Model Ambiguity versus Model Misspecification in Dynamic Portfolio Choice and Asset Pricing *

Anne G. Balter† Pascal J. Maenhout‡ Hao Xing§

February 2023

Abstract

We study aversion to model ambiguity and misspecification in dynamic portfolio choice. Investors with relative risk aversion $\gamma > 1$ fear return persistence, while risk-tolerant investors ($0 < \gamma < 1$) fear return mean reversion, to confront model misspecification concerns when facing a model with IID returns. The intuition is that risk-averse (risk-tolerant) investors who are keen to hedge (speculate) intertemporally worry about an endogenous worst-case misspecification where returns persist (mean-revert) so that hedging (speculation) is impossible. A log investor is myopic and unaffected by model misspecification, therefore only worrying about model ambiguity among IID models. Rather than the multiplier approach of Hansen and Sargent (2001) we utilize a constraint approach, preserving homotheticity and tractability. Our model can explain evidence for the experience hypothesis, for non-participation in equity markets, as well as for extrapolative return expectations. In equilibrium, we show that model misspecification, unlike model ambiguity, can generate excess volatility, even with IID fundamentals.

*We thank Harjoat Bhamra, John Campbell, Gary Chamberlain, Bernard Dumas, and seminar participants at INSEAD for insightful comments and discussions. All errors are ours.

†Tilburg University and Netspar, a.g.balter@tilburguniversity.edu.

‡INSEAD, pascal.maenhout@insead.edu.

§Boston University, haoxing@bu.edu.

The seminal work of Hansen and Sargent (2001) underscores the importance of model uncertainty for economic agents. Indeed, when solving dynamic problems, they face complex model selection and estimation challenges, not unlike econometricians fearing misspecification of the models they estimate. It therefore seems plausible that households and investors seek dynamic decisions that perform well not only when the models or probability laws used to derive them are specified correctly, but also when there is some suspicion of misspecification. The robust control literature initiated by Hansen and Sargent captures this preference for robustness or aversion to model uncertainty by introducing a fictitious malevolent agent who, in the mind of the decision-maker, distorts the model in an adverse way, namely in order to minimize the resulting expected utility.[1] In this approach, the agent seeking robustness entertains a family of models constructed as a neighborhood around a benchmark model and optimizes against the worst-case model within this family.

In recent work, Hansen and Sargent (2022) emphasize the difference between *model ambiguity* and *model misspecification*. Robustness against model ambiguity involves guarding against a worst-case alternative within a set of *structured* models. In contrast, decision-makers fearing model misspecification consider a wider set of alternatives, including *unstructured* models. We study aversion to model ambiguity and misspecification in dynamic portfolio choice problems. The distinction between model ambiguity and model misspecification is particularly intuitive in a two-period binomial tree model with IID returns, which we use to study the problem in its simplest setting. An investor concerned with model ambiguity is uncertain about the probability of the high return outcome, but does not doubt the IID specification. Model misspecification concerns on the other hand also induce doubts about the IID nature of the return process and make the investor consider non-IID return specifications.

We use a constraint approach, rather than the typical multiplier approach of Hansen and Sargent (2001), extended in Hansen and Sargent (2022). The constraint approach preserves homotheticity or scale-invariance, which is economically appealing for the portfolio choice and finance applications that we are interested in. Crucially, the homotheticity or scale-invariance of our approach brings tractability and allows us to obtain explicit and intuitive results. This is the first contribution of the paper. In our constraint approach the difference between model ambiguity and model misspecification in Hansen and Sargent (2022) concerns the way in which entropy constraints are imposed to construct the uncertainty sets entertained by the investor. Imposing only a single intertemporal entropy constraint leads to the widest uncertainty set, namely the set of unstructured models, and deals with model misspecification. With this single entropy constraint, a CRRA investor facing IID returns fears a worst-case model where returns

[1] See also Anderson et al. (2003) and Hansen et al. (2006), among many others, or Hansen and Sargent (2008) for a textbook treatment.

are not IID and expected returns are endogenously time-varying. In particular, an investor who is more (less) risk-averse than log guards against a worst-case model with procyclical (countercyclical) expected returns, resulting in negative hedging (speculative) demands. Despite returns being IID and the investor having CRRA preferences, concerns of model misspecification lead to optimal portfolios that are state-dependent, except for a log investor, who is myopic and only fears model ambiguity. Model ambiguity involves a narrower uncertainty set (the set of structured models) obtained when imposing two pathwise entropy constraints and results in behavior that is identical to Chen and Epstein (2002)'s $\kappa-$ignorance.

