The London School of Economics and Political Science

Essays in Financial Economics

Marcus Vinicius Fernandes Gomes de Castro

Thesis submitted to the Department of Finance of the
London School of Economics and Political Science for the degree of
Doctor of Philosophy

May 2025

Declaration

I certify that the thesis I have presented for examination for the PhD degree of the London School of Economics and Political Science is solely my own work other than where I have clearly indicated that it is the work of others (in which case the extent of any work carried out jointly by me and any other person is clearly identified in it).

Statement of co-authored work. I confirm that Chapter 2 is jointly co-authored with Carlos Carvalho and Ruy Ribeiro – I have contributed 33.3% of this work. I confirm that chapter 3 is co-authored with Carlos Carvalho – I have contributed 50% of this work.

I declare that my thesis consists of 39,068 words.

Marcus Vinicius Fernandes Gomes de Castro

Acknowledgements

Completing this PhD would not have been possible without the unwavering support and guidance I received throughout this journey. I am profoundly grateful to my advisors, Dimitri Vayanos and Ian Martin, whose insight, encouragement, and intellectual generosity have shaped this project from its earliest stages. Their work is part of what inspired me to pursue a PhD in the first place, and they remain academic role models. Their contributions to finance are timeless — I have no doubt they will continue to inspire anyone entering the PhD programme at LSE. Their mentorship has deepened not only my understanding of finance, but also the way I think about problems, communicate ideas, and approach research with clarity and purpose.

I also extend my sincere thanks to the faculty members of the Finance and Economics Departments at the LSE. Their thoughtful feedback, stimulating courses, and intellectually vibrant environment provided the foundation for my academic development. I am especially thankful for the open doors and open minds I encountered along the way.

This journey would have been far more difficult — and far less joyful — without the camaraderie and support of my fellow PhD students. I am particularly grateful to Benjamin Chen, Chris Greiner, Davide La Cara, Diogo Zuzarte de Campos, Elizaveta Smorodenkova, Nicolas Garrido Sureda, Niels Wagner, Oliver Ashtari, Otavio Bitu, Robert Rogers, Vahid Rostamkhani, Xinchen Ma, and Yang Guo. Your friendship, humour, and solidarity turned long nights, tough problem sets, and research frustrations into shared moments of progress and laughter. You made this process more human — I will always carry these memories

with me.

Above all, I am indebted to my family. To my mother, Vera de Melo Fernandes, whose unconditional love and strength have always been a source of courage. To my uncle, José Carlos de Melo Fernandes, whose support has been a lifeline through the most difficult moments. And to my late father, Marco Antonio Gomes de Castro — your absence was deeply felt throughout this journey. Losing you during my first year, just weeks before exams, was one of the most painful and intense periods of my life. This work is, in part, a tribute to your belief in me.

To all of you — mentors, friends, and family — thank you. I could not have done this without you.

Abstract

This thesis explores foundational questions at the intersection of economics and finance, combining theoretical modelling with applied econometrics to address long-standing issues in asset pricing, international macro-finance, and monetary economics.

The first chapter develops a general equilibrium model with multi-asset international financial intermediaries to study jointly uncovered interest parity (UIP), uncovered equity parity (UEP), and the hedging role of exchange rates. By allowing intermediaries to take positions in both global bond and equity markets, the model offers a unified framework where exchange rate dynamics reflect balance sheet exposures across asset classes. It clarifies the joint determination of currency risk premia and equity returns, highlighting how currencies may or may not hedge global portfolios. The model delivers novel testable implications for the behaviour of exchange rates relative to international equity and bond flows.

The second chapter focuses on financial econometrics, revisiting empirical puzzles surrounding the Euler equation and intertemporal substitution. It shows that official consumption data—typically smoothed and filtered—can severely distort estimations of the Euler equation, often yielding implausibly low or negative values for the slope. The chapter develops a flexible method to recover unfiltered consumption data, yielding more stable and economically reasonable estimations across specifications, data types, and asset-holder groups. These findings have direct implications for macro-finance models.

The third chapter addresses empirical challenges in estimating the Phillips curve, a key

component in macro and finance models with nominal dynamics. It proposes a multi-sector framework that incorporates heterogeneity in price stickiness, enabling more realistic nominal frictions. By leveraging sectoral variation that is orthogonal to monetary policy shocks and imposing cross-equation restrictions, the paper delivers robust and meaningful slope estimates. The results reconcile macro and micro evidence on price rigidity and offer insights into the transmission of monetary policy to nominal variables.

Contents

Declaration **1**

Acknowledgements **2**

Abstract **4**

1 Uncovered Interest and Equity Parities: A Unified Theory of Exchange Rates under Global Multi-Asset Intermediation **12**

1.1 Introduction 13

1.2 Model 19

1.3 Global Intermediation with Multi-Asset Exposure 28

1.3.1 UIP and UEP: A Unifying Theory 28

1.3.2 Decomposing Deviations Further 30

1.3.3 Exchange Rate Hedging Roles 32

1.4 Equilibrium Characterisation 34

1.4.1 Financial Autarky Case 34

1.4.2 Bond-Only Intermediation 37

1.4.3 Equity-Only Intermediation 45

1.4.4 Full Model – Joint Bond and Equity Intermediation 46
1.5 Empirical Evidence 50
1.5.1 Further Empirical Patterns 62
1.6 Comparative Statics: Simulating the Full Model 64
1.7 Conclusion 67

2 Intertemporal Substitution with Unfiltered Consumption **71**
2.1 Introduction 72
2.2 Model 77
2.2.1 Consumption Volatility 82
2.2.2 Adjusting Asset Returns 82
2.2.3 Calibration 83
2.3 EIS Estimates with Unfiltered Consumption Data 88
2.3.1 Homoscedastic Framework 91
2.3.2 Heteroscedastic Framework 98
2.4 EIS, Limited Participation and Unfiltered Consumption 100
2.4.1 Results 106
2.5 Conclusion 111

3 Identification of the Phillips Curve using Sectoral Data **113**
3.1 Introduction 114
3.2 Implications of Heterogeneity in Price Stickiness for Estimation of the NKPC 118
3.2.1 Optimal Policy: Ameliorating the Simultaneity Bias 118

3.2.2 Strategic Complementarities in Price Setting: More Sensible Degrees of Stickiness 124
3.3 Inflation Dynamics in a Multi-Sector Model 126
3.3.1 New-Keynesian Phillips curves 128
3.4 Empirical Approach 130
3.5 Data 132
3.6 Results 133
3.6.1 Behind the Scenes 136
3.6.2 Estimator Uncertainty 138
3.6.3 Additional Robustness Checks 143
3.7 Conclusion 145

A Appendix for Chapter 1 **147**
A.1 Households' Optimisation 147
A.2 The Financiers' Optimisation 151
A.2.1 Solution Methodology 152
A.3 Solving the Financial Autarky Model 154
A.4 Solving the Full Model 161
A.4.1 Market-Clearing Conditions for Assets 161
A.4.2 Consumption and Prices of Goods 161
A.4.3 General Representation 164
A.5 Deriving the UIP-UEP-Hedging Relationship 165
A.6 UIP-UEP-Hedging Relationship in a Lucas Orchard 168

A.7 Decomposing UIP, UEP and Hedging Roles 173

A.8 Equilibrium Characterisation: Discussion and Proofs 177

A.8.1 Equity-Only Intermediation Case: Discussion and Proofs 180

A.8.2 Full Model: Proofs 182

A.8.3 Full Model: Equilibrium Properties 184

A.9 Asymmetric Trade Preferences and the Full Model 187

A.10 Data Construction and Alternative Approaches 191

A.11 Unconditional Model, 3-Month Horizon 194

A.11.1 3-Month Horizon, Excluding the Global Financial Crisis 194

A.11.2 3-Month Horizon, Full Sample 197

A.12 Alternative Sample Windows, 12-Month Horizon 200

A.12.1 Excluding the Global Financial Crisis 200

A.12.2 2012 Onwards 207

A.13 Complementary Charts 210

A.14 Complementary Tables 213

B Appendix for Chapter 2 **217**

B.1 Data 217

B.1.1 Section 3: More Details 217

B.1.2 Section 4: More Details 218

B.1.3 A Few Notes on Data Sources involving Consumption, Frequencies and Measurement Errors 218

B.2 Fully Specified Quasi-Differenced Filter Model 220

B.2.1 Adjusting Unfiltered Consumption for Time-Aggregation Bias 228
B.2.2 Adjusting the Timing of Asset Returns 229
B.2.3 Mapping Unfiltered Quasi-Differenced Consumption Back onto Unfiltered Consumption 230
B.2.4 Consumption Volatility 231
B.2.5 Homoscedastic Counterpart and Proof of Proposition 1 231
B.2.6 Initialising the Model 233
B.3 Epstein-Zin Preferences Framework 238
B.3.1 Euler Equations 238
B.4 K-Class Estimators and Critical Values 239
B.5 Consumption: Nondurables Only 241
B.6 Implied EIS from Estimates of Its Reciprocal 246
B.7 Results Using Raw Returns for Both Reported and Unfiltered Consumption . 249
B.8 Results with Restricted Sample 254
B.9 Results for Annual Data when Measurement Errors are Persistent 254

C Appendix for Chapter 3 **262**

C.1 The Heterogeneous Economy 262
C.1.1 Fully Specified Model 262
C.1.2 Optimal Policy: Derivation 266
C.2 Additional Details on the Dataset 268
C.3 Results Perturbing the Instrument Set 269
C.4 Estimated Sectoral Infrequencies Under the Alternative Calibrations 276

C.5 Removing the Aggregate NKPC from the System 278

C.6 Perturbing Starting Values 280

C.7 Expanded Sample and Alternative Calibrations 282

Chapter 1

Uncovered Interest and Equity Parities: A Unified Theory of Exchange Rates under Global Multi-Asset Intermediation

Marcus Vinicius Fernandes Gomes de Castro[1]

Abstract: *This paper develops a unified theory of exchange rate determination that links uncovered interest parity (UIP) and uncovered equity parity (UEP) through the lens of global financial intermediation. Existing frameworks typically analyse bond and equity flows in isolation, overlooking the endogenous interaction between interest rates, stock returns, and exchange rates. I show that international intermediaries with heterogeneous exposures across asset classes and currencies play a central role in shaping deviations from both UIP and UEP. The model delivers a novel decomposition of UEP deviations into three intuitive components: (i) the UIP deviation, (ii) a hedging motive term reflecting the exchange rate's ability to hedge equity risk, and (iii) a volatility asymmetry term driven by cross-country differences in stock market risk. This relationship is not a model-specific artefact,*

[1] I am especially grateful to my advisor, Dimitri Vayanos, whose guidance was essential in the development of this paper. I also thank Dmitry Mukhin, Hélène Rey, Ian Martin, Liliana Varela, and Walker Ray for their valuable feedback and support.

but rather a general property that also emerges in distinct theoretical environments. Conditional on a global intermediation portfolio, I derive closed-form asset prices, hedging covariances, and exchange rate dynamics in a tractable general equilibrium, and show that equity volatility, international equity co-movement, and intermediaries' balance sheet composition jointly determine whether a currency exhibits a hedging role. Empirically, I test the model's predictions both cross-sectionally and in the time series, using monthly data for 15 currencies. The results confirm that UIP and UEP deviations are tightly connected, FX hedging roles are stronger for currencies with safer and less correlated equity markets with respect to the U.S., and carry-trade funding currencies tend to provide better downside protection endogenously. This paper clarifies the multi-asset transmission mechanisms that shape global currency valuations and offers new foundations for understanding capital flow asymmetries in financially integrated markets.

1.1 Introduction

Deviations from the uncovered interest parity (UIP) have long puzzled economists, with a large empirical literature documenting persistent and systematic return differentials on risk-free bonds across countries — see, for example, Fama (1984), Vermeulen et al. (2007) and Brunnermeier, Nagel, and Pedersen (2009). In contrast, the uncovered equity parity (UEP) — a theoretical negative correlation between a country's exchange rate and the relative performance of its equity market, whereby the home currency tends to *depreciate* when domestic equities *outperform* foreign ones, first formalised by Hau and Hélène Rey (2006) — remains less well understood, despite its structural similarity to the UIP. On the one hand, existing theories of the UIP typically centre on bond markets, abstracting from equity flows and multi-asset exposure — Gabaix and Maggiori (2015) being a prominent example. On the other hand, UEP theories often assume trivial interest rates and incomplete hedging, limiting their ability to speak to the full cross-asset implications of exchange rate movements. A general framework that jointly accounts for bond and equity flows, and for the risk and hedging channels that link them to exchange rates, has remained out of reach.

This paper develops a unified theory of exchange rate determination grounded in global multi-asset intermediation. I introduce a tractable general equilibrium model in which international financiers take joint positions in bonds and equities across countries. In this environment, deviations from both UIP and UEP — along with the hedging properties of exchange rates — emerge endogenously from optimal portfolio decisions. A central result is the *UIP–UEP–Hedging relationship*, which shows that UEP deviations can be decomposed into three intuitive components: (i) a UIP deviation, (ii) a hedging term reflecting how exchange rates covary with local equity returns, and (iii) a volatility asymmetry term driven by differences in equity risk across countries. This decomposition, while novel, proves to be surprisingly general — it arises in other theoretical environments and is also strongly confirmed empirically.

In doing so, the paper brings together strands of the literature that have largely been treated in isolation. On the bond side, Gabaix and Maggiori (2015) explain UIP deviations through the lens of intermediary risk-bearing constraints. On the equity side, Hau and Hélène Rey (2006), Curcuru et al. (2014), Cenedese et al. (2015) and Camanho, Hau, and Hélène Rey (2022) highlight the role of equity returns in driving currency movements. Other interesting work such as Corte, Riddiough, and Sarno (2016) explores broader connections between global imbalances and currency risk premia. My contribution synthesises these perspectives into a single coherent framework, in which UIP and UEP are jointly determined by a common set of forces — and are shown to obey a common, but highly non-trivial, empirical structure.

It is worth emphasising that a large share of currency transactions is concentrated in the hands of a few international banks. For instance, Gabaix and Maggiori (2015) report that in 2014, the top 10 banks accounted for 80% of all FX flows, with Citigroup and Deutsche Bank alone responsible for nearly one-third. Despite this, the literature has paid limited attention to how the size and composition of these institutions' balance sheets — particularly their positions in equities — influence exchange rate dynamics. Most models abstract from the broader asset exposures of intermediaries and the implications these have for pricing currency risk. A few exceptions study exchange rates alongside other asset prices, such as Pavlova and Rigobon (2007) and Martin (2011), but their focus is on relative prices rather than capital flows. More closely related is Pavlova, Dahlquist, et al. (2022), who examine capital flows through the lens of wealth redistribution, yet under complete markets and without frictions on cross-border asset positions. In contrast, this paper develops a framework in which the global balance sheet composition of intermediaries — spanning both bond and equity markets — plays a central role in shaping currency risk premia and hedging dynamics.

The model delivers three core empirical predictions. First, UIP deviations — defined as the expected dollar return from a carry trade that is long in U.S. bonds and short in foreign bonds — should increase with the sensitivity of U.S. equity returns (in dollars) to foreign equity returns (in local currency). This reflects the relative risk exposure embedded in cross-border portfolio flows. Second, a currency is more likely to exhibit FX hedging properties — i.e., *appreciate* when its local equity market *underperforms* — if it is associated with an equity market that is both *less* volatile than the U.S. and *less* correlated with it. Third, UEP deviations — defined as the expected dollar return from a zero-investment long–short international equity strategy with U.S. equity in the long leg and foreign equity in the short leg — should increase with both the UIP deviation and the strength of the foreign currency's hedging role, with volatility asymmetries amplifying or attenuating this relationship depending on relative equity risk.

These predictions are tested empirically using data on 15 currencies over a period spanning May 2007 to October 2024. I construct measures of UIP and UEP deviations, FX hedging roles, and equity market parameters (correlation and volatility) at both 12-month and 3-month investment horizons. I then examine their relationships across countries (unconditionally) and over time within each U.S.–foreign bilateral pair (conditionally).

The evidence strongly supports the model. UIP deviations correlate positively with relative *equity* sensitivity, confirming that carry trades favour currencies associated with more reactive equity markets. FX hedging roles emerge for currencies whose equity market exhibits relatively low volatility and co-movement with the U.S. equity market. UEP deviations align with UIP and FX hedging roles in terms of sign, and a decomposition analysis confirms the joint contribution of all three model-implied channels. These results hold across econometric strategies, sample windows, and horizon lengths. In particular, cross-sectional patterns persist even when excluding the Global Financial Crisis, and panel regressions strongly confirm the relevance of fixed effects linked to volatility asymmetries.

Additionally, the model provides novel insights into the equilibrium behaviour of carry trades and currency hedging. It explains why funding currencies in carry trades tend to appreciate in bad times, while investment currencies fail to hedge, and shows how these roles are determined by the joint portfolio composition of financiers.

It is worth noting that the term "parity deviation" is used heuristically, to align with existing terminology in the literature. There is no a priori reason to expect an uncovered equity parity to hold, and even UIP — despite its empirical prominence — has no claim to validity beyond a risk-neutral benchmark, as emphasised by Kremens and Martin (2019). Throughout, I adopt these labels for clarity and comparability, but they can be interpreted more simply as expected excess returns on international bond and equity positions.

The inherent link between UIP, UEP, and FX hedging roles — though previously unrecognised — is, in hindsight, conceptually natural. Existing strands of the literature point in this direction. On the one hand, Kremens and Martin (2019) shows that UIP deviations can be understood through (risk-neutral) covariances between exchange rates and equity returns, particularly from the perspective of dollar-based investors. On the other hand, the UEP literature, as in Hau and Hélène Rey (2006), interprets currency movements through the lens of equity return differentials and rebalancing motives, but again, these mechanisms operate through return covariances between equity and FX. My framework integrates these insights:

it shows that the same hedging covariances that enter traditional UIP and UEP theories also serve as the structural conduit between them. Once global intermediaries operate across asset classes, FX risk premia must jointly reflect bond- and equity-linked exposures, and the parity deviations that arise become a function of shared hedging constraints. In this sense, the UIP–UEP–hedging relationship emerges as a natural consequence of multi-asset intermediation — an organising principle that, while new in formalisation, is rooted in mechanisms long embedded in the literature.

In sum, the paper advances a unified and testable theory of exchange rate behaviour under global intermediation, bridging longstanding gaps between the UIP, UEP, and FX hedging literatures. It offers a new perspective on how currencies relate to international capital flows, and lays the groundwork for future studies of global risk sharing, macro-financial spillovers, and the evolution of safe-haven currencies.

Related Literature. This paper connects to several strands of research in international finance and macroeconomics. On the bond side, Gabaix and Maggiori (2015) explain UIP deviations through limited arbitrage by financial intermediaries, while Lustig and Verdelhan (2007a), Menkhoff et al. (2012) and Liao and Zhang (2020) emphasise the role of risk premia and global volatility in shaping currency returns. Similar to this paper, the last also emphasises the importance of hedging motives. On the equity side, Hau and Hélène Rey (2006), Curcuru et al. (2014), Cenedese et al. (2015) and Camanho, Hau, and Hélène Rey (2022) study how equity returns drive exchange rate dynamics through rebalancing, tactical reallocations, or hedging motives. My model synthesises these views by explicitly modelling intermediaries' joint bond and equity exposures across borders, with a hedging motive term arising naturally as a result of the UEP decomposition.

Empirically, the structure of my UIP–UEP–Hedging decomposition is also related to the findings of Della Corte, Riddiough, and Sarno (2016), who document how currency premia relate to global imbalances, and to the empirical asset pricing work of Colacito and Croce

(2011), who examine exchange rate behaviour in a long-run risk framework. In terms of methodology, my approach shares similarities with the structure in Martin (2011), while the econometric and risk decomposition insights are closely related to ideas in Kremens and Martin (2019), who show that exchange rate risk premia and UIP deviations reflect (risk-neutral) covariances with equity returns. More broadly, the paper also resonates with Gourinchas and Helene Rey (2007) and Gourinchas, Helene Rey, and Truempler (2012), who emphasise the interdependence of capital flows, asset returns, and currency adjustment in global equilibrium. My contribution here lies in formalising a joint mechanism that endogenises UIP and UEP deviations as complementary outcomes of multi-asset international intermediation, offering a unified structure that clarifies the transmission of macro-financial risk across borders.

This paper also relates closely to recent work exploring how financial intermediaries transmit risk across borders through their balance sheets. Matteo Maggiori (2017) develop a general equilibrium model in which international dealers hold both bonds and equities, and face capital constraints that affect exchange rate determination. Their emphasis is on reserve currency pricing and asymmetries in global demand for safe and risky assets. My model complements theirs by highlighting how the interaction between bond and equity positions drives deviations from both UIP and UEP, and by offering a decomposition that isolates hedging covariances and volatility asymmetries as key determinants of currency pricing. Sauzet (2023) takes a different approach, studying how wealth distribution and heterogeneity in investor risk tolerance affect cross-border asset flows and exchange rates. While his framework relies on recursive preferences and consumption dynamics, my model provides a more transparent link between expected asset returns and exchange rate behaviour through the lens of intermediary portfolio choice. Both papers underscore the central role of portfolio composition in global macro-financial transmission.

The paper is organised as follows. Section 1.2 presents the model setup, with global financial intermediaries optimising over bond and equity positions. Section 1.3 introduces a theoretical breakthrough: the general "UIP–UEP–Hedging" decomposition that links deviations from parity to risk premia, hedging covariances, and volatility asymmetries. Section 1.4 studies equilibrium outcomes under financial autarky, bond-only and equity-only intermediation, and the full model. Section 1.5 tests the model's predictions, validating three empirical implications regarding UIP deviations, UEP deviations, and FX hedging roles. Section 1.6 uses simulations to explore comparative statics under the full model. Section 1.7 concludes.

1.2 Model

Time is indexed by $t = 0, 1$, and the world consists of two countries of unit mass, denoted by $j \in \{H, F\}$ for Home and Foreign. For concreteness, think of these as the United States (H) and Japan (F). In each country resides a unit mass of households, who consume, trade, and invest. Investment opportunities include local risk-free bonds and local equities, both denominated in the respective domestic currency.[2]

Households trade goods in a frictionless international market. Their standard consumption-investment decisions give rise to capital flows, which are intermediated by international financiers — potentially at a premium. Since these financiers operate across a wide range of asset classes, the premium they charge reflects not only asset-specific risk, but also the exposure of their balance sheets to broader financial conditions.

Goods

Each country produces two goods: a distinct non-tradeable good (NT) and a tradeable good. The tradeable goods are labeled H and F for the home and foreign countries, respectively. Non-tradeable goods serve as the local numéraires, and their prices are normalized to 1 in

[2] All variables in the model are real. Hence, the term *currency* refers to claims on the numéraire of each country. Similarly, *dollar-denominated* and *yen-denominated* quantities are claims to the local numéraires. The *exchange rate* should therefore be interpreted as the real exchange rate.

domestic currency.

Assets and Shocks

Each country has two financial assets: a local risk-free bond and a risky asset, referred to henceforth as equity or stock. Both bonds are in zero net supply, while households in each country are endowed with one unit of their own country's stock.

The U.S. risk-free bond delivers a gross return of R dollars in period 1 for each dollar invested at $t = 0$, and the Japanese bond pays R^* yen for each yen invested. These bonds trade at prices q_b dollars and q_b^* yen, respectively. The U.S. equity is traded at a price q_s (in dollars), while the Japanese equity is priced at q_s^* (in yen). Each equity pays a period-1 dividend denominated in local non-tradeable goods — that is, in local currency. The U.S. stock pays a dividend $Y_{NT,1}$ (in dollars), while the Japanese stock pays $Y_{NT,1}^*$ (in yen).

To facilitate the exposition of the results that follow, I assume that dividends follow a joint log-normal distribution:

$$\begin{bmatrix} \log Y_{NT,1} \\ \\ \log Y_{NT,1}^* \end{bmatrix} \sim \mathcal{N}\left(\begin{bmatrix} \mu \\ \\ \mu^* \end{bmatrix}, \begin{bmatrix} \sigma^2 & \rho\sigma\sigma^* \\ \\ \rho\sigma\sigma^* & \sigma^{*2} \end{bmatrix} \right).$$

Here, μ and μ^* denote the expected log dividends, σ and σ^* their standard deviations, and ρ the correlation between log dividends across countries. Throughout the paper, I focus on the empirically relevant case where $\rho \in [0, 1]$, reflecting the typically positive correlation between international equity returns.

Households

U.S. households derive utility from consumption according to:

$$\vartheta_0 \ln C_0 + \beta\, \mathbb{E}[\vartheta_1 \ln C_1], \tag{1.1}$$

where the consumption aggregator is:

$$C_t = \left(C_{NT,t}^{\chi_t} C_{H,t}^{a_t} C_{F,t}^{\iota_t}\right)^{1/\vartheta_t}. \tag{1.2}$$

Here, $C_{NT,t}$, $C_{H,t}$, and $C_{F,t}$ denote consumption of the non-tradeable good, the domestic tradeable, and the foreign tradeable, respectively.[3] The weights χ_t, a_t, and ι_t sum to ϑ_t, and are deterministic (though potentially time-varying).[4]

Households can trade both tradeable goods across borders in frictionless markets but must consume non-tradeables domestically. They choose consumption, bond holdings b, and equity holdings x to maximize (1.1), subject to:

$$q_s + P_{H,0}Y_{H,0} = C_{NT,0} + P_{H,0}C_{H,0} + P_{F,0}C_{F,0} + b + q_s x, \tag{1.3}$$

$$P_{H,1}Y_{H,1} + Rb + Y_{NT,1}x = C_{NT,1} + P_{H,1}C_{H,1} + P_{F,1}C_{F,1}, \tag{1.4}$$

where all values are in dollars. Households are endowed with a deterministic stream of domestic tradeables $Y_{H,0}$, $Y_{H,1}$. For tractability, I assume the initial non-tradeable endowment satisfies $Y_{NT,0} = \chi_0$, ensuring marginal utility is stable at $t = 0$ and emphasising shocks in period 1, where asset payoffs are determined.

In Appendix A.1, I show that U.S. households' optimisation yields the dollar value of

[3] Since the non-tradeable is the numéraire, its price is always 1 in local currency: $P_{NT,t} = 1$.

[4] Gabaix and Maggiori (2015) consider stochastic consumption weights. Here, stochasticity in exchange rates arises endogenously via equity risk, so no such assumption is required.

imports as:

$$
\begin{aligned}
P_{F,0}C_{F,0} &= \iota_0, \\
P_{F,1}C_{F,1} &= \frac{\iota_1}{\chi_1} Y_{NT,1}.
\end{aligned}
$$

The second expression shows that positive shocks to the U.S. stock market increase the dollar value of imports, as they affect relative prices.[5]

Japanese households solve an analogous problem, with preferences

$$
\vartheta_0^* \ln C_0^* + \beta^* \mathbb{E}[\vartheta_1^* \ln C_1^*],
$$

and consumption aggregator

$$
C_t^* = \left({C_{NT,t}^*}^{\chi_t^*} {C_{H,t}^*}^{\xi_t} {C_{F,t}^*}^{a_t^*} \right)^{1/\vartheta_t^*}, \quad \text{where} \quad \vartheta_t^* = \chi_t^* + \xi_t + a_t^*.
$$

Their optimisation (detailed in Appendix A.1) implies yen values of U.S. exports given by:

$$
\begin{aligned}
P_{H,0}^* C_{H,0}^* &= \xi_0, \\
P_{H,1}^* C_{H,1}^* &= \frac{\xi_1}{\chi_1^*} Y_{NT,1}^*.
\end{aligned}
$$

Define the exchange rate E_t as the number of dollars per yen, so a higher E_t implies a weaker dollar. Then, the dollar value of U.S. exports becomes:

$$
\text{Period 0:} \quad E_0 P_{H,0}^* C_{H,0}^* = E_0 \xi_0, \quad \text{Period 1:} \quad E_1 P_{H,1}^* C_{H,1}^* = E_1 \frac{\xi_1}{\chi_1^*} Y_{NT,1}^*.
$$

[5]Suppose a positive shock hits U.S. equity at $t = 1$. Since equity dividends are paid in non-tradeables, which become less scarce, the relative price of foreign tradeables (in dollars) rises. The ratio ι_1/χ_1 amplifies this response, reflecting the household's propensity to substitute foreign tradeables for non-tradeables.

Hence, U.S. net exports (in dollars) are:

$$NX_0 = E_0 \xi_0 - \iota_0,$$

$$NX_1 = E_1 \underbrace{\frac{\xi_1}{\chi_1^*} Y_{NT,1}^*}_{=P_{H,1}^* C_{H,1}^*} - \underbrace{\frac{\iota_1}{\chi_1} Y_{NT,1}}_{=P_{F,1} C_{F,1}}. \tag{1.5}$$

All else equal, if U.S. equity receives a positive shock while Japanese equity is hit negatively, then U.S. imports become more valuable in dollar terms, and exports less so — leading to a U.S. current account deficit in period 1.

Asset Pricing

Households price domestic assets, while international intermediaries (introduced later) are small and behave competitively. As shown in Appendix A.1, the U.S. household's optimisation yields Euler equations for bonds and equity:

$$q_b \equiv \frac{1}{R} = \mathbb{E}\left[\frac{\beta \chi_1}{Y_{NT,1}}\right], \tag{1.6}$$

$$q_s = \beta \chi_1. \tag{1.7}$$

With log utility, the price of equity is independent of its payoff, as the marginal utility of consumption fully adjusts. However, the U.S. risk premium — the expected excess return on U.S. equity over the U.S. risk-free rate — is non-zero, and given by:

$$\mathbb{E}[\mathcal{R} - R] = \frac{1}{\beta \chi_1} \left(\mathbb{E}[Y_{NT,1}] - \frac{1}{\mathbb{E}\left[\frac{1}{Y_{NT,1}}\right]} \right), \tag{1.8}$$

where $\mathcal{R} = Y_{NT,1}/q_s$ denotes the gross return on U.S. equity in dollars.

The same logic applies to Japanese households, whose optimisation implies:

$$q_b^* \equiv \frac{1}{R^*} = \mathbb{E}\left[\frac{\beta^* \chi_1^*}{Y_{NT,1}^*}\right], \tag{1.9}$$

$$q_s^* = \beta^* \chi_1^*. \tag{1.10}$$

Financial Intermediation

International financial intermediaries — henceforth, financiers — absorb imbalances in trade or portfolio flows by taking positions in bonds and equities across countries, earning a premium in the process. A unit mass of financiers operates competitively, maximising profits in dollars without initial capital. For tractability, as in Gabaix and Maggiori (2015), I assume that financiers rebate their period-1 profits to Japanese households in a lump-sum manner.[6]

At time $t = 0$, each financier allocates a fixed *dollar* amount to each asset:

θ_b : U.S. bonds, θ_s : U.S. equities,

θ_b^* : Japanese bonds, θ_s^* : Japanese equities.

These allocations satisfy the zero-investment constraint:

$$\theta_b + \theta_b^* + \theta_s + \theta_s^* = 0,$$

with negative positions corresponding to short sales.

Carry Trade in Bonds. Consider first a bond-only strategy ($\theta_s = \theta_s^* = 0$). A long position θ_b in U.S. bonds must be financed by a short position in Japanese bonds, of size $\theta_b^* = -\theta_b$. Since Japanese bonds are denominated in yen, the financier's liability is θ_b/E_0 yen, where E_t

[6]This assumption simplifies the model significantly without affecting the key equilibrium mechanisms.

is the dollar–yen exchange rate (dollars per yen). The expected dollar return in period 1 is:

$$\mathbb{E}\left[R - R^* \frac{E_1}{E_0}\right] \theta_b \equiv \Omega^{\text{UIP}} \theta_b, \tag{1.11}$$

where Ω^{UIP} captures the deviation from uncovered interest parity (UIP) and represents the financier's expected carry-trade return per dollar invested.

Carry Trade in Equities. Now consider a stock-only strategy ($\theta_b = \theta_b^* = 0$). A long position θ_s/q_s in U.S. equities is offset by a short position $\theta_s^* = -\theta_s$, corresponding to $\theta_s/(E_0 q_s^*)$ units of Japanese equities. In period 1, the financier earns $Y_{NT,1}$ dollars per U.S. stock and owes $Y_{NT,1}^* E_1$ dollars per Japanese stock. Hence, the expected dollar return is:

$$\mathbb{E}\left[\mathcal{R} - \mathcal{R}^* \frac{E_1}{E_0}\right] \theta_s \equiv \Omega^{\text{UEP}} \theta_s, \tag{1.12}$$

where Ω^{UEP} captures the deviation from the uncovered equity parity (UEP) — an "equity carry trade" analogue to the UIP — and $\mathcal{R}^* = Y_{NT,1}^*/q_s^*$ denotes the gross return on Japanese equity in yen.

If these strategies are conducted independently (i.e., $\theta_b = -\theta_b^*, \theta_s = -\theta_s^*$), the financier's expected profit is:

$$\Omega^{\text{UIP}} \theta_b + \Omega^{\text{UEP}} \theta_s.$$

Their ability to profit hinges on equilibrium deviations from parity conditions.

On the Term "UEP Deviation". The term "UEP deviation" is used heuristically. Hau and Hélène Rey (2006) define the UEP as a negative correlation between foreign stock excess returns (over home stocks) and foreign exchange rate appreciation. That correlation remains central to this paper, though the microfoundation differs. Whereas Hau and Hélène Rey (2006) rely on incomplete hedging and portfolio rebalancing effects, here the hedging role of exchange rates arises from financiers' tactical equity allocations, consistent with more recent evidence documented in Curcuru et al. (2014).

Financiers' Problem. More generally, financiers treat bond and equity positions jointly. Their optimisation problem — detailed in Appendix A.2 — is:

$$\max_{\boldsymbol{\theta}} \mathbb{E}[\Pi_1] - \frac{a}{2}\mathrm{Var}[\Pi_1], \tag{1.13}$$

subject to the zero-investment constraint $\theta_b + \theta_b^* + \theta_s + \theta_s^* = 0$, where $\boldsymbol{\theta} = [\theta_b, \theta_b^*, \theta_s, \theta_s^*]^T$ and Π_1 denotes the dollar value of the financier's portfolio at time $t = 1$. Since they begin with no initial wealth, $\Pi_0 = 0$. Aggregating over the unit mass of financiers yields total demand functions $\boldsymbol{\Theta} = [\Theta_b, \Theta_b^*, \Theta_s, \Theta_s^*]^T$.

Exchange Rates

In period 0, U.S. households export tradeables worth ξ_0 yen, which they convert into dollars — creating a demand for $E_0\xi_0$ dollars. Conversely, Japanese households receive ι_0 dollars from importing U.S. tradeables and convert this into yen. The net trade balance implies a net period-0 demand for dollars (against yen) of $E_0\xi_0 - \iota_0$.

Including financiers' asset purchases, the period-0 dollar market-clearing condition becomes:[7]

$$\underbrace{E_0\xi_0 - \iota_0}_{\text{Current account}} + \underbrace{\Theta_b + \Theta_s}_{\text{Capital account}} = 0. \tag{1.14}$$

That is, a U.S. current account surplus (deficit) must be offset by a capital account deficit (surplus), implemented through financiers' net purchases of U.S. assets. The composition of this adjustment — between bonds (Θ_b) and equities (Θ_s) — depends on their portfolio preferences.

[7] Alternatively, using the zero-investment condition, the identity may be written as $E_0\xi_0 - \iota_0 - \Theta_b^* - \Theta_s^* = 0$.

In period 1, dollar demand reflects trade flows and financiers' sales of U.S. assets. The market-clearing condition is:

$$E_1 \frac{\xi_1}{\chi_1^*} Y_{NT,1}^* - \frac{\iota_1}{\chi_1} Y_{NT,1} - R\Theta_b - \mathcal{R}\Theta_s = 0. \tag{1.15}$$

The first two terms represent U.S. net exports, while the last two terms capture financiers converting the proceeds from their U.S. bond and equity holdings into yen.[8]

Illustration: Financial Autarky. In the absence of financiers ($\boldsymbol{\Theta} = \mathbf{0}$), trade must balance period-by-period. Then:

$$E_1 = \frac{\iota_1}{\xi_1} \cdot \frac{\chi_1^*}{\chi_1} \cdot \frac{Y_{NT,1}}{Y_{NT,1}^*} \propto \frac{Y_{NT,1}}{Y_{NT,1}^*}. \tag{1.16}$$

A relatively strong U.S. stock market (higher $Y_{NT,1}$) leads to dollar depreciation. With no financial intermediation, households bear all equity risk, so their income is fully exposed to stock market fluctuations. Higher U.S. income leads to a greater dollar value of imports from Japan, both through relative prices and income effects.[9] This effect is amplified when ι_1 is large and χ_1 is small, reflecting a strong substitution towards foreign tradeables. Similarly, lower ξ_1 and higher χ_1^* reduce Japanese imports of U.S. goods, reinforcing the depreciation.

Connection to the UEP. Equation (1.16) provides a transparent illustration of the UEP relationship documented by Hau and Hélène Rey (2006): a relatively stronger equity market corresponds to a weaker currency. While derived under autarky, this inverse relationship between equity performance and exchange rate remains a key force in the model — even when financiers actively shape capital flows and parity conditions become endogenous. I

[8]Since financiers rebate profits to Japanese households, they must liquidate their U.S. assets in dollars and convert the proceeds. Bonds yield $R\Theta_b$ dollars; equities yield $\mathcal{R}\Theta_s$, where $\mathcal{R} = Y_{NT,1}/q_s$.

[9]With no financiers, households hold all domestic equity. Stock market shocks directly affect their income and, under homothetic preferences, their consumption. A positive U.S. stock shock raises the dollar value of imports from Japan through both higher relative prices (as foreign tradeables become relatively scarcer) and higher consumption (as U.S. households become wealthier).

return to this broader result in the analysis that follows.

1.3 Global Intermediation with Multi-Asset Exposure

1.3.1 UIP and UEP: A Unifying Theory

The model introduced above allows UIP and UEP to be analysed jointly within a unified framework — something not possible under the formulations of Hau and Hélène Rey (2006) or Gabaix and Maggiori (2015). In what follows, I show that UEP deviations can be expressed as a linear function of UIP deviations and a covariance that captures exchange rate hedging. To simplify exposition, I adopt the following assumption:

Assumption 1. *Let* $\beta = \beta^* = \chi_t = \chi_t^* = \xi_t = 1$*, for all* t*.*

This assumption clarifies the algebra without altering the economic substance of the model.[10] Under Assumption 1, the following proposition — proved in Appendix A.5 — defines UIP and UEP deviations:

Proposition 1. *The model implies:*

$$\Omega^{\text{UIP}} = \frac{1}{\mathbb{E}\left[\frac{1}{\mathcal{R}}\right]} - \frac{1}{\mathbb{E}\left[\frac{1}{\mathcal{R}^*}\right]}\mathbb{E}\left[\frac{E_1}{E_0}\right], \tag{1.17}$$

$$\Omega^{\text{UEP}} = \mathbb{E}[\mathcal{R}] - \mathbb{E}[\mathcal{R}^*]\mathbb{E}\left[\frac{E_1}{E_0}\right] - \text{Cov}\left(\mathcal{R}^*, \frac{E_1}{E_0}\right). \tag{1.18}$$

If $\mathbb{E}[1/y] = 1/\mathbb{E}[y]$, then the first two terms in (1.17) and (1.18) would coincide, and the UEP deviation would simply equal the UIP deviation plus a hedging term. In practice, Jensen's inequality implies an additional adjustment. In Appendix A.5, I also show that up to a second-order approximation, the following relationship holds:

[10] Trade imbalances are governed by asymmetries in ι_t and ξ_t. With Assumption 1, I focus on the former. For instance, a rise in ι_1 relative to ι_0 suffices to induce dollar depreciation pressures via increased U.S. import demand. A similar assumption appears in Gabaix and Maggiori (2015), though it is more restrictive in their setting due to the absence of financial risk, which forces $R = R^* = 1$ in their paper.

Proposition 2 (UIP-UEP-Hedging Relationship)**.** *UEP deviations can be rewritten as:*

$$\Omega^{\mathrm{UEP}} = \Omega^{\mathrm{UIP}} \cdot e^{\sigma^2} - \mathrm{Cov}\left(\mathcal{R}^*, \frac{E_1}{E_0}\right) + \mathbb{E}[\mathcal{R}^*]\mathbb{E}\left[\frac{E_1}{E_0}\right](e^{\sigma^2 - \sigma^{*2}} - 1). \tag{1.19}$$

Equation (1.19) shows that UEP deviations are shaped by three components: (i) a UIP deviation, scaled by U.S. equity volatility (since financiers maximise in dollars); (ii) a hedging motive term, capturing the extent to which the foreign currency fails (or succeeds) in hedging its own equity risk; and (iii) a volatility differential term, reflecting how asymmetries in equity risk affect the pricing of cross-border equity positions. Remarkably, this structure — previously unexplored in this form — is not unique to the present model. A similar relationship emerges in the two-tree framework of Martin (2011), despite substantial differences in microfoundations and asset market frictions. I discuss these parallels — and key distinctions — in greater detail in Appendix A.6.

Why are UEP deviations linked to UIP? The connection follows directly from how exchange rates are priced (with further details to follow). If the exchange rate risk premium leads to a *negative* UIP deviation — implying that the yen is expected to appreciate (or depreciate less) than what interest rate differentials alone would suggest — then this same expectation affects the relative pricing of equities. Japanese stocks become more attractive in dollar terms, while U.S. stocks become relatively more expensive. As a result, the expected return on a long–short equity strategy that is long Japanese and short U.S. stocks increases, *lowering* Ω^{UEP}. In other words, the same exchange rate risk premium that induces a UIP deviation also pushes the UEP deviation in the same direction. This logic is reflected in the term $\Omega^{\mathrm{UIP}} \cdot e^{\sigma^2}$ in equation (1.19).

Why does a hedging motive matter? The UEP deviation measures the expected return on an international long-short equity strategy: long U.S. stocks, short Japanese. If Japanese equities tend to perform poorly when the yen depreciates, then they offer poor insurance to dollar-based investors. Formally, a positive covariance between Japanese equity returns and dollar depreciation ($\mathrm{Cov}(\mathcal{R}^*, E_1/E_0) > 0$) signals a weak yen hedging role. This increases perceived risk, requiring a higher premium for holding Japanese equities — *lowering* Ω^{UEP}. Conversely, when the foreign currency strengthens in bad times, it provides insurance, reducing required returns.

Why does cross-country volatility matter? If U.S. equities are more volatile than Japanese equities ($\sigma^2 > \sigma^{*2}$), then a long–short strategy — long U.S., short Japanese — entails greater risk. In equilibrium, this additional risk must be compensated by a higher expected return, which requires the exchange rate to adjust accordingly. The result is a larger UEP deviation, captured by the volatility differential term in equation (1.19), which is positive if $\sigma^2 > \sigma^{*2}$, and increasing in $\sigma^2 - \sigma^{*2}$.

1.3.2 Decomposing Deviations Further

Above, I showed that the expected return on an international long–short equity strategy (the UEP deviation) can be expressed as a function of expected carry-trade profitability (the UIP deviation) and a currency hedging term. I now go a step further by imposing market clearing and characterising the general equilibrium behaviour of deviations.

Specifically, both UIP and UEP deviations can be decomposed into *weighted averages of three distinct effects*: (i) the impact of bond flows; (ii) the impact of equity flows; and (iii) imbalances in trade preferences, which produce current account pressures that are partially absorbed by exchange rate movements. These three components reflect the key forces that shape exchange rate pricing in general equilibrium: financial positioning, asymmetries in asset risk, and underlying international trade (in goods). The following proposition, proved in Appendix A.7, formalises this decomposition:

Proposition 3 (UIP and UEP Deviations as Weighted Averages)**.** *UIP deviations follow, up to* *a second-order approximation:*

$$\begin{aligned}
\Omega^{\text{UIP}} \approx\ & \underbrace{\frac{\Theta_b}{\Theta_b + \Theta_s - \iota_0}\left[2e^{-\sigma^2}\,\mathbb{E}(Y_{NT,1})\right]}_{\textit{Bond Flow Effect}} \\
& + \underbrace{\frac{\Theta_s}{\Theta_b + \Theta_s - \iota_0}\,\mathbb{E}(Y_{NT,1})\left[e^{-\sigma^2} + e^{-\sigma^{*2}} + 1 - e^{\rho\sigma\sigma^* - \sigma^{*2}}\right]}_{\textit{Equity Flow Effect}} \\
& - \underbrace{\frac{\iota_0}{\Theta_b + \Theta_s - \iota_0}\,\mathbb{E}(Y_{NT,1})\left[e^{-\sigma^2} - \frac{\iota_1}{\iota_0}\left(e^{-\sigma^{*2}} + 1 - e^{\rho\sigma\sigma^* - \sigma^{*2}}\right)\right]}_{\textit{Current Account Impact}}.
\end{aligned} \tag{1.20}$$

That is, the UIP deviation reflects a weighted average of: (*i*) *the bond flow impact* (*first term*); (*ii*) *a wedge generated by equity flows and cross-country asymmetries in asset risk* (*second term*); *and* (*iii*) *imbalances in trade preferences that would otherwise drive current account surpluses or deficits* (*third term*).

Analogously, the UEP deviation satisfies:

$$\begin{aligned}
\Omega^{\text{UEP}} \approx\ & \underbrace{\frac{\Theta_b}{\Theta_b + \Theta_s - \iota_0}\,\mathbb{E}(Y_{NT,1})\left(2 + e^{-\sigma^2} - e^{\sigma^{*2} - \sigma^2}\right)}_{\textit{Bond Flow Effect}} \\
& + \underbrace{\frac{\Theta_s}{\Theta_b + \Theta_s - \iota_0} \cdot 2\,\mathbb{E}(Y_{NT,1})}_{\textit{Equity Flow Effect}} \\
& - \underbrace{\frac{\iota_0}{\Theta_b + \Theta_s - \iota_0}\,\mathbb{E}(Y_{NT,1})\left(\frac{\iota_0 - \iota_1}{\iota_0}\right)}_{\textit{Current Account Impact}}.
\end{aligned} \tag{1.21}$$

The interpretation is analogous: deviations from UEP are shaped by financial exposures through bond and equity markets, as well as underlying trade imbalances, also being a weighted average of the three effects. When trade preferences are symmetric across periods (*i.e.*, $\iota_0 = \iota_1 = 1$), *the final term disappears.*

If financiers do not hold equities ($\Theta_s = 0$), *the equity flow effect vanishes in both expressions. Similarly, if they do not hold bonds* ($\Theta_b = 0$), *the bond flow effect drops out.*

Most UIP-based models attribute deviations to carry-trade incentives and some form of constraint on risk-bearing. In this model, the latter is governed by the parameters σ, σ^*, and ρ, which control the volatility and co-movement of equity returns across countries.[11] These affect how asset values and exchange rates move together, and thus shape financiers' aggregate exposure. This logic is captured in the first term of equation (1.20).[12] The second term reflects the influence of equity flows and asymmetries in asset risk. If financial markets are symmetric (i.e., $\sigma = \sigma^*$ and $\rho = 1$), then this term vanishes. Finally, if trade preferences are also balanced ($\iota_0 = \iota_1 = 1$), the current account term disappears as well. In that case, both UIP and UEP deviations collapse to the minimal structure driven by carry-trade incentives under symmetric fundamentals.

1.3.3 Exchange Rate Hedging Roles

I now turn to how exchange rate hedging roles arise in the model. Following Hau and Hélène Rey (2006), define an *automatic hedging role of exchange rates* as a situation in which:

$$\mathrm{Cov}\left(\mathcal{R} - \mathcal{R}^*, \frac{E_1}{E_0}\right) > 0.$$

That is, when U.S. equities outperform Japanese ones, the dollar tends to depreciate — offsetting return differentials and providing insurance. This behaviour arises naturally in the model via households' goods trade, as seen in equation (1.16).

Building on this, I define two directional hedging roles. The *dollar hedging role* is said to be active when the dollar appreciates during U.S. stock market downturns:

$$\mathrm{Cov}\left(\mathcal{R}, \frac{E_1}{E_0}\right) > 0.$$

[11] Since bond pricing is non-trivial in this model, it will also depend on those parameters.

[12] Formally, σ^{*2} does not appear in the first term because its effect is absorbed through exchange rate adjustments in response to foreign shocks. Similarly, $Y^*_{NT,1}$ does not enter the expression either.

Likewise, the *yen hedging role* is present when the yen appreciates in periods when Japanese equities underperform, i.e.:

$$-\text{Cov}\left(\mathcal{R}^*, \frac{E_1}{E_0}\right) > 0.$$

Note that if both the dollar and yen hedging roles hold, the automatic hedging condition follows by construction.

The following proposition, proved in Appendix A.7, provides analytical expressions for these covariances up to a second-order approximation.

Proposition 4 (Exchange Rate Hedging Covariances)**.** *The covariance determining the automatic hedging role satisfies:*

$$\begin{aligned}\text{Cov}\left(\mathcal{R} - \mathcal{R}^*, \frac{E_1}{E_0}\right) \approx &\left(\frac{\Theta_b}{\iota_0 - \Theta_b - \Theta_s}\right) \mathbb{E}[Y_{NT,1}] \left[\frac{\mathbb{E}[Y_{NT,1}]}{\mathbb{E}[Y^*_{NT,1}]}\left(e^{-\sigma^2} - e^{\rho\sigma\sigma^* - \sigma^2}\right) + e^{\sigma^{*2} - \sigma^2} - e^{-\sigma^2}\right] \\ &+ \left(\frac{\iota_1 + \Theta_s}{\iota_0 - \Theta_b - \Theta_s}\right) \mathbb{E}[Y_{NT,1}] \left[\frac{\mathbb{E}[Y_{NT,1}]}{\mathbb{E}[Y^*_{NT,1}]}\left(e^{\sigma^{*2}} - e^{\rho\sigma\sigma^*}\right) + e^{\sigma^2} - e^{\rho\sigma\sigma^*}\right].\end{aligned} \tag{1.22}$$

The dollar hedging role is characterised by:

$$\begin{aligned}\text{Cov}\left(\mathcal{R}, \frac{E_1}{E_0}\right) \approx &\left(\frac{\Theta_b}{\iota_0 - \Theta_b - \Theta_s}\right) \frac{\mathbb{E}^2[Y_{NT,1}]}{\mathbb{E}[Y^*_{NT,1}]}\left(e^{-\sigma^2} - e^{\rho\sigma\sigma^* - \sigma^2}\right) \\ &+ \left(\frac{\iota_1 + \Theta_s}{\iota_0 - \Theta_b - \Theta_s}\right) \frac{\mathbb{E}^2[Y_{NT,1}]}{\mathbb{E}[Y^*_{NT,1}]}\left(e^{\sigma^{*2}} - e^{\rho\sigma\sigma^*}\right).\end{aligned} \tag{1.23}$$

The yen hedging role corresponds to:

$$\begin{aligned}-\text{Cov}\left(\mathcal{R}^*, \frac{E_1}{E_0}\right) \approx &\left(\frac{\Theta_b}{\iota_0 - \Theta_b - \Theta_s}\right) \mathbb{E}[Y_{NT,1}]\left(e^{\sigma^{*2} - \sigma^2} - e^{-\sigma^2}\right) \\ &+ \left(\frac{\iota_1 + \Theta_s}{\iota_0 - \Theta_b - \Theta_s}\right) \mathbb{E}[Y_{NT,1}]\left(e^{\sigma^2} - e^{\rho\sigma\sigma^*}\right).\end{aligned} \tag{1.24}$$

1.4 Equilibrium Characterisation

This section characterises equilibrium outcomes by progressively building intuition through a sequence of increasingly rich environments. I begin with the case of financial autarky, where no international asset trade occurs and exchange rates adjust purely through goods trade. I then introduce international financiers in stages: first allowing them to trade only bonds, then only equities, and finally both. Emphasis is placed on the bond-only case, which already captures key mechanisms driving UIP deviations and exchange rate adjustment, and is analytically self-contained. The equity-only and full-model cases are addressed more briefly here, with full derivations and additional discussion deferred to Appendix A.8.

1.4.1 Financial Autarky Case

I begin by considering the case of financial autarky, in which international asset markets are shut down: $\Theta_b = \Theta_b^* = \Theta_s = \Theta_s^* = 0$. In this environment, households cannot save or invest across borders, and all trade imbalances must be absorbed by exchange rate movements. The following result characterises UIP and UEP deviations in this setting:

Proposition 5. *Under financial autarky, UIP and UEP deviations are:*

$$\begin{aligned}
\Omega^{\text{UIP}} &\approx \mathbb{E}(Y_{NT,1}) \left[\underbrace{e^{-\sigma^2} - \frac{\iota_1}{\iota_0} e^{-\sigma^{*2}}}_{\text{Interest Rate Differential}} + \underbrace{\frac{\iota_1}{\iota_0} \frac{e^{\rho\sigma\sigma^*}}{e^{\sigma^{*2}}} - 1}_{\text{FX Adjustment}} \right], \\
\Omega^{\text{UEP}} &\approx \mathbb{E}(Y_{NT,1}) \left(\frac{\iota_0 - \iota_1}{\iota_0} \right).
\end{aligned} \tag{1.25}$$

Corollary 1. *If trade preferences are symmetric* ($\iota_0 = \iota_1 = 1$) *and financial markets are symmetric* ($\rho = 1$, $\sigma = \sigma^*$)*, then:*

$$\Omega^{\text{UIP}} = \Omega^{\text{UEP}} = 0.$$

More generally, if trade preferences are symmetric but financial markets are not, then $\Omega^{\text{UEP}} = 0$ *but*

$\Omega^{\mathrm{UIP}} \neq 0$.

Interpretation. When households are fully exposed to domestic equity shocks and cannot smooth consumption intertemporally, exchange rates adjust to balance trade flows. In this case, financial asymmetries alone do not generate UEP deviations, as any equity return differential is offset by an endogenous FX response — a mechanism akin to the automatic hedging role described by Hau and Hélène Rey (2006). While the model lacks full risk sharing, the logic mirrors Backus and G. W. Smith (1993): countries with higher relative consumption face real exchange rate depreciation. Here, that adjustment operates through goods-market clearing rather than asset markets. Asymmetric trade preferences disrupt this mechanism by creating current account imbalances, which the exchange rate must also absorb — introducing a tension that can undermine its hedging role.

The contrast with Gabaix and Maggiori (2015) is instructive: in their model, symmetric preferences are sufficient to shut down UIP deviations. In this model, however, *financial asymmetries can produce UIP deviations* even under balanced trade preferences, as they affect consumption and import behaviour through endogenous equity risk exposure.

The following proposition completes this intuition by addressing hedging roles:

Proposition 6. *Under financial autarky and symmetric financial markets* ($\rho = 1$, $\sigma = \sigma^*$)*:*

$$\mathrm{Cov}\left(\mathcal{R} - \mathcal{R}^*, \frac{E_1}{E_0}\right) = 0.$$

Mechanism. Under financial symmetry, the only source of exchange rate movements is trade. Since domestic equity returns are identical across countries, consumption and import patterns evolve identically, and the exchange rate evolves deterministically: $\mathcal{R} = \mathcal{R}^*$ and $E_1/E_0 = \iota_1/\iota_0$. In this case, exchange rates do not hedge equity exposures — there is no return differential to hedge.

To explore more general conditions, define the relative volatility of U.S. and Japanese equities as:

$$\kappa = \frac{\sigma}{\sigma^*},$$

with $\kappa < 1$ indicating that U.S. equities are *"safer"*. Using this, I characterise the conditions under which exchange rates play a hedging role:

Proposition 7. *Under financial autarky, the exchange rate always hedges the safer equity market. It also hedges the riskier market, provided its volatility is not too extreme. Formally:*

1. *The dollar hedging role is active (i.e.,* $\text{Cov}(\mathcal{R}, E_1/E_0) > 0$*) if* $\kappa < \frac{1}{\rho}$.

2. *The yen hedging role is active (i.e.,* $-\text{Cov}(\mathcal{R}^*, E_1/E_0) > 0$*) if* $\rho < \kappa$.

3. *The automatic hedging role is active (i.e.,* $\text{Cov}(\mathcal{R} - \mathcal{R}^*, E_1/E_0) > 0$*) if:*

$$\rho < \kappa < \frac{1}{\rho}.$$

Economic Intuition. In financial autarky, households hold only domestic equity and cannot smooth consumption. Their imports therefore move in lockstep with equity shocks. When a country's stock market is relatively stable, import behaviour is more predictable, reinforcing the link between equity returns and exchange rate movements — and strengthening the currency's hedging role. By contrast, highly volatile equity markets generate erratic import responses, weakening that link. This creates a tension between the two hedging roles: when international equities are positively correlated, the exchange rate cannot always hedge both markets simultaneously. The currency tied to the riskier equity market typically loses its hedging role, as import behaviour becomes less systematically linked to income shocks.[13]

Visual Summary. Figure 1.1 shows the regions in (κ, ρ) space where the dollar (left panel) and yen (right panel) exhibit a hedging role. Figure 1.2 displays the intersection where the

[13] In the limit, as $\kappa \to \infty$, U.S. equity becomes purely noise. U.S. households experience large swings in wealth, and their imports become unpredictable. Exchange rate movements increasingly reflect Japanese shocks, as Japanese import behaviour is more stable. In this case, the yen exhibits a stronger hedging role.

automatic hedging role is active.

Figure 1.1: Regions for Dollar and Yen Hedging Roles in Financial Autarky

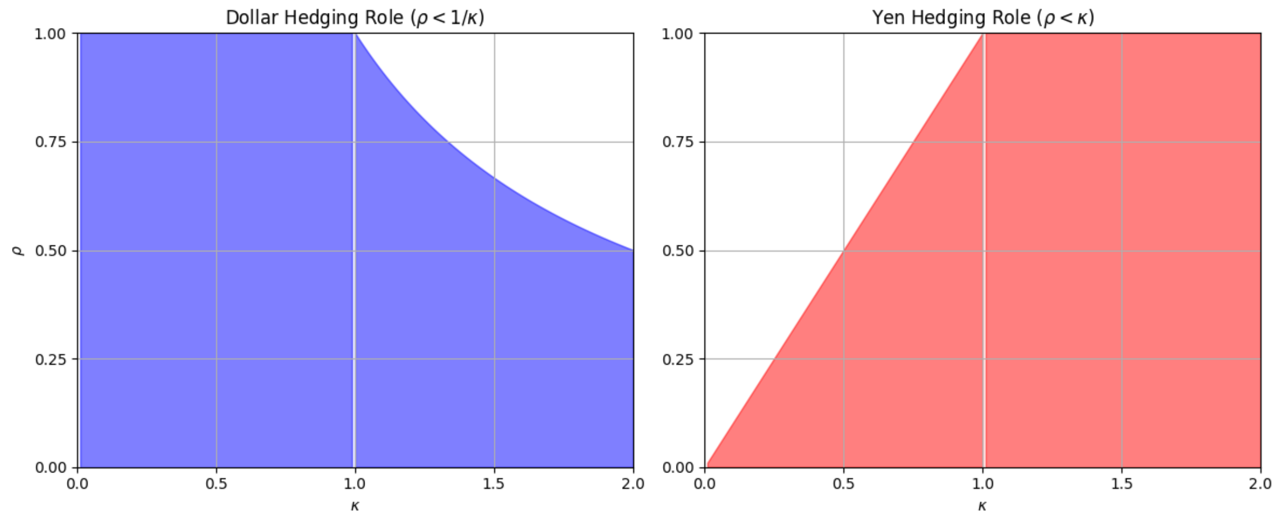

Figure 1.2: Region for Automatic Hedging Role in Financial Autarky

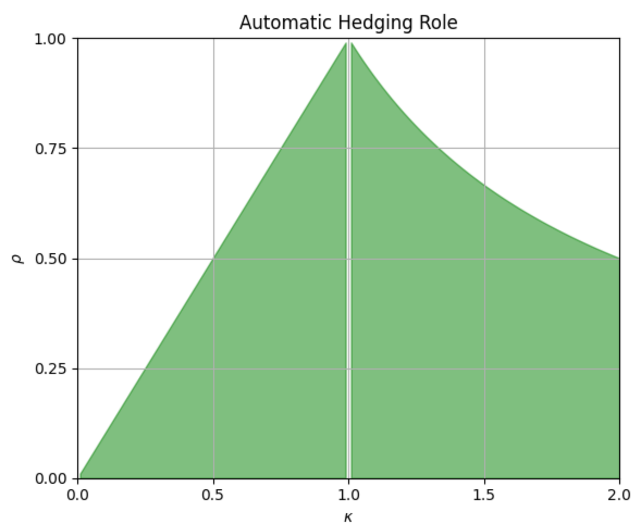

1.4.2 Bond-Only Intermediation

I now introduce international intermediation via bonds, allowing financiers to take positions in U.S. and Japanese risk-free bonds ($\Theta_b \neq 0$, $\Theta_b^* \neq 0$), but not in equity ($\Theta_s = \Theta_s^* = 0$). With zero initial wealth, the bond positions must satisfy $\Theta_b = -\Theta_b^*$. Financiers enter the market only if the expected profitability of the carry trade is positive — that is, if $\Omega^{\text{UIP}} \cdot \Theta_b > 0$. They do not target UEP deviations directly, but may influence them indirectly through their effect on UIP, as described in Section 1.3.1. The same applies to hedging roles.

The following result – whose proof is in Appendix A.8 – establishes a benchmark under symmetry:

Proposition 8 (No Carry Trade under Symmetry)**.** *Suppose trade preferences are symmetric ($\iota_0 = \iota_1 = 1$) and financial markets are symmetric ($\sigma = \sigma^*$, $\rho = 1$). If financiers are restricted to bonds ($\Theta_s = \Theta_s^* = 0$), the only equilibrium is one in which they opt out: $\Theta_b = \Theta_b^* = 0$, and $\Omega^{\text{UIP}} = \Omega^{\text{UEP}} = 0$.*

Asymmetric Markets. When financial markets are asymmetric ($\sigma \neq \sigma^*$, $0 < \rho < 1$), profitable carry trades can emerge. The following result characterises bounds on feasible bond positions:

Proposition 9 (Equilibrium Bond Positions under Asymmetry)**.** *Suppose financiers trade only bonds, and $\Theta_b = -\Theta_b^*$. Then:*

1. *If $\Theta_b > 0$ and $\Omega^{\text{UIP}} > 0$, then:*

$$0 \leq \Theta_b \leq \frac{1}{2}\left[\iota_0 - \iota_1 + \iota_1\left(e^{\rho\sigma\sigma^*} - e^{\sigma^{*2}}\right)\right].$$

2. *If $\Theta_b^* > 0$ and $\Omega^{\text{UIP}} < 0$, then:*

$$0 \leq \Theta_b^* \leq \frac{1}{2}\left[\iota_1 - \iota_0 + \iota_1\left(e^{\sigma^{*2}} - e^{\rho\sigma\sigma^*}\right)\right].$$

Which Currency Financiers Prefer. Under symmetric trade preferences ($\iota_0 = \iota_1 = 1$), the direction of the carry trade depends on which country's stock market is riskier. The following proposition addresses this — see proof in Appendix A.8.

Proposition 10 (UIP, Carry Trade Direction and Relative Equity Risk). *Assume financiers trade only bonds ($\Theta_s = \Theta_s^* = 0$), financial markets are asymmetric ($\sigma \neq \sigma^*$, $0 < \rho < 1$), and trade preferences are symmetric ($\iota_0 = \iota_1 = 1$). Then, the currency associated with the riskier equity market — defined in terms of its responsiveness to foreign equity shocks — always lies on the investment side of the carry trade.*

More specifically, let $\kappa = \sigma/\sigma^*$ and consider the OLS coefficient of a regression of U.S. log-equity returns on Japanese log-equity returns:

$$\textit{OLS coefficient} = \frac{\mathrm{Cov}(\log \mathcal{R}, \log \mathcal{R}^*)}{\mathrm{Var}(\log \mathcal{R}^*)} = \frac{\rho\sigma\sigma^*}{\sigma^{*2}} = \rho\kappa.$$

Then:

1. *If $\rho\kappa > 1$, i.e., a 1% movement in Japanese stocks is associated with more than a 1% response in U.S. stocks, financiers are long in U.S. bonds and short in Japanese bonds: $\Theta_b > 0$, $\Omega^{\mathrm{UIP}} > 0$.*

2. *If $\rho\kappa < 1$, i.e., U.S. stocks react less than one-for-one to Japanese shocks, the opposite holds: financiers are long in Japanese bonds and short in U.S. bonds, $\Theta_b < 0$, $\Omega^{\mathrm{UIP}} < 0$.*

Connection to Hedging. This result also helps clarify which currency is more likely to exhibit a hedging role. When $\rho\kappa > 1$, U.S. equity returns are relatively more sensitive to common shocks than Japanese returns, making the dollar the riskier currency in equilibrium. In this case, financiers fund their carry trades with the yen and invest in dollars. But because U.S. households are fully exposed to domestic equity and experience more volatile income, their import demand becomes erratic. The exchange rate can no longer adjust reliably to hedge U.S. equity risk — undermining the dollar's hedging role. In contrast, the yen, associated with the less volatile equity market, remains more systematically linked to household imports and income, reinforcing its hedging function. As $\rho\kappa$ falls below one, the roles reverse: the safer equity market leads to more predictable trade behaviour, enabling its currency to act as a hedge.

Yen Hedging Role

To understand when exchange rates hedge Japanese equity risk, I now characterise the conditions under which the *yen hedging role* is active. For simplicity, I henceforth assume symmetric trade preferences: $\iota_0 = \iota_1 = 1$.

From Proposition 4, the yen acts as a hedge when:

$$-\mathrm{Cov}\left(\mathcal{R}^*, \frac{E_1}{E_0}\right) > 0 \iff \Theta_b > g(\rho, \kappa, \sigma^*) \equiv \frac{e^{\kappa^2\sigma^{*2}} - e^{\rho\kappa\sigma^{*2}}}{e^{-\kappa^2\sigma^{*2}} - e^{\sigma^{*2}-\kappa^2\sigma^{*2}}}. \tag{1.26}$$

The threshold function $g(\rho, \kappa, \sigma^*)$ captures how much the financiers' bond position must tilt toward U.S. bonds (i.e., be sufficiently positive, or not too negative) in order for the yen to hedge Japanese equity risk. The numerator reflects the effect of equity *co-movements* on exchange rates[14]; the denominator reflects the effect of *relative volatility*, and does not depend on ρ. That is:

$$g(\rho, \kappa, \sigma^*) = \frac{\text{Equity Correlation Effect on FX}}{\text{Equity Relative Volatility Effect on FX}}.$$

Why large Θ_b supports a yen hedging role. When financiers are long U.S. bonds ($\Theta_b > 0$), they are short Japanese bonds ($\Theta_b^* < 0$), and must repurchase yen in period 1 to settle liabilities. This creates asymmetry in how their positions interact with trade-driven exchange rate adjustment:

1. **When Japanese equities underperform**, Japanese households cut imports, leading to yen appreciation. Financiers must buy yen into strength, reinforcing the appreciation — *amplifying the yen hedging role*.

2. **When Japanese equities outperform**, Japanese households increase imports, and the yen depreciates. Financiers again buy yen to repay liabilities, counteracting depreciation — this time *mitigating the yen hedging role*.

[14] Note that the slope coefficient of an OLS regression of $\log(\mathcal{R}^*)$ on $\log(\mathcal{R})$ equals ρ/κ. The sign of the numerator of $g(\cdot)$ is determined by $\rho - \kappa$.

The key asymmetry lies in the strength of the exchange rate response. In both cases, financiers purchase a fixed quantity of yen, but their trades have a larger impact when the yen is already appreciating (case 1) than when it is depreciating (case 2). Since yen liabilities are settled at a stronger yen in bad states for Japanese equity, the dollar cost of repayment is higher, magnifying their effect on the FX market.[15] This is why larger values of Θ_b reinforce the yen's hedging role.

Figure 1.3: Regions for the Yen Hedging Role in Model with Bond-Only Financiers

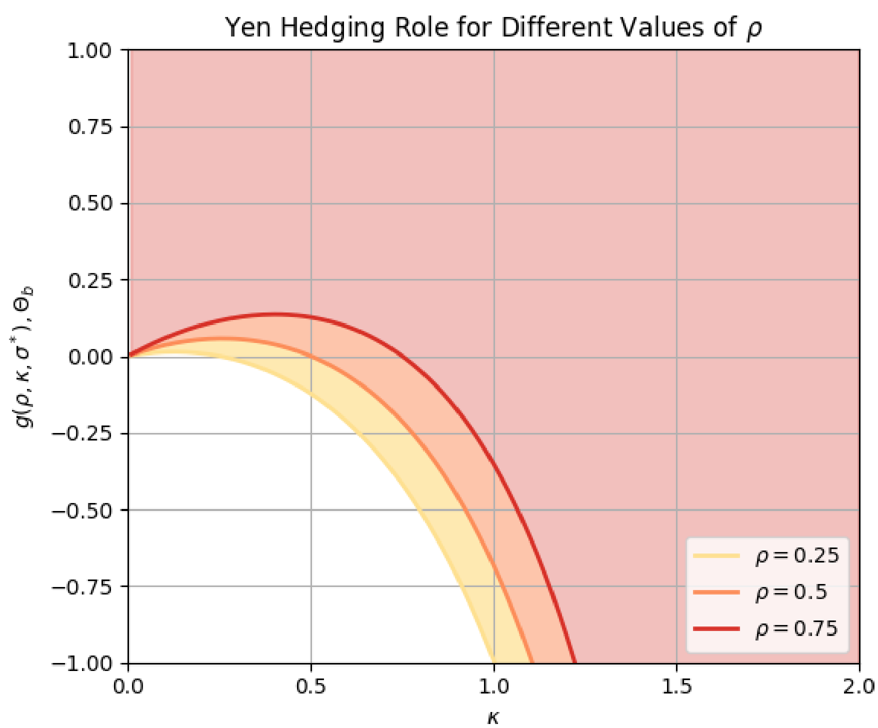

Visual Characterisation. Figure 1.3 plots the regions in (κ, Θ_b) space where the yen hedging role is active, for different values of ρ. Lower ρ expands this region, as less co-movement between equity markets allows FX to adjust more freely. Similarly, higher κ (i.e., safer Japanese equity) also supports the yen hedging role — consistent with the intuition developed under financial autarky. If Japanese equity is safer than U.S. equity ($\kappa > 1$) and correlation is not

[15] In other words, FX market clearing reflects the balance between dollar-denominated demand for yen and yen-denominated demand for dollars. When financiers must post a large dollar amount to obtain a fixed quantity of yen for settlement, the yen tends to appreciate more strongly.

too high, the yen hedging role is always active, regardless of financiers' bond positions.

Dollar Hedging Role

I now turn to the conditions under which the dollar acts as a hedge for U.S. equity. From Proposition 4, and assuming bond-only intermediation ($\Theta_s = 0$) and trade symmetry ($\iota_0 = \iota_1 = 1$), the dollar hedging role is active if and only if:

$$\mathrm{Cov}\left(\mathcal{R}, \frac{E_1}{E_0}\right) > 0 \iff \Theta_b \cdot \psi(\rho, \kappa, \sigma^*) > \phi(\rho, \kappa, \sigma^*), \tag{1.27}$$

where:

$$\begin{aligned}
\psi(\rho, \kappa, \sigma^*) &\equiv e^{-\kappa^2 \sigma^{*2}} - e^{\rho\kappa\sigma^{*2} - \sigma^{*2}}, \\
\phi(\rho, \kappa, \sigma^*) &\equiv e^{\rho\kappa\sigma^{*2}} - e^{\sigma^{*2}}, \\
f(\rho, \kappa, \sigma^*) &\equiv \frac{\phi}{\psi}, \quad \kappa(\rho) \equiv \frac{-\rho + \sqrt{\rho^2 + 4}}{2}.
\end{aligned}$$

Hedging Regions. Whether the dollar serves a hedging role depends on both the volatility ratio $\kappa = \sigma/\sigma^*$ and the financiers' bond position Θ_b. The threshold $\kappa(\rho)$ characterises the value of κ at which the sign of ψ switches from positive to negative. This divides the parameter space into three regimes:

Proposition 11 (Conditions for Dollar Hedging Role)**.** *The dollar acts as a hedge for U.S. equity if one of the following holds:*

1. ***Safe U.S. equity:*** $\kappa \leq \kappa(\rho) < 1$, *with* $f < 0$. *Here, inequality (1.27) always holds regardless of* Θ_b. *Therefore, the dollar hedging role is unconditional.*

2. ***Moderate risk:*** $\kappa(\rho) < \kappa \leq 1/\rho$, *with* $f \geq 0$. *The dollar hedging role is present if* $\Theta_b < f(\rho, \kappa, \sigma^*)$.

3. ***High U.S. equity risk:*** *$\kappa > 1/\rho > 1$, with $f < 0$. A sufficiently negative bond position is required: $\Theta_b < f(\rho, \kappa, \sigma^*) < 0$.*

Sketch of Proof (*See Appendix A.8 for Full Details*). First, determine sign changes in ψ and ϕ across intervals of κ. The root $\kappa(\rho)$ solves the inequality $\kappa^2 + \rho\kappa - 1 = 0$, which determines the sign of ψ. The threshold $\kappa = 1/\rho$ determines the sign of ϕ. Accounting for the sign-flip in the inequality when $\psi < 0$ gives the bounds for Θ_b. □

Visual Intuition. Figure 1.4 illustrates the regions where the dollar hedging role is active. The curves represent the function $f(\rho, \kappa, \sigma^*)$, while the vertical lines denote the threshold $\kappa(\rho)$, plotted for different values of ρ. As in the yen case, the dollar is more likely to serve as a hedge when U.S. equity is relatively safer. However, dollar hedging is generally more robust: for instance, when $\kappa < 1$ and $\rho \in \{0.25, 0.5\}$, the dollar hedging role holds across all feasible values of Θ_b.[16] Everything else equal, the dollar hedging role is more readily obtained in this model — reflecting the fact that financiers maximise their returns in dollars.

Figure 1.4: Regions for the Dollar Hedging Role in Model with Bond-Only Financiers

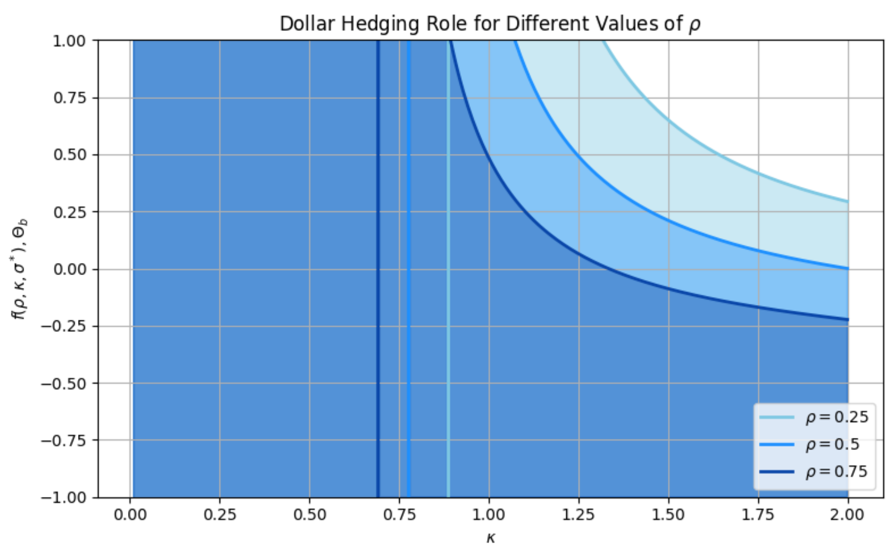

[16] In contrast, for the yen hedging role to be guaranteed under $\kappa > 1$ (i.e., when Japanese equity is safer), this is only true for $\rho = 0.25$, and not for $\rho = 0.5$; see Figure 1.3.

Why a low Θ_b supports dollar hedging. When financiers are short U.S. bonds ($\Theta_b < 0$), they hold long positions in Japanese bonds and must repurchase dollars in period 1 to settle their liabilities. If U.S. equities underperform, households reduce imports, causing the dollar to appreciate. Financiers then buy dollars into strength, reinforcing the appreciation and amplifying the dollar's hedging role. Conversely, if U.S. equities outperform, the dollar depreciates as imports rise, and financiers' dollar purchases partially offset that depreciation. Crucially, the impact is stronger in the former case: financiers face a more appreciated dollar when repurchasing it in downturns, meaning their dollar demand — denominated in yen — is higher. As in the yen case discussed above, this asymmetry makes a dollar hedging more likely when financiers' positions in the U.S. bond are sufficiently low.

UIP and Hedging Roles: General Insights. Financiers' long positions in the bonds of the riskier equity market — U.S. bonds when $\rho\kappa > 1$ ($\Omega^{\mathrm{UIP}} > 0$), or Japanese bonds when $\rho\kappa < 1$ ($\Omega^{\mathrm{UIP}} < 0$) — tend to eliminate that country's currency hedging role. This aligns with empirical findings that carry-trade investment currencies, which are typically riskier, tend to underperform in bad times and fail to appreciate during global downturns (Lustig and Verdelhan 2007b; Brunnermeier, Nagel, and Pedersen 2009; Lustig, Roussanov, and Verdelhan 2011; Menkhoff et al. 2012; Corte, Riddiough, and Sarno 2016). By contrast, when a country's equity market is relatively safe, its currency is more likely to hedge local equity risk — particularly when equity co-movement (ρ) is low. The following proposition, proved in Appendix A.8, characterises the equilibrium selected by financiers through their bond positions.

Proposition 12 (Equilibrium Selection, Carry Trade Direction and Hedging Role Asymmetry)**.** *Suppose financiers can only trade bonds, and households' trade preferences are symmetric* ($\iota_0 = \iota_1 = 1$)*. Then, there are two possibilities :*

- ***Equilibrium A:*** *If $\kappa \geq \frac{1}{\rho}$, financiers choose $\Theta_b \geq 0$ and $\Omega^{\mathrm{UIP}} \geq 0$; the dollar lies on the investment side and the yen on the funding side of the carry trade.*

- ***Equilibrium B:*** *If $\kappa \in (0, \frac{1}{\rho})$, financiers choose $\Theta_b < 0$ and $\Omega^{\mathrm{UIP}} < 0$; the yen lies on the investment side and the dollar on the funding side.*

*In both equilibria, financiers **reinforce** the hedging role of the currency used on the **funding** side of their carry trades. Their positions, however, tend to **eliminate** the hedging role of the currency used on the **investment** side. More specifically:*

- *In **Equilibrium A**, where the dollar is on the investment side, a dollar hedging role is never observed.*

- *In **Equilibrium B**, where the yen is on the investment side, a yen hedging role may or may not emerge, depending on κ and the magnitude of Θ_b. It is always eliminated when $\kappa \in (0, \rho]$, i.e., when Japanese equity is not sufficiently safe.*

1.4.3 Equity-Only Intermediation

In the case where financiers can only trade equity ($\Theta_b = \Theta_b^* = 0$), the model behaves similarly to the financial autarky case. As I show in the appendix, under symmetric trade preferences and absent bond markets, the only possible equilibrium is one in which financiers opt out when financial markets are also symmetric, and the equilibrium hedging patterns replicate those under autarky.

In particular, the presence of equity-only financiers cannot eliminate the exchange rate hedging roles observed in autarky, as their portfolio effects offset each other: while their equity trades dampen the trade balance channel, the unwinding of foreign equity positions generates FX flows that restore it. These results offer a natural transition to the full model with joint bond and equity intermediation, and clarify why the introduction of equity trading does not substantially modify the model's core mechanics.

1.4.4 Full Model – Joint Bond and Equity Intermediation

I now turn to the full model, where financiers trade both bonds and equities across countries. That is, $\Theta_b, \Theta_b^*, \Theta_s, \Theta_s^* \neq 0$. While this richer environment introduces more flexibility in portfolio construction, the exchange rate hedging properties remain largely in line with those already uncovered in the bond-only case. However, the interaction between bond and equity positions adds nuance to the mechanics — particularly regarding how stock market exposure shapes the strength and asymmetry of hedging roles.

As in previous cases, the core determinants of hedging roles are the relative volatilities of equity markets (κ) and their correlation (ρ). The main distinction now is that the bounds determining whether a hedging role emerges also depend on the financiers' equity positions. Specifically, the condition for a dollar hedging role becomes:

$$\mathrm{Cov}\left(\mathcal{R}, \frac{E_1}{E_0}\right) > 0 \iff \Theta_b \cdot \psi(\rho, \kappa, \sigma^*) > \phi(\rho, \kappa, \sigma^*) \cdot (\iota_1 + \Theta_s), \tag{1.28}$$

where ψ and ϕ are as defined in the bond-only model, and Θ_s reflects financiers' exposure to U.S. stocks. This expression generalises the bond-only condition by embedding the equity side of the balance sheet directly into the hedging threshold. The following proposition — which closely mimics Proposition 11, derived under bond-only intermediation — summarises results.

Proposition 13 (Conditions for Dollar Hedging Role in the Full Model)**.** *The dollar acts as a hedge for U.S. equity if one of the following holds:*

1. ***Safe U.S. equity:*** $\kappa \leq \kappa(\rho) < 1$. *Here,* $\psi > 0$, $\phi < 0$, *so* $f < 0$. *Inequality* (1.28) *becomes* $\Theta_b > f(\rho, \kappa, \sigma^*) \cdot (\iota_1 + \Theta_s)$, *with* $f(\rho, \kappa, \sigma^*) \cdot (\iota_1 + \Theta_s) < 0$. *The dollar hedging role is present provided the financier's position in the U.S. bond is not too negative.*

2. ***Moderate risk:*** $\kappa(\rho) < \kappa \leq 1/\rho$, *with* $f \geq 0$. *The dollar hedging role is present if* $\Theta_b < f(\rho, \kappa, \sigma^*) \cdot (\iota_1 + \Theta_s)$, *with* $f(\rho, \kappa, \sigma^*) \cdot (\iota_1 + \Theta_s) \geq 0$.

3. ***High U.S. equity risk:*** *$\kappa > 1/\rho > 1$, with $f < 0$. The dollar hedging role holds if $\Theta_b < f(\rho, \kappa, \sigma^*) \cdot (\iota_1 + \Theta_s) < 0$.*

These results use $\iota_1 + \Theta_s \geq 0$, which always holds — see proof and discussion in Appendix A.8.2.

Figure 1.5 plots the region where the dollar hedging role is active, for two values of ρ. As before, higher correlation across equity markets reduces the scope for hedging. Additionally, when U.S. equity is not too risky ($\kappa < 1/\rho$) and financiers are short in U.S. bonds ($\Theta_b < 0$), the dollar is more likely to hedge local equity shocks.

For a given κ, higher values of Θ_s expand the region where dollar hedging holds — except in the limiting case where U.S. equity is much riskier ($\kappa > 1/\rho$) and financiers are already short in U.S. bonds, in which case the effect reverses. The distinction in how financiers' positions in U.S. stocks affect the dollar's hedging role — depending on the relative riskiness of U.S. equity — arises from the following asymmetry.

Suppose the dollar hedging role is active. Then, holding long U.S. equity, financiers will seek to hedge their dollar exposure by taking a larger long position in U.S. bonds, since the dollar appreciates in downturns and bond payouts retain their value. However, this same long bond position works against the emergence of a dollar hedging role — particularly when U.S. equity is relatively risky — because it reduces the dollar demand that would otherwise arise from short positions in bad states. In other words, the more financiers hedge their equity exposure through bonds, the more they may dilute or eliminate the very exchange rate behaviour that justified the hedge. This dynamic resembles a form of *Goodhart's Law*, whereby optimising behaviour aimed at exploiting the dollar's hedging properties ultimately undermines them through endogenous general equilibrium effects.[17]

[17] This endogenous relationship is made explicit in equation (1.28), where the condition for the dollar to act as a hedge depends jointly on the financier's bond and equity positions. Suppose we are in the third region region, where $\kappa > 1/\rho > 1$ and the dollar hedging role holds with $\Theta_b < f(\rho, \kappa, \sigma^*) \cdot (\iota_1 + \Theta_s) < 0$. As Θ_s increases, the right-hand side of the inequality becomes more demanding — requiring a more negative Θ_b to preserve the hedging role. Thus, attempts to hedge dollar exposure via bond holdings can endogenously tighten or eliminate the very condition that supports dollar hedging.

Conversely, if the dollar does not initially serve as a hedge, financiers may respond by shorting U.S. bonds to offset their long equity exposure. This creates a dollar liability that must be settled in period 1, reinforcing dollar appreciation in downturns and thereby restoring the dollar's hedging role. In this way, their attempt to hedge in response to FX risk can endogenously generate the very hedging property they were seeking.

