Rational Sentiments and Financial Frictions*

Paymon Khorrami
Duke Fuqua

Fernando Mendo
PUC-Rio

March 19, 2025

Abstract

We discover sentiment-driven equilibria in popular macroeconomic models of imperfect risk sharing. In these equilibria, sentiment dynamics behave like uncertainty shocks, in the sense that self-fulfilled beliefs about volatility drive aggregate fluctuations. Because such fluctuations can decouple from the wealth distribution, rational sentiment helps resolve two puzzles plaguing models emphasizing balance sheets: (i) financial crises emerge suddenly, featuring large volatility spikes and asset-price declines; (ii) asset-price booms, with below-average risk premia, predict busts and financial crises. Quantitatively, our sentiment equilibria are able to replicate empirical crisis dynamics for output, volatility, and risk premia, whereas the fundamental equilibrium performs poorly on these dimensions. Methodologically, our contribution is using stochastic stability theory to establish existence of sunspot equilibria.

JEL Codes: E00, E44, G01.

Keywords: financial frictions, sunspot equilibria, self-fulfilling beliefs, sentiment, financial crises, uncertainty shocks.

*Khorrami: paymon.khorrami@duke.edu; Mendo: fmendolopez@gmail.com. We are grateful to Jungsuk Han, Stavros Panageas, Alan Moreira, Zhengyang Jiang (discussants) and to seminar participants at Imperial College, the Central Bank of Chile, USC, Yale SOM, Duke Fuqua, Duke Econ, HKU, NUS, UT Dallas, the CESifo Macro/Money/International conference, the Macro-Finance Society, the EFA, LACEA-LAMES meetings, and in the virtual Junior European Finance series for helpful feedback. A previous version of this paper was titled "Where there is amplification, there are sunspots".

It has by now become commonplace, especially after the 2008 global financial crisis, for macroeconomic models to prominently feature banks, limited participation, imperfect risk-sharing, and other such "financial frictions." Incorporating these features allows macroeconomists to speak meaningfully about financial crises and desirable policy responses. Despite the dramatic growth in this literature, there remain two major sources of disconnect between these models and actual crisis data. For one, standard models have difficulty reproducing the observed severity and suddenness of economic downturns and asset-price dislocations. Secondly, standard models struggle to generate booms that are inherently fragile and prone to bust. To address these shortcomings, some recent contributions add large and sudden bank runs,[1] while others deviate from rational expectations to model booms as episodes of over-optimism.[2]

Figures 1-2 summarize the financial crisis dynamics that motivate our study. Equity market values fall dramatically and suddenly, especially for banks, while credit spreads rise rapidly and substantially. On these dimensions, the US crises displayed in Figure 2 were more severe than the average crisis from the international sample of Figure 1. Those US crises also featured a tremendous rise in asset market volatility, of particular interest for our analysis. In 2008-09, for example, stock market volatility rose from 15% to above 70%, which is far beyond the magnitudes possible in conventional macroeconomic models. In addition to their severity and suddenness, the figures are suggestive of pre-crisis "froth" in the sense that asset prices are high and rising prior to their crash, while measures of risk and risk premia are low and sometimes falling prior to their spike. Related evidence from the literature highlights the predictability of financial crises by large credit and asset price booms that feature below-average credit spreads.[3]

To address these patterns, we embrace *rational sentiment* as an approach complementary to the existing literature. This paper makes two main contributions. First, we uncover a wide variety of novel sentiment-driven sunspot equilibria supported by standard financial friction models. The fluctuations in these equilibria can be self-fulfilling: they only occur because agents coordinate on them. Second, we demonstrate how sentiment fluctuations alleviate some of the shortcomings for this class of models. Rational senti-

[1]For example, Gertler and Kiyotaki (2015) and Gertler et al. (2020) attempt to integrate bank runs into a conventional financial accelerator model, in order to capture additional amplification and non-linearity. Without runs or panic-like behavior, financial accelerator models have a difficult time inducing the financial intermediary leverage needed to generate large amounts of amplification.

[2]For example, Krishnamurthy and Li (2024) and Maxted (2024) build an extrapolative sentiment process on top of a relatively standard financial accelerator model. Agents' excessive optimism in booms lowers risk premia, erodes bank balance sheets, and creates fragility.

[3]See Reinhart and Rogoff (2009); Jordà et al. (2011, 2013, 2015a,b); Mian et al. (2017); López-Salido et al. (2017); Baron and Xiong (2017); Krishnamurthy and Muir (2024).

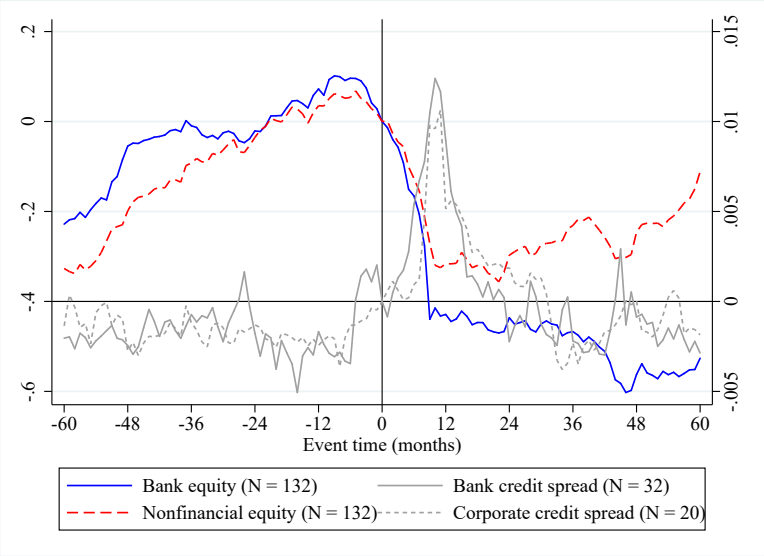


Figure 1: Patterns across banking crises, as defined in Baron et al. (2021), across 46 countries from 1870-2016. This graph reproduces Figure 6A of Baron et al. (2021), showing average equity values (left axis) and credit spreads (right axis) for both banks and non-financial corporations in the 5 years before and after a crisis begins. All variables are normalized to zero in January of the crisis year. The text "N =" in the caption refers to the number of crises for which a particular data series was available.

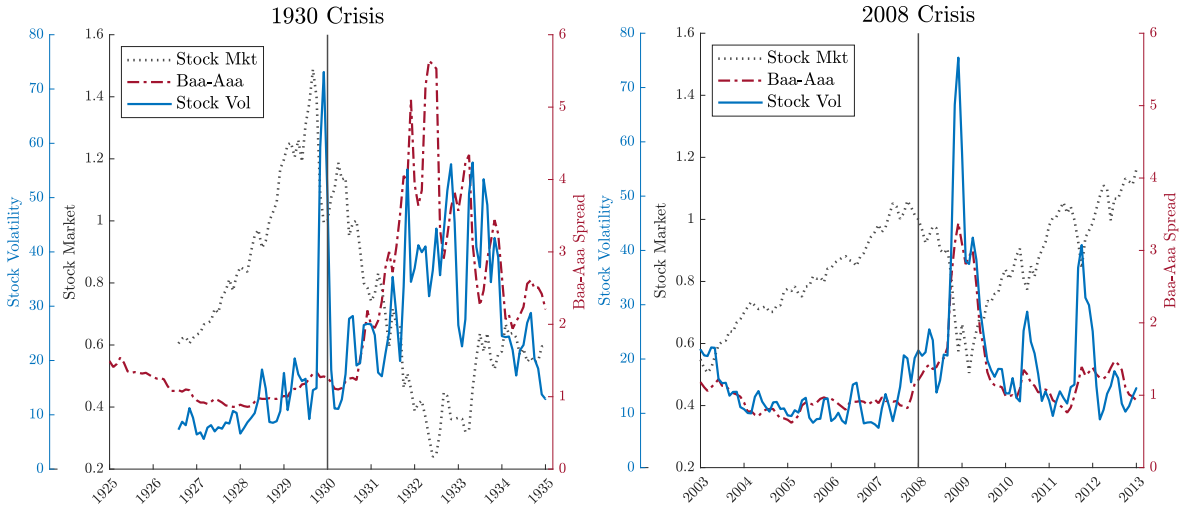


Figure 2: US asset markets surrounding the Great Depression (1930) and Great Financial Crisis (2008). The "Stock Mkt" is the cumulative return on the US value-weighted stock market, normalized to 1 in January of the crisis onset; "Stock Vol" measures monthly realized return volatility, computing using daily returns within the past two months and then annualized; and "Baa-Aaa" refers to the Moody's Baa minus Aaa credit spread index in percentage points.

ment can generate both (i) large and sudden fluctuations, similar to bank runs (footnote 1), and (ii) booms that breed fragility, similar to the "behavioral sentiment" adopted by some recent papers (footnote 2). The model without rational sentiment fails on these dimensions.

Model and mechanism. We study a simple stripped-down model with financial frictions, similar to Kiyotaki and Moore (1997), Brunnermeier and Sannikov (2014), and many others.[4] There are two types of agents ("experts" and "households") with identical preferences but different levels of productivity when managing capital. Heterogeneous productivity means the identity of capital holders matters for aggregate output. Ideally, in a world with complete financial markets, experts would manage all capital and issue sufficient equity to share risks with households. But in our model, incomplete markets prevent agents from sharing those risks, so optimal capital holdings depend to some degree on risk and not only on productivities. There are no other features: no ad-hoc collateral constraints, no default externalities, and no irrational beliefs. And yet, this basic model features a tremendous amount of multiplicity that has been overlooked.

Indeterminacy in this model comes from the combination of incomplete financial markets and heterogeneous productivities. With these features, asset prices today are not pinned down by "fundamentals"— namely the minimal set of state variables—and can also depend on agents' beliefs about the distribution of asset prices tomorrow. Different beliefs deliver different equilibria. Of particular importance in our specific model is the perceived dispersion in future asset prices, or price volatility.

The following story clarifies the mechanics. Suppose agents are *fearful*, anticipating high asset-price volatility. Despite their productivity advantage, experts will only manage a fraction of aggregate capital, as capital price risk cannot be fully shared through markets. Perceived volatility thus causes an inefficient capital allocation, hence low asset prices. On the other hand, if low asset-price volatility is anticipated, experts will hold a large share of capital, and asset prices will be high. Are both of these coordinated volatility perceptions justified? In many models, only one perception of volatility could be consistent with equilibrium, because future paths would otherwise be explosive.

But in our paper, many coordinated beliefs about volatility can satisfy equilibrium conditions and remain non-explosive, mirroring the conventional idea that dynamic stability of equilibrium supports indeterminacy. Here, stability means that asset prices eventually mean-revert, or "bounce back" from extreme values. Supposing the future distribution of asset prices q is characterized by a first and second moment (μ_q, σ_q^2), then a rise in σ_q (fear)—which depresses q—must be accompanied by an eventual rise in μ_q (bounce-back beliefs). In our continuous-time setup, bounce-back beliefs are just boundary conditions on μ_q at extreme states. Such boundary restrictions are both analytically-

[4]We work in continuous time, contributing to a burgeoning literature (He and Krishnamurthy, 2012, 2013, 2019; Moreira and Savov, 2017; Di Tella, 2017, 2019; Klimenko et al., 2017; Drechsler et al., 2018; Caballero and Simsek, 2020; d'Avernas and Vandeweyer, 2023; Silva, 2024). For surveys, see Brunnermeier and Sannikov (2016) and Hansen et al. (2024).

convenient and mild; rich dynamics are admissible away from extreme states.

If volatility is dynamically stable, we can use sunspot shocks to govern agents' beliefs about volatility and create sentiment dynamics. A surprise increase in fear leads to a fire sale, which temporarily depresses asset prices and output. Conversely, a decline in fear raises asset prices, through coordinated purchases. These fear-driven dynamics are sustainable so long as they are expected to eventually subside. A distinctive feature is that sentiment dynamics are always characterized by time-varying *endogenous uncertainty*.

Overview of paper. While explaining our model above, we abstracted from the wealth distribution between experts and households. Typically in the financial frictions literature, this wealth distribution is the key state variable modulating the dynamics. In our analysis, the wealth distribution remains a state variable, but additional "sentiment" state variables naturally arise as potential drivers of equilibrium. Mathematically, we dispense with the assumption that equilibria be Markovian in the wealth distribution, which removes an ad-hoc restriction on agents' beliefs.[5]

Our main theoretical results provide an explicit construction and characterization of a broad class of sentiment-driven equilibria (Section 2). As one might expect from deterministic models, the existence of sunspot equilibria is tied to the stability properties of the equilibrium dynamical system. For many models, such stability questions are settled via linearized spectral analysis near steady state. What is the analog in our stochastic nonlinear environment? To tackle this problem, we leverage tools from the "stochastic stability" literature (analogous to Lyapunov stability for ODE systems). Conveniently, our stability analysis boils down to boundary conditions on our dynamical system.

Sentiment-driven equilibria engender several new insights, related to the shortcomings in existing models (Section 3). First, sentiment permits far higher volatility and risk premia spikes, helping address the suddenness and severity of financial crises. Second, pre-crisis froth can emerge with sentiment: asset-price booms and below-average risk premia predict financial crises. By contrast, the same model without sentiment cannot generate large increases in volatility or risk premia nor any pre-crisis froth. The key flaw with the conventional equilibrium is its dynamics are channeled through expert equity, which is slow-moving; sentiment critically decouples dynamics from expert equity. We summarize such results via counterparts to the crisis event studies in Figures 1-2.

How much better do sentiment equilibria perform on these dimensions? To answer

[5] In a companion paper (Khorrami and Mendo, 2024), we study all equilibria which are Markovian in the wealth distribution. There, we show that any resulting dynamics are approximately identical to the conventional dynamics studied by Brunnermeier and Sannikov (2014) and others. Thus, resolving the literature's puzzles requires us to go beyond the wealth distribution and explore sentiment equilibria.

this, we calibrate a sentiment equilibrium to several "non-crisis moments" and then evaluate its performance on "crisis moments" in the literature. In a model with multiplicity, a natural question is which sentiment dynamics shall be chosen. We take two routes. For some indeterminate objects, we tie our hands by taking them directly from the data, ensuring we do not reverse-engineer sentiment to match crisis dynamics. For others, we perform extensive sensitivity analysis, illustrating that equilibrium dynamics are similar across a range of sentiment specifications. Taken together, our exercises deliver a "possibility result": a host of crisis dynamics can hold in a simple, rational framework, so long as one allows time-varying sentiment.

Related literature. At a high level, the theoretical focus on financial frictions and sunspots is not new to this paper. Several studies show how multiplicity emerges through the interaction between asset valuations and borrowing constraints.[6] Relative to these papers, we study different financial frictions (equity-issuance constraints) that do not feature any mechanical link between prices and constraints.

Our self-fulfilled fluctuations are closest to bank runs, financial panics, and sudden stops.[7] These phenomena similarly rely on financial frictions, are outcomes of coordination, and produce large fluctuations relative to fundamentals. However, whereas bank runs and its cousins are liability-side phenomena, self-fulfilled fire sales are asset-side phenomena. Furthermore, our mechanism does not require asset-market illiquidity or maturity mismatch. Finally, whereas runs are exclusively about downside risk, our sentiment fluctuations also generate interesting boom-bust cycles.

Given the absence of ad-hoc borrowing constraints or runs, our paper illustrates that a much broader class of financial crisis models are subject to sunspots. We also do not rely on more traditional multiplicity-inducing assumptions, such as overlapping generations,[8] non-convex technology,[9] imperfect information,[10] or multiple assets.[11]

Our focus on fear and volatility as drivers of self-fulfilling fluctuations closely relates

[6]For instance, bubbles can relax credit constraints, allowing greater investment and thus justifying the existence of the bubble (Scheinkman and Weiss, 1986; Kocherlakota, 1992; Farhi and Tirole, 2012; Miao and Wang, 2018; Liu and Wang, 2014). Self-fulfilling credit dynamics can also arise with *unsecured* lending as opposed to collateralized (Gu et al., 2013; Azariadis et al., 2016).

[7]Mendo (2020) studies self-fulfilled panics that induce collapse of the financial sector. Gertler and Kiyotaki (2015) and Gertler et al. (2020) study bank runs in a similar class of models.

[8]The classic studies on OLG and multiplicity are Azariadis (1981) and Cass and Shell (1983). A more recent investigation, focusing on wealth redistribution across generations, is Farmer (2018).

[9]For example, see Azariadis and Drazen (1990) for multiplicity under threshold investment behavior. See Farmer and Benhabib (1994) for a multiplicity under increasing returns to scale.

[10]For example, see Benhabib and Wang (2015) and Benhabib et al. (2015, 2016, 2019).

[11]Hugonnier (2012), Gârleanu and Panageas (2021), and Khorrami and Zentefis (2025) all build "redistributive" sunspots that shift valuations among multiple positive-net-supply assets.

to the "self-fulfilling risk panics" of Bacchetta et al. (2012). Benhabib et al. (2020) obtain a similar type of fluctuation by examining economies with either collateral or liquidity constraints, rather than the OLG setup of Bacchetta et al. (2012). Although we do not rely on common multiplicity-inducing features like OLG or collateral constraints, we expound on the deeper connection to these papers in Section 1.3. Also relatedly, Khorrami and Mendo (2025) and Lee and Dordal i Carreras (2024) obtain self-fulfilling risk in New Keynesian models, which feature an endogenous mapping from asset prices to the real economy as in this paper, but replace financial frictions with nominal rigidities.

1 Model

Information structure. Time $t \geq 0$ is continuous. (We also study a discrete-time version of the model in Online Appendix G.) There are two types of uncertainty in the economy, modeled as two independent Brownian motions $Z := (Z^{(1)}, Z^{(2)})$. All random processes will be adapted to Z. As will be clear below, the first shock $Z^{(1)}$ represents a *fundamental shock* in the sense that it directly impacts production possibilities, whereas the second shock $Z^{(2)}$ is a *sunspot shock* that is extrinsic to any economic primitives but nevertheless may impact endogenous objects. Later, we will also consider extrinsic Poisson jumps as part of the information structure.

Technology and markets. There are two goods, a non-durable good (the numéraire, "consumption") and a durable good ("capital") that produces the consumption good. The aggregate supply of capital grows exogenously as

$$dK_t = K_t[gdt + \sigma dZ_t^{(1)}], \tag{1}$$

where g and $\sigma \geq 0$ are exogenous constants. We add endogenous capital investment in Section 2.3. The capital-quality shock $\sigma dZ^{(1)}$ is a standard way to introduce fundamental randomness in technology. Individual capital holdings evolve identically, except that capital may be traded frictionlessly between agents in the market.[12] The relative capital price is q_t and is determined in equilibrium.

[12]Individual capital is thus a choice variable: if an agent holds capital k_t, its law of motion is

$$dk_t = gk_t dt + \sigma k_t dZ_t^{(1)} + d\Omega_t,$$

where the term $d\Omega_t$ corresponds to net purchases. To be clear, both g and $\sigma dZ_t^{(1)}$ affect agents' return-on-capital, whereas the net purchases term $d\Omega_t$ does not.

There are two types of agents, experts and households, who differ in their production technologies. Experts produce a_e units of the consumption good per unit of capital, whereas households' productivity is $a_h \in (0, a_e)$.

Financial markets consist solely of an instantaneously-maturing, risk-free bond that pays interest rate r_t is in zero net supply. The key financial friction: agents cannot issue equity when managing capital. It is inconsequential that the constraint be this extreme. Partial equity issuance, as long as there is some limit, will generate similar results on sunspot volatility (we discuss this further in Section 1.3).

Preferences and optimization. Given the stated assumptions, we can write the dynamic budget constraint of an agent of type $\ell \in \{e, h\}$ (expert or household) as

$$dn_{\ell,t} = \Big[(n_{\ell,t} - q_t k_{\ell,t})r_t - c_{\ell,t} + a_\ell k_{\ell,t}\Big]dt + q_t k_{\ell,t} dR_t, \tag{2}$$

where n_ℓ is the agent's net worth, c_ℓ is consumption, and k_ℓ is capital holdings. The last term $dR_t := \frac{d(q_t K_t)}{q_t K_t}$ is the capital and price appreciation while holding capital.

Experts and households have time-separable logarithmic utility, with discount rates ρ_e and $\rho_h \leq \rho_e$, respectively. All agents have rational expectations and solve

$$\sup_{c_\ell \geq 0,\, k_\ell \geq 0,\, n_\ell \geq 0} \mathbb{E}\Big[\int_0^\infty e^{-\rho_\ell t} \log(c_{\ell,t})dt\Big] \tag{3}$$

subject to (2). Everything in optimization problem (3) is homogeneous in (c, k, n), so we can think of the expert and household as representative agents within their class.

Let us briefly discuss the solvency constraint $n_{\ell,t} \geq 0$ in (3). This constraint says that agents cannot borrow more than the market value of their capital, and since there are no other assets besides capital, one can think of $n_{\ell,t} \geq 0$ as the "natural borrowing limit." We analyze some microfoundations for this solvency in Appendix A, to provide more comfort that the solvency constraint is natural and minimal—i.e., to assure the reader that no ad-hoc borrowing constraint is driving our results.

Finally, to guarantee a stationary wealth distribution, we also allow a type-switching structure: experts retire and become households at rate δ_e, while households retire and become experts at rate δ_h. Technically, the presence of type-switching alters the objective function from (3), but under log utility optimal behavior is still as if solving (3)—we show this in Appendix B.1. To acknowledge the fact that type-switching shifts wealth across agent groups, which does not affect agents' individual net worth evolution, let N_e and N_h denote aggregate expert and household net worth. The dynamics of N_e

and N_h include the effects of type-switching: $dN_e = N_e \frac{dn_e}{n_e} - \delta_e N_e dt + \delta_h N_h dt$ and $dN_h = N_h \frac{dn_h}{n_h} - \delta_h N_h dt + \delta_e N_e dt$. We reiterate that type-switching is unnecessary for our sunspot results and only serves to obtain stationarity in case we set $\rho_e = \rho_h$ (if $\rho_e > \rho_h + \sigma^2$, the wealth distribution will automatically be stationary even without type-switching). For example, the reader may wish to shut down type-switching ($\delta_e = \delta_h = 0$) and instead consider asymmetric discount rates ($\rho_e > \rho_h + \sigma^2$), and this is completely fine.

The definition of competitive equilibrium is standard, following Brunnermeier and Sannikov (2014). To write a formal definition, denote the set of experts by the interval $\mathbb{I} = [0, \nu]$, for some $\nu \in (0, 1)$ and index individual experts by $i \in \mathbb{I}$. Similarly, denote the set of households by $\mathbb{J} = (\nu, 1]$ with index j. If a type-switching structure exists, we necessarily have $\nu = \frac{\delta_h}{\delta_e + \delta_h}$ (i.e., the population size of experts is pinned down by switching rates), and the indexes of retiring experts/households are implicitly swapped with newly entering experts/households.

Definition 1. For any initial capital endowments $\{k^i_{e,0}, k^j_{h,0} : i \in \mathbb{I}, j \in \mathbb{J}\}$ such that $\int_{\mathbb{I}} k^i_{e,0} di + \int_{\mathbb{J}} k^j_{h,0} dj = K_0$, an *equilibrium* consists of stochastic processes—adapted to the filtered probability space generated by $\{Z_t : t \geq 0\}$—for capital price q_t, interest rate r_t, capital holdings $(k^i_{e,t}, k^j_{h,t})$, consumptions $(c^i_{e,t}, c^j_{h,t})$, and net worths $(n^i_{e,t}, n^j_{h,t})$, such that:

(i) initial net worths satisfy $n^i_{e,0} = q_0 k^i_{e,0}$ and $n^j_{h,0} = q_0 k^j_{h,0}$ for $i \in \mathbb{I}$ and $j \in \mathbb{J}$;

(ii) taking processes for q and r as given, each expert $i \in \mathbb{I}$ and household $j \in \mathbb{J}$ solves (3) subject to (2) and their solvency constraint;

(iii) consumption and capital markets clear at all dates, i.e.,

$$\int_{\mathbb{I}} c^i_{e,t} di + \int_{\mathbb{J}} c^j_{h,t} dj = a_e \int_{\mathbb{I}} k^i_{e,t} di + a_h \int_{\mathbb{J}} k^j_{h,t} dj \tag{4}$$

$$\int_{\mathbb{I}} k^i_{e,t} di + \int_{\mathbb{J}} k^j_{h,t} dj = K_t, \tag{5}$$

where K_t follows (1).

1.1 Equilibrium characterization

We present a useful equilibrium characterization that aids all future analysis. First, conjecture the following form for capital price dynamics:

$$dq_t = q_t[\mu_{q,t} dt + \sigma_{q,t} \cdot dZ_t]. \tag{6}$$

There are two potential avenues for random fluctuations. The standard term $\sigma_q \cdot \binom{1}{0}$ represents amplification (or dampening) of fundamental shocks, as in Brunnermeier and Sannikov (2014) and others. By contrast, the second element $\sigma_q \cdot \binom{0}{1}$ measures sunspot volatility that only exists because agents believe in it.

Given log utility and the scale-invariance of agents' budget sets, individual optimization problems are readily solvable. Optimal consumption satisfies the standard formula $c_\ell = \rho_\ell n_\ell$. Optimality conditions for capital holding by experts and households are

$$\frac{a_e}{q} + g + \mu_q + \sigma\sigma_q \cdot \binom{1}{0} - r = \frac{qk_e}{n_e}|\sigma_R|^2 \tag{7}$$

$$\frac{a_h}{q} + g + \mu_q + \sigma\sigma_q \cdot \binom{1}{0} - r \leq \frac{qk_h}{n_h}|\sigma_R|^2 \quad \text{(with equality if } k_h > 0\text{)}, \tag{8}$$

where

$$\sigma_{R,t} := \sigma\binom{1}{0} + \sigma_{q,t} \tag{9}$$

denotes the shock exposure of capital returns. (Note that experts' optimality condition (7) assumes the solution is interior, i.e., $k_e > 0$. But this is clearly required in any equilibrium given experts earn a strictly higher expected return than households.) From these optimality conditions, notice that agents' capital holdings decisions are uniquely determined given the price process for q. The only additional optimality conditions are the transversality conditions

$$\lim_{T\to\infty} \mathbb{E}[e^{-\rho_\ell T}\frac{1}{c_{\ell,T}}n_{\ell,T}] = 0. \tag{10}$$

However, using $c_\ell = \rho_\ell n_\ell$, we see that (10) automatically holds.

Next, we aggregate. Due to financial frictions and productivity heterogeneity, both the distribution of wealth and capital holdings will matter in equilibrium. However, because all experts (and households) make the same scaled consumption c_ℓ/n_ℓ and portfolio choices k_ℓ/n_ℓ, the wealth and capital distributions may be summarized by experts' wealth share and capital share

$$\eta := \frac{N_e}{N_e + N_h} = \frac{N_e}{qK} \quad \text{and} \quad \kappa := \frac{\int_{\mathbb{I}} k_e^i di}{K}.$$

Given agents' solvency and capital short-sales constraints, we must have $\eta \in [0,1]$ and $\kappa \in [0,1]$ in equilibrium. Substitute optimal consumption into goods market clearing (4),

divide by aggregate capital K, and use the definitions of η and κ, to obtain

$$q\bar{\rho} = \kappa a_e + (1-\kappa)a_h, \tag{PO}$$

where $\bar{\rho}(\eta) := \eta\rho_e + (1-\eta)\rho_h$ is the wealth-weighted average discount rate. Equation (PO) connects asset price q to output efficiency κ, which we call a *price-output* relation.

Using the definitions of η and κ, experts' and households' portfolio shares can be written $\frac{qk_e}{n_e} = \frac{\kappa}{\eta}$ and $\frac{qk_h}{n_h} = \frac{1-\kappa}{1-\eta}$. Then, differencing the optimal portfolio conditions (7)-(8), we obtain the *risk-balance* condition

$$0 = \min\left[1-\kappa, \frac{a_e - a_h}{q} - \frac{\kappa-\eta}{\eta(1-\eta)}|\sigma_R|^2\right]. \tag{RB}$$

Either experts manage the entire capital stock ($\kappa = 1$) or the excess return experts obtain over households, $(a_e - a_h)/q$, represents fair compensation for differential risk exposure, $\frac{\kappa-\eta}{\eta(1-\eta)}|\sigma_R|^2$. On the other hand, summing portfolio conditions (7)-(8), weighted by κ and $1-\kappa$, and using (PO), yields an equation for the riskless rate:

$$r = \bar{\rho} + g + \mu_q + \sigma\sigma_q \cdot \binom{1}{0} - \left(\frac{\kappa^2}{\eta} + \frac{(1-\kappa)^2}{1-\eta}\right)|\sigma_R|^2. \tag{11}$$

Finally, by applying Itô's formula to experts' wealth share $\eta = N_e/(N_e + N_h)$, and using agents' net worth dynamics (2) along with contributions from type-switching, wealth share dynamics are given by

$$d\eta_t = \mu_{\eta,t}dt + \sigma_{\eta,t} \cdot dZ_t, \quad \text{given} \quad \eta_0, \tag{12}$$

where

$$\mu_\eta = \eta(1-\eta)(\rho_h - \rho_e) + (\kappa - 2\eta\kappa + \eta^2)\frac{\kappa-\eta}{\eta(1-\eta)}|\sigma_R|^2 + \delta_h - (\delta_e + \delta_h)\eta \tag{13}$$

$$\sigma_\eta = (\kappa-\eta)\sigma_R. \tag{14}$$

The initial wealth distribution $\eta_0 = \frac{\int_{\mathbb{I}} n^i_{e,0}di}{q_0 K_0} = \frac{\int_{\mathbb{I}} k^i_{e,0}di}{K_0}$ is given due being solely a function of the initial endowments of capital.

Lemma 1. *Given $\eta_0 \in (0,1)$, consider a process $(\eta_t, q_t, \kappa_t, r_t)_{t\geq 0}$ with dynamics for q_t and η_t described by (6) and (12), respectively. If $\eta_t \in [0,1]$, $\kappa_t \in [0,1]$, and equations (PO), (RB), (11), (13) and (14) hold for all $t \geq 0$, then $(\eta_t, q_t, \kappa_t, r_t)_{t\geq 0}$ corresponds to an equilibrium of Definition 1. Moreover, any distinct pair of such processes corresponds to distinct equilibria.*

Lemma 1 summarizes the full set of conditions characterizing equilibrium and is proved in Appendix B.2. In the rest of the paper, we use this lemma as a tool to simplify our search for equilibria. Lastly, we make some mild parameter restrictions that will be applicable in the remainder of the paper.

Assumption 1. *Parameters satisfy (i)* $0 < \frac{a_h}{\rho_h} < \frac{a_e}{\rho_e} < +\infty$*; (ii)* $\sigma^2 < \rho_e(1 - a_h/a_e)$*; and (iii) either* $\sigma^2 < \rho_e - \rho_h$*, or* $\delta_e, \delta_h > 0$*.*

Assumption 1 part (i) makes the modest assumption that the capital price is higher if experts control 100% of wealth than if households control 100% of wealth. Part (ii), meant to make the problem interesting, ensures experts sometimes hold all capital, i.e., $\kappa = 1$. If fundamental risk is $\sigma^2 \geq \rho_e(1 - a_h/a_e)$, experts can never hold the entire capital stock, and the economy will always be in the region of inefficiency. Part (iii) ensures household survival: if experts consume at a rate sufficiently higher than households, or some type-switching exists, then experts do not asymptotically hold all wealth.

1.2 Types of equilibria

We categorize our equilibria into two types: fundamental and sunspot. Fundamental equilibria should have the property that only the minimal set of state variables affects observables. Because of financial frictions and productivity heterogeneity, the expert wealth share η is a necessary state variable to summarize economic conditions. Other objects (e.g., q, r, κ) are either prices or control variables, so in that sense η is the minimal state variable needed in this class of models. In other words, a fundamental equilibrium should only depend on η. Sunspot equilibria constitute all other equilibria.

Definition 2. A *Fundamental Equilibrium* (FE) is an equilibrium that is Markov in η. Any other equilibrium is a *Sentiment-driven Brownian Sunspot Equilibrium* (S-BSE).

The literature universally focuses on the FE of this model, e.g., Brunnermeier and Sannikov (2014). We explain these fundamental equilibria in Online Appendix E. The present paper is devoted to the S-BSEs.[13]

[13]We add the qualifier "Sentiment-driven" (i.e., S-BSE) because there actually exists a sunspot equilibrium that is Markovian in η, but only if fundamental risk is absent, $\sigma = 0$. We study this type of equilibrium, which one might refer to as a *Wealth-driven Brownian Sunspot Equilibrium* (W-BSE), in a companion paper (Khorrami and Mendo, 2024). In the context of the present paper, Definition 2 thus perfectly classifies non-sunspot and sunspot equilibria if $\sigma > 0$.

1.3 Benchmarks and discussion

Before proceeding to the main analysis, we analyze two benchmarks—frictionless equity issuance and homogeneous productivities—that clarify the underpinnings of sentiment-driven equilibria.

Frictionless equity issuance. Suppose any agent, when managing capital, could issue unlimited equity to the market. In exchange for taking some exposure to the risk σ_R in capital returns, these outside equity contracts promise an expected excess return $\sigma_R \cdot \pi$ (here, π is the equilibrium risk price vector associated to the two shocks in Z). All agents can participate as buyers in this market. Since equity-issuance is unconstrained, it is straightforward to see that any capital owner must equate her expected excess returns on capital to $\sigma_R \cdot \pi$. (If $\sigma_R \cdot \pi$ were below an agent's expected excess capital returns, unlimited capital purchases financed by unlimited equity issuances would be an arbitrage; if $\sigma_R \cdot \pi$ were above, the agent would prefer to sell all their capital and invest solely in equity securities.) Experts always manage some capital, so

$$\frac{a_e}{q} + g + \mu_q + \sigma\sigma_q \cdot \left(\begin{smallmatrix} 1 \\ 0 \end{smallmatrix}\right) - r = \sigma_R \cdot \pi.$$

However, the analogous equation cannot hold for households, since their productivity is lower, $a_h < a_e$. Households will never manage capital in this economy, so $\kappa_t = 1$ at all times, hence $q_t = a_e/\bar{\rho}(\eta_t)$ by equation (PO). That q is solely a function of η rules out S-BSEs.[14] Thus, it is critical that capital is traded, i.e., κ varies.

For our main results, the friction in equity markets need not be as stark as the baseline model. Indeed, Online Appendix F.1 extends the baseline model to allow "partial equity issuance," subject to a constraint parameterized by $\chi \in [0,1]$. In particular, suppose any agent can issue some equity up to a limit: he/she can offload up to $1-\chi$ fraction of the risk associated to their capital stock as equity to a competitive financial market. The baseline model corresponds to $\chi = 1$ (i.e., zero issuance), while the frictionless model outlined above corresponds to $\chi = 0$ (i.e., unlimited issuance). We show that self-fulfilling volatility is possible for any $\chi > 0$, but as mentioned above, no self-fulfilling volatility is possible if $\chi = 0$.

Homogeneous productivities. Consider our economy with $a_e = a_h = a$. Based on

[14] In fact, q cannot be stochastic at all. Indeed, experts and households share identical risk preferences, so households will purchase the outside equity of experts in an amount that is consistent with perfect risk-sharing, meaning $\sigma_\eta \equiv 0$. Since $q_t = a_e/\bar{\rho}(\eta_t)$ is solely a function of η, which is deterministic, we have $\sigma_q \equiv 0$ as well. Shocks can play no amplifying role with frictionless equity markets.

equation (PO), equal productivities immediately implies $q_t = a/\bar{\rho}(\eta_t)$. Again, q is solely a function of η, which rules out S-BSEs. Critically, sentiment-driven equilibria require real outcomes to depend on κ.

In fact, with equal productivities, equilibrium cannot support any endogenous dependence on shocks, i.e., one can show $\sigma_q \equiv 0$ when $a_e = a_h$.[15] This unveils a more general point about the endogeneity of market incompleteness: one cannot necessarily add unspanned extrinsic shocks to an economy and declare markets incomplete. Even though this equal-productivity economy lacks insurance markets against $Z^{(2)}$ shocks, financial markets are *effectively complete*, in the sense that the economic structure imposes that $Z^{(2)}$ can have no impact on outcomes. What is required is a set of assumptions such that $Z^{(2)}$ has "real effects" in which case financial market incompleteness will have some bite. In our economy, all we require is $a_e > a_h$.[16]

Discussion: imperfect risk-sharing and productivity heterogeneity. Based on the benchmarks above, let us explain the deep reasons why our model admits sentiment-driven equilibria. The fact that we require financial frictions and productivity heterogeneity is not surprising—these features are required even in the "financial accelerator" equilibria of Kiyotaki and Moore (1997) and Brunnermeier and Sannikov (2014). More interestingly, sentiment-driven equilibria require nothing more.

First, with limited equity issuance and lack of markets for insurance against sunspot shocks, capital is traded partly for risk-sharing purposes. In other words, risk can affect the capital ownership distribution (i.e., σ_R can affect κ). Second, productive heterogeneity permits "misallocation": the capital distribution can affect aggregate output, which translates into capital prices (i.e., κ can affect q).

Of course, all these endogenous variables are determined simultaneously, but it may be helpful to visualize, with the symbols of our model, the logic of multiplicity through the following chain of causality:

$$\sigma_R \Longrightarrow \kappa \Longrightarrow q. \tag{15}$$

Financial frictions modulate the first link ($\sigma_R \Rightarrow \kappa$), while productive heterogeneity modulates the second ($\kappa \Rightarrow q$). The current asset price q then depends on the distribution of

[15] Plugging $a_e = a_h$ into equation (RB) implies either $\kappa = \eta$ or $|\sigma_R| = 0$. Either way, $\sigma_\eta = (\kappa - \eta)\sigma_R = 0$. Then, applying Itô's formula to $q_t = a/\bar{\rho}(\eta_t)$, we obtain $q\sigma_q = -\frac{\rho_e - \rho_h}{\bar{\rho}(\eta)} q\sigma_\eta$, which equals zero.

[16] Our model imposes a two-point productivity distribution, in line with the literature, but this is not necessary for our arguments. In fact, there are an even richer set of possible equilibria in a model with a richer productivity distribution. Online Appendix F.2 considers the case of a continuum of productivities, $a \in [\underline{a}, \bar{a}]$, and constructs a sentiment equilibrium.

future asset prices through σ_R. But what determines σ_R? Nothing, as long as we have both financial frictions and productive heterogeneity. S-BSEs, by removing the ad-hoc restriction that equilibria be Markov in η, remove an artificial anchor for σ_R and allow volatility to be coordination-driven.

Chain (15) also suggests a connection to the "self-fulfilling risk panics" of Bacchetta et al. (2012), further analyzed by Benhabib et al. (2020). Bacchetta et al. (2012) emphasize a negative relationship between asset prices and volatility, effectively collapsing the causal chain in equation (15) to $\sigma_R \Rightarrow q$. But digging deeper, Benhabib et al. (2020) explain that the key to risk panic equilibria is a causal dependence of the stochastic discount factor (SDF) on asset prices. Bacchetta et al. (2012) obtain a price-SDF link via OLG (see also Farmer, 2018, and Gârleanu and Panageas, 2021); Benhabib et al. (2020) show how a price-SDF link can also arise due to collateral or liquidity constraints. Our results are deeply connected—our price-output link (PO) necessarily implies a price-SDF link—but distinguished by the fact we do not rely on the common multiplicity-inducing features of OLG or ad-hoc borrowing constraints.

2 Sentiment-driven equilibria

We endeavor here to analyze a rich class of equilibria that are not Markov in η, the S-BSEs. Below, we construct and provide detailed characterization of such equilibria.

Because the capital price q is the critical endogenous object (one may think of q as the "co-state" variable), equilibria which are not Markov in η share the defining characteristic that a variety of different asset prices can prevail for a given wealth distribution. Since η captures all fundamental information in this economy, one can think of "sentiment" as responsible for generating the multiplicity of asset prices corresponding to the same η. This is why Definition 2 refers to this class of equilibria as Sentiment-driven BSEs.

To prepare the reader for our construction, note that it differs substantially from the usual approach. The usual approach first analyzes the non-stochastic equilibria of a model, identifies a fundamental indeterminacy, and then adds sunspot shocks that essentially randomize over the multiplicity of fundamental equilibria (Azariadis, 1981; Cass and Shell, 1983). We cannot take this route because the deterministic equilibrium of our model is *unique*. To see this, assume $\sigma = 0$ and suppose agents coordinate on $\sigma_q = 0$. From equation (RB), we must have $\kappa = 1$ (experts always hold all capital), and from equation (PO), the capital price will be $q = a_e/\bar{\rho}(\eta)$, a unique function of the wealth distribution. Our approach, instead, constructs a stochastic equilibrium and proves all

the equilibrium requirements hold, by using tools from "stochastic stability."[17]

2.1 Construction of S-BSEs

Now, we provide a sketch of an explicit construction of an S-BSE. Remember the goal from Lemma 1: given η_t, we want to find $(\mu_{\eta,t}, \sigma_{\eta,t}, \mu_{q,t}, \sigma_{q,t}, q_t, \kappa_t, r_t)$ satisfying equations (PO), (RB), (11), and (13)-(14) for all $t \geq 0$ and such that $\eta_t, \kappa_t \in [0,1]$.

First, let us first count the number of equations and unknowns. The equations are (PO), (RB), (11), (13), and (14)—these are 6 equations (recall that (14) involves two equations) that hold at each time t. Given η_t at a particular point in time, the unknowns are the wealth share dynamics (μ_η, σ_η), the level and dynamics of capital prices (q, μ_q, σ_q), the capital share κ, and the interest rate r—these are 9 unknowns (recall σ_η and σ_q are 2-by-1 vectors). Thus, we seem to have 3 degrees of freedom in constructing equilibrium. A Fundamental Equilibrium, universally studied by the literature, additionally imposes that equilibria be Markov in η. Such a Markovian restriction eliminates the 3 degrees of freedom: applying Itô's formula to $q(\eta)$ delivers 3 additional conditions for σ_q and μ_q. But in an S-BSE, q_t is not simply a function of η_t, so the 3 Itô conditions are dropped. Instead, (σ_q, μ_q) are determined by coordination.

Domain. The specific construction we outline below has the property that all equilibrium objects are functions of (η_t, q_t), but this is generalized in the formal proof. We are using one degree of freedom in making q a "state variable" in the equilibrium. It turns out that the relevant domain for (η, q) is

$$\mathcal{D} := \{(\eta, q) : 0 < \eta < 1,\ q^L(\eta) < q \leq q^H(\eta)\}, \tag{16}$$

$$\text{where} \quad q^H(\eta) := a_e/\bar{\rho}(\eta)$$
$$q^L(\eta) := [\eta a_e + (1-\eta)a_h]/\bar{\rho}(\eta).$$

Why is this the relevant domain? From the price-output relation (PO), notice that q^H corresponds to the capital price when $\kappa = 1$, whereas q^L corresponds to the capital price when $\kappa = \eta$. Equilibrium must have $\kappa \leq 1$ (Lemma 1) and $\kappa > \eta$, the latter because a solution to equation (RB) will not exist otherwise. These restrictions are captured by ensuring (η, q) remains in $\mathcal{D}$. The shaded region in Figure 3 illustrates this set. For reference, we also place the Markovian Fundamental Equilibrium (FE) with $\sigma = 0.1$.

[17] For a simplified introduction to the mathematics of this method, Online Appendix H presents a reduced-form example where deterministic stability suggests a unique equilibrium but where, nevertheless, a stochastic stability analysis shows that there are multiple stochastic equilibria.

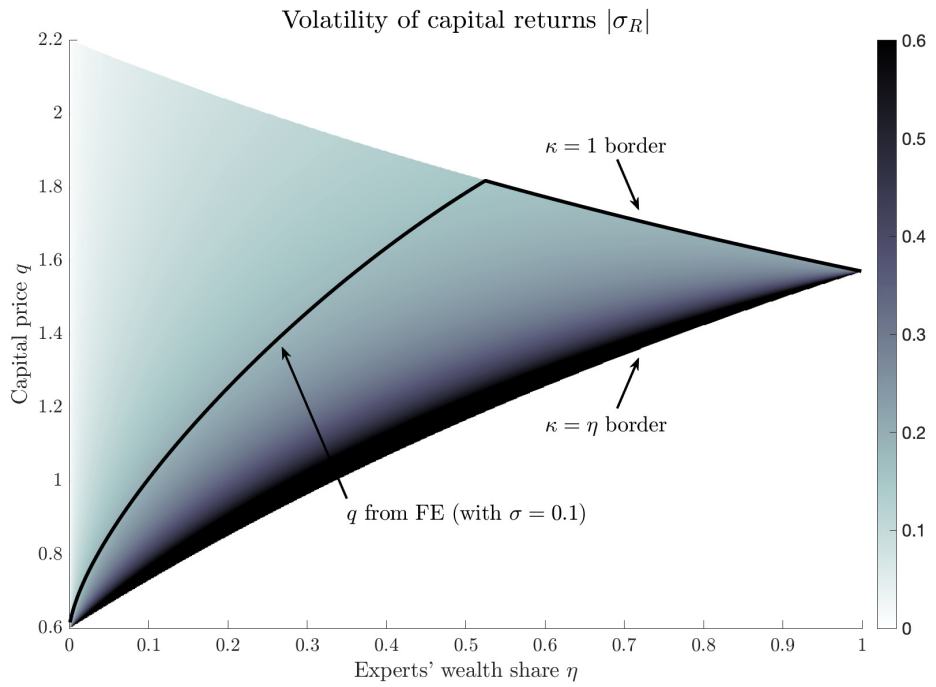


Figure 3: Colormap of volatility $|\sigma_R|$ as a function of (η, q), in the region $\mathcal{D} := \{(\eta, q) : \eta \in (0,1)$ and $\eta a_e + (1-\eta)a_h < q\bar{\rho}(\eta) \leq a_e\}$. Volatility is truncated for aesthetic reasons (since $|\sigma_R| \to \infty$ as $\kappa \to \eta$). Also plotted is the Fundamental Eqm (FE) with $\sigma = 0.1$. Parameters: $\rho_e = 0.07$, $\rho_h = 0.05$, $a_e = 0.11$, $a_h = 0.03$.

Use all the equations. The first step in the construction is to reduce the system. Imagine we know the values of $(\eta, q, \sigma_q, \mu_q)$. Price-output relation (PO) determines κ as a function of (η, q) and nothing else, given by

$$\kappa(\eta, q) := \frac{q\bar{\rho}(\eta) - a_h}{a_e - a_h}. \tag{17}$$

Substituting this result for κ, equation (11) then fully determines r. Equations (13)-(14), after plugging in the result for κ, fully determine (σ_η, μ_η). At this point, given (η, q), the remaining unknowns are (σ_q, μ_q) and the remaining equation is (RB).

The solution for (σ_q, μ_q) depends critically on whether capital is efficiently allocated or not. When capital is efficiently allocated (i.e., $\kappa = 1$), we have $q = q^H(\eta)$ as an explicit function of η. Hence, both σ_q and μ_q are determined by Itô's formula.

But when $q < q^H(\eta)$ (i.e., $\kappa < 1$), we have much more flexibility. Equation (RB) says

$$|\sigma_R| = \sqrt{\frac{\eta(1-\eta)}{\kappa(\eta, q) - \eta} \frac{a_e - a_h}{q}}, \quad \text{if} \quad q < q^H(\eta). \tag{18}$$

In other words, given (η, q), the level of return volatility is pinned down, verifying the

statement made earlier that putting q as a state variable uses one degree of freedom. The level of $|\sigma_R|$ is plotted in Figure 3 via the darkness of the shading; notice that $|\sigma_R|$ and q are inversely related. But notice also that equation (18) only restricts the norm of $\sigma_q = \sigma_R - \sigma\left(\begin{smallmatrix}1\\0\end{smallmatrix}\right)$, not each of its components separately.

Similarly, there is as yet no restriction on μ_q despite using all 6 equilibrium equations. All that remains is to show that $(\eta_t, q_t)_{t\geq 0}$ remains in $\mathcal{D}$ almost-surely, which puts some mild restrictions on μ_q. Proving $(\eta_t, q_t)_{t\geq 0}$ remains in $\mathcal{D}$ is critical to ensure that no optimality or market clearing conditions are violated along the proposed equilibrium path. Specifically, equation (18) is only well-defined for $\kappa > \eta$, or equivalently $q > q^L(\eta)$, while Lemma 1 also requires $\kappa \leq 1$ and $\eta \in [0,1]$. These inequalities only hold on $\mathcal{D}$.

Boundary conditions. To ensure that (η_t, q_t) remains in $\mathcal{D}$, all we need to impose are *boundary conditions* on μ_q. The idea is that (η_t, q_t) can only escape $\mathcal{D}$ through its boundaries, and so μ_q is only restricted at these boundaries. In particular, we only need some force strong enough to push (η_t, q_t) back toward the interior of $\mathcal{D}$. For example, when $q < q^H(\eta)$, we can set μ_q to *any* C^1 function with a boundary condition like, e.g.,

$$\inf_{\eta\in(0,1)} \lim_{q\searrow q^L(\eta)} \left[q - q^L(\eta)\right]\mu_q(\eta, q) = +\infty. \tag{19}$$

Condition (19) says that the drift of q diverges fast enough in order to prevent q from hitting $q^L(\eta)$. But actually the drift μ_q diverges slightly above $q^L(\eta)$. The conditions at the upper boundary $q^H(\eta)$ are slightly more complicated because the economy is actually allowed to visit this upper boundary. However, a similar indeterminacy exists there, namely that (η, q) can remain stuck at the efficient level $(\eta, q^H(\eta))$ for any amount of time before re-entering the interior of $\mathcal{D}$. In short, equilibrium only imposes (some) boundary conditions on μ_q and leaves it indeterminate on most of $\mathcal{D}$. For instance, it is not even required that μ_q be solely a function of (η, q) on $\mathcal{D}$, and in fact we utilize that flexibility in our later examples.

Methodologically, our formal proof employs stochastic stability theory to show that this construction yields a non-degenerate stationary distribution for $(\eta_t, q_t)_{t\geq 0}$. Appendix B.4 states and proves the appropriate version of a stochastic stability lemma that we use. In particular, the key object is the infinitesimal generator $\mathscr{L}$ of the joint process $(\eta_t, q_t)_{t\geq 0}$ induced by equilibrium. And the key task is to find a positive (Lyapunov) function v, which diverges at the boundaries of $\mathcal{D}$, such that $\mathscr{L}v \to -\infty$ at the boundaries of $\mathcal{D}$. This mathematical condition exactly captures the intuition that boundary conditions on the dynamics are sufficient for stationarity. (The ability to leverage

stochastic stability theory to analyze boundary conditions is precisely the simplification offered by our continuous-time setup. That said, Online Appendix G also constructs an example sentiment-driven equilibrium in a discrete-time version of our model.)

