DISCUSSION PAPER SERIES

DP14592

GENERALIZED ROBUSTNESS AND DYNAMIC PESSIMISM

Pascal Maenhout, Andrea Vedolin and Hao Xing

FINANCIAL ECONOMICS

ISSN 0265-8003

GENERALIZED ROBUSTNESS AND DYNAMIC PESSIMISM

Pascal Maenhout, Andrea Vedolin and Hao Xing

Discussion Paper DP14592
Published 09 April 2020
Submitted 08 April 2020

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

GENERALIZED ROBUSTNESS AND DYNAMIC PESSIMISM

Abstract

This paper develops a theory of dynamic pessimism and its impact on asset prices. Notions of time-varying pessimism arise endogenously in our setting as a consequence of agents' concern for model misspecification. We generalize the robust control approach of Hansen and Sargent (2001) by replacing relative entropy as a measure of discrepancy between models by the more general family of Cressie-Read discrepancies. As a consequence, the decision-maker's distorted beliefs appear as an endogenous state variable driving risk aversion, portfolio decisions, and equilibrium asset prices. Using survey data, we estimate time-varying pessimism and find that such a proxy features a strong business cycle component. We then show that using our measure of pessimism helps match salient features in equity markets such as excess volatility and high equity premium.

JEL Classification: F31, G15

Keywords: Cressie Read, Robust control, Subjective beliefs, Pessimism

Pascal Maenhout - pascal.maenhout@insead.edu
INSEAD

Andrea Vedolin - avedolin@bu.edu
Boston University and CEPR

Hao Xing - haoxing@bu.edu
Boston University

Generalized Robustness and Dynamic Pessimism *

Pascal J. Maenhout[†] Andrea Vedolin[‡] Hao Xing[§]

Abstract

This paper develops a theory of dynamic pessimism and its impact on asset prices. Notions of time-varying pessimism arise endogenously in our setting as a consequence of agents' concern for model misspecification. We generalize the robust control approach of Hansen and Sargent (2001) by replacing relative entropy as a measure of discrepancy between models by the more general family of Cressie-Read discrepancies. As a consequence, the decision-maker's distorted beliefs appear as an endogenous state variable driving risk aversion, portfolio decisions, and equilibrium asset prices. Using survey data, we estimate time-varying pessimism and find that such a proxy features a strong business cycle component. We then show that using our measure of pessimism helps match salient features in equity markets such as excess volatility and high equity premium.

Keywords: Cressie Read, robust control, subjective beliefs, pessimism

JEL Classification: F31, G15

This Version: March 2020

*We thank John Campbell, Gary Chamberlain and Lars Hansen for insightful comments and discussions. All errors are ours.

[†]INSEAD, pascal.maenhout@insead.edu.

[‡]Boston University, NBER and CEPR, avedolin@bu.edu.

[§]Boston University, haoxing@bu.edu.

The seminal work of Hansen and Sargent posits that economic agents with an aversion to uncertainty seek robustness by entertaining a family of models constructed as a neighborhood around a benchmark model and optimize against the worst-case within this family.[1] In a Bayesian interpretation the worst-case model that the decision-maker guards against can be viewed as representing endogenously distorted pessimistic beliefs.[2] A crucial modeling assumption throughout the entire literature on Hansen and Sargent robustness is the use of relative entropy, also known as Kullback and Leibler (1951) entropy, to measure the discrepancy between models. Hansen and Sargent (2008) discuss several reasons why relative entropy is appealing. For example, relative entropy is particularly tractable in a continuous-time setting where alternative models are represented by exponential martingales reflecting the drift distortions applied to the state variables of interest to the decision maker. As a consequence, relative entropy often allows for closed-form solutions of difficult dynamic problems. Little is known, however, about the "robustness" of these robust decision rules and corresponding equilibrium quantities and prices to the use of relative entropy. In addition to facilitating analytical tractability, the assumption of relative entropy might not be innocuous, nor without loss of generality. While relative entropy has strong foundations in information theory, work in econometric theory, statistical decision-making, and empirical model selection also considers alternative measures of model discrepancy. In this paper, we consider the family of Cressie-Read divergences (Cressie and Read, 1984), a one-parameter generalization of relative entropy with parameter η, and study deviations from the entropy case and its implications for asset prices in a general equilibrium setting.[3] More specifically, we show that our extension of robustness offers an important generalization of stochastic differential utility. In particular, we derive the conditions under which the agent's pessimistically distorted beliefs form a state variable, which generates endogenous time-variation in pessimism and stochastic effective risk aversion, unlike in the case of entropy, and helps match salient features in equity markets.

We summarize our main contributions as follows. We first carefully construct a generalization of Hansen and Sargent robustness that is designed to be recursive and homothetic. Relative entropy is nested as a special case, when the Cressie-Read parameter η tends to unity. Our generalization of Hansen and Sargent relative-entropy multiplier preferences is based on a novel and carefully constructed discounted divergence measure, used as a penalty or cost function when minimizing distorted expected utility. The penalty is an expected integral of suitably scaled and appropriately weighted discounted Cressie-Read divergence measures. By design, our measure permits a recursive

[1] Notable contributions include Hansen and Sargent (2001), Anderson, Hansen, and Sargent (2003), and Hansen, Sargent, Turmuhambetova, and Williams (2006), among many others, or Hansen and Sargent (2008) for a textbook treatment.

[2] Alternatively, the family of models being considered by the decision-maker can be viewed as the set of non-unique priors in the max-min expected utility of Gilboa and Schmeidler (1989). The work of Chen and Epstein (2002), Epstein and Schneider (2003) and related papers offers another approach to modeling ambiguity aversion in a dynamic setting motivated by the work of Gilboa and Schmeidler (1989). Maccheroni, Marinacci, and Rustichini (2006a) and Maccheroni, Marinacci, and Rustichini (2006b) show that the framework of variational preferences nests these different approaches. Strzalecki (2011) characterizes multiplier preferences axiomatically, extending in an important way the findings in Maccheroni, Marinacci, and Rustichini (2006a). Most recently, Hansen and Sargent (2019) construct a continuous-time extension of variational preferences that combines ambiguity aversion in the sense of Chen and Epstein (2002) with model uncertainty aversion in the sense of Anderson, Hansen, and Sargent (2003).

[3] The family of Cressie-Read divergences nests relative entropy as a limiting case, as well as several other interesting cases such as Hellinger, χ^2, Burg entropy, and others.

formulation, constituting a recursive ambiguity index in the sense of Maccheroni, Marinacci, and Rustichini (2006b), and preserves homotheticity of the resulting preferences, see, e.g., Maenhout (2004). Using the theory of backward stochastic differential equations of El Karoui, Peng, and Quenez (1997) and Kobylanski (2000), we show that these preferences lead to a form of generalized stochastic differential utility of Lazrak and Quenez (2003), which extends Duffie and Epstein (1992), the continuous-time version of the well-known recursive preferences of Epstein and Zin (1989). In particular, the variance multiplier of the utility process is a nonlinear function of the belief distortion state variable. Intuitively, we obtain belief- and state-dependent risk aversion, except when $\eta = 1$, i.e., in the special case of entropy. We show that for $\eta < 1$, risk aversion is countercyclical and declines following positive shocks, while $\eta > 1$ generates procyclical risk aversion. Intuitively, the parameter η governs the desire for intertemporal smoothness reflected in the Cressie-Read penalty that intertemporally aggregates all drift distortions. Low values of η indicate a high desire for intertemporal smoothing of drift distortions by the malevolent agent, or similarly a low willingness to substitute intertemporally inside the Cressie-Read penalty function.

Second, to gain further intuition about the effects at work, we examine a simplified two-stage model, where the belief distortion state variable is frozen at the start of each stage. Solving backwards, we find that in the second stage, the Cressie-Read investor with fixed belief state variable can be interpreted as an entropy investor, but where effective risk aversion depends on the temporarily fixed state variable, i.e., on the state of the economy that prevails at the time when the state variable was frozen. In the first stage, the investor anticipates the fact that future effective risk aversion will be determined by the (from her vantage point) future state of the economy. This gives rise to an endogenous hedging demand against future random market conditions and associated effective risk aversion. Using this intuition, we then illustrate the consequences of these rich dynamics in a simple partial equilibrium portfolio problem. More specifically, we show that generalized robustness produces intertemporal hedging and therefore state- and horizon-dependent portfolios, despite return dynamics being i.i.d. This stands in sharp contrast to the findings for relative entropy, where optimal portfolios are constant and myopic whenever investment opportunities are constant. Intuitively, because of time-varying sentiment, investment opportunities are *perceived* by the investor fearing misspecification as time-varying, even when returns are truly i.i.d. and investment opportunities are constant under the benchmark model.

Third, before delving into a general equilibrium analysis, we examine the dynamics of optimism/pessimism in the data. To this end, we hand collect survey data on macroeconomic aggregates from a large cross-section of professional forecasters using Blue Chip Economic Indicators. In our model, ambiguity averse or robustness seeking agents imply pessimistically distorted beliefs relative to a rational benchmark. To test this hypothesis, we construct belief wedges, i.e., the difference between the agent's subjective (from the survey) and objective beliefs. To measure the latter, we estimate "rational" forecasts using a macroeconomic vector autoregressive model. In line with our theory, we find strong evidence for time-variation in pessimism in the data: wedges are negative on average, indicating pessimism, and contract sharply during recessions, in line with sharply increased pessimism. Since our theory predicts a unique mapping between these wedges and the agent's distorted beliefs, we extract time-varying proxies of distorted beliefs from the data which we use in our calibration exercise.

Finally our fourth contribution concerns the effect of stochastic risk aversion on asset prices in a general equilibrium setting. To this end, we extend our set-up to generalized robustness based on the Cressie-Read penalty introduced earlier, but for an investor with stochastic differential utility or recursive preferences, allowing us to disentangle risk aversion and elasticity of intertemporal substitution even without robustness. This extension plays an important role in the general equilibrium asset pricing analysis with distorted beliefs, see, e.g., Li and Liu (2019), Jin and Sui (2019) and Nagel and Xu (2019) in the case of extrapolative expectations. We show that generalized robustness has several interesting effects. Not surprisingly, the unconditional increase in effective risk aversion helps to generate high risk premia and low risk-free rates due to a stronger precautionary savings motive. This mirrors existing results in the literature for robustness. Additionally, we also obtain endogenous time-variation in risk premia, with Sharpe ratios and equity premia increasing in bad times, along with stochastic volatility. Interestingly, Monte Carlo simulations reveal that the distribution of the state variable is skewed and heavy tailed in bad states of nature, reflecting the higher instantaneous volatility of beliefs in those states of the world.

Related Literature: We contribute to several strands of the literature. Our paper builds on the literature studying pessimistic subjective beliefs such as Hansen and Sargent (2001), Anderson, Hansen, and Sargent (2003), and Hansen, Sargent, Turmuhambetova, and Williams (2006), among many others. Different from this literature, which imposes an entropy penalty, we show that for values of $\eta \neq 1$, we generate time-varying beliefs and risk aversion. Moreover, our model calibration shows that time-varying pessimism induces rich dynamics in asset prices.

Two very recent papers consider extensions of entropy-based robustness in portfolio choice. Chamberlain (forthcoming) offers an excellent survey on empirical methods for robust portfolio choice as an example of econometric issues in decision making. He shows how dynamic $\phi-$divergence preferences can be used for purposes of sensitivity analysis when an investor fears misspecification. $\phi-$divergence measures nest Cressie-Read divergence measures. Balter, Horvath, and Maenhout (2019) show that robustness with a Cressie-Read penalty jettisons the intertemporally myopic behavior of Nature that is imposed by entropy-based robustness. Our work extends theirs by solving explicitly for optimal portfolio choice and dynamic general equilibrium, as well as calibrating the model based on empirical estimates of pessimism in survey data.

A large literature in asset pricing studies time-varying risk aversion. Most prominently, in the habit models of Constantinides (1990), Detemple and Zapatero (1991), and Campbell and Cochrane (1999) time-varying risk aversion is tightly linked to the level of consumption relative to its recent past history. Similarly, Bekaert, Engstrom, and Xing (2009) underscores the importance of time-varying risk aversion to explain the size of the equity premium, variation in equity returns, and long-horizon predictability of equity returns. While in these models time-varying risk aversion is exogenously imposed on the utility function of the representative agent, in our setting, stochastic risk aversion arises endogenously due to the agent's concern for model misspecification.

Another related strand of the literature studies Cressie-Read divergence measures to derive and estimate minimum-dispersion stochastic discount factor (SDF) bounds. In particular, these papers show duality between the optimal portfolios of asset returns and minimum dispersion SDFs, see, e.g., Almeida and Garcia (2017) and Orlowski, Sali, and Trojani (2019). Different from this literature, which does not impose any preference structure to derive SDFs and does not study optimal portfolios in general equilibrium, we derive optimal portfolios in a Merton (1969) setting and study the effect of time-varying risk aversion on equilibrium quantities.

We also contribute to the empirical literature that extracts proxies of optimism and pessimism from survey data. For example, Bhandari, Borovička, and Ho (2019) also model time-varying worst-case drift distortions and argue that macroeconomic survey data comes from the worst-case model. The authors document that US households indeed display pessimism when forecasting unemployment and inflation. This pessimism is then shown to be a major driver of movements in macroeconomic aggregates, in particular in the labor market. Szőke (2019) studies an economy where the agent is concerned about the persistence properties of her baseline model of consumption and inflation are misspecified. Using survey data on interest rates, he then documents that forecasts line up with those predicted from the model. Adam, Matveev, and Nagel (2019) ask whether survey expectations of stock returns reflect ambiguity aversion. To this end, the authors use several different surveys of individual investors, professional investors, and chief financial officers, and find that survey forecasts are often overly optimistic, not pessimistic. We extend this literature along several dimensions. First, we propose a theoretical asset pricing model in which dynamic pessimism is an endogenous outcome. Second, we estimate such pessimism in the data to inform our model calibration and we find that our model matches moments in equity markets well.

Outline of the paper: The rest of the paper is organized as follows. Section 1 provides the theoretical framework and studies the robust utility index. Section 2 examines the partial equilibrium portfolio problem. Section 3 develops a general equilibrium model with an Epstein and Zin (1989) representative agent with Cressie-Read generalized robustness. Section 4 estimates empirical proxies of pessimism that we use in Section 5 for model calibration. Finally, Section 6 concludes. To save space, we collect all proofs and further technical details in a separate appendix.

1 The Model

1.1 The Cressie-Read Penalty Function

The basis of our analysis is the Cressie-Read cost (or penalty) function. We fix a reference probability space $(\Omega, (\mathcal{F}_t)_{t\in[0,T]}, \mathbb{B})$ and a standard d-dimensional Brownian motion $B^{\mathbb{B}}$. Given a bounded d-dimensional vector-valued process u, an alternative model $\mathbb{U}$ parameterized by u is defined as follows. Let

$$Z_t = \exp\Big(- \int_0^t \tfrac{1}{2}|u_s|^2 ds - \int_0^t u_s' dB_s^{\mathbb{B}} \Big), \quad t \in [0, T], \tag{1}$$

be the Radon-Nikodym derivative. Because u is bounded, $\mathbb{E}^{\mathbb{B}}[Z_T] = 1$, and we define $\mathbb{U}$ via

$$\frac{d\mathbb{U}}{d\mathbb{B}}\Big|_{\mathcal{F}_T} = Z_T. \tag{2}$$

For each random outcome $\omega \in \Omega$, $Z_T(w)$ describes the adjustment of the probability for this outcome under $\mathbb{U}$ compared to $\mathbb{B}$. In particular, for an event $A \in \mathcal{F}_T$, its probability under $\mathbb{U}$ is given by $\mathbb{P}^{\mathbb{U}}(A) = \mathbb{E}^{\mathbb{B}}[Z_T 1_A]$. The larger $Z(\omega)$ is, the more weight is put on ω under the adjustment. The process Z also fixes the conditional probabilities. For $A \in \mathcal{F}_s$,

$$\mathbb{P}^{\mathbb{U}}(A|\mathcal{F}_t) = \mathbb{E}_t^{\mathbb{B}}[Z_{t,s} 1_A], \quad \text{for any } t \leq s.$$

Here $\mathbb{E}_t^{\mathbb{B}}[\cdot]$ denotes the conditional expectation $\mathbb{E}^{\mathbb{B}}[\cdot|\mathcal{F}_t]$ and $Z_{t,s} = Z_s/Z_t$.

Given a parameter $\theta > 0$ and two positive processes Ψ and Φ, we define the (conditional) *Cressie-Read divergence* between $\mathbb{U}$ and $\mathbb{B}$ on $[t, T]$ as

$$R_t^{\mathbb{U}} = \frac{1}{\theta \Phi_t} \mathbb{E}_t^{\mathbb{B}}\Big[\int_t^T e^{-\delta(s-t)} \Psi_s dD_{t,s}\Big], \tag{3}$$

where δ is a constant discount rate,

$$D_{t,s} = \phi(Z_{t,s}) \quad \text{and} \quad \phi(z) = \frac{1 - \eta + \eta z - z^{\eta}}{\eta(1-\eta)}, \quad \eta \in \mathbb{R} \setminus \{0, 1\}. \tag{4}$$

Here $\phi(\cdot)$ is the Cressie-Read divergence function (Cressie and Read (1984)). It is convex and satisfies $\phi(1) = 0$. For $\eta \in \{0, 1\}$, the function $\phi(z)$ is defined as the corresponding limit, with $\eta = 1$ being the Kullback and Leibler (1951) divergence (relative entropy), and $\eta = 0$ is known as Burg entropy (Burg (1972)) The class of Cressie-Read divergence functions also includes several other well-known cases, e.g., $\eta = 1/2$ corresponds to the Hellinger (1909) distance and $\eta = 2$ describes the modified χ^2-divergence. Figure 1 plots the Cressie-Read divergence function for a range of different values of η that we consider in this paper.

[Insert Figure 1 here]

In equation (4), $D_{t,s}$ measures the realized divergence between $\mathbb{U}$ and $\mathbb{B}$ on $[t, s]$. Its increment, $D_{t,s+\Delta s} - D_{t,s}$, is weighted by Ψ_s and discounted, before being integrated. The parameter θ describes the strength of the preference for robustness. As we shall see next, Φ is chosen to ensure that $R^{\mathbb{U}}$ has a recursive structure and is thus time-consistent. Moreover, the process Ψ is fixed to ensure the problem remains homothetic and scale-invariant as in Maenhout (2004). In the case when $\Psi = \Phi = 1$, we can rewrite equation (3) using integration-by-parts as

$$R_t^{\mathbb{U}} = \frac{1}{\theta} \mathbb{E}_t^{\mathbb{B}}\Big[\int_t^T \delta e^{-\delta(s-t)} D_{t,s} ds + e^{-\delta(T-t)} D_{t,T}\Big].[4] \tag{5}$$

When $\eta = 1$, $R^{\mathbb{U}}$ in (5) is exactly the entropy divergence introduced by Hansen and Sargent (2001).

