Essays in International Macro-Finance and Asset Pricing

Michael Wulfsohn

New College

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy

Trinity 2025

Acknowledgements

I am sincerely grateful to my primary supervisor, Andrea Ferrero, for his guidance throughout every stage of my thesis, and to my second supervisor, Seung Joo Lee, for his essential advice that shaped my third chapter.

My appreciation extends to the Macroeconomics faculty and the faculty of Saïd Business School at Oxford for providing an intellectually challenging yet supportive environment in which this research could develop. Their feedback, through seminars, conversations, and written comments, has immeasurably strengthened my work. I am similarly indebted to various researchers outside of Oxford who have offered valuable feedback on my work.

Among my peers, I am especially grateful to Jonas Kurle, David Murakami, Lukas Boehnert, Alexander Haas, Matthew McKernan, and Momo Komatsu. I have greatly appreciated their camaraderie and willingness to discuss ideas and opinions over the years.

I have been lucky to stay in regular contact with friends in my home country, particularly Josh Centner and Tristan Van de Vreede, who have kept me grounded throughout.

Finally, I owe more than words can convey to my family. My partner Gyo has shown unwavering patience and support. My parents Vivienne and Sidney, sister Lisa, step-parents Linton and Kellie, and other siblings Emily, Meg, Ben, Jess, and Josh have provided constant encouragement. I also remember my late grandparents Len, Isa, Aubrey, and Dora, whose examples continue to inspire me.

Abstract

This thesis consists of three self-contained chapters, the first two of which address topics in international macro-finance, and the last of which turns to asset pricing.

Chapter 1 develops a two-country model with liquidity constraints to study why a decline in one country's government bond liquidity (or convenience) yield relative to another is typically accompanied by depreciation of the first country's currency. A shock that makes a country's bond less saleable triggers substitution across investment types, lowers total investment, and generates a trade surplus, weakening the exchange rate. The same shock simultaneously drives down the bond's liquidity yield. The mechanism does not depend on any privileged status for U.S. assets.

Chapter 2 analyses loan-to-value ratio macroprudential policy in a two-country setting inspired by the "global saving glut" in the early 2000s, using Germany and Spain as an illustrative pair. The paper considers three policy regimes: internationally cooperative optimisation, a non-cooperative optimisation by national policymakers, and a benchmark of no policy response. The paper compares the policy regimes' ex-ante welfare implications. While international cooperation predictably dominates, the study finds that the non-cooperative regime can, through competitive policy distortions, deliver even lower welfare than doing nothing. Sensitivity checks identify parameter ranges in which this ranking reverses.

Chapter 3 addresses cryptocurrency valuation, particularly the observed rapid price growth. It proposes a model in which investors are uncertain about eventual adoption demand and update their beliefs based on noisy signals over time. Early uncertainty is met with heavily discounted prices. Subsequently, gradual learning raises the price level and lowers the variance of price growth, matching key empirical patterns.

Contents

List of Figures

List of Tables

1

Explaining the link between government bond liquidity yields and exchange rates

Abstract When the government bond liquidity yield (or convenience yield) in one country falls relative to another, the first country's currency also tends to depreciate. Using a two-country model with liquidity constraints, I explore the causal mechanisms behind this link. An adverse shock to the saleable portion of one country's bond forces substitution between investment types, skewing their mix. That country reduces investment and runs a trade surplus, causing exchange rate depreciation. Concurrently, given the bond's reduced liquidity, its liquidity yield reduces. In keeping with the link's observed ubiquity among country pairs, the model does not rely on any special role for US assets.

1.1 Introduction

Recent literature has documented the relative government bond liquidity yield's relationship with exchange rates.[1] The liquidity yield, often called the convenience yield, is a measure of non-pecuniary benefits from holding a bond. In this paper it is calculated as the excess of the swap rate over the government bond yield. A country's relative liquidity yield is then the excess of its yield over another country's. This paper contributes to an emerging literature that attempts to explain this phenomenon and explore its implications. Some insights have been gleaned using models that include bonds in the utility function.[2] However, these reduced form models are not designed to explore the mechanisms behind the effects. Papers that do seek to model liquidity mechanisms often microfound the demand for liquid government bonds. However, to generate relative liquidity yield fluctuations with demand shocks, the bonds must not be perfectly substitutable. These papers therefore often assume a special role for US bonds.[3] Given that the same link exists between country pairs that exclude the US,[4] such explanations are incomplete. This paper completes the story by establishing a microfoundation for bond liquidity that does not rely on a special US role.

The importance of this contribution stems from the literature's difficulty in identifying the drivers of currency exchange rates. Exchange rates are among the most important relative prices for many countries, and are central to international trade and capital flows. By establishing that relative liquidity yields are linked to exchange rates, the literature has made progress on a key facet of the exchange rate disconnect puzzle.[5] The next step, which this paper takes, is to explain the

[1] Jiang et al. (2021) and Engel and Wu (2023)

[2] Engel (2016), Valchev (2020), Engel and Wu (2023), Kekre and Lenel (2024), Jiang et al. (2024a), and Jiang et al. (2024b).

[3] Bianchi et al. (2023) and Devereux et al. (2023).

[4] Engel and Wu (2023)

[5] The puzzle relevant to this paper is the one associated with Meese and Rogoff (1983), who showed that exchange rate models' out of sample performance was no better than a random walk.

empirically observed link, including the mechanisms and causal relationships.

The empirics are summarised in Figure 1.1. Using the Engel and Wu (2023, hereafter E&W) data, and a pairwise implementation of their panel regression specification, I plot the orthogonalised changes in the exchange rate against those of the relative liquidity yield. The robustly negative slopes indicate that, when a country's relative liquidity yield decreases, its currency depreciates. As implied by the figure, and tested more extensively in E&W, the relationship is not exclusive to currency pairs that include the US. It is also not exclusive to liquidity contractions, or periods of crisis.

Fundamentally, the reason liquid assets fetch a premium is because they alleviate liquidity frictions. To study the impact of that mechanism, I use a two-country model in which a certain type of investment (Constrained investment) must be financed by selling government bonds. The most direct motivation for this microfoundation is the role of firms' cash on hand in financing some of their investment. However, the constraint also captures the use of liquid assets by banks in channelling capital to its most productive use.[6]

The paper studies shocks to a country's government bonds' saleability. While government bonds are highly liquid, they are not perfectly so. The shocks capture changes in over-the-counter trading efficiency, leading to shifts in market breadth and depth, and buyer willingness to negotiate. Fluctuations may be large during extremely adverse conditions, such as during financial crises. The notion that government bond liquidity fluctuates, and that different countries' bonds fluctuate differently, is supported by the observed relative liquidity yields. The use of bond-specific shocks is central to the paper's contribution, as they generate relative liquidity yield fluctuations without requiring country-specific assumptions such as a special US role.

[6]For example, the liquidity constraints used are isomorphic to collateral constraints, allowing an alternative interpretation of the model and its frictions.

Figure 1.1: Relationship between relative liquidity yield η_t and exchange rate $\mathcal{E}_t$

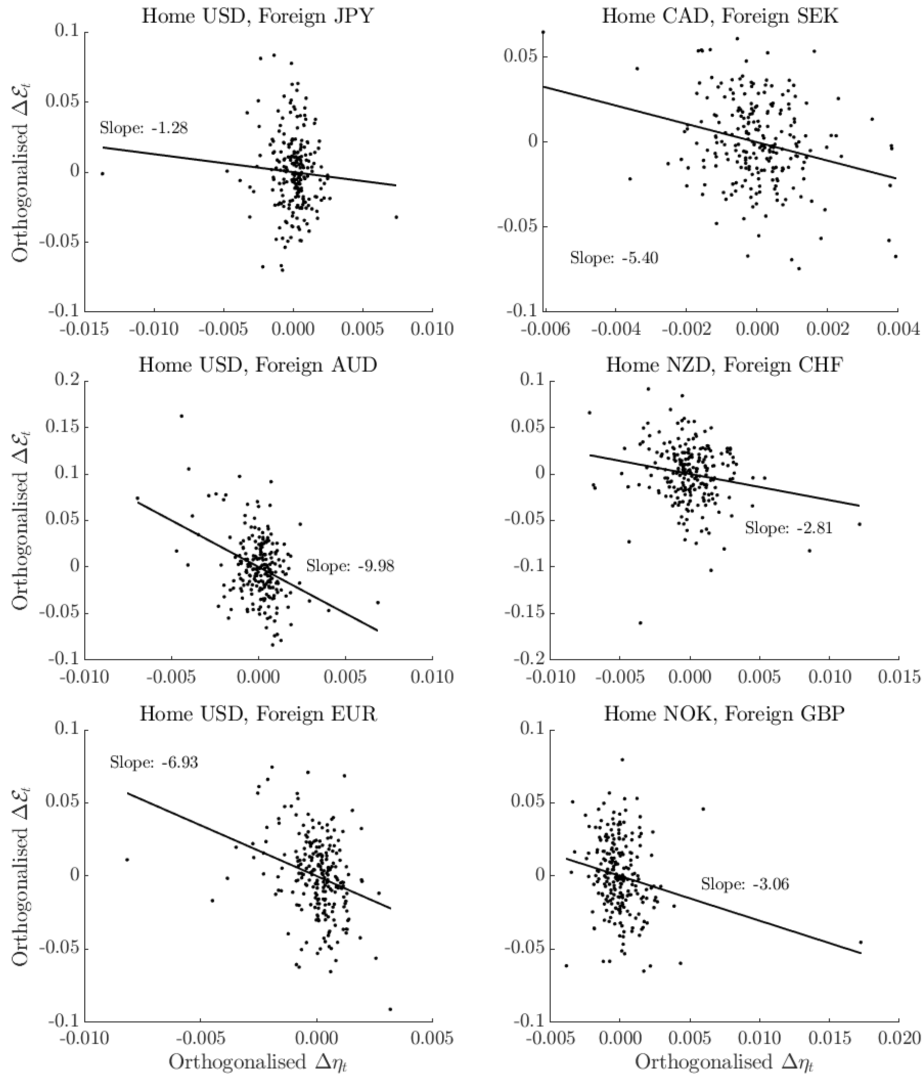

As per the Frisch-Waugh-Lovell theorem (Frisch and Waugh, 1933; Lovell, 1963), the slope of the line of best fit in the plots is the OLS estimate of β_2 in a pairwise implementation of the panel regression specification used by E&W, $\Delta\hat{\mathcal{E}}_t = \alpha + \beta_1 S_{t-1} + \beta_2 \Delta\eta_t + \beta_3 \Delta i_t^R + \beta_4 \eta_{t-1} + \beta_5 i_{t-1}^R + u_t$. $\mathcal{E}_t$ is the nominal exchange rate defined as the Home currency price of Foreign currency, S_t is the real exchange rate, η_t is the relative liquidity yield on government bonds, and i_t^R is the Home minus Foreign government bond interest rate. The variable on the horizontal axes is the residual from the projection of $\Delta\eta_t$ on the other regressors. The variable on the vertical axes is the residual from the projection of $\Delta\hat{\mathcal{E}}_t$ on regressors other than $\Delta\eta_t$. The data is that used by E&W and includes monthly periods from January 1999 to December 2017.

In the model, asymmetric liquidity shocks affect the exchange rate by the following mechanism. When a country's bond experiences a negative shock to its saleable portion, the type of investment facilitated by that bond's liquidity is forced to contract, skewing the mix of investment. As long as both countries hold some of each other's government bonds, both will experience a deterioration in investment efficiency. However, with home bias in bond holdings, the country whose bond was shocked is the most affected. For that country, domestic investment becomes less attractive than running a trade surplus and building up net foreign assets. The overseas country that receives the corresponding capital inflow responds by increasing investment, such that investment spending effectively shifts overseas. Combined with home bias in household preferences, the shift depreciates the shocked country's real exchange rate, with the nominal exchange rate following suit.

The strength of the exchange rate response depends on the extent of home bias of bond holdings. Without home bias, both countries are equally affected, and both are equally keen to reduce investment and run a trade surplus, so these forces offset. However, the stronger the home bias, the larger the difference between the countries' exposure to a particular bond, and the greater the investment reallocation. Thus the depreciation is experienced by the country that is most exposed to the adversely shocked bond.

Importantly, contrary to what intuition may suggest, the exchange rate depreciation is not driven by sales of the adversely shocked bond. In particular, an adversely shocked bond becomes less attractive to all potential holders, so there is no obvious direction of effect on any country's bond portfolio composition. Indeed, in a symmetric calibration of the model, there is no such composition effect. The exchange rate depreciation is instead driven by the reallocation of investment overseas.

Of interest is whether the liquidity yield on the shocked bond reduces at the same time.[7] In this regard, the direction of effect is not obvious, as there are two counteracting effects. The first is a direct effect, since the liquidity yield of a bond is the price of the liquidity that it offers. All else equal, a less liquid bond commands

[7] However, this is not required for the negative regression coefficient on the change in the relative liquidity yield. The regression also controls for relative interest rate changes, and lagged level variables.

a smaller liquidity yield. The second is an indirect effect, operating through the overall supply of liquidity. A negative shock to a bond's liquidity makes liquidity scarcer overall, and the liquidity constraint binds more tightly. All else equal, this raises the liquidity yield, counteracting the direct effect. This indirect mechanism is supported by the negative relationship between convenience yields and the supply of government bonds (Krishnamurthy and Vissing-Jorgensen, 2012), as long as more government bond supply creates more liquidity. In the model, a country's liquidity scarcity effect is measurable by the change in its liquidity constraint multiplier.

Since the relative strength of these two effects is not obvious, it is necessary to allow data to discipline it. The paper draws upon a measurable outcome from liquidity scarcity—the sensitivity of total investment to liquidity. If investment moves little in response to a shock to bond liquidity, then the household has little reason to worry if the constraint is tightened. That means the constraint multiplier also moves little, and the liquidity scarcity effect is weak.

The sensitivity of investment to liquidity in the model is targeted by calibrating the elasticity of substitution between the capital whose investment must be financed by selling government bonds (Constrained capital), and all other capital (Unconstrained capital). Specifically, the model assumes that, each period, the household of each country splits into two types of members – Constrained and Unconstrained. The two capital types correspond to these member types; that is, Constrained members may invest in Constrained capital, and Unconstrained members in Unconstrained capital. The liquidity constraint requires that Constrained members' investment must be financed by selling bonds to Unconstrained members. The greater the ease of substitution between the two capital types, the more easily the liquidity constraint can be circumvented, and the weaker will be the impact of a liquidity shock on the constraint multiplier and thus liquidity scarcity.

I calibrate this elasticity so that the model matches the firm-level marginal effect of cash-on-hand on investment in the US. Specifically, I use empirical evidence from Denis and Sibilkov (2010) to measure the sensitivity of capital expenditure to firms' cash reserves. This indirect approach to calibration avoids the need to measure real-world Constrained and Unconstrained investment, or to rely on proxies. In addition, the strength of the depreciation response is also moderated by the calibrated elasticity, as the more investment reduces from a liquidity shock, the greater the need to reallocate investment overseas. Under the calibrated elasticity, the liquidity scarcity effect is quite small. Thus an adverse shock tends to decrease the bond's liquidity yield, and the relative liquidity yield.

The resulting model successfully replicates the observed link. I quantitatively explore two country pair calibrations—US-Japan and US-Australia—and show that the simulated model output produces regression results similar to E&W.

Turning to the literature, notable progress has been made on the empirical side of the exchange rate disconnect puzzle since the Meese and Rogoff (1983) paper. Aside from the motivating evidence in Figure 1.1, a different paper by Engel and Wu (2024) shows that macroeconomic models are in fact capable of outperforming a random walk when recent data, rather than data from the 1970s, is used. They attribute the change to the increased effectiveness of monetary policy in more recent decades. One of the variables they show to have explanatory power is the change in the relative liquidity yield.

The present paper contributes to the literature seeking to explain exchange rate movements. In particular, it is part of the subset studying the explanatory power of financial frictions. For example, Gabaix and Maggiori (2015) and Itskhoki and Mukhin (2021) show that financial shocks or frictions can cause large exchange rate movements without significant impact on the real economy. As a natural consequence, much exchange rate movement is not readily explainable with real variables, hence the historically observed disconnect.

Within the financial frictions subset of the literature, this paper belongs in the subset which uses financial frictions to explain exchange rates' link to government bond liquidity yields. Many such studies include bonds in the utility function, interpreting the effect as liquidity services, and generating a positive government bond liquidity yield. For example, E&W use such a model to motivate their regression specification. They in turn draw upon Engel (2016), who uses a similar structure to reconcile the contradictory predictions of the effect of risk premia on exchange rates. Similarly, Valchev (2020) shows that bond supply can explain the term structure of Uncovered Interest Rate Parity (UIRP) deviations. Jiang et al. (2024b) combine the approach with incomplete markets and a special role for US Treasury bonds to rationalise the exchange rate volatility puzzle, in addition to the disconnect from macroeconomic fundamentals. Kekre and Lenel (2024) model a time-varying demand for safe US-dollar bonds, which also leads to the US dollar appreciating when the US bond convenience yield increases. Jiang et al. (2024a) assume that there are non-pecuniary benefits to holding safe dollar bonds such as US Treasuries, and show that this can explain a number of observations surrounding the US dollar and exorbitant privilege, including dollar appreciation when the convenience yield on US Treasuries increases. However, models with bonds in the utility function stop short of engaging with the reasons why any non-pecuniary benefits are valued. They are not well positioned to investigate the liquidity mechanisms that are the focus of this paper.

To endogenise the liquidity benefits of government bonds, a typical approach is to create a mechanism by which demand for liquidity (or safety) may fluctuate, but the liquidity of government bonds is assumed fixed, or perfect. Deploying this to explain the impact of liquidity yields on exchange rates, Devereux et al. (2023) build a two-country model in which US government bonds serve as superior collateral on banks' balance sheets. In a crisis, the increase in the convenience yield on US bonds coincides with a repatriation of capital and an appreciation of the US dollar. In Bianchi et al. (2023), in response to dollar settlement frictions, banks endogenously choose their ratio of liquid assets to short-term liabilities.

The “scramble for dollars" in a crisis again appreciates the exchange rate as the dollar convenience yield increases.

The papers that microfound liquidity demand thus assume that US government bonds are the sole provider of liquidity, or safety. This asymmetry is necessary because, with only liquidity demand differing between countries, all fully liquid bonds are perfect substitutes. The bonds will have the same liquidity yield, leading to a constant relative liquidity yield and an inability to study its connection to exchange rates. Many of the bonds-in-the-utility papers also assume a special US role.[8] Such an assumption can yield useful insights due to the US dollar’s unique role in the global economy as a reserve asset and vehicle currency (Gopinath and Stein, 2021).

However, explanations reliant on a special US role are incomplete. E&W robustly document the same link between exchange rates and relative government bond liquidity for country pairs that exclude the US. To be clear, this does not preclude a special role for US government bonds. It does, however, show that the non-pecuniary benefits offered by other countries’ government bonds are sufficient to generate a similar effect. Therefore, assuming a special role for US assets risks obscuring the true nature of the link. However, to the author’s knowledge, no existing paper microfounds the link without a special US role.

The contribution of the present paper is to microfound the link as it exists between a wide range of developed market currency pairs. Rather than focussing on demand for liquidity, I shock the degree of liquidity supplied by individual government bonds.

1.2 Model

In this section I present a dynamic equilibrium model in which there are two types of capital, and where purchases of one type are subject to a liquidity constraint.

[8] Papers already mentioned that assume a special role for US government bonds or the US dollar include Bianchi et al. (2023), Devereux et al. (2023), Jiang et al. (2024a), Jiang et al. (2024b), and Kekre and Lenel (2024).

The model uses a two-country, infinite horizon discrete time setting. Each country consists of a household, a firm, and a government. There is a unit measure of household members worldwide, with $n \in (0, 1)$ members residing in the Home country and $1 - n$ members in the Foreign country. In keeping with the scope of the empirical evidence motivating this study, the modelled countries may correspond to any two developed countries with floating exchange rates.

1.2.1 Notation

The equations presented are for the Home country except where otherwise stated. Foreign variables are denoted by an asterisk. Variables without time (t) subscripts are the steady state value of the corresponding variable. Real quantity variables are in per capita units of the consumption good; this includes consumption, investment, capital, output, bond holdings, government spending, profits, and lump sum tax. A hat over a variable denotes that it is expressed as a percentage deviation from steady state, except in the case of (relative) interest rates and (relative) liquidity yields, in which case a hat denotes percentage point deviations from steady state.

1.2.2 Households

Given the research question of this paper, we require a model in which the liquidity of government bonds affects economic activity. It is therefore natural to have trading in the bonds occur each period. The model motivates the trading by having household members split into two groups each period, trade with each other, then recombine at the end of the period, in the manner of Del Negro et al. (2017).

At the start of each period, before the split, the household holds Home and Foreign government-issued bonds B_{t-1}^{H} and B_{t-1}^{F}, the market bond of their country B_{t-1}^{m}, and Constrained and Unconstrained capital K_{t-1}^{C} and K_{t-1}^{U} which are also domestic to their country. The split randomly assigns members to the Constrained group with probability χ or the Unconstrained group with probability $1-\chi$. All assets are divided evenly between members. Constrained members are able to invest in Constrained capital, and Unconstrained members to invest in Unconstrained capital.

Only Constrained investment is subject to the liquidity constraint. The constraint limits the investment to an exogenously determined saleable portion of each bond held and its real interest earnings, along the lines of Del Negro et al. (2017) and Kiyotaki and Moore (2019). The portion is ϕ_t for Home government bonds and ϕ_t^* for Foreign government bonds. Since the intra-period trading happens between Constrained and Unconstrained members of the same household, it does not affect the household's total holding of either bond.

Thus Constrained investment is defined as investment that is subject to liquidity or financing constraints. The portion of such investment is likely substantial. Doms and Dunne (1998) show that, unlike aggregate national investment, firm-level investment is typically irregular and "lumpy". Unpredictable investment is unlikely to be financeable by revenue or cashflow. Where cash holdings or external finance are needed, the investment becomes exposed to liquidity or financing constraints.

The boundary between Constrained and Unconstrained capital need not be observable. As will be shown in Section 2.3, the model can be calibrated without adopting a real-world approximation of Constrained and Unconstrained capital. However, for the purpose of fixing ideas, examples may include growth and maintenance capital, or tangible and intangible capital, or human and physical capital. However, these do not fully align with the financing-based distinction of interest. For example, growth capital expenditure, such as the construction of new buildings, factories, and machinery, is indeed more likely to be irregular and lumpy at the firm level, and require cash reserve draw-down or external financing. However, some growth capital expenditure may instead be regular and predictable. Likewise, maintenance capital expenditure, such as repairs, renovations, and refurbishment, is more likely to be financed by revenues and cashflow. However, some maintenance expenditure may instead be unpredictable, for example repairs following an accident. Since no such approximation fully aligns with the financing-based distinction of interest, the paper's non-reliance on them is a strength.

An alternative interpretation of the model's structure is that the two capital types are distinguished not by their financing, but by the type of firm within which they are deployed. Under this interpretation there would be two intermediate goods firms, one of which would be financially constrained, for example due to a poor credit rating or small size. Their output would then be combined into a final consumption good. The model's specifications are isomorphic to such a structure.

Lastly, the model can accommodate multiple types of capital (or intermediate goods firms), as long as they can be aggregated into a single Constrained and a single Unconstrained capital stock.

Once members have made their investment decisions I_t^C and I_t^U, and completed their intra-period trading of government bonds, production takes place. The household then recombines its assets, and consumes C_t units of an aggregate consumption good. There is perfect consumption sharing within a household. The household may also adjust their holdings of government bonds at the end of the period by trading with the household of the other country, without intra-period liquidity constraints applying.

Household members own the firm domestic to their country. They earn profits Ξ_t and rent payments at rates r_t^C and r_t^U for capital of each type that they supply to the firm. They pay lump sum tax τ_t.

The optimisation problem of the household is as follows.

$$\max_{<C_t, K_t^C, K_t^U, B_t^m, B_t^H, B_t^F>} \quad \mathbb{E}_t \sum_{s=t}^{\infty} \beta^{s-t} \left(\frac{C_s^{1-\sigma}}{1-\sigma} \right)$$

subject to

$$I_t^C \leq \chi \left(\phi_t B_{t-1}^H \frac{(1+i_{t-1})}{\Pi_t} + \phi_t^* S_t B_{t-1}^F \frac{(1+i_{t-1}^*)}{\Pi_t^*} \right) \tag{1.1}$$

$$C_t + B_t^H + S_t B_t^F + B_t^m + I_t^C + I_t^U \leq B_{t-1}^H \frac{(1+i_{t-1})}{\Pi_t} + S_t B_{t-1}^F \frac{(1+i_{t-1}^*)}{\Pi_t^*} + B_{t-1}^m \frac{(1+i_{t-1}^m)}{\Pi_t} + r_t^C K_{t-1}^C + r_t^U K_{t-1}^U + \Xi_t - \tau_t \tag{1.2}$$

$$K_t^C = I_t^C + (1-\delta) K_{t-1}^C \tag{1.3}$$

$$K_t^U = I_t^U + (1-\delta) K_{t-1}^U \tag{1.4}$$

Equation 1.1 is the liquidity constraint governing intra-period trading. Equation 1.2 is a household-level budget constraint governing beginning- and end-period activity.

Variables not already described include consumption good price inflation $\Pi_t = \frac{P_t}{P_{t-1}}$, utility discount factor β, inverse intertemporal elasticity of substitution σ, capital depreciation rate δ, and real exchange rate S_t of which an increase corresponds to a depreciation of the Home currency.

I study equilibria in which the liquidity constraint binds. The positivity of observed convenience yields supports this assumption, since, in the model, the liquidity yield is only positive if the liquidity constraint binds or is close to binding.

The liquidity constraint is isomorphic to a collateralised borrowing constraint. Under this alternative, instead of selling the bonds, Constrained members post the bonds as collateral for intra-period borrowing from Unconstrained members. The borrowing finances Constrained investment. The loan positions are closed out when the household recombines at the end of the period. The convenience yield on the bonds would then be due to their facility as collateral. Such collateral value can be thought of as arising from both the liquidity and safety components of the convenience yield on government bonds. Krishnamurthy and Vissing-Jorgensen (2012) provide evidence for both components, decomposing the US Treasury bond convenience yield into liquidity and safety components.

ϕ_t and ϕ_t^* are subject to shocks with persistence ρ_ϕ, which affect the portion of the holdings of a particular government bond that can be used to finance Constrained investment. The model abstracts beyond the causes of these shocks, which represent fluctuations in the efficiency of the Over-The-Counter (OTC) markets in which government bonds are typically traded. In general, at any given moment, the amount of money buyers are willing to spend on purchasing government bonds close to the prevailing price (market breadth) or at lower prices (market depth) may be greater or lesser. The availability of buyers will affect the degree to which the sale of a given amount of government bonds will affect the market price. Further, the sale of government bonds relies on negotiation with buyers. During periods of stress, when trust among market participants has been eroded, sellers have less ability to generate buyer interest and bring negotiations to fruition, particularly for large transactions. The speed of transactions—another dimension of liquidity—may also be affected under stressed circumstances. Finally, the term liquidity can refer to the resilience of the market price, or how quickly the market price returns to a "fair" or equilibrium price following a large market-moving transaction (Sarr and Lybek, 2002).

The multiplier on the liquidity constraint, ρ_t, is presented in this paper in per-capita units of the consumption good (by rescaling using the marginal utility of consumption). It is the amount the household is willing to pay to loosen the liquidity constraint by one unit. ρ_t is a measure of liquidity scarcity in the country and influences the liquidity yield.

The household's first-order conditions are as follows.

$$\text{Unconstrained capital } K_t^U \qquad 1 = \mathbb{E}_t \beta \left(\frac{C_{t+1}}{C_t}\right)^{-\sigma} (1 - \delta + r_{t+1}^U) \tag{1.5}$$

$$\text{Constrained capital } K_t^C \qquad 1 = \mathbb{E}_t \beta \left(\frac{C_{t+1}}{C_t}\right)^{-\sigma} \frac{1+\rho_{t+1}}{1+\rho_t} \left(1 - \delta + \frac{r_{t+1}^C}{1+\rho_{t+1}}\right) \tag{1.6}$$

$$\text{Market bonds } B_t^m \quad \frac{1}{1+i_t^m} = \mathbb{E}_t \beta \left(\frac{C_{t+1}}{C_t}\right)^{-\sigma} \frac{1}{\Pi_{t+1}} \tag{1.7}$$

$$\text{Home gov't bonds } B_t^H \qquad \frac{1}{1+i_t} = \mathbb{E}_t \beta \left(\frac{C_{t+1}}{C_t}\right)^{-\sigma} \frac{1}{\Pi_{t+1}} (1 + \chi \rho_{t+1} \phi_{t+1}) \tag{1.8}$$

$$\text{Foreign gov't bonds } B_t^F \qquad \frac{1}{1+i_t^*} = \mathbb{E}_t \beta \left(\frac{C_{t+1}}{C_t}\right)^{-\sigma} \frac{1}{\Pi_{t+1}} \frac{\mathcal{E}_{t+1}}{\mathcal{E}_t} (1 + \chi \rho_{t+1} \phi_{t+1}^*) \tag{1.9}$$

The liquidity constraint multiplier ρ_t drives a wedge between the expected rates of return on the two types of capital. Since Constrained capital is generally less available, its expected return tends to be higher than that of Unconstrained capital.

Comparison of the market bond and home government bond equations 1.7 and 1.8 reveals that the government bond price is raised due to the $\chi \rho_{t+1} \phi_{t+1}$ term. Subtracting one from the other reveals the model's analogue of the liquidity yield studied by E&W.

$$\gamma_t := \frac{1}{1+i_t} - \frac{1}{1+i_t^m} = \mathbb{E}_t \beta \left(\frac{C_{t+1}}{C_t}\right)^{-\sigma} \frac{1}{\Pi_{t+1}} \chi \rho_{t+1} \phi_{t+1} \tag{1.10}$$

In the model, γ_t is defined in terms of bond prices, but to a first-order approximation, it is equal to the E&W definition of $\gamma_t = i_t^m - i_t$.

The liquidity yield γ_t is priced as the expectation of the stochastic discount factor multiplied by the value of liquidity in the next period. The value of liquidity is the product of the prevalence of Constrained household members χ who need liquidity, the scarcity of liquidity in the next period ρ_{t+1}, and the portion of the bond that can be sold during the next period ϕ_{t+1}.

The liquidity yield is the non-pecuniary benefits of holding the government bond. In the model, these exist because holdings of the bond relax the liquidity constraint. However, in practice the government bond liquidity yield may also be described as a convenience yield, in respect of the bond's safety and collateral utility as additional reasons for the lower yield. This corresponds to the aforementioned interpretation of the liquidity constraint as a collateralised borrowing constraint.

The concept of the liquidity yield used in the paper excludes credit or default risk, as the model does not include any default-prone debt. Its calculation of the liquidity yield thus stands in contrast to some other papers which measure the convenience yield as the AAA-rated corporate yield less the government yield, for example Cieslak et al. (2025).

The Home and Foreign government bond equations may be combined and approximated at first order to reveal the model's UIRP condition.

$$\mathbb{E}_t \Delta \hat{\mathcal{E}}_{t+1} = i_t^R + \eta_t \tag{1.11}$$

$$\hat{\mathcal{E}}_t = \mathbb{E}_t \sum_{s=t}^{\infty} \left(-i_s^R - \eta_s \right) + \lim_{T \to \infty} \mathbb{E}_t \hat{\mathcal{E}}_T \tag{1.12}$$

where $\eta_t = \gamma_t - \gamma_t^*$ and $i_t^R = i_t - i_t^*$ as defined in E&W, and where the nominal exchange rate is defined as[9]

$$\mathcal{E}_t = \frac{S_t P_t}{P_t^*} \tag{1.13}$$

As usual, expected Home currency depreciation is equal to the Home minus Foreign interest rate differential, plus a deviation term. In this model where incomplete markets and the liquidity constraints are the only frictions, the relative liquidity yield η_t is the UIRP deviation.

Households' consumption good is a constant elasticity of substitution aggregate of the Home and Foreign goods. The corresponding demand functions are standard, and show the relationships with real goods prices P_t^H and P_t^F, and the roles of home bias parameter ζ_C and elasticity parameter θ_C.

$$C_t = \left(\zeta_C^{\frac{1}{\theta_C}} C_t^{H\frac{\theta_C - 1}{\theta_C}} + (1 - \zeta_C)^{\frac{1}{\theta_C}} C_t^{F\frac{\theta_C - 1}{\theta_C}} \right)^{\frac{\theta_C}{\theta_C - 1}} \tag{1.14}$$

$$C_t^H = \zeta_C P_t^{H^{-\theta_C}} C_t \tag{1.15}$$

$$C_t^F = (1 - \zeta_C) P_t^{F^{-\theta_C}} C_t \tag{1.16}$$

ζ_C, which determines the consumption basket weights to Home and Foreign output, is set according to the size of the Home country, n, and an assumption

[9]Only the proportionate change of $\mathcal{E}_t$ is identified by the model, not its level.

about the degree of openness λ of consumer preferences[10]. The formula for Home country bias to Home goods is

$$\zeta_C = 1 - (1-n)\lambda \tag{1.17}$$

The formula for Foreign country bias to Foreign goods is

$$\zeta_C^* = 1 - n\lambda \tag{1.18}$$

1.2.3 Firms

The Home country has one firm, which takes aggregate capital K_{t-1} as an input and produces an amount Y_t of the Home good. The production function in equation 1.19 exhibits diminishing returns to scale in capital since $\alpha \in (0,1)$. The resulting first-order condition for capital is given in 1.20. Total factor productivity A is exogenous.

$$Y_t = AK_{t-1}^{\alpha} \tag{1.19}$$

$$r_t = \frac{\alpha P_t^H Y_t}{K_{t-1}} \tag{1.20}$$

Decreasing returns to scale leads to profits Ξ_t which are returned to the household. The model is isomorphic to one with fixed unit labour supply from the household and a constant returns to scale production function. The wage payments to labour would then take the place of profit payments to the owner.

$$\Xi_t = (1-\alpha)P_t^H Y_t \tag{1.21}$$

K_{t-1} is determined by aggregation of Constrained and Unconstrained capital provided by the Home household, according to the constant elasticity of substitution

[10]With perfect openness such that $\lambda = 1$, there is no home bias and purchasing power parity holds in the model. If, additionally, $P_t^H = P_t^F$, the consumption shares of Home and Foreign output are equal to the country sizes. The small open economy limit of the model can be obtained by allowing n to approach 0 from above (Home becomes the small open economy) or 1 from below (Foreign becomes the small open economy).

function in equation 1.22, and resulting in demand functions for the individual capital types in equations 1.23 and 1.24.

$$K_{t-1} = \left(\zeta_K^{\frac{1}{\theta_K}} K_{t-1}^{C\,\frac{\theta_K - 1}{\theta_K}} + (1-\zeta_K)^{\frac{1}{\theta_K}} K_{t-1}^{U\,\frac{\theta_K - 1}{\theta_K}} \right)^{\frac{\theta_K}{\theta_K - 1}} \tag{1.22}$$

$$K_{t-1}^C = \zeta_K \left(\frac{r_t^C}{r_t} \right)^{-\theta_K} K_{t-1} \tag{1.23}$$

$$K_{t-1}^U = (1-\zeta_K) \left(\frac{r_t^U}{r_t} \right)^{-\theta_K} K_{t-1} \tag{1.24}$$

Imperfect substitution between K_{t-1}^C and K_{t-1}^U is a key influence on the model's response to the liquidity shock. The binding liquidity constraint forces a reduction in I_t^C in response to an adverse shock. However, the household can partly substitute away from I_t^C to I_t^U, allowing overall investment $I_t = I_t^C + I_t^U$ to decrease by less than the forced decrease in I_t^C. Of course, the shock can also act in the opposite direction, in which case the above movements are reversed.

1.2.4 Government

The government spends a constant amount $\bar{G}$ each period, levies the lump sum tax τ_t on the household, and makes interest payments on and re-issues the government bond, according to the following budget constraint.

$$\bar{G} = \tau_t + B_t^H - B_{t-1}^H \frac{1 + i_{t-1}}{\Pi_t} + \frac{1-n}{n} \left(B_t^{H*} - B_{t-1}^{H*} \frac{1 + i_{t-1}}{\Pi_t} \right) \tag{1.25}$$

1.2.5 Monetary policy

Given that the relationship documented by E&W involves both nominal and real variables, the model must identify certain nominal variables. However, given the model's absence of nominal frictions, monetary policy is designed to be minimalist. It is set using the Taylor rule defined by

$$\hat{i}_t^m = \Gamma_\Pi \hat{\Pi}_t^H \tag{1.26}$$

where nominal domestic good price inflation Π_t^H is related to the real Home good price P_t^H as follows.

$$\Pi_t^H = \frac{P_t^H P_t}{P_{t-1}^H P_{t-1}} = \frac{P_t^H}{P_{t-1}^H}\Pi_t \tag{1.27}$$

The targeting of domestic inflation is in line with findings that optimal monetary policy in open economy models can often be achieved by targeting domestic variables (Galí and Monacelli, 2005; G. Benigno and P. Benigno, 2006).

The model exhibits money superneutrality. In each of the model equations where inflation appears (1.1, 1.2, 1.7, 1.8, 1.9, 1.10, and 1.25), the inflation rate serves only to transform a nominal interest rate into a real interest rate. Defining real counterparts of these interest rates as $1 + i_t^r = \frac{1+i_t}{\Pi_{t+1}}$ for the Home bond interest rate, and similarly for i_t^m, i_t^*, and i_t^{m*}, allows elimination of Π_t and Π_t^* from all model equations except the monetary policy equations (Equation 1.26) for each country. The monetary policy equations can then be understood as deriving inflation rates that are consistent with the real variables' equilibrium values.

Therefore, a nominal frictions extension would be required for analysis of monetary policy implications. However, the model is fit for the present purpose of studying the mechanisms behind the relationship documented by E&W.

1.2.6 Equilibrium

The model's market clearing conditions consist of a goods market clearing equation (1.28), the specification in Equation 1.29 that government bonds are in constant real supply, the specification of zero net supply of the market bond in Equation 1.30, and the law of one price (1.31, 1.32).

$$Y_t = \frac{C_t^H}{C_t}(C_t + I_t^C + I_t^U + \bar{G}) + \frac{1-n}{n}\frac{C_t^{H*}}{C_t^*}(C_t^* + I_t^{C*} + I_t^{U*} + \bar{G}^*) \tag{1.28}$$

$$\overline{B^H} = B_t^H + \frac{1-n}{n}B_t^{H*} \tag{1.29}$$

$$B_t^m = 0 \tag{1.30}$$

$$P_t^H = S_t P_t^{H*} \tag{1.31}$$

$$P_t^F = S_t P_t^{F*} \tag{1.32}$$

The market clearing conditions reflect that consumption goods and government bonds are internationally tradeable, which is in contrast to capital and the market bond.

The below variable definitions complete the model. They define the trade balance (1.33), the current account (1.34), net foreign assets (1.35), the terms of trade (1.36), and the relative liquidity yield (1.37).

$$TB_t = P_t^H Y_t - C_t - I_t^C - I_t^U - \bar{G} \tag{1.33}$$

$$CA_t = TB_t + B_{t-1}^F \left(S_t \frac{1+i_{t-1}^*}{\Pi_t^*} - S_{t-1} \right) - \frac{1-n}{n} B_{t-1}^{H*} \left(\frac{1+i_{t-1}}{\Pi_t} - 1 \right) \tag{1.34}$$

$$\Delta NFA_t = CA_t = \Delta \left(S_t B_t^F \right) - \frac{1-n}{n} \Delta B_t^{H*} \tag{1.35}$$

$$TOT_t = \frac{P_t^F}{P_t^H} \tag{1.36}$$

$$\eta_t = \gamma_t - \gamma_t^* \tag{1.37}$$

The model's dynamic equilibrium is a set of policy rules for the 57 endogenous variables, which are C_t, C_t^H, C_t^F, K_t^C, K_t^U, K_t, I_t^C, I_t^U, Y_t, B_t^H, B_t^F, B_t^m, P_t^H, P_t^F, Π_t, Π_t^H, i_t, i_t^m, r_t, r_t^C, r_t^U, ρ_t, γ_t, τ_t, and Ξ_t, all repeated for the Foreign country, and cross-country variables η_t, S_t, $\mathcal{E}_t$, TB_t, CA_t, NFA_t, and TOT_t. The policy rules take the state variables K_{t-1}, K_{t-1}^C, K_{t-1}^U, K_{t-1}^*, K_{t-1}^{C*}, K_{t-1}^{U*}, B_{t-1}^H, B_{t-1}^F, B_{t-1}^m, B_{t-1}^{H*}, B_{t-1}^{F*}, B_{t-1}^{m*}, i_{t-1}, i_{t-1}^*, i_{t-1}^m, i_{t-1}^{m*}, P_{t-1}^H, and P_{t-1}^{F*} and the shocks ϕ_t and ϕ_t^* as inputs, and are implicitly defined by equations 1.1 to 1.10, 1.13 to 1.16, and 1.19 to 1.37, all repeated for the Foreign country except for cross-country equations 1.13, and 1.31 to 1.37.

1.3 Calibration

The preliminary results in this section use the calibration presented in Table 2.1, which treats the US as the Home country and Japan as the Foreign country.