Theorem 1 (Existence)**.** *Let Assumption 1 hold. Then, there exists a family of S-BSEs in which* $(\eta_t, q_t)_{t\geq 0}$ *remains in* $\mathcal{D}$ *almost-surely and possesses a non-degenerate stationary distribution.*

This family of equilibria is indexed as follows. Let $(y_t)_{t\geq 0}$ *be any sufficiently well-behaved exogenous, stationary, Markov diffusion. Let* $\mathcal{D}^\circ \subset \mathcal{D}$ *be any sub-domain not touching any boundary, i.e., such that* $dist(\mathcal{D}^\circ, \partial\mathcal{D}) > 0$*. Let* $\vartheta(\eta, q, y) \in [0, 1]$ *and* $m(\eta, q, y)$ *be any* C^1 *functions. Then, an S-BSE in this family exists in which*

(i) on $\mathcal{D}^\circ$*, the share of return variance* $|\sigma_R|^2$ *due to the fundamental shock is* $\vartheta(\eta, q, y)$

(ii) on $\mathcal{D}^\circ$*, the drift of* q *is* $m(\eta, q, y)$

Furthermore, this S-BSE can feature an arbitrary exit rate from the efficient region $\{\kappa = 1\}$ *to the inefficient region* $\{\kappa < 1\}$*, in the sense that the expected first exit time* $T(\eta, q, y)$ *can be any solution to equation (B.16), whose coefficients are indeterminate, in the appendix.*

Theorem 1 is formally proved in Appendix B.3 with an explicit S-BSE construction that addresses several of the minor technical issues raised in the preceding discussion. Because the proof is constructive, we obtain a complete characterization of the indeterminacies in equilibrium, which is why the theorem delivers a "family of equilibria."

The first indeterminacy is the *source of volatility*. While $|\sigma_R|$ is pinned down given (η, q), from equation (18), the two components $\sigma_R^{(1)}$ and $\sigma_R^{(2)}$ are not. The reason: when trading, agents only care about total return variance, not its source. Asset prices and economic activity can be either closely linked to fundamentals, or completely decoupled from them, and this decoupling can be time-varying in arbitrary ways.

The second indeterminacy is the *degree of mean-reversion*. While μ_q is pinned down at the boundaries of $\mathcal{D}$, as in equation (19), its values in the interior of $\mathcal{D}$ are indeterminate. The reason: optimal capital holdings are a function of the risk premium. This is clearly visible in the optimal portfolio FOCs (7)-(8), where only the spread $\mu_q - r$ appears. Consequently, even given a price q and diffusion σ_q, only the spread $\mu_q - r$ is pinned down in equilibrium; μ_q and r are not separately determined. For example, asset prices could behave like a random walk (corresponding to $\mu_q \approx 0$ in the interior), with just enough mean-reversion in extreme states to keep things stationary; in such a design, extreme states become arbitrarily close to reflecting boundaries. Alternatively, fluctuations could be much more transitory in nature. The indeterminacy in mean-reversion also manifests

in the efficient region $\{\kappa = 1\}$, because the rate at which the economy re-enters the inefficient region $\{\kappa < 1\}$ is arbitrary, in a sense made precise by the theorem.

A key property of our S-BSEs is that they permit much larger volatility than conventional equilibria. To see this, refer back to Figure 3, paying attention to the level of return volatility $|\sigma_R|$ (indicated by shading). Notice how the Fundamental Equilibrium attains only 10-20% return volatility, whereas the S-BSE can seemingly attain much higher levels when q is very low. In fact, equation (18) shows that $|\sigma_R| \nearrow +\infty$ as $q \searrow q^L(\eta)$ (equivalently, $\kappa \searrow \eta$). The next result summarizes the range of possible q and $|\sigma_R|$.

Corollary 1 (Unbounded volatility)**.** *Given wealth share $\eta \in (0,1)$, let $\mathcal{Q}(\eta)$ denote the set of possible S-BSE values of q, and let $\mathcal{V}(\eta)$ denote the associated set of possible S-BSE values of return variance $|\sigma_R(\eta, q)|^2$. Then, $\mathcal{Q}(\eta)$ is an interval with $\inf \mathcal{Q}(\eta) = q^L(\eta)$ and $\sup \mathcal{Q}(\eta) = q^H(\eta)$, while $\mathcal{V}(\eta)$ consists of at most two intervals, with*

$$\inf \mathcal{V}(\eta) = \min \left[\eta\bar{\rho}(\eta)\frac{a_e - a_h}{a_e}, \sigma^2(\bar{\rho}(\eta)/\rho_e)^2\right]$$
$$\sup \mathcal{V}(\eta) = +\infty.$$

We conclude this section by providing two classes of examples for the indeterminate drift μ_q and then summarizing our results with a remark.

Example 1 (Reflecting boundaries)**.** As a first example, let us introduce upper and lower reflecting barriers $\overline{q}(\eta) \leq q^H(\eta)$ and $\underline{q}(\eta) > q^L(\eta)$, where $\overline{q}$ and $\underline{q}$ can be arbitrary functions. These reflecting barriers keep $q_t \in \left[\underline{q}(\eta_t), \overline{q}(\eta_t)\right]$ almost-surely and affect no other equation in the model.[18] Now, set $\mu_q = 0$ for all $q \in \left(\underline{q}(\eta_t), \overline{q}(\eta_t)\right)$, so that q_t behaves exactly like a random walk until it hits one of the reflecting barriers. This construction constitutes a legitimate S-BSE. An example with three different reflecting lower boundaries $\underline{q}(\eta)$ is displayed in Figure 4. The figure makes clear that the possibility of extreme volatility is strongly influenced by the level of the lower barrier.

[18] Formally, reflection introduces a new term to price dynamics:

$$dq_t = q_t[\mu_{q,t}dt + \sigma_{q,t} \cdot dZ_t + d\underline{P}_t - d\overline{P}_t],$$

where $\underline{P}$ and $\overline{P}$ are the barrier processes that increase only to keep $q_t \geq \underline{q}(\eta_t)$ or $q_t \leq \overline{q}(\eta_t)$, respectively. Let $P := \underline{P} - \overline{P}$. Absence of arbitrage requires the riskless bond return to be $r_t dt + dP_t$, such that the excess return on capital is unaffected by dP_t (c.f., Karatzas and Shreve, 1998, Appendix B). Consequently, agents' FOCs on capital holding remain unaffected, and both the risk-balance condition (RB) and equation (11) for r_t still hold. Finally, the reflections have no impact on the dynamics of η_t, which still take the diffusive form (12). Indeed, excess capital returns feature no components related to dP_t component, so expert and household return-on-wealth contain identical contributions from dP_t, implying $d\eta_t$ contains no dP_t term.

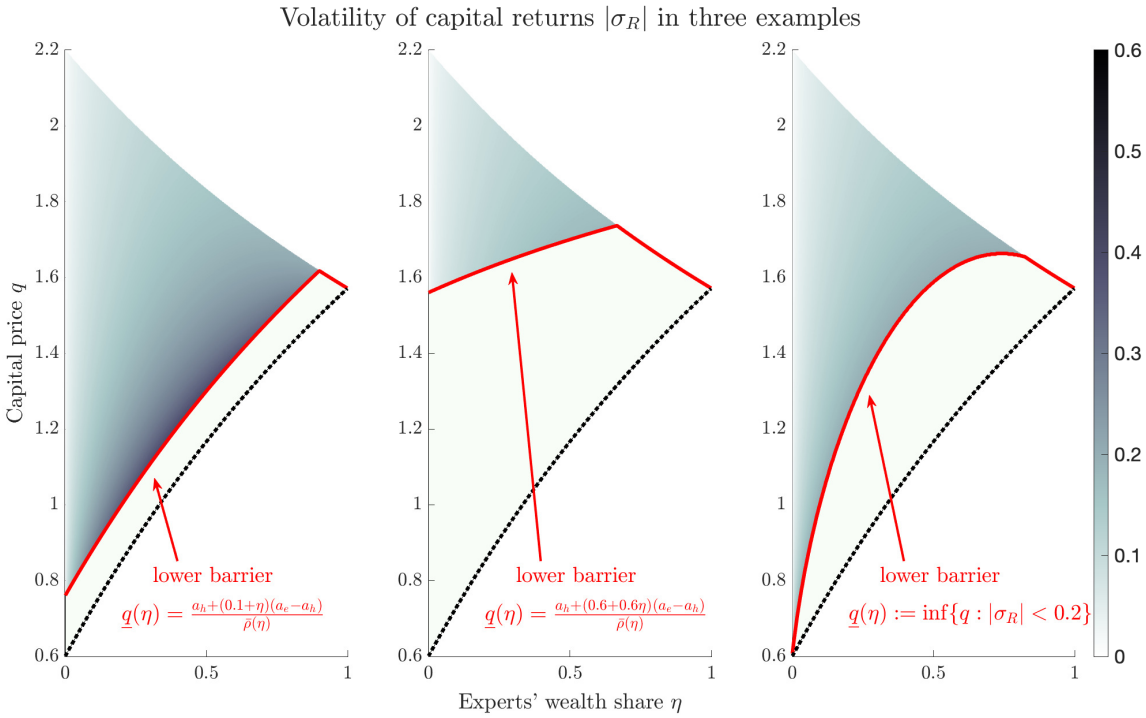


Figure 4: Colormap of volatility $|\sigma_R|$ as a function of (η, q). In each panel, the solid red line denotes the lower reflecting boundary $\underline{q}(\eta)$, and the light pink shaded area denotes the subset of $\mathcal{D}$ which is inaccessible. Parameters: $\rho_e = 0.07$, $\rho_h = 0.05$, $a_e = 0.11$, $a_h = 0.03$.

Example 2 (Interest rate)**.** Our second example sets the drift μ_q via the interest rate r. Recall that these objects are only determinate insofar as their spread $\mu_q - r$ is pinned down in the interior of $\mathcal{D}$. In the inefficient region, let r follow any exogenous, stationary process of the form:

$$dr_t = \mu_r(r_t)dt + \sigma_r(r_t) \cdot dZ_t, \quad \text{when} \quad \kappa_t < 1.$$

In addition, augment the dynamics with a lower reflecting barrier $\underline{q}(\eta)$ as in Example 1. Given r, define μ_q by equation (11) for all $\{(\eta, q) : 0 < \eta < 1, \underline{q}(\eta) < q < q^H(\eta)\}$, i.e., in the inefficient region but above the reflecting barrier. This construction constitutes a legitimate S-BSE. (In the efficient region, by contrast, r is pinned down by equation (11), since μ_q is pinned down by Itô's formula on $q = q^H(\eta)$.)

Remark 1 (Dynamics and indeterminacies)**.** *Summarizing the results, indeterminacies arise because beliefs about capital price dynamics influence real outcomes such as capital allocation. In this model we have two prices—capital price q and interest rate r—and two (non-redundant) market clearing conditions. However, we need to solve not only for current prices but also for future capital price behavior, which is summarized by the diffusion $\sigma_q \in \mathbb{R}^2$ and drift $\mu_q \in \mathbb{R}$*

terms.[19] *Optimality imposes a tight (negative) link between q and $|\sigma_q|$, while long-run stability imposes some mild conditions on μ_q in extreme states. Besides those restrictions, (σ_q, μ_q) are indeterminate. We map these three indeterminacies to total return volatility (Corollary 1), the source of volatility (point (i) of Theorem 1), and the persistence of sentiment fluctuations (point (ii) of Theorem 1).*

2.2 Economic intuition behind S-BSEs

Next, we explain our S-BSEs more intuitively. We first offer an interpretation of our equilibrium as driven by *uncertainty shocks*. Then, we take a dynamical-system perspective to understand why self-fulfilling volatility is possible.

Uncertainty shocks. Given a wealth distribution η and a level of return volatility $|\sigma_R|$, the capital market is equilibrated at each time via the risk-balance condition (RB) and the price-output relation (PO), restated here for convenience:

$$0 = \min\left[1-\kappa,\ \frac{a_e - a_h}{q} - \frac{\kappa-\eta}{\eta(1-\eta)}|\sigma_R|^2\right] \tag{RB}$$

$$q\bar{\rho} = \kappa a_e + (1-\kappa)a_h. \tag{PO}$$

The left panel of Figure 5 shows how the intersection of these two curves determines the capital allocation κ and the capital price q. The downward-sloping risk-balance (RB) can be thought of as experts' relative capital demand: for a fixed level of wealth η and return volatility $|\sigma_R|$, experts will only hold more capital if it is cheaper, thereby offering a higher expected return. (Of course, households also want to buy capital when it is cheaper, but this force is relatively stronger for experts because of their productivity advantage.) The upward-sloping price-output (PO) is a capital supply curve: experts' capital provision raises allocative efficiency and capital valuations.

But whereas η is a state variable that can be rightly treated as fixed in this static sense, return volatility $|\sigma_R|$ is not. The right panel of Figure 5 shows what changes if there is a sudden rise in *fear*, manifested as higher perceived volatility $|\sigma_R|$. Experts, being risk-averse, are less willing to hold capital when volatility is high. This is illustrated as a

[19] The logic in a discrete time model is analogous: the indeterminacies will be associated to the distribution of capital price tomorrow. This distribution is an infinite dimensional object, which makes it challenging to prove the existence of our sentiment-driven equilibria in discrete time models. Online Appendix G provides a discrete-time example of a sentiment-driven equilibrium by specializing to a binomial tree for capital prices. We purposely design this binomial example with a trading interval Δ such that our Brownian model is recovered as $\Delta \to 0$.

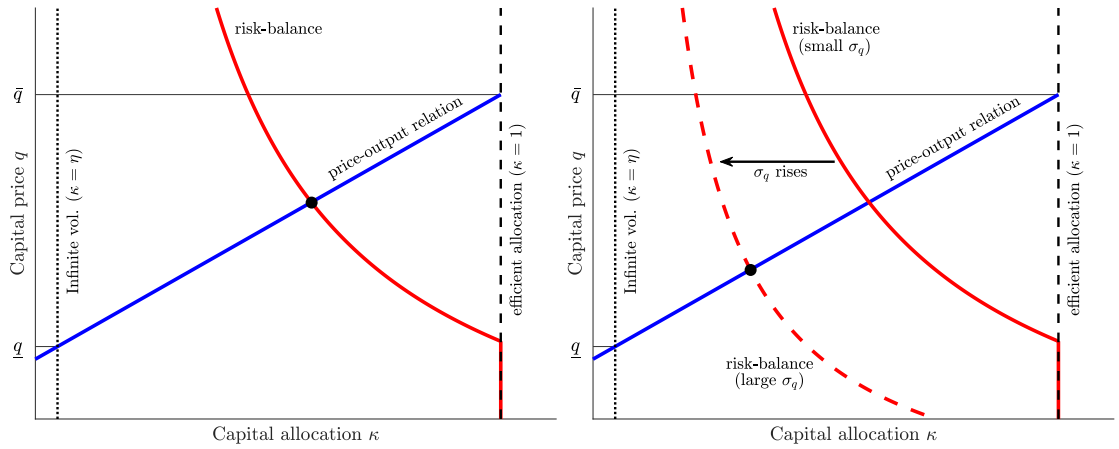


Figure 5: An uncertainty shock. Both panels plot the risk-balance condition (RB) and price-output relation (PO) for a fixed level of $\eta = 0.2$. The horizontal lines labeled $\bar{q}$ and $\underline{q}$ refer to maximal and minimal values of the capital price, corresponding to an efficient capital allocation ($\kappa = 1$) and an infinite-volatility allocation ($\kappa = \eta$), respectively. *Left panel:* equilibrium with $|\sigma_R| = 0.13$. *Right panel:* equilibrium after a shift to $|\sigma_R| = 0.20$. Other parameters: $\rho_e = \rho_h = 0.05$, $a_e = 0.11$, $a_h = 0.03$, and $\sigma = 0.10$.

leftward shift in the risk-balance curve from the solid to the dashed line. After this "fire sale," capital is allocated less efficiently, and asset prices are lower.

So far, nothing rules out this arbitrary rise in fear, and $|\sigma_R|$ appears indeterminate. Mathematically, fixing the state variable η, equations (RB) and (PO) constitute two equations in the three unknowns $(\kappa, q, |\sigma_R|)$. The indeterminacy in $|\sigma_R|$ translates into an indeterminacy in q, which can be seen by combining (RB) and (PO) to eliminate κ and obtain the negative price-variance association:

$$|\sigma_R|^2 = \frac{\eta(1-\eta)(a_e - a_h)^2}{q\bar{\rho}(\eta) - \eta a_e - (1-\eta)a_h} \frac{1}{q} \quad \text{when } \kappa < 1. \tag{20}$$

In our construction leading up to Theorem 1, we treated (η, q) as state variables and determined all other equilibrium objects as functions of (η, q). The preceding story about fear suggests that one can also think of S-BSEs as being driven by uncertainty shocks—time-varying beliefs about volatility $|\sigma_R|$—an interpretation which is supported by the one-to-one mapping between q and $|\sigma_R|$ in equation (20). To see this graphically, refer back to Figure 3: the shading, representing $|\sigma_R|$, darkens monotonically as q falls.

Bounce-back beliefs and dynamic stability. Based on the static conditions (RB) and (PO), equilibrium seems to support a multiplicity of prices q for a fixed η. To understand the beliefs that sustain this multiplicity, it helps to take a dynamical-systems perspective.

Let us think of $(\eta_t, q_t)_{t \geq 0}$ as a stochastic dynamical system. As in deterministic dynamical systems, a pair (η_t, q_t) will only be an equilibrium if it does not lead to explosive

paths. Thus, beliefs must be such that (η_t, q_t) will mean-revert, or bounce back, from extreme states. What does this entail?

To fix ideas, consider the following explosive path. Suppose a fear shock raises volatility $|\sigma_q|$ and lowers asset prices q. Under higher volatility, any subsequent fear shocks would have a larger direct impact on q, further raise volatility $|\sigma_q|$, and so on, ad infinitum. Thus, with enough such fear shocks, we will have $q \searrow q^L(\eta)$ and $|\sigma_q| \nearrow +\infty$.

For this fear-driven path to be an equilibrium, agents must believe that, at least eventually, q will recover and $|\sigma_q|$ will fall. In other words, agents must believe μ_q will increase enough to buoy q from its low level. This is an example of what we label *bounce-back beliefs*. Bounce-back beliefs can be justified because μ_q is indeterminate, as established in Theorem 1.

Translating agents' bounce-back beliefs into specific mathematical conditions on μ_q is straightforward. Because $(\eta_t, q_t)_{t\geq 0}$ evolves in a diffusive fashion, stability criteria conveniently boil down to boundary behavior of the dynamical system. Imposing conditions on μ_q at the boundaries of the domain $\mathcal{D}$ (i.e., the triangle in Figure 3) is sufficient to ensure a stochastically stable system.

In a sense, the mean-reversion embedded in bounce-back beliefs is precisely the mechanism of self-fulfillment in our model. Fear can push asset prices very low precisely because a recovery is expected. Prices can rise in a sentiment-driven boom precisely because agents know the boom will eventually subside. But if the only requirement is that mean-reversion eventually takes hold, there remains significant scope for different types of dynamics. The rise in fear can come from a fundamental or sunspot shock, and it can be very persistent or very transient.

2.3 Enriched model with investment and jumps

In the analysis so far, self-fulfilling fear is captured fully by return variance, and capital growth is completely exogenous. Here, we relax both of these assumptions, allowing sunspot jumps and capital investment. Later, we use this enriched model in a quantitative exercise. More details on the equilibrium characterization for this enriched model are in Online Appendix B.6, where we prove the existence of sentiment equilibria.

We model capital investment with standard "q-theory" assumptions as in Hayashi (1982) Suppose that, when any individual is managing capital, his capital evolves as

$$\frac{dk_t}{k_t} = (g + \iota_t - \delta)dt + \sigma dZ_t^{(1)}, \tag{21}$$

where g is exogenous growth, ι_t is endogenous growth, and δ is the depreciation rate. Let $\Phi(\iota)$ be a convex adjustment cost function, so that $\Phi(\iota)k$ of investment expenditures creates ιk of new capital. For private agents, ι only affects the expected return on capital, so they choose investment optimally to maximize this expected return, i.e., $\max_\iota q\iota - \Phi(\iota)$, leading to the q-theory FOC

$$q_t = \Phi'(\iota_t) \tag{22}$$

In particular, all agents choose the same investment rate by matching its marginal cost to the traded capital price. Because $\Phi(\cdot)$ is convex, let the unique solution to (22) be $\iota_t = \iota(q_t) := (\Phi')^{-1}(q_t)$. Note that $\iota(\cdot)$ is strictly increasing.

Sunspot jumps enter the dynamics of the capital price. Consider a class of solutions taking the form

$$\frac{dq_t}{q_{t-}} = \mu_{q,t-}dt + \sigma_{q,t-} \cdot dZ_t - \ell_{q,t-}dJ_t,$$

where J is a Poisson process with intensity λ_t. We allow the jump intensity to be endogenous and time-varying. For simplicity, we restrict attention to equilibria where the jump size ℓ_q is pre-determined, in particular a function of (η, q, λ) and potentially other variables describing the state of the system, and we focus on adverse jumps with $\ell_q \geq 0$. A similar type of "exogenous liquidity shock" is studied in Krishnamurthy and Li (2024).

These modifications lead to the following changes. First, we now have an endogenous growth rate, with aggregate capital dynamics

$$\frac{dK_t}{K_t} = G(q_t)dt + \sigma dZ_t^{(1)}, \quad \text{where} \quad G(q) := g + \iota(q) - \delta. \tag{23}$$

This turns out to only impact the expression for interest rate r, with g replaced by $G(q)$.

Second, the goods market clearing condition now includes investment, which modifies the price-output relation from (PO) to

$$\bar{\rho}(\eta_t)q_t + \Phi(\iota(q_t)) = \kappa_t a_e + (1-\kappa_t)a_h. \tag{PO-inv}$$

Despite the presence of jumps, log utility agents still consume a constant fraction of wealth, explaining why aggregate consumption per unit of capital is still $\bar{\rho}q$. So long as $\Phi(\cdot)$ is increasing for the relevant set of equilibrium investment rates, we have $\Phi(\iota(\cdot))$ increasing. In that case, equation (PO-inv) is similar to (PO) in that both define increasing mappings from the capital distribution κ to its price q, holding fixed the wealth distri-

bution η. Thus, a core channel behind our results, that coordinated trade moves asset prices, remains qualitatively unaffected.

Third, the risk-balance condition (RB) is modified to read

$$0 = \min\left[1-\kappa,\ \frac{a_e - a_h}{q} - \frac{\kappa-\eta}{\eta(1-\eta)}\left(|\sigma_R|^2 + \frac{\lambda \ell_q^2}{\left(1-\frac{\kappa}{\eta}\ell_q\right)\left(1-\frac{1-\kappa}{1-\eta}\ell_q\right)}\right)\right]. \tag{RBJ}$$

The additional terms involving ℓ_q arise because there is a jump risk premium. By adding a new source of risk, we have an additional degree of freedom. The risk-balance condition disciplines overall risk—the term in parentheses of (RBJ) is pinned down given (η, q)—but the split between the Brownian and Poisson shocks is indeterminate. We thus have tremendous flexibility in our choice of ℓ_q.

To complete the construction of equilibrium, recall that we must ensure that (η, q) remain in the equilibrium domain $\mathcal{D}$ (i.e., the shaded region in Figure 3). But this is easy: just set $\ell_q = 0$ near the boundaries of $\mathcal{D}$. Doing this, the stability analysis remains essentially unchanged from Theorem 1, since near the boundaries the economy behaves as if it is only hit by Brownian shocks. The only other consideration is that a price jump cannot be so large as to exit the equilibrium domain $\mathcal{D}$, i.e., price crashes cannot send experts into bankruptcy or induce such large fire sales that households become the levered entity. We thus prove a generalization of Theorem 1.

Theorem 2. *Let Assumption 1 hold, and suppose Φ is strictly increasing and strictly convex. Then, there exists a family of sentiment equilibria in which $(\eta_t, q_t)_{t\geq 0}$ remains in $\mathcal{D}$ almost-surely and possesses a non-degenerate stationary distribution.*

The jump arrival rate and jump size are also indeterminate in the following sense. Let $(y_t)_{t\geq 0}$ be any sufficiently well-behaved exogenous, stationary, Markov process. Let $\mathcal{D}^\circ \subset \mathcal{D}$ be any subdomain such that κ is bounded away from η and η is bounded away from 0. Let $l(\eta, q, y) \geq 0$ and $h(\eta, q, y) \in [0, \ell_q^{max}(\eta, q, y))$ be any functions, where $\ell_q^{max}(\eta, q, y) > 0$ is defined in equation (B.21). Then, a sentiment equilibrium exists in which

(i) on $\mathcal{D}^\circ$, the jump arrival rate λ is $l(\eta, q, y)$

(ii) on $\mathcal{D}^\circ$, the jump size ℓ_q is $h(\eta, q, y)$

Remark 2 (General preferences)**.** *One interesting generalization would be to solve for sentiment equilibria with non-log preferences. In particular, whereas log preferences are used frequently in this literature, many papers have studied other preference specifications.*[20] *While studying non-log preferences is beyond the scope of this paper, Online Appendix F.3 does put*

[20]Note that the baseline model of Brunnermeier and Sannikov (2014) features linear preferences and a

forth the set of equilibrium equations for a recursive Epstein-Zin preference version of our economy. We show that the same two mechanisms driving our results—namely, price-output and risk-balance conditions—continue to survive. We also show that the equations admit the same degrees of freedom as our log preference economy, suggesting that in principle a sentiment equilibrium construction should be possible. That said, we leave a full construction for future research.

2.4 Discrete-time version

Sentiment equilibria are also possible in a discrete-time version of the model. The key details are identical to the baseline model: agents have log utility, experts are more productive than households, and equity issuance is restricted. We also only consider the special case where aggregate capital does not grow nor have fundamental risk (i.e., $g = \sigma = 0$). The main difference is that agents make choices at the sequence of discrete times $t = 0, \Delta, 2\Delta, \ldots$ More details on this discrete-time environment and its solution are contained in Online Appendix G.

We derive the following four-equation system that fully characterizes the equilibrium variables (η, q, κ, r). First, we obtain a price-output relation, exactly identical to (PO):

$$q_t \bar{\rho}(\eta_t) = \kappa_t a_e + (1 - \kappa_t) a_h. \tag{24}$$

Second, each agent's capital Euler equation holds, which after aggregating yields

$$0 = \mathbb{E}_t \left[\frac{q_{t+\Delta} + a_e \Delta - (1 + r_t \Delta) q_t}{\frac{\kappa_t}{\eta_t} \left(q_{t+\Delta} + a_e \Delta - (1 + r_t \Delta) q_t \right) + (1 + r_t \Delta) q_t} \right] \tag{25}$$

$$0 \geq \mathbb{E}_t \left[\frac{q_{t+\Delta} + a_h \Delta - (1 + r_t \Delta) q_t}{\frac{1-\kappa_t}{1-\eta_t} \left(q_{t+\Delta} + a_h \Delta - (1 + r_t \Delta) q_t \right) + (1 + r_t \Delta) q_t} \right] \tag{26}$$

In expressions (25)-(26), the numerators are each agent's expected excess capital returns multiplied by q_t; the denominators are each agent's net worths at $t + \Delta$, which accounts

non-negativity constraint on consumption, but they also examine a version with log preferences in Section III (see Proposition III.3 and Figure 5). Since then, several studies in this literature have studied log preferences—see, e.g., Adrian and Boyarchenko (2015), Khorrami (2021), Mendo (2020), and Krishnamurthy and Li (2024). Some others have adopted the baseline linear preference specification of Brunnermeier and Sannikov (2014)—see, e.g., Phelan (2016) and Moreira and Savov (2017). Still others have a myopic preference specification for experts, and nevertheless obtain similar macro dynamics—see, e.g., He and Krishnamurthy (2019) and Maxted (2024). Brunnermeier and Sannikov (2016) generalized the model to any CRRA utility, while Hansen et al. (2024) generalized the model to Epstein-Zin preferences (potentially asymmetric across agents). The illustrations in these papers show that the model dynamics are very similar qualitatively across a wide range of preference assumptions: fire sales occur when expert wealth is low, which then causes output declines, asset price declines, and risk premium increases.

for their marginal utility weighting. Finally, from agent's dynamic budget constraints, we obtain the wealth share law of motion:

$$\eta_{t+\Delta} = \frac{1}{1+\rho_e \Delta}\Big(\frac{\kappa_t\big(\frac{a_e\Delta + q_{t+\Delta} - q_t}{q_t} - r_t\Delta\big) + \eta_t(1+r_t\Delta)}{q_{t+\Delta}/q_t}\Big). \tag{27}$$

If we can find (η, q, κ, r) that solve equations (24)-(27) at all times, we have an equilibrium.

Since there is no fundamental risk, there is always a *fundamental equilibrium* in which $\kappa_t = 1$ at all times. But we also can construct a *non-fundamental equilibrium* in which $\kappa_t < 1$ sometimes: a self-fulfilling fire sale. Furthermore, it becomes clear from the construction that the non-fundamental equilibrium is not Markovian in η, analogous to our sentiment equilibria in continuous time.

Our construction assumes binomial uncertainty:

$$q_{t+\Delta} = \begin{cases} u_t q_t, & \text{with probability } 1-\pi_t; \\ d_t q_t, & \text{with probability } \pi_t. \end{cases} \tag{28}$$

The "up" and "down" returns u_t and $d_t \in (0, u_t)$ may be state dependent, as may the probability of a price drop π_t. We will take the "state space" to be the set of possible (η_t, q_t). In other words, (u_t, d_t, π_t, r_t) will be functions of (η_t, q_t). The goal is to find such functions ensuring that (η_t, q_t) remains in the equilibrium domain $\mathcal{D}$ forever. We do this for any Δ sufficiently small in Proposition G.1. In the "interior" (i.e., when κ is sufficiently far from η or 1), we choose (u, d, π) so that $\frac{q_{t+\Delta}-q_t}{\Delta}$ exactly converges to our Brownian dynamics as $\Delta \to 0$. The hard step, as in Theorem 1, is addressing the boundaries. Unlike the continuous-time model, we are unable to provide a general method for this, but we do provide a specific construction that converges to reflecting barriers as $\Delta \to 0$.

3 Resolving puzzles with sentiment

We have just demonstrated that sunspot equilibria, which are endemic to this class of models, in principle support rich dynamics. Now, we solve some concrete examples to illustrate several substantive results. We show how our framework can help resolve the two puzzles outlined in the introduction and shown in Figures 1-2: (i) the severity and suddenness of financial crises; and (ii) pre-crisis frothiness in asset markets. As a comparison, we show how the economy without sentiment performs poorly on these dimensions. Moreover, the qualitative differences in crisis dynamics between our sentiment

equilibria and the fundamental equilibrium hold across a wide variety of specifications for the indeterminate objects. After that, we calibrate our model and compare it to crisis data. We show this framework can—without any information friction, non-rational belief updating, or additional bells and whistles considered in the literature—quantitatively replicate various aspects of pre- and post-crisis dynamics.

3.1 Example economies

We construct example economies to compare our model dynamics to the data. To do this, we must take a stand on a few variables that are not pinned down uniquely. Such flexibility is always a concern in models with multiple equilibria. Our general strategy is to, where possible, make choices that are either minimal or constrained by data. In cases where a data counterpart is unavailable, we perform extensive sensitivity analyses.

We follow the construction outlined prior to Theorem 1, so that the state vector includes (η, q), the expert wealth share and the capital price. This construction pins down the price volatility uniquely ($|\sigma_q|$), but leaves open the source of volatility ($\sigma_q^{(1)}$ versus $\sigma_q^{(2)}$), the price drift (μ_q away from the boundaries of $\mathcal{D}$), and the exit rate from the efficient region. We make the following choices for these objects.

1. Drift in the interior. For the drift μ_q, we follow Example 2. In the interior of $\mathcal{D}$, let r_t follow the exogenous process

$$dr_t = \lambda_r(\bar{r} - r_t)dt + \sigma_r \begin{pmatrix} \theta \\ \sqrt{1-\theta^2} \end{pmatrix} \cdot dZ_t, \quad \text{if} \quad (\eta_t, q_t) \in \text{int}(\mathcal{D}). \tag{29}$$

By forcing r to follow (29), we effectively pick μ_q by equation (11). An advantage of this approach is that we can use data on the interest rate to calibrate the parameters $(\bar{r}, \lambda_r, \sigma_r, \theta)$. We make one of two choices here:

(i) Our baseline sets $\sigma_r = 0$. In this case, $r_t \to \bar{r}$ deterministically at rate λ_r when $\kappa_t < 1$. By minimizing risk-free rate volatility, we can show that our model dynamics do not rely on specific types of time-variation in r or μ_q.

(ii) Alternatively, we pick $(\bar{r}, \lambda_r, |\sigma_r|)$ to match interest rate data. We pick these parameters to match the unconditional mean (0.014), variance (0.023^2), and annual autocorrelation (0.94) of the 3-month US real rate. We perform sensitivity analyses on θ and $\text{sign}(\sigma_r)$, i.e., to what extent interest rates are driven by fundamental shocks and whether real rates are pro- or counter-cyclical.

2. Drift at the boundaries. Near the lower boundary $q^L(\eta)$, we must make a choice for "bounce-back beliefs." We follow Example 1 in assuming a reflecting barrier $\underline{q}(\eta)$, such that $q \geq \underline{q}(\eta)$ always. We set this reflecting barrier by $\underline{q}(\eta) := q^L(\eta) + \underline{\kappa}(1-\eta)\frac{a_e - a_h}{\bar{\rho}(\eta)}$, where $\underline{\kappa}$ is a free parameter. This barrier keeps $\kappa_t \geq \underline{\kappa}$ for all t. Numerically, our baseline sets $\underline{\kappa} = 0.01$, so that the barrier is just above q^L. This choice is, in some sense, a minimal one, because it affects the dynamics infrequently, only in the very extreme states.

The economy is allowed to visit the upper boundary $q = q^H(\eta)$, because this is where capital is efficiently allocated ($\kappa = 1$). However, the expected time until the economy exits the efficient region is indeterminate. As a baseline, we suppose an exit rate such that the economy is efficient approximately 10% of the time. Appendix C.1 explains the exit rate construction. We also perform sensitivity analysis on this exit rate.

3. Source of volatility. Let ϑ be the fraction of return variance from the fundamental shock, i.e., $(\sigma_R^{(1)})^2 = \vartheta|\sigma_R|^2$. Our baseline assumes $\vartheta = 0.5$ in the interior of $\mathcal{D}$ so that the fundamental and sunspot shocks contribute equally to return volatility. We also perform sensitivity on ϑ.

Comparison to Fundamental Equilibrium (FE). We compare the dynamics from an S-BSE to those in a FE. Whenever we make such a comparison, we hold all deep parameters fixed (i.e., $a_e, a_h, \rho_e, \rho_h, \delta_e, \delta_h, \sigma$). However, recall a FE constrains all objects like q, κ, r, etc., to be functions of η, and so this equilibrium lacks the degrees of freedom discussed above for the S-BSE. Online Appendix E provides more detail on the solution to the FE.

Comparison to higher bounce-back beliefs. We also explore an S-BSE where bounce-back beliefs kick in significantly earlier. In particular, continue to assume a reflecting lower barrier $\underline{q}(\eta) = q^L(\eta) + \underline{\kappa}(1-\eta)\frac{a_e - a_h}{\bar{\rho}(\eta)}$, but increase its level to $\underline{\kappa} = 0.6$. Effectively, this prevents fire sales from being too severe (i.e., $\kappa_t \geq 0.6$ for all t).

Enriched model with investment and jumps. Finally, Section 3.3 performs a more quantitative exercise using the enriched model with investment and jumps. We defer these details until that section.

3.2 Financial crisis event studies in the baseline model

We construct model-implied event studies, analogously to the data versions in Figures 1-2. As a baseline definition, a "crisis" is defined as the bottom 3rd percentile of year-to-year log output declines in our simulation. Figure 6 presents the event study results

for the capital price q, the expert wealth share η, the weighted-average risk premium $\kappa\mu_{R,e} + (1-\kappa)\mu_{R,h} - r$, and return volatility $|\sigma_R|$.

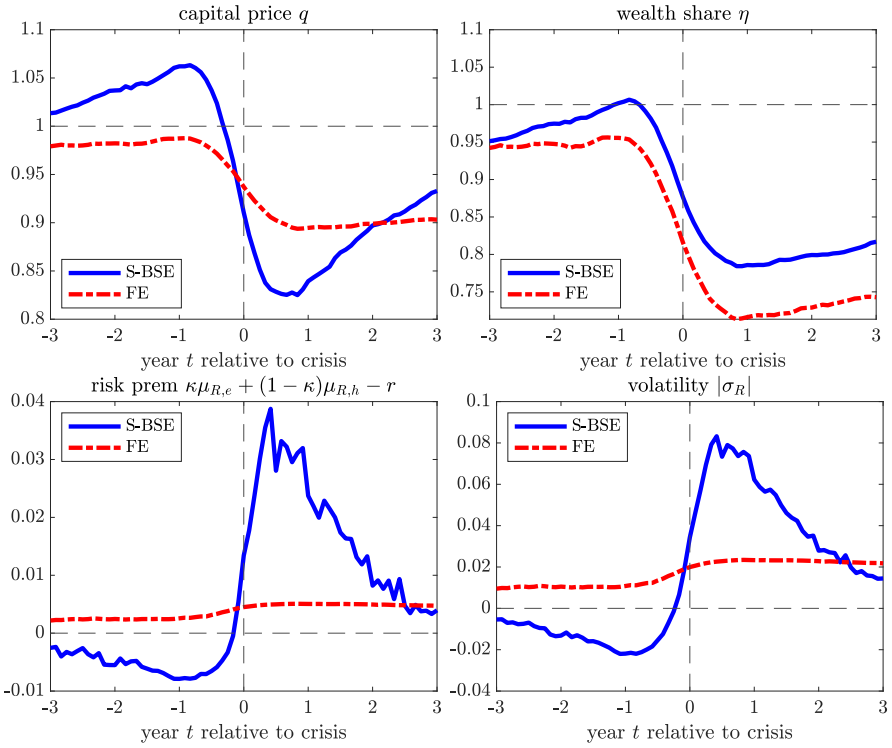


Figure 6: Event studies around financial crises. Crises are defined as the bottom 3rd percentile of year-to-year log output declines, subject to only one crisis occurring in a 6-year window. Data is generated via a 10,000 year simulation at the daily frequency, with the outcomes above then averaged to the monthly level. The solid blue line is the mean path from the S-BSE, while the dot-dashed red line is the mean path from the FE. All variables are re-normalized around their unconditional average (the thin horizontal line). The paths in the top panels (q and η) are rescaled by their $t = 0$ values to be in units of percentage changes, and then shifted so that 1 represents the unconditional mean. The paths in the bottom panels ($\kappa\mu_{R,e} + (1-\kappa)\mu_{R,h} - r$ and $|\sigma_R|$) are in raw units and plotted as deviations from their historical mean. Parameters: $\rho_e = 0.02$, $\rho_h = 0.015$, $a_e = 0.11$, $a_h = 0.03$, $\sigma = 0.05$, $g = 0.01$. Type-switching parameters: $\delta_h = 0.01$ and $\delta_e = 0.015$. Risk-free rate parameters: r_t follows (29) with parameters $\bar{r} = 0.014$, $\lambda_r = -\log(0.94)$, and $\sigma_r = 0$. Variance share parameter: $\vartheta = 0.5$ is the fraction of return variance $|\sigma_R|^2$ due to the fundamental shock (in the interior of $\mathcal{D}$).

The S-BSE (solid blue lines) produces severe, sudden crises with pre-crisis froth. The suddenness and severity are captured by the large swings in all variables at crisis onset ($t = 0$). On average, asset prices q and expert wealth η both drop about 20% in the year of crisis. Conversely, the risk premium $\kappa\mu_{R,e} + (1-\kappa)\mu_{R,h} - r$ and volatility $|\sigma_R|$ rise on average by about 5% and 10%, respectively, in this same year. While our model is particularly simple, these magnitudes are in the ballpark of empirical patterns. Section 3.3 provides a more formal quantitative analysis. Pre-crisis frothiness is captured by the fact that conditions are *better-than-average* and *improving* in the years leading up to crisis. The capital price is above-average and rising; expert's equity is near-average and

also rising; the risk premium and volatility are both below-average and falling pre-crisis. Warning signs of an impending crisis are not visible anywhere.

None of these patterns arise in the standard Fundamental Equilibrium (FE). Prices fall significantly less and in a gradual fashion. Risk premia and volatility rise by an order of magnitude less in the FE than in the S-BSE. This is all despite the fact that expert wealth η declines significantly in the FE. Finally, there is an absence of pre-crisis froth: all pre-crisis indicators are worse-than-average (q and η are below-average, while $\kappa\mu_{R,e} + (1-\kappa)\mu_{R,h} - r$ and $|\sigma_R|$ are above-average).[21]

Importantly, sentiment-driven crises display significant declines in expert wealth η, as happens empirically for banks (Baron et al., 2021). One question is how this finding emerges given our emphasis throughout the paper on the "decoupling" of dynamics from η. To reconcile decoupling with the fact that η crashes in crisis, recall the shock exposure $\sigma_\eta = (\kappa - \eta)\sigma_R$. Experts' balance sheets are always disproportionately exposed to capital return shocks. If a non-fundamental fire sale emerges, for instance due to a "fear" shock, the decline in q causes η to fall as well.

Why can the S-BSE produce dynamics so different from the FE? What is important is that the economy visit extreme states having high volatility and low asset prices. The stationary distribution of the S-BSE, plotted in Figure 7, shows how the economy spends most of its time below the FE (the solid red line), and potentially far below. This is the essential aspect of "decoupling" that matters to deliver the key results. In extensive sensitivity analyses (discussed below), we show that the crisis event study results continue to hold in any economy that has this property

Realistic crisis dynamics seem to depend on reaching extreme values of volatility and risk premia. To confirm this hypothesis, we now implement the alternative bounce-back belief that prevents extreme fire sales. Specifically, we compute an S-BSE where the bounce-back belief is the lower reflecting barrier $\underline{q}(\eta) := (1-\eta)\frac{\underline{\kappa}a_e + (1-\underline{\kappa})a_h}{\rho_h} + \eta\frac{a_e}{\rho_e}$, with $\underline{\kappa} = 0.6$. To visualize how this truncates the equilibrium state space, see the stationary distribution in Figure 8. The S-BSE now resides frequently *above* the FE, a region with higher asset prices and lower volatility. Consequently, the financial crises are much tamer in this equilibrium; now, the FE and S-BSE are not so different, as the event studies in Figure 9 show. Preventing the economy from reaching extreme states rules out severe crises with volatility and risk premium spikes. Online Appendix C.1 shows the CDFs for various equilibrium objects in these two equilibria, showing that the higher bounce-back

[21]The comparison of the FE to S-BSE holds all deep model parameters fixed. A natural question is whether there exist *any parameters* such that the FE generates realistic crisis event studies. In unreported results, we have experimented and found no such parameters.

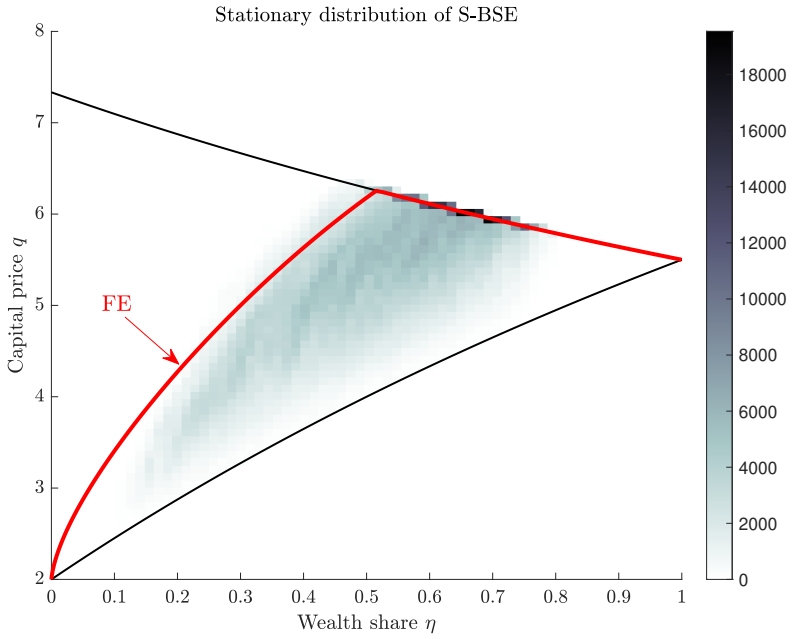


Figure 7: Stationary distribution of (η, q) in the S-BSE. Darker colors represent higher relative frequencies in the histogram. Parameters are the same as Figure 6.

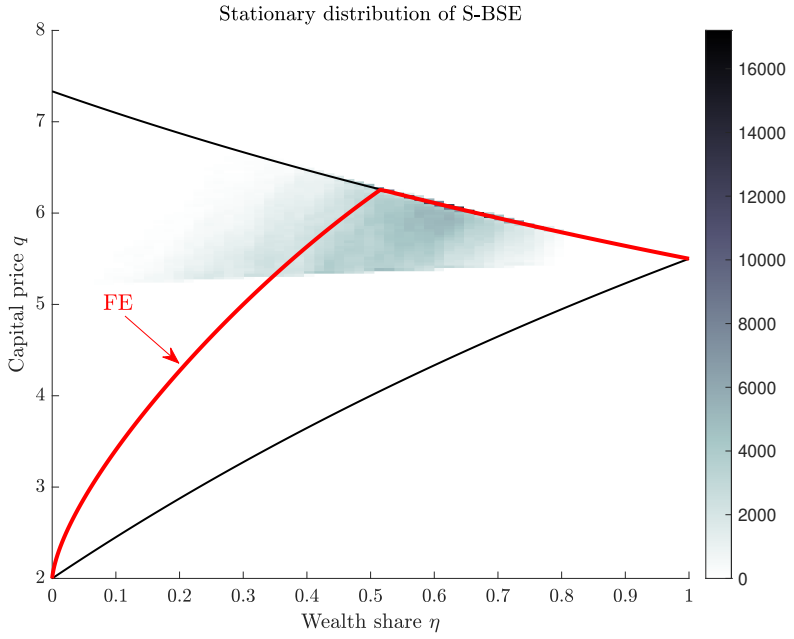


Figure 8: Stationary distribution of (η, q) in the S-BSE *with higher bounce-back beliefs*. Darker colors represent higher relative frequencies in the histogram. The bounce-back belief is a lower reflecting barrier at $q^L(\eta) + 0.6(1-\eta)\frac{a_e - a_h}{\bar{\rho}(\eta)}$. Parameters are otherwise the same as Figure 6.

beliefs eliminate the right tail in volatility and risk premia.

We perform several robustness exercises in Online Appendix C.2. There, we show that the critical success of the S-BSE in reproducing the salient features of a crisis persists

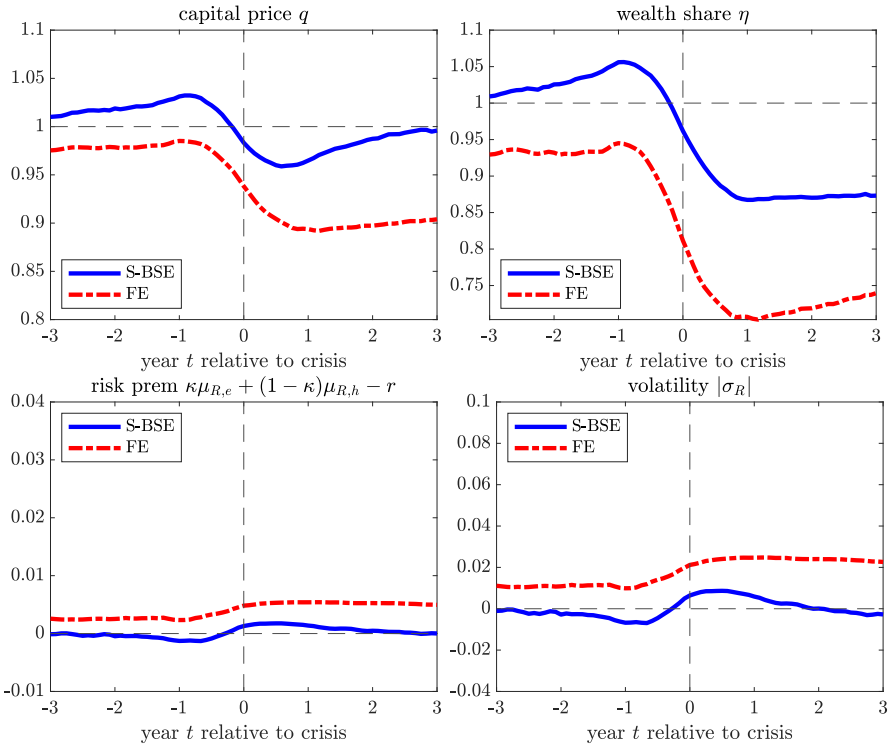


Figure 9: Event studies *with higher bounce-back beliefs*. Crises are defined as the bottom 3rd percentile of year-to-year log output declines, subject to only one crisis occurring in a 6-year window. The bounce-back belief is a lower reflecting barrier at $q^L(\eta) + 0.6(1-\eta)\frac{a_e - a_h}{\bar{\rho}(\eta)}$. All other details are identical to Figure 6.

under several alternatives. First, we define a financial crisis based on large drops in η_t rather than output, more similar to Baron et al. (2021). Second, we alter the calibration of the process for r_t (i.e., we allow a more volatile interest rate, allow it to be either pro- or counter-cyclical, and allow it to be driven by either the fundamental or sunspot shock). Third, we shut down the sunspot shock altogether, so that all dynamics are driven by fundamental shocks. Fourth, we re-calibrate to a slower exit rate from the efficient region $\kappa = 1$, so that the efficient capital allocation emerges more often. None of these lead to significantly different results. Critically, all of these alternative S-BSEs permit volatility and risk premia to reach extreme values that the FE does not allow.

3.3 Calibrated quantitative example with jumps and investment

We now study the generalized version of the model discussed in Section 2.3, which includes capital investment and the possibility of sunspot jump shocks. We use this enriched setting to match some empirical targets. While the goals of this paper are not primarily quantitative, we view the exercise in this section as a "possibility result": we can design a sentiment process to match a litany of crisis and non-crisis moments that

are often viewed jointly as puzzling.