The Cressie-Read divergence in (3) may not be time-consistent for arbitrarily chosen Φ_t, because $D_{t,s}$ depends on t, i.e., the time at which the conditional expectation is computed. For our applications, we need $R^{\mathbb{U}}$ to satisfy a recursive structure. The following lemma characterizes a Φ that ensures the desired recursivity.

[4] Equation (5) shows that it is without loss of generality to restrict alternative models $\mathbb{U}$ to be absolutely continuous with respect to $\mathbb{B}$, otherwise, both $Z_{t,s}$ and $D_{t,s}$ could be infinite with positive probability, hence $R^{\mathbb{U}}$ is ill-defined.

Lemma 1. *When* $\Phi_t = Z_t^{1-\eta}$ *and* $\mathbb{E}^{\mathbb{B}}\big[\int_0^T e^{-\delta s}|\Psi_s|^p ds\big] < \infty$ *for some* $p > 2$*, then* $R^{\mathbb{U}}$ *in* (3) *becomes*

$$R_t^{\mathbb{U}} = \frac{1}{2\theta}\mathbb{E}_t^{\mathbb{U}}\Big[\int_t^T e^{-\delta(s-t)}\Psi_s Z_s^{\eta-1}|u_s|^2 ds\Big] \tag{6}$$

and it satisfies

$$R_t^{\mathbb{U}} = \mathbb{E}_t^{\mathbb{U}}\Big[\int_t^{\tilde{t}} e^{-\delta(s-t)}\frac{1}{2\theta}\Psi_s Z_s^{\eta-1}|u_s|^2 ds + e^{-\delta(\tilde{t}-t)}R_{\tilde{t}}^{\mathbb{U}}\Big], \qquad \textit{for any } \tilde{t} \geq t.$$

Our construction of $R^{\mathbb{U}}$ satisfies Theorem 2 (b) in Maccheroni, Marinacci, and Rustichini (2006b) with $\gamma_t(\mathbb{U}) = \frac{1}{2\theta}\mathbb{E}_t^{\mathbb{U}}\big[\int_t^{t+1} e^{-\delta(s-t)}\Psi_s Z_s^{\eta-1}|u_s|^2 ds\big]$ as the one-period-ambiguity index. Theorem 2 in Maccheroni, Marinacci, and Rustichini (2006b) implies that $\{R_t^{\mathbb{U}}\}$ is a recursive ambiguity index.

Equation (6) shows that the Cressie-Read divergence not only depends on the preference parameter θ, but also on the state-dependent weight $Z^{\eta-1}$ when $\eta \neq 1$. Importantly, notice that Z is the agent's Radon-Nikodym derivative or cumulative distorted belief process. Therefore the Cressie-Read divergence criterion crucially embeds a nonlinear function $Z^{\eta-1}$ of the agent's distorted beliefs. As can easily be seen, when $\eta = 1$, the Cressie-Read divergence becomes the entropy divergence and loses its state-dependence. Beyond the special case of relative entropy, Z therefore matters. When $Z^{\eta-1}$ is large for some states, deviating from the reference model $\mathbb{B}$ is costly for these states. This state-dependence of the penalty function is the key driver of our results and represents the main generalization due to the Cressie-Read specification. We demonstrate below that this is the mechanism generating endogenous time-variation in risk aversion. In particular, depending on the sign of $(\eta-1)$, the nonlinear function $Z^{\eta-1}$ of the agent's distorted beliefs either increases or decreases following a favorable shock, which endogenously affects sentiment and thereby the agent's effective risk aversion, as well as the instantaneous volatility of the agent's beliefs.

For the numerical analysis of our portfolio choice and general equilibrium problems we will consider an approximation of (6)

$$R_t^{\mathbb{U}} = \frac{1}{2\theta}\mathbb{E}_t^{\mathbb{U}}\Big[\int_t^T e^{-\delta(s-t)}\Psi_s Z_{s\wedge\tau}^{\eta-1}|u_s|^2 ds\Big], \tag{7}$$

where $0 < \underline{z} < \overline{z}$ are constants and

$$\tau = \inf\{t \geq 0 : Z_t \leq \underline{z} \text{ or } Z_t \geq \overline{z}\}. \tag{8}$$

The interval $[\underline{z}, \overline{z}]$ contains all plausible adjustments, so that $Z < \underline{z}$ ($Z > \overline{z}$) is regarded as unreasonably underweighted (overweighted) in our model. The constant $\underline{z}$ ($\overline{z}$) is fixed to be sufficiently close to zero (large) so that only extreme weights are excluded and the probability of $\tau < T$ can be made arbitrarily small. When $t > \tau$,

$$R_t^{\mathbb{U}} = \frac{Z_\tau^{\eta-1}}{2\theta}\mathbb{E}_t^{\mathbb{U}}\Big[\int_t^T e^{-\delta(s-t)}\Psi_s|u_s|^2 ds\Big],$$

which can be viewed as an entropy penalty, but with preference for robustness given by $\theta Z_\tau^{1-\eta}$ rather than by θ. This specification helps to pin down boundary conditions for portfolio choice and general equilibrium problems studied later.

1.2 The Utility Index Process with Generalized Robustness: Understanding the Mechanism

The main mechanism generating dynamic pessimism and belief- and state-dependent effective risk aversion that drive the key results in the paper can already be understood from an analysis of the utility index process.

To this end, consider a consumption process c and an intertemporal utility function U. We then define a utility index $\mathcal{U}^c$ for c as

$$\mathcal{U}_t^c = \inf_u \mathbb{E}_t^{\mathbb{U}}\Big[\int_t^T e^{-\delta(s-t)}\delta U(c_s)ds + R_t^{\mathbb{U}} + e^{-\delta(T-t)}\epsilon U(c_T)\Big], \tag{9}$$

where ϵU, with a positive constant ϵ, is the bequest utility and $R^{\mathbb{U}}$ is given in (7). The utility index reflects the alternative measure $\mathbb{U}$, based on an endogenous belief distortion u. The infimization with respect to u addresses the agent's concern of model misspecification and the role of $\{R_t^{\mathbb{U}}\}$ is to constrain the choice of u by penalizing distortions that are too large and deemed unreasonable.

In order to derive the dynamics of $\mathcal{U}^c$ and the associated optimal measure $\mathbb{U}$, consider now for any fixed u the associated utility index

$$\mathcal{U}_t^{c,u} = \mathbb{E}_t^{\mathbb{U}}\Big[\int_t^T e^{-\delta(s-t)}\delta U(c_s)ds + R_t^{\mathbb{U}} + e^{-\delta(T-t)}\epsilon U(c_T)\Big].$$

The martingale representation theorem ensures the existence of a vector-valued process Γ such that

$$d\mathcal{U}_t^{c,u} = \big[\delta\mathcal{U}_t^{c,u} - \delta U(c_t)\big]dt - \Big\{\tfrac{1}{2\theta}\Psi_t Z_{t\wedge\tau}^{\eta-1}|u_t|^2 - \Gamma_t' u_t\Big\}dt + \Gamma_t' dB_t^{\mathbb{B}}.$$

Because $\mathcal{U}^{c,u}$ is subject to the terminal condition $\mathcal{U}_T^{c,u} = \epsilon U(c_T)$, $\mathcal{U}^{c,u}$ satisfies a backward stochastic differential equation (BSDE). The comparison theorem for BSDEs of El Karoui, Peng, and Quenez (1997) implies that $\mathcal{U}_t^c = \inf_u \mathcal{U}_t^{c,u}$ satisfies

$$d\mathcal{U}_t^c = \big[\delta\mathcal{U}_t^c - \delta U(c_t)\big]dt - \inf_u\Big\{\tfrac{1}{2\theta}\Psi_t Z_{t\wedge\tau}^{\eta-1}|u_t|^2 - \Gamma_t' u_t\Big\}dt + \Gamma_t' dB_t^{\mathbb{B}}, \quad \mathcal{U}_T^c = \epsilon U(c_T). \tag{10}$$

The first-order condition in u yields the optimal distortion

$$u_t^* = \frac{\theta\Gamma_t}{\Psi_t Z_{t\wedge\tau}^{\eta-1}}. \tag{11}$$

Finally, combining (10) and (11), we get that

$$d\mathcal{U}_t^c = \Big[\delta\mathcal{U}^c - \delta U(c_t) + \frac{\theta}{2}\frac{1}{\Psi_t Z_{t\wedge\tau}^{\eta-1}}|\Gamma_t|^2\Big]dt + \Gamma_t' dB_t^{\mathbb{B}}, \quad \mathcal{U}_T^c = \epsilon U(c_T). \tag{12}$$

This is a so-called quadratic BSDE and can be seen as a new form of generalized stochastic differential utility introduced by Lazrak and Quenez (2003). When $U(c)$ is bounded and Ψ is bounded away from zero, existence and uniqueness of $\mathcal{U}^c$ is ensured by Kobylanski (2000).[5]

[5] The boundedness assumption is not generally satisfied in our portfolio and general equilibrium problems. Existence of the generalized stochastic differential utility is treated individually there.

In the terminology of Duffie and Epstein (1992), $|\Gamma_t|^2$ enters the utility index via a variance multiplier, which represents a utility penalty as a multiple of the utility volatility. This can be seen from the integral form of (12)

$$\mathcal{U}_t^c = \mathbb{E}_t^{\mathbb{U}}\Big[e^{-\delta(T-t)}\epsilon U(c_T) + \int_t^T e^{-\delta(s-t)}\delta U(c_s)\,ds - \int_t^T e^{-\delta(s-t)}\frac{\theta}{2}\frac{1}{\Psi_s Z_{s\wedge\tau}^{\eta-1}}|\Gamma_s|^2\,ds\Big], \tag{13}$$

where the penalty is determined by the variance multiplier $\frac{\theta}{2}\frac{1}{\Psi Z^{\eta-1}}$. Economically speaking, this factor captures risk aversion, i.e. the desire to smooth across states of nature. The Cressie-Read generalization of robustness makes this variance multiplier belief-dependent and state-dependent, through $Z^{\eta-1}$. To build intuition, we discuss the properties of $\mathcal{U}^c$ in the following proposition.

Proposition 1.

1. *If $Z^{\eta-1} \geq \widetilde{Z}^{\eta-1}$, then $\mathcal{U}^c \geq \widetilde{\mathcal{U}}^c$.*
2. *A component of u^* is positive if and only if the corresponding component of Γ is positive. For a fixed and positive Γ, all components of u^* decrease as $Z^{\eta-1}$ increases.*
3. *When a component of u^* is positive, positive shocks to the corresponding component in $B^{\mathbb{B}}$ decrease Z, hence increase $Z^{\eta-1}$ when $\eta < 1$, or decrease $Z^{\eta-1}$ when $\eta > 1$.*

The first result in the proposition shows that larger $Z^{\eta-1}$ implies a smaller variance multiplier and therefore an increase in the agent's utility. As discussed before, $Z^{\eta-1}$ is the nonlinear function of the agent's distorted belief that introduces time-varying pessimism. To illustrate the intuition, let us focus on a 1-dimensional case, i.e., $d = 1$. To see how we can interpret this as pessimism, we notice from the second result that the optimal distortion u^* is positive when Γ is positive.[6] Under $\mathbb{U}$, $dB_t^{\mathbb{B}} = -u_t^* dt + dB_t^{\mathbb{U}}$ where $B^{\mathbb{U}}$ is a Brownian motion under $\mathbb{U}$. When $u^* > 0$, the expected growth rate of $B^{\mathbb{B}}$ under $\mathbb{U}$ is underestimated relative to the growth rate under $\mathbb{B}$. Therefore, the agent with subjective belief $\mathbb{U}$ is pessimistic compared to the reference model $\mathbb{B}$.

The third result explains how the dynamics of fundamental shocks affect the distorted belief, which in turn affects the variance multiplier, risk aversion, and the volatility of $\log(Z)$. Importantly, the effect crucially depends on whether η is greater than or smaller than 1. Consider $\eta < 1$. Positive fundamental shocks to $B^{\mathbb{B}}$ increase $Z^{\eta-1}$ due to the third result. This decreases the variance multiplier $\frac{\theta}{2}\frac{1}{\Psi Z^{\eta-1}}$, and thereby lowers the agent's risk aversion as well. Furthermore, by the second result, the instantaneous optimal u^* decreases following increasing $Z^{\eta-1}$, making the agent less pessimistic following positive fundamental shocks to $B^{\mathbb{B}}$. From (1), we see that u^* is also the volatility of the (log) belief distortion $\log(Z)$. Therefore, another effect of the positive fundamental shock is to reduce the volatility of $\log(Z)$. On the other hand, negative fundamental shocks to $B^{\mathbb{B}}$ decrease $Z^{\eta-1}$ and hence increase u^*, which exacerbates the agent's pessimism and also increases the volatility of $\log(Z)$. Generalized robustness based on Cressie-Read divergence with $\eta < 1$ can therefore be seen as endogenously generating expectation and belief dynamics that resemble diagnostic expectations Bordalo, Gennaioli, and Shleifer (2018) or extrapolative expectations, where agents extrapolate from

[6] Γ is shown to be positive in our applications later.

recent experience when forming expectations, see, e.g., Hong and Stein (1999), Barberis, Greenwood, Jin, and Shleifer (2015), Barberis (2018), Jin and Sui (2019), Li and Liu (2019), and Nagel and Xu (2019).

In contrast, when $\eta > 1$, the agent can be labeled contrarian. Positive shocks to fundamentals $B^{\mathbb{B}}$ decrease $Z^{\eta-1}$, hence increases u^*, i.e., the agent becomes more pessimistic after positive shocks to fundamentals, while negative fundamental shocks tend to reduce pessimism.

In our applications, we use survey data to investigate empirically how pessimism evolves over the business cycle. Recall that $\eta < 1$ leads to countercyclical risk aversion, pessimism, and volatility of beliefs, whereas $\eta > 1$ is associated with procyclical risk aversion, pessimism, and belief volatility. Table 1 summarizes these results.

[Insert Table 1 here]

Example: The intuition discussed above can be illustrated by the following simplified two-stage example, which allows for an explicit closed-form solution of the BSDE. The first stage is from time 0 to 1; the second stage is from time 1 to 2. Set $\theta = \Psi = 1$ and $\delta = 0$ in (7). To focus on the intuition, we also set $\underline{z} = 0$ and $\overline{z} = \infty$ so that $\tau = \infty$. The agent freezes $Z^{\eta-1}$ in both stages and only updates it at time 1. Therefore, $Z_t^{\eta-1} = Z_0^{\eta-1}$ when $t \in [0, 1)$; $Z_t^{\eta-1} = Z_1^{\eta-1}$ when $t \in [1, 2]$. This leads to the following Cressie-Read penalty function

$$R_0^{\mathbb{U}} = \mathbb{E}^{\mathbb{U}}\Big[\int_0^1 \frac{1}{2} Z_0^{\eta-1}|u_s|^2 ds + \int_1^2 \frac{1}{2} Z_1^{\eta-1}|u_s|^2 ds\Big]. \tag{14}$$

Comparing (14) to (6) with $\eta = 1$ therein, we can view (14) as a combination of two entropy divergences with different Ψ in different stages. In the first stage, Ψ in (6) can be identified as $Z_0^{\eta-1} = 1$; in the second stage, Ψ can be identified as $Z_1^{\eta-1}$. There is no intertemporal utility and the agent consumes only at time 2, where the utility of consumption is $\epsilon U(c_2) = B_2^{\mathbb{B}}$.

In this case, the BSDE given in (12) simplifies to

$$\begin{aligned} d\mathcal{U}_t^c &= \tfrac{1}{2}\Gamma_t^2 dt + \Gamma_t dB_t^{\mathbb{B}}, \quad && t \in [0,1), \\ d\mathcal{U}_t^c &= \tfrac{1}{2} Z_1^{1-\eta}\Gamma_t^2 dt + \Gamma_t dB_t^{\mathbb{B}}, \quad && t \in [1,2], \end{aligned}$$

with terminal condition $\mathcal{U}_2^c = B_2^{\mathbb{B}}$. This BSDE has the following explicit solution

$$\begin{aligned} \mathcal{U}_t^c &= B_t^{\mathbb{B}} - \tfrac{1}{2} Z_1^{1-\eta}(2-t), \quad && \Gamma_t = 1, \quad t \in [1,2], \\ \mathcal{U}_t^c &= B_t^{\mathbb{B}} - \tfrac{1}{2} e^{-\frac{1}{2}\eta(1-\eta)(1-t)} Z_t^{1-\eta} - \tfrac{1}{2}(1-t), \quad && \Gamma_t = 1, \quad t \in [0,1). \end{aligned}$$

From (11), the optimal distortion is

$$\begin{aligned} u_t^* &= 1, \quad t \in [0,1), \\ u_t^* &= Z_1^{1-\eta}, \quad t \in [1,2]. \end{aligned} \tag{15}$$

This shows that u^* is positive in this case. From the explicit solution we can make the following two observations when $\eta < 1$ (all effects are reversed when $\eta > 1$): First, a positive shock to $B_1^{\mathbb{B}}$ decreases Z_1 and also $Z_1^{1-\eta}$ (as $\eta < 1$). Therefore, a positive shock to $B_1^{\mathbb{B}}$ decreases u^* in the second stage; in other words, the agent becomes less pessimistic following positive shocks. Second, Z is a state variable for the utility process when $\eta \neq 1$. The utility does not depend on this state variable when $\eta = 1$.

In summary, we can think of the Cressie-Read divergence as a collection of entropy divergence functions with a dynamic belief- and state-dependent weight $Z^{\eta-1}$. This dynamic weight depends endogenously on the optimal u^*. When $\eta < 1$, the weight increases when fundamentals improve, making the agent less pessimistic.

One way of understanding the intuition behind the effects at work is to view the Cressie-Read penalty function as reflecting a preference of Nature (in this case the fictitious malevolent agent in the max-min expected utility interpretation) for intertemporal smoothing of the process for the instantaneous distortions u^*. The smaller η, the stronger the preference for intertemporal smoothing, such that the decision-maker naturally expects more adverse distortions from Nature in bad times and less adverse distortions in good times. Loosely speaking, when $\eta < 1$, the agent's subjective beliefs exhibit "momentum-like" dynamics, and pessimism gets worse in bad times, while beliefs improve following positive shocks.