This paper is limited by scarcity of data on the holders of government bonds. In particular, there are few pairs of countries that mutually publish the identity of holders of their government bonds in sufficient detail for the calibration of this model. This paper therefore limits its study to the US-Japan and US-Australia pairs. The inclusion of the US in both of these pairs is unfortunate given that the paper aims to give a non-US centric explanation of the liquidity-depreciation link. A productive research avenue would be to document a broader set of cross-holdings of government bonds to see how well the model can explain the data for non-US pairs, and to explore whether the model can replicate differences in regression coefficients between country pairs. Nevertheless, the main contribution of this paper is to explore the underlying mechanism. Since none of the model's features or assumptions are US-specific, the resulting insights are also unlikely to be.

The parameters and steady state values that I calibrate according to country pair data are as follows, for the case of US-Japan. The 2022 GDP figures lead to $n = 0.86$. The countries' average quarterly liquidity yields over the 1999-2018 period lead to $\gamma = 0.0010$ and $\gamma^* = 0.0003$. Steady state bond holdings are calibrated to 2022 figures for domestic non-government holdings of government bonds (94.2% of GDP for the US and 106.8% for Japan), as well as cross-holdings (US holdings of Japanese bonds of 1.6% of US GDP, and Japanese holdings of US bonds of 28.1% of Japan's GDP).

In addition, the paper adopts a calibration target for the marginal effect of liquidity on investment, I_t. The target is derived from Denis and Sibilkov (2010), who measure the marginal effect of cash holdings on firm-level investment. For firms, cash holdings are readily-available sources of liquidity that can finance investment, and are analogous to the model's government bond holdings. Taking a weighted average of the authors' granular estimates[11], I arrive at a figure of 0.036 additional

[11] Table 4 of Denis and Sibilkov (2010) contains estimates broken down by financially constrained and unconstained firms, using four methods to identify whether a firm is constrained. To arrive at an economy-wide estimate, I take a weighted average of both types of firm using the mean firm

Figure 1.2: Marginal effect of liquidity constraint on Home investment, $\frac{dI_t}{dI_t^C}$, for different θ_K

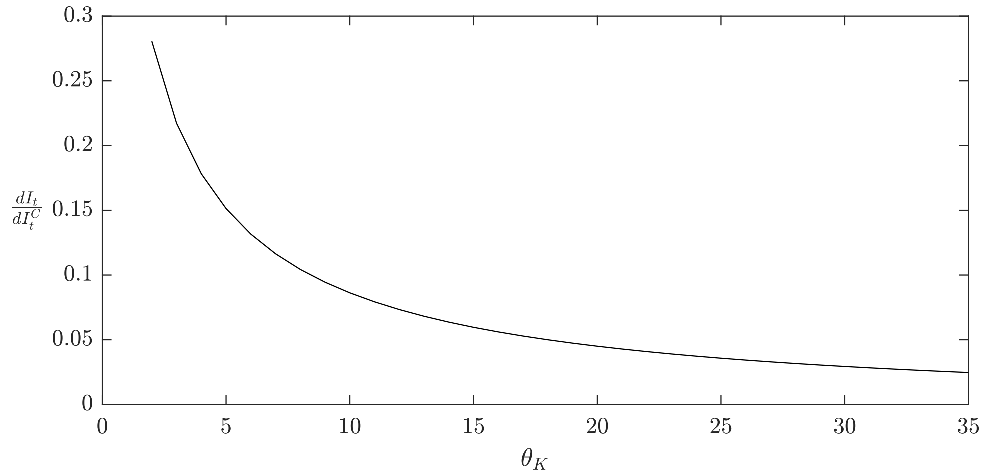

The figure assumes that only the Home bond shock is active.
$\frac{dI_t}{dI_t^C}$ is measured as the coefficient β_2 in the regression $I_t = \beta_1 + \beta_2 I_t^C + u_t$ on simulated model output, where u_t is the error term.
US-Japan calibration.

units of investment per additional unit of cash. Although this target is derived from US firm data, I assume it is representative of all countries. Another limitation is that the target is derived from public firms only. To the extent that public firms have easier access to capital, the target may be an underestimate of the impact of liquidity on firms' investment.

The model analogue for this marginal effect is $\frac{dI_t}{dI_t^C}$, because, since the liquidity constraint is assumed to bind, liquidity may be quantified by the total of Constrained investment, I_t^C. The calibration strategy targets $\frac{dI_t}{dI_t^C} = 0.036$ by changing θ_K, the elasticity of substitution between K_{t-1}^C and K_{t-1}^U. Figure 1.2 illustrates that greater ease of substitution (higher θ_K) allows overall investment to be less sensitive to liquidity (lower $\frac{dI_t}{dI_t^C}$). The resulting values are $\theta_K = 24.8$ for the US and $\theta_K^* = 33.2$ for Japan.

The calibration is to quarterly time periods, hence the discount factor β is set

sizes in Table 2 and the number of firm-years in Table 4, and I take a uniform average over the different methods in Table 4 for classifying firms as constrained and unconstrained.

so as to produce a risk-free rate of return of 2% per annum in steady state. As noted by Krishnamurthy and Vissing-Jorgensen (2012), the yield on a government bond is not the appropriate benchmark for a risk-free rate of return insofar as the bond offers a convenience yield. In the present context which deals with the impact of government bond liquidity, the appropriate risk-free rate is not i_t but i_t^m.

Although not shown in Section 3.2, the quantitative model includes quadratic portfolio adjustment costs for both households (Schmitt-Grohé and Uribe, 2003), in each case applied to deviations of the domestic government bond holding compared to a given steady state bond holding. Without adjustment costs, each household's portfolio mix of Home and Foreign government bonds would be indeterminate, since the bonds are perfect substitutes within the liquidity constraint. The adjustment cost is given a weight of 0.001 in the budget constraints so as not to interfere with business cycle frequency effects.

The remaining parameters are either set to values common in the literature, or determined as a result of the values given to other calibrated parameters and steady state quantities, as indicated in Table 2.1.

A notable quantity that is determined by other calibrated variables is the steady state portion of holdings of government bonds that may be liquidated intra-period by Constrained household members to finance Constrained investment, $\chi\phi$ and $\chi\phi^*$. The model does not require separate identification of χ, the portion of Constrained vs. Unconstrained household members, from ϕ or ϕ^*, the steady-state of the exogenous portion of bonds that may be liquidated in each country. The reason is that these three variables together determine just two multiplicative effects within the liquidity constraint's steady state; one for each bond. However, the ratio of $\frac{\phi}{\phi^*}$ is identified, since $\frac{\phi}{\phi^*} = \frac{\gamma}{\gamma^*}$, from Equation 1.8 for both countries and the fact that $\rho = \rho^*$ in steady state. The ratio $\frac{\phi}{\phi^*}$ is therefore aligned with market data for government bond liquidity yields.

The weights of Constrained capital in the capital aggregation function in each country, ζ_K and ζ_K^*, are also determined by other calibrated values, particularly $\chi\phi$ and $\chi\phi^*$, and the degree of home bias in steady state bond holdings. Considering that the liquidity constraint binds in steady state, the total liquidity available to a household pins down the country's steady state Constrained investment-to-output ratio and thus Constrained capital-to-output ratio. ζ_K and ζ_K^* are allowed to differ so that liquidity does not also pin down the overall capital-output ratio, and thus productivity, in steady state. Instead, the relative importance of steady state Unconstrained capital in production is allowed to adjust accordingly.

Table 1.1: US-Japan base calibration

		Home	Foreign*
Calibrated quantities			
β	Discount factor	0.995	
σ	Relative risk aversion	2	
δ	Capital depreciation rate	0.025	
α	Capital share in production	0.33	
λ	Openness in consumption	0.5	
θ_C	C_t^H vs. C_t^F elasticity of substitution	2	
θ_K	K_t^C vs. K_t^U elasticity of substitution	24.8	33.2
$\frac{B^H+B^F}{4Y}$	Steady state savings-to-GDP	0.96	1.35
ι	Steady state home bias in bond holdings	0.98	0.79
$\frac{\bar{G}}{Y}$	Steady state government expenditure-to-GDP	0.20	
γ	Steady state liquidity yield	0.0010	0.0003
n	Population of Home as portion of world	0.86	
σ_ϕ	Liquidity shock volatility	0.10	
ρ_ϕ	Liquidity shock persistence	0.90	
Γ_Π	Taylor rule inflation coefficient	1.5	
Outworkings of the calibration			
$\frac{K}{4Y}$	Steady state capital-to-GDP	2.71	2.72
$\frac{C}{Y}$	Steady state consumption-to-GDP	0.53	0.53
$\frac{K^C}{K}$	Steady state Constrained capital to total capital	0.67	0.43
ρ	Steady state liquidity constraint multiplier	0.022	
ζ_K	Constrained capital weight in production function	0.775	0.607
$\chi\phi$	Product of Constrained member portion (χ) and steady state saleability (ϕ)	0.0476	0.0148
ζ_C	Home bias in consumption	0.93	0.57

* Foreign column left blank where there is no distinction between the Home and Foreign value. Symbols and formulae given for the Home country are equivalent for the Foreign country. Bondholder home bias ι and ι^* are given by $\frac{B^H}{B^H+B^F}$ and $\frac{B^{F*}}{B^{H*}+B^{F*}}$ respectively. Home = US, Foreign = Japan.

1.4 Results

The purpose of the model is to study an asymmetric shock to liquidity (ϕ_t or ϕ_t^*) and the consequences for the liquidity yields, exchange rates, and other variables. The goals are to replicate pairwise versions of E&W's regressions with simulated model output, elucidate the mechanism by which the relative liquidity yield and the exchange rate are linked, and explore potential causality from a liquidity shock to exchange rate movement. For exposition, I describe an adverse shock to liquidity; however, the shock may also be positive, with the described mechanism applying in reverse.

1.4.1 Replication of empirical results

Table 1.2 shows regression results based on the E&W specification. However, instead of using a panel of all 9 pairs for a given currency, it contains pairwise regressions for US-Japan and US-Australia, for comparison against the model's output for those two country pairs.[12]

The model is able to replicate the sign of coefficients of $\Delta\eta_t$ and Δi_t^R for both calibrations. The US-Australia model calibration leads to greater-magnitude coefficients for these variables than the US-Japan calibration, consistent with the real-world data. The model also replicates the sign of S_{t-1} in both cases.

1.4.2 Effects of adverse liquidity shock on real variables

The E&W regression uses nominal variables, following the original Fama (1984) regression for the forward premium puzzle, along with others in the literature. However the reasons for the link in question are grounded in real effects. I will first cover those, before explaining why the nominal variables follow suit.

[12]The US-Australia calibration (see Appendix) is characterised by Australia being especially small relative to the US ($n = 0.94$, compared to $n = 0.86$ for US-Japan), and having especially low household bond holdings to GDP ($\frac{B^{H*}+B^{F*}}{4Y^*} = 0.12$, compared to $\frac{B^{H*}+B^{F*}}{4Y^*} = 1.35$ for Japan).

Table 1.2: Pairwise regressions

Regression equation:
$\Delta\hat{\mathcal{E}}_t = \alpha + \beta_1 S_{t-1} + \beta_2 \Delta\eta_t + \beta_3 \Delta i_t^R + \beta_4 \eta_{t-1} + \beta_5 i_{t-1}^R + u_t$

	S_{t-1}	$\Delta\eta_t$	Δi_t^R	Observations	R^2
Home currency USD, Foreign currency JPY					
Data†	-0.028**	-1.28	-4.16***	228	0.118
	(0.012)	(1.18)	(0.96)		
Model	-0.070	-4.87	-4.42		
Home currency USD, Foreign currency AUD					
Data†	-0.025*	-9.98***	-5.83***	228	0.257
	(0.013)	(1.61)	(0.90)		
Model	-0.007	-5.28	-5.27		

Standard errors are in parentheses. *, **, and *** denote significance levels 10%, 5%, and 1% respectively. Model standard errors are negligible.
The US-Australia calibration is given in the Appendix.
† Using E&W data.

Figure 1.3: Equilibrium responses to adverse Home liquidity shock, US-Japan

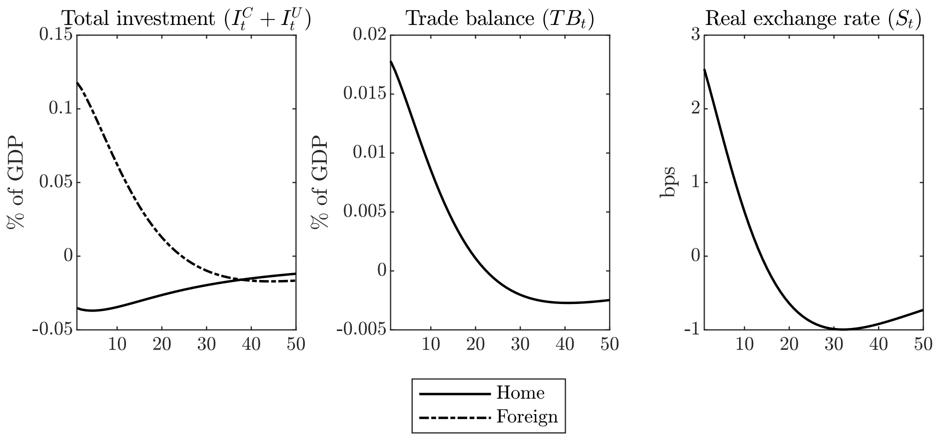

Impulse response functions for a -10% shock to Home liquidity ϕ_t.
The x-axes depict the number of quarters following a shock.

Figures 1.3 and 1.7 present impulse response functions (IRFs) following a negative shock to ϕ_t and ϕ_t^* respectively. Focussing on the shock to the Home (US) bond ϕ_t, although both countries hold US bonds and are affected, the US household holds more US bonds as a portion of its GDP, and is therefore more strongly affected. In response to the shock, Constrained investment is forced downwards in both countries as the shock tightens the already-binding liquidity constraint. Due to substitutability between Constrained and Unconstrained capital, households increase Unconstrained investment. However, the substitutability is not perfect, so the forced rebalancing of investment weakens the marginal product of capital. Total investment falls in the country more exposed to the shock (the US), where households are now comparatively less productive at using capital. The reduced investment manifests as a US trade surplus. Correspondingly, investment in Japan rises. More of the world's output is now being absorbed according to Japanese preferences, which, due to home bias, apply a smaller weight to the US good and a higher weight to the Japanese good. Demand for the US good weakens relative to the Japanese good, and the US real exchange rate depreciates.

The mechanism is an example of a financial shock that affects the ease, or profitability, of investment. In such a scenario, as per the well-known theoretical result, if a country's absorption decreases due to the shock, and if there is home bias in preferences, then the country's home goods prices will also decrease, and its real exchange rate will depreciate. Correspondingly, there are many historical examples of financial shocks that have coincided with such movements of investment and the real exchange rate. For exposition purposes, I give some historical examples, without claiming that the above mechanism was the only or the main influence. Most are negative, such as the Thai currency and financial crisis (1997-99), the Argentinian default and banking crisis (2001-02), the Turkish recent period of financial stress (2018-19), the Russian financial sanctions (2014-15), the Greek debt and banking crisis (2010-15), and the Icelandic banking crisis (2008-09). A positive example is the Polish accession to and financial integration with the European Union (2004-08). In all of these episodes, a rise (fall) in gross fixed capital formation coincided with an appreciation (depreciation) of the real exchange rate.[13]

Coming back to this paper's model, pivotal to the mechanism is the effect of liquidity on overall investment. This effect is moderated by θ_K, the elasticity of

[13]Based on data from World Bank World Development Indicators and Federal Reserve Economic Data (FRED).

substitution between Constrained and Unconstrained capital in production. Easier substitution allows households to replace Constrained investment with Unconstrained following an adverse shock, leading to less overall investment reallocation.

The ease of substitution θ_K also has an important role on the financial side, influencing the liquidity yield γ_t's shock response. As per Equation 1.10, γ_t experiences a reduction due to lower expectations of ϕ_{t+1}. After all, a bond that provides less liquidity earns a lower liquidity yield. However, expectations of liquidity constraint multipliers ρ_{t+1} and ρ^*_{t+1} increase, leading to an offsetting effect. They rise because the shock makes liquidity more scarce, and because Unconstrained investment cannot fully substitute for Constrained. Lower elasticity of substitution θ_K leads to the constraint binding more tightly.

The direction of movement of the liquidity yield γ_t depends on which effect prevails. Figure 1.4 shows that, with sufficiently high elasticity of substitution θ_K, the increase in ρ_t is not enough to overpower the effect of ϕ_t, and the liquidity yield on the shocked bond falls. In effect, a higher elasticity weakens the constraint, and infinite elasticity would nullify the constraint entirely. Using an empirical calibration target to determine θ_K ensures that the balance of the competing effects is grounded in reality.

The resulting θ_K and θ^*_K values result in a liquidity scarcity effect that is small relative to the direct effect of the shock. Thus γ_t falls, and since there is little effect on γ^*_t, η_t falls. While this supports the intuition behind the relationship between $\Delta\mathcal{E}_t$ and $\Delta\eta_t$, a falling η_t is not required for $\Delta\eta_t$'s coefficient to be negative in the Table 1.2 regression, due to the other regressors present.

It is noteworthy that the exchange rate movement is not driven by changes to bond portfolio composition. Firstly, total real borrowings of both governments are assumed constant in this model, so there is no change in total holdings of either bond. Secondly, when a bond's liquidity changes, its attractiveness changes in the same way for all holders, so there is little reason to expect an overall financial

Figure 1.4: Effect of elasticity of substitution between Constrained and Unconstrained capital

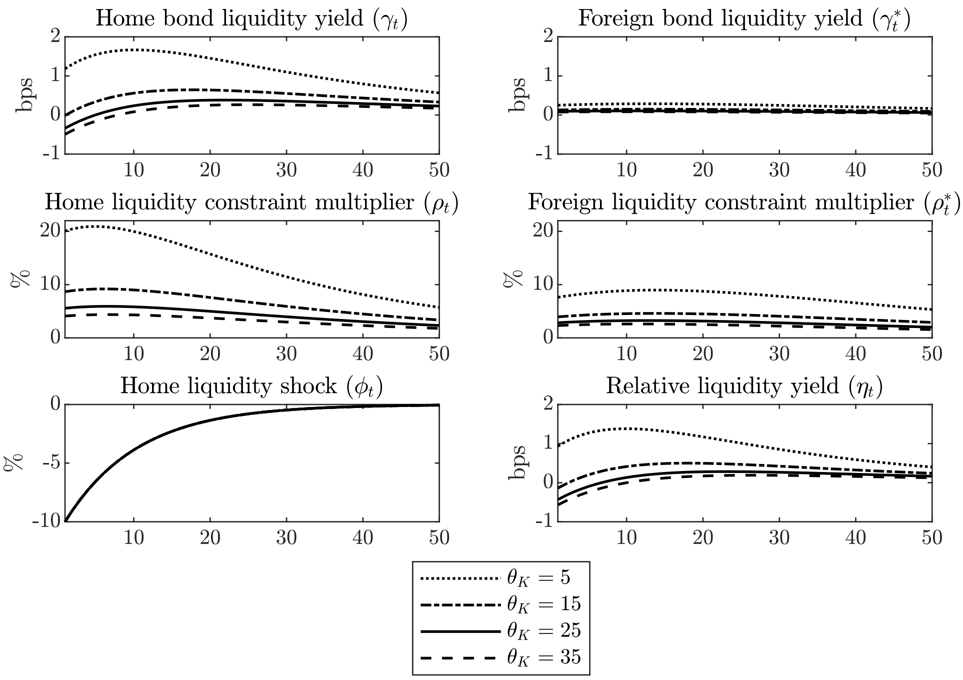

Impulse response functions for a -10% shock to Home liquidity ϕ_t.
The x-axes depict the number of quarters following a shock.
US-Japan calibration.

flow to arise directly from the shock. In this calibration, as shown in Figure 1.10 in the Appendix, when the US bond's liquidity weakens, the US household buys US government bonds and sells Japanese bonds. However, there need not be any portfolio composition change at all; as, for example, in the symmetric model calibration in Figure 1.10. Instead, the exchange rate movement is driven by the investment reallocation channel.

1.4.3 Nominal variables and monetary policy

Although the model's assumed monetary policy has no effect on real variables, it does affect inflation and therefore the nominal variables used in the E&W regression specification, namely exchange rate depreciation and relative interest rates.[14]

Given money superneutrality, analytical expressions can be derived for relative

[14] The relative liquidity yield is a real variable and therefore unaffected.

inflation in terms of the model's real variables. Relative inflation can be shown to be determined by the real relative interest rate and real exchange rate depreciation, and their future expectations.

To see this, combine Equation 1.26 with a first-order approximation of the Fisher equation, $i_t^m = i_t^{m,real} + \mathbb{E}_t \hat{\Pi}_{t+1}$, and iterate forward to obtain

$$\hat{\Pi}_t = \mathbb{E}_t \sum_{s=t}^{\infty} \left(\frac{1}{\Gamma_\Pi^{s-t+1}} \hat{i}_s^{m,real} - \frac{1}{\Gamma_\Pi^{s-t}} \Delta \hat{P}_s^H \right)$$

Inflation is a function of the real interest rate and domestic good price change, and their future expectations.

Subtracting the foreign equivalent to obtain relative inflation gives

$$\hat{\Pi}_t^R = \mathbb{E}_t \sum_{s=t}^{\infty} \left(\frac{1}{\Gamma_\Pi^{s-t+1}} \hat{i}_s^{mR,real} - \frac{1}{\Gamma_\Pi^{s-t}} (\Delta \hat{P}_s^H - \Delta \hat{P}_s^{F*}) \right)$$
$$\hat{\Pi}_t^R = \mathbb{E}_t \sum_{s=t}^{\infty} \left(\frac{1}{\Gamma_\Pi^{s-t+1}} \hat{i}_s^{mR,real} - \frac{1}{\Gamma_\Pi^{s-t}} (-\Delta \hat{TOT}_s + \Delta \hat{S}_s) \right)$$
$$\hat{\Pi}_t^R = \mathbb{E}_t \sum_{s=t}^{\infty} \left(\frac{1}{\Gamma_\Pi^{s-t+1}} \hat{i}_s^{mR,real} + \frac{1}{\Gamma_\Pi^{s-t}} \frac{\lambda}{1-\lambda} \Delta \hat{S}_s \right)$$

The second expression for $\hat{\Pi}_t^R$ above comes from the definition of the terms of trade TOT_t (Equation 1.36) and the law of one price (Equation 1.32). The third expression for $\hat{\Pi}_t^R$ exploits the approximate proportionality of S_t to TOT_t due to the model's absence of non-tradeable goods.[15]

Intuitively, the real relative policy interest rate influence on inflation arises because higher real interest rates are consistent with higher inflation when a policymaker follows the Taylor principle.

The influence of real exchange rate depreciation arises because the inflation measure targeted is not the same as the difference between the nominal and real policy interest rate, which is consumer price inflation. Rather, domestic inflation is used. The two are related by the change in domestic prices, as per Equation

[15] Specifically, approximating the price index equations of both countries to first order gives $\hat{S}_t = (\zeta_C^* - (1 - \zeta_C))\,\hat{TOT}_t = (1-\lambda)\hat{TOT}_t$, where Equations 1.17 and 1.18 are used to express the relationship using openness parameter λ.

1.27. This mismatch results in relative inflation being influenced by the relative change in each country's domestic prices, or equivalently, the difference between terms of trade deterioration and real exchange rate depreciation. In this model, where the terms of trade and the real exchange rate are approximately proportional, the influence can be summarised in terms of real exchange rate depreciation alone.

These influences are evident in the impulse response functions in Figure 1.5. Relative inflation Π_t^R roughly tracks the real relative market rate $i_t^{mR,real}$ plus real exchange rate depreciation ΔS_t. The relative nominal market rate i_t^{mR} is then approximately an amplified version of relative real market rate $i_t^{mR,real}$ plus real exchange rate depreciation ΔS_t, without the on-impact effect since it is expected future inflation that affects the current interest rate. The relative nominal government bond interest rate i_t^R is then similar to the sum of $i_t^{R,real}$, $i_t^{mR,real}$, and ΔS_t, again without the on-impact effect. Nominal exchange rate depreciation $\Delta \mathcal{E}_t$ is approximately an amplified version of ΔS_t plus $i_t^{mR,real}$.[16]

[16]The amplification comes from the fact that $\Delta \mathcal{E}_t = \Delta S_t + \Pi_t^R$, where Π_t^R is also affected by ΔS_t and its future expectations.

Figure 1.5: Effect of adverse Home liquidity shock on nominal variables

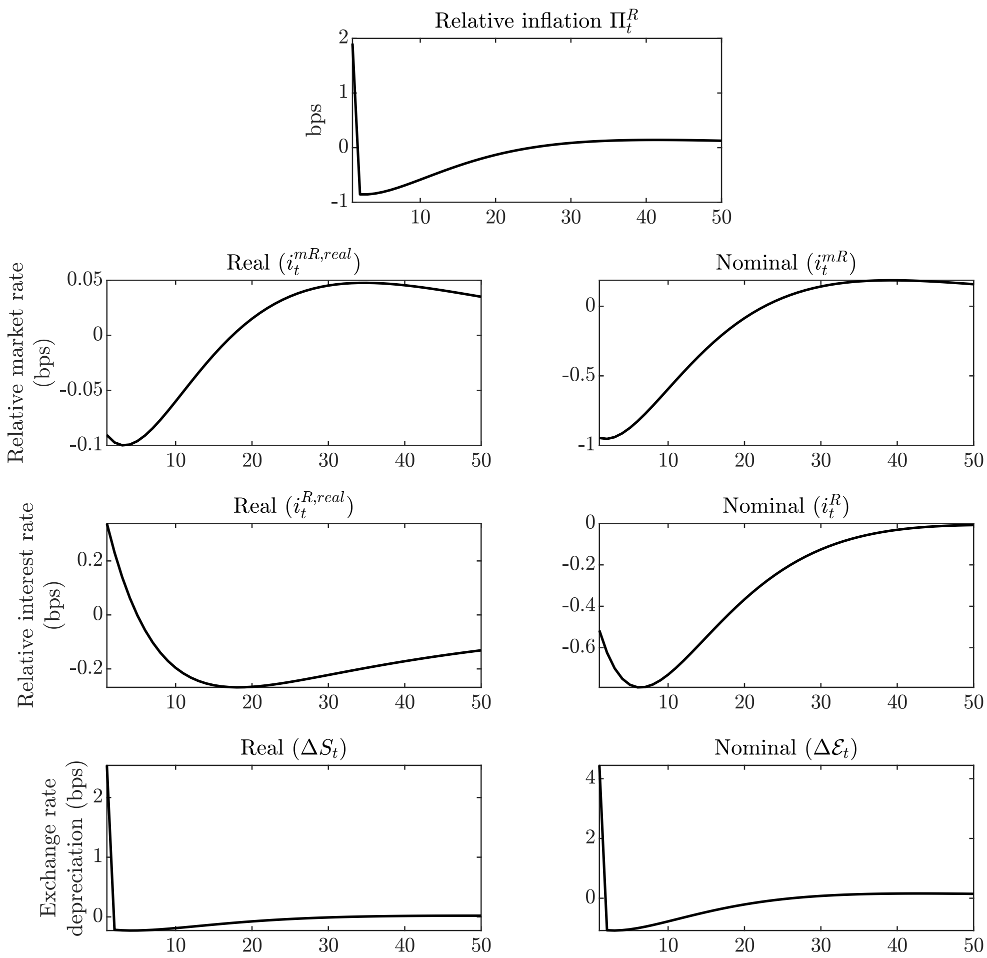

Impulse response to a -10% shock to Home liquidity ϕ_t.
The x-axes depict the number of quarters following a shock.
US-Japan calibration.

1.4.4 A comparison to bonds-in-the-utility models

Bonds-in-the-utility models make the reduced-form assumption that bond holdings contribute directly to utility. When contrasted with a richer model that microfounds the role of liquid assets, it is unsurprising that the reduced form approach omits some phenomena. However, this paper's model can explain why bonds-in-the-utility models have nevertheless been gainfully deployed, and why their reduced form assumption can be reasonable.

In bonds-in-the-utility models, liquid bonds do not alleviate a liquidity friction. Rather, liquidity fluctuations are modelled as a preference shock that changes the marginal utility of holding bonds. Consequently, there is no distinction between two concepts whose interplay is highlighted by this paper—the degree of liquidity provided by a bond, and liquidity scarcity.

Consider a basic bonds-in-the-utility model in which

- B_t^m is a riskless one-period bond with interest rate i_t^m
- B_t is a riskless, liquid one-period bond with interest rate i_t
- B_t's liquidity is modelled by its inclusion in the period utility function, $U(C_t, B_t)$, where C_t is consumption
- First derivatives $U_{C,t}$ and $U_{B,t}$ are positive
- B_t^m and B_t are in fixed supply

In this setup, a negative shock to B_t's liquidity takes the form of a preference shock that lowers $U_{B,t}$.

The first-order conditions for consumer optimisation include

$$U_{B,t} = \gamma_t \mathbb{E}_t \beta U_{C,t+1} \tag{1.38}$$

where the liquidity yield is $\gamma_t := i_t^m - i_t$.

Equation 1.38 equalises the utility gained in period t from liquid bond holdings (left hand side) with the discounted expected utility in period $t+1$ of interest earnings foregone by investing in the liquid bond rather than the market bond (right hand side). The effect of the adverse shock is an unambiguous reduction in γ_t.

Compared to the endogenous liquidity model of this paper, the liquidity shock in the bonds-in-the-utility model is similar to the “direct" effect of a shock, or the influence on γ_t of expectations of ϕ_{t+1} as per Equation 1.10. The liquidity scarcity effect—the influence of expectations of ρ_{t+1}—is absent.

However, the calibration strategy in this paper reveals that the liquidity scarcity effect is normally smaller than the direct effect of the liquidity shock. The result is that, as assumed by bonds-in-the-utility models, the liquidity yield does tend to move in the same direction as the shock. The reduced form assumption of bonds-in-the-utility models can therefore be defended on this basis.

1.4.5 The effect of relative bond exposure

The paper has noted, so far without elaboration, that the effects of the liquidity shock depend on the relative exposure of each country to the shocked bond.

To explain more precisely, the strength of the link between exchange rate depreciation and government bond relative liquidity depends on the exposure gap defined as $\frac{B^H}{4Y} - \frac{B^{H*}}{4Y^*}$, the difference between each country's holdings of the bond as a portion of their GDP. This measure relates closely to the countries' respective home bias in bondholding, ι and ι^*. In the base calibration, ι and ι^* are high, leading to a large exposure gap.

When both ι and ι^* are lowered, the Home country then owns less of the Home bond and the Foreign country owns more of it, and the exposure gap is lessened. Correspondingly, as shown in Table 1.3, the regression coefficient of $\Delta\eta_t$ reduces.

The weakened link from a reduction in both ι and ι^* is also evident in the trade balance's impulse response function to a Home bond adverse liquidity shock (see Figure 1.6). As explained earlier, the key mechanism by which the Home currency depreciates is by the Home country reducing investment and instead running a trade surplus, thus reallocating investment overseas. The extent of the trade surplus and investment reallocation is proportional to the exposure gap.

Notably, there is little effect of asymmetry in bondholder home bias. For example, comparing $\{\iota, \iota^*\} = \{0.85, 0.65\}$ with its symmetric equivalent, there is little change in the $\Delta\eta_t$ regression coefficient and little change in the trade balance response. The reason is that changing ι and ι^* in opposite directions leads to both countries' exposure moving in the same direction. Since the exposure gap is not affected, the strength of the shock response is the same.

Table 1.3: Effect of steady state exposure to the shocked bond

	ι	ι^*	Exposure gap $\frac{B^H}{4Y} - \frac{B^{H*}}{4Y^*}$	Coefficient of $\Delta\eta_t$[1]
Base calibration	0.98	0.79	0.66	-7.01
Lower home bias	0.85	0.65	0.34	-4.52
Symmetric equivalent	0.73	0.73	0.34	-4.80
Still lower home bias	0.75	0.55	0.11	-2.75
Symmetric equivalent	0.63	0.63	0.11	-3.65

1. Foreign shock is shut down when measuring the coefficient of $\Delta\eta_t$ here, to show more clearly the effect of Home bond exposure.
US-Japan calibration.

Figure 1.6: Effect of steady state exposure to the shocked bond

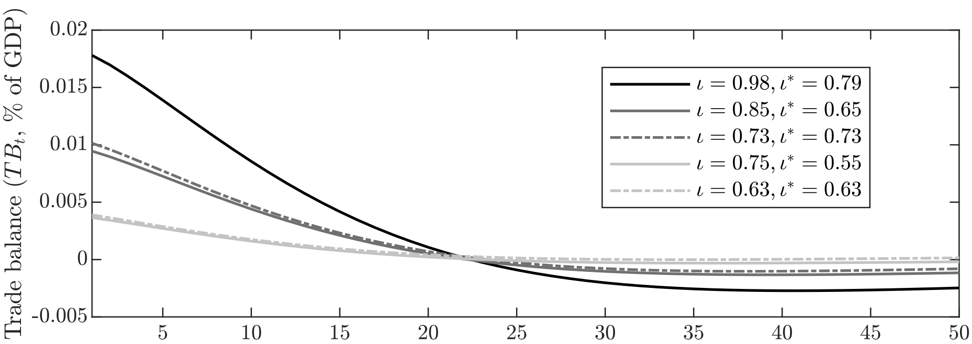

Impulse response to a -10% shock to Home liquidity ϕ_t.
The x-axes depict the number of quarters following a shock.
US-Japan calibration

1.5 Conclusion

This paper seeks to explain the link identified in empirical literature between exchange rates and the relative government bond liquidity yield. An important contribution is to establish microfoundations that do not assume a special role for US government bonds, since the link is present in country pairs that exclude the US. The explanation appeals to the role of liquidity in facilitating investment. An adverse liquidity shock moves investment offshore, depreciating the real exchange rate. Meanwhile, since the shocked bond no longer offers as much liquidity, its liquidity yield falls. I show that model-simulated data is able to produce a negative relationship between currency depreciation and relative liquidity yield, similar to observation.

The paper's quantitative results are limited to the US-Japan and US-Australia country pairs. Extending the analysis to other country pairs would require data on other countries' holdings of each other's government bonds. The primary benefit would be to explore the model's ability to explain the observed differences in the regression coefficients between country pairs. Other research directions may include policy analysis in the presence of the liquidity frictions explored, including the effects of bond issuance or other liquidity provision, and potentially exchange rate targeting. Lastly, there is scope to explore the extent to which different observable categories of investment map onto this paper's unobservable concept of Constrained investment. Such investigation may help explain which sectors or agents contribute the most to exchange rate movements through this channel. I leave these topics to future work.

Appendices to Chapter 1

Table 1.4: US-Australia calibration

		Home	Foreign*
Calibrated quantities			
β	Discount factor	0.995	
σ	Relative risk aversion	2	
δ	Capital depreciation rate	0.025	
α	Capital share in production	0.33	
λ	Openness in consumption	0.5	
θ_C	C_t^H vs. C_t^F elasticity of substitution	2	
θ_K	K_t^C vs. K_t^U elasticity of substitution	18.2	295.0
$\frac{B^H+B^F}{4Y}$	Steady state savings-to-GDP	0.94	0.12
ι	Steady state home bias in bond holdings	1.00	0.70
$\frac{\bar{G}}{\bar{Y}}$	Steady state government expenditure-to-GDP	0.20	
γ	Steady state liquidity yield	0.0010	0.0010
n	Population of Home as portion of world	0.94	
σ_ϕ	Liquidity shock volatility	0.10	
ρ_ϕ	Liquidity shock persistence	0.90	
Γ_Π	Taylor rule inflation coefficient	1.5	
Outworkings of the calibration			
$\frac{K}{4Y}$	Steady state capital-to-GDP	2.70	2.70
$\frac{C}{Y}$	Steady state consumption-to-GDP	0.53	0.53
$\frac{K^C}{K}$	Steady state Constrained capital to total capital	0.53	0.07
ρ	Steady state liquidity constraint multiplier	0.028	
ζ_K	Constrained capital weight in production function	0.645	0.995
$\chi\phi$	Product of Constrained member portion (χ) and steady state saleability (ϕ)	0.0376	0.0380
ζ_C	Home bias in consumption	0.97	0.53

* Foreign column left blank where there is no distinction between the Home and Foreign value. Symbols and formulae given for the Home country are equivalent for the Foreign country. Bondholder home bias ι and ι^* are given by $\frac{B^H}{B^H+B^F}$ and $\frac{B^{F*}}{B^{H*}+B^{F*}}$ respectively. Home = US, Foreign = Australia

Figure 1.7: Equilibrium responses to adverse Foreign liquidity shock, US-Japan

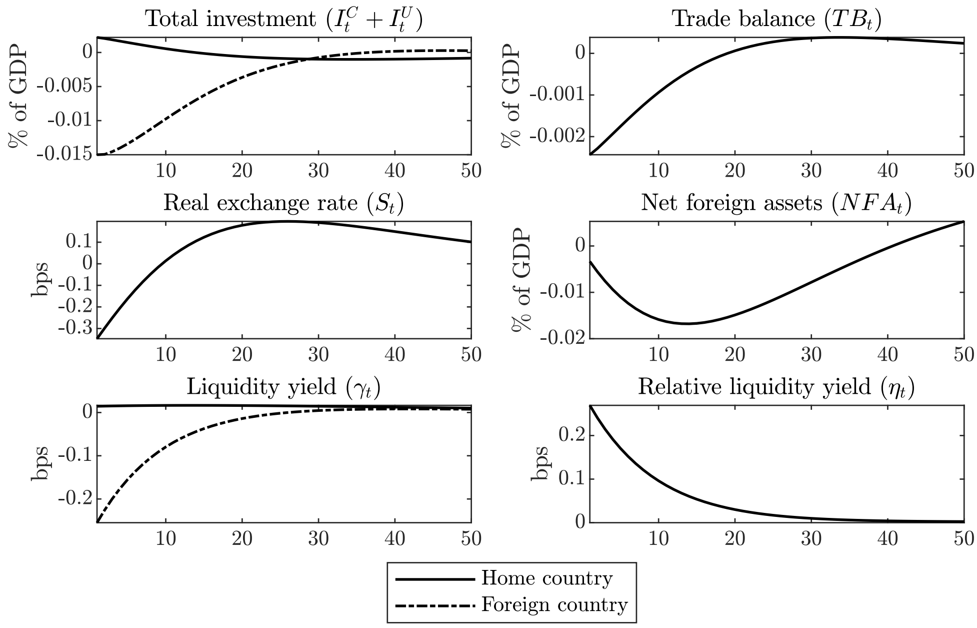

Impulse response functions for a -10% shock to Foreign liquidity ϕ_t^*.
The x-axes depict the number of quarters following a shock.

Figure 1.8: Equilibrium responses to adverse Home liquidity shock, US-Australia

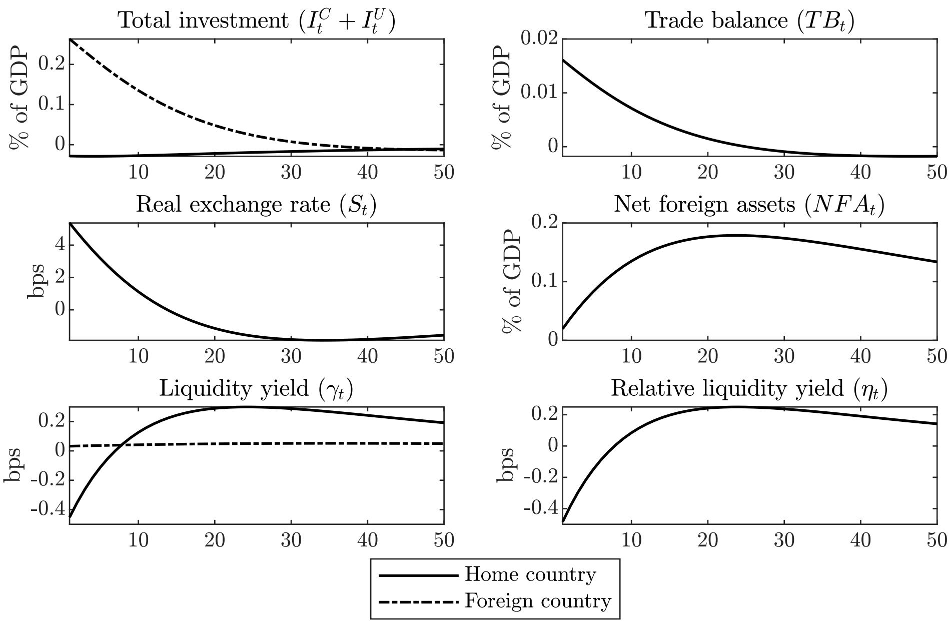

Impulse response functions for a -10% shock to Home liquidity ϕ_t.
The x-axes depict the number of quarters following a shock.

