each of the three assets on offer. The first-order condition for domestic short-term bonds B_t^S yields, after substituting (3.30),

$$\frac{1}{1+i_t} = \mathbb{E}_t\left[\frac{\tilde{m}_{t,t+1}}{\Pi_{t+1}}\right]. \quad (3.33)$$

The first-order condition for domestic long-term bonds B_t^L gives rise to a recursive bond-pricing equation:

$$Q_t^L = \mathbb{E}_t\left[\frac{\tilde{m}_{t,t+1}}{\Pi_{t+1}}\left(1 + \delta\, Q_{t+1}^L\right)\right]. \quad (3.34)$$

Finally, we look at the convenience-bearing foreign bonds B_t^F. The first-order condition on these bonds is given by

$$\frac{\partial \nu}{\partial b_t^F} \cdot \frac{\mathcal{E}_t}{P_t} + \beta\, \mathbb{E}_t[\lambda_{t+1}\, \mathcal{E}_{t+1}] = \lambda_t\, \frac{\mathcal{E}_t}{1+i_t^*}. \tag{3.35}$$

Substituting $\lambda_t = (1-\nu_t)/(C_t P_t)$, $\Psi_t = (\partial \nu / \partial b_t^F)\, C_t$, and $\Phi_t = \Psi_t/(1-\nu_t)$, we obtain a pricing equation relating the foreign interest rate to the convenience flow, the SDF, inflation, and exchange rate movements:

$$\frac{1}{1+i_t^*} = \Phi_t + \mathbb{E}_t\left[\frac{\tilde{m}_{t,t+1}}{\Pi_{t+1}} \cdot \frac{\mathcal{E}_{t+1}}{\mathcal{E}_t}\right]. \tag{3.36}$$

The effective convenience yield Φ_t introduces a wedge, by which the household accepts a lower return $1/(1+i_t^*)$ on foreign bonds because of the convenience yield.

Combining the short-term Home bond Euler (3.33) with the foreign bond Euler (3.36) then yields a slightly modified UIP condition. Define the log exchange rate as standard, $e_t \equiv \log \mathcal{E}_t$. Approximating:

$$i_t - i_t^* = \mathbb{E}_t[\Delta e_{t+1}] + \Phi_t. \tag{3.37}$$

Other than the UIP condition, which relates the returns on home short bonds to foreign short bonds, we can also define the term premium. Combining (3.33) and (3.34), the no-arbitrage condition between short and long home bonds is

$$\mathbb{E}_t\left[\frac{\tilde{m}_{t,t+1}}{\Pi_{t+1}}\, R_{t+1}^L\right] = \frac{1}{1+i_t}, \tag{3.38}$$

where $R_{t+1}^L \equiv \frac{1+\delta\, Q_{t+1}^L}{Q_t^L}$.

Production

The supply side follows a standard New Keynesian specification, which we will briefly summarize here.

Final goods producer: a competitive final good firm aggregates a continuum of differentiated intermediates $j \in [0,1]$ using a Dixit-Stiglitz-type technology:

$$Y_{H,t} = \left[\int_0^1 y_{H,t}(j)^{\frac{\varepsilon-1}{\varepsilon}}\, dj\right]^{\frac{\varepsilon}{\varepsilon-1}}, \qquad \varepsilon > 1. \tag{3.39}$$

Cost minimisation yields a standard demand function for each variety j, $y_{H,t}(j) = \big(p_{H,t}(j)/P_{H,t}\big)^{-\varepsilon} Y_{H,t}$, and the aggregate producer price index at home is equal to $P_{H,t} = \left[\int_0^1 p_{H,t}(j)^{1-\varepsilon}\, dj\right]^{1/(1-\varepsilon)}$.

Intermediate goods producers: each intermediate goods-producing firm j operates the linear technology

$$y_{H,t}(j) = A_t \, n_t(j), \tag{3.40}$$

where A_t is common total factor productivity and $n_t(j)$ is labour input. We postulate a standard Calvo pricing rigidity, by which in each period a fraction $1-\xi$ of firms is randomly selected to reset their price $p^*_{H,t}(j)$, while the remaining fraction ξ keeps its price unchanged. The optimal reset price is chosen to maximise the expected present value of profits, discounted by the household's modified nominal SDF.

Standard manipulations (see Woodford (2011); Galí (2015) yield the Home producer-price Phillips curve

$$\pi_{H,t} = \beta \, \mathbb{E}_t[\pi_{H,t+1}] + \kappa \, \hat{mc}_t, \tag{3.41}$$

where $\pi_{H,t} \equiv \log(P_{H,t}/P_{H,t-1})$ is Home producer price inflation, $\kappa \equiv \frac{(1-\xi)(1-\xi\beta)}{\xi}$ is the Phillips curve slope, and $\hat{mc}_t$ is the log-deviation of real marginal cost from its flexible-price level.

Government

The domestic government consists of both a fiscal and a monetary authority. The fiscal authority must adhere to a flow budget constraint in each period:

$$\frac{B_t^{S,g}}{1+i_t} + Q_t^L \, B_t^{L,g} = B_{t-1}^{S,g} + \left(1 + \delta Q_t^L\right) B_{t-1}^{L,g} + P_t\left(G_t - T_t\right), \tag{3.42}$$

where $B_t^{S,g}$ and $B_t^{L,g}$ are the face value of novel short-term and long-term government bonds outstanding after period-t issuance, respectively. The domestic fiscal authority finances spending G_t via bond issuance, lump-sum payments T_t (which can, however, be negative) and the issuance of both short-term and long-term nominal bonds. Defining $b_t^S \equiv B_t^{S,g}/P_t$, $b_t^L \equiv B_t^{L,g}/P_t$, and $s_t \equiv T_t - G_t$, and dividing (3.42) by P_t, we obtain:

$$\frac{b_t^S}{1+i_t} + Q_t^L \, b_t^L = \frac{b_{t-1}^S}{\Pi_t} + \left(1 + \delta Q_t^L\right) \frac{b_{t-1}^L}{\Pi_t} - s_t. \tag{3.43}$$

Iterating the flow budget constraint forward and imposing standard transversality yields an intertemporal government debt valuation equation. To that goal, we define the nominal SDF as $M_{t,t+k}^n \equiv \prod_{j=1}^k \tilde{m}_{t+j-1,t+j}/\Pi_{t+j}$. Rewrite (3.43) as $d_t = \frac{B_{t-1}^{S,g} + (1+\delta Q_t^L) B_{t-1}^{L,g}}{P_t} - s_t$. Then, the market value d_t can be expressed as $d_t = \mathbb{E}_t[M_{t,t+1}^n \{B_t^{S,g} + (1+\delta Q_{t+1}^L) B_t^{L,g}\}/P_{t+1}]$ after using the household Euler equations. Iterating forward and applying transversality conditions on both types of debt, the debt valuation equation is given by:

$$\frac{B^{S,g}_{t-1} + (1+\delta Q^L_t)\, B^{L,g}_{t-1}}{P_t} = \mathbb{E}_t \sum_{j=0}^{\infty} M^n_{t,t+j}\, s_{t+j}. \tag{3.44}$$

With only short-term debt, the left-hand side of (3.44) would reduce to $B^{S,g}_{t-1}/P_t$, and the price level P_t would be the sole adjustment margin given predetermined for exogenous debt and interest paths. With long-term debt, the bond price Q^L_t can also re-value government debt obligations as in Cochrane (2001, 2025).

To close the government block, we postulate a fiscal rule for surpluses. Each period, the government sets the real primary surplus according to a feedback rule on the real stock of government debt. In log-deviations from steady-state,

$$\hat{s}_t = \psi_b\, \hat{d}_t + \upsilon^s_t, \tag{3.45}$$

where hats denote log-deviations from steady state, $d_t \equiv \frac{b^S_t}{1+i_t} + Q^L_t b^L_t$, $\psi_b \geq 0$ governs the speed of fiscal adjustment to deviations of the value of debt from steady-state, and the exogenous shock altering the surplus process is a standard AR(1) process:

$$\upsilon^s_t = \rho_s\, \upsilon^s_{t-1} + \sigma_s\, \epsilon^s_t, \qquad \epsilon^s_t \sim \text{i.i.d.}(0,1). \tag{3.46}$$

When $\psi_b \geq \bar{r}$, we recover a standard passive-fiscal regime. When $\psi_b < \bar{r}$, the government commits to deficit finance and the fiscal authority is considered active in the language of Leeper (1991).

—

The monetary authority is characterized by an interest rate policy that is akin to a simple Taylor rule:

$$1 + i_t = (1+\bar{\imath}) \left(\frac{\Pi_t}{\bar{\Pi}}\right)^{\phi_\pi} e^{\varepsilon^m_t}, \tag{3.47}$$

where $\bar{\Pi}$ is the inflation target, ϕ_π is the policy feedback coefficient, and ε^m_t is an i.i.d. monetary-policy shock that is AR(1). Log-linearising, we get that

$$\hat{\imath}_t = \phi_\pi\, \hat{\pi}_t + \varepsilon^m_t, \tag{3.48}$$

and the monetary authority is active in the language of Leeper (1991) when $\phi_\pi > 1$, and passive otherwise.

Rest of the world

As we assume that the home economy is an SOE, we postulate that the foreign variables $\{P_t^*, Y_t^*, C_t^*\}$, which capture the rest of the world, are effectively exogenous from the perspective of the SOE. For simplicity, we assume the Foreign economy has zero inflation, $\pi_t^* = 0$, and constant output at its non-stochastic steady-state level unless otherwise specified.

We additionally postulate the foreign nominal interest rate i_t^* to be stochastic. It has a steady state $\bar{\imath}^*$ but fluctuates around it according to an exogenous process. The Home household takes i_t^* as given when choosing its foreign bond position.

Market Clearing and Equilibrium

We require market clearing on the goods market, the labour market, and all bond markets in our SOE.

On the goods market, the Home good is demanded both domestically and abroad. Domestic demand for the Home good is $C_{H,t} = (1-\alpha)(P_{H,t}/P_t)^{-\eta}C_t$, while foreign demand is equal to $C_{H,t}^* = \alpha^*(P_{H,t}^*/P_t^*)^{-\eta^*}C_t^*$. Goods market clearing thus is given by

$$Y_{H,t} = C_{H,t} + C_{H,t}^*. \tag{3.49}$$

The labour market clears at the wage satisfying the household's optimality condition (3.32) and the firms' labour demand $N_t = Y_{H,t}/A_t$.

On the bond market, we have to consider all three savings assets available to the household. For Home-issued bonds, their supply by the government must equal household holdings in every period.

$$B_t^S = B_t^{S,g}, \tag{3.50}$$

$$B_t^L = B_t^{L,g}. \tag{3.51}$$

On the market for foreign bonds, the household's position B_t^F is determined by their optimal portfolio choice, which equates the return on foreign bonds (adjusted for their convenience benefits) to the return on domestic bonds via the modified UIP condition (3.37). The real value of these bonds $b_t^F = \mathcal{E}_t B_t^F/P_t$ is pinned down endogenously through adjustments of the exchange rate and the domestic price level.

Finally, we describe the evolution of the net foreign asset (NFA) position, which evolves dynamically in line with the exchange rate. Combining the household budget constraint (3.29) with the government budget constraint (3.42), the NFA evolves over time as:

$$\frac{\mathcal{E}_t B_t^F}{1+i_t^*} = \mathcal{E}_t\, B_{t-1}^F + P_{H,t}\, Y_{H,t} - P_t\, C_t - P_t\, G_t. \tag{3.52}$$

In real terms, recalling that $b_t^F = \mathcal{E}_t B_t^F / P_t$ and defining the real net export surplus as $nx_t = (P_{H,t}/P_t)\, Y_{H,t} - C_t - G_t$:

$$\frac{b_t^F}{1+i_t^*} = \frac{b_{t-1}^F}{\Pi_t} \cdot \frac{\mathcal{E}_t}{\mathcal{E}_{t-1}} + nx_t. \tag{3.53}$$

Steady-state: we now summarize the non-stochastic steady state. In the following, we focus on the steady-state with zero CPI inflation, $\bar{\Pi} = 1$, and zero inflation in the ROW, $\bar{\Pi}^* = 1$. Overbars denote steady-state values. In line with our previous assumptions, the nominal exchange rate is constant at steady state at $\bar{\mathcal{E}}$.

The nominal interest rate is pinned down directly from the household Euler equation on short home bonds:

$$\frac{1}{1+\bar{\imath}} = \beta \qquad \Longrightarrow \qquad \bar{\imath} = \frac{1}{\beta} - 1. \tag{3.54}$$

The only nominal interest rate consistent with market clearing on the foreign bond market can be determined by using (3.36) and imposing $\tilde{m} = \beta$ and $\Pi = 1$:

$$\frac{1}{1+\bar{\imath}^*} = \bar{\Phi} + \beta \qquad \Longrightarrow \qquad \bar{\imath}^* = \frac{1}{\bar{\Phi}+\beta} - 1. \tag{3.55}$$

Equation (3.34) gives us the steady-state price for the long-term home bond as

$$\bar{Q}^L = \beta\big(1 + \delta\,\bar{Q}^L\big) \qquad \Longrightarrow \qquad \bar{Q}^L = \frac{\beta}{1-\delta\beta}. \tag{3.56}$$

Solution of the model under log-linearisation

Thanks to the existence of long-term consols and the convenience yields on foreign assets, we can still obtain some interesting dynamics under a solution of the model to first-order. We therefore now log-approximate the set of equilibrium conditions and thereafter define an appropriate equilibrium concept. Throughout, $\hat{x}_t \equiv \log(X_t/\bar{X})$ for any variable X_t, and we use the approximation that $\bar{\nu}$, the impact of the convenience flow in steady-state, is small, which is in line with our definition of the convenience function itself. We define the composite parameter

$$\varpi \equiv \frac{\bar{\Phi}}{\beta+\bar{\Phi}} = (1+\bar{\imath}^*)\bar{\Phi}, \tag{3.57}$$

which measures the share of the foreign bond discount factor attributable to the convenience yield.[10]

[10] Note that this is *not* equal to zero: even though the convenience flow can be zero in steady-state, our definition of the convenience *yield* attains a non-zero level for any positive level of foreign bond holdings even in steady-state.

—

We begin by linearizing the expression capturing the convenience yield dynamics. Combining equations (3.27) and (3.28):

$$\hat{\Phi}_t = \hat{\theta}_t - \zeta \hat{b}_t^F, \tag{3.58}$$

i.e., the effective convenience yield rises with the convenience shock θ_t and *falls* with the quantity of real foreign bond holdings b_t^F, in line with Krishnamurthy and Vissing-Jorgensen (2012). Crucially, since $b_t^F = \mathcal{E}_t B_t^F / P_t$, the exchange rate directly enters the convenience yield, as

$$\hat{b}_t^F = \hat{e}_t + \hat{B}_t^F - \hat{P}_t, \tag{3.59}$$

where $\hat{e}_t = \log(\mathcal{E}_t / \bar{\mathcal{E}})$. A depreciation of the Home currency ($\hat{e}_t > 0$) here raises the value of foreign bond holdings, ceteris paribus reducing the marginal convenience yield as it is substituted away into the real value of the underlying bonds.

We can find a standard IS curve by log-linearising the short-term Home bond Euler equation (3.33):

$$\hat{c}_t = \mathbb{E}_t[\hat{c}_{t+1}] - \Big(\hat{\imath}_t - \mathbb{E}_t[\hat{\pi}_{t+1}]\Big). \tag{3.60}$$

To link home output and the consumption levels, we can combine goods-market clearing with the specification and international demand to get that:

$$\hat{y}_{H,t} = (1-\alpha)\hat{c}_t + \alpha\eta\,\hat{\mathcal{S}}_t + \alpha\,\hat{c}_t^*, \tag{3.61}$$

where $\hat{c}_t^*$ is exogenous foreign demand, and the output gap is defined as is standard in deviations from the natural output level, $x_t = \hat{y}_{H,t} - \hat{y}_{H,t}^n$.

The long-term bond price in log-deviations from steady-state can be found by log-linearising (3.34) around $\bar{Q}^L = \beta/(1-\delta\beta)$:

$$\hat{Q}_t^L = -\Big(\hat{\imath}_t - \mathbb{E}_t[\hat{\pi}_{t+1}]\Big) + \delta\beta\,\mathbb{E}_t[\hat{Q}_{t+1}^L]. \tag{3.62}$$

Iterating forward, we can see that the price of the long-term assets is a function of the expected future real interest rate, adjusted for the maturity structure.

$$\hat{Q}_t^L = -\sum_{k=0}^{\infty} (\delta\beta)^k \mathbb{E}_t\Big[\hat{\imath}_{t+k} - \hat{\pi}_{t+k+1}\Big]. \tag{3.63}$$

The UIP condition is relatively standard, but takes into account the possible convenience yield earned on foreign bond holdings.

$$\hat{\imath}_t - \hat{\imath}_t^* = \mathbb{E}_t[\Delta\hat{e}_{t+1}] + \varpi\,\hat{\Phi}_t. \tag{3.64}$$

Following standard derivations in open-economy settings (as in Hyland (2026)), the linearised NKPC in terms of the output gap $x_t \equiv \hat{y}_{H,t} - \hat{y}^n_{H,t}$ is pinning down the PPI in the SOE as

$$\hat{\pi}_{H,t} = \beta\, \mathbb{E}_t[\hat{\pi}_{H,t+1}] + \kappa\Big[(\phi+1)x_t + \alpha\, \hat{\mathcal{S}}_t\Big], \tag{3.65}$$

where $\hat{\mathcal{S}}_t$ is the log-deviation of the terms of trade and $\kappa = (1-\xi)(1-\xi\beta)/\xi$. The CPI and the PPI are consequently linked as:

$$\hat{\pi}_t = (1-\alpha)\,\hat{\pi}_{H,t} + \alpha\,\Delta\hat{\mathcal{S}}_t. \tag{3.66}$$

The terms-of-trade and the exchange rate are linked as is standard in open-economy models. Imposing law of one price:

$$\Delta\hat{\mathcal{S}}_t = \Delta\hat{e}_t + \hat{\pi}^*_t - \hat{\pi}_{H,t}. \tag{3.67}$$

Next, we focus on the dynamics of government debt in our model by linearising the government budget constraint (3.43). Defining the maturity composition as $\omega^S \equiv \bar{b}^S/[(1+\bar{\imath})\bar{d}]$, with $\omega^L = 1-\omega^S$, and noting that at steady-state $\bar{s} = \bar{\imath}\,\bar{d} = \bar{r}\,\bar{d}$, we obtain that the government budget constraint must evolve according to

$$\hat{d}_t = \frac{1}{\beta}\Big[\hat{d}_{t-1} + \omega^S\Big(\hat{\imath}_{t-1} - \hat{\pi}_t\Big) + \omega^L\Big(\hat{R}^L_t - \hat{\pi}_t\Big)\Big] - \frac{\bar{s}}{\bar{d}}\,\hat{s}_t, \tag{3.68}$$

where $\hat{R}^L_t \equiv \log(R^L_t/\bar{R}^L)$ is the log-deviation of the long-bond gross return. The long-bond return itself is equal to

$$\hat{R}^L_t = \delta\beta\,\hat{Q}^L_t - \hat{Q}^L_{t-1}, \tag{3.69}$$

in log-linear form, since $\bar{R}^L = (1+\delta\bar{Q}^L)/\bar{Q}^L = 1/\beta$.

For the purposes of tracking the aggregate budget constraint of the SOE across time, we must also log-linearise the NFA equation (3.53) around the steady state. Using $\bar{b}^F/(1+\bar{\imath}^*)$ on the LHS and $\bar{b}^F + \overline{nx}$ on the RHS, we get:

$$\frac{\bar{b}^F}{1+\bar{\imath}^*}\hat{b}^F_t - \frac{\bar{b}^F}{(1+\bar{\imath}^*)^2}\hat{\imath}^*_t = \bar{b}^F\Big(\hat{b}^F_{t-1} + \Delta\hat{e}_t - \hat{\pi}_t\Big) + \overline{nx}\,\hat{nx}_t,$$

where $\hat{nx}_t$ denotes the log-deviation of real net exports and we have used $\hat{b}^F_{t-1}$ evaluated with the previous period's exchange rate. Dividing through by $\bar{b}^F/(1+\bar{\imath}^*)$,

$$\hat{b}^F_t = (1+\bar{\imath}^*)\Big(\hat{b}^F_{t-1} + \Delta\hat{e}_t - \hat{\pi}_t\Big) + \frac{1}{1+\bar{\imath}^*}\hat{\imath}^*_t + \frac{(1+\bar{\imath}^*)\overline{nx}}{\bar{b}^F}\,\hat{nx}_t. \tag{3.70}$$

Finally, we also provide a linear approximation of the government debt valuation equation. Combining the linearised Euler equations with the Taylor rule and iterating the debt identity forward yields the linearised analogue of (3.44):

$$\hat{d}_{t-1} + \omega^S\Big(\hat{\imath}_{t-1} - \hat{\pi}_t\Big) + \omega^L\Big(\hat{R}^L_t - \hat{\pi}_t\Big) = (1-\beta)\sum_{j=0}^{\infty}\beta^j\,\mathbb{E}_t[\hat{s}_{t+j}]. \tag{3.71}$$

—

To close the model, we specify a standard system of exogenous AR(1) shocks:

$$\hat{\theta}_t = \rho_\theta\,\hat{\theta}_{t-1} + \sigma_\theta\,\epsilon^\theta_t, \qquad \epsilon^\theta_t \sim \text{i.i.d.}(0,1), \tag{3.72}$$

$$\hat{a}_t = \rho_a\,\hat{a}_{t-1} + \sigma_a\,\epsilon^a_t, \qquad \epsilon^a_t \sim \text{i.i.d.}(0,1), \tag{3.73}$$

$$\upsilon^s_t = \rho_s\,\upsilon^s_{t-1} + \sigma_s\,\epsilon^s_t, \qquad \epsilon^s_t \sim \text{i.i.d.}(0,1), \tag{3.74}$$

$$\varepsilon^m_t = \rho_\varepsilon\varepsilon^m_{t-1} + \sigma_m\epsilon^m_t \qquad \epsilon^m_t \sim \text{i.i.d.}(0,1), \tag{3.75}$$

$$\hat{\imath}^*_t = \rho_{i^*}\hat{\imath}^*_{t-1} + \sigma_{i^*}\epsilon^{i^*}_t, \qquad \epsilon^{i^*}_t \sim \text{i.i.d.}(0,1). \tag{3.76}$$

We summarize all the equilibrium conditions in the following equilibrium definition:

Definition 3 (Competitive equilibrium) *A log-linearized competitive equilibrium in the SOE is made up of a sequence of exogenous variables* $\{\hat{\theta}_t, \hat{a}_t, \upsilon^s_t, \varepsilon^m_t, \hat{\imath}^*_t\}_{t\geq 0}$ *and endogenous variables* $\{x_t,\ \hat{\pi}_{H,t},\ \hat{\pi}_t,\ \hat{c}_t,\ \hat{\imath}_t,\ \hat{e}_t,\ \hat{\mathcal{S}}_t,\ \hat{Q}^L_t,\ \hat{R}^L_t,\ \hat{d}_t,\ \hat{b}^F_t,\ \hat{\Phi}_t,\ \hat{s}_t\}_{t\geq 0}$*, such that*

- *the exogenous shock processes are governed by equations* (3.72)*,* (3.73)*,* (3.74)*,* (3.75)*,* (3.76)*,*
- *and the endogenous variables are governed by the equations* (3.45)*,* (3.48)*,* (3.58)*,* (3.62)*,* (3.64)*,* (3.60)*,* (3.65)*,* (3.66)*,* (3.67)*,* (3.68)*,* (3.69)*,* (3.70)*,* (3.71)*.*

In total, the system consists of 18 equilibrium condition in 13 endogenous and 5 exogenous variables when approximated up-to-first-order.

Calibration and parametrization

We parametrize and calibrate the model in accordance with canonical insights from the literature on open-economy models in macroeconomics, following a broad strategy of calibration of a number of parameters externally and close to best-practices of the literature, and estimating a significant number of parameters using long-running macroeconomic data. In appendix C.4, we provide details on the construction of the data series.

Parameter	Value	Source / calibration target
β	0.99	Standard quarterly discount factor
ϕ	2.00	Standard value for inverse Frisch elasticity
α	0.40	Gali and Monacelli (2005)
η	1	Gali and Monacelli (2005)
κ	0.019	Hazell et al. (2022)
$\bar{s}/\bar{d}$	0.02	OBR (2024)
$\bar{i}^*$	0.01	Lower bound of estimates for the U.S. of Lubik et al. (2024)
$\overline{nx}/\bar{b}^F$	0.02	

Table 3.4: Externally calibrated parameters

We estimate the model using long-running time series since 1985 on the British economy as our SOE, postulating that the ROW block is approximated by the behaviour of U.S. time series. We employ a two-step Simulated Method of Moments (SMM) and target in total 14 parameters using 23 moments. For the convenience yield, we target the properties of the shock process θ_t as well as the steady-state convenience yield level $\bar{\Psi}$ using the observed variance of the empirical convenience yield, as well as the covariance of the convenience yield with the bilateral exchange rate, foreign debt supply, and the foreign policy rate. As for the debt structure and debt valuation blocks, we target the properties of the fiscal innovation series, the long-term coupon structure δ, and the overall debt structure ω_S using the variances of fiscal deficits, the overall value of debt, and the return on long-term debt, as well as the covariances between deficits, inflation, and the market value of the stock of government debt. The monetary policy shock processes are identified using the empirical (co-)variances of the respective interest and inflation rates, as well as by the covariance between interest rates and exchange rates. Finally, we target standard moments of the SOE, such as the variation and autocovariance of output, the covariance between output and consumption, the covariance of output and inflation, and the variances of the exchange rate and bond supply.

The parameters that we vary are the shock autocorrelations and standard deviations $(\rho_\theta, \rho_a, \rho_s, \rho_m, \rho_{i^\star}, \sigma_\theta, \sigma_a, \sigma_s, \sigma_m, \sigma_{i^\star})$ as well as the steady-state level of the convenience yield $\bar{\Psi}$, the shape of the CRRA nature of the convenience yield function ζ, and the precise coupon structure of the long-term bond δ.

What we do *not* estimate are the policy parameters, ϕ_π for monetary policy and ψ_b for fiscal policy. Instead, we heuristically consider four possible policy regimes, one strongly active-monetary regime ($\phi_\pi = 2.0, \psi_b = 1.0$), one weakly active-monetary regime ($\phi_\pi = 1.25, \psi_b = 0.75$), one weakly active-fiscal regime ($\phi_\pi = 0.5, \psi_b = 0.0$), and one strongly active-fiscal regime ($\phi_\pi = 0.0, \psi_b = -0.5$). Under the drawback of hard-coding and imposing one of the four policy regimes throughout our estimation period, each estimation is well-defined on its given policy space and yields the model estimate by which the estimated parameters can best replicate the empirical regularities, conditional on the policy regime imposed.[11]

Table 3.5 summarizes the estimated parameters under each of the imposed policy regime.[12] While the overall fit between the data and the model does not vary much, the fiscal-dominance regimes seem to be slightly preferred over time, coming with much higher steady-state convenience yields.

Parameter	ϕ_π=2, ψ_b=1	ϕ_π=1.25, ψ_b=0.75	ϕ_π=0.5, ψ_b=0	ϕ_π=0, ψ_b=0
Fixed policy parameters				
ϕ_π	2.00	1.25	0.50	0.00
ψ_b	1.00	0.75	0.00	0.00
Estimated structural parameters				
$\bar{\Psi}$	0.4828	0.7955	1.4684	1.6441
ζ	1.6976	1.8951	1.6901	1.8333
δ	0.8531	0.9022	0.8998	0.8585
ω_S	0.5423	0.2769	0.0800	0.2291
ρ_θ	0.0016	0.5621	0.5604	0.6853
ρ_a	0.6656	0.8039	0.6075	0.5827
ρ_s	0.5667	0.5403	0.0016	0.0014
ρ_m	0.5440	0.6016	0.6344	0.6723
$\rho_{i^\star}$	0.8235	0.7888	0.8932	0.8913
Estimated shock standard deviations				
σ_θ	0.0059	0.0289	0.0254	0.0204
σ_a	0.1263	0.0872	0.0844	0.0866
σ_s	0.0063	0.0084	0.0264	0.0263
σ_m	0.0025	0.0043	0.0043	0.0044
$\sigma_{i^\star}$	0.0091	0.0096	0.0105	0.0105
Optimization statistics				
Objective function	0.0524	0.0503	0.0496	0.0494
J-stat p-val	0.4006	0.4345	0.4469	0.4493

Standard errors in parentheses. ϕ_π and ψ_b are fixed across estimation runs.

Table 3.5: SMM Estimation Results by Policy Regime

[11] Clearly, we do *not* consider the possibility of time-switching regimes, which has been explored in the context of the U.S. in Davig and Leeper (2007) and Bianchi and Melosi (2017, 2019), as well as the modern exposition of simultaneously coexistent policy regimes in dependence on the precise nature of a given shock in Bianchi et al. (2023). We leave the possibility to explore such regime switches open for future research, noting that chapter 1 of this thesis made some steps moving towards the identification of such regimes in the U.S.

[12] Table C.2 in the appendix details the fit between the data moments and the estimated moments in each of the policy regimes.

Impulse-responses of the estimated model

With our four different policy scenarios at hand, we now compare the model-implied behaviour of macroeconomic aggregates. The impulses we simulate are meant to capture differences in the relative attractiveness of ROW vs. SOE sovereign bonds across the four policy regimes under the various estimated models.

Figure 3.6 considers the response of the economy to a decrease in convenience yields. In the very-active monetary case, the economy effectively stays at steady-state. There are two reasons for this: first, the calibrated convenience yield shock in that scenario features no persistence; second, given the fact that the fiscal authority promises to react to deviations of the value of debt from steady-state one-for-one by adjusting surpluses, any possible windfall gain from substitution to the SOE fiscal asset gets nullified on impact as households are Ricardian.

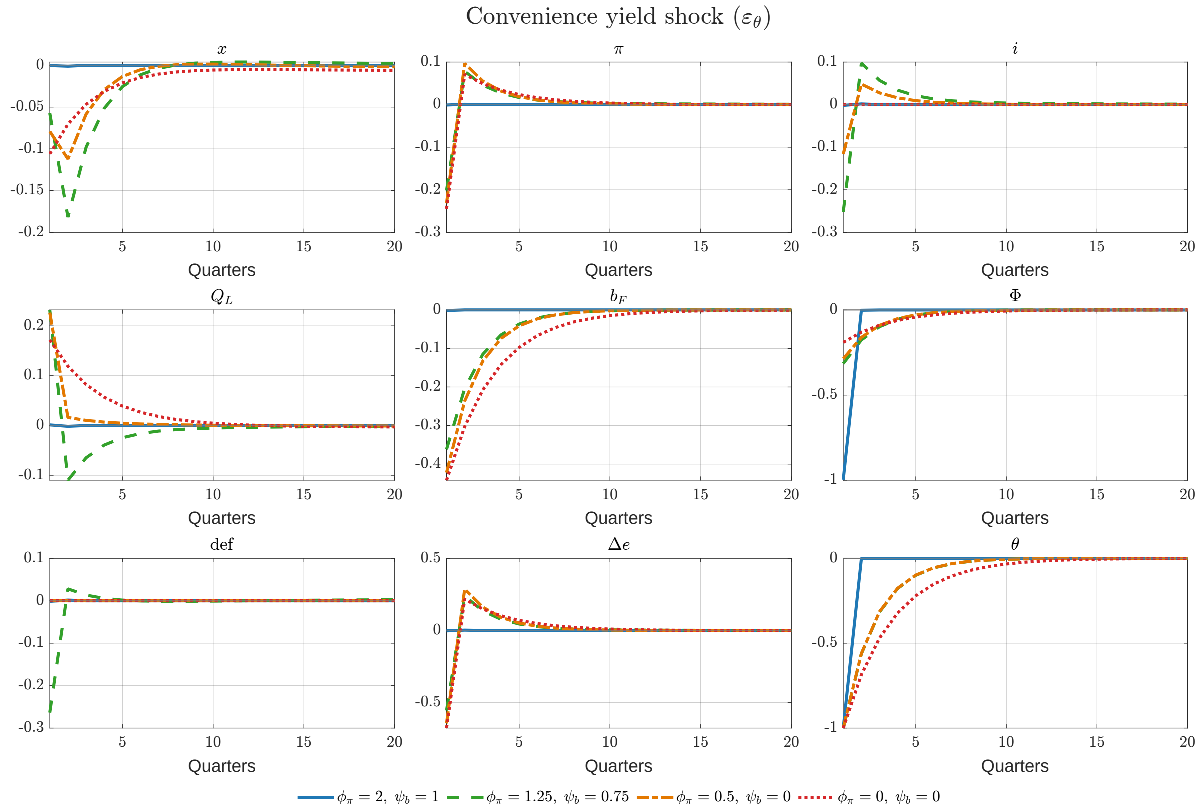

Figure 3.6: IRFs to a convenience yield shock θ_t

In all other scenarios, as convenience yield flows decrease, households themselves substitute away from foreign bond holdings as evidenced by the response of b_t^F, substituting instead towards home bonds whose price increases. Concurrently, the exchange rate appreciates from the perspective of the SOE on impact before gradually returning to steady-state. Due to the increase in bond prices and the concurrent decrease of interest rates that must be paid on SOE debt, households rationally expect monetary and fiscal policy to react in the next period.

Under the scenario with relatively active monetary and quite passive fiscal policy (green line), the substitution towards the SOE savings asset causes the price of long-term bonds to increase and yields a windfall gain for the fiscal authority. However, the overall negative impact for the SOE households from the decrease in convenience yields contributes to a negative deviation in the output gap in the subsequent period. This happens even though the SOE exchange rate appreciates relative to steady-state due to the smaller convenience flows. All this is rooted in the UIP equation (3.64): the change to convenience yield flows and the exchange rate adjustment do not offset one-for-one. Instead, from period 1 onwards, the modified UIP condition requires a tightening of SOE monetary policy, which is consistent with the observed decrease in the output gap as well as the increase in inflation.

Under either scenario with the fiscally-led policy mix, the response of the economy is not too different. The only material change is that the fiscal balance remains silent (consistent with the specification of active fiscal policy) while the price of long-term debt never drops below steady-state, since the smaller monetary tightening from period 1 onwards precludes a deterioration of the fiscal balance under which such a decrease in prices would be model-consistent.

—

While we cannot model the overall budget balance of the ROW in this model (as this would require specifying ROW demand for the ROW sovereign bonds), figure 3.7 attempts to move close to replicating a substitute impacting the budget balance of the ROW: a tightening of the ROW interest rate, i_t^*, making ROW bonds more attractive.[13] For the SOE households, this increase in ROW interest rates makes them more attractive, contributing to an inflow into ROW debt (increase in b_t^F), and outflow of SOE debt (decrease in Q_t^L), a depreciation of the SOE currency, and a windfall gain on impact for the existing bondholders of ROW debt, contributing to the observed real expansion with an increase in the output gap and inflation.

Under the aggressive monetary stance (blue and green lines), the SOE monetary authority reacts by instantaneously tightening their policy stance to avoid an economy that keeps spiralling away from the positive initial wealth effect. Consequently, the attractiveness of SOE bonds increases and the economy converges back to steady-state. The fiscal authority initially must engage in deficit-finance as their bonds are not taken up, and subsequently raises surpluses to ensure budget balance.

[13]Under unchanged surpluses in the ROW, this means that bondholders in the ROW will need to be taxed more to repay the higher interest payments, amounting to a possibly negative wealth effect on behalf of these taxpayers.

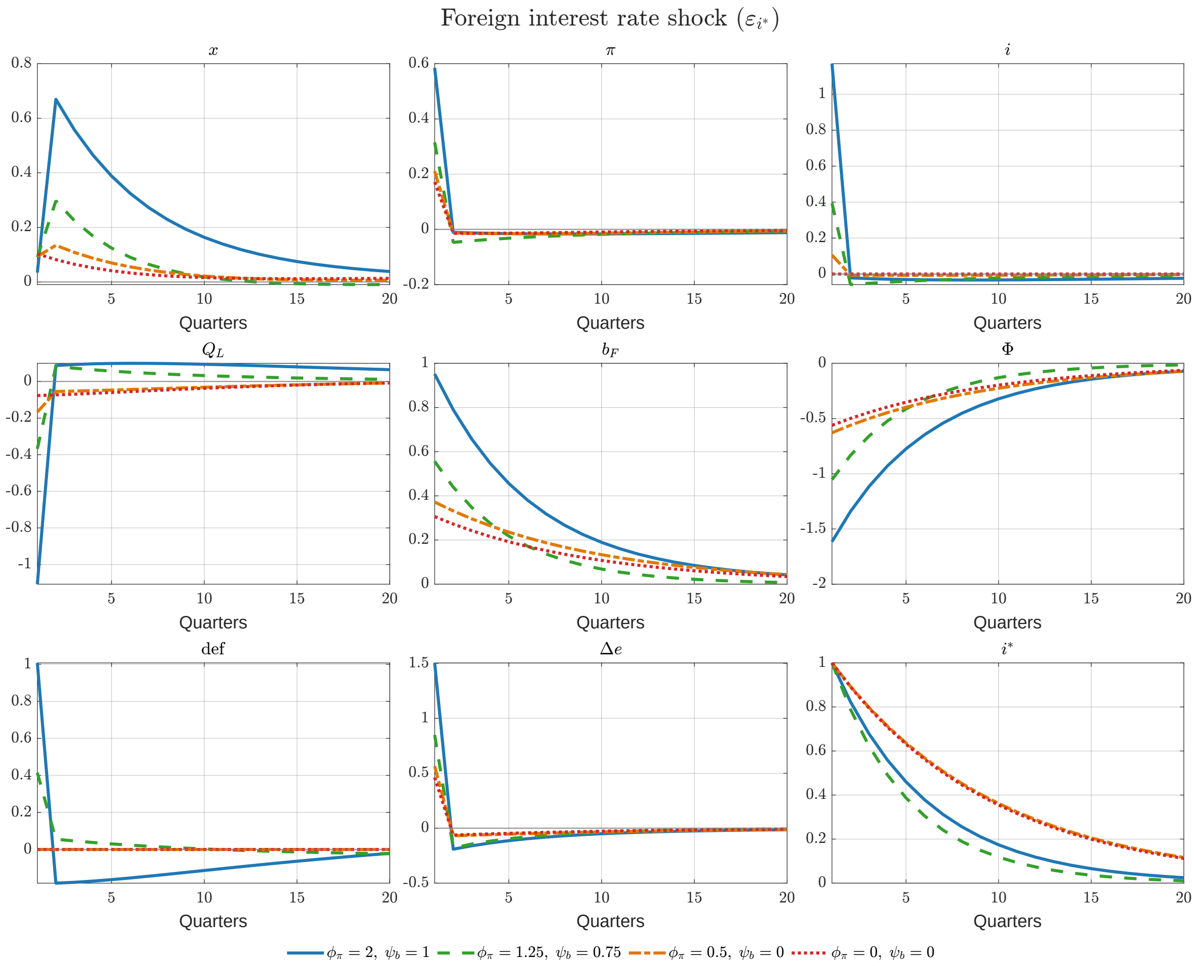

Figure 3.7: IRFs to a foreign interest rate shock i_t^*

Under a fiscally-led policy mix (orange and red lines), the initial real expansion is more muted, even though the monetary tightening of the SOE monetary authority is less pronounced. This is, in turn, rooted in the behaviour of the fiscal authority: as the fiscal stance does not change, the concurrent outflow from the SOE fiscal asset is much smaller. Therefore, the windfall gain experienced by households through substitution into the ROW savings asset is much smaller, contributing to a significantly smaller real and nominal expansion.

—

Finally, what happens if we change the fiscal stance of the SOE? Figure 3.8 shows the answer to this question, embodied through an expansionary government spending impulse.

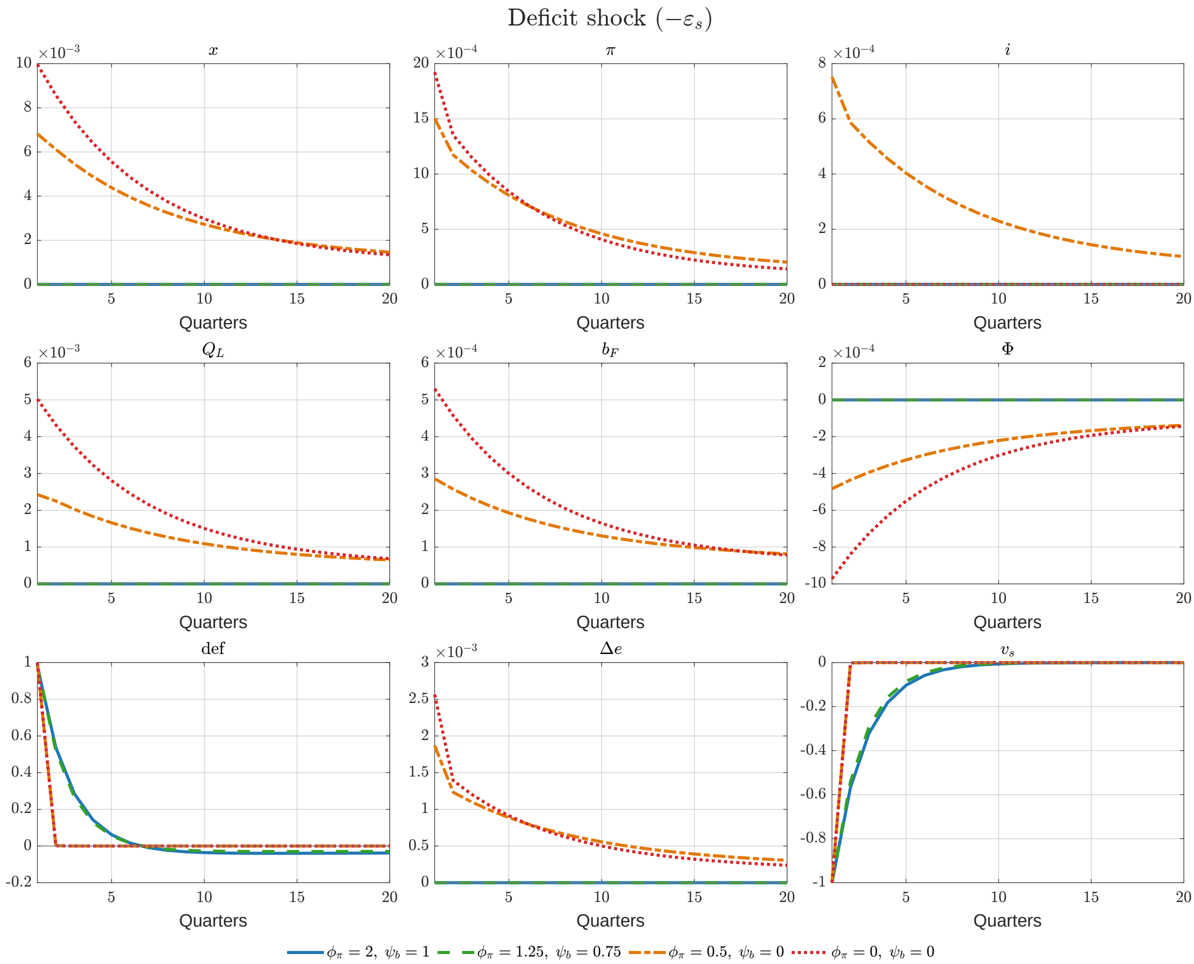

Figure 3.8: IRFs to a deficit shock $-s_t$

Under the blue and green lines, the economy stays at its steady-state. This is not surprising: the deficit may increase, but as all debt is held by the SOE households, they simply expect to repay the additional deficits through taxation, consistent with Ricardian Equivalence.[14] Under the fiscally-led policy mix, the deficit-financed fiscal measure is tantamount to a positive demand shock. Consequently, output and inflation increase, with their persistence outliving the fiscal impulse expense. The price of any debt offered by the SOE sovereign must go up to clear the debt market, as otherwise demand for such savings products would go down. The convenience flows to the SOE household decrease as holdings of SOE debt increase more in relative terms. Due to the worsened overall fiscal stance, the currency of the SOE must depreciate.

[14]This stands in sharp contrast to Hyland (2026), where the two-country nature of the model yields some adjustment in the economy due to home bonds being held abroad, intensifying the international revaluation channel.

3.5 Conclusion

This chapter shows that fiscal sustainability in open economies is jointly disciplined by expected primary surpluses and by the currency-denominated discount factor that global investors apply to those surpluses. This linkage connects fiscal innovations to exchange rates and allows fiscal policy to spill over internationally. Using market-value debt data for the United Kingdom and the United States since 1975, we apportion roughly 50 percent of unexpected year-to-year changes in the U.K. debt ratio to revisions in surplus and deficit expectations, and the remaining 50 percent to discount-rate news, with the latter split across global yields, real exchange rate revisions, and time-varying UIP premia. The revisions in surplus expectations, in turn, are highly correlated with corresponding revisions in observed exchange rates. We have then constructed a discrete-time SOE New Keynesian model that embeds three interacting features: (i) a convenience yield on foreign bonds, derived from an isoelastic specification non-separable from consumption, reflecting the liquidity or safety services of foreign-currency assets; (ii) a geometric consol that introduces long-term domestic debt alongside the standard one-period domestic bond; and (iii) a stochastic foreign interest rate that drives fluctuations in the return on the convenience-bearing asset.

Our headline results shed light on the importance of the fiscal and monetary stances of the respective policy authorities in either mitigating or enhancing the cross-border transmission of shocks. Interestingly enough, and contrary to some conventional wisdom, a fiscally-led policy mix may be enhancing in terms of macroeconomic stability for an SOE, as this can mitigate the impact of foreign shocks onto the SOE, especially of non-pecuniary aspects of utility borne by the SOE households through holdings of ROW debt. Further exploring these insights in a richer model that takes into account the full structure of ROW debt is the logical next step in this research agenda.

Chapter 4

Towards a Bullwhip Theory of Supply Chains*

Abstract

This paper develops a macroeconomic model of supply chains, in which demand, captured by explicit order placement, need not be fulfilled concurrently. Our model gives rise to precautionary inventory accumulation behaviour and the amplification of shocks on supply chains, endogenously giving rise to the 'bullwhip effect' known from the Operations Research literature under additional conditions. The basic model is capable of approximating the relative dynamics of the (inventory-)sales volatility on the upper end of the supply chain relative to the lower end of the supply chain in response to shocks on the production and delivery technology. In addition, the model highlights the importance of price adjustment to ensure intertemporal market clearing when order backlogs become too costly. The model mechanism is highly dependent on demand-side policies, which matter significantly for the extent of the bullwhip effect observable in equilibrium.