As a second contribution, we show how these results carry over to a continuous-time setting where we tackle model misspecification in the classical Merton portfolio problem by applying an idea of Hansen et al. (2006) to use continuation entropy as a state variable. In addition to preserving homotheticity, this approach spotlights the difference between model ambiguity and model misspecification, because the latter is shown to nest the former as a special case. As a consequence, our approach fleshes out the precise economic mechanism through which model misspecification is addressed. Mirroring the intuition of the binomial tree example, model misspecification is addressed by granting Nature the ability to endogenously allocate entropy in a state-dependent fashion, permitted by the single intertemporal entropy constraint in the binomial tree, and disallowed by the pathwise entropy constraints dictated by rectangularity. In continuous time, robustness against model misspecification is achieved when Nature is empowered to set the instantaneous volatility of the shadow price of the entropy state variable to zero, which is accomplished by controlling the instantaneous volatility of continuation entropy. We show that the worst-case model misspecification for a risk-averse investor (with relative risk aversion $\gamma > 1$), who is keen to hedge intertemporally, is that returns persist so that intertemporal hedging is not possible. Similarly, for a risk-tolerant investor ($0 < \gamma < 1$), who likes to speculate intertemporally, the endogenous worst-case misspecification is that stock returns display mean reversion, which precludes intertemporal speculation.

In terms of applications, our model can explain endogenous nonparticipation, also among wealthy investors, as documented extensively in the empirical household finance literature (see Campbell (2018) and Gomes et al. (2021) for a detailed discussion of the evidence). In our model, risk-averse and sufficiently impatient investors who experience a string of negative stock returns become increasingly concerned with model misspecification, to the point where nonparticipation can become optimal due to the endogenous increase in pessimism. Our paper can therefore be seen as providing a dynamic version of the uncertainty aversion explanation for nonparticipation in Dow and Werlang (1992). While Cao et al. (2005), Campanale (2011) and Peijnenburg (2018) also propose dynamic extensions of the uncertainty aversion explanation in Dow and Werlang (1992), ours is the first where the non-participation results endogenously from model misspecification concerns. The dynamic beliefs generated by our analysis are consistent with the experience hypothesis proposed and tested by Malmendier and Nagel (2011), according to which individual risk-taking behavior is shaped by the macroeconomic history (and corresponding historical returns) experienced by the individual. Unlike other work, where the experience hypothesis is typically explained by learning-from-experience dynamics, the mechanism in our paper does not rely on learning, but instead reflects the investor's state-dependent aversion to model uncertainty. Economic shocks directly affect the dynamics of the continuation entropy state variable, and thereby optimal risk-taking and savings behavior, since these are functions of the entropy state variable.

A related implication of our model is that the endogenous beliefs of risk-averse investors concerned with model misspecification can be seen as extrapolative. As explained above, the endogenous worst-case model misspecification for this investor is that returns persist so that intertemporal hedging is not possible. As a result, expected returns increase following high returns and decrease after low returns. A more subtle prediction of our model is that younger investors with shorter life histories are on average more sensitive to recent returns than older investors, in line also with the evidence in Malmendier and Nagel (2011).

On the methodological front, we explain why the log investor, commonly analyzed in the influential Hansen-Sargent multiplier approach, does not in fact deal with model misspecification, but only with model ambiguity, namely because the log investor is myopic and indifferent between intertemporal hedging and intertemporal speculation. In the Merton-Lucas dynamic portfolio model, mirroring the findings in the binomial tree example, CRRA preferences with non-unit risk aversion are required to study model misspecification. Our constraint approach has the advantage of preserving homotheticity and thereby tractability for general CRRA preferences, including $\gamma \neq 1$.