Figure 1.5: Regions for the Dollar Hedging Role Under the Full Model

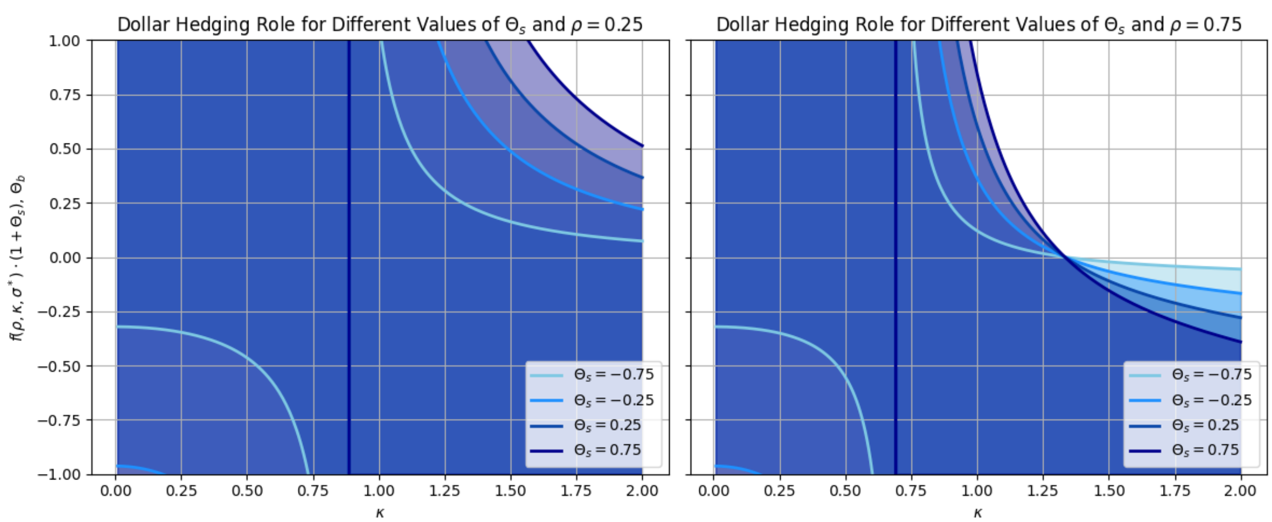

A similar condition characterises the yen hedging role:

$$-\text{Cov}\left(\mathcal{R}^*, \frac{E_1}{E_0}\right) > 0 \iff \Theta_b > g(\rho, \kappa, \sigma^*) \cdot (\iota_1 + \Theta_s). \tag{1.29}$$

Figure 1.6 shows the corresponding region. As in the dollar case, correlation across equities (ρ) compresses the hedging region. But here, larger values of Θ_s have the opposite effect: they tend to shrink the scope for yen hedging when Japanese equity is highly volatile ($\kappa < \rho$), while expanding it when Japanese equity is relatively safe ($\kappa > \rho$).

Figure 1.6: Regions for the Yen Hedging Role Under the Full Model

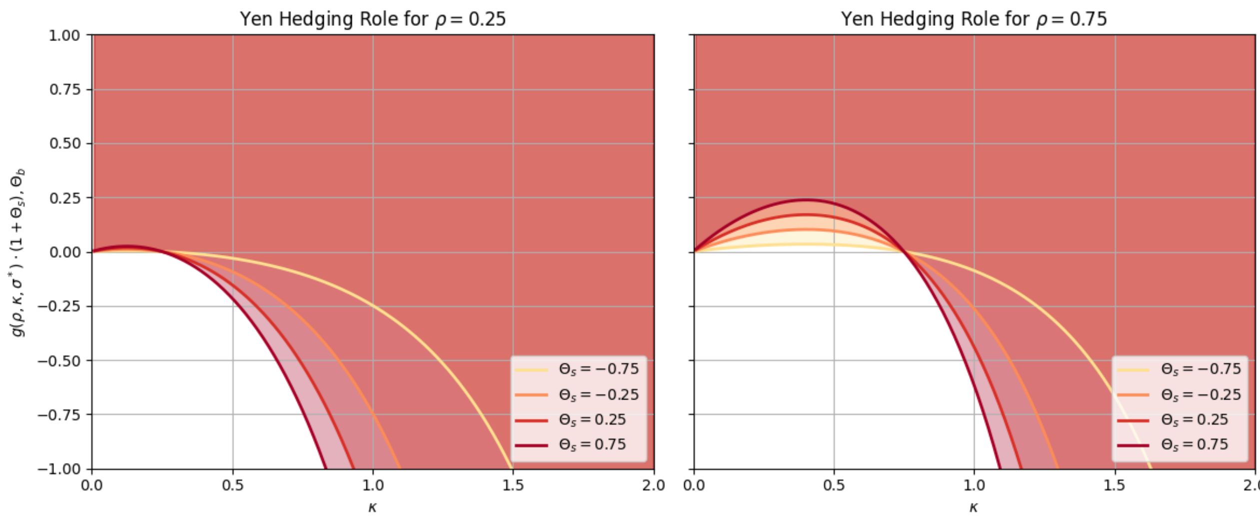

These patterns reflect a fundamental asymmetry introduced by joint bond–equity exposure. When financiers hold positive U.S. equity ($\Theta_s > 0$), their wealth is positively correlated with U.S. stock performance. In good states, they unwind their equity positions by selling dollars to acquire yen, placing downward pressure on the dollar. This tends to complement the dollar's hedging role, as long as U.S. equity is not too risky. By contrast, if they are more exposed to Japanese equity ($\Theta_s < 0$) and Japanese stocks outperform, their wealth increases in yen terms — but they must convert it into dollars. However, the model's equilibrium pricing tends to appreciate the dollar in these same states, muting the hedging effect. This asymmetry arises because financiers maximise in dollars: their balance sheet reacts more forcefully to dollar asset gains than to yen asset gains.[18]

In sum, the full model preserves much of the structure found under bond-only intermediation but introduces richer interactions between bond and equity exposures. The resulting hedging properties depend not just on the direction of trade (funding vs. investment

[18] This asymmetry is clearly visible when comparing Figure 1.5 and Figure 1.6. When stock market volatilities are relatively balanced — i.e., when $\rho < \kappa < 1/\rho$ — an increase in Θ_s expands the region for which the dollar hedging role holds, but a lower Θ_s shrinks the region for which the yen hedging role is active.

currency), but also on how the composition of assets interacts with relative volatility, co-movement, and the financier's currency of account. Further discussion is provided in Appendix A.8.3.

1.5 Empirical Evidence

This section tests the empirical relevance of the model's key mechanisms using international data on bonds, equities, and exchange rates. While the theoretical framework generates several tightly connected predictions, its core empirical implications can be organised into three testable propositions:

- **Prediction 1 (UIP and Relative Equity Sensitivity):** UIP deviations (Ω^{UIP}) should increase with the relative responsiveness of U.S. equity returns (in dollars) to foreign equity returns (in local currency), as captured by the product $\rho\kappa$. That is, we should observe $\uparrow \Omega^{\mathrm{UIP}}$ when $\uparrow \rho\kappa$.

- **Prediction 2 (FX Hedging and Market Asymmetries):** The strength of FX hedging roles should reflect the volatility and co-movement of equity markets. Hedging behaviour is more likely for currencies linked to equity markets that are less volatile than the U.S. (i.e., higher κ) and less correlated with it (lower ρ). In other words, $\uparrow -\mathrm{Cov}(\mathcal{R}^*, E_1/E_0)$ when $\uparrow \kappa$ and $\downarrow \rho$.

- **Prediction 3 (UEP as a Function of UIP and Hedging):** UEP deviations (Ω^{UEP}) should be jointly explained by UIP deviations, FX hedging strength, and volatility differentials. Specifically, Ω^{UEP} should increase in both Ω^{UIP} and the hedging term $-\mathrm{Cov}(\mathcal{R}^*, E_1/E_0)$, in line with the decomposition in equation (1.19). Moreover, currencies in the funding end of carry trades — associated with more positive Ω^{UIP} — should exhibit stronger hedging properties. Finally, the third term in the decomposition — reflecting the impact of volatility asymmetries — should be positive when U.S. equity is more volatile than the other country's, that is, when $\sigma > \sigma^*$, or equivalently, $\kappa > 1$.

More details on variable construction and data sources are provided in Appendix A.10. Throughout the main text, I focus on monthly data and evaluate relationships over a 12-month horizon. I employ two complementary empirical strategies. First, I test each model-implied relationship in the *cross section* of countries by computing moments and deviations over the full time series. This allows me to assess whether the model's comparative statics align with average cross-country patterns. Second, I explore the *time-series* dimension by constructing conditional, time-varying versions of the key variables and testing whether the model's relationships hold dynamically within each bilateral U.S.–foreign pairing. Additional results using a 3-month investment horizon are reported in Appendix A.11.

The countries and regions included in the analysis are: Australia, Brazil, Canada, China, Denmark, the Euro Area, Japan, Korea, Mexico, New Zealand, Norway, Poland, Sweden, Switzerland, the United Kingdom (U.K.), and the United States (U.S.). The baseline dataset covers the period from May 2007 to October 2024. Additional results using alternative sample windows — for instance, excluding the Global Financial Crisis — are reported in Appendix A.12.

Prediction 1: UIP and Relative Equity Sensitivity

Cross-Country Evidence on UIP Deviations. Figure 1.7 provides empirical support for Prediction 1. The left axis plots $\rho\kappa - 1$, where $\rho\kappa$ corresponds to the coefficient from an OLS regression of U.S. equity returns (in dollars) on foreign equity returns (in local currency), as derived in the model section. This captures the relative responsiveness of U.S. equity (compared to foreign equity) to shocks. The right axis shows unconditional 12-month UIP deviations (Ω^{UIP}) between the U.S. and each country in the sample. As predicted, UIP deviations increase with $\rho\kappa$: when U.S. equity reacts relatively more aggressively to shocks, a portfolio with the dollar on the investment side of carry trades delivers higher expected returns. The cross-sectional correlation between the two series is 0.59, consistent with the model's comparative statics and confirming that equity return sensitivity helps explain persistent UIP deviations across countries.

Taking the bond-only intermediation model at face value, the dollar should occupy the investment side of carry trades whenever $\rho\kappa - 1 > 0$; see Proposition 12. Appendices A.11 and A.12 show that this prediction is strongly supported by the data outside the Global Financial Crisis period, for both 3-month and 12-month investment horizons. Additionally, results are maintained when analysing κ (i.e., based on relative equity volatility) instead of $\rho\kappa$ (based on relative equity sensitivity) — see Appendix A.13.

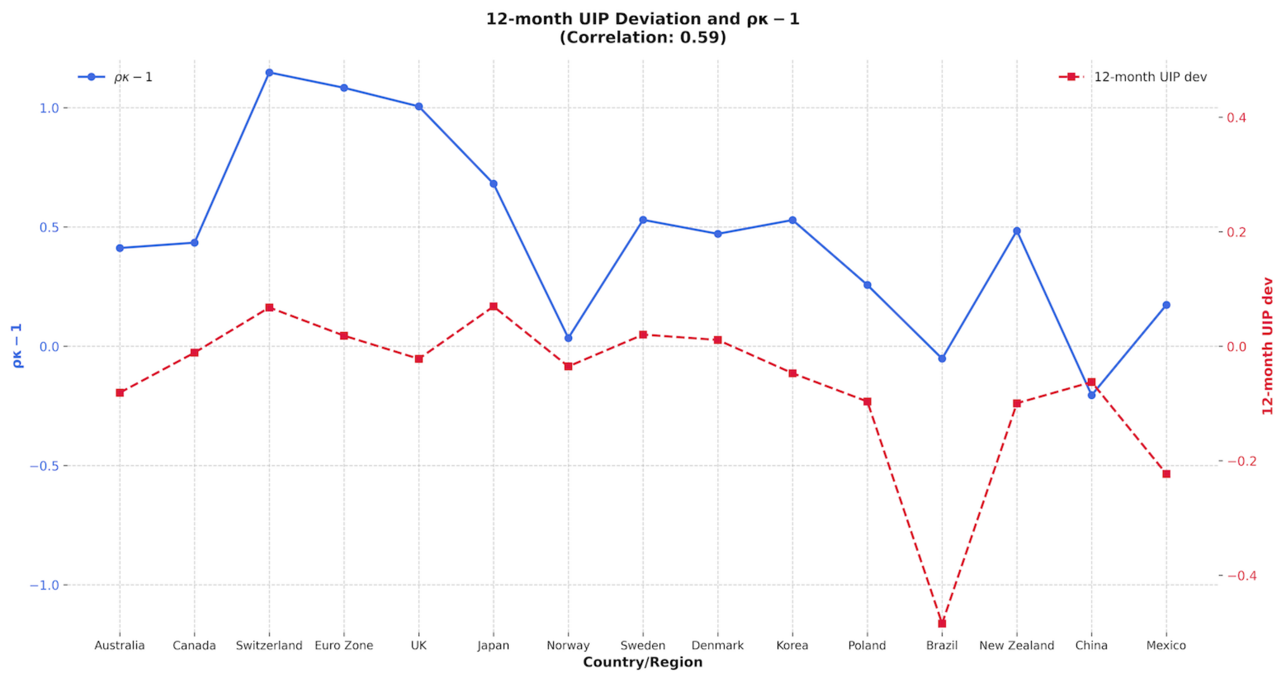

Figure 1.7: 12-month UIP deviations (Ω^{UIP}, right axis) and relative equity sensitivity ($\rho\kappa - 1$, left axis) across countries. The variable $\rho\kappa$ is constructed as the coefficient from an OLS regression of U.S. log-equity returns (in USD) on foreign log-equity returns (in local currency).

Time Series Evidence on UIP Deviations. Figure 1.8 displays the time-series behaviour of UIP deviations alongside variables ρ and κ for each currency pair relative to the U.S. dollar. Consistent with the model's predictions, the UIP deviation is positively correlated with

κ — the relative volatility of U.S. equity returns — in 11 out of 15 currency pairings. This supports Prediction 1 in the time-series dimension. Results for $\rho\kappa - 1$ are comparatively weaker, which is largely due to the absence of clear pattern in equity return correlations (ρ). Appendix A.13 presents the relevant plots and confirms this pattern. Nevertheless, when the sample is restricted to more recent years — as shown in Appendix A.12 — the positive relationship between $\rho\kappa$ and UIP deviations re-emerges more strongly. These findings suggest that volatility asymmetries play a more persistent role in shaping carry trade payoffs, while the contribution of correlation becomes less prominent during specific windows such as global crises or policy shocks.

Prediction 2: FX Hedging and Market Asymmetries

FX Hedging Role and Volatility/Co-Movement: The model implies that a currency is more likely to play a hedging role when its associated equity market is relatively less volatile and less correlated with the U.S. — that is, when κ is high and ρ is low. Figure 1.9 confirms this prediction: FX hedging roles — measured as the (negative) covariance between local equity returns and the exchange rate — correlate positively with κ, with a cross-sectional correlation of 0.64.

Figure 1.10 shows that the hedging role also correlates positively with $\rho\kappa - 1$, consistent with model-implied conditions for the equilibrium selection and the shape of hedging regions under bond-only intermediation. These patterns confirm the model's second prediction in the cross section. As predicted by the model, FX hedging roles correlate more strongly with κ than with $\kappa\rho$.

Conditional Evidence for FX Hedging. Figure 1.11 evaluates the time-series version of Prediction 2 by tracking the relationship between FX hedging strength and the underlying equity market parameters ρ and κ for each U.S.–foreign pairing. Across the 15 countries

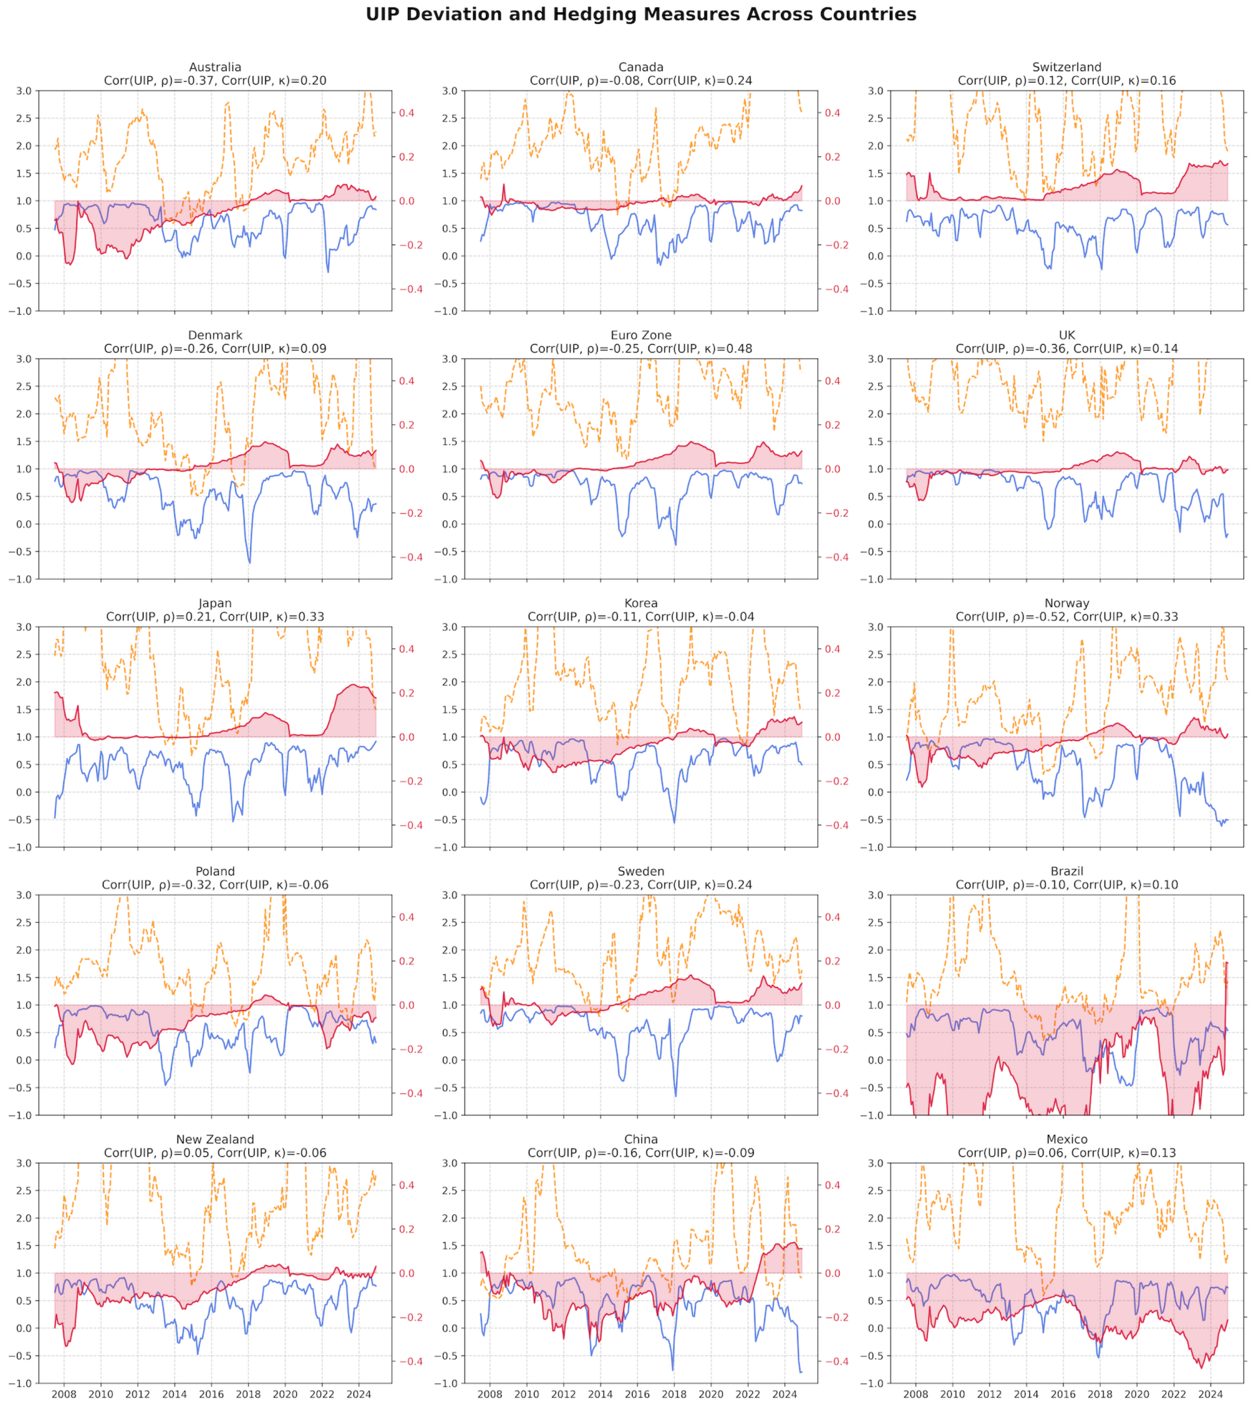

Figure 1.8: UIP Deviation and Hedging Measures Across Countries. Each panel plots the UIP deviation (red), ρ (blue), and relative U.S. equity volatility κ (orange, dashed). Reported correlations are for UIP deviations with ρ and κ, respectively.

examined, the FX hedging role (shown in red) exhibits a *negative* correlation with the equity return correlation ρ in all but one case (Japan), as predicted by the model. This indicates that

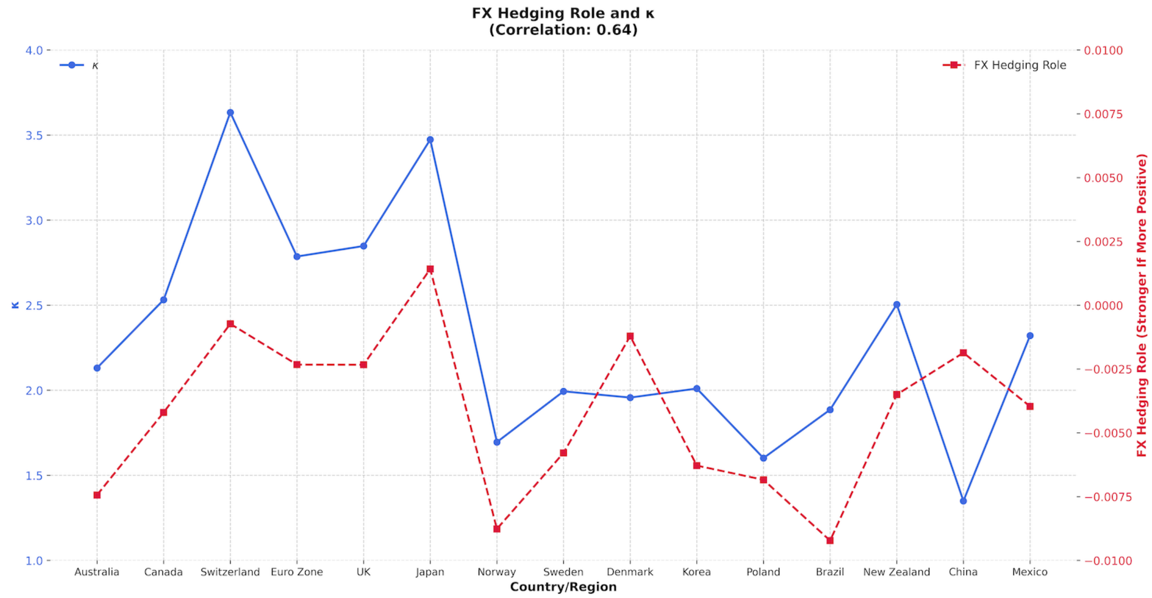

Figure 1.9: FX Hedging Role and κ, obtained for a 12-month investment horizon.

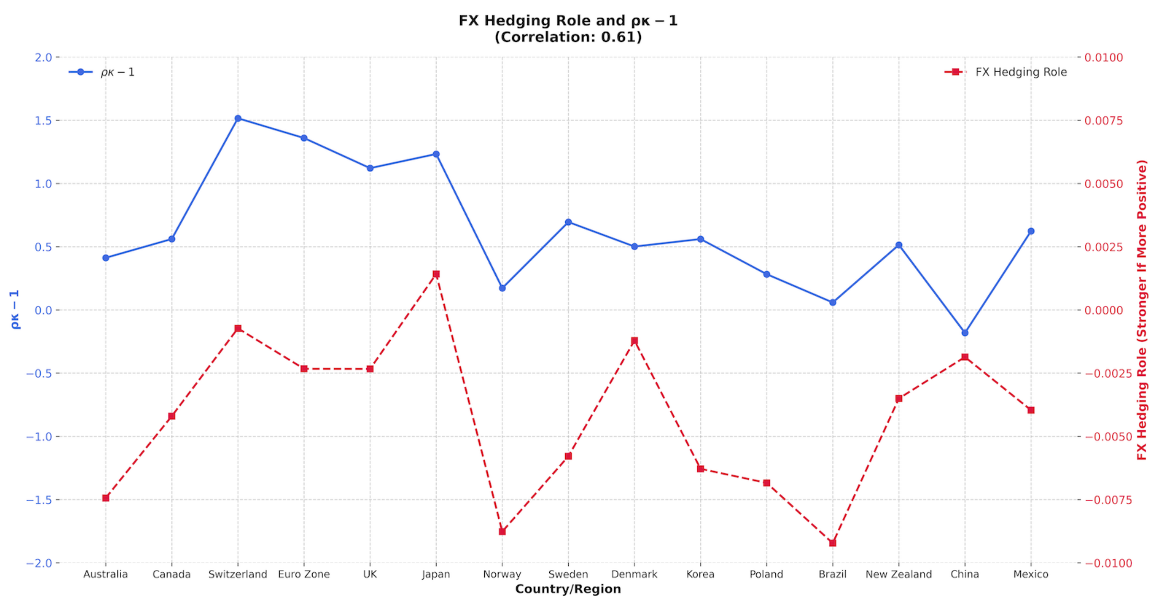

Figure 1.10: FX Hedging Role and $\rho\kappa - 1$, obtained for a 12-month investment horizon.

FX hedging properties tend to weaken when foreign equities co-move more closely with U.S. markets, in line with the mechanism described in Section 1.4. Appendix A.12 confirms this result using alternative samples.

In addition, the FX hedging role correlates *positively* with the relative volatility param-

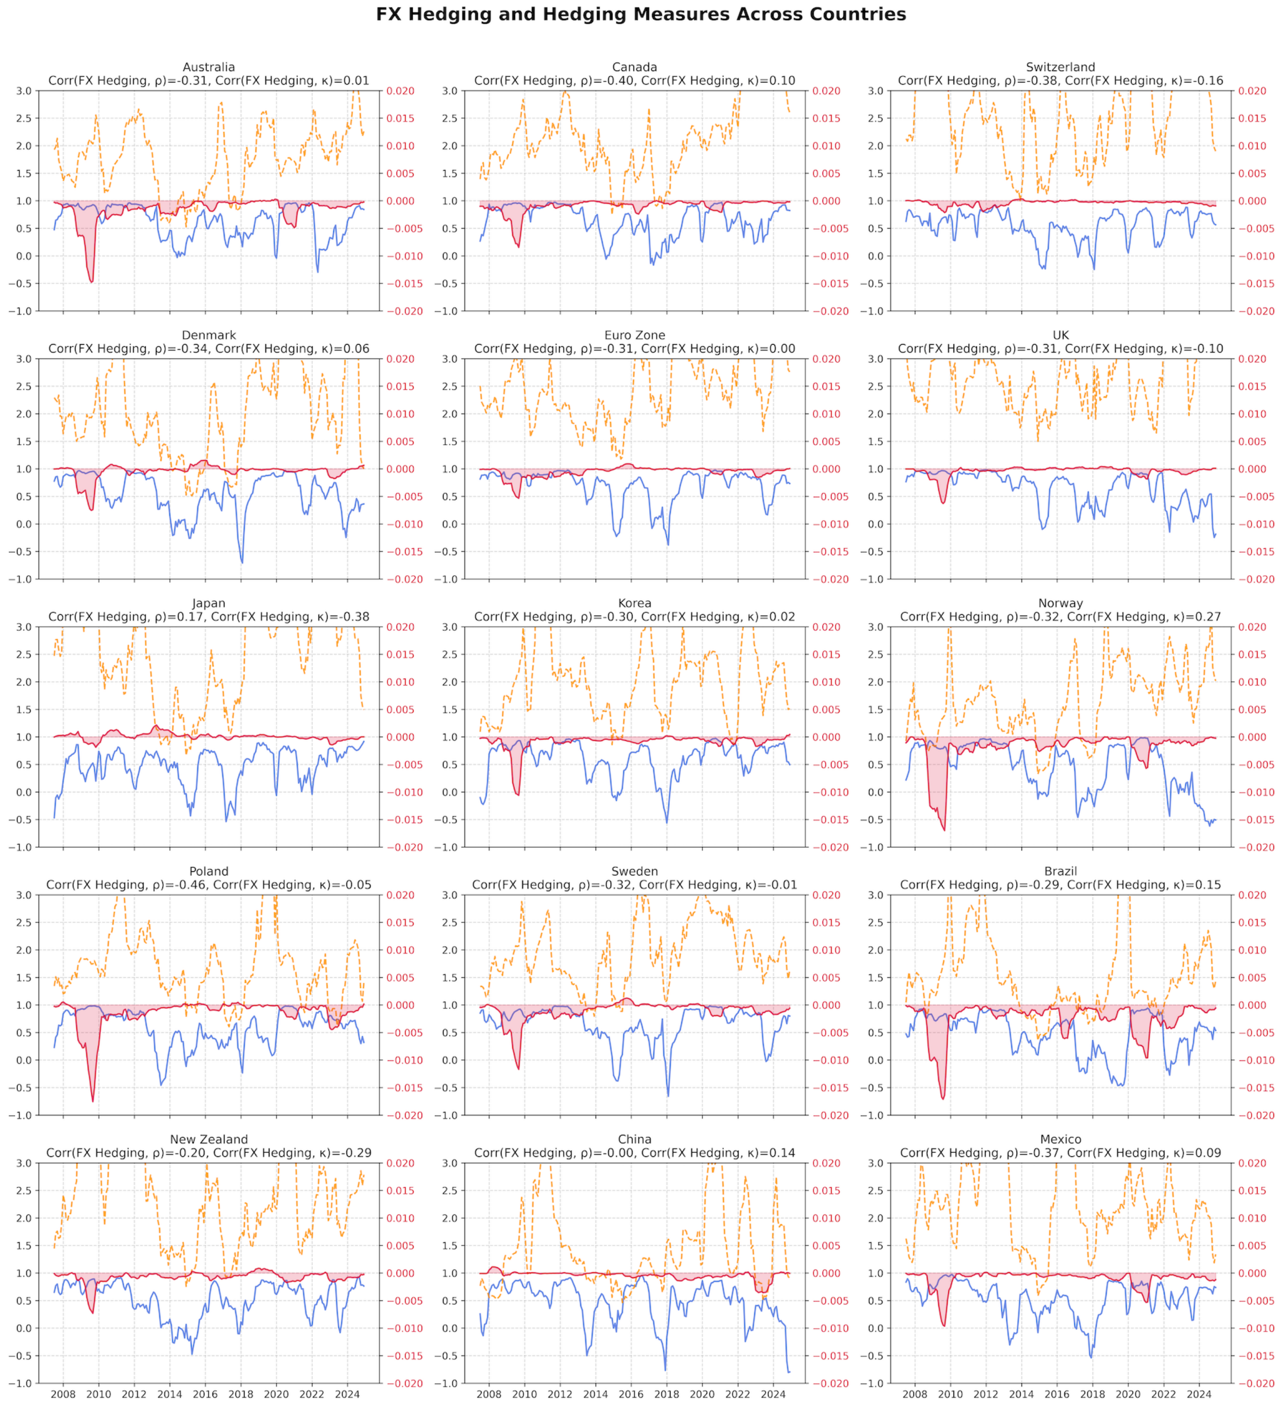

Figure 1.11: FX Hedging Role and Equity Market Parameters Over Time. 12-month investment horizon. Hedging roles are in red, ρ is in blue, and relative U.S. equity volatility κ is represented by the dashed orange line.

eter κ in 9 out of 15 cases. This pattern, though weaker than in the cross section, remains consistent with model-implied comparative statics, which predict stronger hedging effects when foreign equity markets are less volatile than the U.S. market. Once more, Appendix A.12 confirms this result with alternative samples. In Appendix A.13 I display results for the dollar hedging role: we should observe a *negative* correlation between this and κ, which is also verified empirically.

Prediction 3: UIP and FX Hedging

First, prediction 3 implies that FX hedging roles and UIP deviations should move together across countries. Specifically, currencies that lie on the *funding side* of carry trades — characterised by *positive* UIP deviations — should be more likely to act as hedges for domestic equity.

Figure 1.12 confirms this mechanism at the cross-country level. The blue line shows the unconditional 12-month UIP deviation, while the red line captures the FX hedging strength for each country, measured as $-\text{Cov}(\mathcal{R}^*, E_1/E_0)$, so that higher values correspond to stronger hedging role. A clear positive association emerges: the correlation is 0.56.

The dynamic version in Figure 1.13 reinforces this point using the time-series dimension. For the vast majority of countries (11 out of 15), the correlation between the UIP deviation and the FX hedging role is positive over time, consistent with the theoretical mechanism. When a currency becomes more clearly a funding currency (higher Ω^{UIP} for that bilateral relation), its hedging behaviour tends to strengthen.

Prediction 3: Testing the Full UIP–UEP–Hedging Relationship

To close the empirical analysis, I test the full decomposition linking UIP deviations, FX hedging roles, and relative volatility to UEP deviations, as implied by equation (1.19). The model predicts that each component of the decomposition should contribute positively to the UEP deviation, Ω^{UEP}.

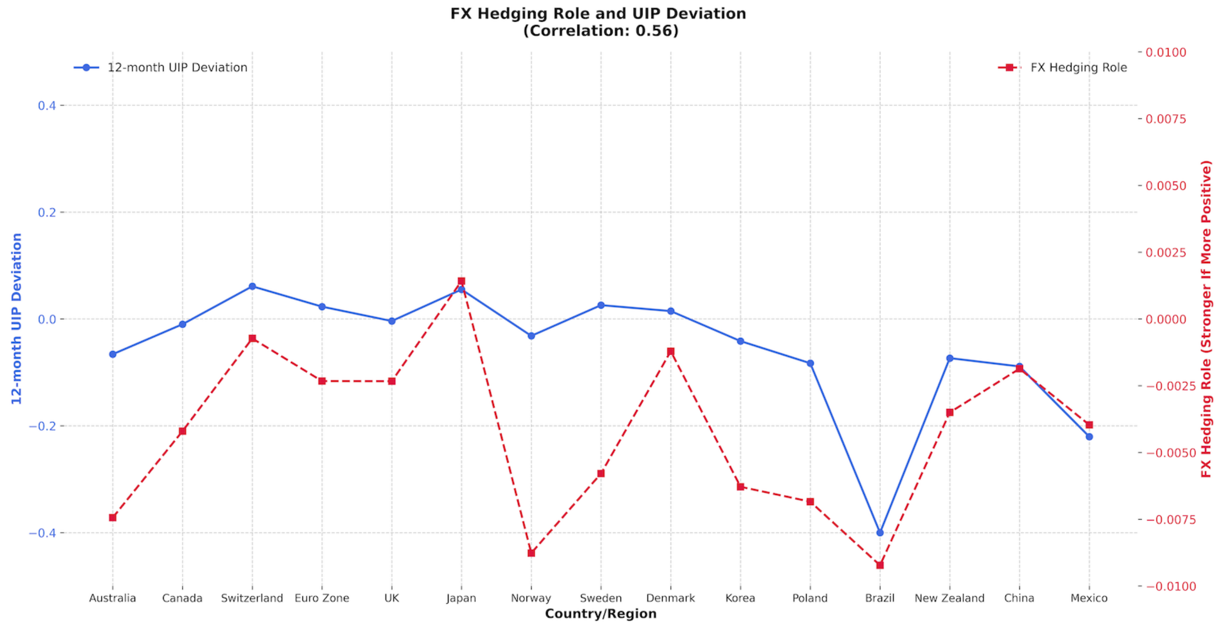

Figure 1.12: 12-month FX Hedging Role and UIP Deviation Across Countries. Full sample.

First, from the charts above, it is clear that $\kappa > 1$ for all countries, so the third term in the decomposition — which reflects relative equity volatility — is unambiguously positive. Second, the coefficient on the UIP deviation is equal to $\exp(\sigma^2) > 0$, and should therefore be positive across all pairings. Finally, the coefficient on the FX hedging term is 1 by construction in the model, and so should also be positive in all cases.

I implement two complementary econometric strategies:

1. **Bilateral Regressions.** For each country, I estimate a time-series regression of Ω^{UEP} on Ω^{UIP}, the FX hedging role and a constant. I check whether the sign of each coefficient aligns with the model: *all* three should be positive for *all* countries .

2. **Panel Regressions with Country Fixed Effects.** Since the coefficient on the UIP term depends only on U.S. equity volatility and the hedging coefficient is equal to 1, the model suggests they should be common across countries. I therefore estimate a panel specification imposing a common slope on these two regressors and absorbing relative volatility effects through country fixed effects. If the model is correct, these fixed effects

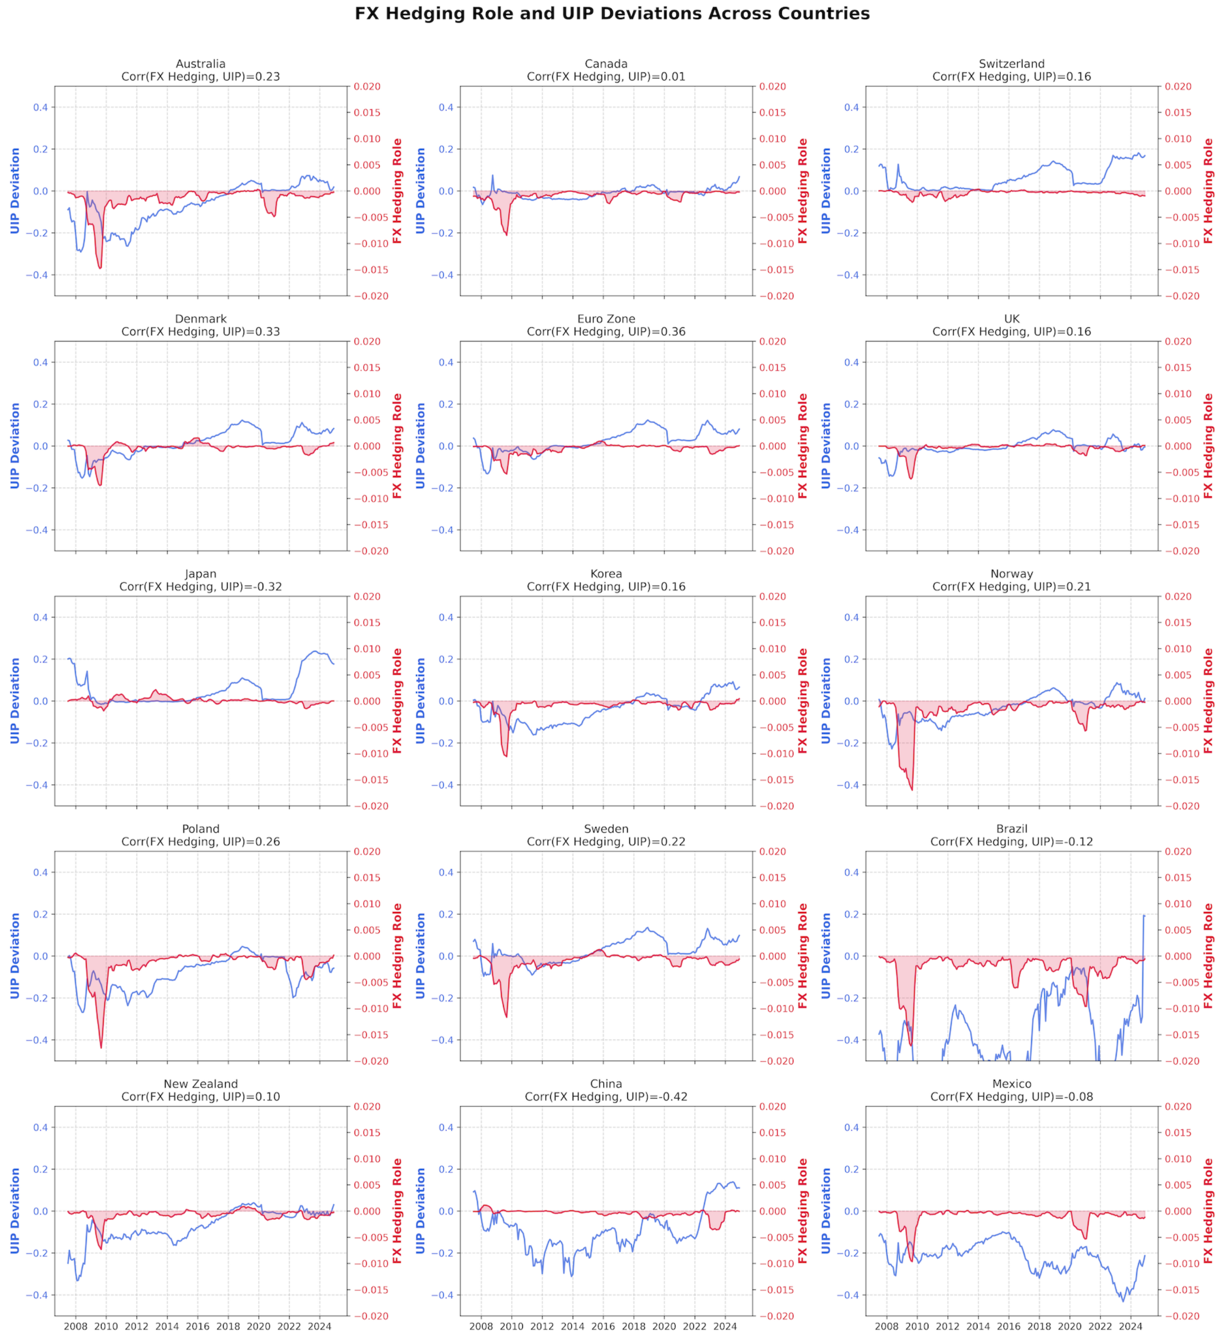

Figure 1.13: 12-month UIP Deviations and FX Hedging Roles Over Time for Each Country. Full sample.

should *all* be strictly positive, reflecting the fact that $\sigma > \sigma^*$ in each case.

Empirical Strategy 1: Bilateral Regressions

Table 1.1 presents the results of bilateral time-series regressions of UEP deviations on UIP deviations and the FX hedging role, testing the decomposition in equation (1.19). Since the data is sampled monthly but constructed over a 12-month investment horizon, observations are overlapping, and the residuals inherit a moving average structure of order 11. To account for this, I use heteroscedasticity- and autocorrelation-consistent (HAC) standard errors with 11 lags.

The results broadly support the model. First, all constant coefficients are positive, as predicted by the model (capturing volatility asymmetries), and 12 out of 15 are statistically significant. Second, the coefficient on the UIP deviation is positive in 12 out of 15 regressions and statistically significant in four cases. The coefficient on the hedging term is more variable — unsurprising given the small magnitudes of the covariances in the data — but still positive in 11 of the 15 regressions. Moreover, although estimates are imprecise, the wide confidence intervals generally do not reject a value of one. Finally, an F-test of joint significance rejects the null that all three coefficients equal zero in 8 of the 15 specifications.

In sum, the evidence from bilateral regressions is consistent with the structure imposed by the model. Results become even stronger when excluding the Global Financial Crisis, as shown in Appendix A.12.

Empirical Strategy 2: Panel Regression with Fixed Effects

Table 1.2 presents the panel estimates of the UIP–UEP–Hedging relationship, imposing common slope coefficients on the UIP deviation and FX hedging role while allowing for country fixed effects to absorb relative volatility differences. As predicted by the model, both slope coefficients are positive and statistically significant: the coefficient on the UIP deviation is 0.46 (and significant), while that on the FX hedging term is 8.81 (and also significant).

All country fixed effects are positive, consistent with the model's implication that relative

Table 1.1: Bilateral Regressions: UEP Deviation on UIP Deviation and Hedging Role

Country	Constant	UIP Deviation	Hedging Role	R^2	F-test p-value
Australia	0.129	0.710	12.439	0.236	0.001
	[0.068, 0.190]	[0.188, 1.233]	[1.146, 23.732]		
Canada	0.119	-0.326	42.054	0.162	0.004
	[0.074, 0.165]	[-2.330, 1.678]	[17.490, 66.618]		
Switzerland	0.076	-0.214	38.921	0.018	0.686
	[-0.018, 0.171]	[-0.980, 0.551]	[-54.576, 132.417]		
Denmark	0.060	0.375	17.017	0.109	0.001
	[0.008, 0.113]	[-0.538, 1.288]	[-3.957, 37.990]		
Euro Zone	0.101	0.267	30.098	0.086	0.115
	[0.016, 0.186]	[-1.049, 1.583]	[-13.276, 73.471]		
UK	0.120	0.573	17.744	0.116	0.051
	[0.090, 0.150]	[-0.286, 1.431]	[-2.337, 37.824]		
Japan	0.089	-0.034	45.417	0.025	0.546
	[-0.006, 0.185]	[-0.787, 0.720]	[-53.799, 144.633]		
Korea	0.115	0.155	20.626	0.138	0.000
	[0.081, 0.148]	[-0.212, 0.523]	[13.305, 27.946]		
Norway	0.143	1.011	-0.232	0.152	0.069
	[0.080, 0.207]	[0.020, 2.003]	[-10.578, 10.114]		
Poland	0.147	0.526	-3.577	0.069	0.384
	[0.077, 0.218]	[-0.220, 1.272]	[-14.242, 7.088]		
Sweden	0.106	0.409	23.648	0.213	0.000
	[0.042, 0.170]	[-0.264, 1.082]	[10.397, 36.899]		
Brazil	0.386	0.643	6.437	0.244	0.001
	[0.240, 0.533]	[0.293, 0.993]	[-10.618, 23.491]		
New Zealand	0.140	0.444	23.560	0.140	0.020
	[0.078, 0.202]	[-0.041, 0.928]	[2.552, 44.568]		
China	0.008	-0.447	-78.756	0.149	0.233
	[-0.147, 0.163]	[-1.266, 0.372]	[-175.653, 18.141]		
Mexico	0.299	0.913	-1.558	0.247	0.001
	[0.198, 0.401]	[0.450, 1.375]	[-31.170, 28.055]		

Notes: Standard Errors are heteroscedasticity and autocorrelation robust (HAC) using 11 lags. 12-month investment horizon. Full sample.

U.S. equity volatility ($\sigma > \sigma^*$) contributes positively to UEP deviations. Thirteen out of fifteen country effects are statistically significant at conventional levels.

Further results are provided in Appendix A.14. The coefficient on the UIP deviation is highly robust across specifications, and fixed effects remain generally positive. As expected, the coefficient on the hedging term is more sensitive to specification, reflecting the low empirical magnitude of equity–FX covariances in the data.

1.5.1 Further Empirical Patterns

Beyond the primary predictions tested above, the empirical analysis reveals additional patterns that are consistent with the model's structure. First, UEP deviations are predominantly positive across both time and countries when measured at a 12-month investment horizon (especially after the Global Financial Crisis). This is evident from Figures 1.15 and 1.14, and the same pattern holds for the 3-month horizon (see Appendix A.11). In other words, a long-short portfolio that buys U.S. equity and sells foreign equity tends to generate positive excess returns on average. These figures also reconfirm that UIP and UEP should exhibit a positive relationship, and the same should apply to the one between UEP and FX hedging roles.

This systematic return asymmetry involving the UEP deviations may reflect the fact that U.S. equity is riskier than its foreign counterparts. In the data, the unconditional standard deviation of U.S. equity returns is approximately 16%, while the average across foreign markets is just 8%. When analysed separately, the model would imply that this volatility differential — captured by a higher κ — should lead to higher UEP deviations, since a more volatile U.S. equity position increases the marginal risk exposure of the long-short portfolio. However, I verify a negative link both unconditionally (in the cross section) and conditionally (in the time series); see Appendix A.13.

Alternatively, persistent UEP deviations might be driven by ρ. In this regard, however,

Table 1.2: Panel Regression: UEP Deviation on UIP Deviation and FX Hedging Role (with Fixed Effects)

Variable	Coefficient	Std. Error	z-stat	p-value	[95% CI]
UIP Deviation	0.457	0.095	4.817	0.000	[0.271, 0.643]
FX Hedging Role	8.812	3.161	2.788	0.005	[2.617, 15.008]
Australia	0.106	0.029	3.664	0.000	[0.049, 0.163]
Brazil	0.318	0.049	6.538	0.000	[0.223, 0.413]
Canada	0.094	0.029	3.197	0.001	[0.036, 0.152]
China	0.130	0.035	3.664	0.000	[0.060, 0.199]
Denmark	0.055	0.021	2.651	0.008	[0.014, 0.096]
Euro Zone	0.084	0.026	3.275	0.001	[0.034, 0.134]
Japan	0.066	0.041	1.603	0.109	[–0.015, 0.146]
Korea	0.116	0.018	6.605	0.000	[0.082, 0.151]
Mexico	0.210	0.031	6.706	0.000	[0.148, 0.271]
New Zealand	0.129	0.021	6.028	0.000	[0.087, 0.171]
Norway	0.144	0.031	4.676	0.000	[0.084, 0.205]
Poland	0.161	0.030	5.461	0.000	[0.103, 0.219]
Sweden	0.087	0.022	4.009	0.000	[0.045, 0.130]
Switzerland	0.024	0.031	0.769	0.442	[–0.037, 0.084]
UK	0.115	0.017	6.851	0.000	[0.082, 0.148]

Model fit: $R^2 = 0.094$, Adjusted $R^2 = 0.089$,.
Notes: Number of observations: 3,105. Robust standard errors (HAC, 11 lags) are used to address autocorrelation due to overlapping 12-month returns. Country dummies absorb cross-sectional heterogeneity. The regression imposes homogeneous slopes on UIP deviations and FX hedging roles across countries, allowing intercepts to vary.

the model predicts that UEP deviations should not be directly explained by the correlation parameter ρ — see Proposition 3. Empirically, this prediction is borne out in the cross section: variation in ρ is not strongly associated with UEP deviations — also documented in Appendix A.13. It is difficult to pin down exactly what drives persistent returns in the long-short equity portfolio — I leave this question for future research.

1.6 Comparative Statics: Simulating the Full Model

To illustrate the model's behaviour across key parameters, I simulate comparative statics based on the full model's numerical solution. The calibration is chosen to reflect average unconditional moments in the data. In particular, the volatility of U.S. equity returns is set to 16%, and the average volatility of foreign equities to 8% — therefore, we have $\kappa = 2$. The average correlation between U.S. and foreign equity markets is calibrated at $\rho = 0.65$. For simplicity and in order to focus on the financial side of the model, I also set $\iota_0 = \iota_1 = 1$. I explore comparative statics over ρ and σ, holding all else fixed. The vertical dashed lines in Figures 1.16 and 1.17 mark $\rho = 0.65$ and $\sigma^* = 0.08$, respectively. Hence, for the second charts, we have $\kappa > 1$ for any level of σ that lies on the right side of that dashed line.

The goal of these simulations is not to match real-world asset positions and main variables in levels — this would be difficult given the two-period nature of the model — but rather to highlight directional movements that align with theoretical predictions. Notably, the simulations confirm three core mechanisms. First, hedging properties deteriorate as international equity correlations rise (higher ρ), reflecting a weakening of the FX channel during global downturns. Second, dollar hedging roles decline with higher U.S. equity volatility (i.e., higher κ). Third, when σ is low (i.e., lower κ), the foreign currency's hedging role also weakens, consistent with the previous discussions.

Interestingly, both the UIP and UEP deviations exhibit non-monotonicities around the point where $\kappa = 1/\rho$, where financiers switch the sign of their positions in foreign bonds

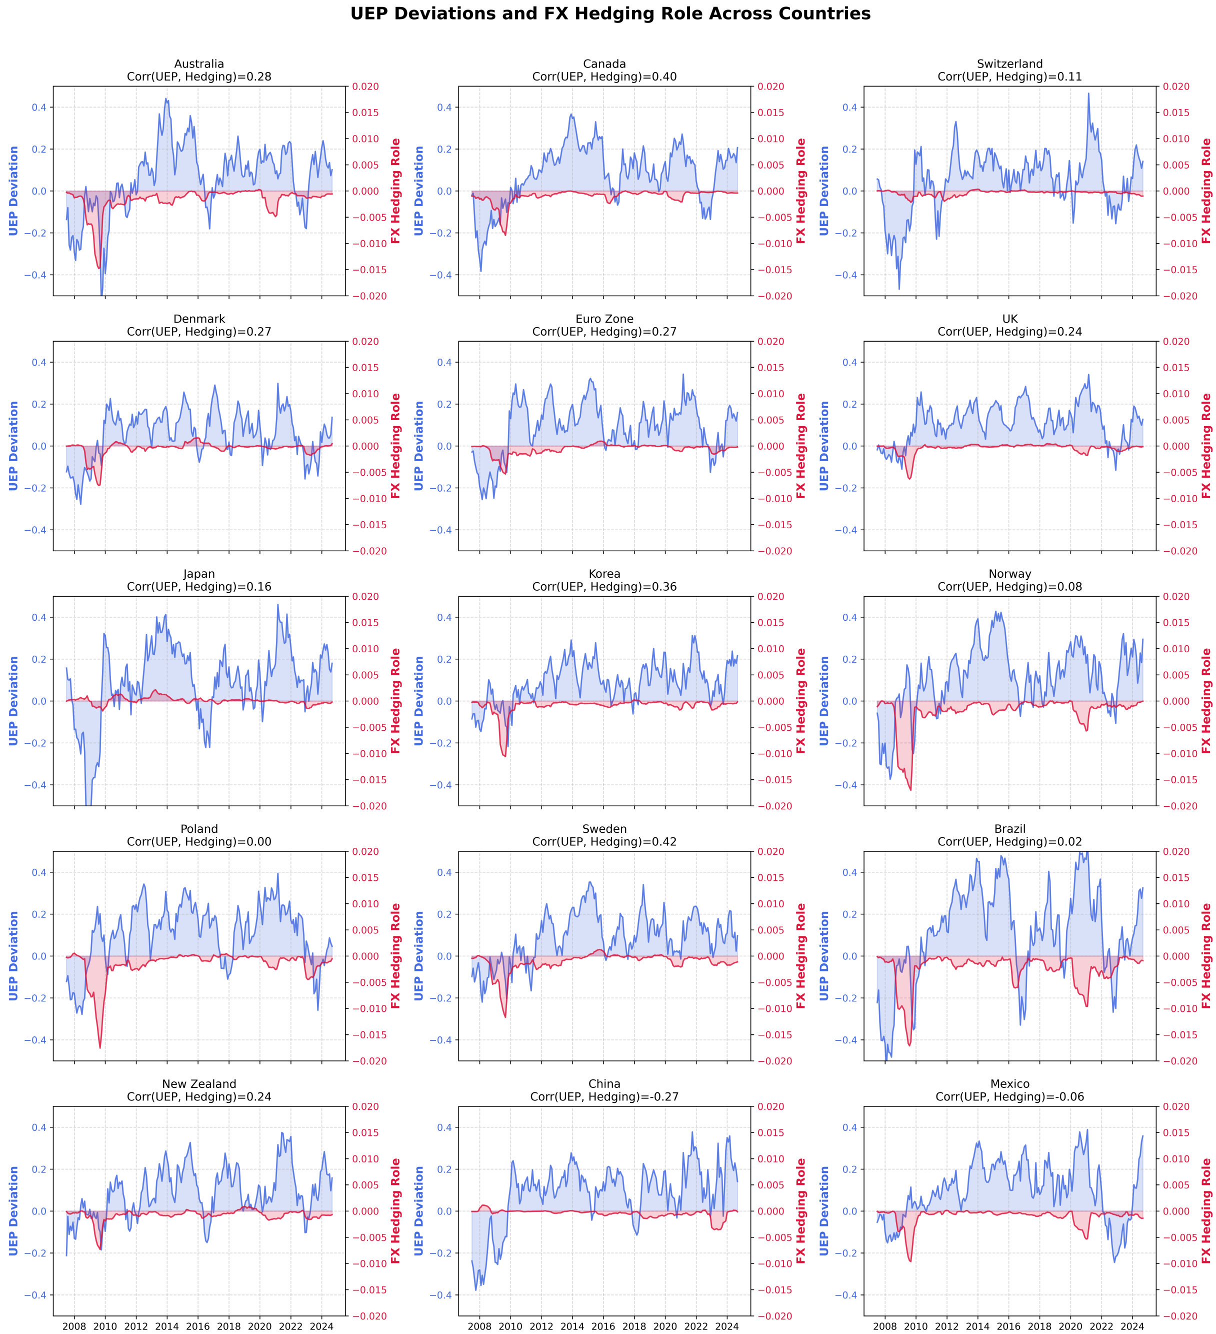

Figure 1.14: 12-month UEP Deviations and FX Hedging Role Across Countries. Full sample.

and equity. This mirrors the equilibrium selection logic derived analytically in the bond-only version of the model (see Proposition 12).

Figure 1.16 also highlights that more volatile exchange rates during global downturns cannot be attributed to increased equity market co-movement alone — a finding consistent

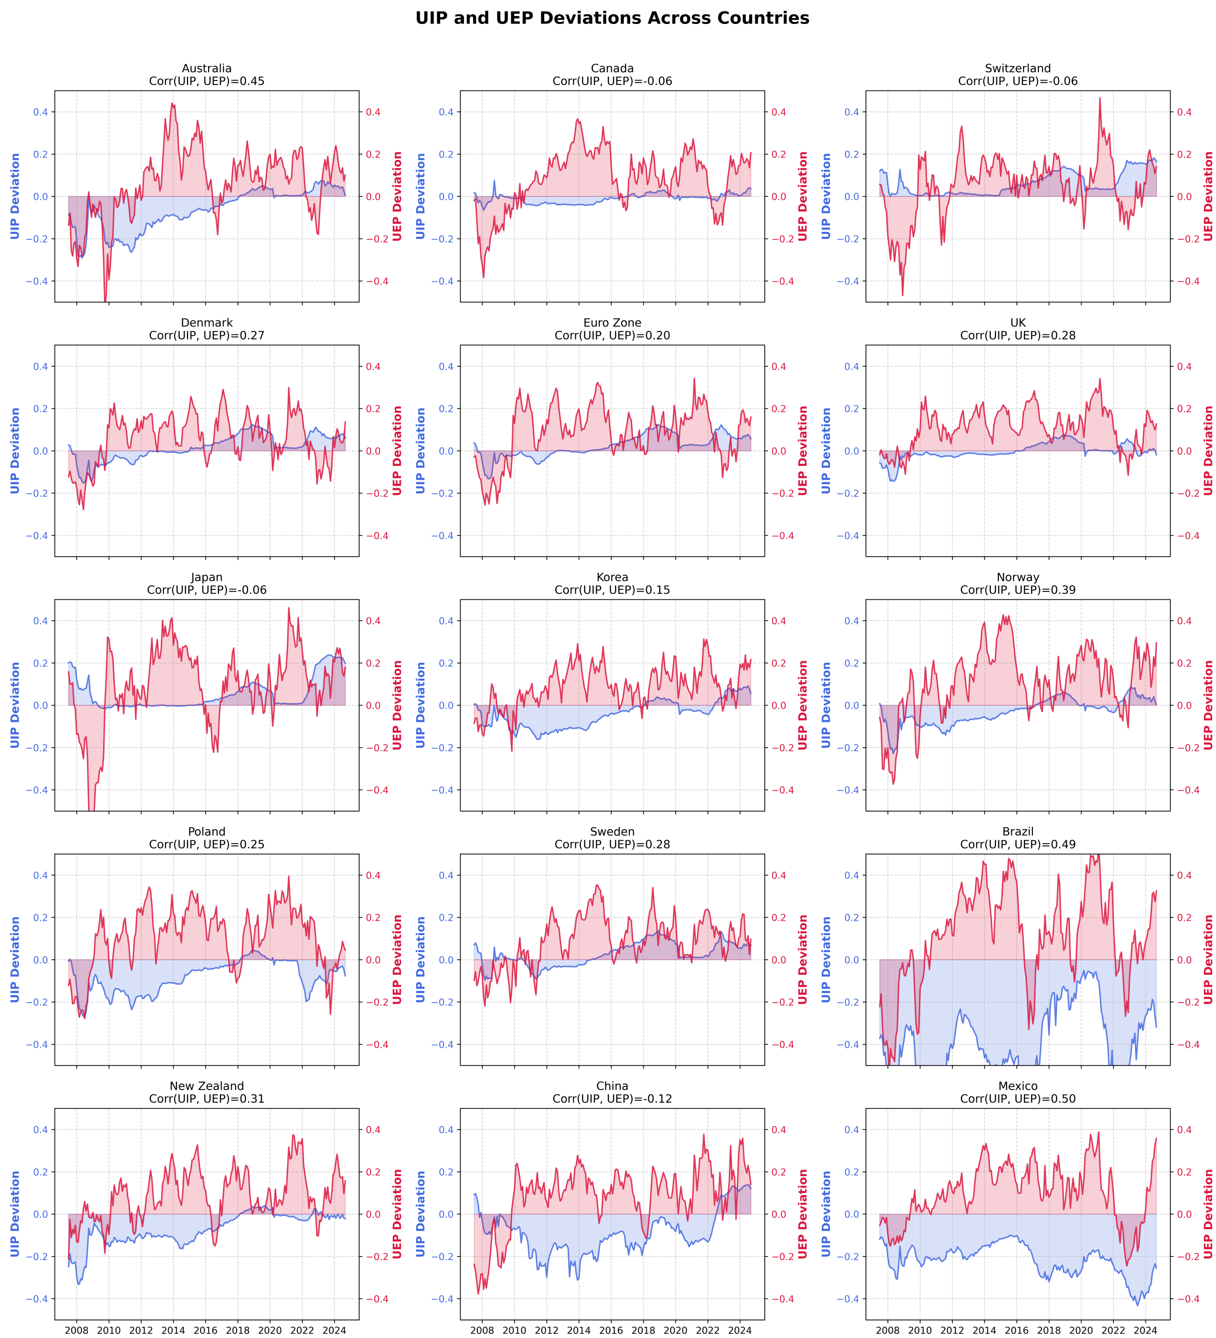

Figure 1.15: 12-month UIP and UEP Deviations Across Countries. Full sample.

with earlier discussions. By contrast, Figure 1.17 suggests that the phenomenon may be explained by spikes in U.S. equity volatility during these episodes. Interestingly, dollar variance displays a "smile" pattern across values of κ: it narrows when idiosyncratic risks are balanced, but increases when one market — either U.S. or foreign — becomes disproportionately risky.

The model also helps rationalise observed dollar appreciations during global downturns. These movements are consistent with shifts in ρ, as shown by the downward slope in "Dollar today" in Figure 1.16. Changes in κ can produce similar effects, but only if U.S. equity becomes markedly safer during crises — a pattern not clearly observed in the cross section. This reinforces the interpretation that increased global co-movement, rather than shifting volatility, is the primary driver of dollar strength in bad times.

Finally, the numerical simulations produce negative UIP deviations and positive UEP deviations for parameter values aligned with the data ($\kappa \approx 2$, $\rho \approx 0.65$). These patterns are in line with the cross-country empirical evidence presented in the previous section.

In Appendix A.9, I demonstrate that the main results are robust to alternative values of the trade preference parameters ι_0 and ι_1. That section also explores how shifts in these household preferences affect the composition of the financiers' balance sheet, thereby generating modest differences in the behaviour of the exchange rate relative to the benchmark case.

1.7 Conclusion

This paper presents a unified framework for understanding deviations from uncovered interest parity (UIP), uncovered equity parity (UEP), and the hedging properties of exchange rates. By introducing global financial intermediaries with joint exposure to international bond and equity markets, the model offers a tractable yet flexible structure capable of generating a rich set of empirical predictions. Central to this approach is the recognition that exchange rates are shaped not only by trade imbalances or monetary policy, but by the balance sheet exposures of a set of globally active investors.

The model delivers a decomposition of UEP deviations into three interpretable components: a UIP deviation scaled by U.S. equity volatility, a hedging motive term reflecting the interaction between currency and equity returns, and a volatility differential term capturing asymmetries in risk across markets. Each of these forces has a distinct empirical fingerprint, and jointly they offer a compelling lens through which to interpret exchange rate dynamics and cross-border asset flows.

Empirically, the model performs well. UIP deviations are directly associated with the relative responsiveness of U.S. equity markets, FX hedging roles emerge in the presence of equity market asymmetries, and UEP deviations are systematically explained by UIP, hedging motives, and volatility differentials. These results are robust across countries, time periods, and investment horizons. The model also generates equilibrium predictions for the direction of carry trades and the presence of currency hedging roles — predictions that are strongly supported empirically.

Finally, numerical simulations show that the comparative statics of the full model are consistent with key empirical features of exchange rates in global downturns. While not intended to match real-world levels, they clarify the equilibrium channels at work and highlight the internal consistency of the framework.

Taken together, these findings suggest that a multi-asset perspective on financial intermediation is essential for understanding global currency dynamics. The framework offers a foundation for future work examining the macro-financial transmission of shocks across borders, and for quantifying the shifting (and hedging) roles of safe-haven currencies in a world of increasingly integrated capital markets.

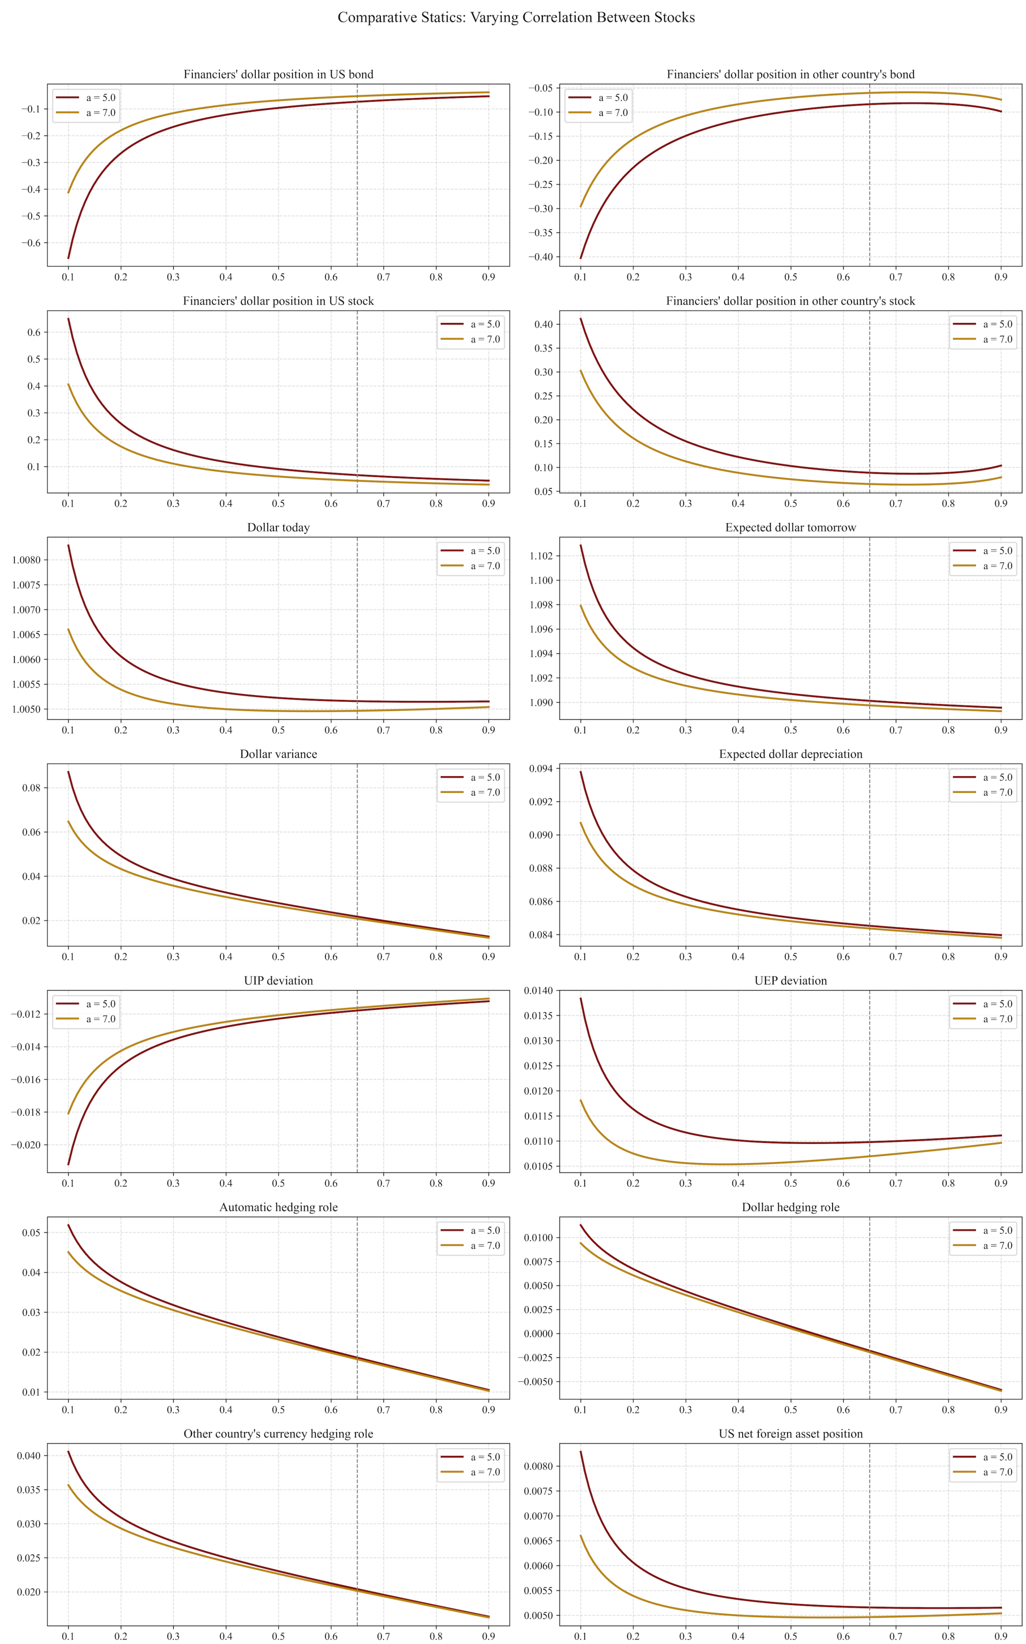

Figure 1.16: Comparative Statics: Varying Equity Correlation (ρ).

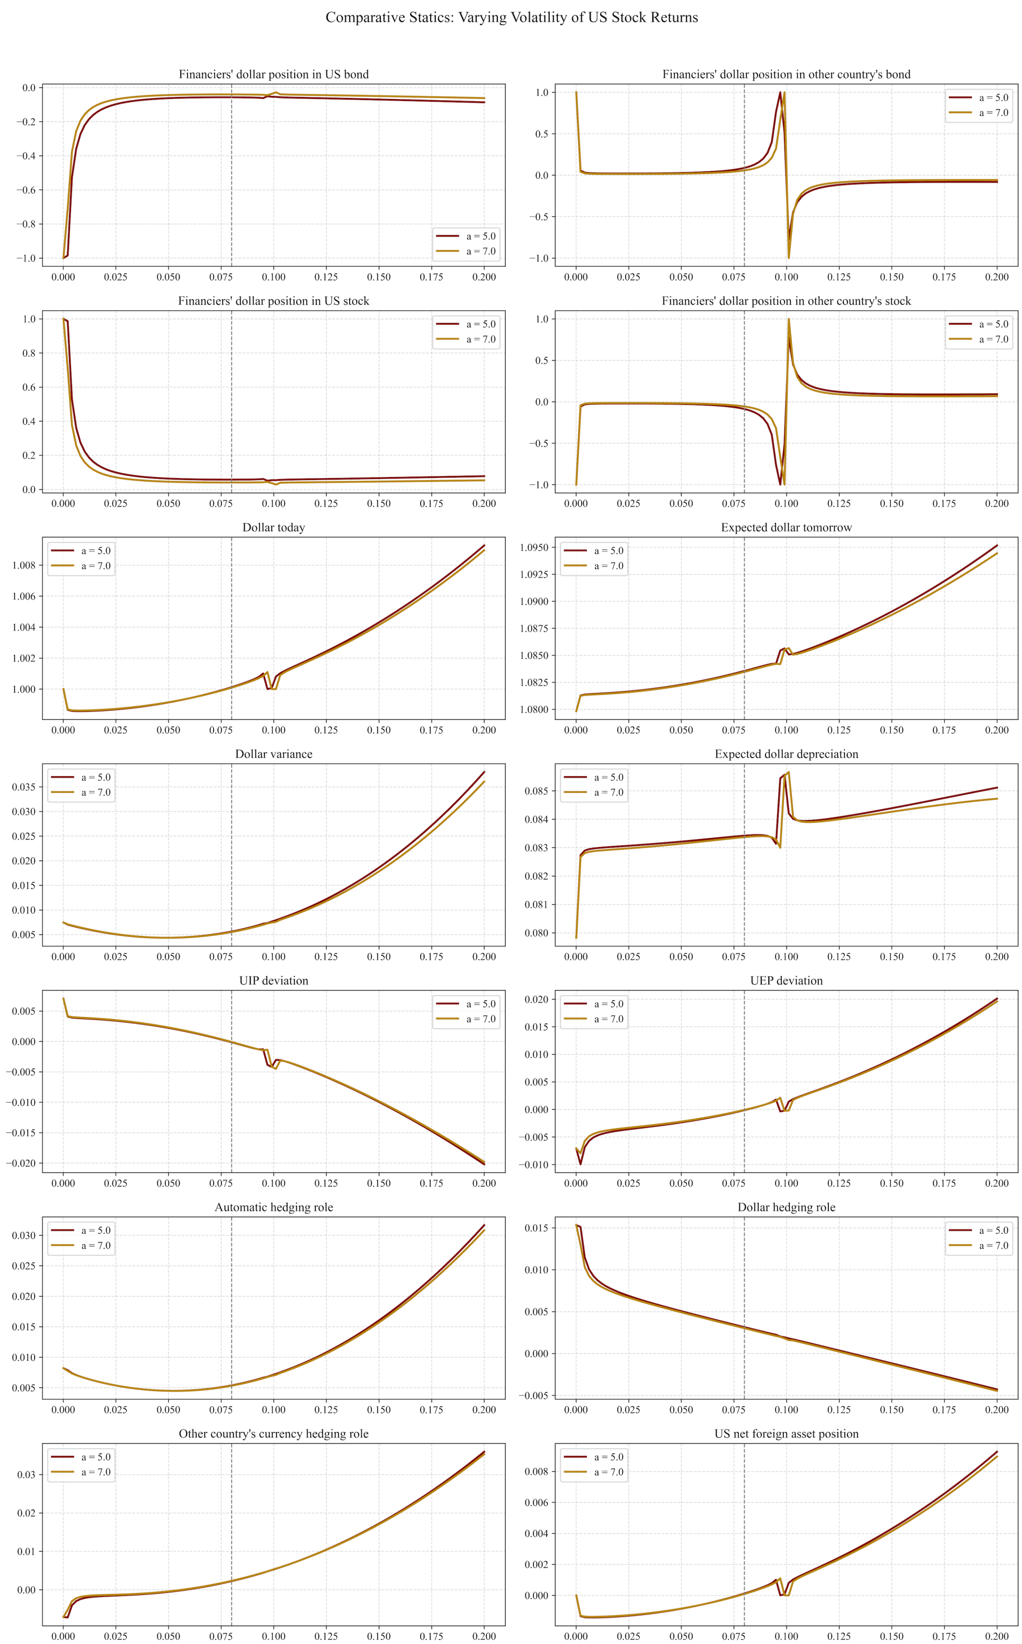

Figure 1.17: Comparative Statics: Varying Volatility of U.S. Stock Returns (σ).

Chapter 2

Intertemporal Substitution with Unfiltered Consumption

Co-authored with Carlos Carvalho and Ruy Ribeiro[1]

Abstract: *Most macro series used in academic research are usually smoothed, filtered and interpolated by official data providers. This paper shows that the use of filtered consumption series may considerably distort estimates of the Euler equation in consumption-based asset pricing models, more specifically of its slope, the elasticity of intertemporal substitution* (*EIS*)*. Once we use unfiltered consumption, we find that point estimates become more similar and confidence intervals can become tighter across different settings, data frequencies, as well as for different types of consumption data – macro and micro. Results also seem less sensitive to the presence of weak instruments, as, for instance, the completely uninformative weak-IV-robust confidence intervals usually found in the literature become rarer. Generally, we find that the EIS is quite low for macro data, albeit not as close to zero as commonly suggested in the literature. In this case, we often obtain values in the interval* $[0, 0.5]$*. For micro data, we estimate Euler equations conditional on the consumption of asset vs. non-asset*

[1]We are grateful to Christian Julliard for his insightful suggestions. We also thank participants of the International Workshop in Financial Econometrics for valuable feedback, especially Tim Bollerslev and Svetlana Bryzgalova. Finally, we are also thankful to Benjamin Moll.

holders. With unfiltered consumption, we do not find enough evidence of a different EIS across these groups. In addition, our estimates for stock holders are positive, but not above 0.3. Estimates for bond holders are higher, but more uncertain, usually from 0.4 to 1. In contrast, reported consumption seems unreliable, consistently returning negative estimates across groups.

2.1 Introduction

The Euler equation and the elasticity of intertemporal substitution (EIS, henceforth) play a central role in models of dynamic choice in macroeconomics and finance, capturing the sensitivity of consumption to the level of expected returns. For example, long-run risks models – Bansal and Yaron (2004), Bansal, Kiku, et al. (2014), Bollerslev, Tauchen, and Zhou (2009) and Bansal and Shaliastovich (2013), for instance – depend on the key assumption of an EIS above 1 to imply a sizeable equity premium, a low and stable risk-free rate and a correct cyclical behaviour of dividend-price ratios. However, the empirical literature has struggled to reach an agreement on a reasonable range for that parameter. Estimates seem heavily influenced by the specification, econometric method, different measures of returns used and the characteristics of the household considered in the studies[2]. It is still not clear what value should be considered as a reasonable guess to calibrate representative-agent models, for example. Early evidence from Hall (1988) had pointed to a "strong conclusion that the elasticity is unlikely to be much above 0.1, and may well be zero", but follow-up papers generally found mixed results and apparently none of them can provide a useful bridge to reconcile its empirical findings with values usually adopted in some macro models.

This paper attempts to improve estimates of the EIS by considering the fact that official consumption data has been filtered, smoothed and interpolated before release. Henceforth,

[2]See Gomes and Paz (2013) for an example of an alternative return measure constructed to capture the representative agent's asset portfolio. Moreover, characteristics of the household are relevant to the extent that they participate differently in local markets – see Vissing-Jorgensen (2002) and Guvenen (2006), for instance. Additionally, Havranek et al. (2013) find that distinct estimates of the EIS for different regions or cohorts seem more related to the specific assets held by different groups and their income than to local preferences.

we refer to these series as *reported* consumption. Kroencke (2017) shows that, in order to mitigate measurement errors, these statistical procedures undermine those series for research purposes by lowering their covariances with returns and introducing non-existent persistence. The former finding can be important to the extent that it implies that estimates based on reported data will be downward biased, what could explain the fact that a reasonable fraction of the papers in the literature state that the EIS is not statistically different from 0. Indeed, in the absence of the effects of such transformations, estimates might be higher and potentially more precise.

Our main findings suggest that *unfiltered* consumption – i.e., adjusted series that eliminate those noisy statistical procedures present in reported data – can significantly improve econometric results when estimating the EIS. These findings are confirmed for several econometric methods, types of consumption data (macro and micro), as well as for data at different frequencies. Unfiltered consumption is also important to obtain more precise estimates of the EIS when considering specific groups of asset holders in the Euler Equation, relevant issue when testing the limited asset market participation hypothesis (henceforth, LAMP), for instance.

First, using aggregate expenditures data from the National Income and Product Accounts (NIPA), we show that estimates of the EIS tend to increase relatively to cases that use reported consumption instead. Point estimates across different classes of estimators and frameworks also become more similar, being roughly from 0 to 0.5 in our baseline specifications, compared to -0.2 to 0.2, when using reported consumption.

Second, those estimates of the EIS seem less affected by the presence of weak instruments. We obtain more stable estimates across econometric methods regardless of whether unfiltered consumption produces lower first-stage F-statistics, compared to its reported analogue.

Third, while completely uninformative weak-instrument-robust confidence intervals[3] are

[3] We define uninformative intervals as either empty sets or ones that cover the whole real line.

quite frequent in the empirical literature of the EIS for macro consumption data, – Yogo (2004), Ait-Sahalia, Parker, and Yogo (2004), Ascari, Magnusson, and Mavroeidis (2016) and Gomes and Paz (2013) –, this paper shows that unfiltered consumption can transform these impractical intervals into more plausible sets. For instance, Yogo (2004) found uninformative sets in 66 percent of specifications estimated with reported macro expenditures data for the US economy. In contrast, for the same framework, we only obtain uninformative sets in 28 percent of our econometric approaches relying on unfiltered consumption.

Lastly, empirical benefits of unfiltered consumption series are also confirmed using micro data from the Consumer Expenditures Survey (CEX, henceforth). With this data, it is also possible to test the LAMP by obtaining distinct estimates of the EIS based on different groups of asset holders. We split households between stock and non-stock holders and between bond and non-bond holders. Unfiltered consumption is once again important. It produces estimates for stock holders that lie in the interval from 0 to 0.3. The EIS for bond holders is more uncertain, albeit higher, generally from 0.3 to 1. In contrast, estimations based on reported consumption exhibit a less clear pattern, consistently returning negative values across different groups of asset holders.

To construct unfiltered consumption data, we adapt the so-called Filter model in Kroencke (2017), so that we can use it to estimate the EIS using several types of data at different frequencies, rather than annual macro data only, as in the original model. According to the method, government statisticians who *only* collect an admittedly noisy observation for consumption opt to use a Kalman filter to estimate the unobserved level of consumed goods as precisely as possible – henceforth, we refer to the latter as *state* consumption[4]. Consequently, reported (or filtered) consumption is then defined by fitted values of this Kalman filter, while unfiltered consumption is obtained by reverse engineering to guess what their first

[4]Kroencke (2017) used the term *true* consumption instead. This reflects the fact that unobserved consumption in the model represents what government statisticians classify as *true* consumption according to their beliefs. We prefer to use the term *state* consumption in order to clarify the fact that this variable is modelled as *true* consumption *only* by those statisticians and should not be confused with the true level of consumed goods in the economy.

(supposedly noisy) observations were before being subject to the procedure. We propose a modification to the original model in order to introduce serially correlated measurement error. Specifically, our variation of the model is possible through the solution of a parallel quasi-differenced Kalman filter. We then map this solution back onto the original model, without changing main assumptions. This modification is essential to the extent that serially correlated measurement error terms are more relevant when either using disaggregated data or data at higher frequencies – see Wilcox (1992), Bell and Wilcox (1993) and the online appendix of Kroencke (2017)[5].

With the modified model in hands, we review identification approaches for the EIS usually adopted in the established literature, importantly those of Yogo (2004) and Vissing-Jorgensen (2002). The former addressed the empirical puzzle that estimates of the EIS are often statistically less than 1, while their reciprocal is not different from 1. He considered L. G. Epstein and Zin (1989) preferences and eleven developed countries while applying weak-identification-robust techniques. Although his final conclusions agree with Hall (1988), he finds point estimates that are rather imprecise across countries, mostly reflecting the presence of weak instruments. This inaccuracy was particularly true for the US economy, addressed in our paper[6]. A sensible explanation for this fact is that limited participation in asset markets may be plaguing results once Euler equations may no longer hold for the representative agent, possibility addressed by Vissing-Jorgensen (2002) and Guvenen (2006). The former used CEX panel data to verify how estimates of the EIS may differ taking into account different types of households in asset markets, as bond vs. stock holders vs. non-asset holders. The latter shows how lower estimates of the EIS are obtained when considering aggregates that ignore the facts that the majority of households do not participate in stock

[5]Consistent with that, we found significantly weaker results at higher frequencies when serially correlated measurement errors were not allowed. This was the case for NIPA consumption at quarterly frequency, for instance. Simply applying the original model of Kroencke (2017) provided such imprecise estimates of the EIS that even official data performed better in comparison.

[6]Yogo (2004) finds many empty and infinite weak-IV-robust confidence intervals for the EIS using US data independent from the data frequency, indicating that his baseline model is entirely rejected for this country. We use his framework in section 3 of this paper.

markets and that most of the wealth is held by a small fraction of population with a high EIS. Similarly, Ait-Sahalia, Parker, and Yogo (2004) review the so-called Equity Premium Puzzle – Mehra and Prescott (1985) – and estimate the EIS using not only the consumption of essential goods, but also that of luxury goods. While Vissing-Jorgensen (2002) finds that the EIS is not the same for stock and bond holders (0.3-0.4 and 0.8-1, respectively), results in Ait-Sahalia, Parker, and Yogo (2004) are somewhat inconclusive, albeit they do mention that the parameter is possibly higher regarding the consumption of luxury goods.

While none of the papers in the EIS empirical literature have considered that inaccurate estimates might be related to the fact that official data are filtered in order to mitigate measurement errors, there are a few papers in the asset pricing literature accounting for this fact. In addition to Kroencke (2017), Savov (2011) addressed the Equity Premium Puzzle and showed that reported consumption performs so poorly in asset pricing models that even the use of garbage data instead provides much better results. In a more complex framework and relying on Bayesian methods, Schorfheide, Song, and Yaron (2018) present a mixed-frequency approach that controls for measurement errors and time-varying volatilities[7]. In general, it is consensus that is quite hard to track true consumption in the data.

Our paper brings up the question about how the use of unfiltered consumption data may generate more reliable and precise estimates of the EIS. Indeed, we present evidence on how filtering out noisy elements present in official releases of consumption data (paradoxically, due to filtering of the original data) can help us to improve econometric results in the estimation of that parameter. Furthermore, we evaluate our findings relying on weak-identification routines. The use of these techniques in the EIS empirical literature does not seem sufficiently disseminated yet. In addition to Yogo (2004), only a few papers address the subject. Ascari, Magnusson, and Mavroeidis (2016), Ait-Sahalia, Parker, and Yogo (2004),

[7]They log-linearise and estimate a state-space representation that simultaneously accounts for consumption and its corresponding measurement errors at different frequencies. We come back to this later, but for now have in mind that the frequency of consumption matters much for researchers interested in asset pricing models that attempt to track implicit/noiseless consumption data.