*This chapter is co-authored with Michael McMahon. We thank participants of the University of Oxford Macroeconomics Work-in-Progress workshop for comments and suggestions.

4.1 Introduction

Inventories have not played a particularly prominent role in macroeconomic research in the last 40 years, despite their importance for the propagation and amplification of shocks in the operations management literature. One reason for this is that the two fields typically model inventories in quite different ways. The latter has since rigorously established the existence of the *bullwhip effect* - the empirical regularity that the volatility of sales and inventories significantly amplifies when progressing up supply chains (Lee et al., 1997a,b; Fransoo and Wouters, 2000; Wang and Disney, 2016), which is related to precautionary inventory management in the face of uncertainty. For any given increase in volatility downstream in a supply chain, a more than one-for-one increase in sales and inventory volatility upstream occurs.

In this paper, we introduce a new macroeconomic theory of inventory behaviour along a supply chain, rationalizing the *bullwhip effect* as a natural, state-dependent amplifier of volatility on supply chains. Rather than including inventories as an input required for production or sales, in our environment firms evaluate the full distribution of exogenous disturbances when undertaking decisions. Heuristically, firms have to undertake production decisions *before* any supply chain-specific stochastic innovation manifests itself, causing firms to consider the entire prior distribution of shocks when formulating production decisions.[1] The resulting propagation mechanism in firm behaviour, we argue, is a direct consequence of the inventory choices. Combining empirical facts on supply chain distortions with our novel theoretical approach, we rationalize the at-times larger upstream volatility of the cyclical component of sales, as well as the time-changing nature of the relative volatility of the inventory-sales ratio upstream versus downstream on supply chains. The degree to which our model replicates the observed negative correlation between inventory-sales ratios and prices remains, however, imperfect and leaves scope for further research on the nature and causes of the bullwhip effect.

We motivate our model derivations using empirical evidence on aggregate bullwhip-type behaviour in a three-stage supply chain, partitioning firms in the economy into manufacturers, who undertake all production decisions of the final good, wholesalers, who engage in labour-intensive relocation of goods, and retailers, who are ultimately concerned with the sale of the good to households. This partition follows exactly the U.S. Census classification and allows us to link aggregate inventory and sales dynamics directly to our model. Isolating the cyclical component of sales, we show that the volatility of sales is principally larger in business cycle-frequency at the upstream end of supply chains relative

[1] Such ideas are not reflected commonly in macroeconomic models of inventories, with notable exceptions such as Kryvtsov and Midrigan (2013).

to the downstream end. In terms of the inventory-sales ratio, the relative magnitude of upstream to downstream volatility changes over time, consistent with a narrative of varying shocks impacting the business cycle since the Great Moderation. By allowing for distinct dynamics in response to upstream and downstream shocks, our model partially rationalizes what plausibly drove cyclical movements in inventory-sales ratios at a given business cycle stage, since each shock implies distinct reactions of upstream and downstream variables.

To understand how the bullwhip effect can be sourced in inventory behaviour on supply chains, we devote significant time to our modelling innovations in this paper. First, we allow for imperfect market clearing at each given stage of the supply chain - despite price adjustments, it is possible for goods demand to be unequal to supply, constraining deliveries of goods in a period. We introduce a *net order book/inventory* variable, which captures pent-up deliveries (or excess delivery capacities) in a single Markovian state variable. Second, to elicit non-zero inventory holdings in steady-state, we introduce what we call *distributional optimisation*: each firm at a given supply chain stage cannot observe its productivity within the same period, such that it takes into account the possibility of production shortfalls, which could dent the quantity of goods that can be delivered in that same period. This gives rise to a precautionary motive for inventory holdings.[2] Both of these frictions ensure that inventories 'make or break' the *propagation* of shocks: inventories can be thought of as 'buffer stock', i.e., as a device dampening demand and supply surprises that would otherwise exacerbate under our assumed delivery frictions.

We then recover the implied policy functions of our multi-stage supply chain model to compare the degree to which exogenous innovations can account for the bullwhip effect in our theory. Upstream productivity innovations clearly replicate the larger volatility in upstream sales and inventory-sales ratios. Downstream productivity shocks usually remain firmly anchored downstream and do not load much on upstream sales and inventory-sales ratio volatility. Finally, we look at government transfer shocks, which proxy pure demand disturbances. In this exercise, the upstream amplification of sales and inventory-sales ratio volatility is obtained more frequently, but the precise scope of this upstream amplification depends very much on the monetary-fiscal policy mix in place. In this exercise, we also get closer towards observing downstream price pressure under some calibrations, and we are able to obtain a conditional negative correlation between prices and inventory-sales ratios upstream whenever the monetary-fiscal policy specification deviates from the commonly assumed baseline.

[2]The majority of models accounting for inventories fail to create endogenous incentives for firms to hold strictly positive levels of inventories even in steady-state, except for cases in which inventories are explicitly required to facilitate sales (Bils and Kahn, 2000), or are a material input into the production function (Iacoviello et al., 2011). In a significant number of industries, such as airplane production, this is certainly not the case.

Since the bullwhip effect is a reflection of the dynamic implications of inventory holdings for other participants of the same supply chain, it has the general attraction that it may help to re-establish the importance of inventories to aggregate macroeconomic fluctuations (see McMahon and Wanengkirtyo (2015) and Wang and Disney (2016) for a related discussion). The role played by such excessive shock propagation has potentially ebbed and flowed over time. For instance, it likely decreased in frequency and importance over time with the emergence of improved supply chain management methods, such as barcodes and RFID technologies that facilitate improved inventory management (McMahon and Wanengkirtyo, 2015). However, in the wake of the supply chain strains during and shortly after the COVID-19 pandemic, the bullwhip effect is believed to have gained significance and traction again.[3]

Literature review - inventories in macroeconomic research

While interest in inventory behaviour can be traced back to post-war times[4], the canonical inaugural model of inventory behaviour in Economics is the production smoothing model (PSM) of Holt et al. (1960). Holt et al. (1960) postulate a firm minimizing the cost of meeting its expected sales, yielding incentives for the firm to hold inventories in an attempt to smooth production efforts. Their model, however, does not seem to align with empirical regularities: as Blinder (1986) pointed out, the PSM fails to align with actual firm behaviour in that it predicts too little volatility in production, as well as a non-positive covariance between sales and inventories.[5]

Subsequent research focusing on the importance of inventories considered the $[S, s]$ model dating back to Arrow et al. (1951), by which firms restock the amount of their goods inventory at a lower bound s, with stock being replenished up to the point S (see Blinder et al. (1981) and Caplin (1985) for details). Closely related to this chapter are a number of contributions made in the 1980s, in particular Christiano and Eichenbaum (1989). As mentioned on pp. 7 of their manuscript, the at-the-time popular inventory dynamics models (such as Blinder (1986) and Eichenbaum (1984)) consolidated the various industry demand curves alongside a given supply chain into a single reduced-form relationship, thus ignoring the frictions occurring between participants of the same supply chain. We address this gap by showing the relevance of within-supply chain frictions for business cycle dynamics.

[3]While the causes of the current economic volatility are subject to heavy debate, supply chain strains frequently attain a prominent role (Antràs and Chor, 2022; Bai et al., 2024; Baqaee and Farhi, 2022; Benigno et al., 2022; Di Giovanni et al., 2022). Carreras-Valle (2026) shows that inventory dynamics have generally become more volatile in the period since the Great Moderation.

[4]Abramovitz (1950) points out the importance of inventory adjustment behaviour for business cycle propagation in the period leading up to the second World War.

[5]Strictly speaking, the PSM itself, too, has an antecedent - namely, the flexible accelerator model, inaugurally proposed by Metzler (1941). The model postulates strictly positive inventory holdings desired by firms, and aims to estimate the relationship between changes in inventories and the distance between inventories and their target level as well as the current-period forecast error in sales.

The subsequent class of stockout avoidance models, popularly associated with Kahn (1992), build on the class of [S,s] models, providing endogenous incentives for firms to retain a permanently positive level of inventories. Their model strikingly matches the empirical regularity of countercyclical mark-ups, which is further supported by Bils and Kahn (2000) and Kryvtsov and Midrigan (2013). A more complete and excellent overview on the burgeoning inventory literature in economics in the 20th century was provided by Ramey and West (1999). Our choice of modelling inventories as a net order book variable while keeping track of possibly unfilled commitments in the form of backorders is directly related to Zarnowitz (1961, 1973) and his characterisation of backlogs in the manufacturing industries.

McMahon (2014) proposed a full-fledged business cycle model, in which inventories take time to move from the producers in the economy to the households, but the temporal delivery lag can be augmented upon paying a 'fast delivery premium', whose marginal cost increases in the share of goods for which fast delivery is commanded. Inventories are thus required to avoid paying the 'fast delivery premium' upon positive changes to demand, although inventory holdings are, notably, subject to an iceberg cost diminishing the value of stored goods. While the model itself is intuitive, it fails to generate the required propagation to match business cycle dynamics following changes in aggregate demand. Sarte et al. (2015) leverage long-term changes in inventory behaviour to inform and estimate a business cycle model that can account for the shifts in cyclical behaviour of the US economy following the stagflationary episodes of the 1970s and 1980s, implementing furthermore a model in which a positive steady-state target for inventories arises in line with the time-to-build technology represented by Kydland and Prescott (1982). Carreras-Valle (2026) finds that the trend towards decreasing inventory holdings reversed since the onset of the Great Financial Crisis, consistent with a mechanism by which firms increase their buffer stock due to higher risk exposure in the face of more complex global supply chains.

The above approaches (with the exception of Sarte et al. (2015)), however, face one drawback: inventory holdings are introduced through mechanical assumptions that justify the need for inventories without dependence on the state of the production environment (such as the $[S, s]$ model assuming mechanical replenishing of inventory stock even under demand shifts, or the PSM postulating an inherent desire for firms to have perfectly smooth production even in the presence of temporary productivity innovations), or through use cases of inventories that are at odds with their empirical role (such as Iacoviello et al. (2011), who require inventories as a production input).

Empirically, the examination of the role of inventories in supply chains (and the related implications for the bullwhip effect) face the major difficulty of requiring explicit knowledge of the production structure of a good: notably, any empirical identification of shock propagation in a good's production environment cannot rely solely on firm aggregates, but instead requires a full-fledged mapping of the production network underlying that good, i.e., it is necessary to trace down all inputs used for a given output, as well as how that output is potentially used as an input for other goods, and so forth. Consequently, most evaluations of the bullwhip effect are limited to sector-specific case studies (Taylor, 1999; Ravichandran, 2008; Wang and Disney, 2016; Zhu et al., 2020). The empirical relevance of inventories for business cycle propagation is, however, very telling. McCarthy and Zakrajšek (2000) provide an account based off Compustat data, while McCarthy and Zakrajšek (2007) and McMahon and Wanengkirtyo (2015) attach significant roles to inventory management in explaining the Great Moderation.

Very much related to this chapter is the concurrently developed contribution by Leng et al. (2025), which addresses the creation of a bullwhip effect in a macroeconomic model of production networks. We view our project as complementary to their excellent characterization of the bullwhip effect as a cascading phenomenon in production networks. Unlike Leng et al. (2025), we place a specific focus on the behaviour of inventories, much in the tradition of the operations research literature that centres around the amplifying role that inventories can have for the bullwhip effect itself.[6] What we share is the important focus on the upstreamness of a given distortion within a production framework as the primary device through which the bullwhip effect takes shape.[7]

Finally, Ferrari (2025) is perhaps the most related to our present study, being also the first macroeconomic framework to account for and create the bullwhip effect as a supply chain-based amplification of demand shocks. The perhaps most important difference between the framework of Ferrari (2025) and ours is that inventories do not serve as a buffer stock to meet within-period demand, since demand can be satisfied generally via just-in-time production, whereas our framework explicitly tracks order-delivery mismatches over time without the imposition of per-period market clearing. That way, we give rise to richer adjustment dynamics in dependence on the quantities of inventories held at each stage of the production process.

In a nutshell, inventory behaviour has not mattered in macroeconomic modelling for decades not because of its irrelevance, but because previous approaches of modelling them were overly restrictive and devoid inventories of their actual role for production, rendering unrealistic and incomplete dynamics. The empirical relevance of inventories

[6] Leng et al. (2025) abstain mostly from modelling inventories, which they mostly view as offering a hedging role in production networks.

[7] Related contributions in the production networks literature that can lead to outcomes akin to the bullwhip effect are Bizzarri et al. (2025) and Schaal and Taschereau-Dumouchel (2025).

for the propagation of shocks, however, remains staggering and motivates the presented idea of incorporating inventories more directly in the production environments modelled by economists.

—

The chapter is structured as follows. Section 4.2 introduces some empirical evidence on the bullwhip effect in aggregate data. Section 4.3 introduces our new theoretical formulation of inventories. Section 4.4 lays out the multi-stage supply chain framework with inventories. In section 4.5, we characterize the economy's general equilibrium and describe the model-implied dynamics in section 4.6, characterizing the bullwhip effect explicitly. We look in greater detail at the role played by fiscal and monetary policy in our model in section 4.7. Finally, section 4.8 concludes.

4.2 Inventories in the Macroeconomy: Empirical Evidence

There are not many empirical studies on the bullwhip effect (Ferrari (2025) is a recent notable exception), and those that do exist are usually limited to sector-specific case studies. This is because explicitly pinning down the bullwhip effect would require *explicit* knowledge on all supply chains that a firm is connected to, identifying precise input measurements used for each good as well as the precise product prices charged. Such data is not readily available for the whole economy.

Nonetheless, even in aggregate data *without explicit information on orders*, a semblance of dynamics consistent with the bullwhip effect can be observed and measured. We present some stylized facts about the aggregate behaviour of sales and inventories on supply chains, following the classification into manufacturers, wholesalers, and retailers endorsed by the U.S. Census. We focus in particular on the relative volatility of sales and inventory-sales ratios on the supply chain, and ultimately uncover novel evidence on a possible link between relative inventory dynamics and downstream price pressure.

Figure 4.1 shows real sales by sector of the supply chain provided by the U.S. Census on a monthly basis. The aggregate data indicate an increase in the absolute size of sales levels for wholesalers and retailers, but less of a clear trend for manufacturers since the early 1990s, reflecting a growing internationalization of supply chains and a relative decline in the contribution of manufacturing to the US economy.

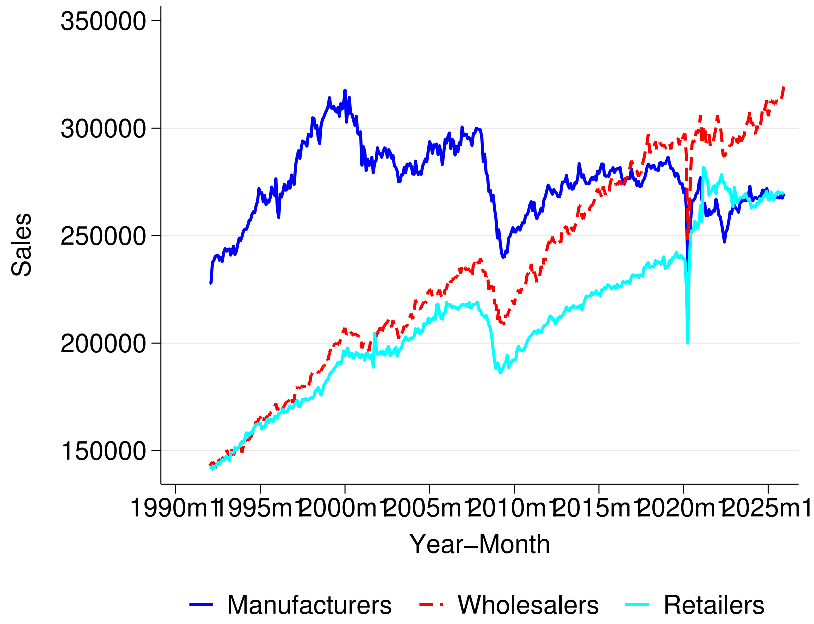

Figure 4.1: The overall evolution of monthly sales in the U.S. economy (in Millions of 1992 USD), deflated using the PPI for manufacturers and wholesalers, and the CPI for retailers.

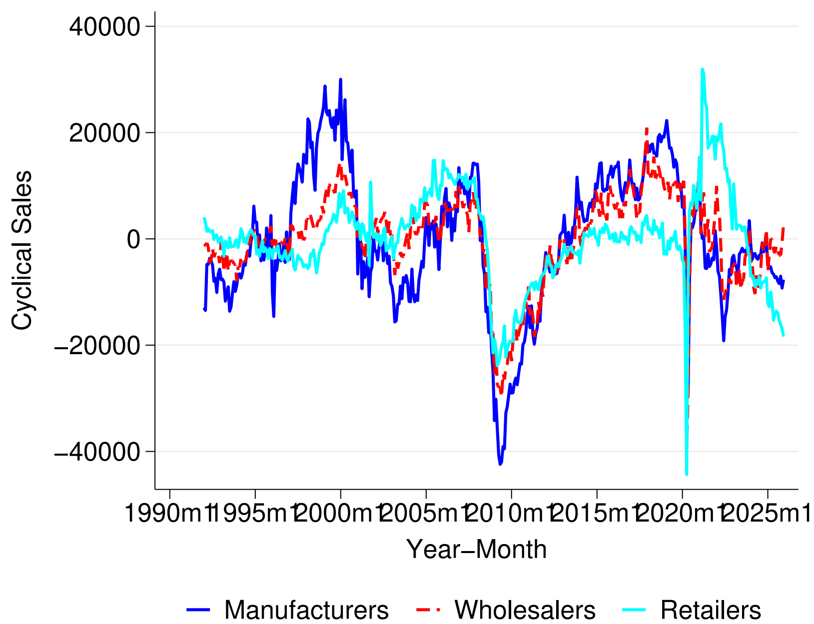

Figure 4.2: The cyclical component in sales at each supply chain stage over time in the U.S. economy, obtained as the difference between the raw sales level and a third-order polynomial trend.

Figure 4.2 plots the cyclical component of sales (derived as the deviation from a third-order polynomial trend in the time series) at each stage of the supply chain.[8] This figure highlights the generally larger cyclical volatility in sales upstream than downstream. This evidence is consistent with the bullwhip effects' prediction of growing volatility as we move up the supply chain from the final consumer. The notable exception is the Covid-induced downstream fluctuations where the cessation of many face-to-face activities led to a marked decline in retail sales revenue.

Other than supply chain stage-specific sales dynamics, our model places a significant weight on inventory dynamics. Following the literature on the macroeconomic effects of inventories, we focus on the *inventory-sales ratio*, which depicts the overall level of inventories relative to sales at a given stage of the supply chain. Figure 4.3 depicts the overall inventory-sales ratio at each supply chain stage since 1995. The early part of the sample is marked by declines in the ratio for each sector consistent with the idea of changes in supply-chain management that were often discussed as a potential source of the Great Moderation (McConnell and Perez-Quiros, 2000). However, this downward trend in the inventory-sales ratio did not persist beyond the Global Financial Crisis, as documented by Carreras-Valle (2026). The figure also shows a clear spike in the ratio in the early-Covid period, especially for retailers, as goods inflows largely continued, but there was temporarily severely depressed demand.

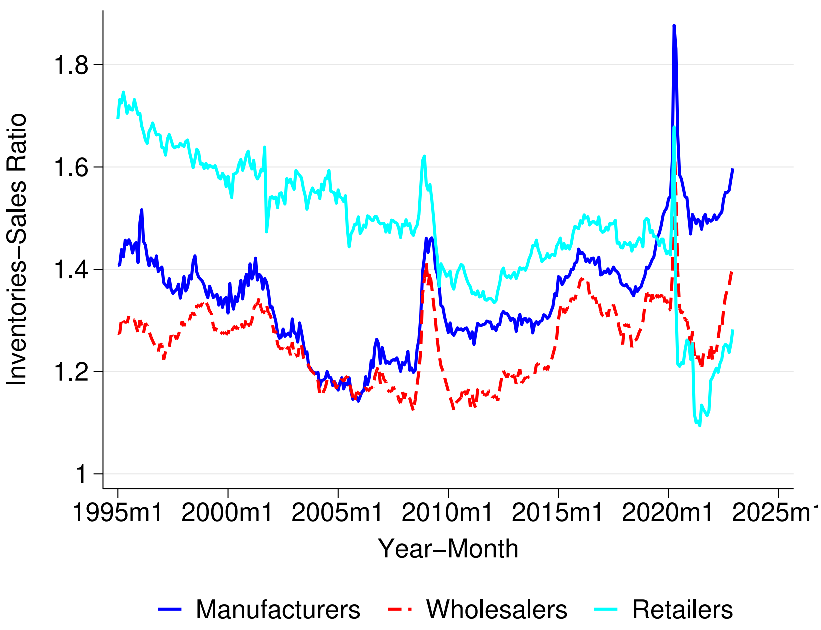

Figure 4.3: The overall evolution of the monthly inventory-sales ratio in the U.S. economy.

[8]Other time-series filtering techniques yield similar results.

Figure 4.4 again isolates the cyclical component of the inventory-sales ratio as the difference between the observed value and the third-order polynomial trend component. Relative to sales levels, the cyclical fluctuations are not clearly more volatile at one end of the supply chain relative to the other. While during the Great Moderation the cyclical variation appears larger at the upper end of the supply chain (i.e., for the manufacturers), we can observe larger cyclical fluctuations at the downstream end in terms of inventory-sales ratios since 2010.

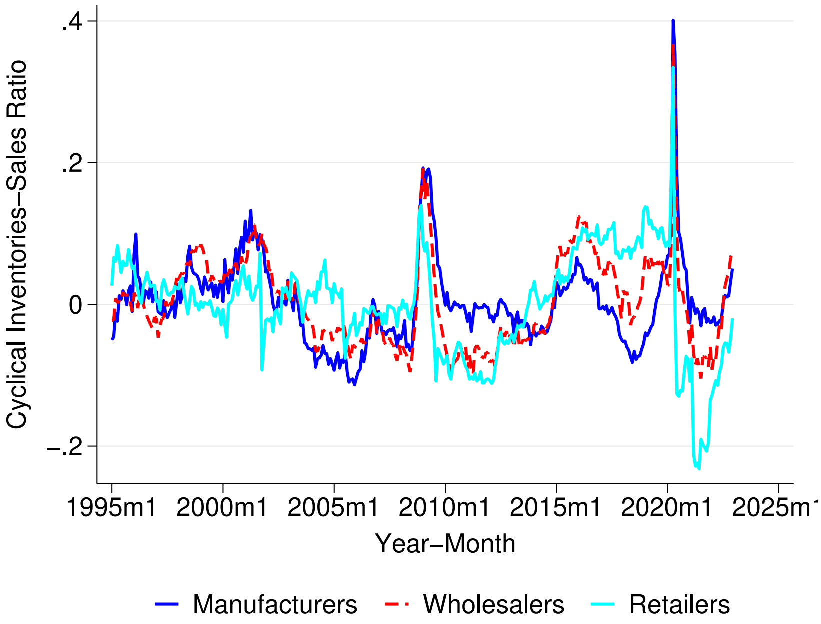

Figure 4.4: The cyclical component of monthly inventory-sales ratios at each supply chain stage in the U.S. economy.

Inventories and prices

To the extent that inventories, or backorders, reflect an imbalance of supply and demand, it is natural to think that a margin of adjustment would be the price mechanism in line with evidence on the relationship between supply chain strains and economic aggregates described by Benigno et al. (2022). The idea is that when a firm has a particularly large holding of inventories, it might reduce prices to sell greater amounts and reduce this. Likewise, a shortfall of supply, reflected in low inventory holdings, should be associated with increasing prices. Such a consideration of dynamic pricing is largely omitted from the discussions of the bullwhip effect in the operations research literature.

While a perhaps obvious line of enquiry, linking firm inventory and sales dynamics to prices is difficult because we don't have direct mapping to firm level data on these elements. To nonetheless provide some motivational evidence on a possible price pass-through, we provide a brief descriptive example of a possible pass-through from inventories to pricing behaviour. For this, we use the aggregate U.S. Census data, as presented above, and retail and producer prices indices.

Figure 4.5 presents the relationship between inventory-sales ratios and pricing indices at the top and the bottom of the supply chain. The manufacturing sector is presented on the left and the retail sector is on the right. In both figures, we use the cyclical variations in all variables – this is generated using the same third-order polynomial trend fitted on the data. These datapoints show a clear negative relationship for the manufacturing sector between the cyclical inventory-sales component and PPI inflation with a highly significant ($p < 0.01$) correlation coefficient of -0.46. The relationship in the retail sector is also negative and highly significant, albeit slightly smaller (-0.39). Therefore, there seems to be some evidence that the pricing margin is used as a corrective device at each supply chain stage, through which excess (or deficient) inventory holdings relative to sales are brought back towards their targeted steady-state level.

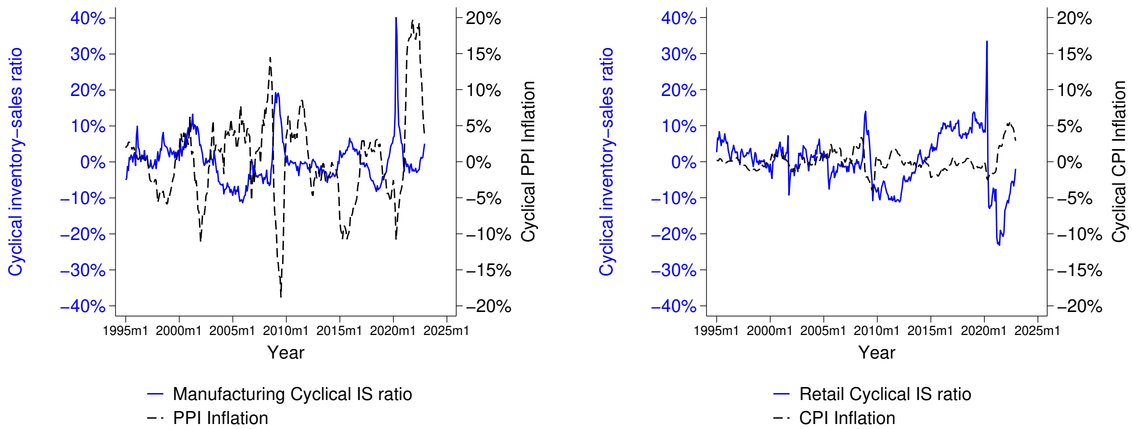

Figure 4.5: The cyclical component of the inventory-sales ratio and the locally relevant price index at the upper end of the supply chain (manufacturers, left panel) and at the lower end of the supply chain (retailers, right panel).

In our quest to establish a model rationalizing the bullwhip effect alongside rational firm pricing behaviour, we therefore take the following set of empirical regularities as a litmus test to the model: first, the model ought to generate upstream volatility from downstream shocks, including significant pricing movements both upstream and downstream. Second, since periods with larger relative upstream volatility in inventory-sales ratios *and* periods with larger relative downstream volatility exist, we want to see whether exogenous shocks can generate larger relative upstream/downstream movements. Third, we want to establish under which innovations we can recover a conditional negative relationship between inventory-sales ratios upstream and downstream.

4.3 Modelling Inventory Relationships on a Supply Chain

We now start building up our novel macroeconomic model of supply chains to explain the inventory and inventory-sales dynamics above. In this section, we first introduce the notation and conventions that describe how we model inventories. The key idea is to use a *net* order book variable, which facilitates analytical tractability and reduces the number of burdensome state variables that must be carried around.

Mapping Inventories, Backorders, and Sales

In a standard inventory model, a sale is either made or it is lost: if demand arrives and no stock is available, the transaction simply does not occur. We depart from this in one important respect. Firms in our model can accept orders that they are not immediately able to fill. These unfilled commitments — *backorders* — sit on the firm's order book as genuine liabilities, to be honoured in future periods (Zarnowitz, 1973). As a result, *received* orders O_t^R and *fulfilled* orders O_t^F need not coincide within a period. A firm may receive ten orders today and fulfil only six, carrying the remaining four as delivery obligations into the next period.[9] This distinction between received and fulfilled orders is absent from standard inventory models and is central to the mechanism we develop.

Given this distinction, the determination of fulfilled orders O_t^F follows directly. Sales in any period are bounded by what the firm can supply and by what buyers demand:

$$O_t^F = \min \left\{ \underbrace{S_t + y_t}_{\text{Supply of goods}} \;, \quad \underbrace{O_t^R + D_t^L}_{\text{Demand for goods}} \right\}, \tag{4.1}$$

S_t are goods in storage left unsold in earlier periods and y_t denotes newly produced output; together they constitute supply of goods available for sale. Demand is the sum of orders received this period *plus* the stock of orders received in prior periods that have not yet been delivered (D_t^L). It is precisely the presence of this second demand term — the carried-over backlog — that gives our state variable its structure, as we now explain.

A unified state variable. One of the challenges with solving models of inventories or backorders is that both are subject to non-negativity constraints ($S_t \geq 0$ and $D_t^L \geq 0$). However, a firm's inventory holdings and its backorder position are essentially two sides of the same coin. A firm holding ten units in storage but owing eight units of outstanding

[9]We rationalise the obligation to honour backorders as soon as possible as optimal firm behaviour by imposing an infinite reputational cost of failing to deliver when physically feasible.

deliveries is economically close to a firm with two units in storage and no backlog. We exploit this symmetry by tracking both in a single net variable:

$$H_t \equiv \underbrace{S_t}_{\substack{\text{goods in}\\ \text{storage}}} - \underbrace{D_t^L}_{\substack{\text{unfilled orders}\\ \text{outstanding}}},$$

The variable H_t is signed, but both sides of zero have clear economic content, as summarised in Table 4.1. As well as the analytical convenience of avoiding the pair of non-negativity constraints, H_t is Markovian — the entire history of inventory accumulation and backorder obligations is summarised in a single state.

	Inventory position	Backorder position
$H_t > 0$	Buffer stock is available for sale.	No backorders outstanding.
$H_t = 0$	No inventories are held.	No backorders outstanding.
$H_t < 0$	No inventories are held.	Unfilled orders are outstanding.

Table 4.1: Economic interpretation of the net inventory/order book variable H_t.

The law of motion. The evolution of H_t follows a natural balance-sheet identity. At the end of any period, the net position equals the beginning-of-period stock, plus newly produced goods, minus the orders received during the period:

$$H_{t+1} = H_t + y_t - O_t^R. \tag{4.2}$$

Note carefully that it is *received* orders O_t^R, not *fulfilled* orders O_t^F, that reduce H_t. The reason is that H_t tracks what a firm is *capable* of delivering, net of the commitments it has already made. The moment an order is received, the firm has incurred a delivery obligation regardless of whether it has yet been filled. Fulfilled orders, by contrast, determine revenues and appear in the profit function, but do not separately enter the state variable.

Figure 4.6 illustrates this with a numerical example. Starting from $H_{t-1} = 50$, the firm produces $y_{t-1} = 30$ units and receives $O_{t-1}^R = 60$ orders; the resulting net position at the start of the next period is $H_t = 50 + 30 - 60 = 20$.

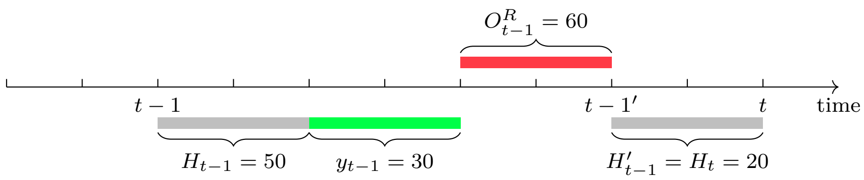

Figure 4.6: Evolution of the net inventory/order book variable.

Connecting back to fulfilled orders. We can now write equation (4.1) precisely in terms of H_t. When $H_t > 0$, the firm holds H_t units of buffer stock and faces no inherited backlog; when $H_t < 0$, the firm holds no stock and carries $|H_t|$ units of unfilled orders. The supply and demand constraints therefore become:

$$\begin{aligned} \text{Supply:} \quad & O_t^F \leq \max\{0, H_t\} + y_t, \\ \text{Demand:} \quad & O_t^F \leq O_t^R + \max\{0, -H_t\}, \end{aligned}$$

and combining these, fulfilled orders are:

$$O_t^F = \min\left\{ \underbrace{\max\{0, H_t\} + y_t}_{\text{Supply of goods}}, \ \underbrace{O_t^R + \max\{0, -H_t\}}_{\text{Demand for goods}} \right\}. \tag{4.3}$$

The max operators simply recover the intuition above: $\max\{0, H_t\}$ extracts available inventory when $H_t > 0$ and returns zero otherwise; $\max\{0, -H_t\}$ extracts the inherited backlog when $H_t < 0$ and returns zero otherwise. Equation (4.3) thus unifies both regimes in a single expression.

Supply Chain Structure

We model a simple supply chain comprised of manufacturers who sell to wholesalers who then sell to retailers who provide goods to the final consumer.[10] We denote each stage of the supply chain by its distance from the final consumer. So final retailers are denoted "1", wholesalers are "2" and manufacturers are "3". The good travels from manufacturer to household through the chain; orders travel in the reverse direction. Figure 4.7 illustrates the structure.

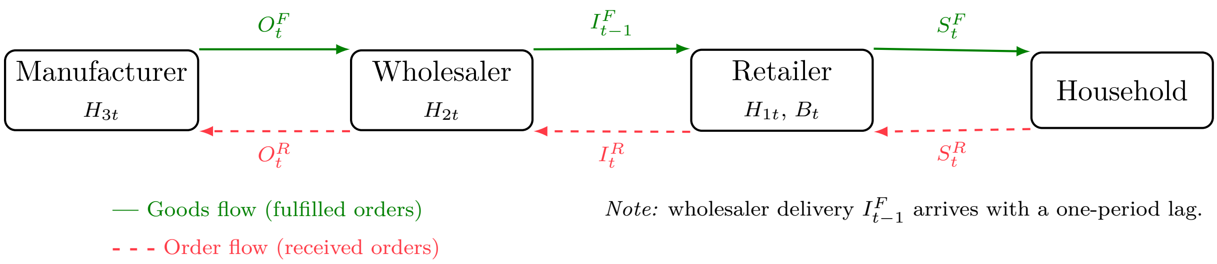

Figure 4.7: Supply chain structure. Green solid arrows indicate the flow of fulfilled orders (goods); red dashed arrows indicate the reverse flow of received orders. State variables H_{jt} denote net inventories at each stage; B_t is the retailer's backroom stock.

Several features of this structure deserve emphasis.

[10] This three-stage decomposition maps directly onto the U.S. Census definitions of "Manufacturing & Trade Inventories & Sales".

Goods flow downstream with a lag at the wholesale stage. Deliveries from the manufacturer to the wholesaler (O_t^F) are available to the wholesaler for onward delivery only in the following period (in the retailer's problem, the same applies to the variable I_t^F). This timing assumption reflects the empirical regularity that wholesale distribution takes time, and it is what prevents contemporaneous upstream productivity shocks from instantaneously affecting downstream delivery capacities.

Orders flow upstream. At each stage, the downstream agent places orders (O_t^R, I_t^R, S_t^R) with its immediate upstream supplier. Crucially, these received orders need not equal fulfilled orders within the same period: a supplier may fulfil an order immediately from inventory, partially fulfil it, or add it to the backorder book. The net inventory variable H_{jt} at each stage tracks this balance, taking positive values when buffer stock is available and negative values when a backlog of unfilled orders has accumulated.

Notation summary. Subscripts $j \in \{1, 2, 3\}$ index the supply chain stage. Superscripts F and R denote fulfilled and received orders respectively. The net inventory variable H_{jt} at each stage follows the law of motion $H_{j,t+1} = H_{jt} + (inflow) - X_{jt}^R$ as in equation (4.2), where X_{jt}^R are the received orders at that stage of the supply chain. The $inflow$ is new production for the manufacturer and deliveries received for the other supply-chain stages. Table 4.2 collects the main flow variables.

Stage	**Agent**	**Fulfilled orders**	**Received orders**
3	Manufacturer	O_t^F	O_t^R (from wholesaler)
2	Wholesaler	I_t^F	I_t^R (from retailer)
1	Retailer	S_t^F	S_t^R (from household)

Table 4.2: Flow variable notation by supply chain stage.

4.4 Our Model of Distributional Optimisation

Our model departs from the standard inventory framework in two related ways. First, firms in each stage of the supply chain are dynamically linked to their neighbours *before* their own idiosyncratic productivity shock is realised. This timing means that firms cannot condition their labour, pricing, and production decisions on the contemporaneous value of their productivity draw; instead, they must optimise over the full distribution of possible realisations. As a result, firms do not simply maximise expected profits — they maximise over the distribution of the profit function, following Kryvtsov and Midrigan

(2013), whose intermediate firm problem we generalise to a multi-stage supply chain network. We call this *distributional optimisation*.[11]

Second, and as a direct consequence, firms hold precautionary inventories even though these are not technologically required for production or sales. A firm that must commit to a production plan before observing its productivity will find it optimal to build a buffer against the possibility that productivity is low and demand cannot be met — just as it will want to avoid excessive overproduction if productivity is high. The asymmetric net inventory cost function $C(H)$ we introduce below formalises this trade-off. Precautionary inventory holding therefore emerges endogenously from the interaction of the supply chain timing structure and the distributional optimisation problem; it is not imposed by assumption.

Section 4.4 explains the distributional optimisation technique and the partial expectation function Ψ that delivers closed-form solutions. Sections 4.4–4.4 then present the economic substance of each firm's problem: the value function, the first-order conditions, and their interpretation. Full derivations — including the integration steps that produce the revenue functions R^3, R^2, R^1 and all envelope conditions — are collected in Appendix D.1.

Distributional Optimisation: The Method

Each firm's revenue in a given period is determined by a $\min(\cdot)$ function that picks the lesser of supply and demand — as established in Section 4.3. Because idiosyncratic productivity a_{jt} is not observed before production decisions are made (see assumption 3 below), expected revenue requires integrating this $\min(\cdot)$ function over the distribution of a_{jt}. This integral is the central computational object of each firm's problem, and its tractability is the methodological payoff of the log-normality assumption.

Concretely, for any firm j the expected revenue function $R^j(\cdot)$ takes the form:

$$R^j(\cdot) = \int_{\xi} \min\Big\{\text{Supply}(a_{jt}, \cdot),\ \text{Demand}(\cdot)\Big\}\, dF^{a_j}(a_{jt}).$$

Because a_{jt} enters only the supply side and is log-normally distributed, there exists a threshold $\bar{a}_{jt}$ below which supply is the binding constraint and above which demand binds. This threshold is a function of the firm's choices and the state variables, and it allows the integral to be split cleanly into two regions. The result involves the partial expectation function

$$\Psi^{a_j}(x) \equiv \int_0^x a\, f^{a_j}(a)\, da = \exp\Big(\frac{\sigma^2_{a_j}}{2}\Big)\, F^{a_j}\Big[x \exp\Big(-\sigma^2_{a_j}\Big)\Big], \tag{4.4}$$

[11] Notice that this does not require risk aversion or other non-linearities in the profit function. Firms are risk-neutral; it is the *timing* of information relative to decisions that generates precautionary behaviour.

where the closed-form expression on the right exploits the log-normal structure. Intuitively, $\Psi^{a_j}(x)$ captures the expected value of productivity *conditional on productivity being below the threshold* x, weighted by the probability of being in that region. It is precisely this function that replaces the $\min(\cdot)$ integral in each firm's first-order conditions, delivering a tractable system of equations.

The resulting revenue functions R^3, R^2, and R^1 are derived in full in Appendix D.1. In the main text, we take them as given and focus on the economic content of the optimality conditions.

Two aspects of the solution are worth highlighting before proceeding. First, the Ψ function appears in the labour demand conditions of every firm in the chain, replacing the standard marginal product of labour. The intuition is that a firm hiring an additional worker raises expected sales only in the states of the world where it is supply-constrained — the probability of which is endogenous to the firm's inventory position and order book. Second, the threshold $\bar{a}_{jt}$ links the firm's inventory position directly to its production incentives: a firm with a large backlog (low H_{jt}) faces a higher probability of being supply-constrained, raising the marginal value of labour and of precautionary inventory accumulation.

The Manufacturer

The manufacturer produces output according to

$$y_t = a_{3t}\, n_{3t}^{\gamma}, \tag{4.5}$$

where $a_{3t} \sim \mathcal{LN}(0, \sigma_{a_3}^2)$ is an idiosyncratic productivity shock, and n_{3t} is labour. The manufacturer chooses labour and its next-period net inventory position $H_{3,t+1}$.

Value function. In recursive form, dropping time subscripts and using primes for next-period values, the manufacturer's problem is:

$$\begin{aligned} V^3(H_3; \vec{X}) = \max_{n_3,\, H_3'} \;\; & p_3\, R^3(H_3, n_3) - n_3\, W(\xi) - C(H_3) \\ & + \beta\, \mathbb{E} \int_{\xi'} V^3(H_3';\, \vec{X}'(\xi'))\, dF(\xi'), \end{aligned} \tag{4.6}$$

subject to the inventory law of motion $H_{3,t+1} = H_{3t} + a_{3t} n_{3t}^{\gamma} - O_t^R$, and ξ is a vector of within-period uncertain exogenous innovations.[12] The inventory holding cost $C(H_3)$ is convex but *asymmetric*, following Varian (1975):

[12]Strictly speaking, the manufacturer chooses *expected* net inventories, since $H_{3,t+1}$ is subject to the uncertainty induced by stochastic productivity a_{3t}. That is, the manufacturer imposes its inventory constraint only in expectation, and does not have full control over $H_{3,t+1}$, but rather over $\mathbb{E}_t H_{3,t+1}$. For simplicity of notation, we simply depict $H_{3,t+1}$ as the manufacturer's choice variable, and we keep that convention for the other supply chain participants as well.

$$C(H_{3t}) \equiv \frac{\exp(\chi_H \omega H_{3t}) - \chi_H \omega H_{3t} - 1}{\omega^2}. \tag{4.7}$$

For $\omega < 0$, costs are asymptotically linear on the positive domain (overstocked warehouses) and exponential on the negative domain (backlogged orders). This asymmetry gives rise to precautionary inventory holdings, as delayed deliveries are more costly than temporary overstocking.

Optimality conditions. The first-order conditions for labour and net inventories are derived in Appendix D.1. Their economic content can be stated compactly. The labour demand condition equates the marginal revenue product of labour — where the marginal product is weighted by Ψ^{a_3}, reflecting the probability of being supply-constrained — to the nominal wage:

$$p_{3t}\,\gamma\, n_{3t}^{\gamma-1}\Psi^{a_3}\left[\frac{O_t^R - H_{3t}}{n_{3t}^{\gamma}}\right] = W_t - \lambda_t^3\,\gamma\exp\left(\frac{\sigma_{a_3}^2}{2}\right) n_{3t}^{\gamma-1}. \tag{4.8}$$

The shadow value λ_t^3 captures the option value of having an additional unit of production available: more output today relaxes tomorrow's inventory constraint, which is valuable when a backlog is likely. The inventory Euler equation equates the current shadow cost of committing to a higher $H_{3,t+1}$ to the discounted expected gain, which depends on whether inventories or backlogs are more probable next period, see appendix D.1 for details.

The Wholesaler

The wholesaler does not produce goods but uses labour to forward them from the manufacturer to the retailer. This creates two important differences from the manufacturer's problem.

First, the **supply of goods available for forwarding** is constrained by two independent factors: the number of goods received from the manufacturer and available from inventory, *and* the labour capacity available for forwarding. These enter the sales function as separate constraints within the min(·):

$$I_t^F = \min\left\{\underbrace{\max\{0, H_{2t}\} + O_{t-1}^F}_{\text{Goods available}},\quad \underbrace{a_{2t}\, n_{2t}^{\gamma}}_{\text{Labour capacity}},\quad \underbrace{I_t^R + \max\{0, -H_{2t}\}}_{\text{Demand}}\right\}. \tag{4.9}$$

Note the one-period lag: goods delivered by the manufacturer in period $t-1$ (O_{t-1}^F) are available for forwarding in period t. This timing prevents upstream productivity shocks from propagating instantaneously downstream.

The next-period net inventory position $H_{2,t+1}$ is effectively determined by the wholesaler's labour choice alone, not by an independent inventory decision, or by prices, which are only indirectly influenced by the wholesaler in line with the bargaining problem we lay out below. The wholesaler's problem therefore reduces to choosing n_{2t}, $H_{2,t+1}$, and its own order quantity O_t^R to maximise:

$$\begin{aligned}\mathbb{V}_t^2(H_{2t}, O_{t-1}^R, a_{2t}) = \max_{n_{2t},\, H_{2,t+1},\, O_t^R} \; & p_{2t}\, R^2(H_{2,t+1}, H_{2t}, n_{2t}, O_{t-1}^F) - n_{2t}\, W_t - C(H_{2t}) \\ & - p_{3t}\, O_t^F(O_t^R, O_{t-1}^R) + \beta\, \mathbb{V}_{t+1}^2(H_{2,t+1}, O_t^R, a_{2,t+1}).\end{aligned} \tag{4.10}$$

The revenue function $R^2(\cdot)$ is derived in Appendix D.1 using the same partial-expectation technique as for the manufacturer, exploiting the nested min structure of (4.9).

Optimality conditions. The labour demand condition has the same structure as the manufacturer's, with Ψ^{a_2} replacing Ψ^{a_3} and goods available from inventory replacing labour as the scale factor. The order placement condition — the wholesaler's choice of O_t^R — equates the marginal cost of an additional order this period (the price p_{3t}, adjusted for the probability of immediate delivery) to the discounted expected benefit of having that unit available next period for further sale:

$$p_{3t}\, \frac{\partial O_t^F}{\partial O_t^R} = \beta\, \mathbb{E}_t\, \partial_{O^R} \mathbb{V}_{t+1}^2(H_{2,t+1}, O_t^R, a_{2,t+1}). \tag{4.11}$$

The derivative $\partial O_t^F / \partial O_t^R$ is an indicator function: an additional order translates one-for-one into a delivery today if and only if the manufacturer has sufficient inventory and production capacity. Otherwise, it enters the backlog and increases the probability of delivery tomorrow. The full expressions are given in Appendix D.1.