In contrast, when $\eta > 1$, Nature does not have a preference for intertemporal smoothing of the distortions. In this case, the agent expects less adverse distortions in bad times and more adverse distortions following positive shocks. In some sense, the agent could be seen as interpreting a recent adverse shock as disguising an adverse distortion $u^\star$ and as thinking that "lightning never strikes twice" i.e., if Nature has just used its ammunition in hitting the agent with a negative shock, it will not do so again immediately. Put differently, now the agent's subjective beliefs exhibit reversal-like dynamics.

The knife-edge case of $\eta = 1$ corresponds to the case where Nature allocates distortions myopically, so that market conditions have no impact on the distortion. The agent's belief distortion state variable Z becomes irrelevant and pessimism is now state-independent. This is the case of entropy. As shown in Balter, Horvath, and Maenhout (2019) the entropy penalty greatly simplifies the analysis by reducing the general Cressie-Read power function to its logarithmic limit, which makes Nature behave myopically. In the following portfolio choice problem, our primary focus is on $\eta < 1$, however, we also compare it to the case where $\eta > 1$.

2 Portfolio Choice

Before solving for asset prices in equilibrium, we build intuition in a partial equilibrium setting in the context of optimal portfolio choice.

2.1 The Consumption and Portfolio Choice Problem

Consider a capital market with a risk-free bond with a constant interest rate r and d risky assets whose prices follow

$$dS_t = \text{diag}(S_t)(\mu dt + \sigma dB_t^{\mathbb{B}}),$$

where μ is a constant d-dimensional vector representing expected returns, σ is a constant $d \times d$-matrix describing the return volatilities, and diag(S) is a d-dimensional diagonal matrix with elements

$\{S^1, \ldots, S^d\}$.

When the agent invests her wealth in the risky assets based on a vector of portfolio weights π and consumes at a rate c, her wealth evolves according to the following stochastic differential equation

$$dW_t = \big[rW_t + W_t\pi_t'(\mu - r) - c_t\big]dt + W_t\pi_t'\sigma dB_t^{\mathbb{B}}. \tag{16}$$

The agent chooses her optimal strategy (π, c) to maximize the expected utility of consumption plus the Cressie-Read cost, but Nature chooses u to minimize the same objective function. The optimization problem is a stochastic differential game

$$\inf_u \sup_{\pi,c} \mathbb{E}_t^{\mathbb{U}}\Big[\int_t^T e^{-\delta(s-t)}\delta U(c_s)ds + \epsilon e^{-\delta(T-t)}U(c_T) + R_t\Big], \tag{17}$$

where R is given in (7). Here, for both intertemporal and bequest utility, we choose isoelastic or CRRA utility $U(c) = \frac{c^{1-\gamma}}{1-\gamma}$ with coefficient of relative risk aversion $0 < \gamma \neq 1$.[7] We denote the optimal value in (17) by V_t.

To maintain homotheticity and scale-invariance of the problem (17), we follow Maenhout (2004) and choose

$$\Psi = (1-\gamma)V. \tag{18}$$

Then problem (17) can be transformed to

$$V_t = \inf_u \sup_{\pi,c} \mathbb{E}_t^{\mathbb{U}}\Big[\int_t^T e^{-\delta(s-t)}\Big(\delta U(c_s) + \frac{1-\gamma}{2\theta}V_s Z_{s\wedge\tau}^{\eta-1}|u_s|^2\Big)ds + \epsilon e^{-\delta(T-t)}U(c_T)\Big], \tag{19}$$

where τ is the stopping time from (8). Because the value V shows up on both sides of (19), it can be considered as an optimization problem for a (generalized) stochastic differential utility.

2.2 A Two-Stage Example

To understand the impact of the Cressie-Read penalty on the agent's portfolio choice, we again consider a simplified two-stage problem: stage 1 starts from time 0 and ends at time 1, stage 2 starts from time 1 until ∞. The agent freezes Z in both stages and only updates Z at time 1. Then the agent chooses the initial u_0 at time 0 and keeps it constant until time 1. At time 1, the agent updates it to u_1 and then keeps it constant forever. Similar to the two-stage problem discussed in Section 1.2, we set $\underline{z} = 0$ and $\overline{z} = \infty$ so that $\tau = \infty$.

We solve the two-stage problem by backward induction. In the second stage, the agent's optimal consumption and investment problem is

$$V_t = \inf_{u_1} \sup_{\pi,c} \mathbb{E}_t^{\mathbb{U}}\Big[\int_t^\infty e^{-\delta(s-t)}\Big(\delta U(c_s) + \tfrac{1-\gamma}{2\theta}V_s Z_1^{\eta-1}|u_1|^2\Big)ds\Big], \quad t \geq 1, \tag{20}$$

subject to (16).

In the problem above, the agent does not update her belief after time 1. The process Z in the Cressie-Read cost function is frozen at Z_1. Problem (20) is equivalent to Maenhout (2004) for an

[7] The case of log utility is discussed in Appendix B.

entropy penalty, but with robustness preference parameter $\theta Z_1^{1-\eta}$ instead of θ. The reason for the equivalence is that even though u is allowed to be stochastic there, and not fixed as here, the optimal u turns out to be constant. Defining $\Sigma = \sigma\sigma'$, the optimal portfolio weight and belief distortion are

$$\pi_1 = \frac{\mu - r}{\Sigma}\frac{1}{\gamma + \theta Z_1^{1-\eta}}, \tag{21}$$

$$u_1 = \frac{\mu - r}{\sigma}\frac{\theta}{\gamma Z_1^{\eta-1} + \theta}. \tag{22}$$

We can see from (21) that the agent's implicit risk aversion is $\gamma + \theta Z_1^{1-\eta}$, where $Z_1 = \exp(-\frac{1}{2}|u_0|^2 - u_0 B_1^{\mathbb{B}})$. Suppose that $u_0 > 0$, we first consider the case $\eta < 1$. After positive return shocks to $B_1^{\mathbb{B}}$, Z_1 decreases, hence $Z_1^{1-\eta}$ decreases as well (due to $\eta < 1$). As a result, the agent's implicit risk aversion $\gamma + \theta Z_1^{1-\eta}$ decreases, which in turn increases the agent's optimal portfolio weight π_1. Meanwhile, the agent's belief distortion u_1 decreases, generating a less pessimistic expected return $\mu - \sigma u_1$ under the subjective measure $\mathbb{U}$. In summary, after a favorable return shock, the agent becomes less pessimistic and less risk-averse. Negative return shocks lead to the opposite phenomenon: the investor becomes more pessimistic and more risk-averse.

Under extremely good market conditions, $Z_1^{1-\eta}$ tends to zero in the limit. The agent's effective risk aversion gets close to its minimal value γ and the optimal u_1 converges to 0. This implies that the agent is no longer pessimistic in the limit. This can also be seen from (20) where the Cressie-Read penalty becomes extremely large for nonzero u, due to the large weight $Z_1^{\eta-1}$, so that it is very costly for the agent to deviate from the reference measure $\mathbb{B}$. Under extremely unfavorable market conditions, $Z_1^{1-\eta}$ explodes to infinity, resulting in an infinitely risk-averse investor who completely shuns the risky asset. In this case $u_1 = \frac{\mu - r}{\sigma}$, which means the equity premium is zero under the agent's extremely pessimistic subjective view.

When $\eta > 1$, but u_0 is still positive, the agent's implicit risk aversion $\gamma + \theta Z_1^{1-\eta}$ increases after positive return shocks, with the investor reducing her portfolio weight in the risky asset. Meanwhile, u_1 increases and the agent becomes more pessimistic after positive return shocks.

We calculate the value for the agent at time 1 to be

$$V_1 = \frac{W_1^{1-\gamma}}{1-\gamma} e^{f_1(Z_1^{1-\eta})},$$

where

$$f_1(Z_1^{1-\eta}) = \gamma \log(\gamma \delta^{\frac{1}{\gamma}}) - \gamma \log\Big(\delta + (\gamma - 1) r + \frac{\gamma - 1}{2}\frac{(\mu - r)^2}{\Sigma}\frac{1}{\gamma + \theta Z_1^{1-\eta}}\Big).$$

Now we turn to the agent's problem in the first stage. Given the agent's continuation utility V_1, the agent's problem in the first stage is

$$V_t = \inf_{u_0}\sup_{\pi, c} \mathbb{E}_t^{\mathbb{U}}\Big[\int_t^1 e^{-\delta(s-t)}\Big(\delta u(c_s) + \frac{1-\gamma}{2\theta} V_s Z_0^{\eta-1}|u_0|^2\Big) ds + e^{-\delta(1-t)}\frac{W_1^{1-\gamma}}{1-\gamma} e^{f_1(Z_1^{1-\eta})}\Big], \quad 0 \le t \le 1. \tag{23}$$

Here $Z_0^{\eta-1} = 1$ and $f_1(Z_1^{1-\eta})$ indicates the dependence of the constant f_1 on $Z_1^{1-\eta}$. If $\eta = 1$, the continuation utility is state-independent, then problem (23) is again equivalent to a problem with the entropy cost and the optimal solution is myopic. When $\eta \neq 1$, the continuation utility is state-dependent, making problem (23) state-dependent as well. The agent has an intertemporal hedging demand in the first stage against market condition fluctuation and her changing belief at time 1.

The intuition gained from this two-stage model survives in the fully dynamic model developed in the next subsection.

2.3 Dynamic Optimal Consumption and Portfolio Choice

We have seen that Z should be the state variable for the problem (19). For ease of interpretation and exposition we now introduce a monotone transformation

$$x_t = -\log Z_t. \tag{24}$$

We call x the *market sentiment variable* and take it as the state variable for problem (19). Because of (1) the dynamics of x are described by the following stochastic differential equation

$$dx_t = \tfrac{1}{2}|u_t|^2 dt + u_t' dB_t^{\mathbb{B}} = -\tfrac{1}{2}|u_t|^2 dt + u_t' dB_t^{\mathbb{U}}. \tag{25}$$

We put a negative sign on the right-hand side of (24) so that positive fundamental shocks to $B^{\mathbb{B}}$ increase the market sentiment variable x, when u is a vector with positive components.[8] The stopping time τ in (8) is reformulated as $\tau = \inf\{t \geq 0 : x_t \leq \underline{x} \text{ or } x_t \geq \overline{x}\}$, where $\underline{x} = -\log \overline{z}$ and $\overline{x} = -\log \underline{z}$.

The choice of Ψ ensures the following decomposition of the optimal value function

$$V_t = \frac{W_t^{1-\gamma}}{1-\gamma} e^{f(t,x_t)}. \tag{26}$$

We obtain the Hamilton-Jacobi-Bellman (HJB) equation satisfied by f using dynamic programming and summarize the agent's optimal investment and consumption strategies in proposition 2.

Proposition 2. *When $\gamma \in (0,1)$, the function f defined in (26) satisfies*

$$\begin{aligned} 0 = \inf_u \sup_{\pi,\tilde{c}} \Big\{ &\partial_t f + \tfrac{1}{2}|u|^2 \Big(\partial_{xx}^2 f - \partial_x f + (\partial_x f)^2\Big) + (1-\gamma)\partial_x f \pi' \sigma u + \delta \tilde{c}^{1-\gamma} e^{-f} \\ &+ (1-\gamma)\big[r + \pi'(\mu - r - \sigma u) - \tilde{c} - \tfrac{1}{2}\gamma \pi' \Sigma \pi\big] - \delta + \tfrac{1-\gamma}{2\theta} e^{(1-\eta)x} |u|^2 \Big\}, \end{aligned} \tag{27}$$

for $(t,x) \in [0,T) \times (\underline{x}, \overline{x})$, with the boundary conditions

$$f(t,\underline{x}) = f_{\underline{x}}^{ent}(t), \quad f(t,\overline{x}) = f_{\overline{x}}^{ent}(t), \quad \text{and} \quad f(T,x) = \log \epsilon. \tag{28}$$

When $\gamma > 1$, the infimum and supremum in (27) are changed to $\sup_u \inf_{\pi,\tilde{c}}$. If $\Sigma = \sigma\sigma'$ is positive definite and $\partial_{xx}^2 f - \partial_x f + (\partial_x f)^2 + \frac{1-\gamma}{\theta} e^{(1-\eta)x} > 0$ when $\gamma \in (0,1)$ (resp. < 0 when $\gamma > 1$), for any $(t,x) \in [0,T) \times (\underline{x}, \overline{x})$, then the agent's optimal belief and strategies are given by

$$\pi^* = \frac{1}{\gamma}\Sigma^{-1}\Big(\mu - r - (1 - \partial_x f)\sigma u^*\Big), \tag{29}$$

$$u^* = \frac{(1-\gamma)(1-\partial_x f)}{\partial_{xx}^2 f - \partial_x f + (\partial_x f)^2 + \frac{1-\gamma}{\theta} e^{(1-\eta)x}} \sigma' \pi^*, \tag{30}$$

$$\tilde{c}^* = \delta^{\psi} e^{-\frac{1}{\gamma} f}. \tag{31}$$

[8] As we see later, u has positive components in our applications.

The function f_x^{ent} in (28) is the value for the problem with an entropy penalty and the robustness parameter $\theta e^{(\eta-1)x}$, as specified in Proposition 5 in Appendix A.

The optimal portfolio weight in (29) generalizes our finding for the two-stage example, and shows the impact of the Cressie-Read penalty function. Without a preference for robustness, CRRA utility as well as Epstein-Zin preferences in combination with i.i.d. returns lead to the well-known myopic portfolio, as first shown by Samuelson (1969) and Merton (1969), namely $\pi^* = \frac{1}{\gamma}\Sigma^{-1}(\mu - r)$. For entropy-based robustness, the main effect in this context is to increase the effective risk aversion, replacing γ by $\gamma + \theta$ (Maenhout (2004)). We obtain two important new effects by considering the Cressie-Read penalty function. First, the investor anticipates future changes in beliefs and the corresponding changes in perceived investment opportunities. This induces a Merton-type intertemporal hedging demand, which is captured by the term $\frac{1}{\gamma}\Sigma^{-1}\partial_x f \sigma u^*$, added to the mean-variance optimal portfolio $\frac{1}{\gamma}\Sigma^{-1}(\mu - r - \sigma u^*)$ under $\mathbb{U}$. Second, importantly, effective risk aversion becomes belief- and state-dependent, driven by the endogenous sentiment state variable. Combining (29) and (30), we obtain

$$\pi^* = \frac{1}{\gamma^{\text{eff}}}\Sigma^{-1}(\mu - r), \quad \text{where } \gamma^{\text{eff}} = \gamma + \frac{(1-\gamma)(1-\partial_x f)^2}{\partial^2_{xx} f - \partial_x f + (\partial_x f)^2 + \frac{1-\gamma}{\theta}e^{(1-\eta)x}} \tag{32}$$

which represents the investor's effective risk aversion. This effective risk aversion is belief- and state-dependent, in line with the discussion of the BSDE for the utility index process in section 1.2 earlier.

We now study the optimal portfolio weight in more detail numerically. We numerically solve the HJB equation (27) together with the associated boundary conditions (28) using finite difference methods with implicit schemes. We focus on a single risky asset ($d = 1$) and use the parameters listed in Table 2.

[Insert Table 2 here]

We start by considering a short horizon T of one year in order to focus on the effect of the Cressie-Read penalty on the myopic portfolio component. Following the logic above, we expect a portfolio allocation close to the entropy case when the state variable x equals 0. In the limit when the investor's horizon shrinks to zero, the intertemporal hedging component vanishes and only the myopic component remains. When the sentiment state variable x is zero, Figure 2 shows that generalized robustness with Cressie-Read divergence produces the same optimal distortion u^* and the same portfolio weight π^*, for any value of η. This corresponds to what we find in the second stage of the two-stage example with $Z = 1$ for u_1 in (22) and for π_1 in (21). Figure 2 panel (a) illustrates the effects when both risk aversion, γ, and preference for robustness, θ, are set to 2. To put the quantitative results in perspective, assuming an equity premium of $\mu - r = 0.07$ and volatility $\sigma = 0.2$, CRRA expected utility ($\theta = 0$) generates a portfolio weight $\pi = 0.775$. Entropy-based robustness ($\eta = 1$) reduces this equity allocation to exactly half, reflecting a pessimistically distorted belief that the equity premium is half its observed value.

[Insert Figures 2, 3 and 4 here]

For $\eta < 1$, when x increases and sentiment improves, effective risk aversion shrinks, as discussed above. Moreover, risk aversion shrinks faster the smaller is η, which explains the steeper slope for $\eta = 0$ than for $\eta = 0.5$. When $\eta > 1$, on the other hand, risk aversion is procyclical, as is apparent from the negatively sloped portfolio rule for $\eta = 1.5$ and a fortiori for $\eta = 2$. All optimal portfolio allocations are supported by the optimal distortions in the left panels of Figure 2. The larger the optimal distortion, the more pessimistic the agent's subjective beliefs and therefore the more cautious the corresponding portfolio allocation.

For higher risk aversion and robustness parameters $\gamma = \theta = 5$, the effects and patterns remain the same as illustrated in Figure 2 panel (b). As can be shown analytically for the entropy case, and as argued in the two-stage example in (22), the optimal distortion does not change when both γ and θ increase by the same factor. This explains why the left panel is unchanged compared to the left panel in panel (a). The portfolio allocation, however, is evidently shifted down when risk aversion and preference for robustness both increase. It is interesting to also study the implications of an increase in the robustness parameter θ without changing risk aversion γ. Figure 2 panel (c) demonstrates that the optimal distortion increases and as a result the portfolio allocation shrinks.