Figure 1.9: Equilibrium responses to adverse Foreign liquidity shock, US-Australia

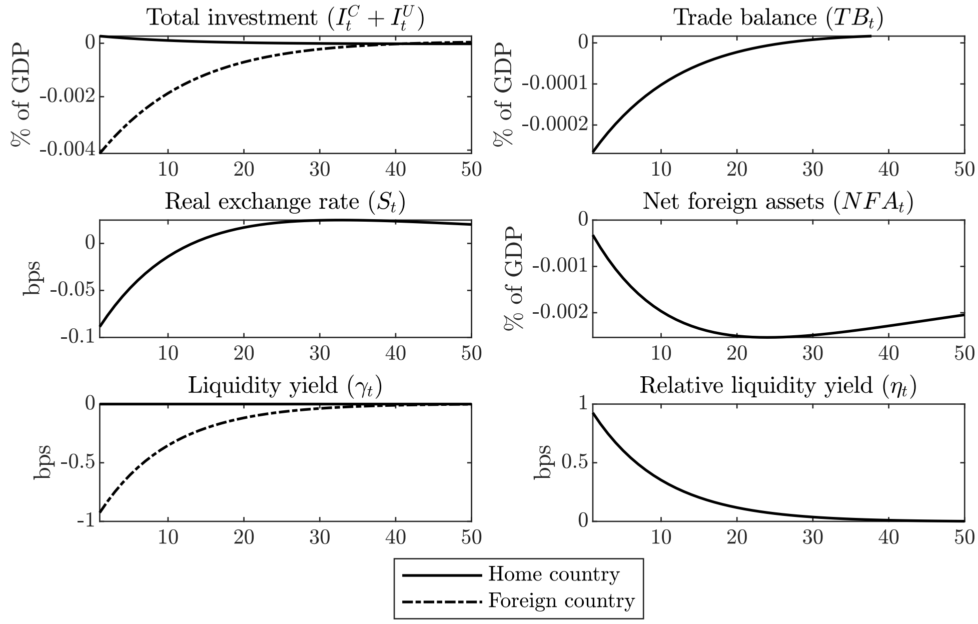

Impulse response functions for a -10% shock to Foreign liquidity ϕ_t^*.
The x-axes depict the number of quarters following a shock.

Figure 1.10: Bond portfolio composition

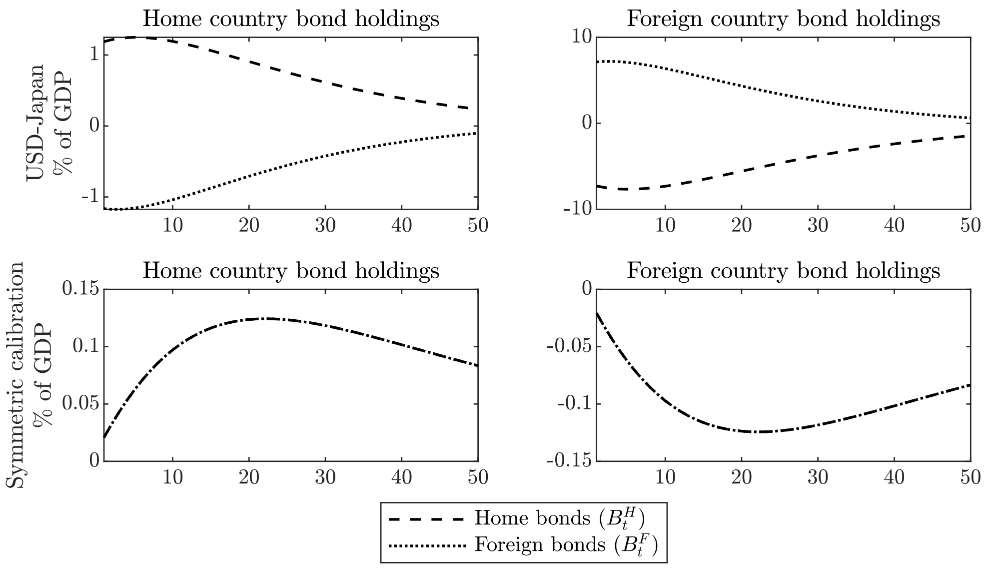

Impulse response functions for a -10% shock to Home liquidity ϕ_t^*. The symmetric calibration models a pair of US countries. Its parameters are the same as for the US-Japan calibration except that $n = 0.5$, $\frac{B^{H*}+B^{F*}}{4Y^*} = \frac{B^H+B^F}{4Y} = 0.96$, $\iota^* = \iota = 0.75$ (arbitrarily chosen), $\gamma^* = \gamma = 0.0010$, and $\theta_K^* = \theta_K = 45.7$ (chosen to target $\frac{dI_t^*}{dI_t^{C*}} = \frac{dI_t}{dI_t^C} = 0.036$). The x-axes depict the number of quarters following a shock.

2

House prices, capital flows, and international cooperation in macroprudential policy

Abstract Following the great recession of 2008, macroprudential policy has gained in popularity. However, with financial integration, it is not obvious that the full benefits can be realised without international cooperation. This paper studies loan-to-value ratio policy in a two-country model, in the context of the "global saving glut" of the early 2000s. Germany and Spain are used as a case study. Optimal, internationally cooperative policy is compared to a Nash equilibrium between national optimising policymakers, and to a status quo of no policy response. Policy regimes are ranked according to the ex-ante expected welfare they generate. As expected, non-cooperation is inferior to cooperation at the global level. However, due to certain competitive policymaker behaviour, non-cooperation is also inferior to no policy response. The paper explores alternative calibrations, some of which overturn this result.

2.1 Introduction

The popularity of macroprudential tools increased following the great recession of 2008 (Galati and Moessner, 2018), given their ability to contain excessive debt build-up. However, with increasing financial integration (Lane and Milesi-Ferretti, 2017), the case for macroprudential policy in a single economy setting amounts to a case for internationally cooperative macroprudential policy in a multi-country setting. When nations set policies independently and in the presence of financial market frictions, there is no guarantee that the non-cooperative equilibrium will be Pareto optimal.

In this paper, I contribute to a growing literature exploring whether the non-cooperative equilibrium can realise the benefits of macroprudential policy, or whether international cooperation is required. Concretely, I focus on the early 2000s period before the 2008 great recession. Bernanke (2005) hypothesised that, during that period, the world experienced an international imbalance characterised by a global saving glut, with house price booms in those countries that were the destination of the capital flows. As per Figure 2.1, countries with negative current account movements in the early 2000s were more likely to simultaneously experience a rise in the price of housing during this period (Bernanke, 2010). The empirical correlation between indebtedness and asset prices naturally evokes the potential need for loan-to-value ratio limits, which are the macroprudential policy instruments studied in this paper.

I use the country pair of Germany and Spain as a case study. As implied by Figure 2.1, Germany acted as a source of capital, and Spain as a recipient. Simultaneously, Spain's housing market experienced a boom in prices, while Germany's house prices remained stagnant.

I develop a two-country currency union model, with each country having a borrower and saver household, and with all households deriving utility from consumption and housing. Borrowing is motivated by borrowers being more impatient than savers. They may only borrow up to a portion of the value of

Figure 2.1: Change in house prices and current account balance from 2000 to 2007

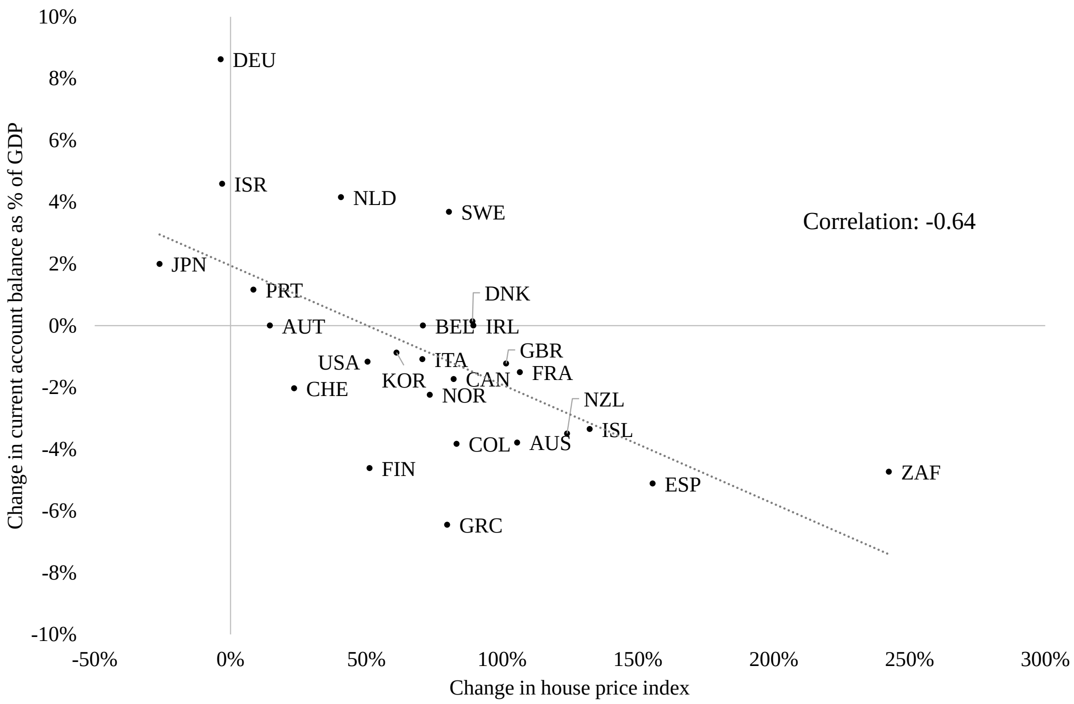

Data sources: World Bank, OECD.

their housing, with this loan-to-value ratio being the instrument of policy for each country. The consumption good is internationally tradable, and households may borrow and save internationally. Given this setting, I study three policy regimes. In the “no policy response” regime, which is intended to model the early and mid 2000s, the loan-to-value ratio of each country is assumed constant through time. In the cooperative regime, a single policymaker sets both loan-to-value ratios in each period so as to maximise global welfare. In the non-cooperative regime, each country’s loan-to-value ratio is controlled by a national policymaker who maximises a utility function that only considers households in their own country.

The shock in the model is a dispersion of preferences shock. The shock is asymmetric, decreasing borrower patience and increasing saver patience in Germany, pushing the economy harder against the borrowing constraint. The shock is interpretable as a fluctuation in the degree of domestic dispersion of agents’ net demand for borrowing or saving.

The model output shows that the no policy response regime is capable of replicating the directions of movement of net foreign assets (NFA) and house prices, given a particular sequence of shocks in Germany. Greater preference dispersion leads to a dilemma for the borrower, who would like to both borrow more and own less housing, since housing is a form of saving. The constraint forces them to reach an unhappy medium in which less additional debt is created than would be needed to satisfy the saver's appetite for safe assets. The German saver responds by lending more to Spanish households. The Spanish saver, faced with a diminished interest rate from debt, demands more housing as saving, pushing its price up. Similarly, the Spanish borrower demands more debt at the lower interest rate, creating further upward pressure on the Spanish house price as long as the borrowing constraint binds. Therefore, as Spain's NFA deficit grows, the price of its housing increases. Meanwhile, the model is calibrated such that German house prices exhibit little change. The model is highly parsimonious, yet replicates the direction of changes to NFA and house prices, as well as the world interest rate and Spain's household debt. A richer model may, in addition, be able to match the magnitudes of these changes.

Using the same calibration, I project the optimal cooperative policy response where a single policymaker maximises global welfare, and the non-cooperative (Nash equilibrium) policy responses where each country's policymaker separately maximises their households' welfare, internalising their impact on the global interest rate, and taking the other policymaker's policy as given. The cooperative policy response tightens loan-to-value ratio policy in both countries in response to the shock, resulting in a significant reduction in borrower debt in both countries relative to the no policy response regime. NFA is not significantly affected by the policy. There is a dampening effect on house prices in both countries, but it is slight. By contrast, the non-cooperative policy regime is characterised by much higher loan-to-value ratios in both Germany and Spain than under cooperative policy.

Correspondingly, household debt is much higher in both countries, and the slight reduction in house prices is reversed.

The paper also evaluates the ex ante welfare effects of the policy regimes. The cooperative policy regime improves global welfare relative to the no policy response regime, which is consistent with the positive expected effect of macroprudential policy in an autarkic context. However, under the non-cooperative policy regime, welfare is not only worse than in the cooperative regime, but also worse than in the no policy response regime. The result obtains for many plausible parameter combinations, but not all. Therefore, to realise the benefits of macroprudential policy in the context that the paper studies, international cooperation may be needed. Non-cooperative macroprudential policy could make matters worse than the status quo, and is always inferior to cooperative policy in the presence of these frictions.

The optimal policy regime is sensitive to whether house prices in the Home country respond to the shock. Although Spain experienced a boom, in Germany, house prices were quite stable, so there was little amplification effect from collateralised borrowing. Therefore, the model estimates that the benefits of macroprudential policy to Germany are small, and would be outweighed by the detrimental effects of the non-cooperative Nash equilibrium. However, for capital-exporting countries where house prices responded more strongly, there may be overall benefits of non-cooperative macroprudential policy relative to no policy response.

The paper's results can be interpreted with reference to the theoretical conditions for beneficial international cooperation presented in Korinek (2016). The conditions, if met, render national policymakers isomorphic to individuals participating in free markets. Then any international policy spillovers are Pareto-efficient, and no scope for beneficial cooperation exists. The "Korinek (2016) conditions" are 1) national policymakers behave competitively in international markets, 2) policy instruments are sufficient for countries to control their external transactions, and 3) international markets are free of imperfections. In particular, domestic inefficiencies, such as the borrowing constraint in the present paper, do not necessarily breach these conditions, even when they are not fully counteracted by domestic policy.

The present paper's model breaches all three of the Korinek (2016) conditions. National policymakers exert market power by internalising their influence on the world interest rate. External instruments are not sufficient in the model, since neither country can tax or subsidise net foreign assets. With a bond being the only internationally traded financial asset, international financial markets are incomplete.

The paper contributes to a strand of literature comparing the internationally cooperative and non-cooperative use of macroprudential policy (including Bengui, 2014; Jeanne, 2014; Rubio, 2014; Fornaro and Romei, 2019; Agénor et al., 2021a; Agénor et al., 2021b; Chen and Phelan, 2021; Agénor and Pereira da Silva, 2022; Intungane, 2023). Invariably, given sufficient financial frictions or other market distortions, gains from international cooperation of macroprudential policy can be found. In particular, this paper is not alone in finding that non-cooperative policy can be worse than a regime of no policy response. Bengui (2014) reaches a similar conclusion from a theoretical examination of international macroprudential liquidity regulation with an international borrowing constraint. Few of these existing papers use a model capable of assessing the effects of optimal time-varying non-cooperative and cooperative policy. None specifically study the context of a global saving glut.

The paper's main contribution is quantitative. For Germany and Spain in the 2000s, and given the capital flows and housing boom associated with the global saving glut, loan-to-value ratio policy may have damaged welfare due to not being internationally cooperative. As a cautious generalisation, non-cooperative policy in a similar context is more likely to be harmful if house prices move little in the capital-exporting country, as was the case in Germany. The result is driven by the incentives for national policymakers to distort domestic policy towards beggar-thy-neighbour ends, and the ability of cooperative policy to achieve a degree of risk sharing despite incomplete international financial markets.

The paper is structured as follows. Section 3.2 outlines the model. Section 2.3 sets out the calibration for the Germany/Spain case. Section 2.4 models a series of shocks designed to simulate the Spanish house price run-up in the early 2000s. Section 2.5 turns to optimal policy, exploring the differing objectives of cooperative and non-cooperative policymakers and the associated outcomes, and the ex ante welfare effects of the policy regimes. Section 3.7 concludes.

2.2 Model

I use a two-country endowment economy model, with the Home country being Germany and the Foreign country being Spain. Each country has a saver and a borrower household, with the impatient borrower having a lower discount factor. Aside from the consumption good, each household also derives utility from housing which is not internationally tradable. The countries are modelled as being in a currency union, with a single, internationally tradable consumption good such that the real exchange rate is always 1. There are no nominal rigidities in the model, implying that monetary policy is able to close any output gap. Each country is the same size.

Two financial frictions are included. The first is that each borrower is subject to a collateralised borrowing constraint, with debt limited to the value of their housing multiplied by a loan-to-value ratio. This friction can be microfounded by the risk of borrower default, from which the model abstracts. When the constraint binds, the endogeneity of the limit leads to accelerator effects (Kiyotaki and Moore, 1997), where greater demand for housing raises its price, relaxes the constraint, and allows more debt-financed housing demand. Borrowers are rational optimisers, but do not internalise their impact on the price of housing and thus their effect on the borrowing limit. The resulting pecuniary externality can lead to over-borrowing (Bianchi, 2011; Stein, 2012; Dávila and Korinek, 2018; Lorenzoni, 2008).

Macroprudential policies that restrict borrowing can therefore improve welfare. In this paper, the loan-to-value ratio is used as the instrument of policy, and is allowed to vary between countries and over time.

The second friction is that financial transactions are limited to purchases and sales of state-independent bonds. The friction operates both between the borrower and saver of a single country, as well as between countries. The lack of state-contingent instruments means that agents are unable to insure against the effects of shocks, or their interaction with the borrowing constraint. Whereas the borrowing constraint prevents equalisation of marginal rates of substitution of different agents, the lack of risk sharing prevents the equalisation of marginal utilities across states.

The competitive equilibrium of the model assumes as “no policy response” policy regime, reflective of Germany and Spain in the 2000s. It is not a laissez-faire regime. Rather, it assumes that both countries do have a loan-to-value ratio limit, but that the limit is unchanging. In Germany and Spain in the 2000s, no substantial time-varying macroprudential policies were active (Arena et al., 2020). However, prevailing microprudential limits on borrowing created a de facto overall limit that was not designed to be responsive to macroprudential considerations.

This section gives the Home equations, noting where the Foreign country equations differ. Foreign variables are denoted by an asterisk.

2.2.1 Household optimisation

Home households maximise the following utility function.

$$\mathbb{W}_t^i = \mathbb{E}_t \sum_{s=0}^{\infty} \prod_{k=t}^{t+s} \beta_k^i \left(\frac{{C_{t+s}^i}^{1-\sigma}}{1-\sigma} + \chi \frac{{H_{t+s}^i}^{1-\sigma_h}}{1-\sigma_h} \right) \tag{2.1}$$

where $i \in \{s, b\}$ for saver and borrower, and C_t^i and H_t^i denote consumption and housing at time t. Parameters σ and σ_h are invariant between Home and Foreign households, whereas χ has a foreign counterpart χ^*.

The discount factors β_t^i in the Home country are defined as follows.

$$\log\left(\beta_t^s\right) = \log\left(\beta^s\right) - \psi \log\left(\frac{C_t^s}{C^s}\right) + \omega u_t \tag{2.2}$$

$$\log\left(\beta_t^b\right) = \log\left(\beta^b\right) - \psi \log\left(\frac{C_t^b}{C^b}\right) - (1-\omega) u_t \tag{2.3}$$

The steady state of the saver's time preference, β^s, is assumed to exceed that of the borrower, β^b. The differing degrees of impatience thus create their respective incentives to lend and borrow. The resulting tension makes the constraint bind in steady state.

The only shock in the model, u_t, acts upon the difference between borrower and saver impatience. The shock has persistence ρ and volatility ν. Its incidence on the saver versus the borrower is governed by $\omega \in [0,1]$.

The shock's impact is asymmetric, moving the agents' discount factors in opposite directions. This precise form of shock is not directly based on any existing literature. However, the use of preference shocks acting on the discount factor in borrower-saver models is in line with the literature (Iacoviello and Neri, 2010; L. Guerrieri and Iacoviello, 2017). Of course, preference shocks are a reduced form source of influence, and do not automatically map to a particular microfoundation. They are less instructive than a microfounded shock, since they abstract from the mechanism of influence. However, they do typically carry an intuitive interpretation. For example, a country-wide discount factor shock commonly signifies a demand shock.

This paper's shock u_t may be interpreted as a shock to the dispersion of domestic saving vs. borrowing preferences. The change in dispersion could be traced to changes in any economic factor whose impact correlates with agents' net borrowing demand. That is, the concept of a dispersion shock abstracts from specific mechanisms, similarly to other preference shocks.

The shock is a natural extension of this two-agent, borrower-saver class of macroeconomic model. In its simplest form, the model studies limits to gross debt.

It generates positive gross debt by introducing dispersion of domestic preferences. Typically, the dispersion is assumed static. The shock in this paper's model allows it to instead fluctuate.

The relevance of this dispersion shock is that, within the model, it creates safe asset shortages, consistent with the global saving glut in the early to mid 2000s. With greater dispersion, a country will want higher gross debt. With a borrowing constraint, an increase in dispersion leads the constraint to bind more tightly, leading to a shortage of domestic safe assets, leading savers to send capital overseas.

Other models have used shocks to the loan-to-value ratio within the borrowing constraint to study the impact of the tightness of the constraint (Jones et al., 2022; V. Guerrieri and Lorenzoni, 2017). In this paper, that ratio is instead the instrument of policy. Fortunately, tightness is determined not only by the constraint itself, but by how strongly agents' desired allocations press against it. The shock in this paper affects the latter. The dispersion shock is therefore a sensible way to use this class of model to study optimal loan-to-value ratio policy in a global saving glut, as per the purpose of this paper.

Despite the novelty of this type of shock, its nature and influence is not the main focus of this paper, implying scope for further work to explore the shock's properties.

The Home households' time preference is assumed to be negatively correlated with consumption deviation from steady state. The assumption is needed to induce stationarity and facilitate the later analysis of optimal policy. The assumption follows Schmitt-Grohé and Uribe (2003), who explore ways of inducing stationarity without materially influencing a model's short-term business cycle characteristics. The strength of the relationship is governed by ψ, which is calibrated so as not to interfere with business cycle frequency effects.

By contrast, the Foreign households' time preference is not exposed to a shock, and does not need to include a stationarity-inducing term. However, they exhibit the same steady-state borrower-saver difference as in Home. Therefore $\beta_t^{s*} = \beta^s$ and $\beta_t^{b*} = \beta^b$.

Savers face the following budget constraint.

$$C_t^s + B_t + Q_t H_t^s (1 + \tau_h) = Y + R_{t-1} B_{t-1} + Q_t H_{t-1}^s - T_{s,t} \tag{2.4}$$

where Q_t is the price of housing, R_t is the world interest rate, and Y is the per-period endowment of the consumption good.

Savings are denoted as B_t. Capital markets are international, so B_t combines savers' lending to both the domestic borrower and overseas.

Savers pay a tax on housing, τ_h, and a lump sum tax $T_{s,t}$, which are used to ensure the efficiency of the steady state for the purposes of welfare analysis. The housing tax revenue is reimbursed through an adjustment to the lump sum tax, as follows.

$$T_{s,t} = T_s - H_t^s Q_t \tau_h \tag{2.5}$$

where T_s is derived in Section 2.2.3.

Borrowers face a budget constraint and a collateralised borrowing constraint.

$$C_t^b - D_t + Q_t H_t^b = Y - R_{t-1} D_{t-1} + Q_t H_{t-1}^b + S_b \tag{2.6}$$

$$R_t D_t \leq \Theta_t Q_t H_t^b \tag{2.7}$$

where D_t is debt, and Θ_t is the loan-to-value ratio limit.

Borrowers receive a lump sum subsidy S_b, derived in Section 2.2.3, which is used to ensure the efficiency of the steady state for the purposes of welfare analysis.

The borrowing limit is a portion Θ_t of the value of housing held by the borrower. The endogeneity of the limit admits accelerator effects (Kiyotaki and Moore, 1997). If additional borrowing is partly spent on housing, this will increase the price of housing Q_t, thereby allowing further borrowing. Importantly, households do not internalise their effect on the price of housing, so their borrowing results in a pecuniary externality. With financial frictions, a pecuniary externality may have welfare effects to the extent that a shock takes the economy away from its Pareto efficient steady state. The externality creates a role for macroprudential policy.

The first order conditions are as follows.

$$\frac{1}{R_t} = \mathbb{E}_t \beta^s_{t+1} \left(\frac{C^s_{t+1}}{C^s_t} \right)^{-\sigma} \tag{2.8}$$

$$\frac{1}{R_t} = \phi_t + \mathbb{E}_t \beta^b_{t+1} \left(\frac{C^b_{t+1}}{C^b_t} \right)^{-\sigma} \tag{2.9}$$

$$(1+\tau_h)\, Q_t C^{s^{-\sigma}}_t = \chi H^{s^{-\sigma_h}}_t + \mathbb{E}_t \beta^s_{t+1} C^{s^{-\sigma}}_{t+1} Q_{t+1} \tag{2.10}$$

$$Q_t C^{b^{-\sigma}}_t = \chi H^{b^{-\sigma_h}}_t + \mathbb{E}_t \beta^b_{t+1} C^{b^{-\sigma}}_{t+1} Q_{t+1} + \Theta_t Q_t \phi_t C^{b^{-\sigma}}_t \tag{2.11}$$

(2.8) and (2.9) are the saver and borrower Euler equations. The borrower equation includes the shadow value of the borrowing constraint, ϕ_t.

(2.10) and (2.11) are the saver and borrower housing demand equations. The left side of each is the utility cost of purchasing an additional unit of housing. The first term on the right is the direct utility benefit directly from additional housing in the current period. The second is the benefit due to the relaxation of the budget constraint in the next period, as housing is a durable good and therefore a form of saving. The borrower's equation has a third term that captures the benefit due to the relaxation of the collateralised borrowing constraint.

Home net foreign assets are defined according to the following.

$$NFA_t = B_t - D_t \tag{2.12}$$

Adding the home saver and borrower budget constraints yields an equation for the current account.

$$CA_t = \Delta NFA_t = TB_t + (R_{t-1} - 1) NFA_{t-1} \tag{2.13}$$

where $TB_t := 2Y - (C^s_t + C^b_t)$.

International financial markets are incomplete, since the only internationally traded asset is a bond. Thus risk markets are absent, and agents are unable to equalise their marginal utility of consumption between possible shock paths. This violates the Korinek (2016) condition of perfect international markets, admitting potential benefits from cooperation. As a result, international policy cooperation can achieve welfare improvements, since policy can be made conditional on the state of the world. Policy can then influence relative prices between states and achieve a greater degree of risk sharing. For example, Spanish policy may be used to help alleviate the effects of the German shock.

The model ignores heterogeneity of real interest rates, which is an alternative explanation for the phenomena of rising house prices and household debt. A single nominal policy rate combined with persistent cross-country inflation differentials leads to lower real rates in higher-inflation members, such as Spain in the early-mid 2000s. Lower real interest rates were associated with softer lending standards and risk-taking by banks (Maddaloni and Peydró, 2011; Jiménez et al., 2014). They have also been linked with housing overvaluation (Hott and Jokipii, 2012).

2.2.2 Equilibrium

An equilibrium consists of decision rules for the 20 endogenous variables C_t^b, C_t^s, C_t^{b*}, C_t^{s*}, H_t^b, H_t^s, H_t^{b*}, H_t^{s*}, D_t, D_t^*, B_t, B_t^*, Q_t, Q_t^*, ϕ_t, ϕ_t^*, R_t, β_t^b, β_t^s, and $T_{s,t}$, given the state of the world at time t indexed by the state variables B_{t-1}, B_{t-1}^*, D_{t-1}, D_{t-1}^*, H_{t-1}^b, H_{t-1}^s, H_{t-1}^{b*}, H_{t-1}^{s*} and the shock u_t. The decision rules solve the equations (2.4), (2.6), (2.7), (2.8), (2.9), (2.10), (2.11) and their foreign equivalents, as well as equations (2.2), (2.3), (2.5), and the following market clearing conditions.

$$4Y = C_t^s + C_t^b + C_t^{s*} + C_t^{b*} \tag{2.14}$$

$$2H = H_t^s + H_t^b \tag{2.15}$$

$$2H = H_t^{s*} + H_t^{b*} \tag{2.16}$$

Equation 2.14 is the goods market clearing condition, allowing overseas lending of the consumption good. By contrast, equations (2.15) and (2.16) reflect the segregation of housing markets along national lines.

2.2.3 Efficient steady state

To rank policy regimes, I measure welfare losses as deviations from a Pareto efficient steady state. It is therefore necessary to ensure this efficiency, which can be done by appropriate choices of the housing and lump sum taxes and subsidies.

An efficient steady state allocation is that which would be chosen by a global planner who is only subject to the market clearing constraints (2.14), (2.15), (2.16). Given a set of Pareto weights $\tilde{\xi}^i$, where $i \in \{s, b, s*, b*\}$, this allocation maximises the following steady-state world utility function.

$$\mathbb{U}^W = \tilde{\xi}^s U(C^s, H^s) + \tilde{\xi}^b U(C^b, H^b) + \tilde{\xi}^{s*} U^*(C^{s*}, H^{s*}) + \tilde{\xi}^{b*} U^*(C^{b*}, H^{b*}) \qquad (2.17)$$

where capital letters without time subscripts denote the steady state of a variable, and where $U(C,H) = \left(\frac{C^{1-\sigma}}{1-\sigma} + \chi \frac{H^{1-\sigma_h}}{1-\sigma_h}\right)$ and $U^*(C,H) = \left(\frac{C^{1-\sigma}}{1-\sigma} + \chi^* \frac{H^{1-\sigma_h}}{1-\sigma_h}\right)$.

The global planner's first order conditions are as follows.

$$\begin{aligned} \tilde{\xi}^s U_{C^s} = \tilde{\xi}^b U_{C^b} = \tilde{\xi}^{s*} U^*_{C^{s*}} = \tilde{\xi}^{b*} U^*_{C^{b*}} &= \mu_c \\ \tilde{\xi}^s U_{H^s} = \tilde{\xi}^b U_{H^b} &= \mu_h \\ \tilde{\xi}^{s*} U^*_{H^{s*}} = \tilde{\xi}^{b*} U^*_{H^{b*}} &= \mu_h^* \end{aligned} \qquad (2.18)$$

Given that the two countries are equally sized, the sum of Pareto weights for each country must be the same. For simplicity it is also assumed that the borrowers are of the same size as the savers, such that we may set $\tilde{\xi}^i = 1 \; \forall \; i$. Further note that $U_C = U_C^*$, since the only country-specific part of the single-period utility function is the housing preference parameter, and the function is additively separable in consumption and housing. Given monotonicity, and combining with the market clearing conditions, the solutions $C^s = C^b = C^{s*} = C^{b*} = Y$, $H^s = H^b = H$, and $H^{s*} = H^{b*} = H$ obtain. That is, in the efficient steady state, consumption and housing are divided evenly within their markets.

The housing and lump sum taxes and subsidy, along with initial debt balances, are chosen to effect the above steady state values for consumption and housing. First, the housing tax that is consistent with both $H^s = H^b$ and $C^s = C^b$ is found by combining the steady state of the borrower and saver housing demand equations.

$$\tau_h = \beta^s - \beta^b - \Theta\phi \tag{2.19}$$

The saver's budget constraint, combined with the efficient steady state consumption and housing amounts, implies the below as the constant component of the lump sum tax. Note that the full lump-sum tax $T_{s,t}$ incorporates the reimbursement of the housing tax, and is specified in Equation 2.5.

$$T_s = B\,(R-1) \tag{2.20}$$

Similarly, the steady state of the borrower's budget constraint implies that the following value of the borrower's lump sum subsidy is consistent with the efficient steady state consumption and housing amounts.

$$S_b = D\,(R-1) \tag{2.21}$$

For simplicity, the model assumes a steady state net foreign assets of $NFA = 0$, meaning that $D = B$. The fact the that borrowing constraint binds in steady state can then be used to derive the steady state value of D consistent with the above.

2.3 Calibration

In the model calibration to Germany and Spain, β_s is set to 0.995 for a real annual interest rate of 2%. β_b is set to 0.99 for greater impatience of borrowers than savers. The further apart the two are, the larger will be the steady state borrowing constraint multiplier. Larger shocks would then be necessary to cause the foreign borrowing constraint to become slack.

The purpose of a positive ψ is to induce stationarity. So as to have a negligible effect on business cycle dynamics, I use the small value of 0.01.

The main combined effect σ and σ_h is on the degree of movement of NFA and the Foreign (Spanish) house price change. More elastic intertemporal substitutability leads to a greater international financial flow and a greater response of Q_t^*. This paper sets σ to 2 and σ_h to 4.

The steady states of the loan-to-value ratio of each country, Θ and $\Theta*$, are set equal to household debt as a portion of household assets in 2001 in Germany and Spain. Similarly, $\frac{D}{8Y}$ and $\frac{D^*}{8Y}$ match the 2001 ratios of household debt to (annual) GDP in each country. These data points combined with the borrower housing demand equations and other variable values pin down the housing preference parameters χ and χ^*.

ω is set at 0.5, which produces a neutral response to the shock by the Home (German) house price. ν is calibrated so as to match the volatility of the changes in Spain's net foreign assets during the early 2000s. The persistence of the shock ρ is set to 0.95, consistent with slow adjustment of preferences.

2.4 Competitive equilibrium

To use the model to simulate the build-up of debt and house prices in Spain, I consider a 30-quarter sequence of shocks of 0.4%, followed by no shock. This has the effect of simulating an outflow of capital from the Home country experiencing the shock, and a run-up in the Foreign house price.

The results in this section are for the "No policy response" macroprudential policy regime. The policy is $\Theta_t = \Theta$ and $\Theta_t^* = \Theta^*$, since both Germany and Spain did not have time-varying macroprudential loan-to-value limits in place

Table 2.1: Parameter values

	Description	Value
β_s	Saver steady state discount factor	0.995
β_b	Borrower steady state discount factor	0.99
σ	Inverse intertemporal elasticity of substitution in consumption	2
σ_h	Inverse intertemporal elasticity of substitution in housing	4
ν	Shock volatility	0.008
ρ	Shock persistence	0.95
ω	Shock incidence proportion on saver	0.5
ψ	Stationarity induction	0.01
$\frac{D}{8Y}$	Germany debt-to-GDP ratio	0.7
$\frac{D^*}{8Y}$	Spain debt-to-GDP ratio	0.55
Θ	Germany steady state LTV ratio	0.44
Θ^*	Spain steady state LTV ratio	0.36
χ	Home housing preference weight	0.0998
χ^*	Foreign housing preference weight	0.1007

in 2001. That is, the prevailing ratios of household debt to assets are assumed to have been the outworkings of other regulations, particularly microprudential lending restrictions which do not take explicit account of macroprudential goals. The assumption of a constant loan-to-value limit is different from assuming no macroprudential policy as the base policy regime.

The model starts from a position in which both Home and Foreign borrowers' borrowing constraints bind. The 30-year shock sequence modelled is insufficient to slacken either borrowing constraint.

Figure 2.2 compares the model results with the motivating data. Overall the model matches the direction, if not the magnitude, of the key variables. Spanish NFA experiences a decline in response to Germans' greater overseas saving imperative. The inflow of capital causes a rise in the Spanish house price, thus replicating the negative correlation between Spanish NFA and house prices. That is accompanied by lower interest rates, mirroring the developments observed in Germany and elsewhere in developed economies at the time, consistent with a global excess of saving. Spanish household debt climbs along with the rise in the value of housing

Figure 2.2: Data vs. modelled shock sequence with no loan-to-value policy response

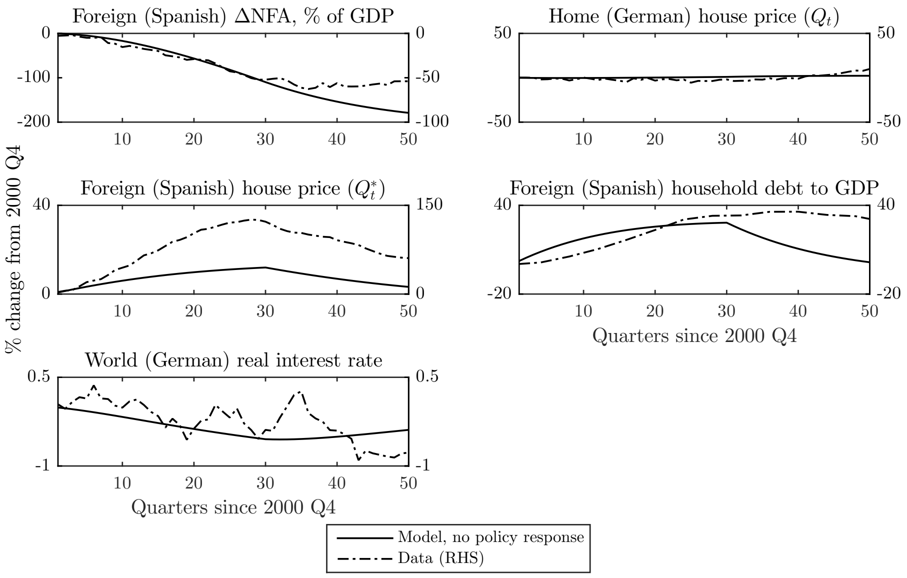

and thus the debt limit, as well as cheaper borrowing. The lack of movement in German house prices is a direct result of the choice of ω.

The model is therefore able to use a single shock to replicate the direction of movements in NFA, house prices in both countries, and Spanish debt. However, a limitation is that the model does not match the magnitude of movements in the Spanish house price and net foreign assets. With a richer model it is possible to more closely match the magnitude of movements; see, for example, Ferrero (2015), which uses two shocks. In the present paper however, I prioritise model parsimony given the objective of optimal policy regime comparison.

The borrowing constraint creates a tension in the German borrower's response to the shock. While he would like to borrow more, he would also like to save less, including reducing his housing stock. Unfortunately for him, as long as the borrowing constraint binds, his housing and debt are proportional to one another. The result is that he does not borrow as much as he would like. The German saver is also affected by the borrower's dilemma. In response to the shock, the saver would like to both lend more to the borrower and purchase more of the domestic housing stock from him. The constraint prevents this, creating a shortage of safe assets in the Home country.

The above mechanism bears similarity to the microfoundations of the global saving glut hypothesis provided in Caballero et al. (2008). In that paper, a shortage of safe assets in emerging economies necessitates their purchase of overseas assets. Similarly in the present paper, the German borrower's desire to own less housing and the saver's desire to own more creates a shortage of domestic debt assets for the German saver to purchase. The more tightly it binds following a shock, the more the country will rely on overseas saving.

There is roughly zero net effect on the Home house price, since the saver's greater desire for housing and the borrower's smaller demand mostly offset. However, due to the net positive change in supply of credit globally, the real interest rate reduces, consistently with the data.

The saving glut emanating from the Home country results in the converse in the Foreign country. The lower interest rate induces the Foreign saver to reduce their bond holdings, and demand more housing as saving. Similarly, the Foreign borrower experiences even stronger desire to increase their debt. However, given the borrowing constraint, they are only able to do so if the value of their housing increases. Both factors mean that a higher Foreign house price is required to clear markets.

In the earlier example, neither of the borrowing constraints ever becomes slack, despite the relatively long period of positive shocks. For the sake of illustrating a scenario of a slack constraint, Figure 2.3 gives the impulse response to a sufficiently large shock, using the Occbin tool (L. Guerrieri and Iacoviello, 2015). The shock size in period 1 is 15%, much larger than the calibrated volatility of the shock of 0.8%. The shock produces a quicker NFA deterioration in Spain than in the previous example. As a result, the upward movement in the Spanish house price is immediate, rather than gradual. This is enough for the value of borrower housing collateral to exceed the borrower's preferred debt level, and the Spanish borrowing constraint

Figure 2.3: Market equilibrium given a 15% shock with no loan-to-value policy response

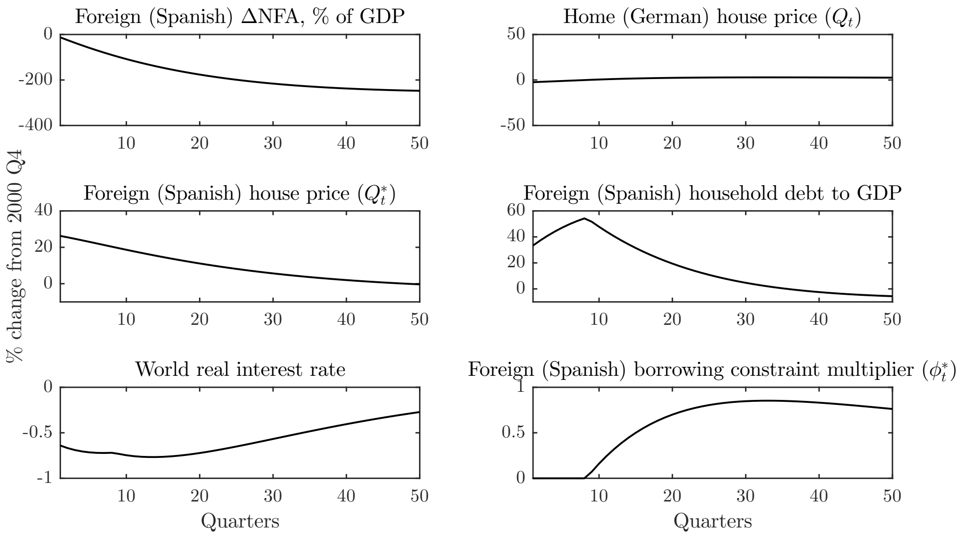

becomes non-binding. This is despite an immediate increase in Spanish borrower debt, in response to the lower world interest rate. Throughout the non-binding period, the Spanish borrower continues to increase their debt. After 8 quarters, this, combined with the decline of the Spanish house price, are sufficient for the constraint to bind again. From that point the Spanish borrower is forced to reduce their debt. The effects of the occasionally binding constraint can be seen in the discontinuity in the world interest rate response, since the constraint, when binding, limits debt demand. Throughout the example, German house prices are largely unaffected, by design through the calibration used.

2.5 Optimal policy

Having established market behaviour under no policy response, I proceed by defining two regimes of optimal macroprudential policy. The first is the global (Ramsey) optimum, in which a single policymaker controls both countries' loan-to-value ratios and aims to maximise global discounted welfare. The Ramsey optimum is interpretable as the cooperative equilibrium, as it is attainable if countries can jointly commit to maximising global welfare. The second is the Nash optimum, in which separate policymakers for each country control the loan-to-value ratio of their own country, and seek only to maximise the discounted welfare of their own households. Each Nash policymaker internalises their impact on the world interest rate, but takes the other policymaker's policy as given, such that they best-respond to each other. The Nash equilibrium is interpretable as a non-cooperative equilibrium, corresponding to the status quo in macroprudential policymaking. The same calibration is used for all policy regimes, resulting in the same steady state.