Gomes and Paz (2013) and J. C. Fuhrer and Rudebusch (2002) are examples, albeit the latter in a more macro-based framework.

This paper is organised as follows. In Section 2.2, we present our modifications of the model in Kroencke (2017). The details on the complete model are available in the appendix. Section 2.3 evaluates how unfiltered consumption affects the estimates of the EIS in a framework of L. G. Epstein and Zin (1989) preferences and log-linearised Euler equations, in spirit of Yogo (2004). In Section 2.4, we feed the methodology into consumption measured by CEX data to verify the potential effects on estimates of the EIS across different types of asset holders, testing the LAMP. The last section concludes.

2.2 Model

This section describes how we adapt the filter model in Kroencke (2017), allowing for serially correlated measurement error. As mentioned earlier, this adaptation makes the model suitable for different types of data at several frequencies – rather than just for macro annual data, as in the original model. For the sake of conciseness, in this section we only cover parts of the model which are modified and that are relevant to a smooth reading of this paper. The complete model and its derivation are presented in the appendix.

We assume that statisticians who prepare the data for release collect a first (primitive) measure of consumed goods y_t, believed to be noisy. They conjecture that y_t is formed by *their belief* of true consumption c_t and an additive measurement error component ξ_t[8]:

$$y_t = c_t + \xi_t. \tag{2.1}$$

Henceforth, we refer to c_t as *state* consumption. We adopt this name to emphasise that the model captures what statisticians *believe* to be true consumption (c_t), rather than the correct

[8] You can see y_t as a first measure of consumption which has not been affected by filtering, smoothing and interpolation procedures. Alternatively, you can think of it as the garbage measure of Savov (2011).

measure of true consumption in the economy[9].

These statisticians model state consumption by a random walk representation:

$$c_t = c_{t-1} + \mu_{c,t} + \sigma_{\eta,t}\eta_t, \tag{2.2}$$

where $\eta_t \sim N(0,1)$ and we assume $\mu_{c,t} = \mu_c = 0$. Equation (2.2) does *not* mean that true consumption follows a random walk process nor that is has a constant drift. Instead, it only implies that government statisticians filter the data considering that true consumption follows that stochastic process, while assuming a constant drift[10]. We later assume that $\sigma_{\eta,t}$ follows a GARCH process.

To generalise the model, we introduce persistent measurement error by relying on the following AR(1) representation:

$$\xi_t = \rho_\xi \xi_{t-1} + \sigma_\nu \nu_t, \tag{2.3}$$

where $\nu_t \sim N(0,1)$. Generally speaking, it is trivial to expand a Kalman filter embedding (2.3)[11]. Harvey, Ruiz, and Sentana (1992) discuss how to model and extend these filters while assuming ARCH or GARCH processes for variance terms. Nonetheless, the model with (2.3) can not be solved in terms of unfiltered consumption using typical procedures. In this regard, we follow E. Anderson et al. (1996), rewriting the state-space representation in terms of a "quasi-difference":

$$\overline{y}_t = y_{t+1} - \rho_\xi y_t, \tag{2.4}$$

where $\overline{y}_t$ represents the "quasi-differenced" counterpart of y_t[12]. Once the solution for $\overline{y}_t$ is

[9]In addition, c_t will be the state variable in a Kalman filter, another reason for that name.

[10]Formally, we remove the mean of the series before calibrating the model, to then add it back to construct unfiltered consumption. These steps were also adopted in the original model and the use of data at different frequencies does not alter this part of the model.

[11]Perhaps the simplest form is to expand the vector of latent variables, now including ξ_t.

[12]Typically, a "quasi-difference" involves lags of the variable. We are following the term used in E. Anderson et al. (1996) here. From (2.4), we have that $[y_{t+1}, y_t, ..., y_0, \hat{c}_0]$ and $[\overline{y}_t, \overline{y}_{t-1}, ..., \overline{y}_0, \hat{c}_0]$ span the same space. By construction, this implies that prediction errors in $\overline{y}_t$ are actually innovations in y_{t+1}. See the appendix and

obtained, we then use (2.4) to map it back onto y_t. This final estimate of the latter ($\hat{y}_t$) is what we later call *unfiltered* consumption.

It is worth emphasising that we are not interested in fitted values of state variables after putting observable data into the filter. Instead, we want to estimate what the original observed data (y_t) were once *all* we have are fitted values of a state variable (c_t) tracked by a Kalman filter. The quasi-differencing approach makes the reverse engineering we have to deal with when inverting the Kalman filter possible without imposing additional complications to the way we solve the model[13]. It also does *not* mean that statisticians consider (2.4) when filtering the data. Instead, they solely consider (2.1), (2.2) and (2.3).

Equations (2.1) to (2.4) form together a quasi-differenced Kalman filter whose solution can be written as[14]:

$$\hat{c}_t = \hat{c}_{t-1} + K_t^c(\overline{y}_{t-1} - (1-\rho_\xi)\hat{c}_{t-1}), \tag{2.5}$$

$$K_t^c = \frac{P_t^c(1-\rho_\xi) + \sigma_{\eta,t}^2}{P_t^c(1-\rho_\xi)^2 + \sigma_{\eta,t}^2 + \sigma_\nu^2}, \tag{2.6}$$

$$P_t^c = P_{t-1}^c(1-(1-\rho_\xi)K_{t-1}^c) + (1-K_{t-1}^c)\sigma_{\eta,t}^2, \tag{2.7}$$

where $\hat{c}_t = E_t[c_t]$ denotes *reported* consumption (conditional time-*t* estimate of true consumption), P_t^c is the conditional variance of c_t and $\sigma_{\eta,t}^2$ denotes the volatility parameter in (2.2). Importantly, K_t^c is the Kalman gain associated with true consumption, what directly governs the persistence of reported consumption. Let $(\overline{y}_{t-1} - (1-\rho_\xi)\hat{c}_{t-1}) = u_t$ be the "re-scaled prediction error", a surprise factor[15]. When K_t^c is relatively high, statisticians attribute more weight to the surprise factor than to their past estimate $\hat{c}_{t-1}$ (reported consumption for the last period). Consequently, $\hat{c}_t$ is less persistent. Kroencke (2017) had also shown that unfiltered consumption exhibits higher covariances with expected returns, what we later confirm

E. Anderson et al. (1996) for more details.

[13]That is, we can solve the model following similar steps as in Kroencke (2017).

[14]Check the appendix for the derivation.

[15]This corresponds to the prediction error of the original model, but the term $(1-\rho_\xi)$ adjusts it for the presence of the quasi-difference $y_t - \rho_\xi y_{t-1}$ instead of y_t.

in our results in terms of the EIS[16].

One way to verify consistency of the filter is to check whether K_t^c increases in periods of economic turbulence (recessions, for example). Intuitively, when the variability of economic shocks is high relatively to the volatility of the measurement error, it is optimal for statisticians to adjust $\hat{c}_t$ taking into account surprising data more than their past estimates, $\hat{c}_{t-1}$. Consequently, K_t^c is higher and reported consumption less persistent.

Algebraically, such mechanism comes from the analogue of (2.6) in the original model of Kroencke (2017)[17]. However, since equation (2.6) is not as simple as his, we must derive parametric conditions under which the derivatives of K_t^c for P_t^c and $\sigma^2_{\eta,t}$ are positive as well (so that more economic turbulence implies a higher Kalman gain).

Fortunately, we find that our model behaves properly in this regard under reasonable parametric conditions. We present these conditions in proposition 1 below. For expository reasons, they are written in terms of a *homoscedastic* version of the model (when $\sigma^2_{\eta,t} = \overline{\sigma}^2_\eta$, while $\overline{K}^c$ and $\overline{P}^c$ are also fixed to steady-state values)[18]. Henceforth, we refer to the baseline model when K_t^c, P_t^c and $\sigma^2_{\eta,t}$ are time-varying as *heteroscedastic*, but we will later present results for its homoscedastic analogue as well. Importantly, bear in mind that heteroscedasticity in the model does not imply the assumption of heteroscedasticity in our estimations. The heteroscedastic version of the Filter solely assumes that statisticians model the volatility of state consumption as time-varying, but it does not impose any restriction whatsoever to moments of unfiltered consumption, used in our regressions.

Proposition 14. *If state consumption is homoscedastic and* $4\frac{\sigma^2_\nu}{\overline{\sigma}^2_\eta} > (1+\rho_\xi)^2\overline{\sigma}^2_\eta - (1-\rho_\xi)^2$*, then its*

[16]When it is the other way around, reported consumption becomes a very persistent and predictable series and its correlation with asset returns normally lowers in comparison.

[17]See Kroencke (2017), p. 54, equation (5).

[18]The same conditions are valid point-to-point in time, but derivatives must hold at any single period. If the filter converges to steady-state values, that should not be a problem. The formal proof as well as more details on the homoscedastic model are exhibited in the appendix.

unconditional variance and Kalman gain follow:

$$\begin{aligned} \overline{P}^c &= \frac{\overline{\sigma}_\eta^2}{2(1-\rho_\xi)} \left([(1-\rho_\xi)^2 \overline{\sigma}_\eta^2 + 4\sigma_\nu^2]^{\frac{1}{2}} - (1+\rho_\xi)\overline{\sigma}_\eta^2 \right), \\ \overline{K}^c &= \frac{\overline{P}^c (1-\rho_\xi) + \overline{\sigma}_\eta^2}{\overline{P}^c (1-\rho_\xi)^2 + \overline{\sigma}_\eta^2 + \sigma_\nu^2}, \end{aligned} \tag{2.8}$$

while:

$$\frac{\partial \overline{K}^c}{\partial \overline{P}^c} > 0; \qquad\qquad \frac{\partial \overline{K}^c}{\partial \overline{\sigma}_\eta^2} > 0 \iff \sigma_\nu^2 - (1-\rho_\xi)\overline{P}^c \rho_\xi > 0. \tag{2.9}$$

Proof. See appendix. □

In practice, a sufficiently small value of ρ_ξ ensures the second part in (2.9). In addition, we find that (2.9) is easily satisfied for different calibrations of the model.

Next, we need to derive a measure of unfiltered consumption that is compatible with the quasi-differenced filter above. Adapting the methods in Kroencke (2017) for our model, one can isolate $\overline{y}_{t-1}$ in (2.5) and conduct simple adjustments that account for time-aggregation bias to find[19]:

$$\hat{\overline{y}}_{t-1} = \frac{\hat{c}_t - (1-(1-\rho_\xi)\Omega_t)\hat{c}_{t-1}}{\Omega_t}, \tag{2.10}$$

where $\Omega_t = \alpha K_t$ and we set $\alpha = 0.8$, as in the original model. Equation (2.10) above represents "quasi-differenced unfiltered consumption". Once $\hat{\overline{y}}_{t-1}$ has been found, we need to transform it back into its primitive, unfiltered consumption, $\hat{y}_t$. Based on (2.4), we do this following:

$$\hat{y}_t = \hat{\overline{y}}_{t-1} + \rho_\xi \hat{y}_{t-1}. \tag{2.11}$$

More details on how we use (2.11), as well as on how we initialise our model are presented in the appendix.

[19]Check the appendix for more details.

2.2.1 Consumption Volatility

We consider a time-varying consumption volatility in (2.2), which follows a GARCH(1,1) stochastic process:

$$\sigma^2_{\eta,t} = a_0 + a_1 \eta^{*2}_{t-1} + a_2 \sigma^2_{\eta,t-1}, \tag{2.12}$$

where $\eta^*_t = \sigma_{\eta,t}\eta_t$[20]. With (2.12), we let the error term capture the inherent dynamics of the data, so that the random walk hypothesis is not an obstacle[21]. Furthermore, since we are ultimately interested in unfiltered consumption, it should not matter much if we have a random walk process with a data-driven model for its variance or a covariance-stationary model with constant volatility, as far as both generate unfiltered series whose moments are sufficiently similar. Indeed, we find that different calibrations for the GARCH specification generate very similar results if moments are relatively matched, so that (2.12) seems to perform well when applied to the data.

2.2.2 Adjusting Asset Returns

We also need time-aggregation-bias adjustments for returns, such that their timing is compatible with that of (2.10). These steps are identical to those in Kroencke (2017) when we use annual data and are unrelated to the presence of serially correlated measurement error, (2.3)[22]. However, we use a (necessary) slight adaptation when working with other data frequencies. Corrections are only performed on the return series when the econometric method applied uses unfiltered consumption. For reported consumption, asset returns need not be corrected for a different timing since we do not adjust that of reported data[23]. It is worth emphasising that these adjustments are not essential to validate the main findings of this

[20] Equation (2.12) is for the (baseline) *heteroscedastic* model. Presumably, $\sigma^2_{\eta,t} = \overline{\sigma}^2_{\eta}$ for the *homoscedastic* version.

[21] By modelling state consumption growth as i.i.d., we let the GARCH component (2.12) absorb the dynamics of the data, so that the choice for the process itself becomes less fundamental.

[22] Recall that the original model only handles annual data.

[23] We follow Kroencke (2017) once more here. Results barely change when we repeat our estimations with reported consumption while correcting the timing of returns.

paper, since similar results are found using raw returns data[24].

2.2.3 Calibration

We follow similar parameterisation techniques to those in Kroencke (2017). However, the quasi-differencing approach demands an additional step. Specifically, with the presence of (2.3) we also need to calibrate ρ_ξ. To the best of our knowledge, Schorfheide, Song, and Yaron (2018) is the only paper that estimates something similar in the literature. Relying on Bayesian methods, the analogue of ρ_ξ with monthly consumption data is estimated as 0.06 in their paper, albeit in a much more complex model. Discrepancies apart, when testing our filter on macro data (NIPA consumption) for different parametric combinations, we found out that it behaves properly for different values of ρ_ξ in a neighbourhood around 0.06 – such that K_t^c increases in recessions, in line with the intuition.

Figure 2.1 below presents how K_t^c varies for ρ_ξ fixed around that neighbourhood, specifically at 0.03, 0.06 and 0.09. Besides, if remaining parameters are calibrated such that benchmarked moments of unfiltered consumption are sufficiently aligned, we find that different values of ρ_ξ in that neighbourhood simply do not matter much[25]. Therefore, even when not using data at monthly frequency, we fix $\rho_\xi = 0.06$ throughout the paper, while adjusting remaining parameters following steps in Kroencke (2017). The only exception is when we estimate the EIS using annual macro data. As already mentioned, serially correlated measurement error terms are less relevant for aggregate data at lower frequencies. Hence, we apply the original model ($\rho_\xi = 0$) in that case.

Based on Bansal and Yaron (2004), we find that values of $\overline{\sigma}_\eta$ in a neighbourhood of $\sqrt{3} \times 0.0078 \approx 1.4\%$ for quarterly NIPA consumption and $\sqrt{12} \times 0.0078 \approx 2.7\%$ for annual NIPA

[24]We repeat our main tables for NIPA consumption (macro data) using raw returns in the appendix.

[25]In fact, correlations between different generated series of unfiltered NIPA consumption are similar once we change ρ_ξ while adjusting other parameters taking into account benchmark moments. Hence, the GARCH component of the model is perfectly adjustable to capture the consumption dynamics even when ρ_ξ is modified. Subsequent unfiltered consumption series are not econometrically distinguishable in terms of estimates of the EIS.

Figure 2.1: Kalman Gain Over Time for Different Values of ρ_ξ

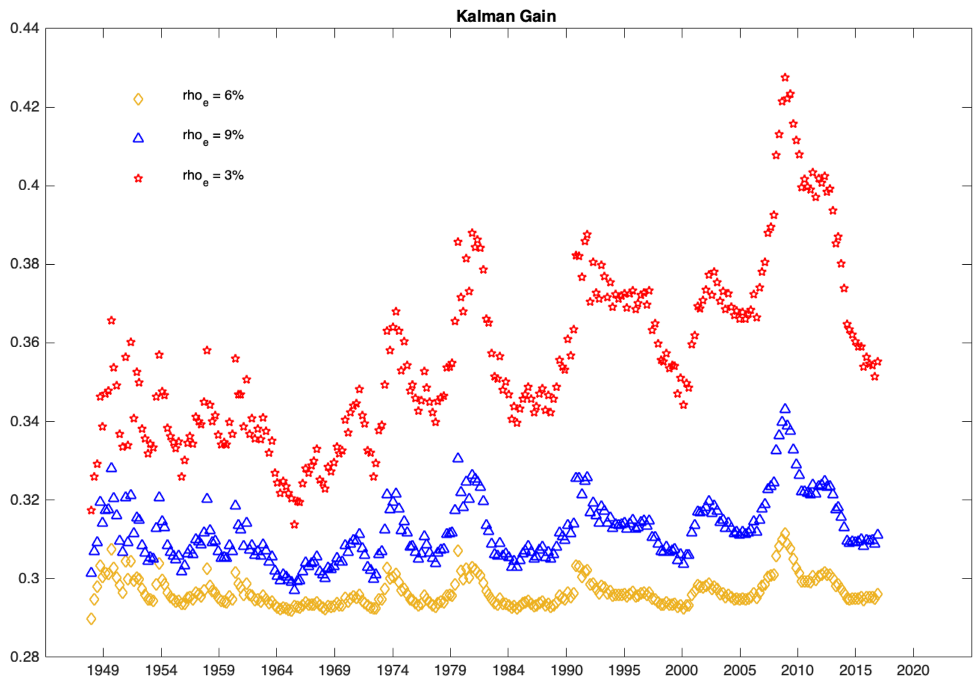

Note: Quarterly Kalman gain K_t^c in our sample for three different values of ρ_ξ (0.03, 0.06 and 0.09). The chart shows that consumption is less filtered (higher K_t^c) during the 2009 crisis, the early 2000's recession and early 90's and 80's recessions, for instance.

consumption match our moment requirements quite well[26]. Kroencke (2017) uses a slightly lower value of $\overline{\sigma}_\eta = 2.5\%$, what agrees with results in Dew-Becker (2016), who finds the same value for the long-run standard deviation of reported NIPA consumption. The small difference does not change moments substantially and, therefore, we set $\overline{\sigma}_\eta = 1.4\%$ and $\overline{\sigma}_\eta = 2.5\%$ for quarterly and annual data, respectively.

It is known that measurement error terms tend to cancel out over longer horizons — e.g. Daniel and Marshall (1996). Because of this, Kroencke (2017) uses a value for σ_ν that matches 6-year standard deviations of simulated and empirical data (garbage, reported and

[26]Those values represent the counterparts of $\overline{\sigma}_\eta$ in the model of Bansal and Yaron (2004), but adjusted for quarterly and annual data instead (they considered the value of 0.0078 at monthly frequency). The connection between their paper and the Filter model is not surprising. In fact, Kroencke (2017) used a modified version of their model to simulate state consumption (referred to as "true" consumption in that paper).

unfiltered as the latter)[27], calibrating the model based on post-war data. We calibrate σ_ν under the exact same method when using annual NIPA data (as in his paper). When using quarterly NIPA data, we accumulate quarterly expenditures to get an implied annual consumption growth measure whose moments satisfy the same procedure. These quarterly consumption data are actually benchmarked to annual counterparts before officially released. Hence, our calibration method is consistent with prevailing procedures conducted on the data[28].

We then choose σ_ν such that the 6-year standard deviation of unfiltered consumption is approximately 1.2 times the value of its reported counterpart, as in Kroencke (2017). This also implies that official statisticians do not make mistakes systematically when filtering the data, so that moments of unfiltered consumption should not considerably exceed those of reported consumption when measured over longer periods. For quarterly NIPA consumption, following this rule returns different calibrations across models: $\sigma_\nu = 3.8\%$ (heteroscedastic) and $\sigma_\nu = 2.5\%$ (homoscedastic). For annual NIPA data, statistical moments do not differ as much regarding the model and $\sigma_\nu = 2.8\%$ is set for both versions.

Table 2.1 compares moments of unfiltered and reported NIPA consumption based on nondurables and services. In the appendix, Table B.1 displays the same information for the consumption of nondurables only[29]. We present other relevant consumption measures shown in Kroencke (2017): *simulated* (he simulates state consumption using a long-run risk model built on Bansal and Yaron (2004)[30]); garbage (as in Savov (2011)), and; unfiltered (for which we simply show results in Kroencke (2017)). Moments are displayed both for the complete sample (1930-2022 for annual and 1947:3-2023-2 for quarterly macro data) and

[27]He uses that horizon as his benchmark based on considerations involving simulated data.

[28]Monthly and quarterly official consumption data are based on the monthly retail trade survey (MRTS), while annual data comes from the annual retail trade survey (ARTS). Since issues of sampling error are more significant in the MRTS, data from the latter are used to mitigate these problems. Hence, calibrating our model for quarterly macro data based on corresponding (implied) moments for annual macro data is consistent with actual steps conducted on the data before release. For more details on how NIPA consumption is generated, see the online appendix of Kroencke (2017) or the official NIPA handbook: BEA (2017).

[29]The calibration for this model is discussed in the appendix.

[30]As mentioned above, he refers to this measure as "true" rather than state consumption.

Table 2.1: Moments for NIPA Consumption of Nondurables and Services

(Implied) Consumption Growth	$E(\Delta C_{year})$	$\sigma(\Delta C_{year})$	$\sigma(\sum_{year=1}^{6} \Delta C_{year})/\sqrt{6}$	$Corr(\Delta C_{year}, \Delta C_{year-1})$
Reported (NIPA)	1.94%	1.41%	1.91%	32.01%
Simulated*	1.90%	2.48%	2.39%	1.54%
Garbage*	1.42%	2.86%	2.44%	-14.26%
Unfiltered - APWG* (1960-14)	1.85%	2.57%	2.44%	0.56%
Unfiltered - APWG* (1928-14)	1.79%	4.07%	3.08%	-10.89%
Unfiltered - Our Model (Quarterly Data)				
Homoscedastic (1960-14)	1.99%	2.88%	2.43%	-3.97%
Heteroscedastic (1960-14)	1.97%	2.30%	2.24%	2.69%
Homoscedastic (1947-23)	1.99%	3.26%	2.36%	-20.31%
Heteroscedastic (1947-23)	1.99%	2.49%	2.15%	-15.59%
Unfiltered - Our Model (Annual Data)				
Homoscedastic (1960-14)	1.30%	3.25%	2.45%	-8.22%
Heteroscedastic (1960-14)	1.91%	2.39%	2.26%	-0.71%
Homoscedastic (1930-22)	1.59%	5.11%	3.34%	-14.41%
Heteroscedastic (1930-22)	2.02%	3.71%	2.50%	-5.76%

Note: Moments of reported and unfiltered consumption (our model). We compare these moments with those of Kroencke (2017) as well: simulated consumption, garbage and unfiltered consumption (APWG stands for "Asset Pricing Without Garbage"). We have simply copied his results here, writing "*" next to variables presented in that paper. Reported and unfiltered consumption are for nondurables and services, from NIPA tables. We consider the quasi-differenced model with serially correlated measurement errors for quarterly data, setting $\rho_\xi = 0.06$. For annual data, the model is the same as in Kroencke (2017).

for the post-war subsample used by Kroencke (2017) to calibrate moments, covering the period 1960-2014. For unfiltered NIPA consumption, we present results for both variants of our model: one where state consumption has constant volatility (homoscedastic version) and another with time-varying volatility (heteroscedastic version, the baseline model).

Regardless of the data frequency, our measures of unfiltered NIPA consumption can reproduce the mean-reversion behaviour exhibited by garbage. In addition, unfiltered NIPA consumption is more autocorrelated in the complete sample than in the period comprehending 1960 to 2014, consistent with results found in Kroencke (2017) – see the first panel in Table 2.1. In contrast to its unfiltered analogue, reported NIPA consumption is quite persistent, consistent with the idea that the data are heavily filtered before release (lower K^c)[31].

Turning to micro (CEX) data, calibrated moments are exhibited in Table 2.2. The same procedures to calibrate the model are applied. The CEX data is subject to a number of statistical procedures before release, many of which relatively similar to those applied on NIPA data – see section 4 and the appendix for a discussion. It is also known that a significant fraction of the CEX consumption categories exhibit a similar behaviour compared to the NIPA analogues. Other categories do measure different things or have similar definitions but exhibit a CEX/NIPA ratio that is too low (high) over time. In terms of the estimation of the EIS, it is fundamental for the Filter model to be able to revert second moments and autocorrelations, as well as to exhibit higher covariation with returns. Inferring how much one source may be overestimating consumption growth relatively to the other is considerably less important. If overall there is no substantial change in how much CEX categories overestimate (underestimate) its NIPA analogues, then one can apply the same method to both sources when calibrating the model. In addition, it is a common procedure to aggregate consumption goods from the CEX taking the NIPA categories as reference[32].

As will become clear in section 4, the CEX data allows us to split households between different types of asset holders – stock holders vs. non-stock holders and bond holders vs. non-bond holders, for instance. Since official statistical procedures do not distinguish between different asset holders, we calibrate the model based on the consumption growth series of *all* households, imposing the resulting parameterisation to the consumption of the corresponding groups.

We have to make one small change to the calibration method when working with CEX data. The time series we construct in section 4 measures semiannual consumption growth, but at monthly frequency. To equalise scale and frequency, we transform the data into

[31] In addition, observable means of all variables are similar. This is intuitive since both measurement errors and treatment procedures made on the data shall cancel over time.

[32] See Attanasio and Weber (1995) and Vissing-Jorgensen (2002), for instance.

monthly consumption growth before calibrating the model, to then revert the scale back into semiannual consumption growth. Since we calibrate the model based on monthly consumption growth at the same frequency, we fix $\overline{\sigma}_\eta = 0.0078$, following Bansal and Yaron (2004), who used monthly data. For the same reason, $\rho_\xi = 0.06$ is set, motivated by results in Schorfheide, Song, and Yaron (2018). Finally, the same rule for $\overline{\sigma}_\nu$ applies, which establishes that the long-run standard deviation of unfiltered consumption is not higher than 1.2 times that of its reported analogue. It returns $\overline{\sigma}_\nu = 2.7\%$.

From Table 2.2, it can be seen that unfiltered CEX consumption (for all households) repeats the same patterns in Table 2.1, for NIPA data. In fact, unfiltered data are again more volatile, exhibiting more mean reversion than reported consumption (it is also the case regardless of the group of asset holders considered). Check the appendix for more details on how we calibrate the model for CEX data.

2.3 EIS Estimates with Unfiltered Consumption Data

In this section, we repeat the estimation approach of Yogo (2004), using unfiltered and reported consumption based on nondurables and services. Our results are also evaluated based on weak-identification methods. Alternative estimations are presented in the appendix, broadly reconfirming our main findings[33].

Under Epstein-Zin preferences, it is possible to derive typical Euler Equations usually used in the literature to estimate the EIS. These connect consumption growth with returns

[33]For instance, those estimations include the consumption of nondurables only, raw data for returns while using unfiltered consumption or applying the quasi-differenced model with serially correlated measurement error for annual data as well. In this section, we remove the first three observations (aiming to exclude the filter's training period) for estimations that use quarterly data. With annual data, we opt to use the entire sample, given the limited number of observations available. In the appendix, we also present additional results based on alternative samples.

Table 2.2: Moments for CEX Consumption: 1982 to 2013

Consumption Growth (Annual, Implied)	$E(\Delta C_{year})$	$\sigma(\Delta C_{year})$	$\sigma(\sum_{year=1}^{6} \Delta C_{year})/\sqrt{6}$	$Corr(\Delta C_{year}, \Delta C_{year-1})$	Observations per Month (Mean)
			NIPA Consumption		
Reported	1.65%	1.23%	1.78%	60.57%	-
Unfiltered	1.65%	2.13%	2.22%	11.51%	
			CEX: All Households		
Reported	1.98%	5.65%	2.64%	-36.40%	246
Unfiltered	2.10%	6.63%	3.17%	-55.45%	
			CEX: Stock Holders		
Reported	3.07%	7.36%	4.37%	-22.30%	49
Unfiltered	3.10%	10.30%	5.22%	-56.66%	
			CEX: Non-Stock Holders		
Reported	1.57%	5.22%	2.50%	-35.42%	197
Unfiltered	1.74%	6.34%	2.97%	-52.36%	
			CEX: Bond Holders		
Reported	3.14%	6.87%	4.01%	-22.69%	70
Unfiltered	3.25%	10.01%	4.96%	-58.49%	
			CEX: Non-Bond Holders		
Reported	1.27%	5.03%	2.18%	-39.34%	176
Unfiltered	1.45%	6.11%	2.65%	-54.42%	

Note: Moments of reported and unfiltered based on CEX data (1982-2013). We also exhibit moments of reported and unfiltered consumption based on macro data (NIPA consumption), calculated for the CEX period. The original CEX data measures semi-annual consumption growth at monthly frequency. We convert these series into monthly consumption growth to calibrate the model but aggregate the data to obtain moments for (implied) annual consumption growth – first column – so that these are comparable with moments in Kroencke (2017). In order to account for the fact that most likely statisticians do not adjust the data splitting by households, we calibrate moments based on all households. The last column provides the mean number of (cross-sectional) observations for each month, measured over the sample. Our CEX sample consists of 90,080 households.

of an asset class i[34]:

$$\Delta c_{k,t+1} = \tau_{i,t} + \psi r_{i,t+1} + \epsilon_{i,t+1}, \tag{2.13}$$

$$r_{i,t+1} = \zeta_{i,t} + \Theta \Delta c_{k,t+1} + \varrho_{i,t+1}, \tag{2.14}$$

where $\tau_{i,t}$ and $\zeta_{i,t}$ encompass mainly terms of second order (conditional variances and covariances), related to consumption growth and returns and $\epsilon_{i,t+1}$ and $\varrho_{i,t+1}$ also include expectational error terms. These are correlated with regressors in (2.13) and (2.14), so that an IV model must be adopted to properly identify their slopes. A standard log-linearisation shows that theoretically we must have $\Theta = 1/\psi$[35]. Nonetheless, it is often hard to show this result when relying on IV methods, regardless of the specification. Yogo (2004) addresses this puzzle, testing whether $\hat{\psi} = 1$ and $\hat{\Theta} = 1$, when individually estimating (2.13) and (2.14). He shows a rejection of the null in the first but not in the second estimation, pointing out that the presence of weak instruments may be substantially affecting these results.

We index consumption growth with k in (2.13) and (2.14) to emphasise the consumption series used. In the tables that follow, k can be *Reported*, *Unf-Hom*, or *Unf-Het*, where the last two refer to unfiltered consumption, constructed from the homoscedastic and heteroscedastic (baseline) models, respectively. The identification approach of this section does not depend on the hypothesis for heteroscedasticity nor on the asset type i[36]. Specifically, we conduct estimations with both stocks and risk-free returns. Second lags of the nominal interest rate, inflation, consumption growth (the measure relevant in the estimation, either reported or unfiltered) and log dividend-price ratio are used instruments. One could use the real interest rate rather than the nominal and inflation as instruments, but we prefer to follow Yogo (2004) strictly to elucidate the comparison. As in that paper, we lag all instruments twice to mitigate concerns of invalid moment conditions under conditional heteroscedas-

[34] We present the recursive form of L. G. Epstein and Zin (1989) preferences and their associated non-linear Euler Equations in the appendix. Structural forms of (2.13) and (2.14) can be seen in Yogo (2004).

[35] See Yogo (2004).

[36] If i is the risk-free rate, for example, only $\tau_{i,t}$ and $\zeta_{i,t}$ change and some of their second-order terms become null.

ticity in (13) and (14). In addition, instruments that are lagged at least twice ensure that problems involving time-aggregation in consumption do not affect estimates, as advised by Hall (1988).

2.3.1 Homoscedastic Framework

Here we present estimates for the EIS (ψ) and its reciprocal ($1/\psi$) using equations (13) and (14), respectively, assuming conditional homoscedasticity. We apply three K-class estimators: TSLS, Fuller-K and LIML[37]. First-stage F-statistics to infer about the relevance of instruments are also reported[38]. Critical values for them under the null hypotheses in Stock and Yogo (2002) are presented in the appendix. In general terms, F-statistics above 10 ensure that the TSLS bias is low enough to be reliable, while the Fuller-K bias is not high enough when that number is above 6. Under conventional first-order asymptotics, all those three estimators should converge to the same limit distribution, with the TSLS being the efficient one. In contrast, under weak instruments, the Fuller-K and LIML are more robust estimators[39]. Here we present estimations for reported and unfiltered consumption series which are constructed from the consumption of nondurables and services component, found in the NIPA tables. Additional estimations for the consumption of nondurables only are provided in the appendix. They produce similar results.

It is worth reemphasising the difference between a homoscedastic filter and conditionally homoscedastic error in the Euler Equation. The former only implies that statisticians *filter* the data assuming a constant volatility parameter for *state* consumption. However, the homoscedastic model does *not* imply whatsoever that errors in the Euler Equations will be

[37] Check the appendix for a better description of those estimators.

[38] If error terms are not serially correlated and homoscedastic, the first-stage F-statistic is a sample analogue of the so-called concentration parameter, that captures how relevant instruments are. When the F-statistic (and the concentration parameter) is sufficiently high, the TSLS is reliable, approximately unbiased and its t-statistic exhibits a proper convergence towards a standard normal.

[39] Under weak instruments, the TSLS can be severely biased, compared with the Fuller-K. Additionally, the Wald test that corresponds to the LIML estimator is less size-distorted than that of the TSLS – see Stock and Yogo (2002), Stock, Wright, and Yogo (2002) and Murray (2006) for more details.

conditionally homoscedastic. The more generalised heteroscedastic Filter model is our baseline, but we also exhibit results for the case unfiltered consumption is constructed using the homoscedastic model.

Table 2.3 below displays results for quarterly data, where unfiltered consumption considers the quasi-differenced Filter model ($\rho_\xi = 0.06$). The first thing to note is that there is much more agreement on estimates of the EIS (ψ) across different estimators when unfiltered consumption is used. Moreover, negative point estimates are completely absent, broadly in line with economic intuition.

Our point estimates for the EIS under the heteroscedastic model lie in the range 0.15-0.38, roughly in line with Hall (1988), Yogo (2004) and L. Epstein and Zin (1991)[40]. In addition, higher point estimates can be obtained using the homoscedastic filter, from 0.20 to 0.57. Bansal and Yaron (2004) show that, when consumption volatility is time-varying, the EIS will be downward biased when regressing consumption growth on the risk-free rate. Hence, it is interesting to note that, once we fix that volatility (homoscedastic model), estimates of the EIS roughly double, in line with their observation.

In contrast, reported consumption shows quite a different picture, with point estimates from -0.20 to 0.07. On the one hand, negative estimates are frequent with reported consumption when we use stocks, and results indicate that a lower first-stage predictability may be explaining this fact. On the other hand, unfiltered consumption produces a very narrow interval of positive point estimates across estimators, roughly from 0.16 to 0.19 under the heteroscedastic model, but with a similar first-stage F-statistic. Standard errors are also more equalised using unfiltered consumption and stocks, suggesting that the former alleviates problems related to weak instruments. However, it is likely that such issue is still partially plaguing estimates, given the low first-stage F-statistic, slightly above 4. When the risk-free is used, these statistics generally exceed critical values of Stock and Yogo (2002), indicating

[40] Our results lie in the lower end of estimates in the latter (0.17-0.87), albeit their results restricted to consumption of nondurables and services are very similar to ours.

that weak identification is not an issue[41].

The lower part of Table 2.3 confirms that the same improvements observed for the EIS under unfiltered consumption are valid for its reciprocal. In fact, our estimates of $1/\psi$ return a notably wide range using reported consumption (from -5.01 to 18.98), regardless of the asset class used. The story is once again quite different for unfiltered consumption, with estimates in the narrower range $0.35 - 5.25$ under the heteroscedastic model, implying an EIS from 0.19 to 2.86[42]. Considering the more robust Fuller-K and LIML estimators, implied EIS estimates using that model and quarterly data are in the range $0.19 - 0.70$, basically in line with direct estimates. Although exhibiting lower first-stage F-statistics, the use of the homoscedastic model returns $1/\psi$ from 0.33 to 4.25. Under Fuller-K and LIML, this implies an EIS in the range from 0.17 to 0.97. Differences between the TSLS and the other two estimators are broadly expected since, as mentioned above, the Fuller-K and LIML estimators are more robust to weak instruments. Since first-stage F-statistics are decreased by a factor of three when unfiltered consumption is used, this justifies the gap[43]. Generally, our estimates for $1/\psi$ under (2.14) agree with what we obtain for the EIS (ψ) using (2.13), considering the more robust estimators[44].

Table 2.4 below presents similar results but for annual data, for which we simply recalibrate the original Filter model in Kroencke (2017)[45]. The big picture is very similar to that of quarterly data. By using unfiltered consumption, we once again get rid of negative

[41]Gomes and Paz (2013) find the same result, conducting similar estimations for the risk-free rate, but using an alternative measure of returns, there argued to better capture the portfolio of the representative agent.

[42]Table B.6 (appendix) exhibits implied estimates of the EIS (ψ) from estimates of its reciprocal ($1/\psi$) using (2.14).

[43]The lower first-stage F-statistic for unfiltered consumption makes sense, once Table 2.1 (appendix) shows that unfiltered consumption is not as serially correlated as its reported analogue and consumption growth is the endogenous regressor in (2.14).

[44]That being said, we could not revert the puzzle that ψ is generally statistically different from 1 but not its reciprocal. In this regard, Table 2.3 provides unclear results, what can indicate that weak-instruments are still affecting estimates when unfiltered consumption is used, even though not as heavily as with reported consumption.

[45]Recall that serially correlated measurement error is not necessary at annual frequency, so we do not use (2.3). Consequently, our model is no longer quasi-differenced, being exactly that of Kroencke (2017). We still present results for annual data imposing the quasi-differenced model ($\rho_\xi = 0.06 \neq 0$) in the appendix. We show that we once more can improve estimates of the EIS, albeit with somewhat weaker results.

Table 2.3: Estimates of the EIS Using K-Class Estimators and Quarterly Data

			K-Class Estimator			
Asset	Estimate	Δc_k	TSLS	Fuller-K	LIML	1S-F
Risk Free	ψ	Reported	0.067^{***}	0.053^{***}	0.053^{***}	22.172
			(0.078)	(0.093)	(0.093)	
	ψ	Unf-Hom	0.527	0.566	0.573	22.335
			(0.467)	(0.481)	(0.484)	
	ψ	Unf-Het	0.346^{*}	0.379^{*}	0.385^{*}	22.474
			(0.336)	(0.350)	(0.352)	
Stocks	ψ	Reported	0.006^{***}	-0.101^{***}	-0.199^{***}	4.575
			(0.017)	(0.096)	(0.213)	
	ψ	Unf-Hom	0.204^{***}	0.221^{**}	0.235^{**}	4.462
			(0.108)	(0.114)	(0.120)	
	ψ	Unf-Het	0.156^{***}	0.178^{***}	0.191^{***}	4.310
			(0.080)	(0.088)	(0.093)	
Risk Free	$\frac{1}{\psi}$	Reported	0.438^{*}	4.953	18.979	6.630
			(0.311)	(4.475)	(33.660)	
	$\frac{1}{\psi}$	Unf-Hom	0.331^{***}	1.031	1.745	1.890
			(0.154)	(0.683)	(1.473)	
	$\frac{1}{\psi}$	Unf-Het	0.349^{***}	1.434	2.599	2.268
			(0.168)	(0.997)	(2.380)	
Stock	$\frac{1}{\psi}$	Reported	0.795	-4.150	-5.014	6.630
			(2.724)	(4.936)	(5.346)	
	$\frac{1}{\psi}$	Unf-Hom	2.884	3.559	4.251	1.890
			(1.372)	(1.739)	(2.159)	
	$\frac{1}{\psi}$	Unf-Het	3.427	4.491	5.247^{*}	2.268
			(1.601)	(2.137)	(2.565)	

Notes: Estimates of the EIS and its reciprocal using (2.13) and (2.14) and quarterly data. Unfiltered consumption extracted relying on the quasi-differenced Filter model whose measurement errors are serially correlated ($\rho_\xi = 0.06$). All consumption series refer to nondurables and services. We apply the same setting of Yogo (2004), using 3 types of K-class estimators and assuming that errors conditionally follow a martingale difference sequence. Reported denotes official consumption data from NIPA tables. Unf-Hom and Unf-Het refer to unfiltered consumption, constructed by the homoscedastic and heteroscedastic models, respectively. When reported consumption is used, asset returns have not been adjusted for time-aggregation. Standard errors are presented in parentheses. The null that the estimated coefficient equals 1 has been tested: ***, ** and * denote rejection of the null hypothesis at 1, 5 and 10 percent significance levels.

estimates of the EIS, obtaining point values in the limited range $0.09 - 0.23$. Such stability is even more surprising considering the fact that first-stage F-statistics are considerably lower with unfiltered consumption and stock returns (slightly above 1), compared to reported consumption in the same situation (around 5). The former generates $\hat{\psi}$ in the narrow interval $0.18 - 0.22$. A higher first-stage predictability does not ensure positive estimates of the EIS using reported consumption: from -0.07 to 0.11. Counter-intuitively, note that estimations with stocks imply values around -0.05, statistically significant at 1%. Our findings for the reciprocal $1/\psi$ approximately repeat those for quarterly data. Unfiltered consumption gives much more precise results, even with first-stage F-statistics that are noticeably lower than those of Table 2.3. Using stocks, these estimates imply an EIS in the very tight range $0.22-0.25$ – see Table B.6 (appendix). Even with substantially low first-stage F-statistics (lower than 2), note that standard errors are quite aligned across estimators. This suggests that the three methods converge to the same limit distribution, as in the case of conventional first-order asymptotics. For the risk-free, implied EIS estimates from $1/\psi$ are from 0.09 to 0.52, when excluding the less robust TSLS estimator, not that far from results for stocks.

Generally, it seems that the connection between first-stage predictability and more precise estimates is not that relevant with unfiltered consumption. Hence, it could be that a sizeable proportion of the econometric difficulties usually attributed to weak instruments corresponds instead to weaknesses involving the consumption time series[46]. In addition, point estimates of the EIS are generally more close to 1, although still not statistically higher than it. Overall, under unfiltered consumption the improvement is expressive enough both quantitatively (higher and more equalised estimates across estimators, none with the wrong sign) and qualitatively (closer to usual choices of values, applied to macro models).

The next step is to evaluate how unfiltered compares with reported consumption using

[46]We still can not rule out that weak instruments are affecting our estimation since estimates and standard errors – even though more equalised – are still different across estimators (recall that in the absence of weak instruments, limit distributions under the three estimators should be approximately the same). This is especially the case for annual data in Table 2.4, as well as for our estimates of $1/\psi$ in both Table 2.3 and Table 2.4, for which first-stage F-statistics are essentially lower.

Table 2.4: Estimates of the EIS Using K-Class Estimators and Annual Data

Asset	Estimate	Δc_k	K-Class Estimator TSLS	Fuller-K	LIML	1S-F
Risk Free	ψ	Reported	0.112^{***}	0.109^{***}	0.108^{***}	11.836
			(0.105)	(0.111)	(0.113)	
	ψ	Unf-Hom	0.099^{***}	0.091^{***}	0.089^{***}	10.727
			(0.190)	(0.205)	(0.208)	
	ψ	Unf-Het	0.097^{***}	0.090^{***}	0.089^{***}	10.727
			(0.188)	(0.203)	(0.205)	
Stocks	ψ	Reported	-0.049^{***}	-0.061^{***}	-0.065^{***}	5.057
			(0.034)	(0.038)	(0.040)	
	ψ	Unf-Hom	0.186^{***}	0.197^{***}	0.226^{***}	1.307
			(0.081)	(0.088)	(0.108)	
	ψ	Unf-Het	0.184^{***}	0.194^{***}	0.222^{***}	1.312
			(0.080)	(0.086)	(0.105)	
Risk Free	$\frac{1}{\psi}$	Reported	1.364	4.592	9.246	1.893
			(0.724)	(3.364)	(9.634)	
	$\frac{1}{\psi}$	Unf-Hom	0.388	1.883	11.147	1.826
			(0.382)	(1.890)	(25.868)	
	$\frac{1}{\psi}$	Unf-Het	0.394	1.916	11.293	1.822
			(0.387)	(0.475)	(26.190)	
Stock	$\frac{1}{\psi}$	Reported	-6.808^{**}	-11.773^{*}	-15.285^{*}	1.893
			(3.885)	(6.818)	(9.365)	
	$\frac{1}{\psi}$	Unf-Hom	4.056^{*}	4.161	4.415	1.826
			(1.835)	(1.915)	(2.106)	
	$\frac{1}{\psi}$	Unf-Het	4.143^{*}	4.241^{*}	4.502	1.822
			(1.863)	(1.937)	(2.131)	

Notes: Estimates of the EIS and its reciprocal using (2.13) and (2.14) and annual data. Unfiltered consumption extracted relying on the Filter model whose measurement errors are not persistent. All consumption series refer to nondurables and services. We apply the same setting of Yogo (2004), using 3 types of K-class estimators and assuming that errors conditionally follow a martingale difference sequence. When reported consumption is used, asset returns have not been adjusted for time-aggregation. Reported denotes official consumption data from NIPA tables. Unf-Hom and Unf-Het refer to unfiltered consumption, constructed by the homoscedastic and heteroscedastic models, respectively. Standard errors are presented in parentheses. The null that the estimated coefficient equals 1 has been tested: ***, ** and * denote rejection of the null hypothesis at 1, 5 and 10 percent significance levels.

robust inference. For this, we invert Moreira (2003) and T. W. Anderson, Rubin, et al. (1949) test statistics, creating 95 percent weak-identification-robust confidence intervals. Table 2.5 below summarises results. When it comes to quarterly data, the first thing to note is that unfiltered consumption effectively reverts an empty set under the Anderson-Rubin statistic. Using the risk-free, the robust confidence interval with the heteroscedastic model in this case is in line with estimates of L. Epstein and Zin (1991): from -0.07 to 0.87. Based on the S-test of Stock and Wright (2000), Ascari, Magnusson, and Mavroeidis (2016) also found an empty interval, using a baseline Euler Equation as (13) and reported consumption. The S-test is a generalisation of the Anderson-Rubin test to a GMM setting, being not only robust to weak instruments but also to heteroscedasticity of arbitrary form. Ascari, Magnusson, and Mavroeidis (2016) test several Euler Equations, derived from many different assumptions, and confidence intervals similar to ours are only obtained relying on internal habit formation. Since habit formation tends to create inertia[47], implicitly flattening the relationship between consumption and returns, our finding is pertinent to the extent that it brings a similar confidence interval to a much simpler Euler Equation, without creating doubts about how habit formation might be implicitly lowering estimates of the EIS. In addition, as emphasised in Yogo (2004), uninformative robust sets are a natural consequence of a very weak IV setting, so that once more our results suggest that unfiltered consumption significantly improves the identification of the EIS. Even though a little wider, our confidence intervals generated by the conditional likelihood ratio test tell a similar story. Using stocks broadly confirms our results with the risk-free, with the additional benefit that it produces narrower intervals and that the homoscedastic and heteroscedastic models return more similar results.

Table 2.5 also shows that our results for annual data are not as impressive. Unfiltered consumption does increase the upper end of intervals, but confidence sets are generally wider. This result is more evident when stock returns are used. In this case, EIS values along the whole real line are possible. It is difficult to infer the reason for this, even though the small

[47] J. C. Fuhrer (2000).

Table 2.5: Weak-IV-Robust CIs for the EIS

		Quarterly Data		Annual Data	
Asset	Δc_k	Anderson-Rubin	Likelihood Ratio	Anderson-Rubin	Likelihood Ratio
Risk Free	Reported	$\emptyset$	$[-0.136, 0.235]$	$[-0.104, 0.316]$	$[-0.131, 0.341]$
	Unf-Hom	$[-0.319, 1.525]$	$[-0.377, 1.591]$	$[-0.272, 0.442]$	$[-0.357, 0.523]$
	Unf-Het	$[-0.077, 0.867]$	$[-0.307, 1.122]$	$[-0.270, 0.438]$	$[-0.352, 0.516]$
Stocks	Reported	$\emptyset$	$(-\infty, +\infty)$	$[-0.245, 0.019]$	$[-0.199, 0.007]$
	Unf-Hom	$[-0.045, 0.921]$	$[0.014, 0.650]$	$(-\infty, +\infty)$	$(-\infty, +\infty)$
	Unf-Het	$[-0.009, 0.710]$	$[0.023, 0.536]$	$(-\infty, +\infty)$	$(-\infty, +\infty)$

Note: Weak-instrument-robust 95% confidence intervals. Sets constructed by inverting statistics of the Anderson-Rubin and Likelihood Ratio tests. Data used both for reported and unfiltered consumption refer to the consumption of nondurables and services. For quarterly data, we use our quasi-differenced Filter model ($\rho_\xi = 0.06$). For annual data, we use the original version ($\rho_\xi = 0$). Reported denotes official consumption data from NIPA tables. Unf-Hom and Unf-Het refer to unfiltered consumption, constructed by the homoscedastic and heteroscedastic models, respectively.

sample size for annual data could be a possible explanation[48]. Moreover, recall that completely uninformative robust intervals for the EIS are frequent in the literature – Yogo (2004), Ait-Sahalia, Parker, and Yogo (2004), Ascari, Magnusson, and Mavroeidis (2016) and Gomes and Paz (2013). Indeed, our results are not particularly surprising in this respect.

2.3.2 Heteroscedastic Framework

Recall that the econometric approach mentioned above is still valid in a heteroscedastic setting. In this framework, GMM is the efficient method. Thus, we now turn to this estimator, using (2.13). In addition to the conventional two-step GMM (2S-GMM), we also present estimates using the continuously updated GMM estimator (CUE-GMM) – L. P. Hansen, Heaton, and Yaron (1996). The latter is less biased, provides confidence intervals with better cover-

[48]In the appendix, we present results when unfiltered consumption at annual frequency is generated by the Filter model with serially correlated measurement errors instead. General findings are broadly in line with those of Table 2.5, suggesting that our hypothesis for measurement error is not causing that problem. Additionally, there we also re-estimate the Euler equations while further restricting the sample (so that substantially more observations are removed for early years, related to the Filter's training period). Results also do not seem sensitive to such choice.

age rates and performs better under weak instruments – L. P. Hansen, Heaton, and Yaron (1996), Stock, Wright, and Yogo (2002) and W. K. Newey and R. J. Smith (2004).

With (2.13), we conduct estimates of the EIS using 2S-GMM and CUE-GMM, with the risk-free as our measure of returns. In this more general GMM setting, ψ can also be identified in joint estimation using both the risk-free and stocks:

$$\Delta c_{k,t+1} = \tau_{f,t} + \psi r_{f,t+1} + \epsilon_{f,t+1}, \qquad \Delta c_{k,t+1} = \tau_{m,t} + \psi r_{m,t+1} + \epsilon_{m,t+1}, \tag{2.15}$$

where indices f and m denote risk-free and market returns, respectively[49]. The system estimation can improve efficiency from exploiting cross-equation correlations in expectational errors included in both innovations. Additionally, weak-instrument-robust confidence intervals constructed by inverting the K-test statistic – Kleibergen (2005) – are presented. This test is robust to weak identification, as well as to autocorrelation and heteroscedastic error terms. It is similar to the S-test – Stock and Wright (2000) – mentioned above, albeit more computationally involved. We choose the K-test against the S-test based on several factors. First, the former applies in the context of non-linear moment conditions. Second, W. K. Newey and Windmeijer (2009) show that the K-test is valid under many weak moment conditions. Third, Andrews and Stock (2005) and Kleibergen and Mavroeidis (2009) specifically recommend it against available alternatives when dealing with heteroscedasticity of arbitrary form.

Table 2.6 below summarises our estimates for the EIS. First, it generally confirms our previous findings by showing higher point values for unfiltered consumption. Second, results with reported consumption are now more in line with those with unfiltered consumption, relatively to the previous tables. This is especially true for quarterly data, where the former no longer generates negative estimates: from 0.0 to 0.2. In this case, estimates with unfil-

[49] Drift terms must be allowed to differ across equations (given different second-order terms in $\tau_{i,t}$ depending on the asset class i) while slopes are restricted to the same value (EIS). Check the appendix for complete specifications in (15).

tered consumption under the heteroscedastic model, for instance, do not differ much: from 0.0 to 0.5. In addition, results with unfiltered consumption and the risk-free rate (first two columns) are broadly in line with those of Table 2.3, suggesting that the homoscedasticity assumption may be less restrictive when stocks are not considered. This is not the case for reported data, with which estimates of the EIS are higher when allowing for a heteroscedastic environment. Unfiltered consumption constructed from the homoscedastic model once more presents higher estimates of the EIS, even though at the cost of higher standard errors for quarterly data. Results for annual data are a little weaker. Comparatively, point estimates lie closer to zero, both for reported and unfiltered consumption. In spite of that, the joint estimation using the former once again reaches a negative value, significant at 10%. Finally, 95% robust intervals still return uninformative sets, suggesting that identification issues related to weak instruments are possibly more relevant at annual frequency, regardless of the homoscedasticity hypothesis for the errors in the Euler equations[50]. In contrast, we again revert completely uninformative robust sets for quarterly data: $[0.18, 9.64]$ based on the homoscedastic model and $[0.05, 8.91]$ for the heteroscedastic analogue[51].

2.4 EIS, Limited Participation and Unfiltered Consumption

The last section demonstrated how important is to account for the fact that macro consumption series are heavily filtered before release, when estimating the EIS. In this section, we aim to verify whether that is again the case when dealing with other types of consumption data series. We use micro data at the level of the household to construct the consumption growth series used in our estimations. These data are extracted from the CEX, a large-scale survey, designed to represent characteristics of the entire US population. We construct measures

[50]Recall the previous tables and the lower first-stage F-statistics obtained for annual data.

[51]In the appendix, we present similar results for the consumption of nondurables only. We again can revert totally uninformative sets into more plausible ones using unfiltered consumption on quarterly data, albeit those intervals still exhibit negative values: $[-0.35, 0.84]$ relying on the homoscedastic model and $[-0.62, 0.26]$ for the heteroscedastic version.

Table 2.6: Heteroscedasticity-Robust Estimates of the EIS

	Quarterly Data				Annual Data			
Δc_k	Two-Step	CUE	SYS	95% CI	Two-Step	CUE	SYS	95% CI
Reported	0.133	0.189^{**}	0.001	$(-\infty, +\infty)$	0.056	0.022	-0.015^{*}	$(-\infty, +\infty)$
	(0.082)	(0.085)	(0.000)		(0.088)	(0.087)	(0.008)	
Unf-Hom	0.601	0.678	0.007	$[0.182, 9.639]$	0.122	0.136	0.067	$(-\infty, +\infty)$
	(0.523)	(0.525)	(0.006)		(0.142)	(0.142)	(0.045)	
Unf-Het	0.448	0.512	0.006	$[0.053, 8.907]$	0.119	0.133	0.066	$(-\infty, +\infty)$
	(0.374)	(0.377)	(0.006)		(0.141)	(0.141)	(0.045)	

Note: 2S-GMM, CUE-GMM and SYS-GMM estimates of ψ (EIS) using equation (13) with the risk-free rate. "SYS" presents estimates of the same coefficient under the joint estimation (15), where market returns are also used (allowing for different drifts across equations). We present 95% confidence intervals that are robust to both heteroscedasticity and a weak-IV setting. These are constructed by inverting the K-statistic of Kleibergen (2005). Consumption series are relative to nondurables and services. Reported denotes official consumption data from NIPA tables. Unf-Hom and Unf-Het refer to unfiltered consumption, constructed by the homoscedastic and heteroscedastic models, respectively. Standard errors are presented in parentheses. The null that the estimated coefficient equals 0 has been tested using conventional t-statistics: ***, ** and * denote rejection of the null hypothesis at 1, 5 and 10 percent significance levels.

of consumption growth following the same procedures in Vissing-Jorgensen (2002). Our sample goes from 1982 to 2013.

The CEX data is subject to a number of statistical transformations, many of which similar to those applied to NIPA consumption series. To cite a few: topcoding; suppression; reallocation, and; imputation procedures. These are all present in the survey. Further, measurement error is so evident in the data that recently the Bureau of Labor Statistics (BLS) implemented the so-called Gemini project to research and develop a complete redesign of the CEX, addressing measurement error and respondent burden issues. Mechanically, statistical procedures in the CEX produce the same effect on final reported data, lowering its variance and diminishing correlations with different measures of return. We detail the procedures used in the CEX in the appendix, also comparing with those performed on NIPA consumption.

An advantage of using CEX data is that we can separate households – and their corresponding measures of consumption growth – based on their asset-holding status. This is relevant to test the LAMP, since it is possible to obtain different estimates of the EIS conditional on different groups of households and their participation in specific asset markets. As noted by Yogo (2004), it is possible that weak instruments are not explaining the whole story of troublesome estimates of the EIS, usually obtained in the literature. Perhaps limited participation in asset markets is plaguing results due to Euler equations that do not hold for the representative agent, as addressed in Vissing-Jorgensen (2002) and Guvenen (2006). Indeed, estimations that rely on an Euler Equation for some asset, but that use households that do not hold a position in that asset, will likely bias estimates of the EIS downwards[52].

First, let us turn to how we use data from the survey. The CEX interviews more than 7500 households per quarter[53]. Each household is interviewed five times, but only the last four interviews are publicly available. Interviews with the same household occur every three months, when they report consumption for the previous three months. In the last (fifth) interview, households report their financial information. We use this information to separate households according to their asset-holding status. Formally, households report holdings for "stocks, bonds, mutual funds and other securities", "US savings bonds", "savings accounts" and "checking accounts, brokerage accounts, and other similar accounts". We use responses for the first two categories to label households as stock (vs. non-stock) or bonds (vs. non-bond) holders.

We classify households by their asset holding status adopting the same criteria in Vissing-Jorgensen (2002). As noted in that paper, it is not possible to perfectly separate households solely by those two categories. There is some overlap between bond and stock holders, but not between asset and non-asset holders for each type of asset (bonds or stocks). Note that imperfect separation should bias *against* finding different estimates of the EIS across these

[52]See Vissing-Jorgensen (1998).

[53]Before 2000, that number was slightly lower, around 5000 per quarter. The programme contains two components, the Interview Survey and the Diary Survey. Each has its own sample. We compile our data set using the former.

groups. Formally, we refer to households with positive responses to "stocks, bonds, mutual funds and other securities" as stock holders, and to those with positive responses to that same category *or* to "US savings bonds" as bond holders.

Additionally, note that an Euler equation that measures consumption from t to $t+1$ should hold at the beginning of first period. Therefore, we shall split households across groups based on their holdings at the beginning of t. To do so, we use two more items in the survey: one question that asks whether households have more, less or the same amount of the asset, relatively to a year ago; and another one which asks the estimated dollar difference in market value of that asset last month, compared to a year ago last month. As in Vissing-Jorgensen (2002), a household is classified as holding asset class i if it: (i) reports the same amount compared to a year ago, holding a positive position in i when interviewed for the last time; (ii) reports lower holdings of the asset, relative to a year ago, or; (iii) reports an increase in its holdings of the asset, but the dollar difference is less than the current value of holdings. Because some of those questions we use to separate households are no longer available after March 2013, this is the last month of consumption observations in our data set.

Our final sample consists of 90,080 households, spread over the period from 1982 to 2013[54]. Amongst these households, 19.6% are classified as stock holders and 29.1% as bond holders[55]. On average, our sample has 246 households each month, of which 70 are bond holders and 49 are stock holders[56]. See Table 2.2.

Our final data set encompasses semiannual consumption growth rates at monthly frequency. To construct consumption growth observations, CEX expenditure categories are carefully aggregated as to mimic definitions of the NIPA consumption of nondurables and

[54]The CEX data is available beginning in 1980. However, we follow Vissing-Jorgensen (2002) in dropping observations for 1980 and 1981. She argues that the quality of the CEX consumption data is considerably lower for that period.

[55]In addition, 2154 households report an increase in holdings of some asset, but not the current value. We classify them as asset holders for the corresponding category – 1593 as stock holders and 561 as bond holders. In addition, a few households report an increase in holdings that exceeds their response for current values. We consider them as non-asset holders in the corresponding category.

[56]Hence, 197 are non-stock holders and 176 are non-bond holders.

services. We exclude three major categories: health care, education costs and housing expenses (except for housing operations). Cash contributions, personal insurance and pensions are also dropped. Categories in the first group are excluded because they exhibit a substantial durable component. In the second, for the same reason or because their definitions are considered out of scope relatively to NIPA consumption – see Garner et al. (2003, p. 12). Major categories in our consumption measure are food (at and away from home), beverages, apparel, tobacco, public and private transportation (including gasoline), personal care services, housing operations, miscellaneous and utilities. Our definition is broadly in line with that given in Attanasio and Weber (1995), being also similar to the one in Vissing-Jorgensen (2002).

For each household h, its consumption growth rate is:

$$\frac{C_{h,m+6} + C_{h,m+7} + C_{h,m+8} + C_{h,m+9} + C_{h,m+10} + C_{h,m+11}}{C_{h,m} + C_{h,m+1} + C_{h,m+2} + C_{h,m+3} + C_{h,m+4} + C_{h,m+5}}.$$

As in Vissing-Jorgensen (2002), the aggregate consumption growth observation is the average of this ratio in the cross section of households of the same group. Since consumption growth is semiannual, groups are classified based on their holdings at the beginning of the relevant period in the Euler equation – i.e., m. We refer to Vissing-Jorgensen (2002) for a formal discussion on how averaging households in the cross section of groups can generate consistent estimates of the EIS in this framework. In addition, closely following that paper, we drop: (i) extreme outliers (observations for which the consumption growth ratio is higher than 5 or less than 0.2); (ii) households that report a change in the age of the household head between two subsequent interviews different from zero or one; (iii) households living in student housing, and; (iv) non-urban households. To construct the semiannual consumption growth ratio above, we need consumption data for all interviews, 2 to 5. Therefore, we also drop households for which any of these interviews are missing. As our last step, we deflate nominal consumption growth observations by the urban CPI for nondurable goods.

We later use the Filter model on the consumption growth series that corresponds to that final sample of 90,080 households. There may be reasons to be sceptical about this, arguing that statistical procedures applied on CEX data probably take those households dropped from our sample into account. Nonetheless, very similar results are found when estimations in this section are repeated applying and calibrating the model for the complete sample (without dropping households), but imposing the calibration to our final data set (which excludes them). Results are also maintained when we calibrate the model *and* estimate the EIS based on the complete sample. Since consumption growth is semiannual but at the monthly frequency, there is an overlap of five months between observations. To calibrate the model, we equalise the scale of consumption growth to its frequency (as in the previous estimations), to then transform it back into semiannual. For a more complete discussion on how we apply the Filter model on the CEX data and different groups of asset holders, check the appendix.

Two types of returns are used when estimating our Euler equations in this section. When differentiating between stock holders vs. non-stock holders, we use the value-weighted return from NYSE, NASDAQ and AMEX. When applying to bond vs. non-bond holders, we use T-bill returns. An important issue relates to how we compute the relevant asset return used in the estimations when the consumption growth data is semiannual. We follow Vissing-Jorgensen (2002), using the middle six months from $(1 + R_m)$ to $(1 + R_{m+10})$: $(1+R_{m+2})(1+R_{m+3})...(1+R_{m+6})(1+R_{m+7})$. In addition, since $C_{h,m}$ is relevant in the consumption growth measure, it follows naturally that lagged instruments are constructed based on $(1 + R_{m-1})(1 + R_{m-2})...(1 + R_{m-5})(1 + R_{m-6})$, also using six months. When we estimate the Euler equations with unfiltered CEX consumption, we conduct similar adjustments as those of the last section on these return series. We better detail them in the appendix.

We use three instrument sets for the log stock return or the log T-bill return in the Euler equations: (i) dividend-price ratio; (ii) dividend-price ratio, lagged stock returns and lagged T-bill returns, and; (iii) dividend-price ratio, lagged corporate bond default premium and lagged government bond horizon premium. The last two take the form $\frac{1+R_t^{\text{long-term corporate bonds}}}{1+R_t^{\text{long-term government bonds}}}$ and $\frac{1+R_t^{\text{long-term government bonds}}}{1+R_t^{\text{short-term government bonds}}}$, respectively. We once more refer to the appendix for more details on these variables. Lastly, returns are deflated by using the urban CPI for total consumption.

As in the last section, we estimate log-linearised Euler equations. For stock holders, for example, the econometric approach we follow is:

$$\frac{1}{H_t^s}\sum_{h=1}^{H_t^s} \Delta lnC_{t+1}^{h,s} = \psi^s ln(1+R_{s,t}) + \frac{1}{H_t^s}\sum_{h=1}^{H_t^s} \Delta ln(\text{family size})_{t+1}^{h,s} + \alpha_1^s D_2 + ... + \alpha_{12}^s D_{12} + u_{t+1}^s, \tag{2.16}$$

where, as before, ψ^s is the EIS for stock holders, D_m are seasonal dummies and H_t^s denotes the number of consumption growth observations for stock holders at time t. Euler equation (16) assumes that seasonality and the family size are multiplicative factors in the utility function[57]. These two variables are included in all the three aforementioned instrument sets[58]. Equation (2.16) holds under the Epstein-Zin framework of last section, as well as under CRRA preferences[59].

2.4.1 Results

Results for two samples are presented. The first encompasses data from 1982 to 1996, the same used in Vissing-Jorgensen (2002). The second uses all the available data, from 1982 to 2013. We do so due to a change in methodology and in the sampling frame around 1996,

[57]Formally, the family size variable is defined as the change in the log average family size for the last two interviews (4 and 5), compared to the first two interviews (2 and 3).

[58]Therefore, it is assumed that family size controls and seasonality factors are exogenous in our estimations.

[59]Regardless of the assumption for the utility function, generally α_m^s involves conditional variances (covariances) of (between) log consumption growth and log returns. In the case of Epstein-Zin preferences, there are also conditional second-order terms relative to wealth returns, based on the total portfolio of households. If some of those conditional second-order terms are not constant, stochastic terms enter u_{t+1}^s – which already included expectational and measurement errors (present in the consumption data). These stochastic terms do not imply inconsistent estimates, as long as they are uncorrelated with instruments used. See Vissing-Jorgensen (1998) and Vissing-Jorgensen (2002) for a formal treatment. Lastly, note that (2.16) can be estimated by instrumental variables methods even when using unfiltered consumption. Since returns are assumed uncorrelated with the measurement error, auto-correlation in the latter does not invalidate lags of the former as instruments.

when the CEX was redesigned[60].

We begin with results for the period 1982-1996, exhibited in Table 2.7[61]. First, note that *all* estimates with reported consumption are negative, many of them statistically significant. In contrast, unfiltered consumption reverts these into more sensible estimates, when considering the Euler Equations for stock and bond holders. It suggests an EIS from 0 to 0.3 for the former, and from 0.4 to 1 for the latter group. Differences between unfiltered and reported consumption are less substantial when considering non-asset holders, but the former still provides estimates of the EIS that are slightly less negative. This is also generally the case when the Euler equation for all households is estimated. Importantly, robust intervals lean towards positive values with unfiltered consumption. These are also substantially narrower when estimating the Euler equation for stock holders, suggesting that weak instruments affect these estimations to a lesser extent. For bond holders, results are more uncertain. Although point estimates seem more precisely estimated with unfiltered consumption, robust intervals are considerably wider.

Unfiltered consumption seems to magnify differences in terms of the EIS between asset and non-asset holders. Such result is consistent with Table 2.2, which shows, for instance, that unfiltering the CEX data introduces more mean reversion and considerably more volatility in the consumption series of bond and stock holders, compared to non-bond and non-stock holders. Robust sets with unfiltered consumption suggest that the EIS is not above 0.6 for stock holders, nor above 0.2 for non-stock holders. Nonetheless, there is considerable overlap between the intervals. Therefore, there seems to be little evidence favouring the limited asset market participation theory in our results. This finding contrasts with Vissing-

[60] However, we calibrate the model based on the entire sample (1982-2013). We need enough observations to calculate the long-run standard deviations of the series. Since these use a horizon of 6 years, calibrating based on the period 1982-1996 gives weaker results.

[61] The use of semiannual consumption growth data at monthly frequency generates overlapping observations for two subsequent months of data. It follows that an MA(5) process enters the error term in (2.16). Therefore, we use Two-Step GMM with a heteroscedasticity and autocorrelation-consistent (HAC) estimator for the covariance matrix. Main results of this section do not change when using CUE-GMM, which is more robust to the presence of weak-instruments.

Jorgensen (2002), who finds substantial differences across those groups for the same period. She does not apply weak-IV-robust methods, though[62].

Table 2.8 presents results for the entire sample, from 1982 to 2013. Again, reported consumption produces negatives estimates of the EIS for *all* groups and instrument sets applied. In contrast, estimations with unfiltered consumption return positive estimates in *all but one* of the cases tested. For all households, the EIS is estimated from 0.05 to 0.1 using stock returns. These values shift to 0.4 to 1.2 with risk-free returns. Note that reported consumption provides counter-intuitive estimates in those cases, from -0.1 to -1.3, depending on the type of return and the instrument set used.

In contrast to Table 2.7, using unfiltered consumption for the entire sample generates estimates of the EIS that are quite alike, comparing asset and non-asset holders. For stock holders, for example, we estimate a coefficient in the narrow interval from 0 to 0.1, not distant from estimates for non-stock holders, around 0.15. Additionally, robust intervals are also similar. As in the previous table, these sets generally shift from showing negative to showing positive numbers, as we replace reported with unfiltered consumption. Once more, there seems to be more uncertainty involving estimations for bond and non-bond holders, as, for instance, robust intervals are again considerably wider. In spite of that, higher point estimates are still obtained with unfiltered consumption. For instrument sets II and III, it produces estimates around 0.4, the lower bound of results with analogous estimations in the previous table.

Generally, estimations in Table 2.8 once more produce limited evidence that the EIS differs substantially between asset and non-asset holders. This is the case even with unfiltered consumption. Recall that, in the previous section, we concluded that unfiltered consumption offered more reliable estimations, which also seemed less plagued by the presence of weak instruments. We then mentioned the possibility that commonly distorted estimates of

[62] Although she does not apply robust methods, our point estimates are considerably distinct from those reported in Vissing-Jorgensen (2002) (for the same period). The data have been revised several times since then, so that these revisions may explain the differences.

Table 2.7: Estimates of the EIS – CEX Data: 1982 to 1996

	A. Estimation with Stocks				B. Estimation with Treasury Bills			
Δc_k	Households	I	II	III	Households	I	II	III
Reported	All	-0.245^{*}	-0.243^{**}	-0.425^{***}	All	-1.065^{**}	-0.933^{*}	-1.680^{***}
		(0.145)	(0.122)	(0.181)		(0.505)	(0.503)	(0.491)
		$[-0.894, -0.101]$	$[-1.048, -0.087]$	$[-1.045, -0.248]$		$[-3.044, 0.821]$	$[-2.547, 1.573]$	$[-4.141, 0.667]$
	Stock Holders	-0.007	-0.224	-0.341	Bond Holders	-0.208	-0.278	-1.064
		(0.197)	(0.172)	(0.228)		(0.690)	(0.693)	(0.676)
		$[-0.703, 0.221]$	$[-0.943, 0.033]$	$[-1.496, -0.085]$		$[-2.969, 3.481]$	$[-2.867, 5.012]$	$[-6.079, 3.753]$
	Non-Stock Holders	-0.295^{*}	-0.287^{**}	-0.513^{**}	Non-Bond Holders	-1.293^{**}	-1.219	-1.966^{***}
		(0.155)	(0.131)	(0.203)		(0.510)	(0.496)	(0.471)
		$[-0.760, -0.130]$	$[-0.807, -0.112]$	$[-0.816, -0.349]$		$[-3.332, 0.604]$	$[-3.504, 1.523]$	$[-4.947, 0.322]$
Unfiltered	All	-0.066	-0.189	-0.135	All	-0.328	-1.017	-0.855
		(0.418)	(0.226)	(0.189)		(2.076)	(1.714)	(1.091)
		$[-0.286, 0.050]$	$[-0.454, -0.102]$	$[-0.345, -0.051]$		$[-4.930, 4.314]$	$[-14.154, 12.089]$	$[-16.033, 8.073]$
	Stock Holders	0.151	0.068	0.323	Bond Holders	1.070	0.417	0.765
		(0.609)	(0.349)	(0.363)		(2.444)	(2.228)	(1.745)
		$[-0.038, 0.311]$	$[-0.338, 0.290]$	$[0.033, 0.663]$		$[-3.719, 8.085]$	$[-19.829, 36.569]$	$[-11.043, 16.187]$
	Non-Stock Holders	-0.091	-0.228	-0.222	Non-Bond Holders	-0.590	-1.244	-1.404
		(0.501)	(0.274)	(0.222)		(2.600)	(2.199)	(1.377)
		$[-0.368, 0.195]$	$[-0.426, 0.235]$	$[-0.391, 0.100]$		$[-5.682, 5.266]$	$[-40.931, 32.163]$	$[-28.711, 8.518]$

Notes: Estimates of the EIS using Euler equation (2.16). The sample encompasses semi-annual consumption growth observations at monthly frequency, from 1982 to 1996. Unfiltered consumption is extracted relying on the quasi-differenced Filter model whose measurement errors are serially correlated. Here we assume that government statisticians filter the data based on our final sample – in which some households are dropped based on conditions described in the main text. Reported uses official CEX data. Unfiltered consumption growth is constructed from the heteroscedastic model. For this case, asset returns are adjusted for time-aggregation issues for any group of asset holders – see appendix. Instrument set I includes the log dividend-price ratio. Set II adds the lagged log real value-weighted return (from NYSE, NASDAQ and AMEX) and the lagged log real T-bill return. Set III replaces the last two by the lagged bond default premium and the lagged bond horizon premium. All these sets include the family size and seasonal controls as instruments (so that these are assumed exogenous). Standard errors are presented in parentheses. 95% confidence intervals that are robust to both heteroscedasticity and a weak-IV setting are shown in brackets. We construct these intervals by inverting the K-test statistic in Kleibergen (2005). The null that the estimated coefficient equals 0 has been tested: ***, ** and * denote rejection of the null hypothesis at 1, 5 and 10 percent significance levels.

Table 2.8: Estimates of the EIS – CEX Data: 1982 to 2013

	A. Estimation with Stocks				B. Estimation with Treasury Bills			
Δc_k	Households	I	II	III	Households	I	II	III
Reported	All	-0.163^{**}	-0.108^{*}	-0.161^{***}	All	-1.340^{**}	-0.389	-0.578^{*}
		(0.065)	(0.060)	(0.056)		(0.519)	(0.352)	(0.324)
		$[-0.246, -0.098]$	$[-0.158, -0.003]$	$[-0.249, -0.094]$		$[-7.884, 2.655]$	$[-2.804, 7.158]$	$[-2.594, 7.302]$
	Stock Holders	-0.107	-0.117	-0.175^{**}	Bond Holders	-0.916	-0.349	-0.532
		(0.103)	(0.103)	(0.078)		(0.646)	(0.475)	(0.408)
		$[-0.350, 0.070]$	$[-0.393, -0.039]$	$[-0.390, -0.045]$		$[-8.532, 9.719]$	$[-8.391, 7.065]$	$[-7.732, 7.839]$
	Non-Stock Holders	-0.161^{***}	-0.101^{**}	-0.149^{***}	Non-Bond Holders	-1.362^{***}	-0.450	-0.687^{**}
		(0.063)	(0.056)	(0.056)		(0.496)	(0.335)	(0.319)
		$[-0.305, -0.084]$	$[-0.201, 0.048]$	$[-0.292, -0.069]$		$[-9.163, 2.868]$	$[-6.837, 4.205]$	$[-3.924, 5.542]$
Unfiltered	All	0.116	0.044	0.065	All	1.181	0.348	0.418
		(0.241)	(0.205)	(0.210)		(2.438)	(1.287)	(0.709)
		$[-0.022, 0.284]$	$[-0.112, 0.186]$	$[-0.094, 0.218]$		$[-0.514, 9.875]$	$[-2.057, 13.297]$	$[-0.628, 9.138]$
	Stock Holders	0.000	0.014	0.035	Bond Holders	-0.337	0.318	0.423
		(0.330)	(0.294)	(0.325)		(2.759)	(1.570)	(1.083)
		$[-0.145, 0.145]$	$[-0.128, 0.188]$	$[-0.107, 0.196]$		$[-6.874, 45.645]$	$[-39.949, 41.042]$	$[-28.484, 47.735]$
	Non-Stock Holders	0.173	0.107	0.143	Non-Bond Holders	2.154	0.967	0.566
		(0.263)	(0.223)	(0.239)		(2.774)	(1.479)	(0.832)
		$[0.029, 0.369]$	$[-0.025, 0.269]$	$[0.004, 0.331]$		$[-0.072, 31.714]$	$[-1.769, 31.042]$	$[-0.658, 26.812]$

Notes: Estimates of the EIS using Euler equation (2.16). Our sample encompasses semi-annual consumption growth observations at monthly frequency, from 1982 to 2013. Unfiltered consumption is extracted relying on the quasi-differenced Filter model whose measurement errors are serially correlated. Here we assume that government statisticians filter the data based on our final sample – in which some households are dropped based on conditions described in the main text. Reported uses official CEX data. Unfiltered consumption growth is constructed from the heteroscedastic model. For this case, asset returns are adjusted for time-aggregation issues for any group of asset holders – see appendix. Instrument set I includes the log dividend-price ratio. Set II adds the lagged log real value-weighted return (from NYSE, NASDAQ and AMEX) and the lagged log real T-bill return. Set III replaces the last two by the lagged bond default premium and the lagged bond horizon premium. All these sets include the family size and seasonal controls as instruments (so that these are assumed exogenous). Standard errors are presented in parentheses. 95% confidence intervals that are robust to both heteroscedasticity and a weak-IV setting are shown in brackets. We construct these intervals by inverting the K-test statistic in Kleibergen (2005). The null that the estimated coefficient equals 0 has been tested: ***, ** and * denote rejection of the null hypothesis at 1, 5 and 10 percent significance levels.

the EIS may be a consequence of an Euler Equation that does not hold for the representative agent because of limited participation in asset markets. First, our estimations with the CEX data indicate that it is not the case. Second, even if it were, accounting for the fact that reported data are statistically treated before release can once again produce more precise estimates of the EIS.

2.5 Conclusion

The empirical literature on the elasticity of intertemporal substitution (EIS) is vast, with considerable variation in estimates depending on the specification, econometric method, sample, and degree of market participation. However, a critical and often overlooked issue lies in the construction of consumption data itself. Most reported consumption series are filtered and interpolated before release, introducing artefacts that can distort identification and weaken the reliability of structural estimates.

This paper addresses that gap by proposing a generalised framework for EIS estimation. Building on the Filter model of Kroencke (2017), we introduce an econometric structure that allows for persistent distortions in the observed series, effectively 'unfiltering' consumption data prior to estimation. This adjustment is non-trivial and yields a flexible methodology that can be applied to different types of data, at varying frequencies, and under different asset market participation assumptions.

Our framework is employed in two distinct settings. First, we construct unfiltered macro series from national accounts data and estimate the EIS using the instrumental variable approach of Yogo (2004), showing that unfiltered consumption stabilises estimates and narrows confidence intervals, even in the presence of weak instruments. Second, we apply the same methodology to disaggregated micro survey data (CEX), distinguishing between asset and non-asset holders. In both cases, the flexibility of our model — particularly its allowance for serial correlation in the measurement error — proves essential for producing plausible and consistent estimates.

Across settings, the use of unfiltered data substantially improves identification. With macro data, weak-IV-robust intervals become informative, and estimates concentrate in a more plausible range. With micro data, the EIS estimates shift upward and become more stable, especially for stock and bond holders, compared to the negative or near-zero values typically found using reported consumption.

Chapter 3

Identification of the Phillips Curve using Sectoral Data

Co-authored with Carlos Carvalho[1]

Abstract: *At the centre of monetary policy mechanism, the Phillips Curve plays a crucial role in macroeconomic models. Yet, its theoretical importance contrasts with a large empirical debate on its stability, robustness, and even existence. This paper takes seriously the theoretical implications of heterogeneity in price stickiness for the estimation of the Phillips Curve. The novel method generates positive, sizeable and stable slope coefficients across different econometric settings, producing degrees of stickiness broadly aligned with the micro evidence, both regarding the entire economy and the cross section of sectors. The degree of robustness exhibited by our estimations contrasts with the empirical literature, which typically struggles with minor changes in the econometric setting.*

[1]We are grateful for the insightful feedback provided by participants at the Econometric Society World Congress, with special thanks to James Stock and Randal Verbrugge. For valuable comments and discussions, we also thank João Victor Issler, Leonardo Rezende, Marcelo Medeiros, Marco Bonomo, Michael McLeay, Ricardo Reis, and Silvana Tenreyro.

3.1 Introduction

The relationship between inflation and economic slack, measured by the Phillips curve, is essential in macroeconomics, at the center of monetary policy transmission mechanism. Laid out back in the 80's and 90's, the New-Keynesian Phillips Curve (henceforth, NKPC) is is currently the most widespread theory in that design. However, its early success and widespread adoption contrast with a series of empirical difficulties, with a pronounced debate on estimated coefficients (Nason and G. W. Smith (2008a), Nason and G. W. Smith (2008b) and Mavroeidis, Plagborg-Møller, and H. Stock (2014)), stability (Stock and Watson (2007), Albuquerque and Baumann (2017), Luengo-Prado, Rao, and Sheremirov (2018), Jorda and Nechio (2018) and Galí and Gambetti (2019)), robustness (Dufour, Khalaf, and Kichian (2010)), or even existence (Hooper, Mishkin, and Sufi (2019)).

Most of the empirical literature focuses on specifications justified on the basis of the simplest New-Keynesian models in Woodford (2003) and Galí (2015), or slight variations that induce a backward-looking component through some form of price-indexation (Christiano, Eichenbaum, and Evans (2005)), staggered wage contracts (J. Fuhrer and Moore (1995)), or some backward-looking rule of thumb (Gali and Gertler (1999) and Galí, Gertler, and López-Salido (2003)). Extensions of the basic theory, however, imply that the NKPC should have additional terms. For example, models with capital accumulation would imply an investment-gap term (e.g., Michael Woodford (2005) and Carvalho and Nechio (2016)), while terms related to trade appear in the NKPC of some open-economy models (e.g., Monacelli (2007) and Zaniboni (2008)).

In this paper, we take seriously the implications of heterogeneity in price stickiness to address empirical difficulties verified by the literature on the NKPC. More precisely, we propose a novel estimation method that relies on a system comprised of an aggregate and sectoral NKPCs, taken from a multi-sector heterogeneous model. An advantage of this framework is that it brings plenty of additional structure from theory, which we exploit through a rich set of cross-equation restrictions in our estimations.

There are more essentially two ways through which a multi-sector heterogeneous economy improves the identification of the NKPC. First, cross-sector variation can potentially mitigate a form of simultaneity bias affecting the equation, which pushes its slope towards zero, or even counter-intuitive negative numbers, caused by the endogenous response of monetary policy.

Fitzgerald, Nicolini, et al. (2014) and McLeay and Tenreyro (2019) show that result for regional data, albeit in the absence of stickiness heterogeneity. We not only demonstrate that the same implications are carried over to the multi-sector case, but also that they are magnified once we also allow for heterogeneity.

Second, price-setting heterogeneity can potentially enhance the reliability of estimations due to the presence of strategic complementarities across sectors. A well known difficulty of the empirical literature on the NKPC is that, even when estimates of the slope are positive, the implied degree of stickiness in the economy is too high, incompatible with the evidence brought by studies using micro (disaggregated) data – e.g., in Bils and Klenow (2004), Nakamura and Steinsson (2008) and Nakamura and Steinsson (2013). Under price-setting heterogeneity, the more sticky sectors exert a disproportionate effect on the aggregate price level – Woodford (2003, Chapter 3). As a result, everything else constant, the price dynamics of the economy is more staggered under the heterogeneous model compared to its homogeneous counterpart, as shown by Carvalho (2006). It follows that the economy behaves *as if* it featured more nominal rigidities, allowing a lower and more sensible degree of stickiness to be recovered from the data.

Some articles attest to the importance of heterogeneity in reconciling macro and micro estimates of degrees of stickiness – e.g., Cagliarini, Robinson, and Tran (2011), for Australia, and Imbs, Jondeau, and Pelgrin (2011), for France. However, contrasting with our paper, the discussion of the former is mostly based on simulated data, while the empirical strategy of the latter focuses on the dynamics of sectoral inflation (aggregate inflation is simulated on the basis of estimates obtained for the sectors). Furthermore, unlike our article, these authors do not estimate Phillips curves for the US economy.

Our method allows us to separate out the effect of each of the two mechanisms – cross-sector variation and heterogeneity – over the estimations. With our baseline model, when we exploit both, it delivers a sizable, stable and positive slope. This coefficient is significant at 1% for *all* of the cases tested in this paper, including a series of robustness checks. Furthermore, the procedure generates a degree of stickiness that approaches levels evidenced by studies using disaggregated data. The model implies an aggregate expected duration of price spells roughly from 7.1 to 8.3 months, while the micro evidence usually delivers estimates in the interval from 4 to 9 months[2]. We also find that estimates of the slope and stickiness behave as predicted by theory as we introduce more or less strategic complementarities in price setting. When it comes to the cross section of sectors, estimated degrees of stickiness are broadly in line with the micro evidence, presenting quite narrow standard errors. For example, the method implies high correlations (around 0.7) between estimated sectoral coefficients that govern stickiness in the model and their micro analogues, mapped from disaggregated data.

