Model Revisions in General Equilibrium*

Aleksei Oskolkov

Princeton University

alekseioskolkov@princeton.edu

September 7, 2026

Abstract

I develop a general equilibrium framework for asset pricing with subjective models. Investors doubt their models of returns and choose subjective models for robustness. Their preferences for robustness change over time, inducing model revisions. Equilibrium is characterized by a representative subjective model. Objective risk premia are more volatile than subjective: within models, investors demand constant risk compensation, and objective risk premia move following changes across models. The representative model and the objective risk premium change due to individual model revisions and due to wealth redistribution between agents with conflicting models. I show analytically how data on subjective expectations identify these two channels. In US data, 70% of variation in risk prices at quarterly frequency are due to model revisions. The variance share of wealth redistribution increases in horizon.

Key Words: risk premium, subjective beliefs, aggregation, robustness

*I thank Fernando Alvarez, Manuel Amador, Fernando Arce, Christina Arellano, Omar Barbiero, Markus Brunnermeier, Ricardo Caballero, V.V. Chari, Roberto Corrao, Sebastian Di Tella, Stelios Fourakis, Lars Peter Hansen, Jonathan Heathcote, Zhen Huo, Magnus Irie, Elena Istomina, Oleg Itskhoki, Rohan Kekre, Arvind Krishnamurthy, Moritz Lenel, Ellen McGrattan, Alp Simsek, Hyun Song Shin, Chris Tonetti, Nico Trachter, Kaushik Vasudevan, and all seminar participants for helpful comments.

1 Introduction

Subjective expectations of excess returns coming from investor surveys differ from forecasts available to econometricians. Piazzesi, Schneider, and Salomao (2015) document this for bond markets, showing that subjective expectations of investors are substantially less volatile and less cyclical than the econometricians' expectations. Nagel and Xu (2023) extend the exercise to stocks, currencies, and commodity futures, and explain their findings through the lens of learning: investors gradually learn about the model and revise expectations with the arrival of new data.

Accounting for subjective models is important for interpreting asset price data. Diversity of models and reallocation of market activity across agents with different models matters as well: wealth distribution across investors with heterogeneous risk-taking capacity is a key determinant of asset prices in a large class of macro-financial models, dating back at least to the seminal contributions of He and Krishnamurthy (2013) and Brunnermeier and Sannikov (2014).

In this paper, I propose an equilibrium asset pricing model with dynamic subjective models. My framework allows for heterogeneous agents and achieves clean aggregation, allowing me to decompose the dynamics of risk premia into individual model revisions and wealth redistribution between agents with different beliefs. The objective-subjective gap in expectations and forecast errors identify the evolution of subjective models and of the moments of the wealth distribution that determine the market price of risk. I show identification analytically and decompose the historical evolution of US equity premia into the model revision and wealth redistribution channels. Model revisions contribute 70% of the variation in risk premia at short horizons, and wealth redistribution between agents with heterogeneous beliefs becomes more important at lower frequencies.

I start with a model of subjective model revisions. Investors choose portfolios of risky assets. They are concerned about model misspecification when it comes to the return process. Looking for robustness, they choose alternative statistical models for returns, benchmarking them against the baseline physical measure. The key primitive is a robustness parameter that penalizes deviations from the baseline. This robustness parameter varies over time, making investors choose models with different levels of pessimism. Unlike in the standard robust control setup, such as in Hansen and Sargent (2001) and closely related papers Bhandari, Borovička, and Ho (2025) and Szőke (2022), investors are not concerned about their model of the evolution of states, questioning instead the mapping from states to returns, which leads to simple portfolio choice.

I next show that this specification of investors' problem is isomorphic to one with exogenous risk limits of the value-at-risk type. Through the lens of my framework, changes in subjective models of returns look exactly like changes in self-imposed risk limits. Barbiero, Bräuning, Joaquim, and Stein (2024) use supervisory data on desk-level risk limits at large dealer banks and show that these constraints are subject to recurrent shocks orthogonal to commonly studied factors and market conditions. My framework explains how shocks to internal risk limits might emerge from model revisions in the dealers' risk departments that are transmitted to trading desks.

I then study an economy with multiple agents with time-varying robustness parameters and report two main results. The first is that the model displays aggregation: a single statistic, the wealth-weighted average of robustness parameters, determines the risk-free rate and the evolution of asset prices. It acts as the economy's effective risk tolerance and maps into the market price of risk. This objective price of risk is more volatile than the subjective one. The reason is that, given a model, investors with logarithmic utility choose mean-variance portfolios, and the required risk compensation within their models stays constant. Subjective risk premia follow the quantity of risk, consistent with results of Nagel and Xu (2023), and the subjective risk price is fixed. The objective price moves because of changes across models instead of within: either shocks to robustness parameters trigger individual model revisions or wealth redistribution changes the market's effective model.

The second result is that these two channels can be separately identified. Despite its large dimensionality, the evolution of the wealth distribution is driven by a single stochastic process: the forecast error in the representative subjective model of the total market return. Positive surprises in growth redistribute toward less pessimistic agents, who lever up to take on aggregate risk. These agents also become richer deterministically because the representative forecast error is biased when the market is pessimistic on average. Importantly, redistribution is fully driven by real shocks. Individual model revisions do not trigger wealth revaluation. Instead, they change agents' exposures to aggregate risk, and the realization of aggregate risk then redistributes between them. This clear division of labor resembles the result in Barbiero, Bräuning, Joaquim, and Stein (2024), who document that idiosyncratic contractions in risk limits lead to other intermediaries stepping in and taking over the risk.

I use these two results to show that data on subjective and objective beliefs and forecast errors identify model revision shocks and redistributive changes in risk premia. First, the wealth-weighted average of the economy's robustness parameters, which maps to the market price of risk, is identified by the ratio between subjective and objective expected excess returns on stocks. Second, the redistribution-induced changes in this average are proportional to the representative subjective forecast error. Model revisions, the other component of its evolution, are conditionally orthogonal to the forecast error. As an overidentifying restriction, I use the fact that the drift of the forecast error itself depends on the average too and co-moves with the market price of risk.

I follow the approach of Bhandari, Borovička, and Ho (2025) and Szőke (2022), who use survey evidence on beliefs and interpret reported expectations as coming from subjective models chosen by robust agents. I take survey data from Nagel and Xu (2023) and replicate their approach to constructing objective risk premia. With objective and subjective expectations and forecast errors in hand, I reconstruct the evolution of the market price of risk from 1987 to 2021 and decompose it into model revision shocks and redistribution. At the quarterly frequency, the decomposition attributes 70% of the variation in the market price of risk to model revision shocks. The variance share of redistribution increases with horizon. Validating the model, I find that the redistribution component closely follows the changes in the intermediary capital ratio from He, Kelly, and Manela (2017), showing a stable rolling-window correlation between 0.55 and 0.85 throughout the sample. Through the lens of my model, intermediaries are agents with less pessimistic models, who take on aggregate risk in equilibrium and whose wealth share is negatively related to the market price of risk. In contrast, the model revision component shows correlation around 0.1 with changes in the intermediaries' capital ratio, which I interpret as evidence supporting my identification results.

1.1 Related literature

My paper is related to a few literatures. First, I contribute to the literature working on the gap between subjective and objective expectations. Piazzesi, Schneider, and Salomao (2015) document that subjective expectations of bond returns are less volatile than expectations from their econometric framework, respond less to business-cycle variables, and contain a lower-frequency cycle related to inflation. Nagel and Xu (2023) extend the measurement to multiple asset classes and combine multiple surveys. In Nagel and Xu (2022), they offer a learning model with fading memory that explains this pattern: the model generates cyclical and volatile objective premia with very stable subjective ones. I offer a complementary framework with dynamic model revisions. The difference is that I endogenize model choice with misspecification concerns instead and add general equilibrium analysis with model heterogeneity. A related theory is Maenhout, Vedolin, and Xing (2025), who use a different specification for robustness concerns and arrive at endogenous waves of pessimism. I endogenize the dynamics of the market's pessimism through heterogeneity and identify model revisions separately from wealth redistribution in a heterogeneous-agent setup.

Second, I contribute to four strands of the literature studying dynamics of risk premia. One of them uses models with heterogeneous agents where risk premia change because wealth gets redistributed. This is characterized in He and Krishnamurthy (2013) and Brunnermeier and Sannikov (2014) and the long literature that follows their general equilibrium frameworks. Most closely related to my setup are Caballero and Simsek (2020) and Kekre and Lenel (2022). Caballero and Simsek (2020) develop a risk-centric macroeconomic model, where growth shocks redistribute between optimistic and pessimistic agents and through that affect risk premia, which then impact aggregate demand. Kekre and Lenel (2022) provide a clear exposition of this force in a simple model and then study it in a quantitative model of the US economy. A special case of my model provides a clear exposition of the redistributive force in these two papers.

The other class of models with dynamic risk premia studies dynamic preferences, technologies, models, or risk-taking constraints. I contribute to it by developing a theory of dynamic risk limits. Specifically, I show that model revisions generate exactly the same portfolio choice as portfolio constraints of the value-at-risk type and can explain where changes in these constraints come from. Value-at-risk measures are widely used in banking. Sizova (2023) collects information on banks updating their models from financial reports and shows that banks often mention "value-at-risk models" in their reports. Barbiero, Bräuning, Joaquim, and Stein (2024) collect data on risk limits faced by the largest US banks. They show that 165 of 167 trading desks whose activities are related to currency trading face value-at-risk constraints. Adrian and Shin (2010) and Adrian and Shin (2014) measure value-at-risk constraints in the largest US banks and provide a foundation for them using the contractual framework from Holmstrom and Tirole (1997). The use of these constraints in macroeconomics dates back to at least Basak and Shapiro (2001), Danielsson, Shin, and Zigrand (2012), and Adrian and Boyarchenko (2018). Hofmann, Shim, and Shin (2022) use them to show the impact of dollar appreciation on portfolio inflows in emerging markets. Empirically, Coimbra, Kim, and Rey (2022) estimate value-at-risk parameters from bank balance sheet data and find substantial heterogeneity in cross-section. Coimbra and Rey (2024) develop a general equilibrium model with heterogeneous value-at-risk constraints on intermediaries. Bräuning and Stein (2024) find that limit changes also affect the functioning of the Treasury market.

Third, a large set of models rely on mean-variance preferences. Vayanos and Vila (2021), Gourinchas, Ray, and Vayanos (2022), Ray (2019), Kamdar and Ray (2024), and Greenwood, Hanson, Stein, and Sunderam (2023) endow arbitrageurs with mean-variance preferences, mentioning that this can capture value-at-risk constraints in reduced form. A recent literature in international macro uses risk aversion shocks of myopic intermediaries to achieve time-varying risk premia. Itskhoki and Mukhin (2021) and Kekre and Lenel (2024) augment the framework developed by Gabaix and Maggiori (2015) with time variation in the risk aversion parameters to achieve desirable properties of exchange rates. Kekre and Lenel (2025) replace risk aversion shocks with monetary shocks that destroy intermediary wealth with similar results.

Fourth, my model revision setup is related to a large literature on robust control and model misspecification. The classical reference for this framework is Hansen and Sargent (2001). Hansen, Szőke, Han, and Sargent (2020) endogenize model choice by fixing structured alternatives and allowing for clouds of models about them. Oskolkov (2026) links model revisions in global financial intermediaries to the global financial cycles. The main difference between my model and theirs is that I separate model misspecification for returns and aggregate states. The agent explores alternative probability measures for shocks driving excess returns but ignores the implications of these alternatives for her views on aggregate states. This allows me to eliminate hedging motives and reduce portfolio choice to mean-variance, breaking away from the equivalence between robustness concerns and Duffie and Epstein (1992) recursive utility.

2 Model

I start by describing the setup and the portfolio problem. Time is continuous and runs forever. There is an infinitely-lived investor. Her state variable is a d-dimensional vector x_t that evolves as

$$dx_t = \mu_X(x_t)dt + \sigma_X(x_t)dZ_t$$

Here $\{Z_t\}_{t\geq 0}$ is a standard b-dimensional Brownian process, so $\sigma_X(x_t)$ is a $(d \times b)$-dimensional matrix, and $\mu_X(x_t)$ is a vector of length d. The investor has access to a riskless asset that pays an instantaneous return $r(x_t)dt$ and a collection of k risky assets with a vector of instantaneous excess returns dR_t given by

$$dR_t = \mu_R(x_t)dt + \sigma_R(x_t)dZ_t$$

Here $\sigma_R(x_t)$ is a $(k \times b)$-dimensional matrix and $\mu_R(x_t)$ is a vector of length k. Returns and exogenous states do not necessarily load on all shocks, so some columns of $\sigma_R(x_t)$ and $\sigma_X(x_t)$ may be zero. I only require $\sigma_R(x_t)\sigma_R(x_t)'$ to have full rank for all x_t. The investor's wealth evolves as

$$dw_t = (r(x_t)w_t - c_t)dt + w_t\theta_t' dR_t$$

Here c_t is consumption and θ_t is a k-dimensional vector of portfolio weights on risky assets.

The investor entertains alternative models of returns. Specifically, she thinks that

$$dR_t = \mu_R(x_t)dt + \sigma_R(x_t)dB_t,$$

where $\{B_t\}_{t\geq 0}$ is a separate b-dimensional shock process that does not coincide with $\{Z_t\}_{t\geq 0}$. Under the investor's model, which is formalized below, all components of dB_t are conditionally perfectly correlated with the corresponding components of dZ_t. However, she doubts that $\{B_t\}_{t\geq 0}$ is truly Brownian, thinking that a truly Brownian process is $\{\tilde{B}_t\}_{t\geq 0}$ with $d\tilde{B}_t = dB_t + h_t dt$. The b-dimensional model correction process $\{h_t\}_{t\geq 0}$ is a choice variable. She wishes to protect her portfolio from potential misspecification of the statistical model of returns.

At the same time, the investor does not doubt the model for the state x_t. This is the key difference compared to the standard setup in Hansen and Sargent (2001). Effectively, my assumption on the investor's doubt means that she questions the mapping from states to returns, while being sure she knows how the states evolve.

Formally, let the original process $\{B_t\}_{t\geq 0}$ be a standard b-dimensional Brownian motion on the filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t\geq 0}, \mathbb{P})$. Let $\{h_t\}_{t\geq 0}$ be a process adapted to $\{\mathcal{F}_t\}_{t\geq 0}$. By Girsanov's theorem, under a condition on $\{h_t\}_{t\geq 0}$, the process $\{\tilde{B}_t\}_{t\geq 0}$ given by $\tilde{B}_0 = B_0$

and $d\tilde{B}_t = dB_t + h_t dt$ is a standard b-dimensional Brownian motion on another statistical basis $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t\geq 0}, \mathbb{Q})$. I call $\mathbb{Q}$ the subjective measure. It satisfies $\mathbb{E}^{\mathbb{Q}}[\varphi_t] = \mathbb{E}^{\mathbb{P}}[M_t \varphi_t]$ for all $t \geq 0$ and for all bounded processes $\{\varphi_t\}_{t\geq 0}$ adapted to $\{\mathcal{F}_t\}_{t\geq 0}$. The process $\{M_t\}_{t\geq 0}$ is adapted to $\{\mathcal{F}_t\}_{t\geq 0}$ and given by $M_0 = 1$ and $dM_t = -M_t h_t' dB_t$. This re-weighting process is a likelihood ratio, and its logarithm $m_t = \log(M_t)$ evolves as

$$dm_t = -\frac{|h_t|^2}{2} dt - h_t' dB_t = \frac{|h_t|^2}{2} dt - h_t' d\tilde{B}_t$$

The investor uses this log-likelihood ratio to discipline her choices. One restriction that using log-likelihood ratios imposes on the environment is that she cannot entertain models with different quantities of risk. With Brownian shocks, volatility and correlation are instantly learnable, and hence create infinite log-likelihood ratios. This limits model adjustments to drift corrections h_t.

The problem of the investor is of the "multiplier" type:

$$V(w_t, x_t) = \max_{\{c_t, \theta_t\}_{t\geq 0}} \inf_{\mathbb{Q}} \mathbb{E}^{\mathbb{Q}} \left[\rho \int_0^\infty e^{-\rho t} \log(c_t) dt + \int_0^\infty e^{-\rho t} \psi_t dm_t \right]$$

She chooses the subjective measure $\mathbb{Q}$ with potentially large losses and then maximizes utility over consumption and portfolio vectors that are robust to these losses. The discipline is provided by the cost of deviating from the original measure as measured by $\psi_t dm_t$ for each increment of time. Small values of the multiplier ψ_t make it easier to select pessimistic probability measures $\mathbb{Q}$, and the investor underestimates expected excess returns, which makes risky assets less attractive.

Proposition 1. *Consider a robust investor with a cost process $\{\psi_t\}_{t\geq 0}$. Her consumption and portfolio choice are*

$$\begin{aligned} c(w_t, x_t) &= \rho w_t \\ \theta(w_t, x_t) &= \frac{\psi_t}{\psi_t + 1} \cdot [\sigma_R(x_t)\sigma_R(x_t)']^{-1} \mu_R(x_t) \end{aligned}$$

The drift corrections she chooses for instantaneous returns are

$$h(w_t, x_t) = \frac{1}{\psi_t + 1} \cdot \sigma_R(x_t)' [\sigma_R(x_t)\sigma_R(x_t)']^{-1} \mu_R(x_t)$$

Her value function is separable over wealth and exogenous states: $V(w_t, x_t) = \log(w_t) + \eta(x_t)$, where $\eta(\cdot)$ solves a second-order partial differential equation.