Model specification. We adopt quadratic adjustment costs $\Phi(\iota) = \iota - \gamma + \frac{\chi}{2}(\iota - \delta)^2$. In that case, the optimal investment decision satisfying (22) is $\iota(q) = \delta + \frac{q-1}{\chi}$.

For jumps, the two key objects are the jump arrival rate process λ_t and the jump size $\ell_{q,t}$. We assume λ_t passes through three states: *normal*, *quiet*, and *panic*. In the *normal* and *quiet* states, the arrival rate takes constant values of λ^{normal} and λ^{quiet}, respectively. In the *panic* state, λ_t takes a value in the interval $[\lambda_L^{panic}, \lambda_H^{panic}]$, depending on the capital price drop generated when the process entered the panic state. In particular, we assign λ_L^{panic} to the lowest possible jump triggering a panic and λ_H^{panic} to the highest possible such jump, linearly interpolating in between. Then, λ_t remains fixed during all the span of each visit to the panic state. The transitions between normal, quiet, and panic states occur at constant rates, independently of all other random variables, with one exception: the transition from *quiet to panic* happens whenever the capital price falls in the quiet state due to a sunspot jump shock.

The specification of jump sizes is described mathematically in Online Appendix C.3 but is easily understood verbally. For each (η, q, λ), there is a maximum theoretical jump size $\ell_q^{max}(\eta, q, \lambda)$. Beyond that constraint, the model permits tremendous freedom. We specify jumps ℓ_q to be some exogenous fraction of this maximum jump size, assuming the result is bigger than a minimum level, and assuming κ is above some minimal level. This jump fraction, the minimal jump size, and the minimal κ are all permitted to depend on whether the latent state is *normal*, *quiet*, or *panic*.

Finally, we introduce two assets that proxy for liquidity and credit risk, in order to speak more closely to existing empirical studies. Both assets are assumed to be in zero net supply and only traded by experts, and so they do not alter any aspect of equilibrium.

The first asset is a *liquidity asset*. It loses fraction α of its value when a "liquidity event" is realized, defined as a jump of $\ell_{q,t-}dJ_t > 0$ that triggers a drop in capital price. In equilibrium, the premium such an assets pays over the risk free rate would be

$$LiqPrem = 1_{\{\ell_q > 0\}} \frac{\lambda \alpha}{1 - \frac{\kappa}{\eta}\ell_q} \tag{30}$$

We calibrate the loss to $\alpha = 0.05$, which corresponds to the effective liquidation loss in the data of the 2008 crisis (Krishnamurthy and Li, 2024).

The second asset is a *credit asset*. Conditional on a "liquidity event", this asset "defaults" with exogenous probability π on a large fraction of its value (incurring a loss-given-default of $m_0 + m_1\ell_q$) and experiences a small decline otherwise (losing m_2 fraction

of its value). In equilibrium, the spread this asset pays over the risk-free rate is

$$CredSpread = 1_{\{\ell_q>0\}} \frac{\lambda}{1-\frac{\kappa}{\eta}\ell_q}\Big[\pi(m_0+m_1\ell_q)+(1-\pi)m_2\Big] \tag{31}$$

In Online Appendix C.3, we calibrate the parameters π, m_0, m_1, and m_2 to various data on average default rates, average loss-given-default, and the 2008 financial crisis loss-given-default data in the US, following Krishnamurthy and Li (2024).

Calibration. The full model calibration is contained in Online Appendix C.3. There, we also show how the model performs on various unconditional macro-financial moments in the US, including data on consumption, investment, output, liquidity premia, stock returns, bank equity returns, and bank leverage. While our calibration features excess consumption and output volatility, relative to the data, this is needed to generate more realistic outcomes for stock returns, which we match relatively well. For example, the average return and return volatility of financial stocks, proxied by experts' equity, are 10.5% and 39% in our model, respectively, compared to 8.8% and 37% in the data. Because the model matches return data relatively well, it also matches investment growth and volatility well.

The most important tests of our model relate to pre- and post-crisis data on GDP, credit-to-GDP, and credit spreads, as well as the predictability of crises. Quantitatively, the model matches these dynamics quite well (Tables C.4-C.5). To give the reader a sense of the magnitudes for these outcomes, Figure 10 plots crisis event studies for various objects of interest. Visible in the figures are the severe downturn in economic activity, sharply negative returns, and spikes in risk premia, volatility, and credit spreads. In addition, we can see evidence of "frothy" behavior pre-crisis, as we elaborate on below.

The top left panel shows GDP growing faster than average in the years prior to a crisis, with a sharp drop upon the arrival of a crisis, and a slow recovery afterward. The 11% drop in the crisis year is close to the 9.3% observed empirically (Reinhart and Rogoff, 2009, p.230). Model-implied GDP is 7% below trend 2-3 years post-crisis, compared to 6.5% empirically (from averages of results in Krishnamurthy and Muir, 2024; Sufi and Taylor, 2022; Schularick and Taylor, 2012). A similar pattern is seen in the credit-to-GDP ratio, defining credit as experts' liabilities, i.e., $Credit = (\kappa-\eta)qK$, following Krishnamurthy and Li (2024). (This credit-to-GDP series is plotted as standard deviations away from its non-crisis mean, and its behavior is sharper than in the data, e.g., Baron et al., 2021, likely because our model measure is purely short-term debt.)

The third panel illustrates the behavior of financial and non-financial (log) returns.

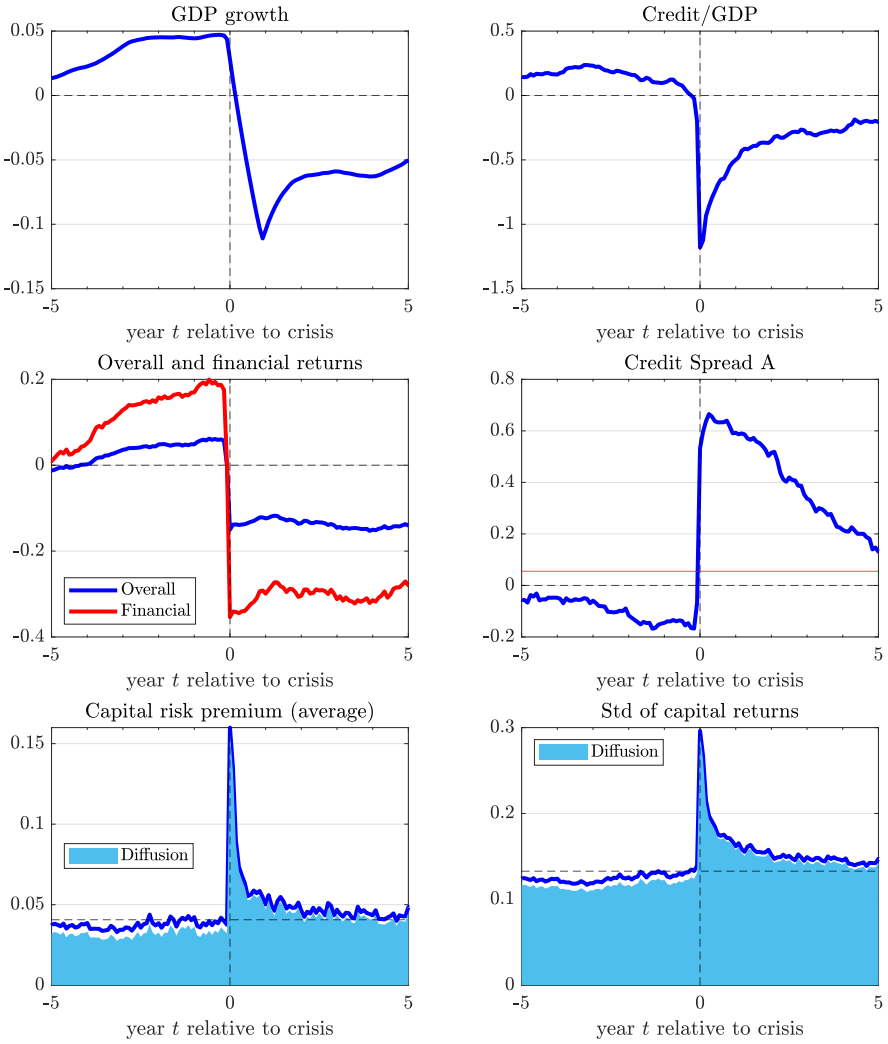


Figure 10: Crisis event studies. Crises are defined as the bottom 3.5th percentile of month-to-month log output declines, conditional on no other crisis in the previous 7 years. In all panels, the dotted line represents the unconditional mean (outside of crisis window observations). The calibration is in Table C.1.

In the model, we proxy financial returns with the return on expert's equity and non-financial return with the return on capital (using the capital-weighted average of dividend yield between experts and households). We plot these returns as deviations from their unconditional averages outside of crisis windows, and normalize the resulting difference to equal zero in the month before crisis. The boom-bust pattern is clear, with large positive pre-crisis returns, followed by a -50% return at crisis onset, and a long-lasting depression (minimal recovery after 5 years). This closely matches data: according to Baron et al. (2021), replicated in our Figure 1, the average financial crisis features a -46.2% peak-to-trough bank equity decline, with no bank equity rebound after 5 years. Furthermore, our event studies show a financial equity decline that fully materializes even as GDP just begins to drop, in line with their empirical findings.

The fourth panel illustrates the behavior of the credit spread. (Not displayed, the liquidity premium measure displays similar behavior.) The 1-year spread increase around crisis is 0.73 standard deviations in the model, compared to between 0.48 and 0.68 in the data (Krishnamurthy and Muir, 2024, Table VII). Some pre-crisis frothiness of credit spreads is present, but slightly more modest than in the data; as one example, in the two years before crisis, spreads are -0.14 standard deviations below their mean in the model, with the corresponding numbers between -0.15 and -0.43 in the data (Krishnamurthy and Muir, 2024, Table VII). That said, the pre-crisis froth is a robust phenomenon, and importantly *froth predicts crises*. For example, using below-median credit spreads as the predictor variable, the model implies a 17% higher probability of crisis in the next 5 years, with a corresponding number being 21% in the data (Krishnamurthy and Muir, 2024, Table VII). Crisis predictability rises if both high credit and low spreads happen at the same time, both in the model and data (Schularick and Taylor, 2012; Baron and Xiong, 2017; Krishnamurthy and Muir, 2024). After the crisis, spreads tend to mean-revert at a rate in line with the data (half-life of 3.1 years in the model versus between 2.5-3.5 years in the data, according to Muir, 2017, and Krishnamurthy and Muir, 2024).

The fifth and sixth panels display the model-implied local risk premium (again using the capital-weighted average of expert and household dividends) and capital return volatility (defined as $\sqrt{|\sigma_R|^2 + \lambda \ell_q^2}$). The shaded areas correspond to the risk premium and volatility, respectively, associated to Brownian risk. There is a very large and sudden spike in these objects upon crisis—the risk premium rises to exceed 15% per annum, while the volatility rises from 13% to 30% per annum. This is about half the magnitude of the volatility spike in Figure 2. Pre-crisis froth is also evident, with both measures slightly below their mean before crises. Paymon: Is it like the others in that the crisis window is excluded for the dashed line mean?

4 Conclusion

We have shown that macroeconomic models with financial frictions may inherently permit sunspot volatility. The types of models we study are extremely common in macroeconomics, so this phenomenon cannot be ignored.

On the bright side, our paper demonstrates how a fully-rational notion of "sentiments" can be a powerful input into macro-finance dynamics. Time-varying uncertainty drives all dynamics in our sentiment-driven fluctuations. Sharp volatility spikes and belief-driven boom-bust cycles are among the many interesting possibilities raised by our framework. While ours is not a full-blown quantitative analysis, we show that rational sentiment can bring the model closer to data on these dimensions.

On the hazier side, our results suggest a modicum of caution. Many researchers employ numerical techniques to solve and analyze DSGE models that are built upon the core frictions in our paper—these procedures implicitly select an equilibrium, without any explicit justification. By demonstrating the necessary assumptions for our self-fulfilling fluctuations (heterogeneous productivity and incomplete aggregate risk-sharing), we pinpoint subsets of the literature which are subject to the multiplicity we identify.[22] A deeper analysis of refinements, perhaps leveraging global-games approaches or adaptive learning, still remains to be done.

What about policy? Caveated by the need for further study on refinements, we can offer some initial thoughts. Some traditional policies become less effective in sunspot equilibria. For example, deposit insurance has less bite because run-like behavior can occur solely due to fire-sale coordination, i.e., on the asset side rather than the liability side. Sunspot equilibria also decouple financial crises from bank balance sheets and

[22] We can provide a "taxonomy" of the literature into papers which would be subject to our sentiment equilibria, and those likely not. Models which include both heterogeneous productivity (hence, the possibility of fire sales) and the incomplete aggregate risk-sharing would be subject to our sentiment equilibria. These include Brunnermeier and Sannikov (2014), Brunnermeier and Sannikov (2016), Adrian and Boyarchenko (2015), Phelan (2016), Moreira and Savov (2017) [this paper features two types of capital, one of which is less productive, which is a stand-in for a fire-sale to a less productive agent], Krishnamurthy and Li (2024), Hansen et al. (2024), Mendo (2020), Khorrami (2021). Several papers additionally impose some type of borrowing or collateral constraint a la Kiyotaki and Moore (1997). A few recent examples include Gertler and Kiyotaki (2015) and Gertler et al. (2020). Because these papers feature the core assumptions of heterogeneous productivity and restricted risk-sharing, we think the same multiplicity plagues them. Although these latter papers are written in discrete time, this alone should not affect the existence of a sentiment equilibrium; in particular, we have shown in Appendix G that similar sentiment equilibria exist in discrete time. Finally, we should make special note of two classes of papers that would likely not be subject to our sentiment equilibria. First, the models in Di Tella (2017) and Di Tella (2019) include frictionless derivative markets for hedging all aggregate shocks, which means that aggregate risks can (in principle) be shared perfectly. Second, the intermediary asset-pricing models of He and Krishnamurthy (2012, 2013) do not feature any real effects. Paymon: We can consider moving this elsewhere if needed.

wealth, which defangs capital requirements, bailouts, and the like. On the other hand, policies that manipulate beliefs can be effective. Future research might better explain which policy designs have the power to manipulate beliefs in this way. Given the framework we study relies on fire sales, asset purchases (or future commitments to them) are one interesting candidate.

References

Tobias Adrian and Nina Boyarchenko. Intermediary leverage cycles and financial stability. *Working Paper*, 2015.

S Rao Aiyagari. Uninsured idiosyncratic risk and aggregate saving. *The Quarterly Journal of Economics*, 109(3):659–684, 1994.

Costas Azariadis. Self-fulfilling prophecies. *Journal of Economic Theory*, 25(3):380–396, 1981.

Costas Azariadis and Allan Drazen. Threshold externalities in economic development. *The Quarterly Journal of Economics*, 105(2):501–526, 1990.

Costas Azariadis, Leo Kaas, and Yi Wen. Self-fulfilling credit cycles. *The Review of Economic Studies*, 83(4):1364–1405, 2016.

Philippe Bacchetta, Cédric Tille, and Eric Van Wincoop. Self-fulfilling risk panics. *The American Economic Review*, 102(7):3674–3700, 2012.

Nicholas Barberis, Robin Greenwood, Lawrence Jin, and Andrei Shleifer. X-capm: An extrapolative capital asset pricing model. *Journal of Financial Economics*, 115(1):1–24, 2015.

Matthew Baron and Wei Xiong. Credit expansion and neglected crash risk. *The Quarterly Journal of Economics*, 132(2):713–764, 2017.

Matthew Baron, Emil Verner, and Wei Xiong. Banking crises without panics. *The Quarterly Journal of Economics*, 136(1):51–113, 2021.

Jess Benhabib and Pengfei Wang. Private information and sunspots in sequential asset markets. *Journal of Economic Theory*, 158:558–584, 2015.

Jess Benhabib, Pengfei Wang, and Yi Wen. Sentiments and aggregate demand fluctuations. *Econometrica*, 83(2):549–585, 2015.

Jess Benhabib, Xuewen Liu, and Pengfei Wang. Sentiments, financial markets, and macroeconomic fluctuations. *Journal of Financial Economics*, 120(2):420–443, 2016.

Jess Benhabib, Xuewen Liu, and Pengfei Wang. Financial markets, the real economy, and self-fulfilling uncertainties. *The Journal of Finance*, 74(3):1503–1557, 2019.

Jess Benhabib, Xuewen Liu, and Pengfei Wang. Self-fulfilling risk panics: An expected utility framework. Unpublished working paper, 2020.

Jaroslav Borovička. Survival and long-run dynamics with heterogeneous beliefs under recursive preferences. *Journal of Political Economy*, 128(1):206–251, 2020.

Markus K Brunnermeier and Yuliy Sannikov. A macroeconomic model with a financial sector. *The American Economic Review*, 104(2):379–421, 2014.

Markus K Brunnermeier and Yuliy Sannikov. Macro, money, and finance: A continuous-time approach. *Handbook of Macroeconomics*, 2:1497–1545, 2016.

Ricardo J Caballero and Alp Simsek. A risk-centric model of demand recessions and speculation. *The Quarterly Journal of Economics*, 135(3):1493–1566, 2020.

David Cass and Karl Shell. Do sunspots matter? *Journal of Political Economy*, 91(2): 193–227, 1983.

Pierre Andre Chiappori and Roger Guesnerie. Sunspot equilibria in sequential markets models. *Handbook of Mathematical Economics*, 4:1683–1762, 1991.

Adrien d'Avernas and Quentin Vandeweyer. Treasury bill shortages and the pricing of short-term assets. *The Journal of Finance*, 2023.

Sebastian Di Tella. Uncertainty shocks and balance sheet recessions. *Journal of Political Economy*, 125(6):2038–2081, 2017.

Sebastian Di Tella. Optimal regulation of financial intermediaries. *The American Economic Review*, 109(1):271–313, 2019.

Itamar Drechsler, Alexi Savov, and Philipp Schnabl. A model of monetary policy and risk premia. *The Journal of Finance*, 73(1):317–373, 2018.

Philip H Dybvig and Chi-Fu Huang. Nonnegative wealth, absence of arbitrage, and feasible consumption plans. *The Review of Financial Studies*, 1(4):377–401, 1988.

Emmanuel Farhi and Jean Tirole. Bubbly liquidity. *The Review of Economic Studies*, 79(2): 678–706, 2012.

Roger EA Farmer. Pricing assets in a perpetual youth model. *Review of Economic Dynamics*, 30:106–124, 2018.

Roger EA Farmer and Jess Benhabib. Indeterminacy and increasing returns. *Journal of Economic Theory*, 63:19–41, 1994.

Nicolae Gârleanu and Stavros Panageas. What to expect when everyone is expecting: Self-fulfilling expectations and asset-pricing puzzles. *Journal of Financial Economics*, 140(1):54–73, 2021.

Mark Gertler and Nobuhiro Kiyotaki. Banking, liquidity, and bank runs in an infinite horizon economy. *The American Economic Review*, 105(7):2011–43, 2015.

Mark Gertler, Nobuhiro Kiyotaki, and Andrea Prestipino. A macroeconomic model with financial panics. *The Review of Economic Studies*, 87(1):240–288, 2020.

Chao Gu, Fabrizio Mattesini, Cyril Monnet, and Randall Wright. Endogenous credit cycles. *Journal of Political Economy*, 121(5):940–965, 2013.

Lars Peter Hansen, Paymon Khorrami, and Fabrice Tourre. Comparative valuation dynamics in production economies: Long-run uncertainty, heterogeneity, and market frictions. *Annual Reviews of Financial Economics*, 2024.

Fumio Hayashi. Tobin's marginal q and average q: A neoclassical interpretation. *Econometrica: Journal of the Econometric Society*, pages 213–224, 1982.

Zhiguo He and Arvind Krishnamurthy. A model of capital and crises. *The Review of Economic Studies*, 79(2):735–777, 2012.

Zhiguo He and Arvind Krishnamurthy. Intermediary asset pricing. *The American Economic Review*, 103(2):732–70, 2013.

Zhiguo He and Arvind Krishnamurthy. A macroeconomic framework for quantifying systemic risk. *American Economic Journal: Macroeconomics*, 11(4):1–37, 2019.

Julien Hugonnier. Rational asset pricing bubbles and portfolio constraints. *Journal of Economic Theory*, 147(6):2260–2302, 2012.

Òscar Jordà, Moritz Schularick, and Alan M Taylor. Financial crises, credit booms, and external imbalances: 140 years of lessons. *IMF Economic Review*, 59(2):340–378, 2011.

Òscar Jordà, Moritz Schularick, and Alan M Taylor. When credit bites back. *Journal of Money, Credit and Banking*, 45(s2):3–28, 2013.

Òscar Jordà, Moritz Schularick, and Alan M Taylor. Betting the house. *Journal of International Economics*, 96:S2–S18, 2015a.

Òscar Jordà, Moritz Schularick, and Alan M Taylor. Leveraged bubbles. *Journal of Monetary Economics*, 76:S1–S20, 2015b.

Òscar Jordà, Moritz Schularick, and Alan M Taylor. Macrofinancial history and the new business cycle facts. *NBER Macroeconomics Annual*, 31(1):213–263, 2017.

Òscar Jordà, Katharina Knoll, Dmitry Kuvshinov, Moritz Schularick, and Alan M Taylor. The rate of return on everything, 1870–2015. *The Quarterly Journal of Economics*, 134(3): 1225–1298, 2019.

Òscar Jordà, Björn Richter, Moritz Schularick, and Alan M Taylor. Bank capital redux: solvency, liquidity, and crisis. *The Review of Economic Studies*, 88(1):260–286, 2021.

Ioannis Karatzas and Steven E Shreve. *Methods of Mathematical Finance*, volume 39. Springer Science & Business Media, 1998.

Rafail Khasminskii. *Stochastic stability of differential equations*, volume 66. Springer Science & Business Media, 2011.

Paymon Khorrami. The risk of risk-sharing: diversification and boom-bust cycles. Unpublished working paper. Imperial College London, 2021.

Paymon Khorrami and Fernando Mendo. Dynamic self-fulfilling fire sales. Unpublished working paper, 2024.

Paymon Khorrami and Fernando Mendo. Fear and volatility at the zero lower bound. Unpublished working paper, 2025.

Paymon Khorrami and Alexander K Zentefis. Segmentation and beliefs: A theory of self-fulfilling idiosyncratic risk. *Journal of Economic Theory*, 223:105954, 2025.

Nobuhiro Kiyotaki and John Moore. Credit cycles. *Journal of Political Economy*, 105(2): 211–248, 1997.

Nataliya Klimenko, Sebastian Pfeil, and Jean-Charles Rochet. A simple macroeconomic model with extreme financial frictions. *Journal of Mathematical Economics*, 68:92–102, 2017.

Narayana R Kocherlakota. Bubbles and constraints on debt accumulation. *Journal of Economic Theory*, 57(1):245–256, 1992.

Leonid Kogan, Stephen A Ross, Jiang Wang, and Mark M Westerfield. Market selection. *Journal of Economic Theory*, 168:209–236, 2017.

Arvind Krishnamurthy and Wenhao Li. Dissecting mechanisms of financial crises: Intermediation and sentiment. *Journal of Political Economy (forthcoming)*, 2024.

Arvind Krishnamurthy and Tyler Muir. How credit cycles across a financial crisis. *Journal of Finance (forthcoming)*, 2024.

Arvind Krishnamurthy and Annette Vissing-Jorgensen. The impact of treasury supply on financial sector lending and stability. *Journal of Financial Economics*, 118(3):571–600, 2015.

Nicolai V Krylov. Diffusion on a plane with reflection. construction of the process. *Siberian Mathematical Journal*, 10(2):244–252, 1969.

Nicolai V Krylov. On weak uniqueness for some diffusions with discontinuous coefficients. *Stochastic Processes and their Applications*, 113(1):37–64, 2004.

Seung Joo Lee and Marc Dordal i Carreras. Self-fulfilling volatility, risk-premium, and business cycles. May 28 2024. SSRN working paper 4461453.

Zheng Liu and Pengfei Wang. Credit constraints and self-fulfilling business cycles. *American Economic Journal: Macroeconomics*, 6(1):32–69, 2014.

David López-Salido, Jeremy C Stein, and Egon Zakrajšek. Credit-market sentiment and the business cycle. *The Quarterly Journal of Economics*, 132(3):1373–1426, 2017.

Peter Maxted. A macro-finance model with sentiment. *Review of Economic Studies*, 91(1): 438–475, 2024.

Fernando Mendo. Risky low-volatility environments and the stability paradox. Unpublished working paper, 2020.

Atif Mian, Amir Sufi, and Emil Verner. Household debt and business cycles worldwide. *The Quarterly Journal of Economics*, 132(4):1755–1817, 2017.

Jianjun Miao and Pengfei Wang. Asset bubbles and credit constraints. *The American Economic Review*, 108(9):2590–2628, 2018.

Alan Moreira and Alexi Savov. The macroeconomics of shadow banking. *The Journal of Finance*, 72(6):2381–2432, 2017.

Tyler Muir. Financial crises and risk premia. *The Quarterly Journal of Economics*, 132(2): 765–809, 2017.

Gregory Phelan. Financial intermediation, leverage, and macroeconomic instability. *American Economic Journal: Macroeconomics*, 8(4):199–224, 2016.

Carmen M Reinhart and Kenneth S Rogoff. *This time is different: Eight centuries of financial folly*. Princeton University Press, 2009.

Carmen M Reinhart and Kenneth S Rogoff. Recovery from financial crises: Evidence from 100 episodes. *The American Economic Review*, 104(5):50–55, 2014.

Jose A Scheinkman and Laurence Weiss. Borrowing constraints and aggregate economic activity. *Econometrica*, 54(1):23–45, 1986.

Moritz Schularick and Alan M Taylor. Credit booms gone bust: monetary policy, leverage cycles, and financial crises, 1870–2008. *The American Economic Review*, 102(2):1029–1061, 2012.

Dejanir Silva. The risk channel of unconventional monetary policy. 2024. Unpublished working paper.

Amir Sufi and Alan M Taylor. Financial crises: A survey. *Handbook of international economics*, 6:291–340, 2022.

Online Appendix:
Rational Sentiments and Financial Frictions

Paymon Khorrami and Fernando Mendo

March 19, 2025

A Solvency constraint as the natural borrowing limit

A.1 Solvency constraint

Here, we discuss the solvency constraint $n_t \geq 0$, which serves as the natural borrowing limit in our framework. The idea of a natural borrowing limit is that agents can borrow at most the present-value of their future income if they want to consume non-negative amounts and also not run a Ponzi scheme (see, e.g., Aiyagari, 1994). In our context, the only asset is capital, and the stream of its future dividends represents future income. Thus, if the income stream is valued at $q_t k_t$ for k_t units of capital holdings, it is sensible that an agent should be able to borrow at most this amount: $b_t \leq q_t k_t$. Since net worth is defined as assets minus liabilities, $n_t = q_t k_t - b_t$, this implies $n_t \geq 0$.

Below, we explore two microfoundations for the solvency constraint $n_t \geq 0$, which clarifies that this constraint is "natural" in some sense. We allow the possibility of zero fundamental volatility, $\sigma = 0$, for generality. Our two microfoundations assume that unsecured debts must be repaid *eventually*. That is, an asymptotic No-Ponzi condition holds, as well as a condition that rules out infinite indebtedness along the way.

To set up the environment and the constraints, consider an agent with net worth n_t who may choose any consumption and trading strategy $\{c_t, k_t\}_{t\geq 0}$ that satisfies appropriate mild integrability conditions. The dynamic budget constraint of this agent takes the form

$$dn_t = \Big[r_t n_t - c_t + q_t k_t(\mu_{R,t} - r_t)\Big]dt + q_t k_t \sigma_{R,t} \cdot dZ_t, \quad n_0 \text{ given}, \tag{A.1}$$

where $\mu_{R,t}$ is that agent's expected return on capital (which differs between experts and households). Given these trading opportunities, let M_t be the state-price density faced by this agent:

$$M_t = \exp\Big[-\int_0^t \Big(r_s + \frac{1}{2}|\pi_s|^2\Big)ds - \int_0^t \pi_s \cdot dZ_s\Big], \tag{A.2}$$

$$\text{where} \quad \sigma_{R,t} \cdot \pi_t = \mu_{R,t} - r_t. \tag{A.3}$$

Note that equation (A.3) defines π_t as the agent's market price of risk process, which again is agent-specific in our model. Because we will refer to it very often, define the exponential local martingale

$$\tilde{M}_t := \exp\Big[-\frac{1}{2}\int_0^t |\pi_s|^2 ds - \int_0^t \pi_s \cdot dZ_s\Big]. \tag{A.4}$$

The process $\tilde{M}_t$, provided it is a true martingale, will be used to define the risk-neutral probability measure $\tilde{\mathbb{P}}$. (In an infinite-horizon model, there is some additional subtlety to the construction of the risk-neutral measure, which we will explain in the proof of Lemma A.2 below.)

Given this environment, we consider two different formulations of the asymptotic No-Ponzi condition. In the first formulation, we assume that agents must obey

$$\liminf_{T\to\infty} M_T n_T \geq 0 \quad \mathbb{P}\text{-almost-surely.} \tag{NPC-1}$$

(this is weaker than the condition $\liminf_{T\to\infty} n_T \geq 0$ because of the fact that $M_T > 0$). In the second formulation, we assume that agents obey

$$\liminf_{T\to\infty} e^{-\int_0^T r_t dt} n_T \geq 0 \quad \tilde{\mathbb{P}}\text{-almost-surely,} \tag{NPC-2}$$

where $\tilde{\mathbb{P}}$ denotes the risk-neutral probability measure. The intuitive idea behind constraints (NPC-1) and (NPC-2) is as follows. By taking expectations of (NPC-1) and (NPC-2), we have that $\mathbb{E}_t[M_\infty n_\infty] \geq 0$ and $\tilde{\mathbb{E}}_t[e^{-\int_0^\infty r_t dt} n_\infty] \geq 0$, respectively. Therefore, these constraints imply that the present-value of unsecured debts must vanish eventually, ruling out arbitrarily large debts asymptotically. However, by themselves, neither (NPC-1) nor (NPC-2) is sufficient to induce the solvency constraint $n_t \geq 0$.

We impose, in addition, a uniform lower bound on net worth, but with two different functional forms. In the first formulation, we impose a lower bound on the present-value of net worth,

$$M_t n_t \geq -\underline{n}, \tag{NLB-1}$$

where $\underline{n}$ can be arbitrarily large but finite. In the second microfoundation, we impose a lower bound on net worth directly,

$$e^{-\int_0^t r_s ds} n_t \geq -\underline{n}, \tag{NLB-2}$$

where again $\underline{n}$ can be arbitrarily large but finite. Allowing $\underline{n}$ to be arbitrarily large permits any trading strategy that doesn't leave the agent infinitely indebted. Constraints (NLB-1)-(NLB-2) are examples of the requirement that portfolios be "tame" (see Karatzas and Shreve, 1998, Chapter 1, Definition 2.4). In dynamic trading models, the point of tame portfolios is to rule out certain trivial arbitrage opportunities like "doubling strategies" (c.f., Karatzas and Shreve, 1998, Chapter 1, Example 2.3). Thus, no equilibrium could exist without a requirement like (NLB-1) or (NLB-2), which is why we view these constraints as a minimal requirement.[23]

Now, we provide two proofs that the solvency constraint holds.

Lemma A.1. *Let (NPC-1) and (NLB-1) hold. Then, every agent must obey $n_t \geq 0$.*

Lemma A.2. *Let (NPC-2) and (NLB-2) hold. Suppose $\tilde{M}_t$ is a martingale. Then, every agent must obey $n_t \geq 0$.*

Remark 3. *We make a brief remark about the assumption that $\tilde{M}_t$ be a martingale in the latter lemma. This assumption should be regarded as relatively minor. Indeed, a sufficient condition for $\tilde{M}_t$ to be a martingale is that $\sup_t |\pi_t| < \infty$, i.e., risk prices be uniformly bounded. It is straightforward to verify that equilibrium risk prices only diverge at the boundary where $\eta \to 0$ and $\kappa/\eta \to +\infty$, so what we need is for state dynamics prevent the economy from approaching this boundary.[24] This can be done: an example of such an equilibrium construction is presented in Proposition D.1, in which risk prices are indeed uniformly bounded.*

Proof of Lemma A.1. The general strategy of the proof is to derive a static budget constraint, and then use this budget constraint to prove that $n_t \geq 0$.

Apply Itô's formula to the process

$$H_t := M_t n_t + \int_0^t M_s c_s ds,$$

[23] An alternative constraint that achieves the same result as (NLB-2) is to impose an integrability condition on the trading strategies agents can do:

$$\tilde{\mathbb{E}}\Big[\int_0^\infty e^{-2\int_0^t r_s ds}(q_t k_t)^2 |\sigma_{R,t}|^2 dt\Big] < \infty,$$

where $\tilde{\mathbb{E}}$ represents the risk-neutral expectation in the model. Dybvig and Huang (1988), Theorems 4 and 5, prove that the lower bound (NLB-2) and the integrability condition above are essentially equivalent in this environment: they both rule out arbitrage and permit essentially the same trading strategies. We work with the uniform net worth lower bound because it translates better into our infinite-horizon proofs.

[24] Indeed, (squared) expert risk prices are given by $|\pi|^2 = (\frac{\kappa}{\eta})^2 |\sigma_R|^2$, which after using the equilibrium value of $|\sigma_R|^2$ when $\kappa < 1$ gives us $|\pi|^2 = (\frac{\kappa}{\eta})^2 \frac{\eta(1-\eta)}{\kappa-\eta} \frac{a_e - a_h}{q}$. This is bounded except at the boundary $\eta \to 0$ and $\kappa \to \bar{\kappa} > 0$. At this boundary, the risk price behaves like $|\pi|^2 \sim \eta^{-1}\bar{C}$, where $\bar{C} := \frac{\bar{\kappa}(a_e - a_h)}{a_h + \bar{\kappa}(a_e - a_h)} \rho_h$.

then use the dynamic budget constraint (A.1) and equation (A.3) for π_t, to obtain

$$H_T - H_t = M_T n_T - M_t n_t + \int_t^T M_s c_s ds = \int_t^T M_s \Big(q_s k_s \sigma_{R,s} - n_s \pi_s \Big) \cdot dZ_s. \tag{A.5}$$

This shows that H_t is a local martingale. Furthermore, the lower bound (NLB-1) and the non-negativity of consumption imply $H_t \geq -\underline{n}$ and so H_t is a super-martingale. Taking time-t expectations of (A.5), we thus have

$$\mathbb{E}_t \Big[M_T n_T \Big] + \mathbb{E}_t \Big[\int_t^T M_s c_s ds \Big] \leq M_t n_t. \tag{A.6}$$

Because consumption is non-negative, the monotone convergence theorem implies

$$\lim_{T \to \infty} \mathbb{E}_t \Big[\int_t^T M_s c_s ds \Big] = \mathbb{E}_t \Big[\int_t^\infty M_s c_s ds \Big].$$

For the terminal wealth term, the lower bound (NLB-1) implies $(M_T n_T)_{T \geq \infty}$ is a uniformly lower-bounded family of random variables, so by Fatou's lemma we have

$$\liminf_{T \to \infty} \mathbb{E}_t \Big[M_T n_T \Big] \geq \mathbb{E}_t \Big[\liminf_{T \to \infty} M_T n_T \Big].$$

Using asymptotic No-Ponzi condition (NPC-1), the right-hand-side term is non-negative. Using these limiting results in (A.6), we have

$$\mathbb{E}_t \Big[\int_t^\infty M_s c_s ds \Big] \leq M_t n_t. \tag{A.7}$$

Equation (A.7) is the usual "static" budget constraint. From (A.7), the fact that consumption is non-negative, and the fact that the state-price density is strictly positive, we immediately obtain $n_t \geq 0$. Since time t was arbitrary, this must hold for all times. □

Proof of Lemma A.2. This proof proceeds slightly differently than Lemma A.1. Indeed, since there is no obvious lower bound that can be applied to $M_T n_T$ in equation (A.6), the proof becomes more technical and complex. The general strategy is to examine the dynamics of $e^{-\int_0^t r_s ds} n_t$, which is lower-bounded, rather than $M_t n_t$.

There are two complications. First, to continue to use martingale methods, we must examine the dynamics of $e^{-\int_0^t r_s ds} n_t$ under the risk-neutral measure $\tilde{\mathbb{P}}$ rather than the true probability $\mathbb{P}$. This is where the assumption that $\tilde{M}_t$ is a martingale, hence a valid change-of-measure, comes into play. Second, because our model is infinite-horizon, $\tilde{\mathbb{P}}$

and $\mathbb{P}$ may be mutually singular asymptotically on the limiting sigma-algebra $\mathcal{F}_\infty$, even though $\tilde{\mathbb{P}}$ and $\mathbb{P}$ are equivalent on every finite horizon. For this reason, the No-Ponzi condition (NPC-2) is written purposefully under $\tilde{\mathbb{P}}$.

First, we define a probability measure $\tilde{\mathbb{P}}$ following the recipe of Chapter 1.7 in Karatzas and Shreve (1998). Using $\tilde{M}_t$ as a change-of-measure, we set

$$\tilde{\mathbb{P}}(A) := \mathbb{E}[\tilde{M}_T \mathbf{1}_A]; \quad A \in \mathcal{F}_T, \quad 0 \leq T < \infty. \tag{A.8}$$

As proven in Chapter 1.7, Proposition 7.4 of Karatzas and Shreve (1998), the probability $\tilde{\mathbb{P}}$ is equivalent to $\mathbb{P}$ on $\mathcal{F}_T$ for each $T \geq 0$ (i.e., a set in $\mathcal{F}_T$ is a $\tilde{\mathbb{P}}$-null set if and only if it is a $\mathbb{P}$-null set). Furthermore, the process

$$\tilde{Z}_t := Z_t + \int_0^t \pi_s ds$$

is a Brownian motion on under $\tilde{\mathbb{P}}$.

Consider now the process

$$H_t := e^{-\int_0^t r_s ds} n_t + \int_0^t e^{-\int_0^s r_u du} c_s ds,$$

which follows

$$dH_t = e^{-\int_0^t r_s ds} \left(q_t k_t \sigma_{R,t} \right) \cdot d\tilde{Z}_t. \tag{A.9}$$

By the non-negativity of consumption and the lower bound (NLB-2), we have that $H_t \geq -\underline{n}$, so H_t is a $\tilde{\mathbb{P}}$-super-martingale. Taking time-t risk-neutral expectations of $H_T - H_t$, we thus have

$$\tilde{\mathbb{E}}_t \left[e^{-\int_0^T r_s ds} n_T \right] + \tilde{\mathbb{E}}_t \left[\int_t^T e^{-\int_0^s r_u du} c_s ds \right] \leq e^{-\int_0^t r_s ds} n_t. \tag{A.10}$$

Because consumption is non-negative, the monotone convergence theorem implies

$$\lim_{T \to \infty} \tilde{\mathbb{E}}_t \left[\int_t^T e^{-\int_0^s r_u du} c_s ds \right] = \tilde{\mathbb{E}}_t \left[\int_t^\infty e^{-\int_0^s r_u du} c_s ds \right].$$

For the terminal wealth term, the lower bound (NLB-2) implies $(e^{-\int_0^T r_s ds} n_T)_{T \geq \infty}$ is a uniformly lower-bounded family of random variables. Because $\tilde{\mathbb{P}}$ and $\mathbb{P}$ are equivalent on all finite horizons, the almost-sure lower-bound holds both under $\tilde{\mathbb{P}}$ and $\mathbb{P}$, so by

Fatou's lemma we have

$$\liminf_{T\to\infty} \tilde{\mathbb{E}}_t\left[e^{-\int_0^T r_s ds} n_T\right] \geq \tilde{\mathbb{E}}_t\left[\liminf_{T\to\infty} e^{-\int_0^T r_s ds} n_T\right].$$

Using asymptotic No-Ponzi condition (NPC-2), the right-hand-side term is non-negative. Using these limiting results in (A.10), we have

$$\tilde{\mathbb{E}}_t\left[\int_t^\infty e^{-\int_0^s r_u du} c_s ds\right] \leq e^{-\int_0^t r_s ds} n_t. \tag{A.11}$$

Equation (A.11) is the usual "static" budget constraint. From (A.11), and the fact that consumption is non-negative, we immediately obtain $n_t \geq 0$. □

A.2 Application: zero fundamental uncertainty case

One of the most striking results we will present is that non-fundamental equilibria can emerge even if $\sigma = 0$. While one could regard this as a limiting case as $\sigma \to 0$, some readers may expect a discontinuity in the results when σ literally equals 0. With no borrowing frictions, the riskless bond market seems to be enough to make financial markets complete when $\sigma = 0$, and so the First Welfare Theorem holds. Under the First Welfare Theorem, we would have generic equilibrium uniqueness.

For our economy, whether or not the financial market is complete or incomplete is actually *endogenous* and depends on whether asset prices q_t are volatile. Imagine an individual expert operating in a world where $\sigma_q \neq 0$. For him, equity-issuance constraints matter because outside equity is the only way to hedge capital price shocks. As stated by Chiappori and Guesnerie (1991), "the existence of a complete set of initial markets is not enough...Insurance markets against sunspot should also be introduced to allow full insurance."

But is this statement vacuous? Why can't a researcher take any economic model and make its financial markets incomplete by simply conjecturing its asset price dynamics depend on some extrinsic shocks? The answer, suggested by our discussion in Section 1.3, is that the structure of most economies rules out any dependence of asset prices on extrinsic shocks. For example, we showed that q cannot be stochastic with $a_e = a_h$. In such cases, even if extrinsic shocks are strictly speaking uninsurable, markets are *effectively complete* because equilibrium cannot support extrinsic shocks to asset prices.

An alternative line of thinking suggests agents should ignore shocks to q when $\sigma = 0$. Whereas fundamental shocks directly impact capital, extrinsic shocks to prices only affect net worth "on paper." For example, consider the following buy-and-hold strategy:

borrow using the riskless bond market; use the proceeds to purchase capital; use the cash flows from capital to repay debts over time; ignore any capital price fluctuations and never sell the capital; and consume after all debts are repaid. Assuming no exogenous growth ($g = 0$) for simplicity, this trading strategy has cash flows $\{a_e - r_t b_t\}_{t \geq 0}$, where the debt balance b_t satisfies $db_t = -(a_e - r_t b_t - c_t)dt$ with $b_0 = q_0$. The consumption associated with this strategy is $c_t = \mathbf{1}_{t>\tau} a_e$, where $\tau := \inf\{t : b_t \geq 0\}$ is the time when all debts are repaid. Since this consumption is non-negative, and zero initial investment was made, this is an arbitrage if it is feasible. Furthermore, if all experts behaved in this way, capital prices would not be volatile or ever fall below their efficient value.

The general problem with such strategies that "ignore market prices" is that debts can become arbitrarily large. When the interest rate rises, the example strategy above produces negative cash flows. Agents must increase their borrowing to continue holding capital. With positive probability, this happens so often and for so long that either debts approach infinity, or default occurs eventually. If markets impose the requirements that net worth remains lower bounded and all debts are eventually repaid, such a strategy is ruled out. This is the content of the previous section, where we showed more generally that a net worth lower bound and a No-Ponzi constraint are equivalent to a solvency constraint $n_t \geq 0$ that rules out all arbitrage trades. In other words, the "ignore market prices" trade is not feasible, so sentiment equilibria are not ruled out even if $\sigma = 0$.

B Proofs for Sections 1-2

B.1 Irrelevance of type-switching for optimal behavior

The objective function with type-switching technically differs from (3), because agents understand that at a future exponentially-distributed time, they will switch occupations. Mathematically, the objective functions and indirect utilities satisfy the recursions, for each type-j (expert or household) agent

$$V_{j,t} = \sup_{c_j \geq 0, k_j \geq 0, n_j \geq 0} \mathbb{E}\Big[\int_0^{T_j} e^{-\rho_j s} \log(c_{j,t+s}) ds + e^{-\rho_j T} V_{-j,t+T_j}\Big], \quad T_j \sim \exp(\delta_j)$$

Standard homogeneity arguments imply that indirect utilities take the additively-separable form $V_{j,t} = \rho_j^{-1} \log(n_{j,t}) + \xi_{j,t}$, for processes $\xi_{j,t}$ that only depend on aggregates (i.e., not on individual net worth). Write $d\xi_{j,t} = \mu_{\xi,j,t} dt + \sigma_{\xi,j,t} \cdot dZ_t$. Then, the HJB equations

associated with these equations are

$$\rho_j V_j = \max_{c,k\geq 0} \log(c) + (\partial_n V_j)[rn - c + qk(\mu_{R,j} - r)] + \frac{1}{2}(\partial_{nn} V_j)(qk)^2|\sigma_R|^2 + \mu_{\xi,j} + \delta_e[V_{-j} - V_j]$$

where $\mu_{R,j}$ is the expected returns on capital for type j. Using the form of V_j, the HJB equations become

$$\log(n) + \rho_j \xi_j = \max_{c,k\geq 0} \log(c) + \rho_j^{-1}[r - \frac{c}{n} + \frac{qk}{n}(\mu_{R,j} - r)] - \frac{1}{2}(\frac{qk}{n})^2|\sigma_R|^2 + \mu_{\xi,j} + \delta_e[\xi_{-j} - \xi_j]$$

Optimal choices take the familiar log-utility forms: consumptions are $c_j = \rho_j n_j$; portfolios are $\frac{qk_j}{n_j} = [\frac{\mu_{R,j}-r}{|\sigma_R|^2}]^+$. Most importantly, these choices are independent of the switching parameters δ_j. To fully verify that this is correct, we must substitute the optimality conditions back into the HJB equations and check that we recover equations for ξ_e and ξ_h that only depend on aggregate variables (e.g., capital price q, interest rate r, etc.). Doing this, we obtain

$$\rho_j \xi_j = \log(\rho_j) + \rho_j^{-1}[r - \rho_j + \frac{1}{2}(\frac{[\mu_{R,j} - r]^+}{|\sigma_R|})^2] + \mu_{\xi,j} + \delta_j[\xi_{-j} - \xi_j],$$

which verifies the conjecture, as all terms either pertain to the ξ processes or aggregate variables. □

B.2 Proof of Lemma 1

We are given η_0 and conditions (PO), (RB), (11), and (13)-(14). We need to check conditions (i)-(iii) of Definition 1. Condition (i) holds by the definition of η_0.

For condition (ii), note that standard martingale techniques can be applied to verify that individual optimality, subject to the dynamic budget constraint (2), is equivalent to the following conditions holding: $c_\ell = \rho_\ell n_\ell$; the portfolio conditions (7)-(8); and the transversality conditions in (10). We must verify that these conditions hold. Given q_t, η_t, κ_t, and individual net worths $n^i_{e,t}$ and $n^j_{h,t}$, let us set

$$c^i_{e,t} = \rho_e n^i_{e,t} \quad \text{and} \quad k^i_{e,t} = \frac{\kappa_t}{q_t \eta_t} n^i_{e,t}, \quad \text{for } i \in \mathbb{I} \tag{B.1}$$

$$c^j_{h,t} = \rho_h n^j_{h,t} \quad \text{and} \quad k^j_{h,t} = \frac{1-\kappa_t}{q_t(1-\eta_t)} n^j_{h,t}, \quad \text{for } j \in \mathbb{J}. \tag{B.2}$$

If we do this, then clearly the optimal consumption-wealth ratio holds. Similarly, after

substituting the suggested capital holdings from (B.1)-(B.2), the optimal portfolio conditions (7)-(8) become a linear transformation of equations (RB) and (11)—i.e., equation (RB) is the difference between (7) and (8), while (11) is the sum of κ times (7) plus $1-\kappa$ times (8). Thus, given (RB) and (11), equations (7)-(8) hold as well. Finally, after substituting the proposals in (B.1)-(B.2) into the transversality conditions in (10), we see that these hold automatically.

For condition (iii), note that $\kappa \in [0,1]$ automatically implies capital market clearing (5). Similarly, substituting $c_\ell = \rho_\ell n_\ell$ and the definitions of κ and η into (PO), we obtain goods market clearing (4).

Thus, we have constructed an equilibrium of Definition 1. Note that (13)-(14) have not been used in this construction, but they are direct consequences (by Itô's formula) of η's definition. The final statement of the lemma is clearly true. Indeed, (q_t, r_t) are directly involved in Definition 1, while the objects (η_t, κ_t) constitute two summary statistics of the distribution of net worth and capital $\{n^i_{e,t}, n^j_{h,t}, k^i_{e,t}, k^j_{h,t} : i \in \mathbb{I},\, j \in \mathbb{J}\}$. Thus, two distinct values of $(\eta_t, q_t, \kappa_t, r_t)_{t\geq 0}$ cannot correspond to the same equilibrium of Definition 1. □

B.3 Proof of Theorem 1

Step 0: Reduce the system. We will start by eliminating $(r, \kappa, \sigma_\eta, \mu_\eta)$ from the system of endogenous objects, given $(\eta, q, \sigma_q, \mu_q)$. First, price-output relation (PO) determines κ as a function of (η, q) and nothing else, given by

$$\kappa(\eta, q) := \frac{q\bar{\rho}(\eta) - a_h}{a_e - a_h}. \tag{B.3}$$

Second, substituting this result for κ, equation (11) fully determines r, given knowledge of $(\eta, q, \sigma_q, \mu_q)$. Third, equations (13)-(14), after plugging in the result for κ, fully determine (σ_η, μ_η), given knowledge of (η, q, σ_q). Thus, given (η, q), the choice of (σ_q, μ_q) needs to ensure that (RB) holds and that the dynamics of (η_t, q_t) remain inside the domain $\mathcal{D} := \{(\eta, q) : 0 < \eta < 1,\, q^L(\eta) < q \leq q^H(\eta)\}$, as defined by (16) in text.

The remainder of the proof is entirely devoted to addressing the boundaries of $\mathcal{D}$. Indeed, given $(\eta, q) \in \mathcal{D}^\circ$ (the interior of $\mathcal{D}$), we can set σ_q according to (B.6) below and set μ_q to any real number. This is not to suggest that the boundary points are inconsequential; on the contrary, without ensuring that the system $(\eta_t, q_t)_{t\geq 0}$ remains in $\mathcal{D}$, the solution constructed in the interior $\mathcal{D}^\circ$ would not be part of an equilibrium. Unfortunately, the choice of (σ_q, μ_q) is more delicate at the boundary $\partial\mathcal{D}$. Furthermore, verifying that $(\eta_t, q_t)_{t\geq 0}$ remains in $\mathcal{D}$ is non-trivial and requires a detailed analysis.