We now turn to a long horizon with T up to 100 years in order to study intertemporal hedging and resulting horizon effects. The seminal work of Merton (1973) explains how non-myopic investors tilt their portfolio in order to hedge against future changes in the investment opportunity set. For example, if returns on a risky asset are contemporaneously negatively correlated with expected returns on that asset, it becomes less risky to hold over longer horizons, inducing investors with longer horizons to increase their holdings. Despite returns being i.i.d. in our setting, so that expected returns are constant, intertemporal hedging appears because of the investor's distorted beliefs. Even though expected returns are constant under the reference measure $\mathbb{B}$, they are not constant under $\mathbb{U}$ when $\eta \neq 1$. In particular, because of time-varying sentiment, belief distortion u reacts to return innovations. When pessimism is countercyclical ($\eta < 1$), positive return shocks are associated with improvements in sentiment, which reduces pessimism and increases the expected return on the risky asset. In other words, for $\eta < 1$, returns and expected returns under $\mathbb{U}$ are positively correlated, which leads to negative intertemporal hedging demands that grow with the horizon. This effect is clear and quite pronounced in Figure 3 for $\eta = 0.5$, and especially for $\eta = 0$ (Note that all results are reported for $x = 0$). With the same logic, and remembering that for $\eta > 1$ an increase in the state variable x as a result of a positive return shock leads to an increase in the optimal distortion (see left panel of Figure 3) and therefore a reduction in the (distorted) expected return, we now have a negative correlation between returns and expected returns. This explains the positive intertemporal hedging demand for $\eta > 1$, which naturally grows with the horizon. The fact that the hedging demand is negative for $\eta < 1$ can also be understood from (29), by observing numerically that the partial derivative $\partial_x f < 0$ for $\eta < 1$, and vice versa for $\eta > 1$.

Figure 4 reports the economically relevant range of the sentiment variable. The top left panel presents the distribution of the sentiment variable at 30 years. The distribution displays a heavy left tail which reflects increased sentiment volatility in an adverse market environment (see Table 1). All other panels in Figure 4 are centered at the mean of sentiment and range from two standard deviations below to two standard deviations above.

3 General Equilibrium

We now extend our analysis to equilibrium asset pricing. To this end, we consider an economy with a single Lucas tree. Its dividend, which must be consumed by the representative agent, is given by the following dynamics

$$\frac{dc_t}{c_t} = \mu^c dt + \sigma^c dB_t^{\mathbb{B}}, \tag{33}$$

where μ^c and σ^c are constants. We first establish the equilibrium price P of this Lucas tree, which follows

$$\frac{dP_t + c_t dt}{P_t} = \mu_t^P dt + \sigma_t^P dB_t^{\mathbb{B}}. \tag{34}$$

The expected return μ^P, the volatility σ^P, and the interest rate r will be determined endogenously. We take x as the state variable and conjecture μ^P, σ^P and r are all functions of time and the state.

Given P and r, the representative agent, with generalized robustness based on a Cressie-Read penalty, invests and consumes with the optimal investment strategy π^* and the optimal consumption strategy c^*. In an attempt to match empirical asset pricing evidence quantitatively, we extend our earlier analysis and endow the representative agent with Epstein-Zin preferences in addition to Cressie-Read generalized robustness.[9] We denote the agent's elasticity of intertemporal substitution by $0 < \psi \neq 1$ and define $\nu = \frac{1-\gamma}{1-\frac{1}{\psi}}$.

Definition 1. $(r, \mu^P, \sigma^P, \pi^*, c^*)$ is an equilibrium if

1. The financial market clears, i.e., $\pi^* \equiv 1$;
2. the aggregate resource constraint holds, i.e., $c^* \equiv c$.

We can now calculate equilibrium quantities which are presented in the following proposition.

Proposition 3. *Let f be a solution of the following equation*

$$\begin{aligned} 0 =& \partial_t f + \tfrac{1}{2}|u^*|^2 \Big(\partial_{xx} f - \partial_x f + \psi(\partial_x f)^2\Big) + (1-\gamma)\partial_x f u^* \sigma^c - \delta \tfrac{\nu}{\psi} \\ &+ \tfrac{1-\gamma}{\psi}\big(\mu^c - u^*\sigma^c - \tfrac{1}{2}\gamma(\sigma^c)^2\big) + \delta^\psi \tfrac{\nu}{\psi} e^{-\frac{\psi}{\nu} f} + \tfrac{1-\gamma}{2\psi\theta} e^{(1-\eta)x}|u^*|^2, \end{aligned} \tag{35}$$

for $(t,x) \in [0,T) \times (\underline{x}, \overline{x})$, with boundary conditions

$$f(t,\underline{x}) = F_{\underline{x}}^{ent}(t), \quad f(t,\overline{x}) = F_{\overline{x}}^{ent}(t), \quad f(T,x) = \log\epsilon. \tag{36}$$

Moreover u^* in (35) is given by

$$u^* = \frac{(1-\gamma)(1-\partial_x f)\sigma^c}{\partial_{xx}^2 f - \psi\partial_x f + \psi(\partial_x f)^2 + \frac{1-\gamma}{\theta}e^{(1-\eta)x}}. \tag{37}$$

[9]The details of the preferences and the associated optimal consumption-investment problem can be found in Appendix A.

In (36)*,* F_x^{ent} *is the value in an equilibrium where the representative agent has an entropy-based preference for robustness with robustness parameter* $\theta e^{(\eta-1)x}$ *and* F_x^{ent} *satisfies an ODE in* (A.18)*. Then the equilibrium expected return, volatility, and risk-free interest rate are given by*

$$\mu^P - r = \gamma^{eff}(\sigma^P)^2, \tag{38}$$

$$\sigma^P = \sigma^c + \tfrac{\psi}{\nu}\partial_x f u^*, \tag{39}$$

$$r = \delta + \tfrac{1}{\psi}\mu^c - \tfrac{1}{2}\gamma^{ent}\big(1+\tfrac{1}{\psi}\big)(\sigma^c)^2 + DA, \tag{40}$$

where

$$\gamma^{eff} = \gamma + \frac{(1-\gamma)(1-\partial_x f)^2}{\partial^2_{xx} f - \partial_x f + (\partial_x f)^2 + \frac{1-\gamma}{\theta}e^{(1-\eta)x}}$$

is the effective state-dependent risk aversion and $\gamma^{ent} = \gamma + \theta e^{(\eta-1)x}$ *is the effective risk aversion in the entropy case with the parameter for robustness frozen at* $\theta e^{(\eta-1)x}$*. DA is the dynamic adjustment,*

$$\begin{aligned} DA = &- \gamma^{dyn}(\sigma^c)^2 + \tfrac{\psi\gamma}{\nu}\partial_x f u^*\sigma^c + (1-\tfrac{1}{\psi})u^{dyn}\sigma^c \\ &+ \big[\tfrac{\psi}{\nu}\partial_x f + \tfrac{\psi^2}{2\nu^2}\big(1-\nu-2\gamma^{imp}\big)(\partial_x f)^2\big]|u^*|^2 \\ &- \tfrac{1-\frac{1}{\psi}}{2\theta}e^{(1-\eta)x}\big[2u^{ent}u^{dyn} + (u^{dyn})^2\big], \end{aligned}$$

where

$$\begin{aligned} \gamma^{eff} =& \gamma^{ent} + \gamma^{dyn}, \\ u^* =& u^{ent} + u^{dyn} = \theta e^{(\eta-1)x}\sigma^c + u^{dyn}. \end{aligned}$$

The equilibrium equity premium in (38) is given by a Consumption CAPM relationship, where the key innovation is the time-varying price of risk. γ^{eff} is precisely the belief- and state-dependent effective risk aversion that determines the optimal portfolios in the partial-equilibrium analysis. The term simplifies to γ when the agent has no preference for robustness ($\theta = 0$) and becomes $\gamma + \theta$ when $\eta = 1$, i.e. in the entropy case, since all spatial derivatives are zero in this case. Generalized robustness instead allows us to produce rich dynamics, despite the extremely stylized and simple underlying dynamics in this lognormal i.i.d. economy. Applying the insights from Table 1, we expect a countercyclical price of risk for $\eta < 1$, since this is the situation where the investor's risk aversion is countercyclical, due to increases in pessimism following adverse shocks. This mechanism is driven entirely by time-varying pessimism and the stochastic beliefs it generates endogenously. The increase in risk aversion in bad states of nature bears some resemblance to models of habit formation such as Campbell and Cochrane (1999).

The second key contribution of generalized robustness to equilibrium asset pricing is that we obtain excess volatility, driven by time-varying sentiment. With entropy, $\sigma^P = \sigma^c$ and the well-known excess volatility puzzle emerges, as is common in standard asset pricing models with lognormal dynamics. Inspection of equation (39) reveals that our model is able to generate excess volatility when $\frac{1}{\nu}\partial_x f > 0$ given that $u^* > 0$. For $\eta < 1$, which is needed for a countercyclical price of risk, $\partial_x f < 0$, so that $\nu < 0$ (or $\psi > 1$ when $\gamma > 1$) is required. This explains why we extended our analysis to Epstein-Zin preferences, as clearly this condition would be violated for CRRA utility.[10] It

[10]The necessity of Epstein-Zin preferences was also pointed out by Jin and Sui (2019) in their model of asset pricing with extrapolative expectations. These preference parameter restrictions are also common in the long-run-risk literature.

is also intuitive that the excess volatility generated in our model is proportional to u^*, since this is the instantaneous volatility of the (log) state variable Z capturing belief distortions. Because u^* is stochastic and increases in bad times, we also expect equilibrium stock price volatility to be stochastic and to increase in bad times, although this cannot be concluded decisively from (39) without solving the PDE for f.

Finally, for the equilibrium risk-free rate we present a decomposition to facilitate comparison with existing results in the literature so as to flesh out our contributions most clearly. We find that Cressie-Read generalized robustness adds a rich dynamic adjustment to the equilibrium risk-free rate that obtains in an economy with entropy-based robustness. The first three terms in (40) represent the effect of the usual determinants of savings behavior on equilibrium interest rates, namely the rate of time preference δ, intertemporal substitution based on the investor's EIS ψ and expected consumption growth μ^c, and precautionary savings. Our Cressie-Read robustness adds a dynamic adjustment to these standard determinants. These additional terms are all related to precautionary savings reflecting the higher effective risk aversion and the higher volatility in the economy due to time-varying sentiment. For the entropy case with the parameter for robustness set to $\theta e^{(\eta-1)x}$, all these additional terms vanish because $u^{dyn} = \gamma^{dyn} = \partial_x f = 0$, reduces (40) exactly to the case in Maenhout (2004). Signing the additional terms requires solving the HJB (35), which we do in the calibration later.

Before turning to the calibration, we show how to use these equilibrium results to price the stock market, i.e., a risky asset with dividends that are not necessarily perfectly correlated with the representative agent's consumption. Consider dividend dynamics given by

$$\frac{dD_t}{D_t} = \mu^D dt + \sigma^D(\rho dB_t^{\mathbb{B}} + \sqrt{1-\rho^2}dB_t^{\perp}), \tag{41}$$

where μ^D and σ^D are constants representing the dividend growth rate and volatility respectively, and $B^{\perp}$ is a Brownian motion independent of $B^{\mathbb{B}}$. The constant ρ is the instantaneous correlation between consumption growth and the dividend growth.

We consider the stock as an asset in zero net supply with a shadow price determined in equilibrium. To find its equilibrium (shadow) price, we first identify the state price density M for the representative agent. Because markets are complete, M follows a stochastic differential equation of the form

$$\frac{dM_t}{M_t} = -r_t dt - \xi_t dB_t^{\mathbb{U}}, \quad M_0 = 1, \tag{42}$$

where r is the equilibrium risk-free rate in the Lucas tree economy. To determine the market price of risk ξ, the sum of discounted wealth and consumption (i.e. the process $M_t W_t + \int_0^t c_s M_s ds$) must be a $\mathbb{P}^{\mathbb{U}}$-martingale. This leads to the following market price of risk

$$\xi_t = \lambda_t - u_t. \tag{43}$$

where $\lambda_t = \frac{\mu_t^P - r_t}{\sigma_t^P}$ is the equilibrium Sharpe ratio of the Lucas tree. Combining (42) and (43), we have

$$M_t = e^{-\int_0^t r_s ds}\mathcal{E}\Big(-\int (\lambda_s - u_s)dB_s^{\mathbb{U}}\Big)_t, \tag{44}$$

where $\mathcal{E}\big(-\int \xi_s dB_s^{\mathbb{U}}\big)_t = \exp\big(-\int_0^t \frac{1}{2}|\xi_s|^2 ds - \int_0^t \xi_s dB_s^{\mathbb{U}}\big)$ is a stochastic exponential. Define a measure $\mathbb{Q}$ via

$$\frac{d\mathbb{Q}}{d\mathbb{P}^{\mathbb{U}}}\Big|_{\mathcal{F}_T} = \mathcal{E}\Big(-\int \big(\lambda_s - u_s\big) dB_s^{\mathbb{U}}\Big)_T.$$

Then $B^{\mathbb{Q}}$, defined via

$$B_t^{\mathbb{Q}} = B_t^{\mathbb{U}} + \int_0^t \lambda_s - u_s ds = B_t^{\mathbb{B}} + \int_0^t \lambda_s ds,$$

is a Brownian motion under $\mathbb{Q}$. The previous dynamics indicate that $\mathbb{Q}$ is the risk-neutral measure.

Using the state price density, the stock can be priced as follows

$$S_t = \frac{1}{M_t}\mathbb{E}_t^{\mathbb{U}}\Big[\int_t^T M_s D_s ds\Big]. \tag{45}$$

Suppose that S follows the dynamics

$$\frac{dS_t + D_t dt}{S_t} = \mu_t^S dt + \sigma_t^S dB_t^{\mathbb{B}} + \sigma_t^{S,\perp} dB_t^{\perp}. \tag{46}$$

Define $\ell = S/D$ as the price-dividend ratio. The following result presents the equilibrium stock return and volatility.

Proposition 4. *Let ℓ be the solution to the following equation*

$$\partial_t \ell + \frac{1}{2}|u^*|^2 \partial_{xx}^2 \ell + \Big(\frac{1}{2}|u^*|^2 - u^*\lambda + \rho\sigma^D u^*\Big)\partial_x \ell + \big(\mu^D - \rho\sigma^D\lambda - r\big)\ell + 1 = 0, \tag{47}$$

with the terminal condition $\ell(T,\cdot) \equiv 0$ and u^* coming from (37). Then

$$\begin{aligned}
\mu_t^S =& \frac{\partial_t \ell}{\ell} + \frac{1}{2}|u^*|^2 \frac{\partial_{xx}^2 \ell}{\ell} + \frac{1}{2}|u^*|^2 \frac{\partial_x \ell}{\ell} + \mu^D + u^*\sigma^D \rho \frac{\partial_x \ell}{\ell} + \frac{1}{\ell},\\
\sigma_t^S =& u^* \frac{\partial_x \ell}{\ell} + \overline{\sigma}\rho,\\
\sigma_t^{S,\perp} =& \overline{\sigma}\sqrt{1-\rho^2}.
\end{aligned} \tag{48}$$

Moreover, the following CAPM relation is satisfied:

$$\mu_t^S = r_t + \lambda_t \sigma^S.$$

The results are intuitive and extend our earlier findings to the case of an asset that pays dividends that are less than perfectly correlated with the consumption stream of the representative agent. The risk premium on the stock is given by the standard Consumption CAPM, but with a time-varying price of risk generated by our model. We also obtain excess volatility. Without robustness consideration or with entropy-based robustness, the price-dividend ratio is trivially constant in a lognormal economy, resulting in the equilibrium stock volatility being equal to the dividend volatility. The Cressie-Read divergence measure injects time-varying pessimism, inducing a dynamic price-dividend ratio. Equation (48) shows that excess volatility emerges when the price-dividend ratio is procyclical. This condition will be verified numerically in our model calibration later for the case $\eta < 1$.

4 Estimating Pessimism

In the following, we estimate a measure of time-varying sentiment from the data. In particular, we would like a proxy for the time variation in the worst-case drift distortion. To measure subjective beliefs, we make use of an extensive survey on several macroeconomic quantities. In addition to subjective beliefs, we also need a measure of objective beliefs since the wedge between the subjective and objective beliefs will help us back out the optimal distortion. Objective beliefs are calculated from a vector autoregression (VAR) from which we infer realized macroeconomic variables.

SUBJECTIVE BELIEFS: We hand collect survey data from Blue Chip Economic Indicators. Blue Chip is a survey of panelists from around 40 major financial institutions. The names of institutions and forecasters are disclosed. The survey is conducted around the beginning of each month and is released on the tenth of each month for responses based on information for the previous month. We use Blue Chip forecasts from the end-of-quarter month survey (i.e., March, June, September, and December) and construct a consensus (mean) estimate from the cross-section of individual forecasts on year-to-year changes of real GDP, inflation, industrial production, and unemployment.

OBJECTIVE BELIEFS: To get a proxy for objective beliefs, we estimate a VAR with two lags on real GDP, inflation, and unemployment and use forecasts from this VAR. In our estimation, we also include the level of the three-month Treasury bill.[11]

[Insert Figure 5 and Table 3 here]

The wedge is then defined as the difference between the subjective and objective beliefs about future macroeconomic activity. In particular, it measures the amount of pessimism or optimism in the economy. Figure 5 plots the wedge for our four macroeconomic variables and Table 3 reports some summary statistics.[12] We notice a very strong component between the four series. Indeed, as shown in Table 3, correlations are between 50% and up to almost 100% between IP and GDP growth. More importantly, the estimates indicate that subjective beliefs are on average pessimistic as indicated by the negative mean in the top row of Table 3. The exception is inflation but the positive average is entirely driven by the sharp increase after the 2008 recession. The most important finding concerns the business cycle dynamics of the wedges: wedges peak before a recession and suffer sharp contractions during recessions. This echoes our theoretical findings which predict that bad shocks lead agents to become more pessimistic.

Finally, we can use our estimates of the GDP wedge to inform us about the dynamics of the optimal distortion $u^\star$. Recall that in our model the wedge is defined as $-u^* \sigma^P$. Estimates of $u^\star$ are therefore estimated by dividing the wedge by the sample standard deviation of the stock.

[11] These variables are used in standard New Keynesian models for forecasting, see, e.g., Cogley and Sargent (2005), Primiceri (2005) or Chauvet and Potter (2013) for a survey. Our results remain unchanged whether we include interest rates or not.

[12] We multiply the wedges for inflation and unemployment by minus 1 to facilitate comparison with GDP and IP growth. Higher unemployment and inflation indicate bad times, while the opposite holds for GDP and IP growth.

5 Calibration

We now turn to the calibration in order to explore the ability of our equilibrium model with Cressie-Read robustness in a stylized Lucas economy with i.i.d. consumption growth and lognormality to qualitatively and quantitatively match salient features of asset prices together with the wedge dynamics. To this end, we numerically solve equations (35) and (47) together with their associated boundary conditions using finite difference methods with implicit schemes. Importantly, we impose discipline in terms of the free parameters governing the strength of the preference for robustness by calibrating to the empirical estimates of the wedges estimated from the data in order to have realistic dynamics for our key new state variable related to sentiment and pessimism.