Choosing alternative measures, or models, $\mathbb{Q}$ makes the agent act like a mean-variance investor with effective risk-tolerance $\psi_t/(\psi_t + 1)$. This portfolio choice is myopic: it features no hedging motives. Achieving this simplicity in portfolio choice often requires working with short-lived agents.

However, my investor is forward-looking and solves a consumption-savings problem too, which is crucial for endogenizing asset prices and the interest rate $r(x_t)$. The fact that consumption is a constant fraction of wealth allows for simple linear aggregation, as is always the case with a unit elasticity of intertemporal substitution.

For all positive ψ_t, portfolio weights are below those of a regular log investor, meaning that robustness adjustments always induce pessimism. Effective risk-tolerance increases in ψ_t because a higher ψ_t penalizes choosing adverse models more. The case $\psi_t = 0$ corresponds to infinite effective risk aversion, while $\psi_t \longrightarrow \infty$ corresponds to a standard log investor who does not make any model adjustments because it is prohibitively costly.

Importantly, within models the investor chooses exact mean-variance portfolios due to log utility. Given a model, she demands a constant risk compensation, in fact, equal to one. Changes in effective risk tolerance come from shifting across models instead. The shape of the mean-variance tradeoff is preserved, up to scaling, because the drift correction $h_t dt$ applies to the underlying shock dB_t, and hence picks up the loading $\sigma_R(x_t)$ when the investor converts it to expected returns.

In Appendix A.1, I show that the tractable portfolios and fixed consumption shares survive adding taxes, including stochastic ones, that are proportional to wealth.

Relation to standard robust preferences. Divorcing the shocks to returns $dB_t - h_t dt$ from aggregate shocks dB_t is key to obtaining tractable mean-variance portfolios. In the standard setup, shocks driving states and returns are required to coincide in all alternative models:

$$\begin{aligned} dR_t &= (\mu_R(x_t) - \sigma_R(x_t)h_t)dt + \sigma_R(x_t)dB_t \\ dx_t &= (\mu_X(x_t) - \sigma_X(x_t)h_t)dt + \sigma_X(x_t)dB_t \end{aligned}$$

The optimal choice of h_t picks up the evolution of x_t and its impact on the investor's value through the gradient $V_{x'}(w,x)$. Since h_t affects the optimal choice of θ_t, the optimal portfolio picks up $V_{x'}(w,x)$ too, and this introduces the hedging motive in portfolio choice: optimal shares θ_t start to depend on $V_{x'}(w,x)$. The value function is still separable: $V(w_t,x_t) = \log(w_t) + \eta(x_t)$. Using this, portfolio choice can be written with two parts:

$$\theta(w_t,x_t) = \underbrace{\frac{\psi_t}{1+\psi_t} \cdot [\sigma_R(x_t)\sigma_R(x_t)']^{-1}\mu_R(x_t)}_{\text{mean-variance}} - \underbrace{\frac{1}{(1+\psi_t)}[\sigma_R(x_t)\sigma_R(x_t)']^{-1}\sigma_R(x_t)\sigma_X(x_t)'\eta_{x'}(x_t)}_{\text{hedging motive}}$$

The mean-variance part coincides with the portfolio in Proposition 1, but the hedging motive part requires solving a second-order PDE to obtain $\eta(x_t)$. I provide more detail in Appendix A.2.

Hedging motives disappear in the limit of infinite robustness parameter, $\psi_t \longrightarrow \infty$, where the agent becomes a standard log investor. Another case where they disappear is one with

$\sigma_R(x_t)\sigma_X(x_t)' = 0$, for example if returns and aggregate states load on different shocks: some columns of $\sigma_R(x_t)$ and $\sigma_X(x_t)$ are zero, and the sets of their non-zero columns do not intersect. This happens when returns are purely idiosyncratic, and aggregates do not load on idiosyncratic shocks due to a large population. An example is the model of Di Tella, Malgieri, and Tonetti (2024). In their economy, entrepreneurs face idiosyncratic productivity risk. The hiring decision is akin to an investment problem with a risky asset. There is a continuum of entrepreneurs, so idiosyncratic shocks do not drive aggregate dynamics. Drift corrections applied to idiosyncratic shocks would not appear in any process for aggregates, and standard robust preferences would produce mean-variance portfolios with effective risk tolerance given by the robustness parameter.

3 Equivalent setup with risk limits

I now describe an alternative problem that induces exactly the same consumption and portfolio choice. The investor does not doubt her model. Instead, she faces a portfolio constraint on risk-taking. The natural origin of such risk limits is regulation. However, my equivalence result shows that, when interpreting recurrent changes in risk limits in the data, tracing them back to changing regulation is not necessary. What looks like a tightening constraint might be a shift in a subjective model: it is impossible to tell the two apart in consumption and portfolio data. In environments with delegated portfolio choice, like in large financial firms, changes in the internal model at the top might be reflected in changing risk limits that are passed downstream to portfolio managers, since communicating an entire model might be harder than setting an equivalent constraint.

As before, the state, returns, and investor's wealth evolve as

$$\begin{aligned} dx_t &= \mu_X(x_t)dt + \sigma_X(x_t)dZ_t \\ dR_t &= \mu_R(x_t)dt + \sigma_R(x_t)dZ_t \\ dw_t &= (r(x_t)w_t - c_t)dt + w_t\theta_t' dR_t \end{aligned}$$

The pair (c_t, θ_t) are now the only control variables.

The problem of the investor is to solve

$$\begin{aligned} V(w_t, x_t) = \max_{\{c_t,\theta_t\}_{t\geq 0}} \mathbb{E}\left[\rho \int_0^\infty e^{-\rho t}\log(c_t)dt\right] \\ \text{s.t. } \mathbb{V}_t[\theta_t' dR_t] \leq \gamma_t \cdot \mathbb{E}_t[\theta_t' dR_t] \end{aligned} \tag{1}$$

The value-at-risk constraint equation (1) is the key feature. It is imposed continuously on incremental returns and caps the variance of returns by a multiple of expected profits. Both variance and expectation are of order dt, so dt cancels out. The multiplier $\gamma_t \in (0, 1]$ is one of the components of

x_t. It is exogenous and stochastic, potentially driving returns and other macroeconomic outcomes. In Appendix A.3, I provide a heuristic explanation for why this constraint limits value-at-risk.

PROPOSITION 2. *Investor's consumption and portfolio choice are*

$$c_t(w_t, x_t) = \rho w_t$$
$$\theta_t(w_t, x_t) = \gamma_t \cdot [\sigma_R(x_t)\sigma_R(x_t)']^{-1}\mu_R(x_t)$$

Investor's value function is separable over wealth and exogenous states: $V(w_t, x_t) = \log(w_t) + \eta(x_t)$, *where* $\eta(\cdot)$ *solves a second-order partial differential equation.*

Risk limits formulated in equation (1) lead to exactly the same consumption and portfolio choice as robustness concerns and model revisions in Section 2 with the correspondence

$$\gamma_t = \frac{\psi_t}{\psi_t + 1}$$

Imposing this risk limit on the agent does not break the constant-share consumption rule. The time-varying limit γ_t acts as risk tolerance, which in Section 2 is an increasing function of ψ_t.

Relation to other value-at-risk constraints and recursive preferences. A literal interpretation of value-at-risk in the literature is a cap on the probability of losses exceeding a certain proportion of wealth. Examples of this approach include Danielsson, Shin, and Zigrand (2011), Danielsson, Shin, and Zigrand (2012), and Adrian and Boyarchenko (2018). In practice, with normally distributed shocks, this leads to equation (1) with another upper bound for variance instead of expected returns. The upper bound is usually a stock variable, such as wealth, instead of a flow variable, such as expected returns, making it less interpretable with Brownian shocks, which make any realistic losses too small compared to total capital. Interestingly, this formulation leads to the same mean-variance portfolios as in Proposition 2, except the effective risk tolerance γ_t is endogenous and includes Lagrange multipliers that require additional solution steps. Specifically,

$$\theta(w_t, x_t) = \lambda(w_t, x_t) \cdot [\sigma_R(x_t)\sigma_R(x_t)']^{-1}\mu_R(x_t)$$

Here the effective risk tolerance $\lambda(\cdot)$ is a function of both individual wealth and exogenous states, which makes this formulation of risk limits empirically appealing, since it can deliver cyclical leverage, but challenging for aggregation and computation.

In Appendix A.4, I also show how exogenous risk limits are related to recursive preferences of Kreps and Porteus (1978) and Duffie and Epstein (1992). The relationship is the same as the one between my model revision setup in Section 2 and the standard robustness setup of Hansen and Sargent (2001). Recursive preferences with a unit elasticity of intertemporal substitution and a constant risk aversion of $1/\gamma$ lead to two-part portfolios:

$$\theta(w_t, x_t) = \underbrace{\gamma \cdot [\sigma_R(x_t)\sigma_R(x_t)']^{-1}\mu_R(x_t)}_{\text{mean-variance}} + \underbrace{\frac{\gamma - 1}{\eta(x_t)}[\sigma_R(x_t)\sigma_R(x_t)']^{-1}\sigma_R(x_t)\sigma_X(x_t)'\eta_{x'}(x_t)}_{\text{hedging motive}}$$

Here $\eta(x_t)$ is a part of the value function, which factorizes: $V(w_t, x_t) \propto (w_t\eta(x_t))^{1-1/\gamma}$. The hedging motive disappears if $\gamma \longrightarrow 1$, which corresponds to the case of an infinite robustness parameter, $\psi_t \longrightarrow \infty$, and makes the agent a log investor.

Finally, in Appendix A.5, I show that this formulation of the problem with risk limits survives adding external income and taxes, provided that they are proportional to wealth, just like in the setup of Section 2.

Discussion. Three things are worth emphasizing. First, risk limits formulated in equation (1) lead to simple mean-variance portfolios with long-lived agents. A large class of models uses short-lived agents to obtain these simple portfolios, which rules out endogenizing the interest rate. Prominent examples include models in the preferred habitat tradition of Vayanos and Vila (2021), where the authors suggest that emergence of mean-variance preferences in their setup could be linked to risk limits. I formalize this link.

Second, the exact equivalence of portfolio choice between the risk limit setup of this section and the subjective model setup of Section 2 allows for explaining changing risk limits with internal model revisions. Barbiero et al. (2024) document shocks to desk-level risk limits at large dealer banks. They find recurring transitory shocks orthogonal to common factors and market conditions. Moreover, they find idiosyncratic shocks, which rules out common changes in regulation as their cause. The theory of subjective models provides an alternative cause for these shocks.

Third, for the purpose of model discrimination, a model with exogenous risk limits cannot be distinguished from a model with exogenous robustness parameters with portfolio data alone. A direct way to tell the two sets of primitives apart would be to look at subjective beliefs of investors. I do this in Section 5 below.

4 Aggregation

I now construct an economy with multiple agents who hold potentially different subjective models. This economy aggregates well, mapping the market price of risk into a single scalar that summarizes the wealth distribution: the wealth-weighted average of robustness parameters. This gives rise to a "representative" agent and a "representative" subjective model. I qualify the degree of aggregation and explain exactly why I put quotation marks around "representative" and why the entire wealth distribution still matters. I then show that the evolution of the distribution is governed by a single process: the forecast error in the representative subjective model. Last come identification results that show how to map data on subjective and objective expectations of market returns and forecast errors to the unobservable price of risk and the two forces that contribute to its evolution: revision of individual models and redistribution of wealth across agents with heterogeneous models.

Suppose there are n investors indexed by i. Their robustness parameters $\{\psi_{it}\}$ are potentially different and stochastic. For ease of exposition, I make a monotone change of variables to

$$\gamma_{it} = \frac{\psi_{it}}{\psi_{it} + 1}$$

I will hereafter refer to γ_{it} as i's robustness parameter, understanding that an increase in γ_{it} is an increase in ψ_{it} and hence in the cost of choosing a pessimistic subjective model. The vector of parameters $\boldsymbol{\gamma}_t$ evolves as

$$d\,\boldsymbol{\gamma}_t = \mu_\gamma(\boldsymbol{\gamma}_t)dt + \sigma_\gamma(\boldsymbol{\gamma}_t)dZ_t^\gamma$$

Here dZ_t^γ is an increment of a g-dimensional standard Brownian motion, $\mu_\gamma(\boldsymbol{\gamma}_t)$ is an $n \times 1$ column vector, and $\sigma_\gamma(\boldsymbol{\gamma}_t)$ is an $n \times g$ matrix of loadings.

The rest of the exogenous states are the dividends of risky assets. There are k risky assets indexed by j. The supply s_j of each asset j is fixed, with $s_j > 0$ for at least one j. The dividends are denoted by $\{y_{jt}\}$, and the vector of dividends $\boldsymbol{y}_t$ evolves as

$$d\,\boldsymbol{y}_t = \mu_y(\boldsymbol{y}_t)dt + \sigma_y(\boldsymbol{y}_t)dZ_t^y$$

Here dZ_t^y is an increment of a d-dimensional standard Brownian motion independent of dZ_t^γ, the drift $\mu_y(\boldsymbol{y}_t)$ is a $k \times 1$ column vector, and $\sigma_y(\boldsymbol{y}_t)$ is a $k \times d$ matrix of loadings. The state space is completed by the vector of wealth. Investor i's wealth is w_{it}, and the state vector is $x_t' = (\boldsymbol{\gamma}_t' \ \boldsymbol{y}_t' \ \boldsymbol{w}_t')$.

Investors solve the problem from Section 2:

$$\max_{\{c_{it}, \boldsymbol{\theta}_{it}\}_{t\geq 0}} \ \inf_{\mathbb{Q}} \ \mathbb{E}^{\mathbb{Q}} \left[\rho \int_0^\infty e^{-\rho t} \log(c_{it}) dt + \int_0^\infty e^{-\rho t} \frac{\gamma_{it}}{1 - \gamma_{it}} dm_t \right]$$

Their choice of portfolio shares $\boldsymbol{\theta}_{it} = \{\theta_{ijt}\}_{j=1}^k$ corresponds to a choice of share holdings in risky assets $\boldsymbol{m}_{it} = \{m_{ijt}\}_{j=1}^k$ with $p_{jt} m_{ijt} = \theta_{ijt} w_{it}$. The risk-free bond holdings $b_{it} = (1 - \boldsymbol{\theta}_{it}' 1_k) w_{it}$ take up the complementary portfolio share.

Given initial holdings, an equilibrium is a collection of processes for prices $\{p_{jt}, r_t\}$, wealth $\{w_{it}\}$, and quantities $\{c_{it}, \boldsymbol{m}_{it}, b_{it}\}$ adapted to the filtration generated by $\{x_t\}_{t\geq 0}$ and satisfying the following conditions. First, quantities are chosen optimally by agents, who take prices as given. Second, the evolution of wealth is consistent with portfolio and consumption choices. Third, markets for all assets and consumption goods clear:

$$
\begin{aligned}
s_j &= \sum\nolimits_{i=1}^{n} m_{ijt} \text{ for all } t \geq 0 \text{ and all } j \in \{1, ...k\} \\
0 &= \sum\nolimits_{i=1}^{n} b_{it} \text{ for all } t \geq 0 \\
\sum\nolimits_{j=1}^{k} s_j y_{jt} &= \sum\nolimits_{i=1}^{n} c_{it} \text{ for all } t \geq 0
\end{aligned}
$$

I will characterize equilibrium prices $r(x_t)$ and $\boldsymbol{p}(x_t) = \{p_{jt}(x_t)\}$ as functions of aggregate states. These aggregate states generally include the wealth vector $\boldsymbol{w}_t = \{w_{it}\}$, and the epistemic assumption is that agents treat these wealth processes as exogenous, not realizing that their own wealth dynamics affect prices. This is what "taking prices as given" means.