2.5.1 Welfare function approximation

The second-order approximations to the welfare functions of the different policymakers, derived here, are useful in explaining policymaker behaviour.

The global policymaker optimises a discounted sum of successive single-period welfare functions of the following form, which is a time-subscripted version of (2.17).

$$\mathbb{U}_t^R = \tilde{\xi}^s U(C_t^s, H_t^s) + \tilde{\xi}^b U(C_t^b, H_t^b) + \tilde{\xi}^{s*} U^*(C_t^{s*}, H_t^{s*}) + \tilde{\xi}^{b*} U^*(C_t^{b*}, H_t^{b*}) \quad (2.22)$$

Again setting the Pareto weights to 1, and incorporating the market clearing constraints for housing and consumption, this function can be approximated to second order as follows (derivation is in the Appendix).

$$\mathbb{U}_t^R - \mathbb{U}^R \approx -\frac{1}{2}\left[\Lambda_c \left(\tilde{c}_t^2 + (\tilde{c}_t^*)^2 + \frac{1}{2}\left(\tilde{c}_t^W\right)^2\right) + \Lambda_h \left(\chi \tilde{h}_t^2 + \chi^* \left(\tilde{h}_t^*\right)^2\right)\right] \quad (2.23)$$

where $\Lambda_c = \frac{1}{2}\sigma C^{1-\sigma}$, $\Lambda_h = \frac{1}{2}\sigma_h H^{1-\sigma_h}$. In Equation 2.23, $\tilde{c}_t = c_t^b - c_t^s$ denotes the Home consumption gap, $\tilde{h}_t = h_t^b - h_t^s$ denotes the Home housing gap, $\tilde{c}_t^W = c_t^b + c_t^s - \left(c_t^{b*} + c_t^{s*}\right)$ denotes the World consumption gap, and lower case variables indicate percentage deviations from steady state such that $x_t = \frac{X_t - X}{X}$.

Equation 2.23 therefore shows that optimal cooperative policy aims to minimise a weighted sum of inequality measures. Policymaker welfare depends negatively on both within-country and international inequality, measuring these by squared deviations of the 'gap' variables from their perfect-equality steady state. Deviations from the within-country equality are measured by the consumption gaps $\tilde{c}_t$ and $\tilde{c}_t^*$, and by the housing gaps $\tilde{h}_t$ and $\tilde{h}_t^*$. International inequality is measured by the world consumption gap $\tilde{c}_t^W$. The weighted sum that the policymaker minimises is influenced by the marginal utilities of consumption and housing.

By contrast, the Nash Home policymaker maximises only Home agent utility $\mathbb{U}_t^N$, and the Nash Foreign policymaker maximises only Foreign agent utility $\mathbb{U}_t^{N*}$.

$$\mathbb{U}_t^N = \tilde{\xi}^s U(C_t^s, H_t^s) + \tilde{\xi}^b U(C_t^b, H_t^b) \tag{2.24}$$

$$\mathbb{U}_t^{N*} = \tilde{\xi}^{s*} U(C_t^{s*}, H_t^{s*}) + \tilde{\xi}^{b*} U(C_t^{b*}, H_t^{b*}) \tag{2.25}$$

$$\mathbb{U}_t^N - \mathbb{U}^N \approx \Lambda_c \left[\frac{2}{\sigma}(\hat{c}_t^b + \hat{c}_t^s) - \frac{1}{2}\left(\tilde{c}_t^2 + \frac{1}{4}\left(\tilde{c}_t^W\right)^2 \right) \right] - \frac{1}{2}\Lambda_h \chi \tilde{h}_t^2 \tag{2.26}$$

$$\mathbb{U}_t^{N*} - \mathbb{U}^{N*} \approx \Lambda_c \left[\frac{2}{\sigma}(\hat{c}_t^{b*} + \hat{c}_t^{s*}) - \frac{1}{2}\left((\tilde{c}_t^*)^2 + \frac{1}{4}\left(\tilde{c}_t^W\right)^2 \right) \right] - \frac{1}{2}\Lambda_h \chi^* \left(\tilde{h}_t^*\right)^2 \tag{2.27}$$

A key difference is that, whereas the cooperative policymaker aims to minimise squared deviations from steady state, non-cooperative policymakers seek to increase domestic consumption ($\hat{c}_t^b + \hat{c}_t^s$ or $\hat{c}_t^{b*} + \hat{c}_t^{s*}$), within a single period. Of course, to consume more in a given period, a nation must borrow more or save less in that period, and this is not internalised in the national policymaker single-period welfare functions. By contrast, the cooperative policymaker faces fixed total resources in a single period, with no way to defer consumption or bring it forward.

The linear terms in the national (non-cooperative) welfare functions take on special relevance in a scenario of a shock that lowers the interest rate. The Home policymaker will try to borrow more, by increasing the world consumption gap in the near term. Similarly, the Foreign policymaker will try to lower the world consumption gap. Loan-to-value ratio policy is the only channel in the model through which this tug-of-war can play out. Prudent loan-to-value ratio policy will therefore be sacrificed.

The welfare functions thereby illustrate the relevance of the violation of the Korinek (2016) condition of sufficient external instruments. With a tax/subsidy on NFA in each country, loan-to-value ratios would be free to target the domestic consumption and housing gaps, while the Nash policymakers would use the NFA tax to try to target the world consumption gap in their favour. Such beggar-thy-neighbour behaviour may result in a different equilibrium interest rate, but the outcome could still be Pareto efficient, with no improvements available from international cooperation. By contrast, with insufficient international instruments, distortion of domestic policies such as the loan-to-value ratio is a likely outcome.

2.5.2 Cooperative and noncooperative policy

Due to the presence of the first-order term in the Nash welfare function approximations, I take a quantitative approach to optimal policy modelling. I use the method and toolbox described in Bodenstein et al. (2019). The method computes an open-loop Nash equilibrium, where policymakers choose full paths of the instrument at time $t = 0$. To avoid time inconsistencies, the method uses the timeless perspective. The result is a fixed point of best responses; each policymaker maximises its country's discounted summed welfare taking the other policymaker's path as given. The toolbox analytically derives the policymaker first-order conditions for each policymaker in the case of the Nash equilibrium (non-cooperative policy regime), and for the global policymaker in the case of the Ramsey equilibrium (cooperative policy regime). The first-order conditions maximise the welfare of the relevant policymaker(s) subject to the implementability constraints given by the competitive equilibrium conditions. The policymaker first-order conditions, combined with the competitive equilibrium conditions, form the system of equations that describe the endogenous variables' equilibrium responses to shocks over time under the policy regime to which the method has been applied. The solution to the system can then be found by a quantitative method.

Welfare losses must be discounted using a single discount factor β^p, which I set equal to the saver's steady state discount factor β^s in both the cooperative and non-cooperative policy regimes. It is noteworthy that the assumed policymaker discount factor differs from the steady state of the borrower discount factor. In effect, it is assumed that the policymaker adopts a prescriptive stance towards the more impatient households. Alternatively, the policymaker behaviour is interpretable as an approximation along the lines of P. Benigno et al. (2020) which equalises β^b and β^s for the purpose of policy analysis. The policymaker discount factor also differs from German household discount factors in that the latter are subject to both the preference shock and the endogenous component that induces stationarity, as per equations (2.2) and (2.3). It is deliberate that policymakers are not affected by the preference shock, which is intended as a reduced form method of inducing the financial flows that are the subject of this study. Similarly, stationarity induction does not require an endogenous policymaker discount factor.

Figure 2.4 plots variable movements in response to the same sequence of shocks to the German $\beta_t^s - \beta_t^b$ differential used in Section 2.4. The figure compares the cooperative and non-cooperative policy regimes with that of no policy response which was displayed earlier.

As expected, cooperative macroprudential policy offers improvements upon the no policy response regime. Instead of allowing borrower debt to accumulate, the policymaker lowers Spain's loan-to-value ratio, and subsequently Germany's as well. There is little effect on net foreign assets or house prices in either country. However, Spanish borrower's debt is significantly reduced. Overall, both countries contribute successfully to macroprudential policy's goal of mitigating the overborrowing caused by the house price externality.

The non-cooperative regime, on the other hand, results in an equilibrium in which loan-to-value ratios are higher than in the cooperative regime. In fact, with the exception of Spain in the earlier part of the boom, the two loan-to-value ratios are even higher than in the no policy response regime, suggesting the non-cooperative regime may be worse than no policy response. The loose loan-to-value ratios are

Figure 2.4: Responses to a sequence of shocks under different policy regimes

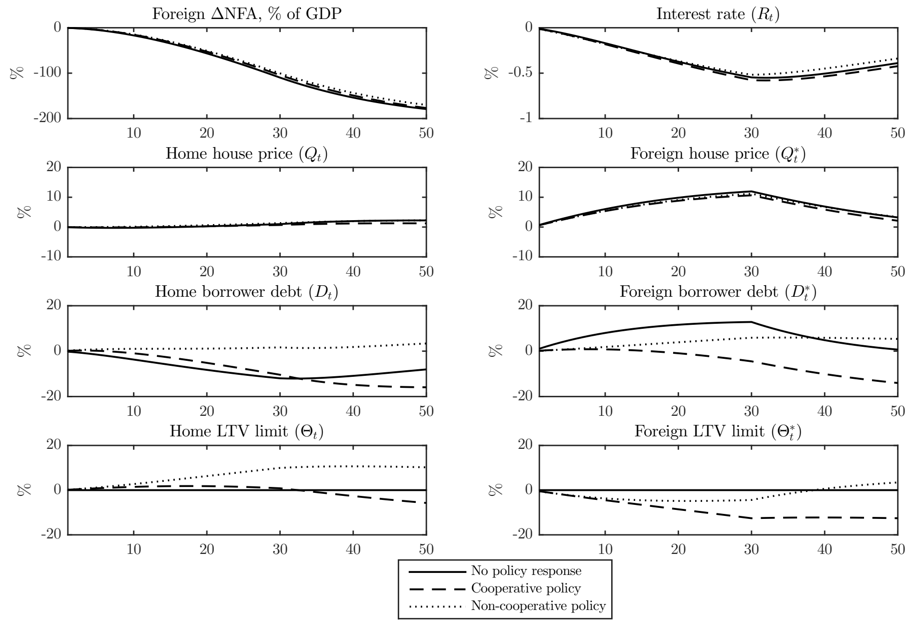

Charts show percentage deviations from steady state.

unable to achieve the normal aims of macroprudential policy. Instead of German and Spanish borrower debt declining as under cooperative policy, they increase. By the end of the boom they are both higher than under the no policy response regime.

Although the results in Figure 2.4 suggest benefits of cooperation, a ranking of policy regimes cannot be devised by examining a single shock path. Rather, it is necessary to measure the expectation of the discounted sum of future welfare under each policy regime. Table 2.2 shows the effect of cooperative and non-cooperative policy on summed discounted welfare globally and in each country, relative to no policy response.

The welfare analysis assumes that both borrowing constraints always bind. This is unproblematic in the results in Figure 2.4 where neither constraint needs to become slack. However, the simulations used to produce the results in Table 2.2 do contain instances where the multipliers ϕ_t and ϕ_t^* become negative. Given the additional challenges of combining the nonlinear optimal policy solution with occasionally

Table 2.2: Welfare effects of cooperative and non-cooperative policy regimes.

	Welfare improvement due to active policy		Preference ordering		
	Cooperative	Non-cooperative	1st	2nd	3rd
Global	2.38%	-0.41%	Cooperative	No policy response	Non-cooperative
Germany	0.78%	0.71%	Cooperative	Non-cooperative	No policy response
Spain	10.36%	-7.25%	Cooperative	No policy response	Non-cooperative

Welfare effects of active policy (i.e. the cooperative and non-cooperative regimes) are expressed as a portion of the welfare effect of the shock under the No policy response regime.

binding constraints, refinement of this aspect of the model is left to further work.

At the global level, as expected, benefits exist from cooperation. Further, the global welfare gains from cooperation are shared, as both countries prefer the cooperative policy regime to the others.

More notably however, the no policy response regime is better for global welfare than the non-cooperative regime. The competitive behaviour of policymakers under the non-cooperative regime is detrimental enough to more than offset the benefits of having an available macroprudential policy instrument.

The result can be analysed in light of the violation of the three Korinek (2016) conditions required for international cooperation to be unnecessary. Firstly, the condition of complete external instruments highlights that non-cooperative policymakers distort loan-to-value ratios to pursue their target for external variables. For example, from Equations (2.26) and (2.27), in response to a shock that lowers the interest rate, each Nash policymaker would like to tilt the world consumption gap in their favour in the short term. This means greater borrowing, which can be achieved by a loan-to-value ratio increase, since that reduces the need to save internationally by creating more domestic debt assets. Of course, one country's borrowing is the other's saving, so this is a zero-sum game. The unfortunate side effect is that the instrument used, loan-to-value ratios, are moved away from the optimal level for responding to the domestic borrowing constraint externalities.

Secondly, the condition of competitive policymaker behaviour is violated by their power over the interest rate. Market power in this model leads to a country's desire to move its NFA toward zero, regardless of the sign. After all, less demand for capital by a net borrower will lower the interest rate they pay, and less supply of capital by a net saver will increase the interest rate they receive. With the shock in question, Germany becomes the saver and Spain the borrower. Therefore non-competitive behaviour would tend to increase the German loan-to-value ratio and decrease the Spanish loan-to-value ratio in the non-cooperative regime, relative to cooperative. The absence of external instruments is again relevant, as loan-to-value ratios are distorted by policymakers' pursuit of monopoly rents.

Thirdly, the condition of perfect international markets is violated by the absence of risk markets; nations may only trade bonds internationally. This is primarily a problem for Germany, who bears the brunt of the shock studied. Cooperation can alleviate this through a form of burden sharing in which Spain's policy instrument is used to help respond to the German shock. In effect, cooperative policy is able to differentially influence outcomes according to which state of the world is realised. A degree of risk sharing is thereby achieved that is unattainable in the non-cooperative and no policy response regimes. Germany's rational expectation of greater risk sharing leads to greater ex ante welfare.

The result that no policy response is preferred to the non-cooperative regime is particularly sensitive to the calibrated value of ω. Table 2.3 shows that when ω is not close to 0.5, the non-cooperative policy regime is preferred to the regime of no policy response.

Concretely, from equations (2.2) and (2.3), while the shock increases the gap between β_t^s and β_t^b, ω governs how much of the increased gap is due to a higher β_t^s versus a lower β_t^b. If $\omega > 0.5$ then the shock increases overall patience in the Home country. Since patience leads to demand for housing as a saving vehicle, house prices will increase if $\omega > 0.5$, and decrease if $\omega < 0.5$, as shown in Figure 2.5.

Table 2.3: Sensitivity of global preference ordering to ω.

ω	Welfare improvement due to active policy		Preference ordering		
	Cooperative	Non-cooperative	1st	2nd	3rd
0	32.36%	31.74%	Cooperative	Non-cooperative	No policy response
0.2	16.81%	16.45%	Cooperative	Non-cooperative	No policy response
0.4	3.78%	1.84%	Cooperative	Non-cooperative	No policy response
0.5	2.38%	-0.41%	Cooperative	No policy response	Non-cooperative
0.6	3.39%	-0.19%	Cooperative	No policy response	Non-cooperative
0.8	7.41%	3.05%	Cooperative	Non-cooperative	No policy response
1	11.14%	6.43%	Cooperative	Non-cooperative	No policy response

Welfare effects of active policy (i.e. the cooperative and non-cooperative regimes) are expressed as a portion of the welfare effect of the shock under the No policy response regime.

Figure 2.5: Effect of shock incidence proportion on saver vs. borrower (ω parameter)

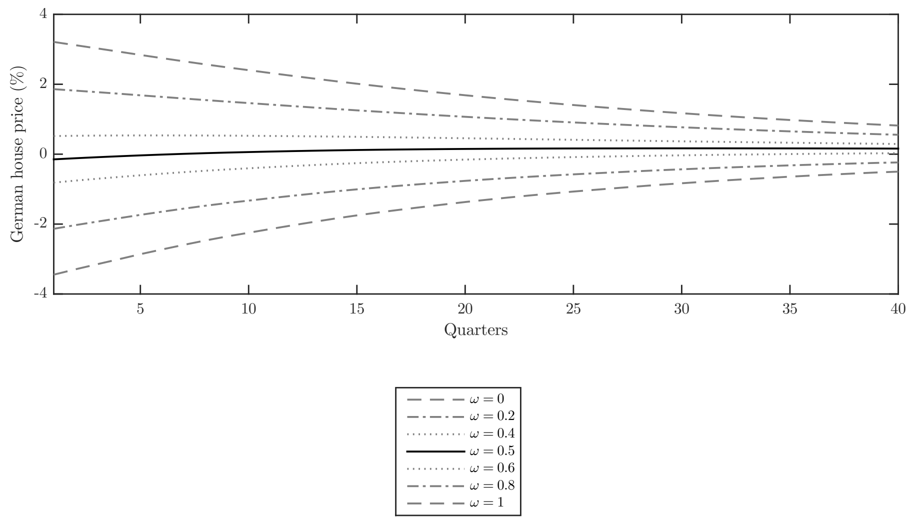

Percentage deviations from steady state in response to a 1-standard-deviation shock.

Since ω governs the response of house prices to the shock, it determines the extent of the pecuniary externality from the dependence of the borrowing constraint on house prices. The more severe the externality, the more benefit available from macroprudential policy, increasing the attraction of the non-cooperative solution despite its drawbacks.

In the baseline calibration, ω was set to 0.5 since German house prices changed little during the early 2000s. However, referring back to Figure 2.1, other "capital-providing" countries with positive current accounts experienced greater movement in

Table 2.4: Robustness of global preference ordering over policy regimes

	Baseline	Changed	Welfare improvement due to active policy: Cooperative	Non-cooperative	Preference ordering: 1st	2nd	3rd
β_b	0.99	0.995	2.46%	-0.15%	Cooperative	No policy response	Non-cooperative
β_b	0.99	0.98	4.25%	1.18%	Cooperative	Non-cooperative	No policy response
β_s, β_b	0.995, 0.99	0.99, 0.98	4.1%	1.08%	Cooperative	Non-cooperative	No policy response
σ	2	0.5	0.97%	-4.14%	Cooperative	No policy response	Non-cooperative
σ	2	1	1.35%	-2.79%	Cooperative	No policy response	Non-cooperative
σ	2	4	4.06%	2.57%	Cooperative	Non-cooperative	No policy response
σ	2	10	6.23%	5.97%	Cooperative	Non-cooperative	No policy response
σ_h	4	0.5	5.63%	5.46%	Cooperative	Non-cooperative	No policy response
σ_h	4	1	1.92%	1.28%	Cooperative	Non-cooperative	No policy response
σ_h	4	2	1.1%	-0.44%	Cooperative	No policy response	Non-cooperative
σ_h	4	10	6.87%	2.01%	Cooperative	Non-cooperative	No policy response
ρ_c	0.95	0.9	1.91%	0.96%	Cooperative	Non-cooperative	No policy response
ρ_c	0.95	0.98	1.94%	-3.02%	Cooperative	No policy response	Non-cooperative
ν	0.0082	0.003	2.45%	0%	Cooperative	No policy response	Non-cooperative
ν	0.0082	0.013	2.42%	-0.43%	Cooperative	No policy response	Non-cooperative
Θ	0.44	0.2	1.22%	0.06%	Cooperative	Non-cooperative	No policy response
Θ	0.44	0.7	4.59%	-0.1%	Cooperative	No policy response	Non-cooperative
Θ^*	0.36	0.1	1.22%	-0.97%	Cooperative	No policy response	Non-cooperative
Θ^*	0.36	0.6	3.85%	0.82%	Cooperative	Non-cooperative	No policy response
$\frac{D}{8Y}$	0.7	0.175	4.51%	-0.51%	Cooperative	No policy response	Non-cooperative
$\frac{D}{8Y}$	0.7	1.4	2.15%	0.56%	Cooperative	Non-cooperative	No policy response
$\frac{D^*}{8Y}$	0.55	0.1375	0.68%	-2.48%	Cooperative	No policy response	Non-cooperative
$\frac{D^*}{8Y}$	0.55	1.1	5.57%	3.33%	Cooperative	Non-cooperative	No policy response
ω	0.5	0	32.36%	31.74%	Cooperative	Non-cooperative	No policy response
ω	0.5	0.2	16.81%	16.45%	Cooperative	Non-cooperative	No policy response
ω	0.5	0.4	3.78%	1.84%	Cooperative	Non-cooperative	No policy response
ω	0.5	0.6	3.39%	-0.19%	Cooperative	No policy response	Non-cooperative
ω	0.5	0.8	7.41%	3.05%	Cooperative	Non-cooperative	No policy response
ω	0.5	1	11.14%	6.43%	Cooperative	Non-cooperative	No policy response

Welfare effects of active policy (i.e. the cooperative and non-cooperative regimes) are expressed as a portion of the welfare effect of the shock under the No policy response regime.

house prices. Such countries may experience more amplification effects and over- or under-borrowing, and therefore greater benefits of a macroprudential policy response.

The policy regime rankings are also sensitive to other aspects of the model's calibration. Table 2.4, and Tables 2.5, and 2.6 in the Appendix, provide the same results for a number of alternative parameter values. Table 2.4 shows that, at the global level, the result that no policy response is better than non-cooperation is obtained in many cases, but not all.

2.6 Conclusion

The paper uses the setting of Germany and Spain as a case study of a global saving glut, such as that experienced in the early and mid 2000s, where capital account surpluses in some countries coincided with asset bubbles, particularly housing bubbles, in others. The two country model developed tracks one way in which financial frictions and incomplete markets can generate a shortage of safe assets, leading to this global saving glut. With these global imbalances, it is not clear that macroprudential policy goals can be achieved without international cooperation.

The result of the study is that, in the context of Germany and Spain's loan-to-value ratio policy, a non-cooperative policy regime would likely have been worse than a static regime of no policy response. Harmful competition among national policymakers would likely have outweighed the macroprudential benefits. The model results indicate that this is more likely to happen where asset prices in the capital-providing country move little in response to the shock, as was the case for German house prices.

The analysis highlights that, given incomplete markets and external instruments, non-cooperative policymakers will experience pressure to use loan-to-value ratio policy to pursue beggar-thy-neighbour objectives. In the modelled scenario, they competitively loosen their ratio to attract capital in the short term. Policymakers are then unable to optimally counteract the pecuniary externality from the collateralised borrowing constraint. While some macroprudential benefits may remain, in the calibration studied these are not enough to offset the harmful effects.

It is plausible that the results may extend to other countries insofar as the capital exporting countries experience little house price movement. While all could have benefited from cooperative macroprudential policy instead of no policy response, some may have found that non-cooperative policy's harms were outweighed by the macroprudential benefits.

The study does not establish clear predictions about other forms of borrower-targeted macroprudential policy. However, it may be worthwhile for future work to investigate whether similar mechanisms operate on, for example, countercyclical capital buffer policy. Similar results may be obtained if other such policies have the potential to cause outflows or inflows of capital, thereby creating beggar-thy-neighbour incentives for national policymakers. In particular, if the welfare impact of that behaviour outweighs that of correcting the overborrowing externality, then cooperation is likely to be needed to realise benefits from macroprudential policy.

The model's parsimony, and consequent use of a single shock, limits its ability to match the magnitude of Spanish house prices and net foreign assets during the early to mid 2000s. However, the model does match the direction of movement of these and other economic variables, and it is on this basis that the results are obtained. Further work could incorporate extra shocks or frictions to improve on this dimension.

Appendices to Chapter 2

2.A Derivation of welfare second-order approximations

Global (Ramsey) policymaker

The global policymaker's period utility function is

$$\mathbb{U}_t^R = \sum_{i \in \{b,s,b^*,s^*\}} U^i(C_t^i, H_t^i)$$

where $U^b(C,H) = U^s(C,H) = \frac{C^{1-\sigma}}{1-\sigma} + \chi \frac{H^{1-\sigma_h}}{1-\sigma_h}$ and $U^{b*}(C,H) = U^{s*}(C,H) = \frac{C^{1-\sigma}}{1-\sigma} + \chi^* \frac{H^{1-\sigma_h}}{1-\sigma_h}$.

The function may be approximated to second order around the steady state as follows.

$$\mathbb{U}_t^R - \mathbb{U}^R \approx \sum_{i \in \{b,s,b^*,s^*\}} U_C^i(C_t^i - C) + U_H^i(H_t^i - H) + \frac{1}{2} U_{CC}^i (C_t^i - C)^2 + \frac{1}{2} U_{HH}^i (H_t^i - H)^2$$

The first order terms vanish due to the model's market clearing conditions (Equations 2.14, 2.15, and 2.16).

The first and second derivatives of the function $U^i(C,H)$, taken at the steady state, are[1]

$$\begin{aligned}
U_C^i &= C^{-\sigma} \\
U_H^i &= \chi^i H^{-\sigma_h} \\
U_{CC}^i &= -\sigma C^{-\sigma-1} = -\sigma C^{1-\sigma} \frac{1}{C^2} \\
U_{HH}^i &= -\sigma_h \chi^i H^{-\sigma_h - 1} = -\sigma_h \chi^i H^{1-\sigma_h} \frac{1}{H^2}
\end{aligned}$$

[1] χ^i varies by country only, such that $\chi^i = \chi$ if $i \in \{s,b\}$ and $\chi^i = \chi^*$ if $i \in \{s^*,b^*\}$.

Plugging in and expressing the consumption and housing variables as deviations from steady state:

$$\mathbb{U}_t^R - \mathbb{U}^R \approx \sum_{i\in\{b,s,b^*,s^*\}} -\frac{1}{2}\sigma C^{1-\sigma}(\hat{c}_t^i)^2 - \frac{1}{2}\sigma_h \chi^i H^{1-\sigma_h}(\hat{h}_t^i)^2 \tag{2.28}$$

Let the gap variables $\tilde{c}_t$, $\tilde{h}_t$, $\tilde{c}_t^*$, $\tilde{h}_t^*$, and $\tilde{c}_t^W$ be defined as in Section 2.5.1. Also note that, due to the first order approximation of the housing market clearing conditions, $\tilde{h}_t = 2\hat{h}_t^b = 2\hat{h}_t^s$ and therefore $\frac{1}{4}\tilde{h}_t^2 = (\hat{h}_t^b)^2 = (\hat{h}_t^s)^2$. Thus the housing deviations from steady state in Equation 2.28 can be expressed in terms of the domestic housing gap.

Now consider the sum of consumption deviations, and note that the following identity holds.

$$(\hat{c}_t^b)^2 + (\hat{c}_t^s)^2 = \frac{1}{2}\left((\hat{c}_t^b + \hat{c}_t^s)^2 + (\hat{c}_t^b - \hat{c}_t^s)^2\right)$$

Adding this equation to its foreign equivalent produces

$$\sum_{i\in\{b,s,b^*,s^*\}} (\hat{c}_t^i)^2 = \frac{1}{2}\left((\hat{c}_t^b + \hat{c}_t^s)^2 + (\hat{c}_t^{b*} + \hat{c}_t^{s*})^2 + \tilde{c}_t^2 + (\tilde{c}_t^*)^2\right)$$

Now, from the market clearing condition, $(\hat{c}_t^b+\hat{c}_t^s)^2+(\hat{c}_t^{b*}+\hat{c}_t^{s*})^2 = 2(\hat{c}_t^b+\hat{c}_t^s)^2$. Further,

$$\begin{aligned}
\tilde{c}_t^W &= \hat{c}_t^b + \hat{c}_t^s - (\hat{c}_t^{b*} + \hat{c}_t^{s*}) \\
&= 2(\hat{c}_t^b + \hat{c}_t^s) \text{ again from market clearing} \\
\frac{1}{2}(\tilde{c}_t^W)^2 &= 2(\hat{c}_t^b + \hat{c}_t^s)^2
\end{aligned}$$

Combining all the gap variable results and plugging in to Equation 2.28, along with the definitions of Λ_c and Λ_h (given in Section 2.5.1), gives the second-order approximation to the global policymaker's welfare function (Equation 2.23).

National (Nash) policymaker

The national policymaker's period utility function is, for the Home country,

$$\mathbb{U}_t^N = \sum_{i\in\{b,s\}} U^i(C_t^i, H_t^i)$$

The second order approximation is

$$\mathbb{U}_t^N - \mathbb{U}^N \approx \sum_{i \in \{b,s\}} U_C^i (C_t^i - C) + U_H^i (H_t^i - H) + \frac{1}{2} U_{CC}^i (C_t^i - C)^2 + \frac{1}{2} U_{HH}^i (H_t^i - H)^2$$

This time, only the housing first-order terms disappear. Applying the derivatives of $U^i(C, H)$ and converting to steady state deviations gives

$$\mathbb{U}_t^N - \mathbb{U}^N \approx C^{1-\sigma}(\hat{c}_t^b + \hat{c}_t^s) - \frac{1}{2}\sigma C^{1-\sigma}\left((\hat{c}_t^b)^2 + (\hat{c}_t^s)^2\right) - \frac{1}{2}\sigma_h H^{1-\sigma_h}\left((\hat{h}_t^b)^2 + (\hat{h}_t^s)^2\right)$$

The sum of the squared housing deviations is equal to $\frac{1}{2}\tilde{h}_t^2$. Further,

$$\begin{aligned}(\hat{c}_t^b)^2 + (\hat{c}_t^s)^2 &= \frac{1}{2}\left((\hat{c}_t^b + \hat{c}_t^s)^2 + (\hat{c}_t^b - \hat{c}_t^s)^2\right) \text{ by identity} \\ &= \frac{1}{2}\left(\left(\frac{1}{2}\tilde{c}_t^W\right)^2 + \tilde{c}_t^2\right)\end{aligned}$$

Combining these results produces the second-order approximation to the Home policymaker's welfare function (Equation 2.26). The Foreign equivalent can be similarly derived.

Using the first-order approximated market clearing conditions, the linear term could be expressed as $\hat{c}_t^b + \hat{c}_t^s = \frac{1}{2}\tilde{c}_t^W$. However, this would introduce second-order errors, which would invalidate the equation as a second-order approximation to the national policymaker's welfare function. The above workings use first-order approximations only to derive squared terms, so that any second-order errors become higher-order.

2.B Robustness of preference orderings

Table 2.5: Robustness of Germany's preference ordering over policy regimes

			Welfare improvement due to active policy		Preference ordering		
	Baseline	Changed	Cooperative	Non-cooperative	1st	2nd	3rd
β_b	0.99	0.995	-0.4%	0%	No policy response	Non-cooperative	Cooperative
β_b	0.99	0.98	1.83%	1.6%	Cooperative	Non-cooperative	No policy response
β_s, β_b	$0.995, 0.99$	$0.99, 0.98$	0.92%	0.79%	Cooperative	Non-cooperative	No policy response
σ	2	0.5	2.16%	2.79%	Non-cooperative	Cooperative	No policy response
σ	2	1	1.21%	1.37%	Non-cooperative	Cooperative	No policy response
σ	2	4	0.75%	0.75%	Cooperative	Non-cooperative	No policy response
σ	2	10	1.06%	1.06%	Cooperative	Non-cooperative	No policy response
σ_h	4	0.5	7.19%	7.08%	Cooperative	Non-cooperative	No policy response
σ_h	4	1	3.51%	3.33%	Cooperative	Non-cooperative	No policy response
σ_h	4	2	1.7%	1.5%	Cooperative	Non-cooperative	No policy response
σ_h	4	10	0.46%	0.54%	Non-cooperative	Cooperative	No policy response
ρ_c	0.95	0.9	0.45%	0.45%	Cooperative	Non-cooperative	No policy response
ρ_c	0.95	0.98	1.74%	2.34%	Non-cooperative	Cooperative	No policy response
ν	0.0082	0.003	1.06%	1.06%	Cooperative	Non-cooperative	No policy response
ν	0.0082	0.013	0.73%	0.68%	Cooperative	Non-cooperative	No policy response
Θ	0.44	0.2	1.48%	1.12%	Cooperative	Non-cooperative	No policy response
Θ	0.44	0.7	-0.88%	-0.24%	No policy response	Non-cooperative	Cooperative
Θ^*	0.36	0.1	0.21%	0.91%	Non-cooperative	Cooperative	No policy response
Θ^*	0.36	0.6	0.86%	0.43%	Cooperative	Non-cooperative	No policy response
$\frac{D}{8Y}$	0.7	0.175	4.37%	3.89%	Cooperative	Non-cooperative	No policy response
$\frac{D}{8Y}$	0.7	1.4	-0.51%	-0.06%	No policy response	Non-cooperative	Cooperative
$\frac{D^*}{8Y}$	0.55	0.1375	-1.24%	-1.31%	No policy response	Cooperative	Non-cooperative
$\frac{D^*}{8Y}$	0.55	1.1	2.22%	2.22%	Cooperative	Non-cooperative	No policy response
ω	0.5	0	27.35%	27.35%	Cooperative	Non-cooperative	No policy response
ω	0.5	0.2	17.14%	17.02%	Cooperative	Non-cooperative	No policy response
ω	0.5	0.4	3.66%	3.57%	Cooperative	Non-cooperative	No policy response
ω	0.5	0.6	-0.27%	-0.22%	No policy response	Non-cooperative	Cooperative
ω	0.5	0.8	0.32%	0.44%	Non-cooperative	Cooperative	No policy response
ω	0.5	1	1.69%	1.85%	Non-cooperative	Cooperative	No policy response

Welfare effects of active policy (i.e. the cooperative and non-cooperative regimes) are expressed as a portion of the welfare effect of the shock under the No policy response regime.

Table 2.6: Robustness of Spain's preference ordering over policy regimes

	Baseline	Changed	Welfare improvement due to active policy Cooperative	Non-cooperative	Preference ordering 1st	2nd	3rd
β_b	0.99	0.995	23.12%	-0.58%	Cooperative	No policy response	Non-cooperative
β_b	0.99	0.98	11.11%	-4.35%	Cooperative	No policy response	Non-cooperative
β_s, β_b	$0.995, 0.99$	$0.99, 0.98$	18.51%	0.97%	Cooperative	Non-cooperative	No policy response
σ	2	0.5	-6.21%	-26.07%	No policy response	Cooperative	Non-cooperative
σ	2	1	-0.63%	-18.95%	No policy response	Cooperative	Non-cooperative
σ	2	4	43.86%	24.56%	Cooperative	Non-cooperative	No policy response
σ	2	10	237.5%	225%	Cooperative	Non-cooperative	No policy response
σ_h	4	0.5	-38.89%	-40%	No policy response	Cooperative	Non-cooperative
σ_h	4	1	-23.39%	-29.03%	No policy response	Cooperative	Non-cooperative
σ_h	4	2	-6.79%	-17.9%	No policy response	Cooperative	Non-cooperative
σ_h	4	10	32.39%	6.57%	Cooperative	Non-cooperative	No policy response
ρ_c	0.95	0.9	20%	6.67%	Cooperative	Non-cooperative	No policy response
ρ_c	0.95	0.98	-1.09%	-19.56%	No policy response	Cooperative	Non-cooperative
ν	0.0082	0.003	7.69%	-7.69%	Cooperative	No policy response	Non-cooperative
ν	0.0082	0.013	9.88%	-7.61%	Cooperative	No policy response	Non-cooperative
Θ	0.44	0.2	-3.91%	-14.06%	No policy response	Cooperative	Non-cooperative
Θ	0.44	0.7	25.76%	0.87%	Cooperative	Non-cooperative	No policy response
Θ^*	0.36	0.1	6.25%	-12.5%	Cooperative	No policy response	Non-cooperative
Θ^*	0.36	0.6	20.79%	2.81%	Cooperative	Non-cooperative	No policy response
$\frac{D}{8Y}$	0.7	0.175	-3.86%	-18.95%	No policy response	Cooperative	Non-cooperative
$\frac{D}{8Y}$	0.7	1.4	28.68%	6.62%	Cooperative	Non-cooperative	No policy response
$\frac{D^*}{8Y}$	0.55	0.1375	12.38%	-5.45%	Cooperative	No policy response	Non-cooperative
$\frac{D^*}{8Y}$	0.55	1.1	27.4%	8.22%	Cooperative	Non-cooperative	No policy response
ω	0.5	0	600%	522.22%	Cooperative	Non-cooperative	No policy response
ω	0.5	0.2	-31.58%	-36.84%	No policy response	Cooperative	Non-cooperative
ω	0.5	0.4	-4.39%	-19.3%	No policy response	Cooperative	Non-cooperative
ω	0.5	0.6	20.07%	1.02%	Cooperative	Non-cooperative	No policy response
ω	0.5	0.8	32.44%	11.59%	Cooperative	Non-cooperative	No policy response
ω	0.5	1	40.2%	18.4%	Cooperative	Non-cooperative	No policy response

Welfare effects of active policy (i.e. the cooperative and non-cooperative regimes) are expressed as a portion of the welfare effect of the shock under the No policy response regime.

3

Tracing the learning curve: On cryptocurrency prices, volatility, and eventual adoption

Abstract Rapid growth of cryptocurrency prices has polarised opinions about their fundamental value. This paper's model uses standard financial theory to replicate this trend. In the model, a cryptocurrency's price includes a discount for uncertainty about the extent of eventual adoption. However, investors gradually learn about eventual adoption by observing noisy signals. Each signal represents emerging information about the cryptocurrency's development and the economy's reactions to it. As signals accumulate, the perceived uncertainty reduces, and the discount diminishes. Moreover, with more information, further signals have less price impact, reducing variance over time, which is also consistent with observation. The model offers an intuitive explanation of why the high risk of cryptocurrency investment has been rewarded. New information leads to revisions to estimates of eventual adoption, which may raise or lower the price (risk). However, the new information also reduces the uncertainty discount, which always raises the price (reward).

3.1 Introduction

> "The world ultimately will have a single currency, the internet will have a single currency. I personally believe that it will be Bitcoin."
>
> — *Jack Dorsey, The Times of London, 21 March 2018*

> "Crypto is the mother or father of all scams and bubbles"
>
> — *Nouriel Roubini, U.S. Senate hearing, 11 October 2018*

Meteoric cryptocurrency price growth has arguably catalysed the rise of the blockchain industry. However, it has also polarised opinions. Standard models struggle to explain why Bitcoin's price grew from $0.07 on 31 July 2010 to $104,638 on 31 May 2025. Previous studies have appealed to fundamentals like gradual adoption, blockchain security assurances, and technology growth (Athey et al., 2016; Bolt and van Oordt, 2020; Cong et al., 2021; Wei and Dukes, 2021; Pagnotta, 2022; Biais et al., 2023; Karau and Moench, 2023; Prat et al., 2025). However, these models have not replicated the key long-term price trends; specifically, high but decreasing average and variance of price growth. In this paper I present an asset pricing model motivated by the vast public disagreement on the potential of cryptocurrencies and blockchains. The model allows for substantial uncertainty about eventual adoption. The mechanism captures the learning process by which the market develops more precise estimates of eventual adoption demand.

The importance is twofold. First, the model is able to capture the long-term price trends of many major cryptocurrencies. It shows that their rapid price growth can be grounded in standard finance theory, without appealing to departures from fundamental value. In the absence of such an explanation, many have appealed to the hypothesis of long-term bubbles, as per one of the introductory quotes. Although the model does not rule bubbles out, it provides an alternative explanation. Second, the model provides conceptual guidance to cryptocurrency investors. It suggests that they are rewarded for exposure to the resolution of adoption uncertainty. The source of risk is that the price moves proportionally to any revision to the market-estimated eventual adoption price. The source of reward is that any arriving information reduces the perceived variance of the eventual adoption price, which reduces the discount for risk embedded in the current price.

The long-term trends in question are evident in Figure 3.1 for market leaders Bitcoin and Ether.[1] The price graph illustrates the magnitude of price growth. However, it is evident from the slope of the price, and from the graph of average price change, that the growth rate has decreased over time. That is matched by a decrease in the variance of monthly log price changes. While the comovement is unsurprising, the decline warrants further explanation.

These patterns are typical. Per Figure 3.2, most of the 15 cryptocurrencies studied in this paper display falling variance. The magnitude of variance depends not on calendar time but age, which can be considered a proxy for maturity.

The model in the paper builds on the cryptocurrency pricing literature wherein a cryptocurrency's price may switch between being investment demand-driven and adoption demand-driven (Bolt and van Oordt, 2020; Prat et al., 2025; van Oordt, 2025). During an adoption regime, the price is simply the adoption demand to hold the cryptocurrency (perhaps for transactional, store of value, or staking purposes) per unit of supply. An investment regime occurs whenever the expected growth in adoption demand exceeds investors' required return, leading some of the units to be held for investment purposes. Thus, during an investment regime, the price at the time of the next switch to an adoption regime provides an anchor for the current price. Prices after that point have no influence.

To isolate the explanatory power of the learning mechanism, I use a parsimonious model spanning the time until the next switch. It assumes that, each period, an investor allocates her wealth between a cryptocurrency and a risk-free asset. The switch to an adoption phase is modelled as a a final period, during which the

[1] Bitcoin and Ether were the two largest cryptocurrencies by market capitalisation as at 31 May 2025.