Step 1: Define perturbed domain. To facilitate analysis, it will be convenient to analyze a slightly modified system instead of (η, q), and on a perturbed domain. The purpose of this perturbation will be threefold. First, as q approaches the lower boundary of $\mathcal{D}$, volatility σ_q necessarily grows without bound; by perturbing this boundary slightly upward, we prevent unbounded volatilities, allowing us to use standard diffusion theory. Second, as q approaches the upper boundary of $\mathcal{D}$, there will exist a wealth level η^* such that $\kappa = 1$ cannot possibly occur on $\{\eta \leq \eta^*\}$ but can occur on $\{\eta > \eta^*\}$; by rotating this upper boundary around any wealth share above η^*, we streamline our arguments. Third, our perturbed domain will be an open set, which is easier to work with. See Figure B.1 below for a visual of the domain perturbation. By the end of this step, it will become clear that if our modified system (η, x) remains in perturbed domain $\mathcal{X}$, then the original system (η, q) remains in the original domain $\mathcal{D}$. Furthermore, after constructing an equilibrium in this perturbed domain, it will be clear that we are able to consider the limit of a sequence of such equilibria as the perturbations vanish, and so we can also construct an equilibrium on the full domain $\mathcal{D}$ (although this is not what Theorem 1 requires us to prove).

First, define the following auxiliary functions that we use to perturb the domain boundaries. Fix $\epsilon_\alpha, \epsilon_\beta \in (0, \frac{a_e - a_h}{\rho_h})$. Let $\alpha(\cdot)$ be an increasing, continuously differentiable function such that $\alpha(0) = 0$, $\alpha'(0) > 0$, and $\alpha(1) = \epsilon_\alpha$. Let $\beta(\cdot)$ be an increasing, continuously differentiable function such that $\beta(\eta) = 0$ for all $\eta \leq \eta^*$ and $\beta(1) = \epsilon_\beta / \eta^*$, where

$$\eta^* := \frac{\rho_h}{\rho_e}\left(\frac{1 - a_h/a_e}{\sigma^2}\rho_e - 1 + \frac{\rho_h}{\rho_e}\right)^{-1}. \tag{B.4}$$

Note that $\eta^* < 1$ by Assumption 1, part (ii). This threshold η^* is the one mentioned above, where equilibrium does not permit $\kappa = 1$ for any $\eta \leq \eta^*$.

Next, recall the following upper and lower bounds for the capital price,

$$\begin{aligned} q^H(\eta) &:= a_e / \bar{\rho}(\eta) \\ q^L(\eta) &:= \bar{a}(\eta) / \bar{\rho}(\eta), \end{aligned}$$

where $\bar{a}(\eta) := \eta\rho_e + (1 - \eta)\rho_h$. Using (B.3), one notices that q^H corresponds to the capital price when $\kappa = 1$, whereas q^L corresponds to the capital price when $\kappa = \eta$. Construct

the following perturbed upper and lower bounds by putting

$$q^H_\beta(\eta) := q^H(\eta) + \beta(\eta)$$
$$q^L_\alpha(\eta) := q^L(\eta) + \alpha(\eta).$$

Using these functions, define the perturbed domain (which is an open set)

$$\mathcal{X} := \Big\{(\eta, x) \,:\, \eta \in (0,1) \quad \text{and} \quad q^L_\alpha(\eta) < x < q^H_\beta(\eta)\Big\}.$$

Note that, boundaries aside, $\mathcal{X}$ will coincide with $\mathcal{D}$ as $\epsilon_\alpha \to 0$ and $\epsilon_\beta \to 0$. For reference, the perturbed domain $\mathcal{X}$ is displayed in Figure B.1.

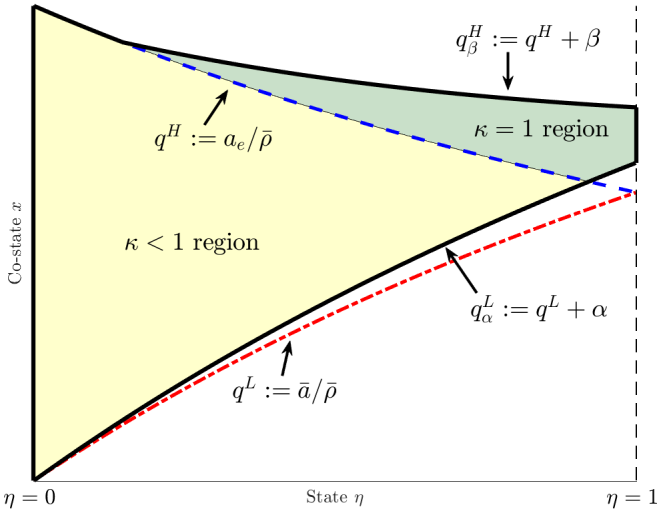


Figure B.1: The perturbed domain $\mathcal{X}$ is shown as the shaded region surrounded by solid black lines. The original domain $\mathcal{D}$ is the region defined by the dashed lines. The perturbation functions α and β are chosen to be linear functions, with $\epsilon_\alpha = 0.2$ and $\epsilon_\beta = 0.05$. Parameters: $\rho_e = 0.07$, $\rho_h = 0.05$, $a_e = 0.11$, $a_h = 0.03$, $\sigma = 0.1$.

We will define a stochastic process x_t such that the capital price q coincides with x when it lies below q^H, i.e.,

$$q_t = \min\Big[x_t,\, q^H(\eta_t)\Big]. \tag{B.5}$$

By (B.5), we may analyze the dynamical system $(\eta_t, x_t)_{t\geq 0}$ rather than $(\eta_t, q_t)_{t\geq 0}$. Furthermore, to prove the claim that $(\eta_t, q_t)_{t\geq 0}$ remains in $\mathcal{D}$ almost-surely, it suffices to prove $(\eta_t, x_t)_{t\geq 0}$ remains in $\mathcal{X}$ almost-surely (Step 6 below).

Step 2: Allow auxiliary state variables. We introduce some auxiliary state variables here, because several of the indeterminate objects to follow can depend arbitrarily on them. We assume y_t satisfies an exogenous, stationary Markov diffusion $dy_t = \mu_y(y_t)dt + \sigma_y(y_t) \cdot dZ_t$ on the N-dimensional domain $\mathcal{Y}$. Moreover, we assume the process is sufficiently

well-behaved in the sense that it possesses an appropriate "Lyapunov" function. To write this condition, first define the infinitesimal generator $\mathscr{L}^y$ of y_t, where for any C^2 function $f : \mathcal{Y} \mapsto \mathbb{R}$,

$$\mathscr{L}^y f = \mu_y \partial_y f + \frac{1}{2}|\sigma_y|^2 \partial_{yy} f.$$

Assume there exists a non-negative C^2 function $v_y : \mathcal{Y} \mapsto \mathbb{R}_+$, such that $\liminf_{y \to \partial\mathcal{Y}} v_y = +\infty$ and such that $\mathscr{L}^y v_y \leq 0$ on $\mathcal{Y}$. The implication of this assumption is that $(y_t)_{t\geq 0}$ is a recurrent process (see Lemma 3.9 of Khasminskii (2011)). Together, the endogenous and exogenous states (η_t, x_t, y_t) exist in the domain $\mathcal{X} \times \mathcal{Y}$.

Step 3: Construct σ_q so that (RB) is satisfied. First consider $\{x < q^H(\eta)\}$ so that $q = x$. Note that this case corresponds to $\kappa < 1$. Let $\vartheta(\eta, x, y) : \mathcal{X} \times \mathcal{Y} \mapsto [0, 1]$ be any C^1 function, whose dependence on y vanishes for all (η, x) close enough to the boundary $\partial\mathcal{X}$. Put

$$\sigma_q = \begin{bmatrix} \sqrt{\vartheta \frac{\eta(1-\eta)}{\kappa - \eta} \frac{a_e - a_h}{q}} - \sigma \\ \sqrt{(1-\vartheta) \frac{\eta(1-\eta)}{\kappa - \eta} \frac{a_e - a_h}{q}} \end{bmatrix}, \quad \text{if } x < q^H(\eta). \tag{B.6}$$

Substituting (B.6), one can verify that the second term of condition (RB) is zero. Importantly, the definitions of q^L_α and q^H_β imply that σ_q is bounded on $\mathcal{X} \cap \{x < q^H(\eta)\}$. Indeed, because of $\alpha'(0) > 0$, the slowest possible rate that $\kappa \to 0$ as $\eta \to 0$ is lower-bounded away from 1, i.e., $\liminf_{\eta \to 0, (\eta,x) \in \mathcal{X}} \kappa/\eta > 1$. And because $\alpha(1) > 0$, we have $\kappa = 1$ for all η near enough to 1; thus η is bounded away from 1 on $\{x < q^H(\eta)\}$.

Next consider $\{x \geq q^H(\eta)\}$ so that $q = q^H(\eta)$. Note that this case corresponds to $\kappa = 1$. Since q is an explicit function of η, we use Itô's formula to compute $\binom{1}{0} \cdot \sigma_q = -\sigma_\eta \bar{\rho}'/\bar{\rho}$, which after substituting equation (14) for σ_η delivers

$$\sigma_q = \begin{bmatrix} -\frac{(1-\eta)(\rho_e - \rho_h)/\bar{\rho}}{1 + (1-\eta)(\rho_e - \rho_h)/\bar{\rho}} \sigma \\ 0 \end{bmatrix}, \quad \text{if } x \geq q^H(\eta). \tag{B.7}$$

Note that (B.7) will be consistent with (RB) as long as $(\eta_t, x_t)_{t\geq 0}$ remains in $\mathcal{X}$ almost-surely, which will be verified in Step 6.[25]

Note finally that σ_q defined in (B.6)-(B.7) is solely a function of (η, x, y), so sometimes

[25] Plugging $q = a_e/\bar{\rho}$ into the second term of equation (RB), we require $|\sigma_R|^2 \leq \eta\bar{\rho}(\eta)(1 - a_h/a_e)$. On the other hand, (B.7) implies $|\sigma_R|^2 = \sigma^2(\bar{\rho}/\rho_e)^2$. Combining these, we require $\eta \geq \eta^*$ when $x \geq q^H(\eta)$, where η^* is defined in (B.4). Therefore, for all $\eta < \eta^*$, equilibrium insists $x < q^H(\eta)$. As long as $(\eta, x) \in \mathcal{X}$, this will hold, because of the way $\mathcal{X}$ is defined as an open set.

we will write $\sigma_q(\eta, x, y)$. Similarly, with σ_q in hand, we now have μ_η and σ_η as functions of (η, x, y) alone. However, notice that, by the assumption made on ϑ, the dependence of $(\sigma_q, \sigma_\eta, \mu_\eta)$ on y vanishes as (η, x) approaches $\partial\mathcal{X}$.

Step 4: Construct μ_q. Similar to σ_q, separately consider $\{x < q^H(\eta)\}$ and $\{x \geq q^H(\eta)\}$. On $\{x \geq q^H(\eta)\}$, since $q = q^H(\eta)$ is an explicit function of η, we set μ_q via Itô's formula. On $\{x < q^H(\eta)\}$, we have no equilibrium considerations restricting μ_q. Thus, we will put $\mu_q = m_q$, where m_q is any function in class $\mathcal{M}$, defined as follows. Let $\epsilon > 0$ be small enough. A function $m : \mathcal{X} \times \mathcal{Y} \mapsto \mathbb{R}$ is a member of $\mathcal{M}$ if m is C^1; if m is independent of $y \in \mathcal{Y}$ for all (η, x) close enough to the boundary $\partial\mathcal{X}$; and if m possesses the following boundary conditions:

$$\inf_{\eta\in(0,1)} \lim_{x\searrow q^L_\alpha(\eta)} \Big(x - q^L_\alpha(\eta)\Big) m(\eta, x, y) = +\infty \tag{B.8}$$

$$\sup_{\eta\in(0,1)} \lim_{x\nearrow q^H_\beta(\eta)} \Big(q^H_\beta(\eta) - x\Big) m(\eta, x, y) = -\infty \tag{B.9}$$

$$\text{for any } x \in (q^L_\alpha(0), q^H_\beta(0)), \quad \lim_{\eta\searrow 0} |m(\eta, x, y)| < +\infty \tag{B.10}$$

$$\text{for any } x \in (q^L_\alpha(1), q^H_\beta(1)), \quad \lim_{\eta\nearrow 1} |m(\eta, x, y)| < +\infty. \tag{B.11}$$

Because of the condition that m is independent of y close enough to the boundary $\partial\mathcal{X}$, the boundary conditions above automatically apply for all possible $y \in \mathcal{Y}$. Collecting these results

$$\mu_q(\eta, x, y) = \begin{cases} m_q(\eta, x, y), & \text{if } x < q^H(\eta); \\ \frac{\rho_e - \rho_h}{\bar\rho(\eta)^2}[-\bar\rho(\eta)\mu_\eta(\eta, x) + |\sigma_\eta(\eta, x)|^2], & \text{if } x \geq q^H(\eta). \end{cases} \tag{B.12}$$

Note that μ_q is indeterminate almost everywhere in the inefficient region ($\kappa < 1$), in the following sense: m_q only needs to satisfy the boundary conditions (B.8)-(B.11), but these boundary conditions constitute a zero-measure subset of $\mathcal{X}$. Furthermore, we clearly can allow the dependence of $m_q(\eta, x, y)$ on y to extend to an arbitrarily large subset of $\mathcal{X}$, so long as the dependence vanishes at $\partial\mathcal{X}$.

Step 5: Construct dynamics of x in the efficient region. The dynamics of x_t are specified as follows. Denote its diffusion and drift coefficients by $(x\sigma_x, x\mu_x)$, where σ_x and μ_x are functions of (η, x, y) to be specified shortly. By (B.5), in the region when $x < q^H(\eta)$, we must put $\sigma_x = \sigma_q$ and $\mu_x = \mu_q$. But when $x \geq q^H(\eta)$, then (σ_x, μ_x) are unrestricted. We set them arbitrarily, subject to the constraint that they induce stationarity.

To this end, let $\tilde{\sigma}_x : \mathcal{X} \times \mathcal{Y} \mapsto \mathbb{R}_+$ be any positive, bounded, C^1 function, whose dependence on y vanishes as (η, x) approaches the boundary of $\mathcal{X}$. Put

$$\sigma_x(\eta, x, y) = \begin{cases} \sigma_q(\eta, x, y), & \text{if } x < q^H(\eta); \\ \tilde{\sigma}_x(\eta, x, y), & \text{if } x \geq q^H(\eta). \end{cases}$$

Note that σ_x is bounded (recall σ_q is bounded, and $\tilde{\sigma}_x$ is assumed bounded).

Similarly, for the drift, let $m_x : \mathcal{X} \times \mathcal{Y} \mapsto \mathbb{R}$ be any function in class $\mathcal{M}$ defined above (note: m_x need not coincide with m_q above). Put

$$\mu_x(\eta, x, y) = \begin{cases} \mu_q(\eta, x, y), & \text{if } x < q^H(\eta); \\ m_x(\eta, x, y), & \text{if } x \geq q^H(\eta). \end{cases}$$

Thus, μ_x satisfies boundary conditions (B.8)-(B.11) on all boundaries of $\mathcal{X}$. At this point, let us observe the following important property: while $(\sigma_x, \sigma_\eta, \mu_x, \mu_\eta)$ can all potentially depend on the auxiliary state variable y, this dependence vanishes as (η, x) approaches the boundary $\partial\mathcal{X}$.

Step 6: Verify stationarity. We demonstrate the time-paths $(\eta_t, x_t)_{t\geq 0}$ remain in $\mathcal{X}$ almost-surely and admit a stationary distribution.

Corresponding to the SDEs induced by $(\sigma_\eta, \sigma_x, \sigma_y, \mu_\eta, \mu_x, \mu_y)$, define the infinitesimal generator $\mathscr{L}^{\eta,x,y}$, where for any C^2 function $f : \mathcal{X} \times \mathcal{Y} \mapsto \mathbb{R}$,

$$\begin{aligned} \mathscr{L}^{\eta,x,y} f = \mu_\eta \partial_\eta f + (x\mu_x)\partial_x f + \frac{1}{2}|\sigma_\eta|^2 \partial_{\eta\eta} f + \frac{1}{2}|x\sigma_x|^2 \partial_{xx} f + x\sigma_x \cdot \sigma_\eta \partial_{\eta x} f \\ + \mu_y \partial_y f + \frac{1}{2}|\sigma_y|^2 \partial_{yy} f + \sigma_\eta \cdot \sigma_y \partial_{\eta y} f + x\sigma_x \cdot \sigma_y \partial_{xy} f \end{aligned} \tag{B.13}$$

The key to the remainder of the proof will be to analyze the behavior of $\mathscr{L}^{\eta,x,y}$ near the boundary of the domain $\mathcal{X} \times \mathcal{Y}$.

The first observation is that, using standard arguments, we can construct a process $(\eta_t, x_t, y_t)_{0\leq t\leq \tau}$ up until the "first exit time" τ from the domain $\mathcal{X} \times \mathcal{Y}$.[26] Our goal is to

[26]This procedure goes as follows. Let $\{\mathcal{X}_n \times \mathcal{Y}_n\}_{n\geq 1}$ be an increasing sequence of open sets, whose closures are contained in $\mathcal{X} \times \mathcal{Y}$, such that $\cup_{n\geq 1} \mathcal{X}_n \times \mathcal{Y}_n = \mathcal{X} \times \mathcal{Y}$. Note that $(\sigma_\eta, \sigma_x, \sigma_y, \mu_\eta, \mu_x, \mu_y)$ are bounded on $\mathcal{X}_n \times \mathcal{Y}_n$ for each n. Consequently, despite the (potential) discontinuity in $(\sigma_\eta, \sigma_x, \mu_\eta, \mu_x)$ at the one-dimensional subset $\{x = q^H(\eta)\}$, there exists a unique weak solution $(\tilde{\eta}^n_t, \tilde{x}^n_t, \tilde{y}^n_t)_{0\leq t\leq \tau_n}$, up to first exit time $\tau_n := \inf\{t : (\eta_t, x_t, y_t) \notin \mathcal{X}_n \times \mathcal{Y}_n\}$, to the SDEs defined by the infinitesimal generator $\mathscr{L}^{\eta,x,y}$. See Krylov (1969, 2004) for weak existence and uniqueness in the presence of drift and diffusion discontinuities. We thus define $(\eta_t, x_t, y_t)_{0\leq t\leq \tau}$ up to exit time $\tau := \lim_{n\to\infty} \tau_n$, by piecing $(\tilde{\eta}^n_t, \tilde{x}^n_t, \tilde{y}^n_t)_{0\leq t\leq \tau_n}$ together for successive n. In other words, $(\eta_t, x_t, y_t) = (\tilde{\eta}^n_t, \tilde{x}^n_t, \tilde{y}^n_t)$ for $0 \leq t \leq \tau_n$, each n.

show (a) $\tau = +\infty$ a.s., so the process never exits its domain; and (b) the resulting process $(\eta_t, x_t, y_t)_{t\geq 0}$ possesses a non-degenerate stationary distribution on $\mathcal{X} \times \mathcal{Y}$. These will be proved if we can obtain a Lyapunov function V satisfying Lemma B.1 below.

Define the positive Lyapunov function V by

$$V(\eta, x, y) := \underbrace{\frac{1}{\eta^{1/2}} + \frac{1}{1-\eta} + \frac{1}{x - q_\alpha^L(\eta)} + \frac{1}{q_\beta^H(\eta) - x}}_{:=v(\eta,x)} + v_y(y), \tag{B.14}$$

where recall that v_y is the function from Step 2 satisfying $\mathscr{L}^y v_y \leq 0$. Note that V diverges to $+\infty$ at the boundaries of $\mathcal{X} \times \mathcal{Y}$, so assumption (i) of Lemma B.1 is proved.

Next, let us note the following useful property that will help in proving assumptions (ii)-(iii) of Lemma B.1. Due to the form of V in (B.14), there are no cross derivatives and so $\mathscr{L}^{\eta,x,y} V = \mathscr{L}^{\eta,x} v + \mathscr{L}^y v_y$, where $\mathscr{L}^{\eta,x}$ is the generator of (η, x), i.e.,

$$\mathscr{L}^{\eta,x} = \mu_\eta \partial_\eta f + (x\mu_x)\partial_x f + \frac{1}{2}|\sigma_\eta|^2 \partial_{\eta\eta} f + \frac{1}{2}|x\sigma_x|^2 \partial_{xx} f + x\sigma_x \cdot \sigma_\eta \partial_{\eta x} f \tag{B.15}$$

for any C^2 function $f : \mathcal{X} \mapsto \mathbb{R}$. Using $\mathscr{L}^y v_y \leq 0$, we then obtain $\mathscr{L}^{\eta,x,y} V \leq \mathscr{L}^{\eta,x} v$. Consequently, if we can show that v satisfies assumptions (ii)-(iii) on the domain $\mathcal{X}$, those assumptions will automatically hold for V on the overall domain $\mathcal{X} \times \mathcal{Y}$.

If assumption (iii) of Lemma B.1 holds for v (which we will prove below), then $\mathscr{L}^{\eta,x} v < 0$ at all points (η, x) sufficiently close to $\partial\mathcal{X}$. Furthermore, for every subset bounded away from this boundary, we have that $\mathscr{L}^{\eta,x} v$ is bounded. Consequently, we can find a constant c large enough such that $\mathscr{L}^{\eta,x} v \leq cv$ on all of $\mathcal{X}$, which verifies part (ii) of Lemma B.1.

It remains to prove that assumption (iii) of Lemma B.1 holds for v, namely that $\mathscr{L}^{\eta,x} v \to -\infty$ as $(\eta, x) \to \partial\mathcal{X}$. We will examine the boundaries of $\mathcal{X}$ one-by-one. In the following, we use the notation $g(x) = o(f(x))$ if $g(x)/f(x) \to 0$ as $x \to 0$, and the notation $g(x) = O(f(x))$ if $g(x)/f(x) \to C$ as $x \to 0$, where C is a finite constant.

Step 6a: boundary as $\eta \to 0$. As $\eta \to 0$ (and x bounded away from $q_\alpha^L(0)$ and $q_\beta^H(0)$, such that κ is bounded away from 0 and 1, the latter due to the definition of q_β^H), we have

$$\mu_\eta = \delta_h + \frac{a_e - a_h}{x}\kappa + \eta[\rho_h - \rho_e - \delta_e - \delta_h] + o(\eta) \quad \text{and} \quad |\sigma_\eta|^2 = \eta(\kappa - \eta)\frac{a_e - a_h}{x} + o(\eta)$$
$$\mu_x = O(1) \quad \text{and} \quad |\sigma_x|^2 = O(1).$$

We used condition (B.10) to obtain μ_x bounded. Thus,

$$\mathscr{L}^{\eta,x}v = -\frac{1}{2\eta^{3/2}}[\delta_h + \frac{1}{4}\frac{a_e - a_h}{x}\kappa] + o(\eta^{-3/2}) \to -\infty,$$

irrespective of $\delta_h > 0$ or $\delta_h = 0$.

Step 6b: boundary as $\eta \to 1$. As $\eta \to 1$ (and x bounded away from $q^L_\alpha(1)$ and $q^H_\beta(1)$; note that $\kappa = 1$ at this boundary), we have

$$\begin{aligned} \mu_\eta = -\delta_e - (\rho_e - \rho_h)(1-\eta) + o(1-\eta) \quad \text{and} \quad |\sigma_\eta|^2 = (1-\eta)^2\sigma^2 \\ \mu_x = O(1) \quad \text{and} \quad |\sigma_x|^2 = O(1). \end{aligned}$$

We used condition (B.11) to obtain μ_x bounded. Thus,

$$\mathscr{L}^{\eta,x}v = -(1-\eta)^{-2}\delta_e - (1-\eta)^{-1}[\rho_e - \rho_h - \sigma^2] + o((1-\eta)^{-1}) \to -\infty,$$

irrespective of δ_e, due to Assumption 1 part (iii).

Step 6c: boundary as $x \to q^L_\alpha$. We separately calculate the limit $x \to q^L_\alpha(\eta)$ (with η bounded away from 0) in the two cases $\{x < q^H(\eta)\}$ and $\{x \geq q^H(\eta)\}$, since $\kappa < 1$ in the first case, and $\kappa = 1$ in the second case. Still, we find that in both cases,

$$\begin{aligned} \mu_\eta = O(1) \quad \text{and} \quad |\sigma_\eta|^2 = O(1) \\ \mu_x = o((x - q^L_\alpha)^{-1}) \quad \text{and} \quad |\sigma_x|^2 = O(1). \end{aligned}$$

We used condition (B.8) to obtain the order of μ_x. Thus,

$$\mathscr{L}^{\eta,x}v = -(x - q^L_\alpha)^{-2}x\mu_x + O((x - q^L_\alpha)^{-3}) \to -\infty.$$

Step 6d: boundary as $x \to q^H_\beta$. Similarly, we separately calculate the limit $x \to q^H_\beta(\eta)$ (with η bounded away from 0) in the two cases $\{x < q^H(\eta)\}$ and $\{x \geq q^H(\eta)\}$. Again, we find that in both cases,

$$\begin{aligned} \mu_\eta = O(1) \quad \text{and} \quad |\sigma_\eta|^2 = O(1) \\ \mu_x = (-1) \times o((q^H_\beta - x)^{-1}) \quad \text{and} \quad |\sigma_x|^2 = O(1). \end{aligned}$$

We used condition (B.9) to obtain the order of μ_x. Thus,

$$\mathscr{L}^{\eta,x}v = (q^H_\beta - x)^{-2}x\mu_x + O((q^H_\beta - x)^{-3}) \to -\infty.$$

Step 6e: boundary as $(\eta, x) \to (0, q^L_\alpha(0))$. Finally, all the corners of $\mathcal{X}$ can be analyzed in a straightforward way by combining the cases above, with the exception of $(\eta, x) = (0, q^L_\alpha(0)) = (0, a_h/\rho_h)$. Approaching this corner, we must take a particular path of $x \to a_h/\rho_h$ as $\eta \to 0$. Denote this path by $\hat{x}(\eta)$ and denote the asymptotic slope by $\hat{x}'(0) \in (\frac{d}{d\eta}q^L_\alpha(0), +\infty)$, where $\frac{d}{d\eta}q^L_\alpha(0) = [\frac{a_e}{a_h} - \frac{\rho_e}{\rho_h}]\frac{a_h}{\rho_h} + \alpha'(0) > 0$, by Assumption 1, part (i), and the fact that $\alpha'(0) > 0$. Denote the associated path of κ by $\hat{\kappa}(\eta)$ and the corresponding asymptotic slope by $\hat{\kappa}'(0) = \frac{1}{a_e - a_h}[\hat{x}'(0)\rho_h + (\rho_e - \rho_h)a_h/\rho_h]$. Substituting in, we find $\hat{\kappa}'(0) \in (1 + \frac{\alpha'(0)}{a_e - a_h}, +\infty)$. When computing $\mathscr{L}^{\eta,x}v$, we will take the supremum over all possible paths, meaning over $\hat{x}'(0)$ and $\hat{\kappa}'(0)$. Using similar calculations from the initial $\eta \to 0$ case, but using these paths, we obtain

$$\mu_\eta = \delta_h + \eta[\frac{a_e - a_h}{\hat{x}}\hat{\kappa}' + \rho_h - \rho_e - \delta_e - \delta_h] + o(\eta) \quad \text{and} \quad |\sigma_\eta|^2 = \eta^2[\hat{\kappa}' - 1]\frac{a_e - a_h}{\hat{x}} + o(\eta)$$
$$\mu_x = o((\hat{x} - q^L_\alpha)^{-1}) \quad \text{and} \quad |\sigma_x|^2 = O(1)$$
$$\text{and} \quad \sigma_x \cdot \sigma_\eta = \eta[\frac{a_e - a_h}{\hat{x}} - \sigma(\vartheta(\hat{\kappa}' - 1)\frac{a_e - a_h}{\hat{x}})^{1/2}] + o(\eta).$$

Since $\hat{x} \geq O(\eta)$ and $\hat{\kappa} \geq O(\eta)$ (in the sense that both could be $+\infty$), we may treat terms like $(\hat{x} - q^L_\alpha)^{-1}$ as smaller than η^{-1}, asymptotically. This identifies the dominant terms as those associated to μ_η, $|\sigma_\eta|^2$, and μ_x. Thus,

$$\mathscr{L}^{\eta,x}v = -\frac{1}{2\eta^{3/2}}\delta_h + \frac{1}{2\eta^{1/2}}[\rho_e - \rho_h + \delta_e + \delta_h - \frac{a_e - a_h}{\hat{x}} - \frac{a_e - a_h}{\hat{x}}(\hat{\kappa}' - 1)/4] + o(\eta^{-3/2})$$
$$- (\hat{x} - q^L_\alpha)^{-2}x\mu_x + O((\hat{x} - q^L_\alpha)^{-3}) \to -\infty,$$

irrespective of δ_h, because $\rho_e - \rho_h - \frac{a_e - a_h}{a_h/\rho_h} = \rho_h[\rho_e/\rho_h - a_e/a_h] < 0$ by Assumption 1, part (i), and because $\inf\{\hat{\kappa}'(0)\} > 1$.

This completes the verification that $\mathscr{L}^{\eta,x}v \to -\infty$ as $(\eta, x) \to \partial\mathcal{X}$, which proves stationarity by Lemma B.1 below. This proves that the construction above is an equilibrium.

Step 7: Indexing the equilibria. We conclude by summarizing the indeterminacies, corresponding to the "equilibrium indexing" referenced in the statement of the theorem:

(i) In Step 3, we allowed for an arbitrary share $\vartheta(\eta, x, y)$ of capital return variance $|\sigma_R|^2$ to arise from the fundamental shock $Z^{(1)}$, in the region $\{x < q^H\}$. Since $q = x$

in this region, we write this function as $\vartheta(\eta, q, y)$ in the theorem's statement.

(ii) In Step 4, we allowed for an arbitrary drift $\mu_q(\eta, x, y) = m_q(\eta, x, y)$ in the region $\{x < q^H\}$, subject to boundary conditions that hold at $\partial\mathcal{X}$. Note that this implies μ_q can be any function of (η, q, y) almost everywhere in the interior of $\mathcal{D}$.

(iii) In Step 5, we allowed for arbitrary dynamics $\sigma_x(\eta, x, y) = \tilde{\sigma}_x(\eta, x, y)$ and $\mu_x(\eta, x, y) = m_x(\eta, x, y)$ in the region $\{x \geq q^H\}$. Together, these dynamics can be chosen to engineer an exit rate from $\{x \geq q^H\} = \{\kappa = 1\}$, in the following sense. Given $(\eta_0, x_0, y_0) = (\eta, x, y)$, define the first passage time to inefficiency by $\tau^\circ := \inf\{t > 0 : x_t < q^H(\eta_t)\}$ and its expectation $T(\eta, x, y) := \mathbb{E}[\tau^\circ \mid (\eta_0, x_0, y_0) = (\eta, x, y)]$. Then, $T(\eta, x, y)$ is determined by solving the PDE[27]

$$\mathscr{L}^{\eta,x,y} T = -1 \quad \text{on} \quad \{(\eta, x) : x > q^H(\eta)\} \times \mathcal{Y} \quad \text{s.t.} \quad T(\eta, q^H(\eta)) = 0 \tag{B.16}$$

By choosing $\tilde{\sigma}_x(\eta, x, y)$ and $m_x(\eta, x, y)$, we can thus obtain various solutions to the PDE (B.16), hence various exit rates.

This completes the proof. □

B.4 Stochastic stability: a useful lemma

To prove the stationarity claims of Theorem 1 and Proposition D.1, we need the following lemma, which is a slight generalization of Theorems 3.5 and 3.7 of Khasminskii (2011), in the sense that weaker conditions are imposed on the coefficients α and β. Indeed, any coefficients (α, β) are permissible as long as they admit existence of a weak solution to the SDE system. The other generalization is that we allow the domain to be any open domain $\mathcal{D}$ rather than $\mathbb{R}^l$ (see also Remark 3.5 and Corollary 3.1 in Khasminskii (2011)).

Lemma B.1. *Suppose $(X_t)_{0 \leq t \leq \tau}$ is a weak solution to the SDE $dX_t = \beta(X_t)dt + \alpha(X_t)dZ_t$ in an open connected domain $\mathcal{D} \subset \mathbb{R}^l$, where Z is a d-dimensional Brownian motion and $\tau := \inf\{t : X_t \notin \mathcal{D}\}$ is the first exit time from $\mathcal{D}$. Define the infinitesimal generator $\mathscr{L}$ by (for any C^2 function f)*

$$\mathscr{L} f = \sum_{i=1}^{n} \beta_i \frac{\partial f}{\partial x_i} f + \frac{1}{2} \sum_{i,j=1}^{n} (\alpha_i \cdot \alpha_j) \frac{\partial^2 f}{\partial x_i \partial x_j}.$$

Suppose there is a non-negative C^2 function $v : \mathcal{D} \mapsto \mathbb{R}_+$ such that (i) $\liminf_{x \to \partial\mathcal{D}} v(x) = +\infty$; (ii) $\mathscr{L} v \leq cv$ for some constant $c \geq 0$; and (iii) $\limsup_{x \to \partial\mathcal{D}} \mathscr{L} v(x) = -\infty$. Then,

[27] This standard PDE is a consequence of the Feynman-Kac theorem.

(a) $\tau = +\infty$ *almost-surely;*

(b) the distribution of X_0 *can be chosen such that* $(X_t)_{t\geq 0}$ *is stationary.*

Proof of Lemma B.1. Let $\{\mathcal{D}_n\}_{n\geq 1}$ be an increasing sequence of open sets, whose closures are contained in $\mathcal{D}$, such that $\cup_{n\geq 1}\mathcal{D}_n = \mathcal{D}$. Let $\tau_n := \inf\{t : X_t \notin \mathcal{D}_n\}$, and note that $\tau = \lim_{n\to\infty}\tau_n$ is the monotone limit of these exit times. Define $w(t,x) := v(x)\exp(-ct)$, which satisfies $\mathscr{L}w \leq 0$ by assumption (ii). Using Itô's formula, we have

$$\mathbb{E}[v(X_{\tau_n\wedge t})e^{-c(\tau_n\wedge t)} - v(X_0)] = \mathbb{E}\int_0^{\tau_n\wedge t}\mathscr{L}w(u,X_u)du \leq 0.$$

Since $(\tau_n \wedge t) \leq t$ and $v \geq 0$, we obtain

$$\mathbb{E}[v(X_{\tau_n\wedge t})] \leq e^{ct}\mathbb{E}[v(X_0)].$$

Because $\mathbb{E}[v(X_{\tau_n\wedge t})] \geq \mathbb{P}[\tau_n \leq t]\inf_{x\in\mathcal{D}\setminus\mathcal{D}_n} v(x)$, we thus have

$$\mathbb{P}[\tau_n \leq t] \leq \frac{e^{ct}\mathbb{E}[v(X_0)]}{\inf_{x\in\mathcal{D}\setminus\mathcal{D}_n} v(x)}.$$

Taking the limit $n \to \infty$, we obtain

$$\mathbb{P}[\tau \leq t] \leq \frac{e^{ct}\mathbb{E}[v(X_0)]}{\liminf_{x\to\partial\mathcal{D}} v(x)} = 0.$$

Thus, taking $t \to \infty$, we prove (a).

Next, since $\tau = +\infty$ a.s., we may consider $(X_t)_{t\geq 0}$ that is now defined for all time. Using Itô's formula,

$$\mathbb{E}[v(X_{\tau_n\wedge t}) - v(X_0)] = \mathbb{E}\int_0^{\tau_n\wedge t}\mathscr{L}v(X_u)du.$$

Note that $\min(\inf_t \mathbb{E}[v(X_t) - v(X_0)], \inf_n \mathbb{E}[v(X_{\tau_n}) - v(X_0)]) \geq b_1$ for some constant b_1, given assumption (i) and $v \geq 0$. Also note that $\sup_{x\in\mathcal{D}}\mathscr{L}v(x) \leq b_2$ for some constant b_2, given assumptions (i)-(iii) and the fact that v is C^2. (b_1 and b_2 are both independent of t and n.) Using these bounds, plus the following obvious inequality

$$\mathscr{L}v(X_u) \leq \mathbf{1}_{\{X_u\in\mathcal{D}\setminus\mathcal{D}_k\}}\sup_{x\in\mathcal{D}\setminus\mathcal{D}_k}\mathscr{L}v(x) + \sup_{x\in\mathcal{D}}\mathscr{L}v(x),$$

we get

$$-\sup_{x\in\mathcal{D}\setminus\mathcal{D}_k}\mathscr{L}v(x)\mathbb{E}\int_0^{\tau_n\wedge t}\mathbf{1}_{\{X_u\in\mathcal{D}\setminus\mathcal{D}_k\}}du\leq tb_2-b_1.$$

Given the proof of (a), we may take the limit $n\to\infty$ (so that $\tau_n\to+\infty$), then apply Fubini's theorem, and then rearrange to obtain

$$\lim_{t\to\infty}\frac{1}{t}\int_0^t\mathbb{P}[X_u\in\mathcal{D}\setminus\mathcal{D}_k]du\leq\frac{b_2}{-\sup_{x\in\mathcal{D}\setminus\mathcal{D}_k}\mathscr{L}v(x)}.$$

Taking $k\to\infty$ and using assumption (iii), we obtain

$$\lim_{k\to\infty}\lim_{t\to\infty}\frac{1}{t}\int_0^t\mathbb{P}[X_u\in\mathcal{D}\setminus\mathcal{D}_k]du\leq 0.$$

Applying Theorem 3.1 of Khasminskii (2011), there exists a stationary initial distribution for X_0. The process $(X_t)_{t\geq 0}$ augmented with this initial distribution is clearly stationary by definition. □

B.5 Proof of Corollary 1

Start from the construction of S-BSE in Theorem 1, and note that we can make ϵ_α arbitrarily small such that the lower boundary converges to its lowest possible level: $q_\alpha^L\to\bar{a}/\bar{\rho}$. Hence, an S-BSE can be constructed such that the set of prices q matches $\mathcal{Q}(\eta)$ arbitrarily closely. The result on the minimal return variance comes from the following two objects: (i) take the limit $\kappa\to 1$ in (B.6) to obtain $|\sigma_R|^2=\eta\bar{\rho}(\eta)\frac{a_e-a_h}{a_e}$; (ii) use (B.7) to obtain $|\sigma_R|^2=\sigma^2\frac{\bar{\rho}(\eta)^2}{\rho_e^2}$. When $\eta\geq\eta^*$, it is clear that the minimal variance is the smaller of (i)-(ii). When $\eta<\eta^*$, we cannot have $\kappa=1$, so the minimal variance is simply the result from (i); however, the equation for η^* in (B.4) can be rearranged to show that $\eta<\eta^*$ is equivalent to $\eta\bar{\rho}(\eta)\frac{a_e-a_h}{a_e}<\frac{\bar{\rho}(\eta)^2}{\rho_e^2}\sigma^2$. Finally, the form of $\mathcal{V}$ being at most two intervals comes from the fact that the variance when $\kappa<1$ is $|\sigma_R|^2=\frac{\eta(1-\eta)}{\kappa-\eta}\frac{a_e-a_h}{q}$, which is continuous in q. □

B.6 Proof of Theorem 2

Step 0: Write the equilibrium equations and domain. To prove Theorem 2, we first document all the equilibrium equations with investment and sunspot jumps. First, note that optimal consumption is still (due to log utility) $c_{j,t}=\rho_j n_{j,t}$. Second, recall that all agents make the same scaled investment decisions, namely $\iota_t=\iota(q_t):=(\Phi')^{-1}(q_t)$. Putting

these facts together, we immediately obtain the price-output relation (PO-inv) from the goods market clearing condition, repeated here for reference:

$$\bar{\rho}(\eta_t)q_t + \Phi(\iota(q_t)) = \kappa_t a_e + (1-\kappa_t)a_h. \tag{PO-inv}$$

At this stage, we can define the equilibrium domain for (η, q) as

$$\mathcal{D} := \{(\eta, q) : 0 < \eta < 1,\, q^L(\eta) < q \leq q^H(\eta)\},$$

where the minimal permissible capital price q^L is such that $\kappa = \eta$, while the maximal capital price q^H is such that $\kappa = 1$. Using (PO-inv), these two extremal prices are

$$\begin{aligned} q^L(\eta) &:= \{q : \bar{\rho}(\eta)q + \Phi(\iota(q)) = \eta a_e + (1-\eta)a_h\} \\ q^H(\eta) &:= \{q : \bar{\rho}(\eta)q + \Phi(\iota(q)) = a_e\}. \end{aligned}$$

Since Φ is strictly increasing and strictly convex, we have that $\Phi(\iota(q)) = \Phi((\Phi')^{-1}(q))$ is strictly increasing in q, hence the extremal price functions q^L, q^H are unique.

Next, recall that each agent, via their investment decisions, obtains a capital growth rate of $G(q) := g + \iota(q) - \delta$. Also recall that our jumps ℓ_q are assumed to occur randomly but have a known size, given observables. Therefore, optimal portfolio FOCs are

$$\begin{aligned} \frac{a_e - \Phi(\iota(q))}{q} + G(q) + \mu_q + \sigma\left(\begin{smallmatrix}1\\0\end{smallmatrix}\right)\cdot\sigma_q - r &= \frac{\kappa}{\eta}|\sigma_R|^2 + \frac{\lambda\ell_q}{1-\frac{\kappa}{\eta}\ell_q} \\ \frac{a_h - \Phi(\iota(q))}{q} + G(q) + \mu_q + \sigma\left(\begin{smallmatrix}1\\0\end{smallmatrix}\right)\cdot\sigma_q - r &\leq \frac{1-\kappa}{1-\eta}|\sigma_R|^2 + \frac{\lambda\ell_q}{1-\frac{1-\kappa}{1-\eta}\ell_q}. \end{aligned}$$

Combining these two equations, we obtain (RBJ), restated here for reference:

$$0 = \min\left[1-\kappa,\, \frac{a_e - a_h}{q} - \frac{\kappa-\eta}{\eta(1-\eta)}\left(|\sigma_R|^2 + \frac{\lambda\ell_q^2}{\left(1-\frac{\kappa}{\eta}\ell_q\right)\left(1-\frac{1-\kappa}{1-\eta}\ell_q\right)}\right)\right]. \tag{RBJ}$$

We can determine the other equilibrium objects similarly to before. The riskless rate is given by, after aggregating the two Euler equations with weights κ and $1-\kappa$, and then using the price-output relation (PO-inv) to replace $\frac{\kappa a_e + (1-\kappa)a_h - \Phi(\iota(q))}{q} = \bar{\rho}$,

$$r = \bar{\rho} + G(q) + \mu_q + \sigma\left(\begin{smallmatrix}1\\0\end{smallmatrix}\right)\cdot\sigma_q - \left(\frac{\kappa^2}{\eta} + \frac{(1-\kappa)^2}{1-\eta}\right)|\sigma_R|^2 - \lambda\ell_q\left(\frac{\kappa}{1-\frac{\kappa}{\eta}\ell_q} + \frac{1-\kappa}{1-\frac{1-\kappa}{1-\eta}\ell_q}\right).$$

The dynamics of η are now given by $d\eta_t = \mu_{\eta,t-}dt + \sigma_{\eta,t-} \cdot dZ_t - \ell_{\eta,t-}dJ_t$, where

$$\mu_\eta = \eta(1-\eta)(\rho_h - \rho_e) + (\kappa - 2\eta\kappa + \eta^2)\frac{\kappa - \eta}{\eta(1-\eta)}|\sigma_R|^2 + \delta_h - (\delta_e + \delta_h)\eta + \frac{(\kappa-\eta)\lambda\ell_q}{\left(1 - \frac{\kappa}{\eta}\ell_q\right)\left(1 - \frac{1-\kappa}{1-\eta}\ell_q\right)}$$

$$\sigma_\eta = (\kappa - \eta)\sigma_R.$$

The wealth share jump ℓ_η is derived by using knowledge of the jump size in q and noting that agents' portfolios (capital and bonds) are predetermined:[28]

$$\ell_\eta = (\kappa - \eta)\frac{\ell_q}{1 - \ell_q}.$$

This is the full set of equilibrium equations. We also have a set of exogenous state variables $y_t \in \mathcal{Y}$, exactly as in Theorem 1, but here we also allow those exogenous state variables to jump.

Step 1: Restrictions on jump size. What remains are restrictions on the jump size so that (η_t, q_t) remains in $\mathcal{D}$ forever. First, jumps cannot be so large as to send experts into bankruptcy ($\eta \leq 0$), nor can they induce households' leverage to exceed experts' ($\kappa \leq \eta$). These two conditions are

$$\ell_q < \frac{\eta}{\kappa} \tag{B.17}$$

$$\bar{\rho}(\hat{\eta})(1-\ell_q)q + \Phi(\iota((1-\ell_q)q)) > (a_e - a_h)\hat{\eta} + a_h, \tag{B.18}$$

where $\hat{\eta} := \eta - (\kappa - \eta)\frac{\ell_q}{1-\ell_q}$ is the post-jump expert wealth share. Although it is obvious, (RBJ) implies another bound on ℓ_q that arises because of $|\sigma_R| \geq 0$. This implies two different bounds, depending on whether $\kappa < 1$ or $\kappa = 1$:

$$\frac{a_e - a_h}{q} \geq \frac{\kappa - \eta}{\eta(1-\eta)}\frac{\lambda\ell_q^2}{\left(1 - \frac{\kappa}{\eta}\ell_q\right)\left(1 - \frac{1-\kappa}{1-\eta}\ell_q\right)}, \quad \text{if } \kappa < 1 \tag{B.19}$$

$$\frac{a_e - a_h}{q^H} \geq \frac{1}{\eta}\left(\sigma^2 + \frac{\lambda\ell_q^2}{1 - \frac{1}{\eta}\ell_q}\right), \quad \text{if } \kappa = 1 \tag{B.20}$$

[28] The derivation is as follows. Let variables with hats, e.g., "$\hat{x}$", denote post-jump variables. Note $\hat{N}_e = \hat{q}\hat{K}\kappa - B$ and $\hat{N}_h = \hat{q}\hat{K}(1-\kappa) + B$, where B is expert borrowing (and household lending, by bond market clearing). Then, $\hat{\eta} = \hat{N}_e/(\hat{q}\hat{K}) = \kappa - B/(\hat{q}\hat{K})$ and by similar logic the pre-jump wealth share is $\eta = \kappa - B/qK$. Thus, $\ell_\eta = \eta - \hat{\eta} = B[1/(\hat{q}\hat{K}) - 1/(qK)] = qK(\kappa - \eta)[1/(\hat{q}\hat{K}) - 1/(qK)]$. Using the fact that $\hat{K} = K$ and the definition $\ell_q := 1 - \hat{q}/q$, we arrive at $\ell_\eta = (\kappa - \eta)[(1-\ell_q)^{-1} - 1]$. This derivation assumes the presumably risk-free bond price does not jump when capital prices jump. Conceptually, there is no reason why this needs to be true, but it preserves its risk-free conjecture.

We next show that a family of solutions ℓ_q exist satisfying the constraints (B.17)-(B.20).

Note that, fixing (η, q, λ), the RHS of (B.19) is strictly increasing in ℓ_q when $\ell_q \in (0, \frac{\eta}{\kappa})$ while the LHS is constant. Moreover, the inequality is satisfied for $\ell_q = 0$ and violated as $\ell_q \to \frac{\eta}{\kappa}$. Hence, this condition, combined with requirement (B.17), imply that $\ell_q \leq \ell_q^{A,\circ}(\eta, q, \lambda)$, where the upper bound $\ell_q^{A,\circ}(\eta, q, \lambda)$ is the unique value of $\ell_q \in (0, \frac{\eta}{\kappa})$ that makes (B.19) an equality.

On the other hand, the RHS of (B.20) is strictly increasing in ℓ_q when $\ell_q \in (0, \eta)$ while the LHS is constant. Moreover, the inequality is satisfied for $\ell_q = 0$ and violated as $\ell_q \to \eta$. Hence, condition (B.20), combined with requirement (B.17), imply that $\ell_q \leq \ell_q^{A,1}(\eta, \lambda)$, where the upper bound $\ell_q^{A,1}(\eta, \lambda)$ is the unique value of $\ell_q \in (0, \eta)$ that makes (B.20) an equality.[29] Putting these two together, define

$$\ell_q^A(\eta, q, \lambda) := \begin{cases} \ell_q^{A,\circ}(\eta, q, \lambda), & \text{if } q < q^H(\eta); \\ \ell_q^{A,1}(\eta, \lambda), & \text{if } q = q^H(\eta). \end{cases}$$

Thus, if we pick $\ell_q < \ell_q^A$, then (B.17) and (B.19)-(B.20) hold. Note that $\ell_q^A > 0$ on $\mathcal{D}$, so a range of jump sizes satisfy this constraint, for each (η, q, λ).

Next, after some algebra, we can write condition (B.18) as

$$(1-\ell_q)^2 - (1-\ell_q)\Big(\frac{\bar{\rho}(\eta)q + q(\rho_e - \rho_h)(\kappa - \eta) + \Phi(\iota(q)) - \Phi(\iota((1-\ell_q)q))}{\bar{\rho}(\eta)q + q(\rho_e - \rho_h)(\kappa - \eta)}\Big) + \frac{(a_e - a_h)(\kappa - \eta)}{\bar{\rho}(\eta)q + q(\rho_e - \rho_h)(\kappa - \eta)} > 0.$$

For each (η, q), define $\ell_q^B(\eta, q) > 0$ as the smallest positive root of equating the LHS to zero, if such root exist, and let $\ell_q^B = 1$ otherwise. Given that inequality (B.18) holds for $\ell_q = 0$ (because for $\ell_q = 0$, it is equivalent to $\kappa > \eta$, which holds on $\mathcal{D}$), it holds for all $\ell_q < \ell_q^B$. Again, since $\ell_q^B > 0$, a range of jump sizes satisfy this constraint, for each (η, q, λ).