Equilibrium characterization. The evolution of prices is

$$
d\boldsymbol{p}(x_t) = \mu_p(x_t)dt + \sigma_{p,y}(x_t)dZ_t^y + \sigma_{p,\gamma}(x_t)dZ_t^\gamma
$$

This defines the k-dimensional drift $\mu_p(x_t)$, the $(k \times d)$-dimensional matrix of loadings $\sigma_{p,y}(x_t)$, and the $(k \times g)$-dimensional matrix of loadings $\sigma_{p,\gamma}(x_t)$. I look for equilibria in which prices are diffusions. Given the drift and loadings of prices, the j-th component of excess returns is

$$
dR_{jt} = \frac{[\mu_p(x_t) + \boldsymbol{y}(x_t) - r(x_t)\boldsymbol{p}(x_t)]_j}{[\boldsymbol{p}(x_t)]_j}dt + \frac{1}{[\boldsymbol{p}(x_t)]_j}[\sigma_{p,y}(x_t)dZ_t^y]_j + \frac{1}{[\boldsymbol{p}(x_t)]_j}[\sigma_{p,\gamma}(x_t)dZ_t^\gamma]_j
$$

Let $\boldsymbol{s} = \{s_j\}$ be the vector of asset supply and denote the total wealth by $w_t = \boldsymbol{p}(x_t)'\boldsymbol{s}$. Since agents consume a constant multiple of their wealth, as shown in Section 2, total wealth is exogenous: $\rho w_t = \boldsymbol{y}(x_t)'\boldsymbol{s}$. Its evolution depends on $\boldsymbol{y}_t$ and can be written as

$$
\frac{dw_t}{w_t} \equiv \mu_w(\boldsymbol{y}_t)dt + \sigma_w(\boldsymbol{y}_t)dZ_t^y
$$

Two properties of wealth dynamics are important. First, the scalar drift $\mu_w(\boldsymbol{y}_t)$ and the $(1 \times d)$-dimensional row vector of loadings $\sigma_w(\boldsymbol{y}_t)$ directly map into the primitives of the environment. Endogenous dynamics are fully described by wealth shares of agents, making levels redundant. The dimensionality of the endogenous state space is $n - 1$. Second, aggregate wealth does not depend on $\boldsymbol{\gamma}_t$. This is a consequence of assuming a unitary elasticity of intertemporal substitution. Model revisions do not change aggregate wealth.

The final step toward characterizing asset prices is defining wealth shares $\nu_{it} \equiv w_{it}/w_t$ and the wealth-weighted average of robustness parameters Γ_t:

$$
\Gamma_t \equiv \sum\nolimits_{i=1}^{n} \gamma_{it}\nu_{it}
$$

The next proposition characterizes the dynamics of asset prices and the interest rate.

Proposition 3. *Risky asset prices and the interest rate satisfy*

$$r(x_t) = \rho + \mu_w(\boldsymbol{y}_t) - \frac{|\sigma_w(\boldsymbol{y}_t)|^2}{\Gamma_t} \tag{2}$$

$$r(x_t)\boldsymbol{p}(x_t) = \mu_p(x_t) + \boldsymbol{y}_t - \frac{\sigma_{p,y}(x_t)\sigma_w(\boldsymbol{y}_t)'}{\Gamma_t} \tag{3}$$

I explain solving this system in Appendix A.6. Equation (2) shows that the risk premium in this economy is equal to $|\sigma_w(\boldsymbol{y}_t)|^2/\Gamma_t$, which means that $1/\Gamma_t$, the only endogenous object in this equation, is the market price of risk. This variable acts as the market's effective risk aversion, and it equals the wealth-weighted harmonic average of individual effective risk-aversions $1/\gamma_{it}$.

Equation (3) is the main asset pricing equation. One part of the equation is present in pricing models without risk, where $r(x_t)\boldsymbol{p}(x_t) = \mu_p(x_t) + \boldsymbol{y}_t$. The risk-pricing content of equation (3) is the last term, which has two important properties.

First, risk adjustment in prices is given by the covariance of prices with the total wealth. Shocks only generate risk premia if they generate co-movements of prices with aggregate wealth. Importantly, since total wealth does not depend on $\boldsymbol{\gamma}_t$, price loadings on $\boldsymbol{\gamma}_t$ do not enter this expression. Uncertainty coming from model revision shocks is not priced in. This fact, however, does not imply that asset prices are independent of $\boldsymbol{\gamma}_t$. Shocks to any individual model do induce relative price changes with a constraint that $\boldsymbol{p}(x_t)'\boldsymbol{s}$ is constant.

Second, risk pricing only depends on the wealth-weighted average Γ_t of all robustness parameters. The effect of the entire wealth distribution on prices is summarized by this scalar, which acts as the market's risk tolerance coefficient in both equation (3) and equation (2).

With hedging motives, aggregation would be considerably harder, since it would have to account for the heterogeneous marginal value of wealth, which usually requires one extra partial differential equation per agent. The fact that only Γ_t enters equation (3) and equation (2) means that a form of Gorman aggregation obtains in the model. Asset prices would be the same in an economy with a single investor with a robustness parameter equal to Γ_t at all histories. The economy also admits a stochastic discount factor that only depends on Γ_t:

Proposition 4. *Asset prices satisfy*

$$\Lambda_t \, \boldsymbol{p}(x_t) = \mathbb{E}_t \int_t^{\infty} \Lambda_s \, \boldsymbol{y}_s \, ds$$

Here $\Lambda_0 = 1$ *and*

$$d\Lambda_t = -r(x_t)\Lambda_t dt - \frac{1}{\Gamma_t}\Lambda_t \sigma_w(\boldsymbol{y}_t) dZ_t^y$$

This discount factor is not necessarily unique, since markets do not have to be complete for any of the results in this model. Proposition 4 only provides one of possibly many factors. It can be interpreted as belonging to the same fictitious representative agent with a robustness parameter $\gamma_t = \Gamma_t$ at all histories. Notably, the evolution of Λ_t does not load on the shocks to $\boldsymbol{\gamma}_t$ itself, reflecting the fact that $\boldsymbol{\gamma}_t$ does not affect aggregate wealth.

Despite this reduction of dimensionality in local dynamics, Gorman aggregation does not imply that the economy collapses to a representative agent in the strong sense. Even though asset pricing is summarized by one scalar, the evolution of this scalar depends on the entire wealth distribution. Put another way, Γ_t is not a Markov process, and the minimal set of variables that have the Markov property is $x_t = (\boldsymbol{y}_t, \boldsymbol{\gamma}_t, \boldsymbol{\nu}_t)$, where $\boldsymbol{\nu}_t$ is the vector of all but one wealth shares. This allows the model to capture movements in risk premia resulting from wealth redistribution.

4.1 Identification

I now show what identifies the market price of risk and its underlying driving forces in the data. Specifically, the price of risk changes because of individual model revisions (changes in robustness parameters γ_{it}) and redistribution of wealth across agents with different models (changes in wealth shares ν_{it}). Both of these contributors to changes in Γ_t are potentially high-dimensional. However, aggregative properties of the model allow mapping their contributions to $d\Gamma_t$ into just a few observable time series.

As a first step, define aggregate market excess returns dR_t as

$$dR_t = \frac{1}{\sum_j s_j p_j(x_t)} \sum_j s_j p_j(x_t) dR_{jt}$$

Using Proposition 3,

$$dR_t = \frac{1}{\Gamma_t} |\sigma_w(\boldsymbol{y}_t)|^2 dt + \sigma_w(\boldsymbol{y}_t) dZ_t^y$$

Consider now a hypothetical investor with a primitive Hansen and Sargent (2001) multiplier Ψ_t that maps into the aggregate robustness parameter Γ_t:

$$\Psi_t = \frac{\Gamma_t}{1 - \Gamma_t}$$

Suppose this hypothetical investor has access to this one asset, the total market index. Let $dW_t \equiv dR_t - \mathbb{E}^{\mathbb{Q}_t}[dR_t]$ be the forecast error process the hypothetical investor chooses for this asset's return, where $\mathbb{Q}_t$ is the subjective probability measure corresponding to her optimally chosen drift correction. The next proposition characterizes it.

Proposition 5. *The forecast error process dW_t chosen by the robust investor is*

$$dW_t = \frac{1-\Gamma_t}{\Gamma_t}|\sigma_w(\boldsymbol{y}_t)|^2 dt + \sigma_w(\boldsymbol{y}_t) dZ_t^y$$

This is the first identification result. Under the physical measure $\mathbb{P}$, the error term dW_t has a predictable part, a drift correction, reflecting the robust investor's pessimistic choice of the model. This proposition implies that we can compare expected excess returns on the stock market in the objective and representative subjective models as

$$\mathbb{E}_t^{\mathbb{P}}[dR_t] = \frac{|\sigma_w(\boldsymbol{y}_t)|^2}{\Gamma_t} dt$$
$$\mathbb{E}_t^{\mathbb{Q}}[dR_t] = |\sigma_w(\boldsymbol{y}_t)|^2 dt$$

The ratio of the two expectations identifies Γ_t.

These two equations relate to empirical results of Nagel and Xu (2023). Subjective expectations of excess returns do vary over time and can be linked to the survey reports on the quantity of risk. Objective expectations vary much more and are substantially more cyclical: they get an additional multiplier $1/\Gamma_t$ that co-moves with the interest rate and asset prices. The reason is that agents evaluate risk in the same way within all of their models. Their required risk compensation is equal to the quantity of risk, and this variable does not change across their models. The object that does change across models is the gap between objective and subjective first moments of returns, and since subjective expectations must be stable, it is objective expectations that move in equilibrium.

I next explain what identifies the forces driving the evolution of Γ_t: changes in subjective models and wealth redistribution.

Evolution of the wealth distribution. It turns out that in this economy, the hypothetical forecast error process dW_t drives wealth dynamics and fully describes the evolution of individual wealth shares conditional on subjective models. Define the leverage of investor i as $\lambda_{it} = \boldsymbol{\theta}'_{it} 1_k$. The next proposition relates the dynamics of wealth to dW_t.

Proposition 6. *In equilibrium, total leverage of each agent i is $\lambda_{it} = \gamma_{it}/\Gamma_t$. Individual holdings of risky assets are $m_{ijt} = s_j \nu_{it} \lambda_{it}$. Wealth shares $\{\nu_{it}\}$ follow*

$$\frac{d\nu_{it}}{\nu_{it}} = (\lambda_{it} - 1) dW_t \tag{4}$$

In equilibrium, agent i's leverage is given by a simple expression γ_{it}/Γ_t. It can be rewritten as

$$\lambda_{it} = 1 + (1 - \nu_{it}) \frac{\gamma_{it} - \gamma_{-it}}{\Gamma_t}$$

Here γ_{-it} is the wealth-weighted average of all robustness parameters except for i's. Excess leverage is positive if $\gamma_{it} > \gamma_{-it}$, which is equivalent to $\gamma_{it} > \Gamma_t$. Investors whose model is less pessimistic than the market's representative model always borrow in the risk-free asset from the rest. A decrease in γ_{it} makes agent i's subjective model more pessimistic, and she steps back from taking risk, while other agents pick it up. This relates to the findings of Barbiero et al. (2024), who document that risk limit tightenings at individual dealers lead to other dealers taking over and raising their risk exposures. Excess leverage also falls in equilibrium if ν_{it} approaches one because all other agents become small and there is nobody to borrow from.

The most important result in Proposition 6 is that the dynamics of ν_{it} do not load on dZ_t^{γ}, since the forecast error dW_t does not. Shocks to robustness parameters, and hence model revisions, do not induce wealth redistribution. The reason is the following. The share of asset j held by investor i is $m_{ijt}/s_j = \nu_{it}\lambda_{it}$. This does not depend on j: investors have the same portfolios of risky assets up to scale, only differing in the risky-riskless split. This implies that the risky part of everyone's portfolio is a scaled copy of the aggregate market portfolio: each investor assigns the same relative weights to all risky assets as the entire market. Hence, any investor's wealth only changes when aggregate wealth does, and this does not happen since model revision shocks only reshuffle asset prices without affecting the total.

Unsurprisingly, equation (4) shows that, if $\lambda_{it} > 1$, ν_{it} is positively exposed to $\sigma_w(\boldsymbol{y}_t)dZ_t^y$ through the forecast error dW_t. Wealth shares of levered investors are positively correlated with aggregate wealth. The new part in equation (4) compared to the log utility case is that there is a drift in the wealth shares. This drift is positive for levered investors, reflecting the deterministic risk compensation they receive from the rest of the market. Levered investors exploit the market's robustness concerns, or risk limits, and get ahead of others by pocketing the risk premium. With $\Gamma_t = 1$, or $\Psi_t = \infty$, this drift in wealth shares disappears.

Armed with Proposition 6, I can characterize the law of motion for Γ_t.

PROPOSITION 7. *The evolution of Γ_t is*

$$d\Gamma_t = \frac{\Delta_t}{\Gamma_t} dW_t + \boldsymbol{\nu}_t' \, d\,\boldsymbol{\gamma}_t \tag{5}$$

Here Δ_t is the wealth-weighted dispersion of the multipliers:

$$\Delta_t = \sum\nolimits_{i=1}^{n} \nu_{it}\gamma_{it}^2 - \left(\sum\nolimits_{i=1}^{n} \nu_{it}\gamma_{it}\right)^2 \geq 0$$

This is the second identification result. It shows that the redistribution-induced change in the aggregate risk tolerance is driven by the subjective forecast error, with the coefficient of proportionality depending on the past Γ_t. The same applies to the market price of risk, which equals $1/\Gamma_t$, a monotone transformation. The remaining changes in the market price of risk come from model revisions and are conditionally orthogonal to the redistribution-induced part, since the forecast error dW_t is conditionally orthogonal to $d\boldsymbol{\gamma}_t$ according to Proposition 5 and the fact that dZ_t^y and dZ_t^γ are independent.

The impact of redistribution is the first part of equation (5). Aggregate shocks induce wealth redistribution if investors choose heterogeneous leverage, taking different bets on the market. For this to happen, they must have heterogeneous subjective models, or heterogeneous robustness parameters γ_{it}, implying $\Delta_t > 0$. In this case, aggregate risk tolerance Γ_t is positively exposed to shocks to aggregate wealth $\sigma_w(\boldsymbol{y}_t)dZ_t^y$ through the forecast error term dW_t. This is because investors with higher robustness parameters γ_{it}, which correspond to less pessimistic subjective models and higher leverage, become relatively richer after positive shocks, driving up the weighted average. There is a positive drift in the redistributive part of $d\Gamma_t$ as well, reflecting the upward drift in the wealth shares of less constrained investors due to the risk compensation they collect.

The second part of equation (5) is individual model revisions that follow changes in robustness parameters. This part could be non-zero even if everyone had the same γ_{it}, and wealth shares were fixed forever. In this special case, risk premia would still be volatile, but the weighted average Γ_t would only move exogenously. This could be desirable in applications because it would reduce the dimensionality of the problem. Appendix C presents an example economy with one agent operating under a time-varying subjective model. This setup produces a direct link between the agent's robustness parameter and the interest rate. The resulting stochastic process for the interest rate is available in closed form. I show examples of robustness parameter processes that lead to Vasicek (1977) and Cox, Ingersoll, and Ross (1985) models, popular specifications of exogenous interest rate dynamics that have been used extensively in the macro-finance literature.

Conversely, redistribution-induced movements in the market price of risk happen even if all subjective models are fixed. Appendix C presents an analytical example with one risky asset and two agents who hold different fixed subjective models. Positive output shocks redistribute wealth to the less pessimistic investor, which decreases the risk premium and raises the interest rate in the economy. The processes for the interest rate and risk premium are again available in closed form. This special case links to Caballero and Simsek (2020) and Kekre and Lenel (2022), who embed wealth redistribution into New-Keynesian models with risky assets and show how redistributing wealth between agents with different beliefs or risk-taking capacity affects aggregate output through risk premia. My framework illustrates this chain in detail.

5 Estimation

I now estimate the model. Specifically, I recover time series for Γ_t and the two components of its growth characterized in Proposition 7: $d\Gamma_t = d\kappa_{1t} + d\kappa_{2t}$. Here $d\kappa_{1t}$ is the redistribution-induced change and $d\kappa_{2t}$ is the average model revision: $d\kappa_{1t} \equiv \Delta_t/\Gamma_t \cdot dW_t$ and $d\kappa_{2t} \equiv \boldsymbol{\nu}_t' \, d\boldsymbol{\gamma}_t$. I first explain the empirical strategy, then describe my data, and then discuss the results.

Empirical strategy. I conduct analysis at the quarterly frequency using three observable series: f_t^{obj}, f_t^{sub}, and rx_{t+4}. Here f_t^{obj} is the econometrician's expectation of excess returns on the US stock market over a one-year horizon, f_t^{sub} is the subjective expectation, and rx_{t+4} is the realized yearly excess return one year ahead. I use the following three equations:

$$\log(f_t^{\text{sub}}) - \log(f_t^{\text{obj}}) = \log(\Gamma_t) + \alpha + \sigma_1 \epsilon_{1t} \tag{6}$$

$$rx_{t+4} - f_t^{\text{sub}} = \frac{1-\Gamma_t}{\Gamma_t} f_t^{\text{sub}} + \sigma_2 \epsilon_{2t} \tag{7}$$

$$\Gamma_{t+1} - \Gamma_t = (1-\Gamma_t)\frac{rx_{t+4} - f_t^{\text{sub}}}{4} + \mu + \sigma_3 \epsilon_{3t} \tag{8}$$

Equation (6) uses the fact that $\Gamma_t = \mathbb{E}_t^{\mathbb{Q}}[dR_t]/\mathbb{E}_t^{\mathbb{P}}[dR_t]$, which is derived from Proposition 5. I assume that the econometrician's estimate of the expected excess return is noisy, and this noise is $\alpha + \sigma_1 \epsilon_{1t}$, where α is a constant and ϵ_{1t} is standard normal drawn independently across periods and conditionally independent of all other variables.

Equation (7) uses Proposition 5, which states that the drift of the subjective forecast error is equal to $(1-\Gamma_t)/\Gamma_t \cdot |\sigma_w(\boldsymbol{y}_t)|^2$, where I also use the fact that $f_t^{\text{sub}} \equiv \mathbb{E}_t^{\mathbb{Q}}[dR_t] = |\sigma_w(\boldsymbol{y}_t)|^2$. Aggregation to discrete time here treats one-year expected returns as a scaled counterpart of the instantaneous conditional expectation. I assume that the rest of the subjective forecast error is normally distributed, with a standard normal ϵ_{2t} drawn independently across time and conditionally independent of all other variables.