The degree of robustness exhibited by our estimations with the heterogeneous model is at odds with the rest of the literature, which, for instance, typically struggles with minor changes in the econometric setting. We test estimator uncertainties perturbing the parametric setting, as the number and which parameters are estimated, calibrated values for those that are not, the sample, instruments and starting values in the estimation algorithm. The model we propose, although not tailored to solve specific problems, exhibits little to no change in terms of performance under those exercises.

Implications for the degree of stickiness are essentially maintained once we switch off heterogeneity. The slope is reduced by half, but it is still stable, significant and substantially

[2]Values obtained with micro data can vary substantially depending on whether sales are included or not. See Bils and Klenow (2004), Nakamura and Steinsson (2008) and Nakamura and Steinsson (2013), for useful references.

larger compared to the case when we switch off both heterogeneity and cross-sector variation.

In the absence of those two mechanisms, the estimations follow simple approaches in the literature, with a specification based on the canonical New-Keynesian model in textbooks. In this case, it also returns the same puzzling results, namely a counter-intuitive slope, which is roughly estimated at zero, and a degree of stickiness too high and inconsistent with the evidence found for micro data.

Among existing studies, the closest in spirit to our approach are McLeay and Tenreyro (2019) and Hazell et al. (2022). Both papers seek to improve identification of the Phillips curve by leveraging cross-sectional variation: McLeay and Tenreyro (2019) use regional variation across U.S. areas, while Hazell et al. (2022) construct a panel of U.S. states using newly built CPI indices for nontradeables. In contrast, we exploit sectoral heterogeneity in price stickiness, which has stronger micro foundations and can be more directly mapped to structural model parameters. Moreover, while both studies rely on reduced-form regressions, our estimation is based on a fully specified structural system that imposes cross-equation restrictions from a multi-sector New Keynesian model. This distinction allows us not only to improve identification but also to isolate the contribution of sector-level heterogeneity to the slope of the aggregate Phillips curve, under well-defined theoretical assumptions.

This paper differs from most part of the empirical literature in using the output gap as the slack variable of the NKPC. In a seminal paper, Gali and Gertler (1999) showed that estimations of the NKPC using the output gap generally return negative values for the slope, and suggested the use of an alternative specification where that slack variable is replaced with a proxy for marginal costs (generally the labour share). From there on, the vast majority of articles follow that workaround – e.g., Sbordone (2002), Galí, Gertler, and López-Salido (2003) and Cogley and Sbordone (2008), to cite a few. However, such alternative also faces skepticism[3]. Results of our paper suggest that it is perfectly possible to obtain reliable and sensible

[3]Some articles point to a limited empirical evidence that proxies for marginal costs can add any information

estimations of the NKPC relying on the straightforward slack variable, i.e., the output gap.

We discuss the two mechanisms through which our model potentially improve estimations of the Phillips curve in the next section. In Sections 3.3 and 3.4, we discuss the model and the empirical approach, respectively. Section 3.5 briefly describes the data set. In Section 3.6, we present our main findings, also discussing robustness checks. Finally, Section 3.7 concludes.

3.2 Implications of Heterogeneity in Price Stickiness for Estimation of the NKPC

This section illustrates the two main mechanisms through which an estimation strategy that accounts for heterogeneity in price stickiness across sectors may improve empirical results for the NKPC. First, heterogeneity and cross-sector variation create a tension between stabilization objectives of monetary policy, what, to some degree, alleviates a form of simultaneity bias in the estimations, caused by the response of monetary policy. Second, under heterogeneity, strategic interactions between sectors give rise to an interesting theoretical channel that helps us to obtain lower and more sensible estimates of the degree of price stickiness in the economy. We begin by addressing the former point below.

3.2.1 Optimal Policy: Ameliorating the Simultaneity Bias

Contemporary findings suggest that the Phillips curve has considerably flattened in the more recent decades, what has induced many authors to question whether the curve is actually dead – e.g., Stock and Watson (2007), Kuester, Müller, and Stölting (2009), Kleibergen and Mavroeidis (2009) and Galí and Gambetti (2019). However, except for the last, these papers

regarding the dynamics of inflation – e.g., Rudd and Whelan (2005), Rudd and Whelan (2007) and Mazumder (2010).

generally emphasize drawbacks of the econometric setting rather than giving explanations based on macroeconomic theory.

A more recent strand of the literature is convenient in linking difficulties to empirically identify the Phillips curve with theoretical reasons why these complications may arise in the data. Those papers are built on the argument that identification of the Phillips curve is blurred at the aggregate level because monetary policy endogenously attempts to offset sources of exogenous variation that would otherwise help to identify the equation.

Pursuing a goal of minimizing welfare losses, what involves deviations of inflation from target and output from potential, a Central Bank will tend to act to increase inflation when output is below potential, inducing a negative correlation between both stabilization objectives and biasing downwards the slope of the Phillips curve. For example, as illustrated by McLeay and Tenreyro (2019), this point can be seen using the textbook model described in Galí (2015). There, under discretion, the optimal policy that minimizes a loss function subject to the NKPC takes the form:

$$y_t = -\frac{\kappa}{\vartheta}\pi_t, \tag{3.1}$$

where y_t and π_t denote the welfare-relevant output gap and inflation, respectively, κ is the slope of the underlying NKPC and ϑ a function of structural parameters of the model[4].

From (3.1), researchers would not be estimating a pure and steeper curve in the data, but merely obtaining an intersection of the Phillips curve and a monetary policy targeting rule, a classic situation of simultaneity bias. A straightforward way to see this comes from the absence of trade-off between stabilization of inflation and the welfare-relevant output gap in simple New-Keynesian models without real imperfections, referred to in Blanchard and Galí (2007) as "divine coincidence". In this case, perfect stabilization of both objectives is

[4]For expository reasons, equation (3.1) implicitly assumes an efficient steady-state. If a steady-state distortion is sufficiently small, the same negative relationship between output and inflation follows, but a constant term would also appear in (3.1).

possible from (3.1), completely blurring the identification of the Phillips curve in the data. As a result, its slope would be estimated at *zero* under simple linear regression methods.

These econometric issues motivate a number of articles that attempt to circumvent that sort of bias by seeking sources of exogenous variation to identify the Phillips relationship, albeit none of which exploiting heterogeneity in price setting. Jorda and Nechio (2018) rely on an estimation strategy that benefits from the so-called monetary policy *trilemma* in a multi-country framework. These authors aim to dissolve the simultaneity bias using the fact that, under free mobility of capital, countries that rely on a fixed exchange rate regime can not conduct an independent monetary policy. Fitzgerald, Nicolini, et al. (2014), McLeay and Tenreyro (2019) and Hooper, Mishkin, and Sufi (2019) rely on a multi-region setting to argue how the use of these sorts of data can help to mitigate the bias to the extent that they are not influenced by policy. Indeed, if the Central Bank does not respond to idiosyncratic shocks in the regions, McLeay and Tenreyro (2019) show that regional Phillips curves can be useful to identify the slope of the aggregate equation, particularly when the variance of these shocks is sufficiently high.

Alternatively, one can think of the simultaneity bias caused by endogenous policy as in a textbook approach of simultaneous equations of demand and supply. Note that aggregate demand shocks tend to induce a positive relationship between inflation and the output gap, whereas cost-push shocks tend to produce the opposite effect. If demand shocks are completely offset by policy, but cost-push shocks accommodated, estimations of the Phillips curve will inherit properties of the latter, returning an unrealistic negative slope that follows the relationship depicted in (3.1) even when the true coefficient is positive and sizable. Hence, controlling for supply shocks is important to recover the Phillips curve in the data.

This suggests why multi-sector New-Keynesian models can prove handy to mitigate the impact of the endogenous response of policy. Aoki (2001), for example, constructs a simple New-Keynesian framework comprised of a flexible-price sector and a sticky-price sector. He shows how relative prices, terms that measure misalignments of prices between sectors, play an important role in these type of models. First, they arise as an additional objective of policy. Second, they enter as a shift component in the NKPC, what justifies why these terms are commonly used as a proxy for supply shocks – e.g., relative prices of food and energy. However, relative prices may also reflect other factors, as idiosyncratic demand shocks and elasticities of substitution between goods.

Regardless of their underlying transmission mechanism, the presence of relative prices in models with price-setting heterogeneity is not only theoretically meaningful, but also econometrically useful. For example, a second-order approximation of the welfare-relevant loss function associated with the model we later use in this paper takes the form[5]:

$$W \approx -\frac{1}{2}\mathbb{E}\sum_{t=0}^{\infty}\beta^t L_t,$$

where:

$$L_t \equiv \left\{(\sigma+\varphi^{-1})y_t^2 + (1+\varphi^{-1}\epsilon)\sum_{k=1}^{K}\eta_k\frac{1}{\kappa_k}\pi_{k,t}^2 + (1+\varphi^{-1})\sum_{k=1}^{K}\eta_k(p_{k,t}-p_t)^2\right\}. \tag{3.2}$$

Lowercase variables denote log-deviations from a zero-inflation steady state, $\pi_{k,t}$ denotes the inflation rate in sector $k \in \{1,...,K\}$, η_k the weight of that sector in aggregate expenditures and $(p_{k,t}-p_t)$ the *relative price* of that sector, the difference between its price index and the aggregate price level. As seen later, κ_k denotes the slope of the NKPC of sector k. Additionally, σ represents the inverse of the elasticity of intertemporal substitution (EIS) in consumption, φ is the Frisch elasticity of labor supply and $\epsilon > 1$ is the elasticity of substitution between goods.

As in a homogeneous model, monetary policy should minimize fluctuations in output, justifying the first term in (3.2). With price-setting heterogeneity across sectors, price (and

[5]We present derivations in the appendix. For expository reasons, (3.2) corresponds to the loss function of a forward-looking model, i.e., where price indexation is not allowed. We later discuss implications for the more generalized case. In addition, to simplify the algebra, we assume that an appropriate optimal subsidy to employment is in place, τ. This will neutralize the mark-up distortion generated by the market power of firms in steady state.

output) dispersion in each sector explain the second term, here a function of inflation in the sectors, rather than aggregate inflation, as in the homogeneous case. Undesirable fluctuations in hours worked, mapped through sectoral demand functions, give rise to functions of sectoral relative prices. These arise specifically from the presence of heterogeneity and appear in the last term of the loss function, but do not exist in the textbook model used in McLeay and Tenreyro (2019), for example[6].

The presence of the last term in (3.2) generally creates an inability to satisfy all of these objectives simultaneously – see Woodford (2003, Chapter 6). Consequently, it now produces a tension between stabilization of aggregate inflation and output. Such trade-off, not featured in the standard models used in the empirical literature on the NKPC, is econometrically important to the extent that it attenuates the endogenous response of monetary policy at the aggregate level. Thus, it also reduces some of the simultaneity bias affecting the estimations of the aggregate NKPC.

We fix $K = 2$ to illustrate the analogue of (3.1) for the heterogeneous model. In this case and under discretion, it is possible to show that the minimization of (3.2) subject to an aggregate and two sectoral NKPCs produces:

$$y_t = -\,\lambda_\pi \times (\eta_1 \kappa_1^{-1} \pi_{1,t} + \eta_2 \kappa_2^{-1} \pi_{2,t}) + \Lambda_t(\eta_k, \kappa_k, \zeta_{k,t}, \epsilon, \varphi), \tag{3.3}$$

where:

$$\lambda_\pi \equiv \epsilon[\eta_1 \kappa_1 + \eta_2 \kappa_2 + 2\eta_1 \eta_2 (\kappa_1 + \kappa_2)\Theta^{-1}], \qquad \Theta \equiv \frac{\sigma + \varphi^{-1}}{1 + \epsilon \varphi^{-1}}, \tag{3.4}$$

[6]The welfare-relevant loss function in (3.2) is very similar to those derived in Aoki (2001), Benigno (2004) and Eusepi, Hobijn, and Tambalotti (2011). However, the last paper exhibits a loss function that also has a fourth component, a cross term comprised of deviations of output and relative prices. This term does not appear here since, unlike their model, our framework will not feature heterogeneity in production functions across sectors.

and:

$$\Lambda_t(\eta_k, \kappa_k, \zeta_{k,t}, \epsilon, \varphi) \equiv \epsilon(1+\varphi^{-1}\epsilon)^{-1} \times [\zeta_{1,t}\eta_2\kappa_1 + \zeta_{2,t}\eta_1\kappa_2 - 2\eta_1\eta_2\epsilon(\kappa_1+\kappa_2)(\zeta_{1,t}+\zeta_{2,t})], \quad (3.5)$$

a time-varying function of sectoral NKPC slopes, weights, structural parameters and two *Lagrange* multipliers $\zeta_{k,t}$, one for each sectoral NKPC in the minimization problem[7].

Note from the first term in (3.3) that it is now sectoral inflation which affects the outcome of policy. In addition, the presence of Λ_t in that equation makes the relationship between the output gap and sectoral inflation non-trivial. Results of a number of articles that address optimal policy in heterogeneous models, however, imply that the policymaker should respond more strongly to stabilize inflation rates of the more sticky sectors[8]. In fact, those are the sectors with flatter NKPCs, more frictions and larger real distortions[9]. Besides Aoki (2001), already cited, see Benigno (2004), Kösem-Alp (2010) and Eusepi, Hobijn, and Tambalotti (2011), to mention a few.

A direct consequence of the equilibrium relationship in (3.3) is that the complete stabilization of headline inflation is no longer optimal. This attenuates the bias caused by the endogenous response of policy when one estimates the short-term relationship between *aggregate* inflation and the output gap. Particularly, identification improvements are generated by movements in aggregate inflation, which mainly absorbs shifts in the inflation rates of the more flexible sectors (those with higher κ_k), less subject to be offset by policy. In this

[7]We assumed that the policymaker chooses levels of output and inflation *in the sectors* to derive (3.3) to (3.5). This is analogous to the textbook model, where monetary policy is assumed to pick levels of aggregate output and inflation.

[8]Similarly, one could think of a counterpart of (3.1) to the heterogeneous economy as:

$$y_{k,t} = -\Upsilon_\pi(\kappa_k, \kappa_{k'})\pi_{k,t} + \varrho_t, \quad (3.6)$$

where $\Upsilon_\pi(\kappa_k, \kappa_{k'})$ is a positive function, decreasing in κ_k, of NKPC slopes of sectors k and $k' \neq k$.

[9]Because of the aforementioned tension between alternative stabilization objectives, these are often referred to as *second-best* optimal policies. However, it does not mean that the *first-best* optimal policy can not be well approximated by these alternatives. For example, Eusepi, Hobijn, and Tambalotti (2011) quantitatively show how stabilization of a welfare-based price index can produce negligible welfare losses compared to the *first-best* optimal policy.

regard, note that the specific source of the shocks should not matter. Even if demand shocks represent the main reason behind shifts in sectoral inflation rates of the more flexible sectors, they still contribute to a better identification of the NKPC, as long as the policymaker accommodates them.

But what if monetary policy does *not* consider a multi-sector model with price-stickiness heterogeneity? In this case, we are back to (3.1), where monetary policy attempts to stabilize both aggregate inflation and output. Nonetheless, once again the heterogeneous model should improve estimations due to the presence of relative prices in its aggregate NKPC. If the Central Bank is inattentive to welfare-relevant losses related to movements in relative prices, it shall not attempt to stabilize these terms. Hence, they can be useful as an auxiliary source of variation that does *not* relate to policy, improving the identification of the Phillips curve in a similar manner. Additionally, even if that is the case, cross-sector variation can still be used in the estimations to attenuate the simultaneity bias, for the same reasons discussed by McLeay and Tenreyro (2019) for regional data.

3.2.2 Strategic Complementarities in Price Setting: More Sensible Degrees of Stickiness

A common difficulty in the empirical literature is that, even when estimates suggest a positive coefficient for the slope, the implied degree of nominal rigidity in the economy is too high, being primarily inconsistent with the micro evidence. Studying disaggregated data, Bils and Klenow (2004) report that half of prices in their sample last less than 5.5 months when excluding sales. Similarly, Nakamura and Steinsson (2008) report a median duration of prices from 7 to 9 months (also ignoring sales). In contrast, empirical estimations justified on the basis of the simplest New-Keynesian model generally imply an expected duration of price spells of *at least* 15 months, but often higher[10].

[10]Simple models based on Calvo-pricing usually return a high and rather imprecise Calvo parameter, rarely lower than 0.8. This gives an expected duration of price spells of at least 5 quarters, or 15 months.

In contrast to those simpler models used in the empirical literature, the combination of price-setting heterogeneity with the presence of strategic interactions across sectors implies a different price dynamics. Particularly, with *strategic complementarity* in price-setting decisions across sectors, firms in the more flexible sectors avoid setting prices that are too disparate compared to the future aggregate price level[11]. It follows that the more sticky sectors will disproportionately influence aggregate prices – Woodford (2003, Chapter 3).

Everything else constant, the price dynamics would be more staggered in the heterogeneous economy compared to the homogeneous case. Indeed, for similar degrees of strategic complementarity, Carvalho (2006) shows that, to replicate the dynamics of the heterogeneous economy, its identical-firms counterpart requires a frequency of price changes that is up to three times lower than the average of the heterogeneous economy.

Following that rationale, by relying on small variations of the simplest New-Keynesian model, the empirical literature is possibly obtaining estimates of the aggregate degree of nominal rigidity that are biased *upwards*. With a trivial inverse relationship between the degree of stickiness and the slope in those simple models, their estimates of the slope are also biased *downwards*[12]. Some articles confirm these findings empirically, but based on simulated data – e.g., Imbs, Jondeau, and Pelgrin (2011), calibrating for France, and Cagliarini, Robinson, and Tran (2011), for Australia. Piazza (2018) use heterogeneity in sectoral Phillips curves to construct an estimation strategy that relies on idiosyncratic shocks (purged from the effect of aggregate shocks), but he does not estimate degrees of stickiness directly (neither of the economy, nor of the sectors).

Differences between the price dynamics of the heterogeneous model and that of its identical-firms counterpart should be more evident with more strategic complementarity in price set-

[11] The model we use in this paper is similar to that in Carvalho (2006). It can be calibrated either for the presence of strategic complementarity or strategic substitutability in price setting, but the latter is only achievable in this type of model under unrealistic calibration values. See Carvalho (2006) and Carvalho and Nechio (2016) for a more detailed discussion.

[12] As discussed later, the relationship between the aggregate degree of stickiness and the slope is not so simple in multi-sector heterogeneous models. The effect of nominal rigidities on the slope will also depend on how stickiness is distributed across sectors.

ting. We later test this theoretical implication empirically, using the heterogeneous model.

If strategic complementarities are strong enough, the more flexible sectors adjust their prices considerably less than they otherwise would, in response to a demand shock. In the limit, aggregate prices would not change much for a large output effect. A first consequence is that the slope – which measures such sensitivity – would tend to zero. A second is that the average degree of stickiness in the economy need not be as high as often estimated in the literature to provide the same dynamics for inflation. Therefore, the estimated degree of stickiness exhibited by the heterogeneous economy would be lower, potentially approaching values observed for micro data.

3.3 Inflation Dynamics in a Multi-Sector Model

The last section discussed the ways through which an estimation strategy that relies on heterogeneity in price setting across sectors may circumvent estimation difficulties affecting the NKPC. Particularly, it should help to revert a downward bias for the slope *and* an upward bias for the implied degree of stickiness in the economy. In addition, with solid evidence of heterogeneity in the frequency of price changes across sectors, – e.g., Bils and Klenow (2004), Dhyne et al. (2006), Nakamura and Steinsson (2008) and Nakamura and Steinsson (2013) –, it should also be the method to choose.

We now turn to the heterogeneous model we use in the paper. It is similar to those used in Carvalho (2006) and Eusepi, Hobijn, and Tambalotti (2011), and it is fully presented in the appendix, for conciseness reasons. A small difference here is the introduction of a standard price-indexation scheme. Dropping this indexation rule returns the same (purely forward-looking) NKPC featured in Carvalho (2006)[13]. Price indexation should not modify the implications of the analysis of section 2.1, according to findings in Steinsson (2003) and

[13]In terms of findings, the model without indexation produces very similar results for the slope compared to those shown in the paper. Results for the degree of stickiness in the economy would be weaker, though.

Kösem-Alp (2010)[14].

The model features a continuum of firms kj in each sector, each producing the consumption variety of the good $j \in [0, 1]$ of sector k. The weight of each sector in aggregate expenditures is represented by η_k, such that $\sum_{k=1}^{K} \eta_k = 1$. Firms are monopolistically competitive, hiring labor based on a linear technology function. A representative household, which owns these firms, also supplies firm-specific labor to them. This type of segmented labor market is the reason why strategic complementarities in price setting arise in the model. The consumer derives utility from a Dixit-Stiglitz composite of differentiated consumption goods in the economy. It is assumed that each firm kj fixes its price as in Calvo (1983), but heterogeneity arises from a sector-specific probability of a price change, denoted by λ_k (equivalently, the Calvo parameter in sector k is $\theta_k \equiv 1 - \lambda_k$).

Firms which can *not* readjust their prices in any given period set them according to:

$$P_{kj,t} = P_{kj,t-1} \left(\frac{P_{k,t-1}}{P_{k,t-2}} \right)^{\gamma_k}, \tag{3.7}$$

where γ_k governs the degree of persistence in sectoral inflation rates. In what follows, we assume $\gamma_k = \gamma$. This considerably simplifies the form of the NKPC of the economy, reducing non-linearities of the moment conditions we estimate. In addition, we later evaluate the model in the cross section of sectors, comparing estimates of λ_k with implications from the micro evidence for each sector. By normalizing for the same γ, we do not affect correlations in such comparison, regardless of possible bias in $\hat{\gamma}$[15].

An advantage of (3.7) is that it does not introduce noise in the channel related to strategic complementarities, analysed in Section 3.2.2. For example, if we were to assume an in-

[14] In a model comprised of a sector that sets prices à la Calvo (1983) and another sector that fixes prices according to a backward-looking rule of thumb, Steinsson (2003) shows that the main features of optimal policy in the purely forward-looking case carry over to the hybrid case. Kösem-Alp (2010) shows that the sectoral slopes of the NKPCs continue to serve as a guide for optimal policy weights, as illustrated in section 2.1, regardless of how degrees of persistence may be distributed across sectors.

[15] With a sector-specific γ_k, we could be overestimating (underestimating) λ_k due to an upward (a downward) bias in $\hat{\gamma}_k$ for some sector. This would affect correlations between these parameters and those implied by the micro evidence, blurring the evaluation of the model

dexation scheme on the aggregate price index, then sectors with lower frequencies of price changes – which, under no indexation, have a disproportionate effect on aggregate prices – would set prices based on an aggregate index that already reflects price setting decisions of the more flexible sectors. These flexible sectors would then have a higher impact on the aggregate price dynamics, attenuating the central role of the more sticky ones and undermining empirical benefits generated by the model in terms of the NKPC.

3.3.1 New-Keynesian Phillips curves

As shown in the appendix, the log-linearized version of the model produces an aggregate NKPC that takes the form:

$$\pi_t = \underbrace{\frac{\beta}{1+\beta\gamma}}_{\equiv\gamma_f} E_t\pi_{t+1} + \underbrace{\frac{\gamma}{1+\beta\gamma}}_{\equiv\gamma_b}\pi_{t-1} + \underbrace{\psi\overbrace{\left(\frac{\sigma+\varphi^{-1}}{1+\epsilon\varphi^{-1}}\right)}^{\text{Real Rigidities}}}_{\equiv\kappa\text{ (Slope)}} y_t + \underbrace{\frac{\psi}{\epsilon}g_t}_{\text{Shift Term}} + u_t, \tag{3.8}$$

where:

$$\psi = \underbrace{\sum_{k=1}^{K}\eta_k\left[\frac{\lambda_k}{(1-\lambda_k)(1+\beta\gamma)} - \frac{\beta\lambda_k}{(1+\beta\gamma)}\right]}_{\text{Nominal Rigidities}},$$

$$g_t = \sum_{k=1}^{K}\tilde{\eta}_k(y_{k,t}-y_t), \qquad \tilde{\eta}_k = \frac{\frac{\lambda_k}{(1-\lambda_k)(1+\beta\gamma)} - \frac{\beta\lambda_k}{(1+\beta\gamma)}}{\sum_{k=1}^{K}\eta_k\left[\frac{\lambda_k}{(1-\lambda_k)(1+\beta\gamma)} - \frac{\beta\lambda_k}{(1+\beta\gamma)}\right]}\eta_k,$$

where β is the discount factor and terms denoted by $(y_{k,t}-y_t)$ represent *relative gaps*, the difference between sectoral output gaps and the output gap of the economy. In this type of model, these terms are a direct function of relative prices, depicted in (3.2)[16]. Pragmatically,

[16]In this type of model, a similar NKPC featuring relative prices could be derived using sectoral demand functions: $y_{k,t} = y_t - \epsilon(p_{k,t}-p_t)$.

they also play the same role, entering the NKPC as a composite shift term, in $\frac{\psi}{\epsilon}g_t$. Note that (3.8) nests standard forms in the literature, since the shift term would disappear under homogeneity ($\lambda_k = \lambda$). This term is proportional to a weighted average of relative gaps. Each weight $\tilde{\eta}_k$ is a transformation of the original weight η_k, but adjusted for the relative degree of flexibility of the sector compared to that of the entire economy.

An identical-firms economy features a direct inverse relationship between the degree of stickiness and the slope of its NKPC, no longer the case in (3.8). Heterogeneity in price setting produces a slope formed by two components. The first, ψ, summarizes the degree of *nominal* rigidity in the economy, as well as its distribution across sectors. The second term, comprised of $\Theta \equiv \frac{\sigma+\varphi^{-1}}{1+\epsilon\varphi^{-1}}$, relates to the degree of *real* rigidities, directly corresponding to the degree of strategic complementarities in price setting[17]. As shown by Carvalho (2006), compared to a homogeneous economy calibrated for the average frequency of price changes, the former tends to increase the sensitivity of inflation to the output gap, whereas the latter operates in the opposite direction.

Sectoral NKPCs of the model take the form:

$$\begin{aligned}\pi_{k,t} =& \frac{\beta}{1+\beta\gamma}E_t\pi_{k,t+1} + \frac{\gamma}{1+\beta\gamma}\pi_{k,t-1} \\ &+ \left[\frac{\lambda_k}{(1-\lambda_k)(1+\beta\gamma)} - \frac{\beta\lambda_k}{(1+\beta\gamma)}\right]\left(\frac{\sigma+\varphi^{-1}}{1+\epsilon\varphi^{-1}} - \frac{1}{\epsilon}\right)y_t \\ &+ \frac{1}{\epsilon}\underbrace{\left[\frac{\lambda_k}{(1-\lambda_k)(1+\beta\gamma)} - \frac{\beta\lambda_k}{(1+\beta\gamma)}\right]}_{\equiv\psi_k(\lambda_k,\beta,\gamma)} y_{k,t} + v_{k,t}.\end{aligned} \tag{3.9}$$

The term $\psi_k(\lambda_k, \beta, \gamma)$ plays a crucial role. First, it relates to the degree of nominal rigidity in each sector, determining the sectoral slope in (3.9) – i.e., $\kappa_k = \epsilon^{-1} \times \psi_k(\lambda_k, \beta, \gamma)$. Second, it is also relevant to identify both the slope and composite shift term of the aggregate NKPC in (3.8), due to its presence in ψ and g_t.

[17]If $\lambda_k = \lambda$ and $\gamma = 0$, the coefficient that multiplies the output gap becomes $\Theta\left(\frac{\lambda}{1-\lambda} - \beta\lambda\right)$, so that Θ corresponds to the Ball and Romer (1990) coefficient of real rigidities in this model.

Essentially, this assessment implies that coefficients in (3.8) can be identified through (3.9). This is the case even under homogeneity ($\lambda_k = \lambda$), when $\psi = \psi_k$ and sectoral and aggregate NKPC slopes would still be linked by the function $\kappa = \epsilon \times \Theta \times \kappa_k$. In addition, following the discussion of Section 3.2.1, if sectoral NKPCs are not as blurred by the endogenous response of policy as that of the economy (especially those of the more flexible sectors), there would also be manifest advantages in a strategy that attempts to identify (3.8) using (3.9). We benefit from this fact in the estimations of the next section.

It is worth emphasising, however, that estimates of the slope, κ, would still be biased towards zero under any estimation strategy that uses (3.9). This bias arises due to the presence of y_t, and, therefore, of aggregate shocks (u_t) in these sectoral equations. Nonetheless, in the next section we show that our estimations that include (3.9) still produce sizable and statistically significant estimates of the slope.

3.4 Empirical Approach

With a significant number of sectors, estimating the structural form in (3.8) becomes non-trivial. It would combine several endogenous variables with a non-linear setting. An alternative is to calibrate a number of its deep parameters – eventually, λ_k for some sectors – while estimating others. The main disadvantage of this procedure is that it biases the comparison between the estimated degree of stickiness of the heterogeneous economy and actual micro data[18].

Another option – the one we adopt here – is to jointly exploit the structures of the aggregate NKPC in (3.8) and its sectoral analogues in (3.9), gaining efficiency from cross-equation

[18]For example, a method that calibrates some of the λ_k based on the micro evidence would bias the implied degree of nominal rigidities in the economy, $\sum_{k=1}^{K} \eta_k(1 - \lambda_k)$, towards that of actual disaggregated data. Additionally, in such non-linear setting, final estimates for the sectors that are not parameterized could be highly sensitive to the adopted values for the ones that are calibrated, depending on the surface of the underlying likelihood function. In this sense, a technique that estimates *all* the parameters related to the degree of stickiness in the sectors is preferred.

restrictions. Our strategy relies on estimating *all* of the parameters associated with degrees of nominal rigidity in the sectors, λ_k, while calibrating the remaining parameters, σ, φ and ϵ, which directly govern the degree of strategic complementarities in price setting (real rigidities). In addition, we also estimate β and γ[19].

We verify how obtained estimates vary based on three alternative degrees of real rigidities, comparing these results with predictions implied by theory, as motivated in Section 3.2.2. Table 3.1 summarizes these alternatives.

The elasticity of substitution between varieties is set to 8. This value is close to the one used in Carvalho and Nechio (2016) and lies in the middle of the two calibrations used in Carvalho (2006). It is also intermediate between typically low elasticities of the IO literature and comparatively high values found in macroeconomic models. The inverse of the EIS is set to 1 in the baseline calibration. We add 1/2 to this value to reconcile with small estimates of the EIS found in the empirical literature (< 1), or subtract 1/2 to approach values often adopted in the macroeconomic literature (≥ 2). Finally, an elasticity of labor supply from 0.2 to 1.8 meets two opposing observations, from those of the microeconomic literature (typically very low – e.g., see Pencavel (1986)) to those of the macroeconomic literature (often around or higher than 2 – e.g., Carvalho (2006), Eusepi, Hobijn, and Tambalotti (2011), Carvalho and Nechio (2016)). A baseline Frisch elasticity of 1 is motivated by evidence presented by Chang and Kim (2006), who finds such value relying on a rich model that features heterogeneity in the workforce.

We estimate the system comprised of (3.8) and (3.9) using General Method of Moments (GMM). For the aggregate NKPC, the instrument set we adopt is a generalization of the one used in Gali and Gertler (1999), but for the heterogeneous economy[20]. In this case,

[19] It is possible to extend the method by allowing one of the remaining parameters (σ, φ and ϵ) to be free in the estimation. However, the resulting non-linear structure in the term related to real rigidities (Θ) further complicates the method and our estimations when allowing more than one of those parameters to be estimated did not improve our results. In addition, choosing which of those parameters should be estimated is clearly arbitrary. We leave the challenging extension of the method where all structural parameters are estimated for future research.

[20] Since the heterogeneous NKPC encompasses a number of additional endogenous variables, for example,

Table 3.1: Calibration of Parameters Related to Real Rigidities

Parameter	Interpretation	Baseline	↑ Real Rigidity	↓ Real Rigidity
ϵ	elasticity of subst. between varieties	8	8	8
σ	inverse of the EIS	1	0.5	1.5
φ	(Frisch) elasticity of labour supply	1	0.2	1.8

Notes: Different calibrations of the model. Each implies a different degree of strategic complementarities in price setting. The first (baseline) implies $\Theta \approx 0.22$, the second produces $\Theta \approx 0.13$ and the third, $\Theta \approx 0.38$.

the aggregate equation is instrumented by the first two lags of the output gap and sectoral sectoral gaps, the labor share (which proxies aggregate marginal costs), aggregate inflation, the Fed Funds rate, a Treasury spread and inflation rates of wages and commodities[21]. For sectoral NKPCs, we simply apply the first two lags of endogenous variables in (3.9)[22]. In the appendix, we show that results presented in the main paper are maintained for alternative approaches to instruments[23].

3.5 Data

We briefly discuss the data set. Our baseline dataset consists of aggregate and sectoral quarterly data for the U.S. economy during the interval from 1964:2 to 2021:2. This dataset excludes the recent inflationary window, so that is aligned with the empirical literature. In Appendix C.7, we present results while including more recent observations. Whereas identification becomes inherently more challenging in this setting, our main findings are stronger when considering that window.

due to the composite shift term in (3.8).

[21]The instrument set in Gali and Gertler (1999) encompasses the first four lags of those variables, except for sectoral output gaps, which are not included since he estimates a homogeneous model. To be able to estimate the system comprised of (3.8) and (3.9), we reduce the number of lags by half due to the resulting number of moment conditions in the GMM.

[22]As for π_{t-1} in (8), $\pi_{k,t-1}$ in (3.9) is considered predetermined.

[23]As the vast majority of articles in the literature, we adopt a standard approach of exclusion restrictions in the estimations that follow, replacing forward-looking expectations in the NKPCs by their actual values, so that innovations also include an expectational error.

For aggregate data, we use the real gross domestic product (GDP) for output, the producer price index (PPI) for all commodities (to construct commodities inflation), a 5-year Treasury rate spread (over the Fed Funds rate), the effective Fed Funds rate (interest rate), the non-farm labor share and average hourly earnings of production and non-supervisory workers (to construct wage inflation). We use real personal consumption expenditures (PCE) for sectoral output and PCE price indices to construct sectoral inflation[24]. For more details on the data set, see Table C.1 (appendix).

Table 3.2 below details the *fifteen* sectors in the economy. It provides weights, calculated from the share in aggregate expenditures, as well as implied Calvo-pricing probabilities, calculated from micro data in Bils and Klenow (2004). We adopt these infrequencies as a *benchmark* regarding the cross section of sectors to estimations we perform in this paper.

3.6 Results

We turn to results from the estimation of the system comprised of the aggregate NKPC in (3.8) and the fifteen sectoral NKPC depicted in (3.9). The baseline specification allows for price-setting heterogeneity, when λ_k is sector-specific[25].

To evaluate the sensitivity of results to the presence of heterogeneity, we also present estimations for the same system when heterogeneity is switched off, $\lambda_k = \lambda$. This is a more parallel approach to those in Fitzgerald, Nicolini, et al. (2014) and McLeay and Tenreyro (2019). It still collects some of the empirical advantages caused by the attenuation of the endogenous response of policy in the sectors, as discussed in Section 3.2.1, but, in the absence of heterogeneity, not those of Section 3.2.2. Hence, by the comparison between this

[24]In the absence of data for income at sector level, we proxy it by sectoral consumption. Cyclical components of aggregate and sectoral outputs are extracted by the Hodrick-Prescott filter (HP), setting $\lambda_{HP} = 1600$, as advised for quarterly data. To account for a model with no population growth, we also define variables in per capita terms, when appropriate.

[25]Structural parameters λ_k, β and γ theoretically lie in the range from 0 to 1. We discount 10^{-3} in each bound of this interval in order to properly provide the Jacobian matrix to the estimation algorithm.

Table 3.2: Sectors, Weights and Infrequencies Based on Micro Data

Sector	Sectoral Weight	Benchmark Infrequency
Motor Vehicles and Parts	5.34%	0.212
Furnishings and Durable Household Equipment	3.61%	0.484
Recreational Goods and Vehicles	2.92%	0.564
Other Durable Goods	1.59%	0.551
Food and Beverages Purchased for Off-Premises Consumption	11.82%	0.327
Clothing and Footwear	5.41%	0.331
Gasoline and Other Energy Goods	3.75%	0.003
Other Nondurable Goods	8.03%	0.541
Housing and Utilities	18.19%	0.212
Health Care	11.79%	0.857
Transportation Services	3.23%	0.375
Recreation Services	3.02%	0.727
Food Services	6.50%	0.590
Financial Services and Insurance	6.44%	0.781
Other Services	8.34%	0.645

Notes: Benchmark infrequencies are implied Calvo-pricing probabilities based on micro data. These come from a mapping between disaggregated PCE price data and evidence exhibited in Bils and Klenow (2004), both expressed monthly. We convert probabilities into quarterly analogues by compounding them geometrically.

approach and the baseline, described above, we evaluate the importance of the channel related to strategic complementarities to our findings. Note that, with $\lambda_k = \lambda$, the composite shift term completely disappears from (3.8), returning an aggregate NKPC that looks similar to models estimated in the empirical literature.

We also present results for the naive case, the single-equation estimation of (3.8) when heterogeneity is once again switched off. In the absence of both heterogeneity and cross-sector variation, this alternative essentially mimics standard specifications represented in the literature. Results for this model are completely isolated from the analyses of Sections 3.2.1 and 3.2.2. Therefore, differences in terms of results between this approach and the last one are most closely associated with attenuation of the simultaneity bias caused by monetary policy in the latter, discussed in Section 3.2.1.

Table 3.3 summarises our findings. It presents estimates of the aggregate slope, κ, the implied degree of stickiness in the economy, $\theta = \sum_{k=1}^{K} \eta_k \theta_k = \sum_{k=1}^{K} \eta_k (1 - \lambda_k)$, β and γ. Parameters that govern real rigidity (σ, φ and ϵ) are calibrated according to Table 3.1. For the heterogeneous economy, Table 3.3 also shows correlations between estimated Calvo-pricing probabilities, $1 - \hat{\lambda}_k$ and benchmark infrequencies implied by evidence in Bils and Klenow (2004) – see Table 3.2.

As in the literature, the single-equation estimation of a textbook NKPC under homogeneity returns degrees of stickiness that are too high, being primarily incompatible with the micro evidence. For example, the benchmark infrequencies would imply $\theta^{micro} = 0.48$. None of the estimations using the single-equation homogeneous model return a value of θ lower than 0.8. Confidence intervals are also substantially wide, being seemingly uninformative. When considered together, they suggest an aggregate Calvo probability from 0.7 to 1. With a trivial inverse relationship between this parameter and the NKPC slope, the latter is roughly estimated at zero. Indeed, we can not reject the hypothesis that the slope is zero (at 5%) in this case for two out of the three cases tested.

Still under homogeneity, exploiting cross equations restrictions through the introduction of the sectoral NKPCs in (3.9) generates a slope that is *at least* five times higher (depending on calibration) and always statistically significant. Recall that, due to the presence of y_t in sectoral NKPCs, these values are still biased towards zero. The implied degree of stickiness of the economy is now estimated around 0.6, regardless of calibration, approaching values found for disaggregated data. With very low standard errors, confidence intervals for these parameters are also quite narrow. Results also suggest that lower values of $\hat{\theta}$ are *not* obtained through an upward bias in $\hat{\gamma}$, since the latter is also lower, around 0.3, compared to 0.4-0.5 under the naive approach.

Taken together, these findings agree with McLeay and Tenreyro (2019) by suggesting that the use of disaggregated data (in our case, sectoral; in their case, regional), can partially mitigate the simultaneity bias caused by monetary policy. However, Table 3.3 also shows that one obtains even higher estimates of the slope when allowing for price-setting heterogeneity, achieving the full potential of the analyses given in Sections 3.2.1 and 3.2.2. In this case, one obtains sizeable estimates of the slope, between 0.04 and and 0.14, all of these significant at 1%. This nearly doubles estimates under the last approach for two out of the three calibrations applied. Estimates of the remaining parameters maintain their values, with $\hat{\theta}$ around 0.6 and $\hat{\gamma}$ being estimated at 0.3. In addition, note that correlations with the micro benchmark are generally high, around 0.7.

Another interesting finding from Table 3.3 is that, as predicted by theory, we generally estimate a *lower* slope and a *lower* degree of stickiness for the heterogeneous economy as we introduce more strategic complementarities in price setting – see Section 3.2.2[26]. In contrast, estimates under the naive (single-equation) approach are counter-intuitive, since we obtain higher estimates of θ as we introduce more real rigidities in the model.

In Appendix C.3, we evaluate the sensitivity of our results perturbing the instrument set. Firstly, we apply a data-driven instrument selection routine based on regularisation. Then, we test three variations of our baseline instrument set that control for potential pitfalls affecting our estimations, namely time-aggregation issues involving macro data and the number of moment conditions in the GMM (i.e., too many instruments). All of those estimations reconfirm our main findings.

3.6.1 Behind the Scenes

To confirm that the model captures essential information found in micro data, Figure 3.1 compares $\hat{\theta}_k$ with micro benchmarks for the same parameters, based on disaggregated data in Bils and Klenow (2004). The calibration is the baseline in Table 3.1, but very similar findings are generated with other parameterisations of the model. Except for probably two sectors ("motor vehicles and parts" and "financial services"), most part of the parameters are

[26]The only exception is the estimate of $\hat{\theta}$ in the first line of Table 3.3, slightly higher than that with the baseline calibration for the same model.

Table 3.3: Estimates of the Slope and Degree of Stickiness

			Parameters			
Calibration	Model	$Corr(\theta_k, Micro)$	κ	θ	β	γ
↑ Real Rigidity	Heterogeneous (SYS)	0.77	0.037***	0.64	0.97	0.29
			(0.000)	[0.63, 0.64]	(0.003)	(0.002)
	Homogeneous (SYS)	-	0.033***	0.58	0.98	0.30
			(0.000)	[0.58, 0.58]	(0.003)	(0.002)
	Homogeneous (SE)	-	0.002	0.87	0.96	0.44
			(0.003)	[0.70, 1.00]	(0.010)	(0.040)
Baseline	Heterogeneous (SYS)	0.65	0.092***	0.55	0.99	0.30
			(0.001)	[0.54, 0.55]	(0.008)	(0.004)
	Homogeneous (SYS)	-	0.052***	0.58	0.97	0.35
			(0.000)	[0.58, 0.58]	(0.004)	(0.002)
	Homogeneous (SE)	-	0.007*	0.82	0.96	0.45
			(0.003)	[0.74, 0.90]	(0.011)	(0.043)
↓ Real Rigidity	Heterogeneous (SYS)	0.74	0.138***	0.58	0.99	0.29
			(0.002)	[0.58, 0.59]	(0.008)	(0.004)
	Homogeneous (SYS)	-	0.070***	0.62	0.98	0.27
			(0.001)	[0.62, 0.63]	(0.008)	(0.005)
	Homogeneous (SE)	-	0.012***	0.81	0.97	0.49
			(0.004)	[0.77, 0.86]	(0.012)	(0.048)

Notes: The first column refers to the three different calibration sets exhibited in Table 3.1. We test three different estimation methods. "Heterogeneous (SYS)" denotes the baseline model with sector-specific λ_k, being estimated by System-GMM with the aggregate NKPC in (3.8) and the fifteen sectoral NKPCs in (3.9). "Homogeneous (SYS)" uses the same system, but imposes $\lambda_k = \lambda$ for every sector. In such case, the shift term disappears from (3.8). "Homogeneous (SE)" mimics the standard approach in the literature, repeating this last exercise considering solely the aggregate NKPC (3.8), i.e., single-equation estimation. Correlations between estimated and benchmark infrequencies $(1-\lambda_k)$ that come from the micro data in Bils and Klenow (2004) are shown in the column "$Corr(\theta_k, Micro)$". **The micro benchmark implies $\boldsymbol{\theta^{micro} \approx 0.48}$.** κ denotes the aggregate slope in (8), while θ is the implied degree of stickiness in the economy. When λ_k varies across sectors (heterogeneous case), $\theta = \sum_{k=1}^{K} \eta_k \theta_k = \sum_{k=1}^{K} \eta_k (1-\lambda_k)$. Under homogeneity ($\lambda_k = \lambda$), this simplifies to $\theta = (1-\lambda)$. We use a HAC estimator for the covariance matrix. Standard errors are presented in parentheses. As in theory, structural parameters (λ_k, β and γ) can assume values in the interval $[0, 1]$. We test the null hypothesis of $\kappa = 0$: *p<0.1; **p<0.05; ***p<0.01.

aligned with the benchmarks.

Figure 3.2 presents individual confidence sets at 95% for the same estimations, showing that intervals are generally narrow for each sector. Analogous charts are exhibited in appendix D for the remaining calibration sets of Table 3.1.

Figure 3.1: $\hat{\theta}_k$ vs. Micro Benchmarks

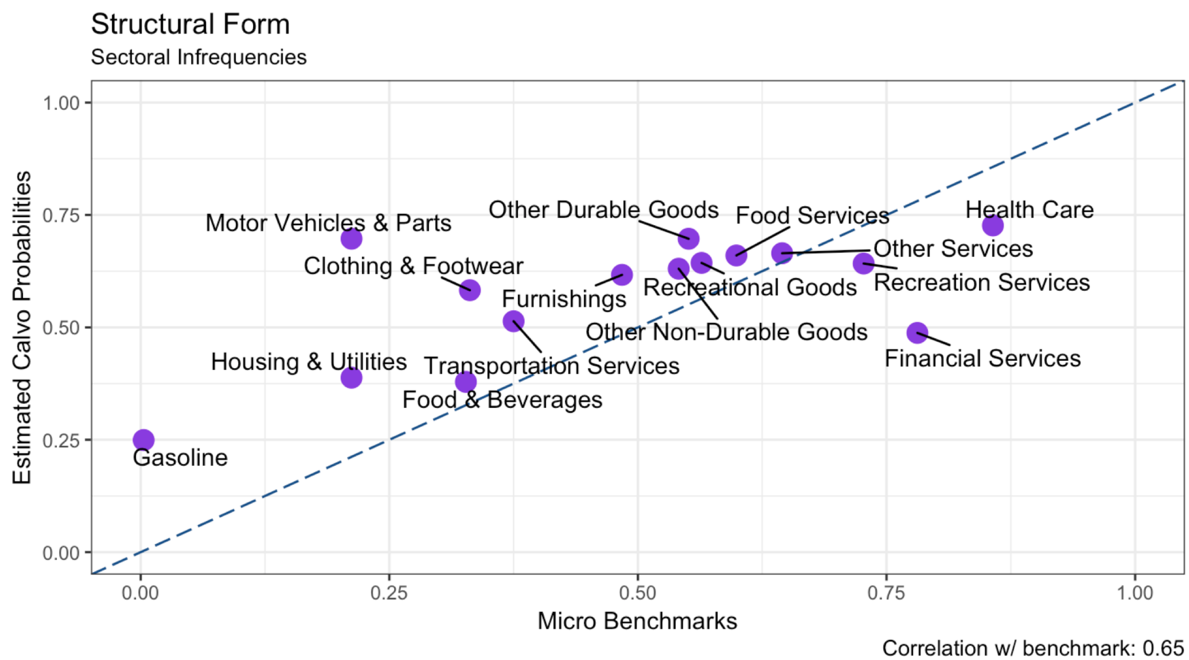

Notes: Estimated Calvo probabilities using the same econometric setting of Table 3.3. Benchmarks are implied probabilities from evidence in Bils and Klenow (2004) – see Table 3.2. We use the baseline calibration of Table 3.1.

3.6.2 Estimator Uncertainty

Parametric Stability

The use of conventional first-order asymptotics in an environment where instruments are potentially weak can be misleading. This pitfall is even more important concerning structural models, especially with non-linear moment conditions.

To address these uncertainties, weak-instrument-robust methods of inference are com-

Figure 3.2: $\hat{\theta}_k$ vs. Micro Benchmarks – Confidence Intervals

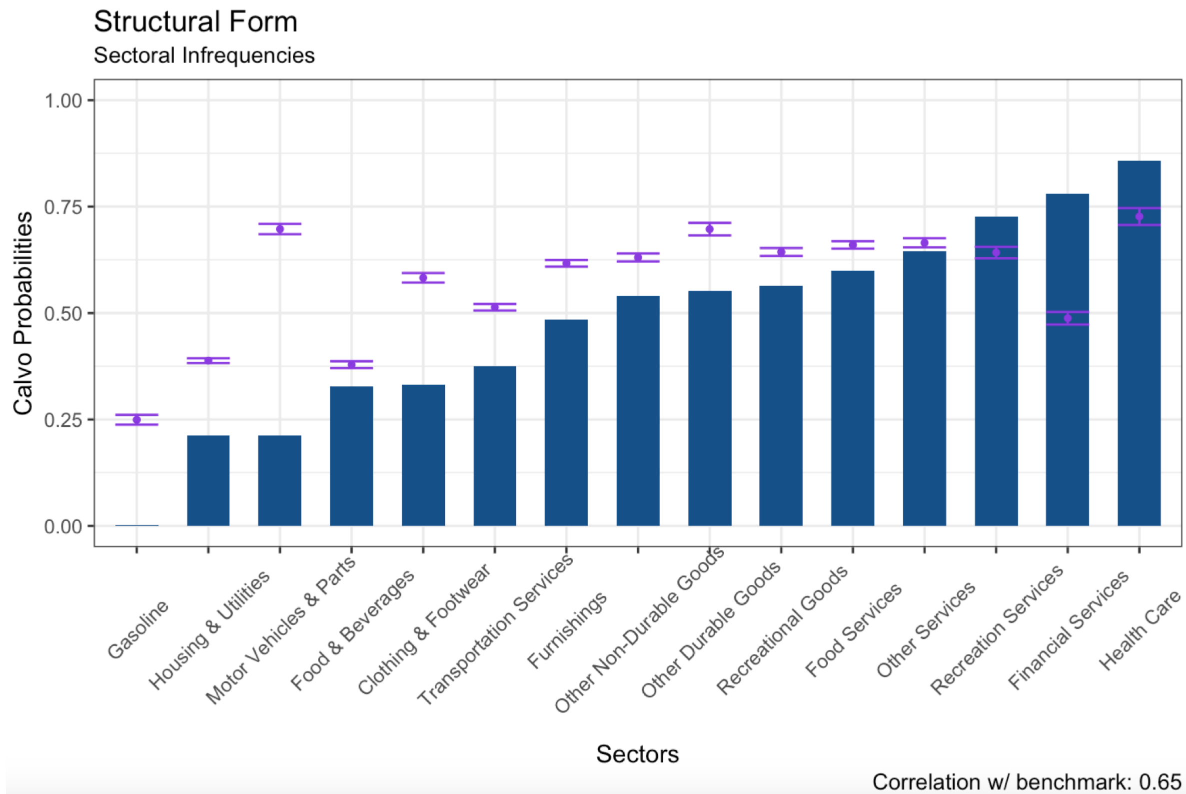

Notes: Estimated Calvo probabilities using the same econometric setting of Table 3.3. Blue bars are micro-based benchmark probabilities implied from evidence in Bils and Klenow (2004) and presented in Table 3.2. For expository purposes, these are sorted according to their degree of flexibility. 95% confidence intervals are shown for each $\hat{\theta}_k$. We use the baseline calibration of Table 3.1.

monly recommended. Two techniques typically used construct robust confidence intervals for deep parameters by inverting the test statistics S (Stock and Wright (2000)) and K (Kleibergen (2005))[27]. These tests are based on the empirical assessment that the estimations of parametric vectors are generally more meaningful than those of individual parameters. Hence, they seek to obtain a robust interval for one parameter while restricting the rest of the parametric space, usually assuming identification for them.

The problem with robust inference techniques in our setting is that, although they can be generalized to the presence of multiple endogenous variables, little is known about their

[27]See Ma (2002), Andrews and Stock (2005), Nason and G. W. Smith (2008a), Kleibergen and Mavroeidis (2009) and Mavroeidis, Plagborg-Møller, and H. Stock (2014), to cite just a few of the examples in the literature.

implications[28]. In such case, however, it is well known that they suffer from poor power[29]. With seventeen parameters we would have to assume that a subvector of sixteen of them is identified to construct robust sets (unlikely). Lastly, in order to invert test statistics, the number of grid points at which we need to evaluate the null hypothesis grows exponentially with the dimension of the parametric vector.

To circumvent those problems, we rely on an alternative method to verify the model based on uncertainties involving the estimation. For each of the $P = 15$ sectoral probabilities λ_k in the system formed by (3.8) and (3.9), we fix $P - q$ of them at the original point estimates, re-estimating q parameters. For $q = 4$, this generates $C_4^{15} = 1365$ re-estimations of the model, with 365 estimates of λ_k for each sector.

Figure 3.3 presents confidence intervals for sectoral Calvo-pricing probabilities, constructed from those restricted estimations. Boxes for each sector represent the interquartile range. Vertical lines provide the 5%ile $- 95\%$ile interval of the distribution. Note that results are quite in line with those exhibited in Figure 3.2. The correlation between the median estimate – horizontal line inside the boxes – of each sector and the corresponding micro-based benchmark is also high, approximately 0.5. Cases for $q < 4$ show similar findings, while repeating the exercise with $q > 4$ produces approximately the same chart.

Subsample Stability

Next, we turn to uncertainties involving the sample. A typical practice has been to divide the sample into two or more periods, commonly splitting it around 1979 – e.g., Gali and Gertler (1999) and Clarida, Gali, and Gertler (2000). However, our GMM features too many moment conditions for such separation to be feasible in our environment.

[28] The use of robust inference in high-dimensional settings is not common. For example, Andrews and Stock (2005) analyze weak-instrument-robust methods covering a sample of studies published in the American Economic Association journals. None of the 230 articles in their sample apply weak-IV methods using more than four endogenous variables in the estimations. Note that there are *sixteen* endogenous variables in our model.

[29] In our case, both the S and the K tests hardly reject the null. This is true for most part of the points in the parametric space – likely because our GMM features too many moment conditions for such tests to be reliable.

Figure 3.3: Confidence Sets Constructed from Restricted Estimations

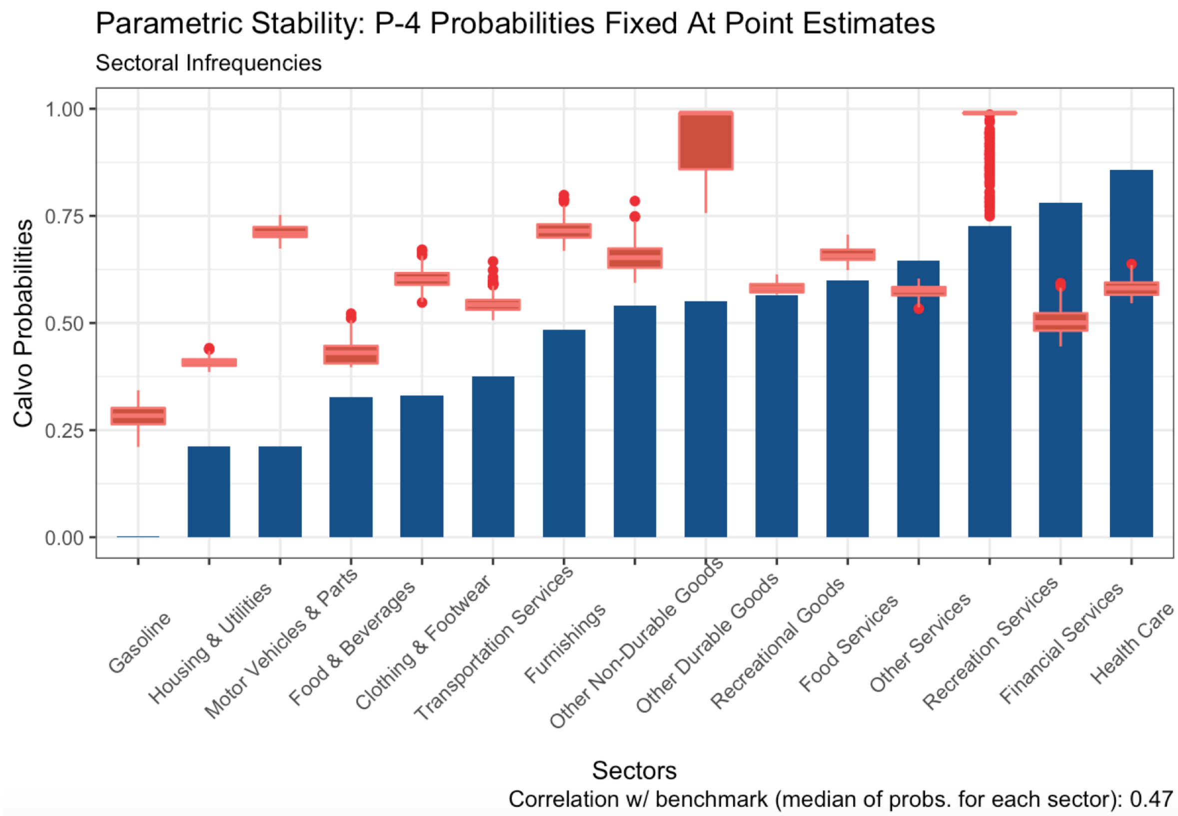

Notes: Parametric stability when 11 sectoral probabilities are fixed. Boxes represent the interval from the 25%ile to the 75%ile of distributions (for each sectoral probability). Horizontal lines are median estimates. 1365 restricted versions of the model are estimated, 365 estimates for each sector in total (vertically positioned). Baseline calibration.

To test the subsample stability of our model, we conduct rolling-GMM estimations with $T = 200$ observations[30]. Figure 3.4 presents similar results to those of Table 3.3 for implied coefficients, such as the aggregate slope, κ, and the aggregate infrequency, θ, based on the heterogeneous model and the system comprised of (3.8) and (3.9)[31]. Point estimates of these coefficients are substantially stable throughout the rolling window, with narrow confidence sets. This is also the case for forward and backward-looking coefficients in the aggregate NKPC (right panels). Figure 3.4 also presents correlations between estimated sectoral in-

[30]To reduce the number of moment conditions in the GMM, we remove the second lag of variables in the instrument set. Thus, instruments are the first lags of the same variables described in Section 3.4.

[31]We exhibit rolling-GMM estimates of all of the structural parameters of the heterogeneous economy in the appendix.

frequencies $(1 - \hat{\lambda}_k)$ and micro benchmarks. These correlations systematically lie around 0.7, approximately their average (horizontal dotted line), agreeing with results exhibited in Table 3.3.

Figure 3.4: Subsample Stability: Implied Coefficients for the Heterogeneous Model

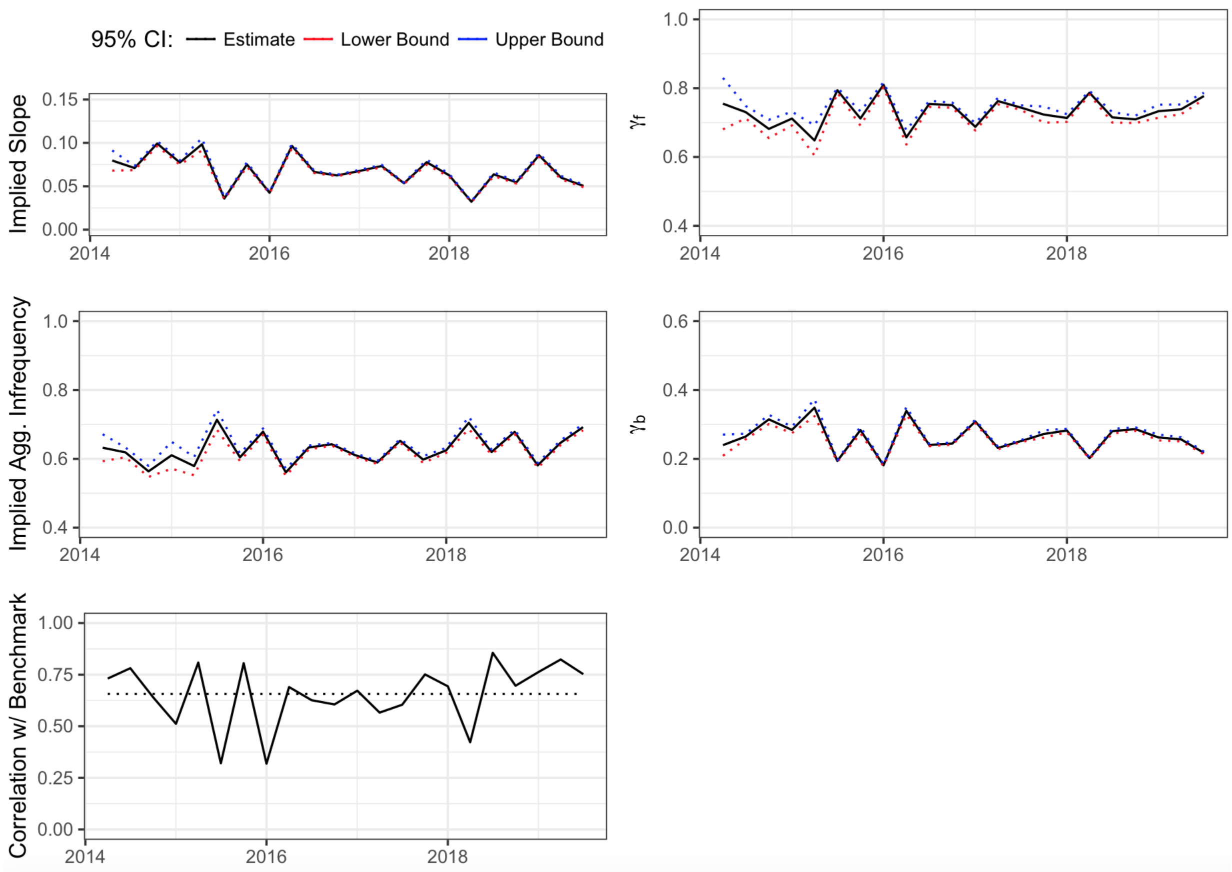

Notes: Implied cofficients from rolling-GMM estimations of the system comprised of the aggregate NKPC in (3.8) and the fifteen sectoral NKPCs in (3.9), under heterogeneity $(\lambda_k \neq \lambda)$. We use a HAC estimator for the covariance matrix. To reduce the number of moment conditions in the GMM, we remove the second lag of variables in the instrument set of estimations in Table 3.3. The horizontal axis measures the date of the last observation in the rolling subsample. The dotted line in the lower-end left panel denotes the time-series average correlation between estimated sectoral Calvo probabilities and the micro benchmark.

In Figure 3.5, we compare estimations using all models of Table 3.3. Consistent with previous results, the left panel indicates that the heterogeneous economy systematically implies a much higher slope for the NKPC. The aforementioned downward bias does not prevent us from obtaining a sizeable value for such coefficient in any of the estimations conducted with subsamples. In contrast, the slope is not statistically different from zero throughout the rolling window when we estimate the homogeneous economy without considering sectoral NKPCs. By including them, however, we mitigate some of the impact of the endogenous response of policy over our estimations, producing a stable slope that lies around 0.03.

For the homogeneous economy, the relationship between the slope and the degree of stickiness is trivial. As a result, a flat NKPC translates into very high estimates of the Calvo probability, as shown in the right panel, for single-equation estimations. However, in line with results of Table 3.3, the presence of sectoral Phillips curves in the estimation pushes that probability towards more reasonable values, approaching estimates based on the heterogeneous economy, as well as the micro evidence.

Figure 3.5: Subsample Stability: Implied Slope (Left) and Infrequency (Right) for All Models

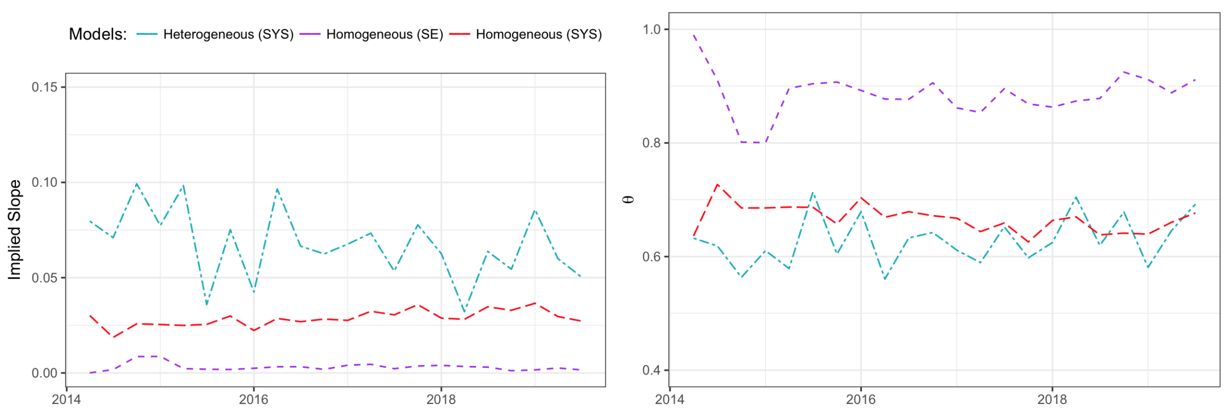

Notes: Implied cofficients from rolling-GMM estimations for each model in Table 3.3. We use a HAC estimator for the covariance matrix. To reduce the number of moment conditions in the GMM, we remove the second lag of variables in the instrument set. The horizontal axis measures the date of the last observation in the rolling subsample.

3.6.3 Additional Robustness Checks

Besides testing different instrument sets, we also conduct two additional robustness checks with the model. For the sake of conciseness, results are exhibited in the appendix.

First, we benefit from the fact that it is possible to identify the aggregate NKPC without directly estimating that equation. Since its parameters are also present in the sectoral NKPCs in (3.9), it is possible to drop the aggregate equation from the system, regardless of the presence of heterogeneity in price setting. Results for the slope would still be biased towards zero – since y_t appears in the sectoral NKPCs. In addition, we would likely lose efficiency by removing important information otherwise exploited through cross-equation correlations in the errors. Nonetheless, all our main findings are reconfirmed – see Appendix C.5.

Second, it is widely known that non-linear estimation methods can be quite sensitive to starting values. In the previous estimations, we relied on implied reset probabilities (λ_k) from the micro evidence as starting values. It could be that the seemingly reliable estimates we show are a direct consequence of that choice. For example, the vector of point estimates could be inherently copying the vector of starting values because of complexities in the moment conditions.

To verify the sensitivity of our model to initial values, in appendix F we re-conduct structural estimates while relying on an "agnostic" routine to generate starting values. Under the heterogeneous model, we individually estimate each sectoral equation, (3.9), and use the resulting $\hat{\lambda}_k$ as initial value for that sector when estimating the system comprised by (3.8) and (3.9). Note that there are several complications involving the single estimation of sectoral NKPCs[32]. Such procedure is *very* conservative, since nothing ensures that resulting estimates – and, then, starting values – are reliable. We show that our main findings are maintained for this strategy. Correlations with the benchmark are slightly lower, but still support the model[33].

[32] First, inflation expectations may not vary much for some of the sectors. Second, we are not using instruments outside of the model to improve the estimations (what could be substantially important for some sectors – e.g., lags of commodity indices instrumenting the NKPC of "gasoline and other energy goods"). Third, we find that estimates of λ_k are very sensitive to the econometric method when sectoral NKPCs are considered individually, what suggests to be sceptical on the reliability of these values. Lastly, note that a single distorted estimate of $\hat{\lambda}_k$ for some sector could further complicate the estimation of the system, if sensitivity to initial values is an issue for the latter.

[33] Additionally, we conduct tests to verify the underlying rank conditions of the system. For example, seventeen parameters are estimated in the system with (3.8) and (3.9), under heterogeneity. We fix fifteen of them, and generate 2×2 combinations with the precision of 10^{-3} for the remaining two in the $[0, 1]^2$ space. The step is continued until all the possible combinations involving deep parameters are exploited – the Jacobian matrix is mapped $C_2^{17} = 136$ times. We do not reject that the Jacobian matrices for those combinations have full rank with a tolerance of 2.331468×10^{-14} – this corresponds to the number of rows in the matrix times the default

3.7 Conclusion

This paper takes seriously the implications of heterogeneity in price setting to address empirical difficulties verified for the Phillips trade-off. Specifically, we propose a novel method that estimates a system of aggregate and sectoral Phillips curves from a heterogeneous multi-sector model. The procedure benefits from additional structure from theory by exploiting a rich set of cross-equations restrictions in the estimations.

First, bringing cross-sector variation to the estimations should improve identification by ameliorating the effect of the endogenous response of optimal monetary policy over coefficients. This is the case regardless of whether the policymaker considers the presence of stickiness heterogeneity across sectors or not. Second, the introduction of heterogeneity in price setting also has the potential to enhance the reliability of estimations. Specifically, an interesting channel arises due to strategic interactions between sectors. It should primarily revert high estimates of the degree of stickiness in the economy found in the literature back to more sensible levels, approaching the micro evidence. Our empirical strategy allows us to separate out the effect each of these two mechanisms has on estimations.

When taking benefits that arise from both cross-sector variation and heterogeneity to their full potential, the novel method delivers a sizeable, statistically significant and stable slope. It also produces degrees of stickiness that substantially approach levels implied by micro evidence, both regarding the entire economy and the cross section of sectors, with narrow standard errors.

The reliability of such approach does not seem affected by estimator uncertainties. We test it perturbing the econometric setting, as the number and which parameters are estimated, calibration, sample, instrument sets and starting values in the algorithm. The model exhibits no significant change in performance in any of those exercises.

Last but not least, this paper shows that it is indeed possible to obtain disciplined esti-

---

epsilon in *Matlab*®.

mations of the NKPC with the straightforward slack variable, the output gap. This contrasts with a tendency in the literature, which normally seeks alternative specifications based on marginal costs (often proxied by the labour share) to be able to obtain more sensible estimates of structural coefficients.

Appendix A

Appendix for Chapter 1

A.1 Households' Optimisation

U.S. Households

The households' optimisation problem can be divided into two separate problems: static and dynamic.

In the static problem, households choose, at a given period and history, how to split their total consumption expenditure between the consumption of different goods. U.S. households in period 1, for instance, will solve the static optimisation:

$$\max_{C_{NT,1},C_{H,1},C_{F,1}} \chi_1 \ln C_{NT,1} + a_1 \ln C_{H,1} + \iota_1 \ln C_{F,1} + \mu_1 \left[CE_1 - C_{NT,1} - P_{H,1}C_{H,1} - P_{F,1}C_{F,1}\right]. \tag{A.1}$$

The static problem for period 0 is defined similarly.

First-order conditions for period 0 are:

$$\frac{\chi_0}{C_{NT,0}} = \mu_0 \implies \frac{\chi_0}{\chi_0} = \mu_0 \implies \mu_0 = 1, \tag{A.2}$$

using market clearing condition $C_{NT,0} = Y_{NT,0}$ and $\chi_t = \chi_t^* = 1,\ \forall t$.

$$\frac{a_0}{C_{H,0}} = \mu_0 P_{H,0} = P_{H,0} \implies a_0 = P_{H,0} C_{H,0}, \tag{A.3}$$

$$\frac{\iota_0}{C_{F,0}} = \mu_0 P_{F,0} = P_{F,0} \implies \iota_0 = P_{F,0} C_{F,0}. \tag{A.4}$$

For period 1, first-order conditions are:

$$\frac{\chi_1}{C_{NT,1}} = \mu_1 \implies \frac{\chi_1}{Y_{NT,1}} = \mu_1, \tag{A.5}$$

$$\frac{a_1}{C_{H,1}} = \mu_1 P_{H,1} \implies a_1 \frac{Y_{NT,1}}{\chi_1} = P_{H,1} C_{H,1}, \tag{A.6}$$

$$\frac{\iota_1}{C_{F,1}} = \mu_1 P_{F,1} \implies \iota_1 \frac{Y_{NT,1}}{\chi_1} = P_{F,1} C_{F,1}, \tag{A.7}$$

using the market clearing condition for the non-tradeable good, $C_{NT,1} = Y_{NT,1}$.

The Lagrangian of the intertemporal problem of a U.S. household is:

$$\begin{aligned}
\mathcal{L} = {} & \chi_0 \ln C_{NT,0} + a_0 \ln C_{H,0} + \iota_0 \ln C_{F,0} + \beta \int_{\mathcal{S}} \pi_s \left[\chi_1 \ln C_{NT,1} + a_1 \ln C_{H,1} + \iota_1 \ln C_{F,1}\right] \\
& + \lambda_0 [q_s \cdot 1 + P_{H,0} Y_{H,0} - C_{NT,0} - P_{H,0} C_{H,0} - P_{F,0} C_{F,0} - b - q_s \cdot x] \\
& + \int_{\mathcal{S}} \lambda_{1,s} \left[P_{H,1} Y_{H,1} + R \cdot b + Y_{NT,1} \cdot x - C_{NT,1} - P_{H,1} C_{H,1} - P_{F,1} C_{F,1}\right],
\end{aligned} \tag{A.8}$$

where π_s denotes the probability of reaching history $s \in \mathcal{S}$ in period 1. Recall that households are choosing consumption levels for each time and history. In other words, they are solving a sequential problem[1]. Therefore, period-1 price, quantities and Lagrange multipliers are

[1] Alternatively, this problem can be rewritten in standard sequential or Arrow-Debreu form. I also could be using (heavier) history-specific notation as, for instance, $C_{NT,1}(s)$. However, I opt to simplify the exposition

history-specific, and the integrals above are over the set of possible paths in period 1.

The first-order conditions associated with the above problem are:

$$
\begin{aligned}
&C_{NT,0}: \; \frac{\chi_0}{C_{NT,0}} = \lambda_0 \implies \frac{\chi_0}{\chi_0} = \lambda_0 \implies \lambda_0 = 1, \\
&C_{H,0}: \; \frac{a_0}{C_{H,0}} = \lambda_0 P_{H,0} = P_{H,0} \implies a_0 = P_{H,0} C_{H,0}, \\
&C_{F,0}: \; \frac{\iota_0}{C_{F,0}} = \lambda_0 P_{F,0} = P_{F,0} \implies \iota_0 = P_{F,0} C_{F,0}, \\
&C_{NT,1}: \; \beta \pi_s \frac{\chi_1}{C_{NT,1}} = \lambda_{1,s}, \text{ for state } s, \\
&C_{H,1}: \; \beta \pi_s \frac{a_1}{C_{H,1}} = \lambda_{1,s} P_{H,1}, \text{ for state } s, \\
&C_{F,1}: \; \beta \pi_s \frac{\iota_1}{C_{F,1}} = \lambda_{1,s} P_{F,1}, \text{ for state } s, \\
&b: \; \lambda_0 = 1 = \int_{\mathbb{S}} \lambda_{1,s} R \implies \int_{\mathbb{S}} \lambda_{1,s} = \frac{1}{R}, \\
&x: \; \lambda_0 \cdot q_s = q_s = \int_{\mathbb{S}} \lambda_{1,s} Y_{NT,1}.
\end{aligned}
\tag{A.9}
$$

Using the market clearing condition for the non-tradeable good:

$$
\beta \pi_s \frac{\chi_1}{Y_{NT,1}} = \lambda_{1,s}.
$$

Integrating this expression over the states we have:

$$
\beta \chi_1 \int_{\mathbb{S}} \frac{\pi_s}{Y_{NT,1}} = \mathbb{E}\left[\frac{\beta \chi_1}{Y_{NT,1}}\right] = \frac{1}{R}, \tag{A.10}
$$

the expression for the U.S. bond price in the main text. For the US stock price, we have:

$$
q_s = \int_{\mathbb{S}} \lambda_{1,s} Y_{NT,1} = \int_{\mathbb{S}} \beta \pi_s \frac{\chi_1}{Y_{NT,1}} Y_{NT,1} = \beta \chi_1, \tag{A.11}
$$

the expression in the main text.

here, also following that of the main text.

Japanese Households

Symmetry implies that we can derive analogous conditions from the static problem of Japanese households:

$$a_0^* = P_{F,0}^* C_{F,0}^*, \tag{A.12}$$

$$\xi_0 = P_{H,0}^* C_{H,0}^*, \tag{A.13}$$

$$a_1^* \frac{Y_{NT,1}^*}{\chi_1^*} = P_{F,1}^* C_{F,1}^*, \tag{A.14}$$

$$\xi_1 \frac{Y_{NT,1}^*}{\chi_1^*} = P_{H,1}^* C_{H,1}^*, \tag{A.15}$$

where the right hand side of (A.15), for example, denotes the yen-value of US exports.

The same applies for asset prices, with:

$$q_b^* \equiv \frac{1}{R^*} = \mathbb{E}\left[\frac{\beta^* \chi_1^*}{Y_{NT,1}^*}\right]. \tag{A.16}$$

for the Japanese bond, and:

$$q_s^* = \beta^* \chi_1^*, \tag{A.17}$$

for the Japanese stock.

A.2 The Financiers' Optimisation

This appendix provides details on how international financiers determine their optimal asset positions in the model. These agents operate competitively, have no initial wealth, and choose portfolios subject to a zero-investment constraint. Their objective is to maximise the expected dollar value of terminal wealth, penalised by its variance.

Problem Statement

Let the vector of tradable assets held by financiers be denoted by:

$$\hat{\boldsymbol{\Theta}} \equiv \begin{bmatrix} \Theta_b^* \\ \\ \Theta_s \\ \\ \Theta_s^* \end{bmatrix}, \tag{A.18}$$

where Θ_b^* is the dollar value invested in Japanese bonds, and Θ_s, Θ_s^* denote the dollar values invested in U.S. and Japanese equities, respectively. Note that the U.S. bond position, Θ_b, is pinned down by the zero-investment condition:

$$\Theta_b = -(\Theta_b^* + \Theta_s + \Theta_s^*). \tag{A.19}$$

Let the corresponding vector of gross returns on these positions be:

$$\hat{\boldsymbol{\mathcal{R}}} \equiv \begin{bmatrix} R^* \cdot \frac{E_1}{E_0} \\ \\ \mathcal{R} \\ \\ \mathcal{R}^* \cdot \frac{E_1}{E_0} \end{bmatrix}, \tag{A.20}$$

where:

- $\mathcal{R} \equiv \frac{Y_{NT,1}}{q_s}$ is the gross return on U.S. equity in dollars,
- $\mathcal{R}^* \equiv \frac{Y^*_{NT,1}}{q^*_s}$ is the gross return on Japanese equity in yen,
- $\frac{E_1}{E_0}$ converts yen payoffs into dollars.

Let R be the U.S. risk-free rate. Then, the dollar value of terminal wealth (i.e., dollar payoff at $t = 1$) is given by:

$$\Pi_1 = R \cdot \Theta_b + \hat{\boldsymbol{\mathcal{R}}}^\top \hat{\boldsymbol{\Theta}}. \tag{A.21}$$

The financier's mean-variance objective is:

$$\max_{\hat{\boldsymbol{\Theta}}} \quad \mathbb{E}[\Pi_1] - \frac{a}{2}\text{Var}[\Pi_1], \tag{A.22}$$

subject to the zero-investment constraint:

$$\Theta_b + \Theta^*_b + \Theta_s + \Theta^*_s = 0. \tag{A.23}$$

A.2.1 Solution Methodology

This optimisation problem is solved numerically via symbolic and numerical computation. All variables are functions of the model's underlying parameters, including the volatility of equity markets, the correlation between equity returns, and the preference parameters of households. Key steps include:

1. Expressing the objective entirely in terms of $\hat{\boldsymbol{\Theta}}$ using the zero-investment constraint to eliminate Θ_b.
2. Substituting for $\mathbb{E}[\Pi_1]$ and $\text{Var}[\Pi_1]$ based on the log-normal distribution of dividends.

3. Solving for the vector $\hat{\boldsymbol{\Theta}}$ that satisfies the first-order conditions of the mean-variance problem, subject to the feasibility constraints implied by such vector — i.e. that exchange rates, consumption levels and prices of tradeable goods are non-negative across the states, evaluated at the financiers' optimal portfolio. To reduce dimensionality associated with this problem, I use the Λ_i's in Appendix A.4.3 as state variables that must lie between 0 and 1 — it can be shown that all consumption levels and prices in this model are positive if, and only if all of these Λ_i's lie inside that interval.

A.3 Solving the Financial Autarky Model

This section solves the model under financial autarky, $\Theta_b = \Theta_b^* = \Theta_s = \Theta_s^*$.

Asset prices are defined by the households' Euler equations and are therefore unaffected by the absence of international portfolio flows:

$$q_b \equiv \frac{1}{R} = \mathbb{E}\left[\frac{\beta \chi_1}{Y_{NT,1}}\right], \tag{A.24}$$

$$q_s = \beta \chi_1, \tag{A.25}$$

$$q_b^* \equiv \frac{1}{R^*} = \mathbb{E}\left[\frac{\beta^* \chi_1^*}{Y_{NT,1}^*}\right], \tag{A.26}$$

$$q_s^* = \beta^* \chi_1^*. \tag{A.27}$$

Exchange rates are such that:

$$\begin{aligned} E_0 &= \frac{\iota_0}{\xi_0}, \\ E_1 &= \frac{\iota_1}{\xi_1} \cdot \frac{Y_{NT,1}}{Y_{NT,1}^*} \cdot \frac{\chi_1^*}{\chi_1}. \end{aligned} \tag{A.28}$$

Using market clearing condition $Y_{NT,0} = C_{NT,0}$, rewrite the U.S. households' budget constraint as:

$$\begin{aligned} P_{H,0} Y_{H,0} &= P_{H,0} C_{H,0} + P_{F,0} C_{F,0} + b + q \cdot (x - 1), \\ &= a_0 + \iota_0 + b + q \cdot (x - 1), \\ &= a_0 + \iota_0 + b + \beta \cdot \chi_1 \cdot (x - 1), \end{aligned} \tag{A.29}$$

where I used (A.3), (A.4), and the price expression for the U.S. stock. For period 1, using the market clearing $C_{NT,1} = Y_{NT,1}$, we can rewrite the period-1 budget constraint as:

$$\begin{aligned} &P_{H,1}Y_{H,1} + R \cdot b + Y_{NT,1} \cdot x \\ =&Y_{NT,1} + P_{H,1}C_{H,1} + P_{F,1}C_{F,1} \\ =&Y_{NT,1} + a_1 \frac{Y_{NT,1}}{\chi_1} + \iota_1 \frac{Y_{NT,1}}{\chi_1} \\ =&Y_{NT,1} + \frac{Y_{NT,1}}{\chi_1} \cdot (a_1 + \iota_1), \end{aligned} \tag{A.30}$$

using (A.6) and (A.7).

Substitute (A.24) into this expression and rearrange to get:

$$\begin{aligned} P_{H,1}Y_{H,1} =&Y_{NT,1} \left[\frac{a_1 + \iota_1}{\chi_1} + 1 - x \right] - \frac{b}{\beta \cdot \chi_1 \cdot \mathbb{E}\left[\frac{1}{Y_{NT,1}}\right]} \\ =&Y_{NT,1} \left[\frac{a_1 + \iota_1}{\chi_1} \right] - \frac{b}{\beta \cdot \chi_1 \cdot \mathbb{E}\left[\frac{1}{Y_{NT,1}}\right]} \end{aligned} \tag{A.31}$$

where I used $x = 1$, the market clearing condition for the U.S. stock in the absence of global financiers (i.e., U.S. households hold all the U.S. stocks).

Now turning to the Japanese problem, symmetry in period 0 implies:

$$P^*_{F,0}Y_{F,0} = \xi_0 + a^*_0 + b^* + \beta^* \chi^*_1 \cdot (x^* - 1). \tag{A.32}$$

In period 1, we have for Japanese households:

$$P^*_{F,1}Y^*_{F,1} + R^* \cdot b^* + Y^*_{NT,1} \cdot x^* + \underbrace{\Pi^*_1}_{\text{Profits (in yen)}} = C^*_{NT,1} + \underbrace{P^*_{H,1}C^*_{H,1}}_{=\frac{\xi_1 \cdot Y^*_{NT,1}}{\chi^*_1}} + \underbrace{P^*_{F,1}C^*_{F,1}}_{=\frac{a^*_1 \cdot Y^*_{NT,1}}{\chi^*_1}} \tag{A.33}$$

Under financial autarky, rebated profits are 0, so we have a similar expression to that for

US households:

$$P^*_{F,1}Y_{F,1} = Y^*_{NT,1}\left[\frac{a^*_1+\xi_1}{\chi^*_1}+1-x^*\right] - \frac{b^*}{\beta^*\chi^*_1\mathbb{E}\left[\frac{1}{Y^*_{NT,1}}\right]}$$
$$= Y^*_{NT,1}\left[\frac{a^*_1+\xi_1}{\chi^*_1}\right] - \frac{b^*}{\beta^*\chi^*_1\mathbb{E}\left[\frac{1}{Y^*_{NT,1}}\right]}, \tag{A.34}$$

using the market clearing condition for Japanese stocks, $x^* = 1$, in the absence of financiers.