Putting the results together, we can construct an upper bound that ensures all required inequalities are satisfied. At this point, we pick $l(\eta, q, y)$ to be any function and

[29] The closed-form expression for this bound is

$$\ell_q^{A,1}(\eta, \lambda) = \frac{1}{2}\Big[-\frac{M}{\lambda\eta} + \sqrt{\Big(\frac{M}{\lambda\eta}\Big)^2 + 4\frac{M}{\lambda}}\Big], \quad \text{where} \quad M(\eta) := \max[0, \eta\frac{a_e - a_h}{q^H(\eta)} - \sigma^2]$$

put the jump arrival rate by $\lambda = l$. Then, we define

$$\ell_q^{max}(\eta, q, y) := \min\{\ell_q^A(\eta, q, l(\eta, q, y)), \ell_q^B(\eta, q)\}. \tag{B.21}$$

Thus, as claimed in the theorem, there exist an entire family of jump sizes: we may pick any non-negative function h such that $h(\eta, q, y) \in [0, \ell_q^{max}(\eta, q, y))$, and we may set $\ell_q = h$ for all $y \in \mathcal{Y}$ (exogenous state variables) and all $(\eta, q) \in \mathcal{D}$. However, to permit a treatment of the boundary conditions on the dynamics, we restrict $h \equiv 0$ for all (η, q) close enough to the $\eta = 0$ and $\kappa = \eta$ boundaries. Hence, this is why the theorem allows us to choose any such function h only on the subset $\mathcal{D}^\circ$.

The remainder of the proof is identical to Theorem 1, and we therefore only sketch it. In particular, we compute the determinate equilibrium objects from the equations above, while setting the indeterminate ones according to exogenous functions. We then check that the dynamics of the equilibrium system are such that (η, q) remain inside the domain $\mathcal{D}$ forever.

Step 2: Compute other equilibrium objects. In the interior of $\mathcal{D}$, we proceed as follows. Use risk-balance condition (RBJ) to solve for $|\sigma_R|^2$, given ℓ_q. Note that this is pinned down uniquely, given ℓ_q. Next, for each (η, q, y), assign $\vartheta(\eta, q, y)$ fraction of the variance to the fundamental Brownian shock, and $1 - \vartheta(\eta, q, y)$ to the sunspot Brownian shock, where $\vartheta \in [0, 1]$ is an exogenous function. Then, solve for other equilibrium objects from the equations above. In this process, we have freedom to set r arbitrarily, but μ_q is pinned down given this choice for r, similar to the S-BSE.

At the top boundary of $\mathcal{D}$ (i.e., when $\kappa = 1$), the economy features no jumps and therefore all objects are identical to those established in Theorem 1 (see Steps 3-4 of the proof). Note that, although this economy with investment features an endogenous growth rate (i.e., $G(q)$ replaces g), it does not enter these expressions.

Step 3: Dynamics at the boundaries. At the top boundary of $\mathcal{D}$ (i.e., when $\kappa = 1$), there is an indeterminacy in the speed at which the economy re-enters the interior and $\kappa < 1$. This can be treated exactly as in Theorem 1, where we "perturb" the top boundary and pick the dynamics of a latent diffusion that governs how fast (η, q) re-enters the interior of $\mathcal{D}$. Near the lower boundary of $\mathcal{D}$ (i.e., when $\kappa \approx \eta$), we can also follow Theorem 1. In particular, given the jumps are shut down near the left and lower boundaries (i.e., when $\eta \approx 0$ or $\kappa \approx \eta$), the behavior there is entirely diffusive, implying the proof of non-accessibility of the boundary is identical to Theorem 1.

C Additional details and results for quantitative exercises

C.1 More details on the baseline model comparisons

Exit rate construction. As mentioned in Section 3.1, we choose a specification at the $\{\kappa = 1\}$ boundary such that the economy spends approximately 10% of its time there. To achieve this, we implement the following specification. We perturb our domain, as in the proof of Theorem 1 (see Figure B.1 for an illustration). In this perturbation, the state variables are actually (η_t, x_t), where x is a latent variable that is allowed to exceed $q^H(\eta) := a_e/\bar{\rho}(\eta)$. We then define $q_t := \min[x_t, q^H(\eta_t)]$. Thus, we control the fraction of time $\kappa_t = 1$ by specifying the dynamics of x_t on the set $\{x_t > q^H(\eta_t)\}$.

For this, we implement the following. First, the top boundary for x is $q^H(\eta) + \beta(\eta)$, where $\beta(\eta) := \max[0, 0.1\frac{\eta^*}{1-\eta^*}\frac{a_e}{\rho_e}]$, where η^* is defined in equation (B.4). This ensures that the top boundary is a line in (η, x) space from $(\eta^*, q^H(\eta^*))$ to $(1, 1.1q^H(1))$. Second, we specify x to have random walk dynamics in the region $\{x > q^H(\eta)\}$, i.e., $dx_t = \sigma_x dZ_t^{(1)}$ with $\sigma_x = 1$ (in sensitivity analyses, we reduce σ_x to increase the fraction of time spent in the efficient region). Third, we implement a reflecting barrier at the top boundary.

Stationary CDFs. Figures C.1-C.2 are CDFs of various equilibrium objects in the two equilibria compared in Section 3.2. In particular, Figure C.1 shows that risk premia and volatility feature a significant right tail in the sentiment equilibrium, much more so than the fundamental equilibrium. Similarly, capital prices and the capital allocation are significantly lower on average in the sentiment equilibrium. By contrast, Figure C.2 shows the same CDFs when the sentiment equilibrium features much higher bounce-back beliefs. In that case, there is no longer a long right tail for risk premia or volatility, and neither capital prices nor the capital allocation ever fall too low.

The marginal distributions of equilibrium variables in the S-BSE and FE, plotted in Figure C.1, highlight the key differences. Compared to the FE, the S-BSE can attain significantly higher risk premia and volatility, precisely what permits the spikes in crisis event studies. This is despite featuring a smoother interest rate and a similar marginal distribution of η.

C.2 Crisis event studies under alternative specifications

This section displays the results from various robustness exercises. In particular, we redo the financial crisis event studies from Section 3.2 under various alternative specifications.

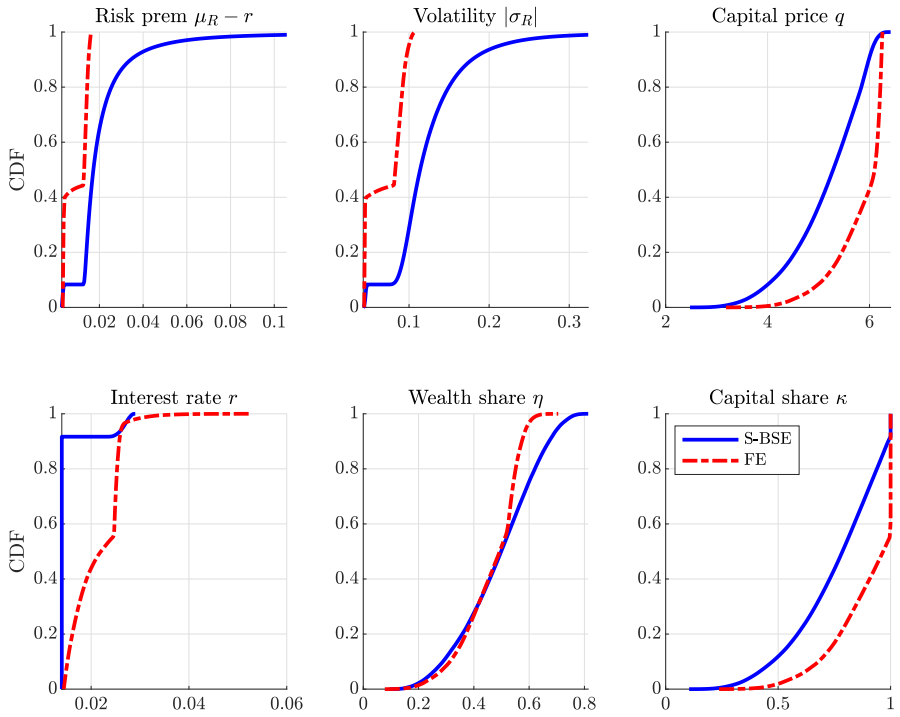


Figure C.1: Stationary marginal CDFs of several objects in the S-BSE. Parameters are the same as Figure 6.

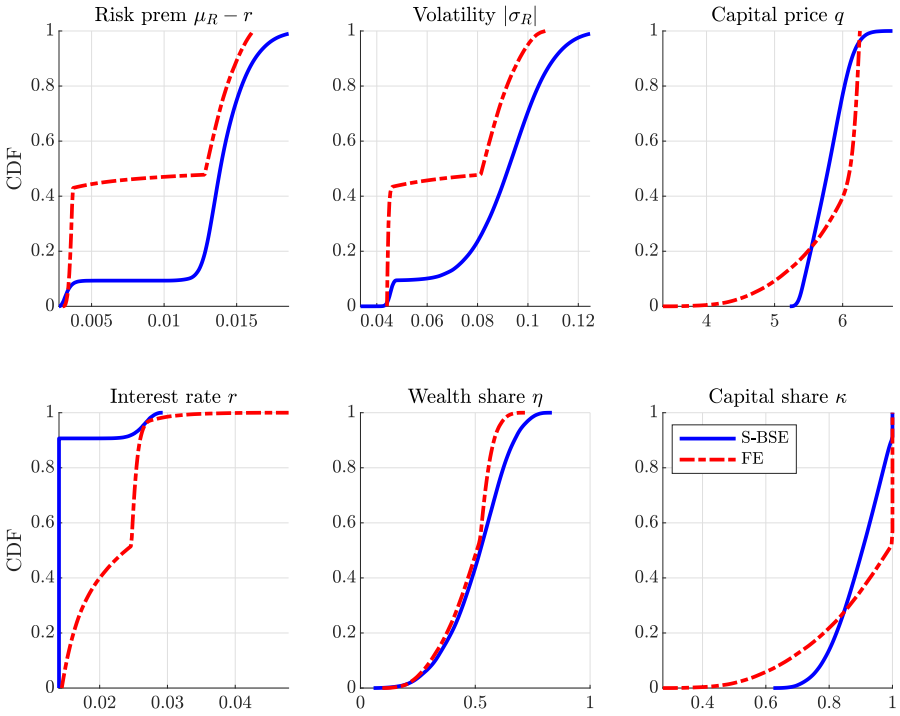


Figure C.2: Stationary marginal CDFs of several objects in the S-BSE *with higher bounce-back beliefs*. The bounce-back belief is a lower reflecting barrier at $q^L(\eta) + 0.6(1-\eta)\frac{a_e - a_h}{\bar{\rho}(\eta)}$. Parameters are the same as Figure 6.

Define crises based on expert wealth. Empirically, financial crises are not defined based on large output drops. Some authors define financial crises based on narrative sources on fire sales and panics (Reinhart and Rogoff, 2009), while others use bank equity declines (Baron et al., 2021). We perform robustness on our baseline by defining crises instead as the bottom 3rd percentile of year-to-year declines in $\log(\eta)$, which we interpret as a proxy for bank equity. Figure C.3 displays the results. One difference from the baseline specification is the presence of a much sharper decline in η. This sharper behavior is to be expected, given the crisis is defined based on η itself. But broadly speaking, the crisis dynamics are similar to the baseline case: the S-BSE still delivers sudden and severe crises with pre-crisis froth, while the FE does not.

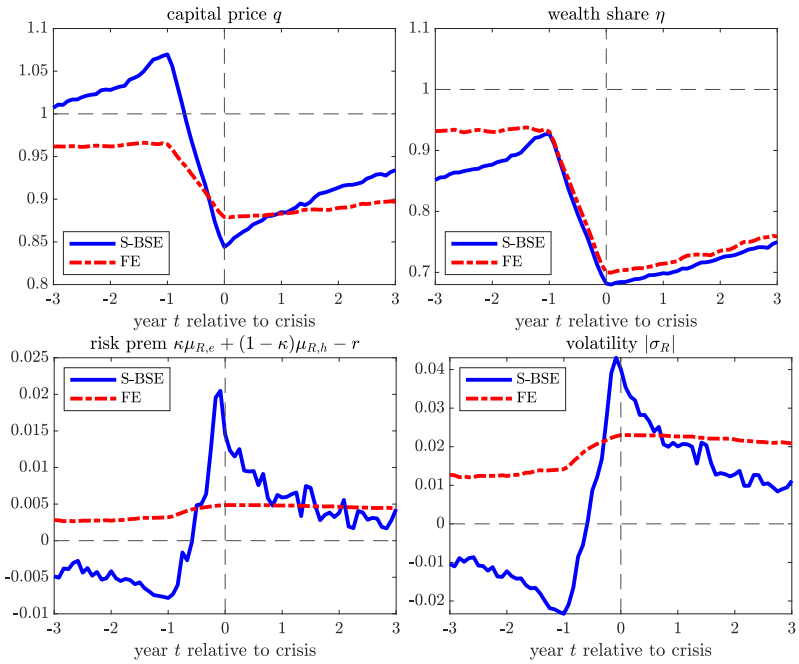


Figure C.3: Event studies around financial crises *defined by expert wealth*. Crises are defined as the bottom 3rd percentile of year-to-year declines in $\log(\eta)$, subject to only one crisis occurring in a 6-year window. All other details are identical to Figure 6.

Shut down sunspot shocks. Now, we perform a particularly stark exercise: we eliminate the sunspot shock $Z^{(2)}$ altogether. The baseline model set $\vartheta = \sigma_R^{(1)}/|\sigma_R| = \sqrt{0.5}$, so that 50% of return volatility was due to the fundamental shock and 50% due to the sunspot shock. Here, we instead set $\vartheta = 1$ so that sunspot shocks contribute nothing. Nevertheless, as Figure C.4 shows, financial crises behave broadly similar to the baseline case and are still much sharper than the FE. The key is that the equilibrium still occasionally visits parts of (η, q) space that are more extreme than the FE, in particular low values of q where volatility is very high. In fact, it turns out that the stationary distribution of

(η, q) is very similar to the baseline specification.

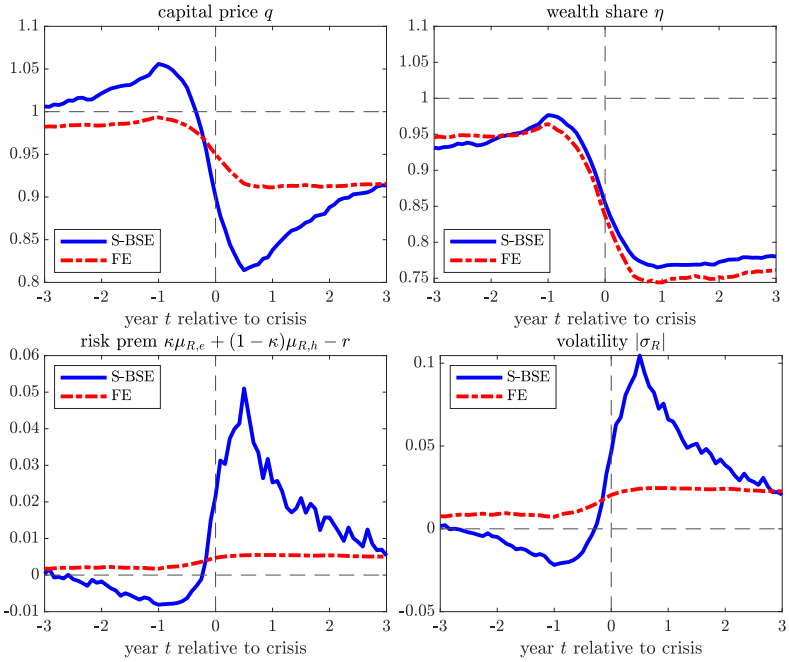


Figure C.4: Event studies around financial crises *without sunspot shocks*. The fraction of return variance $|\sigma_R|^2$ from the sunspot shock is recalibrated to be $\vartheta = 1$. Crises are defined as the bottom 3rd percentile of year-to-year log output declines, subject to only one crisis occurring in a 6-year window. All other details are identical to Figure 6.

Slower exit rate from the efficient region. Recall that the speed of exit from the efficient region with $\kappa = 1$ is not pinned down. In the baseline model, we targeted a 10% unconditional probability of efficiency (Appendix C.1 explains how we implement such a target). Here, we increase this to 30% instead, much closer to the amount of efficiency that arises in the FE. Figure C.5 displays the marginal CDFs from this economy: despite the fact that the economy is efficient significantly more often, there remains a non-trivial chance of extreme volatility and risk premia. Consequently, the crisis event studies, displayed in Figure C.6, are broadly similar to the baseline case.

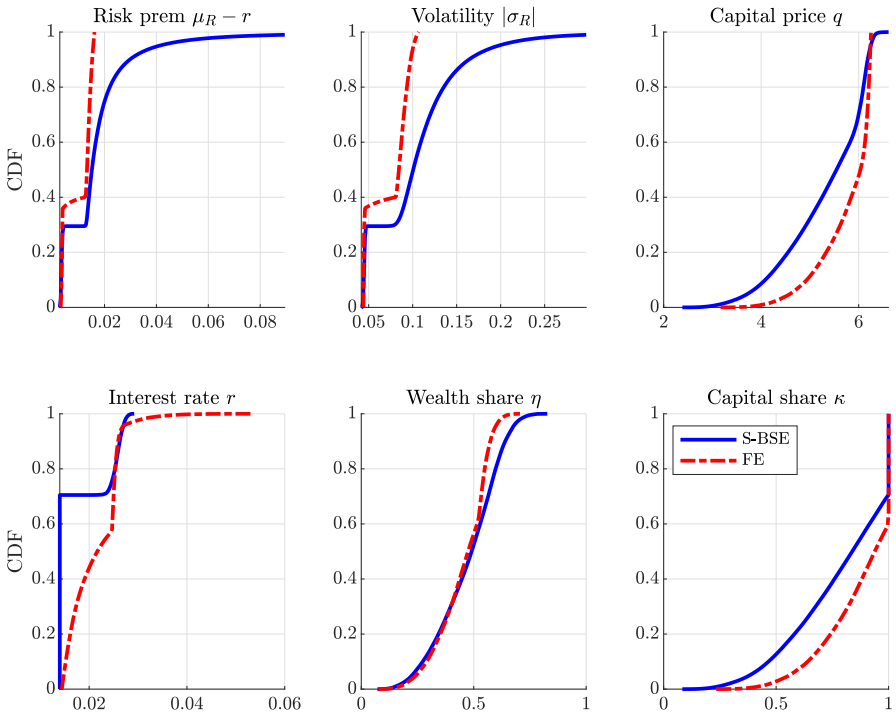


Figure C.5: Stationary marginal CDFs of several objects in the S-BSE *with stickier efficiency*. The exit rate from the efficient region is recalibrated so that $\mathbb{P}\{\kappa_t = 1\} \approx 0.3$ (the baseline probability was 0.1). Other parameters are the same as Figure 6.

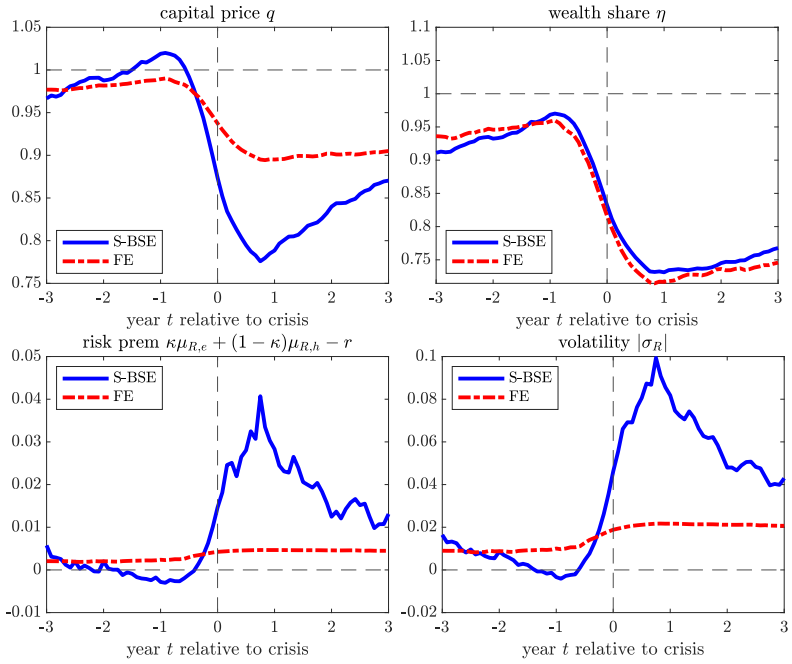


Figure C.6: Event studies around financial crises *with stickier efficiency*. The exit rate from the efficient region is recalibrated so that $\mathbb{P}\{\kappa_t = 1\} \approx 0.3$ (the baseline probability was 0.1). Crises are defined as the bottom 3rd percentile of year-to-year log output declines, subject to only one crisis occurring in a 6-year window. All other details are identical to Figure 6.

Recalibrate interest rate process. Recall that our S-BSE constructions are such that, away from the boundaries, the interest rate follows the exogenous Ornstein-Uhlenbeck process

$$dr_t = \lambda_r(\bar{r} - r_t)dt + \sigma_r \begin{pmatrix} \theta \\ \sqrt{1-\theta^2} \end{pmatrix} \cdot dZ_t. \tag{C.1}$$

In the baseline model, we set $\sigma_r = 0$, so that r deterministically drifts towards its long-run mean $\bar{r}$. Here, we pick these parameters to match the unconditional mean (0.014), variance (0.023^2), and annual autocorrelation (0.94) of the 3-month US real rate. In particular, we set $\bar{r} = 0.014$, $\lambda_r = -\log(0.94)$, and $\sigma_r = |2\log(0.94) \times 0.023^2|$. We perform sensitivity analyses on the sign of σ_r (positive implies procyclical r; negative implies countercyclical r) and the value of θ (whether r is driven by fundamentals or sunspots). The event study results, including r for reference, are displayed in Figures C.7-C.10. Overall, regardless of whether r_t is procyclical, countercyclical, driven by fundamental or sunspot shocks, the crisis event studies are largely similar to the baseline results.

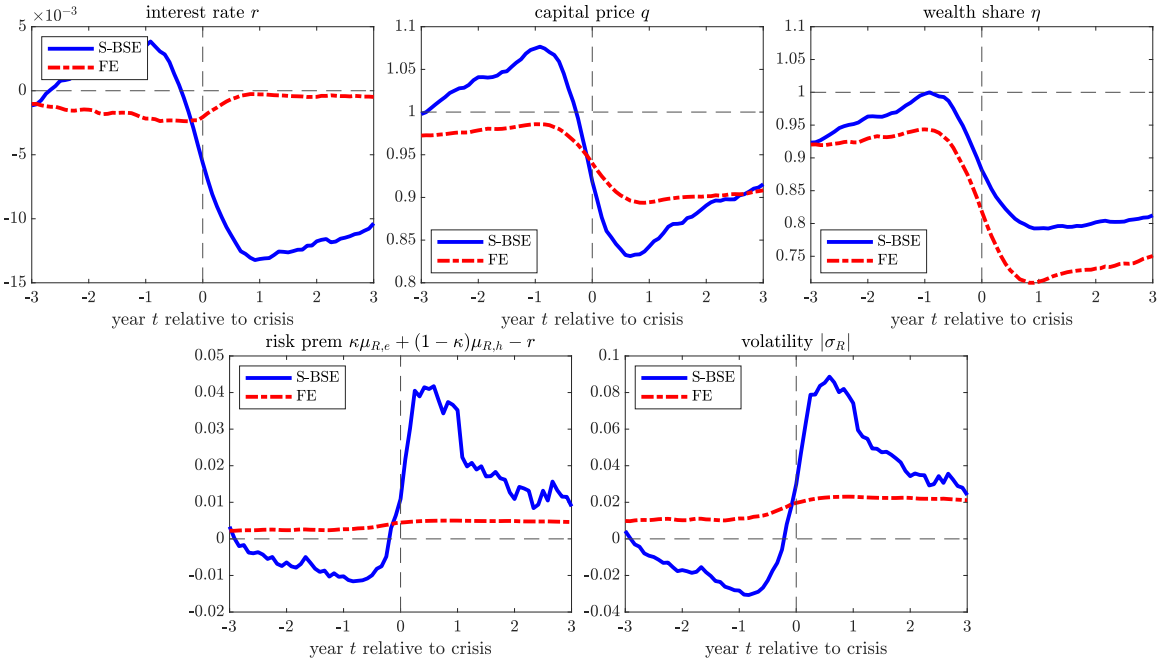


Figure C.7: Event studies around financial crises *with a procyclical interest rate driven by fundamentals*. The interest rate process is (C.1) with $\bar{r} = 0.014$, $\lambda_r = -\log(0.94)$, $\sigma_r = -2\log(0.94) \times 0.023^2$, and $\theta = 1$. Crises are defined as the bottom 3rd percentile of year-to-year log output declines, subject to only one crisis occurring in a 6-year window. All other details are identical to Figure 6.

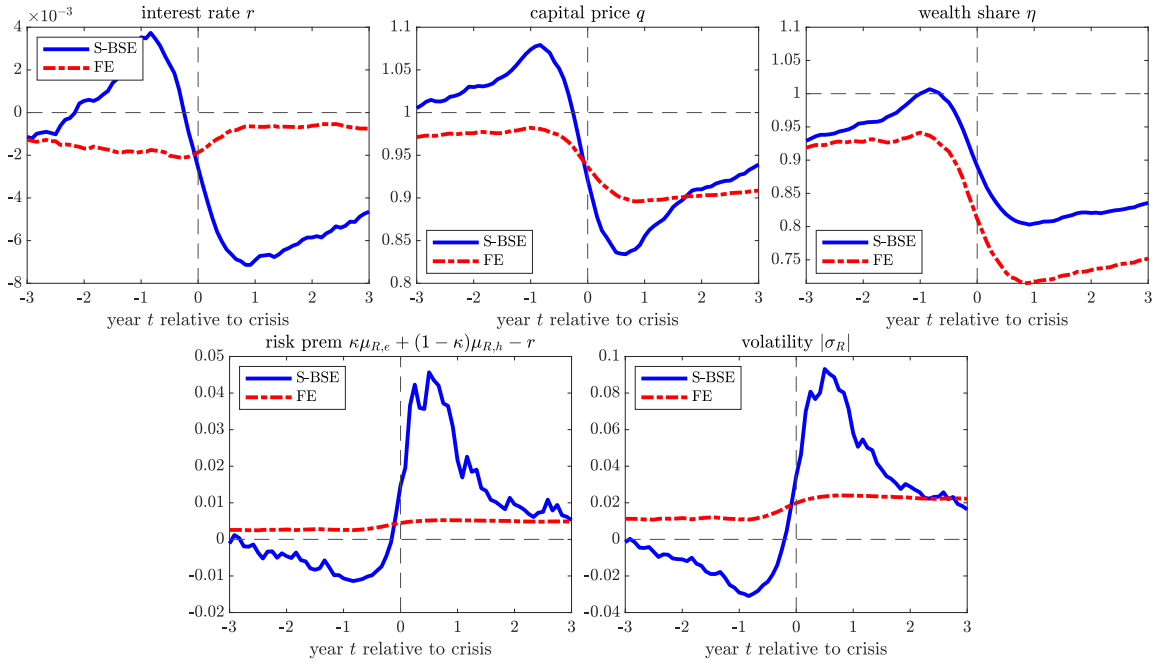


Figure C.8: Event studies around financial crises *with a procyclical interest rate driven by sunspots*. The interest rate process is (C.1) with $\bar{r} = 0.014$, $\lambda_r = -\log(0.94)$, $\sigma_r = -2\log(0.94) \times 0.023^2$, and $\theta = 0$. Crises are defined as the bottom 3rd percentile of year-to-year log output declines, subject to only one crisis occurring in a 6-year window. All other details are identical to Figure 6.

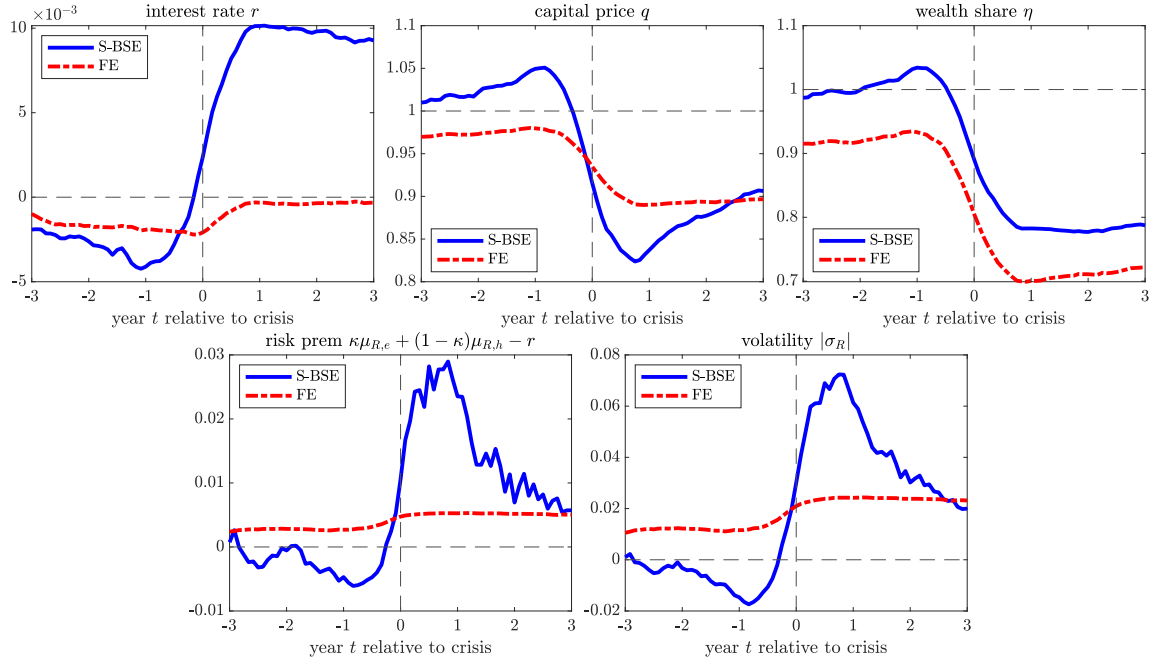


Figure C.9: Event studies around financial crises *with a countercyclical interest rate driven by fundamentals*. The interest rate process is (C.1) with $\bar{r} = 0.014$, $\lambda_r = -\log(0.94)$, $\sigma_r = 2\log(0.94) \times 0.023^2$, and $\theta = 1$. Crises are defined as the bottom 3rd percentile of year-to-year log output declines, subject to only one crisis occurring in a 6-year window. All other details are identical to Figure 6.

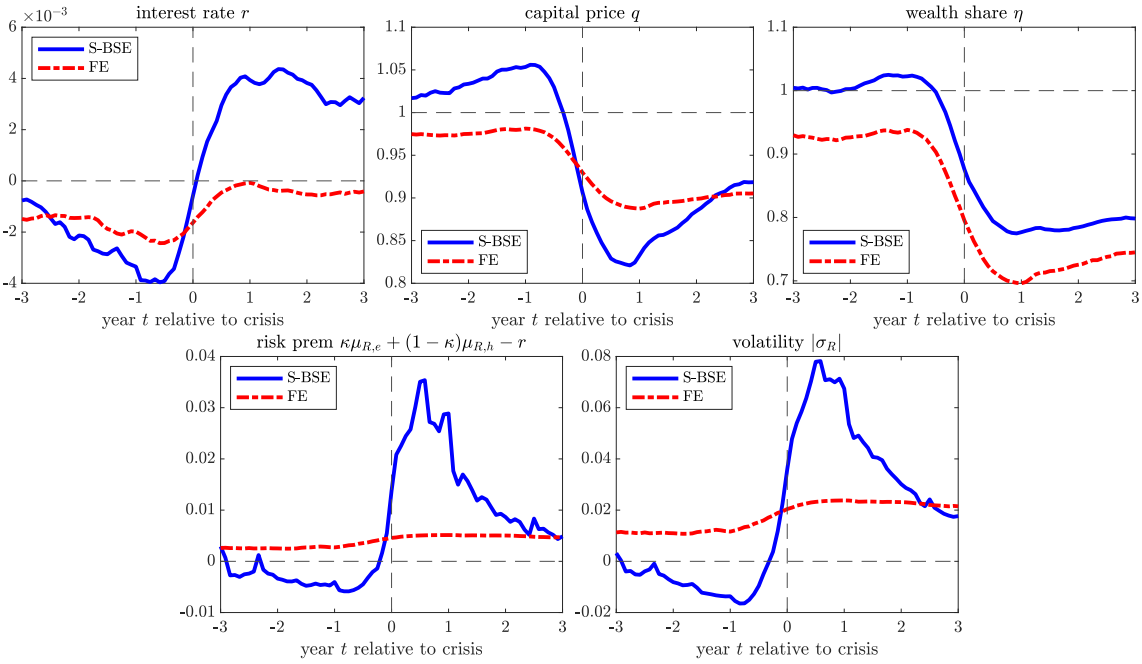


Figure C.10: Event studies around financial crises *with a countercyclical interest rate driven by sunspots*. The interest rate process is (C.1) with $\bar{r} = 0.014$, $\lambda_r = -\log(0.94)$, $\sigma_r = 2\log(0.94) \times 0.023^2$, and $\theta = 0$. Crises are defined as the bottom 3rd percentile of year-to-year log output declines, subject to only one crisis occurring in a 6-year window. All other details are identical to Figure 6.

C.3 Specification of the quantitative model with investment and jumps

We provide more details on the quantitative analysis in Section 3.3, based on the model with capital investment and sunspot jumps. Recall the theoretical characterization of equilibrium for this model is contained in Appendix B.6. Here, we focus on more particular features, such as the jump specification, defining some auxiliary asset-pricing objects, the model calibration, and additional unreported quantitative results.

Jump size. The jump size $\ell_{q,t-}$ is constructed as follows. Because we wish to allow a state-dependent formulation, define $state_\lambda \in \{normal, quiet, panic\}$. We then specify the jump size by the function

$$\ell_q(\eta, q, \lambda, state_\lambda) = \begin{cases} C_{state_\lambda} \ell_q^{max}(\eta, q, \lambda), & \text{if} \quad \kappa > \kappa^{min}_{state_\lambda} \text{ and } C_{state_\lambda} \ell_q^{max}(\eta, q, \lambda) > \ell^{min}_{state_\lambda} \\ 0, & \text{otherwise.} \end{cases}$$

In the formula above, $\ell_q^{max}(\eta, q, \lambda)$ is the maximum possible jump size (an endogenously-determined amount, which is derived in Appendix B.6), $C_{state_\lambda} \in [0, 1)$ is a fraction of this maximum jump size that is realized (a state-dependent parameter), $\kappa^{min}_{state_\lambda}$ is the minimal level of κ such that jumps can affect prices (a state-dependent parameter), and

$\ell^{min}_{state_\lambda}$ is the minimum allowable jump size (a state-dependent parameter).

Some asset pricing objects. We extend the model with two assets that proxy for liquidity and credit premia. Both assets are assumed to be in zero net supply and can only be held by experts, and so they do not alter any aspect of equilibrium (i.e., they are assets that can be priced using an existing equilibrium SDF).

The first asset loses fraction α of its value when a "liquidity event" is realized. We define a liquidity event as a jump of $\ell_{q,t-}dJ_t > 0$ that triggers a drop in capital price. In equilibrium, the premium such an assets pays over the risk free rate would be

$$LiqPrem = 1_{\{\ell_q>0\}}\frac{\lambda\alpha}{1-\frac{\kappa}{\eta}\ell_q} \tag{C.2}$$

We calibrate the loss to 5 percent, i.e., $\alpha = 0.05$, which corresponds to the effective liquidation loss in the data of the 2008 crisis (Krishnamurthy and Li, 2024).

The second asset captures default premia, conditional on the liquidity event. Whenever a liquidity event is realized, with exogenous probability π the asset "defaults" on a fraction of its value and with probability $1-\pi$ it experiences a (small) decline in value. The loss-given-default ℓ_m is assumed to be increasing in the drop of capital price ℓ_q, in particular, $\ell_m = m_0 + m_1\ell_q$. Meanwhile, the price decrease in the non-default scenario is fixed at m_2. In equilibrium, the spread this asset pays over the risk-free rate is

$$CredSpread = 1_{\{\ell_q>0\}}\frac{\lambda}{1-\frac{\kappa}{\eta}\ell_q}\Big[\pi(m_0+m_1\ell_q)+(1-\pi)m_2\Big] \tag{C.3}$$

We target an annual default rate of 4.26%, which corresponds to the difference in default rates of 10-year BAA and AAA bonds in the US, by setting $0.0426 = \pi\mathbb{E}[1_{\{\ell_q>0\}}\lambda_t]$. We target the loss-given-default to that of BAA bonds, which from Moody's data has been 55% on average over the last three decades and rose to 65% during the 2008 crisis. Hence, we calibrate $m_0 = 0.45$ and m_1 to match an average increase of 10% during liquidity events, i.e., $0.10 = m_1\mathbb{E}[\ell_q|\ell_q > 0]$. We set $m_2 = 0.02$. The calibration discussed for the credit spread here follows the one in Krishnamurthy and Li (2024).

Model calibration. Table C.1 presents the values for the parameters of the model. Panel A contains the fundamental parameters of the model. Panel B contains parameters that govern the λ process. Panel C contains parameters associated to the construction of the sentiment equilibrium, including those governing the jump size ℓ_q, the fraction of Brownian risk coming from the sunspot shock, the riskless interest rate, and the

specification for the reflecting boundary at the lower end of the equilibrium domain.

Unconditional moments. Table C.2 presents unconditional moments of the model. In addition to the values coming out of our model, the last column of the table documents the value (or range) in the data. This last column also includes a code (starting with letter "R") that refers to the position in Table C.3, which provides a description of the source of these values in the empirical literature (e.g., R1 refers to moment 1). We briefly discuss the model-implied unconditional moments.

The model's investment-to-capital ratio (0.04) and the GDP-to-capital ratio (0.08) are slightly below their data counterparts (though one should keep in mind that capital is the only productive asset in our model without labor). The model's annual GDP growth (0.034) is slightly above the data (0.025), and the volatility of GDP growth is too large (0.10) compared to data. Similarly, consumption growth (0.034) and in particular its volatility (0.087) are too high compared to data. With log preferences, these high levels of real volatility are needed in order to generate large capital return volatility, which we also discuss below. That said, investment growth (0.034) and the volatility of investment growth (0.125) are close to the corresponding moments in the data (0.026 and 0.127).

In the model, we can define capital return in different ways depending on the dividend considered (experts', households' or their capital-weighted average). We report the return, in excess of the riskless rate, associated to the aggregate dividend: its mean (0.028) and standard deviation (0.1517) are in the ballpark of un-levered US equity returns. We also report the return on experts' net worth: its mean (0.1053) and standard deviation (0.3939) align closely with data on financial stock returns (0.0875 and 0.3653, respectively). Experts' capital ratio (i.e., the inverse of leverage) in the model (0.35) is slightly above the corresponding moment in the data (between 0.20 and 0.27). The model's liquidity premium (0.009) is very close to the data (0.0094).

We also target an annual probability of crisis of 3.5%; evidence from various sources suggests this frequency is between 3%-6%. In our simulations, a crisis is dated in a month with a month-to-month log-change in GDP below percentile 0.035, conditional on no other crisis in the previous 7 years (the length of the crisis event study windows below). Finally, the unconditional probability of a liquidity event is around 13.3% (=7.1% + 1.9% + 4.3%). As mentioned before, the corresponding value in Krishnamurthy and Li (2024) is around 7.2%. While our results is higher, recall that liquidity events during normal times are mild by construction; removing these 7.1%, we would obtain only 6.2% for the "more significant" liquidity events corresponding to quiet and panic states.

Parameter	*Value*	*Description*
A. Fundamental model parameters		
ρ_e	0.04	Discount rate of experts
ρ_h	0.04	Discount rate of households
a_e	0.20	Productivity of expert
a_h	0.10	Productivity of household
σ	0.05	Exogenous volatility of capital shocks
g	0	Exogenous component of growth
χ	30	Adjustment cost of investment
γ	0	Auxiliary parameter for investment function
δ	0.03	Depreciation rate
δ_e	0.10	Retirement rate of experts
δ_h	0.01	Retirement rate of households
B. Poisson processes parameters		
λ^{normal}	0.20	J_t arrival rate during normal times
λ^{quiet}	0.10	J_t arrival rate during quiet times
λ_L^{panic}	0.30	lower bound for J_t arrival rate during panic times
λ_H^{panic}	0.70	upper bound for J_t arrival rate during panic times
T_{max}^{panic}	10	maximum number of years in panic state
$\lambda_{normal \to quiet}$	0.05	arrival rate for transition from normal to quiet times
$\lambda_{quiet \to normal}$	0.00	arrival rate for transition from quiet to normal times
$\lambda_{panic \to normal}$	0.25	arrival rate for transition from panic to normal times
C. Sentiment equilibrium parameters		
C_{normal}	0.40	fraction of max possible q loss (when not zero) during normal times
C_{quiet}	0.75	fraction of max possible q loss (when not zero) during quiet times
C_{panic}	0.50	fraction of max possible q loss (when not zero) during panic times
κ_{normal}^{min}	0.30	κ lower threshold during normal times ($\ell_q = 0$ if κ below it)
κ_{quiet}^{min}	0.90	κ lower threshold during quiet times ($\ell_q = 0$ if κ below it)
κ_{panic}^{min}	0.00	κ lower threshold during panic times ($\ell_q = 0$ if κ below it)
ℓ_{normal}^{min}	0.00	min loss size during normal times ($\ell_q = 0$ if it would be below this)
ℓ_{quiet}^{min}	0.05	min loss size during quiet times ($\ell_q = 0$ if it would be below this)
ℓ_{panic}^{min}	0.00	min loss size during panic times ($\ell_q = 0$ if it would be below this)
ϑ	0.5	share of diffusive return variance via the fundamental shock
$\bar{r}$	0.014	(fixed) interest rate when $\kappa < 1$
$\underline{\kappa}$	0.01	reflecting boundary occurs at a line from $(0, \underline{\kappa})$ to $(1, 1)$ in (η, κ) space

Table C.1: Model parameters

Name in Code	*Value*	*Description*	*Data*
InvCapRatio	0.0409	Investment to capital ratio	0.0879 [R1]
GDPCapRatio	0.0809	GDP to capital ratio	$[0.14, 0.4]$ [R2a,b]
GDPGrowth	0.0341	mean GDP growth (year-to-year)	0.025 [R3a,b,c]
GDPGrowthVol	0.1059	std GDP growth (year-to-year)	0.0368 [R4a,b,c]
InvGrowth	0.0341	mean Investment growth (year-to-year)	0.0261 [R5]
InvGrowthVol	0.1253	std Investment growth (year-to-year)	0.1275 [R6]
ConGrowth	0.0342	mean Consumption growth (year-to-year)	0.0179 [R7]
ConGrowthVol	0.0872	std Consumption growth (year-to-year)	0.0418 [R8]
ExcCapRet	0.0279	Mean of excess log capital return, average dividend (experts and households)	0.0545 [R9]
stdExcCapRet	0.1517	Std of excess log capital return, average dividend (experts and households)	0.2178 [R10]
BanksCapRatio	0.3511	Equity-to-Assets for experts	[0.20,0.27] [R11a,b]
ExpRet	0.1053	Mean log return on experts' wealth (no consumption)	0.0875 [R12]
stdExpRet	0.3939	Std log return on experts' wealth (no consumption)	0.3653 [R13]
LiqPrem	0.0090	Liquidity premium	0.0094 [R14a,b]
annualProbCrisis	0.0350	Annual prob. of a crisis	[0.025,0.06] [R15a,b,c]
LiqEvent_normal	0.0707	Annual prob. of a liquidity event	
LiqEvent_quiet	0.0187	Annual prob. of a liquidity event during quiet times	
LiqEvent_panic	0.0432	Annual prob. of a liquidity event during panic times	

Table C.2: Model unconditional moments. The references (e.g., R#) in the final column correspond to the row number in Table C.3 that provide sources for the data counterparts.

	Name	*Data*	*Description*
R1	Investment capital ratio	0.0879	IMF Investment and Capital Stock Dataset (Adv Economies 1960-2019)
R2a	GDP capital ratio	0.3997	mean in IMF Investment and Capital Stock Dataset (Adv Economies 1960-2019)
R2b	GDP capital ratio	0.14	He and Krishnamurthy (2019)
R3a	GDP growth	0.0199	JST data w/o WWI and WWII periods (real GPP per capita)
R3b	GDP growth	0.0303	JST data w/o WWI and WWII periods (Nominal GDP/CPI)
R3c	GDP growth	0.0284	JST data w/o WWI and WWII periods (real GPP per capita x population)
R4a	std GDP growth	0.0368	JST data w/o WWI and WWII periods (real GPP per capita)
R4b	std GDP growth	0.0470	JST data w/o WWI and WWII periods (Nominal GDP/CPI)
R4c	std GDP growth	0.0377	JST data w/o WWI and WWII periods (real GPP per capita x population)
R5	Investment growth	0.0261	mean in JST data excluding WWI and WWII (real GPP per capita * Inv/GDP)
R6	std Investment growth	0.1275	JST data w/o WWI and WWII (real GPP per capita * Inv/GDP)
R7	Consumption growth	0.0179	JST data w/o WWI and WWII (real C per capita)
R8	std Consumption growth	0.0418	JST data w/o WWI and WWII (real C per capita)
R9	Excess equity return	0.0545	JST data w/o WWI and WWII (equity total return - safe return)
R10	std Excess equity return	0.2178	JST data w/o WWI and WWII (equity total return - safe return)
R11a	Bank capital ratio	0.2650	mean in JST data w/o WWI and WWII (Capital/ (Total Assets - Deposits))
R11b	Bank capital ratio	0.20	Krishnamurthy and Li (2024)
R12	Bank Equity Return	0.0875	Mean bank equity total return index adj for CPI (Baron et al., 2021)
R13	std Bank Equity Return	0.3653	Std bank equity total return index adj for CPI (Baron et al., 2021)
R14a	Liquidity premium	0.0094	spread between P2 rated 3-month commercial paper and 3-month T-bill 1974-2018
R14b	Liquidity premium	0.0075	Krishnamurthy and Vissing-Jorgensen (2015) estimate of liquidity premium of LT bonds vs AAA corporate bonds
R15a	Crisis probability	6.2% = 79/1272	Schularick and Taylor (2012), collecting crisis dates from various sources
R15b	Crisis probability	2.5% = 50/1946	Jordà et al. (2013), collecting crisis dates from various sources
R15c	Crisis probability	$\approx 3\%$	Baron et al. (2021), Figure A12, averaging frequencies across decades; crisis definition: panics & bank equity crashes $\geq 30\%$

Table C.3: Data unconditional moments. JST refers to data from Jordà et al. (2017), Jordà et al. (2019), or Jordà et al. (2021), collectively available online at `https://www.macrohistory.net/database/`. The numbers using these data are the authors' own calculations.

Crisis moments. Table C.4 reports model-implied crisis moments related to GDP, credit, and credit spreads. In addition to the values coming out of our model, the last column of the table documents the value (or range) in the data. This last column also includes a code (starting with letter "RC") that refers to the position in Table C.5, which provides a description of the source of these values in the empirical literature (e.g., RC1 refers to crisis moment 1).

Regarding GDP growth, the model-implied drop around the crisis, and in the years post-crisis, closely align with the data. For example, the average drop in log GDP 12-36 months post-crisis, relative to trend, is 7.01% in the model and 6.49% in the data. In this sense, the crises in the model are reasonable in magnitude.

Regarding credit-to-GDP ratios, while we do not have direct data counterparts to our standardized measure, the model generates a modest amount of pre-crisis "frothiness" and a large drop upon crisis.

An emerging finding that crises are predictable is captured well by the model. In particular, whether we predict crises by high credit growth or low credit spreads, our model predicts crises broadly in line with the data. For example, using below-median credit spreads as the predictor variable, the model implies a 17% higher probability of crisis in the next 5 years, with the corresponding number being 21% in the data. In addition, crisis predictability rises (21.7% in the model; between 24-27% in the data) if both high credit and low spreads happen at the same time, captured by the interaction between the two predictors.

Finally, the model implies quantitatively reasonable credit spread behavior around crises. For example, the 1-year spread increase around crisis is 0.73 in the model and between 0.48 and 0.68 in the data (both in units of standard deviations). Some pre-crisis frothiness of credit spreads is present, but slightly more modest than in the data; as one example, in the two years before crisis, spreads are -0.14 standard deviations below their mean in the model, with the corresponding numbers between -0.15 and -0.43 in the data. After the crisis, spreads tend to mean-revert at a rate in line with the data (half-life of 3.1 years in the model and between 2.5-3.5 years in the data).

Name in Structure	*Value*	*Description*	*Data*
GDPtrend_trough	-0.1110	Largest negative deviation after the crisis (lowest value of GDP event path).	-0.093 [RC1,RC2]
GDPtrend_23y	-0.0701	Average deviation between months 12 and 36 after the crisis (average value of GDP event path).	-0.0649 [RC3]
GDPtrend_45y	-0.0592	Average deviation between months 36 and 60 after the crisis (average value of GDP event path).	-0.0771 [RC4]
CreditGDP_peak	0.2310	Peak in the 3 years before the crisis, deviation from mean in std terms	
CreditGDP_trough	-1.1803	Trough in the first year after the crisis, deviation from mean in std terms	
CreditGDP_pre5y	0.1488	Average of the 5 years before the crisis, deviation from mean in std terms	
CreditGDP_pre2y	0.0876	Average of the 2 years before the crisis, deviation from mean in std terms	
pred5_HighCredit3y	0.0708	Regression of a dummy equal to 1 if there is a crisis in the next 5 years on HighCredit3y, which is the 3-year average of a dummy equal to 1 when Credit/GDP is above the 75th percentile (monthly series).	0.18 [RC5]
pred5_HighFrothA1y	0.1718	Regression of a dummy equal to 1 if there is a crisis in the next 5 years on HighFroth1y, which is the 1-year average of a dummy equal to 1 when CredSpreadA is below the median but above zero (monthly series).	0.21 [RC6]
pred5_HFA1y_HCy	0.2173	Previous regression using the interaction of HighFrothA1y and HighCredit3y as regressor.	[0.24,0.27] [RC7a,b]
CredSpreadA_peak	0.6652	Max value above mean during the first year after crisis (in terms of std above the mean w/o crisis events)	[0.5777,1.0137] [RC8a,b]
CredSpreadA_trough	-0.1682	Minimum value (in terms of std below the mean w/o crisis events) in the 2 years prior to the crisis	[-0.56,-0.27] [RC9a,b]
CredSpreadA_pre5y	-0.0959	Mean 5 years prior to crisis (in terms of std, relative to mean w/o crisis events)	[-0.12,-0.46] [RC10a,b, RC11a,b]
CredSpreadA_pre2y	-0.1420	Mean 2 years prior to crisis (in terms of std, relative to mean w/o crisis events)	[-0.15,-0.43] [RC12a,b]
DCredSpreadA	0.8334	Peak - trough	[1.13,1.28] [RC13a,b]
DCredSpreadA_1y	0.7335	Level a year after crisis (mean months 10, 11, 12) - level just before the crisis (mean months -3, -2, -1), in std units	[0.48,0.68] [RC14a,b]
halflife_CredSpreadA	3.0833	Half-life of peak of Credit Spread A (recover half towards mean w/o crisis events)	≈[2.5,3.5] [RC15a,b]

Table C.4: Key crisis moments involving GDP, the credit measure $Credit = (\kappa - \eta)qK$, and the *CredSpread* measure defined in equation (C.3), as well as their data counterparts (final column). The references (e.g., RC#) in the final column correspond to the row number in Table C.5 that provides sources for the data counterparts.