Finally, equation (8) is a quarterly approximation of Proposition 7. For simplicity, I approximate the variance of $\{\gamma_{it}\}$ by $\Delta_t = \Gamma_t(1-\Gamma_t)$. The redistribution part is proportional to the forecast error:

$$\Delta\kappa_{1,t+1} = (1-\Gamma_t)\frac{rx_{t+4} - f_t^{\text{sub}}}{4}$$

Here I allocate the realized one-year forecast error equally across the four quarters. The remainder of the right-hand side is the model revision part:

$$\Delta\kappa_{2,t+1} = \mu + \sigma_3 \epsilon_{3t}$$

Here I assume that μ and σ_3 are constants, and ϵ_{3t} is standard normal drawn independently across time which I treat as conditionally independent of all other variables.[1]

I use quasi maximum likelihood estimation, optimizing over $\{\Gamma_0, \alpha, \mu, \sigma_1, \sigma_2, \sigma_3\}$ with Γ_t as a hidden state that is confined to $(0, 1]$. On rare occasions, outside of the baseline specification, where it crosses the upper boundary, the algorithm clips it at 1. My observation equations are equation (6) and equation (7), while equation (8) is the transition equation. I linearize the measurement equations to use the Kalman filter and smooth the estimates with a backward recursion from Rauch, Tung, and Striebel (1965).

Data. I take data from the replication package of Nagel and Xu (2023). As f_t^{sub}, I use their data on expected year-ahead US stock market returns of individual investors, which combine the UBS/Gallup survey, the Conference Board survey, and the Michigan Survey of Consumers. Following Nagel and Xu (2023), I construct expected excess returns by subtracting the one-year Treasury yield. These series are quarterly. I take 137 consecutive quarters from 1987Q1 to 2021Q1. I follow the timing convention of Nagel and Xu (2023): the respondents see price data with a one-month lag and macro data with a two-month lag. The series for f_t^{sub} enters both observation equations and the transition.

My measure rx_{t+4} is the realized yearly return four quarters ahead on the CRSP value-weighted index in excess of the one-year Treasury yield. The last observation on this series is dated 2019Q4, which is where my decomposition of Γ_t ends, while the recovered Γ_t itself continues to 2021Q1. The realized return series enters one of the observation equations and the transition equation.

I follow Nagel and Xu (2023) in constructing the econometrician's expectation f_t^{obj}. The exercise uses the following nine predictors: consumption-wealth ratio from Lettau and Ludvigson (2001), log of the dividend-price ratio adjusted for repurchases, experienced payout growth constructed with weights from Nagel and Xu (2022), net equity issuance, negative of industrial production growth, the term spread, the default spread, the lagged real activity factor of Ludvigson and Ng (2009), and the squared VIX. As the headline specification, I use an equally-weighted combination of single-variable forecasts estimated in expanding windows. For each of the predictors, this procedure forms an expanding-window forecast that only uses past data (periods for which the year-ahead return is already realized) for estimation and then takes the average across the nine.

I additionally report results obtained from a version of these expanding-window forecasts shrunk towards the prevailing historical mean of realized one-year returns, following Nagel and Xu (2023), and for 18 in-sample forecasts: predictions formed by using all available data, including from the future, in nine regressions with individual predictors and trailing one-year excess returns. These nine forecasts are computed in two variants: over the entire sample when the data are available

[1] Conditional independence of shocks $\{\epsilon_{1t}, \epsilon_{2t}, \epsilon_{3t}\}$ is an approximation: the use of quarterly data and year-ahead forecasts creates an overlap in observation windows, so these residuals could be correlated with each other.

and over the subsample where the subjective forecast is also present. Coefficient estimates in these regressions are bias-adjusted following Amihud and Hurvich (2004) and Amihud, Hurvich, and Wang (2009). The econometrician's expectation series f_t^{obj} enters one of the observation equations: equation (6). I treat quarters when expected excess returns go negative as missing observations for this equation. This never happens in my headline specification and only happens a few quarters with some of the in-sample predictions.

Finally, I use the intermediary capital ratio series from He, Kelly, and Manela (2017) for model validation. This series never enters estimation and only feeds into untargeted moments.

Results. I start by reporting the recovered path of Γ_t and the two driving forces. Figure 1 reports the three observable series in the top panel and the recovered changes in Γ_t from the start of the sample with the two contributing parts in the bottom panel.

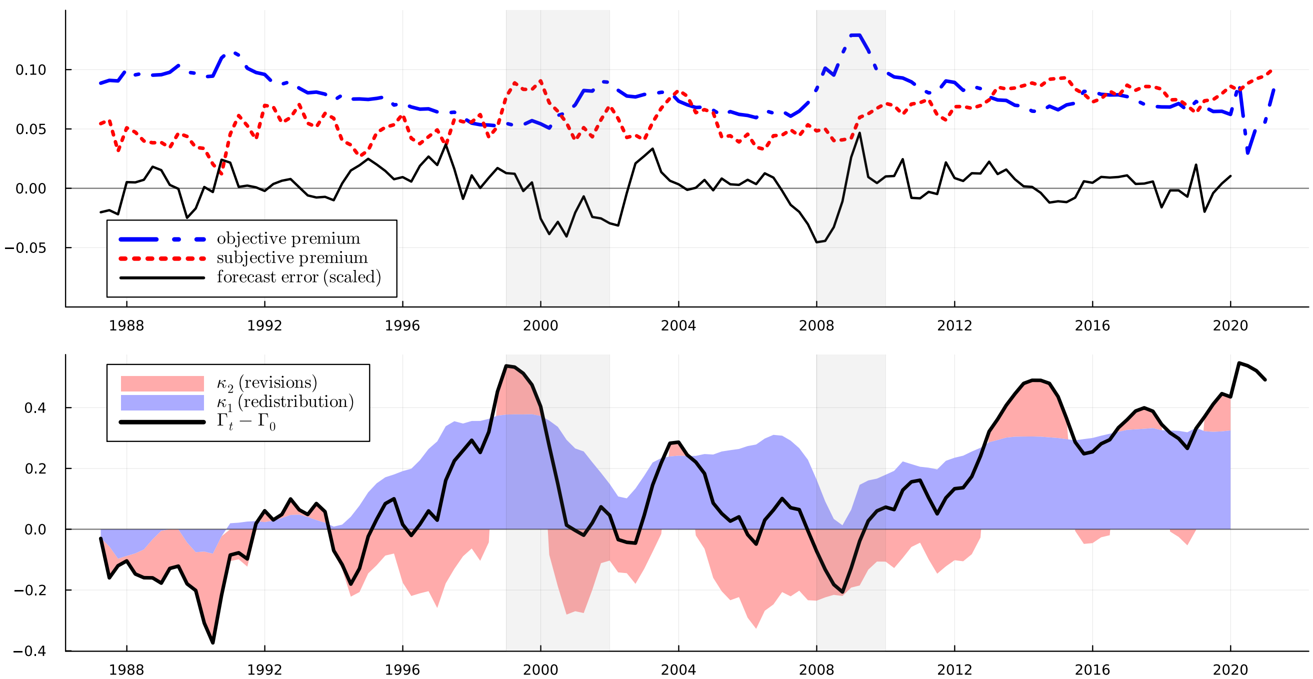

Figure 1: Top panel: objective premium f_t^{obj}, subjective premium f_t^{sub}, and the subjective forecast error $rx_{t+4} - f_t^{\text{sub}}$. The forecast error is scaled to have the same standard deviation as the objective premium. Bottom panel: changes in Γ_t relative to Γ_0 and cumulated $\Delta\kappa_{1t}$ and $\Delta\kappa_{2t}$.

I highlight two episodes on Figure 1: the Dot-com crisis around 2000 and the Global Financial Crisis in 2008. Both of these episodes feature a steep decline in the recovered Γ_t, which maps into a steep increase in the price of risk. These declines, however, have different causes according to my filtering procedure.

The run-up to the Dot-com crisis looks like a large optimistic revision of underlying subjective models. The feature of the data that pushes this interpretation is the fact that subjective expectations of stock market returns overtake the objective premium. Since my framework cannot incorporate subjective premia being higher than objective ones, this episode converts the data into a recovered Γ_t approaching 1 from below. At the turn of the year 2000, the ordering of subjective and objective premia reverts, which pushes the recovered Γ_t into a sharp decrease. At the same time, the negative forecast error right before the reversion did not redistribute wealth towards more pessimistic agents: Γ_t was close to one, and the impact of forecast error on redistribution is proportional to $1 - \Gamma_t$. The episode is hence attributed to model revisions.

In contrast, the decline in Γ_t during the crisis of 2008 is attributed to redistribution away from investors with less pessimistic models. The feature of the data that points to a low Γ_t, and hence a high price of risk, at the height of the crisis is the fact that the objective premium is much higher than the subjective one. The feature of the data that points to redistribution as the way the economy got to a low Γ_t is a large negative forecast error right before. Optimistic agents had taken on large quantities of risk, and the realized negative returns redistributed wealth toward more pessimistic investors, raising the market price of risk.

The variance decomposition of changes in Γ_t indicates that 71% of the variance at the one-quarter horizon is due to model revisions, while 15% is due to redistribution. The rest is attributed to the sample correlation between the two terms: they should be conditionally uncorrelated in continuous time but show non-zero correlation after time aggregation to quarters. The variance share of revisions falls to 64% at four quarters and 61% at eight. That of redistribution increases to 20% at four quarters and 22% at eight.

As a way to externally validate the model, I compare the recovered redistribution dynamics to the dynamics of financial intermediary capital ratios from He, Kelly, and Manela (2017). This ratio is defined as the market value of equity divided by the sum of the market value of equity and book debt. It tracks capitalization of the financial intermediary sector, the agents with high risk-taking capacity that map into investors with optimistic models in my framework.

Figure 2 shows the two series, the redistribution-induced part of changes in Γ_t, cumulated to year-over-year changes, and the year-over-year change in the intermediary capital ratio. I remove a local-linear trend estimated in two-sided 40-quarter windows from the capital ratio. The right panel plots the rolling-window correlation between the two series. The correlation oscillates around 0.7, indicating a stable relationship. What the filter reads as redistribution towards more pessimistic agents coincides with declines in capitalization of financial intermediaries.

This result allows me to relate my framework to a large class of models with financial intermediaries, or specialists, where the net worth or wealth share of these agents is the key state variable determining asset prices. This literature dates back at least to Xiong (2001) and started to develop explosively after seminal contributions of He and Krishnamurthy (2012), He and Krishnamurthy (2013), and Brunnermeier and Sannikov (2014). In my framework, wealth redistribution does not

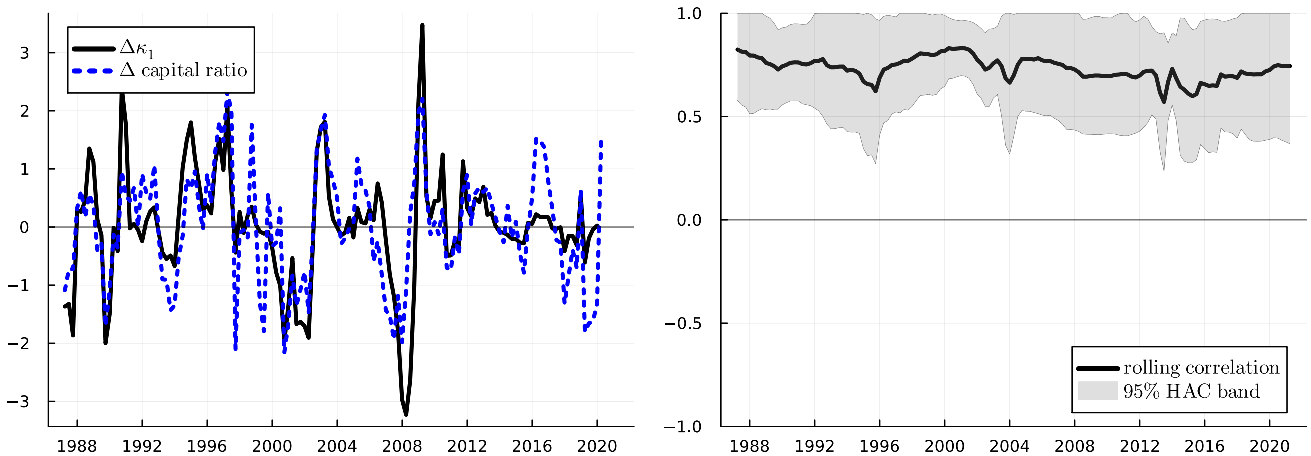

Figure 2: Left panel: $\Delta\kappa_1 \equiv \kappa_{1,t+4} - \kappa_{1,t}$, the recovered redistribution-driven cumulated one-year changes in Γ_t, and the detrended year-over-year changes in the intermediary capital ratio from He, Kelly, and Manela (2017), both standardized. Right panel: rolling-window correlation between them computed in centered 40-quarter windows, with 95% HAC-robust confidence bands.

fully characterize the movements in risk premia: it is accompanied by model revisions, an equally important channel according to the empirical results above. However, only asset price movements that result from wealth redistribution are priced in: model revision shocks are orthogonal to shocks to the total wealth and do not command a risk premium. Redistribution shocks are growth shocks, and hence require a premium, which is consistent with the finding of He, Kelly, and Manela (2017) that movements in the intermediary capital ratio command a positive risk price.

Alternative measures. I now study how alternative ways of constructing the econometrician's expectation f_t^{obj} affect the results.

Table 1 compares alternative specifications to the baseline. The "median shrunk" specification refers to the median of the real-time forecasts of stock returns that apply shrinkage toward the historical mean, following Nagel and Xu (2023). Of the 9 predictors, 6 survive the shrinkage procedure with substantial coefficients. The estimated shrinkage weight on industrial production growth and the default spread is zero, and that on VIX squared is very to close to zero. When the weight is zero, the forecast is the prevailing historical mean alone. The third row reports this "no predictability" benchmark separately. The third row shows that the results do not change substantially even if the econometrician does not see any predictability in excess returns. The "in-sample" specifications return 18 estimates: 9 predictive regressions that each use one of the 9 predictors described above along with trailing one-year returns, each run over 2 samples, the full time window when the data are available and the window of the main analysis. In each column, I present the median of the 18 numbers and the two quartiles.

The first six columns of Table 1 show the correlations between series recovered in the baseline

Table 1: Alternative objective-premium specifications

Specification	Γ_t	$\Delta\Gamma_t$	κ_{1t}	$\Delta\kappa_{1t}$	κ_{2t}	$\Delta\kappa_{2t}$	corr($\Delta\kappa_1, \Delta$ratio)
baseline	—	—	—	—	—	—	0.70
median shrunk	0.81	0.76	0.97	0.97	0.66	0.76	0.71
no predictability	0.86	0.82	0.97	0.97	0.72	0.82	0.71
in-sample median	0.78	0.71	0.97	0.97	0.56	0.60	0.69
in-sample Q25	0.68	0.64	0.93	0.95	0.37	0.50	0.67
in-sample Q75	0.85	0.84	0.99	0.98	0.73	0.81	0.70

Notes: "median shrunk" is the column-wise median across the nine single-predictor real-time shrinkage specifications. For two predictors, the estimated shrinkage weight is zero: their forecasts coincide with the prevailing historical mean of realized one-year returns. The row "no predictability" reports the specification with this prevailing mean as the forecast. The three "in-sample" rows correspond to the eighteen in-sample specifications (nine predictors, two sample conventions). Columns 1 to 6 report correlations of the recovered $\{\Gamma_t, \Delta\Gamma_t, \kappa_{1t}, \Delta\kappa_{1t}, \kappa_{2t}, \Delta\kappa_{2t}\}$ with those in the baseline specification. Column 7 reports correlations of the four-quarter redistribution flow with the four-quarter change in the He, Kelly, and Manela (2017) intermediary capital ratio. Quantiles are computed column by column.

and their counterparts in the alternative models. The correlations are generally above 50% for both levels and changes. The last column shows time-series correlations I compute within models as validation. The correlation between $\Delta\kappa_{1t}$, which captures redistribution-induced changes in Γ_t, and the intermediary capital ratio from He, Kelly, and Manela (2017) stays high across models. The results from these alternative specifications are not drastically different from the baseline.

6 Conclusion

I develop a model of dynamic subjective beliefs that admits easy aggregation in general equilibrium. Subjective models follow from robustness concerns. Shocks to investors' robustness preferences lead to subjective model revisions. Crucially, they only revise models of returns, not aggregate states. This induces simple mean-variance portfolios with time-varying effective risk tolerance. These portfolios then allow for aggregation in economies with heterogeneous subjective models. In equilibrium, a representative subjective model determines asset prices, while the forecast error under that model determines the dynamics of the wealth distribution.

The aggregate risk tolerance changes for two reasons: subjective model revisions and wealth redistribution between investors with different models. I provide an identification result that separates the two. Redistribution-induced changes are proportional to the forecast error under the representative model, while revision-induced changes are conditionally orthogonal to it. Estimating the model on the survey data, I find that model revisions contribute three quarters of the variation in risk premia on the one-quarter horizon. The redistribution-induced changes in the market price of risk are closely related to the capital of financial intermediaries.