We also have the market clearing conditions for tradeable goods:

$$\begin{aligned} C_{H,0} + C^*_{H,0} &= Y_{H,0}, \\ C_{F,0} + C^*_{F,0} &= Y^*_{F,0}, \\ C_{H,1} + C^*_{H,1} &= Y_{H,1}, \\ C_{F,1} + C^*_{F,1} &= Y^*_{F,1}, \end{aligned} \tag{A.35}$$

while the law of one price implies:

$$\begin{aligned} P_{H,0} &= P^*_{H,0}E_0, \\ P_{F,0} &= P^*_{F,0}E_0, \\ P_{H,1} &= P^*_{H,1}E_1, \\ P_{F,1} &= P^*_{F,1}E_1. \end{aligned} \tag{A.36}$$

Use market clearing conditions $b = 0$ and $x = 1$ in (A.29):

$$P_{H,0}Y_{H,0} = a_0 + i_0, \tag{A.37}$$

implying:

$$P_{H,0} = \frac{a_0 + \iota_0}{Y_{H,0}}. \tag{A.38}$$

Then (A.3) becomes:

$$C_{H,0} = \frac{a_0}{a_0 + \iota_0} Y_{H,0}. \tag{A.39}$$

With $b^* = 0$ and $x^* = 1$, symmetry in the Japanese household's problem implies:

$$P^*_{F,0} = \frac{\xi_0 + a^*_0}{Y_{F,0}}. \tag{A.40}$$

Then use the second equation in (A.36) to get:

$$P_{F,0} = P^*_{F,0} \cdot E_0 = \frac{\xi_0 + a^*_0}{Y_{F,0}} \cdot \frac{\iota_0}{\xi_0}. \tag{A.41}$$

Next, (A.4) implies:

$$C_{F,0} = \frac{\iota_0}{P_{F,0}} = \frac{\xi_0}{\xi_0 + a^*_0} Y^*_{F,0}. \tag{A.42}$$

Then:

$$P_{H,0} = P^*_{H,0} \cdot E_0 \implies P^*_{H,0} = \frac{P_{H,0}}{E_0} = \frac{a_0 + \iota_0}{Y_{H,0}} \cdot \frac{\xi_0}{\iota_0}. \tag{A.43}$$

And:

$$C^*_{H,0} = \frac{\xi_0}{P^*_{H,0}} = \frac{\iota_0}{a_0 + \iota_0} Y_{H,0}, \tag{A.44}$$

Finally, (A.35) gives:

$$C_{F,0} + C^*_{F,0} = Y^*_{F,0} \implies \frac{\xi_0}{\xi_0 + a^*_0} Y^*_{F,0} + C^*_{F,0} = Y^*_{F,0} \implies C^*_{F,0} = \frac{a^*_0}{\xi_0 + a^*_0} Y^*_{F,0}. \tag{A.45}$$

Use $b = 0$ in (A.31):

$$P_{H,1} = \left[\frac{a_1 + \iota_1}{\chi_1}\right] \cdot \frac{Y_{NT,1}}{Y_{H,1}}. \tag{A.46}$$

Fixed the supply of the home tradeable good, states in period 1 for which the US stock pays better dividends ($\uparrow Y_{NT,1}$) inflate the price of the home tradeable good (which will become relatively more scarce).

Next, use (A.6):

$$\begin{aligned} C_{H,1} &= \frac{a_1}{\chi_1} \frac{Y_{NT,1}}{P_{H,1}} = \frac{a_1}{\chi_1} \cdot Y_{NT,1} \cdot \frac{Y_{H,1}}{Y_{NT,1}} \cdot \frac{\chi_1}{a_1 + \iota_1} \\ &= \frac{a_1}{a_1 + \iota_1} \cdot Y_{H,1}. \end{aligned} \tag{A.47}$$

Symmetry implies:

$$P^*_{F,1} = \left[\frac{\xi_1 + a^*_1}{\chi^*_1}\right] \frac{Y^*_{NT,1}}{Y^*_{F,1}}, \tag{A.48}$$

$$C^*_{F,1} = \frac{a^*_1}{\xi_1 + a^*_1} \cdot Y^*_{F,1}. \tag{A.49}$$

Using the law of one price (A.36):

$$
\begin{aligned}
P^*_{H,1} = \frac{P_{H,1}}{E_1} =& \frac{1}{E_1} \cdot \left[\frac{a_1 + \iota_1}{\chi_1}\right] \cdot \frac{Y_{NT,1}}{Y_{H,1}} \\
=& \frac{\xi_1}{\iota_1} \cdot \frac{Y^*_{NT,1}}{Y_{NT,1}} \cdot \frac{\chi_1}{\chi^*_1} \cdot \left[\frac{a_1 + \iota_1}{\chi_1}\right] \cdot \frac{Y_{NT,1}}{Y_{H,1}} \\
=& \frac{\xi_1}{\iota_1} \cdot \left[\frac{a_1 + \iota_1}{\chi^*_1}\right] \cdot \frac{Y^*_{NT,1}}{Y_{H,1}}.
\end{aligned}
\tag{A.50}
$$

Then (A.15) gives:

$$
\begin{aligned}
C^*_{H,1} &= \frac{\xi_1}{\chi^*_1} \cdot \frac{Y^*_{NT,1}}{P^*_{H,1}} \\
&= \frac{\xi_1}{\chi^*_1} \cdot Y^*_{NT,1} \cdot \frac{\iota_1}{\xi_1} \cdot \left[\frac{\chi^*_1}{(a_1 + \iota_1)}\right] \cdot \frac{Y_{H,1}}{Y^*_{NT,1}} \\
&= \left[\frac{\iota_1}{a_1 + \iota_1}\right] \cdot Y_{H,1}.
\end{aligned}
\tag{A.51}
$$

Similarly:

$$
C_{F,1} = \frac{\xi_1}{\xi_1 + a^*_1} \cdot Y^*_{F,1}, \tag{A.52}
$$

$$
P_{F,1} = \frac{\iota_1}{\xi_1} \cdot \left[\frac{(\xi_1 + a^*_1)}{\chi_1}\right] \cdot \frac{Y_{NT,1}}{Y^*_{F,1}}. \tag{A.53}
$$

Equilibrium Conditions. The equilibrium is characterised by (together with $b = b^* = 0$ and $x = x^* = 1$):

Exchange Rates:

$\mathrm{E}_0 = \frac{\iota_0}{\xi_0}, \quad E_1 = \frac{\iota_1}{\xi_1} \cdot \frac{Y_{NT,1}}{Y^*_{NT,1}} \cdot \frac{\chi^*_1}{\chi_1},$

Period 0: Non-tradeables

$\mathrm{C}_{NT,0} = Y_{NT,0}, \quad C^*_{NT,0} = Y^*_{NT,0},$

Period 0: Home Tradeables (H), Prices and Quantities for U.S. Households

$\mathrm{P}_{H,0} = \frac{a_0+\iota_0}{Y_{H,0}}, \quad C_{H,0} = \frac{a_0}{a_0+\iota_0} Y_{H,0},$

Period 0: Foreign Tradeables (F), Prices and Quantities for U.S. Households

$\mathrm{P}_{F,0} = \frac{(\xi_0+a_0^*)\cdot\iota_0}{\xi_0 Y_{F,0}}, \quad C_{F,0} = \frac{\xi_0}{\xi_0+a_0^*} Y_{F,0},$

Period 0: Prices (in Yen) and Quantities for Japanese Households

$\mathrm{P}^*_{H,0} = \frac{(a_0+\iota_0)\cdot\xi_0}{\iota_0 Y_{H,0}}, \quad C^*_{H,0} = \frac{\iota_0}{a_0+\iota_0} Y_{H,0},$
$P^*_{F,0} = \frac{\xi_0+a_0^*}{Y_{F,0}}, \quad C^*_{F,0} = \frac{a_0^*}{\xi_0+a_0^*} Y_{F,0},$

Period 1: Non-tradeables

$\mathrm{C}_{NT,1} = Y_{NT,1}, \quad C^*_{NT,1} = Y^*_{NT,1},$

Period 1: Home Tradeables (H), Prices and Quantities for U.S. Households

$\mathrm{P}_{H,1} = \frac{(a_1+\iota_1)}{\chi_1} \cdot \frac{Y_{NT,1}}{Y_{H,1}}, \quad C_{H,1} = \frac{a_1}{a_1+\iota_1} Y_{H,1},$

Period 1: Foreign Tradeables (F), Prices and Quantities for U.S. Households

$\mathrm{P}_{F,1} = \frac{\iota_1}{\xi_1} \cdot \frac{\xi_1+a_1^*}{\chi_1} \cdot \frac{Y_{NT,1}}{Y^*_{F,1}}, \quad C_{F,1} = \frac{\xi_1}{\xi_1+a_1^*} Y^*_{F,1},$

Period 1: Prices (in Yen) and Quantities for Japanese Households

$\mathrm{P}^*_{H,1} = \frac{\xi_1}{\iota_1} \cdot \frac{(a_1+\iota_1)}{\chi_1^*} \cdot \frac{Y^*_{NT,1}}{Y_{H,1}}, \quad C^*_{H,1} = \frac{\iota_1}{a_1+\iota_1} Y_{H,1},$
$P^*_{F,1} = \frac{(\xi_1+a_1^*)}{\chi_1^*} \cdot \frac{Y^*_{NT,1}}{Y^*_{F,1}}, \quad C^*_{F,1} = \frac{a_1^*}{\xi_1+a_1^*} Y^*_{F,1}.$

A.4 Solving the Full Model

A.4.1 Market-Clearing Conditions for Assets

With financiers, market clearing conditions for bonds of the two countries become:

$$b + \Theta_b = 0, \tag{A.1}$$

$$b^* + \frac{\Theta_b^*}{E_0} = 0. \tag{A.2}$$

And for stocks we have the following:

$$x + \frac{\Theta_s}{q_s} = 1, \tag{A.3}$$

$$x^* + \frac{\Theta_s^*}{E_0 \cdot q_s^*} = 1. \tag{A.4}$$

A.4.2 Consumption and Prices of Goods

Solving for the financiers' optimal portfolio numerically, the remaining part of the model admits closed-form solution, being solved in a similar fashion to the case of financial autarky. Using market clearing conditions for U.S. bonds and stocks above, we can rewrite (A.31) as:

$$P_{H,1} = P_{H,1} = \frac{Y_{NT,1}\left[\frac{a_1+\iota_1}{\chi_1} + \frac{\Theta_s}{q_s}\right] + \Theta_b \cdot R}{Y_{H,1}}. \tag{A.5}$$

Holding the supply of home tradeables constant, $Y_{H,1}$, positive shocks to the U.S. stock market, $\uparrow Y_{NT,1}$, will push prices of these goods upwards to the extent that tradeables become relatively more scarce compared with non-tradeables. This effect will be stronger the larger the share of total U.S. stocks held by financiers in equilibrium – i.e., the lower the fraction held by the U.S. households, the weaker this relative price mechanism to the extent that a lower fraction of stock dividends will be retained domestically.

From the U.S. households' first-order conditions, it then follows that:

$$C_{H,1} = \frac{a_1}{\chi_1} \cdot \frac{Y_{NT,1}}{P_{H,1}} = \frac{a_1}{\chi_1} \cdot \frac{Y_{NT,1}}{Y_{NT,1}\left[\frac{a_1+\iota_1}{\chi_1} + \frac{\Theta_s}{q_s}\right] + \Theta_b \cdot R} \cdot Y_{H,1} \equiv \Lambda_H(Y_{NT,1}, \boldsymbol{\Theta}) \cdot Y_{H,1}. \tag{A.6}$$

And by using the market clearing condition for U.S. tradeables one can obtain their consumption level by Japanese households:

$$\begin{aligned} C^*_{H,1} = Y_{H,1} - C_{H,1} = &\left[1 - \frac{a_1}{\chi_1} \cdot \frac{Y_{NT,1}}{Y_{NT,1}\left[\frac{a_1+\iota_1}{\chi_1} + \frac{\Theta_s}{q_s}\right] + \Theta_b \cdot R}\right] \cdot Y_{H,1} \\ \equiv &(1 - \Lambda_H(Y_{NT,1}, \boldsymbol{\Theta})) \cdot Y_{H,1}. \end{aligned} \tag{A.7}$$

Next, the history-specific budget constraint of Japanese households in period 1, with asset market clearing conditions, determines $P_{F,1}$:

$$P^*_{F,1} = \frac{Y^*_{NT,1}\left[\frac{a^*_1+\xi_1}{\chi^*_1} + \frac{\Theta^*_s}{E_0 \cdot q^*_s}\right] + R^* \cdot \frac{\Theta^*_b}{E_0} - \Pi^*_1}{Y_{F,1}}, \tag{A.8}$$

where Π^*_1 denotes the financiers' profits in yen (which are rebated to Japanese households). This is equal to:

$$\begin{aligned} \Pi^*_1 = &\frac{\Pi_1}{E_1} = \frac{1}{E_1}\left[-\Omega^{UIP} \cdot \Theta^*_b + \left(\frac{Y_{NT,1}}{q_s} - R\right) \cdot \Theta_s + \Theta^*_s\left(\frac{Y^*_{NT,1}}{q^*_s}\frac{E_1}{E_0} - R\right)\right] \\ = &\frac{\Pi_1}{E_1} = \frac{1}{E_1}\left[-\Omega^{UIP} \cdot \Theta^*_b + (\mathcal{R} - R) \cdot \Theta_s + \Theta^*_s\left(\mathcal{R}^*\frac{E_1}{E_0} - R\right)\right] \\ = &\frac{1}{E_1}\hat{\boldsymbol{\Theta}}^T(\hat{\boldsymbol{\mathcal{R}}} - \mathbf{1}R), \end{aligned} \tag{A.9}$$

where:

$$\hat{\mathbf{\Theta}} \equiv \begin{bmatrix} \Theta_b^* \\ \\ \Theta_s \\ \\ \Theta_s^* \end{bmatrix}, \qquad\qquad \hat{\mathbf{\mathcal{R}}} \equiv \begin{bmatrix} R^* \frac{E_1}{E_0} \\ \\ \mathcal{R} \\ \\ \mathcal{R}^* \frac{E_1}{E_0} \end{bmatrix}. \tag{A.10}$$

To show that the Japanese household's budget constraint binds state by state in period 1, start by rewriting (A.8) using the results above. After some tedious manipulation, it follows that:

$$\begin{aligned} P_{F,1}^* \cdot Y_{F,1} =& Y_{NT,1}^* \left[\frac{a_1^* + \xi_1}{\chi_1^*} + \frac{\Theta_s^*}{E_0 \cdot q_s^*} \right] + R^* \cdot \frac{\Theta_b^*}{E_0} - \Pi_1^* \\ =& Y_{NT,1}^* \cdot \frac{a_1^* + \xi_1}{\chi_1^*} - R \cdot \frac{\Theta_b}{E_1} - \mathcal{R} \cdot \frac{\Theta_s}{E_1}. \end{aligned} \tag{A.11}$$

Multiply both sides of this equation by E_1:

$$E_1 \cdot P_{F,1}^* \cdot Y_{F,1} = E_1 \cdot \frac{\xi_1}{\chi_1^*} \cdot Y_{NT,1}^* + E_1 \cdot \frac{a_1^*}{\chi_1^*} \cdot Y_{NT,1}^* - R \cdot \Theta_b - \mathcal{R} \cdot \Theta_s. \tag{A.12}$$

Combine this with the dollar-flow equation for period 1:

$$E_1 \cdot P_{F,1}^* \cdot Y_{F,1} = E_1 \cdot \frac{a_1^*}{\chi_1^*} \cdot Y_{NT,1}^* + \frac{\iota_1^*}{\chi_1} \cdot Y_{NT,1} = E_1 \cdot P_{F,1}^* \cdot C_{F,1}^* + P_{F,1} \cdot C_{F,1}. \tag{A.13}$$

Finally, using the law of one price for foreign tradeables in (A.36), we can rewrite it as:

$$P_{F,1} \underbrace{\left(Y_{F,1} - C_{F,1}^* \right)}_{= C_{F,1}} = P_{F,1} \cdot C_{F,1}, \tag{A.14}$$

which always holds, given the market-clearing condition for foreign tradeables in (A.35).

A.4.3 General Representation

Using these results, we can succinctly represent the rules the determine prices and consumption of goods in this model by reduced-form parameters. For example, for consumption, we have:

$$\begin{aligned} C_{H,1} &= \Lambda_H(Y_{NT,1}, \mathbf{\Theta}) \cdot Y_{H,1}, \\ C^*_{H,1} &= (1 - \Lambda_H(Y_{NT,1}, \mathbf{\Theta})) \cdot Y_{H,1}, \\ C_{F,1} &= (1 - \Lambda_F(Y_{NT,1}, Y^*_{NT,1}, \mathbf{\Theta})) \cdot Y_{F,1}, \\ C^*_{F,1} &= \Lambda_F(Y_{NT,1}, Y^*_{NT,1}, \mathbf{\Theta}) \cdot Y_{F,1}, \end{aligned} \tag{A.15}$$

where $\Lambda_i \in [0,1]$, for $i = \{H, F\}$. I use these boundary conditions on the implied Λ_i's to guarantee that the numerical solutions found for the financiers' optimality conditions are feasible and well-behaved.

A.5 Deriving the UIP-UEP-Hedging Relationship

For expository reasons, consider $\beta = \beta^* = \chi_1 = \chi_1^*$. Rewrite the UIP deviation as:

$$\begin{aligned}\Omega^{\text{UIP}} &\equiv \mathbb{E}\left[R - R^* \cdot \mathbb{E}\left(\frac{E_1}{E_0}\right)\right] = \mathbb{E}\left(\frac{1}{\mathbb{E}\left[\frac{1}{Y_{NT,1}}\right]}\right) - \mathbb{E}\left(\frac{\mathbb{E}\left[\frac{E_1}{E_0}\right]}{\mathbb{E}\left[\frac{1}{Y^*_{NT,1}}\right]}\right) = \\ &= \frac{1}{\mathbb{E}\left[\frac{1}{Y_{NT,1}}\right]} - \frac{1}{\mathbb{E}\left[\frac{1}{Y^*_{NT,1}}\right]}\mathbb{E}\left[\frac{E_1}{E_0}\right].\end{aligned} \tag{A.16}$$

Next, rewrite the UEP deviation as:

$$\begin{aligned}\Omega^{\text{UEP}} &\equiv \mathbb{E}\left[\mathcal{R} - \mathcal{R}^*\frac{E_1}{E_0}\right] = \mathbb{E}[Y_{NT,1}] - \mathbb{E}\left[Y^*_{NT,1}\frac{E_1}{E_0}\right] \\ &= \mathbb{E}\left[Y_{NT,1}\right] - \mathbb{E}[Y^*_{NT,1}] \cdot \mathbb{E}\left[\frac{E_1}{E_0}\right] - \text{Cov}\left(Y^*_{NT,1}, \frac{E_1}{E_0}\right).\end{aligned} \tag{A.17}$$

To account for the Jensen inequality, note that:

$$\begin{aligned}\mathbb{E}\left[\frac{1}{Y_{NT,1}}\right] &\approx \mathbb{E}\left[\frac{1}{\mathbb{E}[Y_{NT,1}]} - \frac{1}{\mathbb{E}[Y_{NT,1}]^2}(Y_{NT,1} - \mathbb{E}[Y_{NT,1}])\right. \\ &\quad \left. + \frac{1}{\mathbb{E}[Y_{NT,1}]^3}(Y_{NT,1} - \mathbb{E}[Y_{NT,1}])^2\right] \\ &= \mathbb{E}\left[\frac{1}{\mathbb{E}[Y_{NT,1}]} + \frac{1}{\mathbb{E}[Y_{NT,1}]^3}(Y_{NT,1} - \mathbb{E}[Y_{NT,1}])^2\right] \\ &= \frac{1}{\mathbb{E}[Y_{NT,1}]} + \frac{\text{Var}[Y_{NT,1}]}{\mathbb{E}[Y_{NT,1}]^3} \\ &= \frac{1}{\mathbb{E}[Y_{NT,1}]}\left[1 + \frac{\text{Var}[Y_{NT,1}]}{\mathbb{E}[Y_{NT,1}]^2}\right] \\ &= \frac{1}{\mathbb{E}[Y_{NT,1}]} \cdot \frac{\mathbb{E}[Y^2_{NT,1}]}{\mathbb{E}[Y_{NT,1}]^2}.\end{aligned} \tag{A.18}$$

The term $\frac{\mathbb{E}[Y^2_{NT,1}]}{\mathbb{E}[Y_{NT,1}]^2} = 1 + \frac{\text{Var}[Y_{NT,1}]}{\mathbb{E}[Y_{NT,1}]^2}$ on the right hand side of (A.18) is simple for some distributions of the exponential family:

- If $Y_{NT,1} \sim N(\mu, \sigma^2)$, then $1 + \frac{\text{Var}[Y_{NT,1}]}{\mathbb{E}[Y_{NT,1}]^2} = 1 + \frac{\sigma^2}{\mu^2} = \frac{\mu^2+\sigma^2}{\mu^2}$.

- If $\log(Y_{NT,1}) \sim N(\mu, \sigma^2)$, then $1 + \frac{\text{Var}[Y_{NT,1}]}{\mathbb{E}[Y_{NT,1}]^2} = 1 + e^{\sigma^2} - 1 = e^{\sigma^2}$.
- If $Y_{NT,1} \sim exp(\lambda)$, then $1 + \frac{\text{Var}[Y_{NT,1}]}{\mathbb{E}[Y_{NT,1}]^2} = 1 + \frac{\frac{1}{\lambda^2}}{(\frac{1}{\lambda})^2} = 1 + 1 = 2$.
- If $Y_{NT,1} \sim Beta(\alpha, \beta)$, then $1 + \frac{\text{Var}[Y_{NT,1}]}{\mathbb{E}[Y_{NT,1}]^2} = \frac{1+\alpha}{\alpha} \frac{\alpha+\beta}{\alpha+\beta+1}$.

Therefore, using log-normality, for example, we can approximate the UIP deviation formula by:

$$\begin{aligned} \Omega^{\text{UIP}} &\approx \frac{1}{\mathbb{E}\left[\frac{1}{Y_{NT,1}}\right]} - \frac{1}{\mathbb{E}\left[\frac{1}{Y^*_{NT,1}}\right]} \mathbb{E}\left[\frac{E_1}{E_0}\right] \\ &\approx \mathbb{E}\left[Y_{NT,1}\right] \frac{\mathbb{E}[Y_{NT,1}]^2}{\mathbb{E}[Y^2_{NT,1}]} - \mathbb{E}\left[Y^*_{NT,1}\right] \mathbb{E}\left[\frac{E_1}{E_0}\right] \frac{\mathbb{E}[Y^*_{NT,1}]^2}{\mathbb{E}[{Y^*_{NT,1}}^2]}. \end{aligned} \tag{A.19}$$

We can further rewrite (A.19) as:

$$\begin{aligned} \Omega^{\text{UIP}} &\approx \mathbb{E}\left[Y_{NT,1}\right] \frac{\mathbb{E}[Y_{NT,1}]^2}{\mathbb{E}[Y^2_{NT,1}]} - \mathbb{E}\left[Y^*_{NT,1}\right] \mathbb{E}\left[\frac{E_1}{E_0}\right] \frac{\mathbb{E}[Y^*_{NT,1}]^2}{\mathbb{E}[{Y^*_{NT,1}}^2]} \\ &= \left\{\mathbb{E}\left[Y_{NT,1}\right] - \mathbb{E}\left[Y^*_{NT,1}\right] \mathbb{E}\left[\frac{E_1}{E_0}\right]\right\} \frac{\mathbb{E}[Y_{NT,1}]^2}{\mathbb{E}[Y^2_{NT,1}]} \\ &\quad + \mathbb{E}\left[Y^*_{NT,1}\right] \mathbb{E}\left[\frac{E_1}{E_0}\right] \left\{\frac{\mathbb{E}[Y_{NT,1}]^2}{\mathbb{E}[Y^2_{NT,1}]} - \frac{\mathbb{E}[Y^*_{NT,1}]^2}{\mathbb{E}[{Y^*_{NT,1}}^2]}\right\} \\ &= \left\{\Omega^{\text{UEP}} + \text{Cov}\left(\mathcal{R}^*, \frac{E_1}{E_0}\right)\right\} \frac{\mathbb{E}[Y_{NT,1}]^2}{\mathbb{E}[Y^2_{NT,1}]} \\ &\quad + \mathbb{E}\left[Y^*_{NT,1}\right] \mathbb{E}\left[\frac{E_1}{E_0}\right] \left\{\frac{\mathbb{E}[Y_{NT,1}]^2}{\mathbb{E}[Y^2_{NT,1}]} - \frac{\mathbb{E}[Y^*_{NT,1}]^2}{\mathbb{E}[{Y^*_{NT,1}}^2]}\right\}, \end{aligned} \tag{A.20}$$

or equivalently:

$$\begin{aligned} \Omega^{\text{UEP}} \approx\ & \Omega^{\text{UIP}} \frac{\mathbb{E}[Y^2_{NT,1}]}{\mathbb{E}[Y_{NT,1}]^2} - \text{Cov}\left(\mathcal{R}^*, \frac{E_1}{E_0}\right) - \mathbb{E}\left[Y^*_{NT,1}\right] \cdot \mathbb{E}\left[\frac{E_1}{E_0}\right] \cdot \left\{1 - \frac{\mathbb{E}[Y^2_{NT,1}]}{\mathbb{E}[Y_{NT,1}]^2} \cdot \frac{\mathbb{E}[Y^*_{NT,1}]^2}{\mathbb{E}[{Y^*_{NT,1}}^2]}\right\} \\ &= \Omega^{\text{UIP}} \cdot \frac{\mathbb{E}[\mathcal{R}^2]}{\mathbb{E}[\mathcal{R}]^2} - \text{Cov}\left(\mathcal{R}^*, \frac{E_1}{E_0}\right) - \mathbb{E}\left[\mathcal{R}^*\right] \cdot \mathbb{E}\left[\frac{E_1}{E_0}\right] \cdot \left\{1 - \frac{\mathbb{E}[\mathcal{R}^2]}{\mathbb{E}[\mathcal{R}]^2} \cdot \frac{\mathbb{E}[\mathcal{R}^*]^2}{\mathbb{E}[\mathcal{R}^*]^2}\right\}. \end{aligned} \tag{A.21}$$

Finally, in log-normal terms:

$$\Omega^{\text{UEP}} \approx \Omega^{\text{UIP}} \cdot e^{\sigma^2} - \text{Cov}\left(\mathcal{R}^*, \frac{E_1}{E_0}\right) + \mathbb{E}\left[\mathcal{R}^*\right] \cdot \mathbb{E}\left[\frac{E_1}{E_0}\right] \cdot \left(e^{\sigma^2 - \sigma^{*2}} - 1\right). \tag{A.22}$$

A.6 UIP-UEP-Hedging Relationship in a Lucas Orchard

In this section, I demonstrate how the UIP-UEP-hedging relationship, as outlined in the main text and derived above, holds more generally. Specifically, I focus on Lucas Orchard models, which rely on a completely different set of assumptions and pricing mechanisms compared to the model exhibited in the main text. However, it is important to emphasise that the generality of this relationship does not imply that it has been explicitly addressed in the existing literature. In my view, this relationship represents a foundational result – one that has always existed within the theoretical framework, but has not been directly targeted or fully explored in prior studies.

Martin (2011) investigates the behaviour of asset prices and exchange rates in a continuous-time two-tree model built on the Lucas Orchard framework of Martin (2013). In his model, each country is populated with local households that behave in a hand-to-mouth manner, therefore not affecting prices. A central group denoted "jetsetters" can consume goods from both countries.

Jetsetters are the marginal agents pricing both countries' output claims and riskless bonds. Therefore, asset prices will exhibit a more significant correlation than outputs, as shocks will be transmitted from one market to the other through their impact on their marginal utility. In this sense, jetsetters in Martin (2011) exert a similar role as intermediaries in this paper. The two goods are viewed as substitutes by jetsetters, so that changes in relative prices will impact his marginal utilities with respect to each good. Martin (2011), then, defines exchange rates as a ratio of marginal utilities of this central agent with respect to each good.

In contrast to the model of the main text, the model in Martin (2011) allows prices of assets to change depending on the positions of jetsetters (financiers here). However, it does so at the cost of exchange rates that are purely defined by relative prices. In other words, in that model, there does not exist a flow equation that determines the dollar in equilibrium. I will now show how a similar relationship between UIP, UEP and hedging arises in this type of model, despite key differences involving channels of exchange rate determination.

Jetsetters in Martin (2011) optimise based on:

$$\mathbb{E}\int_0^\infty e^{-\rho t}\frac{C_t^{1-\gamma}}{1-\gamma}, \tag{A.23}$$

with:

$$C_t = \left[w^{\frac{1}{\eta}} D_{1t}^{\frac{\eta-1}{\eta}} + (1-w)^{\frac{1}{\eta}} D_{2t}^{\frac{\eta-1}{\eta}}\right]^{\frac{\eta}{\eta-1}}, \tag{A.24}$$

and where e here denotes the exponential, w governs the relative importance of goods 1 and 2 for jetsetters, D_{1t} and D_{2t} denote dividends from trees of countries 1 and 2 and η measures intra-temporal substitution between the two goods. Writing:

$$v(D_{1t}, D_{2t}) \equiv \frac{C_t^{1-\gamma}}{1-\gamma}, \tag{A.25}$$

as the jetsetters' felicity function and $v_i(D_{1t}, D_{2t})$ as their marginal utility with respect to good $i = 1, 2$, Martin (2011) defines exchange rates as:

$$E_t \equiv \frac{v_1(D_{1t}, D_{2t})}{v_2(D_{1t}, D_{2t})}. \tag{A.26}$$

Martin (2011) treats dividends processes as exogenous. More precisely, he defines

$$y_{it} - y_{i0} \equiv \log D_{it} - \log D_{0t} \tag{A.27}$$

as a Levy process, for each i.

Next, define a cumulant-generating function (CGF), $\boldsymbol{c}(\theta_1, \theta_2)$ as:

$$\boldsymbol{c}(\theta_1, \theta_2) = \log \mathbb{E}e^{\theta_1(y_{1,t+1}-y_{1,t})+\theta_2(y_{2,t+1}-y_{2,t})} = \log \mathbb{E}\left[\left(\frac{D_{1,t+1}}{1,t}\right)^{\theta_1}\left(\frac{D_{2,t+1}}{2,t}\right)^{\theta_2}\right]. \tag{A.28}$$

Consider, for expository reasons, that dividend processes follow log-normal processes. In addition, suppose that countries' fundamentals are sufficiently linked in such a way that the CGFs of their dividend processes are supermodular. This rules out the case where dividend processes are negatively correlated, so that one is a hedge for the other, and is a weak assumption. Formally:

$$\frac{\partial^2 \boldsymbol{c}(\theta_1, \theta_2)}{\partial\theta_1 \partial\theta_2} \geq 0. \tag{A.29}$$

Under these assumptions and relying on a small-country limit where $i = 1$ is small and $i = 2$ is large, the expressions derived in Martin (2011) imply[2]:

$$XS_2^* - XS_1 = XS_{B,2}^* + \gamma\sigma^2(1-\kappa) - 2\sigma^2(1-\kappa)(1-\chi). \tag{A.30}$$

In (A.30), XS_2^* denotes the excess return on output claims of country $i = 2$ denominated in foreign currency. This is his paper's equivalent to $\mathbb{E}\left[\mathcal{R}^*\frac{E_1}{E_0} - R\right]$ here. XS_1 measures excess returns on output claims of country $i = 1$ denominated in their own currency. The analogue of this variable in the model of the main text would be $\mathbb{E}\left[\mathcal{R} - R\right]$. Hence, we have that $XS_2^* - XS_1$ in Martin (2011) corresponds to $-\Omega^{\text{UEP}}$ here.

$XS_{B,2}^*$ denotes the risk premium on a perpetuity bond of country $i = 2$ in foreign currency. This measures a failure of the UIP to the extent that it denotes the expected return with a carry-trade strategy that goes long in country $i = 2$. Therefore, in terms of the model of the main text, this terms would represent $-\Omega^{\text{UIP}}$. The variable $\chi \equiv \frac{\eta-1}{\eta}$ is just a convenient form of write so that χ tends to 1 when $\eta \to \infty$ (perfect substitution) and to 0 when $\eta \to 1$ (Cobb-Douglas). The variable σ^2 denotes the variance of dividend processes (assumed the same), while κ denotes the correlation between dividend processes of the two countries (captured by ρ in the model of the main text).

[2] See pages 24 and 25 of that paper.

Let us use these results to rearrange (A.30) while substituting $\tilde{\Omega}^{\text{UEP}} \equiv XS_1 - XS_2^*$ and $\tilde{\Omega}^{\text{UIP}} \equiv -XS_{B,2}^*$, where I am using a tilde to emphasise that these expressions are his model's analogues of the parity deviations derived in the main text. Then:

$$\tilde{\Omega}^{\text{UEP}} = \tilde{\Omega}^{\text{UIP}} - \gamma\sigma^2(1-\kappa) + 2\sigma^2(1-\kappa)(1-\chi). \tag{A.31}$$

Relationship (A.31) shows that in the framework of Martin (2011), similarly to what was shown in A.22, UEP deviations are a direct (and positive) linear function of UIP deviations.

The term $2\sigma^2(1-\kappa)(1-\chi)$ measures the interaction between how dividends and exchange rates operate in the model. To see how, consider a negative shock on D_{2t}. The negative shock in the dividends of country $i = 2$ makes its good relatively more scarse, or in lower supply, what tends to push its price upwards compared to the good of country $i = 1$. This will produce a force towards *appreciation* of country $i = 2$ when it faces negative stock market shocks, similarly to the role of exchange rates as an automatic hedge for stocks in Hau and Hélène Rey (2006). The parameter χ will control this impact of stock market shocks on exchange rate movements to the extent that it controls the jetsetters' willingness to substitute between goods. The lower χ (or the lower η), the less jetsetters want to substitute between goods, the less they will react to negative shocks on D_{2t} by consuming more of good of country $i = 1$. As a consequence, the larger the price effect, the more severe exchange rates will move as a response.

Since $\kappa \in [0, 1]$ from the assumption of linked fundamentals, it follows that $2\sigma^2(1-\kappa)(1-\chi) > 0$, and will be larger the smaller χ is. In this sense, this term captures the same intuition as $-\text{Cov}\left(\mathcal{R}^*, \frac{E_1}{E_0}\right)$ in (A.22). In Martin (2011), a lower substitution of jetsetters between assets linked to the outputs of the two countries will tend to accentuate the negative relation between stock market shocks in a country and appreciations of its currency, and this impact will be more severe if dividend processes are less correlated – also aligned with my results in the main text.

Finally, the term $-\gamma\sigma^2(1-\kappa)$ in (A.31) captures common movements between the two stocks. This term becomes more negative in situations where stocks move in a more distinct manner ($\downarrow \kappa$). In my model, this is captured by the term $\mathbb{E}\left[\mathcal{R}^*\right] \cdot \mathbb{E}\left[\frac{E_1}{E_0}\right] \cdot \left(e^{\sigma^2-\sigma^{*2}} - 1\right)$.

A.7 Decomposing UIP, UEP and Hedging Roles

Let me begin with the dollar market clearing conditions under assumption 1. For period 0:

$$
\begin{aligned}
E_0 &= \iota_0 - \Theta_b - \Theta_s = \iota_0 - (-\Theta_b^* - \Theta_s - \Theta_s^*) - \Theta_s \\
&= \iota_0 + \Theta_b^* + \Theta_s^*,
\end{aligned} \tag{A.1}
$$

showing that the time-0 dollar market clearing condition can be either written in terms of dollar positions (in bonds and stocks) or in terms of yen positions (in bonds and stocks). For period 1:

$$
\begin{aligned}
E_1 &= \frac{\iota_1 Y_{NT,1} + R\Theta_b + \mathcal{R}\Theta_s}{Y_{NT,1}^*} \\
&= \frac{\iota_1 Y_{NT,1} - R\Theta_b^* - R\Theta_s^* + (\mathcal{R} - R)\Theta_s}{Y_{NT,1}^*}.
\end{aligned} \tag{A.2}
$$

Note that I used $0 = \Theta_b + \Theta_s + \Theta_b^* + \Theta_s^*$, assuming zero-investment. Combining both conditions:

$$
\begin{aligned}
\frac{E_1}{E_0} =& \frac{\iota_1 Y_{NT,1} + R\Theta_b + \mathcal{R}\Theta_s}{(\iota_0 - \Theta_b - \Theta_s)Y_{NT,1}^*} \\
=& \frac{R\Theta_b}{\iota_0 - \Theta_b - \Theta_s}\frac{1}{Y_{NT,1}^*} + \frac{\iota_1 + \Theta_s}{\iota_0 - \Theta_b - \Theta_s}\frac{Y_{NT,1}}{Y_{NT,1}^*}.
\end{aligned} \tag{A.3}
$$

Next, use that $\Omega^{UIP} \approx \mathbb{E}(Y_{NT,1})\frac{\mathbb{E}^2(Y_{NT,1})}{\mathbb{E}(Y_{NT,1}^2)} - \mathbb{E}(Y_{NT,1}^*)\mathbb{E}\left(\frac{E_1}{E_0}\right)\frac{\mathbb{E}^2(Y_{NT,1}^*)}{\mathbb{E}({Y_{NT,1}^*}^2)}$. Under log-normality:

$$
\Omega^{UIP} \approx \mathbb{E}(Y_{NT,1})e^{-\sigma^2} - \mathbb{E}(Y_{NT,1}^*)\mathbb{E}\left(\frac{E_1}{E_0}\right)e^{-\sigma^{*2}}.
$$

Using the expression for an inter-period dollar depreciation above:

$$
\mathbb{E}\left(\frac{E_1}{E_0}\right) \approx \frac{\Theta_b}{\iota_0 - \Theta_b - \Theta_s}\frac{\mathbb{E}(Y_{NT,1})}{\mathbb{E}(Y_{NT,1}^*)}(e^{\sigma^{*2} - \sigma^2}) + \frac{\iota_1 + \Theta_s}{\iota_0 - \Theta_b - \Theta_s}\frac{\mathbb{E}(Y_{NT,1})}{\mathbb{E}(Y_{NT,1}^*)}(e^{\sigma^{*2}} - e^{\rho\sigma\sigma^*} + 1). \tag{A.4}
$$

Collecting these results, we get the expression for UIP deviations shown in the main text:

$$
\begin{aligned}
\Omega^{UIP} \approx & \underbrace{\frac{\Theta_b}{\Theta_b+\Theta_s-\iota_0}\left[2e^{-\sigma^2}\mathbb{E}(Y_{NT,1})\right]}_{\text{Bond Effect}} \\
& \underbrace{+\frac{\Theta_s}{\Theta_b+\Theta_s-\iota_0}\mathbb{E}(Y_{NT,1})\left[e^{-\sigma^2}+e^{-\sigma^{*2}}+1-e^{\rho\sigma\sigma^*-\sigma^{*2}}\right]}_{\text{Equity Flow Effect}} \\
& \underbrace{-\frac{\iota_0}{\Theta_b+\Theta_s-\iota_0}\mathbb{E}(Y_{NT,1})\left[e^{-\sigma^2}-\frac{\iota_1}{\iota_0}\left(e^{-\sigma^{*2}}+1-e^{\rho\sigma\sigma^*-\sigma^{*2}}\right)\right]}_{\text{Current Account Impact}}.
\end{aligned} \tag{A.5}
$$

To derive the expression for the UEP deviation, start with the following UIP-UEP relationship, proved in the last section:

$$
\Omega^{UEP} \approx \Omega^{UIP}e^{\sigma^2} - \mathrm{Cov}\left(\mathcal{R}^*, \frac{\mathrm{E}_1}{\mathrm{E}_0}\right) + \mathbb{E}(\mathcal{R}^*)\mathbb{E}\left(\frac{E_1}{E_0}\right)(e^{\sigma^2-\sigma^{*2}}-1), \tag{A.6}
$$

Let us work out each of the terms. First:

$$
\begin{aligned}
\Omega^{UIP}e^{\sigma^2} \approx & \frac{\Theta_b}{\iota_0-\Theta_b-\Theta_s}\left[2\mathbb{E}(Y_{NT,1})\right] \\
& + \frac{1}{\iota_0-\Theta_b-\Theta_s}\mathbb{E}(Y_{NT,1})\left[(\Theta_s-\iota_0)+(\iota_1+\Theta_s)\left(e^{\sigma^2-\sigma^{*2}}+e^{\sigma^2}-e^{\rho\sigma\sigma^*+\sigma^2-\sigma^{*2}}\right)\right].
\end{aligned} \tag{A.7}
$$

Second:

$$
\begin{aligned}
-\mathrm{Cov}\left(\mathcal{R}^*, \frac{E_1}{E_0}\right) \approx & \left(\frac{\Theta_b}{\iota_0-\Theta_b-\Theta_s}\right)\mathbb{E}[Y_{NT,1}]\left(e^{\sigma^{*2}-\sigma^2}-e^{-\sigma^2}\right) + \\
& \left(\frac{\iota_1+\Theta_s}{\iota_0-\Theta_b-\Theta_s}\right)\mathbb{E}[Y_{NT,1}]\left(e^{\sigma^2}-e^{\rho\sigma\sigma^*}\right).
\end{aligned} \tag{A.8}
$$

Finally, the last is:

$$
\begin{aligned}
\mathbb{E}(\mathcal{R}^*)\mathbb{E}\left(\frac{E_1}{E_0}\right)(e^{\sigma^2-\sigma^{*2}}-1) \approx & \frac{\Theta_b}{\iota_0-\Theta_b-\Theta_s}\mathbb{E}(Y_{NT,1})\left(1-e^{\sigma^{*2}-\sigma^2}\right) \\
+ \frac{\iota_1+\Theta_s}{\iota_0-\Theta_b-\Theta_s}\mathbb{E}(Y_{NT,1}) & \left(e^{\sigma^2}-e^{\sigma^{*2}}+e^{\rho\sigma\sigma^*}-e^{\rho\sigma\sigma^*+\sigma^2-\sigma^{*2}}+e^{\sigma^2-\sigma^{*2}}-1\right).
\end{aligned} \tag{A.9}
$$

Collecting all of these results and using them in (A.6), we get to:

$$
\begin{aligned}
\Omega^{UEP} \approx & \underbrace{\frac{\Theta_b}{\Theta_b + \Theta_s - \iota_0}\mathbb{E}(Y_{NT,1})\left(2 + e^{-\sigma^2} - e^{\sigma^{*2} - \sigma^2}\right)}_{\text{Bond Effect}} \\
& \underbrace{+ \frac{\Theta_s}{\Theta_b + \Theta_s - \iota_0} 2\mathbb{E}(Y_{NT,1})}_{\text{Equity Flow Impact}} \\
& \underbrace{- \frac{\iota_0}{\Theta_b + \Theta_s - \iota_0}\mathbb{E}(Y_{NT,1})\left(\frac{\iota_0 - \iota_1}{\iota_0}\right)}_{\text{Current Account Impact}},
\end{aligned}
\tag{A.10}
$$

the expression for UEP deviations in the main text.

Next, I shall derive the expressions for hedging roles in the model. Using the dollar market clearing conditions, it follows that:

$$
\begin{aligned}
\mathrm{Cov}\left(\mathcal{R} - \mathcal{R}^*, \frac{E_1}{E_0}\right) = & \left(\frac{R\Theta_b}{\iota_0 - \Theta_b - \Theta_s}\right)\left(\mathbb{E}\left[\frac{Y_{NT,1}}{Y^*_{NT,1}}\right] - \mathbb{E}[Y_{NT,1}]\mathbb{E}\left[\frac{1}{Y^*_{NT,1}}\right]\right) \\
& + \left(\frac{\iota_1 + \Theta_s}{\iota_0 - \Theta_b - \Theta_s}\right)\left(\mathbb{E}\left[\frac{Y^2_{NT,1}}{Y^*_{NT,1}}\right] - \mathbb{E}[Y_{NT,1}]\mathbb{E}\left[\frac{Y_{NT,1}}{Y^*_{NT,1}}\right]\right) \\
& - \left(\frac{R\Theta_b}{\iota_0 - \Theta_b - \Theta_s}\right)\left(1 - \mathbb{E}[Y^*_{NT,1}]\mathbb{E}\left[\frac{1}{Y^*_{NT,1}}\right]\right) \\
& - \left(\frac{\iota_1 + \Theta_s}{\iota_0 - \Theta_b - \Theta_s}\right)\left(\mathbb{E}\left[Y_{NT,1}\right] - \mathbb{E}[Y^*_{NT,1}]\mathbb{E}\left[\frac{Y_{NT,1}}{Y^*_{NT,1}}\right]\right).
\end{aligned}
\tag{A.11}
$$

Based on this expression, keep relying on second-order approximations around dividend means to derive:

$$
\mathbb{E}\left[\frac{Y_{NT,1}}{Y^*_{NT,1}}\right] - \mathbb{E}[Y_{NT,1}]\mathbb{E}\left[\frac{1}{Y^*_{NT,1}}\right] \approx \frac{\mathbb{E}[Y_{NT,1}]}{\mathbb{E}[Y^*_{NT,1}]}(1 - e^{\rho\sigma\sigma^*}),
\tag{A.12}
$$

$$
\mathbb{E}\left[\frac{Y^2_{NT,1}}{Y^*_{NT,1}}\right] - \mathbb{E}[Y_{NT,1}]\mathbb{E}\left[\frac{Y_{NT,1}}{Y^*_{NT,1}}\right] \approx \frac{\mathbb{E}^2[Y_{NT,1}]}{\mathbb{E}[Y^*_{NT,1}]}(1 + e^{\sigma^2} + e^{\sigma^{*2}} - 2e^{\rho\sigma\sigma^*}),
\tag{A.13}
$$

$$1-\mathbb{E}[Y^*_{NT,1}]\mathbb{E}\left[\frac{1}{Y^*_{NT,1}}\right]\approx 1-e^{\sigma^{*2}}, \tag{A.14}$$

$$\mathbb{E}\left[Y_{NT,1}\right]-\mathbb{E}[Y^*_{NT,1}]\mathbb{E}\left[\frac{Y_{NT,1}}{Y^*_{NT,1}}\right]\approx \mathbb{E}[Y_{NT,1}]\left(e^{\rho\sigma\sigma^*}-e^{\sigma^2}\right). \tag{A.15}$$

Additionally, use: $R=\mathbb{E}\left[\frac{1}{Y_{NT,1}}\right]\approx\frac{\mathbb{E}[Y_{NT,1}]}{e^{\sigma^2}}$. Collecting all results and simplifying terms, we can derive the following covariances that govern exchange rate hedging roles:

$$\begin{aligned}
&\mathrm{Cov}\left(\mathcal{R}-\mathcal{R}^*,\frac{E_1}{E_0}\right)\approx\\
&\left(\frac{\Theta_b}{\iota_0-\Theta_b-\Theta_s}\right)\mathbb{E}[Y_{NT,1}]\left[\frac{\mathbb{E}[Y_{NT,1}]}{\mathbb{E}[Y^*_{NT,1}]}\left(e^{-\sigma^2}-e^{\rho\sigma\sigma^*-\sigma^2}\right)+e^{\sigma^{*2}-\sigma^2}-e^{-\sigma^2}\right]+\\
&\left(\frac{\iota_1+\Theta_s}{\iota_0-\Theta_b-\Theta_s}\right)\mathbb{E}[Y_{NT,1}]\left[\frac{\mathbb{E}[Y_{NT,1}]}{\mathbb{E}[Y^*_{NT,1}]}\left(e^{\sigma^{*2}}-e^{\rho\sigma\sigma^*}\right)+e^{\sigma^2}-e^{\rho\sigma\sigma^*}\right],
\end{aligned} \tag{A.16}$$

$$\begin{aligned}
\mathrm{Cov}\left(\mathcal{R},\frac{E_1}{E_0}\right)\approx&\left(\frac{\Theta_b}{\iota_0-\Theta_b-\Theta_s}\right)\frac{\mathbb{E}^2[Y_{NT,1}]}{\mathbb{E}[Y^*_{NT,1}]}\left(e^{-\sigma^2}-e^{\rho\sigma\sigma^*-\sigma^2}\right)+\\
&\left(\frac{\iota_1+\Theta_s}{\iota_0-\Theta_b-\Theta_s}\right)\frac{\mathbb{E}^2[Y_{NT,1}]}{\mathbb{E}[Y^*_{NT,1}]}\left(e^{\sigma^{*2}}-e^{\rho\sigma\sigma^*}\right),
\end{aligned} \tag{A.17}$$

$$\begin{aligned}
\mathrm{Cov}\left(\mathcal{R}^*,\frac{E_1}{E_0}\right)\approx&\left(\frac{\Theta_b}{\iota_0-\Theta_b-\Theta_s}\right)\mathbb{E}[Y_{NT,1}]\left[-e^{\sigma^{*2}-\sigma^2}+e^{-\sigma^2}\right]+\\
&\left(\frac{\iota_1+\Theta_s}{\iota_0-\Theta_b-\Theta_s}\right)\mathbb{E}[Y_{NT,1}]\left[-e^{\sigma^2}+e^{\rho\sigma\sigma^*}\right].
\end{aligned} \tag{A.18}$$

A.8 Equilibrium Characterisation: Discussion and Proofs

Proof of Proposition 5. Substitute $\Theta_b = \Theta_s = 0$ into equations (1.20) and (1.21). This eliminates the bond and equity flow terms, leaving only the current account contributions, which simplify as stated. $\square$

Proof of Proposition 6. Use the expression derived in (1.22) with $\rho = 1$ and $\sigma = \sigma^*$. $\square$

Proof of Proposition 7. Plug $\Theta_b = \Theta_s = 0$ into equations (1.22), (1.23), and (1.24), and evaluate the sign conditions as functions of κ and ρ. $\square$

Proof of Proposition 8. Substitute $\Theta_s = \Theta_s^* = 0$ and $\Theta_b = -\Theta_b^*$ into equation (1.20). Under symmetry, we obtain:

$$\Omega^{\text{UIP}} \approx \left(\frac{\Theta_b^*}{1+\Theta_b^*}\right) 2\mathbb{E}(Y_{NT,1})e^{-\sigma^2},$$

with $E_0 = 1 + \Theta_b^* > 0$. Thus, $\text{sign}(\Omega^{\text{UIP}}) = \text{sign}(\Theta_b^*) = -\text{sign}(\Theta_b)$, implying $\Omega^{\text{UIP}} \cdot \Theta_b < 0$. No profitable trade is possible, and financiers opt out. $\square$

Proof of Proposition 9. Use the sign condition on $\Omega^{\text{UIP}} \cdot \Theta_b$ and the UIP expression under $\Theta_s = 0$. For an equilibrium to exist, we must have $\Omega^{\text{UIP}} \cdot \Theta_b > 0$. This produces the stated bound on $\Theta_b > 0$ if $\Omega^{\text{UIP}} > 0$, or on $\Theta_b^* = -\Theta_b > 0$ if $\Omega^{\text{UIP}} < 0$. $\square$

Proof of Proposition 10. Proposition 9 gives the implicit link between UIP deviations and the direction of financiers' carry trades in equilibrium. The result follows from plugging $\iota_0 = \iota_1 = 1$ and writing $\kappa = \sigma/\sigma^*$ in the provided bounds, noting how signs of Θ_b will be linked to those of Ω^{UIP}. $\square$

Proof of Proposition 11. From condition (1.27), we need to find values of κ for which the signs of ψ and ϕ change. First, it is straightforward to verify that $\psi \geq 0$ if, and only if $\kappa^2 + \rho\kappa - 1 \leq 0$ (otherwise, $\psi < 0$). This represents a quadratic inequality whose only positive root is $\kappa(\rho)$. Additionally, note that $\kappa(\rho)$ is decreasing in ρ, with $\kappa(0) = 1$. If $\kappa > \kappa(\rho)$, then $\kappa^2 + \rho\kappa - 1 > 0$, and $\psi < 0$. If $\kappa \leq \kappa(\rho)$, then $\kappa^2 + \rho\kappa - 1 \leq 0$, implying $\psi \geq 0$. Second, we have that $\phi \geq 0$ when $\kappa \geq \frac{1}{\rho}$ and $\phi < 0$ otherwise.

Collecting the results while considering the signs of ψ and ϕ for each interval for κ, it follows that $f = \frac{\phi}{\psi} \leq 0$ when $\kappa \geq \frac{1}{\rho}$, and that $f = \frac{\phi}{\psi} > 0$ when $\kappa < \frac{1}{\rho}$. Next, we need to have in mind that once we solve condition (1.27) for Θ_b, the inequality flips for $\psi < 0$, i.e., when $\kappa > \kappa(\rho)$. Therefore, for κ in this range, we have an *upper* bound rather than a *lower* bound for the financiers' position in the US bond. Note, however, that for the region in which the inequality does not flip sign, the lower bound for Θ_b is irrelevant. This happens since $f(\rho, \kappa, \sigma) < -1$ in that region and, by construction, $\Theta_b \in [-1, 1]$. When stocks are too correlated — e.g., global downturns — the region for Θ_b under which we should observe a dollar hedging role shrinks. □

Proof of Proposition 12. Assume symmetric trade preferences, $\iota_0 = \iota_1 = 1$, and that financiers can only trade bonds ($\Theta_s = 0$). Recall from Proposition 10 that the sign of Ω^{UIP} and the direction of the carry trade depend on whether $\rho\kappa$ is greater or less than one:

- If $\rho\kappa > 1$, then $\Theta_b > 0$, $\Omega^{\mathrm{UIP}} > 0$: the dollar is on the investment side (Equilibrium A).
- If $\rho\kappa < 1$, then $\Theta_b < 0$, $\Omega^{\mathrm{UIP}} < 0$: the yen is on the investment side (Equilibrium B).

I now characterise the hedging roles under each equilibrium.

Funding Currency Hedge (General). From the hedging covariance expressions in Proposition 4, the sign of each covariance depends on the underlying parameters (κ, ρ) and the magnitude of Θ_b. In both equilibria, financiers settle liabilities in the funding currency, purchasing it in period 1. This demand reinforces currency appreciation in local downturns — when the funding country's equity underperforms and imports fall — systematically amplifying the funding currency's hedging role. This pattern is evident in Figures 1.3 and 1.4.

In Figure 1.3, the function $g(\cdot)$ crosses the horizontal axis at $\kappa = \rho$. When the yen is on the funding side ($\Theta_b > 0$), its hedging role is only eliminated if $\kappa \in (0, \rho)$. But Equilibrium A, which financiers select when $\kappa > 1/\rho$, always lies to the right of this interval. Hence, the yen hedging role remains active throughout this region.

In Figure 1.4, the function $f(\cdot)$ intersects the axis at $\kappa = 1/\rho$. When the dollar is the funding currency ($\Theta_b < 0$), the dollar hedging role is only eliminated if $\kappa > 1/\rho$. Yet Equilibrium B — characterised by $\Theta_b < 0$ — is only selected when $\kappa < 1/\rho$. Thus, the condition for elimination is again never satisfied, and the dollar hedging role remains active.

Investment Currency Hedge. Whether the investment currency exhibits a hedging role depends on model parameters — and is asymmetric across equilibria. For the dollar (Equilibrium A), the hedging role is never observed; for the yen (Equilibrium B), it may or may not emerge depending on volatility and correlation.

- In **Equilibrium A** ($\Theta_b > 0$), financiers are long U.S. bonds and short yen bonds. According to equation (1.27) and Proposition 11, the dollar hedging role requires $0 < \Theta_b < f(\rho, \kappa, \sigma^*)$. However, when $\kappa \geq \frac{1}{\rho}$, we have $f(\rho, \kappa, \sigma^*) < 0$, so the condition is never satisfied. A dollar hedging role is therefore ruled out in this region.

- In **Equilibrium B** ($\Theta_b < 0$), financiers are long Japanese bonds and short U.S. bonds. The condition for a yen hedging role is given in equation (1.26). This role is eliminated when $\Theta_b < g(\rho, \kappa, \sigma^*)$, and Figure 1.3 shows that $g(\cdot) \geq 0$ for $\kappa \in (0, \rho]$. Hence, when Japanese equity is not sufficiently safe (i.e., $\kappa \leq \rho$), the yen hedging role is also eliminated.

In both equilibria, financiers' positions consistently support the hedging role of the funding currency, while tending to undermine that of the investment currency. $\square$

A.8.1 Equity-Only Intermediation Case: Discussion and Proofs

This appendix provides the detailed analysis of the model in which financiers can trade equity but not bonds — that is, where $\Theta_b = \Theta_b^* = 0$ and $\Theta_s \neq 0$, $\Theta_s^* \neq 0$. As I show below, this case closely mirrors the financial autarky setting. In particular, the presence of equity-only financiers does not materially alter the hedging behaviour of exchange rates qualitatively. I begin with a benchmark result analogous to Proposition 8.

Proposition 15. *If financiers cannot trade bonds ($\Theta_b = \Theta_b^* = 0$), trade preferences are symmetric ($\iota_0 = \iota_1 = 1$), and financial markets are symmetric ($\sigma = \sigma^*$, $\rho = 1$), then there does not exist an equilibrium in which financiers hold non-zero equity positions and make positive expected profits. The only possible equilibrium is one in which financiers opt out: $\Theta_s = \Theta_s^* = 0$, and $\Omega^{\mathrm{UIP}} = \Omega^{\mathrm{UEP}} = 0$.*

Proof. This follows analogously to the proof of Proposition 8, but applied to the UEP expression. Under full symmetry and no bonds, the optimal position that equates marginal expected returns to marginal risk is $\Theta_s = 0$, and expected profitability is zero. As with the bond-only case, this implies both $\Omega^{\mathrm{UIP}} = 0$ and $\Omega^{\mathrm{UEP}} = 0$ in equilibrium. □

The next result establishes that the hedging roles of exchange rates under equity-only intermediation are identical to those in the financial autarky case. In particular, these properties hold regardless of the exact composition of financiers' equity portfolios.

Proposition 16. *In a model where financiers cannot trade bonds but can trade equity ($\Theta_b = 0, \Theta_s \neq 0$), the automatic, dollar, and yen hedging roles coincide with those described in Proposition 7. Hence, Figures 1.1 and 1.2 apply directly to this case.*

Proof. In this model, Θ_s is bounded by ι_1, due to market clearing for U.S. equity and the U.S. households' period-1 budget constraint.[3] For example, if $\iota_1 = 1$, then $-1 < \Theta_s < 1$, which implies $\iota_1 + \Theta_s = 1 + \Theta_s > 0$. Since $E_0 = \iota_0 - \Theta_b - \Theta_s = \iota_0 - \Theta_s$, and $\Theta_b = 0$, it follows

[3] Alternatively, the bound can be derived from the market clearing condition for Japanese equity and the Japanese households' period-1 budget constraint.

that the ratio $\frac{\iota_1+\Theta_s}{E_0}$ is strictly positive for any feasible Θ_s. Therefore, the signs of the hedging covariances depend only on the exponential terms governed by (σ, σ^*, ρ) — as in the autarky case — and are unaffected by the specific value of Θ_s. □

Economic Intuition. The result above is driven by the fact that, when financiers are restricted to trading only equity, their impact on exchange rate dynamics mimics that of households' import behaviour. Recall that financiers enter the equity market only when $\Omega^{\text{UEP}} \cdot \Theta_s > 0$. Suppose that $\iota_0 > \iota_1$, which generates a force for dollar *appreciation* between periods via declining U.S. import demand. In the absence of bonds, financiers absorb this intertemporal imbalance through equity, setting $\Theta_s > 0$. This can be shown formally by contradiction: if financiers instead opt out ($\Theta_s = 0$), then plugging $\Theta_s = 0$ and $\iota_0 > \iota_1$ into equation (1.21) yields $\Omega^{\text{UEP}} > 0$, which implies a profitable trade opportunity — contradicting the opt-out decision.

As financiers scale up their position — long U.S. equity, short Japanese equity — the second term in equation (1.21) becomes increasingly negative, eventually driving expected profits to zero, tension that determines the equilibrium level of Θ_s.[4]

Offsetting Effects on Hedging. In financial autarky ($\Theta_s = 0$), households fully hold their domestic equity, so when U.S. equity outperforms, they become wealthier and increase imports. This produces a trade-driven depreciation of the dollar, supporting a dollar hedging role. When financiers enter with $\Theta_s > 0$, they partially displace households from their domestic equity, thereby reducing the income channel that drives import changes — weakening the trade-based hedging mechanism.

However, their equity positions also create a need to unwind their foreign holdings in period 1. For example, if $\Theta_s > 0$, financiers are short Japanese equity and must buy yen

[4]Recall that $E_0 = \iota_0 - \Theta_b - \Theta_s > 0$, so with $\Theta_b = 0$, it follows that $\Theta_s < \iota_0$. This makes the denominator of the UEP expression negative, meaning that the second term becomes more negative as Θ_s becomes more positive.

using dollars to settle liabilities. This generates additional dollar depreciation in good states, reinforcing the same effect created by households' imports. In other words, the two channels — trade balance and portfolio unwinding — move in opposite directions, and their effects cancel out in aggregate. This explains why equity-only intermediation preserves the same exchange rate hedging patterns as in the financial autarky case.

A.8.2 Full Model: Proofs

Proof of Proposition 13. The starting point is condition (1.28), which states that a dollar hedging role is observed when:

$$\Theta_b \cdot \psi(\rho, \kappa, \sigma^*) > \phi(\rho, \kappa, \sigma^*) \cdot (\iota_1 + \Theta_s).$$

As in the bond-only case, the signs of ψ and ϕ depend on the relative volatility measure $\kappa = \sigma/\sigma^*$ and the correlation ρ. The sign of ψ is governed by the inequality:

$$\psi(\rho, \kappa, \sigma^*) \geq 0 \iff \kappa^2 + \rho\kappa - 1 \leq 0.$$

This is a quadratic inequality in κ, whose unique positive root is $\kappa(\rho) = \frac{-\rho+\sqrt{\rho^2+4}}{2}$. Hence:

- $\psi > 0$ if $\kappa < \kappa(\rho)$;
- $\psi = 0$ if $\kappa = \kappa(\rho)$;
- $\psi < 0$ if $\kappa > \kappa(\rho)$.

In turn, $\phi \geq 0$ if and only if $\kappa \geq 1/\rho$, and $\phi < 0$ otherwise. Since $f \equiv \phi/\psi$, we obtain the following:

- If $\kappa \leq \kappa(\rho)$, then $\psi > 0$ and $\phi < 0 \implies f < 0$;
- If $\kappa(\rho) < \kappa < 1/\rho$, then $\psi < 0$, $\phi < 0 \implies f > 0$;

- If $\kappa \geq 1/\rho$, then $\psi < 0$, $\phi > 0 \implies f < 0$.

Now, solving the inequality for Θ_b gives:

$$\Theta_b \begin{cases} > f(\rho, \kappa, \sigma^*) \cdot (\iota_1 + \Theta_s) & \text{if } \psi > 0, \\ < f(\rho, \kappa, \sigma^*) \cdot (\iota_1 + \Theta_s) & \text{if } \psi < 0. \end{cases}$$

This leads to the three cases in Proposition 13 (for simplicity, consider $\iota_0 = \iota_1 = 1$):

1. When $\kappa \leq \kappa(\rho)$, the bound is negative. In the bond-only case, the condition was always satisfied because $\Theta_b \in [-1, 1]$ and the threshold was less than -1. In the full model, however, the cutoff is scaled by $\iota_1 + \Theta_s$, which remains non-negative but can reduce the magnitude of the bound. As a result, the dollar hedging role may no longer hold unconditionally — financiers' portfolios now matter.

2. For $\kappa \in (\kappa(\rho), 1/\rho)$, the cutoff is positive, but the sign of ψ is negative, so the inequality reverses: Θ_b must lie below the threshold.

3. When $\kappa \geq 1/\rho$, the sign of ψ is still negative, so the inequality is the same as above. However, the bound is now negative, and a sufficiently negative Θ_b is required for the dollar to hedge U.S. equity.

The condition $\iota_1 + \Theta_s \geq 0$ always holds for all feasible Θ_s. This happens because bounds for Θ_s are the same as those for Θ_b, and also governed by ι_1. As previously discussed for Θ_b, one can combine the budget constraint in period 1 for U.S. households with the market asset clearing condition (now for the U.S. stock) to get feasible bounds for Θ_s, under the condition that period-1 consumption and price of the home tradeable cannot be negative. □

A.8.3 Full Model: Equilibrium Properties

Figures 1.5 and 1.6 offer a detailed visualisation of the regions where dollar and yen hedging roles are active under the full model. These areas are defined by the interaction between financiers' portfolio positions and the structural parameters governing equity volatility (κ), international comovement (ρ), and asset composition (through Θ_b, Θ_s). Below, I unpack these results and highlight how they connect to model equilibria and previous cases.

Dollar Hedging Role and Equilibrium Structure. Each panel in Figure 1.5 plots the region where the inequality in equation (1.28) is satisfied — that is, where the dollar acts as a hedge for U.S. equity risk. The cutoff curves trace values of Θ_b implied by the threshold $f(\rho, \kappa, \sigma^*) \cdot (\iota_1 + \Theta_s)$. Blue regions indicate where the dollar hedging role is active. As expected:

- Higher ρ reduces the dollar hedging region (compare left and right panels), as stronger international equity co-movement weakens the asymmetric response of exchange rates to country-specific equity shocks.
- Higher Θ_s (darker lines) increases the likelihood of dollar hedging when $\kappa < 1/\rho$, but the effect reverses when $\kappa > 1/\rho$.

This reversal reflects the asymmetry in the financier's balance sheet: when they are long both U.S. bonds and stocks, their dollar exposure increases, and unless the bond position becomes more negative, the dollar hedging role may vanish. As discussed in the main text, this reflects an endogenous tension between optimal hedging behaviour and the persistence of hedging properties in equilibrium.[5]

Equilibrium Selection and Hedging Roles. From an equilibrium perspective, if we assume that financiers' positions in the U.S. bond continue to align with the sign of Ω^{UIP} — as

[5]See equation (1.28). The left-hand side reflects the hedging strength; the right-hand side links to the portfolio composition. A shift in equity exposure raises the required bond offset.

they did under bond-only intermediation — then:[6]

- For $\kappa > 1/\rho$, we expect $\Theta_b > 0$, consistent with equilibrium A. In this region, the model predicts no dollar hedging role, which aligns with the figure.
- For $\kappa < 1/\rho$, we expect $\Theta_b < 0$, as in equilibrium B. Here, the dollar typically acts as a hedge — particularly when the bond short position offsets equity exposure.

Yen Hedging Role and Mirrored Asymmetry. Figure 1.6 plots the analogous region for a yen hedging role, defined by the inequality in equation (1.29). Red regions indicate the set of (κ, Θ_b) pairs for which the yen hedging role is active.

This region again shrinks with higher ρ, as co-movement undermines the segmentation needed for currency-specific hedging. But compared to the dollar case, the effect of Θ_s is reversed:

- When $\kappa > \rho$, larger Θ_s supports yen hedging.
- When $\kappa < \rho$, larger Θ_s works against it.

This reflects a fundamental asymmetry driven by the financier's currency of account. Since they maximise in dollars, a positive shock to Japanese equity — while increasing their yen asset value — does not necessarily translate into dollar wealth gains. The model equilibrium naturally counteracts this by appreciating the dollar and depreciating the yen, neutralising the hedge.

Summary of Equilibrium Dynamics. To synthesise, the full model resembles the dynamics verified for the bond-only intermediation case:

[6] In the bond-only case, the sign of Θ_b is directly determined by the sign of Ω^{UIP}. In the full model, where financiers also hold equity, this relationship becomes more complex. However, the qualitative pattern appears to hold under a wide range of parameter values — as discussed in Section 1.6.

- **Funding currency hedging roles** (i.e., the currency that financiers must purchase to settle their bond liabilities) are robust across equilibria. As shown before for the bond-only intermediation case, the act of unwinding a bond liability in bad states reinforces the appreciation of the funding currency. This always holds if the equity market associated with the funding currency for bonds is not too risky, regardless of the financiers' bond positions.

- **Investment currency hedging roles** depend on whether the financiers' bond positions neutralise or reinforce the trade balance mechanism. Again, as previously illustrated, when the investment-side currency belongs to the riskier equity market, financiers' bond holdings often dilute its hedging properties.

A.9 Asymmetric Trade Preferences and the Full Model

I now test the model's simulated properties under asymmetric trade preferences, setting $\iota_0 = 1$ and $\iota_1 = 2$. This introduces a dollar depreciation pressure between periods 0 and 1, as U.S. households exhibit a stronger taste for foreign goods, generating trade-driven current account imbalances. Figures A.1 and A.2 confirm that the core patterns for hedging roles documented in the main text remain intact: hedging behaviour deteriorates with greater equity market correlation (ρ), while an increase in κ (i.e., relatively riskier U.S. equity) improves the foreign currency's hedging role and weakens the dollar's.

The dollar depreciation pressure induced by asymmetric preferences lowers both UIP and UEP deviations. These are now generally negative — consistent with investors favouring long bond and equity positions in the foreign market — and larger in magnitude than those found under balanced trade preferences. This shift is intuitive and in line with the model's closed-form expressions (see equation 1.20).

There are two notable differences in exchange rate behaviour under $\iota_0 = 1, \iota_1 = 2$:

1. In Figure A.1, the dollar is not only more depreciated in expectation, but becomes *increasingly* depreciated as equity market correlation rises — the opposite pattern from the baseline case.

2. In Figure A.2, expected dollar depreciation now *declines* with κ, whereas it increased under symmetric preferences.

These differences stem from the altered incentive structure faced by intermediaries. With stronger import preferences ($\iota_1 > \iota_0$), financiers have greater incentive to hold yen-denominated risk. Their balance sheets exhibit increased exposure to Japanese assets — as confirmed by the top panels of both figures, which show net long positions in both Japanese bonds and equities.

As equity correlation increases, financiers reduce their exposure to Japanese stocks (due to diminished diversification benefits) and shift toward Japanese bonds. Yet, because they price risk in dollars, they also retrench globally. This retrenchment lowers the risk premium on the yen, strengthening the yen today and depreciating the dollar — explaining point (i) above.

Finally, since financiers are generally long in yen bonds and short in dollar bonds, their balance sheet increasingly contributes to a dollar hedging role — as explained in the main paper. When U.S. equity becomes especially risky (i.e., κ rises), this position sustains a dollar appreciation in bad states more robustly than in the symmetric benchmark. As a result, the dollar tends to appreciate more sharply in downturns and is expected to appreciate more tomorrow — explaining point (ii).

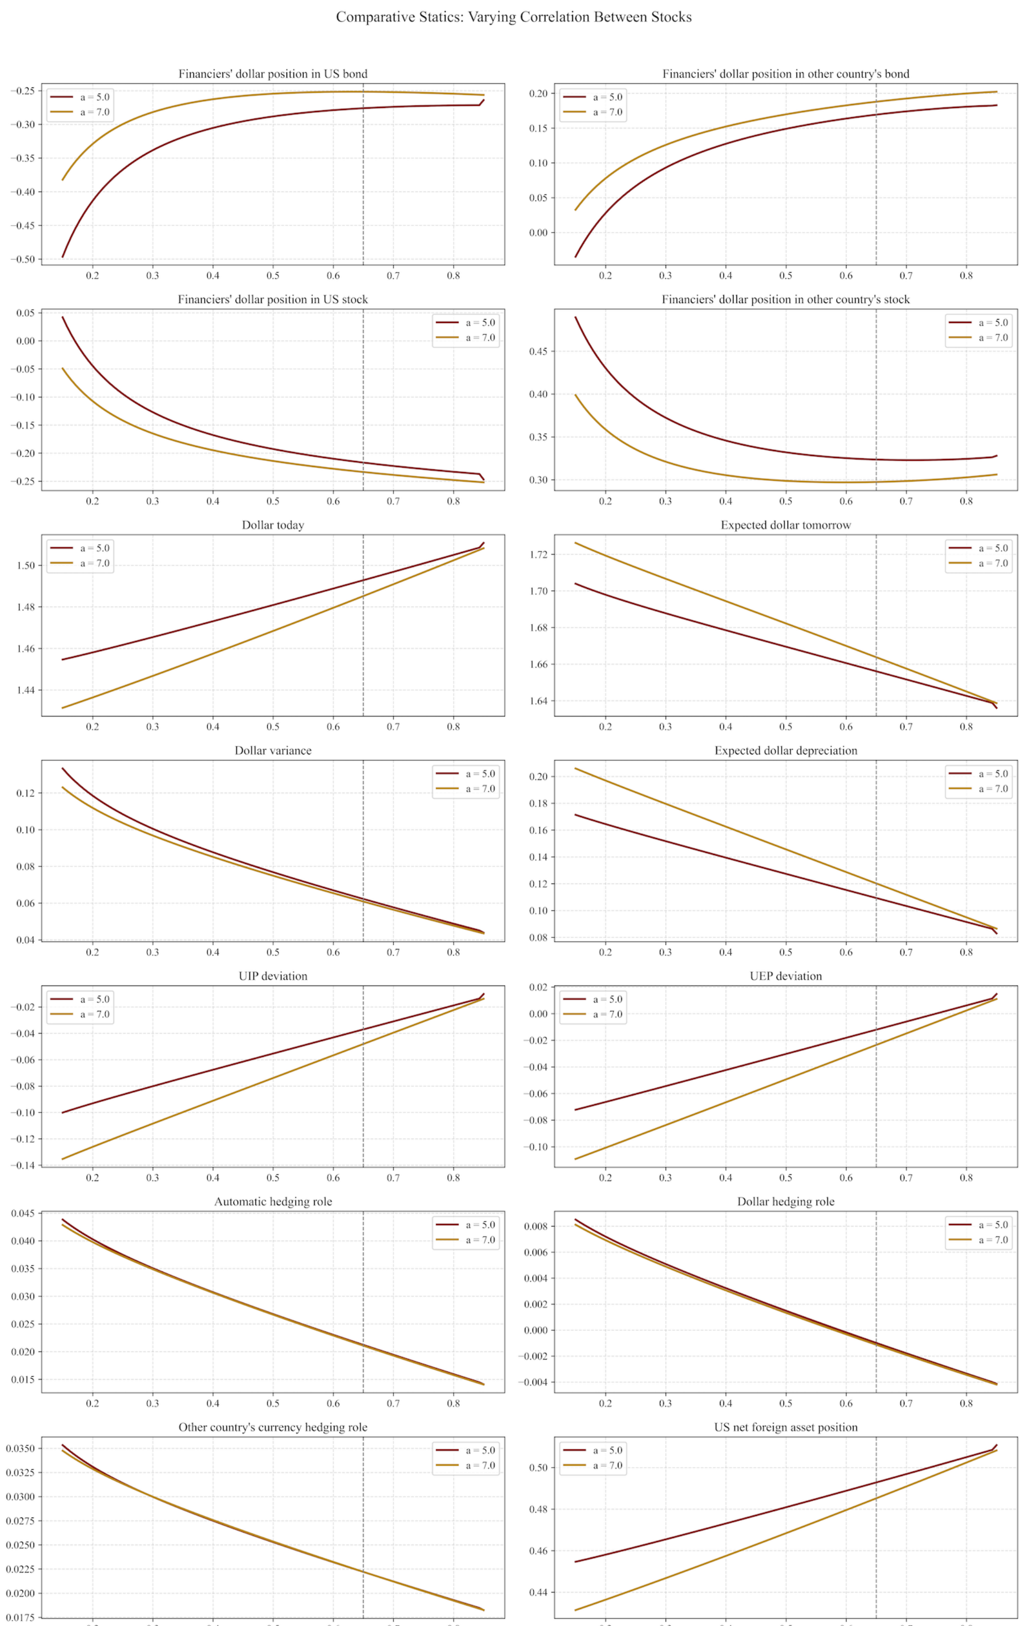

Figure A.1: Comparative Statics: Varying Equity Correlation (ρ). $\iota_0 = 1$ and $\iota_1 = 2$. Remaining parameters follow the same calibration of the main paper.

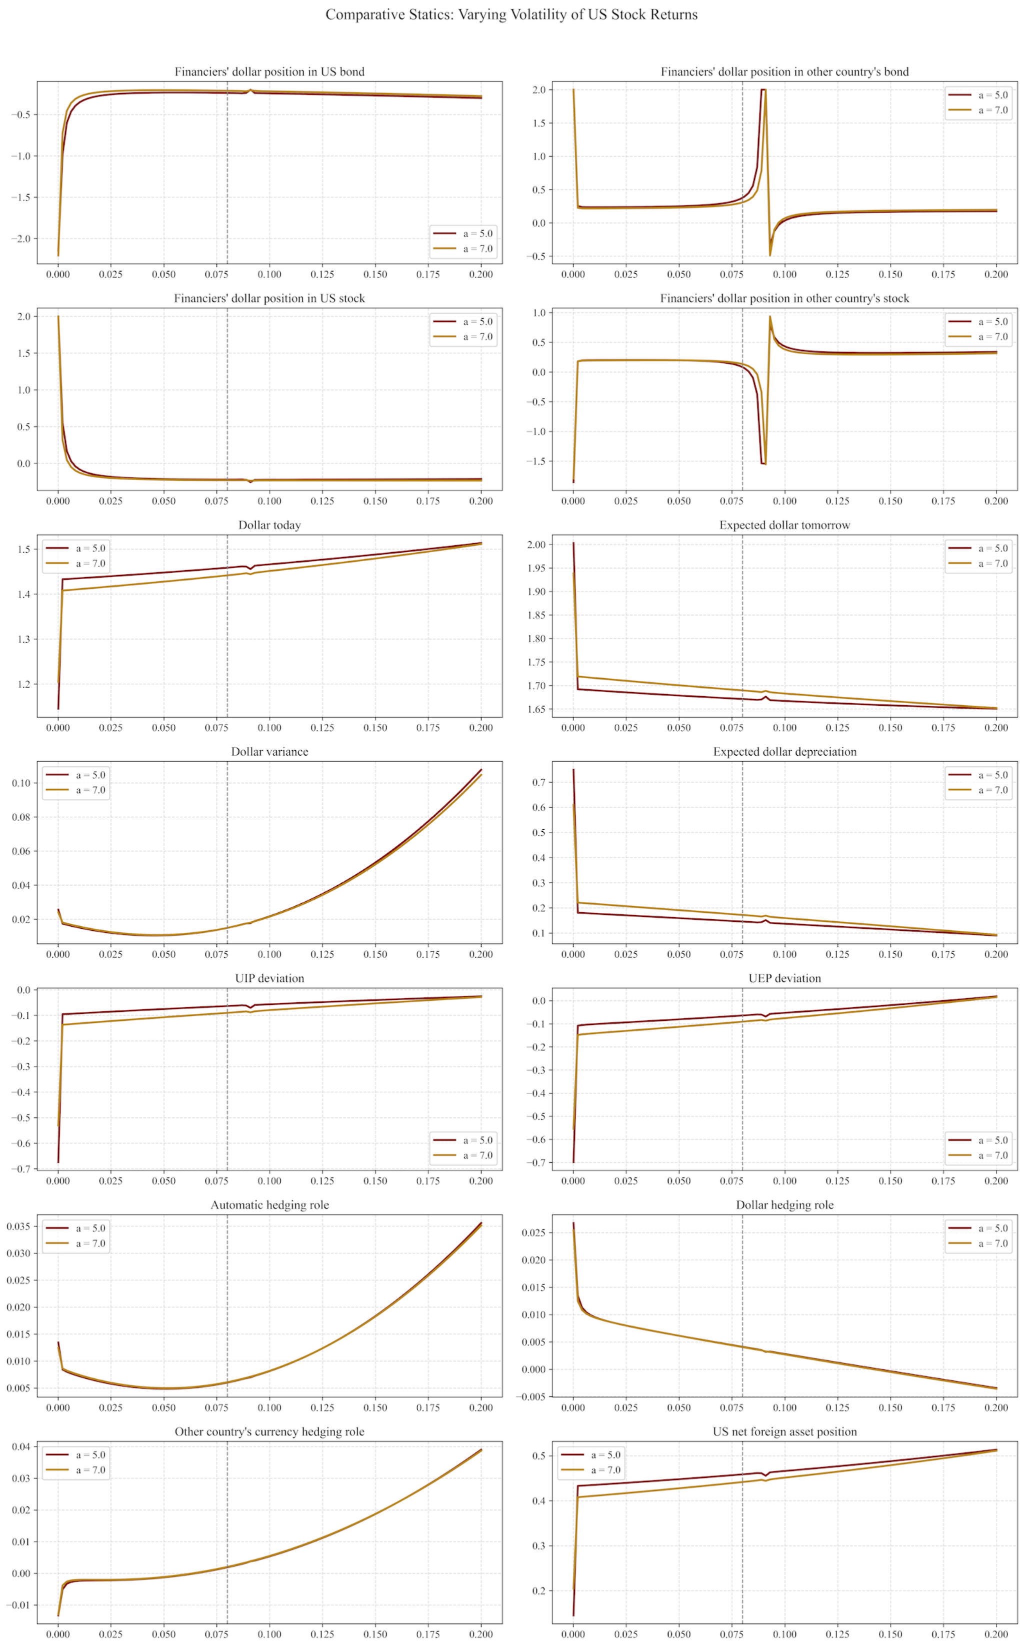

Figure A.2: Comparative Statics: Varying Volatility of U.S. Stock Returns (σ). $\iota_0 = 1$ and $\iota_1 = 2$. Remaining parameters follow the same calibration of the main paper.

A.10 Data Construction and Alternative Approaches

This appendix provides additional information on the construction and treatment of the dataset used throughout the empirical analysis.

Equity Returns

To ensure cross-country comparability, I use MSCI total return equity indices for all countries and regions. These indices reflect local equity market performance, inclusive of dividends, and are constructed in local currency terms.

Monthly log returns are computed using end-of-month index values. For 3-month and 12-month investment horizons, returns are constructed as rolling cumulative log returns over the corresponding horizons.

Interest Rates

Short-term interest rates are taken as effective interbank rates, using both 3-month and 12-month maturities to match the investment horizons. For the U.S., I use the effective Federal Funds Rate and U.S. Treasury yields. For other countries, I use central bank or IMF-sourced interbank lending rates, with a preference for consistently available series. These interest rates are used to compute interest differentials for the construction of UIP deviations.

Exchange Rates

Nominal bilateral spot exchange rates are sourced from Bloomberg, expressed as U.S. dollars per unit of the foreign currency. All FX rates are taken at the last trading day of each month to align with equity and interest rate observations.

To assess robustness, I also constructed real exchange rates using CPI indices. All main results were replicated using real returns, and differences were negligible in terms of sign, magnitude, and statistical significance.

Alternative FX Construction: Isolating Trade Effects

Because exchange rates can reflect both trade and financial flows, I implement a complementary approach to isolate the portfolio-driven component of exchange rates.

I use the BIS's broad trade-weighted exchange rate indices (TWIs) to filter out the trade-driven component of each country's FX movement. For each country, I regress the nominal spot exchange rate on its BIS trade-weighted index:

$$\text{Spot FX}_t = \alpha + \beta \cdot \text{TWI}_t + \varepsilon_t.$$

The residuals ε_t capture deviations in FX that are orthogonal to the trade-weighted component. I construct residual-based FX indices from the cumulative sum of these residuals and then use bilateral ratios of these indices to derive pairwise exchange rate series that aim to strip out trade-induced movements. These filtered FX series are then used to compute UIP and UEP deviations.

I use this method as an (imperfect) proxy for isolating portfolio effects, which is the main focus on the paper. It is also guided by the model's decomposition of exchange rate drivers — where trade preferences directly affect FX, for example showing in the expressions derived for UIP and UEP deviations in the main text — see Proposition 3. I evaluate this approach considering the cross section of countries, unconditionally. The key empirical relationships (UIP deviations, FX hedging roles, and the UIP–UEP–hedging link) remained robust under this alternative FX construction, reinforcing the validity of the main findings.

Sample and Frequency

All data are sampled at a monthly frequency. The full sample spans from May 2007 to October 2024 — it is possible to extent the sample substantially focusing on only 12 instead of the 15 countries presented in the main text, but little change was observed in terms of findings. The investment horizons used throughout the analysis are 3 and 12 months. Time-series relationships are estimated using overlapping observations (e.g., monthly updating 12-month returns), and standard errors are adjusted accordingly to account for induced serial correlation.

Countries and Regions

The analysis covers the following economies: Australia, Brazil, Canada, China, Denmark, the Euro Area, Japan, Korea, Mexico, New Zealand, Norway, Poland, Sweden, Switzerland, the United Kingdom (U.K.), and the United States (U.S.).

A.11 Unconditional Model, 3-Month Horizon

I now display results considering a 3-month investment horizon. I start by first focusing on the window that excludes the Global Financial Crisis.

A.11.1 3-Month Horizon, Excluding the Global Financial Crisis

The results below confirm those of the main text. First, UIP deviations are positively associated with a U.S. equity that is relatively more sensitive. Second, the same happens when consider relative U.S. equity volatility only. Third, foreign FX hedging roles are also positively related to U.S. equity sensitivity (and volatility), what directly implies that UIP deviations and foreign FX hedging roles are equally positively related — in this case, the correlation between UIP deviations and foreign FX hedging roles is 84%.

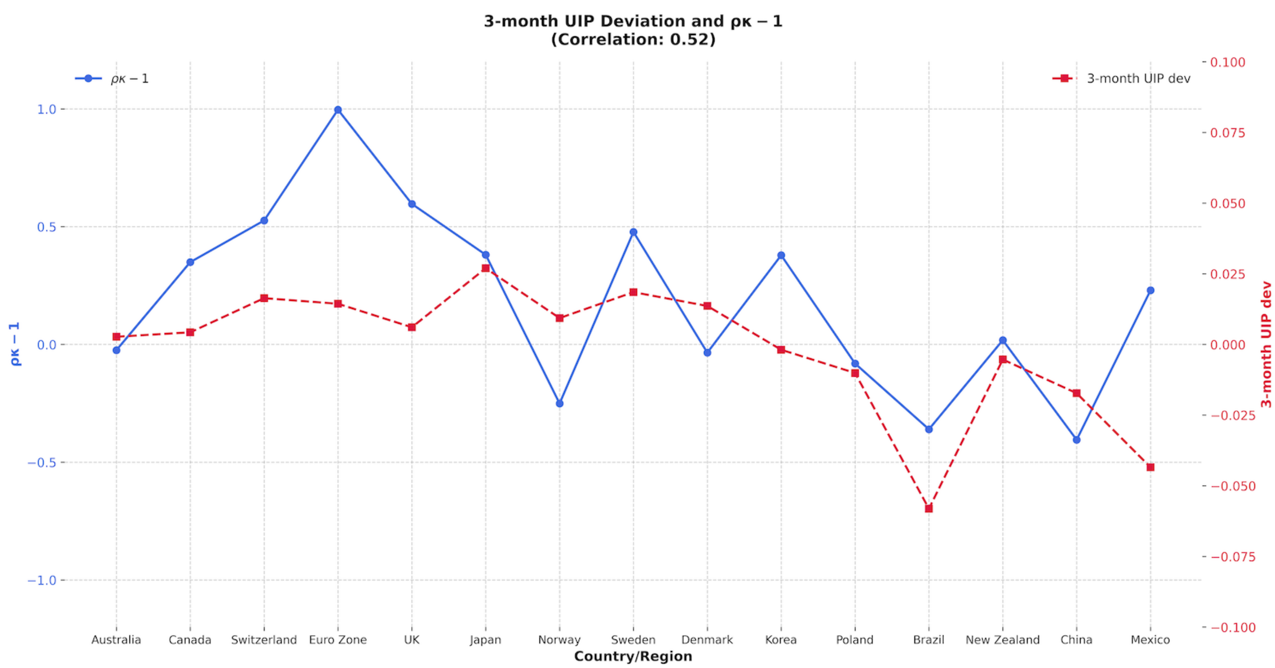

Figure A.3: 3-month UIP deviations (Ω^{UIP}, right axis) and relative equity sensitivity ($\rho\kappa - 1$, left axis) across countries, using data that excludes the Global Financial Crisis (2007–2009). The variable $\rho\kappa$ is constructed as the coefficient from an OLS regression of U.S. log-equity returns (in USD) on foreign log-equity returns (in local currency).