Figure 3.1: Long-term price trends

(a) Bitcoin (BTC)

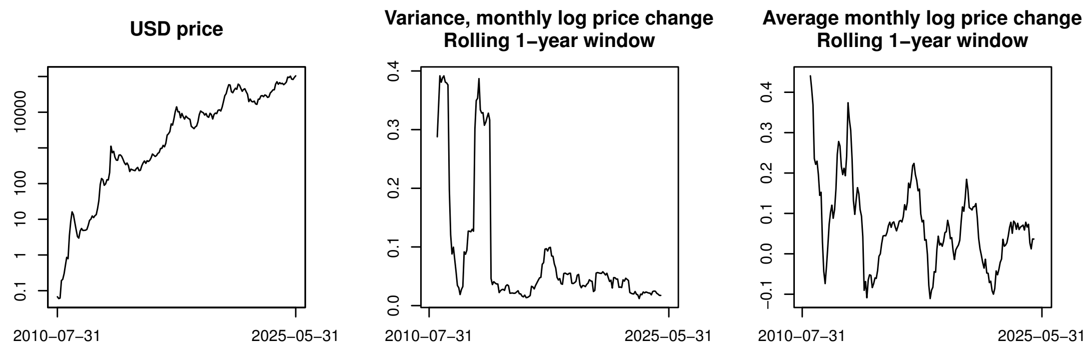

(b) Ether (ETH)

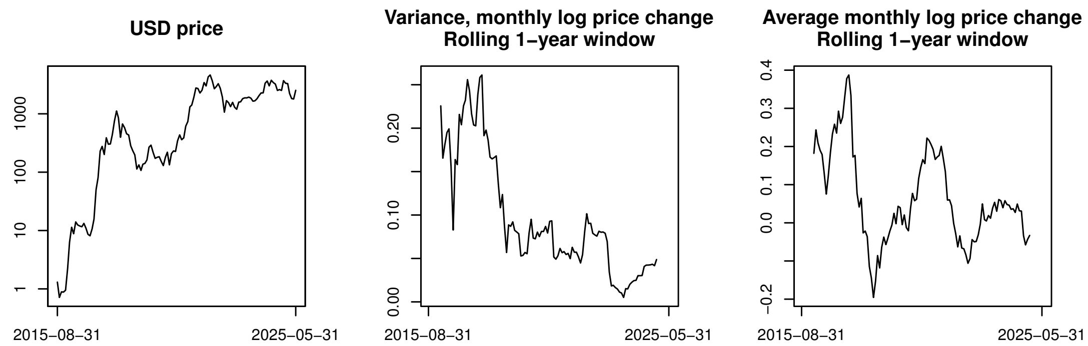

Data sources include Yahoo Finance and coinmetrics.io. Each point on the variance (middle) charts displays the variance of the monthly changes in the log price during the one-year window centred on the date of that point. Similarly, each point on the average (rightmost) charts displays the average of the monthly changes in the log price during the one-year window centred on the date of that point.

investor liquidates her portfolio and consumes the proceeds. The eventual, final-period cryptocurrency price equals exogenous adoption demand (in fiat terms) divided by known supply. However, the investor does not know the eventual adoption price ahead of time. Instead, she receives an informative but noisy signal each period, centred on the log of the true eventual adoption price. The signal represents the process by which society observes the development of a cryptocurrency, the development of its competitors, and the reactions to the cryptocurrency of potential users, the financial sector, regulators, and the general public. The investor

Figure 3.2: Variance by age

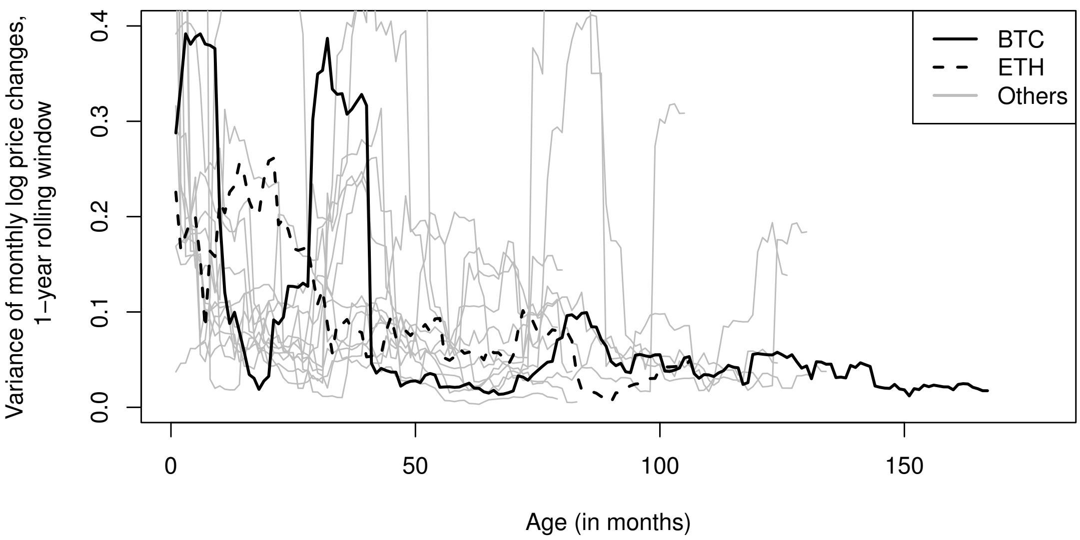

"Others" include BNB, SOL, XRP, DOGE, TRX, ADA, BCH, DOT, LTC, XLM, EOS, DASH, and MAID. Age is the number of months since the start of the series of available data.

applies Weighted Least Squares (WLS) to the series of signals, deriving estimates of the eventual adoption price and her perceived uncertainty about it. Thus her uncertainty is gradually resolved as signals accumulate. The model produces a closed-form solution for the cryptocurrency price. After adjustment for an interest rate differential, the price is the estimated eventual adoption price minus a discount for remaining eventual adoption uncertainty.

According to the model, this discounted eventual adoption price is the fundamental value of cryptocurrency. Fundamental value exists because investors can eventually sell to adopters, such as those who hold the cryptocurrency for transaction purposes.

This study does not take a stand on sources of eventual adoption demand. After all, market prices are primarily concerned with aggregate adoption. That is fortunate; if there is substantial uncertainty about aggregate eventual adoption, that is likely due to even greater uncertainty about individual sources. Nevertheless, to fix ideas, I provide the following examples of potential sources. Bitcoin is envisaged by some as a global currency, alleviating reliance on payment networks controlled by nation-states, facilitating cross-border transactions, and acting as a reserve asset and store of value. A separate category comprises the native cryptocurrencies of Turing-complete virtual machine blockchains such as Ethereum. Transactional demand for such cryptocurrencies may originate from the use of smart contracts that provide alternatives to existing institutions and financial infrastructure. Transactional demand would also accompany any on-chain assets denominated in the cryptocurrency. Another form of demand is service providers' need for staked cryptocurrency as collateral or to provide liquidity.

The model's theoretical assumptions are supported by its quality of fit to the data. For most cryptocurrencies studied, the model successfully captures the long-term decline in the mean and variance of price growth. Additionally, for Bitcoin, Ether, and many other major cryptocurrencies, price levels roughly conform to the model's price mode. This indicates that markets' early estimates of eventual adoption did not stray too far from the current estimate. For some cryptocurrencies, the comparison indicates prolonged periods of over-optimism.

Literature and contribution

The paper contributes to a growing literature developing theoretical models for cryptocurrency or token prices. These papers appeal to fundamental factors to explain prices, particularly gradual growth in transactional usage. Athey et al. (2016) combines gradual usage increases with learning about the likelihood of total collapse. Bolt and van Oordt (2020) and Prat et al. (2025) study the implications of the effective zero lower bound on investor demand. Cong et al. (2021) study usage growth that is driven by technological enhancements and network effects. The authors compare the adoption of tokenised networks to non-tokenised. Wei and Dukes (2021) study network effects in the presence of crash risk. They focus on the mutual reinforcement of adoption and bubbles. Pagnotta (2022) and Karau and Moench (2023) focus on the joint dynamics of prices and Proof of Work blockchain security. Biais et al. (2023) attempt to separate fundamental influences on the Bitcoin price from sunspot shocks.

Some existing literature recognises the centrality of eventual adoption demand. Athey et al. (2016) note that the presence of investors decouples prices from the dynamics of pre-steady state usage. van Oordt (2025) refines the insight, showing that cryptocurrency pricing hinges on the next peak point of usage demand.[2]

However, existing literature mostly ignores uncertainty about the eventual adoption price. The uncertainty is pivotal because investors are risk averse, and the eventual adoption price is their payoff. Nevertheless, the literature has limited its modelling of such uncertainty to a binomial risk of collapse (Athey et al., 2016; Wei and Dukes, 2021; Pagnotta, 2022; Karau and Moench, 2023; Biais et al., 2023). This is unrealistic. Even conditional on the absence of total collapse, eventual adoption demand is highly uncertain.

This paper's model incorporates uncertainty about the eventual adoption price, and the gradual resolution of the uncertainty through learning. These features are sufficient to replicate the high but declining average and variance of price growth (Figure 3.1). The model's ability to fit these trends surpasses previous literature. Also in contrast to existing studies, the model shows general applicability to many cryptocurrencies.

The paper also contributes to a literature explaining asset prices using learning mechanisms. Pastor and Veronesi (2009) provide a review. Most of the body of work focusses on non-cryptocurrency assets. However, Han and Makarov (2021) create a bounded rationality model in which investors imperfectly learn the Sharpe ratio of a speculative asset. The authors show that their model matches various short-term characteristics of cryptocurrency prices, including the emergence of bubbles.

The remainder of the paper is structured as follows. Section 3.2 sets out the model. Section 3.3 explains the estimation method and assumptions, and presents

[2]This is an approximate description of the concept. See the paper for a formal definition.

analytical results from likelihood maximisation. Section 3.4 describes the data used, including the choice of 15 cryptocurrencies for the study and the treatment of interest rates. Section 3.5 assesses and interprets the model fit, and projects the future Bitcoin price distribution. Section 3.6 discusses some points arising from the study, including how other fundamental factors may fit alongside adoption uncertainty and learning. Section 3.7 concludes.

3.2 Model

To explain the model's structure, I start from recent theoretical literature in which two types of demand to hold cryptocurrency are considered: adoption demand and investment demand (Bolt and van Oordt, 2020; Prat et al., 2025; van Oordt, 2025). Adoption demand is that which is relatively insensitive to expected return, and may include transactional demand, store of value demand, and staking demand. Without investors, investment demand is zero, and the price of a cryptocurrency at any time is the adoption demand per unit of supply. With investors, investment demand may still be zero if expected adoption demand growth is low enough. However, if future adoption demand will support a sufficiently high price, then investment demand will be positive, and the current price will switch to being determined by investment factors, rather than adoption. The regime switching is embodied by an occasionally binding constraint that ensures that investment demand is non-negative. Bitcoin and most of the other cryptocurrencies in this study have likely never entered the adoption regime, and may not for many years.[3]

I therefore define eventual adoption as the next point when investment demand reaches zero, such that the non-negativity constraint starts to bind. At this point, the price switches to a regime where it is determined as adoption demand per unit of supply.

[3]Survey evidence on reasons for holding cryptocurrency indicates that investment and speculation are the most common reasons (Steinmetz, 2021; US Federal Reserve, 2024).

In the regime where investment demand is positive, cryptocurrency prices reflect expectations of the eventual adoption price. Conceptually, eventual adoption represents the "cashflow" that investors in cryptocurrency receive, when investors sell to adopters. If the eventual adoption price and timing were known with certainty, then the cryptocurrency would be priced as a risk free bond. Instead, with highly uncertain eventual adoption, cryptocurrencies trade at large discounts to their eventual adoption price.

I capture this in a model with a finite number of periods. The final period, τ, is the point of eventual adoption. In the model, an investor active in period t invests a portion of her wealth ω_t in a cryptocurrency, with the rest invested in a risk-free bond. In the final period τ, she liquidates her portfolio and consumes the proceeds. I use the reduced form assumption that adoption demand prior to the final period is zero. This simplification avoids the need for an explicit occasionally binding constraint.

The eventual adoption price, $\mathcal{E}_\tau$, is the price of the cryptocurrency in fiat currency terms. It is determined as the ratio of eventual adoption demand (in fiat currency) to supply at the point of eventual adoption (in units). Since both of these quantities can be considered exogenous, I assume that the eventual adoption price is exogenous.

The supply of cryptocurrency units is much more predictable than demand. Supply is often predetermined (as for Bitcoin), though it may be also determined formulaically by blockchain conditions such as congestion or the amount staked by validators (as for Ether).

By contrast, eventual adoption demand to hold cryptocurrency is highly unpredictable. In particular, possible sources of demand—the ultimate role of cryptocurrency—have been a topic of wide speculation. The paper does not need to take a stand on which sources will become dominant, since market prices are mainly concerned with their aggregate. Nevertheless, for the sake of fixing ideas, I outline some examples. A first category concerns money transfer-focussed cryptocurrencies such as Bitcoin. Often dubbed "digital gold", Bitcoin could serve as a store of value, including as a reserve asset. Transactional demand to hold Bitcoin may result from its use as a medium of exchange, particularly for large, time-insensitive transactions, or by those who value payment rails free from control by national governments. As a second category, Ethereum and similar blockchains can execute computer code on-chain. They constitute a platform technology. By offering an ecosystem of on-chain applications and services, they could eventually provide an alternative to existing financial infrastructure. Such use would contribute to transactional holdings of the native cryptocurrency of the blockchain, to pay for block space. Further transactional demand would arise from trading in assets denominated in the native cryptocurrency, including digital assets issued on that blockchain. Staking creates another source of demand for cryptocurrency. Staked cryptocurrency functions as an input to the production of a service, such as exchange liquidity or blockchain network security.

Turning to the model, the framework used is mean-variance optimisation. A continuum of investors of measure one faces the following portfolio choice problem between cryptocurrency and a risk-free bond.

$$\max_{\omega_t} \ \mathbb{E}_t U(C_\tau) \qquad \text{s.t.} \qquad C_\tau = W_\tau = W_t e^{\omega_t \left(\log \frac{\mathcal{E}_\tau}{\mathcal{E}_t} + r^*_{t,\tau}\right) + (1-\omega_t) r_{t,\tau}} \tag{3.1}$$

In the above, C_τ is consumption in final period τ, U is a continuous, twice-differentiable utility function with $U' > 0$ and $U'' < 0$ for all $C_\tau > 0$, W_t is wealth in period t, ω_t is a weight variable, $\mathcal{E}_t$ is the price (exchange rate) of the cryptocurrency in terms of fiat currency at time t, $r_{t,\tau} = \sum_{s=t}^{\tau-1} r_s$ is the sum of an exogenous fiat-denominated risk-free return for each period between time t and τ, and $r^*_{t,\tau} = \sum_{s=t}^{\tau-1} r^*_s$ is the sum of an exogenous cryptocurrency-denominated risk-free return for each period between time t and τ. A low-risk, if not risk-free, example of r^*_t is the native staking yield for Proof-of-Stake cryptocurrencies.[4] Expectations

[4]Native staking yields are often not fixed in advance, and carry other risks such as punishment for misbehaviour, or protocol risk. Notably however, in recent years there have been fixed-for-floating interest rate swap transactions incorporating the Ether staking yield as the floating leg,

are taken over investors' perceived probability distribution of $\log \mathcal{E}_\tau$, which I will specify shortly. Inflation and consumption growth are implicitly assumed to be zero; as such, I rely on normalised prices to estimate the model.

Using a second-order approximation to the utility function, I express the problem in terms of the mean and variance of the log portfolio excess return (further detail in Section 3.B).

$$\max_{\omega_t} \ \omega_t \mathbb{E}_t \left(\log \frac{\mathcal{E}_\tau}{\mathcal{E}_t} + r^d_{t,\tau} \right) - \frac{1}{2} \delta \omega_t^2 \mathrm{Var}_t \left(\log \frac{\mathcal{E}_\tau}{\mathcal{E}_t} + r^d_{t,\tau} \right)$$

where $r^d_{t,\tau} = r^*_{t,\tau} - r_{t,\tau}$ is the interest rate differential, and δ is the risk aversion parameter. The first-order condition is then

$$\omega_t = \frac{\mathbb{E}_t \log \frac{\mathcal{E}_\tau}{\mathcal{E}_t} + r^d_{t,\tau}}{\delta \mathrm{Var}_t \log \mathcal{E}_\tau} \tag{3.2}$$

I impose zero net supply of the risk-free bond, such that:

$$\omega_t = 1 \tag{3.3}$$

This market clearing condition underscores a fundamental distinction between the primary source of near-riskless yields denominated in each type of currency. While riskless fiat-denominated bonds are in zero net supply, staked cryptocurrency is locked up in return for the staking rewards. In this regard, staked cryptocurrency is similar to installed physical capital. The distinction extends to fiat currencies when compared to cryptocurrencies themselves. Each unit of fiat currency is a liability for the issuer and an asset for the holder. By contrast, cryptocurrencies do not represent claims on an issuer.[5] In this regard, cryptocurrencies resemble physical assets. Consequently, as land, oil, and gold are in positive net supply, so too are cryptocurrencies.

---

indicating potential further development of this market. The use of cryptocurrency risk-free rates for estimation of the model is addressed in Section 3.4.1.

[5]This is true of a blockchain's native cryptocurrency. There do, however, exist cryptoassets such as stablecoins or wrapped tokens that represent claims on an underlying asset, which itself could be in zero net supply.

Combining the first-order condition and market clearing condition reveals Equation 3.4 for the fundamental value of a cryptocurrency. Interpreting the right-hand side, the expected eventual log price is the long-term payoff. The price is subject to a discount proportional to the variance of the eventual log price. The resulting discounted payoff is then adjusted by the interest rate differential for the intervening period.

$$\log \mathcal{E}_t = \mathbb{E}_t \log \mathcal{E}_\tau + r^d_{t,\tau} - \delta \mathrm{Var}_t \log \mathcal{E}_\tau \tag{3.4}$$

This same result can be obtained for a model in which investors consume in every period, as shown in Section 3.B. Note that the fundamental value can be arbitrarily close to zero, depending in particular on the cryptocurrency's adoption prospects and thus the rationally expected value of $\log \mathcal{E}_\tau$.

Drawing from the fundamental value equation, the model's uncovered interest rate parity (UIRP) condition is as follows.[6]

$$\begin{aligned}
\log \frac{\mathcal{E}_{t+1}}{\mathcal{E}_t} &= \mathbb{E}_{t+1} \log \mathcal{E}_\tau - \mathbb{E}_t \log \mathcal{E}_\tau - r^d_t + \delta(\mathrm{Var}_t \log \mathcal{E}_\tau - \mathrm{Var}_{t+1} \log \mathcal{E}_\tau) \\
\mathbb{E}_t \log \frac{\mathcal{E}_{t+1}}{\mathcal{E}_t} &= -r^d_t + \mu_{t+1} \qquad (3.5)
\end{aligned}$$

where $r^d_t = r^d_{t,t+1}$ is the single-period interest differential, and where $\mu_{t+1} = \delta(\mathrm{Var}_t \log \mathcal{E}_\tau - \mathrm{Var}_{t+1} \log \mathcal{E}_\tau)$ is the UIRP deviation. μ_{t+1} is positive as long as the variance is falling.

Since r^d_t is the excess of the cryptocurrency interest rate over the fiat interest rate, a higher value is typically balanced by lower expected cryptocurrency price growth. The UIRP deviation is the risk premium earned due to the investors' exposure to the portion of eventual adoption uncertainty that is resolved over the period.

So far in this section, I have deviated little from standard mean-variance portfolio optimisation. The contribution of this paper stems from its model for the evolution of the expectation and variance terms in Equation 3.4, based on adaptive learning.

[6] The result obtains due to the Law of Iterated Expectations and the assumption that investors know all signal variances at all times.

The model assumes that $\log \mathcal{E}_\tau$, while exogenous, is unknown to investors. Investors gradually learn about its value by observing an exogenous information signal each period of the following form.

$$x_t = \log \mathcal{E}_\tau + \tilde{x}_t \tag{3.6}$$

where

$$\mathbb{E}(\tilde{x}_t) = 0$$

$$\mathrm{Var}(\tilde{x}_t) = \sigma_{xt}^2$$

$$\mathrm{Cov}(\tilde{x}_t, \tilde{x}_s) = 0 \quad \text{for} \quad s \neq t$$

Within x_t, $\log \mathcal{E}_\tau$ is the informative, static part of the signal, and $\tilde{x}_t$ is the noise. Investors have full knowledge of all variances σ_{xt}^2 in every period; that is, for $s \in [1, \tau]$, the investor active in period s knows $\{\sigma_{xt}^2\}_{t=1}^{\tau}$. There is no private information in the model, as every investor receives the same signal, and there are no other participants in the cryptocurrency market.

To learn $\log \mathcal{E}_\tau$, investors use weighted least squares to estimate the mean of the signals, with a diagonal weight matrix with element in row t, column t set to σ_{xt}^{-2}[7]. The weighted least squares estimator is then the best linear unbiased estimator.[8] The investor therefore calculates her perceived expectation and variance of $\log \mathcal{E}_\tau$ as follows.

$$\mathbb{E}_t \log \mathcal{E}_\tau = \sum_{s=1}^{t} x_s w_{st} \tag{3.7}$$

$$\mathrm{Var}_t \log \mathcal{E}_\tau = \Omega_t^{-1} \tag{3.8}$$

In the above, $\Omega_t := \sum_{s=1}^{t} \sigma_{xs}^{-2}$ is the sum of weights up to time t, and $w_{st} := \frac{\sigma_{xs}^{-2}}{\Omega_t}$ is the proportional weight given to x_s.

[7] It would also be possible to use Bayesian, rather than frequentist, learning to obtain similar behaviour and pricing.

[8] It is also the maximum likelihood estimator in the case of normally distributed $\tilde{x}_t$.

Therefore, although $\log \mathcal{E}_\tau$ is exogenous and static, the investors' estimate of it is not. With the arrival of the signals over time, the investors' estimate adjusts and becomes more precise.

The time-subscripted expectation and variance operators above pertain to the investors' information set at time t, which includes the signals up to that time, all interest rate differentials, and all parameters.

I define a separate, "structural" information set, under which $\mathcal{E}_\tau$ is known, but no signals or prices are yet observed. I use the structural information set to assess the model's ability to capture the trends motivating the paper.

To define both, let $\theta := \{\delta, \varphi, \{\sigma^2_{xs}\}_{s=1}^{\tau-1}\}$ denote the parameter vector, where φ is a vector of any parameters of the distribution of $\tilde{x}_t$ that are not identified by specifying its first two moments.[9] Let $\mathcal{G} := \{\theta, \{r^d_s\}_{s=1}^{\tau-1}\}$ contain all information common to both information sets. Then

$$f_t(X) := f(X|\{x_s\}_{s=1}^t, \mathcal{G}) \qquad \text{investor probability distribution} \tag{3.9}$$

$$f_0^{\mathcal{E}_\tau}(X) := f(X|\mathcal{E}_\tau, \mathcal{G}) \qquad \text{structural probability distribution} \tag{3.10}$$

where X is a model variable. Thus the structural expectation and variance are denoted $\mathbb{E}_0^{\mathcal{E}_\tau}$ and $\mathrm{Var}_0^{\mathcal{E}_\tau}$. Note that the investor information set up to time t confers knowledge of prices $\{\log \mathcal{E}_s\}_{s=1}^t$, since all terms of the right-hand side of Equation 3.4 are known.

With normally distributed signals, investors' expectation of $\log \mathcal{E}_\tau$ (Equation 3.7) is a sufficient statistic for the model's state.[10] Thus, although investors notionally apply weighted least squares estimation to the entire signal history each period, they actually only need to know the last period's expected eventual adoption price (Equation 3.7) and the current period's signal. This state variable can be expressed

[9] If the normal distribution is used, φ is empty.

[10] Detemple (1986) explores a general class of models in which the expectation of unobserved variables, conditional on exogenous informative signals, is a sufficient statistic for the state.

in recursive form, with deterministic gain w_{tt}, as follows.

$$\mathbb{E}_t \log \mathcal{E}_\tau = \mathbb{E}_{t-1} \log \mathcal{E}_\tau + w_{tt} \left(x_t - \mathbb{E}_{t-1} \log \mathcal{E}_\tau \right) \tag{3.11}$$

The learning in the model, where the source of information is exogenous, contrasts with other settings in which agents learn from endogenous variables such as prices. In the latter style of model, common in the macroeconomic learning literature (Evans and Honkapohja, 2001), agents' learning guides their actions, which in turn affect the variables from which they learn. This circularity, and the notion of E-stability within such a system, is avoided in this model. However, models in which agents learn from observing cryptocurrency prices have been shown to be useful in modelling short-term bubbles (Han and Makarov, 2021). This paper's model abstracts from these influences, hypothesising that they are not necessary to capture the long-term trends of interest to this study.

The signals and learning process provide functional form to the expectation and variance terms in the fundamental value equation. Substituting in produces the following asset pricing equation for cryptocurrency.

$$\log \mathcal{E}_t = \log \mathcal{E}_\tau + r^d_{t,\tau} - \delta \Omega_t^{-1} + \sum_{s=1}^{t} \tilde{x}_s w_{st} \tag{3.12}$$

Here, the expected long-term payoff is replaced by the actual payoff ($\log \mathcal{E}_\tau$) plus fluctuation caused by deviations of information signals from the true log eventual adoption price (the summation term). The variance of the long-term payoff is expressed as a function of signal variances.

The WLS results also give functional form to components of the UIRP condition, as follows.

$$\begin{aligned} \log \frac{\mathcal{E}_{t+1}}{\mathcal{E}_t} &= w_{t+1,t+1}(x_{t+1} - \mathbb{E}_t \log \mathcal{E}_\tau) - r^d_t + \delta(\Omega_t^{-1} - \Omega_{t+1}^{-1}) \\ \mathbb{E}_t \log \frac{\mathcal{E}_{t+1}}{\mathcal{E}_t} &= -r^d_t + \delta(\Omega_t^{-1} - \Omega_{t+1}^{-1}) \end{aligned} \tag{3.13}$$

The model thus establishes the eventual adoption price as a fundamental driver of cryptocurrency prices. This price can itself be driven by various underlying factors not addressed by the model, such as factors specific to a particular cryptocurrency or its role, or potentially broader macroeconomic conditions at the time of eventual adoption.

The model is limited to pricing cryptocurrencies during the investment-driven regime. It is not applicable to the adoption-driven regime, when the zero lower bound binds on investor demand, as explored in Bolt and van Oordt (2020), Prat et al. (2025), and van Oordt (2025). An extension might incorporate a model of the adoption demand-driven regime.

Another possible extension would be to admit uncertainty and learning about the timing of τ, rather than only about $\log \mathcal{E}_\tau$.

Another limitation relates to the transition from period $\tau - 1$ to τ. The model assumes that, in the final period, all remaining uncertainty about eventual adoption will be resolved. This may correspond approximately to adoption being an event driven by network effects, where a critical mass threshold in usage is reached, causing briefly exponential growth. However, adoption may be more gradual and predictable, without any particularly large revisions to eventual adoption estimates just as period τ is reached. Future research might study alternative signal processes that lead $\text{Var}_t \log \mathcal{E}_\tau$ to vanish as $t \to \tau$.

The model does not take a stand on whether a particular eventual adoption price corresponds to the demise and replacement of other technologies or services, such as fiat currency. These possibilities are captured only indirectly.

In theory, the model could be applied to any asset with a single, uncertain future payoff, of which investors gradually learn the amount. In practice, learning appears to be closely tied to asset infancy. Thus cryptocurrencies are unique in their clean expression of this profile. In particular, their distinctive pattern of declining variance suggests that learning is the dominant price influence. The absence of declining variance in most other assets' returns suggests that other factors are dominant. For example, the nature of gold, and some technology stocks, may bear similarities to cryptocurrency. But after centuries of mining and using gold, most learning about its long-run demand is likely complete. Likewise, once a company is mature enough to list on a stock exchange, the most substantive part of the learning about its business model and prospects may be complete.

3.3 Estimation

For the purpose of estimating the model for a particular cryptocurrency, I treat the current period's weighted signal noise, $\tilde{x}_t w_{tt}$, as the error term, with the previous periods' signals forming part of the conditional expectation. Thus Equation 3.12 can be expressed as a regression equation as follows.

$$\log \mathcal{E}_t = \log \mathcal{E}_\tau + r^d_{t,\tau} - \delta \Omega_t^{-1} + \sum_{s=1}^{t-1} \tilde{x}_s w_{st} + u_t \tag{3.14}$$
$$\text{where } u_t = \tilde{x}_t w_{tt}, \quad \text{and} \quad \text{Var}(u_t) = \sigma^2_{xt} w^2_{tt}$$

Each $\tilde{x}_t$ value can be calculated for the given set of candidate parameter values using the price $\log \mathcal{E}_t$ and the calculated values of the signals prior to t.

I use maximum likelihood estimation since the variance of the error term u_t and the conditional expectation depend on common parameters, and since the equation is nonlinear in parameters. I include $\log \mathcal{E}_\tau$ as a parameter for the purpose of estimation. The parameters to be estimated are therefore δ, $\log \mathcal{E}_\tau$, any parameters necessary to specify σ^2_{xt}, and the set of parameters φ necessary to govern the probability distribution of the $\tilde{x}_t$ variables beyond their first two moments.

3.3.1 Analytical results for maximum likelihood estimator

Some insights are available from certain analytical results which can be derived from the first-order conditions of the maximum likelihood for δ and $\log \mathcal{E}_\tau$.

The log likelihood function for observation t takes the following form.

$$
\begin{aligned}
\log L_t &\left(\delta, \log \mathcal{E}_\tau, \{\sigma^2_{xs}\}_{s=1}^T, \varphi | \{\log \mathcal{E}_s\}_{s=1}^T, \{r_s^d\}_{s=1}^{\tau-1}\right) \\
&= \log g(u_t) \\
&= \log g\left(\log \mathcal{E}_t - \log \mathcal{E}_\tau - r_{t,\tau}^d + \delta\Omega_t^{-1} - \sum_{s=1}^{t-1} \tilde{x}_s w_{st}\right) \qquad (3.15)
\end{aligned}
$$

where g is the probability density function of $u_t = \tilde{x}_t w_{tt}$.

The following first-order conditions, for δ and $\log \mathcal{E}_\tau$ respectively, obtain for the general case.

$$0 = g'(u_1) \qquad (3.16)$$

$$0 = \sum_{t=1}^T \frac{g'(u_t)}{g(u_t)} w_{tt} \qquad (3.17)$$

Equation 3.16 indicates that the MLE procedure will ensure that u_1 is equal to the mode of its distribution.[11]

In the special case of normally distributed signal errors $\tilde{x}_t$ and thus normally distributed u_t, the first-order conditions for δ and $\log \mathcal{E}_\tau$ become the following, respectively.

$$0 = \tilde{x}_1 \qquad (3.18)$$

$$0 = \sum_{s=1}^T \tilde{x}_s w_{sT} \qquad (3.19)$$

Combining these first-order conditions with Equation 3.12 reveals that, in the normal distribution case, the estimation sets the log price in these periods equal to the structurally expected price.

$$\log \mathcal{E}_1 = \mathbb{E}_0^{\mathcal{E}_\tau} \log \mathcal{E}_1 = \log \mathcal{E}_\tau + r_{1,\tau}^d - \delta\Omega_1^{-1} \qquad (3.20)$$

$$\log \mathcal{E}_T = \mathbb{E}_0^{\mathcal{E}_\tau} \log \mathcal{E}_T = \log \mathcal{E}_\tau + r_{T,\tau}^d - \delta\Omega_T^{-1} \qquad (3.21)$$

where $\mathbb{E}_0^{\mathcal{E}_\tau} \log \mathcal{E}_1$ and $\mathbb{E}_0^{\mathcal{E}_\tau} \log \mathcal{E}_T$ are computed by taking the structural expectation of the log price as per Equation 3.12.

[11] Or one of its modes, if not unimodal. Equivalently, $\tilde{x}_1$ will be set to a mode of its distribution, since $w_{11} = 1$.

This is particularly intuitive in the case of period T, where the modelled investor has access to the same information as the likelihood maximisation process, and where, due to the normal distribution of the errors, both the investor and modeller would use the same estimator for $\log \mathcal{E}_\tau$ given the signals. In the normal distribution case, the modeller has no basis to predict over or undervaluation of the cryptocurrency in the final period.

The equations, when combined, produce Equation 3.22, which shows that estimated δ functions as an ex post measure of reward for risk. The numerator is the total log return over the data period, and the denominator is the variance of the price change over the data period (this is shown in Section 3.C).

$$\delta = \frac{\log \mathcal{E}_T + r^d_{1,T} - \log \mathcal{E}_1}{\Omega_1^{-1} - \Omega_T^{-1}} = \frac{\log \mathcal{E}_T + r^d_{1,T} - \log \mathcal{E}_1}{\mathrm{Var}_1 \log \mathcal{E}_\tau - \mathrm{Var}_T \log \mathcal{E}_\tau} \tag{3.22}$$

δ is best understood in the context of the model's market clearing condition, Equation 3.3, which forces the investor to hold all of her portfolio in cryptocurrency. In effect, δ becomes the balancing item in the estimation. Its estimated value is the risk aversion parameter that would make the investor choose to bear the realised risk for the realised reward.

δ can be negative. In the normal distribution case, the sign of δ matches the sign of the log total return in the numerator, because the denominator is always positive.[12]

Negative δ does not require risk seeking behaviour. For example, an initially promising but unsuccessful blockchain project would have negative δ. In such cases, overall downward movements in the price are caused by downward revisions to investors' estimate of $\log \mathcal{E}_\tau$. Therefore it is possible that investors still received positive rewards for the reduction in eventual adoption uncertainty over the data period, even when the rewards were more than offset by lower than expected signal values.

The ex post nature of the δ estimate means that the fitted model projects that the overall price trend during the data period will continue. This is appropriate if

[12]The denominator will be positive as long as at least one of the signal variances $\{\sigma^2_{xt}\}^T_{t=2}$ is not infinite.

the price trend is caused by factors that will likely persist; for example, unmodelled influences like non-pecuniary benefits of holding the cryptocurrency or a particular degree of correlation to other assets. However, if temporary influences were strong during the data period, then the model's projections of future prices may be inaccurate. For example, negative δ could indicate a promising cryptocurrency which was initially over-hyped, but where investors now hold more realistic views of its future prospects. In such a case, the model would predict further price decline when in fact investors expect growth. The model projections and parameter estimates may therefore be sensitive to short sample periods, particularly those dominated by high or low price growth.

3.3.2 Signal variance functional form

To complete the model, a specification of the σ^2_{xt} function is needed. The function governs how quickly the market learns about eventual adoption of a cryptocurrency. The lower the value of σ^2_{xt} in period t, the more precise the signal and the more information is gained in that period.

I adopt the following specification of σ^2_{xt}.

$$\sigma^2_{xt} = \begin{cases} \sigma^2_{x1}, & \text{if } t = 1 \\ \sigma^2_{x}, & \text{if } t \geq 2 \end{cases}$$

The variance is assumed constant, except for an independent parameter for the first period variance. That is needed because, by the time of the first price observation of a cryptocurrency, much price-relevant information is already available. For example, the envisioned role and function of the cryptocurrency, the team responsible for the launch and subsequent development, any new technologies being deployed, any predetermined ownership or distribution of coins, and the associated white paper providing various other details. The amount of this information is likely to be much greater than the additional information obtained on subsequent individual days. Therefore, it is likely that $\sigma^2_{x1} \ll \sigma^2_{x}$. Of course, given enough time, the total quantity of information arriving after time $t = 1$ may well exceed the "launch information", as measured by their weights in the investors' weighted least squares estimate. That is, it remains possible that $\sigma_{x1}^{-2} < \sum_{s=1}^{t} \sigma_{xs}^{-2}$ for some t.

3.3.3 Parameters governing the probability distribution

The completed model's parameter vector contains $\log \mathcal{E}_\tau, \delta, \sigma_{x1}^2, \sigma_x^2$, and φ. If a normal distribution is assumed, then φ is empty. Further, Equations 3.20 and 3.21 can be used to create a concentrated maximum likelihood estimation problem by substituting out for any two of the parameters in the likelihood function, thereby simplifying the estimation process. However, many of the cryptocurrencies I fit to result in imputed information signals that depart from normality. I therefore use the skewed t distribution. The additional parameters, which are assumed to be the same for every $\tilde{x}_t$, are $\varphi = \{\alpha, \nu\}$. Broadly, α governs skewness, and ν governs kurtosis. These parameters are estimated within the same likelihood maximisation. The other parameters of the skewed t distribution, ξ (location) and ω (scale), are identified by the mean and variance of the signals as specified by the model, and can vary between the $\tilde{x}_t$ variables.

Finally, I include in the optimisation a penalty for the squared distance between the theoretical and the empirical skewness and kurtosis of $\tilde{x}_t$. The penalties prevent these moments from taking extreme values for some of the cryptocurrencies, ensuring that any projections based on simulation are not distorted by outliers.

3.4 Data

The model requires a series of prices for each cryptocurrency, or token (hereafter cryptocurrency). I use Yahoo Finance as the source of monthly USD price data from September 2014 onwards for Bitcoin and Litecoin, and from November 2017 onwards for other cryptocurrencies. Earlier prices are from coinmetrics.io and coingecko.com.

I fit the model once for each of 15 different cryptocurrencies. I focus on the largest cryptocurrencies by market capitalisation, since these are the most likely to be influential. However, to avoid limiting the analysis to successful cryptocurrencies, I aim to include cryptocurrencies that were once influential but have since stagnated or declined. To select them, I compile a database of the top 10 ex-stablecoin cryptocurrencies' market capitalisation each day since 28 April 2013 according to coinmarketcap.com. For each cryptocurrency that appears in the database, I calculate the number of days it was included in the top 10. I select the 15 cryptocurrencies with the most days. The selection process results in the inclusion of some previously-influential cryptocurrencies, the study of which reveals important insights about the model's behaviour.

The selected cryptocurrencies are summarised in Table 3.1.

Table 3.1: Summary of selected cryptocurrencies

			Price data	
Ticker	Name	Days in top 10*	Start date	End date
BTC	Bitcoin	4331	31/07/2010	31/05/2025
XRP	XRP	4233	31/08/2013	31/05/2025
ETH	Ether	3508	31/08/2015	31/05/2025
LTC	Litecoin	3181	30/04/2013	31/05/2025
ADA	Cardano	2346	31/10/2017	31/05/2025
DOGE	Dogecoin	2336	31/12/2013	31/05/2025
BNB	Binance Coin	2230	31/07/2017	31/05/2025
BCH	Bitcoin Cash	1452	31/08/2017	31/05/2025
DASH	Dash	1319	28/02/2014	31/05/2025
XLM	Stellar	1300	30/09/2015	31/05/2025
SOL	Solana	1290	30/04/2020	31/05/2025
TRX	Tron	1107	30/11/2017	31/05/2025
DOT	Polkadot	1101	31/08/2020	31/05/2025
EOS	EOS	927	30/11/2017	31/05/2025
MAID	MaidSafeCoin	867	30/04/2014	31/12/2021

* Between 28 April 2013 and 18 March 2025

From January 2022, MaidSafeCoin trading volume has been low enough that there are some days without any trades, leading to inflated price volatility. Since this volatility is caused by illiquidity rather than information arrival, the period was omitted.

Since the eventual price may be realised far in the future, to assist interpretation I detrend the price series by the average nominal growth of GWP over the period 2010-2023, sourced from the World Bank World Development Indicators database.

3.4.1 Risk-free interest rates

Due to practical constraints, I assume an interest rate differential of zero for all cryptocurrencies and time periods. This is isomorphic to comparing holding cryptocurrency with holding fiat currency, without lending or staking either of them to earn the associated interest rate or staking reward. While incorporating the interest rate differential conceptually aligns the model with existing theories of fiat currency exchange rates, I show in Section 3.5.4 that the model's results are not sensitive to its exclusion.

The main practical difficulty is in obtaining measures of a risk-free interest rate on cryptocurrency holdings. Risk-free rates are needed because any credit risk or illiquidity premia will distort the UIRP deviation. However, markets in cryptocurrency futures, lending, and staking are still immature, and prices likely still incorporate various risk premia. I provide further detail in Section 3.D.

3.5 Results

3.5.1 Graphical results for selected cryptocurrencies

To explore the fit of the model, this subsection presents charts for five cryptocurrencies' price, price change variance, and price growth, comparing each to the structural distribution of model variables.

Specifically, I compare the data to the model's structural price mode,[13] structural

[13] I use the structural price mode because this is the the model's view of the most likely price path realisation. For example, the maximum likelihood estimator sets the model's first-period mode equal to the first-period price, as per the first-order condition in Equation 3.16.