The constraint approach we use also naturally lends itself to a calibration of a 'reasonable' amount of initial entropy via detection-error probabilities. As an additional contribution, we show that existing approaches based on homothetic robustness are nested and equivalent to a homothetic approach to model ambiguity. We therefore provide an alternative and rigorous justification for the homothetic scaling proposed by Maenhout (2004) and commonly used in the literature[2], as well as a recipe for extending that work to the more general case of model misspecification.

As a final contribution, we solve a general equilibrium problem for an investor with Epstein-Zin preferences who worries about model ambiguity and model misspecification. We show that

[2] See, e.g., Skiadas (2003), Uppal and Wang (2003), Hansen (2004), Trojani and Vanini (2004), Liu et al. (2005), Maenhout (2006), Anderson et al. (2009), Liu (2010), Branger et al. (2013), Branger and Larsen (2013), Jin et al. (2019) and Luo et al. (2022).

the equilibrium equity premium is state-dependent and stochastic, even with IID fundamentals. The equity premium consists of three components: the traditional price of risk, the price of model ambiguity and the price of model misspecification. Model misspecification also plays a key role in generating excess volatility in equilibrium.

Related Literature: Our paper builds on the robust control literature initiated by Hansen and Sargent, using key results from a series of seminal papers, namely Hansen and Sargent (2001), Anderson et al. (2003), Hansen et al. (2006), and Hansen and Sargent (2022), among many others. Different from this literature, which uses the multiplier approach subsequently axiomatized by Maccheroni et al. (2006) and Strzalecki (2011), we follow the constraint approach based on an endogenous continuation entropy state variable as suggested in Hansen et al. (2006), in order to analyze model ambiguity and model misspecification. This approach guarantees homotheticity, and was also used by Drechsler (2013) and Aït-Sahalia and Matthys (2019), who both study model ambiguity with jump-diffusions. Instead, we study model misspecification as well as model ambiguity, focusing on diffusions without jumps. Our work nests and provides a rigorous justification for the 'homothetic scaling' of Maenhout (2004), which we show only deals with model ambiguity. Our results cast new light on several important findings in the robustness literature, and especially on the impact and economic implications of different ways of modelling entropy constraints and uncertainty sets in a dynamic setting.

Most closely related to our work are Bidder and Dew-Becker (2016) and Bidder and Smith (2018) who use very different approaches, but also find that models with non-IID dynamics can endogenously arise as worst-case models for pessimistic agents. They focus on macrofinance equilibrium models without explicit portfolio choice, and where agents have distorted beliefs about consumption rather than returns. In contrast, our approach allows a detailed study of dynamic portfolio choice and belief distortions about returns. This enables us to connect directly with recent empirical work in household finance about portfolio choice and biased beliefs, such as limited stock market participation and the experience hypothesis. Another advantage of our analysis is its recursive formulation and direct relationship to the literature on robustness, and in particular the influential multiplier approach of Hansen and Sargent (2001), Hansen et al. (2006), and Hansen and Sargent (2022), as explained in the previous paragraph.

A number of recent papers build on and extend the Hansen-Sargent robustness framework in a variety of ways. Hansen and Miao (2018) construct novel continuous-time formulations of smooth ambiguity preferences and investigate their asset pricing implications in Hansen and Miao (2022). Christensen (2019) studies identification and estimation of dynamic models with multiplier or constraint preferences, and presents results on the identification of the decision-maker's value function, worst-case belief distortion and continuation entropy. Bhandari et al.

(2022) develop a theory of model-consistent notions of pessimism and optimism based on the Hansen-Sargent robustness framework. Cerreia-Vioglio et al. (2022a) is the first paper to offer an axiomatic treatment of model misspecification concerns. Ai et al. (2022) present a dynamic model of endogenous information acquisition to explain the pre-FOMC announcement drift. They provide a recursive utility interpretation of the homothetic scaling of Maenhout (2004) in the context of the generalized risk sensitivity preferences of Ai and Bansal (2018). Bhamra and Uppal (2019) develop a model with differences in familiarity across assets based on the modeling approach of Uppal and Wang (2003) to study social welfare implications of familiarity biases.