	Name	*Data*	*Description*
RC1	max GDP contraction	-0.09	Maximum GDP deviation from trend (assuming on trend at $t = -6$) during the 5 years following a crisis (measured at $t = 5$) (Krishnamurthy and Muir, 2024)
RC2	GDP peak-to-trough	-0.093	Average peak-to-trough decline. Only systemic banking crises. Only 2 year span for peak-to-trough. Reinhart and Rogoff (2009), p.230.
RC3	GDP contraction (mean $t = 2, 3$)	-0.0649	Average deviation from trend, using statistics from Krishnamurthy and Muir (2024), Sufi and Taylor (2022), Schularick and Taylor (2012) (post-WW2)
RC4	GDP contraction (mean $t = 4, 5$)	-0.0771	Same as above
RC5	Crisis predictability: by credit growth	0.18	Linear probability model for a crisis in the next 5 years. Predictor = HighCredit = 3-year average of a dummy equal to one if credit growth (change in Credit/GDP) is above the median (Krishnamurthy and Muir, 2024, Table VII).
RC6	Crisis predictability: by credit spread	0.21	Same as above, but Predictor = HighFroth = 5-year average of a dummy equal to one if the credit spread is below the median
RC7a	Crisis predictability: by credit spread * credit growth	0.27	Same as above, but Predictor = HighCredit*HighFroth
RC7b	Crisis predictability: by credit spread * credit growth	0.24	Same as above, but including year fixed effects
RC8a	Credit spread peak	1.0137	Max value of credit spreads (Krishnamurthy and Muir, 2024, Table VII)
RC8b	Credit spread peak	0.5777	Same as above, but including year fixed effects
RC9a	Froth 2-years before crisis (min)	-0.27	Minimum of credit spread in the 2 years prior to crisis
RC9b	Froth 2-years before crisis (min)	-0.56	Same as above, but including year fixed effects
RC10a	Froth 5-years to crisis	-0.12	Average credit spread 5 years previous to crisis (from $t = -5$ to $t = -1$)
RC10b	Froth 5-years to crisis	-0.36	Same as above, but including year fixed effects
RC11a	Froth 5-years before crisis	-0.23	Average credit spread 5 years previous to crisis (from $t = -6$ to $t = -2$)
RC11b	Froth 5-years before crisis	-0.46	Same as above, but including year fixed effects
RC12a	Froth 2-years before crisis	-0.158	Average credit spread 2 years previous to crisis (from $t = -3$ to $t = -2$)
RC12b	Froth 2-years before crisis	-0.43	Same as above, but including year fixed effects
RC13a	Credit spread max change	1.28	max change in credit spreads (from $t = -2$ to $t = 0$)
RC13b	Credit spread max change	1.1336	Same as above, but including year fixed effects
RC14a	Credit spread change at crisis	0.68	Change in credit spread from $t = -1$ to $t = 0$
RC14b	Credit spread change at crisis	0.4796	Same as above, but including year fixed effects
RC15a	Half-life credit spread post-crisis	3.5	Years for the value at $t = 0$ to return halfway towards the mean
RC15b	Half-life credit spread post-crisis	2.5	Same as above, but including year fixed effects

Table C.5: Crisis event moments. Note that all the moments from RC5 to RC15 are from Krishnamurthy and Muir (2024), Table VII. For all the credit spread moments (RC8-RC15), the units are in standard deviations from the mean (calculated without crisis events). To transform from their original table into standard deviation units, we multiply by 1.09.

Online Appendix 2
(not for publication):
Rational Sentiments and Financial Frictions

Paymon Khorrami and Fernando Mendo

March 19, 2025

D An alternative construction: sentiment state variable

In contrast to the main paper, where (η, q) was the state variable, here we implement our sentiment-driven equilibria with an auxiliary state variable s and with q as a function of η and s. Being explicit about a sentiment state variable is useful for several reasons. First, this equilibrium construction will be pedagogically more familiar to the literature on sunspots. Second, the sentiment state dynamics can be modeled as locally uncorrelated with fundamental shocks, which brings some clarity. Third, this setting happens to facilitate building sunspot equilibria in which experts fully de-lever as their wealth shrinks, i.e., $\kappa \to 0$ as $\eta \to 0$, for which there are natural justifications.

D.1 Explicit equilibrium with a sentiment state variable

Let s be a pure sunspot that is irrelevant to economic fundamentals and loads on only the second shock (recall $Z^{(1)}$ affects capital and $Z^{(2)}$ does not):

$$ds_t = \mu_{s,t}dt + \sigma_{s,t}\left(\begin{smallmatrix}0\\1\end{smallmatrix}\right) \cdot dZ_t, \quad s_t \in \mathcal{S}. \tag{D.1}$$

(Online Appendix D.5 solves additional examples with sentiment correlated to fundamentals, i.e., with $ds = \mu_s dt + \sigma_s^{(1)} dZ^{(1)} + \sigma_s^{(2)} dZ^{(2)}$.) We will also find some use in introducing auxiliary state variables that can affect the drift $\mu_{s,t}$. This is possible to do in a very flexible way, due to the drift indeterminacy result of Corollary **??**. Let $x_t \in \mathcal{X}$ be an arbitrary bounded diffusion,

$$dx_t = \mu_x(x_t)dt + \sigma_x(x_t) \cdot dZ_t,$$

which (only) affects the sentiment drift, through $\mu_{s,t} = \mu_s(\eta_t, s_t, x_t)$.

Definition 3. A *Markov S-BSE* in states $(\eta, s, x) \in (0,1) \times \mathcal{S} \times \mathcal{X}$ consists of functions $(q, \kappa, r, \sigma_\eta, \mu_\eta, \sigma_s) : (0,1) \times \mathcal{S} \mapsto \mathbb{R}$, and $\mu_s : (0,1) \times \mathcal{S} \times \mathcal{X} \mapsto \mathbb{R}$, all C^2 almost-everywhere, such that the process $(\eta_t, q(\eta_t, s_t), \kappa(\eta_t, s_t), r(\eta_t, s_t))_{t \geq 0}$ is an S-BSE.

Remark 4 (Endogenous sentiment dynamics)**.** *Note that the statement of Definition 3 allows* (σ_s, μ_s) *to be endogenous, in the sense that they could depend on the wealth distribution* η*. Our examples in this section purposefully entertain this endogeneity, partly because we think of this as the more interesting and realistic situation. Why? As shown in Section 2, dynamics depend explicitly on* q *in an S-BSE. Thus, it is completely sensible for agents in our S-BSEs to use asset prices directly in forecasting; in particular, sentiment dynamics* (σ_s, μ_s)*—which are nothing but belief dynamics—themselves should condition on* q*. But* q *will depend on both* s *and* η*, implying sentiment dynamics* (σ_s, μ_s) *depend on* η *too, through* q*. That said, Online Appendix D.6 verifies that similar types of sunspot equilibria can be constructed with exogenous sentiment dynamics, i.e.,* (σ_s, μ_s) *are only functions of* s*, not* η*.*

The Markov assumption in Definition 3 allows us to specialize equilibrium conditions. By applying Itô's formula to $q(\eta, s)$, we obtain the capital price volatility σ_q in terms of σ_η, namely

$$q\sigma_q = \sigma_\eta \partial_\eta q + \sigma_s \partial_s q.$$

From equation (14), we also have σ_η in terms of σ_q. Solving this two-way feedback, we obtain

$$\sigma_q = \frac{\binom{1}{0}(\kappa - \eta)\sigma \partial_\eta \log q + \binom{0}{1}\sigma_s \partial_s \log q}{1 - (\kappa - \eta)\partial_\eta \log q}. \tag{D.2}$$

Using (D.2) in (RB), we obtain the following equation linking capital prices, the capital distribution, and sentiment volatility:

$$0 = \min\left[1 - \kappa, \frac{a_e - a_h}{q} - \frac{\kappa - \eta}{\eta(1-\eta)}\left(\frac{\sigma^2 + (\sigma_s \partial_s \log q)^2}{(1 - (\kappa - \eta)\partial_\eta \log q)^2}\right)\right]. \tag{D.3}$$

Our strategy to find a Markov S-BSE is to guess a capital price function $q(\eta, s)$ and then use equation (D.3) to "back out" the sunspot volatility σ_s that justifies it. We will perform a construction such that sunspots only increase volatility relative to the fundamental equilibrium, to highlight their potential for resolving puzzles. For this reason, we sometimes refer to s as *rational pessimism*.

More specifically, suppose a fundamental equilibrium, where sunspots do not matter, exists with equilibrium capital price q^{FE} (see Online Appendix E for details on the fundamental equilibria). We will think of q^{FE} as the "best-case" capital price, because despite featuring amplification, q^{FE} inherits no sunspot volatility. Conversely, think of the capital price q^∞ associated to an infinite-volatility equilibrium as the "worst-case" capital price (substitute $|\sigma_R| \to \infty$ into (20) to find $q^\infty := \frac{\eta a_e + (1-\eta)a_h}{\bar{\rho}}$).

Our strategy is essentially to treat the sentiment variable s as a device to shift continuously between the best-case q^{FE} and the worst-case q^{∞}. Mathematically, we conjecture a capital price that is approximately a weighted average of q^{FE} and q^{∞}, with weights $1-s$ and s. The novelty of our approach here is to then use equation (D.3) to solve for sunspot volatility σ_s, which will generically depend on experts' wealth share η. In terms of Figure 3, the economy will live in the sub-region bounded by the solid FE line and the $\kappa = \eta$ border (and notice this implies the full-deleveraging condition $\kappa \to 0$ as $\eta \to 0$). In the proposition below, we verify that such a construction is indeed an equilibrium.

Proposition D.1. *Let Assumption 1 hold, and assume a fundamental equilibrium exists for each $\sigma \geq 0$ small enough. Then, for all $\sigma \geq 0$ small enough, there exists a Markov S-BSE with capital prices arbitrarily close to $(1-s)q^{FE}(\eta) + sq^{\infty}(\eta)$. In this equilibrium, μ_s is indeterminate except near the boundaries of $(0,1) \times \mathcal{X} \times \mathcal{S}$.*

We construct a numerical example closely following Proposition D.1, which we will use in subsequent sections. The left panel of Figure D.1 shows the capital price function. A rise in rational pessimism s reduces the capital price q, independently of wealth share η (although η will also endogenously respond to s-shocks).

The middle panel of Figure D.1 displays capital return volatility, which can be substantially greater than in the fundamental equilibrium. Implied by capital return volatility is an underlying sunspot shock size σ_s, which is displayed in the right panel of Figure D.1. Sunspot dynamics become more volatile both as experts become poor (η shrinks) and as the economy approaches the worst-case equilibrium (s rises). The dependence of σ_s on η is the notion of endogenous beliefs that can occur in S-BSEs.

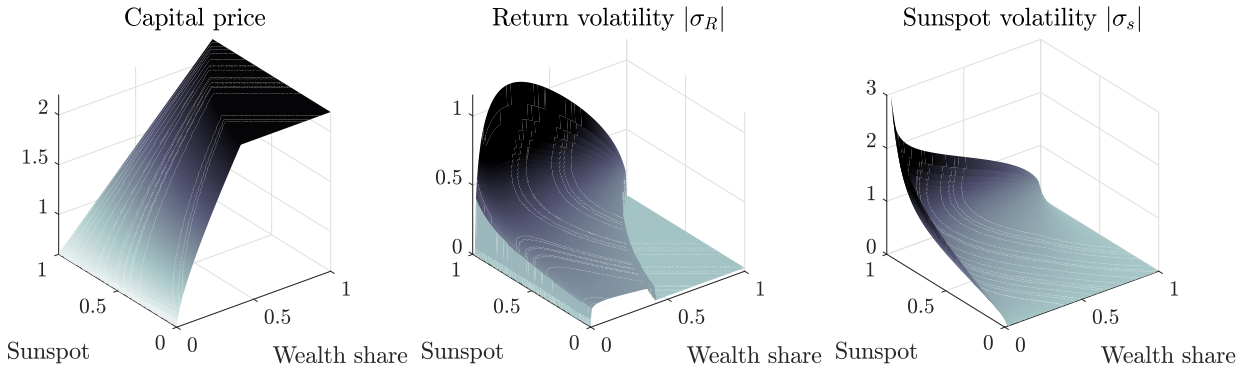


Figure D.1: Capital price q, volatility of capital returns $|\sigma_R|$, and sunspot shock volatility σ_s. Parameters: $\rho_e = \rho_h = 0.05$, $a_e = 0.11$, $a_h = 0.03$, $\sigma = 0.025$.

D.2 Non-fundamental crises and large amplification

We now show how our model with sentiment shocks can help resolve some empirical issues related to financial crises and recoveries.

First, Figure D.2 compares impulse responses to a large negative balance-sheet shock (i.e., decline in η) versus a wave of pessimism (i.e., increase in s). The shock sizes are chosen so that the initial drop in capital price $q_0 - q_{0-}$ is roughly the same. "Balance-sheet recessions" (decline in η) feature a modest increase in volatility followed by relatively slow recoveries, as experts can only rebuild their balance sheets by earning profits over time. By contrast, "pessimism crises" (increase in s) feature large temporary volatility spikes and can have accelerated recoveries (depending on the choice of μ_s). The dynamics after a pessimism shock—both the rise in volatility and speed of recovery—are closer to empirical evidence.[30] Our results on recovery speeds are related to Maxted (2024), who shows how extrapolative beliefs can help this class of models match such evidence, but with our rational sentiment in place of his behavioral sentiment.

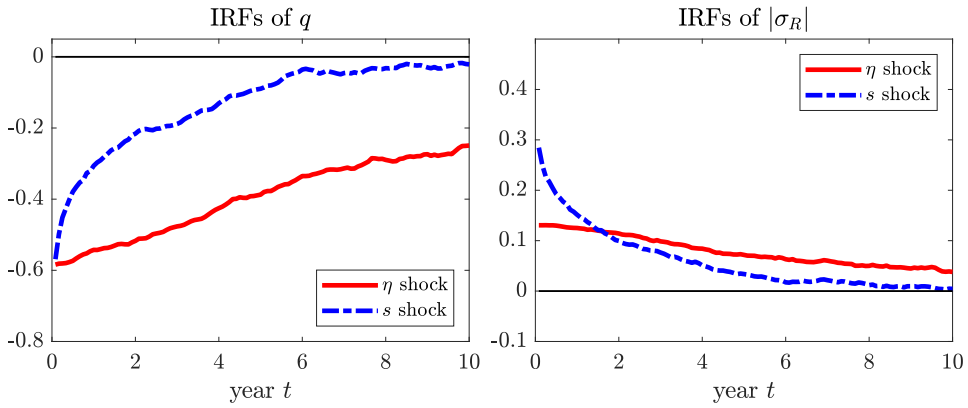


Figure D.2: Bust IRFs of capital price q and return volatility $|\sigma_R|$. The IRFs labeled "η shock" are responses to a decrease in η from $\eta_{0-} = 0.5$ to $\eta_0 = 0.25$, holding s_0 fixed at 0.1. The IRFs labeled "s shock" are responses to an increase in s from $s_{0-} = 0.1$ to $s_0 = 0.8$, holding η_0 fixed at 0.5. These shock sizes are chosen such that the initial response of q are approximately equal. Note that η_0 would respond to an "s shock," since σ_η has a non-zero second element, but we keep it fixed here. IRFs are computed as averages across 500 simulations at daily frequency, with the outcomes above then averaged to the monthly level. Parameters: $\rho_e = \rho_h = 0.05$, $a_e = 0.11$, $a_h = 0.03$, $\sigma = 0.025$. Type-switching parameters: $\delta_h = 0.004$ and $\delta_e = 0.036$. In this example, we set the sunspot drift $\mu_s = 0.0002s^{-1.5} - 0.0002(s_{\max} - s)^{-1.5}$, where $s_{\max} = 0.95$. This choice ensures $s_t \in (0, s_{\max})$ with probability 1.

[30] During the 2008 financial crisis and 2020 COVID-19 episode in the US, implied volatility from option markets spiked by magnitudes on the order of 60%. For a rough idea of what the data says about crisis recoveries, see Jordà et al. (2013) and Reinhart and Rogoff (2014) for output, and see Muir (2017) and Krishnamurthy and Muir (2024) for credit spreads and stock prices. Across these many measures, and using broad-based international panels, crisis recovery times tend to range from 4-6 years on average.

Of course, note that η responds to s-shocks, i.e., σ_η has a non-zero second component. Thus, a true sentiment-driven crisis features dynamics that are a blend of the two IRFs in Figure D.2. Figure D.2 shows a pure shock to s, without the endogenous co-movement in η, for theoretical clarity.

To establish some more confidence in these results, we present the following two propositions which together show that amplification can be arbitrarily high (Proposition D.2) as long as sentiment shocks are the source (Proposition D.3). Given the literature's struggle to identify a "smoking gun" (e.g., TFP shocks, capital efficiency shocks) for financial crises, we view this set of results as a helpful insight. The importance of sentiment s, relative to experts' wealth share η, also echoes the empirical results suggesting financial crises are not associated with pre-crisis levels of bank capital (Jordà et al., 2021).

Proposition D.2 (Arbitrary volatility)**.** *Given a target variance $\Sigma^* > 0$ and any parameters satisfying the assumptions of Proposition D.1, there exists a Markov S-BSE with stationary average return variance exceeding the target, i.e., $\mathbb{E}[|\sigma_R|^2] > \Sigma^*$.*

Proposition D.3 (Decoupling)**.** *In the Markov S-BSEs of Proposition D.1, the fraction of return volatility due to sentiments $|\binom{0}{1} \cdot \sigma_R|/|\sigma_R|$ and total return volatility $|\sigma_R|$ increase with s.*

D.3 Booms predict crises

We now use the same framework to cast light on empirical findings suggesting that financial crises are predictable, in particular by large credit and asset price booms (Reinhart and Rogoff, 2009; Jordà et al., 2011, 2013, 2015a,b; Mian et al., 2017) that feature below-average credit spreads (Krishnamurthy and Muir, 2024; López-Salido et al., 2017; Baron and Xiong, 2017).

To do this, we make use of the auxiliary variable x that can affect the sentiment drift. Following some models of extrapolative beliefs (Barberis et al., 2015; Maxted, 2024), define an exponentially-declining weighted average of sentiment shocks:

$$x_t := x_0 + \sigma_x \int_0^t e^{-\beta_x(t-u)} dZ_u^{(2)}. \tag{D.4}$$

The variable x measures the stock of past pessimism. Assume the drift of s depends on x via

$$\mu_{s,t} = b_x x_t + \hat{\mu}_s(s_t) \quad \text{with} \quad b_x \leq 0.$$

Similar to Section D.2, the term $\hat{\mu}_s$ will be designed to induce stationarity in s_t. The new term $b_x x$ induces the following dynamics: after a wave of optimism ($dZ_t^{(2)} < 0$), s_t and x_t will be low, but this raises $\mu_{s,t}$ and shifts up the conditional distributions of future pessimism s_{t+h}. If the constant b_x is large enough, the shift can generate dynamics reminiscent of "overshooting," in which an optimism-driven boom raises bust

probabilities. Differently from the extrapolative belief literature, the beliefs implied by these sentiment dynamics are completely rational.

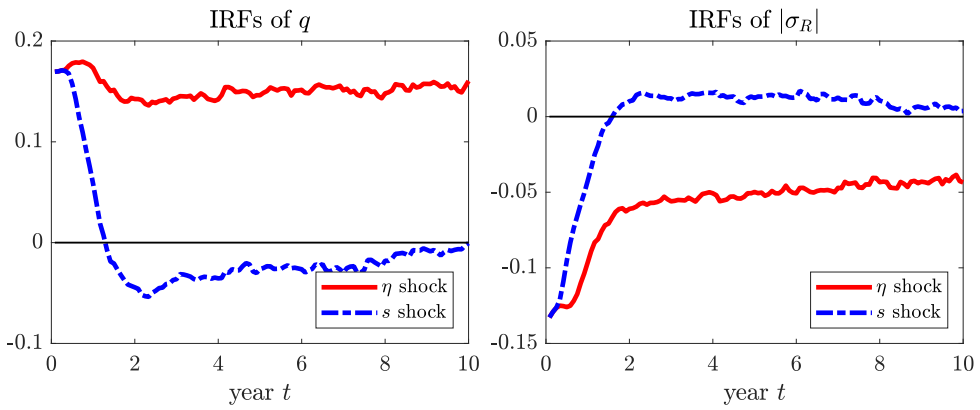


Figure D.3: Boom IRFs of capital price q and return volatility $|\sigma_R|$. The IRFs labeled "η shock" are responses to an increase in η from $\eta_{0-} = 0.5$ to $\eta_0 = 0.7$, holding s_0 fixed at 0.4. The IRFs labeled "s shock" are responses to a decrease in s from $s_{0-} = 0.4$ to $s_0 = 0.1$, holding η_0 fixed at 0.5. These shock sizes are chosen such that the initial response of q are approximately equal. Note that η_0 would respond to an "s shock," since σ_η has a non-zero second element, but we keep it fixed here. IRFs are computed as averages across 2000 simulations at daily frequency, with the outcomes above then averaged to the monthly level. Parameters: $\rho_e = \rho_h = 0.05$, $a_e = 0.11$, $a_h = 0.03$, $\sigma = 0.025$. Type-switching parameters: $\delta_h = 0.004$ and $\delta_e = 0.036$. In this example, we set the sunspot drift $\mu_s = b_x x + 0.0001 s^{-1.5} - 0.0001(s_{\max} - s)^{-1.5}$, where $s_{\max} = 0.95$, $b_x = -25$, $\beta_x = 0.1$, and $\sigma_x = 0.025$. The parameters (β_x, σ_x) are approximately the values used for the mean-reversion and volatility of the diagnostic belief process in Maxted (2024).

Figure D.3 displays IRFs consistent with this overshooting logic. Sentiment-driven booms predict future busts: an optimism shock raises asset prices and lowers volatility for 1-2 years, but predicts lower prices and higher volatility afterward. (This number of years depends on b_x.) By contrast, a boom driven by expert wealth counterfactually predicts high prices, lower volatility, and lower fragility at all horizons.

To connect to the empirical literature, we conduct an event study in Figure D.4. We simulate our model (which thus features contributions from both fundamental and sunspot shocks), identify crises in the simulated data, and plot average outcomes in the years before and after crisis. Crises are identified as the worst 3rd percentile of yearly output drops; other tail outcomes will produce similar graphs. We see that conditions are improving up to 1 year before the crisis, with risk premia below average and *declining*. The crisis emerges suddenly and features spikes in all variables. Although we do not report it here, such dynamics cannot be produced in the non-sunspot equilibria of the model.

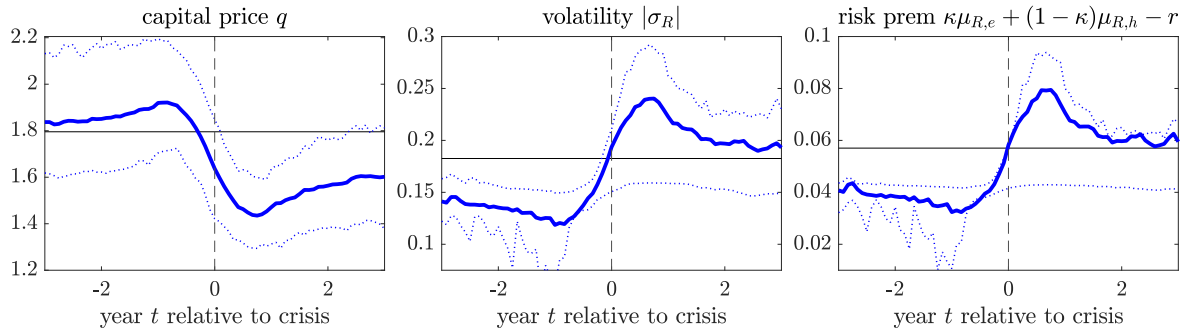


Figure D.4: Event studies around financial crises. Crises are defined as the bottom 3rd percentile of year-to-year log output declines. Data is generated via a 10,000 year simulation at the daily frequency, with the outcomes above then averaged to the monthly level. The solid blue line is the mean path, and the dotted blue lines represent the 25th and 75th percentiles. The thin horizontal line represents the unconditional average. Parameters: $\rho_e = \rho_h = 0.05$, $a_e = 0.11$, $a_h = 0.03$, $\sigma = 0.025$. Type-switching parameters: $\delta_h = 0.004$ and $\delta_e = 0.036$. In this example, we set the sunspot drift $\mu_s = b_x x + 0.0002 s^{-1.5} - 0.0002(s_{\max} - s)^{-1.5}$, where $s_{\max} = 0.95$, $b_x = -25$, $\beta_x = 0.1$, and $\sigma_x = 0.025$. The parameters (β_x, σ_x) are approximately the values used for the mean-reversion and volatility of the diagnostic belief process in Maxted (2024).

D.4 Proofs of Propositions D.1-D.2-D.3

Proof of Proposition D.1. We provide a sketch the proof, which is similar to Theorem 1. Essentially, we want to construct an upper bound for the price based on the fundamental equilibrium, and the lower bound for the price based on a small perturbation of the worst-case price (we want to include this perturbation since volatility explodes when the price approaches its worst-case value). For notation, recall that $\bar{\rho} := \eta\rho_e + (1-\eta)\rho_h$. By analogy, define $\bar{a} := \eta a_e + (1-\eta)a_h$.

Upper and lower bounds for price. Let $(\hat{q}^0, \hat{\kappa}^0)$ be the solution to the fundamental equilibrium (which exists by assumption), and let $\eta^0 := \inf\{\eta : \hat{\kappa}^0 \geq 1\}$. By Lemma E.1 part (v), if σ is small enough then $\eta^0 < 1$, which we assume to be the case. Then, define

$$q^0(\eta) := \begin{cases} \hat{q}^0(\eta), & \text{if} \quad \eta < \eta^0; \\ \hat{q}^0(\eta) + \varphi(\eta), & \text{if} \quad \eta \geq \eta^0, \end{cases} \tag{D.5}$$

where φ is a C^2 function with the properties $\varphi(\eta^0) = 0$ and $\varphi' > (\bar{a}/\bar{\rho})' - (a_e/\bar{\rho})'$ for all η. In words, q^0 is equal to the fundamental equilibrium price $\hat{q}^0$ whenever $\hat{\kappa}^0 \leq 1$ and above it when $\hat{\kappa}^0 = 1$. For the other extremal function, use the "worst-case" price

$$q^1(\eta) := \bar{a}(\eta)/\bar{\rho}(\eta). \tag{D.6}$$

Importantly, we have $q^0 > q^1$ for all η.

Candidate price. We proceed to combine these two extremal functions according to the following convex combination, where $\alpha \in (0,1)$ is fixed:

$$\tilde{q}(\eta, s) := (1 - \alpha s) q^0(\eta) + \alpha s q^1(\eta), \quad (\eta, s) \in \mathcal{D} = (0,1) \times \mathcal{S}. \tag{D.7}$$

where $\mathcal{S} = (0,1)$ is the domain for the sunspot state s. For each $s \in \mathcal{S}$, define $\eta^*(s) := \inf\{\eta : \tilde{q}(\eta, s) \geq a_e/\bar{\rho}\}$, which can be shown is strictly increasing.[31] Put

$$q(\eta, s) := \begin{cases} \tilde{q}(\eta, s), & \text{if} \quad \eta < \eta^*(s) \\ a_e/\bar{\rho}(\eta), & \text{if} \quad \eta \geq \eta^*(s) \end{cases} \quad \text{and} \quad \kappa := \frac{\bar{\rho} q - a_h}{a_e - a_h}.$$

By construction, the pair (q, κ) satisfy equation (PO).

Volatility. Gven the fact that $\alpha < 1$ in (D.7), the resulting capital price is always bounded away from the worst-case price, except as $\eta \to 0$. Thus, the resulting equilibrium volatility will remain bounded for the exact same reasons as in the construction of Theorem 1 (which used a small perturbation of the state space to keep capital prices away from their worst-case value). We omit the construction of this return volatility $|\sigma_R|$, since it is identical to Theorem 1. Given the value of $|\sigma_R|$ and the identity $|\sigma_R|^2 = \frac{\sigma^2 + (\sigma_s \partial_s \log q)^2}{[1-(\kappa-\eta)\partial_\eta \log q]^2}$, we obtain σ_s by inverting this identity. Some technical checks are required to ensure that the resulting σ_s is real, but this can be done. (If $\sigma = 0$, this is guaranteed.)

Sunspot drift and stationarity. Having determined q, κ, and σ_s, we define μ_η and σ_η by (13)-(14). It remains to determine μ_s. We will pick $\mu_s(\eta, s) = m(\eta, s)$, where m is a C^2 function with the following properties: $\partial_s m < 0$, and for some $0 \leq s^0 < s^1 \leq 1$

[31] Indeed, note that $\tilde{q}$ is C^2 on $(\eta^0, \eta^1) \times \mathcal{S}$, which implies η^* is C^1. Then, use the fact that η^* is C^1 to differentiate $\tilde{q}(\eta^*(s), s) = a_e/\bar{\rho}(\eta^*(s))$ with respect to s, and use the fact that $\partial_s \tilde{q} = q^1 - q^0$, and finally rearrange to obtain

$$(\eta^*)'(s) \left[\partial_\eta \tilde{q}(\eta^*(s), s) + \frac{a_e}{\bar{\rho}(\eta^*(s))} \frac{\rho_e - \rho_h}{\bar{\rho}(\eta^*(s))} \right] = q^0(\eta^*(s)) - q^1(\eta^*(s)).$$

If at any point s, we had $(\eta^*)'(s) = 0$, we would necessarily have $q^0(\eta^*(s)) = q^1(\eta^*(s))$. But this contradicts the that $q^0 > q^1$. Thus, $(\eta^*)'(s) \neq 0$ for all s. We can also rule out $(\eta^*)'(s) < 0$ by the fact that $\eta^*(0+) = \eta^0$ and $\eta^*(s) \geq \eta^0$ for all s. Thus, $(\eta^*)'(s) > 0$ for all s.

thresholds,

$$(\text{if } s^0 > 0) \qquad \inf_{\eta \in (0,1)} \lim_{s \searrow s^0} (s - s^0) m(\eta, s) = +\infty \tag{D.8}$$

$$(\text{if } s^0 = 0) \qquad \inf_{\eta \in (0,1)} \lim_{s \searrow s^0} m(\eta, s) > 0 \tag{D.9}$$

$$\sup_{\eta \in (0,1)} \lim_{s \nearrow s^1} (s^1 - s) m(\eta, s) = -\infty. \tag{D.10}$$

Given this choice, we need to demonstrate the time-paths $(\eta_t, s_t)_{t \geq 0}$ remain in $\mathcal{D}$ almost-surely and admit a stationary distribution. This step is very similar to the stochastic stability step in Theorem 1 and is therefore omitted. We simply note that the Lyapunov function to use in this step is $v(\eta, s) := \frac{1}{\eta^{1/2}} + \frac{1}{1-\eta} + \frac{1}{1-s} + \frac{1}{s}$. □

Proof of Proposition D.2. Fix any $\Sigma^* > 0$. The proof is a simple consequence of the fact that σ_q must be unbounded as κ approaches η, which is as q approaches the worst-case price q^1. We fill in the technical details below.

We construct a sequence of equilibria—indexed by (α, ζ)—as follows. Recall the capital price construction in Proposition D.1:

$$q = (1 - \alpha s) q^0 + \alpha s q^1, \quad \text{when} \quad \kappa < 1,$$

where $\alpha < 1$ is a parameter, q^0 is the fundamental equilibrium price, and $q^1 = \bar{a}/\bar{\rho}$ is the worst-case price. Based on the discussion in the text, we may choose μ_s such that equilibrium concentrates on any particular value of s. Thus, pick μ_s such that $s_t \geq \zeta$ almost-surely. Clearly, the choice of μ_s depends on α, but such a choice can always be made for any parameters.

Let $p_{\text{low}} > 0$, $p_{\text{high}} > 0$ be given with $p_{\text{low}} + p_{\text{high}} < 1$. First, note that there exist α^*, ζ^*, ϵ^* such that $\mathbb{P}[\eta_t \leq \epsilon \cap \kappa_t < 1] < p_{\text{low}}$ and $\mathbb{P}[\eta_t \geq 1 - \epsilon \cap \kappa_t < 1] < p_{\text{high}}$ for all $\alpha > \alpha^*$, $\zeta > \zeta^*$, and $\epsilon < \epsilon^*$. This is a consequence of the fact that in any stationary distribution, we have $\lim_{x \to 0} \mathbb{P}[\eta_t < x] = \lim_{x \to 1} \mathbb{P}[\eta_t > x] = 0$ and the fact that $\lim_{\alpha \to 1} \lim_{s \to 1} \kappa(\eta, s) < 1$ for all η.

At this point, fix such an $\epsilon < \epsilon^*$. Let a constant $M > 0$ be given satisfying

$$M \leq (1 - p_{\text{low}} - p_{\text{high}}) \frac{(a_e - a_h)^2}{\rho_e a_e / \rho_h} \frac{\epsilon(1 - \epsilon)}{\Sigma^*}. \tag{D.11}$$

Note that

$$\lim_{\alpha\to 1}\lim_{s\to 1}\sup_{\eta\in(\epsilon,1-\epsilon)}\left|q(\eta,s)-\bar{a}(\eta)/\bar{\rho}(\eta)\right|=0.$$

Consequently, we may pick $\alpha > \alpha^*$ close enough to 1 and $\zeta > \zeta^*$ close enough to 1 such that

$$\sup_{s\in(\zeta,1)}\sup_{\eta\in(\epsilon,1-\epsilon)}\left|q(\eta,s)-\bar{a}(\eta)/\bar{\rho}(\eta)\right|\le M.$$

Finally, using equation (D.3) and substituting $\kappa < 1$ from (PO), we have $|\sigma(\binom{1}{0}) + \sigma_q|^2 = \frac{(a_e-a_h)^2}{q}\frac{\eta(1-\eta)}{\bar{\rho}q-\bar{a}}$. Note also that $q \le a_e/\rho_h$ and $\bar{\rho} \le \rho_e$ are upper bounds. Then,

$$\mathbb{E}[|\sigma(\tbinom{1}{0}) + \sigma_{q,t}|^2] > (1 - p_{\text{low}} - p_{\text{high}})\frac{(a_e - a_h)^2}{\rho_e a_e/\rho_h}\frac{\epsilon(1-\epsilon)}{M}.$$

Using (D.11), we obtain $\mathbb{E}[|\sigma(\binom{1}{0}) + \sigma_{q,t}|^2] > \Sigma^*$. □

Proof of Proposition D.3. First, we prove that $|\sigma_R|$ is increasing in s. From (D.3), we obtain $|\sigma_R|^2 = \frac{(a_e-a_h)^2}{q}\frac{\eta(1-\eta)}{\bar{\rho}q-\bar{a}}$ on $\{\kappa < 1\}$. Differentiating with respect to s, and using $\partial_s q = \alpha(q^1 - q^0) < 0$, we obtain

$$\partial_s|\sigma_R|^2 = -\eta(1-\eta)\frac{(a_e - a_h)^2}{q(\bar{\rho}q - \bar{a})}\Big[\frac{1}{q} + \frac{\bar{\rho}}{\bar{\rho}q - \bar{a}}\Big]\partial_s q > 0.$$

Next, revisiting the proof of Proposition D.1, we compute on $\{\kappa < 1\}$,

$$\partial_s[(\kappa-\eta)\partial_\eta \log q] = \alpha\Big[(\kappa-\eta)\frac{(q^1)' - (q^0)'}{q} + \frac{\bar{a}(q^1-q^0)}{(a_e-a_h)q^2}\partial_\eta q\Big] < 0.$$

The inequality uses the properties of the φ function in (D.5) to say $(q^1)' - (q^0)' < 0$, along with the obvious facts $q^1 - q^0 < 0$ and $\partial_\eta q > 0$. Using $|(\binom{1}{0}\cdot\sigma_R| = \frac{\sigma}{1-(\kappa-\eta)\partial_\eta \log q}$, we obtain $\partial_s|(\binom{1}{0}\cdot\sigma_R| < 0$.

Using the two claims just proved, and the identity $|\sigma_R|^2 = |(\binom{0}{1}\cdot\sigma_R|^2 + |(\binom{1}{0}\cdot\sigma_R|^2$, we see that $|(\binom{0}{1}\cdot\sigma_R|$ is increasing in s on $\{\kappa < 1\}$. For the same reason, namely $|(\binom{0}{1}\cdot\sigma_R|^2$ is both smaller and increasing faster than $|\sigma_R|$, we have that $|(\binom{0}{1}\cdot\sigma_R|/|\sigma_R|$ increasing in s on $\{\kappa < 1\}$. □

D.5 Correlation between sentiment and fundamentals

What happens if sentiment shocks are correlated with fundamental shocks? To model this, we allow

$$ds_t = \mu_{s,t}dt + \sigma_{s,t}^{(1)}dZ_t^{(1)} + \sigma_{s,t}^{(2)}dZ_t^{(2)}.$$

In Section D.1, we restricted attention to $\sigma_{s,t}^{(1)} = 0$. Without this assumption, equations (D.3) and (D.2) are modified to read:

$$0 = \min\left[1-\kappa,\ \frac{a_e - a_h}{q} - \frac{\kappa-\eta}{\eta(1-\eta)}\left(\frac{(\sigma + \sigma_s^{(1)}\partial_s \log q)^2 + (\sigma_s^{(2)}\partial_s \log q)^2}{(1-(\kappa-\eta)\partial_\eta \log q)^2}\right)\right]$$

$$\sigma_q = \frac{\binom{1}{0}(\kappa-\eta)\sigma\partial_\eta \log q + \sigma_s \partial_s \log q}{1-(\kappa-\eta)\partial_\eta \log q}.$$

The rest of the equilibrium restrictions are identical.

For the present illustration, we additionally assume that $\sigma_{s,t}^{(2)} = 0$, i.e., sentiment shocks *only* load on fundamental shocks. What emerges is the possibility that sentiment shocks "hedge" fundamental shocks: we can have $\sigma_s^{(1)}\partial_s \log q < 0$, which lowers return volatility and raises asset prices. In one extreme, if $\sigma_s^{(1)}\partial_s \log q \to -\sigma$, the price function converges to that of a Fundamental Equilibrium with vanishing fundamental risk $\sigma \to 0$; call this FE(0). At the other end, if $\sigma_s^{(1)}\partial_s \log q \to 0$, the economy resembles the Fundamental Equilibrium with positive fundamental shocks; call this FE(σ). Thus, by constructing our conjectured capital price function as a convex combination of FE(0) and FE(σ), with weights $1-s$ and s, we can ensure that $\sigma_s^{(1)}\partial_s \log q$ endogenously emerges negative. Figure D.5 displays the equilibrium constructed this way.

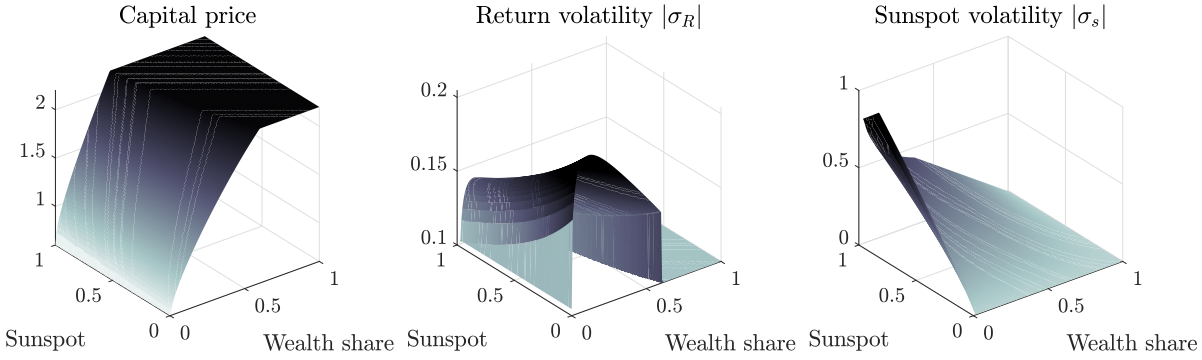


Figure D.5: Capital price q, volatility of capital returns $|\sigma_R|$, and sunspot shock volatility $|\sigma_s|$. Parameters: $\rho_e = \rho_h = 0.05$, $a_e = 0.11$, $a_h = 0.03$, $\sigma = 0.10$.

D.6 Exogenous sunspot dynamics

In Section D.1, we solved for a Markov S-BSE that featured endogenous sunspot dynamics, i.e., (σ_s, μ_s) could potentially depend on η. Here, we show that sunspot equilibria can be built on top of *exogenous* sunspot dynamics as well. As we will show, this construction can be naturally viewed as the limit of equilibria in which the variable s has a vanishing contribution to fundamentals. With that in mind, we actually start from a more general setting in which s can impact fundamental volatility, and then we take the limit as this impact becomes vanishingly small.

Consider the following stochastic volatility model:

$$\begin{aligned}\frac{dK_t}{K_t} &= gdt + \sigma\sqrt{1+\omega s_t}dZ_t \\ ds_t &= \mu_s(s_t)dt + \vartheta\sqrt{1+\omega s_t}dZ_t\end{aligned}$$

where $\vartheta > 0$ is an exogenous parameter and $\omega \in \mathbb{R}$ measures the impact of s_t on capital growth volatility. Thus, the diffusion of s_t, namely $\sigma_s(s) := \vartheta\sqrt{1+\omega s}$, is specified exogenously. Also, $\mu_s(s)$ is an exogenous function that is specified to ensure that $s_t \in (s_{\min}, s_{\max})$, for some pre-specified interval satisfying $s_{\min} \geq 0$ and $\omega s_{\max} > -1$. Such a choice can always be made, e.g., by putting $\mu_s(s) = -(s_{\max} - s)^{-(1+\beta)} + (s - s_{\min})^{-(1+\beta)}$. Note that s_t becomes a sunspot when $\omega = 0$. When $\omega < 0$, the state s_t is an inverse measure of capital's volatility.

For simplicity, we assume there is a single aggregate shock, i.e., Z is a one-dimensional Brownian motion; this can easily be generalized to multiple shocks. Also for simplicity of expressions, we assume here that $\rho_e = \rho_h = \rho$. Then, an equilibrium capital price function $q(\eta, s)$ must satisfy the PDE defined by the following system

$$\begin{aligned}\rho q &= \kappa a_e + (1-\kappa)a_h \\ 0 &= \min\left[1-\kappa,\ \frac{a_e - a_h}{q} - \frac{(\kappa-\eta)(1+\omega s)}{\eta(1-\eta)}\left(\frac{\sigma + \vartheta\partial_s \log q}{1-(\kappa-\eta)\partial_\eta \log q}\right)^2\right].\end{aligned}$$

Technically, the multiplicity arises from the selection of the boundary conditions on $q(\eta, s_{\min})$ and $q(\eta, s_{\max})$, which are not pinned down by any equilibrium restriction.

We perform two exercises. First, we show that there are multiple equilibria for a given set of parameters. We use $\omega < 0$ here, along with $s_{\min} = 0$ and $s_{\max} = 2$. In this case, the "natural" and intuitive solution is for q to increase with s, because volatility decreases. In Figure D.6, we pick a "low" boundary condition for $q(\eta, 0)$ and the solution follows

this intuition.[32]

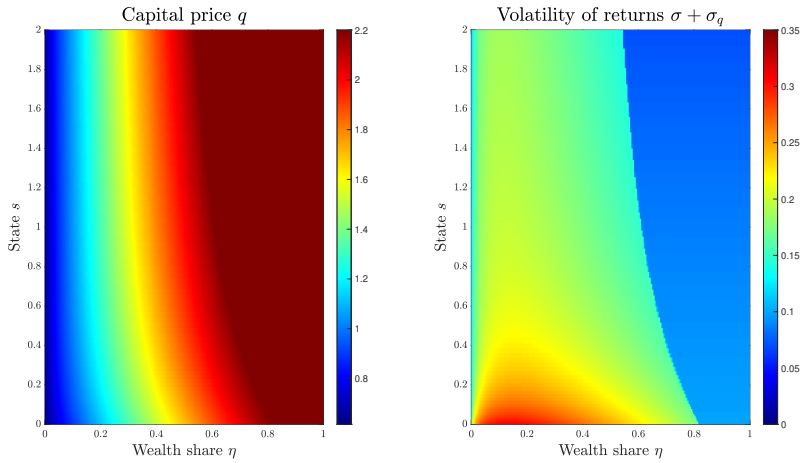


Figure D.6: Equilibrium with $\omega = -0.25$, and the "low" boundary condition for $q(\eta, 0)$, which is a 50% weighted-average of the fundamental equilibrium and the infinite-volatility equilibrium. Other parameters: $\rho_e = \rho_h = 0.05$, $a_e = 0.11$, $a_h = 0.03$, $\sigma = 0.1$, $\vartheta = 0.25$. The boundary condition at $\eta = 0$ is set so that $\kappa(0, s) = 0.01$ for all s.

However, agents could equally well coordinate on a "high" boundary condition, which results in the solution of Figure D.7.[33] Notice the capital price and return volatility exhibit a non-monotonicity in s. At low values of s, q is decreasing in s, while return volatility increases. The very different behavior in Figures D.6 and D.7 is made possible by coordination on the different boundary conditions.

Our second exercise considers the limit $\omega \to 0$. Figure D.8 shows the solution for $\omega = -10^{-6}$, again equipped with the "low" boundary condition for $q(\eta, 0)$. There remains a tremendous amount of variation in the equilibrium as s varies, illustrating convergence to a sunspot equilibrium. Thus, as promised, we are able to construct sunspot equilibria even if the dynamics (σ_s, μ_s) are specified exogenously. In fact, it appears that the amount of price volatility is relatively insensitive to the real effects s has (i.e., the size of ω), which is reminiscent of the "volatility paradox" of Brunnermeier and Sannikov (2014) but one level deeper. Their paradox is that total volatility is only modestly sensitive to exogenous fundamental volatility; our paradox is that total volatility is only modestly sensitive to the *exogenous impact of s on fundamental volatility*.

[32] This "low" boundary condition is a weighted average between the solution with infinite volatility and the fundamental equilibrium solution. The fundamental equilibrium, which is the capital price solution that keeps $s = 0$ fixed forever, is discussed in Online Appendix E. The infinite-volatility solution has $\kappa = \eta$, hence $q = (\eta a_e + (1 - \eta) a_h) / \bar{\rho}(\eta)$.

[33] This "high" boundary condition is a weighted average between $\lim_{v \to 0} \text{FE}(v)$ and $\text{FE}(\sigma)$, where $\text{FE}(\sigma)$ denotes the Fundamental Equilibrium solution with exogenous risk σ.

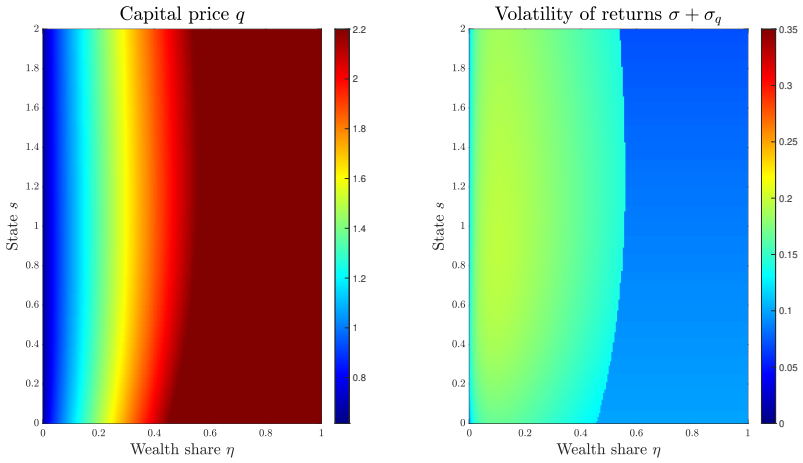


Figure D.7: Equilibrium with $\omega = -0.25$, and the "high" boundary condition for $q(\eta, 0)$, which is a 50% weighted-average of FE(σ) and $\lim_{v \to 0}$ FE(v), where FE(σ) denotes the fundamental equilibrium solution with fundamental risk σ. Other parameters: $\rho_e = \rho_h = 0.05$, $a_e = 0.11$, $a_h = 0.03$, $\sigma = 0.1$, $\vartheta = 0.25$. The boundary condition at $\eta = 0$ is set so that $\kappa(0, s) = 0.01$ for all s.

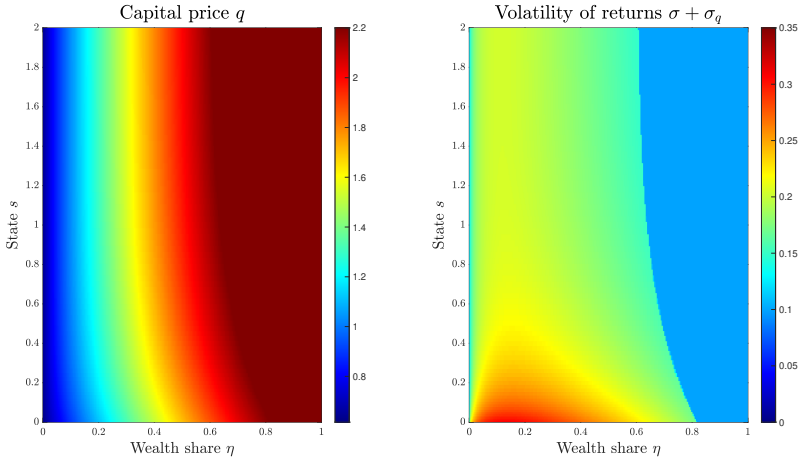


Figure D.8: Equilibrium with near-sunspot $\omega = -10^{-6}$ and the "low" boundary condition for $q(\eta, 0)$, which is a 50% weighted-average of FE(σ) and the infinite-volatility equilibrium (which has $\kappa = \eta$). Other parameters: $\rho_e = \rho_h = 0.05$, $a_e = 0.11$, $a_h = 0.03$, $\sigma = 0.1$, $\vartheta = 0.25$. The boundary condition at $\eta = 0$ is set so that $\kappa(0, s) = 0.01$ for all s.