Figure 3.3: Bitcoin (BTC) model fit

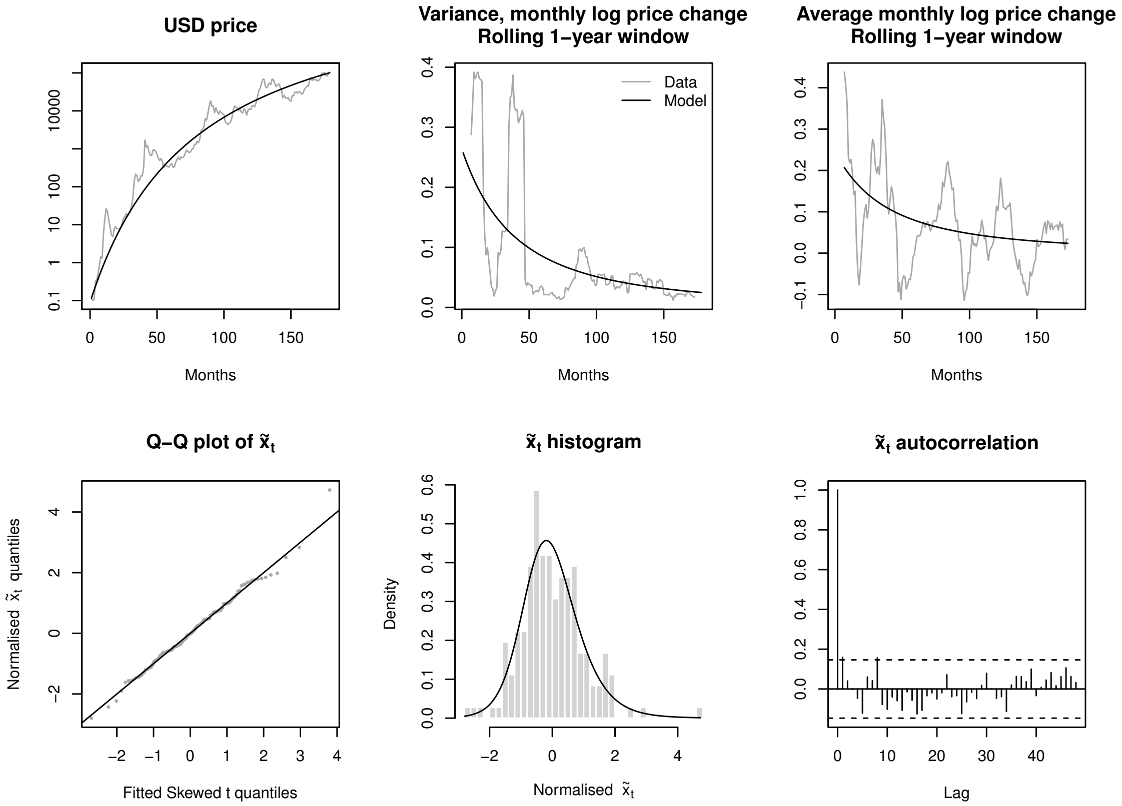

variance of log price changes, and structural average of price changes. The formulae for structural variance and expectation, respectively, are given in Equations 3.23 and 3.24 below, derived in Section 3.C. However, the structural price mode has no closed-form solution, so I use an approximation.

$$\mathrm{Var}_0^{\mathcal{E}_\tau} \log \frac{\mathcal{E}_{t'}}{\mathcal{E}_t} = \Omega_t^{-1} - \Omega_{t'}^{-1} \tag{3.23}$$

$$\mathbb{E}_0^{\mathcal{E}_\tau} \log \frac{\mathcal{E}_{t'}}{\mathcal{E}_t} = -r_{t,t'}^d + \delta\left(\Omega_t^{-1} - \Omega_{t'}^{-1}\right) \tag{3.24}$$

Figures 3.3 and 3.4 show that the model achieves a high quality fit for the two largest cryptocurrencies by market capitalisation, Bitcoin and Ether. The key measures of fit are the rolling annual variance and the rolling average of log price changes. The model-predicted variance of price changes traces the downward trend in the data, and the rolling annual average price changes tell the same story,

Figure 3.4: Ether (ETH) model fit

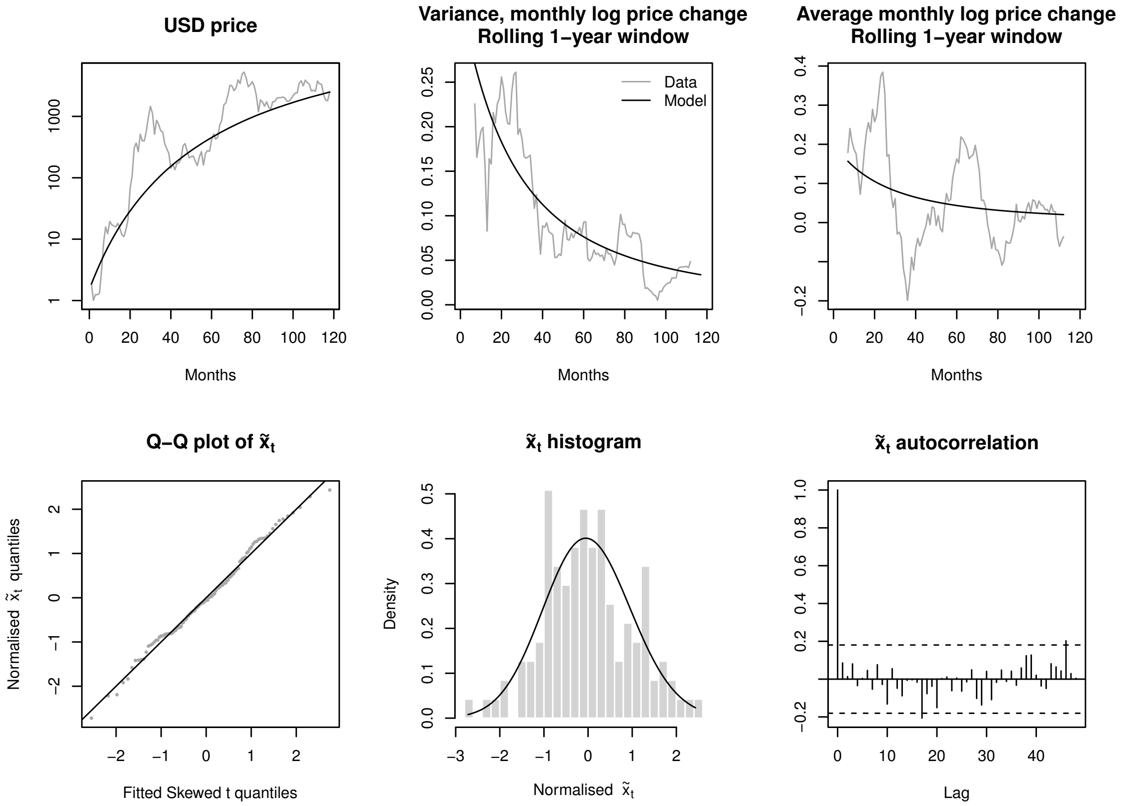

supporting the model's assumption of constant proportional reward for variance (δ). In addition, the signal diagnostics show that the skewed t distribution assumption for the signals is appropriate. The assumption of no autocorrelation may be violated in the case of Bitcoin; however, this is not corroborated by the optimal number of lags shown later in Table 3.2. The model fits most of the other cryptocurrencies similarly well, although many show greater deviations from the model's structural price mode than that seen with Bitcoin. The graphs are given in Section 3.E.

The leftmost panel in the top row shows that the fitted model's structural price mode also traces the same curve as the data. However, the distance of the price from the model's structural price mode is not the right measure of the model's fit. Rather, it indicates the accuracy of early information signals for the given cryptocurrency. In the case of Bitcoin and Ether, according to the information signals we have seen

Figure 3.5: Dash (DASH) model fit

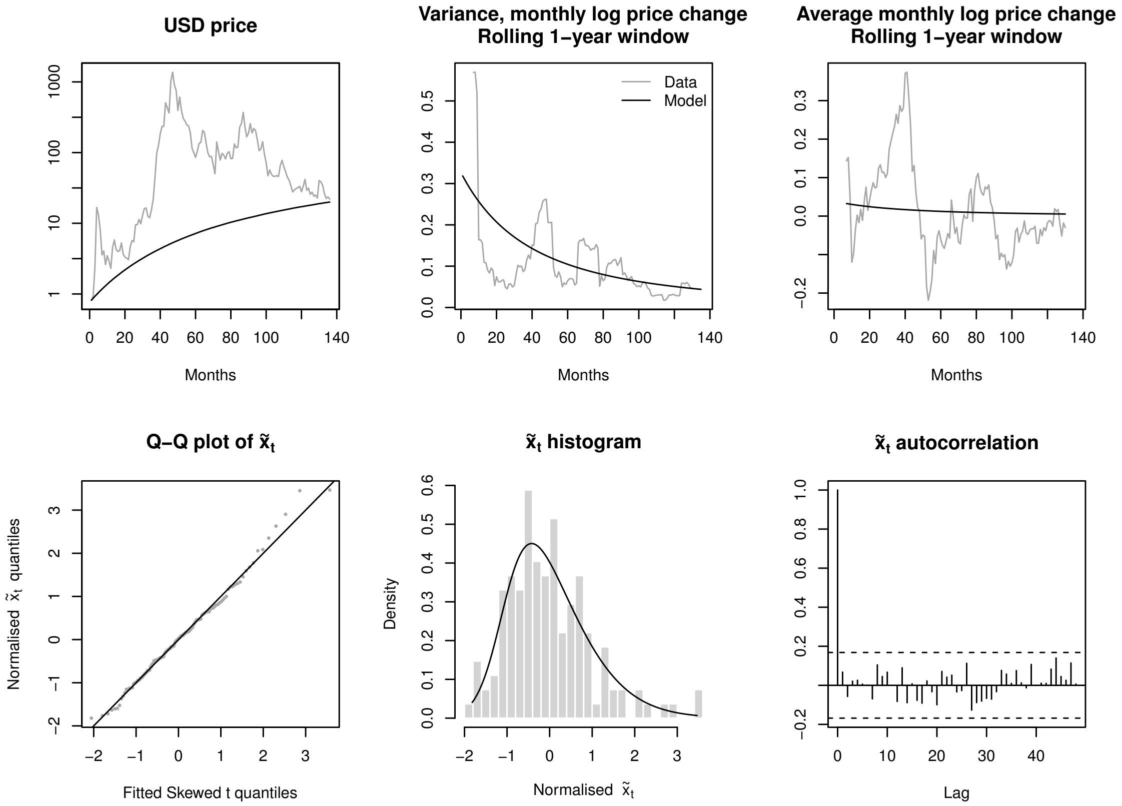

to date, early information signals appear to have been fairly accurate, deviating only for relatively short periods from the eventual adoption price.

Visually, it appears that departures from the model's structural price mode path are followed by reversions to it. However, this trend reversion does not necessarily offer predictive power. It arises because the investors at time $t < T$ possess less information than our estimation process. At time $t = T$, any remaining prediction of over or undervaluation arises from the nonlinear maximum likelihood estimation producing a different result to the investors' WLS estimation.

An example of a cryptocurrency price showing sustained deviation from the structural price mode path is that of Dash. Figure 3.5 shows that the declining variance structure is still present and well approximated by the model, as is the declining price growth rate. However, the market price has mostly exceeded the

Figure 3.6: Privacy coins price history

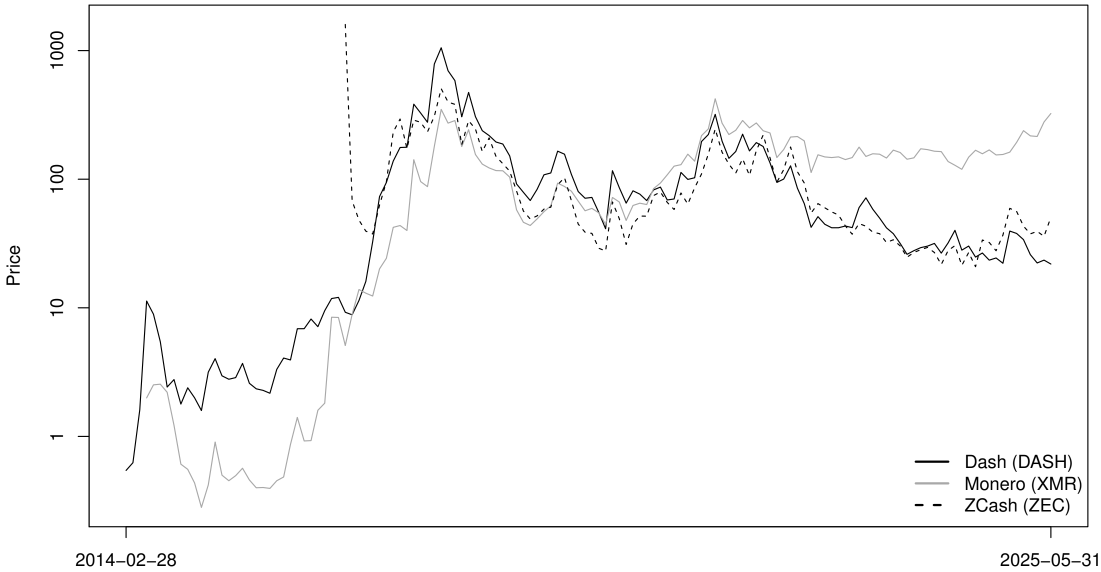

Source: Yahoo Finance and coinmetrics.io

structural mode, when fit to present data. This indicates that some information signals during the observation period were too optimistic, compared to both the earliest signals and the later signals, and have since been corrected. A possible reason is Dash's facilitation of private transactions. Opposition to transaction privacy by authorities has led to the recent de-listing of DASH from many cryptocurrency exchanges, along with other privacy coins—notably Monero (XMR) and Zcash (ZEC) (Kaiko Research, 2024). Indeed, since the end of 2017, all three of these cryptocurrencies' prices have either stagnated or declined (Figure 3.6). Interpreting these developments using the model, there may have been gradual revelation that the use of privacy-centric cryptocurrencies would be curtailed by authorities, thereby limiting their eventual adoption.

The case of $\delta < 0$, which tends to correspond to a declining price over the observation period, was discussed theoretically in Section 3.3.1. One example is Bitcoin Cash (see Figure 3.7). The model's interpretation of the gradual price fall is that, although a positive price effect was generated by declining price variance over

Figure 3.7: Bitcoin Cash (BCH) model fit

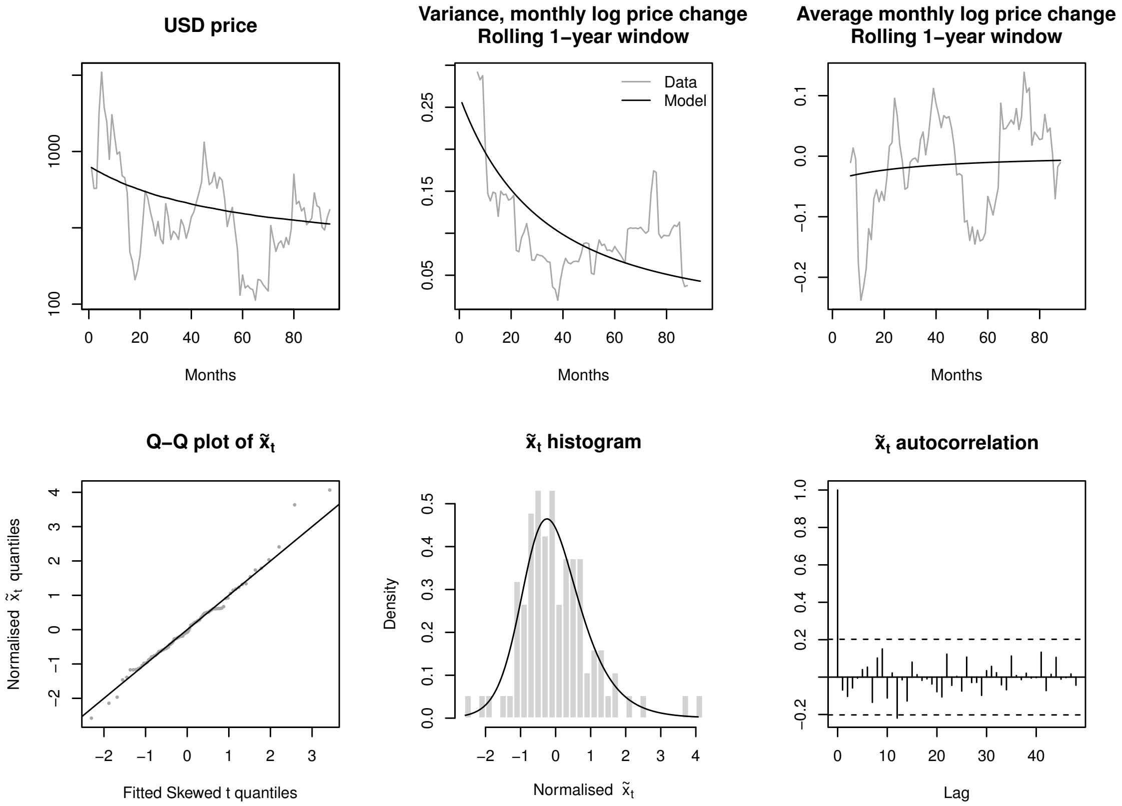

the period, negative revisions to eventual adoption outweighed it. As background, Bitcoin Cash is a separate cryptocurrency from Bitcoin, with its own blockchain. Bitcoin Cash was created by a hard fork of Bitcoin, following disagreements among the Bitcoin community. Specifically, the proponents of Bitcoin Cash preferred a larger block size so as to better facilitate Bitcoin's role as medium of exchange. The negative price movement implies again that the relative success of Bitcoin Cash was overestimated by early signals.

3.5.2 Regression results

Table 3.2 shows the parameter estimates for all 15 cryptocurrencies. It also shows some diagnostic statistics of estimated $\tilde{x}_t$, the noise component of the signals.

As expected, σ^2_{x1} is in every case much smaller than σ^2_x, reflecting the greater amount of information available at launch compared to the amount of information that emerges on an average day.

δ is positive for most cryptocurrencies, however it is approximately zero for Cardano and MaidSafeCoin, and negative for Bitcoin Cash, Polkadot, and EOS. Negative values indicate that the downward revisions to the estimate of eventual adoption outweighed the effect of adoption uncertainty resolution. Bitcoin has the highest δ, followed by many of the other high market capitalisation cryptocurrencies. The correspondence between high δ and high market capitalisation is no coincidence, as both result from strong historical price growth.

The interpretability of the $\log \mathcal{E}_\tau$ estimates is limited. They do give the model's estimate for the mode of the probability distribution of the eventual log price of the cryptocurrency, normalised for GWP growth. However, it is not clear that this is well identified by the model, given that it represents a far future quantity. The estimated value fluctuates depending on the most recent data point, and may be similarly sensitive to structural assumptions and the estimation method. Further, there is no corresponding estimate of τ, the time at which the estimated price $\mathcal{E}_\tau$ would prevail. While the model is well suited to explain price trends and project into the near future, robust identification of $\log \mathcal{E}_\tau$ may require extension.

One measure of goodness of fit is the variance of the signals, σ_{x1}^2 and σ_x^2. There is no natural threshold for concern; however, high values can indicate that little information about the cryptocurrency's eventual adoption level is available from imposing the model's structure on the price history. Dogecoin and MaidSafeCoin stand out in this regard.

The penalties for deviations between the theoretical and empirical kurtosis and skewness of $\tilde{x}_t$ result in these measures matching fairly closely, and without materially compromising the quality of the model fit to the price level, growth rate and variance. However, some cryptocurrencies show a mismatch between the theoretical and empirical variance of $\tilde{x}_t$, particularly for Dogecoin, XRP, and Stellar.

The diagnostic analysis of the $\tilde{x}_t$ series reveals that the use of a non-normal distribution is indeed necessary for most of the cryptocurrencies. Only three produce an insignificant Anderson-Darling test statistic for normality (at 10% confidence). All have positive skewness, and all except Ether and Solana have positive excess kurtosis. The optimal number of autoregressive lags according to the Bayes Information Criterion is zero for all, suggesting negligible autocorrelation of signals.

Table 3.2: Regression results and diagnostics

		Parameter estimates						$\tilde{x}_t$ diagnostics				
Ticker	Periods	$\log \mathcal{E}_\tau$	δ	σ^2_{x1}	σ^2_x	α	ν	Variance	Skewness (theoretical/ empirical)	Excess kurtosis (theoretical/ empirical)	Normality p-value	AR(p) order
BTC	179	17.62	0.91	20.7	1639.8	1.132	7.90	1676.4	0.62 / 0.66	2.16 / 2.14	0.019	0
ETH	118	11.12	0.56	18.2	975.9	0.870	2000.00	993.1	0.10 / 0.09	0.04 / -0.19	0.354	0
BNB	95	7.63	0.60	14.3	226.9	1.329	5.47	281.0	1.17 / 1.73	6.92 / 6.84	0.000	0
SOL	62	7.49	0.45	16.9	481.4	0.598	2000.00	497.5	0.04 / 0.06	0.01 / -0.47	0.843	0
XRP	142	2.21	0.12	35.7	2642.8	2.418	6.45	3777.0	1.40 / 1.73	5.62 / 5.42	0.000	0
DOGE	138	1.14	0.14	41.8	5374.9	2.178	5.82	7223.2	1.50 / 2.13	7.34 / 7.23	0.000	0
TRX	91	-0.82	0.43	10.2	108.5	1.282	14.15	113.3	0.44 / 0.42	0.90 / 0.71	0.148	0
ADA	92	0.41	0.07	18.9	775.7	3.099	7.15	960.6	1.42 / 1.76	4.84 / 4.71	0.000	0
BCH	94	5.23	-0.15	16.5	1043.6	1.399	7.66	1171.2	0.79 / 0.96	2.63 / 2.47	0.004	0
DOT	58	-0.33	-0.41	12.6	1195.6	7.327	40.40	1205.6	1.01 / 1.01	1.16 / 1.13	0.012	0
LTC	146	6.49	0.12	26.7	3536.2	1.446	5.12	3888.8	1.39 / 1.98	9.78 / 9.76	0.000	0
XLM	117	1.36	0.12	32.8	3652.4	3.949	5.85	4523.1	1.86 / 2.38	9.08 / 9.01	0.000	0
EOS	91	-2.52	-0.36	15.0	1369.1	2.608	16.46	1399.1	0.83 / 0.88	1.34 / 1.35	0.045	0
DASH	136	4.87	0.12	25.4	1994.6	2.690	17.60	2061.1	0.83 / 0.91	1.28 / 1.22	0.002	0
MAID	93	3.41	0.08	42.3	8620.8	4.442	1999.99	8430.5	0.82 / 0.69	0.68 / 0.60	0.181	0

α and ν are shape parameters of the Skewed t distribution, controlling skewness and kurtosis respectively. ν was restricted to be between 4.1 and 2000. Anderson-Darling normality test p-values are shown. AR(p) order is the order of autocorrelation p in the auxiliary regression equation $\tilde{x}_t = \sum_{i=1}^{p} \beta_i \tilde{x}_{t-i} + \epsilon_t$ that minimises the Bayes Information Criterion for that equation. Variance, empirical skewness, empirical kurtosis are that of $\{\tilde{x}_t\}_{t=2}^{T}$. Theoretical skewness and kurtosis are based on estimated α and ν.

3.5.3 Projections

Projections of the probability distribution of asset classes' returns are central to portfolio optimisation. Since the model shows an ability to replicate the core trends of cryptocurrency price growth and volatility, it is able to offer practical support to investors by projecting these forward.

The relevant information set is that of the model's investor in the last observation period. It produces projections that incorporate information embodied in the current price, the remaining uncertainty about $\log \mathcal{E}_\tau$, and the impact of the noise component of the signal.

Closed-form solutions are available for investors' expectation and variance of the price and price change over a future period. These are given in Equations 3.25 to 3.28 and are derived in Section 3.C.

Let $t < t' < t''$, and let $\mathbb{V}_{t,t'} := \Omega_t^{-1} - \Omega_{t'}^{-1}$ be the reduction in the remaining uncertainty about $\log \mathcal{E}_\tau$ between the two subscripted periods. Then:

$$\text{Var}_t \log \mathcal{E}_{t'} = \mathbb{V}_{t,t'} \tag{3.25}$$

$$\text{Var}_t \log \frac{\mathcal{E}_{t''}}{\mathcal{E}_{t'}} = \mathbb{V}_{t',t''} \tag{3.26}$$

$$\mathbb{E}_t \log \mathcal{E}_{t'} = \log \mathcal{E}_t - r^d_{t,t'} + \delta \mathbb{V}_{t,t'} \tag{3.27}$$

$$\mathbb{E}_t \log \frac{\mathcal{E}_{t''}}{\mathcal{E}_{t'}} = -r^d_{t',t''} + \delta \mathbb{V}_{t',t''} \tag{3.28}$$

These equations condition on prices up to time t. The results in this section set $t = T$ to use the latest information.

Equation 3.26 highlights that investors' perceived variance of the price change over a given interval is the portion of eventual adoption uncertainty that will be resolved during the interval. Intuitively, the more information received during a period, the greater the investor's potential update to their estimate of $\log \mathcal{E}_\tau$, and therefore the greater the price change variance over that period. Further, price change variances add up, such that $\mathbb{V}_{t,t'} + \mathbb{V}_{t',t''} = \mathbb{V}_{t,t''}$.

Figure 3.8: Bitcoin price projection

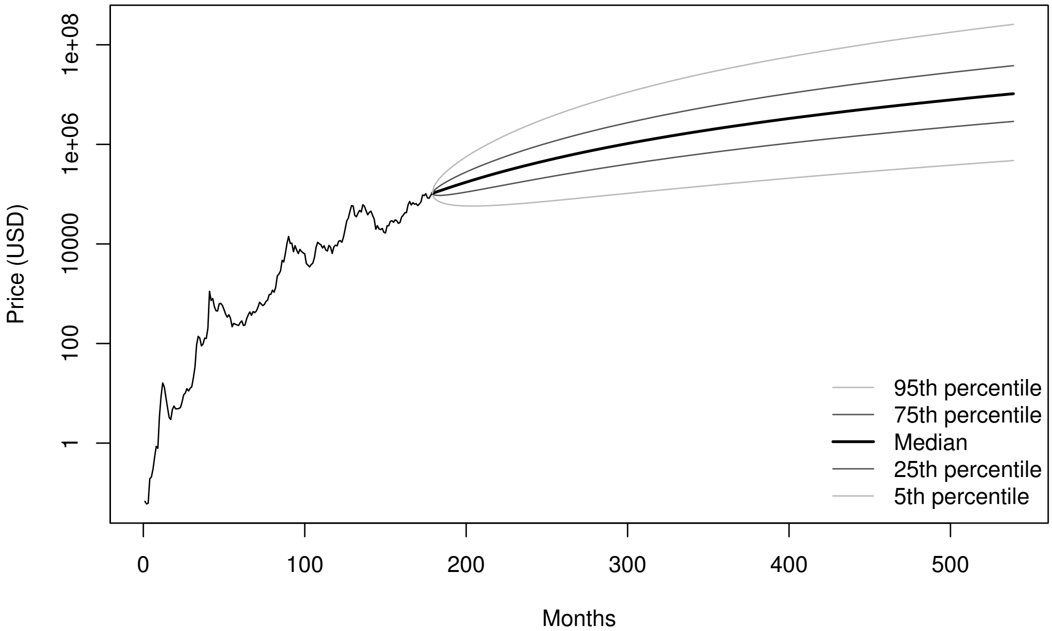

Price in USD, without normalisation. Projections start at the 180th month, which is June 2025.

For normally distributed signals, these closed-form solutions are sufficient to identify the probability distribution of the price and price change in every future period. However, for distributions such as skewed t, simulation is needed.

Figure 3.8 plots confidence intervals of the price (without normalisation) for the next 30 years for Bitcoin. The median price approximately continues the existing growth trend in the Bitcoin price. However, the figure shows substantial variation around the median. For example, taking the price on 31 May 2025 as the starting point, the model anticipates a 5% chance that the price will be lower in nominal terms after 10.2 years.

The expectation and standard deviation of Bitcoin price changes are presented in Figure 3.9. For the purpose of calculating a Sharpe ratio, I report these projections for the quantity $\frac{\mathcal{E}_{t''}}{\mathcal{E}_{t'}} - 1$, where $t'' = t' + 12$ is the last month of the projection year in question. For normally distributed signals, this quantity is lognormally distributed, and closed-form solutions for its expectation and variance exist (Equations 3.29 and

Figure 3.9: Bitcoin annual expected price growth and volatility

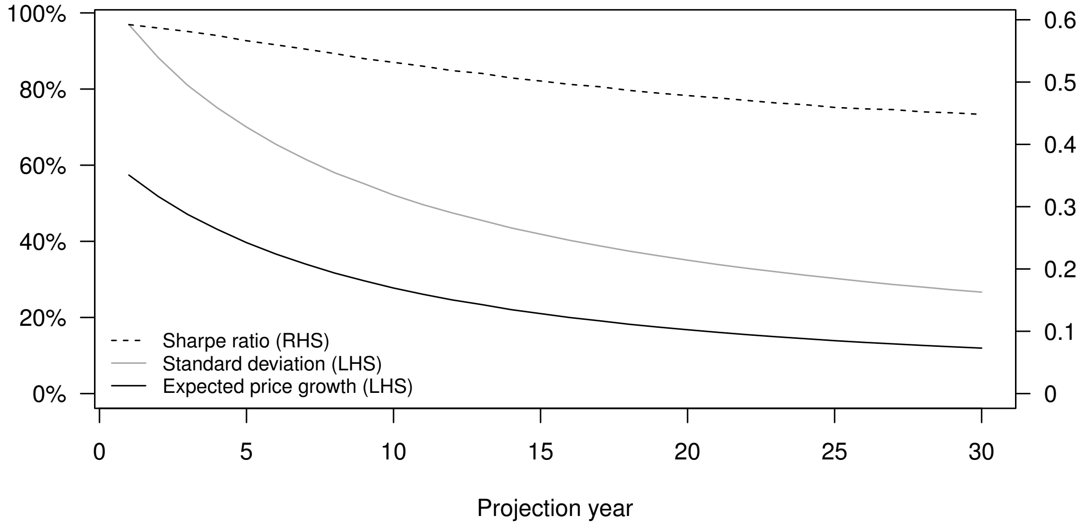

Price changes are not normalised for inflation or GWP growth. The first projection year starts from June 2025.

3.30). However, for skewed t distributed signals, it is again necessary to simulate.

$$\mathbb{E}_t \frac{\mathcal{E}_{t''}}{\mathcal{E}_{t'}} - 1 = e^{\left(\delta+\frac{1}{2}\right)\mathbb{V}_{t',t''} - r^d_{t',t''}} - 1 \tag{3.29}$$

$$\text{Var}_t \frac{\mathcal{E}_{t''}}{\mathcal{E}_{t'}} = e^{(2\delta+1)\mathbb{V}_{t',t''} - 2r^d_{t',t''}} \left(e^{\mathbb{V}_{t',t''}} - 1\right) \tag{3.30}$$

The expectation and standard deviation of price changes are projected to continue to fall, as more and more information about Bitcoin's eventual adoption price arrives. Since the distribution is non-stationary, portfolio optimisation may recommend a changing allocation over time. Note that the expectation and variance of a price change is the same under the structural information set, since they depend only on δ and the signal variances of future periods, which are assumed to be known under both information sets.

Section 3.F provides price and variance projections for all 15 cryptocurrencies.

A caveat to the projections in this section arises from the model's focus on the pre-eventual adoption period. The projections are conditional on the investment demand non-negativity constraint not binding. The model is not equipped to detect whether the constraint binds, or to forecast the timing of the next point at which it will bind.

Figure 3.10: Staking and Treasury bond yields

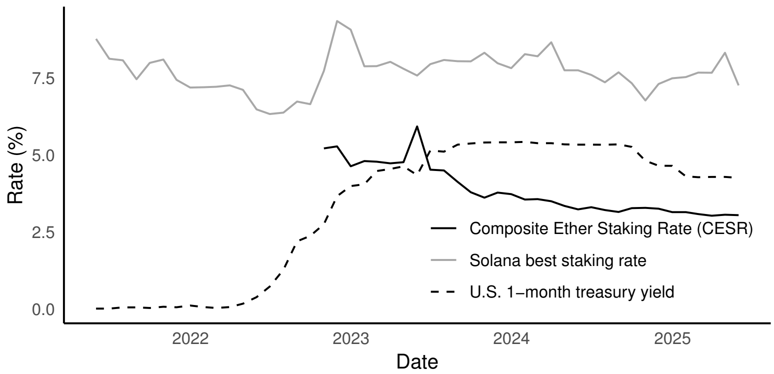

Source: Federal Reserve Economic Data (FRED), CoinDesk, and Solana Compass.

3.5.4 Sensitivity to zero interest rate differential assumption

In this section I relax the assumption that $r_t^d = 0 \ \forall \ t$ for Bitcoin, Ether, and Solana. I use the US 3-month Treasury bill rate for r_t, the fiat denominated risk-free investment.[14] For the Bitcoin denominated rate, I assume that $r_t^* = 0$ for $t \leq T$. For Ether, I assume that $r_t^* = 0$ during Ethereum's proof-of-work period, and use the Composite Ether Staking Rate (CESR; see Figure 3.10) from September 2022 up to time T. For Solana, I calculate the median "best" monthly staking rate. I assume that the future interest rate differentials are zero.

The results are given in Figure 3.11 and Table 3.3. In short, there is negligible impact. The model fit lines for both scenarios (with and without the interest rate differential) are approximately equal. It is possible that an assumption of a non-zero interest rate differential after period T is needed to have a material impact on estimated $\log \mathcal{E}_\tau$, though identification would require an assumed value for τ. A cryptocurrency with a more volatile interest differential may experience greater impact of its inclusion.

[14]Source: FRED

Figure 3.11: Comparison of fit with and without interest rate differential

(a) Bitcoin (BTC)

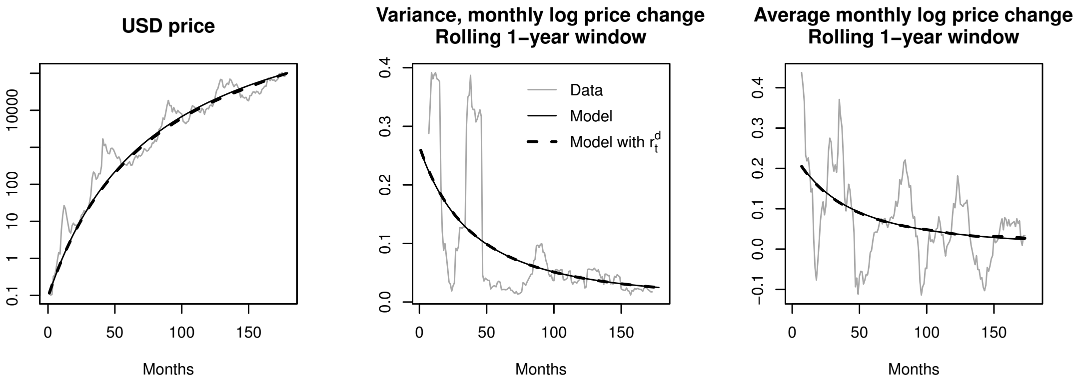

(b) Ether (ETH)

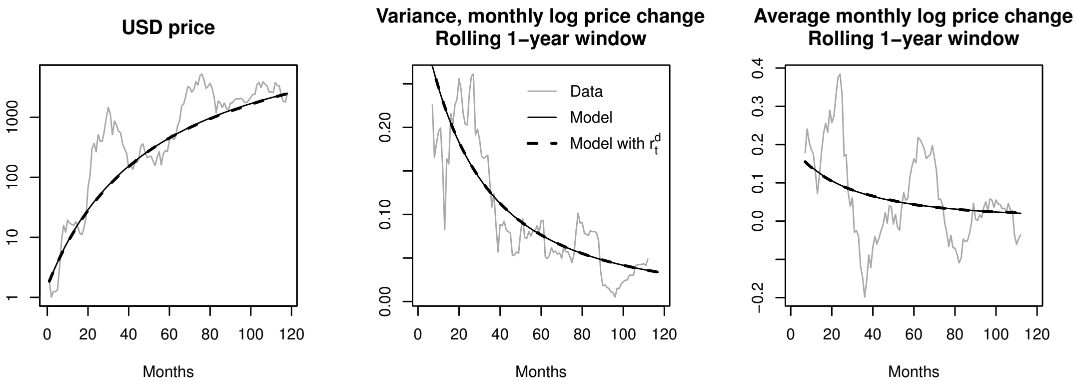

(c) Solana (SOL)

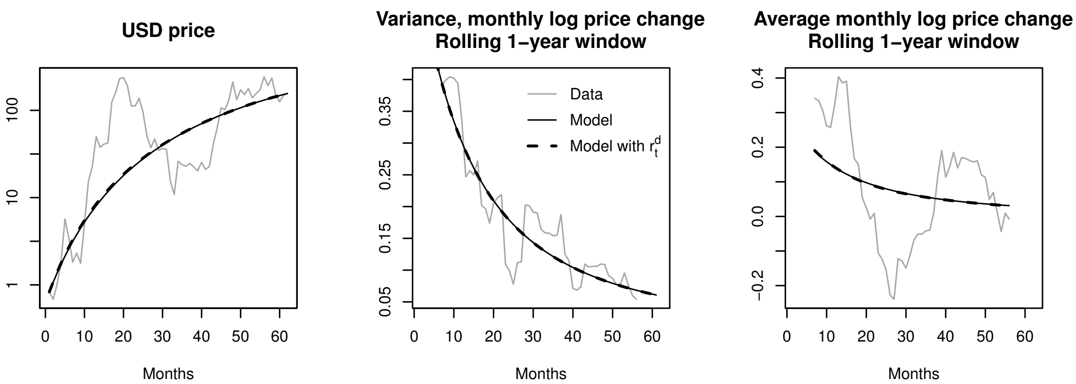

Table 3.3: Comparison of parameter values with and without interest rate differential

	$\log \mathcal{E}_\tau$	δ	σ^2_{x1}	σ^2_x	α	ν
Bitcoin	17.62	0.91	20.67	1639.8	1.132	7.90
Bitcoin with r^d_t	17.49	0.90	20.64	1623.5	1.134	7.89
Ether	11.12	0.56	18.22	975.9	0.870	2000.00
Ether with r^d_t	11.08	0.56	18.24	976.2	0.866	2000.00
Solana	7.49	0.45	16.92	481.4	0.598	2000.00
Solana with r^d_t	7.61	0.47	16.90	480.9	0.604	2000.00

3.6 Discussion

3.6.1 Other influences on cryptocurrency prices

Many think of cryptocurrency prices as being driven by fundamental factors such as the extent of a blockchain's non-investment usage, security, technological capabilities, relevant regulation, and the attractiveness of alternative payment systems. Existing applied theoretical literature has focussed on the hypothesised role of such factors (see Section 3.1).

However, the relationship is not straightforward. As per Athey et al. (2016) and van Oordt (2025), the presence of investors decouples the price from current fundamentals, instead linking them to the next point where prices will be determined by adoption demand. Accordingly, through its empirical validation, the model in this paper introduces evidence of the centrality of eventual adoption and its perceived probability distribution.

However, fundamental factors can still influence the current price. After all, the eventual adoption price depends on the fundamental factors prevailing at that point. Interim changes in fundamental factors may therefore constitute information signals, informing investors about their likely state at eventual adoption.

Any price impact of a given fundamental development depends not only on its nature, but on market knowledge prior to the development. Even a highly positive development may produce a negative price movement, or no price movement. For example, the implementation of Ethereum's transition to Proof-of-Stake on 15 September 2022 was a major, but long-expected, technological enhancement.

There was little discernible price impact (Lin et al., 2025). If the effects of its implementation were understood well in advance, then any price impact coincided with the understanding, not the implementation.

In practice however, new information often does accompany developments in fundamental factors. For example, the development may carry information about its precise nature, timing, and stakeholder reactions. At the least, once a development occurs, it dispels any remaining doubt about whether it will occur.

Therefore, without a valuation model similar to that in this paper, the price impact of any information signal surrounding developments may be mistaken for direct influence of the factor. I explore the explanatory potential of some commonly posited factors in the literature later in this section.

Separately from the effect of revisions to eventual adoption estimates, the model highlights the price impact of any reduction in eventual adoption uncertainty. For example, in 2024, many cryptocurrencies' prices increased in the weeks following the election of Donald Trump for a second U.S. presidential term. Observers tended to conclude that Trump's more supportive stance towards cryptocurrency increased prices. The model accommodates this view, since Trump's election can be viewed as an optimistic information signal. However, the model implies that prices could also have been lifted by the election outcome's revelation, independent of the outcome itself. The knowledge reduced investors' perceived variance of eventual adoption, reducing the uncertainty discount.

Of course, the model interprets all price movements as the result of information arrival. In reality, there are many other influences, particularly short-term impacts such as investor sentiment, idiosyncratic shocks to demand to hold the cryptocurrency, and short-term price bubbles. The model's success in describing long term price trends may mean that the learning mechanism is a strong influence at that timescale.

3.6.2 Technological capability improvements as a driver of cryptocurrency prices

As noted in the Introduction, technological advancement of a cryptocurrency's blockchain has been proposed by some of the literature as a key determinant of cryptocurrency prices. After all, much discussion and news within the industry is focussed on the latest such developments, and comparisons of blockchains often focus on differences in capabilities.

A contribution of this paper is to show that modelling technological enhancements is neither necessary nor sufficient to explain long-term price and volatility trends. The reason is that the model fits well to cryptocurrencies regardless of the associated blockchain's rate of technological progress.

Concretely, while the developer communities of blockchains like Ethereum and Solana target continual improvement, Bitcoin remains strongly focussed on safeguarding its existing degree of decentralisation, security, and scarcity (VanEck, 2025). The Bitcoin community is generally unwilling to compromise Bitcoin's track record in this regard by allowing modifications.[15] Despite this, Bitcoin has the highest δ value, i.e. price growth per reduction in eventual adoption uncertainty, among the cryptocurrencies studied (Table 3.2).

This is not to deny the importance of technological capability as a factor influencing eventual adoption. The more a blockchain's eventual adoption price depends on its functionality, the more the market will respond to information about development effort and achievement.