We follow Jin and Sui (2019) and use the empirical asset pricing evidence reported in Campbell and Cochrane (1999) and in Beeler and Campbell (2012) as objectives to match. In addition, we also match the first three moments of the empirical wedges from our empirical analysis, reported in Table 3. The rest of the model parameters used for the calibration are summarized in Panel B of Table 2.

[Insert Figure 6 here]

Figure 6 reports the equilibrium quantities produced in our model for the following preference parameters: $\delta = 0.04$, $\gamma = \theta = 7$, $\psi = 1.25$ and $\eta = 0.65$. The Cressie-Read parameter is crucial and our value lies between entropy ($\eta = 1$) and Hellinger ($\eta = 0.5$). Because the key contribution of our model is time-varying sentiment and dynamic pessimism, we report all equilibrium quantities for a relevant range of the sentiment state variable x, obtained from Monte Carlo simulation. In the figures, we plot two standard deviations below and above the mean across 10^4 paths from a 50 year Monte Carlo simulation.

Consistent with the results in Section 3, the optimal distortion, equilibrium Sharpe ratio, equilibrium volatility and wedges are all countercyclical, reflecting the countercyclical pessimism generated by the model. The price-dividend ratio is procyclical, as is the equilibrium risk-free rate. Quantitatively, Table 4 shows that the model performs well in generating a sizeable risk premium and a realistic Sharpe ratio. The quantitative success in producing excess volatility is more limited.

Importantly, the wedge between the objective and the subjective risk premium is reasonable and less than one percent at the mean of the state space, while it grows to just over 150 basis points in bad states of nature. As Chamberlain (forthcoming) and Hansen and Sargent (2019) both point out, a central idea in robust Bayesian analysis based on classical work of Good (1952) is to judge the plausibility of a min-max model by examining how reasonable the worst-case measure $\mathbb{U}$ is that is supporting the equilibrium.

Turning to the moments for the wedges produced by the model versus the ones we estimate from the data, our model produces excessive skewness. This is likely due to the heavy-tailed distribution for sentiment due to increased sentiment volatility under adverse market conditions, echoing the findings reported in Figure 4.

It is worth mentioning that the same model and same parameters, but with relative entropy produces a higher risk-free rate of 3.59 percent, a very low and constant risk premium of 1.81 percent, a constant return volatility equal to the volatility of dividends, and a constant Sharpe Ratio of 0.1064.

We conclude by noting that our Cressie-Read extension improves substantially on the quantitative front, in addition to generating meaningful time-variation at the business cycle frequency.

6 Conclusions

Our paper makes the following contributions. First, we extend the Hansen-Sargent robustness setting to the family of Cressie-Read divergences. We show that this generalization has important implications for the nature and source of fluctuations in risk aversion. In particular, we show that the agent's pessimistically distorted beliefs form a state variable that generates endogenous time-variation in pessimism and stochastic effective risk aversion. This stands in sharp contrast to the more standard case applied in the literature relying on entropy as a divergence measure between models and where effective risk aversion is constant when fundamentals are i.i.d.

Second, we illustrate the implications of deviations from entropy using a simple partial equilibrium portfolio choice framework as well as in a general equilibrium Lucas economy. Our choice of a simple model with i.i.d. returns is very deliberate, as we seek to understand the contribution of Cressie-Read divergences which in more complicated models may be harder to uncover. As a first exercise, we derive parameter conditions under which beliefs are state-dependent. We find that when the Cressie-Read parameter η is smaller than one, endogenous effective risk aversion is countercyclical and declines following positive shocks, while the opposite happens whenever $\eta > 1$. In our portfolio choice problem, this induces intertemporal hedging demands and therefore both horizon- and state-dependent portfolios, despite returns being i.i.d. This example nicely contrasts to the entropy case where optimal portfolios are constant and myopic.

Third, the premise of our model posits that the agent's concern for model misspecification leads to pessimistic beliefs. It is obviously an empirical question whether this is true in the data. To test this hypothesis in the data, we collect survey responses about future economic activity. Since our model predicts a unique mapping between the agent's belief distortion and subjective and objective beliefs, we can estimate those distortions from the data. With these estimates in hand, we calibrate our general equilibrium model. We find that our simple model is able to match not only the equity premium and Sharpe ratio, but also produces reasonable values for interest rates and especially for the worst-case beliefs supporting this equilibrium.

It is well-known that macroeconomic fundamentals such as consumption feature fat tails, which might be due to a small probability of a disaster, see, e.g., Barro (2006). Canonical models in asset pricing such as the rare disaster models study the implications of these large negative shocks for asset prices usually within the context of representative agent models with Epstein-Zin preferences, see, e.g.,Tsai and Wachter (2015) for a review. A setting with tail risk and agents featuring robustness concerns is a natural extension of our framework, which we leave for future research.

Figures

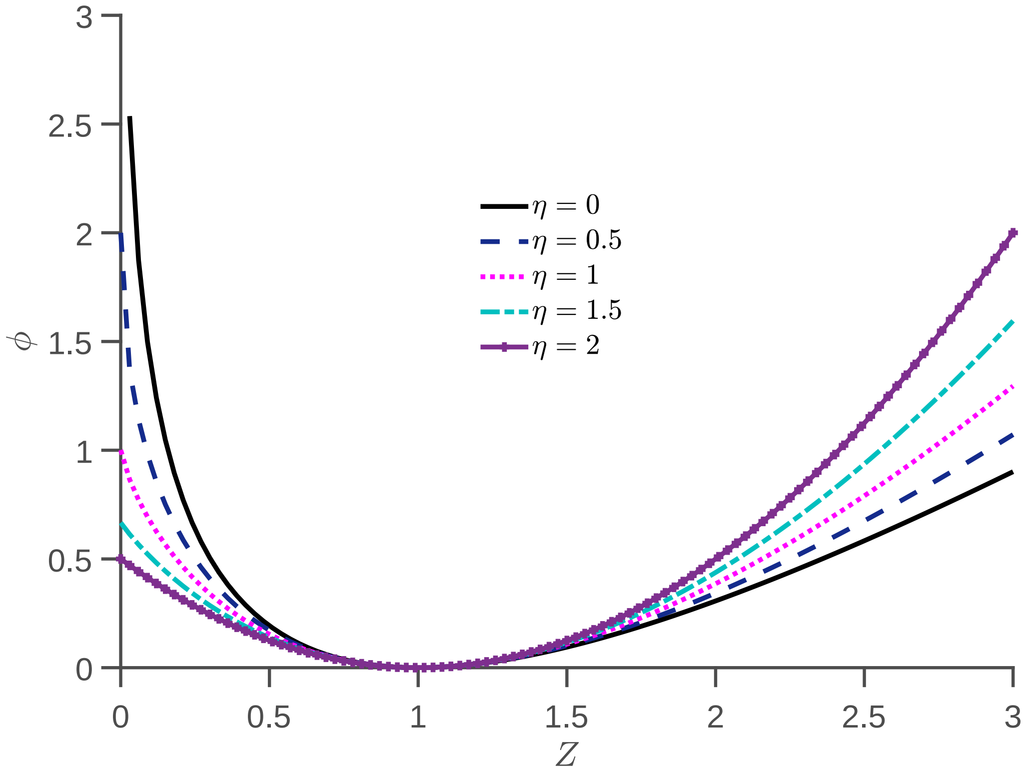

Figure 1. **Cressie-Read Divergence for different values of** η

Notes: This figure plots Cressie-Read divergence for different values of η. $\eta = 0$ corresponds to Burg (1972) entropy, 1/2 to Hellinger (1909) divergence, 1 to Kullback and Leibler (1951) distance, and 2 to modified χ^2 distance.

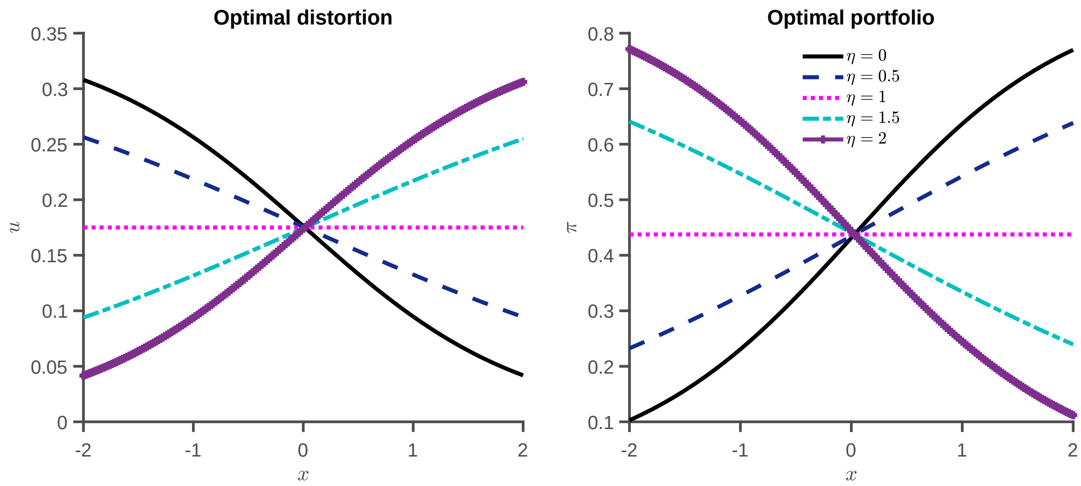

(a) $\gamma = \theta = 2$

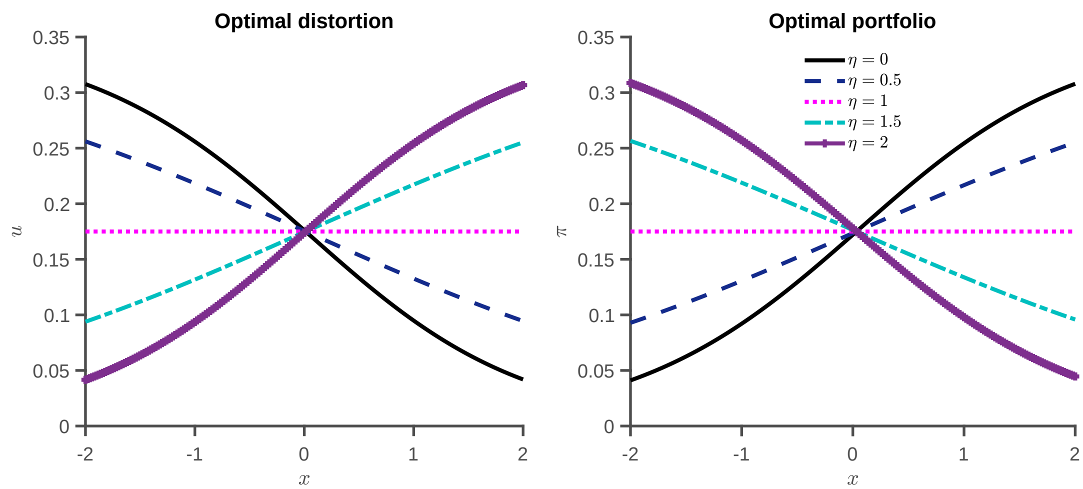

(b) $\gamma = \theta = 5$

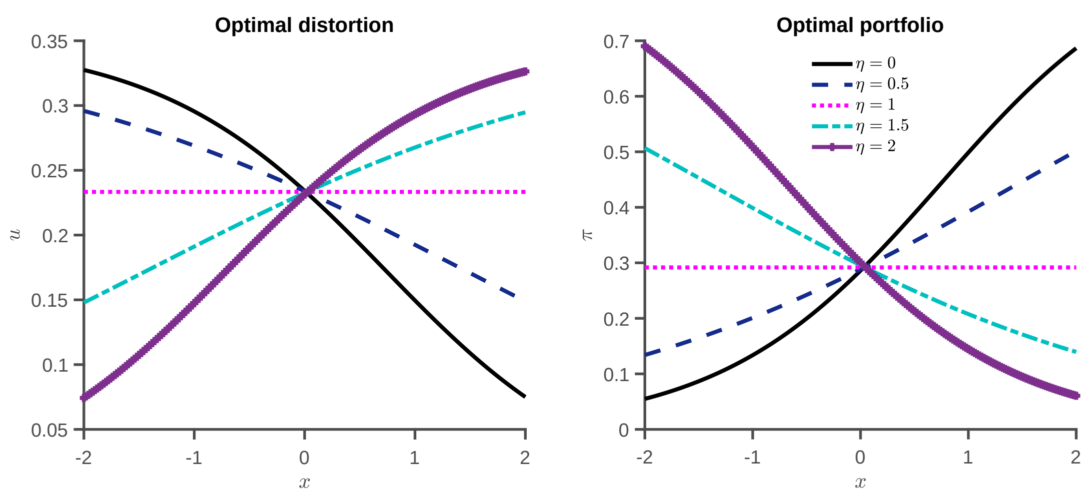

(c) $\gamma = 2$ and $\theta = 4$

Figure 2. **Optimal distortions and portfolios for different values of η**

Notes: This figure plots optimal distortions and portfolios for different values of risk aversion (γ) and preference for robustness (θ). Parameters used are summarized in Panel A of Table 2 and the time horizon is $T = 1$ year.

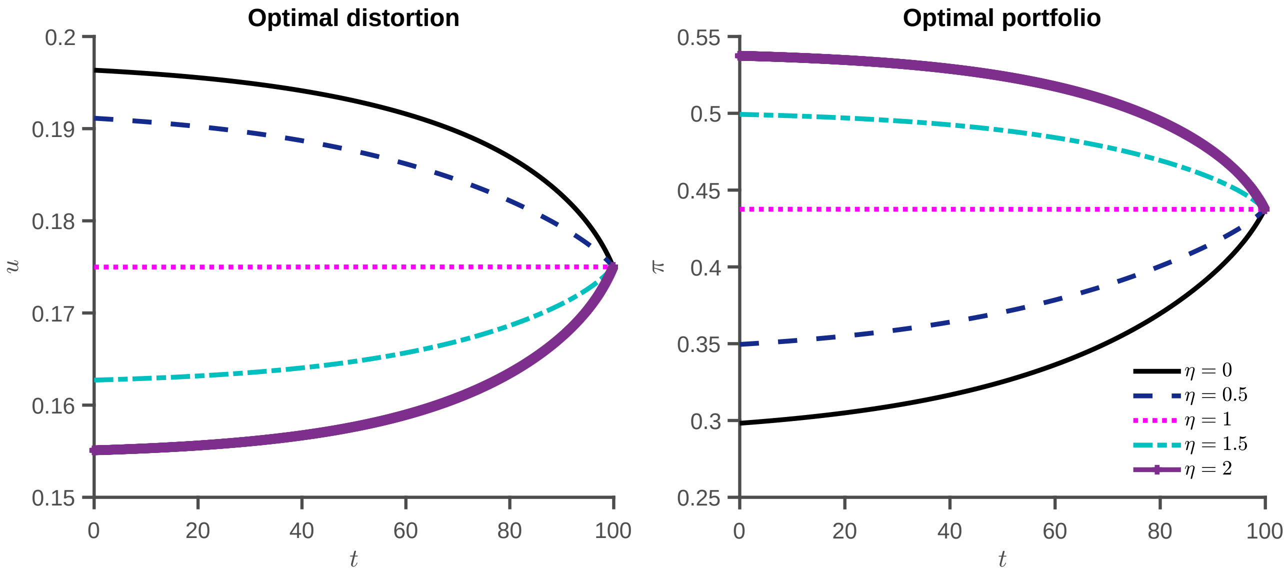

Figure 3. **Optimal distortions and portfolios**

Notes: This figure plots optimal distortions and portfolios for risk aversion (γ) equal to 2, preference for robustness (θ) equal to 2, and $x = 0$. Parameters used are summarized in Panel A of Table 2 and the time horizon is $T = 100$ years.

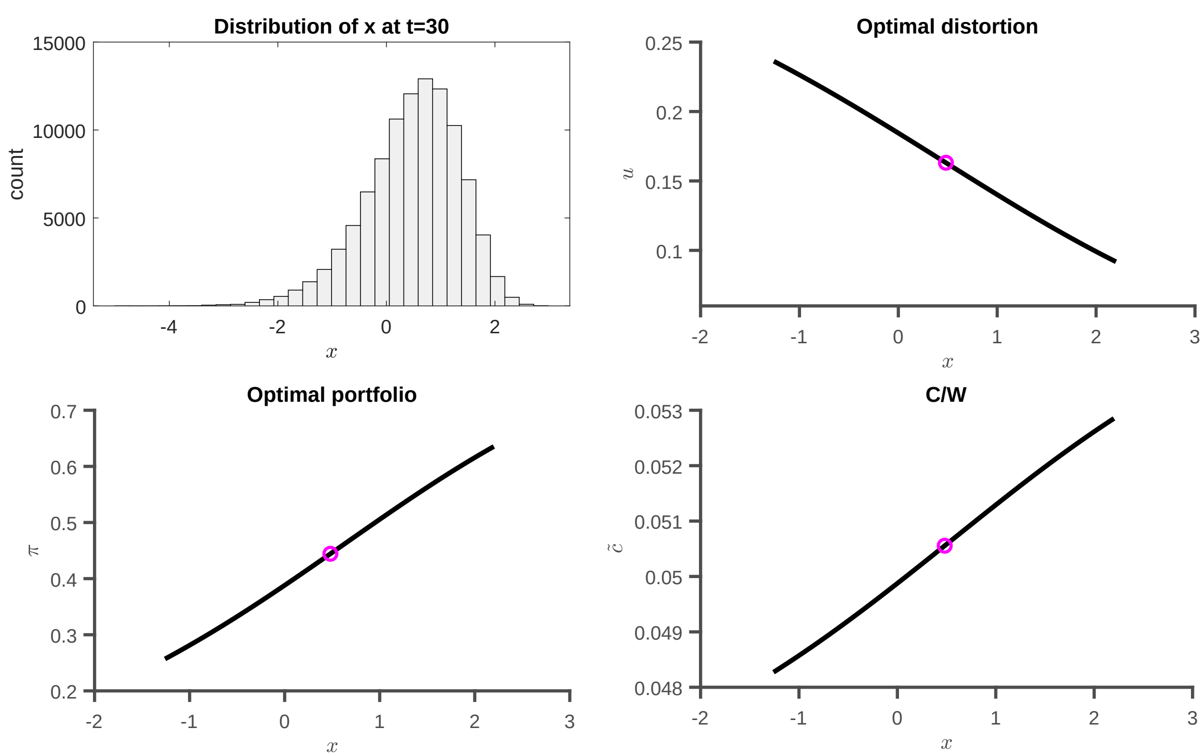

Figure 4. **Optimal distortions and portfolios**

Notes: This figure plots the distribution of the sentiment variable x, optimal distortions, portfolios, and consumption wealth ratio for risk aversion (γ) equal to 2, preference for robustness (θ) equal to 2, and $\eta = 0.5$. Parameters used are summarized in Panel A of Table 2 and time horizon is $T = 60$ years. All figures present quantities at 30 years. They all center at the mean of the sentiment variable and span from two standard deviation below to two standard deviation above.