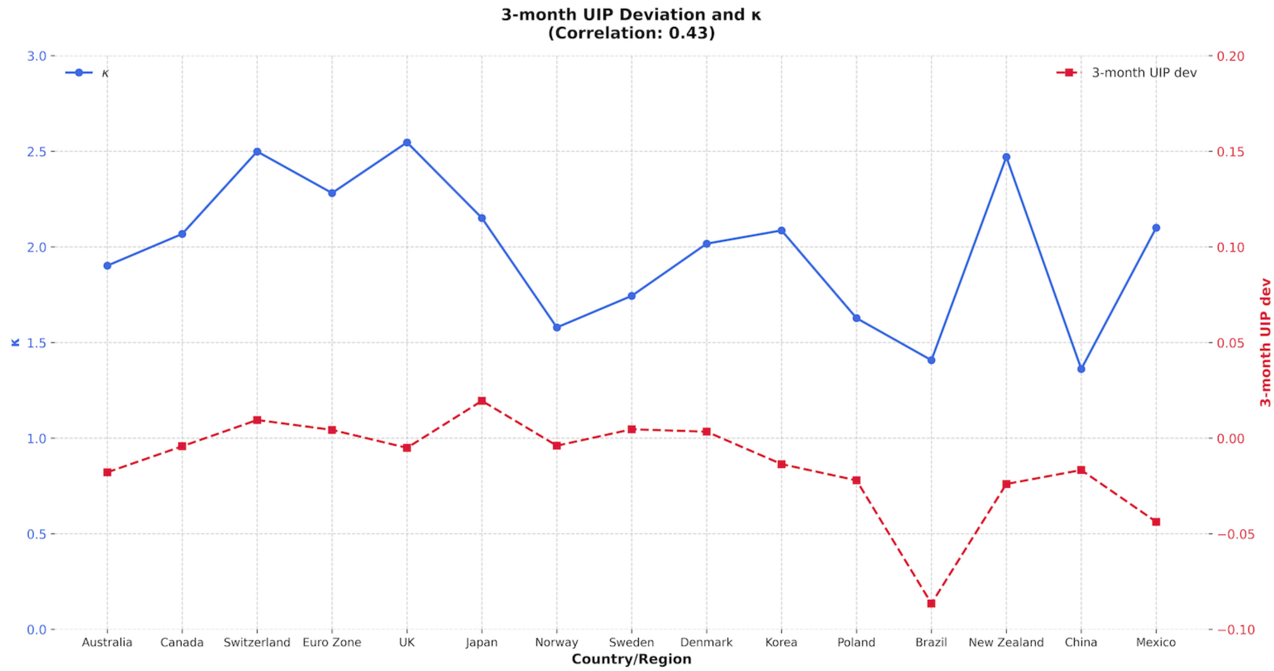

Figure A.4: 3-month UIP deviations (Ω^{UIP}, right axis) and relative equity volatility (κ, left axis) across countries, using data that excludes the Global Financial Crisis (2007–2009).

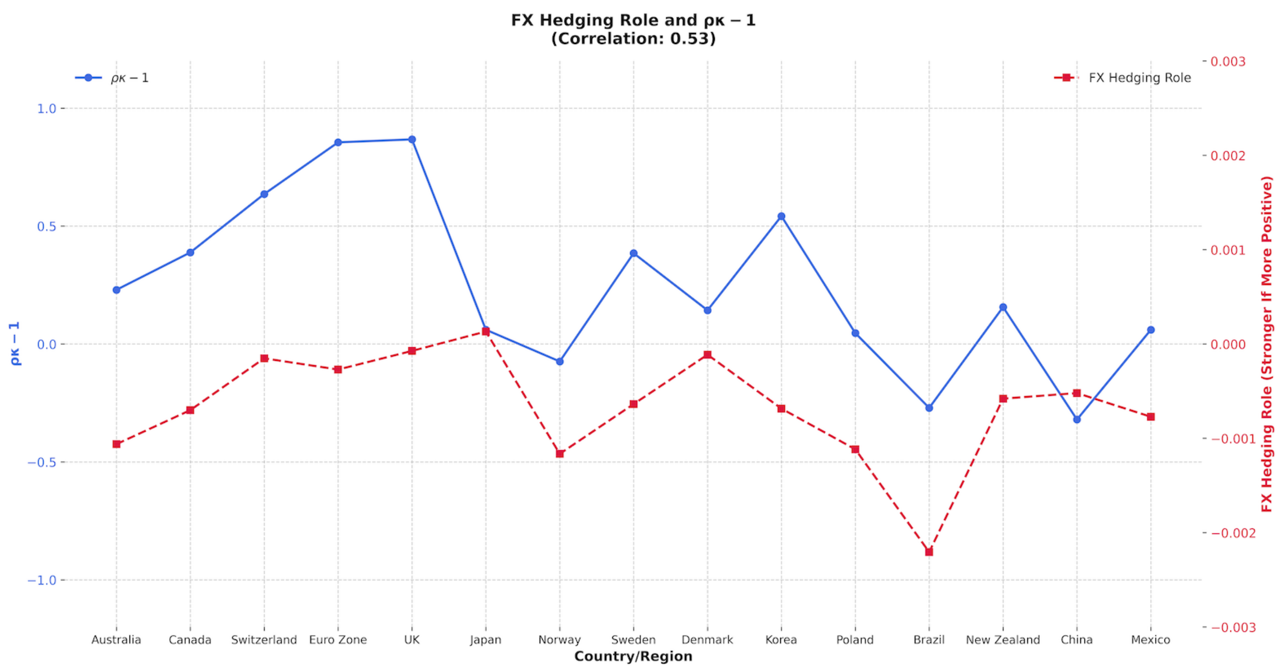

Figure A.5: 3-month FX hedging roles (right axis) and relative equity sensitivity ($\rho\kappa - 1$, left axis) across countries, using data that excludes the Global Financial Crisis (2007–2009). The variable $\rho\kappa$ is constructed as the coefficient from an OLS regression of U.S. log-equity returns (in USD) on foreign log-equity returns (in local currency).

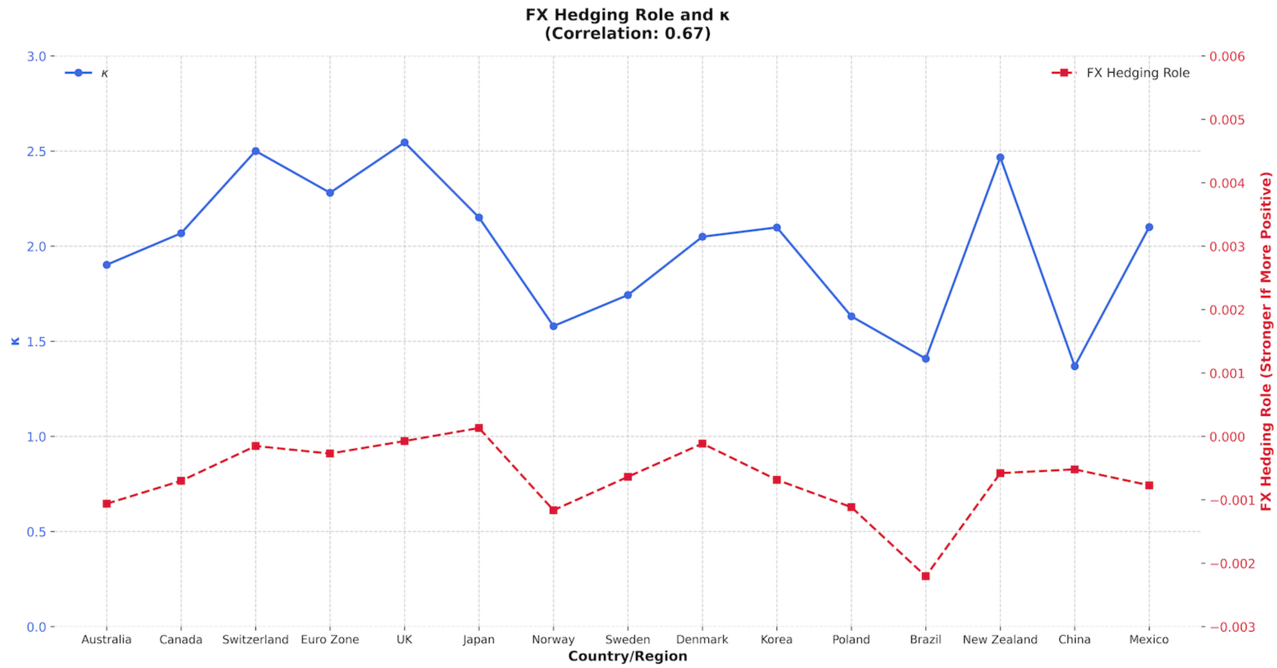

Figure A.6: 3-month FX hedging roles (right axis) and relative equity volatility (κ, left axis) across countries, using data that excludes the Global Financial Crisis (2007–2009).

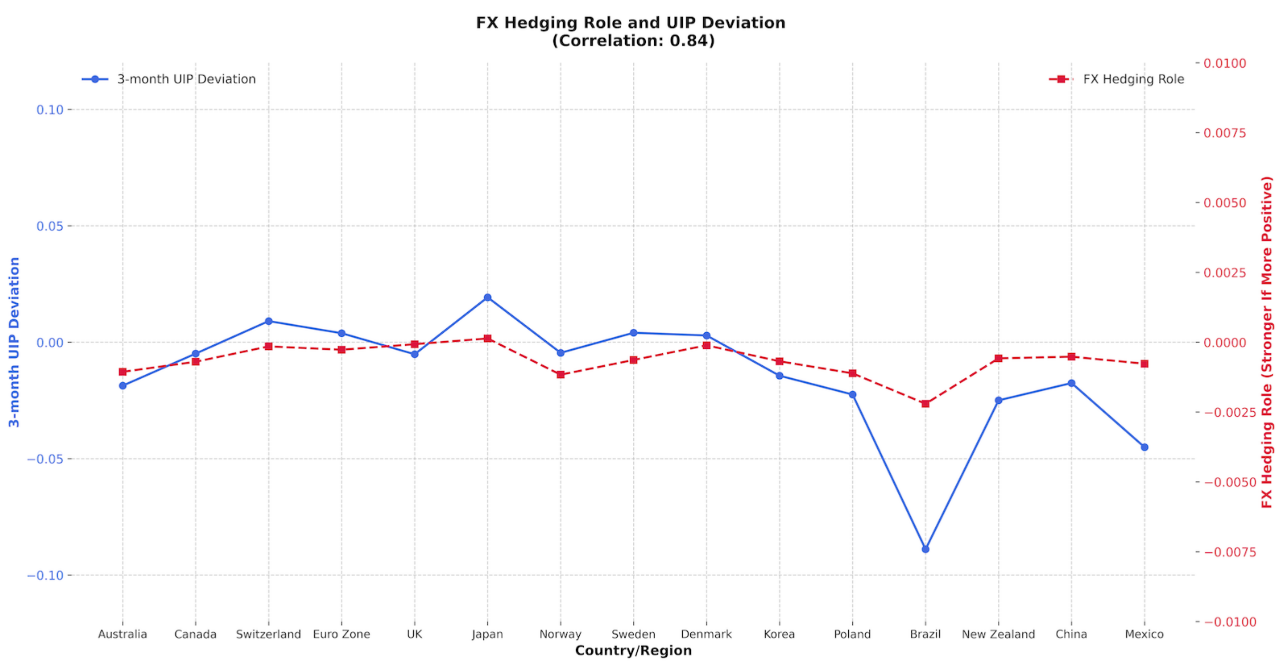

Figure A.7: 3-month FX hedging roles (right axis) and UIP deviations (left axis) across countries, using data that excludes the Global Financial Crisis (2007–2009).

A.11.2 3-Month Horizon, Full Sample

The exact same findings as above are reconfirmed when the full sample is considered.

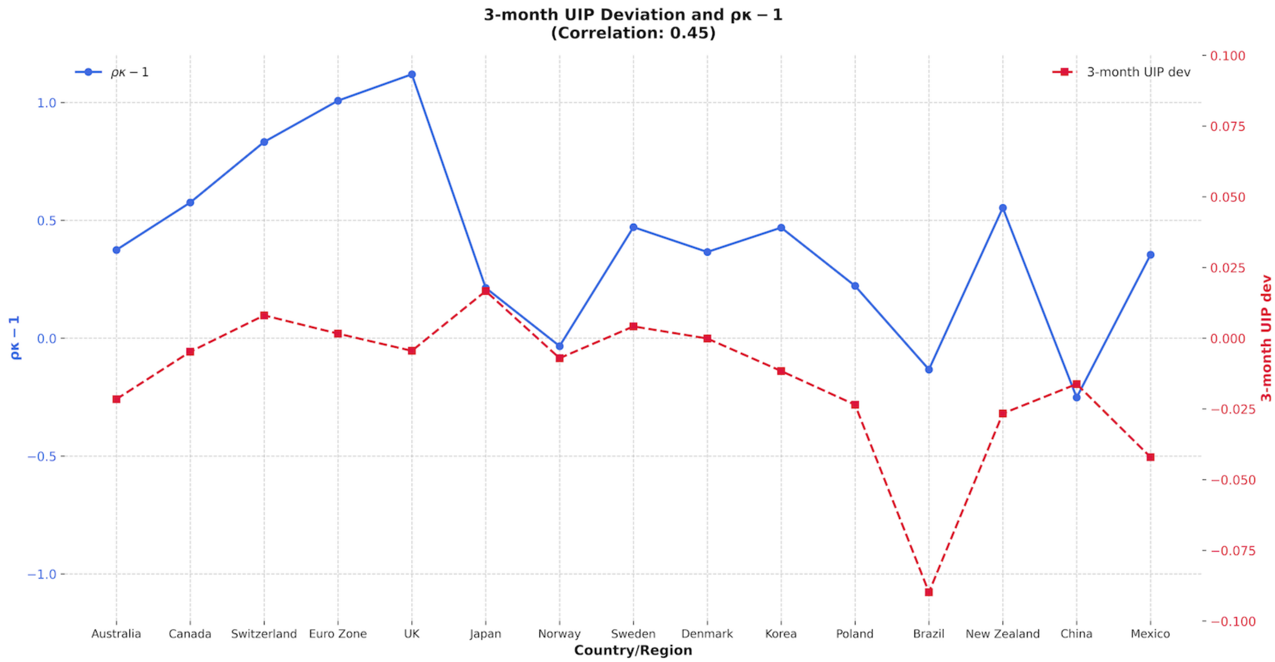

Figure A.8: 3-month UIP deviations (Ω^{UIP}, right axis) and relative equity sensitivity ($\rho\kappa - 1$, left axis) across countries, using the full sample. The variable $\rho\kappa$ is constructed as the coefficient from an OLS regression of U.S. log-equity returns (in USD) on foreign log-equity returns (in local currency).

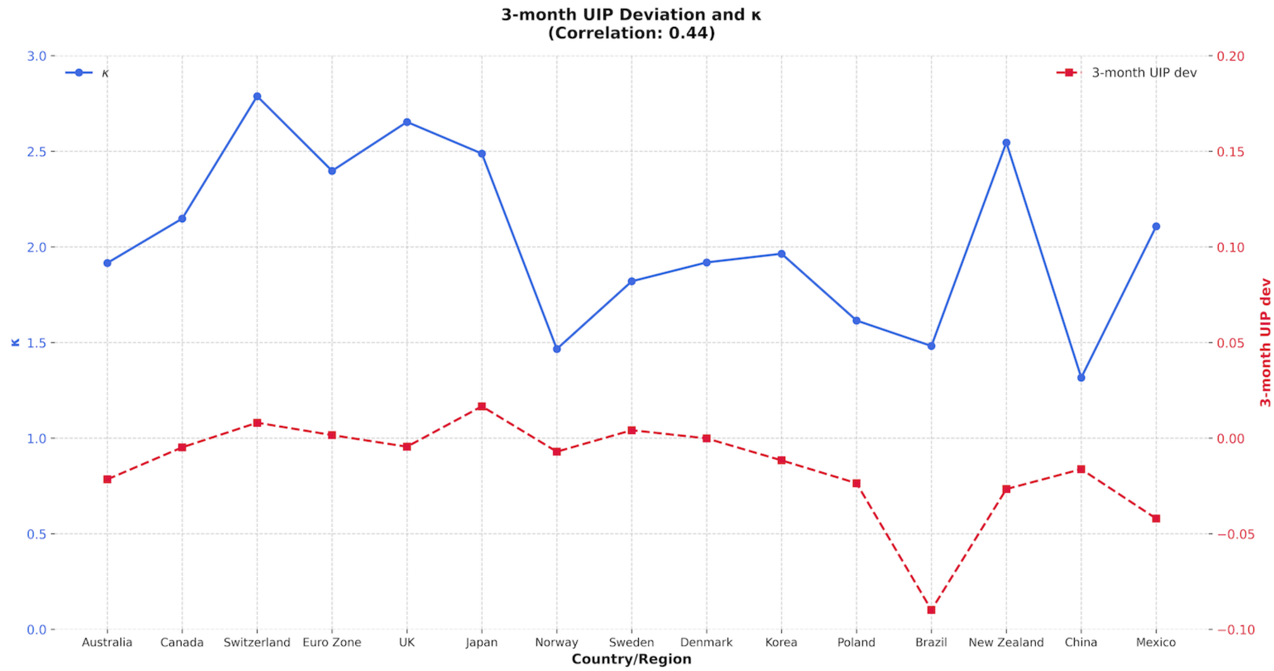

Figure A.9: 3-month UIP deviations (Ω^{UIP}, right axis) and relative equity volatility (κ, left axis) across countries, using the full sample.

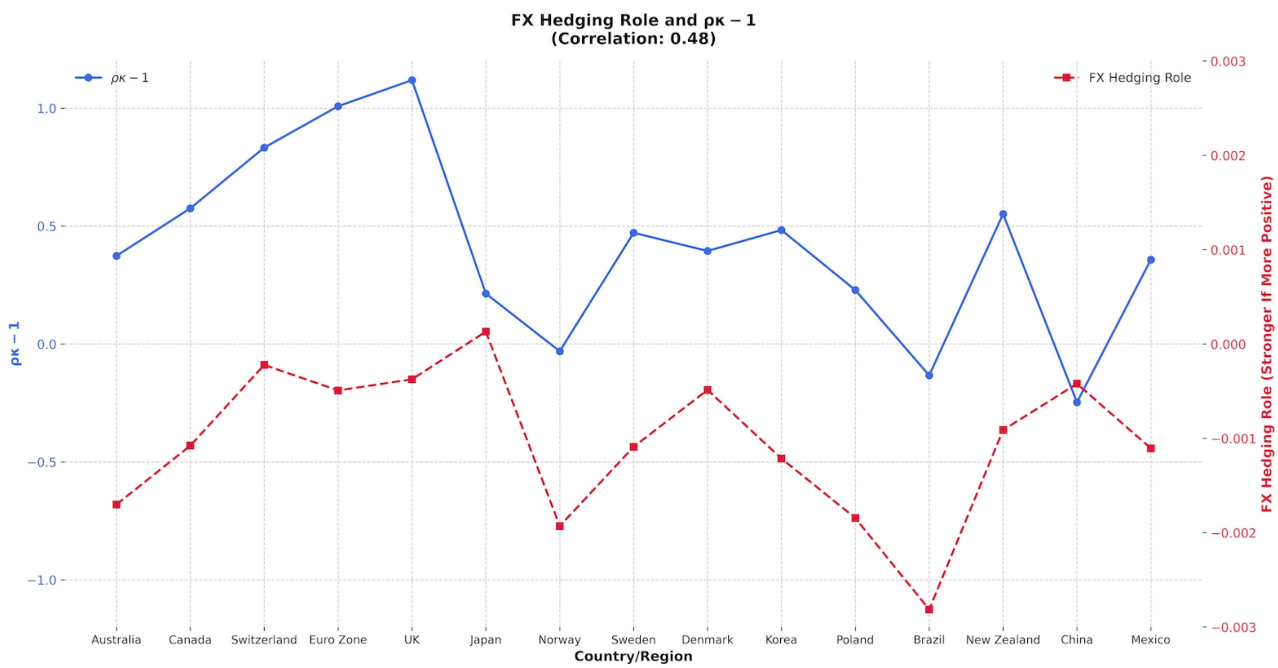

Figure A.10: 3-month FX hedging roles (right axis) and relative equity sensitivity ($\rho\kappa - 1$, left axis) across countries, using the full sample. The variable $\rho\kappa$ is constructed as the coefficient from an OLS regression of U.S. log-equity returns (in USD) on foreign log-equity returns (in local currency).

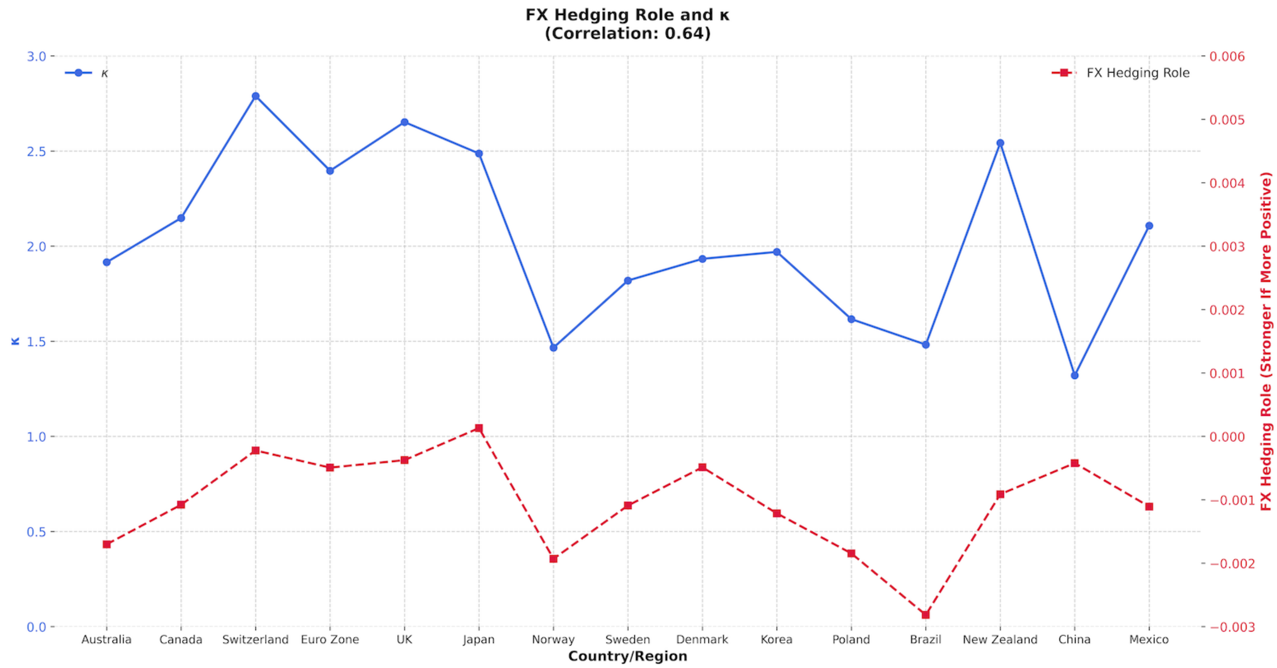

Figure A.11: 3-month FX hedging roles (right axis) and relative equity volatility (κ, left axis) across countries, using the full sample.

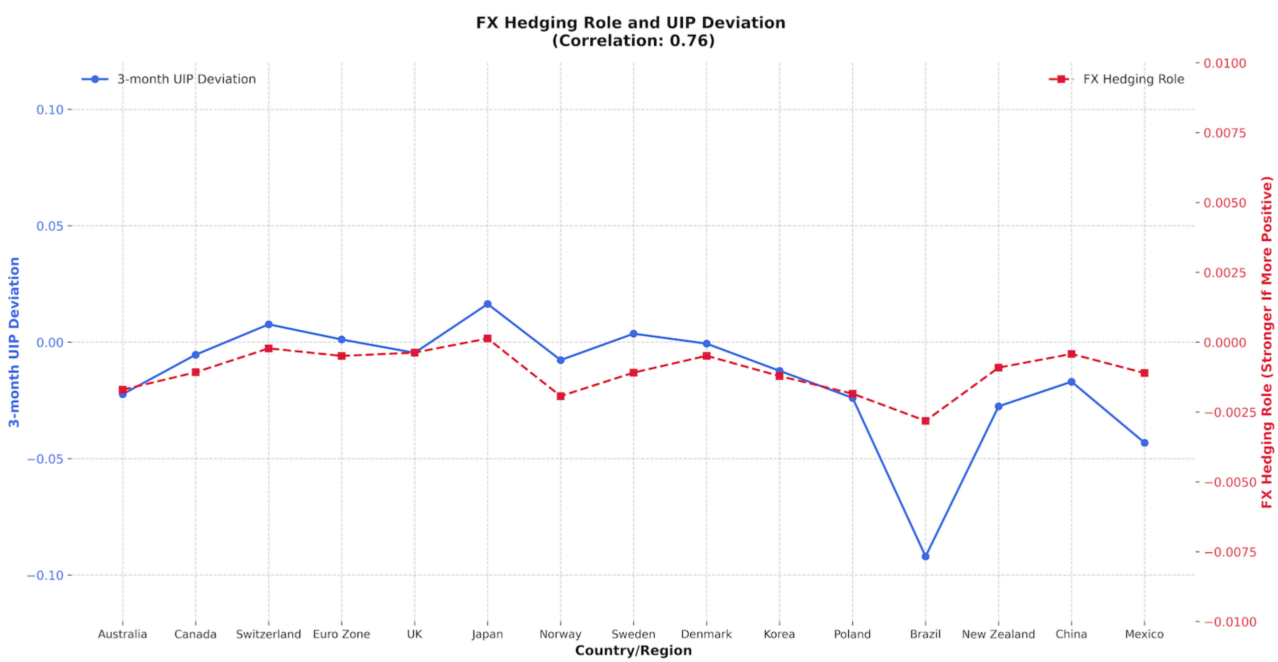

Figure A.12: 3-month FX hedging roles (right axis) and UIP deviations (left axis) across countries, using the full sample.

A.12 Alternative Sample Windows, 12-Month Horizon

In this section, I present results based on alternative sample windows. As in the main text, I consider 12-month investment horizons.

A.12.1 Excluding the Global Financial Crisis

Unconditional Model: Exploring the Cross Section

First, I focus on the sample that excludes the Global Financial Crisis (2007-2009). The following figures show results for the unconditional model, evaluating the model in the cross section of currencies. All the results shown in the main paper are maintained.

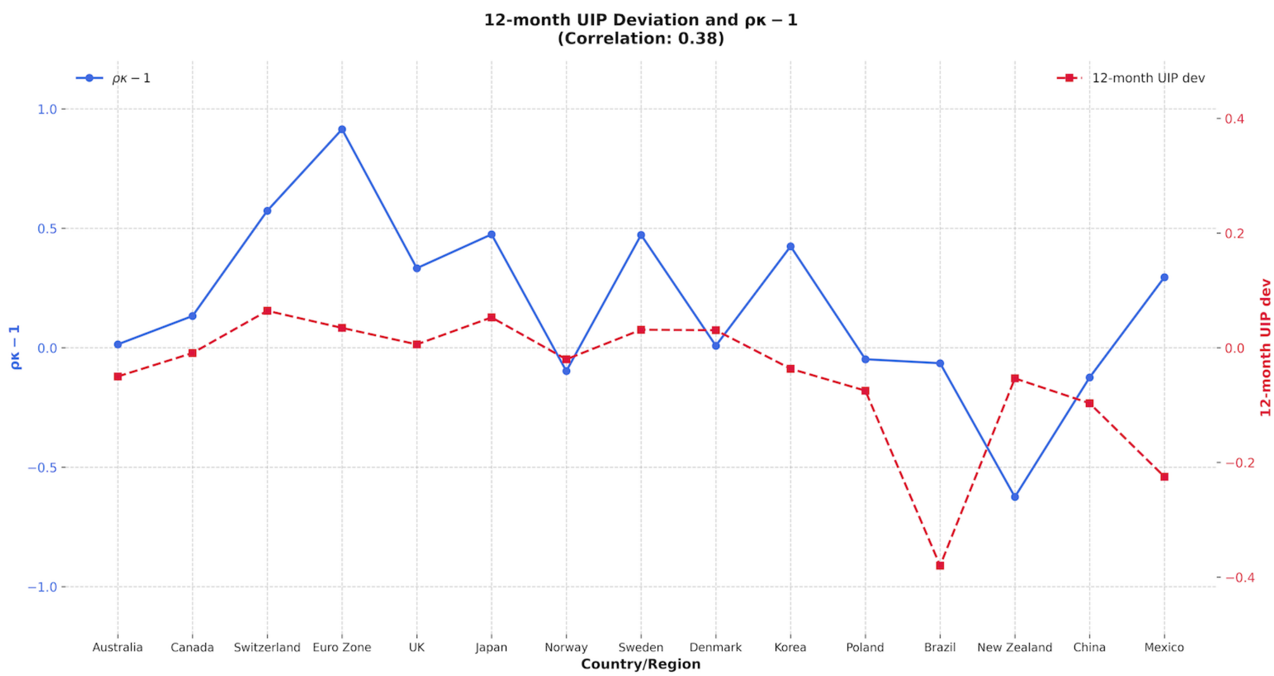

Figure A.13: 12-month UIP deviations (Ω^{UIP}, right axis) and relative equity sensitivity ($\rho\kappa - 1$, left axis) across countries, using data that excludes the Global Financial Crisis (2007–2009). The variable $\rho\kappa$ is constructed as the coefficient from an OLS regression of U.S. log-equity returns (in USD) on foreign log-equity returns (in local currency).

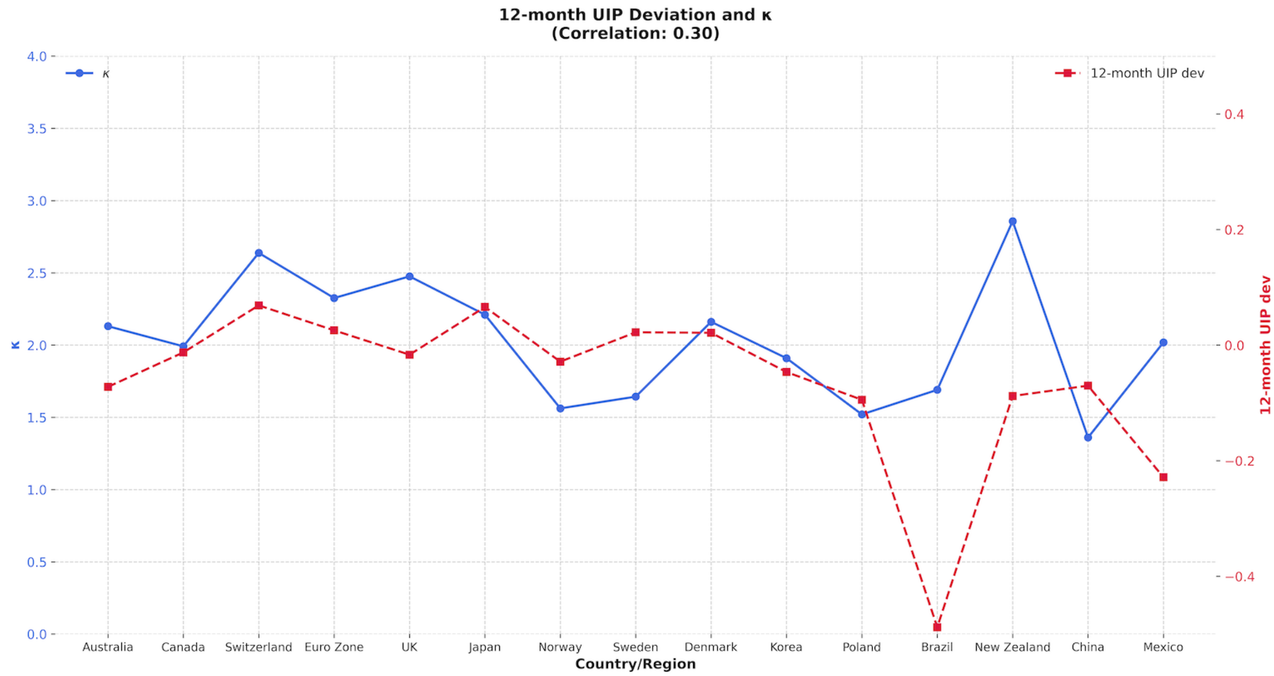

Figure A.14: 12-month UIP deviations (Ω^{UIP}, right axis) and relative equity volatility (κ, left axis) across countries, using data that excludes the Global Financial Crisis (2007–2009).

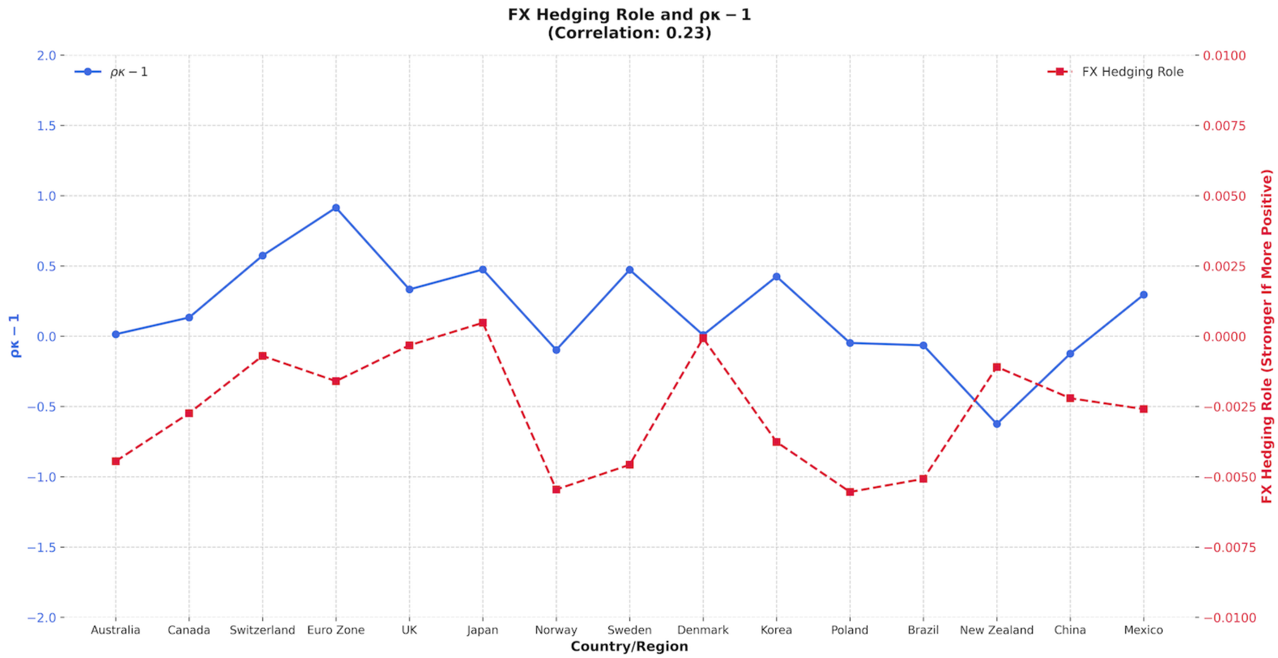

Figure A.15: 12-month FX hedging roles (right axis) and relative equity sensitivity ($\rho\kappa - 1$, left axis) across countries, using data that excludes the Global Financial Crisis (2007–2009). The variable $\rho\kappa$ is constructed as the coefficient from an OLS regression of U.S. log-equity returns (in USD) on foreign log-equity returns (in local currency).

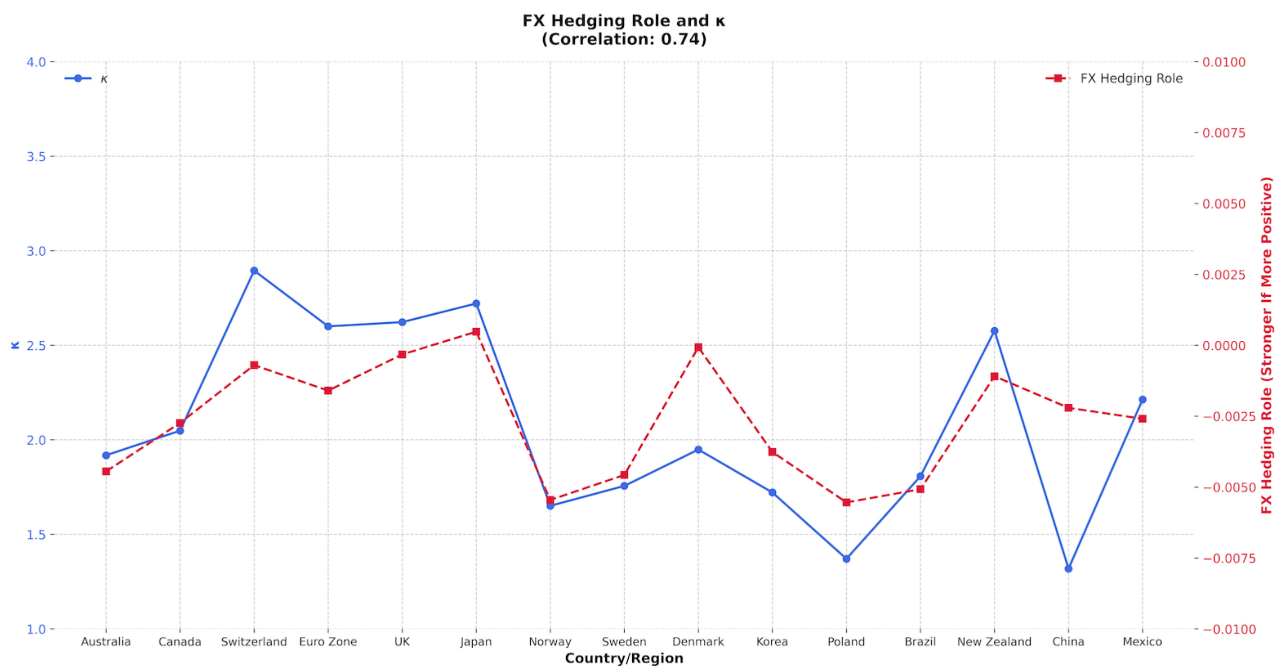

Figure A.16: 12-month FX hedging roles (right axis) and relative equity volatility (κ, left axis) across countries, using data that excludes the Global Financial Crisis (2007–2009).

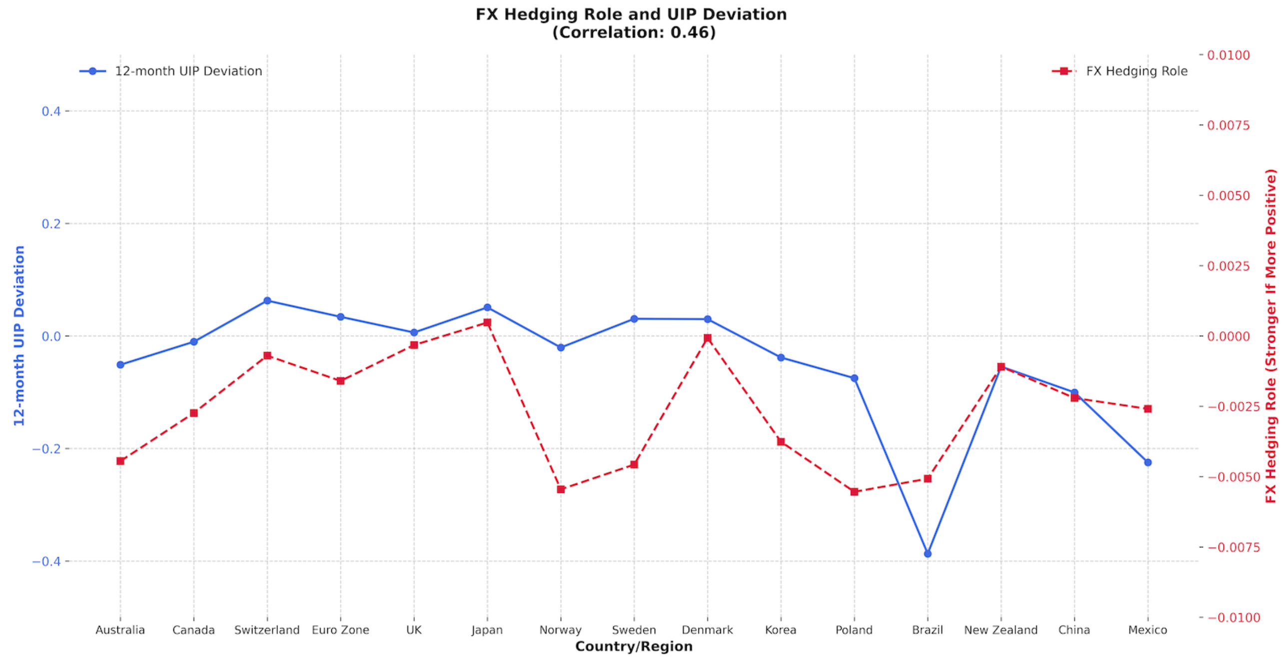

Figure A.17: 12-month FX hedging roles (right axis) and UIP deviations (left axis) across countries, using data that excludes the Global Financial Crisis (2007–2009).

Conditional Model: Exploring the Time Series

I now present results for the same sample, but focusing on the time series behaviour of exchange rates. The following figures confirm results exhibited in the main paper while evaluating the model conditionally.

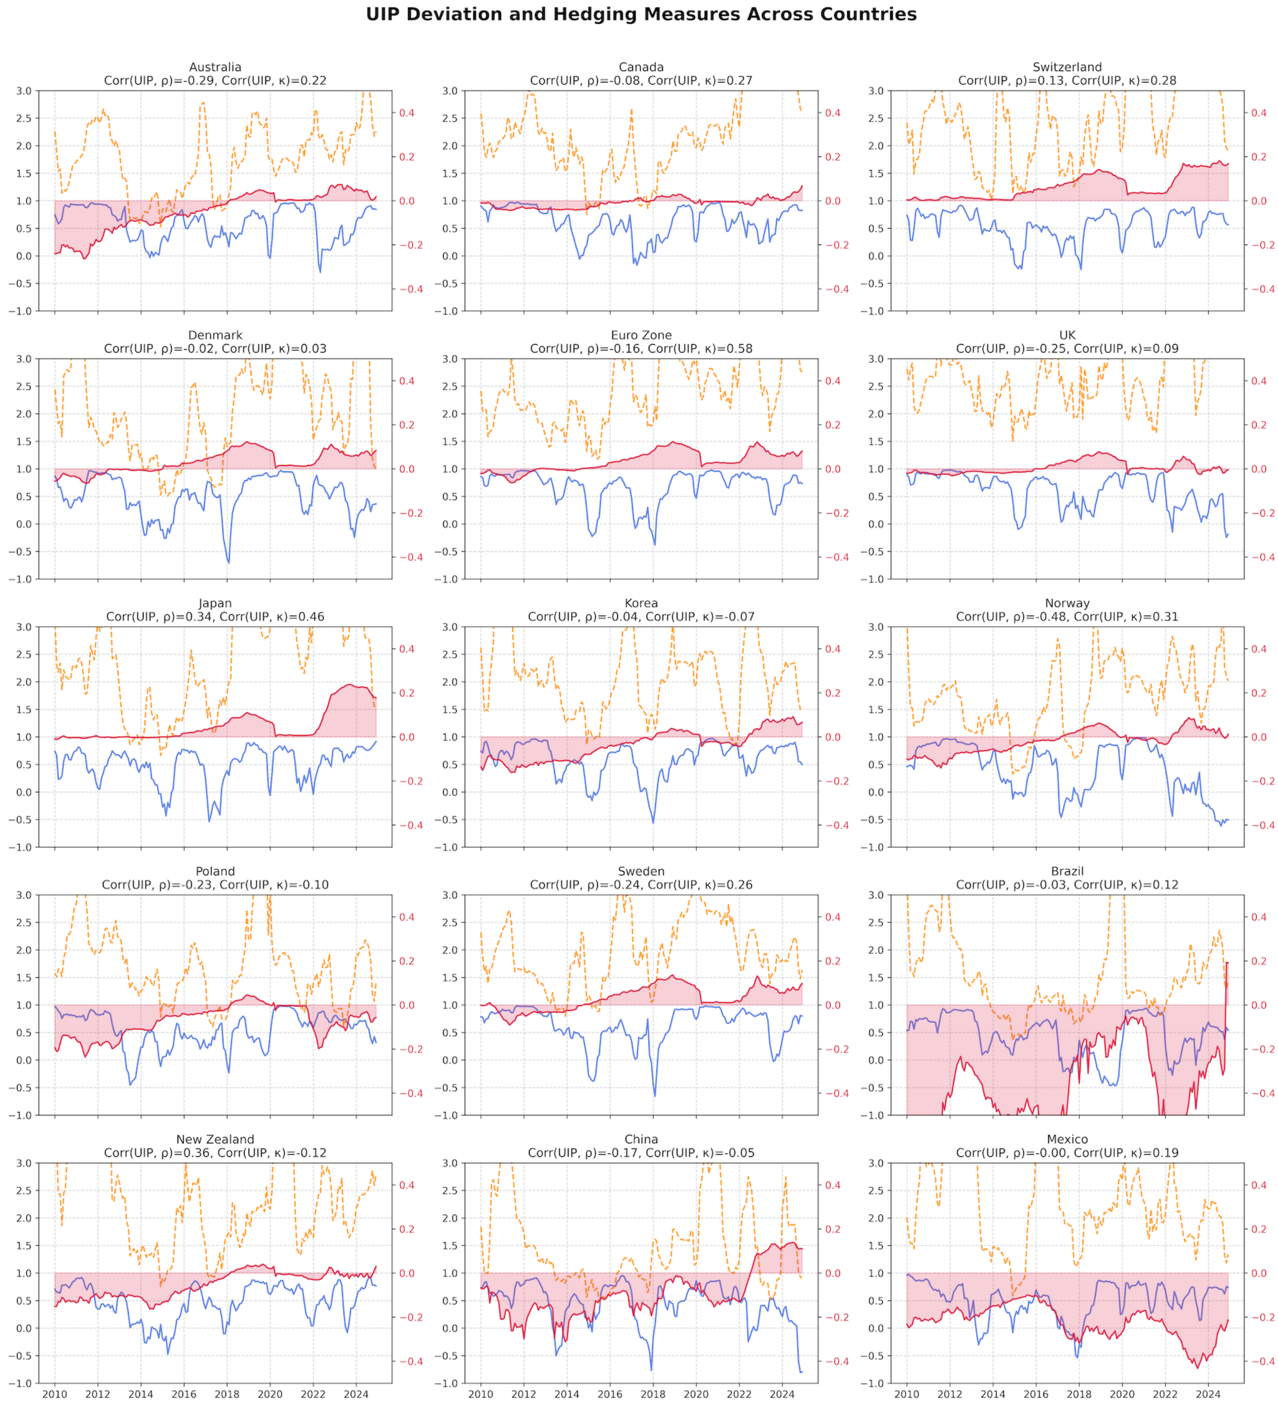

Figure A.18: UIP Deviation and Hedging Measures Across Countries. Each panel plots the UIP deviation (red), ρ (blue) and κ (orange, dashed) for 12-month investment horizons. Sample excludes the Global Financial Crisis (2007-2009).

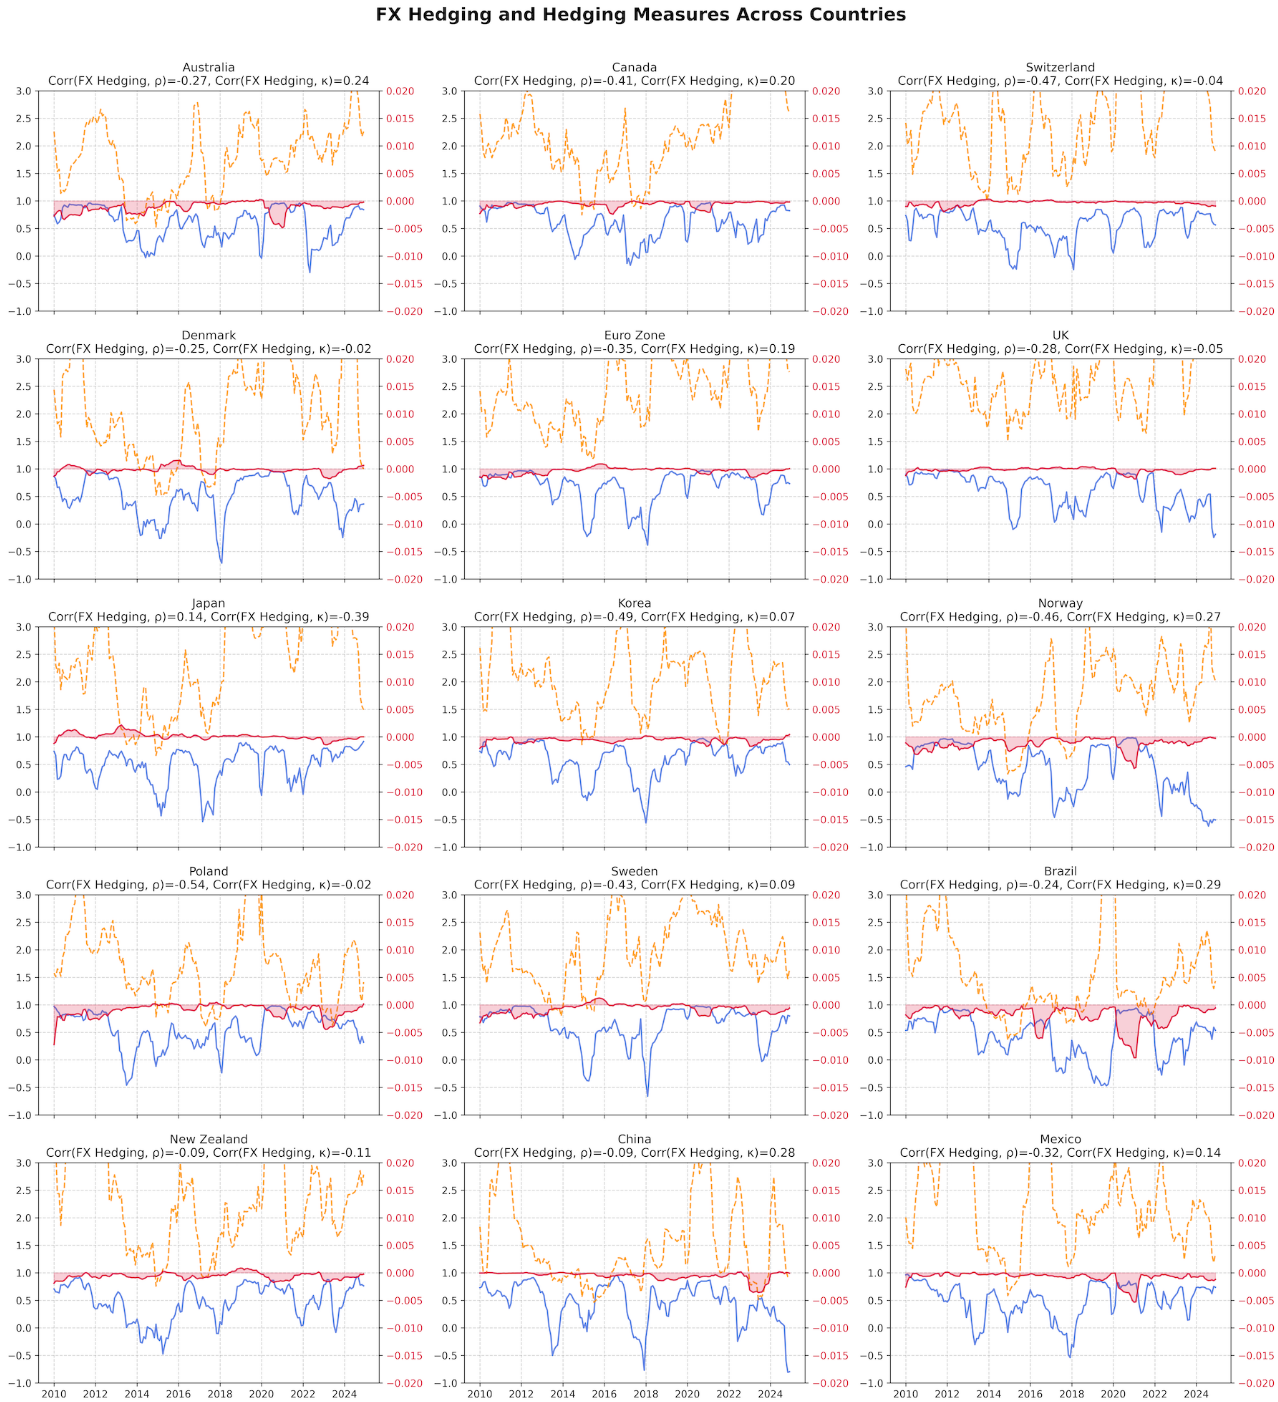

Figure A.19: FX Hedging and Hedging Measures Across Countries. Each panel plots the FX hedging role (red), ρ (blue) and κ (orange, dashed) for 12-month investment horizons. Sample excludes the Global Financial Crisis (2007-2009).

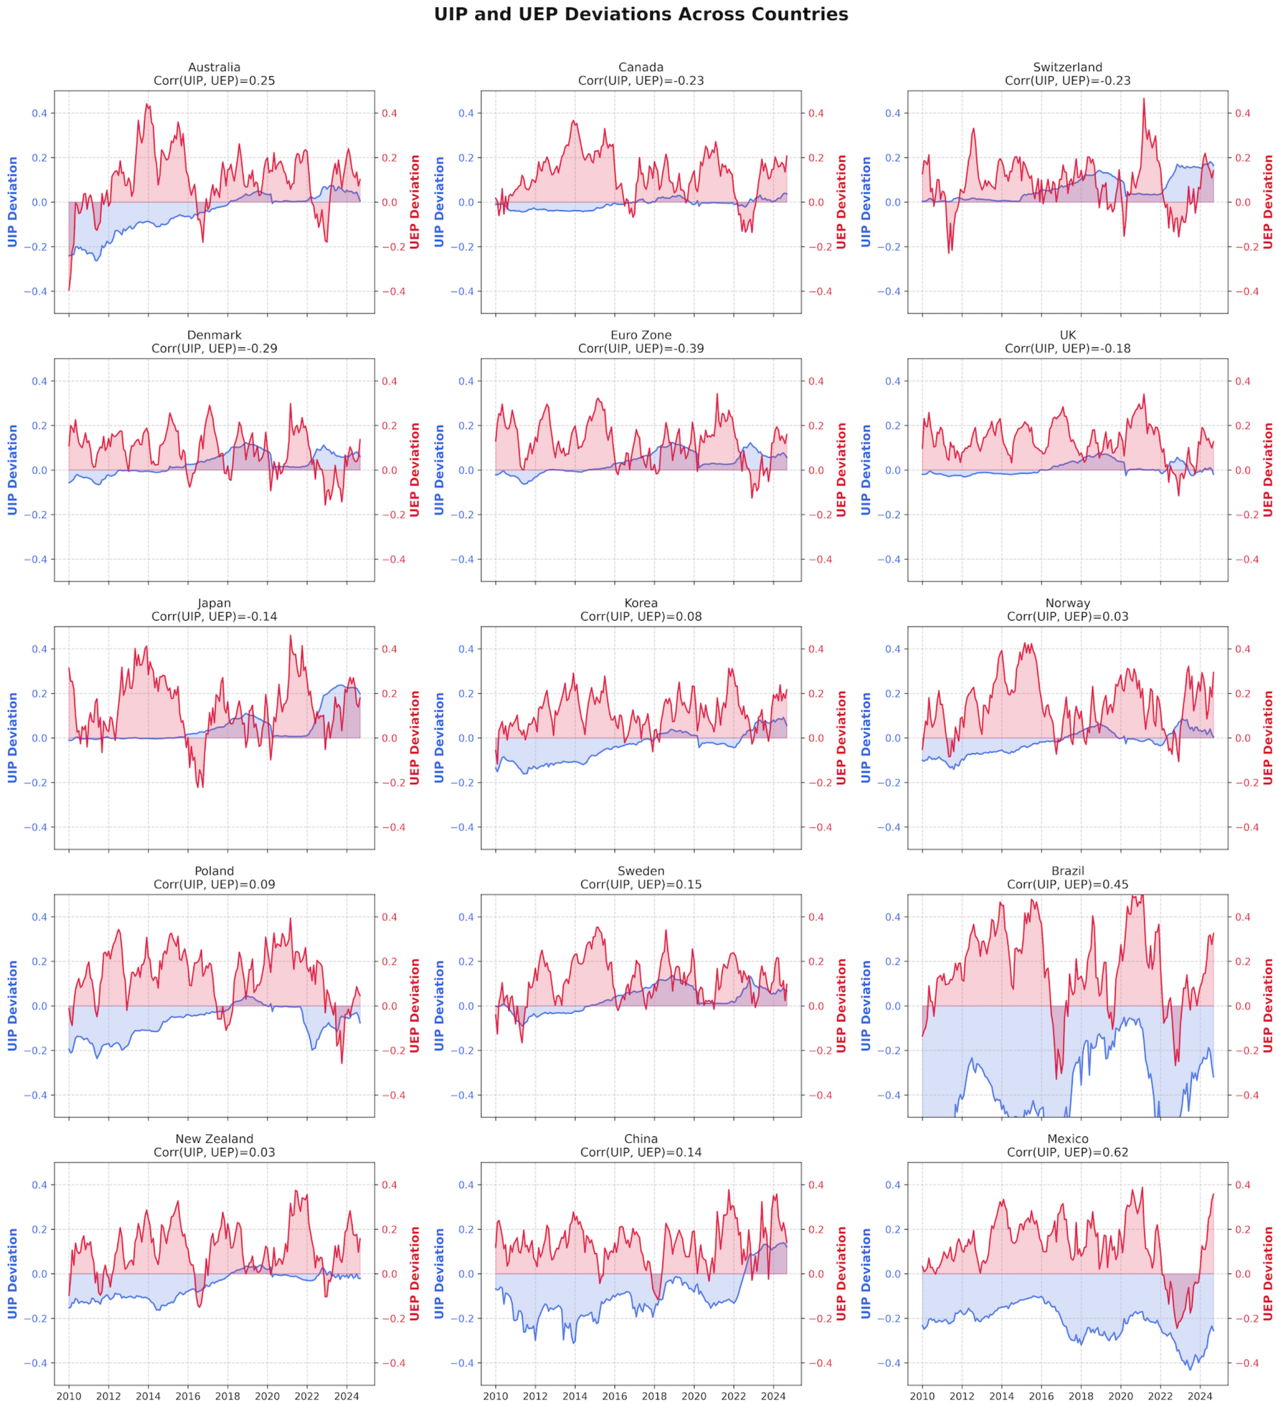

Figure A.20: UIP and UEP Deviations Across Countries. Each panel plots the UEP deviations (red) alongside the UIP deviations (blue) for 12-month investment horizons. Sample excludes the Global Financial Crisis (2007-2009).

A.12.2 2012 Onwards

I now present additional results for a more recent sample, for which correlations between equity markets and exchange rate have been more substantial — in line with what is documented in Lilley and Rinaldi (Working Paper). All my previous findings are also maintained when focusing on this window.

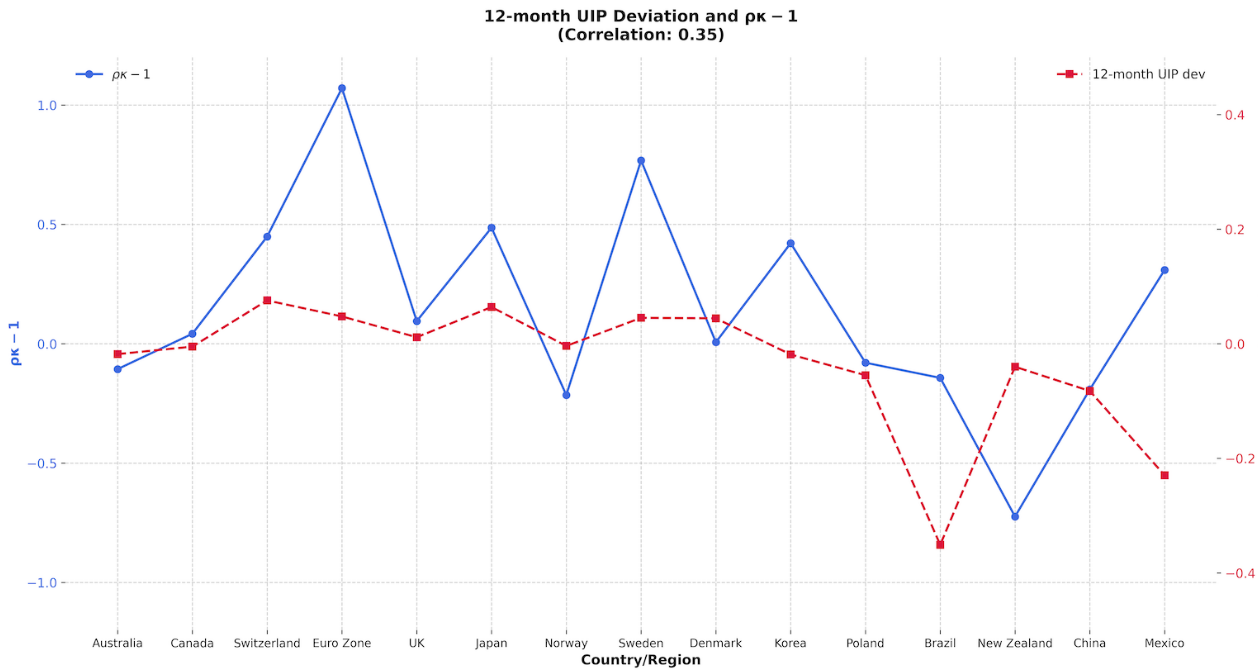

Figure A.21: 12-month UIP deviations (right axis) and relative equity sensitivity ($\rho\kappa - 1$, left axis) across countries. The variable $\rho\kappa$ is constructed as the coefficient from an OLS regression of U.S. log-equity returns (in USD) on foreign log-equity returns (in local currency). Sample starts in January 2012.

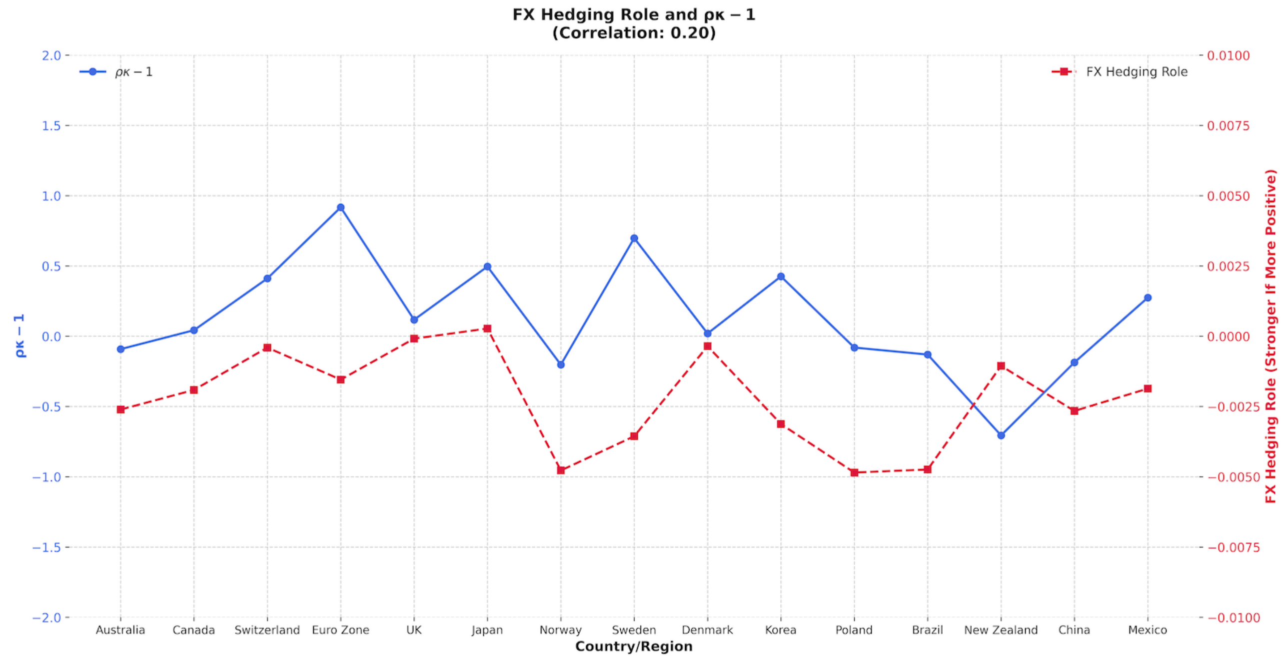

Figure A.22: FX hedging roles (right axis) and relative equity sensitivity ($\rho\kappa - 1$, left axis) across countries. The variable $\rho\kappa$ is constructed as the coefficient from an OLS regression of U.S. log-equity returns (in USD) on foreign log-equity returns (in local currency). Investment horizon is 12 months. Sample starts in January 2012.

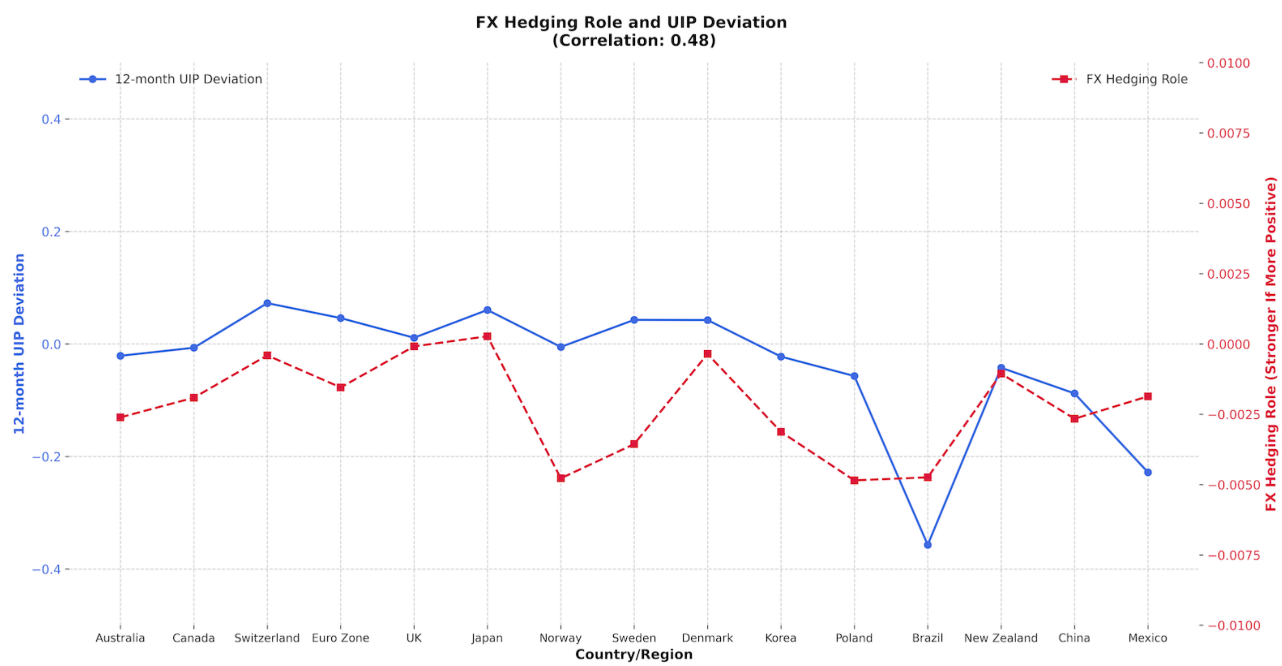

Figure A.23: FX hedging roles (right axis) and UIP deviations (left axis) across countries. Investment horizon is 12 months. Sample starts in January 2012.

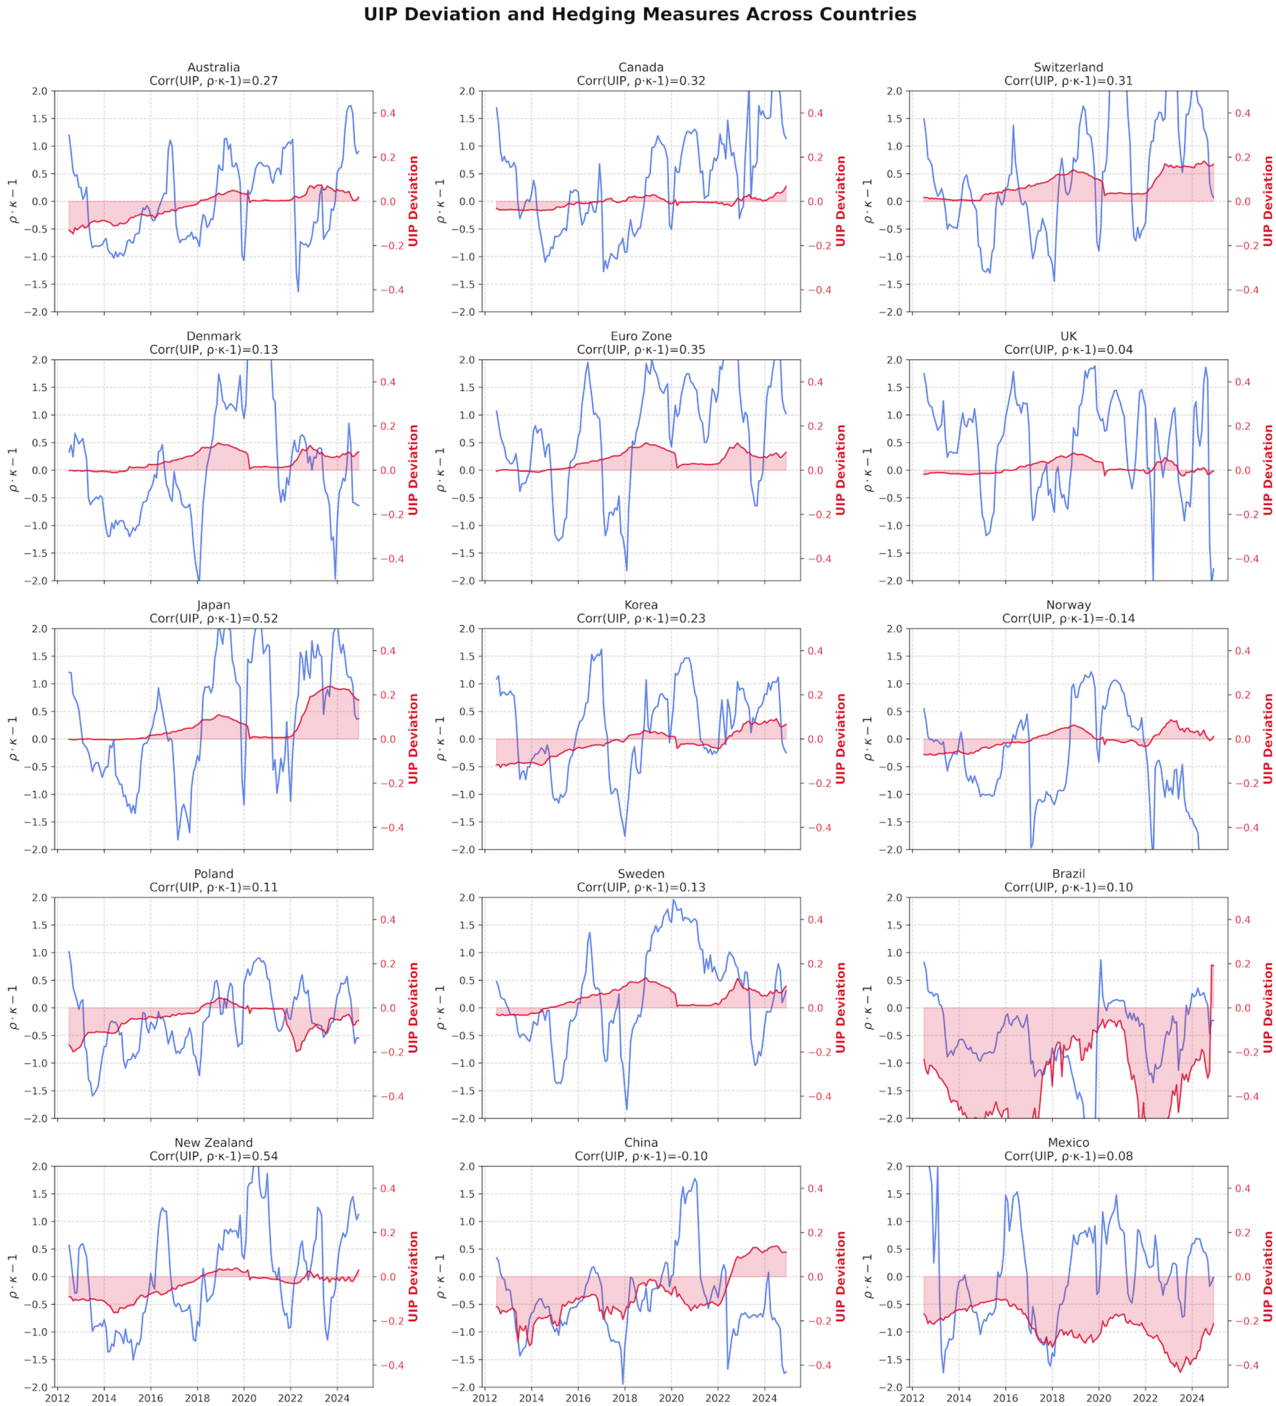

Figure A.24: UIP Deviation and Hedging Measures Across Countries. Each panel plots the UIP deviation (red) and $\rho\kappa - 1$ (blue) for 12-month investment horizons. Sample starts in January 2012.

A.13 Complementary Charts

The main paper shows the relationship between UIP deviations and relative equity sensitivity. Figure A.25 below complements that result, demonstrating that the same result also applies in terms of relative equity volatility.

Figure A.26 confirms that UIP is also generally directly related to relative equity sensitivity in the time series. This once more confirms predictions of the main text.

Finally, Figure A.27 shows that the same intuition of the main text applied to foreign FX hedging roles also applies to the dollar hedging role — here, we should observe a *negative* relationship between κ and the dollar hedging role, in the sense that a *safer* U.S. equity market contributes to a *stronger* dollar hedging role.

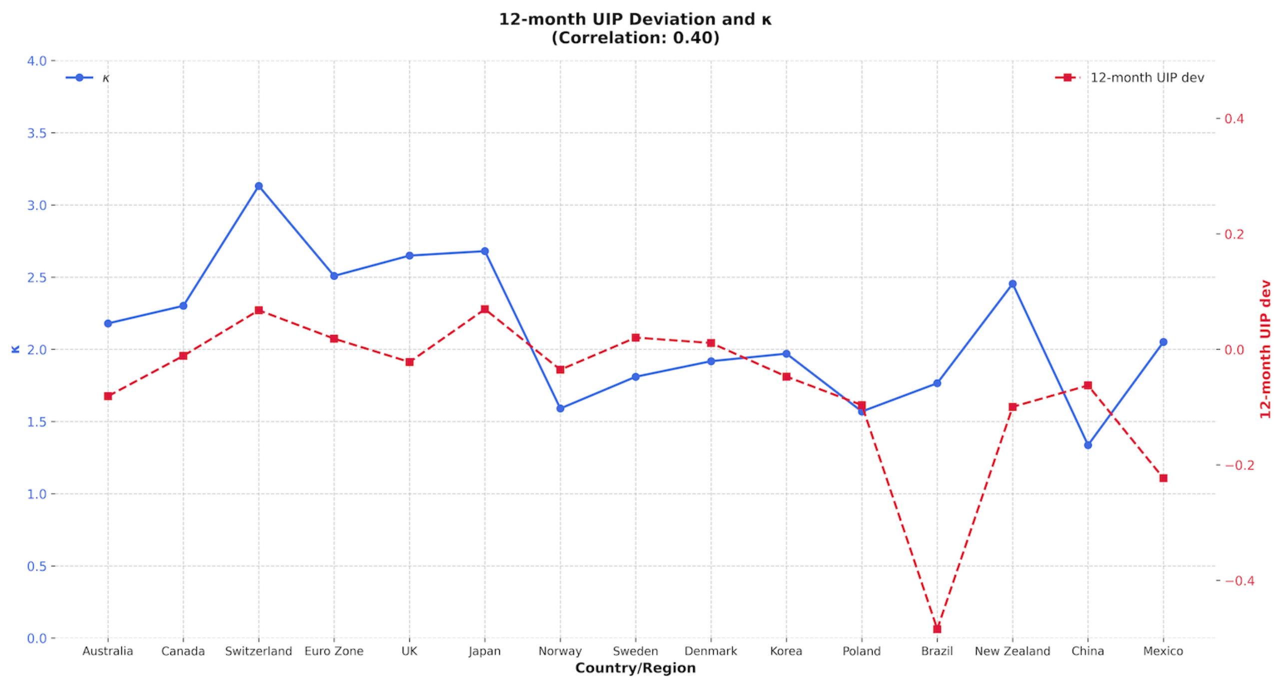

Figure A.25: 12-month UIP deviations (Ω^{UIP}, right axis) and relative equity volatility (κ, left axis) across countries. Full sample.

UIP Deviation and Hedging Measures Across Countries

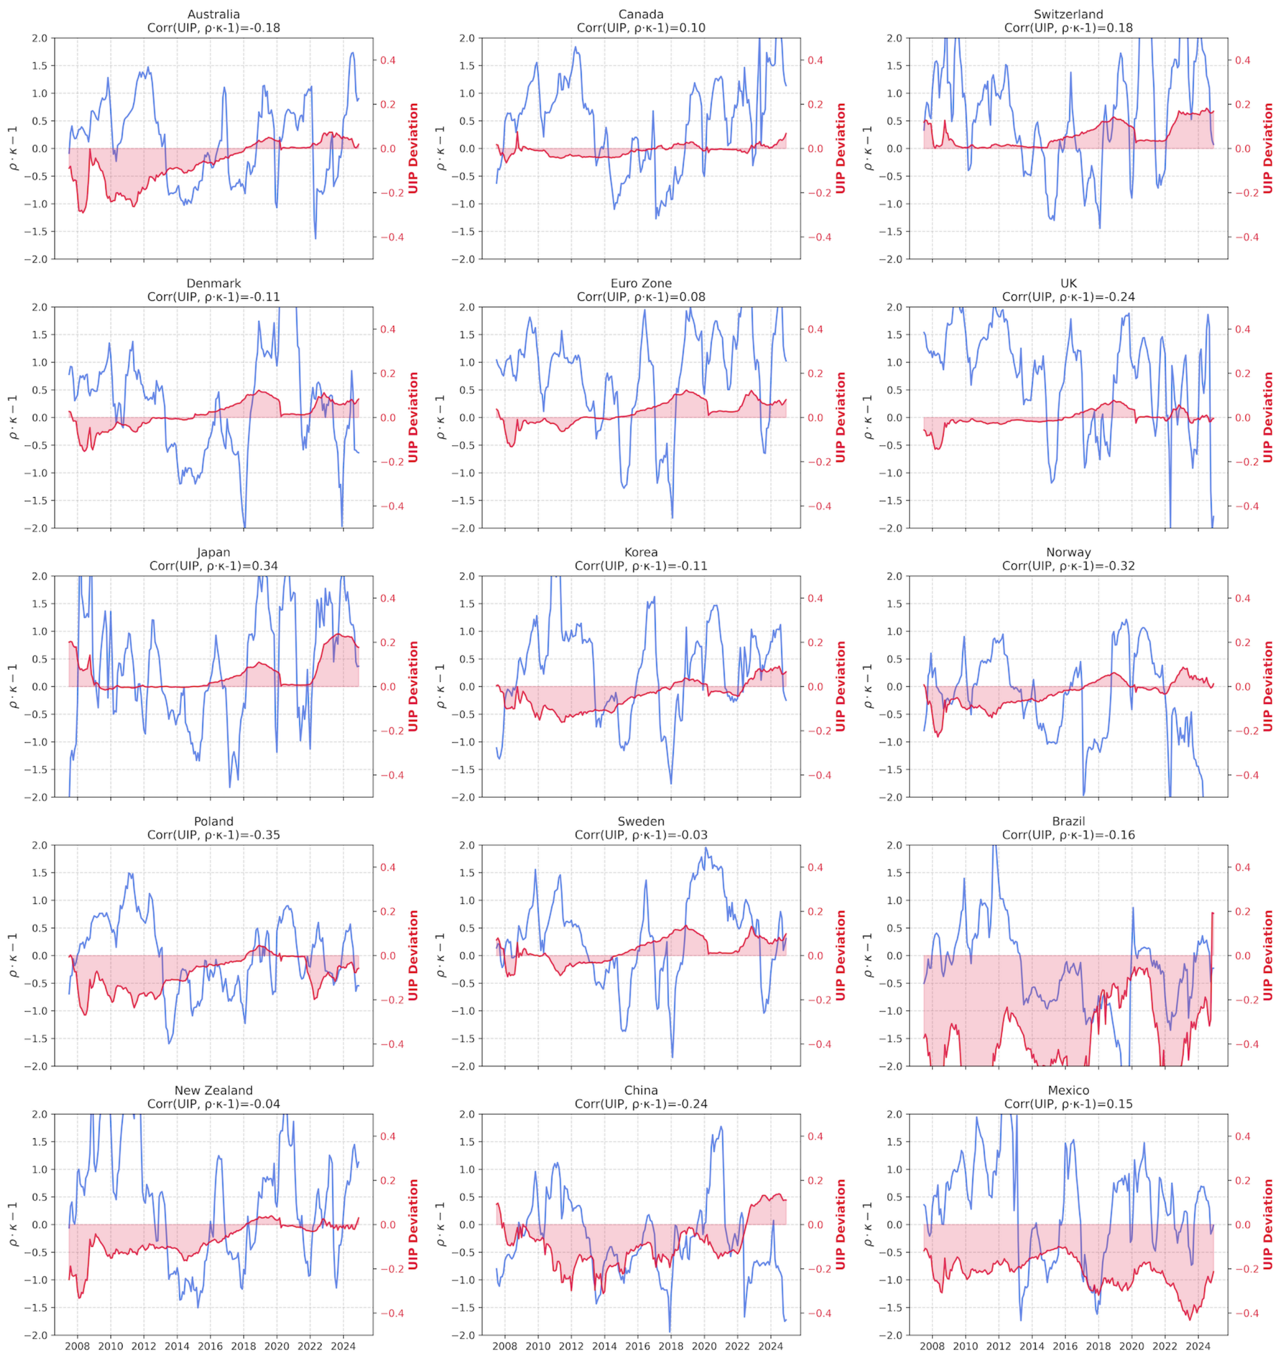

Figure A.26: UIP Deviation and Hedging Measures Across Countries. Each panel plots the UIP deviation (red) and $\rho\kappa - 1$ (blue) for 12-month investment horizons. Full sample.

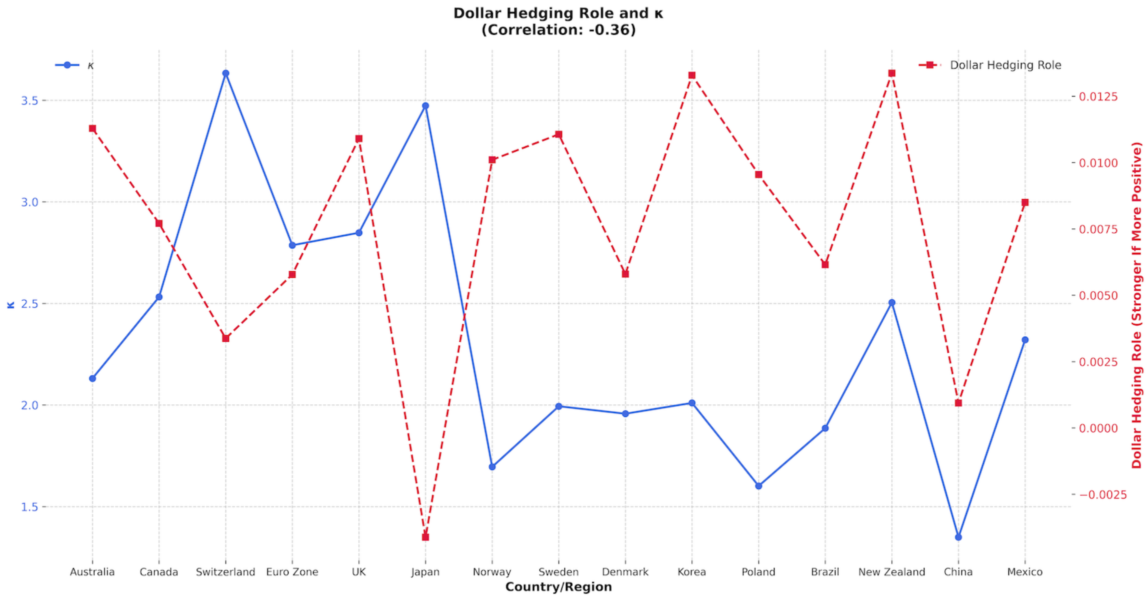

Figure A.27: Dollar Hedging Role and κ, obtained for a 12-month investment horizon.