The Retailer

The retailer introduces a further innovation absent from the manufacturer and wholesaler problems: a distinction between backroom inventories B_t and shelf inventories H_{1t}, reflecting the empirical regularity that goods must be stocked and processed before being available for sale (Eroglu et al., 2013; Mou et al., 2018). Goods delivered by the wholesaler (I_{t-1}^F) arrive in the backroom; the retailer uses labour to move them to the shelves, where the capacity for shelf-stocking is:

$$U_t = a_{1t}\, n_{1t}^{\gamma}\, (B_t + I_{t-1}^F), \tag{4.12}$$

where $B_t = B_{t-1} + I_{t-2}^F - U_{t-1}$ tracks the backroom stock. Realized sales are then:

$$S_t^F = \min\left\{ \underbrace{S_t^R + \max\{0, -H_{1t}\}}_{\text{Demand}},\ \underbrace{U_t + \max\{0, H_{1t}\}}_{\text{Supply}} \right\}. \tag{4.13}$$

The retailer therefore has two state variables governing its inventory position: H_{1t} (net shelf inventory) and B_t (backroom stock). This gives rise to two shadow values in the optimality conditions — one for each inventory constraint — and makes the retailer's Euler equations richer than those at the upstream stages. The two inventory laws of motion are:

$$H_{1,t+1} = H_{1t} + a_{1t}\, n_{1t}^{\gamma}\,(B_t + I_{t-1}^F) - S_t^R, \tag{4.14}$$

$$B_{t+1} = \left(1 - a_{1t}\, n_{1t}^{\gamma}\right)(B_t + I_{t-1}^F). \tag{4.15}$$

Equation (4.15) makes the role of labour transparent: a higher labour input n_{1t} depletes the backroom faster (by moving goods to shelves), so the backroom stock carried into the next period is smaller. The retailer's value function and full set of first-order conditions follow the same distributional optimisation logic as the upstream stages; they are stated in full in Appendix D.1.

Households and Policy

Households

A key departure from standard DSGE frameworks is that the household derives utility from *delivered* consumption S_t^F, not from orders placed S_t^R. Because deliveries may be delayed when the retailer's shelf inventory is insufficient, the household must form rational expectations over the likely delay between order placement and fulfilment. The household's problem is therefore forward-looking in a non-standard way: an order placed today may not be delivered today, depending on the retailer's net inventory state.

The representative household solves:

$$\max_{\{S_t^R, n_t, d_{t+1}\}_{t=0}^{\infty}} \mathbb{E}_0 \left\{ \sum_{t=0}^{\infty} \beta^t \left[\frac{(S_t^F)^{1-\sigma} - 1}{1-\sigma} - \zeta_t \frac{n_t^{1+\frac{1}{\varphi}}}{1+\frac{1}{\varphi}} \right] \right\} \tag{4.16}$$

subject to the budget constraint

$$p_{1t}\, S_t^R + d_{t+1} = w_t\, p_{1t}\, n_t + (1 + i_{t-1})\, d_t - p_{1t}\, t_t, \tag{4.17}$$

and the delivery constraint

$$S_{t+1}^F(S_{t+1}^R, S_t^R) = \min\left\{ S_{t+1}^R + \max\{0, S_t^R - H_{1t} - U_t\},\ U_{t+1} + \max\{0, H_{1t} + U_t - S_t^R\} \right\}. \tag{4.18}$$

The delivery constraint says that consumption next period equals next period's orders plus any undelivered orders from this period, subject to the retailer's shelf availability. Only the most recent period's unfilled orders enter, since the net inventory variable is Markovian. The optimality conditions yield the standard labour supply condition and a bond Euler equation:

$$\lambda_t^{HH} = \beta\, \mathbb{E}\Big[\lambda_{t+1}^{HH}(1+i_t)\Big]\,. \tag{4.19}$$

A more detailed description of the household optimality conditions is provided in appendix D.1.

Government

The government issues one-period nominal bonds d_t, levies lump-sum taxes t_t, and makes transfers g_t. Its budget constraint is:

$$p_{1t}\, t_t + d_{t+1} = (1+i_{t-1})\, d_t + p_{1t}\, g_t. \tag{4.20}$$

Monetary policy sets the nominal interest rate via a Taylor rule responding to retail inflation $\Pi_{1t} \equiv p_{1t}/p_{1,t-1}$ and output:

$$1+i_t = (1+\bar{\imath})^{1-\rho_m}(1+i_{t-1})^{\rho_m}\left[\Pi_{1t}^{\phi_\pi}\left(\frac{S_t^F}{\bar{S}^F}\right)^{\phi_y}\right]^{1-\rho_m}. \tag{4.21}$$

Fiscal policy stabilises the debt-to-steady-state ratio through a simple feedback rule:

$$\frac{t_t}{\bar{t}} = \left(\frac{d_{t-1}}{\bar{d}}\right)^{\psi_d}\frac{g_t}{\bar{g}}. \tag{4.22}$$

Price determination through Nash bargaining

To close the model, we must impose some form of price determination, which is cumbersome given that each of the three markets in the model (between two successive supply chain stages) does not need to clear exactly, with the net inventory/order book variable being able to absorb slack. We therefore posit that all three goods prices are determined through a mechanism leaning onto bilateral Nash bargaining over the common surplus from trade on the supply chain.

Embedding such bilateral bargaining as a market clearing mechanism is computationally burdensome. We therefore adopt the simplifying convention that the bargaining is settled *prior* to shocks on the supply chain: prices $\{p_{1t}, p_{2t}, p_{3t}\}$ are determined conditional on the aggregate state, but prior to the realization of firm-specific productivity innovations

(a_{1t}, a_{2t}, a_{3t}). Firms and households take the bargained price as parametric when solving their optimisation problems and remain naive over the distribution of productivity innovations of the other firms on the supply chain.[13]

To find the solution of such bargaining situations, we must clarify the outside option when bargaining fails. For each party i involved in a bilateral trade with counterparty j, the outside option $\bar{V}_t^{i,j}$ denotes the value of failing to trade with j in period t, which we claim sets the relevant within-period trade flow to zero. In that case, each supply chain participant continues with their inherited state variables and conducts other trades at their own bargained prices. Since the wholesaler and the retailer each participate in two bilateral relationships, they have separate upstream (↑) and downstream (↓) outside options.

For the manufacturer, the outside option corresponds to producing but not shipping to the wholesaler (i.e., $O_t^F = 0$):

$$\bar{V}_t^3 = \max_{n_3, H_3'} \left\{ -n_3 W_t - C(H_{3t}) + \beta \mathbb{E}_t V^3(H_{3,t+1}; \vec{X}_{t+1}) \right\}, \tag{4.23}$$

subject to $H_{3,t+1} = H_{3t} + a_{3t} n_3^{\gamma}$.

For the wholesaler, the two outside options are:

$$\bar{V}_t^{2,\uparrow} = \max_{n_2, H_{2,t+1}} \left\{ p_{2t} R^2(H_{2,t+1}|_{O_t^F=0}, H_{2t}, n_2, 0) - n_2 W_t - C(H_{2t}) + \beta \mathbb{E}_t \mathbb{V}_{t+1}^2 \right\}, \tag{4.24}$$

with $O_t^F = 0$ (the wholesaler still fulfils downstream orders from inherited inventories only) and

$$\bar{V}_t^{2,\downarrow} = \max_{n_2, H_{2,t+1}, O_t^R} \left\{ -p_{3t} \mathbb{E}_t[O_t^F] - n_2 W_t - C(H_{2t}) + \beta \mathbb{E}_t \mathbb{V}_{t+1}^2 \right\}, \tag{4.25}$$

with $I_t^R = 0$ (the wholesaler still procures upstream but makes no sales to the retailer). In both cases, the inventory law of motion $H_{2,t+1} = H_{2t} + O_{t-1}^F - I_t^R$ holds with the appropriate flow set to zero.

Analogously, for the retailer:

$$\bar{V}_t^{1,\uparrow} = \max_{n_1, H_{1,t+1}, B_{t+1}} \left\{ p_{1t} R^1(H_{1t}, n_1, B_t) \Big|_{I_t^F=0} - n_1 W_t - C(H_{1t}) - C(B_t) + \beta \mathbb{E}_t V_{t+1}^1 \right\}, \tag{4.26}$$

with $I_t^R = 0$, and

[13] This timing preserves assumption 3 and keeps all revenue integrals R^1, R^2, R^3 one-dimensional in the firm's own shock. Were prices to be bargained after idiosyncratic shocks realized, the decomposition would be at odds with the simplification of the revenue functions.

$$\bar{V}_t^{1,\downarrow} = \max_{n_1, H_{1,t+1}, B_{t+1}, I_t^R} \left\{ -p_{2t}\mathbb{E}_t[I_t^F] - n_1 W_t - C(H_{1t}) - C(B_t) + \beta \mathbb{E}_t V_{t+1}^1 \right\}, \tag{4.27}$$

with $S_t^R = 0$, each subject to the inventory laws (4.14) and (4.15).

For the household, the outside option corresponds to placing no orders but continuing to supply labour and choose bond holdings:

$$\bar{V}_t^{HH} = \max_{n_t, d_{t+1}} \left\{ u(S_t^F)\Big|_{S_t^R=0} - \zeta_t v(n_t) + \beta \mathbb{E}_t \mathbb{V}_{t+1}^{HH} \right\}, \tag{4.28}$$

subject to the household budget constraint (4.17) with $S_t^R = 0$. Note that $S_t^F|_{S_t^R=0}$ is not necessarily zero, as the household may still receive deliveries on backlogged past orders whenever $H_{1,t-1} < 0$.

With that information, we are ready to pin down the bargaining solutions. Let $\eta_3, \eta_2, \eta_1 \in [0,1]$ denote the bargaining weights of, respectively, the wholesaler (in the upstream bargain with the manufacturer), the retailer (in the bargain with the wholesaler), and the household (in the bargain with the retailer). The three sharing conditions, derived from maximization of the Nash-bargaining style products are given by:

$$\eta_3 \left[\mathbb{V}^3(H_{3t}) - \bar{V}_t^3 \right] = (1-\eta_3) \left[\mathbb{V}_t^2(H_{2t}, O_{t-1}^R) - \bar{V}_t^{2,\uparrow} \right], \tag{4.29}$$

$$\eta_2 \left[\mathbb{V}_t^2(H_{2t}, O_{t-1}^R) - \bar{V}_t^{2,\downarrow} \right] = (1-\eta_2) \left[\mathbb{V}^1(H_{1t}, B_t) - \bar{V}_t^{1,\uparrow} \right], \tag{4.30}$$

$$\eta_1 \left[\mathbb{V}^1(H_{1t}, B_t) - \bar{V}_t^{1,\downarrow} \right] = (1-\eta_1) \frac{1}{\lambda_t^{HH}} \left[\mathbb{V}_t^{HH}(S_{t-1}^R, d_t, \xi_t) - \bar{V}_t^{HH} \right], \tag{4.31}$$

where the factor $1/\lambda_t^{HH}$ in (4.31) converts the household's utility surplus into period-t monetary units, following the standard convention for Nash bargaining in which one party's payoff is measured in utility units.[14]

4.5 General Equilibrium

Market clearing and definition of equilibrium

Labour market: The labour market simply clears in every period by equating labour supply and total labour demand, with the wage acting as the market-clearing price of labour.

$$n_t = n_{1t} + n_{2t} + n_{3t}. \tag{4.32}$$

[14] Without this rescaling, (4.31) would be dimensionally inconsistent, and the weight η_1 would not be directly comparable in magnitude to η_2, η_3.

Inventory markets: Inventory accumulation is determined by firms' dynamic inventory choices rather than by a separate market-clearing condition. Firms trade off the expected benefit of inventories, through looser current and future delivery constraints, against inventory holding costs.

Goods markets: Goods markets do not have simple clearing conditions, as prevailing market prices need not induce the demand for a good to be fulfilled, or its supply to be taken up. We close the model by positing that each price is the outcome of a bilateral Nash bargain between the two parties on either side of the corresponding exchange, as specified in section 4.4. The three conditions (4.29)-(4.31) are solved jointly as a fixed-point system in (p_{1t}, p_{2t}, p_{3t}) each period, together with the six outside-option value functions $\{\bar{V}_t^3, \bar{V}_t^{2,\uparrow}, \bar{V}_t^{2,\downarrow}, \bar{V}_t^{1,\uparrow}, \bar{V}_t^{1,\downarrow}, \bar{V}_t^{HH}\}$ defined in (4.23)-(4.28). We formally define the equilibrium with Nash-bargained goods prices of this economy in appendix D.1.

Parametrization and computational solution of the model

Param.	Description	Value
Inventory cost		
χ_B	Backroom inventory cost	20.0
χ_H	Frontroom inventory cost	0.10
ω	Shape parameter on inventory cost	-2.0
Aggregate shock processes on supply chain		
ρ_{a_1}	Persistence of retailer productivity shocks	$\{0.75, 0.25\}$
ρ_{a_2}	Persistence of wholesaler productivity shocks	$\{0.75, 0.25\}$
ρ_{a_3}	Persistence of manufacturer productivity shocks	$\{0.75, 0.25\}$
ρ_ζ	Persistence of labour-disutility shocks	$\{0.75, 0.25\}$
$\sigma_{\varepsilon_{a_1}}$	Std. dev. of retailer productivity shock	0.088
$\sigma_{\varepsilon_{a_2}}$	Std. dev. of wholesaler productivity shock	0.088
$\sigma_{\varepsilon_{a_3}}$	Std. dev. of manufacturer productivity shock	0.088
$\sigma_{\varepsilon_\zeta}$	Std. dev. of labour-disutility shock	0.088

Table 4.3: Inventory and shock parameters of the model economy

Table 4.3 describes the parameters related to the inventory cost functions and the shock processes in this economy.[15] As for the inventory cost function, we assume that the slope of the inventory cost is principally steeper for the backroom than for the frontroom. The reason for this is the shelf-filling technology of the retailer described by equation (4.12), which yields the law of motion for backroom inventories (4.15): simply put, B

[15]The remaining parametrization of the model, which is broadly standard, can be found in appendix D.3. One caveat must be mentioned: we postulate that $\gamma = 1$, meaning that each supply chain stage-specific production technology is linear in labour. However, we cannot deduce that the production technology is linear in labour in aggregate, given the existing delivery lags and the delivery structure for the wholesaler and the retailer, since additional labour supplied to the wholesaler might not provide *any* additional forwarding of the good, and since the shelf-stocking function of the retailer (4.12) exhibits dependence on the total size of available goods.

cannot be positive in steady-state, but it can be negative, which would be counterfactual. To avoid this, we penalize large values of inventories away from zero through a large value of χ_B, effectively rendering $B \approx 0$ in steady-state. The frontroom inventory cost parameter χ_H implies moderate costs of holding inventories relative to output. The value for ω is *negative*, implying that net inventory costs rise exponentially when negative, which induces precautionary positive steady-state inventory holdings at both the top and the bottom of the supply chain.

We assume that the standard deviation of each supply chain-specific exogenous innovation is fairly large, at 0.088, since $\exp(0.088^2/2) \approx 1$. In our baseline parametrization, the shock persistence is around 0.75, and we assume one period to be equal to one month.

Given the still relatively large number of state variables *and* the fact that the model insights depend on analysing in particular non-linear behaviour, solving this model is computationally challenging. We therefore need to devote significant resources to obtaining an approximately accurate solution.

We first calculate the steady-state twice; once using a brute-force value function iteration of individual equilibrium conditions over a possible set of market prices, and once using a high-dimensional non-linear solver (Kanzow et al., 2004). The resulting errors are of a magnitude of 10^{-6} in a weighted second-order norm.

The steady-state is then used as an input in calculating an initial guess of the dynamic policy functions. The calculation of the policy functions uses a multi-layered neural network, which preserves the rich non-linearities in the projection-based approximations of the true policy functions. Details on this procedure are provided in appendix D.2.

4.6 Results: The Bullwhip Effect in General Equilibrium

We now explore the dynamics of our novel macroeconomic model of inventories, focussing on how various shocks are consistent with our empirical narrative of time-varying relative volatility in sales and inventory-sales ratios upstream vs. downstream, and the possible negative link between inventory-sales movements and pricing decisions. We consider three different sources of macroeconomic disturbances: (i) upstream productivity shocks (a_3), (ii) downstream productivity shocks (a_1), and (iii) fiscal transfers (g), and three adjustment margins: (i) aggregate labour and savings, (ii) supply-chain quantities, and (iii) relative prices. All quantities report percentage deviations from steady-state output, while all prices are log-deviations from their own steady-state unless stated otherwise.

Increased Manufacturer Productivity

We start by considering an exogenous increase to the productivity of the manufacturer, a_3. We do this because it is the standard shock used in macroeconomic models of the business cycle and will help to determine whether the addition of the supply chain naturally amplifies shocks away from their source. Figure 4.8 illustrates the effect on variables that are not directly on the supply chain itself and figure 4.9 shows the response of each sector of the supply chain. First, we focus on the impact of the persistent shock ($\rho = 0.75$) as depicted by the solid lines in the figures.

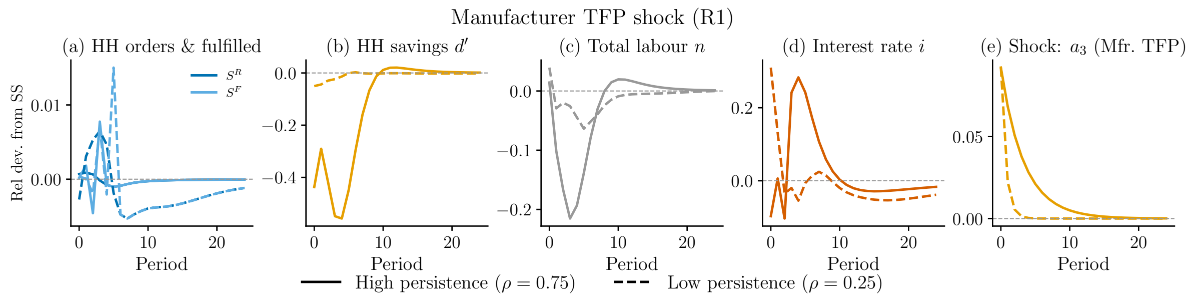

Figure 4.8: Impulse-responses of variables outside of the direct supply chain following a one standard deviation shock to manufacturer productivity, a_3.

The upstream productivity increase does not necessarily imply that there is a direct increase in household consumption of goods because of the natural delivery delays. Households, adhering to their Euler equation, try to smooth consumption over time - S^R increases somewhat for five periods. However, realized deliveries S^F are more volatile, and reveal an aggregate increase in goods deliveries around 5 periods after the initial shock. This is caused by information frictions within the supply chain on upstream goods delivery behaviour, which we assume: under this assumption, retailers initially withhold some deliveries before releasing a larger quantity once wholesaler deliveries (which we look at below) increase above steady-state. Most notably, households decrease their overall labour supply persistently for about 8 periods, drawing down on their savings to finance consumption in place of using flow labour income.

The main responses along the supply chain occur in the manufacturing sector. Faced with higher productivity, manufacturers increase production and build up inventories. The price of the manufactured good decreases to entice the wholesaler to increase their orders (O^R in the first row of panel (b)). This means that the manufacturer's inventory-sales ratio only increases slightly for about 3 periods. The wholesaler receives additional deliveries from the manufacturer, increasing wholesaler inventories persistently. The wholesaler, in turn, sees an initial slight decrease of orders received I^R below steady-state, before the shortfall is made up by subsequent later orders which draw down the stock of inventories. The upstream supply increase does not change downstream demand by a large magnitude, as most of the improvements in TFP accrue to firm profits through our assumption of supply chain stage-specific Nash bargaining over surpluses. Consequently, downstream pricing (column (e) of figure 4.9) remains essentially unchanged, especially relative to the upstream price change.

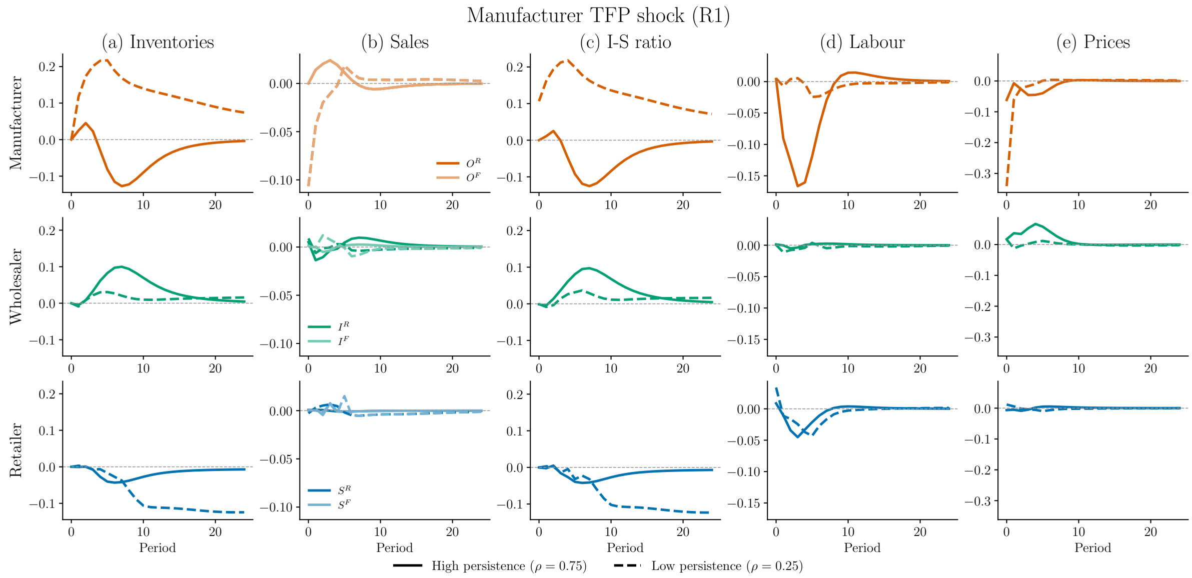

Figure 4.9: Impulse-responses of variables on the supply chain following a one standard deviation shock to manufacturer productivity, a_3.

While the traditional focus of the literature on the bullwhip effect emphasises downstream demand shocks propagating and amplifying upstream along the supply chain, a potential concern is that supply chains are frequently specified such that shocks always amplify to sectors further away along the supply chain. Our exercise here shows that this is not the case in our framework: upstream productivity shocks, despite some downstream propagation, retain most variation upstream. There is slight inventory and sales volatility propagation to the wholesaler, indicating that there is some quantity-related downstream pass-through in our model. But retailer propagation is minimal, echoing the vast periods of relatively larger inventory and sales volatility (especially prior to the GFC) in figures 4.2 and 4.4. Therefore, this exercise places the reaction of the firms on the supply chain in response to the manufacturer productivity shock as being reminiscent of the observed pre-GFC dynamics. In terms of prices, the effects are similar to a positive supply shock in standard DSGE models: upstream prices fall.

An important source of behaviour along the supply chain is labour. Given there is a disutility from supplying labour, but labour is needed at all stages of the supply chain, there is no crowding in of labour to the manufacturer. Instead, upstream labour decreases as the household aims to reduce its labour supply while the persistent nature of the shock is revealed. Because of the increased supply of the wholesaler to the retailer, downstream labour demand also decreases due to higher shelf-stocking efficiency in line with equation (4.12).

Overall, the productivity increase yields a temporary increase in upstream production. Given our calibration that entails sufficient inventories to smooth out such a temporary disturbance, this increase in productivity is spread out and buffered across time especially through inventory management according to our supply chain model. We also observe some downstream propagation in inventory and sales dynamics, consistent with some of the time series behaviour pre-GFC in section 2.5: the model-implied downstream propagation, however, is small. In numbers, the conditional variation in the inventory-sales ratio differs substantially between upstream manufacturers and downstream retailers: the ratio is about 8. This contrast is even more pronounced for prices, where the corresponding ratio of variances is approximately 39.

The role of persistence

We have so far discussed the persistent shock ($\rho = 0.75$), which generates a strong intertemporal substitution motive; the persistent positive upstream TFP shock led to persistently reduced labour supply. To highlight the importance of this mechanism, we examine the impact of a less persistent shock using $\rho = 0.25$.[16] The dotted lines in Figures 4.8 and 4.9 illustrate the IRFs in response to this shock.

The differences in propagation are especially stark in terms of labour supply, evidenced by panel (c) of figure 4.8. When the manufacturer TFP increase is short-lived, there is less of a crowding out of labour. The reason is that the short-lived shock does not allow as much of a persistent decrease in upstream labour supply to maintain the same output level. In general, even under the short-lived shock, the effects on the inventory, sales, and labour adjustment margins all continue to feature some persistence, but household savings and prices all return to steady-state more quickly, while inventory management tools absorb the more persistent aspects of variation in goods supply upstream versus downstream. To minimize distractions, and to remain in line with much of the RBC literature, we will focus on persistent shocks in the following.

[16] Since the model is specified in monthly frequency, this is equivalent to about 1.6% of the shock remaining after one quarter.

Changing Demand via Fiscal Transfers

The traditional bullwhip effect comes via a downstream demand shock. In our model, a clean way to introduce such a disturbance to aggregate demand is decreasing household transfers to the government. This is what we do in the next exercise, depicted by figures 4.10 and 4.11.

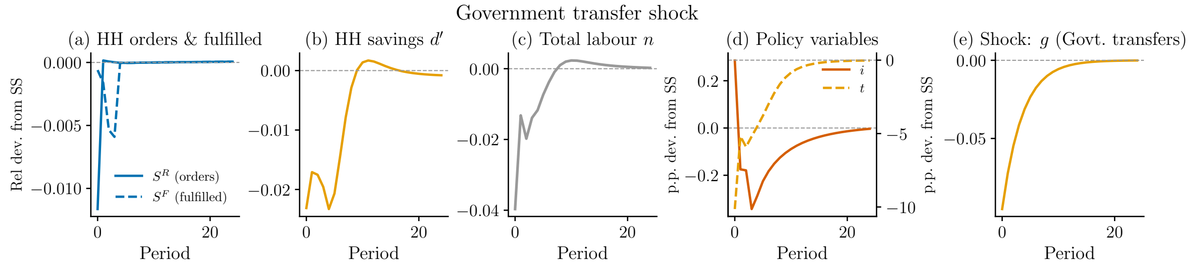

Figure 4.10: Impulse-responses of variables outside of the direct supply chain following a one standard deviation shock to government transfers, g.

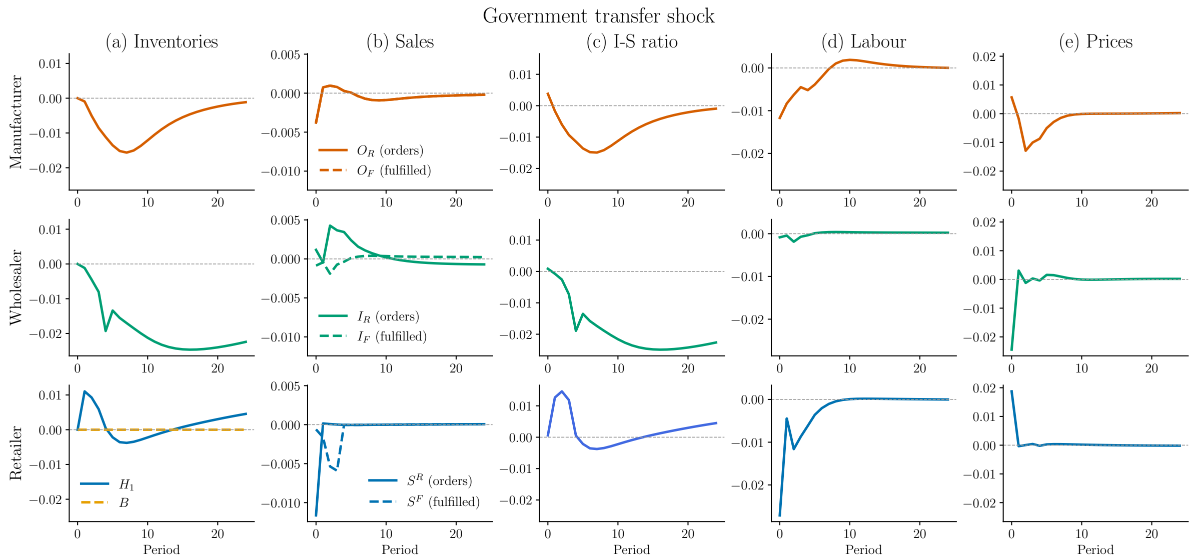

Figure 4.11: Impulse-responses of variables on the supply chain following a one standard deviation shock to government transfers, g.

Fundamentally, the change to government transfers does not imply that the production environment allows for more output, unless it leads to a crowding in of labour. It turns out that households do not do that, reducing in fact their labour supply for 8 periods, since the decrease in transfers allows for a smaller hit to disposable income when reducing labour supply. Under the baseline fiscal rule, the current fiscal impulse is largely offset by expected future adjustments in transfers and asset returns. Hence the shock does not operate as a simple positive disposable-income shock. Instead, households reduce labour supply and downstream demand falls in equilibrium. Consequently, retail inventories (bottom row of column (a) in figure 4.11) increase, as do downstream prices. The downstream impulse to sales and inventories clearly propagates through the supply chain, as the lack of downstream demand entices wholesalers to decrease their orders with the manufacturer O^R, consistent with the overall reduction of labour demand upstream (column (d) of figure 4.11).

Therefore, the change to government transfers mainly manifests through an intertemporal adjustment of production: on impact, labour is crowded out across the board to a degree that goes above the actual demand reduction, as evidenced by the drawing down of inventories upstream. The largest adjustment occurs with downstream labour, which is reduced sharply in the face of temporarily lower household demand. After about 10 periods, there is a slight increase in upstream labour, used to replenish the foregone inventories slowly. Given the observed magnitude and persistence of the volatility of prices, deliveries, and inventories, this intertemporal adjustment is not delivering a bullwhip-type propagation: the demand-side exogenous disturbance propagates up through the supply chain, but the magnitude of adjustment in the inventory-sales ratios is not larger upstream than downstream. In detail, inventories are about 0.9 times as volatile upstream relative to downstream in response to the government transfer shock; in terms of the inventory-sales ratio, upstream volatility is about 0.8 times the downstream variation. The relative price volatility is larger upstream over time, being about 1.2 times the size of downstream price volatility. In addition, transfer shocks punctually replicate the negative correlation between inventory-sales ratios and price movements that we observe in figure 4.5: for the retailer, the correlation between the two lies at -0.59, however, we cannot match the empirical evidence for the manufacturer in response to the upstream TFP shock (where the correlation coefficient is 0.281).

Increased Retailer Productivity

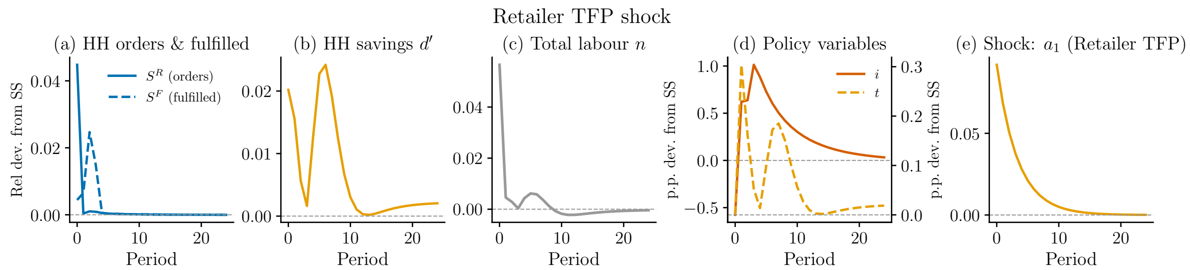

Figure 4.12: Impulse-responses of variables outside of the direct supply chain following a one standard deviation shock to retailer productivity, a_1.

We now consider the consequences of an innovation to the productivity at the *lower* end of the supply chain: a positive innovation to the productivity of the retailer. This shock is not a standard productivity shock in DSGE models; rather, it implies a decrease in the cost of shifting inventories from the backroom to the front of the store. Figures 4.12 and 4.13 illustrate the findings.

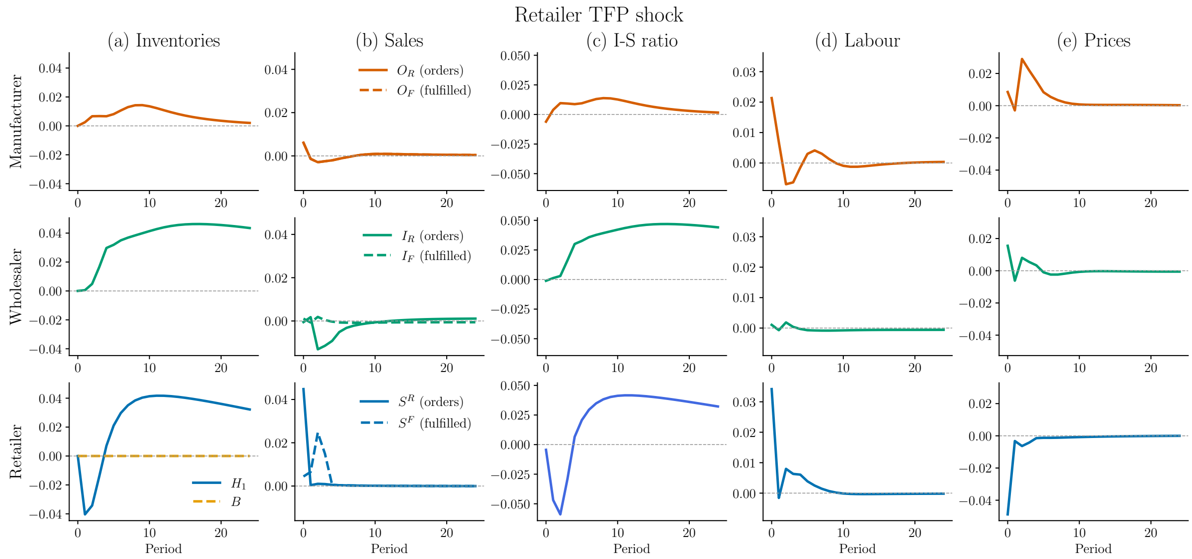

Figure 4.13: Impulse-responses of variables on the supply chain following a one standard deviation shock to retailer productivity, a_1.

It is instructive to note that increasing a_1 does not expand necessarily the number of goods supplied upstream but only makes shifting goods from the backroom to the front of store cheaper (more efficient). Overall goods supply changes are determined by how much labour is crowded in upstream (i.e., in the top panel of column (d) in figure 4.13). The increase in retailer TFP makes shifting goods for delivery to the frontroom cheaper. This is a decrease in the marginal cost of goods supply for the retailer. Consequently, retailer prices go down (bottom panel in row (e) of figure 4.13). Labour gets crowded *in* to the retailer to make use of the good production environment. Consequently, household orders in panel (a) of figure 4.12 increase, and the additional deliveries occur with a time delay of 1-3 periods. The increase in downstream TFP manifests in general equilibrium as a positive demand shock on behalf of the household.

Upstream supply expands slightly because labour is increased in manufacturing (n_3 in the top panel of column (d) of figure 4.13), which is needed to replenish foregone inventories at the bottom end of the supply chain. To entice more upstream production, upstream prices increase, as visible in column (e) of figure 4.12. Given our assumption of Nash bargaining over surpluses, this is fully consistent with a type of profit-sharing motive from the increased efficiency in the inventory management process downstream.

Just as the upstream TFP shock and the government transfer shock failed to deliver bullwhip-type behaviour,[17] the downstream TFP shock delivers upstream propagation, but it cannot be considered amplifying in a bullwhip-type sense, as evidenced by sales (column (b)), inventory-sales ratios (column (c)), and price adjustments (column (e)) alike. In numbers, the volatility in the inventory-sales ratio and sales levels is only 0.05 times as large upstream relative to downstream, while price volatility upstream is 0.7 times the size of downstream price volatility.

While it is tricky to disentangle the precise channels at work in general equilibrium, it seems that the nature of the downstream delivery technology, paired with the fundamental upstream delivery constraints, is what drives the observed result: households, who feel wealthier on the decrease of downstream prices, increase their demand. However, there is a smoothing channel internal to the retailer, given that technology and labour allow for within-retailer substitution in net inventories. Insofar as the productivity shock yields an increase in household net wealth, any additional upstream demand, however, must be smoothed across supply chain stages, which is done through adjustments in inventories, sales, and the pricing margin especially at the lower end of the supply chain, which is connected closest to the source of the demand change. Therefore, while a given downstream change in the production environment causes some propagation upstream through both quantities and prices, it is not an amplification. Larger retailer volatility in sales and inventory-sales ratios matches the Covid and post-Covid evidence in figures 4.2 and 4.4. In addition, the implied dynamics of downstream prices are also consistent with some link between changes to CPI inflation and inventory volatility, although the persistence of the downstream price changes in response to retailer TFP shocks remains limited, especially compared to upstream producer price variation that increases for 8 periods. Finally, we observe a mildly negative correlation only for the wholesaler (-0.148), but the downstream TFP shock does not replicate the empirical evidence for the retailer and the manufacturer in figure 4.5.

4.7 The Role of the Monetary-Fiscal Policy Mix - Mitigation or Amplification?

So far, our discussion paid limited attention to the role of the exact policy specification in use. Implicitly, and as became especially clear in the analysis of the transfer shock, we assumed in the above section that households themselves behave close to Ricardian with respect to the savings asset they have and the policies induced by the monetary and fiscal authorities: we effectively restrained the behaviour of the fiscal transfer to always nullify wealth effects coming from changes to the quantity and value of the savings

[17] As we will see in section 4.7, the policy mix plays a crucial role here.

asset. Concurrently, we assumed monetary policy to induce quite significant movements in the policy rate as a demand stabilization tool. While we nonetheless observed real and nominal effects from both productivity and transfer shocks due to the rich supply-side frictions, we now analyse the role of the policy authorities in greater detail. In particular, we analyse the differences in the equilibrium reactions of each entity within the model once we allow for (a) insufficient demand stabilization in response to technological shocks, and (b) fiscal policy to induce changes to total wealth as perceived by households. In a nutshell, we now compare the above results on the supply chain-based amplification of shocks with alternative scenarios under which monetary policy commits only imperfectly to demand stabilization through restrictive policy, while fiscal policy is allowed to commit to deficit-finance, altering the real present value of debt held by households when a shock hits.

Upstream manufacturer shocks under different policy regimes

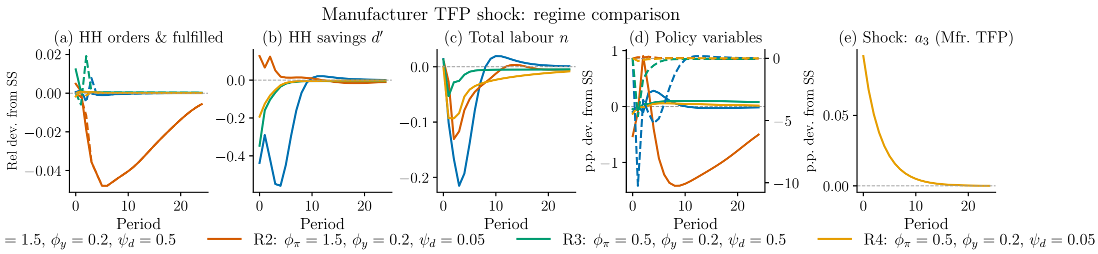

Figure 4.14: The response of the non-supply chain-side of the economy following an impulse in manufacturer productivity, a_3, in dependence on the monetary and fiscal policy regimes.

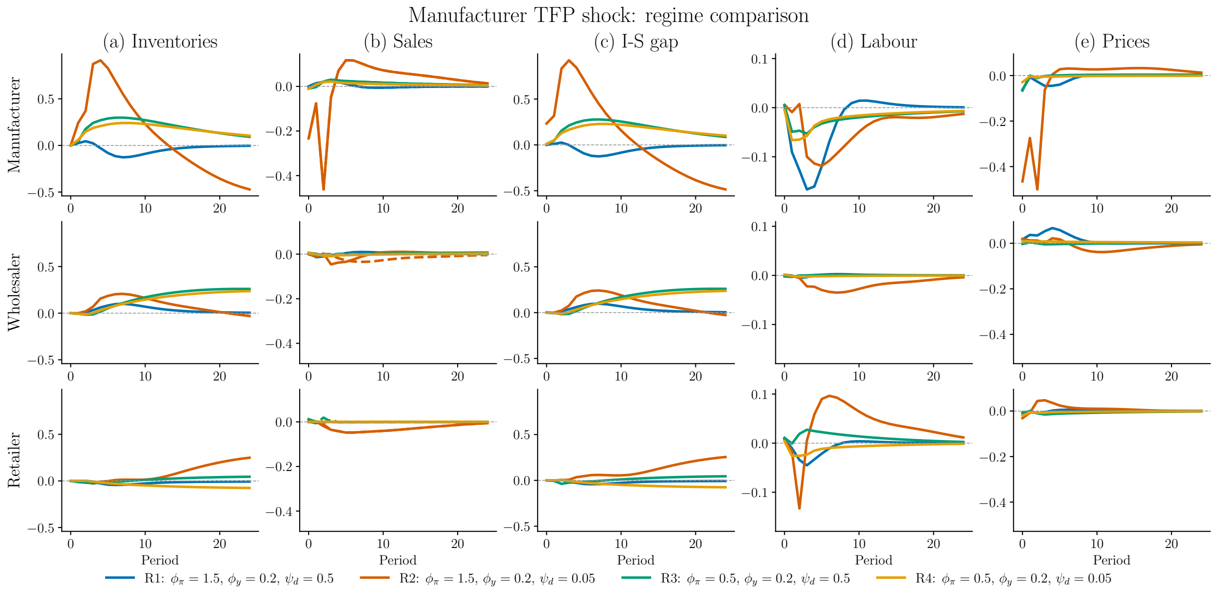

Figure 4.15: The response of the supply chain of the economy following an impulse in manufacturer productivity, a_3, in dependence on the monetary and fiscal policy regimes.

We first consider how the policy specification affects the downstream propagation of shocks by inducing a shock to upstream manufacturer productivity, depicted in figures 4.14 and 4.15. All impulse-responses depict the same persistent increase in manufacturer TFP as in figures 4.8 and 4.9. The blue line shows the base case from above. Under the orange line, fiscal policy does not commit to using transfers to stabilize its budget in equivalent real terms while monetary policy continues to commit to strongly inflation-stabilizing policy, and the green line depicts a less restrictive monetary policy authority paired with a fiscal authority that maintains budget balance through corrective taxation more stringently. Under the yellow line, fiscal policy commits to deficit-finance while the monetary authority does not commit to possibly strongly restrictive monetary policy.

The differences in the propagation are punctually stark. Most importantly, in the pairing of a deficit-tolerant fiscal authority and a rather inflation-averse monetary authority (orange lines), the shock loading on inventories, sales, and prices is much larger, paired with an expansion of household savings that is financed with a persistent decrease in household demand. Clearly, the temporary increase in the return on the savings asset paired with muted movements in future taxation makes the savings asset fairly attractive, contributing to a crowding in to the savings asset from the additional aggregate wealth created by the increase in manufacturer TFP.

This becomes clear across the entire supply chain, but most notably upstream, where the expected lack of downstream demand contributes to a much sharper decrease in prices and sales, paired with a concurrent increase in upstream inventories. This leads also to an adjustment in downstream labour demand (bottom panel of column (d) in figure 4.15), which decreases sharply below zero 1-4 periods after the shock, followed by a temporary increase of labour demand beyond steady-state as household demand recovers.

The other alternative calibrations, where a less inflation-averse monetary authority is paired with a budget-stabilizing fiscal authority (green lines) and a less stabilizing fiscal authority (yellow lines), respectively, also alter the model dynamics quite significantly: in general, there is less crowding out in upstream labour, contributing to an *increase* in manufacturer inventories and, with some delay, in wholesaler inventories, as evidenced by column (a) of figure 4.15. All these responses indicate that the behaviour of the monetary and fiscal authorities matters not only for the downstream variables close to the household, but also for upstream supply dynamics, which are intricately tied together in general equilibrium through changes to overall labour supply and goods demand. Only the qualitative direction of price movements remains broadly stable across the various policy regimes, with upstream prices generally decreasing in response to the TFP shock and downstream prices marginally increasing, if at all.

Even though the policy stance changes some of the propagation of the upstream TFP shocks, the fundamental implications regarding the bullwhip behaviour are qualitatively unchanged: the downstream propagation in sales and inventory-sales ratios is limited. To replicate the observed empirical negative correlation between inventory-sales ratios and prices from figure 4.5, the policy stance in the model matters quite clearly: for the retailer, the base case (blue line) is closest at replicating the negative correlation in resonse to upstream shocks (-0.59), while for the manufacturer, the shock to a_3 can induce a negative correlation between prices and inventory-sales ratios especially when the fiscal policy authority commits to deficit-finance (R2: -0.195, R4: -0.254).

In sum, the change in demand management through the fiscal and monetary policy-makers matters for the propagation of upstream supply shocks across the economy. The crucial mechanism is the steering of the relative attractiveness of the savings asset in relation to current household demand, which depends crucially on the monetary policy stance, which influences the attractiveness of the savings asset directly and, additionally, also impacts the future tax stance of the government that allows it to maintain a balanced budget.

Downstream demand shocks under different policy regimes

Since this section adds considerations of paths of fiscal policy that would be unsustainable on their own, we now consider how the effects of downstream demand shocks induced by the fiscal authority itself are altered by differences in taxation and monetary stabilization schemes. Figures 4.16 and 4.17 illustrate the effects of a reduction in transfers from the household to the fiscal authority.

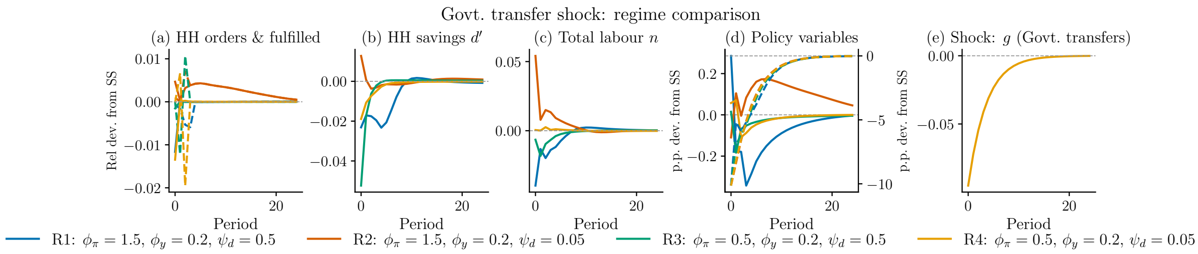

Figure 4.16: The response of the non-supply chain-side of the economy following an impulse in household transfers to the government, g, in dependence on the monetary and fiscal policy regimes.