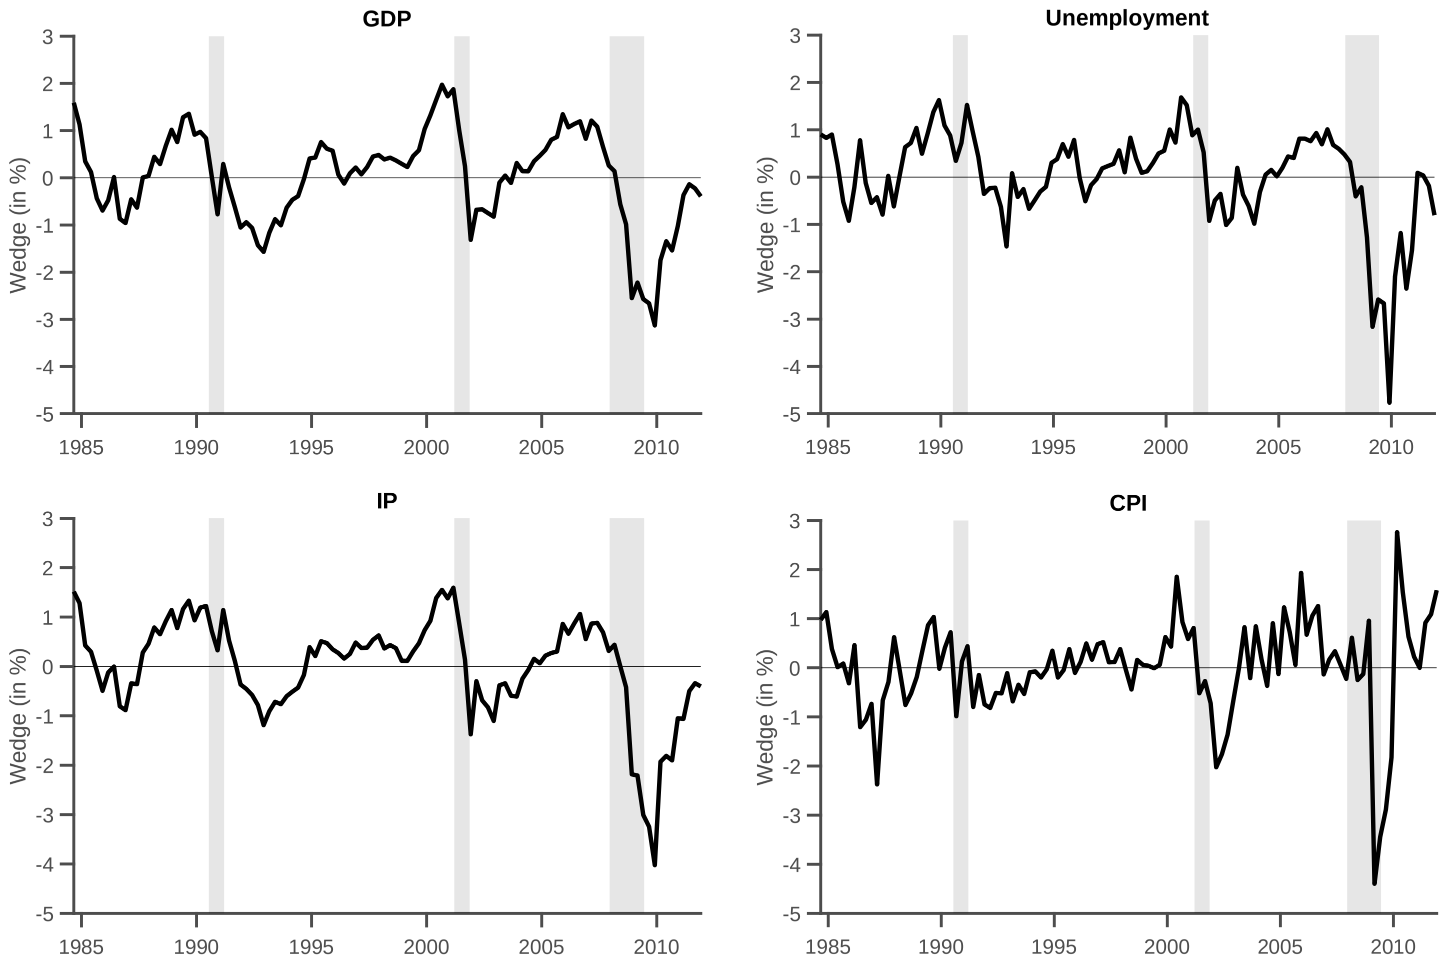

Figure 5. **Wedge: Difference between Subjective and Objective Beliefs**

Notes: This figure plots wedges for real GDP, unemployment, inflation, and industrial production, defined as the difference between the mean one-year-ahead forecasts from Blue Chip Economic Indicators and corresponding statistical (objective) forecasts from a VAR. Gray shaded bar indicate recessions as defined by the NBER. Data is quarterly and running from 1984 to 2012.

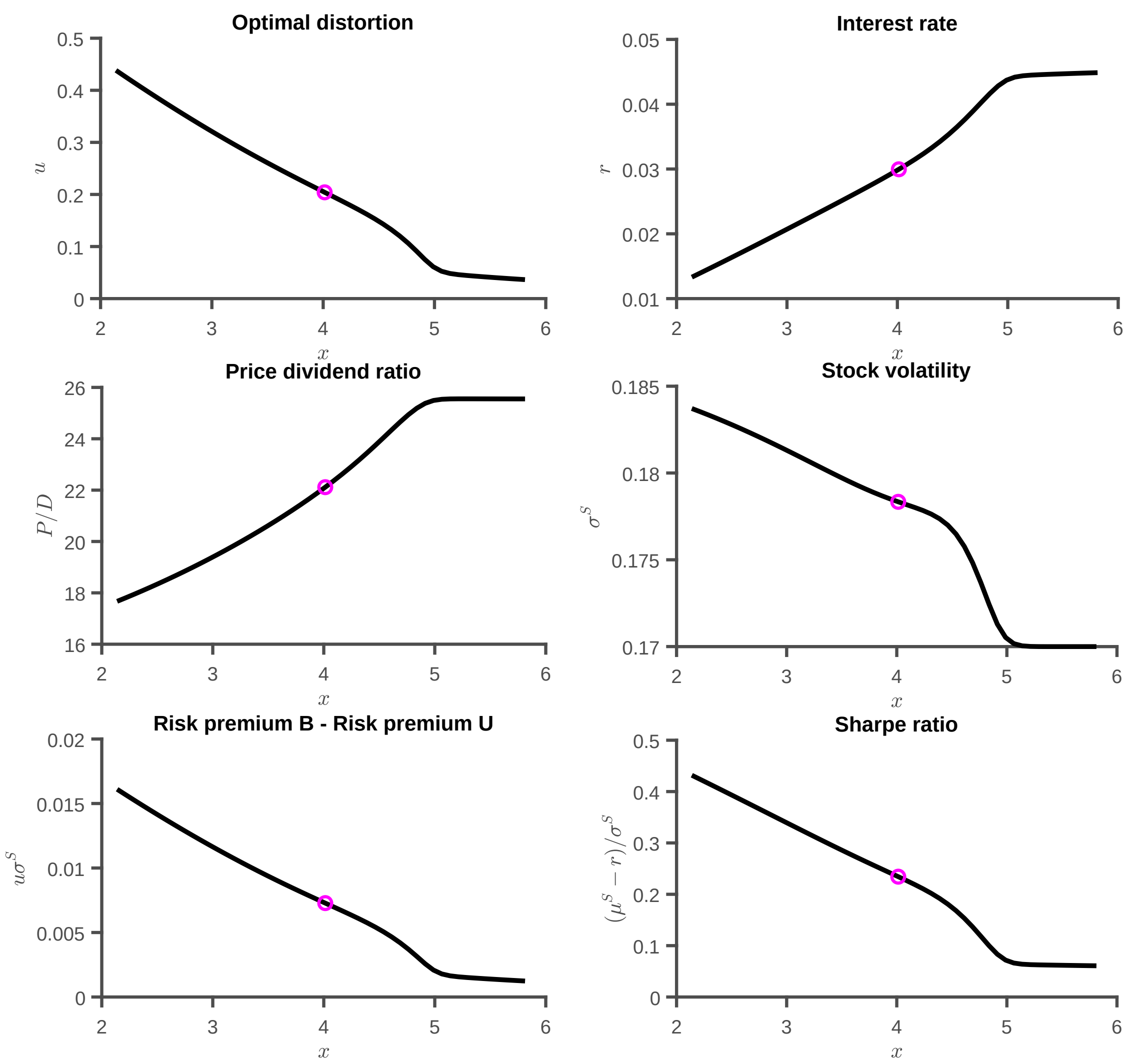

Figure 6. **General Equilibrium Results**

Notes: This figure plots the optimal distortion, interest rate, price-dividend ratio, stock volatility, difference in stock risk premia between measures, and the Sharpe ratio in equilibrium. Parameters used are summarized in Panel B of Table 2 and time horizon is $T = 100$ years. All figures present quantities at 50 years. They all center at the mean of the sentiment variable and span from two standard deviation below to two standard deviation above.

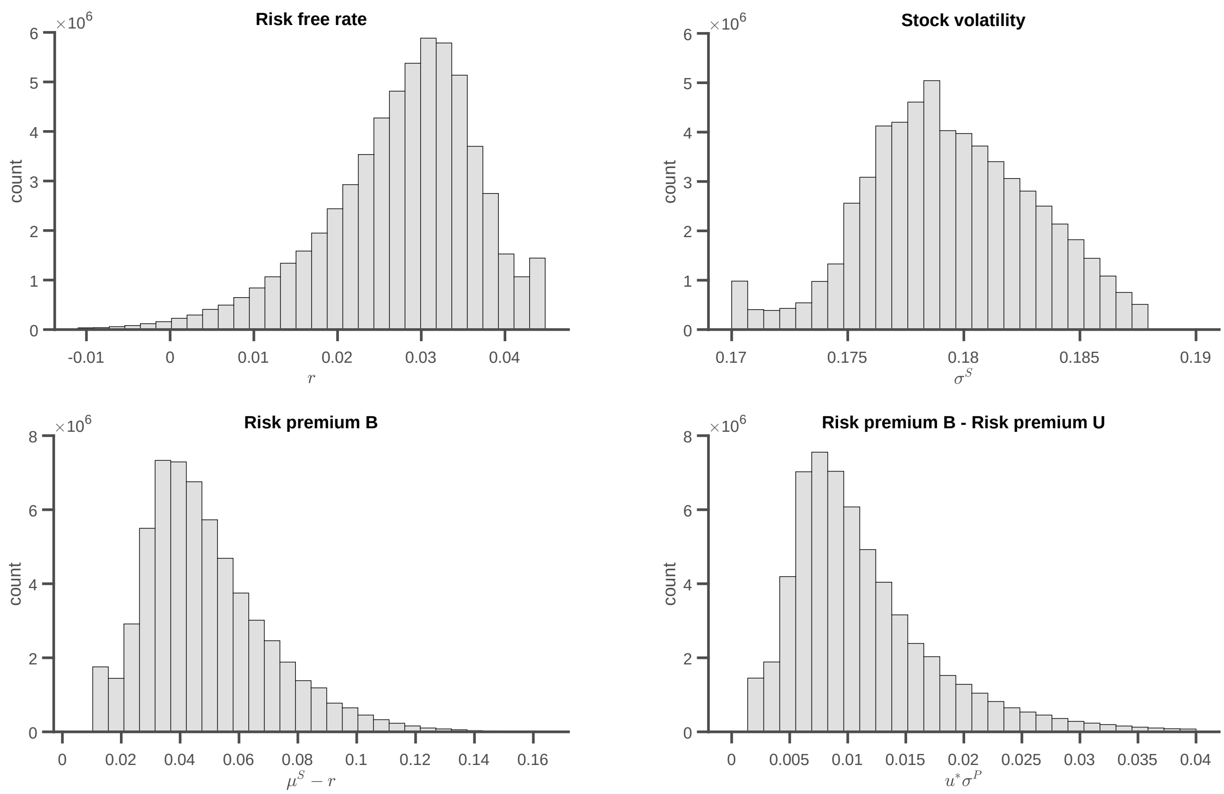

Figure 7. **Distribution of Equilibrium Quantities**

Notes: This figure plots the distributions of interest rate, stock volatility, risk premium $\mathbb{B}$, and difference in stock risk premia between measures along equilibrium path generated via Monte Carlo between year 20 and year 50. Parameters used are summarized in Panel B of Table 2 and time horizon is $T = 100$ years.

Tables

Table 1. **Responses to Fundamental Shocks** $B^{\mathbb{B}}$

This table studies responses to fundamental shocks discussed in the theoretical model.

$\Delta B^{\mathbb{B}}$	Optimal Distortion	Risk Aversion	Sentiment	Volatility Sentiment
		$\eta < 1$		
positive	$\downarrow$	$\downarrow$	less pessimistic	$\downarrow$
negative	$\uparrow$	$\uparrow$	more pessimistic	$\uparrow$
		$\eta > 1$		
positive	$\uparrow$	$\uparrow$	more pessimistic	$\uparrow$
negative	$\downarrow$	$\downarrow$	less pessimistic	$\downarrow$

Table 2. **Parameter Values**

This table reports parameter values used for simulations. In addition to the variables defined below we use $\epsilon = 1$.

Parameter	Variable	Value
Panel A: Partial Equilibrium		
r	Risk-free interest rate	0.03
δ	Discount rate	0.03
σ	Stock volatility	0.20
μ	Expected stock return	0.10
Panel B: General Equilibrium		
δ	Discount rate	0.04
μ^c	Consumption growth rate	0.0191
σ^c	Consumption volatility	0.038
μ^D	Dividend growth rate	0.0245
σ^D	Dividend volatility	0.17
γ	Risk aversion	7
θ	Preference for robustness	7
η	Cressie-Read parameter	0.65
ψ	EIS	1.25
ρ	Correlation dividend and consumption	0.2

Table 3. **Summary Statistics Wedges**

This table reports summary statistics (mean, standard deviation, and skewness) for the wedges for CPI, Unemployment, GDP and IP growth. The wedge is defined as the difference between the mean one-year forecast of each variable from Blue Chip Economic Indicators and the one year ahead forecast using a VAR with two lags. We multiply the wedge for CPI and Unemployment by -1. Data is quarterly and runs from January 1984 to December 2011.

	CPI	**Unemp**	**GDP**	**IP**
Mean	0.11%	-0.17%	-1.40%	-1.59%
Stdev	1.13%	0.61%	1.64%	2.91%
Skewness	-0.012	-0.017	-0.660	-0.013
	Correlations			
CPI	1.00			
Unemp	0.43	1.00		
GDP	0.52	0.86	1.00	
IP	0.48	0.92	0.93	1.00

Table 4. **Summary Statistics of Equilibrium Quantities**

This table reports moments about equilibrium quantities of the calibrated model. The model is disciplined by the mean, standard deviation and skewness of the wedge on GDP. The empirical values of equilibrium quantities are obtained from Campbell and Cochrane (1999) and Beeler and Campbell (2012). The theoretical values are moments of equilibrium quantities between year 20 to 50 obtained by Monte Carlo simulation with 10^4 paths. Parameters used are summarized in Panel B of Table 2. The time horizon is $T = 100$ years.

Statistic	**Calibrated value**	**Empirical value**
Mean wedge ($\mathbb{E}^{\mathbb{B}}[-u^* \sigma^P]$)	-1.15%	-1.40%
Stdev wedge ($\sigma(-u^* \sigma^P)$)	0.68%	1.64%
Skewness wedge (skewness$(-u^* \sigma^P)$)	-1.77	-0.66
	Theoretical value	**Empirical value**
Equity premium ($\mathbb{E}^{\mathbb{B}}[\mu^S - r]$)	4.89%	3.90%
Stock volatility (σ^S)	17.9%	18.0%
Sharpe ratio ($\mathbb{E}^{\mathbb{B}}[\mu^S - r]/\sigma^S$)	0.27	0.22
Interest rate ($\mathbb{E}^{\mathbb{B}}[r]$)	2.76%	2.92%
Interest rate volatility ($\sigma(r)$)	0.91%	2.89%
Price-dividend ratio ($\mathbb{E}^{\mathbb{B}}(P/D)$)	21.7	21.1

References

ADAM, K., D. MATVEEV, AND S. NAGEL (2019): "Do Survey Expectations of Stock Returns Reflect Risk Adjustments?," Bank of Canada Staff Working Paper 2019-11.

ALMEIDA, C., AND R. GARCIA (2017): "Economic Implications of Nonlinear Pricing Kernels," *Management Science*, 63(10), 3361–3380.

ANDERSON, E. W., L. P. HANSEN, AND T. J. SARGENT (2003): "A Quartet of Semigroups for Model Specification, Robustness, Prices of Risk, and Model Detection," *Journal of the European Economic Association*, 1(1), 68–123.

BALTER, A., F. HORVATH, AND P. J. MAENHOUT (2019): "An Exploration of Dynamic Robust Control with Cressie-Read Divergence," *Working Paper, INSEAD*.

BARBERIS, N. (2018): "Psychology-Based Models of Asset Prices and Trading Volume," *Douglas Bernheim, Stefano DellaVigna, and David Laibson, eds., Handbook of Behavioral Economics*, 1, 79–175.

BARBERIS, N., R. GREENWOOD, L. JIN, AND A. SHLEIFER (2015): "X-CAPM: An Extrapolative Capital Asset Pricing Model," *Journal of Financial Economics*, 115, 1–24.

BARRO, R. J. (2006): "Rare Disasters and Asset Markets in the Twentieth Century," *Quarterly Journal of Economics*, 121, 823–866.

BEELER, J., AND J. Y. CAMPBELL (2012): "The long-run risks model and aggregate asset prices: an empirical assessment," *Critical Finance Review*, pp. 141–182.