E Fundamental Equilibria

In this section, we investigate properties of equilibria where sunspot shocks $Z^{(2)}$ are irrelevant and experts' wealth share η serves as the only state variable, i.e., fundamental equilibria. The key equations describing FEs are:

$$q\bar{\rho} = \kappa a_e + (1-\kappa)a_h \tag{E.1}$$

$$0 = \min\left[1-\kappa, \frac{a_e - a_h}{q} - \frac{\kappa-\eta}{\eta(1-\eta)}(\sigma+\sigma_q)^2\right]. \tag{E.2}$$

$$\sigma_q = \frac{(\kappa-\eta)q'/q}{1-(\kappa-\eta)q'/q}\sigma. \tag{E.3}$$

Equation (E.1) just restates (PO). Equation (E.2) is the risk-balance condition (RB) when there is only the fundamental shock $Z^{(1)}$. Equation (E.3) comes from resolving the two-way feedback between wealth share volatility $\sigma_\eta = (\kappa-\eta)(\sigma+\sigma_q)$ and asset-price volatility $\sigma_q = \frac{q'}{q}\sigma_\eta$, which arises from Itô's formula. Finally, wealth share dynamics are given in (13)-(14), restated here for convenience:

$$\begin{aligned}\mu_\eta = {} & -\eta(1-\eta)(\rho_e - \rho_h) + (\kappa - 2\kappa\eta + \eta^2)\frac{\kappa-\eta}{\eta(1-\eta)}\left(\frac{\sigma}{1-(\kappa-\eta)q'/q}\right)^2 \\ & + \delta_h - (\delta_e + \delta_h)\eta\end{aligned} \tag{E.4}$$

$$\sigma_\eta = (\kappa-\eta)(\sigma+\sigma_q). \tag{E.5}$$

We define a fundamental equilibrium as follows, analogously to Lemma 1.

Definition 4. Given $\eta_0 \in (0,1)$, a *Markov fundamental equilibrium* consists of adapted processes $(\eta_t, q_t, \kappa_t, r_t)_{t\geq 0}$ such that $q_t = q(\eta_t)$ for some function $q(\cdot)$, such that (E.1)-(E.3) and (11) hold, and such that (E.4)-(E.5) describe dynamics of η_t.

Note that the interest rate r_t can be simply set from (11), given the other variables, and it affects no other equilibrium equation. Similarly, the dynamics of η_t are set from (E.4)-(E.5), and they affect none of (E.1)-(E.3). Finally, κ_t can be obtained from (η_t, q_t) directly from equation (E.1). Hence, the critical object in a fundamental equilibrium is the function q.

We document some properties of fundamental equilibria, where we additionally impose the standard full-deleveraging boundary condition $\kappa(0) = 0$. Khorrami and Mendo (2024) shows in their online appendix that this boundary condition is the only one that survives a simple refinement based on a vanishingly-small limited commitment friction.

Lemma E.1. *Assuming it exists, suppose q is a fundamental equilibrium in the sense of Definition 4. Assume $\kappa(0+) = 0$. Consider only equilibria with $\sigma + \sigma_q \geq 0$. Define $\eta^* := \inf\{\eta : \kappa = 1\}$. Then, the following hold:*

(i) Free boundary problem for (q, η^*):

$$(\bar{\rho}q - \eta a_e - (1-\eta)a_h)\frac{q'}{q} = a_e - a_h - \sigma\sqrt{q\frac{\bar{\rho}q - \eta a_e - (1-\eta)a_h}{\eta(1-\eta)}} \quad \text{for all} \quad \eta \in (0, \eta^*).$$

(ii) $\eta a_e + (1-\eta)a_h < \bar{\rho}q < a_e$, for all $\eta \in (0, \eta^*)$.

(iii) $\frac{q'(0+)}{q(0+)} = \frac{a_e}{a_h} - \frac{\rho_e}{\rho_h} + \rho_h\left(\frac{a_e - a_h}{\sigma a_h}\right)^2$.

(iv) If σ is sufficiently small, then $q' > \frac{a_e - a_h}{\bar{\rho}}$, for $\eta \in (0, \eta^*)$.

(v) If σ is sufficiently small, then $\frac{\rho_h}{\rho_e}\left(\frac{1 - a_h/a_e}{\sigma^2} - 1 + \frac{\rho_h}{\rho_e}\right)^{-1} < \eta^* < 1$.

(vi) If σ is sufficiently small, then the equilibrium is unique.

Proof of Lemma E.1. Recall $\bar{\rho} := \eta\rho_e + (1-\eta)\rho_h$. By analogy, let $\bar{a} := \eta a_e + (1-\eta)a_h$.

(i) Start from equation (E.2), plug in (E.1) and (E.3), and rearrange to obtain the result, where we have selected the solution with $1 > (\kappa - \eta)\frac{q'}{q}$ to make sure $\sigma + \sigma_q \geq 0$ as mentioned.

(ii) The first inequality, which is equivalent to $\kappa > \eta$, is a direct implication of equation (E.2). The second inequality, equivalent to $\kappa < 1$, is a restatement of the definition of η^*.

(iii) Start from equation (E.2). Taking the limit $\eta \to 0$, and using $\kappa(0+) = 0$, delivers an equation for $\kappa'(0+)$. Differentiating (E.1), we may then substitute $\kappa'(0+) = \frac{\rho_h q'(0+) + (\rho_e - \rho_h)q(0+)}{a_e - a_h}$. Rearranging, we obtain the desired result.

(iv) By part (iii), there exists $\eta^\circ > 0$ and $\bar{\sigma} > 0$ such that uniformly for all $\sigma < \bar{\sigma}$, we have $q' > \frac{a_e - a_h}{\bar{\rho}}$ on the set $\{\eta < \eta^\circ\}$. On the set $\{\eta^\circ \leq \eta < \eta^*\}$, we know that $\kappa - \eta$ is bounded away from zero, uniformly for all $\sigma < \bar{\sigma}$. Using the expression in part (i), the fact that q is bounded by $a_e/\bar{\rho}$ uniformly for all σ, and the previous fact about $\kappa - \eta = \bar{\rho}q - \bar{a}$, we can write

$$q' = \frac{a_e - a_h}{\bar{\rho}q - \bar{a}}q - o(\sigma), \quad \eta \in (\eta^\circ, \eta^*).$$

Therefore,

$$q' + o(\sigma) = \frac{a_e - a_h}{\bar{\rho} q - \bar{a}} q = \frac{a_e - a_h}{\bar{\rho}} \frac{q}{q - \bar{a}/\bar{\rho}} > \frac{a_e - a_h}{\bar{\rho}}, \quad \eta \in (\eta^\circ, \eta^*),$$

where the last inequality is due to $\bar{\rho} q > \bar{a}$ [part (ii)]. Taking σ is small enough implies the result on (η°, η^*), which we combine with the result on $(0, \eta^\circ)$ to conclude.

(v) Consider the function $\tilde{q} := \bar{a}/\bar{\rho}$, whose derivative is $\tilde{q}' = \frac{a_e - a_h}{\bar{\rho}} - \frac{\bar{a}}{\bar{\rho}} \frac{\rho_e - \rho_h}{\bar{\rho}} < \frac{a_e - a_h}{\bar{\rho}}$. Combining this result with part (iv), we obtain $q' > \tilde{q}'$. If $\tilde{q}$ was the capital price, then equation (E.1) implies the associated capital share $\tilde{\kappa} = \eta$. On the other hand, the fact that $q' > \tilde{q}'$ implies $\kappa' > \tilde{\kappa}' = 1$, which implies $\eta^* < 1$.

Next, consider $\eta \in (\eta^*, 1)$ so that $\kappa = 1$ (see Lemma A.6 of Khorrami and Mendo, 2024 for a proof that we must have $\kappa = 1$ for all $\eta > \eta^*$). By equation (E.2), with $q = a_e/\bar{\rho}$, we must have

$$\sigma^2 \leq \eta \bar{\rho} \frac{a_e - a_h}{a_e} \left(1 + (1 - \eta) \frac{\rho_e - \rho_h}{\bar{\rho}}\right)^2, \quad \eta \geq \eta^*.$$

This is equivalent to

$$1 \leq \eta \frac{\rho_e}{\rho_h} \left(\frac{a_e - a_h}{a_e \sigma^2} \rho_e - 1 + \frac{\rho_h}{\rho_e}\right), \quad \eta \geq \eta^*.$$

Substituting $\eta = \eta^*$, and rearranging, we obtain the first inequality. There is no contradiction with $\eta^* < 1$, due to the assumption that σ can be made small enough.

(vi) See Lemma A.7 of Khorrami and Mendo (2024). □

F Model extensions

F.1 Partial equity issuance

We extend the model to allow some equity issuance by capital holders, subject to a constraint. In particular, at any point of time, agents managing capital can issue some equity to the market, but the issuer must keep at least $\chi \in [0,1]$ fraction of their capital risk—this is a so-called "skin-in-the-game" constraint. In other words, if experts and households retain χ_e and χ_h of their capital risk, respectively, it must be the case that

$$\chi_{\ell,t} \geq \chi, \quad \ell \in \{e,h\}. \tag{F.1}$$

Thus, the frictionless model corresponds to $\chi = 0$, while our baseline model corresponds to $\chi = 1$. Outside equity contracts are risky, having risk exposure σ_R (the endogenous capital return volatility), so they must promise an excess return $\sigma_R \cdot \pi$, where π is the equilibrium risk price vector that applies to securities tradable by both experts and households.

Agents' dynamic budget constraints are now given by

$$\begin{aligned} dn_{\ell,t} = &\Big[(n_{\ell,t} - q_t k_{\ell,t})r_t - c_{\ell,t} + a_\ell k_{\ell,t}\Big]dt + d(q_t k_{\ell,t}) \\ &+ [\theta_{\ell,t} n_{\ell,t} - (1-\chi_{\ell,t}) q_t k_{\ell,t}]\sigma_{R,t} \cdot (\pi_t dt + dZ_t). \end{aligned} \tag{F.2}$$

The second line of (F.2) contains the new terms pertaining to equity-issuance: $\theta_{\ell,t} \geq 0$ denotes purchases of equity contracts in the market, per unit of wealth, while $\chi_{\ell,t}$ denotes the fraction of capital risk. Notice that it will be without loss of generality to assume $\chi_{\ell,t} = \chi$ at all times and for all agents, because the purchase variable $\theta_{\ell,t}$ is available as a control. For example, an agent with a slack equity-issuance constraint ($\chi_\ell > \chi$) could issue equity to the constraint (F.1) and then buy back such equity by increasing their θ_ℓ control. Going forward, we simply assume $\chi_{e,t} = \chi_{h,t} = \chi$. The presence of a public equity market implies an additional market clearing condition for equity securities, namely

$$\theta_{e,t} N_{e,t} + \theta_{h,t} N_{h,t} = (1-\chi) q_t K_t. \tag{F.3}$$

At this point, we may solve for equilibrium.

Model solution. The introduction of equity issuance changes nothing about optimal consumption choices, so the price-output relation (PO) still holds.

Optimal portfolio choice now implies the following four FOCs:

$$\mu_{R,e} - (1-\chi)\sigma_R \cdot \pi - r = \chi\Big(\frac{\chi q k_e}{n_e} + \theta_e\Big)|\sigma_R|^2 \tag{F.4}$$

$$\mu_{R,h} - (1-\chi)\sigma_R \cdot \pi - r \leq \chi\Big(\frac{\chi q k_h}{n_h} + \theta_h\Big)|\sigma_R|^2, \quad \text{with equality if } k_h > 0 \tag{F.5}$$

$$\Big(\frac{\chi q k_e}{n_e} + \theta_e\Big)|\sigma_R|^2 \geq \sigma_R \cdot \pi, \quad \text{with equality if } \theta_e > 0 \tag{F.6}$$

$$\Big(\frac{\chi q k_h}{n_h} + \theta_h\Big)|\sigma_R|^2 \geq \sigma_R \cdot \pi, \quad \text{with equality if } \theta_h > 0 \tag{F.7}$$

where $\mu_{R,\ell} := \frac{a_\ell}{q} + g + \mu_q + \sigma\sigma_q \cdot \left(\begin{smallmatrix}1\\0\end{smallmatrix}\right)$ is the expected return on capital for agent ℓ. Equations (F.4)-(F.5) are the FOCs for capital holdings, and (F.6)-(F.7) are the FOCs for equity purchases. Note that the equality in (F.4) assumes $k_e > 0$, which is easy to verify must always be the case in equilibrium, exactly as in the baseline model.

We can derive a new "risk-balance" condition, analogously to the baseline model. If in addition to $k_e > 0$ we have $k_h > 0$, then we cannot simultaneously have $\theta_e > 0$, as this would contradict $\mu_{R,e} > \mu_{R,h}$. Thus, $\theta_e = 0$ whenever $k_h > 0$, and so we may difference (F.4)-(F.5) and use the market clearing condition (F.3) to substitute $\theta_h = \frac{1-\chi}{1-\eta}$, which leads to

$$0 = \min\Big[1-\kappa,\ \frac{a_e - a_h}{q} - \chi\frac{\chi\kappa - \eta}{\eta(1-\eta)}|\sigma_R|^2\Big]. \tag{RBE}$$

In addition to (RBE), equation (F.7) must hold with equality and (F.6) with inequality when $\kappa < 1$. By (F.7) and the derived expression $\theta_h = \frac{1-\chi}{1-\eta}$, we have $\sigma_R \cdot \pi = \frac{1-\chi\kappa}{1-\eta}|\sigma_R|^2$, for which a viable solution is

$$\pi = \frac{1-\chi\kappa}{1-\eta}\sigma_R, \quad \text{if } \kappa < 1. \tag{F.8}$$

Using this expression for π, (F.6) requires $\chi\kappa \geq \eta$, which holds by equation (RBE).

By contrast, when $k_h = 0$ (so $\kappa = 1$), equations (F.6)-(F.7) imply

$$\pi = \min\Big(1, \frac{1-\chi}{1-\eta}\Big)\sigma_R, \quad \text{if } \kappa = 1. \tag{F.9}$$

To prove this, combine the two possible cases:

(i) Suppose $\theta_e > 0$. Note that $\theta_h = 0$ cannot occur, as $\theta_e > 0$ implies $\sigma_R \cdot \pi > 0$ while $k_h = \theta_h = 0$ implies the opposite. Thus, we may combine (F.6)-(F.7), both evaluated under equality, to obtain $\theta_h = \theta_e + \frac{\chi}{\eta}$. Plugging this result into market clearing (F.3)

yields $\theta_e = 1 - \chi/\eta$ and $\theta_h = 1$. Using $\theta_h = 1$ back in (F.7), we obtain $\sigma_R \cdot \pi = |\sigma_R|^2$, for which a viable solution is $\pi = \sigma_R$. Note that $\theta_e = 1 - \chi/\eta > 0$ if and only if $\eta > \chi$.

(ii) Suppose $\theta_e = 0$. Note that market clearing (F.3) implies $\theta_h = \frac{1-\chi}{1-\eta} > 0$ in this case. By (F.7), we have $\sigma_R \cdot \pi = \frac{1-\chi}{1-\eta}|\sigma_R|^2$, for which a viable solution is $\pi = \frac{1-\chi}{1-\eta}\sigma_R$. Using the expression for π, (F.6) requires $\eta \leq \chi$.

Putting the results of (F.8)-(F.9) together, we have that

$$\pi = \min\left(1, \frac{1-\chi\kappa}{1-\eta}\right)\sigma_R. \tag{F.10}$$

Finally, the riskless interest rate can be derived as always, by summing a $(\kappa, 1-\kappa)$-weighted-average of equations (F.4)-(F.5) to get

$$\begin{aligned} r = \frac{\kappa a_e + (1-\kappa)a_h}{q} + g + \mu_q + \sigma\sigma_q \cdot \left(\begin{smallmatrix}1\\0\end{smallmatrix}\right) - (1-\chi)\sigma_R \cdot \pi \\ - \chi\left[\kappa\left(\frac{\chi\kappa}{\eta} + \theta_e\right) + (1-\kappa)\left(\frac{\chi(1-\kappa)}{1-\eta} + \theta_h\right)\right]|\sigma_R|^2. \end{aligned} \tag{F.11}$$

We can simplify this equation using the following facts. First, from the discussion above, $\theta_h > 0$ always holds, so that (F.7) holds with equality, hence $\theta_h = \frac{\sigma_R \cdot \pi}{|\sigma_R|^2} - \frac{\chi(1-\kappa)}{1-\eta}$. Next, we may use the market clearing condition (F.3) to obtain $\theta_e = \frac{1-\chi}{\eta} - \frac{1-\eta}{\eta}\theta_h$. We use these two facts to eliminate θ_e and θ_h from (F.11), then we substitute the solution for π from (F.10), and finally we simplify the result to obtain

$$r = \frac{\kappa a_e + (1-\kappa)a_h}{q} + g + \mu_q + \sigma\sigma_q \cdot \left(\begin{smallmatrix}1\\0\end{smallmatrix}\right) - |\sigma_R|^2 - \left(\frac{\chi\kappa}{\eta} - 1\right)\max\left(0, \frac{\chi\kappa - \eta}{1-\eta}\right). \tag{F.12}$$

This completes the derivation of equilibrium.

Properties of equilibrium. For any $\chi > 0$, we can construct S-BSEs using a similar procedure as the baseline model, i.e., by solving equation (PO) for κ as a function of (η, q), and then substituting this into (RBE) to also solve for $|\sigma_R|$ as a function of (η, q). Importantly, any solution to equation (RBE) requires $\chi\kappa \geq \eta$, and so the effect of lower equity issuance frictions (lower χ) is to reduce the range of possible fluctuations of κ, hence q, for any given η. This effect is depicted in Figure F.1, which shows that the range of possible fluctuations for price q is unambiguously shrinking as χ falls. However, the ergodic set is $\{\eta \leq \chi\}$, and so the resulting equilibrium dynamics actually look very

similar for any $\chi > 0$. On the other hand, if $\chi = 0$, no sunspot equilibrium can exist, as shown in the main text.

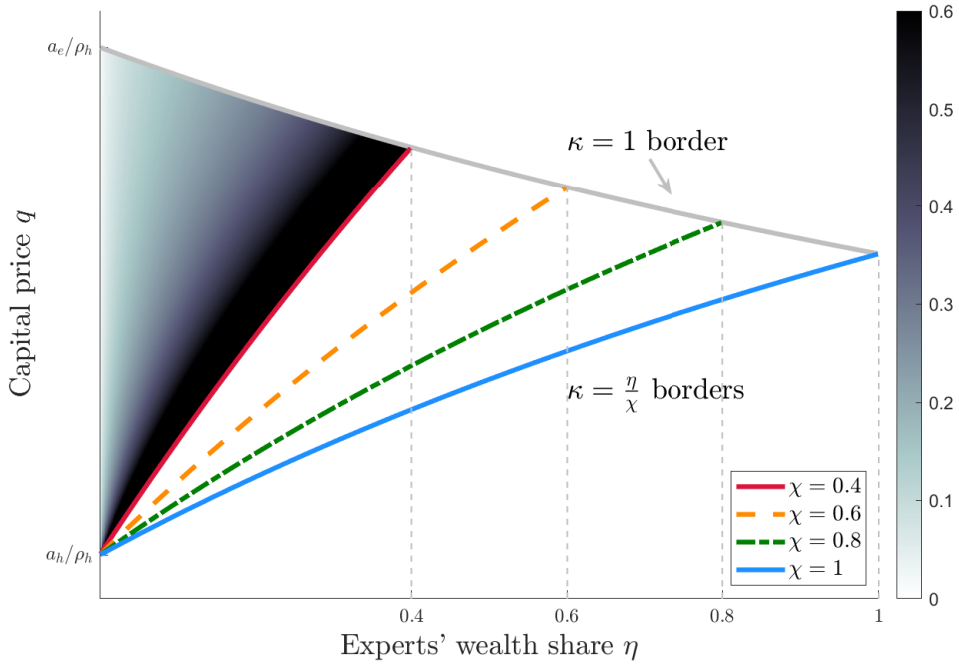


Figure F.1: Colormap of volatility $|\sigma_R|$ as a function of (η, q), in the region $\mathcal{D} := \{(\eta, q) : \eta \in (0,1)$ and $(\eta/\chi)a_e + (1 - \eta/\chi)a_h < q\bar{\rho}(\eta) \leq a_e\}$. Volatility is truncated for aesthetic purposes (because $|\sigma_R| \to \infty$ as $\kappa \to \eta/\chi$). Parameters: $\rho_e = 0.07$, $\rho_h = 0.05$, $a_e = 0.11$, $a_h = 0.03$.

F.2 Continuum of productivities

Here, we generalize the productivity distribution to a continuum. While the analysis below is rather involved algebraically, we essentially prove that a sentiment equilibrium exists.

Let agents' productivities be denoted $a \in [\underline{a}, \bar{a}]$. Without loss of generality, we may index agent types by their productivity, letting agents' discount rates be denoted $\rho(a)$. Let $\eta(a)$ and $\kappa(a)$ be the wealth and capital shares of agents with productivity a, so η and κ are functions mapping $[\underline{a}, \bar{a}]$ into the $\mathbb{R}_+$. We have $\int \eta(a)da = 1$ and $\int \kappa(a)da = 1$, so in fact η and κ are density functions. Any equilibrium features a threshold a^* such that agents with $a \geq a^*$ are holding capital, while agents with $a < a^*$ hold no capital.[34] The threshold a^* will be time-varying.

[34] If $\kappa(a') = 0$ while $\kappa(a) > 0$ for $a' > a$, we would arrive at a contradiction by comparing these two agents' Euler equations.

Euler equations and equilibrium capital shares. In this continuum model, the Euler equations, differenced between agent $\bar{a}$ and agent a, say

$$\frac{\bar{a}-a}{q} = \Big(\frac{\kappa(\bar{a})}{\eta(\bar{a})} - \frac{\kappa(a)}{\eta(a)}\Big)|\sigma_R|^2, \qquad \text{if} \quad a \geq a^*; \tag{F.13}$$

$$\frac{\bar{a}-a}{q} > \frac{\kappa(\bar{a})}{\eta(\bar{a})}|\sigma_R|^2, \qquad \text{if} \quad a < a^*. \tag{F.14}$$

Combine two of the differenced Euler equations, eliminating $q|\sigma_R|$:

$$\frac{1}{a'-a}\Big(\frac{\kappa(a')}{\eta(a')} - \frac{\kappa(a)}{\eta(a)}\Big) = \frac{1}{a''-a'}\Big(\frac{\kappa(a'')}{\eta(a'')} - \frac{\kappa(a')}{\eta(a')}\Big), \quad \text{for} \quad a, a', a'' > a^* \tag{F.15}$$

Consequently, if we define $x(a) := \kappa(a)/\eta(a)$, we obtain the differential equation $\frac{d^2}{da^2}x(a) = 0$ for all $a > a^*$, i.e.,

$$x(a) = \alpha + \beta(a - a^*), \tag{F.16}$$

for some α, β. A requirement is that $\beta > 0$ so that $x(a)$ is an increasing function, by (F.13).

The partial efficiency case is when $a^* > \underline{a}$. If so, then we know that $x(a^*) = 0$. Applying this boundary condition to the solution for $x(a)$, we have $\alpha = 0$. We also have the integration condition $1 = \int \kappa(a)da = \int x(a)\eta(a)da = \beta \int_{a^*}^{\bar{a}}(a - a^*)\eta(a)da$. Hence,

$$\kappa(a) = \frac{\max[0, (a-a^*)\eta(a)]}{\int_{a^*}^{\bar{a}}(a'-a^*)\eta(a')da'} \quad \text{if} \quad a^* > \underline{a}. \tag{F.17}$$

We see that κ depends on the wealth distribution η and whatever drives the threshold a^*, which may be η but may also include sentiment variables. The economic intuition conveyed by (F.17) clearly delineates that self-fulfilling fire sales come from a^*, as the only other endogenous object in the formula is the wealth distribution η.

The inefficient case is when $a^* = \underline{a}$. If so, then applying the integration condition $1 = \int \kappa(a)da = \int x(a)\eta(a)da = \alpha \int \eta(a)da + \beta \int (a - \underline{a})\eta(a)da = \alpha + \beta \int (a - \underline{a})\eta(a)da$ yields

$$\kappa(a) = \alpha\eta(a) + (1-\alpha)\frac{(a-a^*)\eta(a)}{\int (a'-a^*)\eta(a')da'} \quad \text{if} \quad a^* = \underline{a}. \tag{F.18}$$

Here, the driver of self-fulfilling behavior is the object α, which is not yet pinned down. We must have $\alpha > 0$ so that $\kappa(\underline{a}) > 0$, as conjectured for this inefficient case. We also

must have $\alpha < 1$ so that $\beta = \frac{1-\alpha}{\int (a-\underline{a})\eta(a)da} > 0$ and $x(a)$ is an increasing function, as required by (F.13).

Introduce sunspot state driving dynamics. To nest both cases, we let s be an endogenous object, to be determined. Think of this as the self-fulfilling state variable. We suppose s fluctuates between -1 and 1, according to a diffusion

$$ds_t = \mu_{s,t}dt + \sigma_{s,t} \cdot dZ_t, \quad s_t \in (-1, 1). \tag{F.19}$$

We then set α and a^* as follows Let

$$\alpha(s) = \max(0, s). \tag{F.20}$$

We also set a^* according to a linear function of s when $s < 0$, and set $a^* = \underline{a}$ when $s \geq 0$:

$$a^*(s) = \begin{cases} \underline{a}, & \text{if } s \geq 0; \\ \underline{a} - (\bar{a} - \underline{a})s, & \text{if } s < 0. \end{cases} \tag{F.21}$$

With these two objects, we can write the general solution for $\kappa(a)$ that for any s:

$$\kappa(a; \eta, s) = \alpha(s)\eta(a) + (1 - \alpha(s))\frac{(a - a^*(s))\eta(a)}{\int_{a^*(s)}^{\bar{a}} (a' - a^*(s))\eta(a')da'}. \tag{F.22}$$

Notice that (F.22) nests both cases (F.17)-(F.18). In this construction, the latent variable s_t captures self-fulfilling dynamics. When s_t is negative, we have partial efficiency: the least productive agents are out of the market, and as s_t falls towards -1, capital gets increasingly concentrated in the hands of agents with maximal productivity. When s_t is positive, we have inefficiency: all agents are holding capital, and as s_t rises towards 1, the capital shares in (F.22) converge to wealth shares, which we show below is the most inefficient situation where volatility is infinite.

Determine capital price and return volatility. Given $(\eta(a))_{a\in[\underline{a},\bar{a}]}$ and s, hence the capital shares $(\kappa(a))_{a\in[\underline{a},\bar{a}]}$, the price-output relation determines q via

$$q = \frac{\int \kappa(a)ada}{\int \rho(a)\eta(a)da}. \tag{F.23}$$

Equation (F.23) defines a functional $q(\eta, s)$ for the capital price. For completeness, let us

write out this functional explicitly for the two cases $s < 0$ and $s > 0$:

$$q(\eta, s) = \frac{\int_{a^*(s)}^{\bar{a}} a(a - a^*(s))\eta(a)da}{(\int_{a^*(s)}^{\bar{a}}(a' - a^*(s))\eta(a')da')(\int \rho(a)\eta(a)da)}, \quad \text{if} \quad s < 0 \tag{F.24}$$

and

$$q(\eta, s) = s\frac{\int a\eta(a)da}{\int \rho(a)\eta(a)da} + (1-s)\frac{\int a(a - \underline{a})\eta(a)da}{(\int (a' - \underline{a})\eta(a')da')(\int \rho(a)\eta(a)da)}, \quad \text{if} \quad s > 0 \tag{F.25}$$

Given this functional $q(\eta, s)$, we may use Itô's formula to obtain the drift μ_q and diffusion σ_q, once we have fully determined the dynamics of η and s (see below). With q in hand, the Euler equation (F.13) for the marginal agent ($a = a^*$) then determines volatility as

$$|\sigma_R|^2 = \frac{\bar{a} - a^*}{q(\frac{\kappa(\bar{a})}{\eta(\bar{a})} - \frac{\kappa(a^*)}{\eta(a^*)})} \tag{F.26}$$

Notice, as claimed above, that when $\kappa(a)/\eta(a) \to 1$ (which is when $s \to 1$), volatility diverges. Summarizing, given $(\eta(a))_{a\in[\underline{a},\bar{a}]}$ and s, we pin down $(\kappa(a))_{a\in[\underline{a},\bar{a}]}$, q, and $|\sigma_R|$.

Dynamics of wealth shares. We have already specified the dynamics of s above in (F.19), subject to some restrictions to be described shortly. On the other hand, the dynamics of wealth shares are determined as usual by applying Itô's formula to the definition $\eta_t(a) := n_t(a)/\int n_t(a')da'$, and using the dynamic budget constraint. The result is[35]

$$\begin{aligned}\frac{d\eta_t(a)}{\eta_t(a)} = &\Big[\int \rho(a')\eta_t(a')da' - \rho(a) + |\sigma_{R,t}|^2\Big(\int \eta_t(a')\Big((\frac{\kappa_t(a)}{\eta_t(a)})^2 - (\frac{\kappa_t(a')}{\eta_t(a')})^2\Big)da' - \frac{\kappa_t(a) - \eta_t(a)}{\eta_t(a)}\Big)\Big]dt \\ &+ \frac{\kappa_t(a) - \eta_t(a)}{\eta_t(a)}\sigma_{R,t} \cdot dZ_t\end{aligned} \tag{F.27}$$

[35]The dynamic budget constraint, after substituting optimal consumption $c_t(a) = \rho(a)n_t(a)$, re-writing the optimal capital holdings $qk_t(a)/n_t(a)$ as $\kappa_t(a)/\eta_t(a)$, and using the capital complementary slackness condition $\kappa_t(a)(\mu_{R,t}(a) - r_t) = \frac{\kappa_t(a)^2}{\eta_t(a)}|\sigma_{R,t}|^2$, says

$$dn_t(a) = n_t(a)\Big[r_t - \rho(a) + (\frac{\kappa_t(a)}{\eta_t(a)})^2|\sigma_{R,t}|^2\Big]dt + n_t(a)\frac{\kappa_t(a)}{\eta_t(a)}\sigma_{R,t} \cdot dZ_t$$

Aggregating this, and using Fubini's theorem for stochastic integration, we have

$$d\Big(\int n_t(a')da'\Big) = \Big(\int n_t(a')\Big[r_t - \rho(a') + (\frac{\kappa_t(a')}{\eta_t(a')})^2|\sigma_{R,t}|^2\Big]da'\Big)dt + \Big(\int n_t(a')\frac{\kappa_t(a')}{\eta_t(a')}da'\Big)\sigma_{R,t} \cdot dZ_t$$

We then use Itô's quotient rule, along with the definition of $\eta_t(a) := n_t(a)/\int n_t(a')da'$ and the aggregation rules $\int \eta_t(a)da = 1$ and $\int \kappa_t(a)da = 1$, to obtain result (F.27).

Notice, importantly, that the drift and diffusion are completely pinned down given $(\eta_t(a))_{a\in[\underline{a},\bar{a}]}$ and s_t, with the exception of the fraction of return variance $|\sigma_R|^2$ that comes from the fundamental versus sunspot shocks. In particular, as in the baseline model we may choose any functional $\vartheta(\eta, s) \in [0,1]$ and set $\sigma_R^{(1)} = \sqrt{\vartheta(\eta,s)}|\sigma_R|$, where $|\sigma_R|$ has already been pinned down by (F.26).

Interest rate. Finally, the interest rate is determined as usual: aggregate each agent's capital Euler equation, weighted by κ, and use the price-output relation, to obtain

$$r = \int \rho(a)\eta(a)da + g + \mu_q + \sigma\sigma_q \cdot \left(\begin{smallmatrix}1\\0\end{smallmatrix}\right) - |\sigma_R|^2 \int \frac{\kappa(a)^2}{\eta(a)}da \tag{F.28}$$

As always, the drift μ_q is indeterminate, here because it depends on the dynamics of s, which are indeterminate. We will see this explicitly shortly.

Restrictions on sunspot dynamics. Let us close the model by determining the mapping between the s dynamics and the q dynamics. Given the functional $q(\eta, s)$, we apply an infinite-dimensional version of Itô's formula. To do this, first compute the derivatives of q as

$$\partial_s q(\eta,s) = \frac{1}{\int \rho(a)\eta(a)da} \begin{cases} \frac{(\bar{a}-\underline{a})\int_{a^*}^{\bar{a}} a\eta(a)da}{\int_{a^*}^{\bar{a}}(a'-a^*)\eta(a')da'} - \frac{(\bar{a}-\underline{a})(\int_{a^*}^{\bar{a}} a'(a'-a^*)\eta(a')da')(\int_{a^*}^{\bar{a}}\eta(a')da')}{(\int_{a^*}^{\bar{a}}(a'-a^*)\eta(a')da')^2} & \text{if } s<0; \\ \int a\eta(a)da - \frac{\int a(a-\underline{a})\eta(a)da}{\int(a'-\underline{a})\eta(a')da'} & \text{if } s>0. \end{cases} \tag{F.29}$$

$$\partial_{\eta(a)} q(\eta,s) = -q\frac{\rho(a)}{\int \rho(a')\eta(a')da'} \tag{F.30}$$
$$+ \frac{1}{\int \rho(a')\eta(a')da'} \begin{cases} \frac{a(a-a^*)}{\int_{a^*}^{\bar{a}}(a'-a^*)\eta(a')da'} - \frac{(a-a^*)\int_{a^*}^{\bar{a}} a'(a'-a^*)\eta(a')da'}{(\int_{a^*}^{\bar{a}}(a'-a^*)\eta(a')da')^2} & \text{if } s<0; \\ sa + (1-s)[\frac{a(a-\underline{a})}{\int(a'-\underline{a})\eta(a')da'} - \frac{(a-\underline{a})\int a'(a'-\underline{a})\eta(a')da'}{(\int(a'-\underline{a})\eta(a')da')^2}] & \text{if } s>0. \end{cases}$$

and so on for $\partial_{ss}q$, $\partial_{\eta(a)\eta(a')}q$, and $\partial_{s\eta(a)}q$. Then, plug these into to Itô's formula to get

$$\sigma_q = \sigma_s\partial_s q + \int(\sigma_{\eta(a)}\partial_{\eta(a)}q)da \tag{F.31}$$

and

$$\mu_q = \mu_s \partial_s q + \frac{1}{2}\sigma_s \cdot \partial_{ss} q + \int (\mu_{\eta(a)} \partial_{\eta(a)} q) da \\ + \frac{1}{2} \iint (\sigma_{\eta(a)} \cdot \sigma_{\eta(a')} \partial_{\eta(a)\eta(a')} q) da da' + \int (\sigma_s \cdot \sigma_{\eta(a)} \partial_{s\eta(a)} q) da \tag{F.32}$$

where $\mu_{\eta(a)}$ and $\sigma_{\eta(a)}$ are the drift and diffusion of $\eta(a)$ in equation (F.27), while μ_s and σ_s are the drift and diffusion of s in equation (F.19).

Using these mappings, we can identify restrictions on the s dynamics. From equation (F.31), and using the definition $\sigma_R = \sigma(\begin{smallmatrix}1\\0\end{smallmatrix}) + \sigma_q$, we have that

$$\sigma_R = \frac{\sigma(\begin{smallmatrix}1\\0\end{smallmatrix}) + \sigma_s \partial_s q}{1 - \int (\kappa(a) - \eta(a)) \partial_{\eta(a)} q da}. \tag{F.33}$$

Comparing (F.33) and (F.26), and using the arbitrary functional $\vartheta(\eta, s) = (\sigma_R^{(1)})^2 / |\sigma_R|^2$ described earlier to parameterize the fraction of return variance from the fundamental shock, we see that we require the following restrictions on σ_s:

$$(\sigma + \sigma_s^{(1)} \partial_s q)^2 = \vartheta \frac{\bar{a} - a^*}{q(\frac{\kappa(\bar{a})}{\eta(\bar{a})} - \frac{\kappa(a^*)}{\eta(a^*)})} \left(1 - \int (\kappa(a) - \eta(a)) \partial_{\eta(a)} q da\right)^2 \tag{F.34}$$

$$\sigma^2 + |\sigma_s|^2 (\partial_s q)^2 + 2\sigma \sigma_s^{(1)} \partial_s q = \frac{\bar{a} - a^*}{q(\frac{\kappa(\bar{a})}{\eta(\bar{a})} - \frac{\kappa(a^*)}{\eta(a^*)})} \left(1 - \int (\kappa(a) - \eta(a)) \partial_{\eta(a)} q da\right)^2 \tag{F.35}$$

Equation (F.34) determines $\sigma_s^{(1)}$, and then equation (F.35) determines $\sigma_s^{(2)}$. Choosing σ_s this way will implement exactly the desired return exposure ϑ from the fundamental shock. Finally, notice from equation (F.32) that, as claimed earlier, the price drift μ_q is indeterminate, since μ_s is indeterminate. We may pick an arbitrary functional $m(\eta, s)$, subject to certain boundary conditions as $s \to \pm - 1$, and set $\mu_s = m$. The boundary conditions need to keep $s_t \in (-1, 1)$ as conjectured in the above construction. This can be achieved by picking any function that satisfies $\sup_\eta \lim_{s \to -1} \frac{m(\eta,s)}{(1+s)^\nu} = \sup_\eta \lim_{s \to +1} \frac{-m(\eta,s)}{(1-s)^\nu} = +\infty$ for some ν sufficiently large.

F.3 Epstein-Zin preferences

Our paper assumes agents have symmetric logarithmic preferences, which provides significant tractability. While demonstrating the existence of sentiment equilibria with non-log preferences is beyond the scope of the paper, we provide here the set of equations

such an equilibrium must satisfy. These are useful to at least formalize the argument for why we conjecture such a sentiment equilibrium would exist for any preference specification.

Why should multiplicity survive? Intuitively, the key forces driving multiplicity are (i) the link between real capital allocation and capital price (what we have called the price-output relation); and (ii) the link between the capital allocation and risk considerations (what we have called the risk-balance condition). These mechanisms do not depend on preferences and should survive beyond our log assumption. Analytically, we show below that non-log preferences preserve the same the degrees of freedom as in the logarithmic case, suggesting that a sentiment equilibrium construction should be possible.

Let us now provide the equations for the case of recursive preferences of the Duffie-Epstein-Zin type (risk aversion of γ and IES of ψ^{-1}). Let the value function of agent $i \in \{e, h\}$ be written as $V_i = \frac{(\xi_i n_i)^{1-\gamma}}{1-\gamma}$, where n_i is individual net worth and where ξ_i summarizes the information of aggregate states. We will only study the economy with Brownian shocks, so we let ξ_i follow

$$d\xi_i = \xi_i \left(\mu_{\xi,i} dt + \sigma_{\xi,i} \cdot dZ_t\right) \tag{F.36}$$

The drift and diffusion $\mu_{\xi,i}$ and $\sigma_{\xi,i}$ must be determined in equilibrium. The HJB equation provides a restriction linking these objects. It is, for $i \in \{e, h\}$,

$$\begin{aligned} 0 = \frac{\rho_i}{1-\psi}\left(\left(\frac{c_i}{n_i}\right)^{1-\psi}\xi_i^{\psi-1} - 1\right) + r - \frac{c_i}{n_i} + \frac{\mu_{R,i} - r}{|\sigma_R|^2}\sigma_{n,i} \cdot \sigma_R - \frac{\gamma}{2}|\sigma_{n,i}|^2 - (\gamma - 1)\sigma_{n,i} \cdot \sigma_{\xi,i} \\ + \mu_{\xi,i} - \frac{\gamma}{2}|\sigma_{\xi,i}|^2, \end{aligned} \tag{F.37}$$

where c_i/n_i denotes the consumption-wealth ratio, $\sigma_{n,i} := \frac{qk_i}{n_i}\sigma_R$ is the diffusion of the net worth growth rate dn_i/n_i, and $\mu_{R,i}$ is the agent-specific expected return on capital holdings. Maximizing over consumption c_i, we obtain

$$\frac{c_i}{n_i} = \rho_i^{\psi^{-1}} \xi_i^{1-\psi^{-1}} \tag{F.38}$$

Maximizing over $\sigma_{n,i} = \frac{qk_i}{n_i}\sigma_R$, subject to the no-shorting constraint $k_i \geq 0$, we obtain

$$\frac{qk_i}{n_i} \geq \frac{\mu_{R,i} - r}{\gamma|\sigma_R|^2} + \frac{1-\gamma}{\gamma}\frac{\sigma_R \cdot \sigma_{\xi,i}}{|\sigma_R|^2}, \tag{F.39}$$

with equality when $k_i > 0$. Plugging these back into the HJB equation delivers:

$$0 = \frac{\psi}{1-\psi}\rho_i^{\psi^{-1}}\xi_i^{1-\psi^{-1}} - \frac{\rho}{1-\psi} + r + \left(\gamma\sigma_{n,i} + (\gamma-1)\frac{\sigma_R\cdot\sigma_{\xi,i}}{|\sigma_R|^2}\sigma_R\right)\cdot\sigma_{n,i} - \frac{\gamma}{2}|\sigma_{n,i}|^2 - (\gamma-1)\sigma_{n,i}\cdot\sigma_{\xi,i} + \mu_{\xi,i} - \frac{\gamma}{2}|\sigma_{\xi,i}|^2. \tag{F.40}$$

One should think of (F.40) as an equation for ξ_i.

Now, we obtain the key aggregate conditions analogous to our baseline model. First, using equation (F.38) in the goods market clearing condition, we obtain the price-output relationship

$$q\left[\eta\rho_e^{\psi^{-1}}\xi_e^{1-\psi^{-1}} + (1-\eta)\rho_h^{\psi^{-1}}\xi_h^{1-\psi^{-1}}\right] = \kappa a_e + (1-\kappa)a_h \tag{F.41}$$

Second, using equation (F.39) for $i \in \{e,h\}$ delivers a risk-balance condition and an equation for the riskless rate:

$$0 = \min\left\{1-\kappa, \frac{a_e - a_h}{q} - \gamma\frac{\kappa-\eta}{\eta(1-\eta)}|\sigma_R|^2 - (\gamma-1)(\sigma_{\xi,e}-\sigma_{\xi,h})\cdot\sigma_R\right\} \tag{F.42}$$

and

$$r = \frac{\kappa a_e + (1-\kappa)a_h}{q} + g + \mu_q + \sigma\sigma_q\cdot\binom{1}{0} - \gamma\left(\frac{\kappa^2}{\eta} + \frac{(1-\kappa)^2}{1-\eta}\right)|\sigma_R|^2 - (\gamma-1)\left(\kappa\sigma_{\xi,e} + (1-\kappa)\sigma_{\xi,h}\right)\cdot\sigma_R \tag{F.43}$$

In the log case, these were the 3 equations that, together with global stability, characterize equilibrium. Now, we have additional unknowns related to agents' value function $(\xi_i, \mu_{\xi,i}, \sigma_{\xi,i})_{i\in\{e,h\}}$ that are coupled together with this system.

Let us illustrate a particular construction where the aggregate states are (η, q), which are the fewest possible for a sentiment-driven equilibrium (e.g., Theorem 1 allows for additional exogenous states y that we do not use here). In this case, we have the following Itô conditions for ξ_i, i.e.,

$$\xi_i\sigma_{\xi,i} = \xi_{i,q}(q\sigma_q) + \xi_{i,\eta}\sigma_\eta \tag{F.44}$$

$$\xi_i\mu_{\xi,i} = \xi_{i,q}(q\mu_q) + \xi_{i,\eta}\mu_\eta + \frac{1}{2}(\xi_{i,qq}|q\sigma_q|^2 + \xi_{i,\eta\eta}|\sigma_\eta|^2) + \xi_{i,q\eta}q\sigma_q\cdot\sigma_\eta \tag{F.45}$$

where $\xi_{i,x}$, $\xi_{i,xy}$ denote the partial derivatives of ξ_i with respect to a variable x, or x and y, respectively. Thus, given the choice of aggregate states, we pin down the dynamics

of $\xi_{i,e}$ and $\xi_{i,h}$. Substituting these conditions back into (F.40) yields a pair of partial differential equations (PDEs) for $\xi_{i,e}$ and $\xi_{i,h}$. The boundary conditions for these PDEs depend on the boundary conditions induced by the state dynamics $(\mu_q, \sigma_q, \mu_\eta, \sigma_\eta)$. We return to these PDEs below.

It is straightforward to derive equations for the dynamics of the wealth share η, using its definitions and the Euler equations (F.39) above:

$$\begin{aligned}\mu_\eta = {} & \eta(1-\eta)(\rho_e - \rho_h) + \eta(1-\eta)\Big(\gamma(|\sigma_{n,e}|^2 - |\sigma_{n,h}|^2) + (\gamma-1)(\sigma_{n,e}\cdot\sigma_{\xi,e} - \sigma_{n,h}\cdot\sigma_{\xi,h})\Big) \\ & - \eta(1-\eta)\Big(\eta|\sigma_{n,e}|^2 - (1-\eta)|\sigma_{n,h}|^2 + (1-2\eta)\sigma_{n,e}\cdot\sigma_{n,h}\Big) + \delta_h - (\delta_e + \delta_h)\eta \qquad \text{(F.46)}\end{aligned}$$

$$\sigma_\eta = (\kappa - \eta)\sigma_R \qquad \text{(F.47)}$$

This pins down (μ_η, σ_η) in terms of the other objects.

Hence, assuming transversality conditions hold, we can characterize a sentiment equilibrium recursive in (η, q) as functions for prices (μ_q, σ_q, r), the capital allocation κ, wealth share dynamics (σ_η, μ_η) and investment opportunities $(\xi_i, \mu_{\xi,i}, \sigma_{\xi,i})_{i\in\{e,h\}}$ such that (F.41)-(F.47) hold, and $\eta_t \in [0,1]$ and $\kappa_t \in [0,1]$ for all $t \geq 0$.

Counting carefully, the system that characterizes the recursive equilibrium in (η, q) has 2 degrees of freedom. Why? As in the case with logarithmic utility, we do not have enough equations to discipline the dynamics of asset prices (μ_q, σ_q), i.e., there are 3 degrees of freedom. We use one degree of freedom, by specifying that q be a state. That leaves indeterminate the fraction of return variance $|\sigma_R|$ due to the sunspot versus fundamental shocks, as well as the price drift μ_q. Although the setting with recursive preferences adds the HJB equation (F.40) and the Itô conditions (F.44)-(F.45) for the value function dynamics, these 8 conditions are only enough to pin down the 8 unknowns $(\xi_i, \mu_{\xi,i}, \sigma_{\xi,i})_{i\in\{e,h\}}$. In other words, the remaining 2 degrees of freedom in the q dynamics are not eliminated. Therefore, we expect that a sentiment equilibrium should be possible to construct, though highly non-trivial.

G Discrete-time model

The following discrete-time model is exactly analogous to our continuous-time model. We model each decision on a time-step of Δ (it will turn out that the decision interval Δ cannot be arbitrarily large).

Technology. For simplicity, we assume that aggregate capital K is fixed, i.e., there is no fundamental uncertainty. Note nevertheless that individual positions on capital are not predetermined since agents can trade capital.

Individual agent problem. An individual can hold two assets, riskless bonds b_t and capital k_t, and decides consumption c_t. The individual net worth, just before consuming, is $n_t = b_t + q_t k_t$, where q_t is the market price of capital. The one-period return on bonds is $R_t^f = 1 + r_t\Delta$, and the return-on-capital is $R_{t+\Delta}^k := \frac{a\Delta}{q_t} + \frac{q_{t+\Delta}}{q_t}$, where a is the agent's productivity per unit of time while holding capital. Then, the agent's dynamic budget constraint is[36]

$$n_{t+\Delta} = q_t k_t (R_{t+\Delta}^k - R_t^f) + (n_t - c_t) R_t^f. \tag{G.1}$$

Each agent takes q_t, R_t^f, and $R_{t+\Delta}^k$ as given and chooses (c, k, n) to maximize

$$\mathbb{E}\left[\sum_{i=0}^{\infty} \left(\frac{1}{1+\rho\Delta}\right)^i \log(c_{i\Delta})\right], \tag{G.2}$$

subject to (G.1), subject to the no-shorting constraint $k_t \geq 0$, and subject to the solvency constraint $n_t \geq 0$.

The first-order optimality conditions are the standard Euler equations

$$1 = \frac{1}{1+\rho\Delta} R_t^f \mathbb{E}_t\left[\frac{c_t}{c_{t+\Delta}}\right] \tag{G.3}$$

$$0 \geq \frac{1}{1+\rho\Delta} \mathbb{E}_t\left[\frac{c_t}{c_{t+\Delta}} (R_{t+\Delta}^k - R_t^f)\right], \tag{G.4}$$

where (G.4) holds with equality when $k_t > 0$ is chosen.

[36] To derive (G.1), proceed as follows. First, note that the bond market account next period, before adjusting the portfolio of bonds and capital, will have value $b'_{t+\Delta} = R_f^f(b_t - c_t) + ak_t\Delta$—that is, after consumption expenditures are made, the residual earns the interest rate, and the cash flows from holding capital are also added at the end of the period. Second, the capital holdings k_t will have value $q_{t+\Delta}k_t$ next period. Adding these two quantities must equal tomorrow's net worth $n_{t+\Delta}$. Hence, $n_{t+\Delta} = R_f^f(b_t - c_t) + ak_t\Delta + q_{t+\Delta}k_t$. Using the definition $n_t = b_t + q_t k_t$ gives the result (G.1).

In addition, it is straightforward to show that optimal consumption satisfies the standard log utility formula[37]

$$c_t = \frac{\rho\Delta}{1+\rho\Delta} n_t. \tag{G.5}$$

Using this fact, plus the budget constraint (G.1) in (G.3)-(G.4), we obtain

$$1 = \frac{1}{1+\rho\Delta} R_t^f \mathbb{E}_t \left[\frac{1}{\theta_t (R_{t+\Delta}^k - R_t^f) + (1+\rho\Delta)^{-1} R_t^f} \right] \tag{G.6}$$

$$0 \geq \frac{1}{1+\rho\Delta} \mathbb{E}_t \left[\frac{R_{t+\Delta}^k - R_t^f}{\theta_t (R_{t+\Delta}^k - R_t^f) + (1+\rho\Delta)^{-1} R_t^f} \right], \quad \text{with equality if } \theta_t > 0 \tag{G.7}$$

where $\theta_t := \frac{q_t k_t}{n_t}$ is the share of wealth allocated to capital. At this point, one can prove that (G.6) holds automatically if (G.7) holds.[38] Therefore, we can drop the bond Euler equation (G.6) from the remainder of the analysis, i.e., (G.5) and (G.7) fully characterize the agent's optimal choices.