3.6.3 Gradual adoption as a driver of cryptocurrency prices

Past studies have frequently invoked gradual adoption processes to explain cryptocurrency prices. These admit adoption during the investment-driven phases, in

[15] The block size wars arose from tension surrounding this imperative. Indeed, the more change-resistant (BTC) side of the resulting fork experienced faster subsequent price growth than Bitcoin Cash (BCH).

contrast to the model of this paper. There is indeed substantial value in studying the effects of gradual adoption. For example, Bolt and van Oordt (2020) and Prat et al. (2025) point out that, if adopters demand an amount of the native cryptocurrency that is fixed in fiat currency terms, then these adopters act as shock absorbers for the price. They respond to any reduction in price by purchasing more. As the user base grows, price volatility reduces. This argues that gradual adoption is a factor whose developments matter directly for prices, in addition to informing investors about eventual adoption.

3.6.4 Fundamental value, active investment, and bubbles

Cryptocurrency fundamental value exists because investors can eventually sell to adopters. It relies on eventual adoption demand that is relatively insensitive to expected price appreciation, such as for transactions or staking. It is dissimilar to the fundamental value of equity of a company, or other cashflow-generating assets. However, cryptocurrency transactional demand is similar to the demand for conventional currency, for relieving a cash-in-advance constraint. Similarly, cryptocurrency staking demand is similar to the demand for a commodity such as land, in that they both have fundamental value as inputs to production.

Fundamental analysis, or active investment, would be profitable for investors who have better access to information than others. Although the model does not incorporate asymmetric information, it does imply a process for fundamental analysis. The process would require estimating the probability distribution of the eventual adoption price, the sum of future interest rate differentials, and δ, and then using Equation 3.4 at $t = T$ to derive a target price consistent with those estimates. Deviations of the market price from the target may indicate mispricing, to the extent that the analyst's estimates are superior.[16]

[16]This process does not rely on the model's learning component. However, learning could be used to estimate the probability distribution of the eventual adoption price. This would entail estimating all the model's parameters, including $\log \mathcal{E}_\tau$, and then using Equation 3.12 at $t = T$ to

Some postulate that major cryptocurrencies' prices are far above their fundamental value, as illustrated by the introductory quote from Nouriel Roubini concerning bubbles. Under this view, the entire history of prices of major cryptocurrencies, in particular Bitcoin, constitutes a bubble, and the price will eventually return to nearly zero. Importantly, this is distinct from the claim that there have been shorter-term price bubbles, such as may be represented by the peaks in the Bitcoin price in June 2011, December 2013, December 2017, and November 2021. Contrasting with the long-term bubble hypothesis, others expect substantial appreciation of the Bitcoin price.[17]

The model cannot directly assess whether long-term cryptocurrency prices are driven mainly by bubble dynamics, or otherwise deviate from fundamental value.[18] However, the model re-frames the debate about a long-term bubble as a debate about eventual adoption. Specifically, the model's large estimates of the variance of the eventual adoption price leave room for widely divergent yet plausible beliefs. Under this interpretation, those who are on the lower tail of the distribution of beliefs would calculate a fundamental value that is much lower than the market price, and would conclude that a bubble exists. Those on the upper tail would do the opposite. The disagreement underlines the salience of eventual adoption uncertainty and its resolution via learning.

3.6.5 Long-term investment

The model provides conceptual guidance to investors considering holding cryptocurrency. Specifically, the model offers an answer to the question of the nature

obtain an estimate of $\sum_{s=1}^{T} \tilde{x}_s w_{sT}$ as the market's pricing error. Deviations of this quantity from zero may indicate mispricing if the analyst's parameter estimates are superior to those implied by the market.

[17]For example, Larry Fink, CEO of BlackRock, an investment management firm (Bambrough, 2025).

[18]The scope for bubble and Ponzi scheme dynamics in cryptocurrency prices is explored theoretically by van Oordt (2025), using a model allowing rational bubbles alongside any fundamental value, and assessing the effects of net issuance.

of the risk exposure, and whether the risk is rewarded. The model interprets cryptocurrency price growth as the investors' reward for accepting exposure to eventual adoption uncertainty. The source of risk is that, whenever a signal arrives that is lower than the current estimate of the eventual adoption price, the current price will fall. Likewise, in response to high information signals, the price will rise. The source of reward is that any signal reduces the variance of the eventual adoption price estimate (Ω_t^{-1}), and hence reduces the risk discount in the pricing equation (Equation 3.12). Gradual reduction of the discount accounts for the significant price increases of some cryptocurrencies since their inception.

The model also provides practical guidance. In making portfolio allocation decisions, investors normally require estimates of the variance and expectation of future returns. The model's learning process that governs cryptocurrency price changes can be projected into the future. The resulting distribution of future prices (see Section 3.5.3) can serve as an input into portfolio construction.

3.7 Conclusion

This paper aims to explain the rapid growth of many major cryptocurrencies' prices. Although the growth has arguably underpinned the rise of the blockchain industry, existing literature and standard models have provided little guidance. This knowledge gap may partly explain the divergence of views on cryptocurrency fundamental value.

The model provides an explanation that is intuitive, empirically validated, and grounded in rational choice. Eventual adoption uncertainty, and its resolution through learning, create a declining price discount. As signals accumulate, prices become less sensitive to further information, reducing the average and variance of price growth. The model projects future expectations and variance of price growth, facilitating portfolio construction.

The model's combination of parsimony and quality of fit to many cryptocurrencies is instructive. First, it underscores the power of the learning mechanism. Second, it reframes the role of fundamental factors such as gradual adoption, network security, and technological enhancements. These factors do affect eventual adoption, and changes to investor knowledge of eventual adoption affect current prices. However, any changes to fundamental factors only affect current prices to the extent that they convey new information to investors about eventual adoption. Modelling of fundamental factors may therefore be most useful when incorporated into a framework of gradual learning about eventual adoption.

The paper leaves many avenues for extension and enhancement. The model only addresses the investment-driven regime, and could be expanded to also model the adoption-driven regime. Generalising the uncertainty to the entire time path of adoption demand may reveal investor views of the timing of adoption. Alternative signal processes, and additional sources of shock, may improve the model's fit. Modelling cryptocurrency supply, alongside some of the above enhancements, may enable a rigorous estimation of the probability distribution of eventual adoption demand, revealing investor views of a cryptocurrency's future prominence. Expansion to multiple cryptocurrencies would enable the study of their signal correlation. I leave these topics and others to future research.

Appendices to Chapter 3

3.A Maximum likelihood estimation first order conditions

Following from the likelihood equation (3.15), since u_t is a function of the previous $\tilde{x}_s$ values, where $s < t$, and since their values depend on the parameter values when the data is taken as given, it is necessary to compute the derivatives of $\tilde{x}_s$ with respect to δ and $\log \mathcal{E}_\tau$ in order to proceed.

$\tilde{x}_t$ is defined by the following rearrangement of the asset pricing equation (3.12).

$$\tilde{x}_t = w_{tt}^{-1}(\log \mathcal{E}_t - \log \mathcal{E}_\tau - r^d_{t,\tau}) + \frac{\delta}{\sigma_{xt}^{-2}} - \sum_{s=1}^{t-1} \tilde{x}_s \frac{\sigma_{xs}^{-2}}{\sigma_{xt}^{-2}} \tag{3.31}$$

Starting with δ, I will prove by induction that

$$\frac{\partial \tilde{x}_s}{\partial \delta} = 0, s \geq 2 \tag{3.32}$$

First, note that, for $s = 1$ and using Equation 3.31, $\frac{\partial \tilde{x}_1}{\partial \delta} = \Omega_1^{-1} = \sigma_{x1}^2$.

For $s = 2$, I must again calculate the derivative manually, as follows.

$$\frac{\partial \tilde{x}_2}{\partial \delta} = \frac{1}{\sigma_{x2}^{-2}} - \sigma_{x1}^2 \frac{\sigma_{x1}^{-2}}{\sigma_{x2}^{-2}} = 0$$

Now, for the induction step, assume that Equation 3.32 is true for all $s \geq 2$.

Then

$$\begin{aligned}\frac{\partial \tilde{x}_{t+1}}{\partial \delta} &= \frac{1}{\sigma_{x,t+1}^{-2}} - \frac{\partial \tilde{x}_1}{\partial \delta} \frac{\sigma_{x1}^{-2}}{\sigma_{x,t+1}^{-2}} \\ &= \frac{1}{\sigma_{x,t+1}^{-2}} - \sigma_{x1}^{-2} \frac{\sigma_{x1}^{-2}}{\sigma_{x,t+1}^{-2}} = 0\end{aligned}$$

Therefore Equation 3.32 holds for $s \geq 2$, so $\frac{\partial \tilde{x}_t}{\partial \delta} = \sigma_{x1}^2$ for $t = 1$, and 0 otherwise.

Now, for $\log \mathcal{E}_\tau$, I will prove by induction that

$$\frac{\partial \tilde{x}_s}{\partial \log \mathcal{E}_\tau} = -1, s \geq 1 \tag{3.33}$$

In the base case, $s = 1$. Note from Equation 3.31 that $\frac{\partial \tilde{x}_1}{\partial \log \mathcal{E}_\tau} = -1$, since $w_{11} = 1$. Assume that Equation 3.33 is true for all $s \geq 1$. The induction step is as follows.

$$\begin{aligned} \frac{\partial \tilde{x}_{t+1}}{\partial \log \mathcal{E}_\tau} &= -w_{t+1,t+1}^{-1} - \sum_{s=1}^{t} (-1) \frac{\sigma_{xs}^{-2}}{\sigma_{x,t+1}^{-2}} \\ &= -\frac{\Omega_{t+1}}{\sigma_{x,t+1}^{-2}} + \frac{\Omega_{t+1} - \sigma_{x,t+1}^{-2}}{\sigma_{x,t+1}^{-2}} = -1 \end{aligned}$$

Therefore Equation 3.33 holds for $s \geq 1$.

Using these equations, the partial derivatives of $\log L_t$ can be derived as follows.

$$\begin{aligned} \frac{\partial \log L_t}{\partial \delta} &= \frac{\partial \log g(u_t)}{\partial u_t} \frac{\partial u_t}{\partial \tilde{x}_t} \frac{\partial \tilde{x}_t}{\partial \delta} \\ &= \frac{g'(u_t)}{g(u_t)} w_{tt} \sigma_{x1}^2 \text{ for } t = 1; 0 \text{ otherwise} \\ \frac{\partial \log L_t}{\partial \log \mathcal{E}_\tau} &= \frac{\partial \log g(u_t)}{\partial u_t} \frac{\partial u_t}{\partial \tilde{x}_t} \frac{\partial \tilde{x}_t}{\partial \log \mathcal{E}_\tau} \\ &= \frac{g'(u_t)}{g(u_t)} w_{tt}(-1) \end{aligned}$$

where, as earlier, g is the probability density function of $u_t = \tilde{x}_t w_{tt}$.

Summing across $t \in [1, T]$ and setting to zero yields the first-order conditions given in Equations 3.16 and 3.17 in the paper.

In the case of the normal distribution, both first-order conditions simplify. For that of δ, note that the mode of the normal distribution is its mean, and $\mathbb{E}\tilde{x}_t = 0$, so the mode condition is satisfied when $\tilde{x}_1 = 0$. For that of $\log \mathcal{E}_\tau$, the workings proceed from Equation 3.17 as follows.

$$\begin{aligned} g(u_t) &= \frac{1}{\sqrt{2\pi w_{tt}^2 \sigma_{xt}^2}} \exp\left(-\frac{1}{2} \frac{u_t^2}{w_{tt}^2 \sigma_{xt}^2}\right) \\ g'(u_t) &= -\frac{u_t}{w_{tt}^2 \sigma_{xt}^2} g(u_t) \\ \frac{g'(u_t)}{g(u_t)} &= -w_{tt} \tilde{x}_t w_{tt}^{-2} \sigma_{xt}^{-2} \end{aligned}$$

$$= -\tilde{x}_t \sigma_{xt}^{-2} w_{tt}^{-1}$$

Plugging back into Equation 3.19,

$$0 = \sum_{s=1}^{T} -\tilde{x}_s \sigma_{xs}^{-2}$$
$$= \sum_{s=1}^{T} \tilde{x}_s w_{sT}$$

where the last step divides by $-\Omega_T$, giving the required result.

Finally, from Equation 3.12, note that

$$\log \mathcal{E}_1 = \log \mathcal{E}_\tau + r_{t,\tau}^d - \delta \Omega_1^{-1} + \tilde{x}_1$$
$$\log \mathcal{E}_T = \log \mathcal{E}_\tau + r_{t,\tau}^d - \delta \Omega_T^{-1} + \sum_{s=1}^{T} \tilde{x}_s w_{sT}$$

When using a normal distribution, the MLE first-order conditions set the last term in each of these equations to zero, leaving only the structurally expected price on the right-hand side.

3.B Derivation of the pricing equation

In this section I derive the main asset pricing equation for both the case where the investor only consumes in the final period and the case where there is consumption in every period. For both I include the interest rates native to both cryptocurrency and fiat currency, although in the base model used for estimation in this paper, the interest differential is set to zero in all periods.

Starting with the case in which consumption is only in the final period, the investor solves the following problem.

$$\max_{\omega_t} \ \mathbb{E}_t U(C_\tau) \qquad \text{s.t.} \qquad C_\tau = W_\tau = W_t e^{\omega_t \left(\log \frac{\mathcal{E}_\tau}{\mathcal{E}_t} + r_{t,\tau}^d \right) + r_{t,\tau}}$$

Redefine the utility function such that $\hat{U}(\log C_\tau) = U(C_\tau)$ and use a second-order approximation of the expectation of the redefined utility function, such that

$$\mathbb{E}_t \hat{U}(\log C_\tau) = \mathbb{E}_t \log C_\tau - \frac{1}{2} \delta \text{Var}_t \log C_\tau$$

where risk aversion coefficient δ is assumed constant and treated as a parameter. Then the investors' problem can be defined in terms of the expectation and variance of the log return.

$$\begin{aligned}
&\max_{\omega_t}\ \mathbb{E}_t\left(\log W_t + \omega_t\left(\log\frac{\mathcal{E}_\tau}{\mathcal{E}_t} + r^d_{t,\tau}\right) + r_{t,\tau}\right)\\
&\qquad -\frac{1}{2}\delta\mathrm{Var}_t\left(\log W_t + \omega_t\left(\log\frac{\mathcal{E}_\tau}{\mathcal{E}_t} + r^d_{t,\tau}\right) + r_{t,\tau}\right)\\
=&\max_{\omega_t}\ \log W_t + \omega_t\left(\mathbb{E}_t\log\frac{\mathcal{E}_\tau}{\mathcal{E}_t} + r^d_{t,\tau}\right) + r_{t,\tau} - \frac{1}{2}\delta\omega_t^2\mathrm{Var}_t\log\mathcal{E}_\tau
\end{aligned}$$

The first-order condition then gives

$$\log\mathcal{E}_t = \mathbb{E}_t\log\mathcal{E}_\tau + r^d_{t,\tau} - \delta\omega_t\mathrm{Var}_t\log\mathcal{E}_\tau$$

which, when combined with the market clearing condition $\omega_t = 1$, and the investors' estimates of $\log\mathcal{E}_\tau$ and its variance, gives the required pricing equation.

Moving to the case of consumption in every period, consider now an agent who maximises the following value function each period by choosing C_t and the allocation to cryptocurrency ω_t

$$V(W_t) = U(C_t) + \beta\mathbb{E}_t V(W_{t+1})$$

subject to

$$W_{t+1} = (W_t + Y_t - C_t)e^{\omega_t\left(\log\frac{\mathcal{E}_{t+1}}{\mathcal{E}_t} + r^*_t\right) + (1-\omega_t)r_t}$$

where Y_t is exogenous income and W_t is wealth.

Now redefine the value function such that

$$\hat{V}(\log W_t) = V(W_t)$$

Also, use a second-order approximation of the expectation of the redefined value function, such that

$$\mathbb{E}_t\hat{V}(\log W_{t+1}) \approx \mathbb{E}_t\log W_{t+1} - \frac{1}{2}\delta\mathrm{Var}_t\log W_{t+1}$$

The value function can then be approximated as follows:

$$\begin{aligned}V(W_t) &\approx U(C_t) + \beta\Big[\log(W_t + Y_t - C_t) + \omega_t\Big(\mathbb{E}_t \log\frac{\mathcal{E}_{t+1}}{\mathcal{E}_t} + r_t^*\Big) + (1-\omega_t)r_t \\ &\quad -\frac{1}{2}\delta \mathrm{Var}_t\Big(\log(W_t + Y_t - C_t) + \omega_t\Big(\log\frac{\mathcal{E}_{t+1}}{\mathcal{E}_t} + r_t^*\Big) + (1-\omega_t)r_t\Big)\Big] \\ &\approx U(C_t) + \beta\Big[\log(W_t + Y_t - C_t) + \omega_t\Big(\mathbb{E}_t \log\frac{\mathcal{E}_{t+1}}{\mathcal{E}_t} + r_t^*\Big) + (1-\omega_t)r_t \\ &\quad -\frac{1}{2}\delta\omega_t^2 \mathrm{Var}_t \log\mathcal{E}_{t+1}\Big]\end{aligned}$$

The first-order condition for ω_t is

$$0 = \beta\left(\mathbb{E}_t \log\frac{\mathcal{E}_{t+1}}{\mathcal{E}_t} + r_t^d - \delta\omega_t \mathrm{Var}_t \log\mathcal{E}_{t+1}\right) \tag{3.34}$$

The remainder of this section will show by induction that, combined with the market-clearing condition of $\omega_t = 1$ and the exogenous cryptocurrency price in the final period $\log\mathcal{E}_\tau$, Equation 3.34 produces the same asset pricing equation as before, which identifies the price of the cryptocurrency in each period as a function of the model's parameters and the signals. Although it is possible to also derive an approximate first-order condition governing consumption choices in this setting, these choices do not affect the cryptocurrency prices. Instead, prices are determined by investors' perception of the probability distribution of eventual adoption, combined with their risk aversion.

Set $\omega_t = 1$ and conjecture that

$$\log\mathcal{E}_t = \mathbb{E}_t \log\mathcal{E}_\tau + r_{t,\tau}^d - \delta \mathrm{Var}_t \log\mathcal{E}_\tau \tag{3.35}$$

for all $t < \tau$, noting that $r_{t,\tau}^d = \sum_{s=t}^{\tau-1} r_s^d$. This is satisfied in the case where $t = \tau - 1$, as per Equation 3.34.

For the induction step, consider the previous period log price according to the first-order condition, and substitute in the conjecture.

$$\begin{aligned}\log\mathcal{E}_{t-1} &= \mathbb{E}_{t-1}\log\mathcal{E}_t + r_{t-1}^d - \delta \mathrm{Var}_{t-1}\log\mathcal{E}_t \\ &= \mathbb{E}_{t-1}\Big(\mathbb{E}_t \log\mathcal{E}_\tau + r_{t,\tau}^d - \delta \mathrm{Var}_t \log\mathcal{E}_\tau\Big) + r_{t-1}^d\end{aligned}$$

$$- \delta \text{Var}_{t-1} \left(\mathbb{E}_t \log \mathcal{E}_\tau + r^d_{t,\tau} - \delta \text{Var}_t \log \mathcal{E}_\tau \right)$$

$$= \mathbb{E}_{t-1} \log \mathcal{E}_\tau + r^d_{t-1,\tau} - \delta (\text{Var}_t \log \mathcal{E}_\tau + \text{Var}_{t-1} \mathbb{E}_t \log \mathcal{E}_\tau)$$

where the last line uses the law of iterated expectations, and the fact that, since all variances are known ahead of time, the variance of the variance is zero.

To obtain an expression for $\text{Var}_{t-1}\mathbb{E}_t \log \mathcal{E}_\tau$, apply the investors' weighted least squares estimates for the expectation and variance of the log eventual adoption price.

$$\begin{aligned}
\text{Var}_{t-1}\mathbb{E}_t \log \mathcal{E}_\tau &= \text{Var}_{t-1} \sum_{s=1}^{t} x_s w_{st} \\
&= \text{Var}_{t-1}(x_t w_{tt}) \\
&= w_{tt}^2 \text{Var}_{t-1}(\log \mathcal{E}_\tau + \tilde{x}_t) \\
&= \sigma_{xt}^{-4} \Omega_t^{-2} (\Omega_{t-1}^{-1} + \sigma_{xt}^2) \\
&= \sigma_{xt}^{-2} \Omega_t^{-1} (\Omega_{t-1}^{-1} \Omega_t^{-1} \sigma_{xt}^{-2} + \Omega_t^{-1}) \\
&= (\Omega_t - \Omega_{t-1}) \Omega_t^{-1} (\Omega_{t-1}^{-1} \Omega_t^{-1} (\Omega_t - \Omega_{t-1}) + \Omega_t^{-1}) \\
&= (\Omega_t - \Omega_{t-1}) \Omega_t^{-1} (\Omega_{t-1}^{-1} - \Omega_t^{-1} + \Omega_t^{-1}) \\
&= (\Omega_t - \Omega_{t-1}) \Omega_t^{-1} \Omega_{t-1}^{-1} \\
&= \Omega_{t-1}^{-1} - \Omega_t^{-1} \\
&= \text{Var}_{t-1} \log \mathcal{E}_\tau - \text{Var}_t \log \mathcal{E}_\tau
\end{aligned}$$

Substituting these variance formulae back in,

$$\begin{aligned}
\log \mathcal{E}_{t-1} &= \mathbb{E}_{t-1} \log \mathcal{E}_\tau + r^d_{t-1,\tau} - \delta (\text{Var}_t \log \mathcal{E}_\tau + \text{Var}_{t-1} \log \mathcal{E}_\tau - \text{Var}_t \log \mathcal{E}_\tau) \\
&= \mathbb{E}_{t-1} \log \mathcal{E}_\tau + r^d_{t-1,\tau} - \delta \text{Var}_{t-1} \log \mathcal{E}_\tau
\end{aligned}$$

which matches the conjecture. Again, the investors' estimates of $\log \mathcal{E}_\tau$ and its variance can be plugged in to obtain the pricing equation.

3.C Derivation of formulae for expectation and variance of prices and price changes under investor and structural information sets

In this section I derive the formulae set out in Table 3.4. The investor and structural information sets are defined in Equations 3.9 and 3.10, and explained by the surrounding text. All derivations in this section start from the asset pricing equation (3.12), and refer to time periods $t < t' < t''$.

Table 3.4: Formulae for price level and change

Information set	Expectation	Variance
Structural	$\mathbb{E}_0^{\mathcal{E}_\tau} \log \mathcal{E}_t = \log \mathcal{E}_\tau + r^d_{t,\tau} - \delta\Omega_t^{-1}$	$\text{Var}_0^{\mathcal{E}_\tau} \log \mathcal{E}_t = \Omega_t^{-1}$
	$\mathbb{E}_0^{\mathcal{E}_\tau} \log \frac{\mathcal{E}_{t'}}{\mathcal{E}_t} = -r^d_{t,t'} + \delta(\Omega_t^{-1} - \Omega_{t'}^{-1})$	$\text{Var}_0^{\mathcal{E}_\tau} \log \frac{\mathcal{E}_{t'}}{\mathcal{E}_t} = \Omega_t^{-1} - \Omega_{t'}^{-1}$
Investor	$\mathbb{E}_t \log \mathcal{E}_{t'} = \log \mathcal{E}_t - r^d_{t,t'} + \delta(\Omega_t^{-1} - \Omega_{t'}^{-1})$	$\text{Var}_t \log \mathcal{E}_{t'} = \Omega_t^{-1} - \Omega_{t'}^{-1}$
	$\mathbb{E}_t \log \frac{\mathcal{E}_{t''}}{\mathcal{E}_{t'}} = -r^d_{t',t''} + \delta(\Omega_{t'}^{-1} - \Omega_{t''}^{-1})$	$\text{Var}_t \log \frac{\mathcal{E}_{t''}}{\mathcal{E}_{t'}} = \Omega_{t'}^{-1} - \Omega_{t''}^{-1}$

3.C.1 Derivation of structural expectations

Price level

$$\mathbb{E}_0^{\mathcal{E}_\tau} \log \mathcal{E}_t = \log \mathcal{E}_\tau + r^d_{t,\tau} - \delta\Omega_t^{-1} \quad \text{since } \mathbb{E}_0^{\mathcal{E}_\tau} \tilde{x}_t = 0$$

Price change (first given in Equation 3.24)

$$\begin{aligned}\mathbb{E}_0^{\mathcal{E}_\tau} \log \frac{\mathcal{E}_{t'}}{\mathcal{E}_t} &= \mathbb{E}_0^{\mathcal{E}_\tau} \log \mathcal{E}_{t'} - \mathbb{E}_0^{\mathcal{E}_\tau} \log \mathcal{E}_t \\ &= r^d_{t',\tau} - r^d_{t,\tau} + \delta\left(\Omega_t^{-1} - \Omega_{t'}^{-1}\right) \\ &= -r^d_{t,t'} + \delta\left(\Omega_t^{-1} - \Omega_{t'}^{-1}\right)\end{aligned}$$

3.C.2 Derivation of structural variances

Price level

$$\begin{aligned}\text{Var}_0^{\mathcal{E}_\tau} \log \mathcal{E}_t &= \text{Var}_0^{\mathcal{E}_\tau} \sum_{s=1}^{t} \tilde{x}_s w_{st} \\ &= \sum_{s=1}^{t} \sigma_{xs}^2 w_{st}^2\end{aligned}$$

$$
\begin{aligned}
&= \sum_{s=1}^{t} \sigma_{xs}^{-2}\Omega_t^{-2} \\
&= \Omega_t^{-2}\Omega_t \\
&= \Omega_t^{-1} \qquad (3.36)
\end{aligned}
$$

Price change (first given in Equation 3.23)

$$
\begin{aligned}
\mathrm{Var}_0^{\mathcal{E}_\tau} \log \frac{\mathcal{E}_{t'}}{\mathcal{E}_t} &= \mathrm{Var}_0^{\mathcal{E}_\tau} \left(\sum_{s=1}^{t'} w_{st'}\tilde{x}_s - \sum_{s=1}^{t} w_{st}\tilde{x}_s \right) \\
&= \mathrm{Var}_0^{\mathcal{E}_\tau} \left(\sum_{s=t+1}^{t'} w_{st'}\tilde{x}_s + \sum_{s=1}^{t}(w_{st'} - w_{st})\tilde{x}_s \right) \\
&= \sum_{s=t+1}^{t'} w_{st'}^2\sigma_{xs}^2 + \sum_{s=1}^{t}(w_{st'} - w_{st})^2\sigma_{xs}^2 \\
&= \Omega_{t'}^{-2}(\Omega_{t'} - \Omega_t) + \Omega_t(\Omega_{t'}^{-1} - \Omega_t^{-1})^2 \\
&= \Omega_t^{-1} - \Omega_{t'}^{-1}
\end{aligned}
$$

3.C.3 Derivation of investor expectations

For investor expectations, it is convenient to begin from the below expression of the pricing equation, which does not separate x_t into its components:

$$
\log \mathcal{E}_t = \sum_{s=1}^{t} x_s w_{st} + r_{t,\tau}^d - \delta\Omega_t^{-1}
$$

Price level (first given in Equation 3.27)

$$
\begin{aligned}
\mathbb{E}_t \log \mathcal{E}_{t'} &= \mathbb{E}_t \sum_{s=1}^{t'} x_s w_{st'} + r_{t',\tau}^d - \delta\Omega_{t'}^{-1} \\
&= \mathbb{E}_t\mathbb{E}_{t'} \log \mathcal{E}_\tau + r_{t',\tau}^d - \delta\Omega_{t'}^{-1} \\
&= \mathbb{E}_t \log \mathcal{E}_\tau + r_{t',\tau}^d - \delta\Omega_{t'}^{-1} \quad \text{by law of iterated expectations} \\
&= \sum_{s=1}^{t} x_s w_{st} + r_{t',\tau}^d - \delta\Omega_{t'}^{-1} \\
&= \log \mathcal{E}_t - r_{t,\tau}^d + \delta\Omega_t^{-1} + r_{t',\tau}^d - \delta\Omega_{t'}^{-1} \\
&= \log \mathcal{E}_t - r_{t,t'}^d + \delta(\Omega_t^{-1} - \Omega_{t'}^{-1})
\end{aligned}
$$

Price change (first given in Equation 3.28)

$$
\mathbb{E}_t \log \frac{\mathcal{E}_{t''}}{\mathcal{E}_{t'}} = \mathbb{E}_t \log \mathcal{E}_{t''} - \mathbb{E}_t \log \mathcal{E}_{t'}
$$

$$= -r^d_{t',t''} + \delta(\Omega_{t'}^{-1} - \Omega_{t''}^{-1})$$

3.C.4 Derivation of investor variances

Price level (first given in Equation 3.25)

$$\begin{aligned}
\mathrm{Var}_t \log \mathcal{E}_{t'} &= \mathrm{Var}_t \left(\sum_{s=1}^{t'} x_s w_{st'} - \delta \Omega_{t'}^{-1} \right) \\
&= \mathrm{Var}_t \left(\sum_{s=t+1}^{t'} x_s w_{st'} \right) \\
&= \mathrm{Var}_t \left(\log \mathcal{E}_\tau \sum_{s=t+1}^{t'} w_{st'} + \sum_{s=t+1}^{t'} \tilde{x}_s w_{st'} \right) \\
&= \Omega_t^{-1} \left(\Omega_{t'}^{-1}(\Omega_{t'} - \Omega_t) \right)^2 + \sum_{s=t+1}^{t'} \sigma_{xs}^2 \left(\frac{\sigma_{xs}^{-2}}{\Omega_{t'}} \right)^2 \\
&= \Omega_t^{-1} \left(\Omega_{t'}^{-1}(\Omega_{t'} - \Omega_t) \right)^2 + \Omega_{t'}^{-2}(\Omega_{t'} - \Omega_t) \\
&= \Omega_t^{-1} - \Omega_{t'}^{-1}
\end{aligned}$$

Price change (first given in Equation 3.26)

$$\begin{aligned}
\mathrm{Var}_t \log \frac{\mathcal{E}_{t''}}{\mathcal{E}_{t'}} &= \mathrm{Var}_t \left(\sum_{s=t+1}^{t''} w_{st''} x_s - \sum_{s=t+1}^{t'} w_{st'} x_s \right) \\
&= \mathrm{Var}_t \Bigg[\log \mathcal{E}_\tau \left(\sum_{s=t+1}^{t''} w_{st''} - \sum_{s=t+1}^{t'} w_{st'} \right) + \sum_{s=t'+1}^{t''} w_{st''} \tilde{x}_s \\
&\qquad + \sum_{s=t+1}^{t'} (w_{st''} - w_{st'}) \tilde{x}_s \Bigg] \\
&= \Omega_t^{-1} \left(\sum_{s=t+1}^{t''} w_{st''} - \sum_{s=t+1}^{t'} w_{st'} \right)^2 + \sum_{s=t'+1}^{t''} w_{st''}^2 \sigma_{xs}^2 \\
&\qquad + \sum_{s=t+1}^{t'} (w_{st''} - w_{st'})^2 \sigma_{xs}^2 \\
&= \Omega_t^{-1} \left(\Omega_{t''}^{-1}(\Omega_{t''} - \Omega_t) - \Omega_{t'}^{-1}(\Omega_{t'} - \Omega_t) \right)^2 + \Omega_{t''}^{-2}(\Omega_{t''} - \Omega_{t'}) \\
&\qquad + \sum_{s=t+1}^{t'} \sigma_{xs}^{-2} (\Omega_{t''}^{-1} - \Omega_{t'}^{-1})^2 \\
&= \Omega_t^{-1} \left(\Omega_{t''}^{-1}(\Omega_{t''} - \Omega_t) - \Omega_{t'}^{-1}(\Omega_{t'} - \Omega_t) \right)^2 + \Omega_{t''}^{-2}(\Omega_{t''} - \Omega_{t'}) \\
&\qquad + (\Omega_{t''}^{-1} - \Omega_{t'}^{-1})^2 (\Omega_{t'} - \Omega_t) \\
&= \Omega_{t'}^{-1} - \Omega_{t''}^{-1}
\end{aligned}$$

3.D Cryptocurrency-native risk-free interest rates

In this section I explore the appropriateness of yields on cryptocurrency lending and staking as risk-free yields, and the futures basis as a risk-free interest differential, for the purpose of fitting the model.

In the case of borrowing and lending of cryptocurrencies, it is unlikely that lenders perceive this activity to be risk-free. To the author's knowledge, no governments or highly creditworthy entities have issued cryptocurrency-denominated bonds. Lending is primarily organised through lending exchange services, either by off-chain companies, or on-chain smart contracts. Lending orchestrated off-chain generally exposes lenders to counterparty risk of the lending exchange, and the credit risk of those to whom they in turn lend. In late 2022 and early 2023, many such lending platforms failed.[19] On-chain lending using smart contracts, with over-collateralisation of loans and automatic liquidation, can materially reduce credit risk. However, counterparty risk is replaced by the potential for protocol vulnerability due to coding errors, and the possibility of de-peg events if the borrowing relies on a cross-chain solution. Newer or unaudited smart contracts are more exposed. Importantly, cryptocurrency-denominated lenders are not covered by any lender of last resort.

Turning to staking, the most widely used staking contracts tend to be those associated with the consensus and security mechanisms of major proof-of-stake blockchains (see Figure 3.10). However, the associated staking rates have likely included risk premia. Taking Ethereum as an example, any protocol risk, or the risk of errors made by stakers while running a validator, could result in loss of funds. The Ether staking rate, before the inception of the Composite Ether Staking Rate (CESR), declined from an initial 20% yield (Rasmussen et al., 2022). Although these early rates reflect in part that Ether staking was fully illiquid until September 2022, the continued decline evident in Figure 3.10 could indicate that

[19] For example, BlockFi, Celsius, Voyager Digital, and Genesis.

the market is still maturing. Further, staking yields are not generally known in advance, unlike risk-free yields. Fixed yields have occasionally been arranged for Ether staking through fixed-for-floating interest rate swaps, with the floating leg being the staking rate (Lido, 2024). If this market flourishes in future, generating a risk-free interest rate time series, and exposing market expectations of future interest rates, the model could be fit to this data. Ether staking is not fully liquid; however, liquid staking solutions do exist.[20]

The futures basis is another potential measure of the historical interest rate differential. However, even for cryptocurrencies with liquid futures markets, this tends to be highly volatile. Schmeling et al. (2023) attribute this to a volatile futures convenience yield rather than the interest rate differential.

Aside from the data issues for historic interest differentials, the model requires an assumption about the future interest rate differential after the end of the data period, as per the asset pricing equation (3.12). The influence of the assumption may be material given it would apply for $t \in [T+1, \tau-1]$. Further, if non-zero future r_t^d were assumed, τ would become an estimated parameter. Also, constant $r_t^d = \overline{r^d} \neq 0$ could not be assumed for future periods, since this would prevent separate identification of $\log \mathcal{E}_\tau$ and τ. That is, the asset pricing equation would include $\log \mathcal{E}_\tau + \tau\overline{r^d}$ as the only appearance of these parameters, such that any maximised log likelihood sum could be produced by an infinite number of combinations of these parameters' values. The issues surrounding future interest rate differentials and τ estimation are avoided if the model is transformed so that only price changes, rather than levels, are fitted to data. However, this would not produce an estimate of either $\log \mathcal{E}_\tau$ or τ.

Overall, an assumption of a zero interest rate differential is necessary given the impracticality of measuring the historical risk-free interest rate differential of the various cryptocurrencies in this study, and given the need to fit the model to price levels, not only changes.

[20]As of April 2025, liquid staking services are provided by Lido, Binance, Coinbase, and others.

3.E Further graphical results

This section provides graphical model results for XRP, BNB, SOL, DOGE, TRX, ADA, DOT, LTC, XLM, EOS, and MAID, in addition to the results presented in the body of the paper for BTC (Figure 3.3), ETH (Figure 3.4), DASH (Figure 3.5), and BCH (Figure 3.7).

These charts show the data against the model's structural price mode (no closed-form solution), structural log price change variance (Equation 3.23), and structural expected log price change (Equation 3.24).