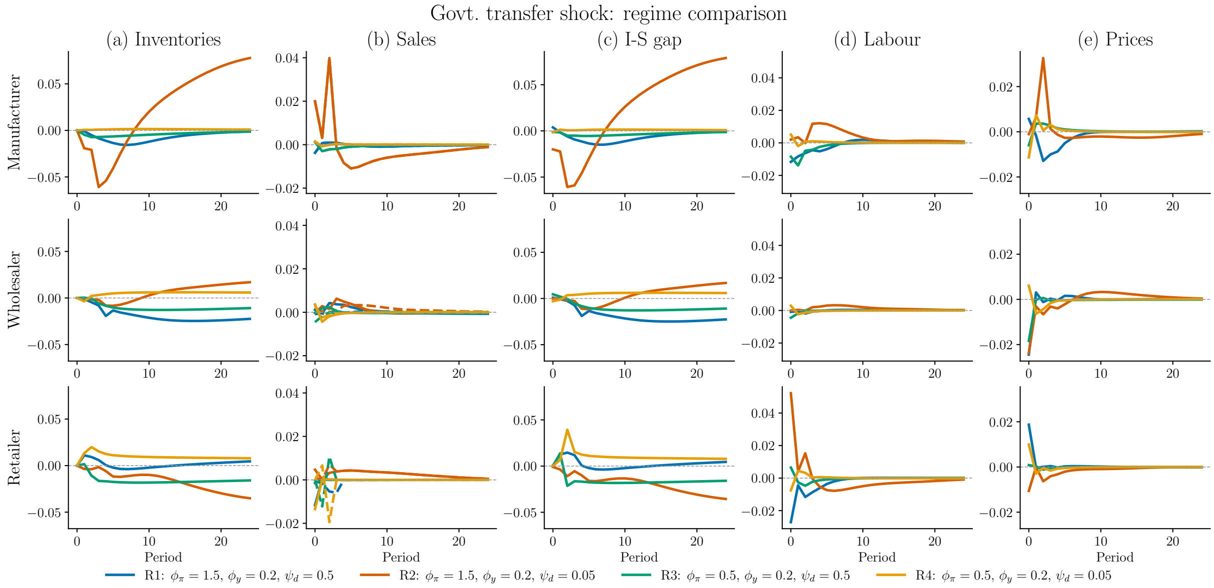

Figure 4.17: The response of the supply chain-related variables in the economy following an impulse in household transfers to the government, g, in dependence on the monetary and fiscal policy regimes.

This exercise is close in spirit to Ferrari (2025) in terms of focussing on the effects of demand-side distortionary policy actions. The reduction in household-to-government transfer payments shock has quite different effects depending on the policy specification chosen. As in the previous analysis of the upstream TFP shock, the combination of an inflation-averse monetary authority and a deficit-tolerant fiscal authority displays the starkest differences: only in that case, we observe an increase in household savings alongside a *persistent* increase in household demand, as well as a crowding in of aggregate labour. The differences in the equilibrium outcomes are again not only limited to the source of the shock (here, downstream), but we also observe a major upstream adjustment margin through all of sales, inventories, and prices, as evidenced by the top row of figure 4.17. While the dynamics in the scenarios R3 (less active monetary, stabilizing fiscal - green lines) and R4 (less active monetary, deficit-tolerant fiscal - yellow lines) are also interesting, we do not discuss them in detail now to limit the length of the manuscript.

The increase in downstream demand in the orange case is clearly related to the slight decrease in downstream prices in equilibrium (column (e) of figure 4.17), which crowds in that additional downstream demand. Clearly, under this policy specification households have the largest wealth effect from the fiscal impulse, as implied by standard DSGE frameworks: on the one hand, the fiscal policy authority does not commit to future tax increases that would recuperate their temporary losses in equivalent real terms, and on the other hand, the monetary authority ends up tightening the policy rate in the future, delivering a gain on household lending to the government. All this bolsters household final goods demand. Consequently, upstream inventories (column (a) of figure 4.17) remain below steady-state for quite a while to replace the additional goods moved downstream. In these cases, we also have what ends up the clearest evidence for bullwhip-type upstream propagation: the variance of upstream sales is 4.5 times as large relative to downstream sales, and the relative upstream-to-downstream volatility of inventory-sales ratios and prices lie at 5 and 7, respectively.

Pinning down the bullwhip as a volatility phenomenon

Our discussion of the model behaviour alluded to measures of conditional volatility (where the conditioning is effectively on the type of exogenous shock). We provide a systematic review of this type of shock-dependent volatility here. In particular, we look at the relative difference of upstream to downstream volatility across the different policy regimes and shocks considered. Table 4.4 illustrates the results for inventory-sales ratios and inventories, while table 4.5 shows the results for sales and prices.

	Shock to a_3		Shock to a_1		Shock to g	
Regime	I/S	Inv	I/S	Inv	I/S	Inv
R1: ϕ_π=1.5, ψ_d=0.5	7.335	7.832	0.039	0.042	0.792	0.891
R2: ϕ_π=1.5, ψ_d=0.05	6.504	5.975	5.421	4.888	5.125	4.859
R3: ϕ_π=0.5, ψ_d=0.5	11.755	13.820	0.002	0.003	0.036	0.055
R4: ϕ_π=0.5, ψ_d=0.05	4.975	5.612	0.011	0.015	0.009	0.013

Table 4.4: Manufacturer-to-retailer variance ratios ($\mathbb{V}_{\text{mfr}}/\mathbb{V}_{\text{ret}}$) implied by the impulse response functions. Each entry reports the ratio of the manufacturer's to the retailer's conditional variance of the inventory-sales (I/S) ratio or inventory level, computed from deviations relative to steady state. Values above one indicate upstream variance amplification.

	Shock to a_3		Shock to a_1		Shock to g	
Regime	Sales	Price	Sales	Price	Sales	Price
R1: ϕ_π=1.5, ψ_d=0.5	23.579	38.689	0.087	0.707	1.018	1.232
R2: ϕ_π=1.5, ψ_d=0.05	0.554	45.668	0.588	6.101	0.577	6.927
R3: ϕ_π=0.5, ψ_d=0.5	14.488	3.946	0.031	0.287	0.091	30.044
R4: ϕ_π=0.5, ψ_d=0.05	1306.823	1.426	0.054	0.139	0.007	1.744

Table 4.5: Manufacturer-to-retailer variance ratios ($\mathbb{V}_{\text{mfr}}/\mathbb{V}_{\text{ret}}$) for sales and prices, implied by the impulse response functions. Sales volatility is the ratio of the manufacturer's (O_F) to the retailer's (S_F) conditional variance, computed from level deviations relative to steady state. Price volatility is the ratio of the manufacturer's (p_3) to the retailer's (p_1) conditional variance, computed from log-deviations relative to steady state. Values above one indicate upstream variance amplification.

Larger variation upstream occurs for each entry larger than one - clearly, this is not a global result baked into the model. Volatility in inventory-sales ratios, inventories, sales, and prices is larger upstream *only* when shocks originate upstream (in the two leftmost columns) *or* when downstream shocks are paired with appropriate policy regimes (R1 for sales, R2 for prices) in table 4.5. The strength of the amplification of downstream shocks to upstream varies substantially with the specification of monetary and fiscal policy. This is useful in light of the empirical evidence: the data do not point to a permanent upstream dominance in every dimension, but rather to time-varying relative volatility across stages.

The evidence in section 2.5 points at the possibility of time-varying propagation of shocks into relative price differences upstream and downstream. In particular, there are some periods in which price variation is larger downstream, which our previous discussion did not really account for. Looking at table 4.5, the leading candidate to rationalize such episodes are downstream TFP shocks, since the variance ratio for prices in response to shocks to a_1 is frequently less than one.

Finally, table 4.5 drives home the important point that the bullwhip mechanism implied by the model is not purely real. The same shocks that generate movements in inventories, sales, and inventory-sales ratios also induce sizeable movements in relative prices along the supply chain. But, as alluded to in section 4.6 in particular, our model at times struggles to match the negative correlation between prices and inventory-sales ratios for every shock, especially for the manufacturer.

4.8 Conclusion

This chapter has developed a new macroeconomic theory of inventory behaviour on supply chains. We incorporate a simple form of imperfect information that requires firms to commit to production and inventory decisions prior to the manifestation of exogenous innovations, building endogenous incentives for firms to hold precautionary inventories even in a steady-state without aggregate risk. Our derivations rely on a novel formulation of inventories as net order book variables, which captures the difference between the momentary capacity to deliver goods and the orders demanded by buyers of the firm's good that have yet not been delivered.

In our model, the *bullwhip effect* can arise as a consequence of delivery lags, intertemporal optimisation, and incentives for manufacturers at the upper end of the supply chain to over-produce in a bid to retain buffer stock when exogenous shocks manifest. The responses to upstream and downstream productivity shocks are able to partly rationalize observed empirical regularities on relatively larger upstream sales volatility and the time-varying relative magnitude of volatility in inventory-sales ratios, respectively.

Finally, we have also shown that monetary and fiscal authorities can act as stabilizing anchors *or* as destabilizing amplifiers depending on their interest- and transfer-setting behaviours, interacting supply-side amplification with demand-side policies. That behaviour, in turn, matters for the propagation of relative inventory(-sales) volatility to downstream price pressures, hinting at the possible relevance of this mechanism for the post-Covid price pressures and the general relevance of demand policies for bullwhip-type behaviour on supply chains.

While we made some steps in relating our model to the rich empirical evidence on the bullwhip effect (including amplification of inventory/sales and price dynamics), we cannot account for all of the time series evidence in section 2.5. Future work must look at a greater fit of the model to the empirics, such as by evaluating the relative importance of each of the structural shocks of the model over time, as well as by disciplining the model more rigorously with empirical moments. This is especially important for the steady-state, since the model is hugely dependent on the reference steady-state through the non-linearities embedded.

Additionally, the rich information structure and the consideration of distributions of innovations is prone to give rise to equilibrium multiplicity. We briefly referred to this possibility above, but did not explore it within the confines of our solution method, which is left open for future research.

Bibliography

Abramovitz, M. (1950). Inventories and Business Cycles. *National Bureau of Economic Research - Studies in Business Cycles 1950*.

Acemoglu, D. and M. K. Jensen (2015). Robust Comparative Statics in Large Dynamic Economies. *Journal of Political Economy 123*(3), 587–640.

Alberola, E., C. Cantu Garcia, P. Cavallino, and N. Mirkov (2021). Fiscal Regimes and the Exchange Rate. *Bank for International Settlements WP no. 950*.

Andreasen, M. M., J. H. Christensen, and S. Riddell (2021). The TIPS Liquidity Premium. *Review of Finance 25*(6), 1639–1675.

Angeletos, G.-M., C. Lian, and C. Wolf (2024a). Can Deficits Finance Themselves? *Econometrica 92*(5), 1351–1390.

Angeletos, G.-M., C. Lian, and C. Wolf (2024b). Deficits and Inflation: HANK meets FTPL. Technical report, NBER Working Paper No. w33102.

Antràs, P. and D. Chor (2022). Global Value Chains. In *Handbook of International Economics: International Trade, Volume 5*, pp. 297–376. Elsevier.

Antsaklis, P. J. and A. N. Michel (2006). *Linear Systems.* Springer.

Arrow, K. J., T. Harris, and J. Marschak (1951). Optimal Inventory Policy. *Econometrica*, 250–272.

Ascari, G., P. Beck-Friis, A. Florio, and A. Gobbi (2023). Fiscal Foresight and the Effects of Government Spending: It's All in the Monetary-Fiscal Mix. *Journal of Monetary Economics 134*, 1–15.

Ascari, G., D. Bonam, L. Mori, and A. Smadu (2024). Fiscal Policy and Inflation in the Euro Area. *DNB Working Paper No. 820*.

Atkeson, A. and A. Burstein (2008). Pricing-to-market, Trade Costs, and International Relative Prices. *American Economic Review 98*(5), 1998–2031.

Auclert, A., B. Bardóczy, M. Rognlie, and L. Straub (2021). Using the Sequence-Space Jacobian to Solve and Estimate Heterogeneous-Agent Models. *Econometrica 89*(5), 2375–2408.

Auclert, A., R. Rigato, M. Rognlie, and L. Straub (2024). New Pricing Models, Same Old Phillips Curves? *The Quarterly Journal of Economics 139*(1), 121–186.

Auclert, A., M. Rognlie, and L. Straub (2023). Determinacy and Existence in the Sequence Space. *Manuscript*.

Auclert, A., M. Rognlie, and L. Straub (2024a). Fiscal and Monetary Policy with Heterogeneous Agents. *NBER Working Paper No. w32991*.

Auclert, A., M. Rognlie, and L. Straub (2024b). The Intertemporal Keynesian Cross. *Journal of Political Economy 132*(12), 4068–4121.

Auerbach, A. J. and Y. Gorodnichenko (2012). Measuring the Output Responses to Fiscal Policy. *American Economic Journal: Economic Policy 4*(2), 1–27.

Backus, D. K. and G. W. Smith (1993). Consumption and Real Exchange Rates in Dynamic Economies with Non-traded Goods. *Journal of International Economics 35*(3–4), 297–316.

Bahaj, S., R. Czech, S. Ding, and R. Reis (2023). The Market for Inflation Risk. *Bank of England Staff Working Paper No. 1028*.

Bai, X., J. Fernández-Villaverde, Y. Li, and F. Zanetti (2024). The Causal Effects of Global Supply Chain Disruptions on Macroeconomic Outcomes: Evidence and Theory. *NBER Working Paper No. w32098*.

Banerjee, R., V. Boctor, A. N. Mehrotra, and F. Zampolli (2022). *Fiscal Deficits and Inflation Risks: The Role of Fiscal and Monetary Regimes.* Bank for International Settlements, Monetary and Economic Department.

Baqaee, D. and E. Farhi (2022). Supply and Demand in Disaggregated Keynesian Economies with an Application to the Covid-19 Crisis. *American Economic Review 112*(5), 1397–1436.

Barr, D. G. and J. Y. Campbell (1997). Inflation, Real Interest Rates, and the Bond Market: A Study of UK Nominal and Index-Linked Government Bond Prices. *Journal of Monetary Economics 39*(3), 361–383.

Barro, R. (2003). Optimal Management of Indexed and Nominal Debt. *Annals of Economics and Finance 4*(1), 1–15.

Barro, R. J. and F. Bianchi (2025). Fiscal Influences on Inflation in OECD Countries, 2020-2023. *The Economic Journal*, ueaf066.

Bassetto, M., L. Benzoni, and J. Hall (2024). On the Mechanics of Fiscal Inflations. Technical report, Working Paper.

Bassetto, M. and W. Cui (2018). The Fiscal Theory of the Price Level in a World of Low Interest Rates. *Journal of Economic Dynamics and Control 89*, 5–22.

Bassetto, M. and D. S. Miller (2025). A Monetary-Fiscal Theory of Sudden Inflations. *The Quarterly Journal of Economics 140*(3), 1–42.

Bayer, C., B. Born, and R. Luetticke (2024). Shocks, Frictions, and Inequality in US Business Cycles. *American Economic Review 114*(5), 1211–1247.

Bayer, C. and R. Luetticke (2020). Solving Discrete Time Heterogeneous Agent Models with Aggregate Risk and Many Idiosyncratic States by Perturbation. *Quantitative Economics 11*(4), 1253–1288.

Benigno, G., J. Di Giovanni, J. J. Groen, and A. I. Noble (2022). The GSCPI: A New Barometer of Global Supply Chain Pressures. *FRB of New York Staff Report* (1017).

Benigno, G., B. Hofmann, and D. Sandri (2024). Quo vadis, r*? The Natural Rate of Interest After the Pandemic. *BIS Quarterly Review 4*, 17–30.

Benigno, P. and G. B. Eggertsson (2023). It's Baaack: The Surge in Inflation in the 2020s and the Return of the Non-Linear Phillips Curve. *NBER Working Paper No. w31197*.

Benigno, P. and A. Missale (2004). High Public Debt in Currency Crises: Fundamentals Versus Signaling Effects. *Journal of International Money and Finance 23*(2), 165–188.

Bergin, P. R. (2000). Fiscal Solvency and Price Level Determination in a Monetary Union. *Journal of Monetary Economics 45*(1), 37–53.

Bianchi, F., R. Faccini, and L. Melosi (2023). A Fiscal Theory of Persistent Inflation. *The Quarterly Journal of Economics 138*(4), 2127–2179.

Bianchi, F. and L. Melosi (2017). Escaping the Great Recession. *American Economic Review 107*(4), 1030–1058.

Bianchi, F. and L. Melosi (2019). The Dire Effects of the Lack of Monetary and Fiscal Coordination. *Journal of Monetary Economics 104*, 1–22.

Bianchi, F. and L. Melosi (2022). Inflation as a Fiscal Limit. *FRB of Chicago Working Paper No. WP 2022-37*.

Bianchi, J. and S. Bigio (2022). Portfolio Choice and Settlement Frictions: A Theory of Endogenous Convenience Yields. *NBER Working Paper No. 34474*.

Bigio, S., N. Caramp, and D. Silva (2024). Sticky Inflation: Monetary Policy when Debt Drags Inflation Expectations. *NBER Working Paper No. w33190*.

Bils, M. and J. A. Kahn (2000). What Inventory Behavior Tells us About Business Cycles. *American Economic Review 90*(3), 458–481.

BIS (2024). Bank for International Settlements Debt Securities Statistics. Data retrieved from `https://data.bis.org/topics/DSS/tables-and-dashboards`.

Bizzarri, M., M. Pangallo, and F. Queirós (2025). Production Networks, Time to Build and Endogenous Oscillations. *Available at SSRN 5166813*.

Blanchard, O. J. (1985). Debt, Deficits, and Finite Horizons. *Journal of Political Economy 93*(2), 223–247.

Blanchard, O. J. and C. M. Kahn (1980). The Solution of Linear Difference Models under Rational Expectations. *Econometrica*, 1305–1311.

Blinder, A. S. (1986). Can the Production Smoothing Model of Inventory Behavior Be Saved? *The Quarterly Journal of Economics 101*(3), 431–453.

Blinder, A. S., M. C. Lovell, and L. H. Summers (1981). Retail Inventory Behavior and Business Fluctuations. *Brookings Papers on Economic Activity 1981*(2), 443–520.

Bordo, M. D., O. Bush, and R. Thomas (2022). "Muddling Through or Tunnelling Through?" UK Monetary and Fiscal Exceptionalism During the Great Inflation. *Bank of England Staff Working Paper No 1135*.

Braun, R., S. Miranda-Agrippino, and T. Saha (2024). Measuring Monetary Policy in the UK: The UK Monetary Policy Event-Study Database. *Journal of Monetary Economics*, 103645.

Brunnermeier, M. K., S. Merkel, and Y. Sannikov (2022). The Fiscal Theory of Price Level with a Bubble. *NBER Working Paper No. w27116*.

Brunnermeier, M. K., S. Merkel, and Y. Sannikov (2024). Safe Assets. *Journal of Political Economy 132*(11), 3603–3657.

Cairns, A. and D. Wilkie (2026). British Government Securities Database. Data retrieved from the Heriot-Watt/ Institute and Faculty of Actuaries/ESCoE British Government Securities Database, `https://www.macs.hw.ac.uk/~andrewc/gilts/`.

Campbell, J. Y. and R. J. Shiller (1988). The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors. *The Review of Financial Studies 1*(3), 195–228.

Campbell, J. Y. and R. J. Shiller (1996). A Scorecard for Indexed Government Debt. *NBER Macroeconomics Annual 11*, 155–197.

Campos, R. G., J. Fernández-Villaverde, G. Nuño, and P. Paz (2024). Navigating by Falling Stars: Monetary Policy with Fiscally Driven Natural Rates. *NBER Working Paper No. w32219*.

Caplin, A. S. (1985). The Variability of Aggregate Demand with (S, s) Inventory Policies. *Econometrica*, 1395–1409.

Caramp, N. and D. H. Silva (2023). Fiscal Policy and the Monetary Transmission Mechanism. *Review of Economic Dynamics 51*, 716–746.

Caramp, N. and D. H. Silva (2026). Monetary Policy and Wealth Effects: The Role of Risk and Heterogeneity. *The Journal of Finance 81*(2), 1011–1052.

Carreras-Valle, M.-J. (2026). Increasing Inventories: The Role of Delivery Times. *American Economic Journal: Macroeconomics, forthcoming*.

Chen, X., E. M. Leeper, and C. Leith (2022). Strategic Interactions in US Monetary and Fiscal Policies. *Quantitative Economics 13*(2), 593–628.

Christiano, L. J. and M. Eichenbaum (1989). Temporal Aggregation and the Stock Adjustment Model of Inventories. In *The Rational Expectations Equilibrium Inventory Model: Theory and Applications*, pp. 70–108. Springer.

Cloyne, J., J. Martinez, H. Mumtaz, and P. Surico (2023). Do Tax Increases Tame Inflation? In *AEA Papers and Proceedings*, Volume 113, pp. 377–381.

Cochrane, J. H. (2001). Long-Term Debt and Optimal Policy in the Fiscal Theory of the Price Level. *Econometrica 69*(1), 69–116.

Cochrane, J. H. (2011). Determinacy and Identification with Taylor Rules. *Journal of Political Economy 119*(3), 565–615.

Cochrane, J. H. (2019). The Value of Government Debt. NBER Working Paper No. w26090, National Bureau of Economic Research.

Cochrane, J. H. (2022a). Fiscal Histories. *Journal of Economic Perspectives 36*(4), 125–46.

Cochrane, J. H. (2022b). The Fiscal Roots of Inflation. *Review of Economic Dynamics 45*, 22–40.

Cochrane, J. H. (2023). *The Fiscal Theory of the Price Level*. Princeton University Press.

Cochrane, J. H. (2025). Inflation Dynamics with a Generalized Lucas Phillips Curve. NBER Working Paper No. w33888, National Bureau of Economic Research.

Corsetti, G., L. Dedola, and S. Leduc (2008). International Risk Sharing and the Transmission of Productivity Shocks. *The Review of Economic Studies 75*(2), 443–473.

Corsetti, G. and B. Maćkowiak (2006). Fiscal Imbalances and the Dynamics of Currency Crises. *European Economic Review 50*(5), 1317–1338.

Corsetti, G. and B. Maćkowiak (2024). Gambling to Preserve Price (and Fiscal) Stability. *IMF Economic Review 72*(1), 32–57.

Daniel, B. C. (2007). The Fiscal Theory of the Price Level and Initial Government Debt. *Review of Economic Dynamics 10*(2), 193–206.

Davig, T. and E. M. Leeper (2007). Generalizing the Taylor Principle. *American Economic Review 97*(3), 607–635.

Davoodi, H., P. Elger, A. Fotiou, D. Garcia-Macia, A. Lagerborg, R. Lam, and S. Pillai (2022). *Fiscal Rules Dataset: 1985-2021*. International Monetary Fund, Washington, D.C.

Di Giovanni, J., S. Kalemli-Özcan, A. Silva, and M. A. Yildirim (2022). Global Supply Chain Pressures, International Trade, and Inflation. *National Bureau of Economic Research Working Paper no. 30240*.

Di Tella, S., B. M. Hébert, P. Kurlat, and Q. Wang (2023). The Zero-Beta Interest Rate. Technical report, NBER Working Paper No. w31596.

Ding, X. and Z. Jiang (2024). Monetary-Fiscal Coordination with International Hegemon. Technical report, NBER Working Paper No. w33123.

Du, W., R. Keerati, and J. Schreger (2025). Decoupling Dollar and Treasury Privilege. Technical report, International Finance Discussion Paper No. 1427.

Dupor, B. (2000). Exchange Rates and the Fiscal Theory of the Price Level. *Journal of Monetary Economics 45*(3), 613–630.

Dupraz, S. and A. R. Picco (2025). Fiscal Requirements for Price Stability when Households are Not Ricardian. *European Central Bank Working Paper no. 2025/3038*.

Eichenbaum, M. S. (1984). Rational Expectations and the Smoothing Properties of Inventories of Finished Goods. *Journal of Monetary Economics 14*(1), 71–96.

Ellison, M. and A. Scott (2020). Managing the UK National Debt: 1694–2018. *American Economic Journal: Macroeconomics 12*(3), 227–257.

Engel, C. and K. D. West (2005). Exchange Rates and Fundamentals. *Journal of Political Economy 113*(3), 485–517.

Eroglu, C., B. D. Williams, and M. A. Waller (2013). The Backroom Effect in Retail Operations. *Production and Operations Management 22*(4), 915–923.

Fama, E. F. (1984). Forward and Spot Exchange Rates. *Journal of Monetary Economics 14*(3), 319–338.

Fernández-Villaverde, J. (2026). Deep Learning for Solving Economic Models. *Journal of Economic Literature* (forthcoming).

Fernández-Villaverde, J. and O. Levintal (2018). Solution Methods for Models with Rare Disasters. *Quantitative Economics 9*(2), 903–944.

Fernández-Villaverde, J., J. F. Rubio-Ramírez, and F. Schorfheide (2016). Solution and Estimation Methods for DSGE Models. In *Handbook of Macroeconomics*, Volume 2, pp. 527–724. Elsevier.

Ferrari, A. (2025). Inventories, Demand Shocks Propagation and Amplification in Supply Chains. *arXiv preprint arXiv:2205.03862*.

Fischer, S. (1975). The Demand for Index Bonds. *Journal of Political Economy 83*(3), 509–534.

Fleming, M. J. and N. Krishnan (2012). The Microstructure of the TIPS Market. *FRBNY Economic Policy Review*.

Fransoo, J. C. and M. J. Wouters (2000). Measuring the Bullwhip Effect in the Supply Chain. *Supply Chain Management: An International Journal 5*(2), 78–89.

Gagliardone, L., M. Gertler, S. Lenzu, and J. Tielens (2023). Anatomy of the Phillips Curve: Micro Evidence and Macro Implications. *NBER Working Paper No. w31382*.

Galí, J. (2015). *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework and its Applications*. Princeton University Press.

Gali, J. and T. Monacelli (2005). Monetary Policy and Exchange Rate Volatility in a Small Open Economy. *The Review of Economic Studies 72*(3), 707–734.

Garcia, J. A. and A. A. van Rixtel (2007). Inflation-Linked Bonds from a Central Bank Perspective. *ECB Occasional Paper No. 62*.

Giannitsarou, C., A. Scott, and E. M. Leeper (2006). Inflation Implications of Rising Government Debt. In *NBER International Seminar on Macroeconomics*, Number 1, pp. 393–439. The University of Chicago Press Chicago, IL.

Goncalves, S., A. M. Herrera, L. Kilian, and E. Pesavento (2024). State-Dependent Local Projections. *Journal of Econometrics 244*(2), 105702.

Gourinchas, P.-O. and H. Rey (2007). International Financial Adjustment. *Journal of Political Economy 115*(4), 665–703.

Graham, L. and S. Wright (2007). Nominal Debt Dynamics, Credit Constraints and Monetary Policy. *The B.E. Journal of Macroeconomics 7*(1).

Greenwood, R., S. G. Hanson, and J. C. Stein (2015). A Comparative-Advantage Approach to Government Debt Maturity. *The Journal of Finance 70*(4), 1683–1722.

Gürkaynak, R. S., B. Sack, and J. H. Wright (2010). The TIPS Yield Curve and Inflation Compensation. *American Economic Journal: Macroeconomics 2*(1), 70–92.

Hagedorn, M. (2021). A Demand Theory of the Price Level. CEPR Discussion Paper no. DP11364.

Hagedorn, M. (2024). The Failed Theory of the Price Level. CEPR Discussion Paper no. DP17816.

Hall, G. J., J. Payne, T. J. Sargent, et al. (2018). *US Federal Debt 1776-1960: Quantities and Prices*. NYU Stern, Department of Economics.

Hall, G. J. and T. J. Sargent (2011). Interest Rate Risk and Other Determinants of Post-WWII US Government Debt/GDP Dynamics. *American Economic Journal: Macroeconomics 3*(3), 192–214.

Hall, G. J. and T. J. Sargent (2025). Fiscal Consequences of the US War on COVID. *International Economic Review 66*(5), 1753–1780.

Hamilton, J. D. (2018). Why You Should Never Use the Hodrick-Prescott Filter. *Review of Economics and Statistics 100*(5), 831–843.

Hazell, J., J. Herreno, E. Nakamura, and J. Steinsson (2022). The Slope of the Phillips Curve: Evidence from US States. *The Quarterly Journal of Economics 137*(3), 1299–1344.

Hazell, J. and S. Hobler (2024). Do Deficits Cause Inflation? A High Frequency Narrative Approach. *mimeo*.

Hilscher, J., A. Raviv, and R. Reis (2022). Inflating Away the Public Debt? An Empirical Assessment. *The Review of Financial Studies 35*(3), 1553–1595.

Holt, C. C., F. Modigliani, J. F. Muth, and H. A. Simon (1960). Production Planning, Inventories, and Workforce. *Englewood Cliffs, N]: Prentice Hall*.

Hyland, C. C. (2026). Convenience Yields as Fiscal Limits. Technical report, mimeo.

Iacoviello, M., F. Schiantarelli, and S. Schuh (2011). Input and Output Inventories in General Equilibrium. *International Economic Review 52*(4), 1179–1213.

Itskhoki, O. (2021). The Story of the Real Exchange Rate. *Annual Review of Economics 13*(1), 423–455.

Jacobson, M. M., E. M. Leeper, and B. Preston (2023). Recovery of 1933. Technical report, NBER Working Paper No. w25629.

Jiang, Z. (2021). US Fiscal Cycle and the Dollar. *Journal of Monetary Economics 124*, 91–106.

Jiang, Z. (2022). Fiscal Cyclicality and Currency Risk Premia. *The Review of Financial Studies 35*(3), 1527–1552.

Jiang, Z., A. Krishnamurthy, and H. Lustig (2021). Foreign Safe Asset Demand and the Dollar Exchange Rate. *The Journal of Finance 76*(3), 1049–1089.

Jiang, Z., A. Krishnamurthy, H. Lustig, and J. Sun (2024). Convenience Yields and Exchange Rate Puzzles. Technical report, NBER Working Paper No. w32092.

Jiang, Z., H. Lustig, S. Van Nieuwerburgh, and M. Z. Xiaolan (2024). The US Public Debt Valuation Puzzle. *Econometrica 92*(4), 1309–1347.

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

Jordà, Ò. and A. M. Taylor (2025). Local Projections. *Journal of Economic Literature 63*(1), 59–110.

Judd, K. L., L. Maliar, S. Maliar, and R. Valero (2014). Smolyak Method for Solving Dynamic Economic Models: Lagrange Interpolation, Anisotropic Grid and Adaptive Domain. *Journal of Economic Dynamics and Control 44*, 92–123.

Kahn, J. A. (1992). Why is production more volatile than sales? Theory and Evidence on the Stockout-Avoidance Motive for Inventory-Holding. *The Quarterly Journal of Economics 107*(2), 481–510.

Kanzow, C., N. Yamashita, and M. Fukushima (2004). Levenberg–Marquardt Methods with Strong Local Convergence Properties for Solving Nonlinear Equations with Convex Constraints. *Journal of Computational and Applied Mathematics 172*(2), 375–397.

Kaplan, G. (2025a). A Note On Uniqueness of Equilibrium in New Keynesian Models With and Without Ricardian Households. *mimeo*.

Kaplan, G. (2025b). Implications of Fiscal-Monetary Interaction from HANK Models. *Econometric Society World Congress semi-plenary lecture*.

Kaplan, G., B. Moll, and G. L. Violante (2018). Monetary Policy According to HANK. *American Economic Review 108*(3), 697–743.

Kaplan, G., G. Nikolakoudis, and G. L. Violante (2024). Price Level and Inflation Dynamics in Heterogeneous Agent Economies. *National Bureau of Economic Research WP no. w31433*.

Krishnamurthy, A. and Y. Ma (2025). The Demand and Supply of Convenience Assets. *Annual Review of Financial Economics 17*.

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

Kryvtsov, O. and V. Midrigan (2013). Inventories, Markups, and Real Rigidities in Menu Cost Models. *Review of Economic Studies 80*(1), 249–276.

Kwicklis, N. (2025). Active vs. Passive Policy and the Trade-Off Between Output and Inflation in HANK. *Journal of Monetary Economics 151*, 103732.

Kydland, F. E. and E. C. Prescott (1982). Time to Build and Aggregate Fluctuations. *Econometrica 50*(6), 1345–1370.

Lee, H. L., V. Padmanabhan, and S. Whang (1997a). Information Distortion in a Supply Chain: The Bullwhip Effect. *Management Science 43*(4), 546–558.

Lee, H. L., V. Padmanabhan, and S. Whang (1997b). The Bullwhip Effect In Supply Chains. *Sloan Management Review 38*(3), 93–102.

Leeper, E. M. (1991). Equilibria Under 'Active' and 'Passive' Monetary and Fiscal Policies. *Journal of Monetary Economics 27*(1), 129–147.

Leeper, E. M. (2023). Fiscal Dominance: How Worried Should We Be? *Mercatus Policy Brief Series*.

Leeper, E. M. and C. Leith (2016). Understanding Inflation as a Joint Monetary–Fiscal Phenomenon. In *Handbook of Macroeconomics*, Volume 2, pp. 2305–2415. Elsevier.

Leeper, E. M., N. Traum, and T. B. Walker (2017). Clearing up the Fiscal Multiplier Morass. *American Economic Review 107*(8), 2409–2454.

Leng, Y., E. Liu, Y. Ren, and A. Tsyvinski (2025). The Bullwhip: Time-to-Build and Sectoral Fluctuations. Technical report, NBER Working Paper No. w33638.

Liemen, M. O. and O. Posch (2022). FTPL and the Maturity Structure of Government Debt in the New-Keynesian Model. *Available at SSRN: 4157228*.

Lubik, T. A., B. Merone, and N. Robino (2024). Stargazing: Estimating r* in Other Countries. *Federal Reserve Bank of Richmond Economic Brief*.

Lucas, R. E. and N. L. Stokey (1983). Optimal Fiscal and Monetary Policy in an Economy without Capital. *Journal of Monetary Economics 12*(1), 55–93.

McCarthy, J. and E. Zakrajšek (2000). Microeconomic Inventory Adjustment: Evidence from US Firm-Level Data. *FRB of New York Staff Report* (101).

McCarthy, J. and E. Zakrajšek (2007). Inventory Dynamics and Business Cycles: What Has Changed? *Journal of Money, Credit and Banking 39*(2-3), 591–613.

McConnell, M. M. and G. Perez-Quiros (2000). Output Fluctuations in the United States: What Has Changed since the Early 1980's? *American Economic Review 90*(5), 1464–1476.

McKay, A. and C. Wolf (2022). Optimal Policy Rules in HANK. Working Paper, FRB Minneapolis.

McMahon, M. (2014). Inventories in Motion: A New Approach To Inventories Over The Business Cycle. Technical report, University of Warwick mimeo.

McMahon, M. and B. Wanengkirtyo (2015). Beyond Inventory Management: The Bullwhip Effect and the Great Moderation. Technical report, University of Oxford mimeo.

Meese, R. A. and K. Rogoff (1983). Empirical Exchange Rate Models of the Seventies: Do They Fit out of Sample? *Journal of International Economics 14*(1-2), 3–24.

Metzler, L. A. (1941). The Nature and Stability of Inventory Cycles. *The Review of Economics and Statistics 23*(3), 113–129.

Miao, J. and D. Su (2024). Fiscal and Monetary Policy Interactions in a Model with Low Interest Rates. *American Economic Journal: Macroeconomics 16*(4), 35–76.

Mierzwa, S. (2024). Spillovers from Tax Shocks to the Euro Area. *Oxford Economic Papers*, gpae024.

Mou, S., D. J. Robb, and N. DeHoratius (2018). Retail Store Operations: Literature Review and Research Directions. *European Journal of Operational Research 265*(2), 399–422.

Müller, K., C. Xu, M. Lehbib, and Z. Chen (2025). The Global Macro Database: A New International Macroeconomic Dataset. Working Paper 33714, National Bureau of Economic Research.

Nagel, S. (2016). The Liquidity Premium of Near-Money Assets. *The Quarterly Journal of Economics 131*(4), 1927–1971.

Nagel, S. and Z. Yan (2022). Inflation Hedging on Main Street? Evidence from Retail TIPS Fund Flows. Technical report, NBER Working Paper No. w30692.

Nakamura, E., V. Riblier, and J. Steinsson (2025). Beyond the Taylor Rule. *Presented at the Jackson Hole Economic Policy Symposium.*

Neumeyer, P. A. and J. P. Nicolini (2025). The Incredible Taylor Principle: A Comment. *Journal of Political Economy 133*(4), 1382–1399.

Niepelt, D. (2004). The Fiscal Myth of the Price Level. *The Quarterly Journal of Economics 119*(1), 277–300.

NIESR (2022). An Independent Assessment of the Mini-Budget - 23 September 2022. available online at `https://niesr.ac.uk/wp-content/uploads/2022/09/September-2022-Mini-Budget-Response.pdf`.

OBR (2024). Office of Budget Responsibility Historical Official Forecasts Database. Data retrieved from `https://obr.uk/data/`.

Pinter, G. (2023). An Anatomy of the 2022 Gilt Market Crisis. *Bank of England Working Paper No. 1019*.

Rachel, L. and M. Ravn (2025). Brothers in Arms: Monetary-Fiscal Interactions Without Ricardian Equivalence. *CEPR Discussion Paper no. 20445*.

Ramey, V. A. (2016). Macroeconomic Shocks and Their Propagation. *Handbook of Macroeconomics 2*, 71–162.

Ramey, V. A. (2019). Ten Years After the Financial Crisis: What Have We Learned from the Renaissance in Fiscal Research? *Journal of Economic Perspectives 33*(2), 89–114.

Ramey, V. A. and K. D. West (1999). Inventories. In *Handbook of Macroeconomics*, Volume 1, pp. 863–923. Elsevier.

Ramey, V. A. and S. Zubairy (2018). Government Spending Multipliers in Good Times and in Bad: Evidence from US Historical Data. *Journal of Political Economy 126*(2), 850–901.

Ravichandran, N. (2008). Managing Bullwhip Effect: Two Case Studies. *Journal of Advances in Management Research 5*(2), 77–87.

Reichlin, L., G. Ricco, and M. Tarbé (2023). Monetary–Fiscal Crosswinds in the European Monetary Union. *European Economic Review 151*, 104328.

Reis, R. (2022). Debt Revenue and the Sustainability of Public Debt. *Journal of Economic Perspectives 36*(4), 103–124.

Reiter, M. (2009). Solving Heterogeneous-Agent Models by Projection and Perturbation. *Journal of Economic Dynamics and Control 33*(3), 649–665.

Rogoff, K. (1996). The Purchasing Power Parity Puzzle. *Journal of Economic Literature 34*(2), 647–668.

Romelli, D. (2024). Trends in Central Bank Independence: A De-Jure Perspective. *BAFFI CAREFIN Centre Research Paper* (217).

Sargent, T. J. and N. Wallace (1981). Some Unpleasant Monetarist Arithmetic. *Federal Reserve Bank of Minneapolis Quarterly Review 5*(3), 1–17.

Sarte, P.-D., F. Schwartzman, and T. A. Lubik (2015). What Inventory Behavior Tells Us About How Business Cycles Have Changed. *Journal of Monetary Economics 76*, 264–283.

Schaal, E. and M. Taschereau-Dumouchel (2025). Echoes and Delays: Time-to-build in Production Networks. mimeo.

Schmid, L., V. Valaitis, and A. T. Villa (2026). Government Debt Management and Inflation with Real and Nominal Bonds. *Journal of Financial Economics, forthcoming*.

Schmitt-Grohé, S. and M. Uribe (2003). Closing Small Open Economy Models. *Journal of International Economics 61*(1), 163–185.

Sims, C. A. (1994). A Simple Model for Study of the Determination of the Price Level and the Interaction of Monetary and Fiscal Policy. *Economic Theory 4*, 381–399.

Sims, C. A. (1997). Fiscal Foundations of Price Stability in Open Economies. *Available at SSRN 75357*.

Sims, C. A. (1999). The Precarious Fiscal Foundations of EMU. *De Economist 147*(4), 415–436.

Sims, C. A. (2011). Stepping on a Rake: The Role of Fiscal Policy in the Inflation of the 1970s. *European Economic Review 55*(1), 48–56.

Sims, C. A. (2013). Paper Money. *American Economic Review 103*(2), 563–584.

Smets, F. and R. Wouters (2024). Fiscal Backing, Inflation and US Business Cycles. *CEPR Discussion Paper no. DP19791*.

Taylor, D. H. (1999). Measurement and Analysis of Demand Amplification Across the Supply Chain. *The International Journal of Logistics Management*.

Valchev, R. (2020). Bond Convenience Yields and Exchange Rate Dynamics. *American Economic Journal: Macroeconomics 12*(2), 124–166.

Van Binsbergen, J. H., W. F. Diamond, and M. Grotteria (2022). Risk-Free Interest Rates. *Journal of Financial Economics 143*(1), 1–29.

Varian, H. R. (1975). A Bayesian Approach to Real Estate Assessment. In *Studies in Bayesian Econometrics and Statistics in Honor of Leonard J. Savage*, pp. 195–208. North Holland.

Wang, X. and S. M. Disney (2016). The Bullwhip Effect: Progress, Trends and Directions. *European Journal of Operational Research 250*(3), 691–701.

Woodford, M. (1995). Price-Level Determinacy Without Control of a Monetary Aggregate. In *Carnegie-Rochester Conference Series on Public Policy*, Volume 43, pp. 1–46. Elsevier.

Woodford, M. (1998). Control of the Public Debt: A Requirement for Price Stability? In *The Debt Burden and its Consequences for Monetary Policy: Proceedings of a Conference held by the International Economic Association at the Deutsche Bundesbank, Frankfurt, Germany*, pp. 117–158. Springer.

Woodford, M. (2011). *Interest and Prices: Foundations of a Theory of Monetary Policy.* Princeton University Press.

Zarnowitz, V. (1961). The Timing of Manufacturers' Orders during Business Cycles. In *Business Cycle Indicators, Volume 1*, pp. 420–484. Princeton University Press.

Zarnowitz, V. (1973). Orders, Production, and Inventory Investment. In *Orders, Production, and Investment: A Cyclical and Structural Analysis*, pp. 344–408. NBER.

Zhu, T., J. Balakrishnan, and G. J. da Silveira (2020). Bullwhip Effect in the Oil and Gas Supply Chain: A Multiple-Case Study. *International Journal of Production Economics 224*, 107548.

Appendices

Chapter A

Appendix to: Debt Indexation, Determinacy, and Inflation

A.1 Additional Empirical Evidence

The distribution of government debt holdings across households

Publicly available microdata reinforces the idea of an unequal distribution of indexed debt in household portfolios. This brief section focuses on the U.S. due to the superior availability of household-level data on asset holdings.

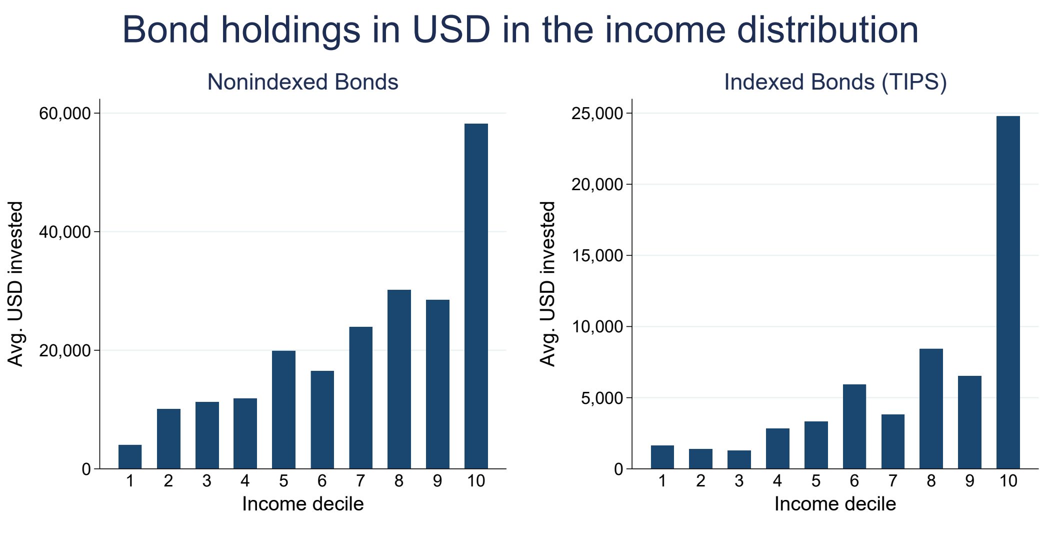

Figure A.1: Distribution of indexed and non-indexed debt holdings across household income deciles, denoted in real (2017) USD. Data source: Survey of Consumer Finances (U.S.); sample period: 2014-2019.

Figure A.1 plots the real (2017) Dollar value of nonindexed and indexed government debt holdings of households questioned in the U.S. Survey of Consumer Finances (SCF), separated by income deciles.[1] The left-hand panel of figure A.1 reflects the well-known left-skew of bond holdings of households in the income distribution, by which households at the upper end of the income distribution hold a significantly larger share of sovereign bonds. The right-hand panel of figure A.1 reflects a less well-known observation: this left skew is *vastly* more pronounced for indexed sovereign bonds, with the top income decile holding almost 40% of such bonds in the sample.

Figure A.2 provides further evidence in favour of the distribution implied by the model, which has not been a targeted moment. Just as in the model, the density of the asset distribution of both indexed and non-indexed bonds exhibits a significant skew, which is more pronounced overall for inflation-indexed debt. In particular, the size of the bins, even if not exactly matched, broadly reflects the distribution of the model very well.

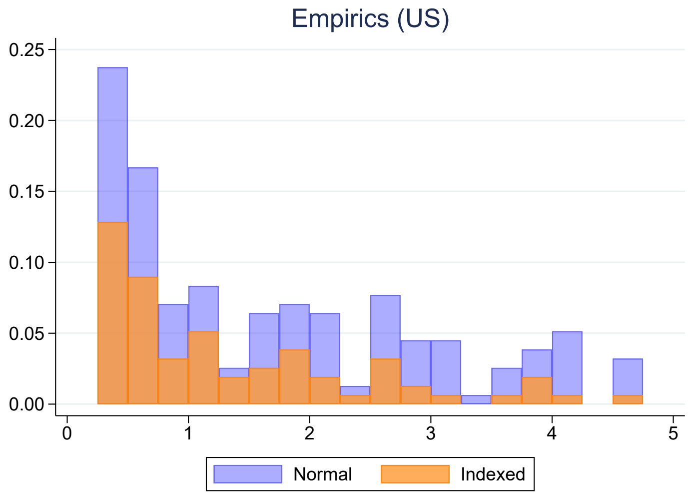

Figure A.2: Density of indexed and non-indexed debt holdings in the U.S. Survey of Consumer Finances; snapshot from October 2019.

Additional evidence on the September 2022 'mini-budget' shock

In this part of the appendix, we provide additional details on the fiscal policy measure dubbed "The Growth Plan" in September 2022 in the United Kingdom, commonly known as the 'mini-budget', and the reaction on financial markets to its announcement. To do so, we utilize ticker-frequency data on market prices (and thus yields), debt quantities, and risk perceptions. The data comes from the Bank of England (BoE), the U.K. Treasury, and Bloomberg Financial Services.