References

Adrian, T. and N. Boyarchenko (2018). Liquidity policies and systemic risk. Journal of Financial Intermediation 35, 45–60.

Adrian, T. and H. S. Shin (2010). Liquidity and leverage. Journal of financial intermediation 19(3), 418–437.

Adrian, T. and H. S. Shin (2014). Procyclical leverage and value-at-risk. The Review of Financial Studies 27(2), 373–403.

Amihud, Y. and C. M. Hurvich (2004). Predictive regressions: A reduced-bias estimation method. Journal of Financial and Quantitative Analysis 39(4), 813–841.

Amihud, Y., C. M. Hurvich, and Y. Wang (2009). Multiple-predictor regressions: Hypothesis testing. Review of Financial Studies 22(1), 413–434.

Barbiero, O., F. Bräuning, G. Joaquim, and H. Stein (2024). Dealer risk limits and currency returns. Available at SSRN.

Basak, S. and A. Shapiro (2001). Value-at-risk-based risk management: optimal policies and asset prices. The review of financial studies 14(2), 371–405.

Bhandari, A., J. Borovička, and P. Ho (2025). Survey data and subjective beliefs in business cycle models. Review of Economic Studies 92(3), 1375–1437.

Bräuning, F. and H. Stein (2024). The effect of primary dealer constraints on intermediation in the treasury market. Available at SSRN 4862212.

Brunnermeier, M. K. and Y. Sannikov (2014). A macroeconomic model with a financial sector. American Economic Review 104(2), 379–421.

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

Coimbra, N., D. Kim, and H. Rey (2022). Central bank policy and the concentration of risk: Empirical estimates. Journal of Monetary Economics 125, 182–198.

Coimbra, N. and H. Rey (2024). Financial cycles with heterogeneous intermediaries. Review of Economic Studies 91(2), 817–857.

Cox, J. C., J. E. Ingersoll, and S. A. Ross (1985). A theory of the term structure of interest rates. Econometrica 53(2), 385–407.

Danielsson, J., H. S. Shin, and J.-P. Zigrand (2011). Balance sheet capacity and endogenous risk.

Danielsson, J., H. S. Shin, and J.-P. Zigrand (2012). Endogenous and systemic risk. In Quantifying systemic risk, pp. 73–94. University of Chicago Press.

Di Tella, S., C. Malgieri, and C. Tonetti (2024). Risk markups.

Duffie, D. and L. G. Epstein (1992). Stochastic differential utility. Econometrica: Journal of the Econometric Society, 353–394.

Gabaix, X. and M. Maggiori (2015). International liquidity and exchange rate dynamics. The Quarterly Journal of Economics 130(3), 1369–1420.

Gourinchas, P.-O., W. Ray, and D. Vayanos (2022). A preferred-habitat model of term premia, exchange rates, and monetary policy spillovers. Technical report, National Bureau of Economic Research.

Greenwood, R., S. Hanson, J. C. Stein, and A. Sunderam (2023). A quantity-driven theory of term premia and exchange rates. The Quarterly Journal of Economics 138(4), 2327–2389.

Hansen, L. P. and T. J. Sargent (2001). Robust control and model uncertainty. American Economic Review 91(2), 60–66.

Hansen, L. P., B. Szőke, L. S. Han, and T. J. Sargent (2020). Twisted probabilities, uncertainty, and prices. Journal of Econometrics 216(1), 151–174.

He, Z., B. Kelly, and A. Manela (2017). Intermediary asset pricing: New evidence from many asset classes. Journal of Financial Economics 126(1), 1–35.

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

He, Z. and A. Krishnamurthy (2013). Intermediary asset pricing. American Economic Review 103(2), 732–770.

Hofmann, B., I. Shim, and H. S. Shin (2022). Risk capacity, portfolio choice and exchange rates. Available at SSRN 4028446.

Holmstrom, B. and J. Tirole (1997). Financial intermediation, loanable funds, and the real sector. the Quarterly Journal of economics 112(3), 663–691.

Itskhoki, O. and D. Mukhin (2021). Exchange rate disconnect in general equilibrium. Journal of Political Economy 129(8), 2183–2232.

Kamdar, R. and W. Ray (2024). Optimal macro-financial stabilization in a new keynesian preferred habitat model.

Kekre, R. and M. Lenel (2022). Monetary policy, redistribution, and risk premia. Econometrica 90(5), 2249–2282.

Kekre, R. and M. Lenel (2024). Exchange rates, natural rates, and the price of risk. University of Chicago, Becker Friedman Institute for Economics Working Paper (2024-114).

Kekre, R. and M. Lenel (2025). A model of us monetary policy and the global financial cycle.

Kreps, D. M. and E. L. Porteus (1978). Temporal resolution of uncertainty and dynamic choice theory. Econometrica: journal of the Econometric Society, 185–200.

Lettau, M. and S. Ludvigson (2001). Consumption, aggregate wealth, and expected stock returns. the Journal of Finance 56(3), 815–849.

Ludvigson, S. C. and S. Ng (2009). Macro factors in bond risk premia. The Review of Financial Studies 22(12), 5027–5067.

Maenhout, P. J., A. Vedolin, and H. Xing (2025). Robustness and dynamic sentiment. Journal of Financial Economics 163, 103953.

Nagel, S. and Z. Xu (2022). Asset pricing with fading memory. Review of Financial Studies 35(5), 2190–2245.

Nagel, S. and Z. Xu (2023). Dynamics of subjective risk premia. Journal of Financial Economics 150(2), 103713.

Oskolkov, A. (2026). A heterogeneous-country model of the global financial cycle.

Piazzesi, M., M. Schneider, and J. Salomao (2015). Trend and cycle in bond premia.

Rauch, H. E., F. Tung, and C. T. Striebel (1965). Maximum likelihood estimates of linear dynamic systems. AIAA Journal 3(8), 1445–1450.

Ray, W. (2019). Monetary policy and the limits to arbitrage: Insights from a new keynesian preferred habitat model. In 2019 Meeting Papers, Volume 692. Society for Economic Dynamics.

Sizova, E. (2023). Banks' next top model. Available at SSRN.

Szőke, B. (2022). Estimating robustness. Journal of Economic Theory 199, 105225.

Vasicek, O. (1977). An equilibrium characterization of the term structure. Journal of financial economics 5(2), 177–188.

Vayanos, D. and J.-L. Vila (2021). A preferred-habitat model of the term structure of interest rates. Econometrica 89(1), 77–112.

Xiong, W. (2001). Convergence trading with wealth effects: an amplification mechanism in financial markets. Journal of Financial Economics 62(2), 247–292.

A Additional details

A.1 Incorporating external income with model revisions

Assuming that investors only have asset income and consumption expenditures is restrictive. The simple form of consumption and portfolio choice in Proposition 1 extends to some cases when a part of income is not chosen by the agent. Specifically,

$$dw_t = (r(x_t)w_t - c_t)dt + w_t\theta_t' dR_t - w_t\varsigma(x_t)dt - w_t\tau(x_t)'dB_t$$

Here $w_t\varsigma(x_t)dt$ and $w_t\tau(x_t)'dB_t$ are new terms that the agent does not choose. They can represent taxes, with $\varsigma(x_t)$ being a locally deterministic tax and $\tau(x_t)$ being a vector of taxes that load on the exogenous shocks. These taxes can be useful for inducing stationarity in the model through wealth redistribution: they can prevent the least constrained agent from taking over the entire economy. Taxes are lump-sum in the sense that the agent does not directly choose the tax base, which would be the case, for instance, with proportional taxes on portfolio returns. They still scale with wealth, however, and this feature is key to preserving consumption and portfolio choice.

The locally deterministic tax $\varsigma(x_t)$ clearly does not change anything in the agent's decision problem as it is simply isomorphic to a change in the interest rate. The stochastic tax in general affects portfolio choice. For a clearer characterization, I focus on a special case

$$\tau(x_t) = \zeta(x_t) \cdot \sigma_R(x_t)'[\sigma_R(x_t)\sigma_R(x_t)']^{-1}\mu_R(x_t) \tag{A.1}$$

If taxes are set up this way, they generate the same exposure to shocks as the optimal portfolio:

$$\tau(x_t)'dB_t = \frac{\zeta(x_t)(\psi_t + 1)}{\psi_t}\theta(w_t, x_t)'\sigma_R(x_t)dZ_t \equiv \frac{\zeta(x_t)(\psi_t + 1)}{\psi_t}\theta(w_t, x_t)'(dR_t - \mu_R(x_t)dt)$$

Here $\theta(w_t, x_t)$ is the optimal vector of portfolio weights in the baseline model. This allows a potential government to effectively tax away a share $\zeta(x_t)$ of stochastic returns without imposing proportional taxes. I next show the effect of imposing these taxes on the agent.

PROPOSITION 8. *Suppose the stochastic tax rate is given by equation (A.1). Then, consumption choice is* $c(w_t, x_t) = \rho w_t$. *The optimal portfolio is*

$$\theta(w_t, x_t) = \frac{\psi_t(1 + \zeta(x_t))}{1 + \psi_t} \cdot [\sigma_R(x_t)\sigma_R(x_t)']^{-1}\mu_R(x_t)$$

If taxes decrease the agent's exposure to shocks, she is willing to take more risk and increases her risky asset holdings.

A.2 Relation to standard robustness concerns

I now describe the version of my setup with standard robustness concerns of Hansen and Sargent (2001). The recursive representation of the robust problem when both return and aggregate state processes are misspecified is

$$\begin{aligned}\rho V(w,x) = \max_{c,\theta}\ \min_{h}\ &\rho\log(c) + \frac{\psi|h|^2}{2} + (r(x)w - c + w\theta'\mu_R(x) - w\theta'\sigma_R(x)h)V_w(w,x)\\ &+ \frac{\theta'\sigma_R(x)\sigma_R(x)'\theta}{2}w^2V_{ww}(w,x) + \frac{1}{2}\mathrm{tr}[\sigma_X(x)'V_{xx'}(w,x)\sigma_X(x)] + w\theta'\sigma_R(x)\sigma_X(x)'V_{wx'}(w,x)\\ &+ (\mu_X(x)' \underbrace{-h'\sigma_X(x)'}_{\text{new}})V_{x'}(w,x)\end{aligned}$$

The solution $V(w,x) = \log(w) + \eta(x)$ can be guessed and verified. Given this, $c = \rho w$ and

$$\begin{aligned}h &= \frac{1}{\psi}\sigma_X(x)'\eta_{x'}(x) + \frac{1}{\psi}\sigma_R(x)'\theta\\ \theta &= [\sigma_R(x)\sigma_R(x)']^{-1}(\mu_R(x) - \sigma_R(x)h)\end{aligned}$$

The fact that consumption is linear in w and (θ, h) only depend on x verifies the conjecture that $V(w,x)$ is log-separable over w and x. The optimal portfolio θ is given by

$$\theta = \frac{\psi}{1+\psi}[\sigma_R(x)\sigma_R(x)']^{-1}\mu_R(x) - \underbrace{\frac{1}{(1+\psi)}[\sigma_R(x)\sigma_R(x)']^{-1}\sigma_R(x)\sigma_X(x)'\eta_{x'}(x)}_{\text{hedging motive}} \tag{A.2}$$

This expression exactly coincides with the portfolio choice from the setup with recursive utility of Duffie and Epstein (1992), given below by equation (A.5), if $\gamma = \psi/(1+\psi)$.

A.3 Heuristic explanation for value-at-risk

Why is equation (1) a value-at-risk constraint if it restricts variance? The answer relies on the Brownian nature of the shocks. A generic value-at-risk constraint caps the probability of a given level of losses at a certain level. For example, given constants L and α, the constraint is

$$\mathbb{P}\left\{\theta_t' dR_t \le -\sqrt{Ldt}\right\} \le \alpha \tag{A.3}$$

Choosing Ldt instead of $\sqrt{Ldt}$ in this expression would not work. Heuristically, the variance of Brownian shocks is proportional to dt, so the standard deviation of $\theta_t' dR_t$ is of order $\sqrt{dt}$. Hence, as $dt \longrightarrow 0$, the probability on the left of equation (A.3) would just converge to 0 or 1 depending on how $\mathbb{E}[\theta_t' dR_t] = \theta_t' \mu_R(x_t) dt$ compares to Ldt. With $\sqrt{Ldt}$ instead, equation (A.3) becomes

$$\Phi\left(-\frac{\sqrt{Ldt} + \theta_t' \mu_R(x_t) dt}{\sqrt{\theta_t' \sigma_R(x_t) \sigma_R(x_t)' \theta_t dt}}\right) \leq \alpha$$

Here $\Phi(\cdot)$ is the cumulative distribution function of the standard normal distribution. As $dt \longrightarrow 0$, the $O(dt)$ term in the numerator vanishes. For all $\alpha < 1/2$, the limit of this inequality becomes

$$\theta_t' \sigma_R(x_t) \sigma_R(x_t)' \theta_t \leq L \cdot (\Phi^{-1}(\alpha))^{-2}$$

One choice of L is a proportion of net worth w_t, capping the probability of losing a fraction of capital. In continuous time, flow losses are not commensurate with net worth, which is a stock. No amount of instantaneous Brownian losses will make a dent in the intermediary's capital. Instead, one could choose another flow, such as a multiple of expected profits, $L = \hat{\gamma}_t \theta_t' \mu_R(x_t)$, to make a flow-to-flow comparison. With $\gamma_t = \hat{\gamma}_t / (\Phi^{-1}(\alpha))^2$, the constraint takes its final form:

$$\theta_t' \sigma_R(x_t) \sigma_R(x_t)' \theta_t \leq \gamma_t \theta_t' \mu_R(x_t)$$

The last step is realizing that $\mathbb{V}[\theta_t' dR_t] = \theta_t' \sigma_R(x_t) \sigma_R(x_t)' \theta_t dt$ and $\mathbb{E}[\theta_t' dR_t] = \theta_t' \mu_R(x_t) dt$.

A.4 Relation to recursive preferences

Consider recursive preferences of Duffie and Epstein (1992). Given a process for consumption $\{c_t\}_{t \geq 0}$, they define the investor's value process $\{V_t\}_{t \geq 0}$ as follows:

$$V_t = \mathbb{E}\left[\int_t^\infty \varphi(c_s, V_s) ds\right]$$

The problem is to maximize V_t over $\{c_s, \theta_s\}_{s \geq t}$, while keeping $w_s \geq 0$ for all $s \geq t$. To be consistent with consumption choice of log investors, choose a form of $\varphi(\cdot)$ that keeps the elasticity of intertemporal substitution equal to one while allowing for non-unitary relative risk aversion:

$$\varphi(c, V) = \rho(1 - 1/\gamma) V \left[\log(c) - \frac{\log((1 - 1/\gamma) V)}{1 - 1/\gamma}\right]$$

The recursive representation of this problem is

$$\begin{aligned} 0 = \max_{c,\theta} \ & \varphi(c, V(w, x)) + (r(x) w - c + w\theta' \mu_R(x)) V_w(w, x) + \frac{\theta' \sigma_R(x) \sigma_R(x)' \theta}{2} w^2 V_{ww}(w, x) \\ & + \mu_X(x)' V_{x'}(w, x) + \frac{1}{2} \mathrm{tr}[\sigma_X(x)' V_{xx'}(w, x) \sigma_X(x)] + w\theta' \sigma_R(x) \sigma_X(x)' V_{wx'}(w, x) \end{aligned} \quad (A.4)$$

The solution to this problem is usually guessed and verified. The value function is

$$V(w,x) = \frac{(w\eta(x))^{1-1/\gamma}}{1-1/\gamma}$$

Here $\eta(x)$ is the marginal value of wealth that solves a second-order partial differential equation. This implies $c = \rho w$ and $\varphi(c,V) = \rho(w\eta(x))^{1-1/\gamma}[\log(\rho) - \log(\eta(x))]$. Hence, $w^{1-1/\gamma}$ cancels out from equation (A.4), and the conjecture that $V(\cdot)$ is multiplicatively separable over w and other states proves correct. Under a constant γ, the optimal portfolio is

$$\theta = \gamma[\sigma_R(x)\sigma_R(x)']^{-1}\mu_R(x) + \underbrace{\frac{\gamma-1}{\eta(x)}[\sigma_R(x)\sigma_R(x)']^{-1}\sigma_R(x)\sigma_X(x)'\eta_{x'}(x)}_{\text{hedging motive}} \tag{A.5}$$

This expression exactly coincides with equation (A.2) if $\gamma = \psi/(1+\psi)$. The optimal portfolio consists of the mean-variance part $\gamma[\sigma_R(x)\sigma_R(x)']^{-1}\mu_R(x)$, sometimes called the myopic portfolio, and the hedging part. Computing the latter requires knowing the marginal value of wealth $\eta(x)$, which solves a second-order partial differential equation. The presence of this term increases computational complexity and reduces tractability. The only case without the hedging motive is $\gamma = 1$, but fixing γ makes it impossible to use fluctuations in preferences for risk as a driver of dynamics. Eliminating the hedging motive and introducing dynamics to γ is the main advantage of using the mean-variance risk limit over recursive preferences.