Figure 3.12: XRP (XRP) model fit

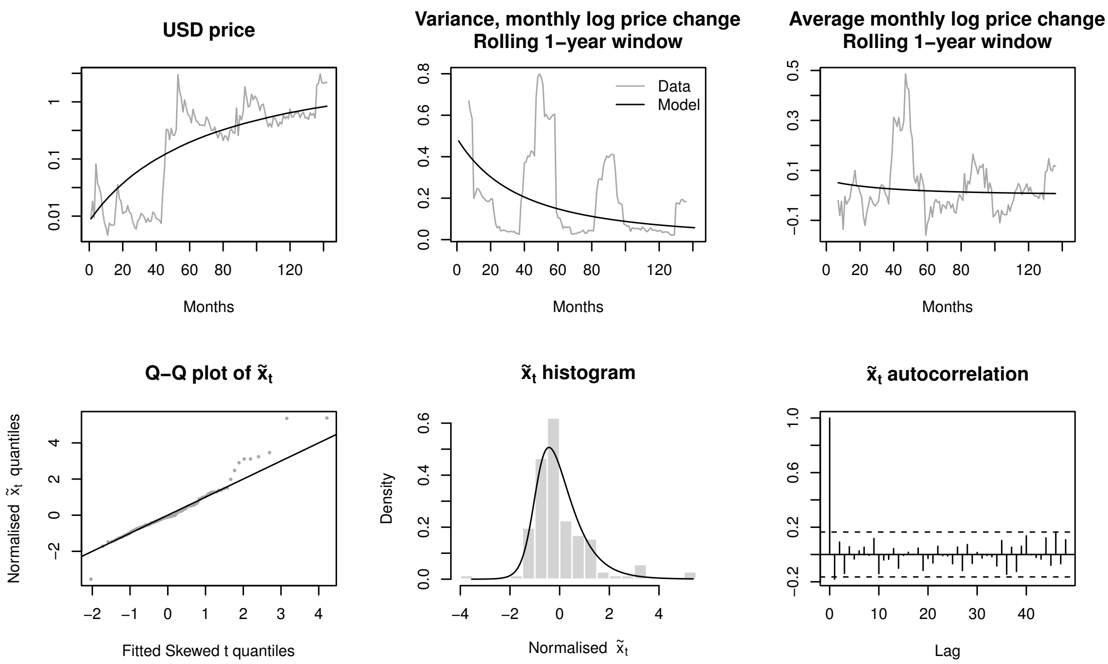

Figure 3.13: Binance Coin (BNB) model fit

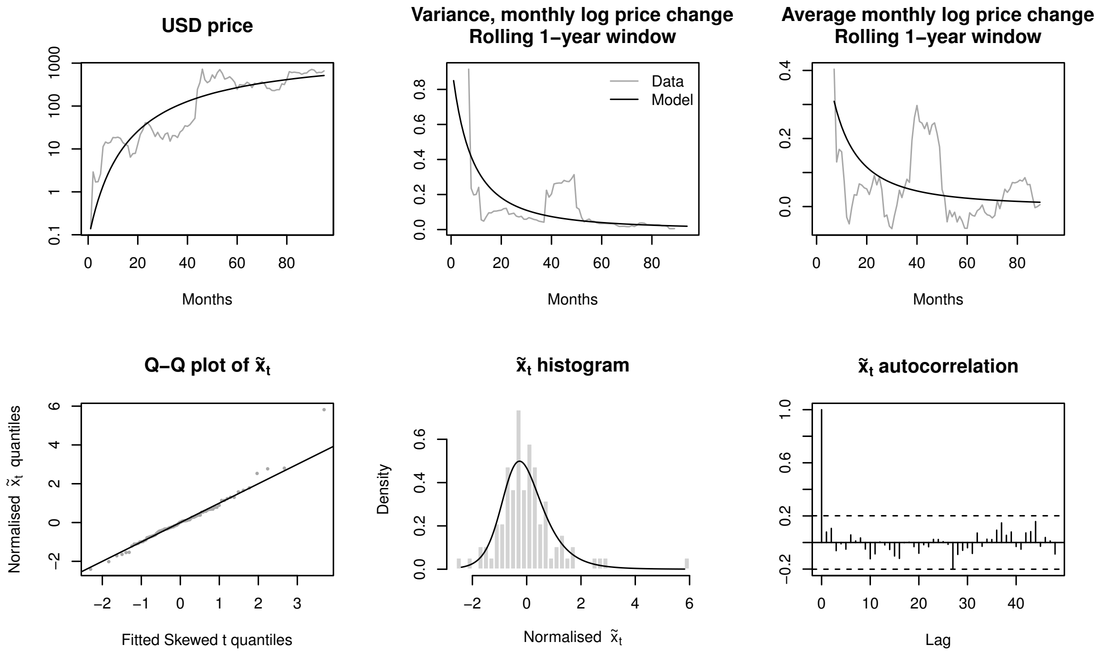

Figure 3.14: Solana (SOL) model fit

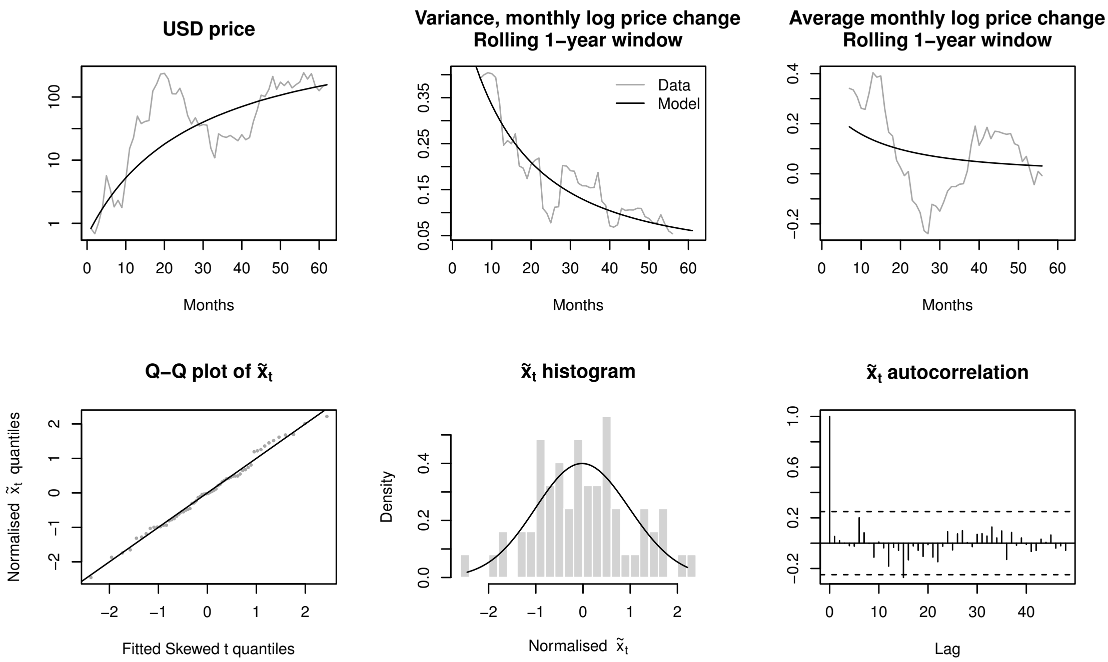

Figure 3.15: Dogecoin (DOGE) model fit

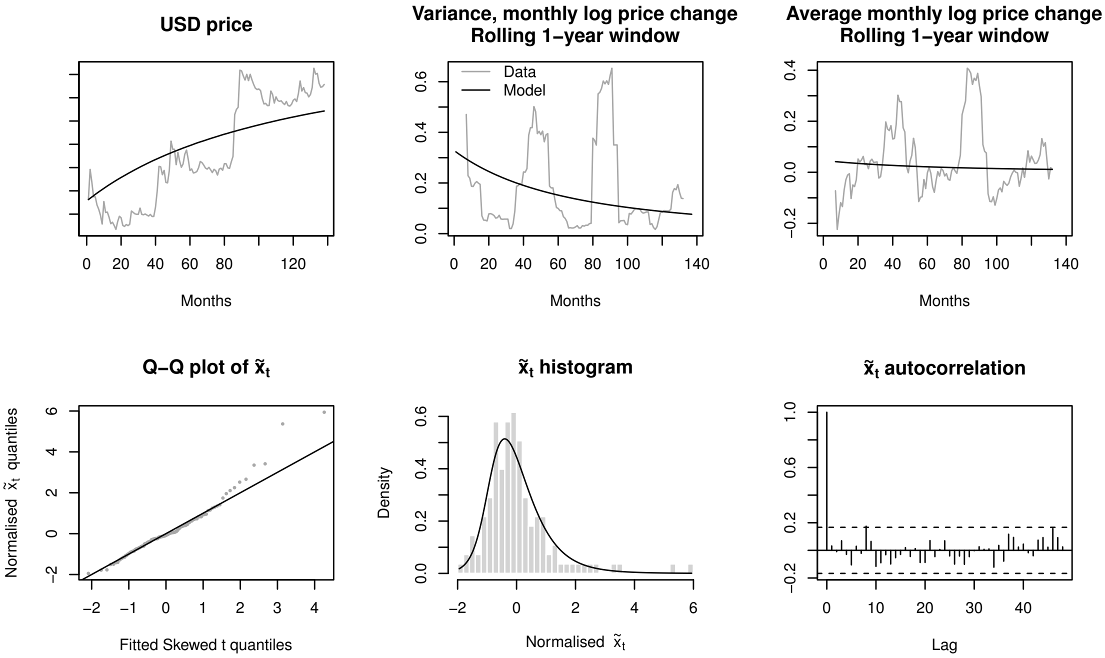

Figure 3.16: Tron (TRX) model fit

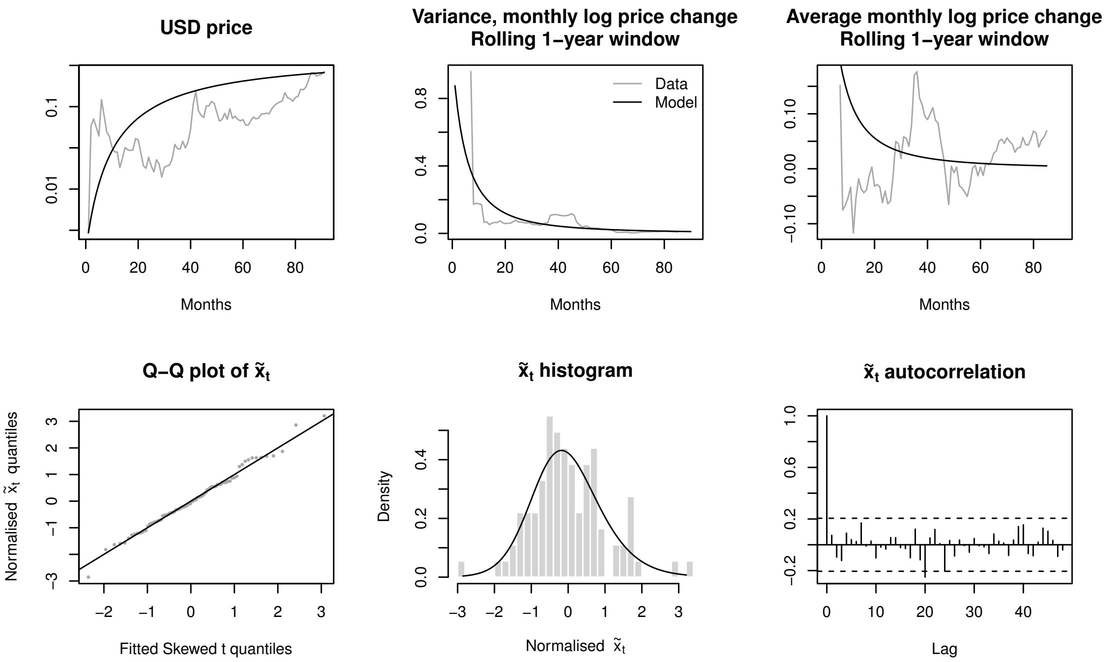

Figure 3.17: Cardano (ADA) model fit

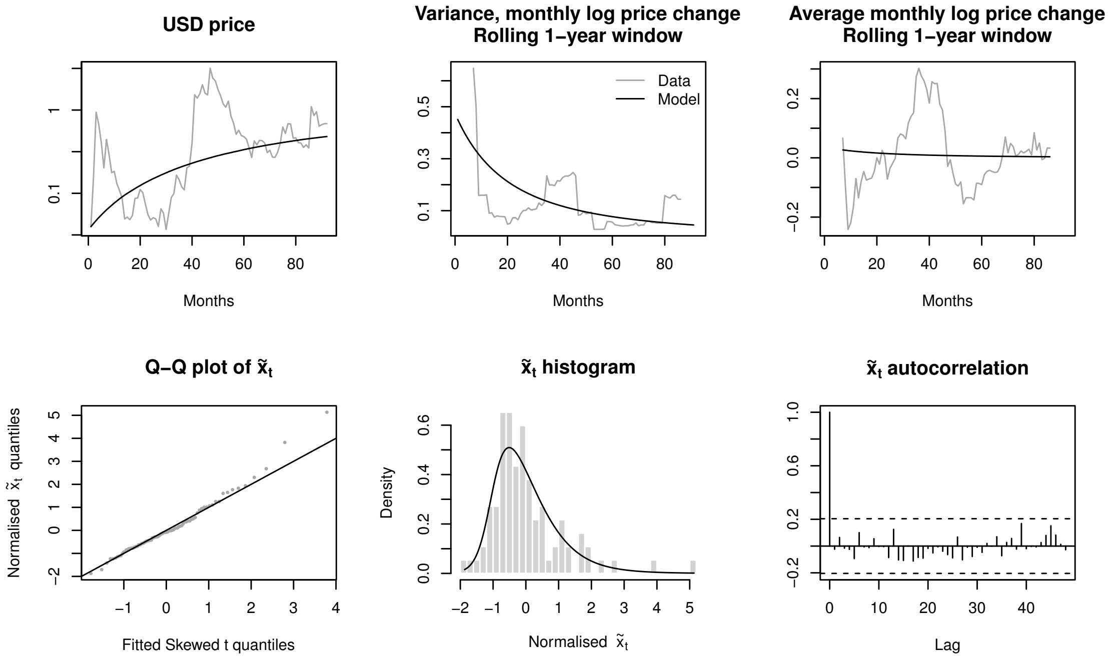

Figure 3.18: Polkadot (DOT) model fit

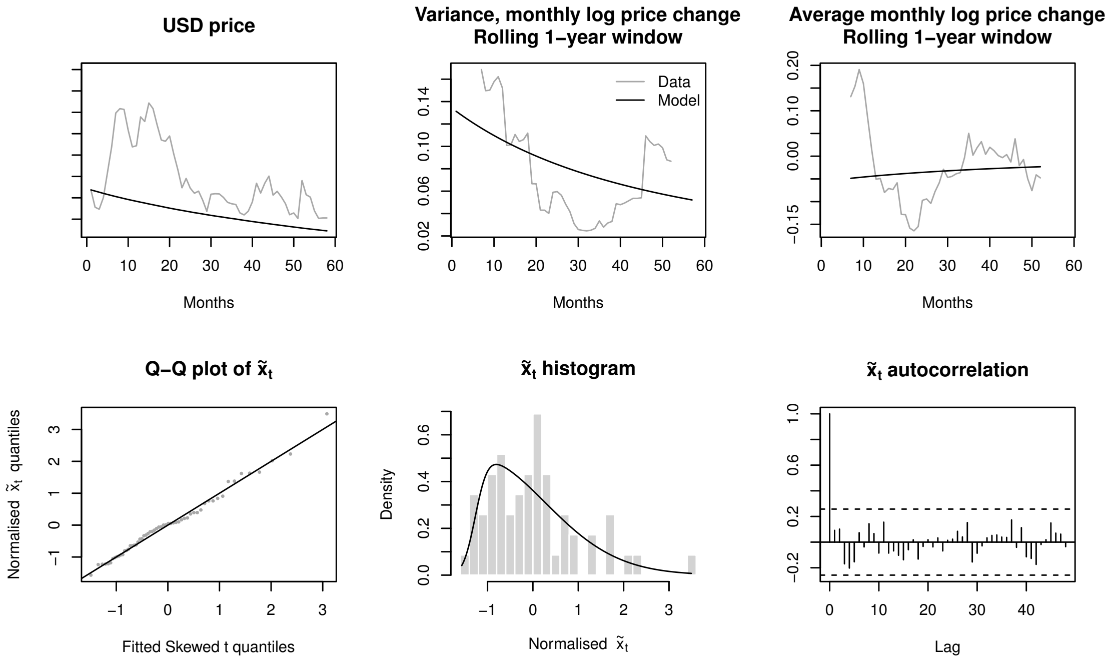

Figure 3.19: Litecoin (LTC) model fit

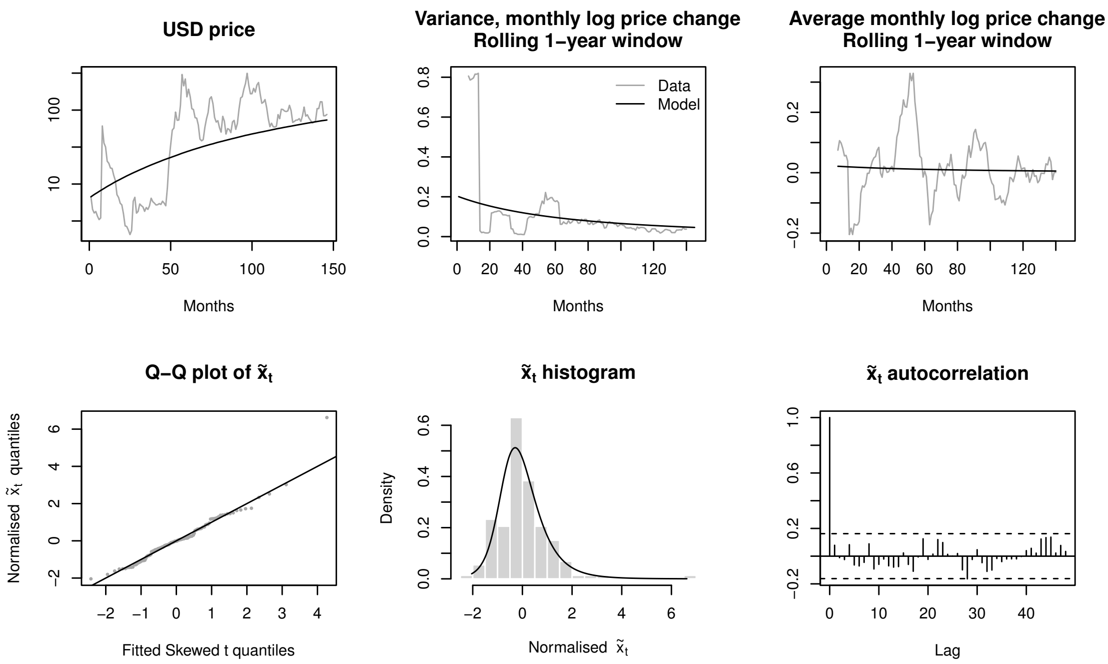

Figure 3.20: Stellar (XLM) model fit

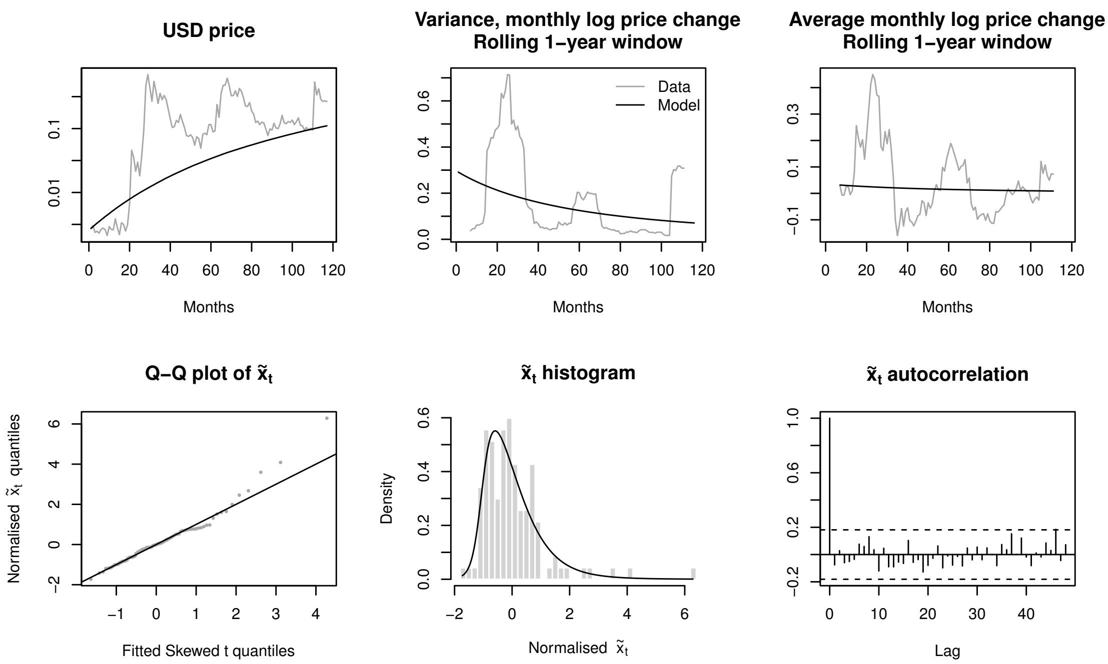

Figure 3.21: EOS (EOS) model fit

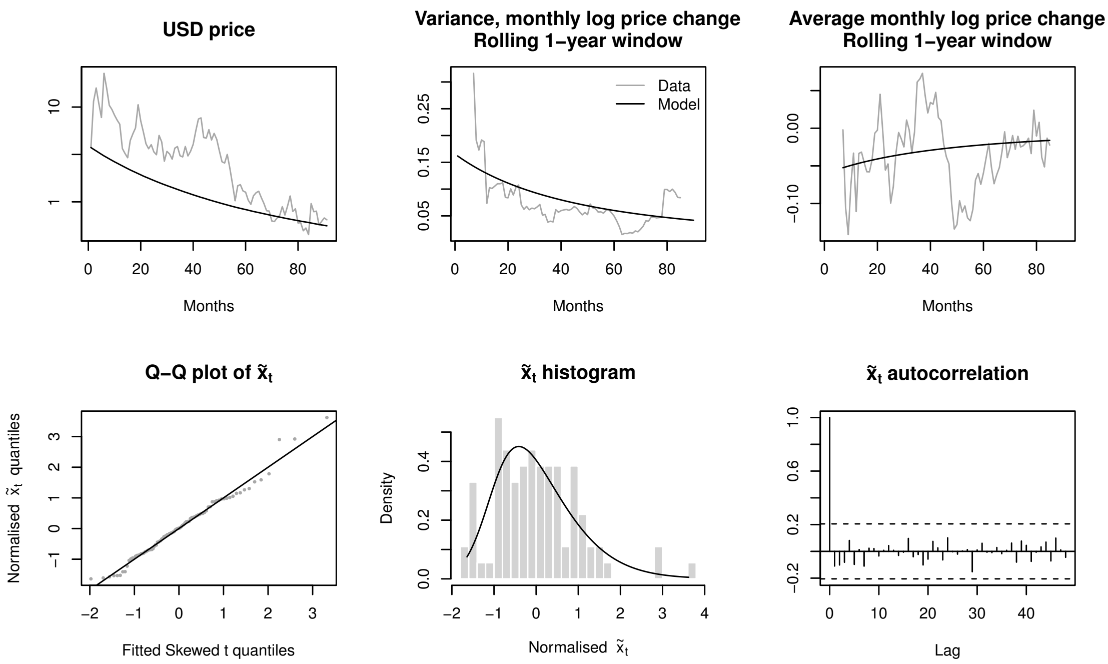

Figure 3.22: MaidSafeCoin (MAID) model fit

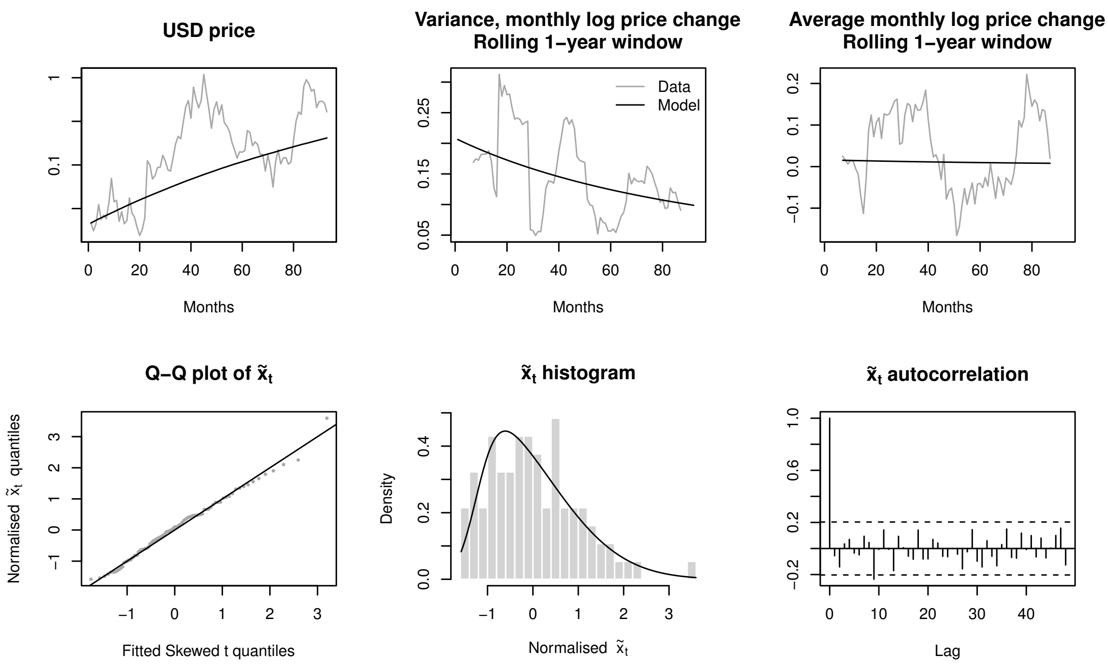

3.F Projections for all 15 cryptocurrencies

In the table in this section, current prices are as at 31 May 2025. Projection periods start as of the following day. Note that expectations and standard deviations of price changes can be elevated for cryptocurrencies with high skewness and kurtosis (see Table 3.2).

Table 3.5: Price and variance projections

			Projection horizon				
		Current price	1-year	5-year	10-year	20-year	30-year
BTC	Median price	104,802.6	140,450.8	390,266.4	1,024,086	3,959,020	10,406,778
	Standard deviation of log price change		53%	110%	142%	175%	193%
	Single-year expected price change		57.4%	39.7%	27.7%	16.8%	11.9%
	Single-year standard deviation		96.9%	70.0%	52.2%	35.1%	26.7%
	Single-year Sharpe ratio		0.59	0.57	0.53	0.48	0.45
ETH	Median price	2,533.06	3,248.17	6,899.26	13,796.88	33,436.26	66,047.69
	Standard deviation of log price change		61%	122%	154%	183%	197%
	Single-year expected price change		55.2%	34.0%	20.6%	11.6%	8.9%
	Single-year standard deviation		107.8%	69.2%	46.5%	30.8%	22.9%
	Single-year Sharpe ratio		0.51	0.49	0.44	0.37	0.39
BNB	Median price	659.54	744.45	1,171.75	1,716.92	2,929.52	4,627.95
	Standard deviation of log price change		45%	85%	104%	119%	126%
	Single-year expected price change		28.5%	16.0%	9.6%	6.4%	4.9%
	Single-year standard deviation		71.7%	40.5%	27.0%	16.9%	12.2%
	Single-year Sharpe ratio		0.40	0.39	0.35	0.38	0.40
SOL	Median price	156.89	219.56	503.24	914.87	1,936.95	3,228.75
	Standard deviation of log price change		80%	147%	176%	198%	208%
	Single-year expected price change		89.4%	34.9%	18.6%	9.2%	6.6%
	Single-year standard deviation		171.3%	76.7%	45.3%	26.1%	18.5%
	Single-year Sharpe ratio		0.52	0.45	0.41	0.35	0.36
XRP	Median price	2.18	2.32	3.44	4.99	8.58	14.24
	Standard deviation of log price change		81%	164%	210%	255%	277%
	Single-year expected price change		59.2%	39.5%	25.4%	14.4%	10.5%
	Single-year standard deviation		200.3%	128.9%	87.1%	52.9%	38.5%
	Single-year Sharpe ratio		0.30	0.31	0.29	0.27	0.27

			Projection horizon				
		Current price	1-year	5-year	10-year	20-year	30-year
DOGE	Median price	0.19	0.21	0.36	0.59	1.37	2.60
	Standard deviation of log price change		94%	193%	251%	310%	341%
	Single-year expected price change		96.6%	64.6%	37.9%	23.3%	16.3%
	Single-year standard deviation		432.9%	219.6%	120.6%	77.7%	56.6%
	Single-year Sharpe ratio		0.22	0.29	0.31	0.30	0.29
TRX	Median price	0.27	0.29	0.39	0.50	0.77	1.16
	Standard deviation of log price change		34%	63%	77%	87%	92%
	Single-year expected price change		16.4%	9.3%	6.5%	4.9%	4.3%
	Single-year standard deviation		41.8%	26.3%	18.2%	11.4%	8.4%
	Single-year Sharpe ratio		0.39	0.35	0.36	0.43	0.51
ADA	Median price	0.69	0.73	1.01	1.39	2.19	3.26
	Standard deviation of log price change		70%	136%	167%	195%	207%
	Single-year expected price change		46.6%	24.4%	14.9%	9.0%	6.0%
	Single-year standard deviation		154.5%	85.7%	51.6%	31.1%	22.0%
	Single-year Sharpe ratio		0.30	0.29	0.29	0.29	0.27
BCH	Median price	415.16	392.20	387.18	398.62	499.00	655.89
	Standard deviation of log price change		69%	136%	170%	201%	216%
	Single-year expected price change		25.0%	16.5%	11.1%	7.1%	5.9%
	Single-year standard deviation		110.5%	71.3%	55.6%	32.8%	24.3%
	Single-year Sharpe ratio		0.23	0.23	0.20	0.22	0.24
DOT	Median price	4.09	3.56	2.83	2.56	2.64	3.16
	Standard deviation of log price change		76%	149%	186%	220%	235%
	Single-year expected price change		24.9%	15.6%	9.0%	6.7%	5.1%
	Single-year standard deviation		126.9%	82.9%	54.9%	35.0%	26.1%
	Single-year Sharpe ratio		0.20	0.19	0.16	0.19	0.20

			Projection horizon				
		Current price	1-year	5-year	10-year	20-year	30-year
LTC	Median price	87.26	94.36	145.32	218.03	420.18	726.62
	Standard deviation of log price change		73%	151%	196%	243%	268%
	Single-year expected price change		54.2%	36.2%	25.5%	15.5%	11.9%
	Single-year standard deviation		339.4%	115.3%	81.9%	53.5%	48.5%
	Single-year Sharpe ratio		0.16	0.31	0.31	0.29	0.25
XLM	Median price	0.26	0.29	0.52	0.90	1.97	3.68
	Standard deviation of log price change		90%	183%	235%	287%	314%
	Single-year expected price change		98.0%	58.3%	36.7%	20.0%	14.0%
	Single-year standard deviation		503.7%	227.2%	233.4%	69.5%	48.5%
	Single-year Sharpe ratio		0.19	0.26	0.16	0.29	0.29
EOS	Median price	0.65	0.57	0.45	0.38	0.38	0.44
	Standard deviation of log price change		68%	137%	173%	207%	224%
	Single-year expected price change		16.2%	10.3%	8.5%	5.7%	4.8%
	Single-year standard deviation		96.2%	70.2%	52.5%	35.5%	26.4%
	Single-year Sharpe ratio		0.17	0.15	0.16	0.16	0.18
DASH	Median price	21.93	24.17	38.09	57.28	107.43	179.96
	Standard deviation of log price change		70%	143%	183%	222%	242%
	Single-year expected price change		47.9%	30.7%	21.2%	12.6%	8.9%
	Single-year standard deviation		136.2%	87.5%	64.9%	42.3%	31.2%
	Single-year Sharpe ratio		0.35	0.35	0.33	0.30	0.28
MAID	Median price	0.41	0.48	1.00	1.84	4.99	10.60
	Standard deviation of log price change		107%	222%	290%	361%	400%
	Single-year expected price change		133.1%	82.2%	56.5%	33.4%	21.8%
	Single-year standard deviation		492.5%	239.4%	155.0%	96.8%	69.2%
	Single-year Sharpe ratio		0.27	0.34	0.36	0.34	0.31

References

Agénor, Pierre-Richard, Timothy Jackson, and Pengfei Jia (2021a). "Macroprudential Policy Coordination in a Currency Union". *European Economic Review* 137, 103791.

Agénor, Pierre-Richard and Luiz A. Pereira da Silva (2022). "Financial spillovers, spillbacks, and the scope for international macroprudential policy coordination". *International Economics and Economic Policy* 19, pp. 79–127.

Agénor, Pierre-Richard, Timothy Jackson, Enisse Kharroubi, Leonardo Gambacorta, Giovanni Lombardo, and Luiz A. Pereira Da Silva (2021b). "Assessing the gains from international macroprudential policy cooperation". *Journal of Money, Credit and Banking* 53.7, pp. 1819–1866.

Arena, Marco, Tingyun Chen, Seung Mo Choi, Nan Geng, Cheikh Anta Gueye, Tonny Lybek, Evan Papageorgiou, and Yuanyan Sophia Zhang (2020). "Macroprudential policies and house prices in europe". *IMF Departmental Paper Series No. 20/03.*

Athey, Susan, Ivo Parashkevov, Vishnu Sarukkai, and Jing Xia (2016). "Bitcoin pricing, adoption, and usage: Theory and evidence". *Working paper.*

Bambrough, Billy (2025). "BlackRock CEO Reveals Sovereign Wealth Fund Bitcoin Talks— Predicts $14 Trillion As Dollar Collapse Fears Boost Price". *Forbes.*

Bengui, Julien (2014). "Macro-prudential policy coordination". *Working paper.*

Benigno, Gianluca and Pierpaolo Benigno (2006). "Designing targeting rules for international monetary policy cooperation". *Journal of Monetary Economics* 53.3, pp. 473–506.

Benigno, Pierpaolo, Gauti B Eggertsson, and Federica Romei (2020). "Dynamic debt deleveraging and optimal monetary policy". *American Economic Journal: Macroeconomics* 12.2, pp. 310–350.

Bernanke, Ben S (2005). "Remarks by Governor Ben S. Bernanke at the Sandridge Lecture, Virginia Association of Economics". *Richmond, Virginia, March 10, 2005.*

— (2010). "Monetary policy and the housing bubble: a speech at the Annual Meeting of the American Economic Association". *Atlanta, Georgia, January 3, 2010.*

Biais, Bruno, Christophe Bisière, Matthieu Bouvard, Catherine Casamatta, and Albert J. Menkveld (2023). "Equilibrium Bitcoin Pricing". *The Journal of Finance* 78.2, pp. 967–1014.

Bianchi, Javier (2011). "Overborrowing and systemic externalities in the business cycle". *American Economic Review* 101.7, pp. 3400–3426.

Bianchi, Javier, Saki Bigio, and Charles Engel (2023). "Scrambling for Dollars: International Liquidity, Banks and Exchange Rates". *NBER working paper no. w29457.*

Bodenstein, Martin, Luca Guerrieri, and Joe LaBriola (2019). "Macroeconomic policy games". *Journal of Monetary Economics* 101, pp. 64–81.

Bolt, Wilko and Maarten R.C. van Oordt (2020). "On the value of virtual currencies". *Journal of Money, Credit and Banking* 52.4, pp. 835–862.

Caballero, Ricardo J., Emmanuel Farhi, and Pierre-Olivier Gourinchas (2008). “An Equilibrium Model of "Global Imbalances" and Low Interest Rates”. *American Economic Review* 98.1, pp. 358–393.
Chen, William and Gregory Phelan (2021). “International coordination of macroprudential policies with capital flows and financial asymmetries”. *Journal of Financial Stability* 56, 100929.
Cieslak, Anna, Wenhao Li, and Carolin E Pflueger (2025). “Inflation and treasury convenience”. *Working paper.*
Cong, Lin William, Ye Li, and Neng Wang (2021). “Tokenomics: Dynamic Adoption and Valuation”. *The Review of Financial Studies* 34.3, pp. 1105–1155.
Dávila, Eduardo and Anton Korinek (2018). “Pecuniary Externalities in Economies with Financial Frictions”. *The Review of Economic Studies* 85.1, pp. 352–395.
Del Negro, Marco, Gauti Eggertsson, Andrea Ferrero, and Nobuhiro Kiyotaki (2017). “The great escape? A quantitative evaluation of the Fed’s liquidity facilities”. *American Economic Review* 107.3, pp. 824–857.
Denis, David J. and Valeriy Sibilkov (2010). “Financial Constraints, Investment, and the Value of Cash Holdings”. *The Review of Financial Studies* 23.1, pp. 247–269.
Detemple, Jérôme B (1986). “Asset pricing in a production economy with incomplete information”. *The Journal of Finance* 41.2, pp. 383–391.
Devereux, Michael B, Charles Engel, and Steve Pak Yeung Wu (2023). “Collateral Advantage: Exchange Rates, Capital Flows and Global Cycles”. *NBER working paper no. w31164.*
Doms, Mark and Timothy Dunne (1998). “Capital Adjustment Patterns in Manufacturing Plants”. *Review of Economic Dynamics* 1.2, pp. 409–429.
Engel, Charles (2016). “Exchange Rates, Interest Rates, and the Risk Premium”. *American Economic Review* 106.2, pp. 436–474.
Engel, Charles and Steve Pak Yeung Wu (2023). “Liquidity and exchange rates: An empirical investigation”. *The Review of Economic Studies* 90.5, pp. 2395–2438.
— (2024). “Exchange Rate Models are Better than You Think, and Why They Didn’t Work in the Old Days”. *NBER working paper no. w32808.*
Evans, George W and Seppo Honkapohja (2001). *Learning and expectations in macroeconomics.* Princeton University Press.
Fama, Eugene F (1984). “Forward and spot exchange rates”. *Journal of Monetary Economics* 14.3, pp. 319–338.
Ferrero, Andrea (2015). “House Price Booms, Current Account Deficits, and Low Interest Rates”. *Journal of Money, Credit and Banking* 47.S1, pp. 261–293.
Fornaro, Luca and Federica Romei (2019). “The paradox of global thrift”. *American Economic Review* 109.11, pp. 3745–3779.
Frisch, Ragnar and Frederick V. Waugh (1933). “Partial Time Regressions as Compared with Individual Trends”. *Econometrica* 1.4, pp. 387–401.
Gabaix, Xavier and Matteo Maggiori (2015). “International Liquidity and Exchange Rate Dynamics”. *The Quarterly Journal of Economics* 130.3, pp. 1369–1420.
Galati, Gabriele and Richhild Moessner (2018). “What Do We Know About the Effects of Macroprudential Policy?” *Economica* 85.340, pp. 735–770.
Galí, Jordi and Tommaso Monacelli (2005). “Monetary Policy and Exchange Rate Volatility in a Small Open Economy”. *The Review of Economic Studies* 72.3, pp. 707–734.

Gopinath, Gita and Jeremy C Stein (2021). "Banking, Trade, and the Making of a Dominant Currency". *The Quarterly Journal of Economics* 136.2, pp. 783–830.

Guerrieri, Luca and Matteo Iacoviello (2015). "OccBin: A toolkit for solving dynamic models with occasionally binding constraints easily". *Journal of Monetary Economics* 70, pp. 22–38.

— (2017). "Collateral constraints and macroeconomic asymmetries". *Journal of Monetary Economics* 90, pp. 28–49.

Guerrieri, Veronica and Guido Lorenzoni (2017). "Credit Crises, Precautionary Savings, and the Liquidity Trap". *The Quarterly Journal of Economics* 132.3, pp. 1427–1467.

Han, Brandon Yueyang and Igor Makarov (2021). "Feedback trading and bubbles". *Working paper.*

Hott, Christian and Terhi Jokipii (2012). "Housing bubbles and interest rates". *Schweizerische Nationalbank working paper.*

Iacoviello, Matteo and Stefano Neri (2010). "Housing Market Spillovers: Evidence from an Estimated DSGE Model". *American Economic Journal: Macroeconomics* 2.2, pp. 125–164.

Intungane, Doriane (2023). "The impact of macroprudential policies on the transmission of shocks across financially integrated countries". *Review of International Economics* 31.1, pp. 249–273.

Itskhoki, Oleg and Dmitry Mukhin (2021). "Exchange rate disconnect in general equilibrium". *Journal of Political Economy* 129.8, pp. 2183–2232.

Jeanne, Olivier (2014). "Macroprudential Policies in a Global Perspective". *NBER working paper no. w19967.*

Jiang, Zhengyang, Arvind Krishnamurthy, and Hanno Lustig (2021). "Foreign Safe Asset Demand and the Dollar Exchange Rate". *The Journal of Finance* 76.3, pp. 1049–1089.

— (2024a). "Dollar safety and the global financial cycle". *Review of Economic Studies* 91.5, pp. 2878–2915.

Jiang, Zhengyang, Arvind Krishnamurthy, Hanno Lustig, and Jialu Sun (2024b). "Convenience Yields and Exchange Rate Puzzles". *NBER working paper no. w32092.*

Jiménez, Gabriel, Steven Ongena, José-Luis Peydró, and Jesús Saurina (2014). "Hazardous Times for Monetary Policy: What Do Twenty-Three Million Bank Loans Say About the Effects of Monetary Policy on Credit Risk-Taking?" *Econometrica* 82.2, pp. 463–505.

Jones, Callum, Virgiliu Midrigan, and Thomas Philippon (2022). "Household Leverage and the Recession". *Econometrica* 90.5, pp. 2471–2505.

Kaiko Research (2024). "Market shrugs off longshot XRP filing". *https://research.kaiko.com/ insights/market-shrugs-off-longshot-xrp-filing.*

Karau, Sören and Emanuel Moench (2023). "Is It Really Mine? Mining Shocks and their Effects on Bitcoin Valuations". *Working paper.*

Kekre, Rohan and Moritz Lenel (2024). "The Flight to Safety and International Risk Sharing". *American Economic Review* 114.6, pp. 1650–1691.

Kiyotaki, Nobuhiro and John Moore (1997). "Credit Cycles". *Journal of Political Economy* 105.2, pp. 211–248.

— (2019). "Liquidity, Business Cycles, and Monetary Policy". *Journal of Political Economy* 127.6, pp. 2926–2966.

Korinek, Anton (2016). "Currency wars or efficient spillovers? A general theory of international policy cooperation". *NBER working paper no. w23004.*

Krishnamurthy, Arvind and Annette Vissing-Jorgensen (2012). “The Aggregate Demand for Treasury Debt”. *Journal of Political Economy* 120.2, pp. 233–267.

Lane, Philip R. and Gian Maria Milesi-Ferretti (2017). “International Financial Integration in the Aftermath of the Global Financial Crisis”. *IMF Working Papers WP/17/115*.

Lido (2024). “Research: stETH-based Swaps using ERC-6123”. *https://research.lido.fi/t/research-steth-based-swaps-using-erc-6123/8825*.

Lin, Min-Bin, Ștefan Găman, Ruting Wang, and Daniel Traian Pele (2025). “Market Responses to Ethereum Development Milestones: An Event Study Approach”. *Working paper*.

Lorenzoni, Guido (2008). “Inefficient credit booms”. *The Review of Economic Studies* 75.3, pp. 809–833.

Lovell, Michael C. (1963). “Seasonal Adjustment of Economic Time Series and Multiple Regression Analysis”. *Journal of the American Statistical Association* 58.304, pp. 993–1010.

Maddaloni, Angela and José-Luis Peydró (2011). “Bank Risk-taking, Securitization, Supervision, and Low Interest Rates: Evidence from the Euro-area and the U.S. Lending Standards”. *The Review of Financial Studies* 24.6, pp. 2121–2165.

Meese, Richard A and Kenneth Rogoff (1983). “Empirical exchange rate models of the seventies: Do they fit out of sample?” *Journal of International Economics* 14.1-2, pp. 3–24.

Pagnotta, Emiliano S (2022). “Decentralizing Money: Bitcoin Prices and Blockchain Security”. *The Review of Financial Studies* 35.2, pp. 866–907.

Pastor, Lubos and Pietro Veronesi (2009). “Learning in Financial Markets”. *Annual Review of Financial Economics* 1, pp. 361–381.

Prat, Julien, Vincent Danos, and Stefania Marcassa (2025). “Fundamental pricing of utility tokens”. *Management Science* Articles in advance, pp. 1–16.

Rasmussen, Ryan, Anais Rachel, and Gayatri Choudhury (2022). “Staking as a Service: An Emerging Crypto Sector Poised for Growth”. *Bitwise Timely Insights*.

Rubio, Margarita (2014). “Macroprudential policy implementation in a heterogeneous monetary union”. *XXI Encuentro Economía Pública*, p. 14.

Sarr, Abdourahmane and Tonny Lybek (2002). “Measuring Liquidity in Financial Markets”. *IMF working papers WP/02/232*.

Schmeling, Maik, Andreas Schrimpf, and Karamfil Todorov (2023). “Crypto carry”. *BIS Working Papers No 1087*.

Schmitt-Grohé, Stephanie and Martín Uribe (2003). “Closing small open economy models”. *Journal of International Economics* 61.1, pp. 163–185.

Stein, Jeremy C. (2012). “Monetary Policy as Financial Stability Regulation”. *The Quarterly Journal of Economics* 127.1, pp. 57–95.

Steinmetz, Fred (2021). “Behavioural clusters of cryptocurrency users: Frequencies of non-speculative application domains”. *BRL Working Paper Series No. 24*.

US Federal Reserve (2024). “Economic Well-Being of U.S. Households in 2023”. *Board of Governors of the Federal Reserve Board*.

Valchev, Rosen (2020). “Bond Convenience Yields and Exchange Rate Dynamics”. *American Economic Journal: Macroeconomics* 12.2, pp. 124–166.

VanEck (2025). “Bitcoin vs. Ethereum in 2024: Comparison & Outlook”. *https://www.vaneck. com/us/en/blogs/digital-assets/bitcoin-vs-ethereum/*.

Van Oordt, Maarten R.C. (2025). “On bubbles in cryptocurrency prices”. *Working paper*.

Wei, Yanhao (Max) and Anthony Dukes (2021). "Cryptocurrency Adoption with Speculative Price Bubbles". *Marketing Science* 40.2, pp. 241–260.