[1] We chose income deciles due to their clear definition in the survey with a single question. Constructing individual wealth variables is possible with the survey data, albeit this process is subject to particular choices about what to consider as household wealth. For most definitions of wealth, the results continue to hold qualitatively.

We begin by examining the degree to which the policy announcement can be informative about the propensity of a type of 'fiscally-led policy mix' in a wider sense, i.e., whether the policy measures around this particular fiscal shock can be placed in a context at which monetary policy passively adjusts to the fiscal policy measure, taking the fiscal announcement as given.[2] A possible measure that is plausibly related to debt sustainability concerns introduced by the budget announcement as well as to the prospective monetary reaction are expected overnight interest rates. These are the interest rates used for overnight bank lending activities on financial markets, instrumented using swaps on overnight lending between the day at question and the day of the next monetary policy meeting. Normally, these swaps follow the prevailing nominal interest rate closely (with a spread of a couple of basis points), as any other rate would induce arbitrage by the possibility of a risk-free hedge using the current overnight nominal interest rate. As figure A.3 shows, however, the turmoil introduced by the 'mini-budget' caused a remarkable wedge between the two rates:

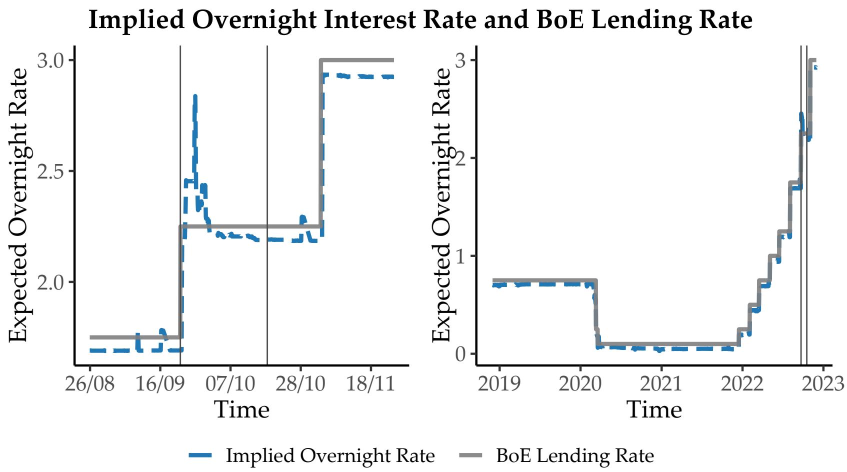

Figure A.3: The prevailing Bank of England Lending Rate and the Implied Overnight Interest Rate, derived by instrumenting overnight interest swaps from today to the expected next meeting of the BoE Monetary Policy Committee.

As can be inferred from the right-hand side panel for a period of five years, and on the left-hand side in more detail for the period of interest, the implied overnight interest rate follows the BoE lending rate closely, exhibiting jumps around meeting dates of the BoE Monetary Policy Committee in alignment with monetary policy decisions.

[2]Determining uniquely whether a given policy announcement, or a given time period, clearly relates to a monetary-led or a fiscally-led policy mix in a narrow sense, i.e., in relation to the respective policy rules and how they inform the stability of the underlying economic system, is generally not possible purely based off time-series data. Simply put, the "Taylor Principle" cannot be tested as its impact on the uniqueness properties surrounding macroeconomic models depends on off-equilibrium threats that cannot be observed under the condition of the Taylor Principle itself holding (Cochrane, 2011; Neumeyer and Nicolini, 2025).

The period of the mini-budget, which commenced one day after a Monetary Policy Committee (MPC) meeting (September 23 and September 22, respectively), induced movements in the expected overnight rates that were not observed at any other point in time - despite no MPC meeting in near sight.[3] Expected overnight rates shot up far beyond the then-prevailing BoE bank lending rate by up to 50 basis points. Such movements can be caused by an array of different possibilities: it could be either that fiscal policy caused a shift in market expectations of monetary policy in the short-term, thus implying that monetary policy was considered to be 'reactive' to the fiscal policy announcement, or that the mini-budget was expected to have such detrimental consequences on inflation that the BoE was required to react immediately, or it might be reflective of liquidity issues in the swap market in the same period.

An important caveat is that reducing the observed dynamics to expected revaluations of bonds and prospective interest rate movements does not capture all aspects related to this fiscal policy announcement. Uncertainty surrounding the proposed policy measures might have also been an important contributor to market reactions. Figure A.4 plots the FTSE 100 IVI Index that can plausibly serve as a proxy for uncertainty.

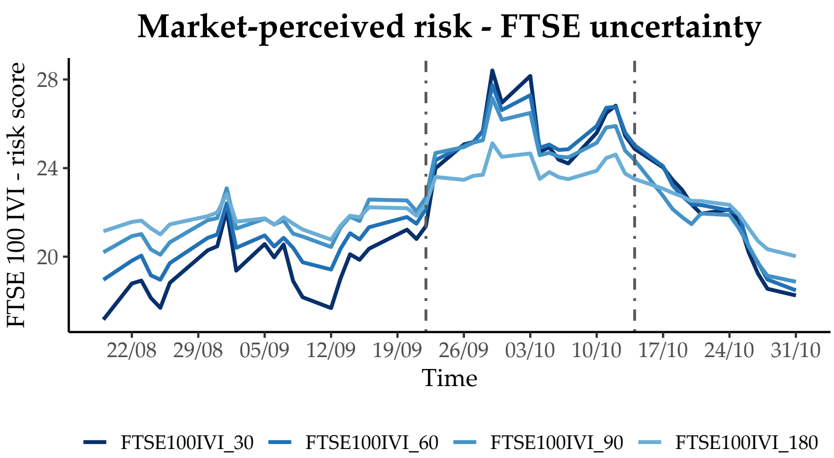

Figure A.4: The uncertainty index over equity of the largest publicly traded British companies. The lines measure implied uncertainty in 30-day, 60-day, 90-day, and 180-day forward-looking windows.

The mini-budget episode coincided with an inversion of the market-perceived risk relative to the forecast horizon: whereas in the periods before and after the mini-budget announcement market risk was perceived to be higher in the medium-term (180 days)

[3] On September 27, then-BoE chief economist Huw Pill stated that the proposed U.K. government budget might require a "significant monetary response", indicating readiness on behalf of the BoE to adjust the monetary stance, but no concrete emergency meeting date had been proposed at that point.

than in the short-term (30 or 60 days), the opposite has been the case during this short-lived fiscal episode.

A final and related aspect is the possibility of elevated default risk, which has been omitted from the model presented in this chapter. In figure A.5, we present the Credit Default Swap spread alongside the quantity of Gilts maturing in the near-term around September 2022, which plausibly indicate near-term fiscal refinancing pressures.

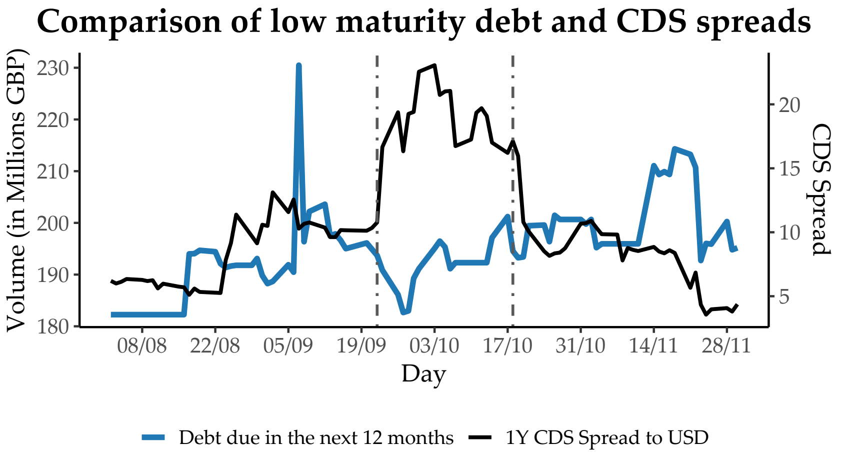

Figure A.5: Government debt maturing over the course of the next 12 months, plotted against the 1-year spread of Credit Default Swaps on U.K. Gilts over USD futures. Data source: Bloomberg.

During the period of the 'mini-budget' announcement, the amount of maturing government debt was at rather low levels, i.e., there was no inherent pressure to refinance a large quantity of maturing obligations around the time of the 'mini-budget' announcement. Yet, the black line, which plots short-term Credit Default Swaps based off U.K. Gilts, was elevated during the 'mini-budget' episode, confirming a disconnect of the total amount of maturing debt from perceived market risk. While the two measures have a correlation of 0.49 in the period leading up to the announcement of the mini-budget, that correlation drops to -0.51 between the middle of September and the middle of October 2022.

All these data points reinforce the idea that the 'mini-budget' announced by the U.K. Treasury in September 2022 was indeed an unexpected fiscal measure that significantly subverted perceived fiscal sustainability, with wide-spread ramifications for expected real returns on government bonds, expected inflation and interest rates, and elevated risk levels.

FAQ: the 'mini-budget shock' episode

In addition to the evidence derived using market data, we here present narrative viewpoints complementing the understanding of the 'mini-budget' episode.

Why did markets possibly reverse the uptick in inflation expectations initially?

- On September 26, around 4.00pm, Kwasi Kwarteng announced to publish a 'medium-term fiscal plan', which possibly indicated greater restraint in fiscal policy: (Bloomberg).
- On September 28, a further plausible shock to perceptions of fiscal sustainability occurred: Moody's explicitly deemed the mini-budget to put U.K. debt sustainability in danger, followed by a same-day increase of inflation expectations: (Reuters).
- Likewise, on September 28, in a reversal of expectations caused by the September 26 statement, the Treasury explicitly rejected for the first time since the initial announcement any idea of reneging on the additional budget shortfall, thereby re-affirming expectations about the fiscal policy measure actually being pushed through. See: (BBC). On the very same day, the Bank of England intervened in bond markets by re-starting long-dated government bond purchases, which, again, plausibly re-affirmed the idea that the Treasury will not back down. This was announced at around 11.20am - see: (X). Swap breakeven rates shot up by 60bps in the three hours after the statement made by the Bank of England.

Why were inflation swaps priced much higher in August 2022 compared to the dynamics occurring in September and October 2022?

- On August 17, 2022, the ONS released a report of CPI inflation being 10.1%, breaking the 10% barrier for the first time in 40 years, also beating the private sector forecasts decisively (Bloomberg). This occurred alongside a significant depreciation of the British Pound (FT). Likewise, implied interest rate raises corresponding to expectations of vastly more aggressive monetary tightening from that period onward alongside a yield curve inversion appeared around August 15 (FT).
- Implied one-year ahead inflation expectations peaked at around 8% in August. This is still *vastly* below the forecasts released in August 2022 by major financial market actors: the Goldman Sachs forecast of one-year ahead inflation amounted to 14.8%, with a 'negative' scenario of 22.1% annual inflation for the U.K. implied in their August 2022 briefings (FT). Relative to that forecast, the change in inflation swaps implied in that policy uncertainty episode was relatively benign. This period of increased inflation expectations also coincided with record prices on natural gas spot markets in the U.K..

Further details on the Local Projection exercise

To shed further light on the evidence presented in section 1.2, we here provide the additional information related to the local projection exercise on U.K. data presented in figure 2.12 and introduce additional evidence using U.S. data in a similar exercise.

Figure A.6 plots the time series of the level and the first-difference of indexed debt in the U.K., showing the secular increase in the share of indexed debt in the sovereign bond portfolio since 1980.

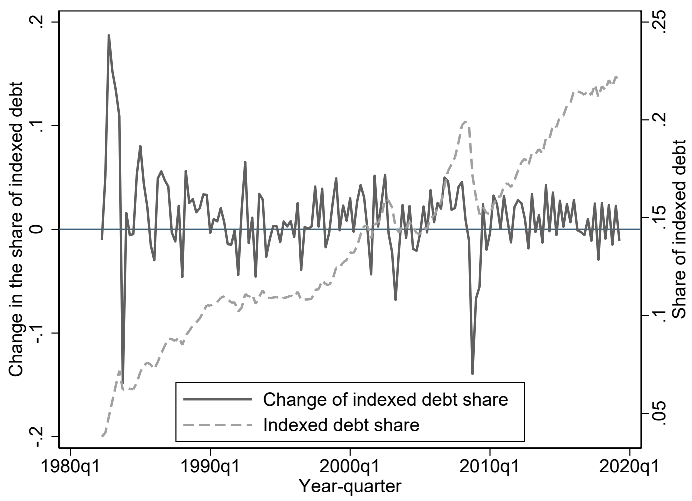

Figure A.6: Time series of the level of the share of inflation-indexed debt in the U.K. as well as the first-difference thereof.

Table A.1 gives the concrete numbers corresponding to figure 2.12, specifying the exact coefficients of the interaction effect of $\Delta\omega_t \, \varepsilon_t$ and the individual effects of the change in the indexed debt share $\Delta\omega_t$ and the exogenous fiscal shock ε_t on the cumulative price level change from the pre-shock period -1 until the period specified.

While the share of indexed debt itself does not impact medium-term inflation, the interaction effect of the share of indexed debt with the identified fiscal shock follows the pattern given in figure 2.12.

To ensure that this mechanism is not idiosyncratic to the U.K., we provide now the results of a similar exercise applied to the U.S. Here, we utilize the U.S. fiscal shock series provided by Mierzwa (2024). We leverage the identified fiscal shocks and estimate the same local projection specification (equation (1.2)) to estimate the role played by inflation-indexed debt in exacerbating the effects of fiscal spending shocks in the U.S. Table A.2 and figure A.7 summarize the results of this exercise using data since 1980, which is the earliest period for which identified fiscal shocks are available.

	Dependent variable: log(Cumulative Inflation)							
Forecast period:	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Fiscal Shock	-0.01**	-0.01	-0.01	-0.01	-0.02	-0.02	-0.02	-0.02
	(0.00)	(0.01)	(0.01)	(0.01)	(0.01)	(0.02)	(0.02)	(0.02)
Index Share	0.02***	0.02**	0.03**	0.02	0.02	0.02	0.03	0.03
	(0.00)	(0.01)	(0.01)	(0.01)	(0.02)	(0.02)	(0.02)	(0.02)
Fiscal Shock × Index Share	0.10*	0.09	0.20	0.26*	0.39**	0.40*	0.60**	0.81***
	(0.06)	(0.10)	(0.13)	(0.14)	(0.18)	(0.21)	(0.24)	(0.27)
Additional Controls	Y	Y	Y	Y	Y	Y	Y	Y
Observations	155	154	153	152	151	150	149	148
R^2	0.412	0.518	0.559	0.630	0.592	0.599	0.575	0.602

Table A.1: Local Projection results for the U.K. The controls include the first four lags of the real GDP growth rate, the Bank of England Bank Rate, the change in the weighted real exchange rate, a same-period recession indicator, year-fixed effects, and the first lag of the price level difference. Standard errors are robust to heteroskedasticity and autocorrelation.

	Dependent variable: log(Cumulative Inflation)							
Lag periods:	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Fiscal Shock	0.00	0.00	-0.00	-0.00	-0.00	-0.00	-0.00	-0.00
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.01)	(0.01)	(0.01)
Index Share	0.00	-0.01	-0.01	-0.01	-0.00	-0.00	-0.01	-0.01
	(0.01)	(0.01)	(0.01)	(0.01)	(0.02)	(0.02)	(0.02)	(0.02)
Fiscal Shock × Index Share	0.02	0.01	0.06	0.15**	0.19***	0.21**	0.22**	0.24***
	(0.03)	(0.08)	(0.08)	(0.06)	(0.07)	(0.09)	(0.10)	(0.09)
Additional Controls	Y	Y	Y	Y	Y	Y	Y	Y
Observations	161	160	159	158	157	156	155	154
R^2	0.324	0.371	0.474	0.531	0.543	0.542	0.559	0.554

Table A.2: Local Projection results for the U.S. The controls include the first four lags of the real GDP growth rate, the Federal Funds Rate, the change in the weighted real exchange rate, a same-period recession indicator, year-fixed effects, and the first lag of the price level difference. Standard errors are robust to heteroskedasticity and autocorrelation.

The results here paint a supporting picture, as the interaction effect between the change in the share of inflation-indexed debt and the identified fiscal shock appears to be statistically significant in the medium-term again, though the level of the effect is smaller than in the U.K.

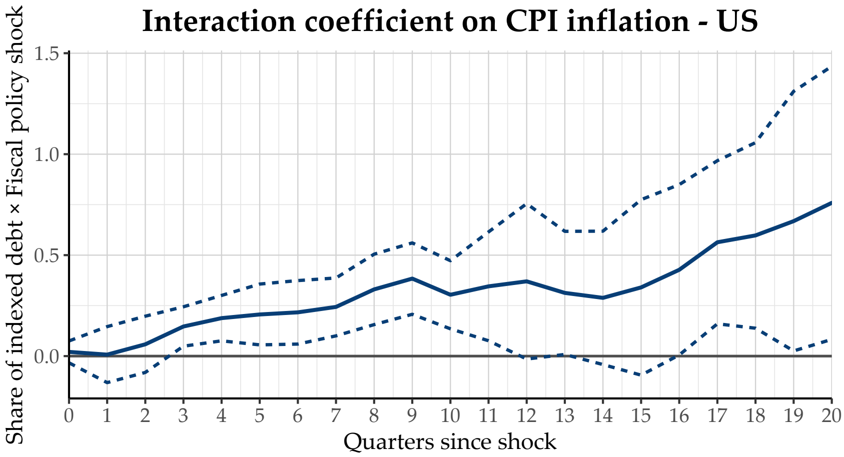

Figure A.7: IRFs implied by a local projection in the style of equation (1.2). The controls include the first four lags of the real GDP growth rate, the Federal Funds Rate, the change in the weighted real exchange rate, a same-period recession indicator, year-fixed effects, and the first lag of the price level difference. Confidence intervals are provided at the 90% level. Sample length: 1980 Q1 - 2019 Q4.

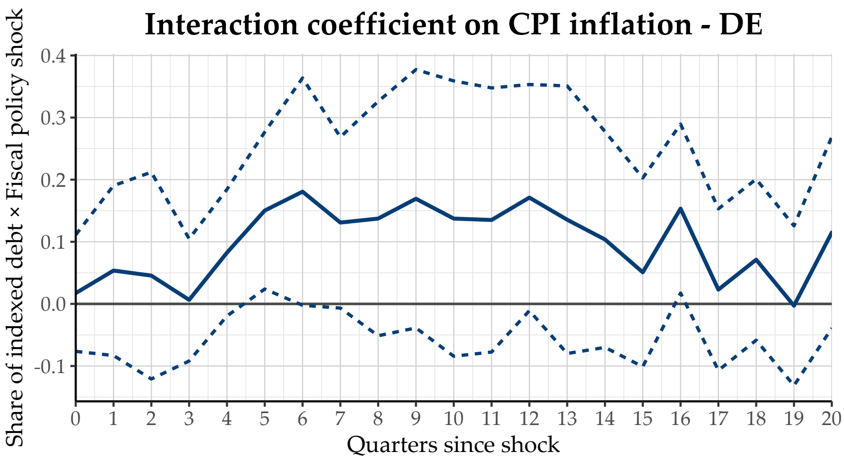

Figure A.8: IRFs implied by a local projection in the style of equation (1.2) for Germany. The controls include the first four lags of the real GDP growth rate, the prevalent policy rate, the change in the weighted real exchange rate, a same-period recession indicator, year-fixed effects, and the first lag of the price level difference. Confidence intervals are provided at the 90% level. Sample length: 1980 Q1 - 2019 Q4.

Finally, we can also leverage the shocks for the third country provided in the analysis of Mierzwa (2024), which is Germany. Germany introduced inflation-linked bonds in 2006, with the share of inflation-indexed debt in the overall sovereign debt portfolio amounting to around 7% at the beginning of 2020. Here, the interaction coefficient on the lagged share of inflation-indexed debt and the identified tax-side fiscal policy shock depicts dynamics similar to the U.K. and the U.S., although the results are overall less significant due to the smaller number of sizeable tax shocks in the years in which a positive quantity of indexed debt had been issued.[4]

A.2 Derivations and Proofs from the Main Text

Derivations from section 1.4

On the transversality condition with uninsurable idiosyncratic risk

We first show why it is not possible to directly arrive at a government debt valuation equation starting from aggregate debt quantities, following the logic of Brunnermeier et al. (2024).

To illustrate this point, start from the government budget constraint.

$$B_{t-1} + \Pi_t b_{t-1} = P_t s_t + Q_t B_t + q_t b_t.$$

This is the government budget constraint given some (real) surplus schedule s_t and bond pricing kernels Q_t, q_t.[5] All elements are multiplied by the unweighted average household SDF $\mathcal{M}_{t,t+1}$ and divided by the current price level P_t to obtain

$$\mathcal{M}_{t,t+1}\frac{B_{t-1}}{P_t} + \mathcal{M}_{t,t+1}\frac{b_{t-1}}{P_{t-1}} = \mathcal{M}_{t,t+1}s_t + Q_t\mathcal{M}_{t,t+1}\Pi_{t+1}\frac{B_t}{P_{t+1}} + q_t\mathcal{M}_{t,t+1}\frac{b_t}{P_t}.$$

Adding and subtracting elements suitably on the right-hand side, we can express this equation as:

$$\begin{aligned}\mathcal{M}_{t,t+1}\frac{B_{t-1}}{P_t} + \mathcal{M}_{t,t+1}\frac{b_{t-1}}{P_{t-1}} = {} & \mathcal{M}_{t,t+1}s_t + \left(Q_t\mathcal{M}_{t,t+1}\Pi_{t+1} - \mathcal{M}_{t+1,t+2}\right)\frac{B_t}{P_{t+1}} \\ & + \left(q_t\mathcal{M}_{t,t+1} - \mathcal{M}_{t+1,t+2}\right)\frac{b_t}{P_t} + \mathcal{M}_{t+1,t+2}\left(\frac{B_t}{P_{t+1}} + \frac{b_t}{P_t}\right).\end{aligned}$$

Iterating on this expression until T, dividing the resulting expression by the SDF, and taking limits $T \to \infty$ ultimately gives:

$$\begin{aligned}\frac{B_{t-1}}{P_t} + \frac{b_{t-1}}{P_{t-1}} = {} & \mathbb{E}_t\left[\sum_{l=0}^{\infty}\frac{\mathcal{M}_{t+l,t+l+1}}{\mathcal{M}_{t,t+1}}s_{t+l} + \frac{Q_{t+l}\mathcal{M}_{t+l,t+l+1}\Pi_{t+l+1} - \mathcal{M}_{t+l+1,t+l+2}}{\mathcal{M}_{t,t+1}}\frac{B_{t+l}}{P_{t+l+1}}\right. \\ & \left. + \frac{q_{t+l}\mathcal{M}_{t+l,t+l+1} - \mathcal{M}_{t+l+1,t+l+2}}{\mathcal{M}_{t,t+1}}\frac{b_{t+l}}{P_{t+l}}\right] + \lim_{T\to\infty}\frac{\mathcal{M}_{T+1,T+2}}{\mathcal{M}_{t,t+1}}\left(\frac{B_T}{P_{T+1}} + \frac{b_T}{P_T}\right).\end{aligned} \tag{A.1}$$

[4] As Germany is part of the European Monetary Union since 1999, this type of analysis is furthermore convoluted by the fact that the monetary union hinders the interpretation of any type of non-indexed debt as purely nominal debt, as argued by Sims (2013).

[5] All debt in this model is single-period. We briefly expose the effects of long-term debt in appendix A.4.

This expression nests the usual debt valuation equation under complete markets, making use of $Q_t \mathcal{M}_{t,t+1}\Pi_{t+1} = \mathcal{M}_{t+1,t+2}$ and $q_t \mathcal{M}_{t,t+1} = \mathcal{M}_{t+1,t+2}$.

It would be a mistaken belief that the current price level is determined by this integrated government budget constraint. This logic would require the last limiting term to vanish and go to zero. However, this is *not* necessarily the case: even though the transversality condition holds on the household level as a consequence of household optimality and a no-Ponzi condition, it *cannot* be aggregated to derive the aggregate transversality condition directly off-the-shelf. The reason is that the unweighted average SDF $\mathcal{M}_{t,t+1}$ is discarding the heterogeneity of underlying consumption (which led to the rise of household-specific discount factors), and thus ignores the possibility of the government possibly earning an excess return on its debt issuance. This can be considered a 'safe asset premium' (Brunnermeier et al., 2024) and is reflective of the inherent value that such debt bears to households in partially overcoming market incompleteness, possibly yielding different 'fundamental' valuations of government debt by the household vis-à-vis the government.

However, assuming additionally perfect insurance of idiosyncratic risk, it would be possible to define a simple average SDF that is consistent across household and government valuations,

$$\bar{\mathcal{M}}_{t,t+1} \equiv \sum_i \frac{B_{it} + b_{it}}{B_t + b_t} \mathcal{M}_{i,t,t+1},$$

under which the final limiting term and all wedges would vanish due to a fair bond pricing valuation, creating a 'standard formulation' of the government debt valuation equation:

$$\frac{B_{t-1}}{P_t} + \frac{b_{t-1}}{P_{t-1}} = \mathbb{E}_t \left[\sum_{l=0}^{\infty} \frac{\bar{\mathcal{M}}_{t+l,t+l+1}}{\bar{\mathcal{M}}_{t,t+1}} s_{t+l} \right].$$

Derivation of equation (1.9) (proof of proposition 1)

This section presents the derivations underlying a *dynamic trading perspective* for asset valuation laid out in Brunnermeier et al. (2024), which avoids fallacies related to a possibly non-existent aggregate transversality condition by defining government debt valuation recursively at the household level and then aggregating across households. This allows the leveraging of household-level transversality conditions to derive an aggregate government debt valuation equation that only holds for one initial candidate price level.

The starting point for the valuation equation of government debt is the household budget constraint, which is given by

$$P_t c_{it} + Q_t B_{it} + q_t b_{it} = \varepsilon_{it}(1 - \tau_{it}) P_t w_t N_t + B_{i,t-1} + \Pi_t b_{i,t-1}$$

for each household i. Following the results derived in the household block, let households price bonds in accordance with their stochastic discount factor:

$$B_{i,t-1} + \Pi_t b_{i,t-1} = \mathbb{E}_t\left(\mathcal{M}_{i,t,t+1}\right) B_{it} + \mathbb{E}_t\left(\Pi_{t+1}\mathcal{M}_{i,t,t+1}\right) b_{it} + P_t\left(c_{it} - \varepsilon_{it} w_t N_t (1-\tau_{it})\right).$$

Dividing all elements by P_t yields

$$\frac{B_{i,t-1} + \Pi_t b_{i,t-1}}{P_t} = \mathbb{E}_t\left(\mathcal{M}_{i,t,t+1}\right)\frac{B_{it}}{P_t} + \mathbb{E}_t\left(\Pi_{t+1}\mathcal{M}_{i,t,t+1}\right)\frac{b_{it}}{P_t} + \left(c_{it} - \varepsilon_{it}(1-\tau_{it}) w_t N_t\right).$$

To obtain an expression that can be iterated correctly, define the real ex-post value of the household's government bond holdings at time t as

$$X_{it} \equiv \frac{B_{i,t-1} + \Pi_t b_{i,t-1}}{P_t}.$$

Using the identity

$$\Pi_{t+1} X_{i,t+1} = \Pi_{t+1}\frac{B_{it} + \Pi_{t+1} b_{it}}{P_{t+1}} = \frac{B_{it}}{P_t} + \Pi_{t+1}\frac{b_{it}}{P_t},$$

the household budget constraint at optimality can be rewritten as

$$X_{it} = \left(c_{it} - \varepsilon_{it}(1-\tau_{it}) w_t N_t\right) + \mathbb{E}_t\left[\mathcal{M}_{i,t,t+1}\Pi_{t+1} X_{i,t+1}\right]. \tag{A.2}$$

Equation (A.2) is the correct recursive representation of the household valuation equation. It is crucial that the inflation term Π_{t+1} remains inside the conditional expectation together with the stochastic discount factor, so that the object being iterated is the full next-period real payoff of the household's government bond portfolio.

Now iterate forward. Substituting out $X_{i,t+1}$ one period ahead gives

$$\begin{aligned} X_{it} = {} & \left(c_{it} - \varepsilon_{it}(1-\tau_{it}) w_t N_t\right) \\ & + \mathbb{E}_t\left[\mathcal{M}_{i,t,t+1}\Pi_{t+1}\left(c_{i,t+1} - \varepsilon_{i,t+1}(1-\tau_{i,t+1}) w_{t+1} N_{t+1}\right)\right] \\ & + \mathbb{E}_t\left[\mathcal{M}_{i,t,t+1}\Pi_{t+1}\mathbb{E}_{t+1}\left(\mathcal{M}_{i,t+1,t+2}\Pi_{t+2} X_{i,t+2}\right)\right]. \end{aligned}$$

Applying the Law of Iterated Expectations and using the identity

$$\mathcal{M}_{i,t,t+k}\mathcal{M}_{i,t+k,t+k+\ell} = \mathcal{M}_{i,t,t+k+\ell}$$

for all admissible t, k, ℓ, together with the analogous multiplicative notation for cumulative gross inflation,

$$\Pi_{t+1,t+k} \equiv \prod_{j=1}^{k} \Pi_{t+j}, \qquad \Pi_{t+1,t} \equiv 1,$$

yields

$$X_{it} = \mathbb{E}_t \left[\sum_{k=0}^{\infty} \mathcal{M}_{i,t,t+k} \Pi_{t+1,t+k} \left(c_{i,t+k} - \varepsilon_{i,t+k}(1-\tau_{i,t+k}) w_{t+k} N_{t+k} \right) \right] + \lim_{T\to\infty} \mathbb{E}_t \left[\mathcal{M}_{i,t,t+T} \Pi_{t+1,t+T} X_{i,t+T} \right]. \tag{A.3}$$

This expression is the integrated household budget constraint at optimality, from which the integrated *government* budget constraint is derived.

Crucially, household optimality implies a household-level transversality condition, while a no-Ponzi condition on household debt holdings rules out explosive debt accumulation. Therefore, the final limit term in (A.3) converges to zero, so that

$$\frac{B_{i,t-1} + \Pi_t b_{i,t-1}}{P_t} = \mathbb{E}_t \left[\sum_{k=0}^{\infty} \mathcal{M}_{i,t,t+k} \Pi_{t+1,t+k} \left(c_{i,t+k} - \varepsilon_{i,t+k}(1-\tau_{i,t+k}) w_{t+k} N_{t+k} \right) \right]. \tag{A.4}$$

The formulation of equation (A.4) is intuitive: the real value of household bond holdings is equal to the expected discounted stream of future net resources that the government can raise from that household. The presence of both non-indexed and inflation-indexed debt enters through the correct inflation-adjusted discounting term $\mathcal{M}_{i,t,t+k}\Pi_{t+1,t+k}$, which prices the full real payoff of the household's bond portfolio.

Aggregating the household-level constraints and making use of the asset market clearing conditions

$$B_t = \sum_i B_{it}, \qquad b_t = \sum_i b_{it},$$

yields

$$\frac{B_{t-1} + \Pi_t b_{t-1}}{P_t} = \sum_i \mathbb{E}_t \left[\sum_{k=0}^{\infty} \mathcal{M}_{i,t,t+k} \Pi_{t+1,t+k} \left(c_{i,t+k} - \varepsilon_{i,t+k}(1-\tau_{i,t+k}) w_{t+k} N_{t+k} \right) \right]. \tag{A.5}$$

Since

$$\frac{\Pi_t b_{t-1}}{P_t} = \frac{b_{t-1}}{P_{t-1}},$$

the left-hand side can be rewritten as

$$\frac{B_{t-1}}{P_t} + \frac{b_{t-1}}{P_{t-1}}.$$

To simplify the integrated government debt valuation equation, define the variable A_{it} as

$$A_{it} \equiv c_{it} - \varepsilon_{it}(1-\tau_{it}) w_t N_t,$$

which captures the surpluses raised by the government from each household i. Define

$$\bar{A}_t \equiv \frac{1}{N_i} \sum_i A_{it}.$$

Then equation (A.5) can be rewritten as

$$\frac{B_{t-1}}{P_t} + \frac{b_{t-1}}{P_{t-1}} = \mathbb{E}_t \left[\sum_{k=0}^{\infty} \left(\sum_i \mathcal{M}_{i,t,t+k} \Pi_{t+1,t+k} \frac{A_{i,t+k}}{\bar{A}_{t+k}} \right) \bar{A}_{t+k} \right].$$

Defining the *household value-weighted stochastic discount factor*

$$\tilde{\mathcal{M}}_{t,t+k} \equiv \sum_i \mathcal{M}_{i,t,t+k} \Pi_{t+1,t+k} \frac{A_{i,t+k}}{\bar{A}_{t+k}}$$

gives the final government debt valuation equation:

$$\frac{B_{t-1}}{P_t} + \frac{b_{t-1}}{P_{t-1}} = \mathbb{E}_t \left[\sum_{k=0}^{\infty} \tilde{\mathcal{M}}_{t,t+k} \bar{A}_{t+k} \right]. \tag{A.6}$$

Equation (A.6) is the government debt valuation equation, at times called the FTPL equation, that informs the price level at time t, given the inherited nominal and indexed debt positions from period $t-1$. The crucial difference relative to the incorrect derivation is that the insurance properties of indexed debt are already embedded in the inflation-adjusted discount factor $\mathcal{M}_{i,t,t+k}\Pi_{t+1,t+k}$, rather than being separated into additional covariance and surprise-inflation terms inside A_{it}.

Proof of proposition 2

We first show that a unique steady-state can be attained even with endogenous real interest rates when indexed debt is present. Complementing this proof, we show how indexed debt translates into a model where taxation is assumed to cover all interest expenses over time on a stationary equilibrium path, following Hagedorn (2021). We therefore maintain a 'true Balanced Growth Path' (BGP) with a constant real value of the debt portfolio thanks to an appropriate taxation schedule.

To apply the ideas of Hagedorn (2021), we start with his steady-state taxation function, but rewrite it to account for possible non-zero steady-state inflation and some positive level of indexed debt, since the presence of both changes the nominal value of taxation over time. The aim of this step is to find an asset demand function that depends only on model primitives and allows the derivation of the asset market equilibrium.[6] Doing so requires pinning down steady-state asset demand under incomplete markets in closed-form, for which we leverage the results of Acemoglu and Jensen (2015).

To find the steady-state level of taxation consistent with the bond issuance schedule that keeps the real value of bonds constant, we begin with an arbitrary per-period

[6] For the sake of completeness, we here specify the approach Hagedorn (2021) takes to determine steady-state taxation. He specifies the per-period government budget constraint as $B_{t+1} = (1+i_t)B_t - T_t \Leftrightarrow T_t = (1+i_t)B_t - B_{t+1}$ to arrive in steady-state at $T_{ss} = i_{ss}S_{ss}$, where S_{ss} is steady-state asset demand. In real terms, $t_{ss} =: \frac{T_{ss}}{P_{ss}} = r_{ss}S_{ss}$.

government budget constraint (setting $G_t = 0$, such that real surpluses are $s_t = t_t$, or, in nominal terms, $P_t s_t = P_t t_t =: T_t$):

$$B_{t-1} + \frac{P_t}{P_{t-1}} b_{t-1} = T_t + Q_t B_t + q_t b_t.$$

Q_t and q_t must be equal to some constant values in steady-state. Without aggregate uncertainty, the bond prices arising through asset demand must solely depend on the offered interest rates, since cross-sectional risks average out. Dividing the budget constraint by P_t to express it in real terms:

$$\frac{B_{t-1}}{P_t} + \frac{b_{t-1}}{P_{t-1}} = t_t + Q_t \frac{B_t}{P_t} + q_t \frac{b_t}{P_t}.$$

Writing $\frac{B_{t-1}}{P_t} = \frac{1}{\Pi_t} \frac{B_{t-1}}{P_{t-1}}$ and evaluating *in steady-state* (where the real values B_{ss}/P_{ss} and b_{ss}/P_{ss} are constant, which we henceforth denote by B_{ss} and b_{ss}, respectively):

$$\frac{1}{\Pi_{ss}} B_{ss} + b_{ss} = t_{ss} + \frac{1}{1+i_{ss}} B_{ss} + \frac{1}{1+r_{ss}} b_{ss}$$
$$\Leftrightarrow t_{ss} = \left(\frac{1}{\Pi_{ss}} - \frac{1}{1+i_{ss}} \right) B_{ss} + \left(1 - \frac{1}{1+r_{ss}} \right) b_{ss}.$$

Using the Fisher equation, $\frac{1}{\Pi_{ss}} - \frac{1}{1+i_{ss}} = \frac{1}{\Pi_{ss}} \left(1 - \frac{\Pi_{ss}}{1+i_{ss}} \right) = \frac{1}{\Pi_{ss}} \left(1 - \frac{1}{1+r_{ss}} \right) = \frac{1}{\Pi_{ss}} \cdot \frac{r_{ss}}{1+r_{ss}} = \frac{r_{ss}}{1+i_{ss}}$, and therefore:

$$t_{ss} = \frac{r_{ss}}{1+i_{ss}} B_{ss} + \frac{r_{ss}}{1+r_{ss}} b_{ss}.$$

Define by $S_t \left(\Omega_t, \{1 + r_l, t_l\}_l^\infty \right)$ the cumulative asset demand function under incomplete markets, which depends on the household distribution of wealth Ω_t, real interest rates $1+r_t$, and tax levels t_t, and is well-defined under standard regularity conditions (Acemoglu and Jensen, 2015). To relate steady-state taxation more clearly to gross asset demand, fix the shares of B_{ss} and b_{ss} of gross asset demand S_{ss} in steady-state. Denoting by θ the share of indexed debt b_{ss} in the steady-state asset portfolio, the taxation term in steady-state becomes

$$t_{ss} = \left[(1-\theta) \frac{r_{ss}}{1+i_{ss}} + \theta \frac{r_{ss}}{1+r_{ss}} \right] S_{ss}.$$

Under such steady-state taxes, the gross asset demand function arising from heterogeneous household demand ($S_{t+1} = \mathcal{S}(\Omega_t; 1+r_t, 1+r_{t+1}, 1+r_{t+2}, ...; t_t, t_{t+1}, ...)$) simplifies to the following mapping in steady-state:

$$S_{ss} = \mathcal{S} \left(\Omega_{ss}; 1+r_{ss}, 1+r_{ss}, 1+r_{ss}, ...; \right.$$
$$\left. \left[(1-\theta) \frac{r_{ss}}{1+i_{ss}} + \theta \frac{r_{ss}}{1+r_{ss}} \right] S_{ss}, \left[(1-\theta) \frac{r_{ss}}{1+i_{ss}} + \theta \frac{r_{ss}}{1+r_{ss}} \right] S_{ss}, ... \right).$$

With i_{ss} being equal to some constant set by the monetary policymaker in steady-state and the taxation function just derived, asset demand is derived by finding the fixed point of the above equation, which yields asset demand as a function of the real interest rate r_{ss}, following Acemoglu and Jensen (2015):

$$\underline{\text{Asset demand:}} \qquad S(r).$$

By the previous derivations, we now directly leverage asset supply in real terms as the left-hand side of the derivations of the asset market equation evaluated in steady-state, such that the stationary asset market equilibrium must be pinned down by

$$S(r) = \frac{B}{\tilde{P}} + \frac{b}{\tilde{P}(1+\pi_{ss})},$$

or, making use of the Fisher equation,

$$S(r) = \frac{B}{\tilde{\tilde{P}}} + \frac{b}{\tilde{\tilde{P}}} \frac{(1+r_{ss})}{(1+i_{ss})}.$$

Pinning down steady-state inflation through BGP consistency on the issuance schedule: Solving the asset market equilibrium requires taking a stance on the source of π_{ss}, the (possibly) non-zero steady-state inflation rate in this economy. The steady-state taxation function above was derived under the requirement that the real bond stocks B_{ss}/P_{ss} and b_{ss}/P_{ss} remain constant, and following Hagedorn (2021) we maintain that real taxes are also constant. Since b_t denotes the nominal face value of indexed debt set at issuance, constant real values of *both* debt instruments require that each nominal stock grows precisely at the gross inflation rate:

$$\frac{B_{t+1}}{B_t} = \frac{b_{t+1}}{b_t} = \frac{P_{t+1}}{P_t} = 1 + \pi_{ss}.$$

In net terms, this is the BGP consistency condition

$$g_B = g_b = \pi_{ss},$$

which we impose as an explicit restriction on the bond-issuance schedule (the assumption $g_B = g_b$ in the proposition statement). Constant real taxes follow as a corollary: with both nominal stocks growing at the gross rate $1 + \pi_{ss}$, nominal taxation $T_t = P_t \cdot t_{ss}$ inherits the same growth rate, so that

$$\frac{T'}{T} = 1 + \pi_{ss}, \qquad \text{equivalently} \qquad \pi_{ss} = \frac{T' - T}{T},$$

where variables with a prime denote next-period values. Note that this condition relates the issuance and inflation rates but does *not* on its own pin down π_{ss}; the steady-state inflation rate will be determined below by combining the policy-set i_{ss} with the asset-market-determined r_{ss} via the Fisher equation.

Using the debt valuation equation to determine the price ***level***: We now invoke the above derivations within the debt valuation equation to pin down the price level uniquely, provided that the real interest rate is recovered through the asset market equilibrium.

The steady-state real interest rate can be recovered from the asset market through household demand, provided that this demand function is invertible, as

$$r_{ss} = S^{-1}\left(\frac{B}{\tilde{P}} + \frac{b}{\tilde{P}(1+\pi_{ss})}\right),$$

which can be inserted in the government budget equilibrium ($\frac{B}{\tilde{P}} + \frac{b}{\tilde{P}(1+\pi_{ss})} = \sum_{j=0}^{\infty}\left(\frac{1}{1+r_{ss}}\right)^j \bar{s}$) with $r_{ss} > 0$ (such that the right-hand side can be rewritten as a geometric sum, $\sum_{j=0}^{\infty}\left(\frac{1}{1+r_{ss}}\right)^j = \frac{1+r_{ss}}{r_{ss}}$) to get the following condition:

$$\frac{B_{ss} + b_{ss}(1+\pi_{ss})}{\tilde{P}} = \bar{s}\frac{1+r_{ss}}{r_{ss}}.$$

The fixed point of this equation pins down the price level uniquely, given asset market optimality. Using the previously derived definition of the surplus process, i.e., $\bar{s} = t_{ss} = \frac{r_{ss}}{1+i_{ss}}B_{ss} + \frac{r_{ss}}{1+r_{ss}}b_{ss}$:

$$\frac{B_{ss} + b_{ss}(1+\pi_{ss})}{\tilde{P}} = \left[\frac{r_{ss}}{1+i_{ss}}B_{ss} + \frac{r_{ss}}{1+r_{ss}}b_{ss}\right]\frac{1+r_{ss}}{r_{ss}}.$$

Using the Fisher equation ($(1+i_{ss}) = (1+r_{ss})(1+\pi_{ss})$), this equilibrium relation pins down the price level consistently as having to move in line with the stationary equilibrium rate of inflation:

$$\tilde{P} = 1 + \pi_{ss}.$$

It remains to identify π_{ss}. Since i_{ss} is set by the monetary policymaker and r_{ss} has been recovered from the asset market through $S^{-1}(\cdot)$, the Fisher equation delivers steady-state inflation directly:

$$1 + \pi_{ss} = \frac{1+i_{ss}}{1+r_{ss}}.$$

The fiscal authority's bond-issuance schedule then sustains this stationary equilibrium provided it complies with the BGP consistency condition derived above, $g_B = g_b = \pi_{ss}$. This closes the system: the policy rate i_{ss} together with the asset market pins down r_{ss} and hence π_{ss}, the debt valuation equation pins down $\tilde{P}$, and the issuance schedule must adjust to sustain the implied common nominal growth rate.

A.3 Further Simulation Results

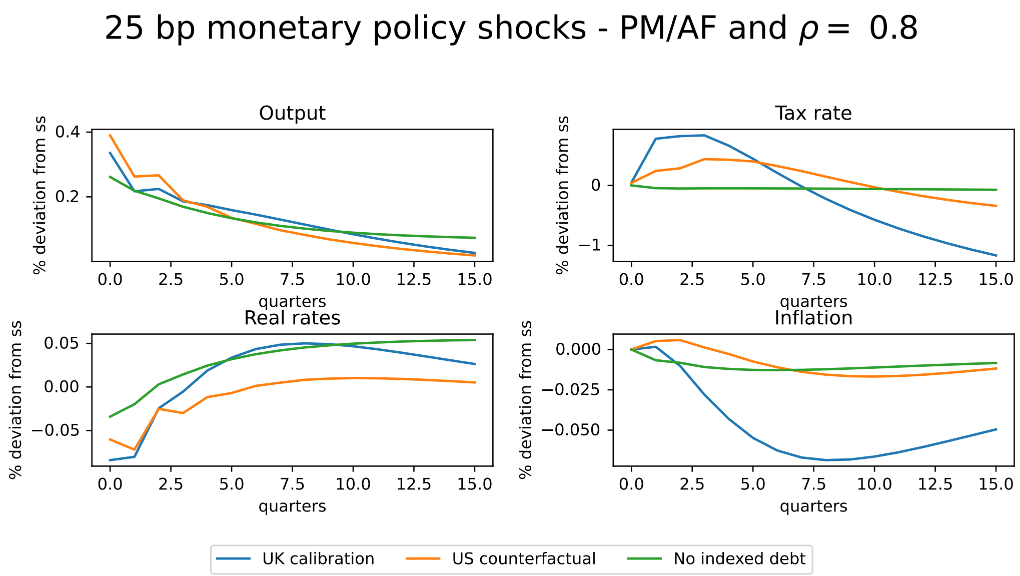

Figure A.9: IRFs to a 25bps expansionary monetary shock - under a fiscally-led policy mix and $\rho = 0.8$.