BEKAERT, G., E. ENGSTROM, AND Y. XING (2009): "Risk, Uncertainty and Asset Prices," *Journal of Financial Economics*, 91(1), 59–82.

BHANDARI, A., J. BOROVIČKA, AND P. HO (2019): "Survey Data and Subjective Beliefs in Business Cycle Models," Working Paper, University of Minnesota.

BORDALO, P., N. GENNAIOLI, AND A. SHLEIFER (2018): "Diagnostic Expectations and Credit Cycles," *Journal of Finance*, 73(1), 199–227.

BURG, J. P. (1972): "The Relationship Between Maximum Entropy Spectra and Maximum Likelihood Spectra," *Geophysics*, 37(2), 375–376.

CAMPBELL, J. Y., AND J. COCHRANE (1999): "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior," *Journal of Political Economy*, 107(2), 205–251.

CHAMBERLAIN, G. (forthcoming): "Robust Decision Theory and Econometrics," *Annual Review of Economics*.

CHAUVET, M., AND S. POTTER (2013): *Forecasting Output*vol. 2, chap. 0, pp. 141–194. Elsevier B.V.

CHEN, Z., AND L. EPSTEIN (2002): "Ambiguity, Risk, and Asset Returns in Continuous Time," *Econometrica*, 70(4), 1403–1443.

COGLEY, T., AND T. J. SARGENT (2005): "Drifts and Volatilities: Monetary Policies and Outcomes in the Post WWII U.S.," *Review of Economic Dynamics*, 8, 262–302.

CONSTANTINIDES, G. M. (1990): "Habit Formation: A Resolution of the Equity Premium Puzzle," *Journal of Political Economy*, 98, 519–543.

CRESSIE, N., AND T. R. C. READ (1984): "Multinomial Goodness-of-Fit Tests," *Journal of the Royal Statistical Society. Series B (Methodological)*, 46(3), 440–464.

DETEMPLE, J. B., AND F. ZAPATERO (1991): "Asset Prices in an Exchange Economy with Habit Formation," *Econometrica*, 59(6), 1633–1657.

DUFFIE, D., AND L. G. EPSTEIN (1992): "Stochastic Differential Utility," *Econometrica*, 60(2), 353–394.

EL KAROUI, N., S. PENG, AND M.-C. QUENEZ (1997): "Backward Stochastic Differential Equation in Finance," *Mathematical Finance*, 7(1), 1–71.

EPSTEIN, L. G., AND M. SCHNEIDER (2003): "Recursive Multiple-Priors," *Journal of Economic Theory*, 113(1), 1–31.

EPSTEIN, L. G., AND S. E. ZIN (1989): "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework," *Econometrica*, 57.

GILBOA, I., AND D. SCHMEIDLER (1989): "Maxmin Expected Utility with Non-Unique Prior," *Journal of Mathematical Economics*, 18(2), 141–153.

GOOD, I. J. (1952): "Rational Decisions," *Journal of the Royal Statistical Society. Series B (Methodological)*, 14(1), 107–114.

HANSEN, L. P., AND T. J. SARGENT (2001): "Robust Control and Model Uncertainty," *American Economic Review*, 91(2), 60–66.

HANSEN, L. P., AND T. J. SARGENT (2008): *Robustness*. Princeton University Press.

HANSEN, L. P., AND T. J. SARGENT (2019): "Structured Uncertainty and Model Misspecification," *Working Paper, University of Chicago*.

HANSEN, L. P., T. J. SARGENT, G. TURMUHAMBETOVA, AND N. WILLIAMS (2006): "Robust Control and Model Misspecification," *Journal of Economic Theory*, 128(1), 45–90.

HELLINGER, E. (1909): "Neue Begründung der Theorie Quadratischer Formen von Unendlichvielen Veränderlichen," *Journal für die Reine und Angewandte Mathematik*, 136, 210–271.

HONG, H., AND J. C. STEIN (1999): "A Unified Theory of Underreaction, Momentum Trading, and Overreaction in Asset Markets," *The Journal of Finance*, 54(6), 2143–2184.

JIN, L. J., AND P. SUI (2019): "Asset Pricing with Return Extrapolation," *Working Paper, California Institute of Technology*.

KOBYLANSKI, M. (2000): "Backward Stochastic Differential Equations and Partial Differential Equations with Quadratic Growth," *The Annals of Probability*, 28(2), 558–602.

KULLBACK, S., AND R. A. LEIBLER (1951): "On Information and Sufficiency," *Annals of Mathematical Statistics*, 22(1), 79–86.

LAZRAK, A., AND M.-C. QUENEZ (2003): "A Generalized Stochastic Differential Utility," *Mathematics of Operations Research*, 28(1), 154–180.

LI, K., AND J. LIU (2019): "Extrapolative Asset Pricing," *Working Paper, University of California, San Diego*.

MACCHERONI, F., M. MARINACCI, AND A. RUSTICHINI (2006a): "Ambiguity Aversion, Robustness, and the Variational Representation of Preferences," *Econometrica*, 74(6), 1447–1498.

——— (2006b): "Dynamic Variational Preferences," *Journal of Economic Theory*, 128(1), 4–44.

MAENHOUT, P. J. (2004): "Robust Portfolio Rules and Asset Pricing," *The Review of Financial Studies*, 17(4), 951–983.

MERTON, R. C. (1969): "Lifetime Portfolio Selection under Uncertainty: The Continous-Time Case," *Review of Economics and Statistics*, 51(3), 247–257.

——— (1973): "An Intertemporal Capital Asset Pricing Model," *Econometrica*, 41(5), 867–887.

NAGEL, S., AND Z. XU (2019): "Asset Pricing with Fading Memory," *NBER Working Paper, University of Chicago.*

ORLOWSKI, P., A. SALI, AND F. TROJANI (2019): "Arbitrage Free Dispersion," Swiss Finance Institute Research Paper No 19-20.

PRIMICERI, G. (2005): "Time-Varying Structural Vector Autoregressions and Monetary Policy," *Review of Economic Studies*, 72, 821–852.

SAMUELSON, P. (1969): "Lifetime Portfolio Selection by Dynamic Stochastic Programming," *Review of Economics and Statistics*, 51(3), 239–246.

STRZALECKI, T. (2011): "Axiomatic Foundations of Multiplier Preferences," *Econometrica*, 79(1), 47–73.

SZŐKE, B. (2019): "Estimating Robustness," Working Paper, New York University.

TSAI, J., AND J. A. WACHTER (2015): "Disaster Risk and Its Implications for Asset Pricing," *Annual Review of Financial Economics*, 7, 219–252.

Appendix A Proofs

Proof of Lemma 1

Using Itô's formula on $D_{t,s}$ defined in equation (4) yields that

$$dD_{t,s} = d\phi(Z_{t,s}) = \tfrac{Z_{t,s}-Z^{\eta}_{t,s}}{1-\eta}(-u'_s)dB^{\mathbb{B}}_s + \tfrac{1}{2}Z^{\eta}_{t,s}|u_s|^2 ds. \tag{A.1}$$

It follows from Hölder's inequality that

$$\mathbb{E}^{\mathbb{B}}_t\Big[\int_t^T e^{-\delta(s-t)}(Z_{t,s}-Z^{\eta}_{t,s})^2\Psi_s^2|u_s|^2 ds\Big] \le C^2\mathbb{E}^{\mathbb{B}}_t\Big[\int_t^T e^{-\delta(s-t)}(Z_{t,s}-Z^{\eta}_{t,s})^{2q}ds\Big]^{\frac{1}{q}}\mathbb{E}^{\mathbb{B}}_t\Big[\int_t^T e^{-\delta(s-t)}\Psi_s^{2p}ds\Big]^{\frac{1}{p}},$$

where $C = \max|u|$ and $1/p + 1/q = 1$. Because $\mathbb{E}^{\mathbb{B}}_t\big[\int_t^T e^{-\delta(s-t)}\Psi_s^{2p}ds\big] < \infty$ with some $p > 1$ by assumption and $\mathbb{E}^{\mathbb{B}}_t\Big[\int_t^T e^{-\delta(s-t)}(Z_{t,s} - Z^{\eta}_{t,s})^{2q}ds\Big] < \infty$ due to the boundedness of u, the process $\{e^{-\delta(s-t)}\Psi_s(Z_{t,s} - Z^{\eta}_{t,s})u_s\}_{s\ge t}$ is square integrable under $\mathbb{B}$. Hence $\int_t^{\cdot} e^{-\delta(s-t)}\Psi_s(Z_{t,s} - Z^{\eta}_{t,s})(-u'_s)dB^{\mathbb{B}}_s$ is a martingale under $\mathbb{B}$. Then we have from (3) and (A.1) that

$$R^{\mathbb{U}}_t = \frac{1}{2\theta\Phi_t}\mathbb{E}^{\mathbb{B}}_t\Big[\int_t^T e^{-\delta(s-t)}\Psi_s Z^{\eta}_{t,s}|u_s|^2 ds\Big].$$

When $\Phi_t = Z_t^{1-\eta}$,

$$R^{\mathbb{U}}_t = \frac{1}{2\theta}\mathbb{E}^{\mathbb{B}}_t\Big[\int_t^T e^{-\delta(s-t)}\Psi_s Z_{t,s}Z_s^{\eta-1}|u_s|^2 ds\Big] = \frac{1}{2\theta}\mathbb{E}^{\mathbb{U}}_t\Big[\int_t^T e^{-\delta(s-t)}\Psi_s Z_s^{\eta-1}|u_s|^2 ds\Big].$$

Then

$$\begin{aligned}R^{\mathbb{U}}_t =&\mathbb{E}^{\mathbb{U}}_t\Big[\int_t^{\tilde t} e^{-\delta(s-t)}\frac{1}{2\theta}\Psi_s Z_s^{\eta-1}|u_s|^2 ds + e^{-\delta(\tilde t-t)}\int_{\tilde t}^T e^{-\delta(s-\tilde t)}\frac{1}{2\theta}\Psi_s Z_s^{\eta-1}|u_s|^2 ds\Big]\\ =&\mathbb{E}^{\mathbb{U}}_t\Big[\int_t^{\tilde t} e^{-\delta(s-t)}\frac{1}{2\theta}\Psi_s Z_s^{\eta-1}|u_s|^2 ds + e^{-\delta(\tilde t-t)}R^{\mathbb{U}}_{\tilde t}\Big].\end{aligned}$$

Proof of Proposition 1

The first result follows from the Comparison Theorem of El Karoui, Peng, and Quenez (1997). The second result is a direct consequence of (11). And the third result is a consequence of (1).

Proof of Proposition 2

In the following, we prove all portfolio choice results for Epstein-Zin utility. Proposition 2 is then a special case. To this end, consider an agent whose preference over consumption streams is described by a continuous time stochastic differential utility of Kreps-Porteus and Epstein-Zin type. Given the discount rate δ, the relative risk aversion $0 < \gamma \neq 1$, and the EIS $0 < \psi \neq 1$, the Epstein-Zin aggregator f (see, e.g., Duffie and Epstein (1992)) is

$$f(c, v) \equiv \delta\frac{c^{1-\frac{1}{\psi}}}{1-\frac{1}{\psi}}\big((1-\gamma)v\big)^{1-\frac{1}{\nu}} - \delta\nu v,$$

where $\nu = \frac{1-\gamma}{1-\frac{1}{\psi}}$. Incorporating the Cressie-Read distance, we introduce the robust Epstein-Zin preference for a consumption stream c as

$$\mathcal{U}_t^c = \inf_u \mathbb{E}_t^{\mathbb{U}}\Big[\int_t^T f(c_s, \mathcal{U}_s^c) + \tfrac{1}{2\theta}\Psi_s Z_{s\wedge\tau}^{\eta-1}|u_s|^2\, ds + \epsilon u(c_T)\Big], \tag{A.2}$$

where $\tau = \inf\{t \geq 0 : Z_t \leq \underline{z} \text{ or } Z_t > \overline{z}\}$.

Following the discussion in Section 1.2, $\mathcal{U}^c$ follows the following BSDE

$$d\mathcal{U}_t^c = \Big[\delta \mathcal{U}_t^c - f(c_t, \mathcal{U}_t^c) + \frac{\theta}{2}\frac{|\Gamma_t|^2}{\Psi_t Z_{t\wedge\tau}^{\eta-1}}\Big]dt + \Gamma_t' dB_t^{\mathbb{B}}, \quad \mathcal{U}_T^c = \epsilon u(c_T),$$

The optimal subjective measure is induced by

$$u_t^* = \frac{\theta\, \Gamma_t}{\Psi_t Z_{t\wedge\tau}^{\eta-1}}.$$

Consider a portfolio choice problem with robust Epstein-Zin preferences. The optimal consumption-investment problem (19) for this agent is then

$$V_t = \inf_u \sup_{\pi, c} \mathbb{E}_t^{\mathbb{U}}\Big[\int_t^T f(c_s, V_s) + \frac{1-\gamma}{2\theta} V_s Z_{s\wedge\tau}^{\eta-1}|u_s|^2\, ds + \epsilon u(c_T)\Big]. \tag{A.3}$$

The choice of Ψ in (18) ensures the following decomposition of the optimal value

$$V_t = \frac{W_t^{1-\gamma}}{1-\gamma} e^{f(t, x_t)}, \tag{A.4}$$

where x follows (25). Using dynamic programming, we obtain the HJB equation satisfied by f and summarize the agent's optimal investment and consumption strategies as follows.

Proposition 5. *When $\gamma \in (0,1)$, the function f defined in (26) satisfies*

$$\begin{aligned} 0 = \inf_u \sup_{\pi, \tilde{c}} \Big\{ &\partial_t f + \tfrac{1}{2}|u|^2\Big(\partial_{xx}^2 f - \partial_x f + (\partial_x f)^2\Big) + (1-\gamma)\partial_x f \pi' \sigma u + \delta \nu \tilde{c}^{1-\frac{1}{\psi}} e^{-\frac{1}{\nu} f} \\ &+ (1-\gamma)\big[r + \pi'(\mu - r - \sigma u) - \tilde{c} - \tfrac{1}{2}\gamma \pi' \Sigma \pi\big] - \delta\nu + \tfrac{1-\gamma}{2\theta} e^{(1-\eta)x}|u|^2 \Big\}, \end{aligned} \tag{A.5}$$

for $(t,x) \in [0,T) \times (\underline{x}, \overline{x})$, with boundary conditions

$$f(t, \underline{x}) = f_{\underline{x}}^{ent}(t), \quad f(t, \overline{x}) = f_{\overline{x}}^{ent}(t), \quad \text{and} \quad f(T, x) = \log \epsilon. \tag{A.6}$$

When $\gamma > 1$, the infimum and supremum in (A.5) are changed to $\sup_u \inf_{\pi, \tilde{c}}$. If Σ is positive definite and $\partial_{xx}^2 f - \partial_x f + (\partial_x f)^2 + \frac{1-\gamma}{\theta} e^{(1-\eta)x} > 0$ when $\gamma \in (0,1)$ (resp. < 0 when $\gamma > 1$), for any $(t,x) \in [0,T) \times (\underline{x}, \overline{x})$, then the agent's optimal belief and strategies are given by

$$\pi^* = \Big(\gamma + \frac{(1-\gamma)(1-\partial_x f)^2}{\partial_{xx}^2 f - \partial_x f + (\partial_x f)^2 + \frac{1-\gamma}{\theta} e^{(1-\eta)x}}\Big)^{-1} \Sigma^{-1}(\mu - r), \tag{A.7}$$

$$u^* = \frac{(1-\gamma)(1-\partial_x f)}{\partial_{xx}^2 f - \partial_x f + (\partial_x f)^2 + \frac{1-\gamma}{\theta} e^{(1-\eta)x}} \sigma' \pi^*, \tag{A.8}$$

$$\tilde{c}^* = \delta^{\psi} e^{-\frac{\psi}{\nu} f}. \tag{A.9}$$

Function f_x^{ent} in (A.6) is the value for the problem with an entropy cost and $\theta(x) = \theta e^{(\eta-1)x}$. It then satisfies the ODE given in (A.12).

Proof. Dynamic programming implies that

$$\tilde{V}_t = V_t + \int_0^t f(c_s, V_s) + \tfrac{1-\gamma}{2\theta} V_s e^{(1-\eta)x_s} |u_s|^2 \, ds,$$

where V is given in (A.3), is a martingale under $\mathbb{U}$ when u, π, c are agent's optimal strategy and $t < \tau$. To calculate the drift of $\tilde{V}$, we use equation (25) and apply Itô's formula to derive

$$de^{f(t,x_t)} = e^{f(t,x_t)} \Big[\partial_t f + \tfrac{1}{2} |u_t|^2 \Big(\partial^2_{xx} f - \partial_x f + (\partial_x f)^2 \Big) \Big] dt + e^{f(t,x_t)} \partial_x f u_t' dB_t^{\mathbb{U}}.$$

Moreover define $\tilde{c} = \frac{c}{W}$ as the consumption-wealth ratio. Then

$$d\frac{W_t^{1-\gamma}}{1-\gamma} = W_t^{1-\gamma} \big[r + \pi'(\mu - r - \sigma u) - \tilde{c} - \tfrac{1}{2}\gamma \pi' \Sigma \pi \big] dt + W_t^{1-\gamma} \pi' \sigma dB_t^{\mathbb{U}}, \tag{A.10}$$

Combining the previous two equations, we obtain the drift of $\tilde{V}$ (divided throughout by $W^{1-\gamma} e^f(t, x_t)$)

$$\begin{aligned} & r + \pi'(\mu - r - \sigma u) - \tilde{c} - \tfrac{1}{2}\gamma \pi' \sigma \pi + \tfrac{1}{1-\gamma} \big[\partial_t f + \tfrac{1}{2}|u|^2 (\partial^2_{xx} f - \partial_x f + (\partial_x f)^2) \big] + \partial_x f \pi' \sigma u \\ & + \tfrac{\delta}{1-\frac{1}{\psi}} \tilde{c}^{1-\frac{1}{\psi}} e^{-\frac{1}{\nu} f} - \delta \tfrac{\nu}{1-\gamma} + \tfrac{1}{2\theta} e^{(1-\eta)x} |u|^2. \end{aligned}$$

Then the HJB equation for f is

$$\begin{aligned} 0 = \inf_u \sup_{\pi, \tilde{c}} \Big\{ & \partial_t f + \tfrac{1}{2}|u|^2 \Big(\partial^2_{xx} f - \partial_x f + (\partial_x f)^2 \Big) + (1-\gamma) \partial_x f \pi' \sigma u + \delta \nu \tilde{c}^{1-\frac{1}{\psi}} e^{-\frac{1}{\nu} f} \\ & + (1-\gamma) \big[r + \pi'(\mu - r - \sigma u) - \tilde{c} - \tfrac{1}{2}\gamma \pi' \Sigma \pi \big] - \delta\nu + \tfrac{1-\gamma}{2\theta} e^{(1-\eta)x} |u|^2 \Big\}, \end{aligned} \tag{A.11}$$

when $\gamma \in (0,1)$. The infimum and supremum changed to $\sup_u \inf_{\pi, \tilde{c}}$ in the previous equation when $\gamma > 1$.