A.14 Complementary Tables

Robustness: Excluding the Global Financial Crisis

Panel Regression with Fixed Effects

Table A.1 presents results from a panel regression identical to the baseline specification, but excluding the Global Financial Crisis period (2007–2009). As before, this subsample is intended to assess whether the baseline findings are disproportionately driven by the extreme dislocations observed during that time. The regression specification remains unchanged, including country fixed effects, and robust standard errors are again computed using HAC estimators with 11 lags.

The results are broadly consistent with those in the main text. The coefficient on the UIP deviation remains *positive and statistically significant* at the 1% level, as expected under the theoretical framework. This reinforces the core prediction that deviations from uncovered interest parity are systematically associated with deviations from uncovered equity parity. The estimated coefficient is slightly lower than in the full-sample regression (0.232 vs. 0.457), but the direction and statistical significance are preserved, suggesting robustness to the exclusion of crisis dynamics.

The coefficient on the hedging term flips sign and becomes *mildly negative*, but it is *not statistically significant*, and the *confidence interval remains wide*. Importantly, the theoretical benchmark value of 1 continues to fall well within the 95% confidence band, which spans from approximately –18 to 14. This suggests that while the precision of the estimate deteriorates in the reduced sample, the result is *not inconsistent* with the model's prediction. The sign flip may simply reflect increased noise due to the exclusion of a high-volatility period that likely contained valuable identifying variation.

All country fixed effects remain *positive and statistically significant*, with tight confidence intervals across the board. These estimates reflect country-specific average deviations in the dependent variable, and their statistical significance is mechanically implied by the inclusion of a full set of country fixed effects in the regression. Since the fixed effects soak up persistent country-level heterogeneity, their significance indicates systematic cross-sectional variation in the level of the UEP deviation not explained by the two main regressors — this is, the third term in 1.19, associated with relative equity volatility.

Overall, these findings provide further support for the robustness of the main results. Excluding the Global Financial Crisis does not alter the qualitative interpretation.

Bilateral Regressions

The bilateral regression results in Table A.2 further corroborate the findings from the panel analysis. While estimates vary across countries, several key patterns are preserved. First, the coefficient on the UIP deviation is positive in many cases, consistent with the theoretical prediction. Second, the coefficient on the FX hedging role often exhibits large standard errors and wide confidence intervals, mirroring the imprecise estimation observed in the panel regressions. Notably, in only one case (Mexico) the model-implied value of one can be statistically rejected based on the 95% confidence bands. This reinforces the conclusion that, despite limited precision in some cases, the results are broadly consistent with the model's predictions. Finally, all regression constant coefficients are *positive and significant*, in line with model predictions.

Table A.1: Panel Regression: UEP Deviation on UIP Deviation and FX Hedging Role (Excluding GFC, with Fixed Effects)

Variable	**Coefficient**	**Std. Error**	**z-stat**	**p-value**	**[95% CI]**
UIP Deviation	0.232	0.088	2.643	0.008	[0.060, 0.404]
FX Hedging Role	–1.836	8.277	–0.222	0.824	[–18.059, 14.387]
Australia	0.111	0.027	4.141	0.000	[0.058, 0.163]
Brazil	0.256	0.043	5.928	0.000	[0.171, 0.341]
Canada	0.122	0.023	5.295	0.000	[0.077, 0.167]
China	0.156	0.017	8.962	0.000	[0.122, 0.190]
Denmark	0.080	0.018	4.571	0.000	[0.046, 0.115]
Euro Zone	0.117	0.020	5.855	0.000	[0.078, 0.156]
Japan	0.128	0.029	4.338	0.000	[0.070, 0.185]
Korea	0.120	0.015	7.918	0.000	[0.090, 0.150]
Mexico	0.175	0.028	6.189	0.000	[0.120, 0.231]
New Zealand	0.125	0.022	5.650	0.000	[0.081, 0.168]
Norway	0.154	0.027	5.763	0.000	[0.102, 0.206]
Poland	0.153	0.026	5.877	0.000	[0.102, 0.204]
Sweden	0.110	0.021	5.352	0.000	[0.070, 0.150]
Switzerland	0.070	0.023	3.102	0.002	[0.026, 0.114]
UK	0.130	0.016	8.118	0.000	[0.099, 0.161]

Model fit: $R^2 = 0.059$, Adjusted $R^2 = 0.054$.
Notes: Number of observations: 2,655. Robust standard errors (HAC, 11 lags) are used to address autocorrelation due to overlapping 12-month returns. Country dummies absorb cross-sectional heterogeneity. The regression excludes the Global Financial Crisis period (2007–2009). Homogeneous slopes are imposed across countries.

Table A.2: Bilateral Regressions: UEP Deviation on UIP Deviation and FX Hedging Role

Country	UIP Deviation	FX Hedging Role	Constant	R^2	p-value (F-test)
Australia	0.404	–7.596	0.113	0.065	0.338
	[–0.135, 0.942]	[–45.181, 29.990]	[0.053, 0.173]		
Canada	–1.169	28.940	0.128	0.074	0.297
	[–2.633, 0.296]	[–50.039, 107.918]	[0.069, 0.186]		
Switzerland	–0.468	13.165	0.119	0.057	0.382
	[–1.173, 0.237]	[–49.216, 75.546]	[0.063, 0.176]		
Denmark	–0.514	16.375	0.105	0.094	0.155
	[–1.033, 0.005]	[–42.567, 75.317]	[0.078, 0.132]		
Euro Zone	–0.931	29.146	0.169	0.182	0.053
	[–1.682, –0.180]	[–23.445, 81.737]	[0.117, 0.222]		
UK	–0.499	3.864	0.136	0.033	0.412
	[–1.357, 0.359]	[–103.429, 111.157]	[0.113, 0.158]		
Japan	–0.142	22.781	0.144	0.026	0.621
	[–0.823, 0.539]	[–77.816, 123.379]	[0.068, 0.219]		
Korea	0.076	29.550	0.133	0.032	0.499
	[–0.275, 0.427]	[–24.194, 83.294]	[0.088, 0.178]		
Norway	0.202	–20.482	0.129	0.029	0.405
	[–0.659, 1.064]	[–50.282, 9.317]	[0.068, 0.190]		
Poland	0.017	18.948	0.158	0.036	0.590
	[–0.603, 0.637]	[–24.553, 62.449]	[0.089, 0.228]		
Sweden	–0.035	47.258	0.155	0.175	0.025
	[–0.690, 0.620]	[10.352, 84.163]	[0.090, 0.220]		
Brazil	0.431	–10.805	0.315	0.212	0.008
	[0.105, 0.757]	[–35.486, 13.877]	[0.160, 0.470]		
New Zealand	0.023	12.474	0.122	0.005	0.892
	[–0.638, 0.684]	[–47.251, 72.199]	[0.062, 0.182]		
China	0.184	14.846	0.160	0.032	0.388
	[–0.116, 0.485]	[–7.247, 36.938]	[0.112, 0.208]		
Mexico	1.063	–47.372	0.327	0.503	0.000
	[0.620, 1.507]	[–61.938, –32.806]	[0.241, 0.413]		

Notes: Each row corresponds to a separate OLS regression of the UEP deviation on the UIP deviation and FX hedging role, estimated at the bilateral country level. Confidence intervals (95%) are reported in brackets. HAC standard errors with 11 lags are used. The table illustrates cross-country heterogeneity and complements the panel regressions by allowing coefficients to vary across countries.

Appendix B

Appendix for Chapter 2

B.1 Data

B.1.1 Section 3: More Details

Returns data are compiled from the Kenneth French's online library, which uses CRSP data for stocks and the one-month Treasury bill rate (from Ibbotson Associates) for risk-free returns. We transform these series into real terms adjusting for CPI inflation (explained below). Adjustments described in the main paper are performed on each of these datasets – see more details below in this appendix. Ignoring observations lost by applying lagged instruments in the estimation, our initial dataset covers the period between 1931 to 2022 for annual and between 1947:3 to 2023:2 for quarterly data. Our original consumption data (or equivalently, reported consumption) come from NIPA tables, available on BEA's website. We use two time series: consumption of nondurables and services and consumption of nondurables (only). As explained in the main paper, unfiltered consumption has been constructed from these two series under distinct calibrations. For annual data, we construct unfiltered consumption from the original model in Kroencke (2017), which does not feature serially correlated measurement errors. For quarterly data, these are relevant, so that we introduce such form of persistence relying on the quasi-differenced filter described in section 2 of chapter 2. For more details on how we calibrate the model, see section “calibration" below.

For other series, inflation rate uses quarter-over-quarter and year-over-year CPI data, the nominal interest rate is the same used above (from Kenneth French’s website) and the dividend-price ratio has been taken from Robert Shiller’s online data source. Recall that we take logs of the latter.

B.1.2 Section 4: More Details

For stocks, we use value-weighted returns that consider NYSE, NASDAQ and AMEX. For T-bill returns, we rely on the same dataset of section 3 of chapter 2. To calculate the bond default premium, we use the Moody’s Seasoned BAA Corporate Bond Yield (which is based on bonds with maturities of 20 years and above) as the long term corporate yield. To compute the bond horizon premium, we rely on 20-year and 1-month T-bill rates (Federal Reserve). From January 1984 to September 1993, the 20-year data are not available, so we use the 10-year analogue instead.

As indicated in the main paper, consumption growth data is constructed from the CEX interviews. These are deflated using the CPI deflator for nondurables considering urban households. To deflate return series mentioned above we use the CPI for total consumption (also for urban households).

B.1.3 A Few Notes on Data Sources involving Consumption, Frequencies and Measurement Errors

The CEX encompasses two major data sources, the interview and the diary surveys. They do not share the same sample. We use the former to construct our consumption data series for the groups of asset holders. Data in the diary survey is much more detailed, and likely more prone to measurement error and misreporting issues. However, the interview survey is also prone to such problems. Indeed, data that are probably misreported is easily verified for several of the questions used to construct consumption in section 4. For instance: households that report two distinct quantities for the consumption of the same item, in the same month of the same year; households that report negative values for some item; or even households that do not report consumption for some quarter (or, some interview), but that report numbers to other questions, not related to consumption (these households are included in our sample).

The BLS has been systematically attempting to change the methodology, so that respondent burden and measurement errors are less significant. The Gemini project, which was mentioned in the main paper, is an example.

Data Sources Comparison: NIPA vs. CEX

There are many methodological differences between consumption measured by the PCE (or NIPA, from BEA) and the CEX (from BLS).

First, consumption measured by the BLS takes into account that the data source varies with the sampling frequency of consumption data. For example, monthly and quarterly data in the PCE are based on the monthly retail trade survey (MRTS), while annual data comes from the annual retail trade survey (ARTS). It is known that sampling errors are more problematic in the former, and the BEA takes this into account. In so far as we can tell, adjustments applied by the BLS are not specific to any major data source in the CEX. Hence, even though the interview and the diary survey have distinct samples, the Filter model could be also applied to the latter.

Second, the BEA benchmarks quarterly and annual data to "the best available source data", which happens to be the quinquennial – see BEA (2017). It follows that quarterly and annual frequency data are interpolated, so that they are compatible with benchmark years. In a second step, quarterly estimates are benchmarked to their annual counterparts. Bell and Wilcox (1993) argue that benchmark procedures reduce measurement errors, inherently affecting the autocorrelation in the consumption series. Contrasting with the PCE, the CEX does not benchmark the data. Instead, the survey is continuously redesigned to circumvent issues of measurement error.

Imputation is present in both the PCE and the CEX data. BEA and BLS rely on statistical models for non-response to predict missing values. These are considerably in the ARTS, about 8%. The CEX did not apply imputation to asset data, but began to do in 2004. Allocation routines are also common across both sources. The CEX applies tabulation corrections *before* and *after* other adjustment routines.

Residual methods are applied by the BEA to measure the consumption of some categories. They use "residuals" from government expenditures to do so. As far as we can tell, these routines are absent in the CEX.

Finally, the BLS applies smoothing techniques over the data. Direct forms of smoothing are absent in the CEX, albeit topcoding routines are present, generating similar effects. Topcoding techniques modify the consumption, positions in assets and the income data of outliers, so that these can not be identified from the public micro-files. Thresholds applied are also constantly under revision, aiming to correctly disentangle "true" outliers from misrespondents or coding errors.

B.2 Fully Specified Quasi-Differenced Filter Model

As mentioned before, when we referred to annual consumption in section 3 of the main text we were using the canonical filter model in Kroencke (2017). The only difference in that case relates to calibration (since different time series are used). Specifically, measurement errors followed a simple white noise stochastic process with no form of persistence introduced. The quasi-differenced filter model (that accounts for such persistence) is only used when handling quarterly data in section 3, and semi-annual data (but at monthly frequency), in section 4. This follows for the reasons described in the main text. In this section, we provide the complete specification of this quasi-differenced version and better detail how we modify the original Filter model in Kroencke (2017).

Recall that a Filter model without persistence in measurement error is not suitable for data at monthly or quarterly frequencies. The more frequent the data (or the smaller the level of aggregation) the more likely accounting for a serially correlated error becomes essential since sampling errors are probably autocorrelated – see the online appendix of Kroencke (2017). In our estimations, unfiltered consumption performed poorly in terms of estimates of the EIS when serially correlated measurement error terms are not considered in the model and the data frequency is either monthly or quarterly. Comparatively, even reported consumption provided more precise estimates in those cases.

Turning to the model, assume a simple state-space representation:

$$x_{t+1} = Fx_t + R\eta_{t+1}, \tag{B.1}$$

$$y_t = Hx_t + \xi_t, \tag{B.2}$$

those representing the state and measurement equations of a simple Kalman filter, respectively. x_t represents a vector of state variables and the last term is its corresponding disturbance. y_t is observed consumption (it can be the garbage measure in Savov (2011), for instance). Unfiltered consumption is our estimate of this time series, while ξ_t represents measurement errors.

By permitting serially correlated measurement error terms we introduce (2.3) of the main text in the system:

$$\xi_t = \rho_\xi \xi_{t-1} + \nu_t,$$

where ν_t is a simple white noise process. Generally, the usual assumption is $E[\xi_t \eta'_{t+h}] = E[\eta_t \eta'_s] = E[\xi_t \xi'_s] = 0$, for $t \neq s$ and $h > 0$ – see Hamilton (1994), for example. We relax this hypothesis in order to introduce the possibility of (2.3). Particularly, let us assume that the innovation of state consumption and that of measurement errors are conditionally normally distributed:

$$\begin{bmatrix} \eta_{t+1} \\ \\ \xi_t \end{bmatrix} \sim N\left(\begin{bmatrix} 0 \\ \\ 0 \end{bmatrix}, \begin{bmatrix} R\sigma^2_{\eta,t+1}R' & R\omega_{\eta,\nu} \\ \\ \omega_{\nu,\eta}R' & \sigma^2_\nu \end{bmatrix} \right), \tag{B.3}$$

where we are not assuming any zeros in the covariance matrix, but we allow for a time-varying element in its upper left corner.

Next, define:

$$P_{t+1|t} = E[(x_{t+1} - \hat{x}_{t+1|t})(x_{t+1} - \hat{x}_{t+1|t})'|\mathcal{F}_t], \tag{B.4}$$

where the information set $\mathcal{F}_t$ tracks data realisations conditional on period t and $P_{t+1|t}$ is the covariance of prediction errors conditional on the same period (*a priori*). With a small abuse of notation, we denoted its first element (relative to consumption) by P^c_t in section 2 of chapter 2 – we turn back to this below. In addition, $\hat{x}_{t+1|t} = E[x_{t+1}|\mathcal{F}_t]$, as usually. Note that $x_{t+1|t-1} \sim N(\hat{x}_{t+1|t-1}, P_{t+1|t-1})$, and:

$$\hat{x}_{t+1|t-1} = F\hat{x}_{t|t-1}, \tag{B.5}$$

by (1). In a similar vein and using $E[x_t \eta'_{t+1}] = E[\hat{x}_{t|t-1}\eta'_{t+1}] = 0$:

$$P_{t+1|t-1} = E[(F(x_t - \hat{x}_{t|t-1}) + R\eta_{t+1})(F(x_t - \hat{x}_{t|t-1}) + R\eta_{t+1})'|\mathcal{F}_{t-1}] = FP_{t|t-1}F' + R\sigma^2_{\eta,t+1}R', \tag{B.6}$$

Likewise:

$$\hat{y}_{t|t-1} = H\hat{x}_{t|t-1}, \qquad S_{t|t-1} = HP_{t|t-1}H' + \sigma^2_\nu, \tag{B.7}$$

where $S_{t|t-1} = E[(y_t - \hat{y}_{t|t-1})(y_t - \hat{y}_{t|t-1})'|\mathcal{F}_{t-1}]$, the covariance of pre-fit prediction errors. Finally, if $\Sigma_{t|t-1} = E[(y_t - \hat{y}_{t|t-1})(x_{t+1} - \hat{x}_{t+1|t})'|\mathcal{F}_{t-1}]$ is the cross-correlation matrix between state and observable variables, then:

$$\Sigma_{t|t-1} = FP_{t|t-1}H' + \omega_{\nu,\eta}R', \tag{B.8}$$

so that:

$$\begin{bmatrix} x_{t+1}|\mathcal{F}_{t-1} \\ \\ y_t|\mathcal{F}_{t-1} \end{bmatrix} \sim N\left(\begin{bmatrix} F\hat{x}_{t|t-1} \\ \\ H\hat{x}_{t|t-1} \end{bmatrix}, \begin{bmatrix} P_{t+1|t-1} & \Sigma_{t|t-1} \\ \\ \Sigma'_{t|t-1} & S_{t|t-1} \end{bmatrix}\right). \tag{B.9}$$

From (B.8), one can thus express the distribution of $x_{t+1}|\mathcal{F}_t$ by marginalising $x_{t+1}|\mathcal{F}_{t-1}$ in terms of $y_t|\mathcal{F}_{t-1}$. Hence, relying on the multivariate normal:

$$\begin{aligned} \hat{x}_{t+1|t} &= F\hat{x}_{t|t-1} + \Sigma_{t|t-1}S^{-1}_{t|t-1}(y_t - H\hat{x}_{t|t-1}) \\ &= F\hat{x}_{t|t-1} + (FP_{t|t-1}H' + R\omega_{\eta,\nu})(HP_{t|t-1}H' + \sigma^2_\nu)^{-1}(y_t - H\hat{x}_{t|t-1}), \end{aligned} \tag{B.10}$$

using (B.6) and (B.7). In a similar fashion:

$$\begin{aligned} P_{t+1|t} &= P_{t+1|t-1} - \Sigma_{t|t-1}S^{-1}_{t|t-1}\Sigma'_{t|t-1} \\ &= FP_{t|t-1}F' + R\sigma^2_{\eta,t+1}R' - (FP_{t|t-1}H' + R\omega_{\eta,\nu})(HP_{t|t-1}H' + \sigma^2_\nu)^{-1}(FP_{t|t-1}H' + \omega_{\eta,\nu}R')' \end{aligned} \tag{B.11}$$

Finally, the Kalman gain is simply:

$$K_t = \Sigma_{t|t-1}S^{-1}_{t|t-1} = (FP_{t|t-1}H' + R\omega_{\eta,\nu})(HP_{t|t-1}H' + \sigma^2_\nu)^{-1} \tag{B.12}$$

One can derive the original filter in Kroencke (2017) with the system above. However, up to this point, we have not allowed for serially correlated measurement errors – equation

(2.3) of the main text – yet. Perhaps the easiest way to introduce this possibility in a Kalman filter is to expand the model so that measurement errors are defined as a new state variable. Harvey, Ruiz, and Sentana (1992) present different forms of modelling that while still relying on ARCH or GARCH processes to express variances (as we do here). However, note that we actually aim to *invert* a Kalman filter. This is, we are not interested in tracking state variables given observables. Instead, we aim to infer what would those observables were once we (researchers) *only* know estimates of state variables that supposedly took those into account to be constructed. As in Kroencke (2017), we do *not* use simple recursions of a Kalman filter but instead their reverse counterparts. With that in mind, re-scaling the system is much simpler than developing alternative methods of "reverse engineering" that support an expanded system that includes (2.3). We then opt to re-scale the original Kalman filter above, expressing it in terms of a quasi-differenced system. This re-scaled filter does not modify the standard interpretation given in Kroencke (2017).

In the following, we blend findings of E. Anderson et al. (1996) with the original Filter model, establishing our quasi-differencing approach. Assume that $\nu_t \sim N(0, \sigma_\nu^2)$, the error term in equation (2.3). Next, we define observables in terms of a quasi-difference:

$$\overline{y}_t \equiv y_{t+1} - \rho_\xi y_t, \tag{B.13}$$

where $\overline{y}_t$ is referred here as "quasi-differenced observable consumption". Note that the state equation (2.1) does not change with this modification, but the measurement equation is transformed into[1]:

$$\overline{y}_t = (HF - \rho_\xi H)x_t + HR\eta_{t+1} + \nu_{t+1} \equiv \overline{H}x_t + \overline{\xi}_t. \tag{B.14}$$

It is worth mentioning that by rewriting the Filter model in terms of a quasi-differenced

[1]Write the measurement equation one period forward. Use the state equation in x_{t+1}. Then, subtract $\rho_\xi \times y_t$ from this, using the measurement equation in the current period as y_t.

observable we are *not* assuming that the raw data first observed by official statisticians is $\overline{y}_t$. Instead, *we are only re-scaling the system in order to solve it, to then mapping back $\overline{y}_t$ onto y_t.* This will probably become more clear below.

By following similar developments, we can rewrite covariances in terms of the new composite error term in (13). Particularly, if $R\eta_{t+1} \equiv \overline{\eta}_{t+1}$, it follows that $\sigma^2_{\overline{\xi}} = HR\sigma^2_{\eta,t+1}R'H' + \sigma^2_{\nu}$ and $\omega_{\overline{\eta},\overline{\xi}} = R\sigma^2_{\eta,t+1}R'H'$. Applying the same algebra as above:

$$\begin{aligned}\hat{x}_{t+1|t} =& F\hat{x}_{t|t-1} \\ & + (FP_{t|t-1}\overline{H}' + R\sigma^2_{\eta,t+1}R'H')(\overline{H}P_{t|t-1}\overline{H}' + HR\sigma^2_{\eta,t+1}R'H' + \sigma^2_{\nu})^{-1}(\overline{y}_t - \overline{H}\hat{x}_{t|t-1})\end{aligned} \tag{B.15}$$

$$\begin{aligned}P_{t+1|t} =& FP_{t|t-1}F' + R\sigma^2_{\eta,t+1}R' \\ & - (FP_{t|t-1}\overline{H}' + R\sigma^2_{\eta,t+1}R'H')(\overline{H}P_{t|t-1}\overline{H}' + HR\sigma^2_{\eta,t+1}R'H' + \sigma^2_{\nu})^{-1} \\ & (\overline{H}P_{t|t-1}F' + HR\sigma^2_{\eta,t+1}R'),\end{aligned} \tag{B.16}$$

with the Kalman gain vector following:

$$K_t = (FP_{t|t-1}\overline{H}' + R\sigma^2_{\eta,t}R'H')(\overline{H}P_{t|t-1}\overline{H}' + HR\sigma^2_{\eta,t}R'H' + \sigma^2_{\nu})^{-1} \tag{B.17}$$

Note that $[\overline{y}_t, \overline{y}_{t-1}, ..., \overline{y}_0, \hat{x}_0]$ and $[y_{t+1}, y_t, ..., y_0, \hat{x}_0]$ span the same space since:

$$\begin{aligned}\overline{y}_t - E[\overline{y}_t|\overline{y}_{t-1}, ..., \overline{y}_0, \hat{x}_0] &= (y_{t+1} - Dy_t) \\ &\quad - E[y_{t+1} - Dy_t|y_t - Dy_{t-1}, y_{t-1} - Dy_{t-2}, ..., y_0, \hat{x}_0] \\ &= y_{t+1} - Dy_t + Dy_t - E[y_{t+1}|y_t, y_{t-1}, y_{t-2}, ..., y_0, \hat{x}_0] \\ &= y_{t+1} - E[y_{t+1}|y_t, y_{t-1}, y_{t-2}, ..., y_0, \hat{x}_0]\end{aligned} \tag{B.18}$$

Equations (B.14) to (B.16) express the algorithm in the quasi-differenced Kalman filter.

The next step is to invert those equations to isolate quasi-differenced unfiltered consumption. Before, let's identify our system in terms of our findings above. State variables are defined as:

$$x_t = \begin{bmatrix} c_t \\ c_{t-1} \\ \eta_t^* \end{bmatrix}, \qquad F = \begin{bmatrix} 1 & 0 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}, \qquad R = \begin{bmatrix} 1 \\ 0 \\ 1 \end{bmatrix} \qquad \eta_t^* = \sigma_{\eta,t}^2 \eta_t, \tag{B.19}$$

while we assume, as in Kroencke (2017), that its variance follows a GARCH(1,1) stochastic process given by:

$$\sigma_{\eta,t}^2 = a_0 + a_1 \eta_{t-1}^{*2} + a_2 \sigma_{\eta,t-1}^2 \tag{B.20}$$

We can express the variance of prediction errors as:

$$\begin{aligned} S_{t|t-1} &= \begin{bmatrix} 1-\rho_\xi & 0 & 0 \end{bmatrix} P_{t|t-1} \begin{bmatrix} 1-\rho_\xi \\ 0 \\ 0 \end{bmatrix} + \sigma_{\eta,t}^2 + \sigma_\nu^2 \\ &= (1-\rho_\xi)^2 P_{t|t-1}^c + \sigma_{\eta,t-1}^2 + \sigma_\nu^2, \end{aligned} \tag{B.21}$$

where $P_{t|t-1}^c$ denotes the element (1,1) – related to the state variable c_t – of covariance matrix $P_{t|t-1}$, and we have used the fact that $H = \begin{bmatrix} 1 & 0 & 0 \end{bmatrix}$ and then $\overline{H} = \begin{bmatrix} 1-\rho_\xi & 0 & 0 \end{bmatrix}$. One can then obtain the first component of the Kalman gain in (16):

$$K_t = \left\{ \begin{bmatrix} 1 \\ 0 \\ 1 \end{bmatrix} \begin{bmatrix} 1 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} \sigma^2_{\eta,t} + \begin{bmatrix} 1 & 0 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} P_{t|t-1} \begin{bmatrix} 1-\rho_\xi \\ 0 \\ 0 \end{bmatrix} \right\} \\ [(1-\rho_\xi)^2 P^c_{t|t-1} + \sigma^2_{\eta,t} + \sigma^2_\nu]^{-1}, \tag{B.22}$$

where the matrix pre-multiplying $P_{t|t-1}$ is F. Solving (B.22) for its first element, we obtain the Kalman gain – relative to consumption – described in the main text, equation (2.6):

$$K^c_t = \frac{(1-\rho_\xi)P^c_{t|t-1} + \sigma^2_{\eta,t}}{(1-\rho_\xi)^2 P^c_{t|t-1} + \sigma^2_{\eta,t} + \sigma^2_\nu},$$

and in section 2 we used the notation P^c_t instead than $P^c_{t|t-1}$.

Next, use (B.16), writing it in terms of our model:

$$P_{t|t-1} = \begin{bmatrix} 1 & 0 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} P_{t-1|t-2} \begin{bmatrix} 1 & 1 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} + \sigma^2_{\eta,t} \begin{bmatrix} 1 & 0 & 1 \\ 0 & 0 & 0 \\ 1 & 0 & 1 \end{bmatrix} \\ - K_{t-1} \left\{ \begin{bmatrix} 1-\rho_\xi & 0 & 0 \end{bmatrix} P_{t-1|t-2} \begin{bmatrix} 1 & 1 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} + \begin{bmatrix} 1 & 0 & 0 \end{bmatrix} \begin{bmatrix} 1 & 0 & 1 \\ 0 & 0 & 0 \\ 1 & 0 & 1 \end{bmatrix} \sigma^2_{\eta,t} \right\}, \tag{B.23}$$

whose first element is:

$$P^c_{t|t-1} = P^c_{t-1|t-2}(1 - (1-\rho_\xi)K^c_{t-1}) + (1 - K^c_{t-1})\sigma^2_{\eta,t},$$

Recall that $P^c_t \equiv P^c_{t|t-1}$, in our notation of section 2.

Our quasi-differenced Filter model is almost completely described now. The next step is to obtain our best guess of $\overline{y}_{t-1}$. In order to do that, write (B.15) for the update phase ($\hat{x}_{t|t}$). Substituting relevant matrices (B.19), the first element of $\hat{x}_{t|t}$ is:

$$\hat{c}_t = \hat{c}_{t-1} + K^c_t(\overline{y}_{t-1} - (1-\rho_\xi)\hat{c}_{t-1}), \tag{B.24}$$

where $\hat{c}_t \equiv E[c_t|\mathcal{F}_t]$. Isolating $\overline{y}_{t-1}$:

$$\underbrace{\hat{\overline{y}}_{t-1} = \frac{\hat{c}_t - (1 - (1-\rho_\xi)K^c_t)\hat{c}_{t-1}}{K^c_t}}_{\text{Unfiltered Quasi-Differenced Consumption}} \tag{B.25}$$

Equation (B.25) is *not* ready to be mapped back onto original unfiltered consumption yet. We still need a few adjustments, described in the following.

B.2.1 Adjusting Unfiltered Consumption for Time-Aggregation Bias

First, we adapt (B.25), so that it accounts for time-aggregation bias. Hall (1988) addressed this using an AR(2) representation. In a similar fashion and adapting adjustments in Kroencke (2017) to your model, we have[2]:

$$\Delta c^{adTA}_t = \frac{[\Delta c^{TA}_t - (1 - (1-\rho_\xi)\alpha)\Delta c^{TA}_{t-1}]}{\alpha}, \tag{B.26}$$

[2]See Kroencke (2017) for more details.

where Δc_t^{adTA} denotes the time-aggregation bias-adjusted estimate of consumption growth and Δc_t^{TA} represents its time-aggregated counterpart. The parameter α guarantees that the second moment of Δc_t^{adTA} is the same of point-to-point consumption. It is set to the same value of Kroencke (2017), 0.8[3]. Adapting (B.25) to our model gives:

$$\hat{\hat{y}}_{t-1} = \frac{\hat{c}_t - (1 - (1 - \rho_\xi)\Omega_t)\hat{c}_{t-1}}{\Omega_t}, \tag{B.27}$$

where $\Omega_t = \alpha K_t$.

B.2.2 Adjusting the Timing of Asset Returns

By construction, the variable in (B.27) has its second moment perfectly compatible with consumption measured point-to-point in time. However, the timing of asset returns is misaligned. We correct for this in section 3 by adapting adjustments in Kroencke (2017) – which was based on Cochrane (1996) – for the use of quarterly series.

First, we sum end-of-month levels $\Pi_{i,m,t+1}$ to obtain a measure of quarterly time-aggregated stock returns:

$$\Delta R_{i,t+1}^{TA} = \frac{\sum_{m=1}^{3} \Pi_{i,m,t+1}}{\sum_{m=1}^{3} \Pi_{i,m,t}} - 1, \tag{B.28}$$

where i represents the asset class and m is the corresponding month of quarter t.

Second, we bring first and second moments of this series back to point-to-point counterparts to make it compatible with (B.27). We conduct this adjustment using returns for the last quarter of the year. The motivation is in Jagannathan and Wang (2007), who argued that investors are more prone to adjust their investment portfolios in the fourth quarter of

[3]It solves $Var(\Delta c_t^{TA}) = \frac{\alpha}{2-\alpha} Var(\Delta c_t^{adTA}) = \frac{2}{3} Var(\Delta c_t^{adTA})$. The former equality is implied from (34) setting $\rho_\xi = 0$ while the latter uses results in Working (1960) and Breeden, Gibbons, and Litzenberger (1989). The approximation for ρ_ξ does not distort results sensibly since we have set that parameter to a value very close to zero in our estimations. The first equality does not change in comparison with Kroencke (2017) since our model still relies on a random-walk representation for state consumption whose conditional variances are modelled through a GARCH(1,1).

the year.

$$\Delta R_{i,t+1}^{adTA} = \frac{\Delta R_{i,t+1}^{TA} - E(\Delta R_{i,t+1}^{TA})}{\sigma(\Delta R_{i,t+1}^{TA})}\sigma(\Delta R_{i,t+1}^{Q4-Q3}) + E(\Delta R_{i,t+1}^{Q4-Q3}), \tag{B.29}$$

where $\Delta R_{i,t+1}^{Q4-Q3} = \Pi_{i,12,t+1}/\Pi_{i,9,t+1}$, returns measured for the last quarter[4].

Step (B.27) is not necessary in section 4 since data were at monthly frequency. This simply implies that $\Delta R_{i,t+1}^{TA}$ is equal to (raw) monthly returns used. When it comes to (27), we do the following, for section 4:

$$\Delta R_{i,t+1}^{adTA} = \frac{\Delta R_{i,m+1}^{TA} - E(\Delta R_{i,m+1}^{TA})}{\sigma(\Delta R_{m,t+1}^{TA})}\sigma(\Delta R_{i,December}) + E(\Delta R_{i,December}), \tag{B.30}$$

where sub-index "December" denotes monthly returns in December of the relevant year. Results of section 4 do not change if we use October or November instead.

Finally, $\Delta R_{i,t+1}^{adTA}$ in (B.29) is aggregated to represent semi-annual returns as described in the main text – this is, we used $R_{i,t+1}$ instead of $\Delta R_{i,t+1}^{adTA}$ for simplicity reasons in section 4.

B.2.3 Mapping Unfiltered Quasi-Differenced Consumption Back onto Unfiltered Consumption

The last step is the simplest one. In order to map $\hat{\overline{y}}_{t-1}$ back onto $\hat{y}_t$ – unfiltered consumption – we perform:

$$\underbrace{\hat{y}_t}_{\text{Unfiltered Consumption}} = \hat{\overline{y}}_{t-1} + \rho_\xi \hat{y}_{t-1} \tag{B.31}$$

[4]Kroencke (2017) conducted a similar adjustment but for annual series, using the first two moments of December-to-December consumption growth as his correction in a similar fashion. Note that by using moments of the last quarter components $E(\Delta R_{i,t+1}^{Q4-Q3})$ and $\sigma(\Delta R_{i,t+1}^{Q4-Q3})$ change every 4 observations (or equivalently, every 4 quarters) - contrasting with corrections in Kroencke (2017), that change for each observation (since described in annual terms). This could imply an unnecessary persistence for the series $\Delta R_{i,t+1}^{adTA}$. However, by comparing return series generated by (28) and (29) with their raw analogues – and repeating the same experiment for annual series, using the method in Kroencke (2017) – we have found that the impact of those modifications for our estimates is minimal. Hence, we evaluate that our adjustments for quarterly data do not perform considerably different from those of the original model. For complete results involving raw returns data, see section 6.4.5 below.

The fact that $\hat{y}_{t-1}$ appears in the right hand side of (B.31) implies we can not identify y_0. We turn back to this and how we initialise the model below.

B.2.4 Consumption Volatility

Motivated by results in Harvey, Ruiz, and Sentana (1992), we use the approximation $\eta^{*2}_{t-1} \approx E_{t-1}(\eta^{*2}_{t-1})$. It follows that $E_t(\eta^{*2}_t)$ – and hence $E_{t-1}(\eta^{*2}_{t-1})$ – is obtained by:

$$\begin{aligned} E_t(\eta^{*2}_t) &= P^{\eta}_{t|t} + \eta^{*2}_{t|t} \\ &= \left(1 - \frac{\sigma^2_{\eta,t-1}}{P^c_{t|t-1}(1-\rho_\xi)^2 + \sigma^2_{\eta,t-1} + \sigma^2_\nu}\right)\sigma^2_{\eta,t-1} \\ &+ \left(\frac{\sigma^2_{\eta,t-1}}{P^c_{t|t-1}(1-\rho_\xi)^2 + \sigma^2_{\eta,t-1} + \sigma^2_\nu}\right)^2 u^2_t, \end{aligned} \tag{B.32}$$

where $u_t \equiv \overline{y}_{t-1} - (1-\rho_\xi)\hat{c}_{t-1}$, the re-scaled prediction error.

B.2.5 Homoscedastic Counterpart and Proof of Proposition 1

It is simple to derive a version of the model that features homoscedasticity in state consumption. This not only implies $\sigma^2_{\eta,t} = \overline{\sigma}^2_\eta$, but also $K^c_t = \overline{K}^c$ and $P^c_{t|t-1} = \overline{P}^c$.

Under homoscedasticity, we have:

$$\overline{P}^c = \overline{P}^c(1 - (1-\rho_\xi)\overline{K}^c) + (1 - \overline{K}^c)\overline{\sigma}^2_\eta, \tag{B.33}$$

and the Kalman gain becomes:

$$\overline{K}^c = \frac{(1-\rho_\xi)\overline{P}^c + \overline{\sigma}^2_\eta}{(1-\rho_\xi)^2\overline{P}^c + \overline{\sigma}^2_\eta + \sigma^2_\nu}. \tag{B.34}$$

By plugging (B.34) in (B.33) and after some algebraic manipulations, one can finally find the second order equation:

$$\{\overline{P}^c\}^2 + \overline{\sigma}_\eta^2 \left(\frac{1+\rho_\xi}{1-\rho_\xi}\right) \overline{P}^c - \frac{\overline{\sigma}_\eta^2 \sigma_\nu^2}{(1-\rho_\xi)^2} = 0 \tag{B.35}$$

After rewriting terms, it is possible to show that the roots of (B.35) are given by:

$$\overline{P}^c = \frac{\overline{\sigma}_\eta^2}{2(1-\rho_\xi)} \left\{ \pm \left[(1-\rho_\xi)^2 \overline{\sigma}_\eta^2 + 4\sigma_\nu^2\right]^{\frac{1}{2}} - (1+\rho_\xi)\overline{\sigma}_\eta^2 \right\}, \tag{B.36}$$

so that the only sensible solution is (recall that $1-\rho_\xi > 0$):

$$\overline{P}^c = \frac{\overline{\sigma}_\eta^2}{2(1-\rho_\xi)} \left\{ \left[(1-\rho_\xi)^2 \overline{\sigma}_\eta^2 + 4\sigma_\nu^2\right]^{\frac{1}{2}} - (1+\rho_\xi)\overline{\sigma}_\eta^2 \right\} \tag{B.37}$$

Given (B.37), our model exhibits the expected behaviour if:

$$4\frac{\sigma_\nu^2}{\overline{\sigma}_\eta^2} > (1+\rho_\xi)^2 \overline{\sigma}_\eta^2 - (1-\rho_\xi)^2. \tag{B.38}$$

Equations (B.33) and (B.34) characterise the homoscedastic model. Those are also useful when initialising the heteroscedastic Filter model – see more details below, in this appendix. For our parameterisation, it follows that the right hand side of (B.38) is negative and even if otherwise a sizeable σ_ν^2 compared with $\overline{\sigma}_\eta^2$ would do the trick. Generally, we have found that (B.37) is met under an ample set of realistic parameterisations.

As mentioned in the main paper, one can evaluate how well a Filter model behaves by checking whether its corresponding Kalman gain increases during a period of economic turbulence (recessions, for instance). This generates an unfiltered series less persistent and probably more connected with movements in the assets market.

The Kalman gain in our model is more complicated than that in Kroencke (2017), but we can still see that a more volatile measurement error lowers $\overline{K}^c$ and expectations of the state of consumption do not change much, in line with the intuition. Now we must ensure that reported consumption is less filtered when $\overline{P}^c$ and(or) $\overline{\sigma}_\eta^2$ are higher. By taking derivatives of the Kalman gain, it is possible to show that:

$$\frac{\partial \overline{K}^c}{\partial \overline{P}^c} = \left\{(1-\rho_\xi)\Phi - (1-\rho_\xi)^2[(1-\rho_\xi)\overline{P}^c + \overline{\sigma}_\eta^2]\right\}\Phi^{-2}, \tag{B.39}$$

$$\frac{\partial K_t}{\partial \overline{\sigma}_\eta^2} = \left\{\Phi - (1-\rho_\xi)\overline{P}^c + \overline{\sigma}_\eta^2\right\}\Phi^{-2}, \tag{B.40}$$

where $\Phi = (1-\rho_\xi)^2\overline{P}^c + \overline{\sigma}_\eta^2 + \sigma_\nu^2 > 0$. It is straightforward to see that (B.39) is always positive while (B.40) is positive for ρ_ξ small enough, so that $\sigma_\nu^2 - (1-\rho_\xi)\overline{P}^c\rho_\xi > 0$. Our calibration – when $\rho_\xi = 0.06$ – easily meets this condition. The fact that (B.39) and (B.40) hold in our model then ensures its validity. We have also found that both conditions are also met in the heteroscedastic model – when derivatives those are time-dependent.

B.2.6 Initialising the Model

We start the model using its homoscedastic analogue described above, so that $P^c_{t=1} = \overline{P}^c$ and $K^c_{t=1} = \overline{K}^c$. In addition, based on the long term representation of the GARCH process in (12): $\sigma^2_{\eta,t=1} = \alpha_0/(1-\alpha_1-\alpha_2)$[5]. As mentioned earlier, by construction we are not able to identify $\hat{y}_0$ when using the quasi-differenced Filter model. Therefore, we initialise the filter assuming that $\Delta y_{t=1} = 0$, then burning the first observations for which consumption growth seems to exhibit an abrupt and unrealistic movement. It follows that we burned the first three observations when dealing with quarterly data in section 3 but none when using annual data – the latter justified by the low number of observations available. Below we repeat methods in section 3, while restricting the sample to 1960:1–2023:2 (1940–2022) for quarterly (annual) data. In section 4, we burned the first 7 months of data – recall that consumption growth is semi-annual but the frequency is monthly.

[5]Being more specific, we do $\sigma^2_{\eta,t=1} = \overline{\sigma}^2_\eta$ and calibrate a_1 and a_2 based on bechmark moments of section 2 such that a_0 is uniquely determined.

Notes on Calibration

Recall that for quarterly data in section 3, we use the quasi-differenced Filter model shown in section 2. In that case we fix $\overline{\sigma}_\eta = 0.0078 * \sqrt{3} \approx 1.4\%$, adapting similar results with monthly data presented in Bansal and Yaron (2004) for that frequency. Kroencke (2017) noted that his model seems little sensitive to choices of a_1 and a_2 (GARCH process), once remaining parameters are correctly calibrated to the same moments. We confirmed the same finding for our model. We choose $a_1 = 0.22$ and $a_2 = 0.5$[6]. Since the services component of consumption is quite more imprecise (and volatile) than its nondurables analogue, we fix different values for $\overline{\sigma}_\nu$ based on each type of consumption: nondurables and services or nondurables only. Specifically, we use $\overline{\sigma}_\nu = 3.8\%$ for the former and $\overline{\sigma}_\nu = 2.2\%$ for the latter when applying the heteroscedastic model. For its homoscedastic analogue, we use $\overline{\sigma}_\nu = 2.5\%$ and $\overline{\sigma}_\nu = 2.0\%$, repectively. These values not only match our benchmark moments quite well but also the difference itself makes sense, given the imprecision of the services component mentioned above (implying a lower value of $\overline{\sigma}_\nu$ when that group is removed from the measure). Finally, recall that we establish $\rho_\xi = 0.06$ regardless of the consumption type.

For results involving annual data in section 3 our model does not feature persistent measurement error. Therefore, we follow the exact same steps in Kroencke (2017), with mere alternations in calibration to account for an updated time series (until 2017 instead of 2014)[7]. There we set $\overline{\sigma}_\eta = 2.5\%$, $a_1 = 0.01$ and $a_2 = 0.85$. The former is the same value used in Kroencke (2017). We do not adapt it to our time series since by following the same logic we use for quarterly data would give a very similar value ($0.0078 * \sqrt{12} \approx 2.7\%$) and very similar results. In addition, we fix $\overline{\sigma}_\nu = 2.8\%$ (nondurables and services) and $\overline{\sigma}_\nu = 1.9\%$ (nondurables only – see section 6.4.3 below) when using annual data. These values do not

[6] That implies $a_0 = 0.00005$.

[7] Technically, we start our Filter model in 1930 and its first observation generated for unfiltered consumption relates to 1931. Kroencke (2017) expanded the original time series provided in NIPA tables to the period that 1927-30, so that it encompasses the Great Depression period. He used data available in Robert Shiller's website for that, with the implicit assumption that the representative statistician does not change the hypothetical Filter model across different datasets. We do not use data from that period, so that all our consumption observations come from BEA (NIPA).

change depending on the model used to unfilter consumption.

As in Kroencke (2017), we ensure that the long-term standard deviation (over 6 years) of annual unfiltered consumption is not much more than 1.2 times that of reported consumption (we impose this condition when calibrating $\overline{\sigma}_{\nu}$). It is intuitive that this gap should not be considerably high since: (i) measurement errors should cancel out when consumption is measured over longer horizons, – see Daniel and Marshall (1996) –, and; (ii) filtering procedures should be smart enough such that in reported data is not considerably more volatile than unfiltered data in the long run – presumably, the implicit algorithm would be otherwise corrected to take new evidence and perceived errors into account. That being said, it is clear that the intuition behind that rule does not change regardless of the nature of the stochastic process for the measurement error. For example, for nondurables and services we have that the ratio between long-run standard deviations of unfiltered and reported consumption are 1.12 and 1.23 for quarterly and annual data, respectively. For completeness, in Table B.1 below we present similar results as those shown in Table 2.1 but for consumption of nondurables only. The 1.2 times rule is still valid.

Table 2.2 summarises semi-annual consumption growth moments for CEX data – see section 4. Although we use data from that survey, we benchmark moments to results obtained for unfiltered NIPA consumption in section 3. We restrict moments of the latter to the available sample period of the former accordingly.

It is well known that a large fraction of CEX consumption categories reproduce a similar behaviour compared to NIPA counterparts. Other categories do measure different things or have similar definitions but exhibit a ratio CEX/NIPA that is too low (high) overtime. However, in terms of the estimation of the EIS, it is more central for the Filter model to be able to revert second moments and auto-correlations than to infer how much one source may be overestimating consumption compared to the other. Indeed, if overall there is no considerable change in how much CEX categories overestimate (underestimate) its NIPA analogues, then one can benchmark CEX aggregates to NIPA counterparts to calibrate the

Table B.1: Calibrated Moments for NIPA Consumption of Nondurables Only

(Implied) Consumption Growth	$E(\Delta C_{year})$	$\sigma(\Delta C_{year})$	$\sigma(\sum_{year=1}^{6} \Delta C_{year})/\sqrt{6}$	$Corr(\Delta C_{year}, \Delta C_{year-1})$
Reported (NIPA)	1.38%	2.60%	2.34%	32.09%
Unfiltered - APWG* (1960-14)	-	2.68%	2.30%	0.77%
Unfiltered - APWG* (1928-14)	-	4.15%	3.12%	-11.31%
Unfiltered - Our Model (Quarterly Data)				
Homoscedastic (1960-14)	1.40%	3.38%	2.13%	-29.93%
Heteroscedastic (1960-14)	1.34%	2.16%	1.88%	-11.53%
Homoscedastic (1947-23)	1.34%	3.75%	2.10%	-32.34%
Heteroscedastic (1947-23)	1.29%	2.30%	1.78%	-21.80%
Unfiltered - Our Model (Annual Data)				
Homoscedastic (1960-14)	1.29%	2.62%	2.23%	-0.66%
Heteroscedastic (1960-14)	1.29%	2.63%	2.23%	-0.66%
Homoscedastic (1930-22)	1.51%	4.10%	2.89%	-7.22%
Heteroscedastic (1930-22)	1.51%	4.02%	2.87%	-6.48%

Note: Moments of reported and unfiltered consumption (our model). We compare these moments with those of unfiltered consumption in Kroencke (2017) as well (APWG stands for "Asset Pricing Without Garbage"). We have simply copied his results here, writing "*" next to variables presented in that paper. Reported and unfiltered consumption are for nondurables only, from NIPA tables. We consider the quasi-differenced model with serially correlated measurement errors of section 2 for quarterly data, while setting $\rho_\xi = 0.06$. For annual data, the model is the same as in Kroencke (2017). See section 2.3 for more details on calibration.

model.

Since official statistical procedures do not differentiate between different asset holders, we calibrate the model based on the consumption growth series for *all* households. Recall that, even though consumption growth is semi-annual, its frequency is monthly. Therefore, to suit its scale for its frequency we convert the series into monthly consumption growth, and then calibrate the model based on the latter. Once unfiltered consumption is obtained, we transform the series back into semi-annual growth terms. This series is then comparable with the original input in the model – reported (CEX) consumption. We repeat this procure imposing the calibration of the model for *all* households to each type of asset holder. This generates unfiltered consumption growth series associated with each group.

Given that we calibrate the model based on monthly consumption growth at the same frequency, we fix $\overline{\sigma}_\eta = 0.0078$ in section 4 – the same value in Bansal and Yaron (2004), for the same scale and frequency. We once more parameterise $\overline{\sigma}_\nu$ such that the long-term (6-years) standard deviation is not greater than 1.2 times that of reported consumption. This rule gives us $\overline{\sigma}_\nu = 2.0\%$. We maintain $\rho_\xi = 0.06$ in section 4[8]. There we also establish $a_1 = 0.20$ and $a_2 = 0.30$[9].

Note that patterns observed in Table 2.2 for CEX data are similar to those of Table 2.1 for NIPA consumption. Unfiltered consumption once more is more volatile and less auto-correlated than official data. In addition, note that this also holds true for most cases shown in Table 2.2 when we split the sample between different types of asset holders. Particularly, unfiltered consumption pictures a strong mean reversion pattern for stock and bond holders.

Our measures of consumption based on CEX data have a window of 10 periods of missing observations. This happens due to a methodological change, at the end of 1985. The BLS replaced the households IDs, so that we can not match households across that change. To construct unfiltered consumption, we need to imput values to those missing observations.

[8]We will test the stability of our model for different values of ρ_ξ in a future version of this paper.
[9]Such that $a_0 = 0.00003$.

We use the final consumption series for each group to estimate the observation for the next period (out of sample), based on a simple AR(1) model. We then apply and calibrate the model, but exclude those observations associated with the window. To isolate the effect of the training period around the imputation, we remove 6 observations before and after the period.

B.3 Epstein-Zin Preferences Framework

This section gives auxiliary algebra and complementary results for section 3.

B.3.1 Euler Equations

Consider L. G. Epstein and Zin (1989) recursive preferences defined by:

$$U_t = \left[(1-\delta)C_t^{\frac{1-\gamma}{\theta}} + \delta(E_t U_{t+1}^{1-\gamma})^{\frac{1}{\theta}}\right]^{\frac{\theta}{1-\gamma}}, \tag{B.41}$$

where $\theta = (1-\gamma)/(1-\psi^{-1})$, δ is the discount factor, γ denotes the relative risk aversion coefficient and $C_{k,t}$ is real consumption of type k (unfiltered or reported) in period t. Denote w as the household's wealth and $1 + R_{w,t+1}$ as the gross real return on wealth. If the representative household combines it with the implicit inter-temporal budget constraint $W_{t+1} = (1+R_{w,t+1})(W_t - C_{k,t})$, it is possible to show that the following Euler Equation holds[10]:

$$1 = E_t\left[\left(\delta\left(\frac{C_{k,t+1}}{C_{k,t}}\right)^{-\frac{1}{\psi}}\right)^{\theta}\left(\frac{1}{1+R_{w,t+1}}\right)^{1-\theta}(1+R_{f,t+1})\right], \tag{B.42}$$

where $1 + R_{f,t+1}$ denotes the gross real returns on risk-free bonds.

Following Campbell (2003) but allowing for time-varying second-order variables as in Yogo (2004) and Campbell, Viceira, Viceira, et al. (2002), if we assume that returns and

[10] See L. Epstein and Zin (1991).

consumption are jointly log-normal, this implies that the riskless real interest rate is[11]:

$$r_{f,t+1} = -log(\delta) + \frac{1}{\psi}E_t[\Delta c_{k,t+1}] + \frac{\theta - 1}{2}Var_t[r_{w,t+1} - E_t r_{w,t+1}] - \frac{\theta}{2\psi^2}Var_t[\Delta c_{k,t+1} - E_t \Delta c_{k,t+1}], \tag{B.43}$$

Equation (B.43) above is used to estimate (2.14) of the main text with the risk-free rate. In addition, as in Yogo (2004), we obtain (2.14) of the main text, for market returns $(i = m)$, by properly defining $r_{i,t+1} - E_t r_{i,t+1} - \frac{1}{\psi}(\Delta c_{k,t+1} - E_t \Delta c_{k,t+1}) \equiv \varrho_{i,t+1}$ under a similar log-linearisation:

$$\begin{aligned} r_{i,t+1} = &-log(\delta) + \frac{1}{\psi}E_t[\Delta c_{k,t+1}] + \frac{\theta - 1}{2}Var_t[r_{w,t+1} - E_t r_{w,t+1}] - \frac{\theta}{2\psi^2}Var_t[\Delta c_{k,t+1} - E_t \Delta c_{k,t+1}] \\ &- \frac{1}{2}Var_t[r_{i,t+1} - E_t r_{i,t+1}] + \frac{\theta}{\psi}Cov_t[r_{i,t+1} - E_t r_{i,t+1}, \Delta c_{k,t+1} - E_t \Delta c_{k,t+1}] \\ &+ (1 - \theta)Cov_t[r_{i,t+1} - E_t r_i, r_{w,t+1} - E_t r_{w,t+1}] \end{aligned} \tag{B.44}$$

Finally, (2.13) of the main text is obtainable from (B.44) by rearranging terms while defining $\Delta c_{k,t+1} - E_t \Delta c_{k,t+1} - \psi(r_{i,t+1} - E_t r_{i,t+1}) \equiv \epsilon_{i,t}$.

B.4 K-Class Estimators and Critical Values

Consider the standard simultaneous equations system[12]:

$$y = Y\beta + X\gamma + u \tag{B.45}$$

$$Y = Z\Pi + X\Phi + V \tag{B.46}$$

[11] Where $r_{f,t} = ln(1 + R_{f,t})$, for instance.

[12] Where y denotes the dependent variable, Y is a matrix constructed from n endogenous variables, X is the matrix of K_1 exogenous regressors and Z has K_2 instruments. All variables have dimension T.

As in Yogo (2004), the three K-class estimators used can be synthesised by[13]:

$$\hat{\beta} = [Y^{\perp\prime}(I - kM_{z^{\perp}}Y^{\perp})Y^{\perp}]^{-1}[Y^{\perp\prime}(I - kM_{z^{\perp}}Y^{\perp})y^{\perp}] \tag{B.47}$$

If $k = 1$, then we have TSLS. If k is the smallest root of $|\overline{Y}'M_x\overline{Y} - k\overline{Y}'M_z\overline{Y}|$, then the equation above is the LIML estimator. Finally, the Fuller-K estimator is obtained when $k = k_{LIML} - 1/(T - K_1 - K_2)$.

For expository reasons, we repeat critical values of Stock and Yogo (2002) for first-stage F-statistics under the following null hypotheses:

1 TSLS bias is a fraction not greater than 10 percent that of the OLS: 10.27

2 Size of the TSLS t-test (5% significance) can not be greater than 10 percent: 24.58

3 Fuller-K bias as a fraction of the OLS bias is not greater than 10 percent: 6.37

4 Size of the LIML t-test (5% significance) can not be greater than 10 percent: 5.44

[13] Where $Y^{\perp} = M_xY$, $\overline{Z} = [X, Z]$, $\overline{X} = [Y, X]$ and $\overline{Y} = [y, Y]$.

B.5 Consumption: Nondurables Only

In this subsection we repeat tables of section 3 but for consumption series constructed from nondurables only (excluding the services component). Our main results are maintained.

Two things stand out in results for the homoscedastic framework (Table B.2, Table B.3 and Table B.4). First, when we estimate $1/\psi$ using (14), first-stage F-statistics are actually twice as high for unfiltered consumption (heteroscedastic model) as for its reported analogue. In contrast, estimates with the former are again similar across estimators, once more suggesting that first-step predictability does not seem especially relevant in generating more sensible estimates[14]. Second, with unfiltered consumption our weak-instrument-robust confidence intervals are mostly in the positive region. Unfortunately, we still have uninformative robust intervals under the AR and LR tests and unfiltered consumption at annual frequency.

There is nothing particularly different in Table B.5. Although we obtain negative estimates for the EIS using quarterly data and unfiltered consumption, there is still improvement relative to reported consumption. Again, barely none of those estimates are statistically different from zero[15]. However, robust intervals from the J-K test are substantially narrower relative to Table 2.6.

[14]Other variables apart, Table B.1 would suggest that unfiltered consumption is a weaker instrument. In absolute terms, its auto-correlation diminishes roughly by a factor of five when we use unfiltered instead of reported consumption, jumping from 32.1% to mere -6.4% (for the complete sample). However, cross-correlations with asset returns are much more definite for unfiltered consumption. Most probably the second effect prevails on net, accounting for the more disciplined estimates obtained.

[15]The only exception is the estimate under SYS-GMM and reported consumption for annual data. However, the value of -0.003 is sufficiently small and not statistically significant at 5%.

Table B.2: Estimates of the EIS Using K-Class Estimators and Quarterly Data

			K-Class Estimator			
Asset	Estimate	Δc_k	TSLS	Fuller-K	LIML	1S-F
Risk Free	ψ	Reported	-0.082^{***}	-0.101^{***}	-0.106^{***}	22.120
			(0.127)	(0.131)	(0.132)	
	ψ	Unf-Hom	0.561	0.585	0.590	22.296
			(0.588)	(0.604)	(0.608)	
	ψ	Unf-Het	0.356^{**}	0.378^{**}	0.382^{**}	22.308
			(0.281)	(0.294)	(0.296)	
Stocks	ψ	Reported	-0.002^{***}	-0.016^{***}	-0.020^{***}	4.530
			(0.028)	(0.034)	(0.036)	
	ψ	Unf-Hom	0.246^{***}	0.263^{***}	0.279^{***}	4.405
			(0.135)	(0.142)	(0.148)	
	ψ	Unf-Het	0.125^{***}	0.158^{***}	0.170^{***}	4.306
			(0.066)	(0.079)	(0.084)	
Risk Free	$\frac{1}{\psi}$	Reported	-0.819^{***}	-3.716	-9.473	1.636
			(0.539)	(3.100)	(11.839)	
	$\frac{1}{\psi}$	Unf-Hom	0.247^{***}	0.831	1.694	1.669
			(0.122)	(0.609)	(1.745)	
	$\frac{1}{\psi}$	Unf-Het	0.435^{***}	1.660	2.620	2.665
			(0.183)	(1.036)	(2.033)	
Stock	$\frac{1}{\psi}$	Reported	-0.311	-11.003	-49.778	1.636
			(3.446)	(13.680)	(90.231)	
	$\frac{1}{\psi}$	Unf-Hom	2.416	2.939	3.584	1.669
			(1.180)	(1.482)	(1.901)	
	$\frac{1}{\psi}$	Unf-Het	3.354	5.061^{*}	5.889^{*}	2.665
			(1.607)	(2.436)	(2.901)	

Notes: Estimates of the EIS and its reciprocal using (13) and (14) and quarterly data. Unfiltered consumption extracted relying on the quasi-differenced Filter model whose measurement errors are serially correlated ($\rho_\xi = 0.06$). All consumption series refer to nondurables only. We apply the same setting of Yogo (2004), using 3 types of K-class estimators and assuming that errors conditionally follow a martingale difference sequence. Reported denotes official consumption data from NIPA tables. Unf-Hom and Unf-Het refer to unfiltered consumption, constructed by the homoscedastic and heteroscedastic models, respectively. Standard errors are presented in parentheses. The null that the estimated coefficient equals 1 has been tested: ***, ** and * denote rejection of the null hypothesis at 1, 5 and 10 percent significance levels.

Table B.3: Estimates of the EIS Using K-Class Estimators and Annual Data

			K-Class Estimator			
Asset	Estimate	Δc_k	TSLS	Fuller-K	LIML	1S-F
Risk Free	ψ	Reported	-0.088^{***}	-0.123^{***}	-0.128^{***}	12.498
			(0.084)	(0.094)	(0.096)	
	ψ	Unf-Hom	0.094^{***}	0.051^{***}	0.043^{***}	10.827
			(0.208)	(0.230)	(0.234)	
	ψ	Unf-Het	0.091^{***}	0.049^{***}	0.042^{***}	10.817
			(0.204)	(0.225)	(0.229)	
Stocks	ψ	Reported	0.030^{***}	0.047^{***}	0.051^{***}	6.100
			(0.024)	(0.030)	(0.032)	
	ψ	Unf-Hom	0.199^{***}	0.228^{***}	0.265^{***}	1.469
			(0.084)	(0.102)	(0.128)	
	ψ	Unf-Het	0.195^{***}	0.222^{***}	0.256^{***}	1.482
			(0.082)	(0.099)	(0.123)	
Risk Free	$\frac{1}{\psi}$	Reported	-1.224^{***}	-5.262^{*}	-7.792	2.947
			(0.775)	(3.399)	(5.811)	
	$\frac{1}{\psi}$	Unf-Hom	0.264^{**}	0.971	23.076	2.160
			(0.307)	(1.195)	(124.367)	
	$\frac{1}{\psi}$	Unf-Het	0.266^{**}	0.973	24.042	2.147
			(0.314)	(1.210)	(132.284)	
Stock	$\frac{1}{\psi}$	Reported	5.308	15.157	19.669	2.947
			(3.448)	(8.954)	(12.325)	
	$\frac{1}{\psi}$	Unf-Hom	3.411	3.587	3.773	2.160
			(1.531)	(1.675)	(1.823)	
	$\frac{1}{\psi}$	Unf-Het	3.521	3.703	3.901	2.147
			(1.571)	(1.715)	(1.868)	

Notes: Estimates of the EIS and its reciprocal using (13) and (14) and annual data. Unfiltered consumption extracted relying on the Filter model whose measurement errors are not persistent. All consumption series refer to nondurables only. We apply the same setting of Yogo (2004), using 3 types of K-class estimators and assuming that errors conditionally follow a martingale difference sequence. When reported consumption is used, asset returns have not been adjusted for time-aggregation. Reported denotes official consumption data from NIPA tables. Unf-Hom and Unf-Het refer to unfiltered consumption, constructed by the homoscedastic and heteroscedastic models, respectively. Standard errors are presented in parentheses. The null that the estimated coefficient equals 1 has been tested: ***, ** and * denote rejection of the null hypothesis at 1, 5 and 10 percent significance levels.

Table B.4: Weak-IV-Robust CIs for the EIS

Quarterly Data			
Asset	Δc_k	Anderson-Rubin	Likelihood Ratio
Risk Free	Reported	$[-0.383, 0.139]$	$[-0.392, 0.145]$
	Unf-Hom	$[-0.611, 1.838]$	$[-0.624, 1.851]$
	Unf-Het	$[0.143, 0.624]$	$[-0.204, 0.995]$
Stocks	Reported	$[-0.124, 0.040]$	$[-0.160, 0.051]$
	Unf-Hom	$[-0.087, 1.135]$	$[-0.004, 0.786]$
	Unf-Het	$[0.015, 0.539]$	$[0.022, 0.503]$
Annual Data			
Risk Free	Reported	$[-0.186, -0.075]$	$[-0.350, 0.045]$
	Unf-Hom	$[-0.236, 0.302]$	$[-0.482, 0.499]$
	Unf-Het	$[-0.238, 0.299]$	$[-0.473, 0.489]$
Stock	Reported	$[0.015, 0.102]$	$[-0.005, 0.152]$
	Unf-Hom	$(-\infty, +\infty)$	$(-\infty, +\infty)$
	Unf-Het	$(-\infty, +\infty)$	$(-\infty, +\infty)$

Notes: Weak-instrument-robust 95% confidence intervals. Sets constructed by inverting statistics of the Anderson-Rubin and Likelihood Ratio tests. Data used for both reported and unfiltered consumption refer to the consumption of nondurables only. For quarterly data, we use our quasi-differenced Filter model ($\rho_\xi = 0.06$) while for annual data we use the canonical version - with no persistence for measurement errors. Reported denotes official consumption data from NIPA tables. Unf-Hom and Unf-Het refer to unfiltered consumption, constructed by the homoscedastic and heteroscedastic models, respectively. The calibration of other parameters described in this section apply in both cases.

Table B.5: Heteroscedasticity-Robust Estimates of the EIS

Δc_k	Two-Step	CUE	SYS	95% CI
Quarterly Data				
Reported	-0.136	-0.151	0.000	$(-\infty, \infty)$
	(0.113)	(0.114)	(0.000)	
Unf-Hom	0.125	0.152	0.015	$[-0.345, 0.841]$
	(0.555)	(0.555)	(0.012)	
Unf-Het	-0.034	-0.046	0.003	$[-0.616, 0.262]$
	(0.268)	(0.268)	(0.003)	
Annual Data				
Reported	-0.079	-0.220	-0.003^{*}	$(-\infty, \infty)$
	(0.091)	(0.103)	(0.001)	
Unf-Hom	0.089	0.094	0.039	$(-\infty, \infty)$
	(0.144)	(0.143)	(0.010)	
Unf-Het	0.085	0.091	0.039	$(-\infty, \infty)$
	(0.142)	(0.141)	(0.027)	

Note: 2S-GMM and CUE-GMM estimates of ψ (EIS) in equation (13) using the risk-free rate. The third column presents estimates of the same coefficient under the joint estimation (15), where market returns are also used, while allowing for different drifts across equations. We present 95% confidence intervals that are robust to both heteroscedasticity and weak-IV settings. These are constructed by inverting the K-statistic of Kleibergen (2005). Consumption series are relative to nondurables only. Reported denotes official consumption data from NIPA tables. Unf-Hom and Unf-Het refer to unfiltered consumption, constructed by the homoscedastic and heteroscedastic models, respectively. Standard errors are presented in parentheses. The null that the estimated coefficient equals 0 has been tested using conventional t-statistics: ***, ** and * denote rejection of the null hypothesis at 1, 5 and 10 percent significance levels.

B.6 Implied EIS from Estimates of Its Reciprocal

This section presents implied EIS estimates from $1/\psi$ (lower part of Table 2.3 and Table 2.4) using equation (14). All consumption data refer to nondurables and services and standard errors have been constructed by Delta method.

Table B.6: Implied EIS from the Estimation of (14) - Nondurables and Services

		K-Class Estimator		
Asset	Δc_k	TSLS	Fuller-K	LIML
Quarterly Data				
Risk Free	Reported	2.285	0.202^{***}	0.053^{***}
		(1.621)	(0.182)	(0.093)
	Unf-Hom	3.021	0.970	0.573
		(1.406)	(0.643)	(0.484)
	Unf-Het	2.866	0.697	0.385^{*}
		(1.376)	(0.484)	(0.352)
Stock	Reported	1.258	-0.241^{***}	-0.199^{***}
		(4.308)	(0.287)	(0.213)
	Unf-Hom	0.550	0.242^{***}	0.168^{***}
		(0.690)	(0.218)	(0.145)
	Unf-Het	0.292^{***}	0.223^{***}	0.191^{***}
		(0.136)	(0.106)	(0.093)
Annual Data				
Risk Free	Reported	0.733	0.218^{***}	0.108^{***}
		(0.389)	(0.160)	(0.113)
	Unf-Hom	2.577	0.531	0.090^{***}
		(2.537)	(0.533)	(0.208)
	Unf-Het	2.539	0.522^{***}	0.089^{***}
		(2.496)	(0.129)	(0.205)
Stock	Reported	-0.147^{***}	-0.085^{***}	-0.065^{***}
		(0.084)	(0.049)	(0.040)
	Unf-Hom	0.247^{***}	0.240^{***}	0.227^{***}
		(0.112)	(0.111)	(0.108)
	Unf-Het	0.241^{***}	0.236^{***}	0.222^{***}
		(0.109)	(0.108)	(0.105)

Note: Implied ψ (EIS) estimates from (14). Consumption series have been constructed taking into account the consumption of nondurables and services. Reported denotes official consumption data from NIPA tables. Unf-Hom and Unf-Het refer to unfiltered consumption, constructed by the homoscedastic and heteroscedastic models, respectively. Standard values are presented in parentheses. The null that the estimated coefficient equals 1 has been tested: ***, ** and * denote rejection of the null hypothesis at 1, 5 and 10 percent significance levels.

Table B.7: Implied EIS from the Estimation of (14) - Nondurables Only

Quarterly Data				
		K-Class Estimator		
Asset	Δc_k	TSLS	Fuller-K	LIML
Risk Free	Reported	-1.221^{***}	-0.269^{***}	-0.106^{***}
		(0.803)	(0.225)	(0.132)
	Unf-Hom	4.049	1.203	0.590
		(2.000)	(0.882)	(0.608)
	Unf-Het	2.298	0.602	0.382^{**}
		(0.965)	(0.376)	(0.296)
Stock	Reported	-3.220	-0.091^{***}	-0.020^{***}
		(35.721)	(0.113)	(0.036)
	Unf-Hom	0.414^{***}	0.340^{***}	0.279^{***}
		(0.202)	(0.172)	(0.148)
	Unf-Het	0.298^{***}	0.198^{***}	0.170^{***}
		(0.143)	(0.095)	(0.084)
Annual Data				
Risk Free	Reported	-0.817^{***}	-0.190^{***}	-0.128^{***}
		(0.518)	(0.123)	(0.096)
	Unf-Hom	3.788	1.030	0.043^{***}
		(4.405)	(1.267)	(0.234)
	Unf-Het	3.766	1.028	0.042^{***}
		(4.455)	(1.278)	(0.229)
Stock	Reported	0.188^{***}	0.066^{***}	0.051^{***}
		(0.122)	(0.039)	(0.032)
	Unf-Hom	0.293^{***}	0.279^{***}	0.265^{***}
		(0.132)	(0.130)	(0.128)
	Unf-Het	0.284^{***}	0.270^{***}	0.256^{***}
		(0.127)	(0.125)	(0.123)

Note: Implied ψ (EIS) estimates from (14). Consumption series have been constructed taking into account the consumption of nondurables only. Reported denotes official consumption data from NIPA tables. Unf-Hom and Unf-Het refer to unfiltered consumption, constructed by the homoscedastic and heteroscedastic models, respectively. Standard values are presented in parentheses. The null that the estimated coefficient equals 1 has been tested: ***, ** and * denote rejection of the null hypothesis at 1, 5 and 10 percent significance levels.

B.7 Results Using Raw Returns for Both Reported and Unfiltered Consumption

In the main text, we adjusted return series for potential time-aggregation bias when conducting estimates with unfiltered consumption – see section 2.2. Recall that returns are never adjusted when reported consumption is used. Here, we present results for the case when we use raw returns with both reported and unfiltered consumption. In general terms, our findings are broadly similar to those of section 3. Note that robust intervals constructed from inverting AR and LR statistics are no longer uninformative with unfiltered consumption and annual data. However, our results for stock returns are somewhat weaker in comparison. In addition, we can not revert uninformative sets in the heteroscedasticity-robust framework (K test).

Table B.8: EIS Using K-Class Estimators and Quarterly Data – Raw Returns

			K-Class Estimator			
Asset	Estimate	Δc_k	TSLS	Fuller-K	LIML	1S-F
Risk Free	ψ	Reported	0.067^{***}	0.053^{***}	0.053^{***}	22.172
			(0.078)	(0.093)	(0.093)	
	ψ	Unf-Hom	0.521	0.559	0.566	22.335
			(0.461)	(0.475)	(0.478)	
	ψ	Unf-Het	0.342^{**}	0.375^{*}	0.380^{*}	22.474
			(0.332)	(0.346)	(0.348)	
Stocks	ψ	Reported	0.006^{***}	-0.101^{***}	-0.199^{***}	4.575
			(0.017)	(0.096)	(0.213)	
	ψ	Unf-Hom	0.215^{***}	0.234^{***}	0.249^{***}	4.462
			(0.114)	(0.121)	(0.127)	
	ψ	Unf-Het	0.165^{***}	0.189^{***}	0.202^{***}	4.310
			(0.084)	(0.093)	(0.099)	
Risk Free	$\frac{1}{\psi}$	Reported	0.438^{*}	4.953	18.979	6.630
			(0.311)	(4.475)	(33.660)	
	$\frac{1}{\psi}$	Unf-Hom	0.334^{***}	1.043	1.765	1.890
			(0.155)	(0.691)	(1.491)	
	$\frac{1}{\psi}$	Unf-Het	0.353^{***}	1.452	2.630	2.268
			(0.170)	(1.009)	(2.408)	
Stock	$\frac{1}{\psi}$	Reported	0.795	-4.150	-5.014	6.630
			(2.724)	(4.936)	(5.346)	
	$\frac{1}{\psi}$	Unf-Hom	2.725	3.362	4.015	1.890
			(1.296)	(1.643)	(2.040)	
	$\frac{1}{\psi}$	Unf-Het	3.237	4.242	4.956	2.268
			(1.512)	(2.019)	(2.423)	

Notes: Estimates of the EIS and its reciprocal using (13) and (14) and quarterly data. Unfiltered consumption extracted relying on the quasi-differenced Filter model whose measurement errors are serially correlated ($\rho_\xi = 0.06$). Return series are *not* adjusted for time-aggregation issues as in section 2.2. All consumption series refer to nondurables and services. We apply the same setting of Yogo (2004), using 3 types of K-class estimators and assuming that errors conditionally follow a martingale difference sequence. Reported denotes official consumption data from NIPA tables. Unf-Hom and Unf-Het refer to unfiltered consumption, constructed by the homoscedastic and heteroscedastic models, respectively. When reported consumption is used, asset returns have not been adjusted for time-aggregation. Standard errors are presented in parentheses. The null that the estimated coefficient equals 1 has been tested: ***, ** and * denote rejection of the null hypothesis at 1, 5 and 10 percent significance levels.

Table B.9: EIS Using K-Class Estimators and Annual Data – Raw Returns

Asset	Estimate	Δc_k	K-Class Estimator TSLS	Fuller-K	LIML	1S-F
Risk Free	ψ	Reported	0.112^{***}	0.109^{***}	0.108^{***}	11.836
			(0.105)	(0.111)	(0.113)	
	ψ	Unf-Hom	0.271^{***}	0.289^{**}	0.293^{**}	11.699
			(0.269)	(0.284)	(0.288)	
	ψ	Unf-Het	0.121^{***}	0.137^{***}	0.140^{***}	11.707
			(0.172)	(0.183)	(0.186)	
Stocks	ψ	Reported	-0.049^{***}	-0.061^{***}	-0.065^{***}	5.057
			(0.034)	(0.038)	(0.040)	
	ψ	Unf-Hom	-0.037^{***}	-0.024^{***}	-0.022^{***}	11.432
			(0.058)	(0.063)	(0.064)	
	ψ	Unf-Het	0.019^{***}	0.034^{***}	0.038^{***}	9.026
			(0.042)	(0.046)	(0.047)	
Risk Free	$\frac{1}{\psi}$	Reported	1.364	4.592	9.246	1.893
			(0.724)	(3.364)	(9.634)	
	$\frac{1}{\psi}$	Unf-Hom	0.574	1.788	3.410	1.689
			(0.328)	(1.303)	(3.345)	
	$\frac{1}{\psi}$	Unf-Het	0.613	2.688	7.127	1.764
			(0.462)	(2.297)	(9.426)	
Stock	$\frac{1}{\psi}$	Reported	-6.808^{**}	-11.773^{*}	-15.285^{*}	1.893
			(3.885)	(6.818)	(9.365)	
	$\frac{1}{\psi}$	Unf-Hom	-1.587^{**}	-5.532	-46.488	1.689
			(1.187)	(5.335)	(137.715)	
	$\frac{1}{\psi}$	Unf-Het	1.715	10.520	26.585	1.764
			(2.085)	(9.642)	(33.559)	

Notes: Estimates of the EIS and its reciprocal using (13) and (14) and annual data. Unfiltered consumption extracted relying on the Filter model whose measurement errors are not persistent. Return series are *not* adjusted for time-aggregation issues as in section 2.2. All consumption series refer to nondurables and services. We apply the same setting of Yogo (2004), using 3 types of K-class estimators and assuming that errors conditionally follow a martingale difference sequence. When reported consumption is used, asset returns have not been adjusted for time-aggregation. Reported denotes official consumption data from NIPA tables. Unf-Hom and Unf-Het refer to unfiltered consumption, constructed by the homoscedastic and heteroscedastic models, respectively. Standard errors are presented in parentheses. The null that the estimated coefficient equals 1 has been tested: ***, ** and * denote rejection of the null hypothesis at 1, 5 and 10 percent significance levels.

Table B.10: Weak-IV-Robust CIs for the EIS – Raw Returns

Quarterly Data			
Asset	Δc_k	Anderson-Rubin	Likelihood Ratio
Risk Free	Reported	$\emptyset$	$[-0.136, 0.235]$
	Unf-Hom	$[-0.315, 1.507]$	$[-0.372, 1.573]$
	Unf-Het	$[-0.076, 0.857]$	$[-0.303, 1.109]$
Stocks	Reported	$\emptyset$	$(-\infty, +\infty)$
	Unf-Hom	$[-0.048, 0.975]$	$[0.014, 0.688]$
	Unf-Het	$[-0.009, 0.751]$	$[0.024, 0.568]$
Annual Data			
Risk Free	Reported	$[-0.104, 0.316]$	$[-0.131, 0.341]$
	Unf-Hom	$[-0.304, 0.934]$	$[-0.292, 0.920]$
	Unf-Het	$[-0.203, 0.510]$	$[-0.233, 0.545]$
Stock	Reported	$[-0.245, 0.019]$	$[-0.199, 0.007]$
	Unf-Hom	$[-0.131, 0.107]$	$[-0.143, 0.125]$
	Unf-Het	$[-0.044, 0.148]$	$[-0.049, 0.158]$

Note: Weak-instrument-robust 95% confidence intervals. Sets constructed by inverting statistics of the Anderson-Rubin and Likelihood Ratio tests. Return series are *not* adjusted for time-aggregation issues as in section 2.2. Data used both for reported and unfiltered consumption refer to the consumption of nondurables and services. For quarterly data, we use our quasi-differenced Filter model ($\rho_\xi = 0.06$) while for annual data we use the canonical version – with no persistence for measurement errors. Reported denotes official consumption data from NIPA tables. Unf-Hom and Unf-Het refer to unfiltered consumption, constructed by the homoscedastic and heteroscedastic models, respectively.

Table B.11: Heteroscedasticity-Robust Estimates of the EIS – Raw Returns

Quarterly Data				
Δc_k	Two-Step	CUE	SYS	95% CI
Reported	0.133	0.189^{**}	0.001	$(-\infty, +\infty)$
	(0.082)	(0.085)	(0.000)	
Unf-Hom	0.594	0.670	0.008	$(-\infty, +\infty)$
	(0.517)	(0.519)	(0.006)	
Unf-Het	0.443	0.506	0.007	$(-\infty, +\infty)$
	(0.370)	(0.372)	(0.001)	
Annual Data				
Reported	0.056	0.022	-0.015^{*}	$(-\infty, +\infty)$
	(0.088)	(0.087)	(0.008)	
Unf-Hom	0.067	0.047	-0.001	$(-\infty, +\infty)$
	(0.261)	(0.261)	(0.000)	
Unf-Het	0.033	0.025	-0.002	$(-\infty, +\infty)$
	(0.172)	(0.172)	(0.001)	

Note: 2S- and CUE-GMM estimates of ψ (EIS) in equation (13) using the risk-free. The third column presents results under the joint estimation (15), where market returns are also used (allowing for different drifts across equations). We present 95% confidence intervals that are robust to both heteroscedasticity and a weak-IV setting. These are constructed by inverting the K-statistic of Kleibergen (2005). Return series are *not* adjusted for time-aggregation issues as in section 2.2. Consumption series are relative to nondurables and services. Reported denotes official consumption data from NIPA tables. Unf-Hom and Unf-Het refer to unfiltered consumption, constructed by the homoscedastic and heteroscedastic models, respectively. Standard errors are presented in parentheses. The null that the estimated coefficient equals 0 has been tested using conventional t-statistics: ***, ** and * denote rejection of the null hypothesis at 1, 5 and 10 percent significance levels.

B.8 Results with Restricted Sample

In this section, we repeat the estimations of section 3 in the main paper while restricting our sample. The motivation here is to remove first observations for which the Kalman filter is still learning.

Our estimations here use the period encompassing 1960:1 to 2023:2 for quarterly and 1940 to 2022 for annual data. Table B.12, Table B.13, Table B.14 and Table B.15 below broadly confirm our findings in the main paper.

B.9 Results for Annual Data when Measurement Errors are Persistent

In the main paper, we showed results for annual data considering the original model in Kroencke (2017), which does not feature serially correlated measurement errors ($\rho_\xi = 0$). For completeness, in this section we exhibit results of section 3 for annual data while $\rho_\xi = 0.06 \neq 0$. It is worth reemphasising that we have found little sensitiveness of unfiltered consumption to different values of ρ_ξ once other parameters have been properly calibrated according to benchmark moments. Hence, we repeat $\rho_\xi = 0.06$ (as we did with quarterly data), but other parameters have been changed: $\overline{\sigma}_\eta = 0.0078 \times \sqrt{12} \approx 2.7\%$, $a_1 = 0.05$, $a_2 = 0.85$ and $\sigma_\nu = 3.3\%$ (heteroscedastic model) or $\sigma_\nu = 2.1\%$ (homoscedastic model)[16]. These parametric conditions ensure that moments of annual unfiltered consumption are not much different from those presented in Table 2.1. All consumption data refer to nondurables and services. Table B.16, Table B.17 and Table B.18 below exhibit the results. These findings are weaker than the ones of the main paper, though. This is probably explained by the fact that persistence in the measurement error is a much more important feature of the model

[16] We lower σ_ν for the homoscedastic model simply to ensure that the long-term standard deviation of unfiltered consumption is not greater than 1.2 times that of reported consumption.

Table B.12: EIS Using K-Class Estimators and Quarterly Data – 1960:2017

Asset	Estimate	Δc_k	K-Class Estimator TSLS	Fuller-K	LIML	1S-F
Risk Free	ψ	Reported	0.120^{***}	0.083^{***}	0.081^{***}	22.810
			(0.075)	(0.108)	(0.110)	
	ψ	Unf-Hom	0.300^{*}	0.321^{*}	0.326^{*}	23.558
			(0.392)	(0.404)	(0.406)	
	ψ	Unf-Het	0.263^{**}	0.281^{**}	0.285^{**}	23.652
			(0.284)	(0.293)	(0.294)	
Stocks	ψ	Reported	0.017^{***}	-0.011^{***}	-3.946	3.821
			(0.016)	(0.069)	(46.662)	
	ψ	Unf-Hom	0.135^{***}	0.152^{***}	0.163^{***}	4.133
			(0.082)	(0.090)	(0.094)	
	ψ	Unf-Het	0.117^{***}	0.126^{***}	0.134^{***}	4.025
			(0.063)	(0.066)	(0.069)	
Risk Free	$\frac{1}{\psi}$	Reported	0.498^{**}	4.978	12.292	12.118
			(0.246)	(4.168)	(16.545)	
	$\frac{1}{\psi}$	Unf-Hom	0.299^{***}	1.216	3.071	1.658
			(0.169)	(0.981)	(3.830)	
	$\frac{1}{\psi}$	Unf-Het	0.472^{**}	1.720	3.514	1.768
			(0.240)	(1.271)	(3.632)	
Stock	$\frac{1}{\psi}$	Reported	1.467	-0.093	-0.253	12.118
			(2.411)	(2.944)	(2.997)	
	$\frac{1}{\psi}$	Unf-Hom	3.475	4.857	6.144	1.658
			(1.819)	(2.652)	(3.560)	
	$\frac{1}{\psi}$	Unf-Het	5.272^{*}	6.210^{*}	7.440^{*}	1.769
			(2.569)	(3.090)	(3.843)	

Notes: Estimates of the EIS and its reciprocal using (13) and (14) and quarterly data. We restrict our sample to the period from 1960:1 to 2017:4. Unfiltered consumption extracted relying on the quasi-differenced Filter model whose measurement errors are serially correlated ($\rho_\xi = 0.06$). All consumption series refer to nondurables and services. We apply the same setting of Yogo (2004), using 3 types of K-class estimators and assuming that errors conditionally follow a martingale difference sequence. Reported denotes official consumption data from NIPA tables. Unf-Hom and Unf-Het refer to unfiltered consumption, constructed by the homoscedastic and heteroscedastic models, respectively. When reported consumption is used, asset returns have not been adjusted for time-aggregation. Standard errors are presented in parentheses. The null that the estimated coefficient equals 1 has been tested: ***, ** and * denote rejection of the null hypothesis at 1, 5 and 10 percent significance levels.