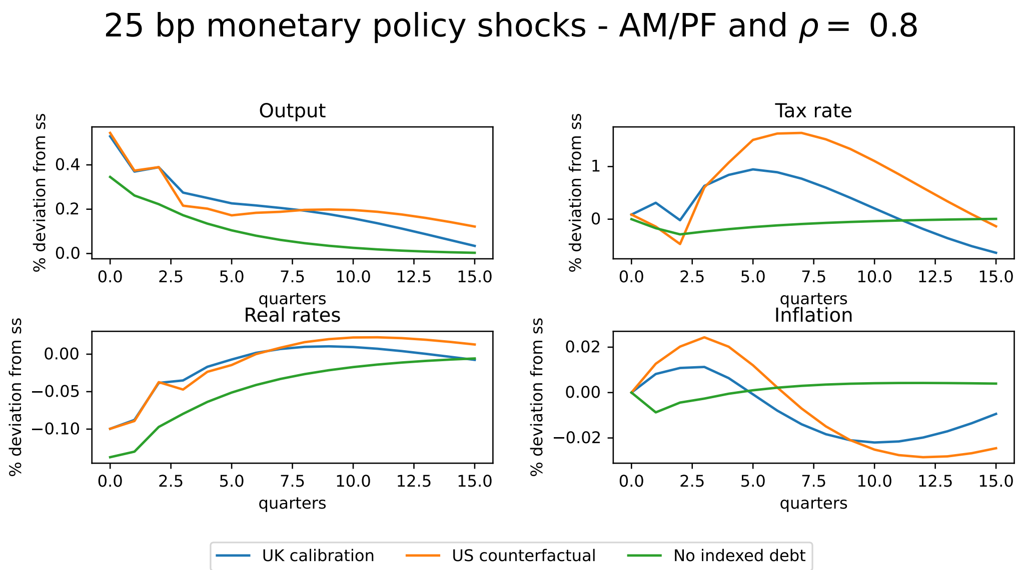

Figure A.10: IRFs to a 25bps expansionary monetary shock - with under a monetary-led policy mix and $\rho = 0.8$.

For the main policy scenario (the 'fiscally-led policy mix'), we furthermore provide additional evidence on changes of quantities directly informing the intertemporal government budget constraint (1.9).

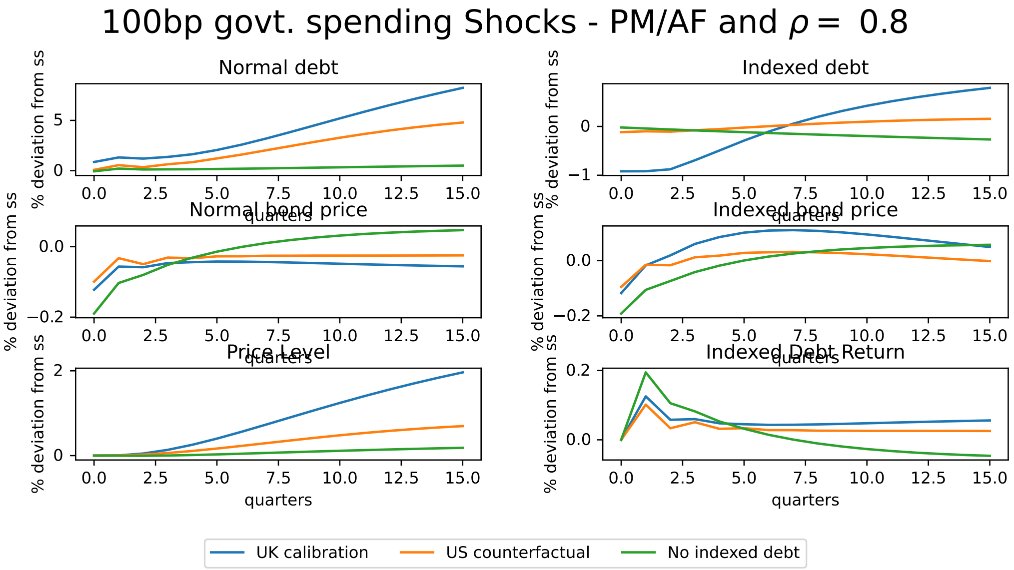

Figure A.11: Further IRFs to a 100bps expansionary fiscal spending shock - under a fiscally-led policy mix and $\rho = 0.8$.

A.4 Long-term Debt and Debt Indexation

In this part of the appendix, we briefly derive the debt valuation equation under complete markets with long-term debt. Due to the assumption of complete markets and following the exposition in the main body of the text, bond pricing kernels for long-term assets maturing at time $(t + j)$ evaluated at time t are given by:

$$Q_t^{(t+j)} = \mathbb{E}_t \left(\beta^j \frac{P_t}{P_{t+j}} \right), \qquad q_t^{(t+j)} = \beta^j,$$

reflecting that inflation-indexed debt always has the same price, as its face value accounts for changes to the price level between issuance and redemption. That being said, indexed debt is *not* fully equivalent to a real claim in the sense that its payout value is not scaled by the prevailing price level.

In this context, the government flow budget condition is given by:

$$B_{t-1}^{(t)} + \Pi_t b_{t-1}^{(t)} = P_t s_t + \sum_{j=1}^{\infty} Q_t^{(t+j)} \left(B_t^{(t+j)} - B_{t-1}^{(t+j)} \right) + \sum_{j=1}^{\infty} q_t^{(t+j)} \left(b_t^{(t+j)} - \Pi_t b_{t-1}^{(t+j)} \right).$$

This condition states that in each period t, the payout of maturing debt (left-hand side) must be equal to the nominal surpluses raised *plus* the possible income from issuing additional debt maturing as a later point in the future (relative to what had already been issued before). Governments can also redeem more bonds than they issue, in which case either of the sums on the right-hand side can also be negative.

That flow condition keeps track of mounting payments on inflation-indexed debt by adjusting the prospective cost of serving indexed debt in each period by the accumulated face value payments, given by $\left(b_t^{(t+j)} - \Pi_t b_{t-1}^{(t+j)}\right)$. In sum, surpluses on the right-hand side of the previous equation get diminished when inflation Π_t from the last period has been high, as that inflation is reflected in the obligations that the government will have as that long-term inflation-indexed debt matures.

Grouping terms in the previous equation yields:

$$\sum_{j=0}^{\infty} Q_t^{(t+j)} B_{t-1}^{(t+j)} + \sum_{j=0}^{\infty} \Pi_t b_{t-1}^{(t+j)} = P_t s_t + \sum_{j=1}^{\infty} Q_t^{(t+j)} B_t^{(t+j)} + \sum_{j=1}^{\infty} q_t^{(t+j)} b_t^{(t+j)}. \tag{A.7}$$

Let the real value of debt now be defined as:

$$V_t := \sum_{j=0}^{\infty} Q_t^{(t+j)} \frac{B_{t-1}^{(t+j)}}{P_t} + \sum_{j=0}^{\infty} q_t^{(t+j)} \frac{b_{t-1}^{(t+j)}}{P_{t-1}}.$$

Focus on the right-hand side of equation (A.7). We now rewrite the two summative terms again to obtain V_{t+1}. Dividing those terms by P_t:

$$\sum_{j=1}^{\infty} Q_t^{(t+j)} \frac{B_t^{(t+j)}}{P_t} + \sum_{j=1}^{\infty} q_t^{(t+j)} \frac{b_t^{(t+j)}}{P_t}.$$

Shifting the index from $j = 1$ to $j = 0$ gives:

$$\sum_{j=0}^{\infty} Q_t^{(t+j+1)} \frac{B_t^{(t+j+1)}}{P_t} + \sum_{j=0}^{\infty} q_t^{(t+j+1)} \frac{b_t^{(t+j+1)}}{P_t}.$$

$q_t^{(t+j+1)} = \beta q_{t+1}^{(t+j+1)}$ by the bond pricing kernels defined previously. Thus, the previous expression becomes

$$\beta \left[\sum_{j=0}^{\infty} Q_{t+1}^{(t+j+1)} \frac{B_t^{(t+j+1)}}{P_{t+1}} + \sum_{j=0}^{\infty} q_{t+1}^{(t+j+1)} \frac{b_t^{(t+j+1)}}{P_t} \right] = \beta V_{t+1}$$

by the definition of V_t. Now, applying a transversality condition of the form

$$\lim_{T \to \infty} \beta^T \left[\sum_{j=0}^{\infty} Q_{t+T}^{(t+j+T)} \frac{B_{t+T}^{(t+j+T)}}{P_{t+T}} + \sum_{j=0}^{\infty} q_{t+T}^{(t+j+T)} \frac{b_{t+T}^{(t+j+T)}}{P_{t+T}} \right] = 0,$$

obtains the debt valuation equation with inflation-indexed debt:

$$\sum_{j=0}^{\infty} Q_t^{(t+j)} \frac{B_{t-1}^{(t+j)}}{P_t} + \sum_{j=0}^{\infty} q_t^{(t+j)} \frac{b_{t-1}^{(t+j)}}{P_{t-1}} = \mathbb{E}_t \sum_{j=0}^{\infty} \beta^j s_{t+j}, \tag{A.8}$$

which is a straightforward generalization of the government debt valuation equation exposed, for instance, in Cochrane (2001).

Chapter B

Appendix to: Inflation-Indexed Debt and the Risks of Fiscal Dominance

B.1 The Fisherian Model under a Standard Taylor Rule

In place of the bond return rate-targeting monetary rule in section 2.2, we here explore the dynamic properties of the Fisherian framework under a standard policy rule defined on the simple nominal interest rate.

The central bank would then follow a simplified monetary rule on the gross nominal interest rate:

$$R_t^n = \frac{1}{\beta}\Pi_t^{\phi}, \tag{B.1}$$

where $R_{n,t} \equiv 1 + i_t$ is the gross nominal interest rate, and ϕ is the monetary policy parameter of interest that will indicate whether the economy operates under a monetary-led policy mix or not. This monetary rule implies a zero-inflation steady-state consistent with household optimality (captured by the bond pricing equations), under which $R^n = \frac{1}{\beta}$. Thus, under the present setting, $Q_t = \frac{1}{R_t^n}$, i.e., the price of the nominal bond must be the inverse of the gross nominal interest rate as all bonds are short-term and the relevant interest rate is under control of the monetary authority.

*Equilibrium: The equilibrium is derived exactly as in the main body of the chapter, but making use of the monetary rule (B.1). Expressing the system in log-deviations from a deterministic zero-inflation steady-state, we obtain the following system of equations:

- Nominal bond prices: using $q_t = \frac{1}{R_t}$ and $q_t = \beta\mathbb{E}_t\left[\Pi_{t+1}^{\theta-1}\right]$, we get, to a first-order approximation,

$$-\hat{R}_t = (\theta - 1)\mathbb{E}_t\hat{\pi}_{t+1}. \tag{B.2}$$

- Monetary rule: taking logs of the monetary rule, we get

$$\hat{R}_t^n = \phi \hat{\pi}_t. \tag{B.3}$$

- Law of motion of debt: as in the main text,

$$(\theta - 1)\hat{\pi}_t + \hat{s}_{t-1} + \hat{R}_{t-1} = (1-\beta)\hat{\tau}_t + \beta \hat{s}_t. \tag{B.4}$$

- Fiscal rule: the fiscal rule can be log-linearized exactly:

$$\hat{\tau}_t = \gamma \hat{s}_{t-1} + \varphi_t. \tag{B.5}$$

A no-arbitrage argument: To induce finite demands on either type of asset (and thereby to close the model under the Taylor Rule on the nominal interest rate), which is necessary for the existence of an equilibrium, we posit that households demand equivalent *real* returns from both assets in log-deviations:

$$\hat{R}_t^n - \mathbb{E}_t \hat{\pi}_{t+1} = \hat{R}_t - (1-\theta)\mathbb{E}_t \hat{\pi}_{t+1}. \tag{B.6}$$

The left-hand side captures the real return on the nominal asset. The right-hand side captures the real return on the government-issued partially-indexed bond, where the inflation adjustment is weighted by $(1-\theta)$, the *non-indexed share* of that bond.

We can combine equations (B.2)+(B.3)+(B.6) to attain the following system of difference equations:

$$\mathbb{E}_t \hat{\pi}_{t+1} = \phi \hat{\pi}_t, \tag{B.7}$$

For the fiscal block, we remain with the same specification as in the main text:

$$\hat{s}_t = (1-\beta)\gamma \hat{s}_t + (1-\beta)\mathbb{E}_t \varphi_{t+1} + \mathbb{E}_t \beta \hat{s}_{t+1}. \tag{B.8}$$

Collecting, the system can be written as

$$\underbrace{\begin{pmatrix} 1 & 0 \\ 0 & \beta \end{pmatrix}}_{\equiv A_0} \mathbb{E}_t \begin{pmatrix} \hat{\pi}_{t+1} \\ \hat{s}_{t+1} \end{pmatrix} = \underbrace{\begin{pmatrix} \phi & 0 \\ 0 & 1-\gamma(1-\beta) \end{pmatrix}}_{\equiv A_1} \begin{pmatrix} \hat{\pi}_t \\ \hat{s}_t \end{pmatrix} + \underbrace{\begin{pmatrix} 0 \\ -(1-\beta) \end{pmatrix}}_{\equiv C} \mathbb{E}_t \varphi_{t+1}. \tag{B.9}$$

The matrix relevant for the determinacy properties is

$$Z = A_0^{-1} A_1 = \begin{pmatrix} \phi & 0 \\ 0 & \dfrac{1-\gamma(1-\beta)}{\beta} \end{pmatrix}. \tag{B.10}$$

Its eigenvalues are ϕ and $\dfrac{1-\gamma(1-\beta)}{\beta}$. A *monetary-led* equilibrium requires $\phi > 1$ and $\gamma > 1$. A *fiscally-led* equilibrium requires $\phi < 1$ and $\gamma < 1$. In sum, inflation-indexed debt does *not* matter for the determinacy properties of this simple model when the Taylor Rule ignores the composition of government debt. This is a direct consequence of the no-arbitrage argument that forces the real return rate of the partially-indexed government bond to be equal to the return rate of the zero-supply outside asset.

B.2 An Argument Ensuring Exactly One Root Inside the Unit Circle

Proof. We here present a brief argument by which we can derive the conditions under which exactly one root of a 2×2 system lies inside the unit circle.

Recall that the characteristic equation of Z is given by

$$p(\lambda) = \lambda^2 - \operatorname{tr}(Z)\lambda + \det(Z).$$

Define the two following expressions for convenience:

$$p(1) \equiv 1 - \operatorname{tr}(Z) + \det(Z), \qquad p(-1) \equiv 1 + \operatorname{tr}(Z) + \det(Z).$$

Now, let the two eigenvalues, as defined by the characteristic equation, be given by λ_1 and λ_2. Then, $p(\cdot)$ evaluated at 1 and -1 yields:

$$p(1) = (1-\lambda_1)(1-\lambda_2), \qquad p(-1) = (1+\lambda_1)(1+\lambda_2),$$

while the determinant and trace can be expressed as $\det(Z) = \lambda_1\lambda_2$, $\operatorname{tr}(Z) = \lambda_1 + \lambda_2$.

Now, combining the insight of Woodford (2011) on the location of both eigenvalues outside the unit circle with results that allow us to express cases when both eigenvalues are inside the unit circle (e.g., Antsaklis and Michel (2006), Theorem 10.10), we can summarize the necessary and sufficient conditions for the existence of both eigenvalues either inside or outside the unit circle as follows:

- *Both inside* ($|\lambda_1|, |\lambda_2| < 1$) iff

$$\det(Z) < 1, \quad p(1) > 0, \quad p(-1) > 0.$$

- *Both outside (Case I)* iff

$$\det(Z) > 1, \quad p(1) > 0, \quad p(-1) > 0.$$

- *Both outside (Case II)* iff
$$p(1) < 0, \quad p(-1) < 0.$$

In addition, we now use a technical boundary to rule out that $p(-1) < 0$. That technical boundary is expressed as the condition under which $p(-1) > 0$, which is given by:

$$\begin{aligned} p(-1) &> 0 \\ \Leftrightarrow 2\left(1 + \frac{1}{\beta}\right) + \phi\sigma\left(1 + \frac{1}{\beta}\right) - \frac{\kappa\sigma\theta}{\beta} &> 0 \\ \Leftrightarrow (2 + \phi\sigma)\left(1 + \frac{1}{\beta}\right) &> \frac{\kappa\sigma\theta}{\beta} \\ \Leftrightarrow \phi &> -\frac{1}{\sigma}\left[2 - \frac{\kappa\sigma\theta}{1+\beta}\right]. \end{aligned}$$

Without indexed debt, this condition is always satisfied under the technical boundary ensuring a positive determinant, $\phi > -\frac{1}{\sigma}$. Here, for high levels of inflation-indexed debt (and a sufficiently high debt-to-GDP ratio or a sufficiently steep Phillips Curve), this might be an additional restriction to consider.

In sum total, we can express a total of eight cases in dependence of $\det(Z)$, $p(1)$, and $p(-1)$, for which six cases are covered so far as necessary and sufficient:

- $\det(Z) < 1, p(1) > 0, p(-1) > 0$: Both roots inside
- $\det(Z) < 1, p(1) > 0, p(-1) < 0$: Ruled out
- $\det(Z) < 1, p(1) < 0, p(-1) > 0$:
- $\det(Z) < 1, p(1) < 0, p(-1) < 0$: Both roots outside (case II)
- $\det(Z) > 1, p(1) > 0, p(-1) > 0$: Both roots outside (case I)
- $\det(Z) > 1, p(1) > 0, p(-1) < 0$: Ruled out
- $\det(Z) > 1, p(1) < 0, p(-1) > 0$:
- $\det(Z) > 1, p(1) < 0, p(-1) < 0$: Both roots outside (case II)

Then, all that must be proven is the restriction under which $p(1) = 1 - \mathrm{tr}(Z) + det(Z) = \left(\frac{1}{\beta} - 1\right)\phi\sigma + \frac{D^{SS}}{Y^{SS}}\frac{\kappa\sigma\theta}{\beta}$ is smaller than zero. These are the remaining cases for which exactly one real root lies inside the unit circle. This restriction is given by:

$$\phi < -\frac{\kappa\theta}{1-\beta} \tag{B.11}$$

Therefore, the region ensuring a saddle-path stable equilibrium under a fiscally-led policy mix restricts the monetary policy parameter ϕ into the following space:

$$\max\left\{-\frac{1}{\sigma}, -\frac{1}{\sigma}\left[2-\frac{\kappa\sigma\theta}{1+\beta}\right]\right\} < \phi < -\frac{\kappa\theta}{1-\beta}. \tag{B.12}$$

For $\theta = 0$, this reverts to the standard restriction for a fiscally-led policy mix in RANK models that $-\frac{1}{\sigma} < \phi < 0$.

■

B.3 Criteria for Monetary Dominance in RANK

Define again the matrix in equation (2.31) as A. Its characteristic polynomial is given by:

$$\begin{aligned}\det(A-\lambda I) = & \left[\frac{1}{\beta}(1-\tau_d)-\lambda\right]\left[\left(1+\phi\sigma-\frac{D^{SS}}{Y^{SS}}\frac{\kappa\sigma\theta}{\beta}-\lambda\right)\left(\frac{1}{\beta}-\lambda\right)+\frac{D^{SS}}{Y^{SS}}\frac{\kappa\theta}{\beta^2}\right] \\ = & \left[\frac{1}{\beta}(1-\tau_d)-\lambda\right]\left[\lambda^2-\lambda\left(\frac{1}{\beta}+1+\sigma\phi-\frac{D^{SS}}{Y^{SS}}\frac{\kappa\sigma\theta}{\beta}\right)+\frac{1}{\beta}(1+\sigma\phi)\right]\end{aligned} \tag{B.13}$$

In order to establish how indexed debt influences the system's stability properties, we consider the boundaries under which conventional monetary dominance is established.

Monetary dominance requires fiscal passivity under RANK for saddle path-stability, which here is ensured by $\tau_d > 1-\beta$, i.e., fiscal policy must be sufficiently reactive to deviations of the value of government debt from its steady-state. In that case, the root implied by the first bracket in equation (B.13) is inside the unit circle, so we require two further roots outside the unit circle.

Analysing those can be done equivalently by reducing our focus to the 2×2 sub-system on (y_t, π_t), characterized by the following matrix Z:

$$Z \equiv \begin{bmatrix} 1+\phi\sigma-\frac{D^{SS}}{Y^{SS}}\frac{\kappa\sigma\theta}{\beta} & \frac{D^{SS}}{Y^{SS}}\frac{\sigma\theta}{\beta} \\ -\frac{\kappa}{\beta} & \frac{1}{\beta}\end{bmatrix}.$$

Following Woodford (2011) and Rachel and Ravn (2025), either of the following sets of conditions is sufficient to guarantee that two roots of Z lie outside the unit circle:

1. First case:

   - $\det(Z) > 1$
   - $\det(Z) - \text{tr}(Z) > -1$
   - $\det(Z) + \text{tr}(Z) > -1$

2. Second case:
    - $\det(Z) - \operatorname{tr}(Z) < -1$
    - $\det(Z) + \operatorname{tr}(Z) < -1$.

We focus on case 1, ruling out case 2 along the way. First, note that:

$$\det(Z) = \frac{1}{\beta}(1 + \sigma\phi)$$
$$\operatorname{tr}(Z) = (1 + \sigma\phi) + \frac{1}{\beta}\left[1 - \frac{D^{SS}}{Y^{SS}}\kappa\sigma\theta\right].$$

Let active monetary policy be characterized by $\phi \geq 0$, as in the main body of the text. Then, $\det(Z) > 1$ is trivially true for any $\beta < 1$. Next, we consider the second condition of the first case:

$$\det(Z) - \operatorname{tr}(Z) = \frac{1}{\beta}\left(\sigma\phi + \frac{D^{SS}}{Y^{SS}}\kappa\sigma\theta\right) - (1 + \sigma\phi) > -1$$
$$\leftrightarrow \qquad \left(\frac{1}{\beta} - 1\right)\sigma\phi + \frac{D^{SS}}{Y^{SS}}\kappa\sigma\theta > 0,$$

which, again, holds trivially true for all values of θ, including $\theta = 0$ (which is the base case of Angeletos et al. (2024b) and Rachel and Ravn (2025)). Therefore, the second set of conditions must not be considered, since the previous equation implies that $\det(Z) - \operatorname{tr}(Z) < -1$ cannot hold true for $\phi > 0$.

Finally, consider the third condition of the first case. Here, we find that:

$$\det(Z) + \operatorname{tr}(Z) = \left(\frac{1}{\beta} + 1\right)(1 + \sigma\phi) + \frac{1}{\beta}\left(1 - \frac{D^{SS}}{Y^{SS}}\kappa\sigma\theta\right) > -1$$
$$\leftrightarrow \qquad (1 + \beta)(1 + \sigma\phi) + \left(1 - \frac{D^{SS}}{Y^{SS}}\kappa\sigma\theta\right) > -\beta$$
$$\leftrightarrow \qquad (1 + \beta)(1 + \sigma\phi) + > \frac{D^{SS}}{Y^{SS}}\kappa\sigma\theta - (1 + \beta)$$
$$\leftrightarrow \qquad 1 + \sigma\phi > \frac{D^{SS}}{Y^{SS}}\frac{\kappa\sigma\theta}{1 + \beta} - 1$$
$$\leftrightarrow \qquad \phi > \frac{1}{\sigma}\left[\frac{D^{SS}}{Y^{SS}}\frac{\kappa\sigma\theta}{1 + \beta} - 2\right],$$

i.e., the admissible bounds for ϕ that ensure monetary dominance *can* tighten when the share of indexed debt is sufficiently high. Tighter determinacy bounds under monetary dominance due to the presence of indexed debt are supported through steep Phillips Curves (high κ), high degrees of risk aversion (high σ), or high debt-to-GDP ratios (high $\frac{D^{SS}}{Y^{SS}}$). Note that without indexed debt ($\theta = 0$), this condition reduces to $\phi > -\frac{2}{\sigma}$, which is trivially true.

B.4 Proof of the Boundaries Ensuring Saddle Path-Stable Equilibria

The trace is defined as:

$$\operatorname{tr}(A) = 1 + \phi\left(\sigma + \frac{D^{SS}}{Y^{SS}}(1-\beta\omega)(1-\tfrac{1}{\omega})\right) - \frac{\kappa\theta}{\beta}\left(\sigma + \frac{D^{SS}}{Y^{SS}}(1-\beta\omega)(1-\tfrac{1}{\omega})\right) \\ -\frac{\tau_y}{\beta}(1-\beta\omega)(1-\tfrac{1}{\omega}) + \frac{2-\tau_d}{\beta}.$$

The determinant, in turn, is given by:

$$\det(A) = a_{11}\left(\tfrac{1}{\beta}\right)\left(\tfrac{1-\tau_d}{\beta}\right) - \left(\tfrac{\theta}{\beta}\left(\sigma + \frac{D^{SS}}{Y^{SS}}(1-\beta\omega)(1-\tfrac{1}{\omega})\right)\right)\left(-\tfrac{\kappa}{\beta}\right)\left(\tfrac{1-\tau_d}{\beta}\right) \\ + \left(\tfrac{(1-\beta\omega)(1-\frac{1}{\omega})}{\beta}(1-\tau_d)\right)\left(-\tfrac{\kappa}{\beta}\right)\left(\tfrac{\theta}{\beta}\tfrac{D^{SS}}{Y^{SS}}\right) - \left(\tfrac{(1-\beta\omega)(1-\frac{1}{\omega})}{\beta}(1-\tau_d)\right)a_{31}\left(\tfrac{1}{\beta}\right).$$

Factoring out $\frac{1-\tau_d}{\beta^2}$, we obtain:

$$\det(A) = \frac{1-\tau_d}{\beta^2}\left[a_{11} + \frac{\theta\kappa}{\beta}\left(\sigma + \frac{D^{SS}}{Y^{SS}}(1-\beta\omega)(1-\tfrac{1}{\omega})\right) - \frac{\theta\kappa}{\beta}\frac{D^{SS}}{Y^{SS}}(1-\beta\omega)(1-\tfrac{1}{\omega}) - (1-\beta\omega)(1-\tfrac{1}{\omega})a_{31}\right].$$

Next, use $\left(\sigma + \frac{D^{SS}}{Y^{SS}}(1-\beta\omega)(1-\frac{1}{\omega})\right) - \frac{D^{SS}}{Y^{SS}}(1-\beta\omega)(1-\frac{1}{\omega}) = \sigma$:

$$\det(A) = \frac{1-\tau_d}{\beta^2}\left[a_{11} + \frac{\theta\kappa}{\beta}\sigma - (1-\beta\omega)(1-\tfrac{1}{\omega})a_{31}\right].$$

Substitute a_{11} and a_{31}:

$$\det(A) = \frac{1-\tau_d}{\beta^2} \\ \Bigg[\underbrace{\left(1 + \phi\left(\sigma + \frac{D^{SS}}{Y^{SS}}(1-\beta\omega)(1-\tfrac{1}{\omega})\right) - \tfrac{\kappa\theta}{\beta}\left(\sigma + \frac{D^{SS}}{Y^{SS}}(1-\beta\omega)(1-\tfrac{1}{\omega})\right) - \tfrac{\tau_y}{\beta}(1-\beta\omega)(1-\tfrac{1}{\omega})\right)}_{a_{11}} \\ + \frac{\theta\kappa}{\beta}\sigma - (1-\beta\omega)(1-\tfrac{1}{\omega})\underbrace{\left(\tfrac{D^{SS}}{Y^{SS}}\phi - \tfrac{\tau_y}{\beta} - \tfrac{\theta\kappa}{\beta}\tfrac{D^{SS}}{Y^{SS}}\right)}_{a_{31}}\Bigg]$$

$$= \frac{1-\tau_d}{\beta^2}\Bigg[1 + \phi\left(\sigma + \frac{D^{SS}}{Y^{SS}}(1-\beta\omega)(1-\tfrac{1}{\omega})\right) - \frac{\kappa\theta}{\beta}\left(\sigma + \frac{D^{SS}}{Y^{SS}}(1-\beta\omega)(1-\tfrac{1}{\omega})\right) \\ - \frac{\tau_y}{\beta}(1-\beta\omega)(1-\tfrac{1}{\omega}) + \frac{\theta\kappa}{\beta}\sigma - (1-\beta\omega)(1-\tfrac{1}{\omega})\left(\tfrac{D^{SS}}{Y^{SS}}\phi\right) + (1-\beta\omega)(1-\tfrac{1}{\omega})\left(\tfrac{\tau_y}{\beta}\right) \\ + (1-\beta\omega)(1-\tfrac{1}{\omega})\left(\tfrac{\theta\kappa}{\beta}\tfrac{D^{SS}}{Y^{SS}}\right)\Bigg].$$

Many elements cancel out:

$$-\frac{\tau_y}{\beta}(1-\beta\omega)(1-\tfrac{1}{\omega})+(1-\beta\omega)(1-\tfrac{1}{\omega})\frac{\tau_y}{\beta}=0,$$

$$\begin{aligned}&-\frac{\kappa\theta}{\beta}\left(\sigma+\frac{D^{SS}}{Y^{SS}}(1-\beta\omega)(1-\tfrac{1}{\omega})\right)+\frac{\theta\kappa}{\beta}\sigma+\frac{\theta\kappa}{\beta}\frac{D^{SS}}{Y^{SS}}(1-\beta\omega)(1-\tfrac{1}{\omega})\\&=-\frac{\kappa\theta}{\beta}\big(\sigma+\tfrac{D^{SS}}{Y^{SS}}(1-\beta\omega)(1-\tfrac{1}{\omega})\big)+\frac{\theta\kappa}{\beta}\sigma+\frac{\theta\kappa}{\beta}\frac{D^{SS}}{Y^{SS}}(1-\beta\omega)(1-\tfrac{1}{\omega})=0.\end{aligned}$$

For the ϕ-terms,

$$\begin{aligned}&\phi\left(\sigma+\frac{D^{SS}}{Y^{SS}}(1-\beta\omega)(1-\tfrac{1}{\omega})\right)-(1-\beta\omega)(1-\tfrac{1}{\omega})\Big(\tfrac{D^{SS}}{Y^{SS}}\phi\Big)\\&=\phi\Big(\left(\sigma+\frac{D^{SS}}{Y^{SS}}(1-\beta\omega)(1-\tfrac{1}{\omega})\right)-\tfrac{D^{SS}}{Y^{SS}}(1-\beta\omega)(1-\tfrac{1}{\omega})\Big)=\phi\sigma.\end{aligned}$$

Hence,

$$\det(A)=\frac{1-\tau_d}{\beta^2}\big(1+\sigma\phi\big).$$

Let the characteristic polynomial be defined in the form:

$$\lambda^3+\Gamma_2\lambda^2+\Gamma_1\lambda+\Gamma_0=0.$$

Then, its elements are given by:

$$\Gamma_2=-\mathrm{tr}(A)=$$

$$-\left(\left(1+\phi\left(\sigma+\frac{D^{SS}}{Y^{SS}}(1-\beta\omega)(1-\tfrac{1}{\omega})\right)-\frac{\kappa\theta}{\beta}\left(\sigma+\frac{D^{SS}}{Y^{SS}}(1-\beta\omega)(1-\tfrac{1}{\omega})\right)-\frac{\tau_y}{\beta}(1-\beta\omega)(1-\tfrac{1}{\omega})\right)+\frac{2-\tau_d}{\beta}\right)$$

$$\begin{aligned}\Gamma_1=\frac{2-\tau_d}{\beta}&\left[1+\phi\left(\sigma+\frac{D^{SS}}{Y^{SS}}(1-\beta\omega)(1-\tfrac{1}{\omega})\right)-\frac{\kappa\theta}{\beta}\left(\sigma+\frac{D^{SS}}{Y^{SS}}(1-\beta\omega)(1-\tfrac{1}{\omega})\right)-\frac{\tau_y}{\beta}(1-\beta\omega)(1-\tfrac{1}{\omega})\right]\\&+\frac{1-\tau_d}{\beta^2}+\frac{\theta\kappa}{\beta^2}\left(\sigma+\frac{D^{SS}}{Y^{SS}}(1-\beta\omega)(1-\tfrac{1}{\omega})\right)-\frac{(1-\tau_d)(1-\beta\omega)(1-\tfrac{1}{\omega})}{\beta}\left[\frac{D^{SS}}{Y^{SS}}\phi-\frac{\tau_y}{\beta}-\frac{\theta\kappa}{\beta}\frac{D^{SS}}{Y^{SS}}\right].\end{aligned}$$

$$\Gamma_0=-\det(A)=-\frac{1-\tau_d}{\beta^2}\big(1+\sigma\phi\big)$$

Then, in line with Rachel and Ravn (2025), the following sets of conditions are sufficient to claim that two roots of the characteristic polynomial lie outside the unit circle:

1. First set:

- $1 + \Gamma_2 + \Gamma_1 + \Gamma_0 < 0$
- $-1 + \Gamma_2 - \Gamma_1 + \Gamma_0 > 0$

2. Second set:

- $1 + \Gamma_2 + \Gamma_1 + \Gamma_0 > 0$
- $-1 + \Gamma_2 - \Gamma_1 + \Gamma_0 < 0$
- $\Gamma_0(\Gamma_0 - \Gamma_2) + \Gamma_1 - 1 > 0$

3. Third set:

- $1 + \Gamma_2 + \Gamma_1 + \Gamma_0 > 0$
- $-1 + \Gamma_2 - \Gamma_1 + \Gamma_0 < 0$
- $\Gamma_0(\Gamma_0 - \Gamma_2) + \Gamma_1 - 1 < 0$
- $|\Gamma_2| > 3$.

We now start to check under which boundaries either of the sets of inequalities holds true. The first set of inequalities, in particular, must normally not be considered given our parametric boundaries on ϕ and τ_d, since the second inequality of the first set is always violated.

For simpler notation, let us define:

$$\mathcal{T} \equiv (1 - \beta\omega)\Big(1 - \tfrac{1}{\omega}\Big), \qquad \mathcal{S} \equiv \sigma + \frac{D^{SS}}{Y^{SS}}\mathcal{T}.$$

$\mathcal{T}$ is smaller then zero when there exists any mortality risk, and zero otherwise. The exact sign of $\mathcal{S}$ is unclear and depends on the relative strength of the risk aversion in relation to the debt-weighted wedge induced by the mortality risk. Usually, we can expect $\mathcal{S} > 0$ for common calibrations. We therefore impose the following assumption.

Assumption 1 *The relative strength of the risk aversion channel is stronger than the effect of imperfect household insurance arising through mortality risk; formally,* $\mathcal{S} > 0$.

Recall that the entries from the matrix are given by:

$$a_{11} = 1 + \phi\mathcal{S} - \frac{\kappa\theta}{\beta}\mathcal{S} - \frac{\tau_y}{\beta}\mathcal{T}, \qquad a_{12} = \frac{\theta}{\beta}\mathcal{S}, \qquad a_{13} = \frac{\mathcal{T}}{\beta}(1 - \tau_d),$$
$$a_{21} = -\frac{\kappa}{\beta}, \qquad a_{22} = \frac{1}{\beta}, \qquad a_{23} = 0,$$
$$a_{31} = \frac{D^{SS}}{Y^{SS}}\phi - \frac{\tau_y}{\beta} - \frac{\theta\kappa}{\beta}\frac{D^{SS}}{Y^{SS}}, \qquad a_{32} = \frac{\theta}{\beta}\frac{D^{SS}}{Y^{SS}}, \qquad a_{33} = \frac{1 - \tau_d}{\beta}.$$

Proof.

Now, define

$$C(\tau_d) := \sigma\Big(\tau_d - (1-\beta)\Big) + \beta\mathcal{T}\frac{D^{SS}}{Y^{SS}}, \qquad \tau_d^0 := (1-\beta) - \frac{\beta\mathcal{T}}{\sigma}\frac{D^{SS}}{Y^{SS}},$$

$$\phi^\star(\tau_d) := \frac{\tau_y\mathcal{T}}{C(\tau_d)} - \frac{\theta\kappa}{1-\beta},$$

and

$$\Delta(\phi,\tau_d) := 1 - S_2(A) + \mathrm{tr}(A)\det(A) - \det(A)^2.$$

Here, $C(\tau_d)$ can be interpreted as the "tax-finance wedge" of additional debt issuance, with τ_d^0 being its zero-bound and $\phi^*(\tau_d)$ being a related cut-off parameter of the monetary rule.

For

$$\tau_d^\sharp := 1 - \left[\beta^2 + \frac{\beta^2}{\sigma}\frac{D^{SS}}{Y^{SS}}\mathcal{T}\right],$$

the equation $\Delta(\phi,\tau_d) = 0$ has two parts $\phi_\pm(\tau_d;\theta)$ whenever $\tau_d > \tau_d^\sharp$, and $\Delta < 0$ if $\phi \in \Big(\phi_-(\tau_d;\theta), \phi_+(\tau_d;\theta)\Big)$. The trace inside that policy space that is useful for our proof for ϕ is restricted by $|\mathrm{tr}(A)| \le 3$. Thus,

$$|\mathrm{tr}(A)| > 3 \iff \phi \notin [\phi_{\mathrm{tr}}^-, \phi_{\mathrm{tr}}^+], \qquad \phi_{\mathrm{tr}}^\pm(\tau_d;\theta) := \frac{\pm 3 - C_0(\tau_d;\theta)}{\mathcal{S}}, \quad C_0(\tau_d;\theta) := 1 + \frac{2-\tau_d}{\beta} - \frac{\kappa\theta}{\beta}\mathcal{S} - \frac{\tau_y}{\beta}\mathcal{T}.$$

Comparative statics in θ. The key boundaries move with θ as follows:

$$\frac{\partial\phi^\star}{\partial\theta} = -\frac{\kappa}{1-\beta} < 0$$

$$\frac{\partial\Delta}{\partial\theta} = \frac{\kappa(1-\tau_d)}{\beta^3}\Big[\beta\sigma - \mathcal{S}(1+\sigma\phi)\Big], \quad \Rightarrow \begin{cases} \text{for small } \phi: \ \Delta = 0 \text{ shifts left,} \\ \text{for large } \phi: \ \Delta = 0 \text{ shifts right,} \end{cases}$$

$$\frac{\partial\phi_{\mathrm{tr}}^\pm}{\partial\theta} = \frac{\kappa}{\beta} > 0 \quad \text{(the trace thresholds shift right one-for-one with } \theta\text{).}$$

Some limits

- **As $\sigma \to 0^+$.** $\det(A) = \frac{1-\tau_d}{\beta^2}\Big(1+\sigma\phi\Big) \to \frac{1-\tau_d}{\beta^2}$, so it is independent of ϕ. Also, $C(\tau_d) \to \beta\mathcal{T}\frac{D^{SS}}{Y^{SS}}$, so

  $$\phi^\star(\tau_d) \to \frac{\tau_y}{\beta\frac{D^{SS}}{Y^{SS}}} - \frac{\theta\kappa}{1-\beta}$$

  and $\Delta(\phi,\tau_d)$ becomes nearly affine in ϕ.

- **As $\kappa \to 0^+$.** The inflation block tends to $\pi_{t+1} = (1/\beta)\pi_t$, yielding one unstable root $\lambda_\pi \to 1/\beta > 1$. Creating the second unstable root becomes difficult unless τ_d is small. Moreover, all θ-slopes are proportional to κ (e.g., $\partial_\theta \phi^\star = -\kappa/(1-\beta)$, $\partial_\theta \phi^\pm_{\text{tr}} = \kappa/\beta$), hence the determinacy boundaries become close to independent of θ as $\kappa \to 0^+$. On the flip side, $\kappa \uparrow$ thus seems to increase the relevance of θ.

We now provide the derivations related to the previous statement.

PART 1: The inequality of interest defined on the elements of the characteristic polynomial is given by:

$$-1 + \Gamma_2 - \Gamma_1 + \Gamma_0 = -\Big(1 + \text{tr}(A) + S_2(A) + \det(A)\Big).$$

Thus, it is sufficient to show that:

$$1 + \text{tr}(A) + S_2(A) + \det(A) > 0,$$

where $S_2(A)$ is defined as the sum of all determinants of the (2×2) minors of the full model matrix:

$$S_2(A) \equiv \underbrace{(a_{11}a_{22} - a_{12}a_{21})}_{M_1} + \underbrace{(a_{11}a_{33} - a_{13}a_{31})}_{M_2} + \underbrace{(a_{22}a_{33} - a_{23}a_{32})}_{M_3}.$$

From the earlier derivations, we know that the determinant is given by:

$$\det(A) = \frac{1-\tau_d}{\beta^2}\Big(1 + \sigma\phi\Big).$$

Therefore, under $\tau_d \in [0,1]$ and $\phi > -1/\sigma$, we have

$$\det(A) \geq 0, \qquad \text{so} \qquad -\det(A) = \Gamma_0 \leq 0.$$

Consider now the 2×2 minor defined on the output-inflation sub-system:

$$M_1 \equiv \det\begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{bmatrix} = a_{11}a_{22} - a_{12}a_{21}.$$

Using the entries above,

$$\begin{aligned} M_1 &= \frac{1}{\beta}\Big(1 + \phi\mathcal{S} - \frac{\kappa\theta}{\beta}\mathcal{S} - \frac{\tau_y}{\beta}\mathcal{T}\Big) - \Big(\frac{\theta}{\beta}\mathcal{S}\Big)\Big(-\frac{\kappa}{\beta}\Big) \\ &= \frac{1}{\beta}\Big(1 + \phi\mathcal{S} - \frac{\tau_y}{\beta}\mathcal{T}\Big), \end{aligned}$$

that is, all terms related to θ cancel out.

Now impose that $\omega \in (0,1]$, which implies $\mathcal{T} = (1-\beta\omega)(1-1/\omega) \le 0$. With $\tau_y \ge 0$, we obtain

$$-\frac{\tau_y}{\beta}\mathcal{T} \ge 0, \qquad \Rightarrow \qquad M_1 \ge \frac{1+\phi\mathcal{S}}{\beta}.$$

Moreover, since $\mathcal{S} = \sigma + \frac{D^{SS}}{Y^{SS}}\mathcal{T} \le \sigma$ when $\mathcal{T} \le 0$, we have

$$\phi > -1/\sigma \quad \Longrightarrow \quad 1+\phi\mathcal{S} \ge 1+\phi\sigma > 0.$$

Therefore, $M_1 > 0$.

We now show directly that $\det(I+A) = 1 + \mathrm{tr}(A) + S_2(A) + \det(A) > 0$ on the parameter set $\phi \in (-1/\sigma, \infty)$, $\tau_d \in [0,1]$, under one additional restriction on the degree of debt indexation.

Assumption 2 *We bound the degree of debt indexation. The degree of debt indexation to inflation satisfies*

$$\theta < \bar{\theta} \equiv \frac{1+\beta}{\kappa\sigma}. \tag{B.14}$$

This bound is mild for standard calibrations. With $\beta = 0.99$, $\kappa = 0.03$, $\sigma = 1$ we obtain $\bar{\theta} \approx 66.3$, which would be far more than 100%. The bound tightens only when the Phillips-curve slope κ and risk aversion σ are both very large relative to the discount factor.

We now directly compute $\det(I+A)$. Define the auxiliary quantity

$$\Sigma(\tau_d) := (1+\beta-\tau_d)\sigma + \beta\frac{D^{SS}}{Y^{SS}}\mathcal{T}.$$

Since $\mathcal{S} > 0$ and $\tau_d \le 1$, we have $\Sigma(\tau_d) \ge \beta\Big(\sigma + \frac{D^{SS}}{Y^{SS}}\mathcal{T}\Big) = \beta\mathcal{S} > 0$.

Recall that $a_{23} = 0$. Expanding $\det(I+A)$ along the third column of the model matrix, the determinant splits into two groups.

The first group is the $(1,3)$ cofactor.

$$a_{13}\Big(a_{21}a_{32} - (1{+}a_{22})a_{31}\Big) = \frac{\mathcal{T}(1{-}\tau_d)}{\beta^2}\Big[\kappa\theta\tfrac{D^{SS}}{Y^{SS}} - (\beta{+}1)\tfrac{D^{SS}}{Y^{SS}}\phi + \tfrac{(\beta+1)\tau_y}{\beta}\Big],$$

obtained by substituting $a_{21} = -\kappa/\beta$, $a_{32} = \theta\frac{D^{SS}}{Y^{SS}}/\beta$, $1{+}a_{22} = (1{+}\beta)/\beta$, and the expression for a_{31}.

The second group is the $(3,3)$ cofactor. Using $1 + a_{33} = (1+\beta-\tau_d)/\beta$ and $a_{12}a_{21} = -\kappa\theta\mathcal{S}/\beta^2$,

$$(1{+}a_{33})\Big[(1{+}a_{11})(1{+}a_{22}) - a_{12}a_{21}\Big] = \frac{\beta{+}1{-}\tau_d}{\beta}\Big[\frac{\beta{+}1}{\beta}\Big(2+\phi\mathcal{S} - \tfrac{\tau_y}{\beta}\mathcal{T}\Big) - \frac{\kappa\theta}{\beta}\mathcal{S}\Big].$$

Adding both groups together, the ϕ-terms combine into $\Sigma(\tau_d)\phi$ (via $\mathcal{S} - \frac{D^{SS}}{Y^{SS}}\mathcal{T} = \sigma$), and the $\kappa\theta$-terms combine into $-\kappa\theta\Sigma(\tau_d)$. The result is:

$$\beta^2 \det(I + A) = (1+\beta)\Big[2(1+\beta-\tau_d) - \tau_y\mathcal{T}\Big] + \Sigma(\tau_d)\Big[(1+\beta)\phi - \kappa\theta\Big]. \tag{B.15}$$

We now show that $\det(I + A)$ is positive.

Expression (B.15) is affine in ϕ with slope $(1+\beta)\Sigma(\tau_d) > 0$. Its minimum on $\phi \in (-1/\sigma, \infty)$ is therefore attained as $\phi \downarrow -1/\sigma$. Evaluating at that boundary:

$$\beta^2 \det(I + A)\Big|_{\phi=-1/\sigma} = (1+\beta)R(\tau_d) - \kappa\theta\Sigma(\tau_d),$$

where

$$R(\tau_d) := 2(1+\beta-\tau_d) - \tau_y\mathcal{T} - \frac{\Sigma(\tau_d)}{\sigma} = (1+\beta-\tau_d) + |\mathcal{T}|\Big(\tau_y + \frac{\beta}{\sigma}\frac{D^{SS}}{Y^{SS}}\Big).$$

Since $1+\beta-\tau_d \geq \beta > 0$, $|\mathcal{T}| \geq 0$, $\tau_y \geq 0$, $D^{SS}/Y^{SS} \geq 0$, and $\sigma > 0$, we conclude that

$$R(\tau_d) \geq \beta > 0 \qquad \text{for all } \tau_d \in [0,1].$$

It therefore suffices to verify that $(1+\beta)R(\tau_d) > \kappa\theta\Sigma(\tau_d)$ for every $\tau_d \in [0,1]$. The ratio $R(\tau_d)/\Sigma(\tau_d)$ is non-decreasing in τ_d. By the quotient rule, the sign of its derivative is determined by

$$R'\Sigma - R\Sigma' \propto |\mathcal{T}|\Big(2\beta\tfrac{D^{SS}}{Y^{SS}} + \sigma\tau_y\Big) \geq 0,$$

so the binding constraint is at $\tau_d = 0$. At that point, $R(0) \geq 1+\beta$ and $\Sigma(0) \leq (1+\beta)\sigma$, yielding the lower bound

$$\frac{(1+\beta)R(0)}{\Sigma(0)} \geq \frac{(1+\beta)^2}{(1+\beta)\sigma} = \frac{1+\beta}{\sigma}.$$

Hence, the assumption $\kappa\theta < (1+\beta)/\sigma$ guarantees

$$(1+\beta)R(\tau_d) > \kappa\theta\Sigma(\tau_d) \qquad \forall\tau_d \in [0,1],$$

and therefore

$$\det(I + A) > 0 \qquad \text{for all } \phi \in (-1/\sigma, \infty), \tau_d \in [0,1].$$

Since $\det(I + A) = 1 + \text{tr}(A) + S_2(A) + \det(A) > 0$, and

$$-1 + \Gamma_2 - \Gamma_1 + \Gamma_0 = -\Big(1 + \text{tr}(A) + S_2(A) + \det(A)\Big) = -\det(I + A) < 0,$$

we conclude that

$$-1 + \Gamma_2 - \Gamma_1 + \Gamma_0 < 0 \qquad \text{for all} \quad \phi \in (-1/\sigma, \infty), \tau_d \in [0,1],$$

provided $\theta < \bar{\theta} \equiv (1+\beta)/(\kappa\sigma)$.