A.5 Incorporating external income with risk limits

Assuming that investors only have asset income and consumption expenditures is restrictive. The simple form of consumption and portfolio choice in Proposition 2 extends to some cases when a part of income is not chosen by the agent. Specifically,

$$dw_t = (r(x_t)w_t - c_t)dt + w_t\theta_t' dR_t - w_t\varsigma(x_t)dt - w_t\tau(x_t)'dZ_t$$

Here $w_t\varsigma(x_t)dt$ and $w_t\tau(x_t)'dZ_t$ are new terms that the agent does not choose. They can represent taxes, with $\varsigma(x_t)$ being a locally deterministic tax and $\tau(x_t)$ being a vector of taxes that load on the exogenous shocks. These taxes can be useful for inducing stationarity in the model through wealth redistribution: they can prevent the least constrained agent from taking over the entire economy. Taxes are lump-sum in the sense that the agent does not directly choose the tax base, which would be the case, for instance, with proportional taxes on portfolio returns. They still scale with wealth, however, and this feature is key to preserving consumption and portfolio choice.

The locally deterministic tax $\varsigma(x_t)$ clearly does not change anything in the agent's decision problem as it is simply isomorphic to a change in the interest rate. The stochastic tax in general affects portfolio choice. For a clearer characterization, I focus on a special case

$$\tau(x_t) = \zeta(x_t)\gamma_t \cdot \sigma_R(x_t)'[\sigma_R(x_t)\sigma_R(x_t)']^{-1}\mu_R(x_t) \tag{A.6}$$

If taxes are set up this way, they generate the same exposure to shocks as the optimal portfolio:

$$\tau(x_t)'dZ_t = \zeta(x_t)\theta(w_t, x_t)'\sigma_R(x_t)dZ_t \equiv \zeta(x_t)\theta(w_t, x_t)'(dR_t - \mu_R(x_t)dt)$$

Here $\theta(w_t, x_t)$ is the optimal vector of portfolio weights in the baseline model. This allows a potential government to effectively tax away a share $\zeta(x_t)$ of stochastic returns without imposing proportional taxes. I next show the effect of imposing these taxes on the agent.

PROPOSITION 9. *Suppose the stochastic tax rate is given by equation (A.6). Let $\zeta(x_t) > -1/\gamma_t$ Then, consumption choice is $c(w_t, x_t) = \rho w_t$. The optimal portfolio is*

$$\theta(w_t, x_t) = \min\{\gamma_t, 1 + \zeta(x_t)\gamma_t\} \cdot [\sigma_R(x_t)\sigma_R(x_t)']^{-1}\mu_R(x_t)$$

The value-at-risk constraint is slack if $\zeta(x_t) < 1 - 1/\gamma_t$ and binds otherwise.

Portfolio choice does not change unless the rate $\zeta(x_t)$ is very negative. The intuition is that taxing away a part of random returns decreases the agent's overall exposure to shocks, making her more willing to take risk. But the value-at-risk constraint is binding already without the taxes, so it continues to bind when $\zeta(x_t)$ is positive and even becomes tighter as measured by the size of the multiplier. A negative $\zeta(x_t)$, on the contrary, increases the agent's exposure to shocks and makes the constraint less tight. If $\zeta(x_t)$ falls below $1 - 1/\gamma_t$, the constraint stops binding, and the tax rate shows up in portfolio choice directly.

A.6 Computation in general equilibrium

I make a brief note on solving for asset prices. Apply Itô's lemma to $p_j(\cdot)$ to find $\mu_p(\cdot)$ and $\sigma_{p,y}(\cdot)$. Omitting the time subscripts, define the drift and shock loadings of the state $x = (\boldsymbol{y}, \boldsymbol{\gamma}, \boldsymbol{\nu})$ as $\mu_x(x)$ and $\sigma_x(x)' \equiv (\sigma_{x,y}(x)' \ \sigma_{x,\gamma}(x)')$, where

$$dx = \mu_x(x)dt + \sigma_{x,y}(x)dZ^y + \sigma_{x,\gamma}(x)dZ^\gamma$$

With this definition, equation (3) can be transformed to

$$\left(\rho + \mu_w(\boldsymbol{y}) - \frac{|\sigma_w(\boldsymbol{y})|^2}{\Gamma(x)}\right)p_j(x) = y_j + \mathcal{D}p_j(x)\left(\mu_x(x) - \frac{\sigma_{x,y}(x)\sigma_w(\boldsymbol{y})'}{\Gamma(x)}\right) + \frac{1}{2}\mathrm{tr}[\mathcal{H}p_j(x)\sigma_x(x)\sigma_x(x)']$$

This formulation is reminiscent of Hamilton-Jacobi-Bellman equations used to solve for value functions. In my setup, the key forward-looking objects are asset prices instead.

Two elements are new compared to risk-neutral pricing or HJB equations: overall discount rates are suppressed due to demand for precautionary savings and expected capital gains are attenuated by risk compensation. The same result obtains in a model with homogeneous investors with log utility, nested as $\Gamma = 1$. The tractability gain of the value-at-risk model is that it generates time-varying effective risk tolerance without losing the convenience of the pricing equations and only expanding the state space just enough to account for agent heterogeneity. The linear partial differential equation for prices can be solved on a grid for $x = (\boldsymbol{y}, \boldsymbol{\gamma}, \boldsymbol{\nu})$ coupled with the exogenous laws of motion for $\boldsymbol{y}$ and $\boldsymbol{\gamma}$ and that for $\boldsymbol{\nu}$ given in Proposition 6. Conveniently, local dynamics of each coordinate of $\boldsymbol{\nu}$ are given only by the coordinate itself and $\Gamma = \boldsymbol{\nu}'\boldsymbol{\gamma}$.

B Proofs

Proof of Proposition 1. Alternative measures $\mathbb{Q}$ are indexed by drift correction processes $\{h_t\}_{t\geq 0}$, so the problem transforms into maximizing value over $\{c_t, \theta_t\}_{t\geq 0}$ after minimizing it over $\{h_t\}_{t\geq 0}$. Moreover, the cost function penalizing the log-likelihood ratio turns into $\psi|h|^2/2$ because $\mathbb{E}^{\mathbb{Q}}[\psi_t dm_t|\mathcal{F}_t] = \psi_t|h_t|^2/2$. The recursive representation of this problem is

$$\begin{aligned}\rho V(w,x) = \max_{c,\theta}\ \min_{h}\ \rho\log(c) &+ \frac{\psi|h|^2}{2} + (r(x)w - c + w\theta'\mu_R(x) - w\theta'\sigma_R(x)h)V_w(w,x) \\ &+ \frac{\theta'\sigma_R(x)\sigma_R(x)'\theta}{2}w^2V_{ww}(w,x) + \mu_X(x)'V_{x'}(w,x) \\ &+ \frac{1}{2}\mathrm{tr}[\sigma_X(x)'V_{xx'}(w,x)\sigma_X(x)] + w\theta'\sigma_R(x)\sigma_X(x)'V_{wx'}(w,x) \end{aligned} \tag{A.7}$$

The solution to equation (A.7) can be guessed and verified. Conjecture $V(w,x) = \log(w) + \eta(x)$. Then, $c = \rho w$ and $V_{wx'}(w,x) = 0$, so the guess that $V(\cdot)$ is separable over w and other states is correct. The model corrections and portfolio weights satisfy

$$\begin{aligned} h &= \psi^{-1}\sigma_R(x)'\theta \\ \theta &= [\sigma_R(x)\sigma_R(x)']^{-1}(\mu_R(x) - \sigma_R(x)h) \end{aligned}$$

This linear system can be solved as

$$\begin{aligned} h &= \frac{1}{1+\psi}\sigma_R(x)'[\sigma_R(x)\sigma_R(x)']^{-1}\mu_R(x) \\ \theta &= \frac{\psi}{1+\psi}[\sigma_R(x)\sigma_R(x)']^{-1}\mu_R(x) \end{aligned}$$

Portfolio choice here coincides with that in equation (A.10), with $\gamma = \psi/(1+\psi) < 1$.

The exogenous part of the value function $\eta(\cdot)$ satisfies the following partial differential equation:

$$\begin{aligned}\rho\eta(x) = \rho\log(\rho) + r(x) - \rho + \frac{\psi}{2(1+\psi)}\mu_R(x)'[\sigma_R(x)\sigma_R(x)']^{-1}\mu_R(x) \\ + \mu_X(x)'\eta_{x'}(x) + \frac{1}{2}\mathrm{tr}[\sigma_X(x)'\eta_{xx'}(x)\sigma_X(x)]\end{aligned}$$

The partial differential equation is similar to one in the value-at-risk setup. The appropriate boundary conditions depend on the process for x. □

Proof of Proposition 2. Take the recursive form of the problem. Let $V(w,x)$ be the value of an agent with wealth w given the aggregate state x. The HJB equation is

$$\begin{aligned}\rho V(w,x) = \max_{c,\theta}\ \rho\log(c) + (r(x)w - c + w\theta'\mu_R(x))V_w(w,x) + \frac{\theta'\sigma_R(x)\sigma_R(x)'\theta}{2}w^2V_{ww}(w,x) \\ + \mu_X(x)'V_{x'}(w,x) + \frac{1}{2}\mathrm{tr}[\sigma_X(x)'V_{xx'}(w,x)\sigma_X(x)] + w\theta'\sigma_R(x)\sigma_X(x)'V_{wx'}(w,x) \qquad \text{(A.8)} \\ \text{s.t. } \theta'\sigma_R(x)\sigma_R(x)'\theta \le \gamma\theta'\mu_R(x) \qquad \text{(A.9)}\end{aligned}$$

Guess and verify the solution to equation (A.8) to be $V(w,x) = \log(w) + \eta(x)$. First, this implies $V_{wx'}(w,x) = 0$. Second, $V_x(w,x)$ and $V_{xx'}(w,x)$ are functions of x only: $V_x(w,x) = \eta_x(x)$ and $V_{xx'}(w,x) = \eta_{xx'}(x)$. Finally, $c = \rho w$. Next, consider portfolio choice. Let $\xi(x,w)$ be the multiplier on equation (A.9). Taking the first-order condition with respect to θ,

$$\theta = \frac{1+\gamma\xi(w,x)}{1+2\xi(w,x)}[\sigma_R(x)\sigma_R(x)']^{-1}\mu_R(x)$$

If the constraint were slack, then $\xi(w,x) = 0$ and θ would be the regular mean-variance portfolio typical for log investors, $\theta = [\sigma_R(x)\sigma_R(x)']^{-1}\mu_R(x)$. But by equation (A.9), this contradicts $\gamma < 1$. Hence, $\xi(w,x) > 0$. Plugging θ into equation (A.9) yields

$$\begin{aligned}\gamma &= \frac{1+\gamma\xi(w,x)}{1+2\xi(w,x)} \\ \theta &= \gamma[\sigma_R(x)\sigma_R(x)']^{-1}\mu_R(x) \qquad \text{(A.10)}\end{aligned}$$

The rest of the investor's value function $\eta(x)$, which only matters for welfare accounting, solves

$$\begin{aligned}\rho\eta(x) = \rho\log(\rho) + r(x) - \rho + \frac{2\gamma-\gamma^2}{2}\mu_R(x)'[\sigma_R(x)\sigma_R(x)']^{-1}\mu_R(x) \\ + \mu_X(x)'\eta_{x'}(x) + \frac{1}{2}\mathrm{tr}[\sigma_X(x)'\eta_{xx'}(x)\sigma_X(x)]\end{aligned}$$

with appropriate boundary conditions. □

Proof of Proposition 3. Write down excess returns in matrix form using prices and dividends. The vector of excess returns $d\boldsymbol{R}_t = \overline{\mu}_R(x_t)dt + \overline{\sigma}_R(x_t)dZ_t$, where $(dZ_t)' = [(dZ_t^y)' \ (dZ_t^\gamma)']$, is

$$d\boldsymbol{R}_t \equiv D(\boldsymbol{p}(x_t))^{-1}[\mu_p(x_t) + \boldsymbol{y}(x_t) - r(x_t)\boldsymbol{p}(x_t)]dt + D(\boldsymbol{p}(x_t))^{-1}\sigma_p(x_t)dZ_t \tag{A.11}$$

Here $D(\boldsymbol{p}(x_t))$ is a diagonal matrix with $\boldsymbol{p}(x_t)$ on the main diagonal and zeros everywhere else, and $\sigma_p(x_t) = [\sigma_{p,y}(x_t) \ \sigma_{p,\gamma}(x_t)]$. Expected excess returns on each asset j consist of capital gains $[\mu_p(x_t)]_j/[\boldsymbol{p}(x_t)]_j$ and dividend yield $[\boldsymbol{y}(x_t)]_j/[\boldsymbol{p}(x_t)]_j$ over and above the risk-free rate $r(x_t)$. The loadings of excess returns on the shocks are the loadings of capital gains $[\sigma_p(x_t)]_{j\cdot}/[\boldsymbol{p}(x_t)]_j$ only.

Portfolio choice of every agent i is

$$\boldsymbol{\theta}_{it} = \gamma_{it}[\overline{\sigma}_R(x_t)\overline{\sigma}_R(x_t)']^{-1}\overline{\mu}_R(x_t) \tag{A.12}$$

Multiplying this by w_{it}, summing across i, and using the market clearing condition for each asset,

$$D(\boldsymbol{p}(x_t))\boldsymbol{s} = \sum\nolimits_{i=1}^n \gamma_{it}w_{it} \cdot [\overline{\sigma}_R(x_t)\overline{\sigma}_R(x_t)']^{-1}\overline{\mu}_R(x_t)$$

Multiply both sides by $\overline{\sigma}_R(x_t)\overline{\sigma}_R(x_t)'$ on the left:

$$\overline{\sigma}_R(x_t)\overline{\sigma}_R(x_t)'D(\boldsymbol{p}(x_t))\boldsymbol{s} = \sum\nolimits_{i=1}^n \gamma_{it}w_{it} \cdot \overline{\mu}_R(x_t)$$

Using the definition of Γ_t and the fact that total wealth is $\boldsymbol{p}(x_t)'\boldsymbol{s}$,

$$\overline{\sigma}_R(x_t)\overline{\sigma}_R(x_t)'D(\boldsymbol{p}(x_t))\boldsymbol{s} = \Gamma_t \cdot \boldsymbol{p}(x_t)'\boldsymbol{s} \cdot \overline{\mu}_R(x_t) \tag{A.13}$$

Multiply both sides by $D(\boldsymbol{p}(x_t))$ and replace $D(\boldsymbol{p}(x_t))\overline{\mu}_R(x_t)$ using equation (A.11):

$$\mu_p(x_t) + \boldsymbol{y}(x_t) - r(x_t)\boldsymbol{p}(x_t) = \frac{1}{\Gamma_t \cdot \boldsymbol{p}(x_t)'\boldsymbol{s}} \cdot \sigma_p(x_t)\sigma_p(x_t)'\boldsymbol{s}$$

This is equation (3) because $\sigma_p(x_t)'\boldsymbol{s}/(\boldsymbol{p}(x_t)'\boldsymbol{s}) = \sigma_w(\boldsymbol{y}_t)'$. Multiplying this by $\boldsymbol{s}'$ on the left,

$$\boldsymbol{s}'\mu_p(x_t) + \boldsymbol{s}'\boldsymbol{y}(x_t) - r(x_t)\boldsymbol{s}'\boldsymbol{p}(x_t) = \frac{1}{\Gamma_t} \cdot \frac{\boldsymbol{s}'\sigma_p(x_t)\sigma_p(x_t)'\boldsymbol{s}}{\boldsymbol{p}(x_t)'\boldsymbol{s}}$$

Using the fact that the total consumption $\boldsymbol{s}'\boldsymbol{y}(x_t)$ is a fraction ρ of total wealth $\boldsymbol{s}'\boldsymbol{p}(x_t)$, divide both sides by $\boldsymbol{s}'\boldsymbol{p}(x_t)$ to get

$$\frac{\boldsymbol{s}'\mu_p(x_t)}{\boldsymbol{s}'\boldsymbol{p}(x_t)} + \rho - r(x_t) = \frac{1}{\Gamma_t} \cdot \frac{\boldsymbol{s}'\sigma_p(x_t)\sigma_p(x_t)'\boldsymbol{s}}{(\boldsymbol{p}(x_t)'\boldsymbol{s})^2}$$

Reorganizing,

$$r(x_t) = \rho + \frac{\boldsymbol{s}'\mu_p(x_t)}{\boldsymbol{s}'\boldsymbol{p}(x_t)} - \frac{1}{\Gamma_t} \cdot \frac{|\sigma_p(x_t)'\boldsymbol{s}|^2}{(\boldsymbol{p}(x_t)'\boldsymbol{s})^2} = \rho + \mu_w(\boldsymbol{y}_t) - \frac{|\sigma_w(\boldsymbol{y}_t)|^2}{\Gamma_t}$$

This is equation (2). □

Proof of Proposition 4. The proof is constructive. Take a function $\Lambda(\cdot)$ that satisfies

$$\Lambda(x_t)\,\boldsymbol{p}(x_t) = \mathbb{E}_t \int_t^{\infty} \Lambda(x_s)\,\boldsymbol{y}_s\,ds$$