Table B.13: EIS Using K-Class Estimators and Annual Data – 1940:2017

Asset	Estimate	Δc_k	TSLS	Fuller-K	LIML	1S-F
			K-Class Estimator			
Risk Free	ψ	Reported	0.012^{***}	0.014^{***}	0.015^{***}	16.600
			(0.094)	(0.098)	(0.099)	
	ψ	Unf-Hom	0.151^{***}	0.149^{***}	0.146^{***}	11.846
			(0.146)	(0.147)	(0.149)	
	ψ	Unf-Het	0.151^{***}	0.150^{***}	0.147^{***}	11.828
			(0.146)	(0.147)	(0.149)	
Stocks	ψ	Reported	-0.060^{***}	-0.077^{***}	-0.088^{***}	2.691
			(0.044)	(0.051)	(0.057)	
	ψ	Unf-Hom	0.051^{***}	0.039^{***}	0.025^{***}	1.500
			(0.056)	(0.066)	(0.077)	
	ψ	Unf-Het	0.051^{***}	0.040^{***}	0.027^{***}	1.506
			(0.056)	(0.065)	(0.076)	
Risk Free	$\frac{1}{\psi}$	Reported	0.208	1.502	68.677	1.694
			(0.713)	(2.376)	(468.18)	
	$\frac{1}{\psi}$	Unf-Hom	2.636	3.370	6.848	0.620
			(1.632)	(2.353)	(6.962)	
	$\frac{1}{\psi}$	Unf-Het	2.622	3.370	6.826	0.626
			(1.620)	(2.352)	(6.921)	
Stock	$\frac{1}{\psi}$	Reported	-5.996^{*}	-8.842^{*}	-11.380	1.694
			(4.000)	(5.681)	(7.443)	
	$\frac{1}{\psi}$	Unf-Hom	4.800	6.468	39.723	0.620
			(3.469)	(5.811)	(121.877)	
	$\frac{1}{\psi}$	Unf-Het	4.867	6.616	37.384	0.626
			(3.469)	(5.855)	(107.139)	

Notes: Estimates of the EIS and its reciprocal using (13) and (14) and annual data. We restrict our sample to the period from 1940 to 2017. Unfiltered consumption extracted relying on the Filter model whose measurement errors are not persistent. All consumption series refer to nondurables and services. We apply the same setting of Yogo (2004), using 3 types of K-class estimators and assuming that errors conditionally follow a martingale difference sequence. When reported consumption is used, asset returns have not been adjusted for time-aggregation. Reported denotes official consumption data from NIPA tables. Unf-Hom and Unf-Het refer to unfiltered consumption, constructed by the homoscedastic and heteroscedastic models, respectively. Standard errors are presented in parentheses. The null that the estimated coefficient equals 1 has been tested: ***, ** and * denote rejection of the null hypothesis at 1, 5 and 10 percent significance levels.

Table B.14: Weak-IV-Robust CIs for the EIS – Restricted Sample

Asset	Δc_k	Anderson-Rubin	Likelihood Ratio
Quarterly Data: 1960:1 – 2017:4			
Risk Free	Reported	$\emptyset$	$[-0.136, 0.283]$
	Unf-Hom	$[-0.446, 1.134]$	$[-0.478, 1.168]$
	Unf-Het	$[-0.262, 0.860]$	$[-0.295, 0.896]$
Stocks	Reported	$\emptyset$	$(-\infty, +\infty)$
	Unf-Hom	$[-0.062, 0.691]$	$[-0.021, 0.509]$
	Unf-Het	$[-0.033, 0.609]$	$[0.005, 0.393]$
Annual Data: 1940 – 2017			
Risk Free	Reported	$[-0.164, 0.196]$	$[-0.185, 0.218]$
	Unf-Hom	$[-0.367, 0.599]$	$[-0.186, 0.452]$
	Unf-Het	$[-0.365, 0.600]$	$[-0.185, 0.453]$
Stock	Reported	$[-2.074, 0.034]$	$[-0.583, 0.015]$
	Unf-Hom	$(-\infty, +\infty)$	$(-\infty, +\infty)$
	Unf-Het	$(-\infty, +\infty)$	$(-\infty, +\infty)$

Note: Weak-instrument-robust 95% confidence intervals inverting statistics of the Anderson-Rubin and Likelihood Ratio tests. We restrict our estimations to the sample 1960:1-2017:4 (1940-2017) for quarterly (annual) data. Data used both for reported and unfiltered consumption refer to the consumption of nondurables and services. For quarterly data, we use our quasi-differenced Filter model ($\rho_\xi = 0.06$) while for annual data we use the canonical version – with no persistence for measurement errors. Reported denotes official consumption data from NIPA tables. Unf-Hom and Unf-Het refer to unfiltered consumption, constructed by the homoscedastic and heteroscedastic models, respectively.

Table B.15: Heteroscedasticity-Robust Estimates of the EIS – Restricted Sample

Δc_k	Two-Step	CUE	SYS	95% CI
Quarterly Data: 1960:1 – 2017:4				
Reported	0.123	0.347	0.008	$(-\infty, +\infty)$
	(0.077)	(0.093)	(0.002)	
Unf-Hom	0.349	0.363	0.002	$[0.022, 2.782]$
	(0.333)	(0.333)	(0.001)	
Unf-Het	0.305	0.312	0.004	$[-0.173, 2.066]$
	(0.240)	(0.240)	(0.001)	
Annual Data: 1940 – 2017				
Reported	0.024	-0.029	-0.008	$(-\infty, +\infty)$
	(0.087)	(0.087)	(0.005)	
Unf-Hom	0.152	0.175	0.015	$(-\infty, +\infty)$
	(0.141)	(0.141)	(0.009)	
Unf-Het	0.152	0.173	0.015	$(-\infty, +\infty)$
	(0.140)	(0.140)	(0.009)	

Note: 2S-GMM and CUE-GMM estimates of ψ (EIS) in equation (13) using the risk-free rate. The third column presents estimates of the same coefficient under the joint estimation (15), where market returns are also used, while allowing for different drifts across equations. 95% confidence intervals that are robust to both heteroscedasticity and a weak-IV setting are also shown in the last column. These are constructed by inverting the K-statistic of Kleibergen (2005). We restrict our estimations to the sample 1960:1-2017:4 (1940-2017) for quarterly (annual) data. Consumption series are relative to nondurables and services. Reported denotes official consumption data from NIPA tables. Unf-Hom and Unf-Het refer to unfiltered consumption, constructed by the homoscedastic and heteroscedastic models, respectively. Standard errors are presented in parentheses. The null that the estimated coefficient equals 0 has been tested using conventional t-statistics: ***, ** and * denote rejection of the null hypothesis at 1, 5 and 10 percent significance levels.

when it is applied to data at higher frequencies.

Table B.16: Estimates of the EIS – Persistent M.E. and Annual Data

			K-Class Estimator			
Asset	Estimate	Δc_k	TSLS	Fuller-K	LIML	1S-F
Risk Free	ψ	Reported	0.112***	0.109***	0.108***	11.836
			(0.105)	(0.111)	(0.113)	
	ψ	Unf-Hom	0.097***	0.089***	0.088**	10.727
			(0.188)	(0.203)	(0.205)	
	ψ	Unf-Het	-0.030***	0.012***	0.020***	11.096
			(0.197)	(0.213)	(0.216)	
Stocks	ψ	Reported	-0.049***	-0.061***	-0.065***	5.057
			(0.034)	(0.038)	(0.040)	
	ψ	Unf-Hom	0.184***	0.194***	0.222***	1.312
			(0.080)	(0.086)	(0.105)	
	ψ	Unf-Het	0.059***	0.124***	0.168***	1.979
			(0.078)	(0.115)	(0.145)	
Risk Free	$\frac{1}{\psi}$	Reported	1.364	4.592	9.246	1.893
			(0.724)	(3.364)	(9.634)	
	$\frac{1}{\psi}$	Unf-Hom	0.394	1.910	11.293	1.822
			(0.387)	(1.916)	(26.190)	
	$\frac{1}{\psi}$	Unf-Het	-0.115***	0.275	49.474	1.739
			(0.358)	(1.043)	(528.708)	
Stock	$\frac{1}{\psi}$	Reported	-6.808**	-11.773*	-15.285*	1.893
			(3.885)	(6.818)	(9.365)	
	$\frac{1}{\psi}$	Unf-Hom	4.143*	4.241*	4.502	1.822
			(1.863)	(1.937)	(2.131)	
	$\frac{1}{\psi}$	Unf-Het	1.818	4.137	5.956	1.739
			(2.282)	(3.734)	(5.130)	

Note: Estimates of the EIS and its reciprocal using (13) and (14) and annual data. We use 3 types of K-class estimators while assuming that errors conditionally follow a martingale difference sequence. When reported consumption is used, we have not adjusted the timing of returns. Reported denotes official consumption data from NIPA tables. Unf-Hom and Unf-Het refer to unfiltered consumption, constructed by the homoscedastic and heteroscedastic models, respectively. The quasi-differenced model with $\rho_\xi = 0.06$ has been used, while adjusting other parameters to the dynamics and benchmark moments of annual data. All consumption measures refer to nondurables and services. Standard errors are presented in parentheses. The null that the estimated coefficient equals 1 has been tested: ***, ** and * denote rejection of the null hypothesis at 1, 5 and 10 percent significance levels.

Table B.17: Weak-IV-Robust CIs for the EIS – Persistent M.E. and Annual Data

Asset	Δc_k	Anderson-Rubin	Likelihood Ratio
Risk Free	Reported	$[-0.104, 0.316]$	$[-0.131, 0.341]$
	Unf-Hom	$[-0.269, 0.438]$	$[-0.352, 0.516]$
	Unf-Het	$[-0.334, 0.430]$	$[-0.395, 0.515]$
Stock	Reported	$[-0.245, 0.018]$	$[-0.199, 0.007]$
	Unf-Hom	$(-\infty, +\infty)$	$(-\infty, +\infty)$
	Unf-Het	$(-\infty, +\infty)$	$(-\infty, +\infty)$

Note: Weak-instrument-robust 95% confidence intervals for annual data. Sets constructed by inverting statistics of the Anderson-Rubin and Likelihood Ratio tests. Data used both for reported and unfiltered consumption refer to consumption of nondurables and services. We set ($\rho_\xi = 0.06$) while adjusting other parameters to align benchmark moments. Reported denotes official consumption data from NIPA tables. Unf-Hom and Unf-Het refer to unfiltered consumption, constructed by the homoscedastic and heteroscedastic models, respectively.

Table B.18: Het-Robust Estimates of the EIS – Persistent M.E. and Annual Data

Δc_k	Two-Step	CUE	SYS	95% CI
Reported	0.056	0.022	-0.015^{*}	$(-\infty, \infty)$
	(0.088)	(0.087)	(0.008)	
Unf-Hom	0.119	0.133	0.066	$(-\infty, \infty)$
	(0.041)	(0.141)	(0.045)	
Unf-Het	-0.099	-0.399	0.000	$(-\infty, \infty)$
	(0.201)	(0.204)	(0.001)	

Note: 2S-GMM and CUE-GMM estimates of ψ (EIS) in equation (13) using the risk-free rate. The third column presents estimates of the same coefficient under the joint estimation (15), where market returns are also used, while allowing for different drifts across equations. We present 95% confidence intervals that are robust to both heteroscedasticity and a weak-IV setting. These are constructed by inverting the K-statistic of Kleibergen (2005). Consumption series are relative to nondurables and services. We set $\rho_\xi = 0.06$ while adjusting other parameters to align benchmark moments. Reported denotes official consumption data from NIPA tables. Unf-Hom and Unf-Het refer to unfiltered consumption, constructed by the homoscedastic and heteroscedastic models, respectively. Standard errors are presented in parentheses. The null that the estimated coefficient equals 0 has been tested using conventional t-statistics: ***, ** and * denote rejection of the null hypothesis at 1, 5 and 10 percent significance levels.

Appendix C

Appendix for Chapter 3

C.1 The Heterogeneous Economy

C.1.1 Fully Specified Model

The economy is a generalisation of the simple New-Keynesian models described in Woodford (2003) and Galí (2015), but with price-stickiness heterogeneity across firms of different sectors. When $\gamma_k = \gamma = 0$ in (7), this framework produces a purely forward-looking model which is similar to those in Aoki (2001), Benigno (2004), Carvalho (2006), Eusepi, Hobijn, and Tambalotti (2011) and Carvalho and Nechio (2016).

The representative household, which provides firm-specific labour, solves:

$$\max E_0 \sum_{t=0}^{\infty} \beta^t \left(\frac{C_t^{1-\sigma} - 1}{1-\sigma} - \sum_{k=1}^{K} \eta_k \int_0^1 \frac{L_{kj,t}^{1+\varphi^{-1}}}{1+\varphi^{-1}} dj \right),$$

subject to:

$$P_t C_t = \sum_{k=1}^{K} \eta_k \int_0^1 L_{kj,t} W_{kj,t} dj + T_t + I_{t-1} B_{t-1} - B_t,$$

where C_t denotes consumption of the composite good, $L_{kj,t}$ labour in firm kj, $W_{kj,t}$ nominal wages related to the latter, P_t is the aggregate price index, T_t are firms' profits distributed by lump sum transfers and B_t denotes bond holdings that collect a gross interest I_t each period. We assume a cashless economy with one-period maturity for those bonds, which are in zero net supply.

Demand for variety j, produced in sector k, takes the form:

$$Y_{kj,t} = \left(\frac{P_{kj,t}}{P_{k,t}}\right)^{-\epsilon} Y_{k,t}, \qquad\qquad Y_{kj,t} = C_{kj,t} = N_{kj,t},$$

where $C_{kj,t}$ and $N_{kj,t}$ denote the consumption of that variety and the specific labour input, respectively. In line with the main text, $Y_{kj,t}$ and $Y_{k,t}$ represent respectively the aggregate output and the output of sector k. A welcome advantage of a linear technology function is that deep parameters related to the curvature in production are absent in the NKPC, what implies we have one less parameter to calibrate in the estimations.

Total demand in sector k and in the economy follow:

$$Y_{k,t} = \eta_k \left[\int_0^1 Y_{kj,t}^{\frac{\epsilon-1}{\epsilon}} dj\right]^{\frac{\epsilon}{\epsilon-1}}, \qquad\qquad Y_t = \left[\sum_{k=1}^{K} \eta_k^{\frac{1}{\epsilon}} Y_{k,t}^{\frac{\epsilon-1}{\epsilon}}\right]^{\frac{\epsilon}{\epsilon-1}}.$$

Usual market clearing conditions in the goods markets imply $Y_t = C_t$, $Y_{k,t} = C_{k,t}$ and $Y_{kj,t} = C_{kj,t}$.

When firm kj can re-optimise, it sets price $X_{kj,t}$ by maximising the following expression for discounted expected future profits:

$$E_t \sum_{s=0}^{\infty} Q_{t,t+s}(1-\lambda_k)^s [X_{kj,t} Y_{kj,t+s} - W_{kj,t+s} N_{kj,t+s}],$$

subject to:

$$Y_{kj,t+s} = N_{kj,t+s}, \qquad Y_{kj,t+s} = \left(\frac{X_{kj,t+s}}{P_{t+s}}\right)^{-\epsilon} Y_{t+s},$$

where, for instance, $W_{kj,t+s}$ is the nominal wage in the firm kj for the period $t+s$, conditional on time-t information, while $Q_{t,t+s} = \beta(C_{t+s}/C_t)^{\sigma}(P_t/P_{t+s})$ is the stochastic nominal discount factor between the periods t and $t+s$. Additionally:

$$P_{k,t} = \left[\int_0^1 P_{kj,t}^{1-\epsilon} dj\right]^{\frac{1}{1-\epsilon}}, \qquad P_t = \left[\sum_{k=1}^{K} \eta_k P_{k,t}^{1-\epsilon}\right]^{\frac{1}{1-\epsilon}},$$

representing sectoral and aggregate price indices, respectively.

Defining $\Pi_{k,\tau+s,\tau} \equiv P_{k,\tau+s}/P_{k,\tau}$, the aforementioned price setting decisions, the price-indexation scheme (7) and first order conditions yield:

$$X_{kj,t} = \frac{\epsilon}{\epsilon - 1} \frac{E_t \sum_{s=0}^{\infty} Q_{t,t+s}(1-\lambda_k)^s P_{t+s}^{\epsilon} Y_{k,t+s} W_{kj,t+s}}{E_t \sum_{s=0}^{\infty} Q_{t,t+s}(1-\lambda_k)^s P_{t+s}^{\epsilon} Y_{k,t+s} \Pi_{k,t+s-1,t-1}^{\gamma}},$$

Finally, the following law of movement for sectoral prices holds:

$$P_{k,t} = [\lambda_k X_{k,t}^{1-\epsilon} + (1-\lambda_k)(P_{k,t-1}\Pi_{k,t-1,t-2}^{\gamma})^{1-\epsilon}]^{\frac{1}{1-\epsilon}}.$$

The aggregate NKPC (8) and its sectoral counterparts (9) can be derived from this model using similar developments to those described in the appendix of Carvalho (2006), for example.

One can log-linearise those equations in terms of deviations from steady-state values. The law of movement becomes:

$$p_{k,t} = \lambda_k(x_{k,t} - p_{k,t-1}) + (1-\lambda_k)p_{k,t-1} + \gamma(1-\lambda_k)\pi_{k,t-1}.$$

Additionally, it follows that:

$$x_{k,t} = (1 - \beta(1 - \lambda_k))E_t \sum_{s=0}^{\infty} \beta^s (1 - \lambda_k)^s [p_{t+s} + \Theta y_{t+s} - \gamma \pi_{k,t+s-1,t-1}],$$

where Θ is defined in the main text. Remaining parts of the model are represented by the following expressions:

$$y_t = E_t y_{t+1} - \sigma^{-1}(i_t - E_t \pi_{t+1})$$
$$y_t = \sum_{k=1}^{K} \eta_k y_{k,t}$$
$$y_{k,t} = y_t - \epsilon(p_{k,t} - p_t)$$
$$y_{kj,t} = y_{k,t} - \epsilon(p_{kj,t} - p_{k,t})$$
$$y_{k,t} = \int_0^1 y_{kj,t} dj$$
$$p_t = \sum_{k=1}^{K} \eta_k p_{k,t}$$

$$p_{k,t} = \int_0^1 p_{kj,t} dj$$
$$w_{kj,t} - p_t = \varphi^{-1} l_{kj,t} + \sigma c_t$$
$$b_t = 0$$
$$y_{kj,t} = c_{kj,t} = n_{kj,t} = l_{kj,t}$$

C.1.2 Optimal Policy: Derivation

Following Rotemberg and Michael Woodford (1997), a second-order approximation of the ex-ante expected utility of such economy yields:

$$W \approx u_c Y \left[y_t + \frac{1-\sigma}{2} y_t^2 - \sum_{k=1}^{K} \eta_k \int_0^1 \left(n_{kj,t} + \frac{1+\varphi^{-1}}{2} n_{kj,t}^2 \right) dj \right],$$

where we used the market clearing condition $y_t = c_t$, u_c denotes the first-order derivative of the expected utility with respect to consumption and Y represents the output level of steady state.

Next, use the market clearing condition for labour and that, up to a second-order approximation, $y_{k,t} \approx E_j(y_{kj,t}) + \frac{1}{2}\left(\frac{\epsilon-1}{\epsilon}\right) Var_j(y_{kj,t})$, to find:

$$\begin{aligned} W &\approx u_c Y \left[y_t + \frac{1-\sigma}{2} y_t^2 - \sum_{k=1}^{K} \eta_k y_{k,t} - \frac{1+\varphi^{-1}}{2} \sum_{k=1}^{K} \eta_k \int_0^1 \left(\frac{1-\epsilon}{\epsilon} Var_j(y_{kj,t}) + y_{kj,t}^2 \right) dj \right] \\ &\approx u_c Y \left[\frac{1-\sigma}{2} y_t^2 - \frac{1+\varphi^{-1}}{2} \sum_{k=1}^{K} \eta_k \int_0^1 \left(\frac{1-\epsilon}{\epsilon} Var_j(y_{kj,t}) + y_{kj,t}^2 \right) dj \right], \end{aligned}$$

where $E_j(y_{kj,t}) = \int_0^1 y_{kj,t} dj$ and $Var_j(y_{kj,t}) = E_j[(y_{kj,t} - y_{k,t})^2]$, the dispersion of output in sector k. We also used $y_t = \sum_{k=1}^{K} \eta_k y_{k,t}$ in the second step above.

Now, use $E_j(y_{kj,t}^2) \approx y_{k,t}^2 + Var_j(y_{kj,t})$ to find:

$$\begin{aligned} W &\approx u_c Y \left[\frac{1-\sigma}{2} y_t^2 + \frac{1}{2} \frac{\epsilon-1}{\epsilon} \sum_{k=1}^{K} \eta_k Var_j(y_{kj,t}) - \frac{1+\varphi^{-1}}{2} \sum_{k=1}^{K} \eta_k \left(y_{k,t}^2 + Var_j(y_{kj,t}) \right) \right] \\ &\approx u_c Y \left[\frac{1-\sigma}{2} y_t^2 - \frac{1}{2}(\epsilon + \varphi^{-1}) \sum_{k=1}^{K} \eta_k Var_j(y_{kj,t}) - \frac{1}{2}(1+\varphi^{-1}) \left(\sum_{k=1}^{K} \eta_k y_{k,t}^2 \right) \right]. \end{aligned}$$

With the demand condition faced by each firm, the output dispersion within the sectors can be written as variance of the prices of the varieties in each sector: $Var_j(y_{kj,t}) =$

$\epsilon^2 Var_j(p_{kj,t})$. Using this:

$$W \approx u_c Y \left[\frac{1-\sigma}{2} y_t^2 - \frac{1}{2}\epsilon(1+\epsilon\varphi^{-1}) \sum_{k=1}^{K} \eta_k Var_j(p_{kj,t}) - \frac{1}{2}(1+\varphi^{-1}) \left(\sum_{k=1}^{K} \eta_k y_{k,t}^2 \right) \right]$$

Additionally, it is possible to show that the measure of sectoral price dispersion, summed over time, can be approximated up to a second order by – e.g., see Woodford (2003, p. 706):

$$\begin{aligned} \sum_{t=0}^{\infty} \beta^t Var_j(p_{kj,t}) &\approx \left[\frac{\lambda_k}{(1-\lambda_k)} - \beta\lambda_k \right]^{-1} \sum_{t=0}^{\infty} \beta_t \pi_{k,t}^2 \\ &\approx \psi_k(\lambda_k, \beta)^{-1} \sum_{t=0}^{\infty} \beta^t \pi_{k,t}^2. \end{aligned}$$

With this and a second-order approximation of the sectoral demand function, it is possible to finally derive (3.2). Simply insert these approximations in the expression for W above. Note that cross terms involving the product of the output gap and sectoral relative prices – e.g., see Eusepi, Hobijn, and Tambalotti (2011, eq. 5) – will cancel out. This happens because, unlike their model, ours does not feature heterogeneity in the production function across sectors.

C.2 Additional Details on the Dataset

Table C.1 below summarises our data set, as well as applied transformations. Sectoral information is presented in the main paper – see Table 3.2.

Table C.1: Dataset for the Heterogeneous Economy

Variable	Source	Aggregate	Sectoral	Literature	HP Filter
Output	BEA	yes	no	yes	yes
Consumption	BEA	no	yes	no	yes
Non-Farm Labour Share	BLS	yes	no	yes	yes
PCE Inflation	BEA	yes	yes	yes	no
PPI Commodities Inflation	BLS	yes	no	yes	no
Effective Fed Funds	FED	yes	no	yes	no
5-Year Treasury Spread	FED	yes	no	yes	no
Avg. Hourly Earnings Inflation	BLS	yes	no	yes	no

Notes: Types of aggregate and sectoral data compiled in our data set. Output and consumption are measured in per capita terms to account for a model with no population growth. The PPI, the 5-year spread, the Fed Funds rate, the labour share as well as the wage inflation are taken from the Federal Reserve Bank of St. Louis' FRED economic database. The second column provides the sources of the variables: BEA stands for U.S. Bureau of Economic Analysis, BLS for U.S. Bureau of Labor Statistics and FED for the Federal Reserve System. The third and fourth columns give the level of aggregation of the data. The fifth column indicate whether that variable is usually used in the empirical literature or not. Finally, the last column shows for which of those variables we extract cyclical components based on the HP filter.

C.3 Results Perturbing the Instrument Set

Given the number of endogenous variables in the NKPCs and the potential sensitivity of our results, we also evaluate the model based on four different instrument sets. The first approach chooses instruments based on a data-driven technique that relies on regularisation. The three remaining approaches test the sensitivity of our findings to the lag structure of the variables used as instruments in the main paper.

First, based on results in Berriel, Medeiros, and Sena (2016), we apply instrument selection in the GMM setting. For each endogenous variable $\mathcal{Y}_t$ in the aggregate NKPC (8), we run the AdaLASSO estimator (Zou (2006)):

$$\hat{\rho} = argmin_{\rho} \|\mathcal{Y} - Z\rho\|_2^2 + \Lambda \sum_{j=1}^{P} w_j \|\rho_j\|_1,$$

where $\|\|_p$ is the ℓ^p norm, ρ is a $P \times 1$ vector of coefficients, Z is a $T \times P$ matrix of instrument candidates, $\{z_1, ..., z_p, ..., z_P\}$. Λ controls the shrinkage whereas $w_j = |\tilde{\rho}_j|^{-\tau}$ is a candidate-specific penalty weight formed by a preliminary LASSO estimator $\tilde{\rho}_j$. Finally, $\tau = 1$ is a common choice. Candidate z_p is selected as instrument if $\hat{\rho}_p \neq 0$.

The matrix of candidates Z is comprised of the first lags of variables used as instruments in the main paper. For each $z_{p,t-1}$ selected, we apply the first two lags of such variable, $\{z_{p,t-1}, z_{p,t-2}\}$. This rule follows the approach with the best results in the framework of Berriel, Medeiros, and Sena (2016) – there, called "AdaLASSO Observables". The authors show how such approach provides more reliable and disciplined estimations of the NKPC of a homogeneous economy. Our selection routine is just a natural extension of theirs to the multi-sector heterogeneous economy[1]. Applying the same rule with similar instruments for sectoral NKPCs often resulted in under-identification of some equations. Hence, we main-

[1] Berriel, Medeiros, and Sena (2016) also choose instruments based on the first lag, but apply the first three lags of selected variables. We apply the first two lags in our setting due to the number of moment conditions in our GMM.

tain the instruments applied in the main paper for those. One could extend the method by including sector-specific candidates (e.g., using oil drilling measures as instrument candidates for the Phillips curve of "gasoline and other energy goods"). We leave these underlying alternatives to future research.

Table C.2 presents our findings under the instrument selection technique. For estimations that exploit variation through sectoral NKPCs, results are broadly similar to those exhibited in the main paper. The data-driven routine seems to improve estimations for the naive single-equation approach, compared to results of the main paper. Nonetheless, the slope is still at least seven times smaller under the naive method, compared to the estimations that consider heterogeneity, and the implied Calvo probability still lies too close to 1 to seem reliable.

Next, we conduct estimations with a small variation of the approach used in the main paper. We apply the second and third lags of the same variables used as instruments in the baseline estimations of Table 3.3, rather than the first two lags. This variation is intended to control for the well known time-aggregation bias in macro data – e.g., Hall (1988). In such case, it is advised to apply instruments that are at least lagged twice. By further lagging the instruments, some loss of precision is expected, due to potentially lower correlations between instruments and endogenous variables. However, Table C.3 shows that results of the main paper are essentially reconfirmed. Compared to results of Table 3.3, changing the instrument set seems to impact more heavily the identification under the naive approach. The Calvo probability in such case lies roughly at 1 for two out of the three calibrations applied.

The remaining two approaches evaluate the sensitivity of our findings to the number of moment conditions in the GMM. They reflect our best attempt to avoid the common pitfall of too many instruments affecting the identification – see Bårdsen, Jansen, and Nymoen (2004), Andrews and Stock (2005) and C. Hansen, Hausman, and W. Newey (2008), for example. It is well known that the use of too many instruments often biases Two-Stage Least Squares (TSLS) estimators towards the OLS limit distribution. This is also more evident the weaker the instruments are.

In Table C.4, we reduce the number of moment conditions in the GMM by dropping the second lags of those variables used as instruments in the main paper. The instrument set is then comprised of the first lags of variables used in the estimations of Table 3.3. Reducing the number of instruments by half potentially imply lower precision. However, Table C.4 displays results that are moderately similar to those of Table 3.3.

Table C.5 presents results for a similar instrument set, but now using the second lag of variables applied as instruments in the main paper. Therefore, the instrument set is comprised of the second lag of those variables. Results do not differ much from those shown in the baseline estimations of Table 3.3. Confidence intervals are slightly wider compared to results of the main paper, potentially related to some loss of precision. However, point estimates suggest that our main findings are maintained.

Table C.2: Estimates of the Slope and Degree of Stickiness using Instrument Selection

			Parameters			
Calibration	Model	$Corr(\theta_k, Micro)$	κ	θ	β	γ
↑ Real Rigidity	Heterogeneous (SYS)	0.79	0.070***	0.48	0.99	0.85
			(0.001)	[0.46, 0.50]	(0.040)	(0.035)
	Homogeneous (SYS)	-	0.036***	0.56	0.97	0.34
			(0.000)	[0.56, 0.56]	(0.003)	(0.002)
	Homogeneous (SE)	-	0.007***	0.78	0.97	0.34
			(0.002)	[0.72, 0.84]	(0.003)	(0.017)
Baseline	Heterogeneous (SYS)	0.58	0.075***	0.56	0.99	0.29
			(0.001)	[0.56, 0.56]	(0.006)	(0.003)
	Homogeneous (SYS)	-	0.057***	0.57	0.95	0.36
			(0.000)	[0.57, 0.57]	(0.004)	(0.003)
	Homogeneous (SE)	-	0.009***	0.81	0.95	0.31
			(0.002)	[0.77, 0.86]	(0.004)	(0.019)
↓ Real Rigidity	Heterogeneous (SYS)	0.30	0.075***	0.66	0.97	0.22
			(0.001)	[0.65, 0.66]	(0.007)	(0.004)
	Homogeneous (SYS)	-	0.066***	0.63	0.98	0.30
			(0.000)	[0.62, 0.63]	(0.006)	(0.003)
	Homogeneous (SE)	-	0.010***	0.86	0.92	0.38
			(0.002)	[0.83, 0.89]	(0.004)	(0.017)

Notes: Results under instrument selection based on AdaLASSO. The first column refers to the three different calibration sets exhibited in Table 3.1. We test three different estimation methods. "Heterogeneous (SYS)" denotes the baseline model with sector-specific λ_k, being estimated by System-GMM with the aggregate NKPC in (8) and the fifteen sectoral NKPCs in (9). "Homogeneous (SYS)" uses the same system, but imposes $\lambda_k = \lambda$ for every sector. In such case, the shift term disappears from (8). "Homogeneous (SE)" mimics the standard approach in the literature, repeating this last exercise considering solely the aggregate NKPC (8), i.e., single-equation estimation. Correlations between estimated and benchmark infrequencies $(1-\lambda_k)$ that come from the micro data in Bils and Klenow (2004) are shown in the column "$Corr(\theta_k, Micro)$". **The micro benchmark implies $\boldsymbol{\theta^{micro} \approx 0.48}$.** κ denotes the aggregate slope in (8), while θ is the implied degree of stickiness in the economy. When λ_k varies across sectors (heterogeneous case), $\theta = \sum_{k=1}^{K} \eta_k \theta_k = \sum_{k=1}^{K} \eta_k (1-\lambda_k)$. Under homogeneity $(\lambda_k = \lambda)$, this simplifies to $\theta = (1-\lambda)$. We use a HAC estimator for the covariance matrix. Standard errors are presented in parentheses. As in theory, structural parameters $(\lambda_k, \beta$ and $\gamma)$ can assume values in the interval $[0, 1]$. We test the null hypothesis of $\kappa = 0$: *p<0.1; **p<0.05; ***p<0.01.

Table C.3: Estimates of the Slope and Degree of Stickiness – Alternative Instrument Set I

			Parameters			
Calibration	Model	$Corr(\theta_k, Micro)$	κ	θ	β	γ
↑ Real Rigidity	Heterogeneous (SYS)	0.16	0.032***	0.62	0.99	0.74
			(0.000)	[0.61, 0.63]	(0.007)	(0.005)
	Homogeneous (SYS)	-	0.036***	0.48	0.95	0.44
			(0.000)	[0.47, 0.48]	(0.004)	(0.002)
	Homogeneous (SE)	-	0.000	0.99	0.97	0.48
			(0.004)	[0.00, 1.00]	(0.011)	(0.073)
Baseline	Heterogeneous (SYS)	0.50	0.101***	0.59	0.99	0.20
			(0.001)	[0.58, 0.59]	(0.009)	(0.005)
	Homogeneous (SYS)	-	0.058***	0.48	0.96	0.42
			(0.000)	[0.44, 0.52]	(0.008)	(0.005)
	Homogeneous (SE)	-	0.000	0.99	0.98	0.54
			(0.003)	[0.00, 1.00]	(0.010)	(0.062)
↓ Real Rigidity	Heterogeneous (SYS)	0.61	0.114***	0.61	0.97	0.19
			(0.002)	[0.60, 0.62]	(0.022)	(0.012)
	Homogeneous (SYS)	-	0.082***	0.58	0.98	0.48
			(0.005)	[0.55, 0.61]	(0.047)	(0.029)
	Homogeneous (SE)	-	0.002	0.91	0.99	0.61
			(0.003)	[0.81, 1.00]	(0.010)	(0.052)

Notes: Results applying the second and third lags of the same variables used as instruments in Table 3.3. The first column refers to the three different calibration sets exhibited in Table 3.1. We test three different estimation methods. "Heterogeneous (SYS)" denotes the baseline model with sector-specific λ_k, being estimated by System-GMM with the aggregate NKPC in (8) and the fifteen sectoral NKPCs in (9). "Homogeneous (SYS)" uses the same system, but imposes $\lambda_k = \lambda$ for every sector. In such case, the shift term disappears from (8). "Homogeneous (SE)" mimics the standard approach in the literature, repeating this last exercise considering solely the aggregate NKPC (8), i.e., single-equation estimation. Correlations between estimated and benchmark infrequencies $(1-\lambda_k)$ that come from the micro data in Bils and Klenow (2004) are shown in the column "$Corr(\theta_k, Micro)$". **The micro benchmark implies $\boldsymbol{\theta^{micro} \approx 0.48}$.** κ denotes the aggregate slope in (8), while θ is the implied degree of stickiness in the economy. When λ_k varies across sectors (heterogeneous case), $\theta = \sum_{k=1}^{K} \eta_k \theta_k = \sum_{k=1}^{K} \eta_k (1-\lambda_k)$. Under homogeneity ($\lambda_k = \lambda$), this simplifies to $\theta = (1-\lambda)$. We use a HAC estimator for the covariance matrix. Standard errors are presented in parentheses. As in theory, structural parameters (λ_k, β and γ) can assume values in the interval $[0, 1]$. We test the null hypothesis of $\kappa = 0$: *p<0.1; **p<0.05; ***p<0.01.

Table C.4: Estimates of the Slope and Degree of Stickiness – Alternative Instrument Set II

Calibration	Model	$Corr(\theta_k, Micro)$	Parameters κ	θ	β	γ
$\uparrow$ Real Rigidity	Heterogeneous (SYS)	0.69	0.033***	0.68	0.99	0.50
			(0.000)	[0.66, 0.70]	(0.039)	(0.024)
	Homogeneous (SYS)	-	0.015***	0.70	0.98	0.15
			(0.000)	[0.70, 0.70]	(0.003)	(0.002)
	Homogeneous (SE)	-	0.000	0.99	0.98	0.36
			(0.006)	[0.00, 1.00]	(0.011)	(0.092)
Baseline	Heterogeneous (SYS)	0.87	0.086***	0.53	0.99	0.89
			(0.006)	[0.45, 0.62]	(0.122)	(0.110)
	Homogeneous (SYS)	-	0.026***	0.68	0.97	0.33
			(0.000)	[0.68, 0.69]	(0.003)	(0.002)
	Homogeneous (SE)	-	0.001	0.93	0.98	0.41
			(0.005)	[0.60, 1.00]	(0.011)	(0.090)
$\downarrow$ Real Rigidity	Heterogeneous (SYS)	0.74	0.173***	0.55	0.99	0.13
			(0.010)	[0.52, 0.57]	(0.036)	(0.018)
	Homogeneous (SYS)	-	0.088***	0.54	0.99	0.42
			(0.018)	[0.39, 0.69]	(0.010)	(0.006)
	Homogeneous (SE)	-	0.001	0.96	0.97	0.47
			(0.004)	[0.70, 1.00]	(0.012)	(0.087)

Notes: Results when the second lags of variables used as instruments in Table 3.3 are dropped from the instrument set. Hence, instruments are the first lags of those variables. The first column refers to the three different calibration sets exhibited in Table 3.1. We test three different estimation methods. "Heterogeneous (SYS)" denotes the baseline model with sector-specific λ_k, being estimated by System-GMM with the aggregate NKPC in (8) and the fifteen sectoral NKPCs in (9). "Homogeneous (SYS)" uses the same system, but imposes $\lambda_k = \lambda$ for every sector. In such case, the shift term disappears from (8). "Homogeneous (SE)" mimics the standard approach in the literature, repeating this last exercise considering solely the aggregate NKPC (8), i.e., single-equation estimation. Correlations between estimated and benchmark infrequencies $(1 - \lambda_k)$ that come from the micro data in Bils and Klenow (2004) are shown in the column "$Corr(\theta_k, Micro)$". **The micro benchmark implies $\boldsymbol{\theta^{micro} \approx 0.48}$.** κ denotes the aggregate slope in (8), while θ is the implied degree of stickiness in the economy. When λ_k varies across sectors (heterogeneous case), $\theta = \sum_{k=1}^{K} \eta_k \theta_k = \sum_{k=1}^{K} \eta_k (1 - \lambda_k)$. Under homogeneity $(\lambda_k = \lambda)$, this simplifies to $\theta = (1 - \lambda)$. We use a HAC estimator for the covariance matrix. Standard errors are presented in parentheses. As in theory, structural parameters (λ_k, β and γ) can assume values in the interval $[0, 1]$. We test the null hypothesis of $\kappa = 0$: *p<0.1; **p<0.05; ***p<0.01.

Table C.5: Estimates of the Slope and Degree of Stickiness – Alternative Instrument Set III

Calibration	Model	$Corr(\theta_k, Micro)$	Parameters κ	θ	β	γ
↑ Real Rigidity	Heterogeneous (SYS)	0.43	0.115^{***}	0.51	0.99	0.30
			(0.021)	[0.19, 0.84]	(0.147)	(0.080)
	Homogeneous (SYS)	-	0.017^{***}	0.71	0.93	0.74
			(0.000)	[0.71, 0.71]	(0.001)	(0.001)
	Homogeneous (SE)	-	0.000	0.99	0.99	0.31
			(0.005)	[0.00, 1.00]	(0.010)	(0.099)
Baseline	Heterogeneous (SYS)	0.50	0.096^{***}	0.57	0.99	0.38
			(0.008)	[0.45, 0.68]	(0.071)	(0.041)
	Homogeneous (SYS)	-	0.029^{***}	0.58	0.93	0.78
			(0.002)	[0.56, 0.61]	(0.033)	(0.029)
	Homogeneous (SE)	-	0.000	0.99	0.99	0.33
			(0.005)	[0.00, 1.00]	(0.010)	(0.100)
↓ Real Rigidity	Heterogeneous (SYS)	0.74	0.173^{***}	0.55	0.99	0.13
			(0.010)	[0.52, 0.57]	(0.036)	(0.018)
	Homogeneous (SYS)	-	0.088^{***}	0.54	0.99	0.42
			(0.018)	[0.39, 0.69]	(0.010)	(0.006)
	Homogeneous (SE)	-	0.000	0.99	0.99	0.35
			(0.005)	[0.00, 1.00]	(0.010)	(0.100)

Notes: Results when the first lags of variables used as instruments in Table 3.3 are dropped from the instrument set. Hence, instruments are the second lags of those variables. The first column refers to the three different calibration sets exhibited in Table 3.1. We test three different estimation methods. "Heterogeneous (SYS)" denotes the baseline model with sector-specific λ_k, being estimated by System-GMM with the aggregate NKPC in (8) and the fifteen sectoral NKPCs in (9). "Homogeneous (SYS)" uses the same system, but imposes $\lambda_k = \lambda$ for every sector. In such case, the shift term disappears from (8). "Homogeneous (SE)" mimics the standard approach in the literature, repeating this last exercise considering solely the aggregate NKPC (8), i.e., single-equation estimation. Correlations between estimated and benchmark infrequencies $(1 - \lambda_k)$ that come from the micro data in Bils and Klenow (2004) are shown in the column "$Corr(\theta_k, Micro)$". **The micro benchmark implies $\boldsymbol{\theta^{\text{micro}} \approx 0.48}$.** κ denotes the aggregate slope in (8), while θ is the implied degree of stickiness in the economy. When λ_k varies across sectors (heterogeneous case), $\theta = \sum_{k=1}^{K} \eta_k \theta_k = \sum_{k=1}^{K} \eta_k (1 - \lambda_k)$. Under homogeneity ($\lambda_k = \lambda$), this simplifies to $\theta = (1 - \lambda)$. We use a HAC estimator for the covariance matrix. Standard errors are presented in parentheses. As in theory, structural parameters (λ_k, β and γ) can assume values in the interval $[0, 1]$. We test the null hypothesis of $\kappa = 0$: *p<0.1; **p<0.05; ***p<0.01.

C.4 Estimated Sectoral Infrequencies Under the Alternative Calibrations

Figure C.1 and Figure C.2 present similar findings to those in Figure 3.2, but for the remaining two calibrations described in the second and third columns of Table 3.1, respectively.

Figure C.1: $\hat{\theta}_k$ vs. Micro Benchmarks – Confidence Intervals

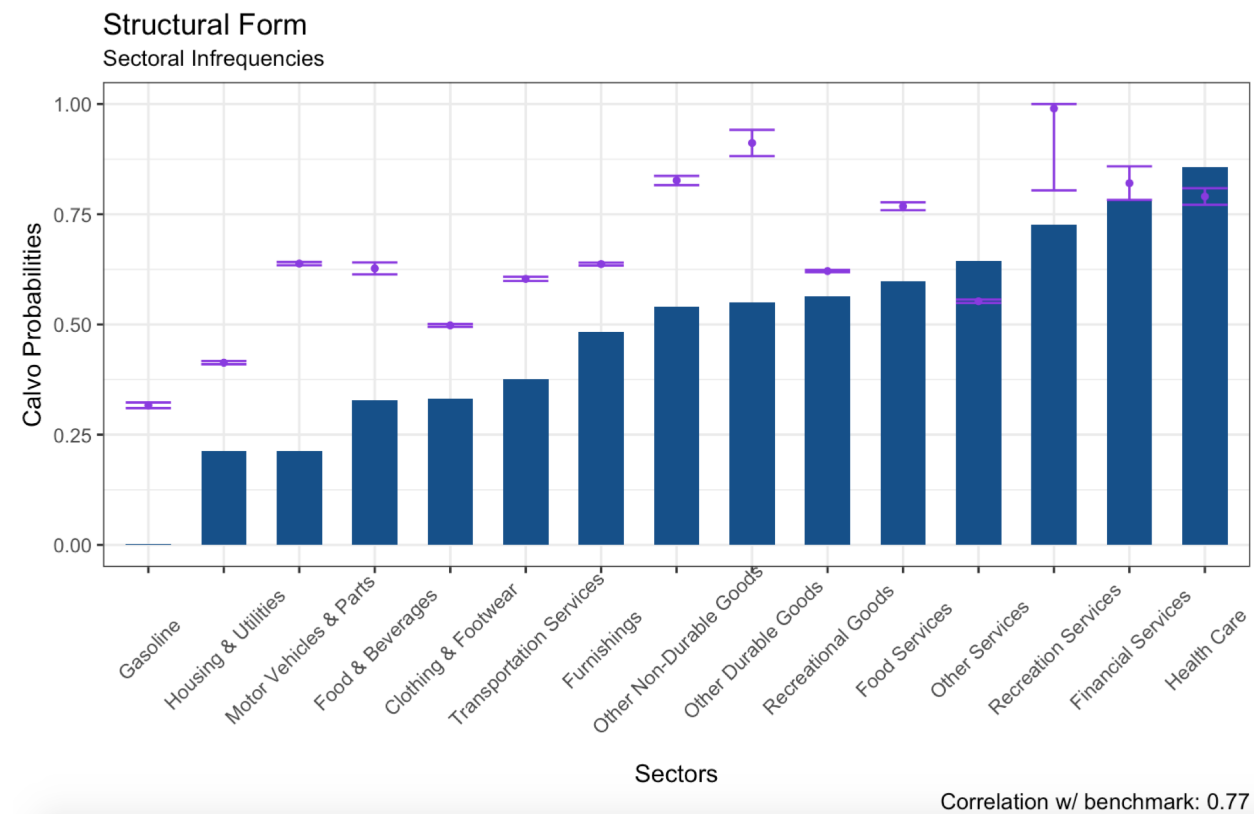

Notes: Estimated Calvo probabilities using the same econometric setting of Table 3.3. Blue bars are micro-based benchmark probabilities implied from evidence in Bils and Klenow (2004) and presented in Table 3.2. For expository purposes, these are sorted according to their degree of flexibility. 95% confidence intervals are shown for each $\hat{\theta}_k$. We use the calibration described in the second column of Table 3.1.

Figure C.2: $\hat{\theta}_k$ vs. Micro Benchmarks – Confidence Intervals

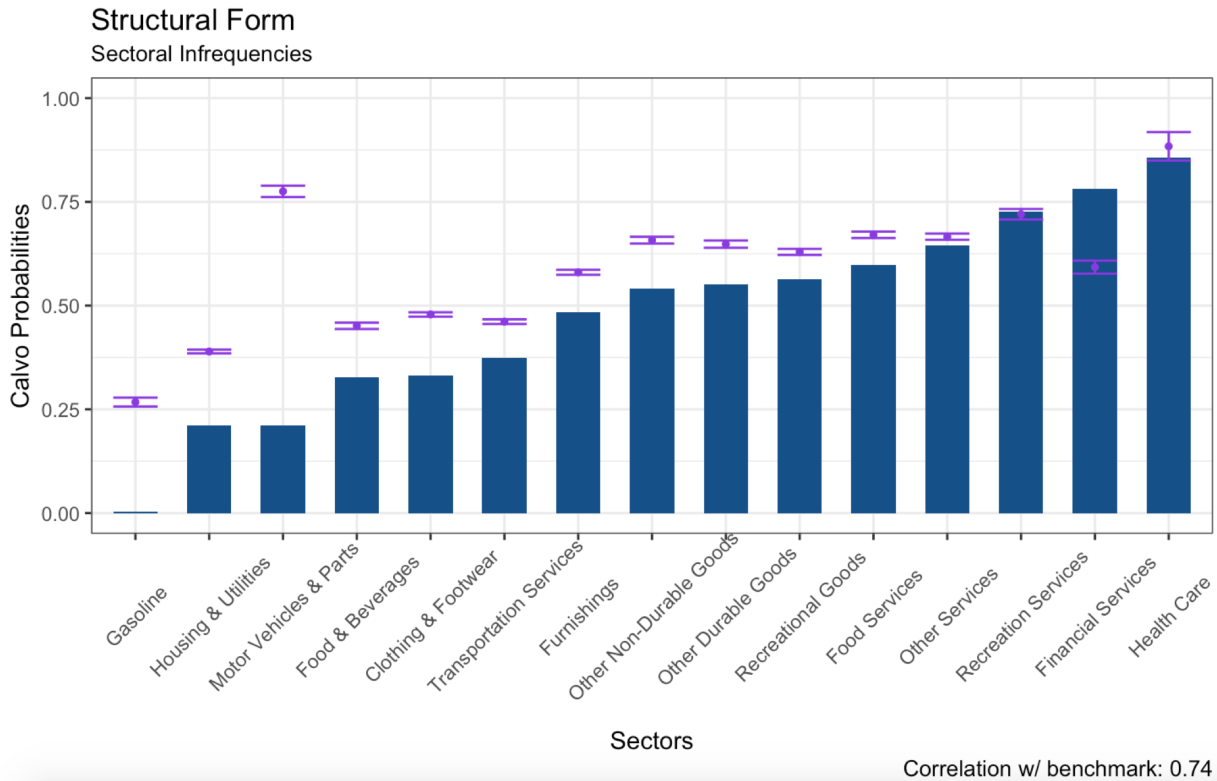

Notes: Estimated Calvo probabilities using the same econometric setting of Table 3.3. Blue bars are micro-based benchmark probabilities implied from evidence in Bils and Klenow (2004) and presented in Table 3.2. For expository purposes, these are sorted according to their degree of flexibility. 95% confidence intervals are shown for each $\hat{\theta}_k$. We use the calibration described in the third column of Table 3.1.

C.5 Removing the Aggregate NKPC from the System

In Table C.6 below, we remove the aggregate NKPC from our estimations – i.e., we show results for the case we estimate the system comprised of the fifteen sectoral NKPCs that take the form in (3.9). We maintain results of the main paper for the naive single-equation approach (based on the aggregate NKPC of the homogeneous economy) to facilitate the comparison. By ignoring information, some loss of precision is expected. However, Table C.6 basically reconfirms findings of the main paper.

Table C.6: Estimates of the Slope and Degree of Stickiness Dropping (3.8) from the System

			Parameters			
Calibration	Model	$Corr(\theta_k, Micro)$	κ	θ	β	γ
↑ Real Rigidity	Heterogeneous (SYS)	0.84	0.054***	0.64	0.98	0.31
			(0.001)	[0.62, 0.66]	(0.008)	(0.004)
	Homogeneous (SYS)	-	0.033***	0.58	0.96	0.30
			(0.000)	[0.58, 0.59]	(0.004)	(0.002)
	Homogeneous (SE)	-	0.002	0.87	0.96	0.44
			(0.003)	[0.70, 1.00]	(0.010)	(0.040)
Baseline	Heterogeneous (SYS)	0.78	0.136***	0.55	0.98	0.26
			(0.011)	[0.49, 0.59]	(0.043)	(0.023)
	Homogeneous (SYS)	-	0.044***	0.65	0.97	0.30
			(0.000)	[0.64, 0.65]	(0.005)	(0.002)
	Homogeneous (SE)	-	0.007*	0.82	0.96	0.45
			(0.003)	[0.74, 0.90]	(0.011)	(0.043)
↓ Real Rigidity	Heterogeneous (SYS)	0.76	0.126***	0.63	0.99	0.21
			(0.008)	[0.50, 0.76]	(0.029)	(0.015)
	Homogeneous (SYS)	-	0.094***	0.59	0.97	0.20
			(0.006)	[0.56, 0.62]	(0.038)	(0.020)
	Homogeneous (SE)	-	0.012***	0.81	0.97	0.49
			(0.004)	[0.77, 0.86]	(0.012)	(0.048)

Notes: Estimates when the aggregate NKPC (3.8) is dropped from the system. The first column refers to the three different calibration sets exhibited in Table 3.1. We test three different estimation methods. "Heterogeneous (SYS)" denotes the baseline model with sector-specific λ_k, being estimated by System-GMM with the fifteen sectoral NKPCs in (3.9). "Homogeneous (SYS)" uses the same system, but imposes $\lambda_k = \lambda$ for every sector. "Homogeneous (SE)" mimics the standard approach in the literature and repeats results of Table 3.3. Correlations between estimated and benchmark infrequencies $(1 - \lambda_k)$ that come from the micro data in Bils and Klenow (2004) are shown in the column "$Corr(\theta_k, Micro)$". **The micro benchmark implies $\boldsymbol{\theta^{micro} \approx 0.48}$.** κ denotes the aggregate slope in (3.8), while θ is the implied degree of stickiness in the economy. When λ_k varies across sectors (heterogeneous case), $\theta = \sum_{k=1}^{K} \eta_k \theta_k = \sum_{k=1}^{K} \eta_k (1 - \lambda_k)$. Under homogeneity ($\lambda_k = \lambda$), this simplifies to $\theta = (1 - \lambda)$. We use a HAC estimator for the covariance matrix. Standard errors are presented in parentheses. As in theory, structural parameters (λ_k, β and γ) can assume values in the interval $[0, 1]$. We test the null hypothesis of $\kappa = 0$: *p<0.1; **p<0.05; ***p<0.01.

C.6 Perturbing Starting Values

Non-linear GMM with complex moment conditions can be quite sensitive to starting values in the algorithm. We then shall evaluate the sensitivity of our results to such values. In the estimations of the main text, we used benchmark probabilities implied from micro data – Bils and Klenow (2004) – as initial values. In Table C.6, we replace that approach with an agnostic routine that finds starting values based on sectoral data. First, we estimate each sectoral NKPC in (3.9) individually – relying on the same instruments of Table 3.3. Second, we apply estimates of this first stage, $\hat{\lambda}^I_k$, as initial values when estimating the system based on (3.8) and (3.9). When heterogeneity is switched off, the initial value is calculated using $\hat{\lambda}^I = \sum_{k=1}^{K} \eta_k \hat{\lambda}^I_k$. For the reasons discussed in the main text, this approach is very conservative. Thus, some loss of precision is expected.

Table C.6 shows that, once more, the heterogeneous model considerably outperforms its homogeneous counterpart. The slope from the former is at least twice that of the latter, while estimated implied infrequencies are also lower for the heterogeneous economy. Correlations with benchmark probabilities are considerably closer to zero, but this is expected due to a potential misalignment between first-stage estimates (initial values) and the data.

Table C.7: Estimates of the Slope and Degree of Stickiness – Alternative Starting Values

			Parameters			
Calibration	Model	$Corr(\theta_k, Micro)$	κ	θ	β	γ
↑ Real Rigidity	Heterogeneous (SYS)	0.07	0.026***	0.69	0.86	0.50
			(0.000)	[0.68, 0.69]	(0.002)	(0.001)
	Homogeneous (SYS)	-	0.011***	0.74	0.96	0.26
			(0.000)	[0.73, 0.74]	(0.001)	(0.001)
	Homogeneous (SE)	-	0.001	0.78	0.96	0.44
			(0.003)	[0.68, 1.00]	(0.010)	(0.041)
Baseline	Heterogeneous (SYS)	-0.14	0.031***	0.75	0.92	0.35
			(0.000)	[0.75, 0.76]	(0.002)	(0.001)
	Homogeneous (SYS)	-	0.004***	0.87	0.96	0.31
			(0.000)	[0.87, 0.87]	(0.007)	(0.004)
	Homogeneous (SE)	-	0.000	0.99	0.96	0.42
			(0.003)	[0.18, 1.00]	(0.010)	(0.041)
↓ Real Rigidity	Heterogeneous (SYS)	-0.13	0.021***	0.83	0.93	0.41
			(0.000)	[0.83, 0.84]	(0.002)	(0.001)
	Homogeneous (SYS)	-	0.013***	0.83	0.95	0.32
			(0.000)	[0.83, 0.83]	(0.001)	(0.001)
	Homogeneous (SE)	-	0.000	0.99	0.96	0.43
			(0.003)	[0.52, 1.00]	(0.010)	(0.040)

Notes: Results perturbing initial values in the algorithm. In a first step, we estimate each sectoral NKPC represented by (8) individually with the baseline instrument set of the main paper. Next, we apply each estimate $\hat{\lambda}_k$ of this first stage regressions as initial value for that parameter when estimating the system comprised of (8) and (9). For the homogeneous models, we calculate the aggregate infrequency implied from estimates of the first stage, using it as the initial value in the algorithm. The first column refers to the three different calibration sets exhibited in Table 3.1. We test three different estimation methods. "Heterogeneous (SYS)" denotes the baseline model with sector-specific λ_k, being estimated by System-GMM with the aggregate NKPC in (8) and the fifteen sectoral NKPCs in (9). "Homogeneous (SYS)" uses the same system, but imposes $\lambda_k = \lambda$ for every sector. In such case, the shift term disappears from (8). "Homogeneous (SE)" mimics the standard approach in the literature, repeating this last exercise considering solely the aggregate NKPC (8), i.e., single-equation estimation. Correlations between estimated and benchmark infrequencies $(1-\lambda_k)$ that come from the micro data in Bils and Klenow (2004) are shown in the column "$Corr(\theta_k, Micro)$". **The micro benchmark implies $\boldsymbol{\theta^{\text{micro}} \approx 0.48}$.** κ denotes the aggregate slope in (8), while θ is the implied degree of stickiness in the economy. When λ_k varies across sectors (heterogeneous case), $\theta = \sum_{k=1}^{K} \eta_k \theta_k = \sum_{k=1}^{K} \eta_k (1-\lambda_k)$. Under homogeneity ($\lambda_k = \lambda$), this simplifies to $\theta = (1-\lambda)$. We use a HAC estimator for the covariance matrix. Standard errors are presented in parentheses. As in theory, structural parameters (λ_k, β and γ) can assume values in the interval $[0, 1]$. We test the null hypothesis of $\kappa = 0$: *p<0.1; **p<0.05; ***p<0.01.

C.7 Expanded Sample and Alternative Calibrations

In this section, I report results for a sample that includes the recent inflationary period beginning in mid-2021. I also provide robustness checks under alternative calibrations of ϵ.

As expected, elevated inflation strengthens the relationship between prices and real activity, resulting in a substantially steeper Phillips curve. These estimates are not directly comparable to those in the existing literature, however, as prior studies typically exclude this more recent inflationary episode.

Figure C.3: $\hat{\theta}_k$ vs. Micro Benchmarks (Extended Sample)

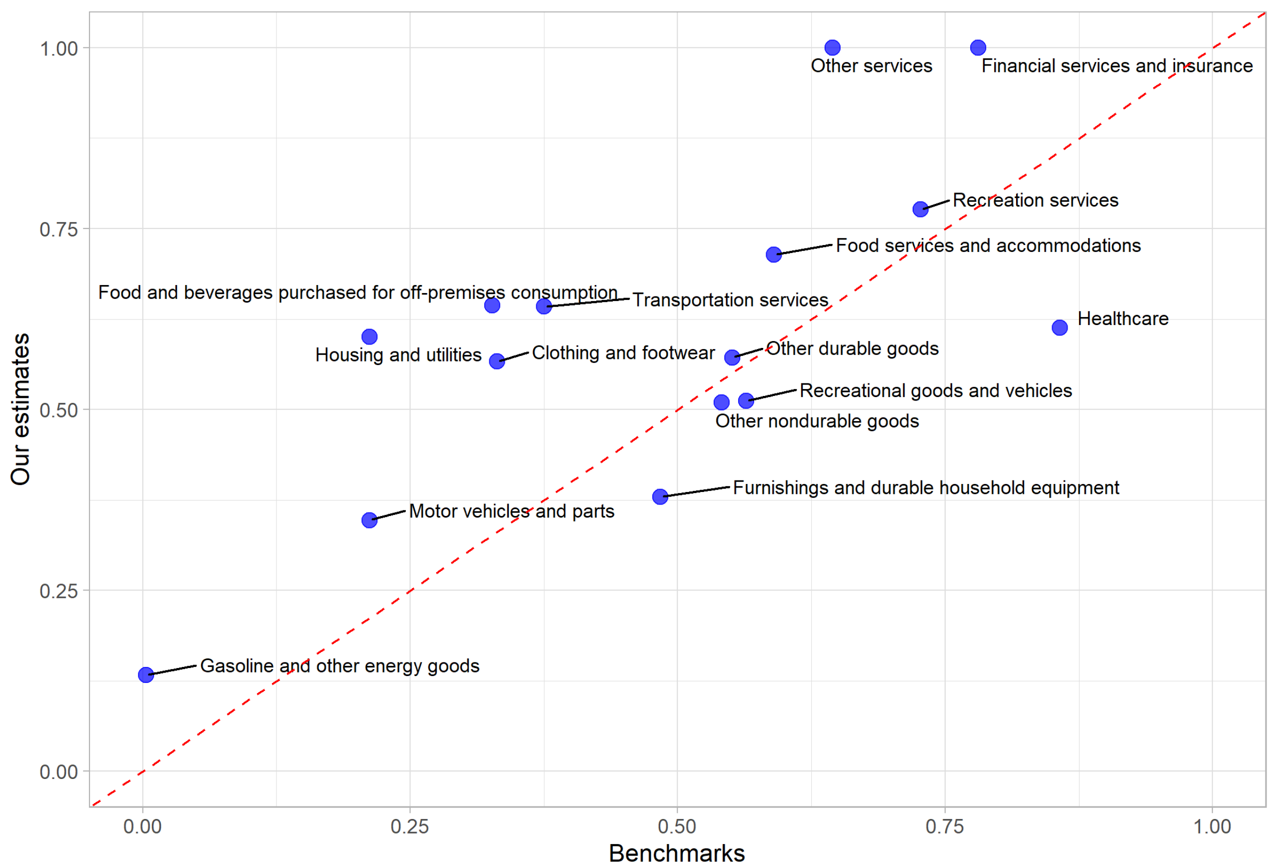

Notes: Estimated Calvo probabilities using the same econometric setting of Table 3.3. Sample is extended until 2024. Benchmarks are implied probabilities from evidence in Bils and Klenow (2004) – see Table 3.2. We use the baseline calibration of Table 3.1. **Correlation with benchmark:** 81%.

Table C.8: Phillips Curve Slope and Correlations for Varying ϵ

	$\epsilon = 3$	$\epsilon = 4$	$\epsilon = 5$	$\epsilon = 8$
Slope of Phillips Curve (κ)	0.2740	0.2740	0.2740	0.2740
Correlation with Benchmarks	0.7108	0.7063	0.7018	0.6890
Correlation (Excluding Food)	0.7591	0.7539	0.7488	0.7341
Standard Error	0.0152	0.0152	0.0152	0.0152

Notes: Implied slope of the aggregate Phillips curve constructed from estimated Calvo probabilities using the same econometric setting of Table 3.3. Sample is extended until 2024. Benchmarks are implied probabilities from evidence in Bils and Klenow (2004) — see Table 3.2. We test alternative values of ϵ — remaining parametric values are those of the baseline calibration in Table 3.1.

Bibliography

Ait-Sahalia, Yacine, Jonathan A Parker, and Motohiro Yogo (2004). "Luxury goods and the equity premium". In: *The Journal of Finance* 59.6, pp. 2959–3004.

Albuquerque, Bruno and Ursel Baumann (2017). "Will US inflation awake from the dead? The role of slack and non-linearities in the Phillips curve". In: *Journal of Policy Modeling* 39.2, pp. 247–271.

Anderson, Evan et al. (1996). "Mechanics of forming and estimating dynamic linear economies". In: *Handbook of Computational Economics*. Ed. by H. M. Amman, D. A. Kendrick, and J. Rust. 1st ed. Vol. 1. Elsevier. Chap. 04, pp. 171–252.

Anderson, Theodore W, Herman Rubin, et al. (1949). "Estimation of the parameters of a single equation in a complete system of stochastic equations". In: *The Annals of Mathematical Statistics* 20.1, pp. 46–63.

Andrews, Donald and James H Stock (2005). *Inference with weak instruments*.

Aoki, Kosuke (2001). "Optimal monetary policy responses to relative-price changes". In: *Journal of monetary economics* 48.1, pp. 55–80.

Ascari, Guido, Leandro Magnusson, and Sophocles Mavroeidis (2016). "Empirical Evidence on the Euler equation for Consumption and Output in the US". In.

Attanasio, Orazio P and Guglielmo Weber (1995). "Is consumption growth consistent with intertemporal optimization? Evidence from the consumer expenditure survey". In: *Journal of political Economy* 103.6, pp. 1121–1157.

Backus, David K and Gregor W Smith (1993). “Consumption and real exchange rates in dynamic economies with non-traded goods”. In: *Journal of International Economics* 35.3–4, pp. 297–316.

Ball, Laurence and David Romer (1990). “Real rigidities and the non-neutrality of money”. In: *The Review of Economic Studies* 57.2, pp. 183–203.

Bansal, Ravi, Dana Kiku, et al. (2014). “Volatility, the macroeconomy, and asset prices”. In: *The Journal of Finance* 69.6, pp. 2471–2511.

Bansal, Ravi and Ivan Shaliastovich (2013). “A long-run risks explanation of predictability puzzles in bond and currency markets”. In: *The Review of Financial Studies* 26.1, pp. 1–33.

Bansal, Ravi and Amir Yaron (Aug. 2004). “Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles”. In: *Journal of Finance* 59.4, pp. 1481–1509.

Bårdsen, Gunnar, Eilev S Jansen, and Ragnar Nymoen (2004). “Econometric evaluation of the New Keynesian Phillips curve”. In: *Oxford Bulletin of Economics and Statistics* 66, pp. 671–686.

BEA (2017). *NIPA Handbook: Concepts and Methods of the U.S. National Income and Product Accounts*. U.S. Bureau of Economic Analysis.

Bell, William R and David W Wilcox (1993). “The effect of sampling error on the time series behavior of consumption data”. In: *Journal of Econometrics* 55.1-2, pp. 235–265.

Benigno, Pierpaolo (2004). “Optimal monetary policy in a currency area”. In: *Journal of international economics* 63.2, pp. 293–320.

Berriel, Tiago, Marcelo Medeiros, and Marcelo J. Sena (2016). “Instrument selection for estimation of a forward-looking Phillips Curve”. In: *Economics Letters* 145.C, pp. 123–125.

Bils, Mark and Pete Klenow (2004). “Some Evidence on the Importance of Sticky Prices”. In: *Journal of Political Economy* 112.5, pp. 947–985.

Blanchard, Olivier and Jordi Galí (2007). “Real wage rigidities and the New Keynesian model”. In: *Journal of money, credit and banking* 39, pp. 35–65.

Bollerslev, Tim, George Tauchen, and Hao Zhou (2009). "Expected stock returns and variance risk premia". In: *The Review of Financial Studies* 22.11, pp. 4463–4492.

Breeden, Douglas T, Michael R Gibbons, and Robert H Litzenberger (1989). "Empirical tests of the consumption-oriented CAPM". In: *The Journal of Finance* 44.2, pp. 231–262.

Brunnermeier, Markus K, Stefan Nagel, and Lasse Heje Pedersen (2009). "Carry trades and currency crashes". In: *NBER Macroeconomics Annual* 23, pp. 313–347.

Cagliarini, Adam, Tim Robinson, and Allen Tran (2011). "Reconciling microeconomic and macroeconomic estimates of price stickiness". In: *Journal of Macroeconomics* 33.1, pp. 102–120.

Calvo, Guillermo A. (Sept. 1983). "Staggered prices in a utility-maximizing framework". In: *Journal of Monetary Economics* 12.3, pp. 383–398.

Camanho, Nelson, Harald Hau, and Hélène Rey (Apr. 2022). "Global Portfolio Rebalancing and Exchange Rates". In: *The Review of Financial Studies*. hhac023. ISSN: 0893-9454. eprint: https://academic.oup.com/rfs/advance-article-pdf/doi/10.1093/rfs/hhac023/43970636/hhac023.pdf.

Campbell, John Y (2003). "Consumption-based asset pricing". In: *Handbook of the Economics of Finance* 1, pp. 803–887.

Campbell, John Y, Luis M Viceira, Luis M Viceira, et al. (2002). *Strategic asset allocation: portfolio choice for long-term investors*. Clarendon Lectures in Economic.

Carvalho, Carlos (Dec. 2006). "Heterogeneity in Price Stickiness and the Real Effects of Monetary Shocks". In: *The B.E. Journal of Macroeconomics* 6.3, pp. 1–58.

Carvalho, Carlos and Fernanda Nechio (2016). "Factor specificity and real rigidities". In: *Review of Economic Dynamics* 22, pp. 208–222.

Cenedese, Gino et al. (July 2015). "What Do Stock Markets Tell Us about Exchange Rates? *". In: *Review of Finance* 20.3, pp. 1045–1080. ISSN: 1572-3097. eprint: https://academic.oup.com/rof/article-pdf/20/3/1045/26318121/rfv032.pdf.

Chang, Yongsung and Sun-Bin Kim (2006). "From individual to aggregate labor supply: A quantitative analysis based on a heterogeneous agent macroeconomy". In: *International Economic Review* 47.1, pp. 1–27.

Christiano, Lawrence J, Martin Eichenbaum, and Charles L Evans (2005). "Nominal rigidities and the dynamic effects of a shock to monetary policy". In: *Journal of political Economy* 113.1, pp. 1–45.

Clarida, Richard, Jordi Gali, and Mark Gertler (2000). "Monetary policy rules and macroeconomic stability: evidence and some theory". In: *The Quarterly journal of economics* 115.1, pp. 147–180.

Cochrane, John H (1996). "A cross-sectional test of an investment-based asset pricing model". In: *Journal of Political Economy* 104.3, pp. 572–621.

Cogley, Timothy and Argia M Sbordone (2008). "Trend inflation, indexation, and inflation persistence in the New Keynesian Phillips curve". In: *American Economic Review* 98.5, pp. 2101–26.

Colacito, Riccardo and Mariano M. Croce (2011). "Risks for the Long Run and the Real Exchange Rate". In: *Journal of Political Economy* 119.1, pp. 153–181. ISSN: 00223808, 1537534X.

Corte, Pasquale Della, Steven J. Riddiough, and Lucio Sarno (2016). "Currency Premia and Global Imbalances". In: *The Review of Financial Studies* 29.8, pp. 2161–2193. ISSN: 08939454, 14657368.

Curcuru, Stephanie E. et al. (2014). "Uncovered Equity Parity and rebalancing in international portfolios". In: *Journal of International Money and Finance* 47, pp. 86–99. ISSN: 0261-5606.

Daniel, Kent and David Marshall (1996). *The equity premium puzzle and the risk-free rate puzzle at long horizons*. Working Paper Series, Issues in Financial Regulation WP-96-4. Federal Reserve Bank of Chicago.

Della Corte, Pasquale, Steven J. Riddiough, and Lucio Sarno (June 2016). "Currency Premia and Global Imbalances". In: *The Review of Financial Studies* 29.8, pp. 2161–2193. ISSN: 0893-

9454. eprint: https://academic.oup.com/rfs/article-pdf/29/8/2161/24453308/hhw038.pdf.

Dew-Becker, Ian (2016). “How risky is consumption in the long-run? benchmark estimates from a robust estimator”. In: *The Review of Financial Studies* 30.2, pp. 631–666.

Dhyne, Emmanuel et al. (2006). “Price changes in the euro area and the United States: Some facts from individual consumer price data”. In: *Journal of Economic Perspectives* 20.2, pp. 171–192.

Dufour, Jean-Marie, Lynda Khalaf, and Maral Kichian (2010). “On the precision of Calvo parameter estimates in structural NKPC models”. In: *Journal of Economic Dynamics and Control* 34.9, pp. 1582–1595.

Epstein, Larry G and Stanley E Zin (July 1989). “Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework”. In: *Econometrica* 57.4, pp. 937–969.

Epstein, Lawrence and Stanley E Zin (1991). “Substitution, risk aversion, and the temporal behavior of consumption and asset returns: An empirical analysis”. In: *Journal of political Economy* 99.2, pp. 263–286.

Eusepi, Stefano, Bart Hobijn, and Andrea Tambalotti (2011). “CONDI: a cost-of-nominal-distortions index”. In: *American Economic Journal: Macroeconomics* 3.3, pp. 53–91.

Fama, Eugene F. (1984). “Forward and spot exchange rates”. In: *Journal of Monetary Economics* 14.3, pp. 319–338. ISSN: 0304-3932.

Fitzgerald, Terry J, Juan Pablo Nicolini, et al. (2014). *Is There a Stable Relationship Between Unemployment and Future Inflation?: Evidence from US Cities*. Federal Reserve Bank of Minneapolis, Research Department.

Fuhrer, Jeff and George Moore (1995). “Inflation Persistence”. In: *The Quarterly Journal of Economics* 110.1, pp. 127–159.

Fuhrer, Jeffrey C. (2000). “Habit formation in consumption and its implications for monetary-policy models”. In: *American Economic Review* 90.3, pp. 367–390.

Fuhrer, Jeffrey C. and Glenn D. Rudebusch (2002). *Estimating the Euler equation for output*. Working Papers 02-3. Federal Reserve Bank of Boston.

Gabaix, X. and M Maggiori (2015). "International Liquidity and Exchange Rate Dynamics". In: *Quarterly Journal of Economics* 130.3, pp. 1369–1420.

Gali, Jordi and Mark Gertler (Oct. 1999). "Inflation dynamics: A structural econometric analysis". In: *Journal of Monetary Economics* 44.2, pp. 195–222.

Galí, Jordi (2015). *Monetary policy, inflation, and the business cycle: an introduction to the new Keynesian framework and its applications*. Princeton University Press.

Galí, Jordi and Luca Gambetti (2019). *Has the US wage phillips curve flattened? A semi-structural exploration*. Tech. rep. National Bureau of Economic Research.

Galí, Jordi, Mark Gertler, and J. David López-Salido (June 2003). *Robustness of the Estimates of the Hybrid New Keynesian Phillips Curve*. Working Papers 44. Barcelona Graduate School of Economics.

Garner, Thesia I et al. (2003). "The Consumer Expenditure Survey in comparison: focus on personal consumption expenditures". In: *Federal Economic Statistics Advisory Committee Meeting, March*. Vol. 17.

Gomes, Fábio Augusto Reis and Lourenço S Paz (2013). "Estimating the elasticity of intertemporal substitution: Is the aggregate financial return free from the weak instrument problem?" In: *Journal of Macroeconomics* 36, pp. 63–75.

Gourinchas, Pierre-Olivier and Helene Rey (2007). "From world banker to world venture capitalist: US external adjustment and the exorbitant privilege". In: *G7 Current Account Imbalances: Sustainability and Adjustment*. University of Chicago Press, pp. 11–66.

Gourinchas, Pierre-Olivier, Helene Rey, and Kai Truempler (2012). "The financial crisis and the geography of wealth transfers". In: *Journal of International Economics*.

Guvenen, Fatih (2006). "Reconciling conflicting evidence on the elasticity of intertemporal substitution: A macroeconomic perspective". In: *Journal of Monetary Economics* 53.7, pp. 1451–1472.

Hall, Robert (Feb. 1988). “Intertemporal Substitution in Consumption”. In: *Journal of Political Economy* 96, pp. 339–57.

Hamilton, James Douglas (1994). *Time series analysis*. Vol. 2. Princeton university press Princeton, NJ.

Hansen, Christian, Jerry Hausman, and Whitney Newey (2008). “Estimation with many instrumental variables”. In: *Journal of Business & Economic Statistics* 26.4, pp. 398–422.

Hansen, Lars Peter, John Heaton, and Amir Yaron (1996). “Finite-sample properties of some alternative GMM estimators”. In: *Journal of Business & Economic Statistics* 14.3, pp. 262–280.

Harvey, Andrew, Esther Ruiz, and Enrique Sentana (Apr. 1992). “Unobserved Component Time Series Models With ARCH Disturbances”. In: *Journal of Econometrics* 52, pp. 129–157.

Hau, Harald and Hélène Rey (2006). “Exchange Rates, Equity Prices, and Capital Flows”. In: *The Review of Financial Studies* 19.1, pp. 273–317. ISSN: 08939454, 14657368.

Havranek, Tomas et al. (2013). *Cross-Country Heterogeneity in Intertemporal Substitution*. Working Papers IES 2013/11. Charles University Prague, Faculty of Social Sciences, Institute of Economic Studies.

Hazell, Jonathon et al. (Feb. 2022). “The Slope of the Phillips Curve: Evidence from U.S. States*”. In: *The Quarterly Journal of Economics* 137.3, pp. 1299–1344. ISSN: 0033-5533. eprint: https://academic.oup.com/qje/article-pdf/137/3/1299/44839862/qjac010.pdf.

Hooper, Peter, Frederic S Mishkin, and Amir Sufi (2019). *Prospects for Inflation in a High Pressure Economy: Is the Phillips Curve Dead or is It Just Hibernating?* Tech. rep. National Bureau of Economic Research.

Imbs, Jean, Eric Jondeau, and Florian Pelgrin (2011). “Sectoral Phillips curves and the aggregate Phillips curve”. In: *Journal of Monetary Economics* 58.4, pp. 328–344.

Jagannathan, Ravi and Yong Wang (2007). “Lazy investors, discretionary consumption, and the cross-section of stock returns”. In: *The Journal of Finance* 62.4, pp. 1623–1661.

Jorda, Oscar and Fernanda Nechio (2018). “Inflation Globally”. In: Federal Reserve Bank of San Francisco.

Kleibergen, Frank (2005). “Testing parameters in GMM without assuming that they are identified”. In: *Econometrica* 73.4, pp. 1103–1123.

Kleibergen, Frank and Sophocles Mavroeidis (2009). “Weak instrument robust tests in GMM and the new Keynesian Phillips curve”. In: *Journal of Business & Economic Statistics* 27.3, pp. 293–311.

Kösem-Alp, Sevim (2010). “Optimal Monetary Policy under Sectoral Heterogeneity in Inflation Persistence”. In: *Ancara: Central Bank of the Republic of Turkey*.

Kremens, Lukas and Ian Martin (Mar. 2019). “The Quanto Theory of Exchange Rates”. In: *American Economic Review* 109.3, pp. 810–43.

Kroencke, Tim A. (Feb. 2017). “Asset Pricing without Garbage”. In: *Journal of Finance* 72.1, pp. 47–98.

Kuester, Keith, Gernot J Müller, and Sarah Stölting (2009). “Is the New Keynesian Phillips Curve Flat?” In: *Economics Letters* 103.1, pp. 39–41.

Liao, Gordon Y. and Tony Zhang (May 2020). *The Hedging Channel of Exchange Rate Determination*. International Finance Discussion Papers 1283. Board of Governors of the Federal Reserve System (U.S.)

Lilley, Andrew and Gianluca Rinaldi (Working Paper). “Currency Betas and Interest Rate Spreads”. In: *2020*.

Luengo-Prado, María José, Nikhil Rao, and Viacheslav Sheremirov (2018). “Sectoral inflation and the Phillips curve: What has changed since the Great Recession?” In: *Economics Letters* 172, pp. 63–68.

Lustig, Hanno, Nikolai Roussanov, and Adrien Verdelhan (2011). “Common risk factors in currency markets”. In: *The Review of Financial Studies* 24.11, pp. 3731–3777.

Lustig, Hanno and Adrien Verdelhan (Mar. 2007a). “The Cross Section of Foreign Currency Risk Premia and Consumption Growth Risk”. In: *American Economic Review* 97.1, pp. 89–117.

— (2007b). “The cross-section of foreign currency risk premia and consumption growth risk”. In: *Journal of Political Economy* 115.3, pp. 279–318.

Ma, Adrian (2002). “GMM estimation of the new Phillips curve”. In: *Economics Letters* 76.3, pp. 411–417.

Maggiori, Matteo (Oct. 2017). “Financial Intermediation, International Risk Sharing, and Reserve Currencies”. In: *American Economic Review* 107.10, pp. 3038–71.

Martin, Ian (Nov. 2011). *The Forward Premium Puzzle in a Two-Country World*. Working Paper 17564. National Bureau of Economic Research.

— (2013). “THE LUCAS ORCHARD”. In: *Econometrica* 81.1, pp. 55–111. ISSN: 00129682, 14680262.

Mavroeidis, Sophocles, Mikkel Plagborg-Møller, and James H. Stock (Mar. 2014). “Empirical Evidence on Inflation Expectations in the New Keynesian Phillips Curve”. In: 52.

Mazumder, Sandeep (2010). “The new Keynesian Phillips curve and the cyclicality of marginal cost”. In: *Journal of Macroeconomics* 32.3, pp. 747–765.

McLeay, Michael and Silvana Tenreyro (2019). *Optimal inflation and the identification of the Phillips curve*. Tech. rep. National Bureau of Economic Research.

Mehra, Rajnish and Edward C Prescott (1985). “The equity premium: A puzzle”. In: *Journal of monetary Economics* 15.2, pp. 145–161.

Menkhoff, Lukas et al. (2012). “Carry trades and global foreign exchange volatility”. In: *Journal of Finance* 67.2, pp. 681–718.

Monacelli, Tommaso (2007). “Comment on" Globalization and Inflation Dynamics: A Structural Perspective"”. In: *International Dimensions of Monetary Policy*. University of Chicago Press, pp. 579–590.

Moreira, Marcelo J (2003). “A conditional likelihood ratio test for structural models”. In: *Econometrica* 71.4, pp. 1027–1048.

Murray, Michael P (2006). “Avoiding invalid instruments and coping with weak instruments”. In: *Journal of economic Perspectives* 20.4, pp. 111–132.

Nakamura, Emi and Jón Steinsson (2008). “Five facts about prices: A reevaluation of menu cost models”. In: *The Quarterly Journal of Economics* 123.4, pp. 1415–1464.

— (Jan. 2013). *Price Rigidity: Microeconomic Evidence and Macroeconomic Implications*. Working Paper 18705. National Bureau of Economic Research.

Nason, James M and Gregor W Smith (2008a). “Identifying the new Keynesian Phillips curve”. In: *Journal of Applied Econometrics* 23.5, pp. 525–551.

— (2008b). “The new Keynesian Phillips curve: Lessons from single-equation econometric estimation”. In: *FRB Richmond Economic Quarterly* 94.4, pp. 361–395.

Newey, Whitney K and Richard J Smith (2004). “Higher order properties of GMM and generalized empirical likelihood estimators”. In: *Econometrica* 72.1, pp. 219–255.

Newey, Whitney K and Frank Windmeijer (2009). “Generalized method of moments with many weak moment conditions”. In: *Econometrica* 77.3, pp. 687–719.

Pavlova, Anna, Magnus Dahlquist, et al. (May 2022). *International Capital Markets and Wealth Transfers*. CEPR Discussion Papers 17334. C.E.P.R. Discussion Papers.

Pavlova, Anna and Roberto Rigobon (2007). “Asset Prices and Exchange Rates”. In: *The Review of Financial Studies* 20.4, pp. 1139–1181. ISSN: 08939454, 14657368.

Pencavel, John (1986). “Labor supply of men: a survey”. In: *Handbook of labor economics* 1, pp. 3–102.

Piazza, Mr Roberto (2018). *Is There a Phillips Curve? A Full Information Partial Equilibrium Approach*. International Monetary Fund.

Rotemberg, Julio J and Michael Woodford (1997). “An optimization-based econometric framework for the evaluation of monetary policy”. In: *NBER macroeconomics annual* 12, pp. 297–346.

Rudd, Jeremy and Karl Whelan (2005). “Does labor’s share drive inflation?” In: *Journal of Money, Credit, and Banking* 37.2, pp. 297–312.

Rudd, Jeremy and Karl Whelan (2007). *Modeling inflation dynamics: a critical review of recent research*. Open Access publications. School of Economics, University College Dublin.

Sauzet, Maxime (Feb. 2023). *Asset Prices, Global Portfolios, and the International Financial System*. https://ssrn.com/abstract=3882255. Proceedings of the EUROFIDAI-ESSEC Paris December Finance Meeting 2023.

Savov, Alexi (2011). "Asset Pricing with Garbage". In: *Journal of Finance* 66.1, pp. 177–201.

Sbordone, Argia M (2002). "Prices and unit labor costs: a new test of price stickiness". In: *Journal of Monetary economics* 49.2, pp. 265–292.

Schorfheide, Frank, Dongho Song, and Amir Yaron (2018). "Identifying Long-Run Risks: A Bayesian Mixed-Frequency Approach". In: *Econometrica* 86.2, pp. 617–654. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.3982/ECTA14308.

Steinsson, Jón (2003). "Optimal monetary policy in an economy with inflation persistence". In: *Journal of Monetary Economics* 50.7, pp. 1425–1456.

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

Stock, James H and Jonathan H Wright (2000). "GMM with weak identification". In: *Econometrica* 68.5, pp. 1055–1096.

Stock, James H, Jonathan H Wright, and Motohiro Yogo (2002). "A survey of weak instruments and weak identification in generalized method of moments". In: *Journal of Business & Economic Statistics* 20.4, pp. 518–529.

Stock, James H and Motohiro Yogo (2002). *Testing for weak instruments in linear IV regression*.

Vermeulen, Philip et al. (2007). "Price setting in the euro area: some stylised facts from individual producer price data". In: *National Bank of Belgium Working Paper* 111.

Vissing-Jorgensen, Annette (1998). "Limited stock market participation and the equity premium puzzle". In.

— (2002). *Limited Asset Market Participation and the Elasticity of Intertemporal Substitution*. NBER Working Papers 8896. National Bureau of Economic Research, Inc.

Wilcox, David W (1992). “The construction of US consumption data: some facts and their implications for empirical work”. In: *The American economic review*, pp. 922–941.

Woodford, M (Dec. 2003). “Interest and Prices: Foundation of a Theory of Monetary Policy”. In.

Woodford, Michael (2005). *Firm-specific capital and the New-Keynesian Phillips curve*. Tech. rep. National Bureau of Economic Research.

Working, Holbrook (1960). “Note on the correlation of first differences of averages in a random chain”. In: *Econometrica: Journal of the Econometric Society*, pp. 916–918.

Yogo, Motohiro (Aug. 2004). “Estimating the Elasticity of Intertemporal Substitution When Instruments Are Weak”. In: *The Review of Economics and Statistics* 86.3, pp. 797–810.

Zaniboni, Nicola (2008). “Globalization and the phillips curve”. In: *V Unpublished manuscript*.

Zou, Hui (Feb. 2006). “The adaptive LASSO ad its oracle properties”. In: *Journal of the American Statistical Association* 101, pp. 1418–1429.