The inequality $-1 + \Gamma_2 - \Gamma_1 + \Gamma_0 > 0$ is therefore never satisfied under these conditions, and we can focus on the second and third sets of inequalities.

—

PART 2: Now, we characterize when $1 + \Gamma_2 + \Gamma_1 + \Gamma_0 > 0$, which is a part of the second and third sets of conditions.

We work on the restricted parameter set $\phi \in (-1/\sigma, \infty)$, $\tau_d \in [0, 1]$, as defined before.

Recall that we defined the variables

$$\mathcal{T} \equiv (1 - \beta\omega)\Big(1 - \tfrac{1}{\omega}\Big), \qquad \mathcal{S} \equiv \sigma + \frac{D^{SS}}{Y^{SS}}\mathcal{T}.$$

Under $\omega \in (0, 1]$ we have $\mathcal{T} \leq 0$, and we additionally assumed that $\mathcal{S} > 0$.

Recall from before that we have

$$\det(I - A) = \frac{1}{\beta^2}\Big(C(\tau_d)[\theta\kappa + (1 - \beta)\phi] - \tau_y\mathcal{T}(1 - \beta)\Big),$$

where

$$C(\tau_d) \equiv \sigma(\tau_d - (1 - \beta)) + \beta\mathcal{T}\frac{D^{SS}}{Y^{SS}} = \sigma\Big(\tau_d - (1 - \beta)\Big) + \beta\mathcal{T}\frac{D^{SS}}{Y^{SS}}.$$

Therefore,

$$1 + \Gamma_2 + \Gamma_1 + \Gamma_0 > 0 \quad \Longleftrightarrow \quad C(\tau_d)[\theta\kappa + (1 - \beta)\phi] > \tau_y\mathcal{T}(1 - \beta).$$

A direct solution for the parametric boundaries cannot be found, but in line with Rachel and Ravn (2025), we can restrict the severity of the conditions in the (τ_d, ϕ)-space.

We now postulate the following 'vertical boundary' on τ_d:

$$\tau_d^0 \equiv (1 - \beta) - \frac{\beta\mathcal{T}}{\sigma}\frac{D^{SS}}{Y^{SS}} = 1 - \beta - \frac{\beta\mathcal{T}}{\sigma}\frac{D^{SS}}{Y^{SS}}.$$

At that line, $C(\tau_d)$ is zero. Since $\mathcal{T} \leq 0$ and $\frac{D^{SS}}{Y^{SS}} \geq 0$, we have $\tau_d^0 \geq (1 - \beta)$. Moreover

$$C(1) = \sigma(1 - (1 - \beta)) + \beta\mathcal{T}\frac{D^{SS}}{Y^{SS}} = \beta\Big(\sigma + \mathcal{T}\frac{D^{SS}}{Y^{SS}}\Big) = \beta\mathcal{S},$$

so if $\mathcal{S} > 0$, then $C(\tau_d)$ is strictly increasing in τ_d, negative for $\tau_d < \tau_d^0$ and positive for $\tau_d > \tau_d^0$, with a single crossing at $\tau_d^0 \in [0, 1]$.

For $C(\tau_d) \neq 0$ we can solve for the zero-boundary for ϕ as a function of τ_d:

$$\phi^\star(\tau_d) = \frac{\tau_y\mathcal{T}}{C(\tau_d)} - \frac{\theta\kappa}{(1 - \beta)}.$$

Note: $\frac{d\phi^\star}{d\tau_d} = -\tau_y\mathcal{T}\sigma\Big/C(\tau_d)^2 \geq 0$ because $\tau_y \geq 0$, $\sigma > 0$, $\mathcal{T} \leq 0$, hence $\phi^\star(\tau_d)$ is *monotone increasing* in τ_d. It has a vertical asymptote at $\tau_d = \tau_d^0$ (since $C(\tau_d^0) = 0$); with $\mathcal{T} < 0$ we have

$$\lim_{\tau_d \nearrow \tau_d^0} \phi^\star(\tau_d) = +\infty, \qquad \lim_{\tau_d \searrow \tau_d^0} \phi^\star(\tau_d) = -\infty.$$

Inequality region: Because the characteristic polynomial is $\frac{1}{\beta^2}$ times an affine function of ϕ with τ_d-dependent slope $C(\tau_d)(1-\beta)$, the sign pattern is:

$$p(1) > 0 \iff \begin{cases} \phi > \phi^\star(\tau_d), & \text{if } C(\tau_d) > 0\ ,\ (\tau_d > \tau_d^0), \\ \phi < \phi^\star(\tau_d), & \text{if } C(\tau_d) < 0\ ,\ (\tau_d < \tau_d^0). \end{cases}$$

On the set $C(\tau_d) = 0$ the characteristic polynomial is

$$-\frac{\tau_y \mathcal{T}(1-\beta)}{\beta^2},$$

which is > 0.

In sum, the set of parameters for which we have a determinate solution is defined through both an *increasing* curve $\phi^\star(\tau_d)$, and the vertical line $\tau_d = \tau_d^0$.

In sum, we can provide the following parametric boundaries implied by the first restriction on the second and third sets of conditions sufficient to postulate that the system is saddle-path stable:

- Let $\tau_d^0 = 1 - \beta - \frac{\beta \mathcal{T}}{\sigma} \frac{D^{SS}}{Y^{SS}}$, $\qquad \phi^\star(\tau_d) = \frac{\tau_y \mathcal{T}}{\sigma(\tau_d - 1 + \beta) + \beta \mathcal{T} \frac{D^{SS}}{Y^{SS}}} - \frac{\theta\kappa}{1-\beta}$.
- On $(\phi, \tau_d) \in \Big(-1/\sigma, \infty\Big) \times [0,1]:\quad 1 + \Gamma_2 + \Gamma_1 + \Gamma_0 > 0 \iff \begin{cases} \phi > \phi^\star(\tau_d), & \tau_d > \tau_d^0, \\ \phi < \phi^\star(\tau_d), & \tau_d < \tau_d^0. \end{cases}$
- Then, as $\mathcal{S} > 0$ and $\mathcal{T} < 0$, then $\phi^\star(\tau_d)$ is increasing with a vertical asymptote at τ_d^0.

—

PART 3: In the last part, two distinctions must be made:

Case 1: $\Gamma_0(\Gamma_0 - \Gamma_2) + \Gamma_1 - 1 > 0$.

We express $S_2(A)$ as

$$S_2(A) = \frac{1}{\beta^2}\Big[\beta \frac{D^{SS}}{Y^{SS}} \mathcal{T}\phi - \mathcal{T}\tau_y - \beta\phi\sigma\tau_d + 2\beta\phi\sigma - \beta\tau_d + 2\beta + \kappa\sigma\theta(\tau_d - 1) - \tau_d + 1\Big].$$

We define now Δ as a function in ϕ and τ_d:

$$\Delta(\phi, \tau_d) \equiv 1 - S_2(A) + \text{tr}(A)\det(A) - \det(A)^2.$$

For each $\tau_d \in [0,1]$, Δ is a quadratic polynomial in ϕ:

$$\Delta(\phi, \tau_d) = \frac{1}{\beta^4}\Big[\alpha(\tau_d)\phi^2 + b(\tau_d;\theta)\phi + c(\tau_d;\theta)\Big]\ ,$$

with the leading coefficient simplifying to

$$\alpha(\tau_d) = \sigma(1-\tau_d)\Big[\beta^2 \frac{D^{SS}}{Y^{SS}} \mathcal{T} + \sigma\Big(\beta^2 - (1-\tau_d)\Big)\Big]\ .$$

Two useful facts follow immediately:

- Since $1-\tau_d \geq 0$ and $\mathcal{T} \leq 0$, there exists a relatively easily reachable threshold

$$\tau_d > \tau_d^\sharp \equiv 1 - \left[\beta^2 + \frac{\beta^2 \frac{D^{SS}}{Y^{SS}} \mathcal{T}}{\sigma}\right]$$

  such that $\alpha(\tau_d) > 0$ for all $\tau_d > \tau_d^\sharp$ (as $\tau_d^\sharp \gtrsim 1-\beta^2$ and is typically small). When $\alpha(\tau_d) > 0$, $\Delta(\phi, \tau_d)$ is convex in ϕ.

- On that range, the boundary $\Delta = 0$ is given by:

$$\phi_\pm(\tau_d;\theta) = \frac{-b(\tau_d;\theta) \ \pm \ \sqrt{b(\tau_d;\theta)^2 - 4\alpha(\tau_d)c(\tau_d;\theta)}}{2\alpha(\tau_d)} \ ,$$

  and the inequality $\Delta(\phi, \tau_d) < 0$ holds between the two curves:

$$\alpha(\tau_d) > 0 \ \Rightarrow \quad \Delta < 0 \iff \phi_-(\tau_d;\theta) < \phi < \phi_+(\tau_d;\theta) \ .$$

For compactness we separate the θ-dependence of the linear and constant coefficients:

$$b(\tau_d;\theta) = b_0(\tau_d) + \theta b_1(\tau_d), \qquad c(\tau_d;\theta) = c_0(\tau_d) + \theta c_1(\tau_d),$$

such that we can express the θ-related slopes as:

$$b_1(\tau_d) = -\beta\kappa\sigma(1-\tau_d)\mathcal{S}, \qquad c_1(\tau_d) = \beta\kappa(1-\tau_d)(\beta\sigma - \mathcal{S})$$

while $b_0(\tau_d)$ and $c_0(\tau_d)$ collect the remaining θ-free terms, which are affine in τ_d and involve $\sigma, \beta, \mathcal{T}, \tau_y, \frac{D^{SS}}{Y^{SS}}$.

We now show that the effect of θ on the region $\Delta < 0$ is monotone. The previously defined function Δ co-moves with θ as follows:

$$\frac{\partial \Delta}{\partial \theta} = \frac{\kappa(1-\tau_d)}{\beta^3}\Big[\beta\sigma - \mathcal{S}\big(1+\sigma\phi\big)\Big].$$

Hence, for any fixed (ϕ, τ_d), we can find the following link between θ and $\phi_0^\star$:

$$\begin{cases} \dfrac{\partial \Delta}{\partial \theta} > 0 \iff \phi < \phi_\theta^\star \ \equiv \ \frac{\beta}{\mathcal{S}} - \frac{1}{\sigma}, \\ \dfrac{\partial \Delta}{\partial \theta} < 0 \iff \phi > \phi_\theta^\star. \end{cases}$$

Therefore, increasing θ shifts the $\Delta = 0$ boundary down in ϕ for large ϕ, and up in ϕ for small ϕ. Equivalently, the values for which $\Delta < 0$ expand toward larger ϕ when $\phi < \phi_\theta^\star$, and shrink when $\phi > \phi_\theta^\star$.

To summarize:

- The inequality $\Gamma_0(\Gamma_0 - \Gamma_2) + \Gamma_1 - 1 > 0$ holds exactly where $\Delta(\phi, \tau_d) < 0$.
- For most τ_d (specifically, when $\tau_d > \tau_d^\sharp$), $\Delta(\cdot, \tau_d)$ is convex in ϕ, so the feasible region is the band $\phi_-(\tau_d; \theta) < \phi < \phi_+(\tau_d; \theta)$.
- That feasible region shifts with θ according to $\partial_\theta \Delta = \frac{\kappa(1-\tau_d)}{\beta^3}\Big[\beta\sigma - \mathcal{S}(1 + \sigma\phi)\Big]$. Therefore, increasing θ loosens (tightens) the inequality for $\phi < \phi_\theta^\star$ ($\phi > \phi_\theta^\star$).

—

Case 2: $\Gamma_0(\Gamma_0 - \Gamma_2) + \Gamma_1 - 1 < 0$ AND $|\Gamma_2| > 3$.

As for the first part, $\Gamma_0(\Gamma_0 - \Gamma_2) + \Gamma_1 - 1 < 0$, the inverse of the previous statement holds true. As for the second part, the restriction that $|\operatorname{tr}(A)| > 3$ is more restrictive. Recall that $\operatorname{tr}(A) = C_0(\tau_d; \theta) + \mathcal{S}\phi$, which is an affine function of ϕ. Thus, define

$$\phi_{\rm tr}^+(\tau_d; \theta) \equiv \frac{3 - C_0(\tau_d; \theta)}{\mathcal{S}}, \qquad \phi_{\rm tr}^-(\tau_d; \theta) \equiv \frac{-3 - C_0(\tau_d; \theta)}{\mathcal{S}}.$$

Since $\mathcal{S} > 0$,

$$|\operatorname{tr}(A)| > 3 \iff \phi > \phi_{\rm tr}^+(\tau_d; \theta) \ \text{ or } \ \phi < \phi_{\rm tr}^-(\tau_d; \theta).$$

Because $C_0(\tau_d; \theta) = C_0(\tau_d; 0) - (\kappa\theta/\beta)\mathcal{S}$,

$$\frac{\partial \phi_{\rm tr}^\pm}{\partial \theta} = \frac{-(\partial C_0/\partial \theta)}{\mathcal{S}} = \frac{\kappa}{\beta} > 0,$$

so both thresholds shift to the right as θ increases.

We now look at the relevant comparative statics in θ. Recall that:

$$\frac{\partial \Delta}{\partial \theta} = \frac{\kappa(1 - \tau_d)}{\beta^3}\Big[\beta\sigma - \mathcal{S}(1 + \sigma\phi)\Big].$$

Thus, for given τ_d,

$$\frac{\partial \Delta}{\partial \theta} \begin{cases} > 0, & \phi < \phi_\theta^\star \equiv \frac{\beta}{\mathcal{S}} - \frac{1}{\sigma}, \\ < 0, & \phi > \phi_\theta^\star. \end{cases}$$

To keep $\Delta > 0$ when θ rises, ϕ must increase; i.e. $\phi_+(\tau_d; \theta)$ shifts right with θ. Since also $\partial_\theta \phi_{\rm tr}^+ = \kappa/\beta > 0$, the boundary $\Phi^\star(\tau_d; \theta)$ moves to the right as θ increases.

To summarize: within $\phi \in (-1/\sigma, \infty)$, $\tau_d \in [0, 1]$ and $\mathcal{S} > 0$, the set where

$$\Gamma_0(\Gamma_0 - \Gamma_2) + \Gamma_1 - 1 < 0 \quad \text{and} \quad |\Gamma_2| > 3$$

is fulfilled is described by a wedge

$$\phi > \Phi^\star(\tau_d; \theta) = \max\{\phi_+(\tau_d; \theta),\ \phi_{\rm tr}^+(\tau_d; \theta),\ -1/\sigma\}, \quad \tau_d > \tau_d^\sharp,$$

and this wedge *shrinks* (shifts to larger ϕ) as θ increases. Simply put, more inflation-indexed debt (higher θ) places additional restrictions on the admissible (ϕ, τ_d) space *from below.*

■

B.5 Additional Empirical Results

Tables B.1 and B.2 show, respectively, the results when ´we do not control for the composition of central bank board members and when we include also countries that never issued any inflation-linked government bonds in the sample period.

Dep. var.	AnyR	ER	BR	DR	AnyR	ER	BR	DR
CB independence	-0.9230	-0.3724	-1.5230	0.2394	-0.4363	6.8147**	-1.2566	-0.2508
	(0.9653)	(1.9799)	(0.9818)	(1.9364)	(1.2777)	(2.3326)	(1.2063)	(2.3983)
Indexed debt share	3.0495***	4.5955***	3.1201***	0.3687	6.2171***	12.0975***	6.1895***	1.9829
	(0.6649)	(0.9761)	(0.6850)	(0.8924)	(1.1265)	(1.5091)	(1.2104)	(1.3227)
Inflation	-0.0136	-0.0173	-0.0721*	-0.1561*	-0.0167	-0.1236***	-0.1319**	-0.2138*
	(0.0144)	(0.0181)	(0.0353)	(0.0651)	(0.0812)	(0.0255)	(0.0443)	(0.0975)
Constant	-3.0944***	-4.7794***	-2.6751***	-4.1030***	-1.2264	-8.0511***	-0.4638	-1.0064
	(0.5391)	(0.8917)	(0.5229)	(1.1886)	(0.9449)	(1.7184)	(0.8612)	(1.5822)
Obs.	652	652	652	652	285	261	285	95
ll	-104.2834	-53.1646	-91.7183	-42.2181	-59.2459	-19.6852	-56.3336	-24.7103
χ^2	24.9541	80.2726	24.3913	12.9324	72.2432	90.1966	71.7882	7.4952
p	0.0000	0.0000	0.0000	0.0048	0.0000	0.0000	0.0000	0.2775
R^2	0.0973	0.2555	0.1078	0.0217	0.3528	0.6571	0.3159	0.0998
Year-FE	No	No	No	No	Yes	Yes	Yes	Yes

Robust standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table B.1: Rule Suspension Regression. The outcome variables are defined as: $AnyR$ = "Any Fiscal Rule Suspended", ER = "Expenditure Rule Suspended", BR = "Budget Balancing Rule Suspended", DR = "Deficit Rule Suspended".

Dep. var.	Central Bank Independence Index (Romelli, 2024)							
CB board index	0.5595***	0.5172***	0.6292***	0.5358***	0.5654***	0.5114***	0.6235***	0.5255***
	(0.0085)	(0.0083)	(0.0075)	(0.0077)	(0.0086)	(0.0084)	(0.0076)	(0.0078)
Fiscal Rule Intensity	0.0164***	0.0096***	0.0198***	0.0095***	0.0139***	0.0091***	0.0202***	0.0092***
	(0.0015)	(0.0015)	(0.0012)	(0.0012)	(0.0015)	(0.0016)	(0.0012)	(0.0012)
Indexed debt share	-0.1972***	-0.2759***	0.4347***	0.1838***	-0.3216***	0.0779	0.4583***	0.1565***
	(0.0389)	(0.0373)	(0.0419)	(0.0292)	(0.0470)	(0.0745)	(0.0523)	(0.0336)
FisRules × IndexDebt			-0.0903***			-0.3081***	-0.0991***	
			(0.0235)			(0.0380)	(0.0256)	
Inflation					0.0000	0.0000	-0.0000	-0.0000
					(0.0000)	(0.0000)	(0.0000)	(0.0000)
IndexDebt × Inflation					0.0123***	0.0013	-0.0013	0.0017
					(0.0029)	(0.0031)	(0.0019)	(0.0016)
Constant	0.3144***	0.3415***	0.2725***	0.2826***	0.3184***	0.3505***	0.2798***	0.2879***
	(0.0044)	(0.0043)	(0.0037)	(0.0065)	(0.0044)	(0.0044)	(0.0038)	(0.0068)
Obs.	5568	5568	5568	5568	5282	5282	5282	5282
R^2	0.5092	0.5558	0.8560	0.8769	0.5166	0.5736	0.8549	0.8769
F	1924.2808	1554.2882	2458.8637	300.8337	1127.4948	768.2458	1580.1035	279.3330
R^2_{adj}	0.5089	0.5525	0.8518	0.8724	0.5161	0.5700	0.8505	0.8723
RMSE	0.1167	0.1114	0.0641	0.0595	0.1160	0.1094	0.0645	0.0596
Year-FE	No	Yes	No	Yes	No	Yes	No	Yes
Country-FE	No	No	Yes	Yes	No	No	Yes	Yes

Robust standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table B.2: Regressions on possible predictors of central bank independence following Romelli (2024). Zero-imputation applied to missing values of IndexedDebtShare, i.e., we use all countries here, including those who never issued any indexed debt.

Rule suspension regressions against lagged levels of indexed debt

Dep. var.	AnyR	ER	BR	DR	AnyR	ER	BR	DR
CB independence	-0.9616	-0.2181	-1.5826	0.1545	-0.4310	7.7258***	-1.2557	-0.2000
	(0.9171)	(1.9145)	(0.9429)	(1.9678)	(1.2830)	(2.2862)	(1.2018)	(2.4362)
L. Indexed debt share	2.9801***	4.5868***	3.1124***	0.5533	6.4393***	12.5273***	6.0911***	2.0002
	(0.6474)	(0.9392)	(0.6645)	(0.8735)	(1.1373)	(1.6212)	(1.1368)	(1.1816)
L. Inflation	-0.0095	-0.0130	-0.0756*	-0.1814***	0.0346	-0.0858**	-0.1113**	-0.1331**
	(0.0198)	(0.0253)	(0.0356)	(0.0541)	(0.1189)	(0.0274)	(0.0409)	(0.0507)
Constant	-3.0050***	-4.8858***	-2.5595***	-3.9959***	-1.1878	-9.2778***	-0.6361	-1.6762
	(0.5224)	(0.8759)	(0.5040)	(1.1948)	(0.9295)	(1.7401)	(0.9002)	(1.5650)
Obs.	632	632	632	632	286	262	286	96
ll	-106.3518	-52.3364	-93.5731	-41.7471	-58.9689	-18.6168	-57.3783	-25.2193
χ^2	24.9385	81.4065	24.3770	24.1283	71.4016	75.5276	79.6409	11.7768
p	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0671
R^2	0.0963	0.2622	0.1109	0.0270	0.3717	0.6760	0.3234	0.0841
Year-FE	No	No	No	No	Yes	Yes	Yes	Yes

Robust standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table B.3: Rule Suspension Regression with lagged regressors. Specification includes *L*.Indexed debt share and *L*.Inflation. Outcomes: *AnyR* (Any Rule Suspended), *ER* (Expenditure Rule), *BR* (Budget-Balance Rule), *DR* (Deficit Rule). Columns 1–4 exclude year fixed effects; columns 5–8 include year fixed effects (coefficients not shown).

Dep. var.	AnyR	ER	BR	DR	AnyR	ER	BR	DR
CB independence	-0.0808	3.7843*	-0.5463	0.2099	-0.0274	8.6885**	-0.7570	0.0139
	(1.1800)	(1.8603)	(1.2741)	(2.1038)	(1.9263)	(2.7690)	(1.8040)	(3.3163)
L. Indexed debt share	1.7687*	3.1557***	1.8284*	0.6844	3.2555**	7.4810***	2.8992**	2.3885
	(0.7415)	(0.9130)	(0.8265)	(0.9016)	(1.1305)	(1.7710)	(1.1188)	(1.7064)
L. Inflation	-0.0045	-0.0055	-0.0376	-0.1959**	0.0430	0.0052	-0.0569*	-0.1041
	(0.0393)	(0.0605)	(0.0283)	(0.0723)	(0.0534)	(0.0298)	(0.0277)	(0.0548)
Lagged Rule	4.2898***	5.2069***	4.5895***	4.9449***	0.0000	0.0000	0.0000	0.0000
	(0.5175)	(0.8183)	(0.5701)	(0.8629)	(.)	(.)	(.)	(.)
Constant	-4.1893***	-8.3652***	-3.9760***	-4.7069***	-3.2621	-9.3682***	-2.5648	-1.8548
	(0.7992)	(1.2691)	(0.8350)	(1.3290)	(1.6873)	(2.0013)	(1.5199)	(1.9630)
Obs.	632	632	632	632	286	48	286	48
ll	-70.8248	-28.8303	-59.8161	-29.0105	-28.9850	-7.1589	-28.0778	-12.5956
χ^2	81.0217	66.7653	75.5655	47.9633	.	29.0637	.	5.1460
p	0.0000	0.0000	0.0000	0.0000	.	0.0000	.	0.2726
R^2	0.3982	0.5936	0.4317	0.3238	0.6912	0.4800	0.6689	0.0852
Year-FE	No	No	No	No	Yes	Yes	Yes	Yes

Robust standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. "LDV" denotes the lag of the column's dependent variable; dots (.) indicate omission.

Table B.4: Rule Suspension Regression with lagged dependent variable (LDV), lagged indexed debt share, and lagged inflation (no interaction terms). Outcomes: *AnyR* (Any Rule Suspended), *ER* (Expenditure Rule), *BR* (Budget-Balance Rule), *DR* (Deficit Rule). Columns 1–4 exclude year fixed effects; columns 5–8 include year fixed effects (coefficients not shown).

Chapter C

Appendix to: The Open-Economy Debt Valuation Equation and International Shock Transmission

C.1 Omitted Derivations of the Simple Model

Solution for equilibrium

Definition 4 *An allocation* $\left\{c_t, c_t^*, L_t, L_t^*, B_{ROW,t}, B_{SOE,t}, B_{ROW,t}^*, B_{SOE,t}^*, B_t, B_t^*\right\}_{t=1,2}$, *nominal prices* $\{P_t, P_t^*, \mathcal{E}_t\}_{t=1,2}$, *fiscal policy* $\{T_t, T_t^*\}_{t=1,2}$, *monetary policy* $\left\{\bar{i}, i_t^*\right\}_{t=1,2}$, *and government spending* $\{G_2\}$ *are called a competitive equilibrium in this economy, when:*

- *all time-t variables are t-measurable;*
- *given* $\{P_t, \bar{i}, T_t, P_t^*, i_t^*, T_t^*\}$ *and* $\{G_2\}$, *the variables* $\{c_t, L_t, B_{ROW,t}, B_{SOE,t}\}_{t=1,2}$ *solve equations (3.13), (3.14), (3.15), and (3.11); and the variables* $\{c_t^*, L_t^*, B_{ROW,t}^*, B_{SOE,t}^*\}_{t=1,2}$ *solve equations (3.16), (3.17), (3.18), and (3.12), with* $B_{ROW,2} = B_{SOE,2} = B_{ROW,2}^* = B_{SOE,2}^* = 0$.
- *The domestic government satisfies its budget constraint and policy rules:*
    - $B_t = B_{t-1}(1+\bar{i}) + P_t(G_t - T_t),\ t = 1, 2,$ *where* $B_2 = 0$.
    - $T_2 = G_2 + \frac{B_1(1+\bar{i})}{P_2^T}$ *in regime ML;* $T_2 = \bar{T}$ *in regime FL. Taxes in period 1 adjust to ensure a constant real debt burden.*
    - *Interest rates are pegged to the value* $\bar{i}$.
- *The SOE government satisfies its budget constraint and policy rules:*
    - $B_t^* = B_{t-1}^*(1+i_{t-1}^*) + P_t^*(G_t^* - T_t^*)\ t = 1, 2, 3$ *and* $B_2^* = 0$.

  - *Taxes in all periods are fixed. In $t = 1$, taxes offset the cost of serving maturing debt, while in $t = 2$, taxes are simply given by some exogenous level $\bar{T}^*$.*
  - *Interest rates i_t^* adjust to make sure that household optimality is achieved as the sole free parameter ensuring that the UIP conditions on both households hold.*

- *Markets clear:*
  - $C_t + G_t = L_t$ *for* $t = 1, 2$,
  - $C_t^* + G_t^* = L_t^*$ *for* $t = 1, 2$,
  - $B_t = B_{ROW,t} + B_{ROW,t}^*$ *for* $t = 1, 2$,
  - $B_t^* = B_{SOE,t} + B_{SOE,t}^*$ *for* $t = 1, 2$.

In total, we have 17 endogenous variables in 17 conditions.

Now, we solve for the equilibrium of this model. We start from the terminal period ($t = 2$) and move backwards in solving the model.

To pin down the terminal exchange rate $\mathcal{E}_2$ without invoking the UIP (as there are no bonds outstanding in period 3), we start by recognizing that the model set-up requires all cross-border bond holdings to unravel at the terminal period. In an accounting sense, cross-border flows of financial goods in the terminal period must allow the ROW country to repatriate the SOE holdings of ROW bonds, and vice versa. Therefore, the terminal exchange rate reflects the price level-adjusted flows that allow cross-country bond holdings to unravel. The logic is that, in equilibrium,

$$P_2(1 + \bar{i})B_{ROW,1}^* = \mathcal{E}_2 P_2^*(1 + i_1^*)B_{SOE,1} \qquad \Leftrightarrow \qquad \mathcal{E}_2 = \frac{P_2}{P_2^*}\frac{(1 + \bar{i})}{(1 + i_1^*)}\frac{B_{ROW,1}^*}{B_{SOE,1}}, \qquad \text{(C.1)}$$

such that the non-domestic bond holdings in period 1 in combination with their respective interest rates are fully informative about the nominal exchange rate in period 2. This assumption allows us to close the outstanding stock of cross-border financial accounts in the terminal period.

In economic terms, this assumption equates the value of SOE debt held by ROW and ROW debt by SOE in the last period. Therefore, it facilitates a *positive* correlation between the value of ROW and SOE debt in the terminal period. Assume a fiscal expansion in ROW, reducing the real value of ROW debt in the terminal period. By this assumption, the real value of SOE debt denoted in the ROW currency must also decrease, allowing for a simple reduced-form international transmission of fiscal shocks.

This is, in a sense, an assumption about absolute hegemony of ROW: since SOE fiscal policy is assumed to be fixed, a fiscal deterioration at ROW must translate to a fiscal deterioration in SOE by depressing the market value of SOE debt. To some extent, this rationalizes the taste for ROW bonds in SOE, manifest in its utility function.

As alluded to previously, terminal bond holdings are zero by household optimality: $B_2 = B_2^* = B_{ROW,2} = B_{SOE,2} = B_{ROW,2}^* = B_{SOE,2}^* = 0$. Additionally, for our further analysis, we denote the measure of fiscal spending in the ROW country conditional on being in the fiscally-led regime as $\mu_1 \equiv \mathbb{E}_1(G_2|FL)$.

Conditional on the fiscal policy regime in place, we are able to pin down the price level in the terminal period at ROW through the specification of the fiscal policy regime in place:

$$P_2 = \begin{cases} P_2^T & \text{if ML,} \\ \frac{B_1(1+\bar{i})}{\bar{T}-G_2} & \text{if FL.} \end{cases} \tag{C.2}$$

Therefore, the level of taxes levied at ROW in the terminal period is given by:

$$T_2 = \begin{cases} \frac{B_1(1+\bar{i})}{P_2^T} + G_2 & \text{if ML,} \\ \bar{T} & \text{if FL.} \end{cases} \tag{C.3}$$

Since $G_1 = 0$ and $u'(c_t) = 1\ \forall t$, we can simplify the Euler equation applying in period 1 in the ROW household, which is:

$$u'(c_1) = \beta(1+\bar{i})\mathbb{E}_t\left[u'(c_2)\frac{P_1}{P_2}\right],$$

by recognizing that household optimality prescribes constant consumption across time, such that the Euler equation becomes:

$$\frac{1}{P_1} = \beta(1+\bar{i})\left[\psi_1\frac{\bar{T}-\mu_1}{B_1(1+\bar{i})} + (1-\psi_1)\frac{1}{P_2^T}\right]. \tag{C.4}$$

Alternatively, we can express also the price level explicitly as:

$$P_1 = \frac{1}{\beta(1+\bar{i})}\frac{1}{\left[\psi_1\frac{\bar{T}-\mu_1}{B_1(1+\bar{i})} + (1-\psi_1)\frac{1}{P_2^T}\right]}, \tag{C.5}$$

which means that we have now pinned down the domestic price level in both period 1 and period 2. It generally does not depend on the exchange rate and SOE fiscal or monetary policy. This is because our assumptions implicitly allowed us to characterize the ROW policy authorities as the 'dominant' ones, imposing their tax and interest policy on the other country through the exchange rate, with little effect vice versa. We are therefore taking here the stark assumption of an 'absolute hegemon' in terms of the monetary and fiscal policy spillovers.

Now, we focus on the SOE country as well as the international dimension. To aid our discussion, we first solve for the SOE price level in period 1. By the SOE Euler condition on SOE bonds,

$$\frac{1}{P_t^*} = \beta(1+i_t^*)\mathbb{E}_t\left[\frac{1}{P_{t+1}^*}\right].$$

We evaluate this at $t = 1$. Since

$$P_2^* = \frac{(1+i_1^*)B_1^*}{\bar{T}^*} \tag{C.6}$$

by the Euler equation of SOE, we find that

$$\frac{1}{P_1^*} = \beta\frac{\bar{T}^*}{B_1^*} \tag{C.7}$$

fully determines the price level. In particular, as P_2^* is fully known conditional on i_1^*, we can pin down the price level in SOE at $t = 1$ as:

$$P_1^* = \frac{1}{\beta}\frac{B_1^*}{\bar{T}^*},$$

where i_0^* was inherited from the initial steady-state.

What is left to determine are $\mathcal{E}_1$, $\mathcal{E}_2$, and i_1^*, which will be jointly given by our two UIP conditions and the terminal exchange rate condition.

Proof of Proposition 5

Proof. We proceed in three steps: first, we derive an expression for the exchange rate $\mathcal{E}_1$ from the ROW household's Euler equation on SOE bonds; second, we derive another expression for $\mathcal{E}_1$ from the SOE household's Euler equation on ROW bonds; third, we equate the two expressions and solve for $(1 + i_1^*)$.

We take the modified UIP condition on SOE bonds, equation (3.15), to get:

$$(1+\phi_{ROW}) = \beta(1+i_1^*)\,\mathbb{E}_1\left[\frac{P_1}{P_2}\frac{\mathcal{E}_2}{\mathcal{E}_1}\right].$$

We now insert the terminal exchange rate condition (C.1) after rewriting it as $\frac{\mathcal{E}_2}{P_2} = \frac{1}{P_2^*}\frac{(1+\bar{i})}{(1+i_1^*)}\frac{B^*_{ROW,1}}{B_{SOE,1}}$. This yields:

$$(1+\phi_{ROW}) = \beta(1+i_1^*)\frac{P_1}{\mathcal{E}_1}\frac{(1+\bar{i})}{(1+i_1^*)}\frac{B^*_{ROW,1}}{P_2^*\,B_{SOE,1}}.$$

Note that the $(1 + i_1^*)$ terms cancel, leaving:

$$(1+\phi_{ROW}) = \frac{\beta\,P_1\,(1+\bar{i})\,B^*_{ROW,1}}{\mathcal{E}_1\,P_2^*\,B_{SOE,1}}.$$

Solving for $\mathcal{E}_1$ and substituting $P_2^* = \frac{(1+i_1^*)B_1^*}{\bar{T}^*}$ from equation (C.6):

$$\mathcal{E}_1 = \frac{\beta\, P_1\, (1+\bar{i})\, \bar{T}^*\, B^*_{ROW,1}}{(1+\phi_{ROW})\,(1+i_1^*)\, B_1^*\, B_{SOE,1}}. \tag{C.8}$$

Now, we take the SOE household's Euler equation on ROW bonds, equation (3.17):

$$1 = \beta(1+\bar{i})\,\mathbb{E}_1\left[\frac{P_1^*}{P_2^*}\frac{\mathcal{E}_1}{\mathcal{E}_2}\right] + \phi_{SOE}\,\mathcal{E}_1.$$

From the terminal exchange rate condition (C.1), we compute:

$$\frac{1}{\mathcal{E}_2} = \frac{P_2^*}{P_2}\frac{(1+i_1^*)}{(1+\bar{i})}\frac{B_{SOE,1}}{B^*_{ROW,1}},$$

so that:

$$\frac{\mathcal{E}_1}{\mathcal{E}_2} = \mathcal{E}_1\frac{P_2^*}{P_2}\frac{(1+i_1^*)}{(1+\bar{i})}\frac{B_{SOE,1}}{B^*_{ROW,1}}.$$

Substituting into the Euler equation:

$$1 = \beta(1+\bar{i})\frac{P_1^*}{P_2^*}\cdot\mathcal{E}_1\cdot\frac{P_2^*}{1}\cdot\frac{(1+i_1^*)}{(1+\bar{i})}\cdot\frac{B_{SOE,1}}{B^*_{ROW,1}}\cdot\mathbb{E}_1\left[\frac{1}{P_2}\right] + \phi_{SOE}\,\mathcal{E}_1.$$

The P_2^* terms cancel, as do the $(1+\bar{i})$ terms, leaving:[1]

$$1 = \beta\, P_1^*\,(1+i_1^*)\,\frac{B_{SOE,1}}{B^*_{ROW,1}}\,\mathbb{E}_1\left[\frac{1}{P_2}\right]\,\mathcal{E}_1 + \phi_{SOE}\,\mathcal{E}_1.$$

Now substitute $P_1^* = \frac{B_1^*}{\beta\bar{T}^*}$ from equation (C.7), and $\mathbb{E}_1\left[\frac{1}{P_2}\right] = \psi_1\frac{\bar{T}-\mu_1}{B_1(1+\bar{i})} + (1-\psi_1)\frac{1}{P_2^T}$ from equations (C.2) and (C.4):

$$1 = \frac{B_1^*}{\bar{T}^*}\,(1+i_1^*)\,\frac{B_{SOE,1}}{B^*_{ROW,1}}\left[\psi_1\frac{\bar{T}-\mu_1}{B_1(1+\bar{i})} + (1-\psi_1)\frac{1}{P_2^T}\right]\mathcal{E}_1 + \phi_{SOE}\,\mathcal{E}_1.$$

Solving for $\mathcal{E}_1$:

$$\mathcal{E}_1 = \frac{1}{\frac{B_1^*}{\bar{T}^*}\,(1+i_1^*)\,\frac{B_{SOE,1}}{B^*_{ROW,1}}\left[\psi_1\frac{\bar{T}-\mu_1}{B_1(1+\bar{i})} + (1-\psi_1)\frac{1}{P_2^T}\right] + \phi_{SOE}}. \tag{C.9}$$

Finally, we now equate (C.8) and (C.9). First, substitute $P_1 = \frac{1}{\beta(1+\bar{i})\left[\psi_1\frac{\bar{T}-\mu_1}{B_1(1+\bar{i})}+(1-\psi_1)\frac{1}{P_2^T}\right]}$ from equation (C.5) into (C.8) to obtain:

[1]Note that P_1^*, P_2^*, and i_1^* are all known in period 1, so we can take them outside the expectations operator. The only random variable under the expectation is $1/P_2$.

$$\mathcal{E}_1 = \frac{\bar{T}^* \, B^*_{ROW,1}}{(1+\phi_{ROW})\,(1+i^*_1)\,B^*_1\,B_{SOE,1}\left[\psi_1 \frac{\bar{T}-\mu_1}{B_1(1+\bar{i})} + (1-\psi_1)\frac{1}{P^T_2}\right]}.$$

This takes the form $\mathcal{E}_1 = \frac{1}{(1+\phi_{ROW})\cdot D}$, where we define:

$$D \equiv \frac{B^*_1}{\bar{T}^*}\,(1+i^*_1)\,\frac{B_{SOE,1}}{B^*_{ROW,1}}\left[\psi_1 \frac{\bar{T}-\mu_1}{B_1(1+\bar{i})} + (1-\psi_1)\frac{1}{P^T_2}\right].$$

Meanwhile, equation (C.9) takes the form $\mathcal{E}_1 = \frac{1}{D+\phi_{SOE}}$. Equating:

$$\frac{1}{(1+\phi_{ROW})\cdot D} = \frac{1}{D+\phi_{SOE}}.$$

Cross-multiplying:

$$D + \phi_{SOE} = (1+\phi_{ROW})\cdot D = D + \phi_{ROW}\cdot D,$$

which simplifies to:

$$\phi_{SOE} = \phi_{ROW}\cdot D = \phi_{ROW}\,\frac{B^*_1}{\bar{T}^*}\,(1+i^*_1)\,\frac{B_{SOE,1}}{B^*_{ROW,1}}\left[\psi_1 \frac{\bar{T}-\mu_1}{B_1(1+\bar{i})} + (1-\psi_1)\frac{1}{P^T_2}\right].$$

Solving for $(1 + i^*_1)$:

$$(1+i^*_1) = \frac{\phi_{SOE}\,\bar{T}^*\,B^*_{ROW,1}}{\phi_{ROW}\,B^*_1\,B_{SOE,1}\left[\psi_1 \frac{\bar{T}-\mu_1}{B_1(1+\bar{i})} + (1-\psi_1)\frac{1}{P^T_2}\right]}. \tag{C.10}$$

Note that for $\phi_{ROW} = 0$ (no friction on the ROW side), the equation $\phi_{SOE} = \phi_{ROW}\cdot D$ reduces to $\phi_{SOE} = 0$, implying that no equilibrium with $\phi_{SOE} > 0$ exists when the ROW household faces no cost of holding SOE bonds. The presence of $\phi_{ROW} > 0$ is therefore necessary for existence. When additionally $\phi_{SOE} = 0$, the expression yields $(1 + i^*_1) = 0$; in that case, without any cross-border preference frictions, the model reverts to standard indeterminacy of i^*_1 absent a separate policy rule, and one would require $i^*_1 = \bar{i}$ for bond markets to clear. ■

Proof of Corollary 2

Proof. Taking the first derivative of equation (3.19) with respect to ψ_1:

$$\frac{\partial(1+i^*_1)}{\partial\psi_1} = -\frac{\phi_{SOE}\,\bar{T}^*\,B^*_{ROW,1}}{\phi_{ROW}\,B^*_1\,B_{SOE,1}}\cdot\frac{1}{\left[\psi_1 \frac{\bar{T}-\mu_1}{B_1(1+\bar{i})} + (1-\psi_1)\frac{1}{P^T_2}\right]^2}\cdot\left(\frac{\bar{T}-\mu_1}{B_1(1+\bar{i})} - \frac{1}{P^T_2}\right).$$

Now, the assumption that $P_2|_{FL} > P_2|_{ML}$ means that $\frac{B_1(1+\bar{i})}{\bar{T}-\mu_1} > P_2^T$, which in turn implies:

$$\frac{\bar{T}-\mu_1}{B_1(1+\bar{i})} < \frac{1}{P_2^T},$$

so the last factor $\left(\frac{\bar{T}-\mu_1}{B_1(1+\bar{i})} - \frac{1}{P_2^T}\right)$ is strictly negative. The leading minus sign then yields a positive overall sign. The prefactor $\frac{\phi_{SOE}\,\bar{T}^*\,B^*_{ROW,1}}{\phi_{ROW}\,B_1^*\,B_{SOE,1}}$ is strictly positive since $\phi_{SOE} > 0$, $\phi_{ROW} > 0$, and all bond holdings and fiscal parameters are positive. The squared term in the denominator is also strictly positive. Therefore:

$$\frac{\partial(1+i_1^*)}{\partial\psi_1} > 0.$$

■

C.2 Additional Details on the Empirical Framework

Further Details on the VAR Framework

This appendix presents the derivations underpinning the VAR model estimated in Section 3.2. The derivations follow Cochrane (2019), which we extend to an open-economy framework. We begin with a narrative of the major UK fiscal episodes visible in Figure C.1, broadly in line with Ellison and Scott (2020), before turning to the backward and forward decompositions and the variance decomposition.

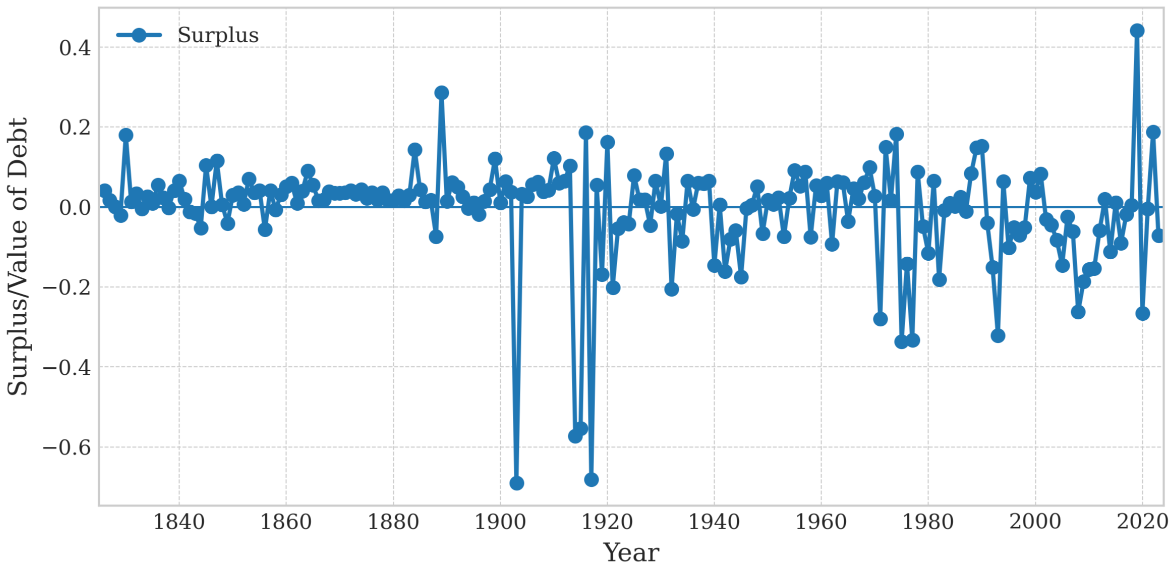

Figure C.1: Implied UK surpluses to ensure the government flow identity holds.

C. Appendix to: The Open-Economy Debt Valuation Equation and International Shock Transmission

The last half-century covered by Figure C.1 divides naturally into two broad regimes, each punctuated by acute crises. The first regime, spanning roughly 1970 to 2000, is characterised by the transition from fiscal dominance to a credible nominal anchor. The March 1972 "dash for growth" budget combined large tax cuts with relaxed credit controls, pushing borrowing requirements to £3.4 billion, then 5% of GDP. The resulting demand boom collided with the 1973 oil shock, sending CPI inflation into double digits and gilt yields to post-war highs. For the first time in three centuries, Britain ran substantial primary deficits outside of wartime. Market participants began to price a fiscal dominance regime: Sterling sold off, gilt investors demanded an inflation premium, and the surplus-over-debt ratio remained below zero until the 1976 IMF programme. The 1980 Medium-Term Financial Strategy then imposed cash limits on departmental spending, turning a structural deficit into a modest surplus within four years, and subsequent positive surpluses persisted through the Thatcher administration. This consolidation was interrupted by the twin blows of a deep domestic recession and "Black Wednesday" in 1992–93: when speculative pressure forced Sterling out of the Exchange Rate Mechanism, the Treasury raised base rates to 15% in a failed defence of the parity, output contracted, and tax receipts collapsed. On the valuation side, the real exchange rate dropped 15% within weeks, instantly re-inflating the domestic-currency value of foreign-currency gilts. The subsequent adoption of inflation targeting stabilised expectations, and by 1995 the surplus ratio had drifted back towards zero. Between 1996 and 2000, robust productivity growth, a high tax base, and restrained departmental spending kept the cash deficit near balance, while falling global real rates reduced the discount factor applied to future surpluses, so that only small primary surpluses were required to keep the ratio flat.