This implies $\Lambda(x_t)\,\boldsymbol{y}_t\,dt = -\mathbb{E}_t[d(\Lambda(x_t)\,\boldsymbol{p}(x_t))]$. Use Itô's lemma:

$$\mathbb{E}_t[d(\Lambda(x_t)\,\boldsymbol{p}(x_t))] = \mu_\lambda(x_t)\Lambda(x_t)\,\boldsymbol{p}(x_t)dt + \mu_p(x_t)\Lambda(x_t)dt + \sigma_{p,y}(x_t)\sigma_{\lambda,y}(x_t)'\Lambda(x_t)dt$$

Here the drift and loadings $\mu_\lambda(x_t)$ and $\sigma_\lambda(x_t)$ are defined as

$$d\Lambda(x_t) = \mu_\lambda(x_t)\Lambda(x_t)dt + \Lambda(x_t)\sigma_{\lambda,y}(x_t)'dZ_t^y$$

Using $\Lambda(x_t)\,\boldsymbol{y}_t\,dt = -\mathbb{E}_t[d(\Lambda(x_t)\,\boldsymbol{p}(x_t))]$,

$$\mu_\lambda(x_t)\,\boldsymbol{p}(x_t) + \mu_p(x_t) + \boldsymbol{y}_t + \sigma_{p,y}(x_t)\sigma_{\lambda,y}(x_t)' = 0$$

One choice of $\mu_\lambda(x_t)$ and $\sigma_\lambda(x_t)$ that satisfies both this and equation (3) is $\mu_\lambda(x_t) = -r(x_t)$ and $\sigma_{\lambda,y}(x_t) = -\sigma_w(\boldsymbol{y}_t)'/\Gamma(x_t)$. Using equation (2), we have

$$\frac{d\Lambda(x_t)}{\Lambda(x_t)} = -r(x_t)dt - \frac{1}{\Gamma(x_t)}\sigma_w(\boldsymbol{y}_t)dZ_t^y$$

This completes the proof. □

Proof of Proposition 5. Using the result of Proposition 2, we can characterize the drift correction H_t the hypothetical investor with a robustness parameter Ψ_t chooses for the aggregate market return $dR_t = \hat{\mu}_R(x_t)dt + \hat{\sigma}_R(x_t)dZ_t^y$:

$$H_t = \frac{1}{\Psi_t + 1}\hat{\sigma}_R(x_t)'[\hat{\sigma}_R(x_t)\hat{\sigma}_R(x_t)']^{-1}\hat{\mu}_R(x_t) \qquad \text{(A.14)}$$

From Proposition 3, we know that $\hat{\sigma}_R(x_t) = \sigma_w(\boldsymbol{y}_t)$ and $\hat{\mu}_R(x_t) = |\sigma_w(\boldsymbol{y}_t)|^2/\Gamma_t$. Hence,

$$H_t = \frac{1}{\Psi_t + 1} \cdot \frac{1}{\Gamma_t}\sigma_w(\boldsymbol{y}_t)' = \frac{1-\Gamma_t}{\Gamma_t}\sigma_w(\boldsymbol{y}_t)' \tag{A.15}$$

The last equality uses $\Psi_t = \Gamma_t/(1-\Gamma_t)$. The final step is recognizing that the drift correction maps to the forecast error in the following way: since $\mathbb{E}_t^{\mathbb{Q}}[dZ_t^y] = -H_t dt$,

$$\begin{aligned} dR_t - \mathbb{E}^{\mathbb{Q}_t}[dR_t] &= dR_t - \hat{\mu}_R(x_t)dt + \hat{\sigma}_R(x_t)H_t dt = \hat{\mu}_R(x_t)dt + \hat{\sigma}_R(x_t)dZ_t^y - \hat{\mu}_R(x_t)dt + \hat{\sigma}_R(x_t)H_t dt \\ &= \frac{1-\Gamma_t}{\Gamma_t}|\sigma_w(\boldsymbol{y}_t)|^2 dt + \sigma_w(\boldsymbol{y}_t)dZ_t^y \end{aligned}$$

This completes the proof. $\square$

Proof of Proposition 6. This proof follows the notation introduced in the proof of Proposition 3. To obtain total leverage λ_{it} of investor i, use equation (A.12) for portfolio choice:

$$\lambda_{it} = \gamma_{it} \cdot 1_k'[\overline{\sigma}_R(x_t)\overline{\sigma}_R(x_t)']^{-1}\overline{\mu}_R(x_t)$$

The market clearing condition for risk-free bonds implies

$$0 = \sum\nolimits_{i=1}^n w_{it}(1-\lambda_{it}) = \sum\nolimits_{i=1}^n w_{it} - \sum\nolimits_{i=1}^n \gamma_{it}w_{it} \cdot 1_k'[\overline{\sigma}_R(x_t)\overline{\sigma}_R(x_t)']^{-1}\overline{\mu}_R(x_t)$$

Dividing by total wealth,

$$1 = \Gamma_t \cdot 1_k'[\overline{\sigma}_R(x_t)\overline{\sigma}_R(x_t)']^{-1}\overline{\mu}_R(x_t)$$

Plugging this back into the expression for leverage,

$$\lambda_{it} = \frac{\gamma_{it}}{\Gamma_t}$$

To get the expression for holdings $\{m_{ijt}\}$, notice that $m_{ijt} = \theta_{ijt}w_{it}/p_{jt}$, which means

$$m_{ijt} = \gamma_{it}w_{it} \cdot [D(\boldsymbol{p}(x_t))^{-1}[\overline{\sigma}_R(x_t)\overline{\sigma}_R(x_t)']^{-1}\overline{\mu}_R(x_t)]_j$$

For every j, holdings m_{ijt} are proportional to $\gamma_{it}w_{it}$ with a common j-specific coefficient of proportionality. Since holdings sum to s_j,

$$h_{ijt} = s_j \cdot \frac{\gamma_{it}w_{it}}{\sum_{i=1}^n \gamma_{it}w_{it}} = s_j \cdot \frac{\gamma_{it}\nu_{it}}{\Gamma_t} = s_j\nu_{it}\lambda_{it}$$

The last remaining result to establish is equation (4) for the dynamics of wealth shares. Start with using equation (A.13) to replace $\mu_R(x_t)$ in equation (A.12):

$$\boldsymbol{\theta}_{it} = \frac{\gamma_{it}}{\Gamma_t \cdot \boldsymbol{p}(x_t)'\boldsymbol{s}} \cdot D(\boldsymbol{p}(x_t))\boldsymbol{s} = \frac{\lambda_{it}}{\boldsymbol{p}(x_t)'\boldsymbol{s}} \cdot D(\boldsymbol{p}(x_t))\boldsymbol{s}$$

This implies

$$\begin{aligned}
\boldsymbol{\theta}_{it}'\mu_R(x_t) &= \frac{\lambda_{it}\boldsymbol{s}'D(\boldsymbol{p}(x_t))\overline{\sigma}_R(x_t)\overline{\sigma}_R(x_t)'D(\boldsymbol{p}(x_t))\boldsymbol{s}}{\Gamma_t(\boldsymbol{p}(x_t)'\boldsymbol{s})^2} = \frac{\lambda_{it}\boldsymbol{s}'\sigma_p(x_t)\sigma_p(x_t)'\boldsymbol{s}}{\Gamma_t(\boldsymbol{p}(x_t)'\boldsymbol{s})^2} = \frac{\lambda_{it}}{\Gamma_t} \cdot \frac{|\sigma_p(x_t)'\boldsymbol{s}|^2}{(\boldsymbol{p}(x_t)'\boldsymbol{s})^2} \\
\boldsymbol{\theta}_{it}'\sigma_R(x_t) &= \lambda_{it} \cdot \frac{\boldsymbol{s}'\sigma_p(x_t)}{\boldsymbol{p}(x_t)'\boldsymbol{s}}
\end{aligned}$$

Here I used the fact that $D(\boldsymbol{p}(x_t))\overline{\sigma}_R(x_t) = \sigma_p(x_t)$. The dynamics of individual wealth are

$$\begin{aligned}
dw_{it} &= (r(x_t) - \rho)w_{it}dt + w_{it}\boldsymbol{\theta}_{it}'\overline{\mu}_R(x_t)dt + w_{it}\boldsymbol{\theta}_{it}'\overline{\sigma}_R(x_t)dZ_t \\
&= (r(x_t) - \rho)w_{it}dt + \frac{\lambda_{it}w_{it}}{\Gamma_t} \cdot \frac{|\sigma_p(x_t)'\boldsymbol{s}|^2}{(\boldsymbol{p}(x_t)'\boldsymbol{s})^2}dt + \lambda_{it}w_{it} \cdot \frac{\boldsymbol{s}'\sigma_p(x_t)}{\boldsymbol{p}(x_t)'\boldsymbol{s}}dZ_t
\end{aligned}$$

Summing this across i and denoting the total wealth by w_t,

$$dw_t = (r(x_t) - \rho)w_t dt + \frac{w_t}{\Gamma_t} \cdot \frac{|\sigma_p(x_t)'\boldsymbol{s}|^2}{(\boldsymbol{p}(x_t)'\boldsymbol{s})^2}dt + w_t \cdot \frac{\boldsymbol{s}'\sigma_p(x_t)}{\boldsymbol{p}(x_t)'\boldsymbol{s}}dZ_t$$

Here I use the fact that the average leverage weighted with wealth shares is one:

$$\sum\nolimits_{i=1}^n \lambda_{it}w_{it} = \sum\nolimits_{i=1}^n \lambda_{it}\nu_{it}w_t = \sum\nolimits_{i=1}^n \frac{\gamma_{it}\nu_{it}}{\Gamma_t}w_t = w_t$$

Now consider the dynamics of $\nu_{it} = w_{it}/w_t$:

$$\begin{aligned}
d\nu_{it} &= \nu_{it}\left(\frac{dw_{it}}{w_{it}} - \frac{dw_t}{w_t} + \frac{(dw_t)^2}{w_t^2} - \frac{dw_t}{w_t}\frac{dw_{it}}{w_{it}}\right) \\
&= \nu_{it}\left(\frac{\lambda_{it} - 1}{\Gamma_t} \cdot \frac{|\sigma_p(x_t)'\boldsymbol{s}|^2}{(\boldsymbol{p}(x_t)'\boldsymbol{s})^2}dt + (\lambda_{it} - 1) \cdot \frac{\boldsymbol{s}'\sigma_p(x_t)}{\boldsymbol{p}(x_t)'\boldsymbol{s}}dZ_t + (1 - \lambda_{it}) \cdot \frac{|\sigma_p(x_t)'\boldsymbol{s}|^2}{(\boldsymbol{p}(x_t)'\boldsymbol{s})^2}dt\right) \\
&= \nu_{it}(\lambda_{it} - 1) \cdot \left[\frac{1 - \Gamma_t}{\Gamma_t} \cdot \frac{|\sigma_p(x_t)'\boldsymbol{s}|^2}{(\boldsymbol{p}(x_t)'\boldsymbol{s})^2}dt + \frac{\boldsymbol{s}'\sigma_p(x_t)}{\boldsymbol{p}(x_t)'\boldsymbol{s}}dZ_t\right] \\
&= \nu_{it}(\lambda_{it} - 1) \cdot \left[\frac{1 - \Gamma_t}{\Gamma_t}|\sigma_w(\boldsymbol{y}_t)|^2 dt + \sigma_w(\boldsymbol{y}_t)dZ_t^y\right] = \nu_{it}(\lambda_{it} - 1)dW_t
\end{aligned}$$

This completes the proof. $\square$

Proof of Proposition 7. Apply Itô's lemma to Γ_t:

$$d\Gamma_t = \sum_{i=1}^{n} d\nu_{it}\gamma_{it} + \sum_{i=1}^{n} \nu_{it} d\gamma_{it} + \sum_{i=1}^{n} d\nu_{it} d\gamma_{it} \tag{A.16}$$

Since wealth shares do not load on dZ_t^γ, which follows from Proposition 6, the last term is zero. The second term is $\boldsymbol{\nu}_t' \, d\,\boldsymbol{\gamma}_t$. The first term is

$$\sum_{i=1}^{n} d\nu_{it}\gamma_{it} = \frac{1}{\Gamma_t}\sum_{i=1}^{n} \nu_{it}\gamma_{it}(\gamma_{it} - \Gamma_t) dW_t = \frac{1}{\Gamma_t}\left(\sum_{i=1}^{n} \nu_{it}\gamma_{it}^2 - \Gamma_t^2\right) dW_t = \frac{\Delta_t}{\Gamma_t} dW_t \tag{A.17}$$

Here Δ_t is the wealth-weighted variance of γ_t. $\square$

Proof of Proposition 8. Take the recursive form of the baseline problem with the new budget constraint including $\tau(x)$. Assume $\varsigma(x) = 0$. Let $V(w, x)$ be the value. The HJB equation is

$$\begin{aligned}
\rho V(w,x) = \max_{c,\theta}\ \min_{h}\ & \rho \log(c) + \frac{\psi |h|^2}{2} + (r(x)w - c + w\theta'\mu_R(x) - w\theta'\sigma_R(x)h)V_w(w,x) \\
&+ \frac{(\theta'\sigma_R(x) - \tau(x)')(\sigma_R(x)'\theta - \tau(x))}{2} w^2 V_{ww}(w,x) + \mu_X(x)' V_{x'}(w,x) \\
&+ \frac{1}{2}\mathrm{tr}[\sigma_X(x)' V_{xx'}(w,x)\sigma_X(x)] + w(\theta'\sigma_R(x) - \tau(x)')\sigma_X(x)' V_{wx'}(w,x)
\end{aligned} \tag{A.18}$$

Like in the proof of Proposition 1, guess and verify the solution to equation (A.18) to be $V(w,x) = \log(w) + \eta(x)$. Again, the consequences are that $V_{wx'}(w,x) = 0$ and that $V_x(w,x)$ and $V_{xx'}(w,x)$ are functions of x only: $V_x(w,x) = \eta_x(x)$ and $V_{xx'}(w,x) = \eta_{xx'}(x)$. Consumption is still a constant fraction of wealth: $c = \rho w$. The model corrections and portfolio weights satisfy

$$\begin{aligned}
h &= \psi^{-1}\sigma_R(x)'\theta \\
\sigma_R(x)(\sigma_R(x)'\theta - \tau(x)) &= \mu_R(x) - \sigma_R(x)h
\end{aligned}$$

Now use the fact that $\tau(x) = \zeta(x)\sigma_R(x)'[\sigma_R(x)\sigma_R(x)']^{-1}\mu_R(x)$ to solve this linear system as

$$\begin{aligned}
h &= \frac{1+\zeta(x)}{1+\psi}\sigma_R(x)'[\sigma_R(x)\sigma_R(x)']^{-1}\mu_R(x) \\
\theta &= \frac{\psi(1+\zeta(x))}{1+\psi}[\sigma_R(x)\sigma_R(x)']^{-1}\mu_R(x)
\end{aligned}$$

The equation that the rest of investor's value function $\eta(x)$ solves changes to

$$\begin{aligned}\rho\eta(x) &= \rho\log(\rho) + r(x) - \rho + \frac{\psi - \zeta(x)^2 + 2\psi\zeta(x)}{2(1+\psi)}\mu_R(x)'[\sigma_R(x)\sigma_R(x)']^{-1}\mu_R(x) \\ &\quad + \mu_X(x)'\eta_{x'}(x) + \frac{1}{2}\mathrm{tr}[\sigma_X(x)'\eta_{xx'}(x)\sigma_X(x)]\end{aligned}$$

with appropriate boundary conditions. □

Proof of Proposition 9. Take the recursive form of the problem with a risk limit with the new budget constraint including $\tau(x)$. Let $\varsigma(x) = 0$. Let $V(w, x)$ be the value. The HJB equation is

$$\begin{aligned}\rho V(w,x) = \max_{c,\theta}\ & \rho\log(c) + (r(x)w - c + w\theta'\mu_R(x))V_w(w,x) \\ & + \frac{(\theta'\sigma_R(x) - \tau(x)')(\sigma_R(x)'\theta - \tau(x))}{2}w^2 V_{ww}(w,x) + w(\theta'\sigma_R(x) - \tau(x)')\sigma_X(x)'V_{wx'}(w,x) \\ & + \mu_X(x)'V_{x'}(w,x) + \frac{1}{2}\mathrm{tr}[\sigma_X(x)'V_{xx'}(w,x)\sigma_X(x)] \qquad \text{(A.19)} \\ \text{s.t. } & \theta'\sigma_R(x)\sigma_R(x)'\theta \leq \gamma\theta'\mu_R(x) \qquad \text{(A.20)}\end{aligned}$$

Like in the proof of Proposition 2, guess and verify the solution to equation (A.19) to be $V(w,x) = \log(w) + \eta(x)$. Again, the consequences are that $V_{wx'}(w,x) = 0$ and that $V_x(w,x)$ and $V_{xx'}(w,x)$ are functions of x only: $V_x(w,x) = \eta_x(x)$ and $V_{xx'}(w,x) = \eta_{xx'}(x)$. Consumption is still a constant fraction of wealth: $c = \rho w$. Moving to portfolio choice, let $\xi(x,w)$ be the multiplier on equation (A.20). The first-order condition with respect to θ is now

$$\theta = \frac{1+\gamma\xi(w,x)}{1+2\xi(w,x)}[\sigma_R(x)\sigma_R(x)']^{-1}\mu_R(x) + \frac{1}{1+2\xi(w,x)}[\sigma_R(x)\sigma_R(x)']^{-1}\sigma_R(x)\tau(x)$$