The first-order condition for π yields

$$\pi^* = \tfrac{1}{\gamma} \Sigma^{-1} \Big(\mu - r - (1 - \partial_x f) \sigma u^* \Big).$$

This is the agent's optimal strategy when Σ is positive definite. The first order condition in u yields

$$u^* = \frac{(1-\gamma)(1 - \partial_x f)}{\partial^2_{xx} f - \partial_x f + (\partial_x f)^2 + \frac{1-\gamma}{\theta} e^{(1-\eta)x}} \sigma' \pi^*.$$

This is the agent's optimal belief choice if $\partial^2_{xx} f - \partial_x f + (\partial_x f)^2 + \frac{1-\gamma}{\theta} e^{(1-\eta)x} > 0$. The agent's optimal choice of consumption wealth is

$$\tilde{c}^* = \delta^{\psi} e^{-\frac{\psi}{\nu} f}.$$

When the state variable x reaches the boundaries $\underline{x}$ and $\overline{x}$, the x is absorbed there, and the problem becomes a problem where the cressie-Read penalty in (A.3) becomes

$$\tfrac{1-\gamma}{2\theta} V_s e^{(1-\eta)x_\tau} |u_s|^2, \qquad \text{for } s \geq \tau.$$

Effectively, this is an entropy penalty

$$\tfrac{1-\gamma}{2\theta(x_\tau)} V_s |u_s|^2, \qquad \text{where } \theta(x_\tau) = \theta e^{(\eta-1)x_\tau}.$$

As a result, the boundary conditions of f at $\underline{x}$ and $\overline{x}$ are specified by the value f^{ent}_x with the robust parameter $\theta(\underline{x})$ or $\theta(\overline{x})$. Setting the spatial derivatives to be zero in (A.5), f^{ent}_x, with $x = \underline{x}$ or $\overline{x}$, satisfies the following ODE

$$
\begin{aligned}
0 =& \partial_t f^{\text{ent}}_x - \delta\nu + (1-\gamma)\big[r + (\pi^{\text{ent}})'(\mu - r - \sigma u^{\text{ent}}) - \tfrac{1}{2}\gamma(\pi^{\text{ent}})'\Sigma\pi^{\text{ent}}\big] \\
&+ \delta^{\psi}\tfrac{\nu}{\psi}e^{-\frac{\psi}{\nu}f^{\text{ent}}_x} + \tfrac{1-\gamma}{2\theta}e^{(1-\eta)x}|u^{\text{ent}}|^2, \\
f^{\text{ent}}_x(T) =& \log\epsilon,
\end{aligned}
\tag{A.12}
$$

where

$$
\pi^{\text{ent}} = \frac{1}{\gamma + \theta e^{(\eta-1)x}}\Sigma^{-1}(\mu - r) \quad \text{and} \quad u^{\text{ent}} = \frac{\theta e^{(\eta-1)x}}{\gamma + \theta e^{(\eta-1)x}}\sigma'\Sigma^{-1}(\mu - r).
$$

□

We revisit the utility formulation in Section 1. The optimal distortion u^* obtained above is a continuous function evaluated on $(t, X_t) \in [0,T]\times[\underline{x},\overline{x}]$. Hence u^* is bounded and satisfies the requirement in Section 1.1. Moreover, by our construction of the HJB equation (A.5), the process defined in (A.4) satisfies the BSDE associated to (A.3). Hence it is a utility index defined in Section 1. Finally, the optimal portfolio weight π^* and consumption-wealth ratio $\tilde{c}^*$ are bounded, because they are bounded functions evaluated on $(t, X_t) \in [0,T]\times[\underline{x},\overline{x}]$. Therefore the associated wealth process W has all finite moments. Thanks to the boundedness of f, Ψ defined in (21) satisfies the integrability assumption required in Lemma 1.

Proof of Proposition 3

When the agent invests in asset P, her optimal consumption and investment problem can be solved as in Section 2. Instead of constants μ and σ, μ^P and σ^P depend on the state variable x. However, when the portfolio choice problem is solved in Proposition 5, state variable x is already taken into account. Therefore, even μ^P and σ^P are now random, no more state variable needs to be introduced and the function f in (26) still solves (A.5) with μ and σ therein replaced by μ^P and σ^P. The optimal belief and strategies are given by (A.7), (A.8), and (A.9).

From consumption market clearing and (A.9),

$$
c_t = \delta^{\psi}e^{-\frac{\psi}{\nu}f(t,x_t)}W_t. \tag{A.13}
$$

Applying Itô's formula on the right-hand side, yields

$$
de^{-\frac{\psi}{\nu}f(t,x_t)} = -\tfrac{\psi}{\nu}e^{-\frac{\psi}{\nu}f(t,x_t)}\Big[\partial_t f + \tfrac{1}{2}|u|^2\big(\partial^2_{xx}f + \partial_x f - \tfrac{\psi}{\nu}(\partial_x f)^2\big)|u|^2\Big]dt - \tfrac{\psi}{\nu}e^{-\frac{\psi}{\nu}f(t,x_t)}\partial_x f u\, dB^{\mathbb{B}}_t.
$$

Then using capital market clearing $\pi^* = 1$ and (A.9), we obtain

$$
\begin{aligned}
de^{-\frac{\psi}{\nu}f(t,x_t)}W_t =& W_t de^{-\frac{\psi}{\nu}f(t,x_t)} + e^{-\frac{\psi}{\nu}f(t,x_t)}dW_t + d\langle e^{-\frac{\psi}{\nu}f(t,x_t)}, W_t\rangle_t \\
=& e^{-\frac{\psi}{\nu}f}W_t\Big[-\tfrac{\psi}{\nu}\partial_t f - \tfrac{\psi}{2\nu}|u|^2\Big(\partial^2_{xx}f + \partial_x f - \tfrac{\psi}{\nu}(\partial_x f)^2\Big)\Big]dt \\
&+ e^{-\frac{\psi}{\nu}f}W_t(\mu^P - \delta^{\psi}e^{-\frac{\psi}{\nu}f})dt - e^{-\frac{\psi}{\nu}f}W_t\tfrac{\psi}{\nu}\partial_x f u\sigma^P\, dt \\
&+ e^{-\frac{\psi}{\nu}f}W_t\big[-\tfrac{\psi}{\nu}\partial_x f u + \sigma^P\big]dB^{\mathbb{B}}_t.
\end{aligned}
$$

Using the previous dynamics and matching the drift and volatility on both sides of (A.13), we obtain

$$\mu^P = \mu^c + \tfrac{\psi}{\nu}\partial_t f + \tfrac{\psi}{2\nu}|u|^2\Big(\partial^2_{xx} f + \partial_x f - \tfrac{\psi}{\nu}(\partial_x f)^2\Big) + \delta^\psi e^{-\frac{\psi}{\nu}f} + \tfrac{\psi}{\nu}\partial_x f u \sigma^P, \tag{A.14}$$

$$\sigma^P = \sigma^c + \tfrac{\psi}{\nu}\partial_x f u. \tag{A.15}$$

Plugging (A.15) into the right-hand side of (A.14), we transform μ^P into

$$\mu^P = \mu^c + \tfrac{\psi}{\nu}\partial_t f + \tfrac{\psi}{2\nu}|u|^2\Big(\partial^2_{xx} f + \partial_x f + \tfrac{\psi}{\nu}(\partial_x f)^2\Big) + \delta^\psi e^{-\frac{\psi}{\nu}f} + \tfrac{\psi}{\nu}\partial_x f u \sigma^c \tag{A.16}$$

Combining (A.7) and (A.16), we obtain from capital market clearing that

$$\mu^P - r = \Big[\gamma + \frac{(1-\gamma)(1-\partial_x f)^2}{\partial^2_{xx} f - \partial_x f + (\partial_x f)^2 + \frac{1-\gamma}{\theta}e^{(1-\eta)x}}\Big](\sigma^P)^2. \tag{A.17}$$

Plug (A.15) and (A.16) back into (A.5) and simplify, we obtain

$$\begin{aligned} 0 =& \partial_t f + \tfrac{1}{2}|u^*|^2\Big(\partial_{xx} f - \partial_x f + \psi(\partial_x f)^2\Big) + (1-\gamma)\partial_x f u \sigma^c - \delta\tfrac{\nu}{\psi} \\ &+ \tfrac{1-\gamma}{\psi}\Big(\mu^c - u^*\sigma^c - \tfrac{1}{2}\gamma(\sigma^c)^2\Big) + \delta^\psi \tfrac{\nu}{\psi} e^{-\frac{\psi}{\nu}f} + \tfrac{1-\gamma}{2\psi\theta}e^{(1-\eta)x}|u^*|^2, \end{aligned}$$

where u^* in (37) is obtained by plugging (A.15) into (A.8) and solving for u^*. Finally, (40) follows from combining (35), (A.16) and (A.17).

When x reaches the boundary $\underline{x}$ or $\overline{x}$, f is specified by F^{ent}_x which is the value function in an equilibrium with an entropy cost. F^{ent}_x satisfies the following ODE

$$\begin{aligned} 0 =& \partial_t F^{\text{ent}}_x - \delta\tfrac{\nu}{\psi} + \tfrac{1-\gamma}{\psi}\big(\mu^c - u^{\text{ent}}\sigma^c - \tfrac{1}{2}\gamma(\sigma^c)^2\big) + \delta^\psi\tfrac{\nu}{\psi}e^{-\frac{\psi}{\nu}F^{\text{ent}}_x} + \tfrac{1-\gamma}{2\psi\theta}e^{(1-\eta)x}|u^{\text{ent}}|^2, \\ F^{\text{ent}}_x(T) =& \log\epsilon, \end{aligned} \tag{A.18}$$

where

$$u^{\text{ent}} = \theta e^{(\eta-1)x}\sigma^c.$$

Proof of Proposition 4

It follows from (45) and the definition of $\mathbb{Q}$ that

$$S_t = \mathbb{E}^{\mathbb{Q}}_t\Big[\int_t^T e^{-\int_t^s r_v dv} D_s ds\Big] \tag{A.19}$$

Let $\ell = S/D$ be the price-dividend ratio. It follows from the previous equation that $\widetilde{D}_t = e^{-\int_0^t r_v dv} D_t \ell(x_t) + \int_0^t e^{-\int_0^s r_v dv} D_s ds$ is a $\mathbb{Q}$-martingale. Given that the dynamics of x and $\overline{D}$ under $\mathbb{Q}$ are

$$\begin{aligned} dx_t =& \big(\frac{1}{2}|u^*_t|^2 - u^*_t\lambda_t\big)dt + u^*_t dB^{\mathbb{Q}}_t, \\ \frac{dD_t}{D_t} =& \big(\mu^D - \rho\sigma^D\lambda_t\big)dt + \sigma^D(\rho dB^{\mathbb{Q}}_t + \sqrt{1-\rho^2}dB^\perp_t). \end{aligned}$$

Equating the drift of $\widetilde{D}$ to be zero, we get equation (47) for ℓ.

To obtain μ^S and σ^S, we apply Itô's formula to $S_t = D_t \ell(t, x_t)$ to obtain

$$dS_t = d\big(D_t\ell(t,x_t)\big) = D_t\ell(t,x_t)\Big(\frac{\partial_t \ell}{\ell} + \frac{1}{2}|u^*|^2\frac{\partial_x \ell}{\ell} + \frac{1}{2}|u^*|^2\frac{\partial^2_{xx}\ell}{\ell} + \mu^D + u^*\sigma^D\rho\frac{\partial_x \ell}{\ell}\Big)dt$$
$$+ D_t\ell(t,x_t)\Big[\big(u^*\frac{\partial_x\ell}{\ell} + \overline{\sigma}\rho\big)dB^{\mathbb{B}}_t + \overline{\sigma}\sqrt{1-\rho^2}dB^{\perp}_t\Big].$$

Adding $D_t dt$ on both sides and dividing by $S_t = D_t\ell$, we obtain

$$\frac{dS_t + D_t dt}{S_t} = \Big(\frac{\partial_t \ell}{\ell} + \frac{1}{2}|u^*|^2\frac{\partial_x \ell}{\ell} + \frac{1}{2}|u^*|^2\frac{\partial^2_{xx}\ell}{\ell} + \mu^D + u^*\sigma^D\rho\frac{\partial_x \ell}{\ell} + \frac{1}{\ell}\Big)dt$$
$$+ \big(u^*\frac{\partial_x\ell}{\ell} + \overline{\sigma}\rho\big)dB^{\mathbb{B}}_t + \overline{\sigma}\sqrt{1-\rho^2}dB^{\perp}_t.$$

Matching the previous equation with equation (46), we obtain μ^S, σ^S, and $\sigma^{S,\perp}$.
Finally, to obtain the CAPM relation, we note that $S_t M_t + \int_0^t M_s D_s ds$ is a $\mathbb{P}^{\mathbb{U}}$-martingale. Then the CAPM relation follows from combining (42), (43), and (46).

Appendix B Log utility

We study the portfolio choice problem for an agent with the log utility by taking the scaling limit of Proposition 2 as $\gamma \to 1$.

For the function f in (26), define h via

$$h(t,x) = \tfrac{f(t,x)}{1-\gamma}.$$

Then the agent's value function (after adding a constant $-\frac{1}{1-\gamma}$) is

$$V = \frac{W^{1-\gamma}e^{(1-\gamma)h} - 1}{1-\gamma} = \frac{e^{(1-\gamma)(\log W + h)} - 1}{1-\gamma}.$$

As $\gamma \to 1$, the right-hand side converges to $\log W + h$, which is the value function in the log utility case with entropy cost in Maenhout (2004), Appendix B.

From (27), we derive the equation satisfied by h:

$$0 = \partial_t h + \tfrac{1}{2}|u|^2\Big(\partial^2_{xx}h - \partial_x h + (1-\gamma)(\partial_x h)^2\Big) - \tfrac{\delta}{1-\gamma} + (1-\gamma)\partial_x h\pi'\sigma u$$
$$+ [r + \pi'(\mu - r - \sigma u) - \tfrac{1}{2}\gamma\pi'\Sigma\pi] + \tfrac{\gamma}{1-\gamma}\delta^{\frac{1}{\gamma}}e^{(1-\gamma)(-\frac{h}{\gamma})} + \tfrac{1}{2\theta}e^{(1-\eta)x}|u|^2 = 0, \tag{B.1}$$

where

$$\pi = \frac{1}{\gamma + \frac{(1-(1-\gamma)\partial_x h)^2}{\partial^2_{xx}h + (1-\gamma)(\partial_x h)^2 - \partial_x h + \frac{1}{\theta}e^{(1-\eta)x}}}\Sigma^{-1}(\mu - r), \tag{B.2}$$

$$u = \frac{1-(1-\gamma)\partial_x h}{\partial^2_{xx}h + (1-\gamma)(\partial_x h)^2 - \partial_x h + \frac{1}{\theta}e^{(1-\eta)x}}\sigma'\pi. \tag{B.3}$$

Here we consider $\epsilon = 1$ in the bequest utility. Then the terminal condition for h is

$$h(T,x) = 0.$$

Now send $\gamma \to 1$ in (B.1), (B.2), and (B.3) to identify the limiting equation satisfied by h. We first consider the limit of

$$-\tfrac{\delta}{1-\gamma} + \tfrac{\gamma}{1-\gamma}\delta^{\frac{1}{\gamma}} e^{(1-\gamma)(-\frac{h}{\gamma})}.$$

When $\gamma \to 1$, up to the first order of $1-\gamma$,

$$-\tfrac{\delta}{1-\gamma} + \tfrac{\gamma}{1-\gamma}\delta^{\frac{1}{\gamma}} e^{(1-\gamma)(-\frac{h}{\gamma})} \approx \frac{-\delta + \gamma\delta^{\frac{1}{\gamma}}(1+(1-\gamma)(-\frac{h}{\gamma}))}{1-\gamma} = \frac{\delta^{\frac{1}{\gamma}} - \delta}{1-\gamma} - \delta^{\frac{1}{\gamma}} - \delta^{\frac{1}{\gamma}} h.$$

By L'Hospital rule, $\lim_{\gamma\to 1} \frac{\delta^{\frac{1}{\gamma}}-\delta}{1-\gamma} = \delta \log \delta$. Then

$$\lim_{\gamma \to 1} -\tfrac{\delta}{1-\gamma} + \tfrac{\gamma}{1-\gamma}\delta^{\frac{1}{\gamma}} e^{(1-\gamma)(-\frac{h}{\gamma})} = -\delta - \delta h + \delta \log \delta.$$

Using the previous identity, we obtain the limit of (B.1), (B.2), and (B.3) as

$$0 = \partial_t h + \tfrac{1}{2}|u|^2\big(\partial^2_{xx} h - \partial_x h\big) - \delta h - \delta + \delta \log \delta + r + \pi'(\mu - r - \sigma u) - \tfrac{1}{2}\pi'\Sigma\pi + \tfrac{1}{2\theta} e^{(1-\eta)x}|u|^2 = 0, \quad \text{(B.4)}$$

where

$$\pi = \frac{\partial^2_{xx} h - \partial_x h + \frac{1}{\theta} e^{(1-\eta)x}}{1 + \partial^2_{xx} h - \partial_x h + \frac{1}{\theta} e^{(1-\eta)x}} \Sigma^{-1}(\mu - r), \quad \text{(B.5)}$$

$$u = \frac{1}{\partial^2_{xx} h - \partial_x h + \frac{1}{\theta} e^{(1-\eta)x}} \sigma'\pi. \quad \text{(B.6)}$$

From (B.5) and (B.6), we observe that the portfolio choice for the log utility agent is still dynamic. The log utility agent still hedges against the future belief variation. This can be also seen from the following identity,

$$\sigma u + \Sigma\pi = \mu - r,$$

which is obtained after combining (B.5) and (B.6). The variation of u drives the variation of π so that the sum of σu and $\Sigma\pi$ is always constant.