The second regime, from the early 2000s onward, is dominated by two large crisis-driven fiscal expansions. The Global Financial Crisis produced the deepest post-war trough in the surplus ratio during 2009–12: GDP fell 4.5% in 2009, policy aimed at stabilising the banking system pushed the headline deficit above 10% of GDP, and automatic stabilisers added a further two percentage points. Simultaneously, Bank of England gilt purchases suppressed real yields to multi-decade lows, which the accounting identity treats as requiring even larger negative surpluses to match the observed market value of debt; public sector net debt jumped from 41% to 81% of GDP. From 2013, the coalition's fiscal consolidation and a stabilisation of real rates nudged the ratio back toward zero, though the path was flat rather than rising: low discount factors meant a balanced-budget stance was sufficient for debt stabilisation without the large primary surpluses of the 1980s. Modest surpluses in 2017–19 briefly lifted the ratio into positive territory, but the Covid-19 pandemic reversed these gains. Emergency spending, furlough schemes, and revenue shortfalls generated a deficit of 14.6% of GDP in 2020, yet in Figure C.1

the 2020 observation *rises* to +0.42. The explanation lies in the valuation channel: gilt real yields turned sharply negative and inflation picked up, generating a large *implied* surplus in present-value terms even as the cash balance deteriorated. This is a Lucas and Stokey (1983) state-contingent financing of debt through surprise inflation. The episode encapsulates the fiscal-theory message: debt dynamics hinge on the path of discount factors and prices as much as on the arithmetic of tax minus spending.

Taken together, Figure C.1 shows that shifts in the surplus-over-debt ratio align closely with the regime breaks and policy narratives documented by Bordo et al. (2022). Episodes that placed acute pressure on public finances, whether war financing, the 1970s stagflation, or the Global Financial Crisis, drive the surplus ratio deep into negative territory, whereas phases of deliberate consolidation or windfall receipts keep it persistently positive. Crucially, the accounting framework highlights that shifts in valuation terms, through unexpected inflation and movements in discount rates, can do as much of the adjustment work as the cash-flow surpluses themselves.

Additional Details on the Backward Decomposition

The present-value identity in Section 3.2 converts a sequence of forward flows into one equilibrium condition that we can take to the data and decompose into economically meaningful components. We now derive this identity. First, we iterate both flow identities forward to arrive at present-value identities for home and foreign debt:

$$v_t^H = \sum_{j=1}^{\infty} \rho_H^{j-1} s_{t+j}^H + \sum_{j=1}^{\infty} \rho_H^{j-1} g_{t+j}^H - \sum_{j=1}^{\infty} \rho_H^{j-1} \Big[r_{t+j}^{n,H} - \pi_{t+j}^H \Big], \tag{C.11}$$

$$v_t^F = \sum_{j=1}^{\infty} \rho_F^{j-1} s_{t+j}^F + \sum_{j=1}^{\infty} \rho_F^{j-1} g_{t+j}^F - \sum_{j=1}^{\infty} \rho_F^{j-1} \Big[r_{t+j}^{n,H} - \Delta e_{t+j} - \varphi_{t+j} - \pi_{t+j}^F \Big]. \tag{C.12}$$

Investors price US and UK bonds jointly. Persistent Dollar appreciation or a sustained fall in the UIP premium lowers the discount factor for US surpluses while raising it for UK surpluses, binding the two budget constraints in expectation. Rewriting the forward identity for the foreign country yields

$$v_t^F = \sum_{j\geq 1} \rho_F^{j-1} s_{t+j}^F + \sum_{j\geq 1} \rho_F^{j-1} g_{t+j}^F - \sum_{j\geq 1} \rho_F^{j-1} r_{t+j}^{n,H} + \sum_{j\geq 1} \rho_F^{j-1} \Delta e_{t+j} + \sum_{j\geq 1} \rho_F^{j-1} \varphi_{t+j} + \sum_{j\geq 1} \rho_F^{j-1} \pi_{t+j}^F. \tag{C.13}$$

Equation (C.13) holds ex ante as well as ex post. Taking expectations conditional on information at time t gives

$$v_t^F = \underbrace{\mathbb{E}_t\left[\sum_{j\geq 1}\rho_F^{j-1}s_{t+j}^F\right]}_{PV_t^{(s)}} + \underbrace{\mathbb{E}_t\left[\sum_{j\geq 1}\rho_F^{j-1}g_{t+j}^F\right]}_{PV_t^{(g)}} - \underbrace{\mathbb{E}_t\left[\sum_{j\geq 1}\rho_F^{j-1}r_{t+j}^{n,H}\right]}_{PV_t^{(r)}}$$
$$+ \underbrace{\mathbb{E}_t\left[\sum_{j\geq 1}\rho_F^{j-1}\Delta e_{t+j}\right]}_{PV_t^{(\Delta e)}} + \underbrace{\mathbb{E}_t\left[\sum_{j\geq 1}\rho_F^{j-1}\varphi_{t+j}\right]}_{PV_t^{(\varphi)}} + \underbrace{\mathbb{E}_t\left[\sum_{j\geq 1}\rho_F^{j-1}\pi_{t+j}^F\right]}_{PV_t^{(\pi)}}, \tag{C.14}$$

where $\rho_F \simeq 0.97$ is the annual Campbell–Shiller discount factor. For convenience we gather the real discount-rate components as

$$PV_t^{(d)} \equiv PV_t^{(r)} - PV_t^{(\pi)}.$$

Equation (C.14) can therefore be written as

$$v_t^F = PV_t^{(s)} + PV_t^{(g)} + PV_t^{(\pi)} - \left[PV_t^{(r)} + PV_t^{(\Delta e)} + PV_t^{(\varphi)}\right]. \tag{C.15}$$

Our goal is to turn the forward-looking identity in (C.15) into measurable objects that can be tracked year by year.

Turning to the trends in the evolution of UK debt visible in Figure 3.1, five distinct episodes emerge, though they are best understood in terms of the interaction between flow and valuation channels rather than in isolation.

The Sterling crisis of 1975–81 illustrates the dominance of valuation adjustment over fiscal retrenchment. The $-s$ line rises rapidly, reflecting large primary deficits, yet the debt ratio v increases only modestly because a cumulative 25% depreciation (the REER series) and double-digit inflation (the $-\pi$ series) erode the real value of gilts. Bordo et al. (2022) date this as the first peacetime episode of fiscal dominance since the 1690s, and the backward decomposition confirms that the adjustment came predominantly through the price level and the exchange rate.

During the subsequent North Sea revenue and disinflation period of 1982–2001, a very different pattern obtains. The $-s$ line flattens as small primary surpluses emerge, while the r^n series climbs as global real rates rise with the Volcker disinflation, so the $r^n - \pi - g$ bundle pulls the debt ratio upward even in the face of favourable flows. Roughly half of the cumulative drift in debt over these two decades is valuation-driven rather than cashflow-driven. Embedded within this period, Sterling's 1992 departure from the ERM produces a sharp kink in the REER line, offset by a drop in r^n as gilts rally; these opposing valuation effects leave the debt path broadly unchanged, masking the cyclical widening and subsequent narrowing of primary balances.

The Global Financial Crisis of 2007–13 creates the largest wedge in the sample. Primary deficits widen sharply, and simultaneously the r^n line falls as world safe real rates plunge; the two forces reinforce each other, pushing the debt series sharply higher. Cochrane (2019) notes that in the US decomposition approximately 52% of the post-2008 rise in debt is discount-rate news; the UK share is similar *only* once exchange rate channels are included, hinting at different debt dynamics for financial hegemons relative to smaller open economies.

Finally, the pandemic shock of 2020–22 generates record cash deficits alongside a large valuation gain: gilt real yields turn deeply negative, inflation re-accelerates, and the UIP curve spikes as investors demand a wider risk premium on Sterling assets. Roughly one half of the pandemic-era debt increase is therefore absorbed by prices rather than by flows.

C.3 Linking Deficits to Exchange Rates

Complementing section 3.2, we provide additional results akin to the 'reverse regression' approach from Engel and West (2005) and Jiang (2021).

$$\frac{s^i_{t+k}}{B^i_{t+k}} = \alpha^{(k)} + \beta^{(k)}\,\Delta\log e_t^{(-i)/i} + \delta_1^{(k)}\,\frac{s^i_t}{B^i_t} + \delta_2^{(k)}\,\Delta\frac{s^i_t}{B^i_t} + \gamma_1^{(k)}\,\frac{s^{-i}_t}{B^{-i}_t} + \gamma_2^{(k)}\,\Delta\frac{s^{-i}_t}{B^{-i}_t} + \varepsilon^i_{t+k}, \quad \text{(C.16)}$$

that is; we test the link between the deficit/debt-ratio at time $t+k$ and the change in the real exchange rate as well as the level and the first-difference of *both* domestic and foreign deficit/debt-ratio levels. The results of this exercise are presented in table C.1 and figure C.2 for the UK.

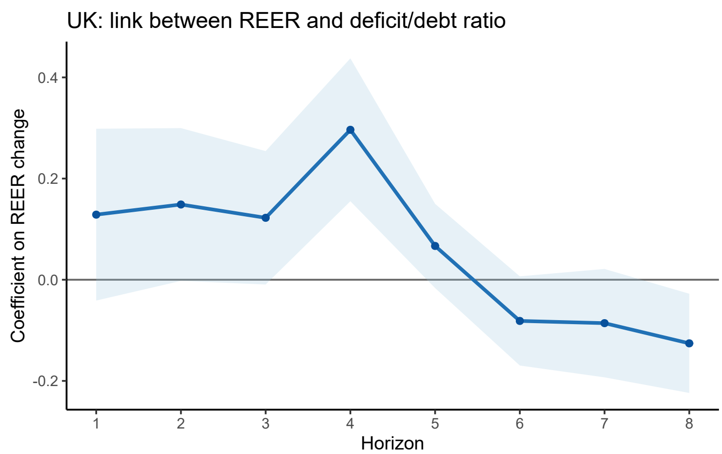

Figure C.2: Engel and West (2005)-style reverse regression for the U.K.

	Dependent variable: deficit-to-debt ratio at t+k							
	k=1	k=2	k=3	k=4	k=5	k=6	k=7	k=8
Δ REER	0.129	0.149	0.122	0.296***	0.067	−0.081	−0.086	−0.126**
	(0.104)	(0.092)	(0.080)	(0.086)	(0.051)	(0.054)	(0.065)	(0.060)
$\frac{s}{B}$-UK	0.187	−0.491**	−0.858***	−1.109***	−0.905**	−0.287	−0.142	−0.219
	(0.256)	(0.201)	(0.224)	(0.200)	(0.410)	(0.260)	(0.289)	(0.478)
$\Delta\frac{s}{B}$-UK	0.184	0.090	0.729**	0.729*	0.689	1.389**	0.604	0.706
	(0.202)	(0.276)	(0.348)	(0.405)	(0.584)	(0.626)	(0.461)	(0.608)
$\frac{s}{B}$-US	0.481**	0.814***	0.827***	0.897***	0.684	−0.005	−0.125	−0.192
	(0.194)	(0.229)	(0.256)	(0.188)	(0.421)	(0.281)	(0.291)	(0.409)
$\Delta\frac{s}{B}$-US	−0.073	−0.018	−0.647*	−0.960	−0.752	−0.824	−0.038	−0.149
	(0.246)	(0.220)	(0.386)	(0.588)	(0.673)	(0.667)	(0.458)	(0.560)
Constant	−0.010	−0.024	−0.048***	−0.058***	−0.067***	−0.095***	−0.097***	−0.109***
	(0.010)	(0.015)	(0.016)	(0.010)	(0.016)	(0.012)	(0.013)	(0.017)
Obs.	29	28	27	26	25	24	23	22
R^2	0.659	0.436	0.365	0.548	0.287	0.354	0.251	0.387
Adj. R^2	0.585	0.308	0.214	0.435	0.099	0.175	0.030	0.195

Table C.1: UK: Engel and West (2005)-style regressions for the deficit-to-debt ratio. Newey-West standard errors in parentheses. *p<0.1; **p<0.05; ***p<0.01.

C.4 Additional Details on the SOE Model

Preparation of data for the estimation

In our estimation, we make use of a number of time series, which we describe here in detail. The full sample runs from 1971 Q3 until 2025 Q4 without any interruptions, corresponding to the post-Bretton Woods period in which the Pound Sterling has been free-floating.

- real output Y_t: FRED data series NGDPRSAXDCGBQ, normalized in log-terms. To isolate cyclical variation, we estimate a sixth-degree polynomial of output over time and take the deviations of observed output from trend, following Ramey and Zubairy (2018).[2]

- GDP deflator: FRED data series NGDPDSAIXGBQ. The raw data series was used to normalize all variables otherwise denoted in nominal terms, such that their units are comparable to the real output definition above.

- CPI π_t: FRED data series GBRCPIALLQINMEI, expressed in log-differences over time. Used as the measure of inflation relevant to household intertemporal decision-making.

[2]In line with the criticisms levied by Hamilton (2018), we abstain from using the Hodrick-Prescott filter to isolate cyclical variation for the purposes of estimation.

- Nominal exchange rate e_t: FRED data series DEXUSUK, expressed in log-differences over time. Used as our headline measure of the nominal exchange rate.
- Return on long-bonds R_t^L: per-period holding returns on the portfolio of UK government debt as created by Ellison and Scott (2020). The analysis was updated for the period 2019-2025 using data from Cairns and Wilkie (2026).
- Returns on the portfolio of foreign (here: US) bonds: data from Hall et al. (2018), available on George Hall's website until December 2025.
- Total market value of UK debt, B_t, split into long-term and short-term bonds: Ellison and Scott (2020). The analysis was updated for the period 2019-2025 using data from Cairns and Wilkie (2026).
- Total market value of foreign (here: US) debt, B_t^F: data from Hall et al. (2018), available on George Hall's website until December 2025. We further adjust for the foreign holdings of treasuries using the FRED time series BOGZ1FL263061145Q and BOGZ1FL263061130Q.
- Estimates on the convenience yield spread between US and UK treasuries, θ_t: Jiang et al. (2021) for the period 1970-2018; since 2018, the data series of Du et al. (2025) (available on Wenxin Du's website) has been used.
- consumption C_t: ONS data series ABJR, normalized in log-terms. To isolate cyclical variation, we estimate a sixth-degree polynomial of consumption over time and take the deviations of observed consumption from trend.
- Net borrowing by home $-s_t$: ONS data series J5II. We normalize deficits by nominal GDP.
- Current Account deficit of the UK ca_t: ONS data series AA6H. We normalize the current account by nominal GDP.
- UK Bank Rate, i_t: FRED data series BOERUKM. To find the relevant variation in i_t relative to the business cycle state, we normalize using the r-star estimate for the UK of Lubik-Merone-Robino.
- Fed Funds Rate, i_t^*: FRED data series FEDFUNDS. To find the relevant variation in i_t^* relative to the business cycle state, we normalize using the r-star estimate for the US of Lubik-Matthes.

Additional estimation details

Moment	Data	ϕ_π=2, ψ_b=1	ϕ_π=1.25, ψ_b=0.75	ϕ_π=0.5, ψ_b=0	ϕ_π=0, ψ_b=0
A. Convenience-yield block					
$\text{Var}(\Phi)$	0.0129	0.0765	0.0394	0.0268	0.0183
$\text{Cov}(\Phi, \Delta e)$	0.0064	-0.0129	0.0036	0.0048	0.0011
$\text{Cov}(\Phi, \Delta b_F)$	0.0038	-0.0077	0.0019	0.0029	0.0005
B. Debt maturity and fiscal-valuation block					
$\text{Var}(R_L)$	0.0948	0.0211	0.0225	0.0145	0.0097
$\text{Var}(d)$	0.0256	0.1368	0.2733	0.0576	0.0567
$\text{Var}(\text{surp})$	0.0774	0.1409	0.1581	0.0728	0.0723
$\text{Cov}(\text{surp}, d)$	0.0182	0.1357	0.2007	-0.0046	-0.0049
$\text{Cov}(\text{surp}, \pi)$	0.0013	-0.0050	-0.0052	0.0008	0.0006
C. Monetary-policy and foreign-rate block					
$\text{Var}(i)$	0.0510	0.0134	0.0112	0.0041	0.0037
$\text{Cov}(i, \pi)$	0.0063	0.0066	0.0073	0.0023	-0.0001
$\text{Var}(\pi)$	0.0048	0.0035	0.0064	0.0053	0.0036
$\text{Var}(i^*)$	0.0502	0.0242	0.0229	0.0504	0.0494
$\text{ACov}_1(i^*)$	0.0491	0.0199	0.0180	0.0448	0.0439
$\text{Cov}(i^*, \Delta e)$	-0.0064	0.0075	0.0045	0.0032	0.0026
$\text{Cov}(\Phi, i^*)$	0.0005	-0.0406	-0.0243	-0.0308	-0.0272
$\text{Cov}(i, \Delta e)$	-0.0040	0.0152	0.0183	0.0062	-0.0006
D. Standard SOE macro block					
$\text{Var}(x)$	0.1169	0.2841	0.2106	0.1244	0.1234
$\text{ACov}_1(x)$	0.0903	0.1830	0.1618	0.0703	0.0653
$\text{Var}(c)$	0.5272	0.0710	0.0322	0.0313	0.0432
$\text{Cov}(c, x)$	0.1872	0.0255	0.0224	0.0077	0.0080
$\text{Var}(\Delta e)$	0.1665	0.0215	0.0427	0.0397	0.0262
$\text{Var}(\Delta b_F)$	0.1472	0.0091	0.0179	0.0168	0.0110
$\text{Cov}(\pi, x)$	0.0053	0.0098	0.0108	0.0047	0.0050

Table C.2: Matched Moments: Data vs Model

Additional Impulse-Responses

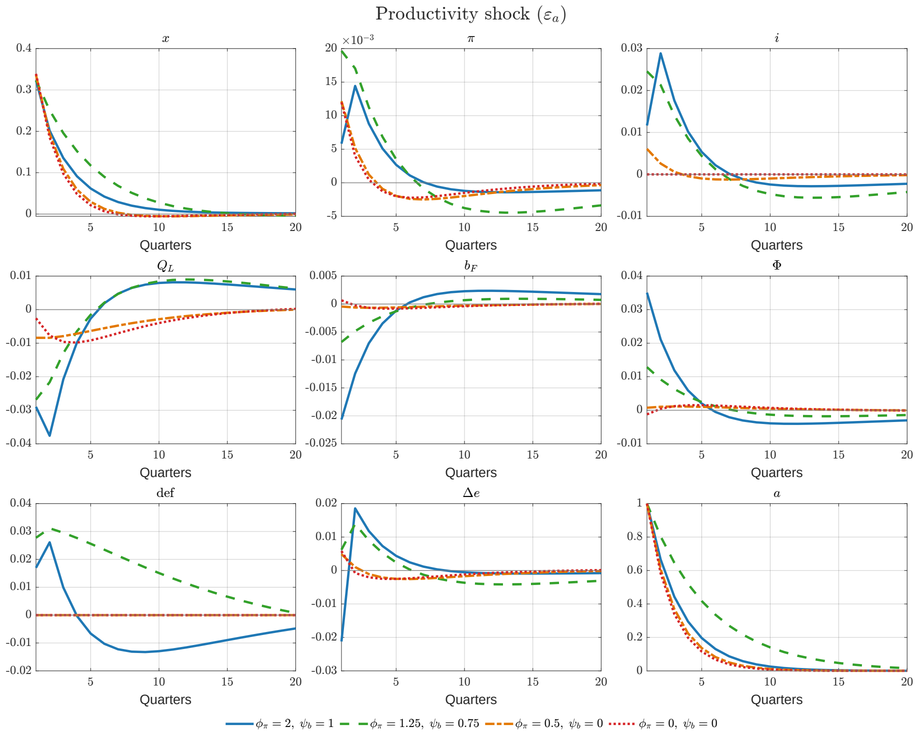

Figure C.3: IRFs to a productivity shock a_t.

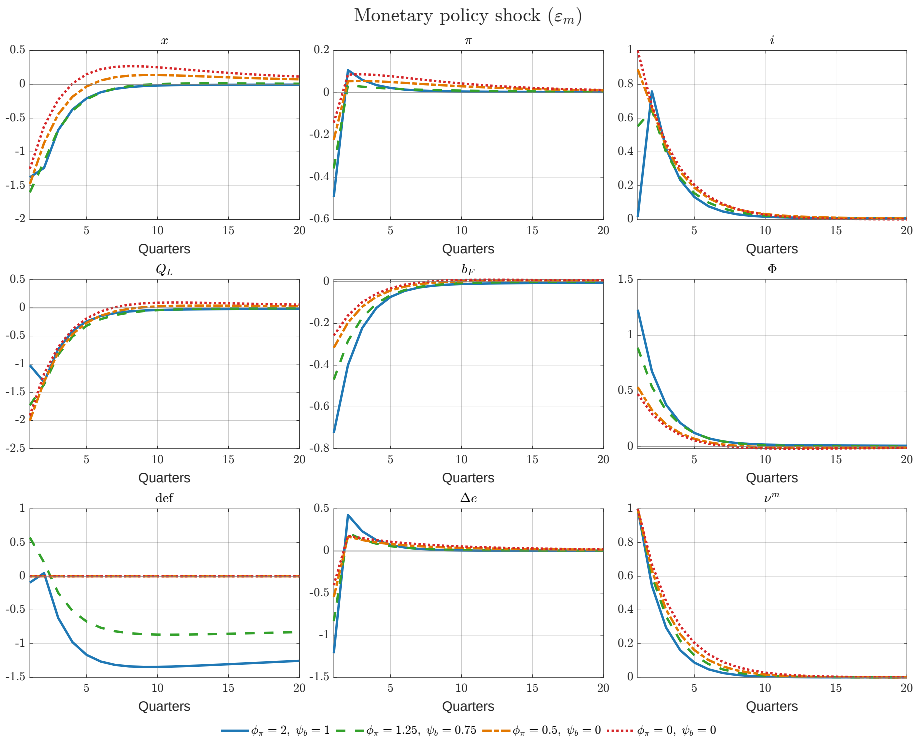

Figure C.4: IRFs to a monetary policy shock ε_t^m.

Chapter D

Appendix to: Towards a Bullwhip Theory of Supply Chains

D.1 Derivations Related to the Model

This appendix collects the full derivations that underpin Section 4.4. It is organised to mirror the structure of the main text: manufacturer first, then wholesaler, then retailer.

In the following, we make use of the structure of the production function to further simplify the problem by postulating that the sales function today *only* depends directly on each firm j's own productivity a_{jt}, but not on the other elements of the exogenous shock vector. As those elements do not enter contemporary production decisions, we can discard them in the context of the revenue function below, since they are orthogonal to the element of the $min(\cdot)$ function that has been chosen with the exception of the direct effect on the choice variables of the manufacturer. Consequently, the integral only needs to be taken over the cdf of the technological shock a_{jt}.

Assumption 3 *Within a given period, the exogenous innovations of firm-specific productivity that are not part of the firm's own production environment are not part of a given firm's information set. Consequently, firms do not internalize those variables, consider them to be unknown states, and form expectations over the distribution of their own firm-specific productivity variable, which manifests itself after its optimal decisions are made. Non-firm specific shocks, however, are known at the beginning of a given period.*

The Manufacturer: Revenue Function and First-Order Conditions

The constraint on fulfilled orders at the manufacturer stage is:

$$O_t^F = \min\left\{\max\{0, H_{3t}\} + a_{3t}n_{3t}^{\gamma}, O_t^R + \max\{0, -H_{3t}\}\right\}. \tag{D.1}$$

Integrating over a_{3t} and using the simplification $\max\{0, -H_{3t}\} - \max\{0, H_{3t}\} = -H_{3t}$, the cutoff below which supply is the binding constraint is:

$$\bar{a}_{3t} = \frac{O_t^R - H_{3t}}{n_{3t}^{\gamma}}. \tag{D.2}$$

The expected revenue function is therefore:

$$\begin{aligned} R^3(H_{3t}, n_{3t}) = &\max\{0, H_{3t}\}\, F^{a_3}[\bar{a}_{3t}] + n_{3t}^{\gamma} \Psi^{a_3}[\bar{a}_{3t}] \\ &+ \Big[O_t^R + \max\{0, -H_{3t}\}\Big]\Big(1 - F^{a_3}[\bar{a}_{3t}]\Big), \end{aligned} \tag{D.3}$$

where F^{a_3} is the marginal cdf of a_3 and $\Psi^{a_3}(\cdot)$ is the partial expectation defined in equation (4.4).

The program described by equation (4.6), subject to (D.3) and the inventory law of motion, is solved by the following first-order conditions.

Labour (n_{3t})**:**

$$p_{3t}\, \gamma\, n_{3t}^{\gamma-1} \Psi^{a_3}[\bar{a}_{3t}] - W_t + \lambda_t^3\, \gamma \exp\left(\frac{\sigma_{a_3}^2}{2}\right) n_{3t}^{\gamma-1} = 0. \tag{D.4}$$

Net inventories $(H_{3,t+1})$**:**

$$\begin{aligned} -\lambda_t^3 + \beta\, \mathbb{E}_t\Big[&p_{3,t+1}\Big(\mathbb{1}_{H_{3,t+1}\geq 0}\, F^{a_3}[\bar{a}_{3,t+1}] - \mathbb{1}_{H_{3,t+1}<0}\Big(1 - F^{a_3}[\bar{a}_{3,t+1}]\Big)\Big) \\ &- C'(H_{3,t+1}) + \lambda_{t+1}^3\Big] = 0. \end{aligned} \tag{D.5}$$

The Wholesaler: Revenue Function and First-Order Conditions

Using the identity $\min\{a, b, c\} = \min\{a, \min\{b, c\}\}$ and then $\min\{x, x-z\} = x - \max\{0, z\}$, the wholesaler's revenue function simplifies to:

$$\begin{aligned} R^2(H_{2,t+1}, H_{2t}, n_{2t}, O_{t-1}^F) &= n_{2t}^{\gamma}\, \Psi^{a_2}\left(\frac{\max\{0, H_{2t}\} + O_{t-1}^F - \max\{0, H_{2,t+1}\}}{n_{2t}^{\gamma}}\right) \\ + \Big(\max\{0, H_{2t}\} + O_{t-1}^F - \max\{0, H_{2,t+1}\}\Big) &\left[1 - F^{a_2}\left(\frac{\max\{0, H_{2t}\} + O_{t-1}^F - \max\{0, H_{2,t+1}\}}{n_{2t}^{\gamma}}\right)\right]. \end{aligned} \tag{D.6}$$

The wholesaler's first-order conditions are given by:

Labour (n_{2t})**:**

$$p_{2t}\, \gamma\, n_{2t}^{\gamma-1} \Psi^{a_2}(\bar{a}_{2t}) - W_t = 0, \tag{D.7}$$

where $\bar{a}_{2t} \equiv (\max\{0, H_{2t}\} + O_{t-1}^F - \max\{0, H_{2,t+1}\})/n_{2t}^{\gamma}$.

Net inventories ($H_{2,t+1}$):

$$-p_{2t}\,\mathbb{1}_{H_{2,t+1}\geq 0}\left[1-F^{a_2}(\bar{a}_{2t})\right]-\lambda_t^2+\beta\,\mathbb{E}_t\,\partial_{H_2}\mathbb{V}_{t+1}^2(H_{2,t+1},O_t^R,a_{2,t+1})=0. \tag{D.8}$$

Orders (O_t^R):

$$-p_{3t}\,\frac{\partial O_t^F}{\partial O_t^R}+\beta\,\mathbb{E}_t\,\partial_{O^R}\mathbb{V}_{t+1}^2(H_{2,t+1},O_t^R,a_{2,t+1})=0. \tag{D.9}$$

Envelope conditions:

$$\partial_{H_2}\mathbb{V}_t^2=p_{2t}\,\mathbb{1}_{H_{2t}\geq 0}\left[1-F^{a_2}(\bar{a}_{2t})\right]-C'(H_{2t})+\lambda_t^2. \tag{D.10}$$

$$\partial_{O^R}\mathbb{V}_t^2=\Big(p_{2t}\left[1-F^{a_2}(\bar{a}_{2t})\right]+\lambda_t^2\Big)\,\frac{\partial O_{t-1}^F}{\partial O_{t-1}^R}-p_{3t}\,\frac{\partial O_t^F}{\partial O_{t-1}^R}. \tag{D.11}$$

Order fulfilment derivatives:

$$\frac{\partial O_t^F}{\partial O_t^R}=\mathbb{1}_{O_t^R+\max\{0,-H_{3t}\}\,\leq\,\max\{0,H_{3t}\}+a_{3t}n_{3t}^{\gamma}}. \tag{D.12}$$

$$\begin{aligned}\frac{\partial O_{t+1}^F}{\partial O_t^R}=&-\mathbb{1}_{H_{3,t+1}\geq 0}\mathbb{1}_{\max\{0,H_{3,t+1}\}+a_{3,t+1}n_{3,t+1}^{\gamma}\leq O_{t+1}^R+\max\{0,-H_{3,t+1}\}}\\&+\mathbb{1}_{H_{3,t+1}<0}\mathbb{1}_{O_{t+1}^R+\max\{0,-H_{3,t+1}\}\,\leq\,\max\{0,H_{3,t+1}\}+a_{3,t+1}n_{3,t+1}^{\gamma}}.\end{aligned} \tag{D.13}$$

The Retailer: Revenue Function and First-Order Conditions

The retailer's integrated revenue function is:

$$\begin{aligned}R^1(H_{1t},n_{1t},B_t)=&\max\{0,H_{1t}\}\,F^{a_1}(\bar{a}_{1t})+n_{1t}^{\gamma}(B_t+I_{t-1}^F)\,\Psi^{a_1}(\bar{a}_{1t})\\&+\Big[S_t^R+\max\{0,-H_{1t}\}\Big]\left[1-F^{a_1}(\bar{a}_{1t})\right],\end{aligned} \tag{D.14}$$

where $\bar{a}_{1t}\equiv(S_t^R-H_{1t})/\Big(n_{1t}^{\gamma}(B_t+I_{t-1}^F)\Big)$.

And the value of the retailer is hence given by

$$\begin{aligned}\mathbb{V}_t^1(H_{1t},B_t,a_{1t})=\max_{n_{1t},H_{1,t+1},B_{t+1},I_t^R}\,&p_{1t}R^1(H_{1,t+1},n_{1t},B_{t+1})-n_{1t}W_t-C(H_{1t})-C(B_t)-p_{2t}I_t^F\\&+\beta\mathbb{V}_{t+1}^1(H_{1,t+1},B_{t+1};a_{2,t+1}),\end{aligned} \tag{D.15}$$

subject to the inventory evolution constraints

$$\begin{aligned}H_{1,t+1}&=H_{1t}+a_{1t}n_{1t}^{\gamma}(B_t+I_{t-1}^F)-S_t^R,\\B_{t+1}&=(1-a_{1t}n_{1t}^{\gamma})(B_t+I_{t-1}^F),\end{aligned}$$

such that the labour demanded by the retailer is effectively determining what share of available backroom inventories are put to the front shelves for sale.

The first-order conditions with respect to n_{1t}, $H_{1,t+1}$, B_{t+1}, and I_t^R are:

$$p_{1t}\,\frac{\partial R_t^1}{\partial n_{1t}} = W_t + \left(\lambda_{2t}^1 - \lambda_{1t}^1\right)\gamma\exp\!\left(\tfrac{\sigma_{a_1}^2}{2}\right) n_{1t}^{\gamma-1}(B_t + I_{t-1}^F), \tag{D.16}$$

$$\lambda_{1t}^1 = \beta\,\mathbb{E}\,\partial_{H_1}\mathbb{V}_{t+1}^1, \tag{D.17}$$

$$\lambda_{2t}^1 = \beta\,\mathbb{E}\,\partial_{B}\mathbb{V}_{t+1}^1, \tag{D.18}$$

$$p_{2t}\,\frac{\partial I_t^F}{\partial I_t^R} = \beta\,\mathbb{E}\,\partial_{I^R}\mathbb{V}_{t+1}^1. \tag{D.19}$$

The relevant partial derivatives of R^1 are:

$$\frac{\partial R_t^1}{\partial n_{1t}} = \gamma\, n_{1t}^{\gamma-1}(B_t + I_{t-1}^F)\,\Psi^{a_1}(\bar a_{1t})\,, \tag{D.20}$$

$$\frac{\partial R_t^1}{\partial H_{1t}} = \mathbb{1}_{H_{1t}\geq 0}\,F^{a_1}(\bar a_{1t}) - \mathbb{1}_{H_{1t}<0}\left[1 - F^{a_1}(\bar a_{1t})\right], \tag{D.21}$$

$$\frac{\partial R_t^1}{\partial B_t} = \frac{\partial R_t^1}{\partial I_{t-1}^F} = n_{1t}^{\gamma}\,\Psi^{a_1}(\bar a_{1t})\,. \tag{D.22}$$

The envelope conditions with respect to H_{1t}, B_t, and I_{t-1}^R are:

$$\partial_{H_1}\mathbb{V}_t^1 = p_{1t}\,\frac{\partial R_t^1}{\partial H_{1t}} - C'(H_{1t}) + \lambda_{1t}^1, \tag{D.23}$$

$$\partial_{B}\mathbb{V}_t^1 = p_{1t}\,\frac{\partial R_t^1}{\partial B_t} - C'(B_t) + \left(\lambda_{1t}^1 - \lambda_{2t}^1\right)a_{1t}\,n_{1t}^{\gamma} + \lambda_{2t}^1, \tag{D.24}$$

$$\partial_{I^R}\mathbb{V}_t^1 = \left[p_{1t}\,\frac{\partial R_t^1}{\partial I_{t-1}^F} + \lambda_{1t}^1\,a_{1t}\,n_{1t}^{\gamma} + \lambda_{2t}^1\,(1 - a_{1t}\,n_{1t}^{\gamma})\right]\frac{\partial I_{t-1}^F}{\partial I_{t-1}^R} - p_{2t}\,\frac{\partial I_t^F}{\partial I_{t-1}^R}. \tag{D.25}$$

The order fulfilment derivatives at the wholesale stage are given by:

$$\frac{\partial I_t^F}{\partial I_t^R} = \mathbb{1}_{I_t^R \leq a_{2t}n_{2t}^{\gamma} - \max\{0, I_{t-1}^R - H_{2,t-1} + O_{t-2}^F\}}\mathbb{1}_{I_t^R \leq O_{t-1}^F + H_{2,t-1} + O_{t-2}^F - I_{t-1}^R}, \tag{D.26}$$

$$\begin{aligned}\frac{\partial I_{t+1}^F}{\partial I_t^R} = {} & \mathbb{1}_{I_t^R > H_{2t} + O_{t-1}^F}\mathbb{1}_{I_t^R \leq H_{2t} + O_{t-1}^F - I_{t+1}^R + a_{2,t+1}n_{2,t+1}^{\gamma}}\mathbb{1}_{I_t^R \leq H_{2t} + O_{t-1}^F - I_{t+1}^R + O_t^F} \\ & - \mathbb{1}_{I_t^R \leq H_{2t} + O_{t-1}^F}\mathbb{1}_{I_t^R > H_{2t} + O_{t-1}^F + O_t^F - a_{2,t+1}n_{2,t+1}^{\gamma}}\mathbb{1}_{I_t^R > H_{2t} + O_{t-1}^F + O_t^F - I_{t+1}^R}.\end{aligned} \tag{D.27}$$

Household optimality conditions

As derived in section 4.4 the household bears utility from delivered consumption goods, taking into account that an order placed today may not materialise concurrently. We only must keep track of the last period's state of order placement, as the net balance of order placement and order deliveries remains Markovian. The maximisation problem is solved by the following conditions on the choice variables:

$$\mathbb{1}_{S_t^R \leq U_t + H_{1,t-1} + U_{t-1} - S_{t-1}^R} u'(S_t^F) = \lambda_t^{HH} p_{1t} + \beta \mathbb{E}\, \partial_{S^R} \mathbb{V}_{t+1}^{HH}(S_t^R, d_{t+1}, \zeta_{t+1}) \tag{D.28}$$

$$\zeta_t v'(n_t) = \lambda_t^{HH} w_t p_{1t} \tag{D.29}$$

$$\lambda_t^{HH} = \beta \mathbb{E}\, \partial_d \mathbb{V}_{t+1}^{HH}(S_t^R, d_{t+1}, \zeta_{t+1}) \tag{D.30}$$

The corresponding envelope conditions on the state variables are given by

$$\partial_{S^R} \mathbb{V}_t^{HH}(S_{t-1}^R, d_t, \zeta_t) =$$
$$u'(S_t^F) \left[\mathbb{1}_{S_{t-1}^R \geq H_{1,t-1} + U_{t-1}} \mathbb{1}_{S_t^R + S_{t-1}^R \leq U_t + H_{1,t-1} + U_{t-1}} - \mathbb{1}_{S_{t-1}^R \leq H_{1,t-1} + U_{t-1}} \mathbb{1}_{S_t^R + S_{t-1}^R \geq U_t + H_{1,t-1} + U_{t-1}} \right]$$
$$\partial_d \mathbb{V}_t^{HH}(S_{t-1}^R, d_t, \zeta_t) = \lambda_t^{HH}(1 + i_{t-1})$$

Equilibrium definition

Definition 5 (Equilibrium with Nash-bargained goods prices) *An equilibrium with Nash-bargained goods prices consists of processes for*

- *quantities* $\{n_t, n_{1t}, n_{2t}, n_{3t}, B_t, H_{1t}, H_{2t}, H_{3t}, k_t, i_{3t}, O_t^R, O_t^F, I_t^R, I_t^F, S_t^R, S_t^F, d_t, y_t\}_{t=0}^{\infty}$,
- *(shadow) prices* $\{p_{1t}, p_{2t}, p_{3t}, w_t, \lambda_t^3, \lambda_t^2, \lambda_t^1, \lambda_t^{HH}\}_{t=0}^{\infty}$,
- *outside-option values* $\{\bar{V}_t^3, \bar{V}_t^{2,\uparrow}, \bar{V}_t^{2,\downarrow}, \bar{V}_t^{1,\uparrow}, \bar{V}_t^{1,\downarrow}, \bar{V}_t^{HH}\}_{t=0}^{\infty}$,
- *policy instruments* $\{i_t, t_t\}_{t=0}^{\infty}$,
- *and stochastic processes* $\{a_{1t}, a_{2t}, a_{3t}, \zeta_t, g_t\}_{t=0}^{\infty}$,

such that:

- *the retailer: delivers in accordance with* (4.13), *has inventories evolve according to* (4.14) *and* (4.15), *and has optimality conditions* (D.16), (D.17), (D.18), (D.19);
- *the wholesaler: delivers in accordance with* (4.9), *has inventories evolve according to* $H_{2,t+1} = H_{2t} + O_{t-1}^F - I_t^R$, *and has optimality conditions* (D.7), (D.8), (D.9);

- *the manufacturer: produces in accordance with* (4.5)*, has inventories evolve according to* $H_{3,t+1} = H_{3t} + y_t - O_t^R$*, has order delivery behaviour as specified in* (4.3)*, and it behaves optimally in accordance with* (D.4) *and* (D.5)*;*

- *the household: adheres to its budget constraint* (4.17) *and has optimality conditions* (D.28)*,* (D.29)*,* (D.30)*;*

- *the government: sets the nominal interest rate according to the Taylor rule* (4.21) *and the tax rate according to the fiscal rule* (4.22)*;*

- *price-setting:* p_{3t} *satisfies the manufacturer-wholesaler Nash sharing condition* (4.29)*,* p_{2t} *satisfies the wholesaler-retailer Nash sharing condition* (4.30)*, and* p_{1t} *satisfies the retailer-household Nash sharing condition* (4.31)*;*

- *outside options:* $\bar{V}_t^3$*,* $\bar{V}_t^{2,\uparrow}$*,* $\bar{V}_t^{2,\downarrow}$*,* $\bar{V}_t^{1,\uparrow}$*,* $\bar{V}_t^{1,\downarrow}$*, and* $\bar{V}_t^{HH}$ *are determined by the no-trade maximization problems* (4.23)*-*(4.28)*;*

- *market clearing: on labour markets by* (4.32)*;*

- *the stochastic processes* $\{a_{1t}, a_{2t}, a_{3t}, \zeta_t, g_t\}$ *are all log-normal AR(1) processes.*

D.2 Computational Appendix

The most challenging task is the recovery of policy functions that reflect the non-linearities embedded in the net inventory variables at each stage of the supply chain. Even higher-order perturbations are not sufficiently reflecting the breaks in optimal policy behaviour at the delivery constraint, ruling out perturbation methods generally. Older projection methods (such as Chebyshev polynomials), in turn, run into exploding computational costs with our relatively large number of state variables (Fernández-Villaverde et al., 2016).

Therefore, we follow the excellent survey of Fernández-Villaverde (2026) and use multi-layered neural networks to approximate the policy functions closely. These policy functions are recovered in a fully non-linear way around the steady-state, which is used as the starting point to train the network.[1]

Intuitively, the multi-layered neural network that we use improves on standard projection methods, such as the aforementioned Chebyshev polynomials, by learning the basis adaptively over time. The approximation of the true policy function does not

[1] We do not claim here to recover the *global* solution to the model, since there is no guarantee that the policy functions are unique in the dynamic equilibrium. The neural network is a global solution method, but does not ensure global uniqueness.

optimize taking the basis of the function approximation as given, but learns precisely that basis iteratively.

We briefly discuss the hyperparameters of the network. As activation functions in each layer, we use $\tanh(x)$ for smooth variables and functions, and ReLU otherwise. We use 256 hidden units (neurons) per layer. On the training side, we use Adaptive Moment Estimation (ADAM) as the optimization algorithm with a learning rate of 10^{-4}, 40 draws per Monte Carlo-chain to form expectations, and 50000 training iterations on four network layers.[2]

The vector of exogenous innovations that are not known in the same period is given by

$$\xi_t \equiv \begin{bmatrix} a_{1t} & a_{2t} & a_{3t} \end{bmatrix}.$$

The code was implemented in Python using PyTorch and takes about 10 hours to run on a laptop with an Intel Core i9-12900H processor with 32GB of RAM.

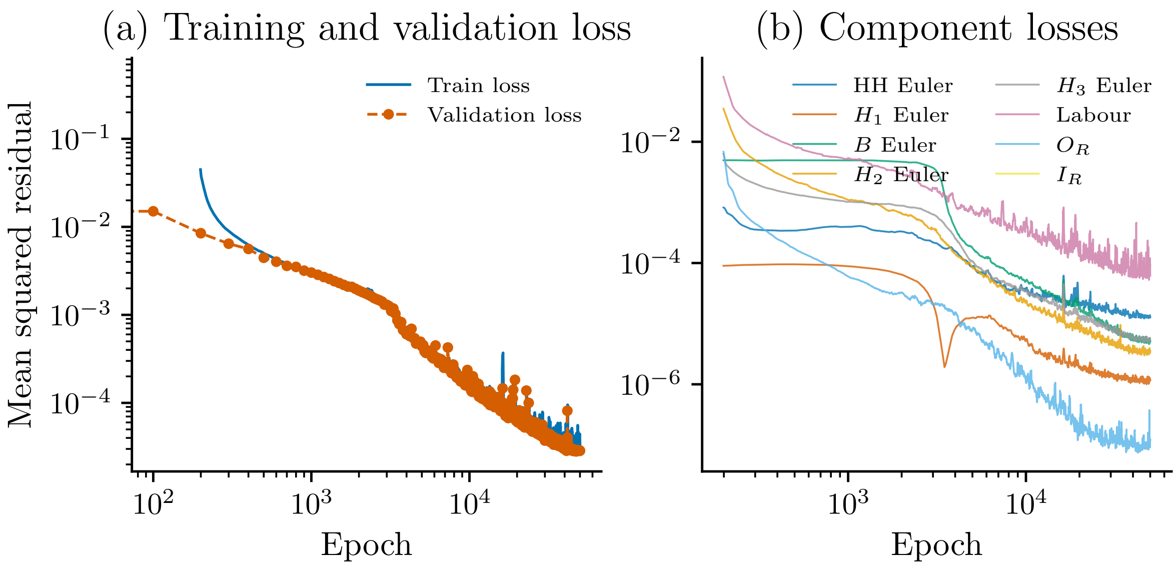

Figure D.1: Losses in the policy function approximation in the neural network throughout the estimation under the baseline policy calibration.

[2] An alternative would be the use of anisotropic Smolyak grids which can handle fairly large state-spaces whose variables are significantly positively correlated (Judd et al., 2014; Fernández-Villaverde and Levintal, 2018). The main advantage of such anisotropic Smolyak grids would be that the convergence properties of the optimization step are understood better, whereas neural networks occasionally require numerical experimentation to accurately solve a model.

D.3 Baseline Parametrization of the Model

Param.	Description	Value
Preferences		
β	Discount factor	0.99
σ	CRRA risk aversion	1
ϕ	Frisch elasticity of labour supply	0.5
ζ	Relative disutility of labour supply	0.25
Technology		
γ	Elasticity of production technology w.r.t. labour	1.0
Price bargaining		
η_1	Bargaining power retailer-household	0.5
η_2	Bargaining power wholesaler-retailer	0.5
η_2	Bargaining power manufacturer-wholesaler	0.5
Policy		
ρ_m	Interest-rate smoothing parameter	0.85
ρ_g	Persistence of log government spending/transfers	0.9
σ_g	Std. dev. of government-spending innovation	0.1
g_{ss}	Steady-state government spending (consumption units)	0.20
t_{ss}	Steady-state lump-sum tax level	0.203

Table D.1: Baseline parametrization of the economy

Regime	ϕ_π	ϕ_y	ψ_d	Interpretation
1	1.5	0.2	0.5	Active monetary policy with debt-stabilizing fiscal rule
2	1.5	0.2	0.05	Active monetary policy with weak debt response in the fiscal rule
3	0.5	0.2	0.5	Passive monetary policy with debt-stabilizing fiscal rule
4	0.5	0.2	0.05	Passive monetary policy with weak debt response in the fiscal rule

Table D.2: Policy regimes