Now use the fact that $\tau(x) = \gamma\zeta(x)\sigma_R(x)'[\sigma_R(x)\sigma_R(x)']^{-1}\mu_R(x)$:

$$\theta = \frac{1+\gamma\xi(w,x)+\gamma\zeta(x)}{1+2\xi(w,x)}[\sigma_R(x)\sigma_R(x)']^{-1}\mu_R(x)$$

Suppose the constraint is slack and $\xi(w,x) = 0$. Then,

$$\theta = (1+\zeta(x)\gamma)[\sigma_R(x)\sigma_R(x)']^{-1}\mu_R(x)$$

Plug this into equation (A.20) to see that slackness is equivalent to

$$1 + \zeta(x)\gamma < \gamma$$

Alternatively, $\zeta(x) < 1 - 1/\gamma$. If $\zeta(x) > 1 - 1/\gamma$, the constraint binds, and

$$\frac{1 + \gamma\xi(w, x) + \gamma\zeta(x)}{1 + 2\xi(w, x)} = \gamma$$

Portfolio choice in this case is hence the same as without taxes:

$$\theta = \gamma[\sigma_R(x)\sigma_R(x)']^{-1}\mu_R(x)$$

The equation that the rest of investor's value function $\eta(x)$ solves changes to

$$\begin{aligned}\rho\eta(x) = \rho\log(\rho) + r(x) - \rho + \frac{2\gamma - (1 - \zeta(x))^2\gamma^2}{2}\mu_R(x)'[\sigma_R(x)\sigma_R(x)']^{-1}\mu_R(x) \\ + \mu_X(x)'\eta_{x'}(x) + \frac{1}{2}\mathrm{tr}[\sigma_X(x)'\eta_{xx'}(x)\sigma_X(x)]\end{aligned}$$

with appropriate boundary conditions. □

Proof of Corollary 1. The statement directly follows Proposition 6 . □

Proof of Corollary 2, Corollary 3, and Corollary 4. Implied by Proposition 7 and applying Itô's lemma to $\pi(\nu_t, \gamma_t) = \sigma^2/\Gamma(\nu_t, \gamma_t)$. □

Proof of Corollary 5. Implied by Proposition 6 and Proposition 7. □

C Analytical examples

I use the simplest possible setup that incorporates both wealth redistribution and model revisions. The claim to total output is the only risky asset. There are two investors with different robustness parameters, and hence different subjective models. In equilibrium, the agent with a less pessimistic model borrows from the other one to bet on output growth. One of the robustness parameters changes exogenously, following a mean-reverting process that generates a cycle in the market's effective risk-tolerance through two channels. First, it directly changes effective risk-tolerance of one of the agents and hence the wealth-weighted average. Second, without instantly affecting wealth shares, it makes agents rebalance their portfolios and take different exposure to aggregate shocks. Output shocks then redistribute wealth and change the risk premium. I characterize the law of motion of wealth shares, the risk premium, and the interest rate analytically.

The first agent's robustness parameter, denoted by γ_t, varies over time stochastically:

$$d\gamma_t = \mu_\gamma(\gamma_t)dt + \sigma_\gamma(\gamma_t)dZ_t^\gamma$$

Here γ_t varies within $[\underline{\gamma}, \overline{\gamma}] \subset (0, 1]$. The second agent's parameter is fixed at $\hat{\gamma} \in [\underline{\gamma}, \overline{\gamma}]$. Output is produced by a Lucas tree in unit supply, $\boldsymbol{s} = 1$. The flow output of the tree is $y_t dt$, where the rate of production evolves as

$$\frac{dy_t}{y_t} = \mu dt + \sigma dZ_t^y$$

The tree price is p_t, and excess returns on the tree are $dR_t = (dp_t + y_t dt)/p_t - r_t dt$. Since the supply is normalized to one, total wealth is equal to the tree's price: $p_t = w_t$. This implies $\rho p_t = y_t$. The price-dividend ratio is constant, and the capital gains process coincides with that of output growth. Excess returns transform into $dR_t = (\rho + \mu - r_t)dt + \sigma dZ_t^y$.

To induce stationarity, I make agents pay small wealth taxes. Agent i's budget constraint is

$$dw_{it} = (r_t w_{it} - c_{it})dt + \theta_{it} w_{it} dR_t - T_{it} w_{it} dt$$

Here, as before, θ_{it} is the share of investor i's portfolio allocated to the risky asset. The tax rate T_{it} is the following function of i's wealth share ν_{it}: $T_{it} = T(\nu_{it})$, where

$$T(\nu_{it}) = \tau \left(\nu_{it} - \frac{1}{2}\right) \sqrt{\frac{1 - \nu_{it}}{\nu_{it}}}$$

Here τ is a small positive parameter. This tax policy balances the budget:

$$\sum\nolimits_{i \in \{1,2\}} T_{it} w_{it} = w_t \cdot \sum\nolimits_{i \in \{1,2\}} T_{it} \nu_{it} = \tau w_t \cdot \sum\nolimits_{i \in \{1,2\}} \left(\nu_{it} - \frac{1}{2}\right) \sqrt{\nu_{it}(1 - \nu_{it})} = 0$$

Importantly, agent i takes T_{it} as given and does not realize how her wealth accumulation affects the tax rate. The fact that agents still perceive dw_{it} as linear in w_{it} preserves the functional forms of consumption and portfolio choice.

There are two state variables in this economy: the endogenous wealth share of the first agent, below denoted by ν_t, and her exogenous robustness parameter γ_t. Denote the risk premium by $\pi(\nu_t, \gamma_t) \equiv \rho + \mu - r(\nu_t, \gamma_t)$. Equation (2) from Proposition 3 leads to

$$\pi(\nu_t, \gamma_t) = \frac{\sigma^2}{\Gamma(\nu_t, \gamma_t)} \tag{A.21}$$

When the agent with a higher γ_{it}, or the less pessimistic model, accumulates more wealth, the market's aggregate risk tolerance rises, and the risk premium falls. Since the price-dividend ratio is fixed, this is achieved through a higher interest rate.

With just one risky asset, every investor's leverage is her risky portfolio share: $\lambda_{it} = \theta_{it}$, and hence $\theta_{it} = \gamma_{it}\pi(\nu_t,\gamma_t)/\sigma^2$ for both $\gamma_{it} \in \{\gamma_t, \hat{\gamma}\}$. Wealth dynamics can be represented as

$$\frac{dw_{it}}{w_{it}} = \left(\mu - \pi(\nu_t,\gamma_t) + \gamma_{it}\frac{\pi(\nu_t,\gamma_t)^2}{\sigma^2}\right)dt + \gamma_{it}\frac{\pi(\nu_t,\gamma_t)}{\sigma}dZ_t^y - T_{it}dt \tag{A.22}$$

The term $\mu - \pi(\nu_t,\gamma_t)$ reflects the consumption-savings trade-off. A high risk premium $\pi(\nu_t,\gamma_t)$ corresponds to a low interest rate, reflecting precautionary motives. The term $\gamma_{it}\pi(\nu_t,\gamma_t)^2/\sigma^2$ is the compensation for risk. Higher γ_{it} leads the investor to take a longer position in the tree and raises the compensation together with exposure of her wealth to dZ_t^y.

I now apply a version of Proposition 6 to find the evolution of the first agent's wealth share.

Corollary 1. *The evolution of the first agent's wealth share ν_t is*

$$d\nu_t = \underbrace{\frac{\nu_t(1-\nu_t)(\gamma_t-\hat{\gamma})}{\Gamma(\nu_t,\gamma_t)}}_{\text{1st agent's excess leverage}} \cdot \left[\underbrace{\frac{1-\Gamma(\nu_t,\gamma_t)}{\Gamma(\nu_t,\gamma_t)}\sigma^2 dt + \sigma dZ_t^y}_{\text{representative forecast error}}\right] - \underbrace{\tau\sqrt{\nu_t(1-\nu_t)}\left(\nu_t - \frac{1}{2}\right)dt}_{\text{wealth tax}} \tag{A.23}$$

The first agent's wealth share is positively exposed to the growth surprises in the representative subjective model if and only if she currently has higher risk-taking capacity, $\gamma_t > \hat{\gamma}$. In this case, she borrows from the other agents and takes on more aggregate risk. Her wealth share also drifts up: she receives deterministic compensation, which makes her relatively richer over time.

I next characterize the dynamics of the risk premium $\pi(\nu_t,\gamma_t)$. Denote

$$\frac{d\pi(\nu_t,\gamma_t)}{\pi(\nu_t,\gamma_t)} = \mu_\pi(\nu_t,\gamma_t)dt + \sigma_{\pi,y}(\nu_t,\gamma_t)dZ_t^y + \sigma_{\pi,\gamma}(\nu_t,\gamma_t)dZ_t^\gamma$$

Applying Itô's lemma to the result of Proposition 7 and supposing $\tau = 0$ for illustration, I get the following evolution of $\pi(\nu_t,\gamma_t)$.

Corollary 2. *The shock loadings of $\pi(\nu_t,\gamma_t)$ are*

$$\sigma_{\pi,y}(\nu_t,\gamma_t) = -\sigma\cdot\frac{\Delta(\nu_t,\gamma_t)}{\Gamma(\nu_t,\gamma_t)^2}$$
$$\sigma_{\pi,\gamma}(\nu_t,\gamma_t) = -\sigma_\gamma(\gamma_t)\cdot\frac{\nu_t}{\Gamma(\nu_t,\gamma_t)}$$
$$\mu_\pi(\nu_t,\gamma_t) = \Gamma(\nu_t,\gamma_t)\cdot\left(\underbrace{\sigma_{\pi,\gamma}(\nu_t,\gamma_t)^2\left[1-\frac{\mu_\gamma(\gamma_t)}{\sigma_\gamma(\gamma_t)^2\nu_t}\right]}_{\text{impact of preferences}} + \underbrace{\sigma_{\pi,y}(\nu_t,\gamma_t)^2\left[1-\frac{(1-\Gamma(\nu_t,\gamma_t))\Gamma(\nu_t,\gamma_t)}{\Delta(\nu_t,\gamma_t)}\right]}_{\text{impact of redistribution}}\right)$$

The second term, impact of redistribution, is always negative.

The risk premium decreases after positive output shocks because they redistribute to the less pessimistic agent, who levers up to bet on output growth. It also decreases following positive shocks to the robustness parameter, which make the representative subjective model less pessimistic. The drift in the risk premium has two parts. The part coming from the evolution of the robustness parameter can take either sign, while that coming from wealth redistribution is always negative, since deterministic risk compensation makes the less pessimistic agent relatively richer over time.

Vasicek (1977) and Cox, Ingersoll, and Ross (1985) models with a single agent. Consider a special case of this model with $\nu_t = 1$: the investor with a time-varying multiplier dominating the market. The state space reduces to one dimension, with γ_t being the only state variable. In this model, the risk premium fluctuates for exogenous reasons, and so does the interest rate. It is possible to obtain any process for the interest rate by choosing an appropriate process for γ_t, subject to the natural restriction $\gamma_t \in (0, 1]$, which implies $r_t \in (-\infty, \rho + \mu - \sigma^2]$. One example is the interest rate process from the Vasicek (1977) model. This Ornstein-Uhlenbeck process has been widely used in models of term structure, including a recent wave of preferred-habitat models in the style of Vayanos and Vila (2021).

Corollary 3. *Let $\kappa_r > 0$ and σ_r be parameters and suppose the process for γ_t is*

$$\frac{d\gamma_t}{\gamma_t} = \left(\kappa_r + \frac{\sigma_r^2}{\sigma^4}\gamma_t^2\right) dt + \frac{\sigma_r}{\sigma^2}\gamma_t dZ_t^\gamma$$

with a reflecting boundary at $\gamma_t = 1$. Then, the process for the interest rate is

$$dr_t = \kappa_r(\rho + \mu - r_t)dt + \sigma_r dZ_t^\gamma$$

with a reflecting boundary at $r_t = \rho + \mu - \sigma^2$.

Another example is the interest rate process from Cox, Ingersoll, and Ross (1985), which adds a volatility factor of the square root of the interest rate to that in Vasicek (1977). This process for the interest rate is non-negative, so γ_t now varies within $[\sigma^2/(\rho + \mu), 1]$, assuming $\rho + \mu > \sigma^2$.

Corollary 4. *Let $\kappa_r > 0$ and σ_r be parameters and suppose the process for γ_t is*

$$\frac{d\gamma_t}{\gamma_t} = \left(\kappa_r + \frac{\sigma_r^2\gamma_t^2}{\sigma^4}\left(\rho + \mu - \frac{\sigma^2}{\gamma_t}\right)\right) dt + \frac{\sigma_r\gamma_t}{\sigma^2}\sqrt{\rho + \mu - \frac{\sigma^2}{\gamma_t}} dZ_t^\gamma$$

with a reflecting boundary at $\gamma_t = 1$. Then, the process for the interest rate is

$$dr_t = \kappa_r(\rho + \mu - r_t)dt + \sigma_r\sqrt{r_t} dZ_t^\gamma$$

with a reflecting boundary at $r_t = \rho + \mu - \sigma^2$.

Caballero and Simsek (2020) model with fixed multipliers. Consider the converse case, where ν_t fluctuates within $(0, 1)$ but γ_t is fixed at $\overline{\gamma}$, the upper bound of the support. For simplicity, let the second investor's multiplier be at the lower bound: $\hat{\gamma} = \underline{\gamma} < \overline{\gamma}$. The state space again reduces to one dimension, but the state variable is now the wealth share ν_t, which is endogenous.

This reduction of the state space to one dimension means that the wealth share ν_t and the risk premium π_t are both Markov processes that can be characterized in the following form:

$$\begin{aligned} d\nu_t &= \mu_\nu(\nu_t)dt + \sigma_\nu(\nu_t)dZ_t \\ \frac{d\pi_t}{\pi_t} &= \mu_\pi(\pi_t)dt + \sigma_\pi(\pi_t)dZ_t \end{aligned}$$

The risk premium fluctuates on $[\underline{\pi}, \overline{\pi}]$, where $\underline{\pi} = \sigma^2/\overline{\gamma}$ and $\overline{\pi} = \sigma^2/\underline{\gamma}$. Corollary 2 implies

Corollary 5. *Supposing $\tau = 0$ for simplicity, the drift and volatility of ν_t are*

$$\begin{aligned} \mu_\nu(\nu_t) &= \nu_t\,(1-\nu_t)\cdot\frac{\sigma^2(\overline{\gamma}-\underline{\gamma})\,(1-\nu_t\overline{\gamma}-(1-\nu_t)\underline{\gamma})}{(\nu_t\overline{\gamma}+(1-\nu_t)\underline{\gamma})^2} \\ \sigma_\nu(\nu_t) &= \nu_t\,(1-\nu_t)\cdot\frac{\sigma(\overline{\gamma}-\underline{\gamma})}{\nu_t\overline{\gamma}+(1-\nu_t)\underline{\gamma}} \end{aligned}$$

with $\mu_\nu(\nu_t) > 0$ and $\sigma_\nu(\nu_t) > 0$ for all interior ν_t. The drift and volatility of the risk premium are

$$\begin{aligned} \mu_\pi(\pi_t) &= \left(\frac{\sigma}{\overline{\pi}\underline{\pi}}\right)^2\cdot(\sigma^2(\overline{\pi}+\underline{\pi}-\pi_t)-\overline{\pi}\underline{\pi})\cdot(\overline{\pi}-\pi_t)(\pi_t-\underline{\pi}) \\ \sigma_\pi(\pi_t) &= -\frac{\sigma}{\overline{\pi}\underline{\pi}}\cdot(\overline{\pi}-\pi_t)(\pi_t-\underline{\pi}) \end{aligned}$$

with $\mu_\pi(\pi_t) < 0$ and $\sigma_\pi(\pi_t) < 0$ for all interior π_t.

Both drift and volatility of the wealth share ν_t depend on the difference $\overline{\gamma} - \underline{\gamma}$, which measures polarization in attitudes to risk. The gap in risk appetite creates different portfolio allocations and ultimately allows output shocks to induce redistribution towards more risk-tolerant agents. This is a simplified version of the mechanism in Caballero and Simsek (2020), who induce two agents to take differential bets on output growth because of disagreement on the underlying process. When the realized output growth is above average, the more optimistic agent becomes relatively richer, which suppresses the risk premium and raises the interest rate. Caballero and Simsek (2020) then mount a New-Keynesian block on this structure, generating additional output expansion if the monetary authority does not correspondingly raise the policy rate.

In my setup, differences in risk tolerance come from value-at-risk multipliers. Heterogeneous multipliers reflect the fact that agents choose different statistical models for asset returns. There is a disagreement over expected output growth just like in Caballero and Simsek (2020). The difference between my setup and Caballero and Simsek (2020) is that here agents have the same reference model, and the fundamental heterogeneity is in robustness preference parameters.