Aggregation and equilibrium conditions. As in the main text, we assume there are two types of agents: experts have productivity a_e and discount rate ρ_e, while households have productivity $a_h < a_e$ and discount rate $\rho_h \leq \rho_e$. Clearly, then, experts have a higher return-on-capital than households: $R_{e,t+\Delta}^k > R_{h,t+\Delta}^k$.

We now aggregate. The market clearing condition for goods, capital, and bonds are given by, respectively,

$$c_{e,t} + c_{h,t} = (a_e k_{e,t} + a_h k_{h,t})\Delta \tag{G.8}$$

$$k_{e,t} + k_{h,t} = K \tag{G.9}$$

$$b_{e,t} + b_{h,t} = c_{e,t} + c_{h,t}. \tag{G.10}$$

Equation (G.10) says that bondholdings just after consuming (which is $b_t - c_t$) sum to

[37] This can be showed by writing out the Bellman equation and guessing-and-verifying that the value function takes the form $v_t = (1-\beta)^{-1}\log(n_t) + f(\Omega_t)$ for $\beta = (1+\rho\Delta)^{-1}$ and some function f that only depends on aggregate states Ω_t. Then, the envelope condition says $c_t^{-1} = \frac{\partial}{\partial n} v_t = (1-\beta)^{-1} n_t^{-1}$, which is the consumption formula.

[38] Indeed, if $\theta_t = 0$ it is obvious that (G.6) holds. If $\theta_t > 0$, then (G.7) holds with equality, so we then have

$$0 = \mathbb{E}_t \Big[\frac{\theta_t (R_{t+\Delta}^k - R_t^f)}{\theta_t (R_{t+\Delta}^k - R_t^f) + (1+\rho\Delta)^{-1} R_t^f} \Big]$$

Adding this expression to equation (G.6), we obtain the identity $1 = 1$.

the zero net supply. By combining (G.10) with the individual net worth definition $n_t = b_t + q_t k_t$, we obtain an alternative statement of bond market clearing that we will use:

$$n_{e,t} + n_{h,t} = q_t K + c_{e,t} + c_{h,t}. \tag{G.11}$$

Definition 5. An *equilibrium* is a collection of stochastic processes for allocations $(k_{j,t\Delta}, n_{j,t\Delta}, c_{j,t\Delta})_{t=0}^{\infty}$ for $j \in \{e, h\}$ with $k_{e,0}$ and $k_{h,0}$ given, and for prices $(q_{t\Delta}, R^f_{t\Delta})_{t=0}^{\infty}$ such that (i) given prices, allocations solve each agent type's problem, and (ii) markets clear.

G.1 Equilibrium characterization

We have already characterized optimal decisions and market clearing conditions. In particular, a collection of stochastic processes for allocations and prices constitute an equilibrium if they satisfy (G.1), (G.5), and (G.7) for each agent type (experts and households), along with equations (G.8), (G.9), and (G.11) at the aggregate level.

We further tighten this characterization and reduce it to four stochastic processes satisfying a set of conditions, exactly as in our continuous-time model. First, to keep track of the distribution of wealth and capital, let $\eta_t := (1 + \rho_e \Delta)^{-1} n_{e,t} / q_t K$ and $\kappa_t := k_{e,t}/K$ denote expert's wealth and capital shares.[39] Whereas κ_t is a "jumpy" variable because it is linked to agent's capital choices, η_t is a "state" variable because it is determined via agent's slow-moving wealths. Using the budget constraint (G.1), we can obtain the dynamics of η_t as

$$\eta_{t+\Delta} = \frac{1}{1 + \rho_e \Delta} \left(\frac{\kappa_t (R^k_{e,t+\Delta} - R^f_t) + \eta_t R^f_t}{q_{t+\Delta}/q_t} \right). \tag{G.12}$$

Next, we aggregate the consumption decisions across these two types. To do this, plug the consumption rules from (G.5) into the goods and bond market clearing conditions (G.8) and (G.11), and combine the results to obtain

$$q_t \bar{\rho}(\eta_t) = \kappa_t a_e + (1 - \kappa_t) a_h, \tag{G.13}$$

where $\bar{\rho}(\eta) := \eta \rho_e + (1 - \eta)\rho_h$ is a wealth-weighted average discount rate. Identical to our continuous-time model, equation (G.13) is a *price-output relation* that links asset values q_t to the efficiency of the capital distribution κ_t. Finally, we aggregate the Euler

[39] Note that the wealth share is defined just after consumption choices are made, i.e., $\eta_t = (n_{e,t} - c_{e,t})/(n_{e,t} + n_{h,t} - c_{e,t} - c_{h,t})$ is the definition we are using.

equations (G.7) within the two types using the fact that experts will always be on the margin (i.e., since $R^k_{e,t+\Delta} > R^k_{h,t+\Delta}$, we have $k_{e,t} > 0$ at all times). We also use the fact that $\theta_{e,t} = \frac{q_t k_{e,t}}{n_{e,t}} = \frac{1}{1+\rho_e \Delta}\frac{\kappa_t}{\eta_t}$ and $\theta_{h,t} = \frac{q_t k_{h,t}}{n_{h,t}} = \frac{1}{1+\rho_h \Delta}\frac{1-\kappa_t}{1-\eta_t}$ to write the results in a more convenient way. The results are

$$0 = \mathbb{E}_t\left[\frac{q_{t+\Delta} + a_e\Delta - R^f_t q_t}{\frac{\kappa_t}{\eta_t}\left(q_{t+\Delta} + a_e\Delta - R^f_t q_t\right) + R^f_t q_t}\right] \tag{G.14}$$

$$0 \geq \mathbb{E}_t\left[\frac{q_{t+\Delta} + a_h\Delta - R^f_t q_t}{\frac{1-\kappa_t}{1-\eta_t}\left(q_{t+\Delta} + a_h\Delta - R^f_t q_t\right) + R^f_t q_t}\right] \tag{G.15}$$

where the latter holds as an equality when households hold capital, i.e., when $\kappa_t < 1$.

Thus, an equilibrium is fully characterized by the collection of stochastic processes $(\eta_{t\Delta}, \kappa_{t\Delta}, q_{t\Delta}, R^f_{t\Delta})_{t=0}^{\infty}$, with $\eta_0 = k_{e,0}/K$ given, such that the two optimality conditions (G.14)-(G.15) hold; the price-output relation (G.13) holds; and the law of motion for η_t is given by (G.12). To establish the analog to our continuous-time model, we also state this characterization as a lemma—notice that the verbiage is almost identical to Lemma 1.

Lemma G.1. *Given $\eta_0 \in (0,1)$, consider stochastic processes $\{\eta_{t\Delta}, q_{t\Delta}, \kappa_{t\Delta}, R^f_{t\Delta}\}_{t=0}^{\infty}$ such that η_t evolution is described by (G.12). If $\eta_t \in [0,1]$, $\kappa_t \in [0,1]$, and equations (G.13), (G.14), and (G.15) hold for all $t \geq 0$, then $\{\eta_{t\Delta}, q_{t\Delta}, \kappa_{t\Delta}, R^f_{t\Delta}\}_{t=0}^{\infty}$ corresponds to an equilibrium.*

Notice from Lemma G.1 that we have as many equations as unknown non-state variables (q_t, κ_t, R^f_t). However, Euler equations (G.14)-(G.15) also depend on the probability distribution of the future asset price $q_{t+\Delta}$, in order to determine the asset price q_t and riskless rate R^f_t today. This will be the key reason why the set of equilibrium conditions above is not enough to pin down q_t uniquely. In the continuous-time model, the distribution of future asset prices was summarized by the drift and the volatility (μ_q, σ_q). Here, the distribution of $q_{t+\Delta}$ could be more general, but we present a binomial example below. We now proceed to analysis of the two types of equilibria: fundamental and non-fundamental.

G.2 Fundamental equilibrium

A *fundamental equilibrium* has $\kappa_t = 1$ for all periods. In such an equilibrium, (G.13) says that the capital price should be

$$q_t = \frac{a_e}{\bar{\rho}(\eta_t)}, \quad \text{if} \quad \kappa_t = 1. \tag{G.16}$$

Substituting this result into the state dynamics (G.12), we have

$$\eta_{t+\Delta} = \frac{1}{1+\rho_e\Delta}\left[1+\bar{\rho}(\eta_{t+\Delta}) - \frac{\bar{\rho}(\eta_{t+\Delta})}{\bar{\rho}(\eta_t)}(1-\eta_t)R_t^f\right], \quad \text{if} \quad \kappa_t = \kappa_{t+\Delta} = 1. \tag{G.17}$$

As the only $(t+\Delta)$-measurable object in (G.17), $\eta_{t+\Delta}$ evolves deterministically in a fundamental equilibrium. Because q_t is solely a function of η_t in (G.16), $q_{t+\Delta}$ is also known as of time t. As a result, experts' return-on-capital must coincide with the riskless rate, i.e., $R_t^f = \frac{a_e\Delta}{q_t} + \frac{q_{t+\Delta}}{q_t}$, or

$$R_t^f = \bar{\rho}(\eta_t) + \frac{\bar{\rho}(\eta_t)}{\bar{\rho}(\eta_{t+\Delta})}, \quad \text{if} \quad \kappa_t = \kappa_{t+\Delta} = 1. \tag{G.18}$$

Combining (G.17) and (G.18), we obtain the solved dynamics

$$\eta_{t+\Delta} = \frac{\eta_t(1+\rho_e\Delta)^{-1}}{\eta_t(1+\rho_e\Delta)^{-1} + (1-\eta_t)(1+\rho_h\Delta)^{-1}}, \quad \text{if} \quad \kappa_t = \kappa_{t+\Delta} = 1. \tag{G.19}$$

Thus, expert's wealth share asymptotically tends toward zero. Intuitively, they earn zero excess capital returns and consume at a higher rate than households.

G.3 Non-fundamental equilibrium

A *non-fundamental equilibrium* has $\kappa_t < 1$ for some t. We proceed with a simple binomial tree example to show that non-fundamental equilibria exist, although more complicated information structures are also likely possible. We conjecture an equilibrium with

$$q_{t+\Delta} = \begin{cases} u_t q_t, & \text{with probability } 1-\pi_t; \\ d_t q_t, & \text{with probability } \pi_t. \end{cases} \tag{G.20}$$

The "up" and "down" returns u_t and $d_t \in (0, u_t)$ may be state dependent, as may the probability of a price drop π_t. As in our baseline model, we will take the "state space" to be the set of possible (η_t, q_t), or equivalently (η_t, κ_t). In other words, (u_t, d_t, π_t) will be functions of (η_t, κ_t), as will the interest rate r_t. The rest of this appendix constructs an example equilibrium under the binomial scheme (G.20). In particular, we will prove the following by construction:

Proposition G.1. *For all Δ sufficiently small, a non-fundamental equilibrium exists.*

To start, we may solve for the optimal portfolios explicitly in this binomial environ-

ment. Using (G.12) and (G.20) in the expert Euler equation (G.14), we have

$$\frac{\kappa_t}{\eta_t} = -R_t^f \frac{(1-\pi_t)u_t + \pi_t d_t + \frac{a_e \Delta}{q_t} - R_t^f}{(u_t + \frac{a_e \Delta}{q_t} - R_t^f)(d_t + \frac{a_e \Delta}{q_t} - R_t^f)}. \tag{G.21}$$

Doing the same for the household Euler equation (G.15), we have

$$\frac{1-\kappa_t}{1-\eta_t} = -R_t^f \min\Big(0, \frac{(1-\pi_t)u_t + \pi_t d_t + \frac{a_h \Delta}{q_t} - R_t^f}{(u_t + \frac{a_h \Delta}{q_t} - R_t^f)(d_t + \frac{a_h \Delta}{q_t} - R_t^f)}\Big). \tag{G.22}$$

Next, note that the price-output relation (G.13) and state dynamics (G.12) are unchanged by the binomial setup, and we repeat them here for convenience:

$$\bar{\rho}(\eta_t) = \frac{\kappa_t a_e + (1-\kappa_t) a_h}{q_t} \tag{G.23}$$

$$\eta_{t+\Delta} = \frac{1}{1+\rho_e \Delta} \frac{\kappa_t (\frac{a_e \Delta}{q_t} + \frac{q_{t+\Delta}}{q_t} - R_t^f) + \eta_t R_t^f}{q_{t+\Delta}/q_t}. \tag{G.24}$$

As mentioned in Lemma G.1, to find an equilibrium we only need to check that we can pick (u_t, d_t, π_t) to satisfy (G.21)-(G.24) at every point in the state space and that the resulting equilibrium dynamics do not cause the dynamical system to "exit the feasible region." To this end, we immediately note that $\eta_t \in (0,1)$ on any equilibrium path, which can be verified by checking the state dynamics (G.24).[40]

To continue, we will specialize below to a particular choice of u and d. Our construction will correspond to an approximation of Brownian motion in the "interior" of the

[40] Examine the state dynamics (G.24) in the down state and substitute (G.21) to obtain

$$d_t \frac{\eta_{t+\Delta}^d}{\eta_t} = \frac{1}{1+\rho_e \Delta} R_t^f \Big(1 - \frac{(1-\pi_t)u_t + \pi_t d_t + \frac{a_e \Delta}{q_t} - R_t^f}{u_t + \frac{a_e \Delta}{q_t} - R_t^f}\Big) > 0.$$

Similarly, mirroring (G.24), the symmetric condition for household's net worth share dynamics is

$$1 - \eta_{t+\Delta} = \frac{1}{1+\rho_e \Delta} \frac{(1-\kappa_t)(\frac{a_h \Delta}{q_t} + \frac{q_{t+\Delta}}{q_t} - R_t^f) + (1-\eta_t) R_t^f}{q_{t+\Delta}/q_t}$$

Examining this condition in the up state and substituting (G.22), we obtain

$$u_t \frac{1-\eta_{t+\Delta}^u}{1-\eta_t} = \frac{1}{1+\rho_h \Delta} R_t^f \Big(1 - \min\Big(0, \frac{(1-\pi_t)u_t + \pi_t d_t + \frac{a_h \Delta}{q_t} - R_t^f}{d_t + \frac{a_h \Delta}{q_t} - R_t^f}\Big)\Big) > 0.$$

Thus, the requirement to keep $\eta_t \in (0,1)$ is automatically satisfied.

state space, with special considerations imposed at the "boundaries" of this state space. More specifically, we define the following regions. First, we have the entire feasible state space

$$\mathcal{D} := \Big\{(\eta, \kappa) : \eta \in (0,1),\, \kappa \in (\eta, 1]\Big\}.$$

The reason why $\kappa > \eta$ is required is because $\kappa \leq \eta$ is inconsistent with the expert and household Euler equations (G.21)-(G.22), since $a_e > a_h$. Next, there will be a region near the top of $\mathcal{D}$, where κ is close to 1, such that positive shocks will just take the economy to the border:

$$\mathcal{D}_{high} := \Big\{(\eta, \kappa) \in \mathcal{D} : \kappa < 1,\, f(\kappa, \eta) < 0\Big\}.$$

for some function f to be defined endogenously below. At the other ends, let us pick some $\epsilon > 0$ and define the lower boundary region:

$$\mathcal{D}^{\epsilon}_{low} := \Big\{(\eta, \kappa) \in \mathcal{D} \backslash \mathcal{D}_{high} : \kappa \leq (1+\epsilon)\eta\Big\}.$$

For reasons that will become clear at the end of the construction, we will impose

$$\epsilon > \frac{a_h \rho_e}{(a_e - a_h)\rho_h}. \tag{G.25}$$

Finally, we will detail a separate method to deal with the top border region

$$\mathcal{D}_1 := \Big\{(\eta, \kappa) \in \mathcal{D} : \kappa = 1\Big\}.$$

The "interior" region is defined by subtracting these boundary regions:

$$\mathcal{D}^{\circ} := \mathcal{D} \backslash (\mathcal{D}_{high} \cup \mathcal{D}^{\epsilon}_{low} \cup \mathcal{D}_1).$$

We explain our construction in each of these regions in sequence.

Brownian approximation in the interior. In the interior region $\mathcal{D}^{\circ}$, we construct a non-fundamental equilibrium by explicitly specifying (u_t, d_t, π_t) to take a form that approximates Brownian motion in the $\Delta \to 0$ limit. In particular, we set

$$u_t = 1 + v_t \sqrt{\Delta} \tag{G.26}$$

$$d_t = 1 - v_t \sqrt{\Delta} \tag{G.27}$$

$$\pi_t = \frac{v_t - m_t \sqrt{\Delta}}{2v_t}, \tag{G.28}$$

for some endogenous variables m_t and v_t. Note that $\pi_t \in (0,1)$ requires $m_t\sqrt{\Delta} \in (-v_t, v_t)$. Of course, we also require $v_t \leq 1/\sqrt{\Delta}$. These constraints on m_t and v_t become arbitrarily loose as $\Delta \to 0$.

One can verify that (G.26)-(G.28) imply that

$$\mathbb{E}_t[\frac{q_{t+\Delta} - q_t}{q_t}] = m_t\Delta.$$

Thus, the interpretation of the variable m_t introduced is as the drift of percentage price changes. Also, we may compute

$$\mathbb{E}_t[(\frac{q_{t+\Delta} - q_t}{q_t})^2] = v_t^2\Delta,$$

so that v_t corresponds roughly to the instantaneous volatility of percentage price changes. Notice that any higher moments of price changes are of order $o(\Delta)$. Similarly, substituting the specification (G.26)-(G.28) into (G.24), one can verify that the state dynamics converge as $\Delta \to 0$ to the continuous-time model. Indeed, examine the conditional mean and second moment of $\eta_{t+\Delta} - \eta_t$:

$$\begin{aligned}\mathbb{E}_t[\eta_{t+\Delta} - \eta_t] &= \Big(\kappa_t\frac{a_e}{q_t} - \eta_t\rho_e + (\kappa_t - \eta_t)(m_t - r_t - v_t^2)\Big)\Delta + o(\Delta)\\ \mathbb{E}_t[(\eta_{t+\Delta} - \eta_t)^2] &= (\kappa_t - \eta_t)^2 v_t^2\Delta + o(\Delta).\end{aligned}$$

Dividing by Δ and taking $\Delta \to 0$, it becomes clear that these moments coincide with those of the continuous-time model.

Now, we determine what m_t and v_t must be to satisfy agents' optimality conditions. In this Brownian approximation, the expert and household Euler equations (G.21)-(G.22) become

$$\frac{\kappa_t}{\eta_t} = (1 + r_t\Delta)\frac{\frac{a_e}{q_t} + m_t - r_t}{v_t^2 - (\frac{a_e}{q_t} - r_t)^2\Delta} \tag{G.29}$$

$$\frac{1-\kappa_t}{1-\eta_t} = (1 + r_t\Delta)\max\Big\{0, \frac{\frac{a_h}{q_t} + m_t - r_t}{v_t^2 - (\frac{a_h}{q_t} - r_t)^2\Delta}\Big\}. \tag{G.30}$$

As $\Delta \to 0$, these two specialized Euler equations (G.29)-(G.30) coincide with the familiar mean-variance portfolio choice. However, to recover the same equations as in our

continuous-time model, let us take the difference between (G.29)-(G.30) to get

$$0 = \min\left\{1-\kappa_t,\ (1+r_t\Delta)\Big[\frac{\frac{a_e}{q_t}+m_t-r_t}{v_t^2-(\frac{a_e}{q_t}-r_t)^2\Delta}-\frac{\frac{a_h}{q_t}+m_t-r_t}{v_t^2-(\frac{a_h}{q_t}-r_t)^2\Delta}\Big]-\frac{\kappa_t-\eta_t}{\eta_t(1-\eta_t)}\right\}. \quad \text{(G.31)}$$

Equation (G.31) clearly coincides with our baseline risk-balance condition as $\Delta \to 0$. Then, summing (G.29)-(G.30), weighted by κ_t and $1-\kappa_t$ respectively, we have

$$\frac{\kappa_t^2}{\eta_t}+\frac{(1-\kappa_t)^2}{1-\eta_t} = (1+r_t\Delta)\Big[\kappa_t\frac{\frac{a_e}{q_t}+m_t-r_t}{v_t^2-(\frac{a_e}{q_t}-r_t)^2\Delta}+(1-\kappa_t)\frac{\frac{a_h}{q_t}+m_t-r_t}{v_t^2-(\frac{a_h}{q_t}-r_t)^2\Delta}\Big]. \quad \text{(G.32)}$$

Again, this coincides with the equation for μ_q in the continuous-time model as $\Delta \to 0$.

To solve the model, first we use the expert Euler equation to solve for v_t^2:

$$v_t^2 = (1+r_t\Delta)\Big[\frac{a_e}{q_t}+m_t-r_t\Big]\frac{\eta_t}{\kappa_t}+(\frac{a_e}{q_t}-r_t)^2\Delta.$$

Then, we use the household Euler equation, when $\kappa_t < 1$, to also solve for v_t^2:

$$v_t^2 = (1+r_t\Delta)\Big[\frac{a_h}{q_t}+m_t-r_t\Big]\frac{1-\eta_t}{1-\kappa_t}+(\frac{a_h}{q_t}-r_t)^2\Delta.$$

Setting these expressions equal gives an equation for m_t, which is

$$m_t = r_t+\frac{(1-\kappa_t)\eta_t}{\kappa_t-\eta_t}\frac{a_e}{q_t}-\frac{\kappa_t(1-\eta_t)}{\kappa_t-\eta_t}\frac{a_h}{q_t}+\frac{\kappa_t(1-\kappa_t)\big[(\frac{a_e}{q_t}-r_t)^2-(\frac{a_h}{q_t}-r_t)^2\big]}{(1+r_t\Delta)(\kappa_t-\eta_t)}\Delta. \quad \text{(G.33)}$$

Substituting back into the equations for v_t^2, we solve for

$$v_t^2 = (1+r_t\Delta)\frac{\eta_t(1-\eta_t)}{\kappa_t-\eta_t}\frac{a_e-a_h}{q_t}+\frac{\kappa_t(1-\eta_t)(\frac{a_e}{q_t}-r_t)^2-\eta_t(1-\kappa_t)(\frac{a_h}{q_t}-r_t)^2}{\kappa_t-\eta_t}\Delta. \quad \text{(G.34)}$$

Given a choice for r_t, we can obtain m_t and v_t^2 from equations (G.33)-(G.34), for any point in the interior of the state space. The only restriction is that we choose r_t so that $m_t\sqrt{\Delta} \in (-v_t, v_t)$ and hence that $\pi_t \in (0,1)$, which leaves a wide range of choices. To be explicit, we will choose r_t such that $m_t = O(\Delta)$, in particular we set

$$r_t = \frac{\kappa_t(1-\eta_t)}{\kappa_t-\eta_t}\frac{a_h}{q_t}-\frac{(1-\kappa_t)\eta_t}{\kappa_t-\eta_t}\frac{a_e}{q_t}. \quad \text{(G.35)}$$

This choice makes it automatic that $m_t\sqrt{\Delta} \in (-v_t, v_t)$ if Δ is also chosen small enough.

As an aside, note that these equations, in the $\Delta \to 0$ limit, are identical to the continuous-time versions (when there is zero fundamental risk and zero growth). Indeed, equation (G.34) says

$$v_t^2 = \frac{\eta_t(1-\eta_t)}{\kappa_t - \eta_t}\frac{a_e - a_h}{q_t} + O(\Delta).$$

Next, by doing some algebra on (G.33), it reads

$$m_t = r_t - \bar{\rho}(\eta_t) + \Big(\frac{\kappa_t^2}{\eta_t} + \frac{(1-\kappa_t)^2}{1-\eta_t}\Big)v_t^2 + O(\Delta).$$

Consequently, m_t and v_t are indeed the discrete-time counterparts to $\mu_{q,t}$ and $\sigma_{q,t}$.

Reflection approximation near the lower boundary. In the lower region $\mathcal{D}_{low}^{\epsilon}$, we proceed with a different construction that ensures the economy never exits $\mathcal{D}$ through its lower border. Luckily, in everything so far, r_t was indeterminate, and this flexibility is what allows us to construct such an equilibrium. In particular, to ensure we always have $\kappa_t \in (\eta_t, 1)$, we impose some rules similar to our "boundary conditions" in continuous time.

In $\mathcal{D}_{low}^{\epsilon}$, we will use the binomial specification

$$u_t = 1 + v_t^2/m_t \tag{G.36}$$

$$d_t = 1 \tag{G.37}$$

$$\pi_t = \frac{v_t^2 - m_t^2\Delta}{v_t^2} \tag{G.38}$$

Equations (G.36)-(G.38) preserve the desired moment properties that $\mathbb{E}_t[\frac{q_{t+\Delta}-q_t}{q_t}] = m_t\Delta$ and $\mathbb{E}_t[(\frac{q_{t+\Delta}-q_t}{q_t})^2] = v_t^2\Delta$. Again, we must have probabilities in between zero and one, so we always require $m_t\sqrt{\Delta} \in (-v_t, v_t)$.

With this specification, the Euler equations become

$$\frac{\kappa_t}{\eta_t} = (1+r_t\Delta)\frac{\frac{a_e}{q_t} + m_t - r_t}{\frac{v_t^2}{m_t}(r_t - \frac{a_e}{q_t}) - (\frac{a_e}{q_t} - r_t)^2\Delta} \tag{G.39}$$

$$\frac{1-\kappa_t}{1-\eta_t} = (1+r_t\Delta)\frac{\frac{a_h}{q_t} + m_t - r_t}{\frac{v_t^2}{m_t}(r_t - \frac{a_h}{q_t}) - (\frac{a_h}{q_t} - r_t)^2\Delta}. \tag{G.40}$$

As before, we may use these two equations to solve for m_t and v_t^2:

$$m_t = r_t + \frac{(1+r_t\Delta)\left[\frac{\eta_t}{\kappa_t}(r_t - \frac{a_h}{q_t})\frac{a_e}{q_t} - \frac{1-\eta_t}{1-\kappa_t}(r_t - \frac{a_e}{q_t})\frac{a_h}{q_t}\right] - (\frac{a_e - a_h}{q_t})(r_t - \frac{a_e}{q_t})(r_t - \frac{a_h}{q_t})\Delta}{(1+r_t\Delta)\left[\frac{1-\eta_t}{1-\kappa_t}(r_t - \frac{a_e}{q_t}) - \frac{\eta_t}{\kappa_t}(r_t - \frac{a_h}{q_t})\right]} \tag{G.41}$$

$$v_t^2 = m_t\Big[\frac{(1+r_t\Delta)\frac{\eta_t}{\kappa_t}(\frac{a_e}{q_t} + m_t - r_t)}{r_t - \frac{a_e}{q_t}} + (r_t - \frac{a_e}{q_t})\Delta\Big] \tag{G.42}$$

Given that the Euler equations hold for this choice of (m_t, v_t^2), we have an equilibrium as long as $m_t\sqrt{\Delta} \in (-v_t, v_t)$ and $\kappa_t > \eta_t$ in all periods.

The condition that $\kappa_t > \eta_t$ is the more complex and restrictive condition. The key issue is that (η_t, κ_t) can jump from $\mathcal{D}^{\epsilon}_{low}$ to a point outside of the feasible region $\mathcal{D}$.[41] Resolving this issue requires us to make particular choices for r_t such that the dynamics of (η_t, κ_t) "point toward the interior" of the state space, i.e., the dynamics starting from $\mathcal{D}^{\epsilon}_{low}$ are such that $(\eta_{t+\Delta}, \kappa_{t+\Delta})$ moves closer to $\mathcal{D}^{\circ}$. Sufficient conditions for this are that $\eta_{t+\Delta} \leq \eta_t$ when $(\eta_t, \kappa_t) \in \mathcal{D}^{\epsilon}_{low}$. Indeed, if $\eta_{t+\Delta} \leq \eta_t$, then the dynamics of q_t are such that $\kappa_{t+\Delta} \geq \kappa_t$. Since the lower-boundary of $\mathcal{D}$ is upward-sloping in (η, κ)-space, the combination of $\eta_{t+\Delta} \leq \eta_t$ and $\kappa_{t+\Delta} \geq \kappa_t$ implies that the new point is further away from exiting $\mathcal{D}$.

Ensuring that $\eta_{t+\Delta} \leq \eta_t$ translates to the following condition on the risk-free rate:

$$\begin{aligned} &r_t \geq \tilde{r}_t, \quad \text{whenever} \quad (\eta_t, \kappa_t) \in \mathcal{D}^{\epsilon}_{low}, \\ &\text{where} \quad \tilde{r}_t := \max\Big[\frac{\kappa_t a_e - \rho_e \eta_t q_t}{q_t(\kappa_t - \eta_t)}, \frac{\kappa_t a_e - \rho_e \eta_t q_t(1 + v_t^2/m_t)}{q_t(\kappa_t - \eta_t)} + \frac{v_t^2}{m_t\Delta}\Big]. \end{aligned} \tag{G.43}$$

Now, the equilibrium values of v_t and m_t in (G.41)-(G.42) depend on r_t, so the comparison between r_t and $\tilde{r}_t$ is not explicit. However, we can show that a valid solution to (G.43) exists if Δ is made small enough.

To see this, let us set

$$r_t = \frac{\kappa_t a_e - \rho_e \eta_t q_t}{q_t(\kappa_t - \eta_t)} + \frac{\alpha_t}{\Delta} + C_r \tag{G.44}$$

for some $\alpha_t > 0$ small enough and some constant C_r. Using equations (G.44) and (G.41)-(G.42), one may conjecture and verify that, as $\Delta \to 0$, the variables (r_t, m_t, v_t^2) obey the

[41] Another potential issue is that (η_t, κ_t) can jump from the interior $\mathcal{D}^{\circ}$ to a point outside of the feasible region $\mathcal{D}$. This issue is removed by choosing small enough Δ, because the step sizes in the interior are proportional to $\sqrt{\Delta}$.

following asymptotic relationships

$$r_t\Delta \to \alpha_t$$
$$m_t\Delta \to \alpha_t$$
$$v_t^2/m_t \to \alpha_t.$$

In that case, we have that $r_t - \tilde{r}_t \sim \frac{\rho_e \eta_t q_t \alpha_t}{q_t(\kappa_t - \eta_t)} + \frac{\alpha_t - v_t^2/m_t}{\Delta} + C_r$ as $\Delta \to 0$. Thus, if we pick $C_r = -\lim_{\Delta\to 0} \Delta^{-1}(\alpha_t - v_t^2/m_t)$, the inequality $r_t \geq \tilde{r}_t$ holds for all small enough Δ. It is easy to see that $\Delta^{-1}(v_t^2/m_t - \alpha_t) = O(1)$ as $\Delta \to 0$ so that C_r will be a finite constant. Furthermore, given that α_t is a free parameter, it may be chosen small enough so that upward percentage step size v_t^2/m_t is small enough. Given that the choice (G.44) is continuous in Δ, and equations (G.41)-(G.42) are continuous in r_t, it follows that for all small enough Δ, a valid r_t exists satisfying (G.43).

The final question is whether or not this choice also satisfies $m_t\sqrt{\Delta} \in (-v_t, v_t)$, such that the probabilities of up- and down-moves are within zero and one. To answer this, we can study

$$\frac{v_t^2}{m_t^2\Delta} = 1 + \frac{\frac{a_e}{q_t} + m_t - r_t}{m_t}\left[\frac{(1 + r_t\Delta)\frac{\eta_t}{\kappa_t}}{r_t\Delta - \frac{a_e\Delta}{q_t}} - 1\right]. \tag{G.45}$$

We can see from equation (G.41) that as $\Delta \to 0$, we have

$$\frac{a_e}{q_t} + m_t - r_t \to \frac{1}{1+\alpha_t}\frac{\kappa_t(1-\kappa_t)}{\kappa_t - \eta_t}\frac{a_e - a_h}{q_t}\left[\alpha_t - (1+\alpha_t)\frac{1-\eta_t}{1-\kappa_t}\right] > 0.$$

In addition, the term in square brackets in equation (G.45) is positive in the $\Delta \to 0$ limit if and only if $\kappa_t/\eta_t < (1+\alpha_t)/\alpha_t$. Therefore, by picking α_t small enough, we ensure that the expression in (G.45) is strictly larger than 1 for all Δ small enough. This shows that $m_t\sqrt{\Delta} \in (-v_t, v_t)$ by choosing α_t and Δ small enough.

Jumps to efficiency. At some points when κ_t is sufficiently close to 1, the Brownian approximation above could potentially make κ_t jump above 1, which is inconsistent with equilibrium. At these points, we must instead design the shocks so that κ_t jumps to 1. Such points will constitute the region earlier denoted by $\mathcal{D}_{high}$, whose border with $\mathcal{D}^\circ$ was previously left unspecified and which we will now make explicit.

First, let us define the binomial scheme by

$$u_t = \frac{a_e}{q_t \bar{\rho}(\eta_t^{max})} \tag{G.46}$$

$$d_t = \text{free parameter} \tag{G.47}$$

$$\pi_t = \frac{u_t - 1 - m_t\Delta}{u_t - d_t}, \tag{G.48}$$

where

$$\eta_t^{max} := \frac{\kappa_t a_e(1+\rho_e\Delta) - (\kappa_t - \eta_t)q_t\rho_h(1+r_t\Delta)}{a_e[1+\rho_e\Delta - \kappa_t(\rho_e-\rho_h)\Delta] + (\kappa_t-\eta_t)q_t(1+r_t\Delta)(\rho_e-\rho_h)} \tag{G.49}$$

is the net worth share that would arise if κ jumps to 1.[42] It is straightforward to check that for Δ small enough, we have $\eta_t^{max} < \kappa_t < 1$, so that η_t^{max} is a valid wealth share. Note also that the setup in (G.46)-(G.48) by construction preserves specification of m_t as the local mean $\mathbb{E}_t[\frac{q_{t+\Delta}-q_t}{q_t}] = m_t\Delta$.

The Euler equations become

$$\frac{\kappa_t}{\eta_t} = -(1+r_t\Delta)\frac{(m_t + \frac{a_e}{q_t} - r_t)\Delta}{(u_t + \frac{a_e\Delta}{q_t} - (1+r_t\Delta))(d_t + \frac{a_e\Delta}{q_t} - (1+r_t\Delta))} \tag{G.50}$$

$$\frac{1-\kappa_t}{1-\eta_t} = -(1+r_t\Delta)\frac{(m_t + \frac{a_h}{q_t} - r_t)\Delta}{(u_t + \frac{a_h\Delta}{q_t} - (1+r_t\Delta))(d_t + \frac{a_h\Delta}{q_t} - (1+r_t\Delta))}. \tag{G.51}$$

We can use the two Euler equations to solve for m_t and d_t as

$$\begin{aligned} m_t = r_t &+ \frac{1}{1+r_t\Delta}\frac{\kappa_t(1-\kappa_t)\frac{a_e-a_h}{q_t}(u_t + \frac{a_h\Delta}{q_t} - (1+r_t\Delta))(u_t + \frac{a_e\Delta}{q_t} - (1+r_t\Delta))}{(\kappa_t-\eta_t)(u_t-(1+r_t\Delta)) + \kappa_t(1-\eta_t)\frac{a_e\Delta}{q_t} - \eta_t(1-\kappa_t)\frac{a_h\Delta}{q_t}} \\ &- \frac{(\kappa_t-\eta_t)\frac{a_e a_h\Delta}{q_t^2} + [\kappa_t(1-\eta_t)\frac{a_h}{q_t} - \eta_t(1-\kappa_t)\frac{a_e}{q_t}](u_t-(1+r_t\Delta))}{(\kappa_t-\eta_t)(u_t-(1+r_t\Delta)) + \kappa_t(1-\eta_t)\frac{a_e\Delta}{q_t} - \eta_t(1-\kappa_t)\frac{a_h\Delta}{q_t}} \end{aligned} \tag{G.52}$$

[42] In particular, if κ_t jumps to $\kappa_{t+\Delta} = 1$, then from (G.23) q_t jumps to $q_{t+\Delta} = a_e/\bar{\rho}(\eta_{t+\Delta})$. But the dynamics of η from (G.24) must also hold, which means that $\eta_{t+\Delta}$ solves

$$\eta_{t+\Delta} = \frac{1}{1+\rho_e\Delta}\frac{\kappa_t[\frac{a_e\Delta}{q_t} + \frac{a_e}{q_t\bar{\rho}(\eta_{t+\Delta})} - (1+r_t\Delta)] + \eta_t(1+r_t\Delta)}{a_e/(q_t\bar{\rho}(\eta_{t+\Delta}))}.$$

We denote the solution by η_t^{max}, given in (G.49).

and

$$d_t = (1 + r_t\Delta)\left[1 - \frac{\eta_t}{\kappa_t}\frac{(m_t + \frac{a_e}{q_t} - r_t)\Delta}{u_t + \frac{a_e\Delta}{q_t} - (1 + r_t\Delta)}\right] - \frac{a_e\Delta}{q_t}. \tag{G.53}$$

To guarantee that this constitutes an equilibrium, we must verify $\pi_t \in (0,1)$ along with $0 < d_t < 1 < u_t$.

To check these conditions explicitly, let us pick $r_t = 0$, and let us consider Δ small. As it will turn out (which we will verify below), when Δ is small the region $\mathcal{D}_{high}$ will be associated with $\kappa_t = 1 - O(\sqrt{\Delta})$, so that our choice implies $m_t = -a_h/q_t + O(\sqrt{\Delta})$ from equation (G.52). Substituting this result into equation (G.53), we see that $0 < d_t < 1$ if Δ is small enough. It is easy to check that $u_t > 1$ holds as long as $\rho_e - \rho_h$ is not too large, which we implicitly assume. Lastly, given these results just discussed, we have $\pi_t \in (0,1)$ automatically when Δ is small enough. This shows that, if Δ is small enough, then $r_t = 0$ is a valid choice, and the other equilibrium conditions all hold.

Finally, we need to specify the boundary between $\mathcal{D}_{high}$ and the interior region $\mathcal{D}^\circ$. The procedure will be to compute v_t associated to $\mathcal{D}^\circ$—from equation (G.34)—and then compare $1 + v_t\sqrt{\Delta}$ to $a_e/(q_t\bar{\rho}(\eta_t^{max}))$. If $1 + v_t\sqrt{\Delta} > a_e/(q_t\bar{\rho}(\eta_t^{max}))$ at a given point $(\eta_t, \kappa_t) \in \mathcal{D}$, then we allocate that point to set $\mathcal{D}_{high}$. Otherwise, the given point (η_t, κ_t) is considered to be part of $\mathcal{D}^\circ$. This proves the result used above that $u_t - 1 = O(\sqrt{\Delta})$, and hence $1 - \kappa_t = O(\sqrt{\Delta})$.

Analysis at $\kappa = 1$ border. Finally, given that $\kappa_t = 1$ sometimes, we must describe how the economy exits this region and re-enters the interior $\mathcal{D}^\circ$. We specify a particularly simple approach that always works, although it is unnecessarily restrictive in general.

We will consider a binomial scheme that either maintains $\kappa_{t+\Delta} = 1$ with some probability and otherwise has $\eta_{t+\Delta} \approx 0$ (i.e., expert near-bankruptcy) with the residual probability. This scheme is

$$u_t = 1 \tag{G.54}$$

$$d_t = 1 - \frac{(\eta_t - \omega_t)(1 + \rho_e\Delta)}{1 - \omega_t(1 + \rho_e\Delta)} \tag{G.55}$$

$$\pi_t = \text{free parameter}, \tag{G.56}$$

along with a particular choice for the riskless rate:

$$r_t = \rho_h. \tag{G.57}$$

Using (G.54), (G.55), and (G.57) in the state dynamics (G.24), one can verify that

$$\eta^u_{t+\Delta} = \eta_t$$
$$\eta^d_{t+\Delta} = \omega_t.$$

In other words, a positive shock keeps (η_t, q_t) in place, while a negative shock drives η down to ω_t.

For this to be a valid construction, we require that $q^d_{t+\Delta} = d_t q_t$ is larger than the minimum possible price at the new wealth share, which is $q_{min}(\eta^d_{t+\Delta}) = q_{min}(\omega_t) = (\omega_t a_e + (1-\omega_t)a_h)/\bar{\rho}(\omega_t)$. Using the fact that $q_t = a_e/\bar{\rho}(\eta_t)$, this validity condition is equivalent to

$$\bar{\rho}(\omega_t)\big[1-\eta_t-\big(\bar{\rho}(\eta_t)-(1-\eta_t)\rho_h\big)\Delta\big]a_e > \bar{\rho}(\eta_t)\big[1-\omega_t(1+\rho_e\Delta)\big]\big(\omega_t a_e+(1-\omega_t)a_h\big).$$

As $\Delta \to 0$, this condition becomes

$$\bar{\rho}(\omega_t)(1-\eta_t)a_e > \bar{\rho}(\eta_t)(1-\omega_t)\big(\omega_t a_e+(1-\omega_t)a_h\big).$$

Taking $\omega_t \to 0$ as well, we have the condition

$$\rho_h(1-\eta_t)a_e > \bar{\rho}(\eta_t)a_h \;\Leftrightarrow\; \eta_t < \frac{(a_e-a_h)\rho_h}{(a_e-a_h)\rho_h + a_h\rho_e} := \eta_{top}.$$

Finally, we use the choice of ϵ in (G.25), which implies that the line $\kappa = (1+\epsilon)\eta$ intersects the horizontal line $\kappa = 1$ at a point $\eta < \eta_{top}$. Consequently, if Δ is chosen small enough, equilibrium paths with $\kappa_t = 1$ in period t will have $\eta_t < \eta_{top}$ in the same period. This implies that if Δ and ω_t are chosen small enough, then we can ensure that $q^d_{t+\Delta} > q_{min}(\eta^d_{t+\Delta})$.

Given that $\kappa_t = 1$ at these points, the household Euler inequality (G.22) must hold with strict inequality. A sufficient condition is that households make negative excess returns when capital price remains constant, i.e.,

$$0 > \frac{a_h\Delta + q_{t+\Delta}}{q_t} - R^f_t = \left[\frac{a_h}{a_e}\bar{\rho}(\eta_t) - \rho_h\right]\Delta$$

which always holds since $\rho_e > \bar{\rho}(\eta)$ and $a_e/\rho_e > a_h/\rho_h$.

It remains to verify that the expert Euler equation (G.21) holds. However, this is

guaranteed if the remaining free parameter π_t takes the particular value

$$\pi_t = \frac{(\bar{\rho}(\eta_t) - \rho_h)\Delta}{1 - d_t} + \frac{(\bar{\rho}(\eta_t) - \rho_h)(d_t - 1 + (\bar{\rho}(\eta_t) - \rho_h)\Delta)\Delta}{\eta_t(1 + \rho_h\Delta)(1 - d_t)}.$$

Plugging in d_t from (G.55), we have

$$\pi_t = \frac{1 - \omega_t(1 + \rho_e\Delta)}{(\eta_t - \omega_t)(1 + \rho_e\Delta)}\Big[1 + \frac{(\bar{\rho}(\eta_t) - \rho_h)\Delta - \frac{(\eta_t - \omega_t)(1+\rho_e\Delta)}{1-\omega_t(1+\rho_e\Delta)}}{\eta_t(1 + \rho_h\Delta)}\Big](\bar{\rho}(\eta_t) - \rho_h)\Delta.$$

Note that $\eta_t > \frac{(\eta_t - \omega_t)}{1-\omega_t}$, so that $\pi_t > 0$ for all Δ small enough. In addition, note that $\pi_t \to 0$ as $\Delta \to 0$. Therefore, for all Δ small enough, we are guaranteed to have $\pi_t \in (0, 1)$.

H Stochastic stability in a simplified reduced-form model

Our equilibrium construction differs from the literature. Sunspot equilibria are often constructed as follows:

1. Analyze the deterministic steady state of your dynamic model (suppose there is only one).
   (a) If the steady state is "unstable" or "saddle path stable," then there is a unique equilibrium, which involves either jumping directly to the steady state or jumping to a unique transition path toward steady state, respectively.
   (b) If the steady state is "stable," typically diagnosed by showing that the linearized transition dynamics have more stable eigenvalues than state variables (pre-determined variables), then there will be multiple deterministic transition paths toward steady state.
2. In the stable case, one can add sunspot shocks, so long as these shocks do not cause the system to leave the "stable region" near steady state. Sunspot shocks essentially randomize over the starting points of the multiplicity of deterministic transition paths.

Classic papers like Azariadis (1981) and Cass and Shell (1983) follow this approach. More recent papers Gârleanu and Panageas (2021) and Khorrami and Zentefis (2025) generalize this same idea to larger classes of models and sunspot shocks.

By contrast, the deterministic version of our model features an unstable steady state; critically, the introduction of volatility flips the stability properties of equilibrium. This distinction may be why our equilibria have gone unnoticed despite the framework being widespread.

To clarify our methodology without forcing the reader deep into our proofs, let us present a simple example. The analogy between this example and our model is not perfect but still illuminates some key issues. Consider an economy with a single endogenous outcome x_t and no fundamental state variables. Suppose x_t follows a process of the form

$$dx_t = \mu_t dt + \sigma_t dZ_t, \tag{H.1}$$

where Z_t is a one-dimensional Brownian motion. To fix ideas, think of x_t as an equilibrium asset price, so that there is no initial condition associated with (H.1). The drift μ_t and volatility σ_t are also endogenous.

Equilibrium is defined as any x_t process that satisfies

$$0 = F(x_t, \mu_t, \sigma_t), \quad \text{where} \quad x_t \in (\underline{x}, \bar{x}), \tag{H.2}$$

where F is some known function. This situation is analogous to our main model, where the future distribution of asset prices (captured by μ_t and σ_t) can influence today's price. Think of F as an equilibrium condition which relates the market interest rate and risk premium to agents' capital and bond holdings. The domain $(\underline{x}, \bar{x})$ arises because of constraints on agents' investment decisions.

To transparently elaborate on our methodology and connect it to our model, let us specialize to the following parametric example:

$$F(x, \mu, \sigma) = \begin{cases} v - x - \sigma, & \text{if } \sigma \neq 0 \\ \phi\mu - x, & \text{if } \sigma = 0, \end{cases} \quad \text{and} \quad \underline{x} < 0 < \bar{x}, \tag{H.3}$$

where $v > 0$, $\phi \geq 0$ are parameters. The key feature of F, shared by our main model, is that it depends in a critical way on whether or not there is volatility. In a stochastic equilibrium, $F = 0$ places no restrictions on the drift μ. In a deterministic equilibrium, $F = 0$ pins down the drift $\mu = \phi^{-1}x$ as an increasing function of the price, provided $\phi > 0$. (In fact, our main model has the property that $\phi = 0$ so that $x = 0$ is immediately pinned down as the unique deterministic equilibrium.) The criticality of σ stems from deeper economic forces that we highlight in our paper.

What are the equilibria in this environment? A conventional analysis would start by examining the deterministic steady state and its stability properties. In this example, the steady state is $x^* = 0$. It is unstable when $\phi > 0$. Instability implies that any deterministic equilibrium has $x_t = 0$ forever: if $x_0 \neq 0$, then x_t eventually exits the domain $(\underline{x}, \bar{x})$. Likewise, our main model has a unique deterministic equilibrium. By contrast, non-fundamental equilibria arise in many papers because of steady state stability and the associated multiplicity of transition paths, which here would require $\phi < 0$.

In stochastic equilibria, volatility is pinned down as a decreasing function of the asset price $\sigma(x) = v - x$. The fact that volatility is pinned down in this reduced-form setup is also a feature shared by our model (a similar phenomenon occurs in Benhabib et al., 2015, where the sentiment distribution is uniquely-determined). Since the drift μ is not pinned down when $\sigma \neq 0$, we may engineer it to ensure that $x_t \in (\underline{x}, \bar{x})$ at all times. For

example, the function

$$\mu(x) = \left(\frac{1}{x - \underline{x}}\right)^2 - \left(\frac{1}{\bar{x} - x}\right)^2 \tag{H.4}$$

diverges sufficiently quickly as x approaches the boundaries of $(\underline{x}, \bar{x})$, enough to ensure that shocks do not push x_t outside of its equilibrium domain. Thus, the dynamics are "stochastically stable" and an equilibrium exists of this type.

The most important methodological point here is that only boundary dynamics matter for stochastic stability. To see this, consider modifying (H.4) to

$$\mu(x) = \left(\frac{1}{x - \underline{x}}\right)^2 \mathbf{1}_{x \leq \underline{x} + \epsilon} - \left(\frac{1}{\bar{x} - x}\right)^2 \mathbf{1}_{x \geq \bar{x} - \epsilon} + \hat{\mu}(x) \mathbf{1}_{x \in (\underline{x} + \epsilon, \bar{x} - \epsilon)}, \tag{H.5}$$

for ϵ small. This modified drift also diverges quickly at the boundaries, but can have almost arbitrary behavior away from the boundaries. Still, x_t is stochastically stable under (H.5), regardless of $\hat{\mu}(x)$. Our analysis reveals that the typical question of local stability near steady state, which concerns the function $\hat{\mu}(x)$, is irrelevant to global stochastic stability. Because only the boundary behavior matters, we must abandon the standard linearized spectral analysis that may be more familiar to many readers. Related analyses of boundary behavior occur in the literature on the "market-selection hypothesis" (Kogan et al., 2017; Borovička, 2020).

To conclude, this simple reduced-form example features a unique deterministic equilibrium ($x_t = 0$ for all t) but a variety of stochastic equilibria. For example, if agents coordinate on a price drift function given by either (H.4) or (H.5), then x_0 can take any value in $(\underline{x}, \bar{x})$ and x_t will never leave this region. Furthermore, it will not be the case (except for knife-edge choices of the parameters $\underline{x}, \bar{x}, v$) that $x_t \to 0$ asymptotically. This economy can stay stochastic forever.

How many of these stochastic equilibria are there? For one, the initial condition x_0 is not pinned down. Additionally, the choice of drift $\mu(x)$ only needs to satisfy certain boundary conditions that keep $x_t \in (\underline{x}, \bar{x})$. This drift can essentially take any functional form away from these boundaries. Thus, there are a great variety of equilibria, but they all share the property that σ_t is pinned down given x_t, as in our baseline model.