Our paper is also related to different strands of the literature that study time-varying beliefs or risk aversion using alternative mechanisms. For example, Ai (2010), Collin-Dufresne et al. (2016), Dew-Becker and Nathanson (2019), and Kozlowski et al. (2020), among many others, generate rich belief dynamics endogenously in asset pricing models with simple fundamentals when agents are allowed to learn. Cogley and Sargent (2008) combine insights from the robust control literature with Bayesian learning to study the impact of a pessimistically distorted prior on the market price of risk. Asset pricing studies that feature time-varying risk aversion include the habit models of Constantinides (1990), Detemple and Zapatero (1991), and Campbell and Cochrane (1999) where time-varying risk aversion is tightly linked to the level of consumption relative to its recent past history. While in these models time-varying risk aversion is exogenously imposed on the utility function of the representative agent, in our setting, the state-dependence of effective risk aversion arises endogenously due to the investor's concern for model misspecification and the dynamics of the continuation entropy state variable.

Motivated by empirical and experimental evidence on belief formation, a vast literature in behavioral economics and finance proposes models of deviations from rational expectations. Prominent examples include diagnostic expectations as, e.g., in Bordalo et al. (2018), Bordalo et al. (2019), and Bordalo et al. (2020), and extrapolative expectations, see, e.g., Hong and Stein (1999), Barberis et al. (2015), Barberis (2018), Jin and Sui (2022), and Li and Liu (2019). Our model endogenously generates similar belief dynamics, in the sense that a risk-averse investor endogenously fears return persistence, as this is indeed the worst-case misspecification for a long-term investor who desires Mertonian intertemporal hedging.

Outline of the paper: The paper is organized as follows. Section 1 studies a two-period binomial tree to gain intuition and to show the main effects at work in the simplest possible setting. Section 2 presents a general continuous-time set-up with Brownian information structures. Section 3 explains the difference between model misspecification and model ambiguity in our setting, connecting with the multiplier approach and highlighting the special case of log utility. Section 4 presents the quantitative implications of our model in a dynamic Merton portfolio problem. Section 5 discusses results for equilibrium asset pricing. Finally, Section 6 concludes.

1 Entropy Constraints in a Two-Period Binomial Tree

To gain intuition about the different ways of modelling asset allocation in the face of both risk and uncertainty, we follow Knox (2002) and consider a two-period binomial tree. Under a baseline measure $\mathbb{B}$, returns on the risky asset are identically and independently distributed (IID), with possible outcomes H and L, which have probability p and $1-p$, resp. Because the agent is uncertain about the baseline measure she seeks a robust investment decision by considering different probabilities according to a set of alternative measures. The uncertainty set is described by a bound on some divergence measure of the probabilities under the alternative measure $\mathbb{U}$ relative to the probabilities under $\mathbb{B}$.

Unlike Knox (2002) who focuses on the multiplier (or penalty) formulation of robust control, we find it useful to zoom in on the constraint formulation for reasons of tractability. This approach guarantees homotheticity and has further advantages as discussed later. We otherwise follow the robust control literature and use the relative entropy (or Kullback-Leibler) divergence to index dissimilarity between probability measures. We discuss the connection to the multiplier approach in detail in the continuous-time analysis in Section 3.1.

Let q denote the probability of an up-move under the alternative measure. Model misspecification concerns induce the investor to explore alternative unstructured models that need not follow the assumption of IID returns. We therefore introduce subscripts to allow for time- and state-dependence: as shown in Figure 1, q_0 denotes the probability of an up-move from $t=0$ to $t=1$, and q_H (q_L) denotes the probability of an up-move from $t=1$ to $t=2$, conditional on being in the up-state (down-state) at $t=1$. Model ambiguity restricts the investor's attention to the set of structured models, i.e., the investor is uncertain about the probability p, but does not doubt the IID nature of the baseline measure.

Besides the risky asset with gross return H in the up-state and L in the down-state, the risk-free asset has a riskless gross return of R_f. No arbitrage requires that $H > R_f > L$.

The investor maximizes CRRA expected utility from terminal wealth with coefficient of relative risk aversion $\gamma > 0$,

$$U(W) = \frac{W^{1-\gamma}}{1-\gamma}, \quad \gamma \neq 1 \tag{1}$$

and $U(W) = \log(W)$ when $\gamma = 1$. As in Knox (2002) we first simplify by ignoring intermediate consumption. The control variables for the investor are the portfolio weight invested in the risky asset at the start of the tree π_0, in the upper middle node π_H and in the lower middle node π_L.

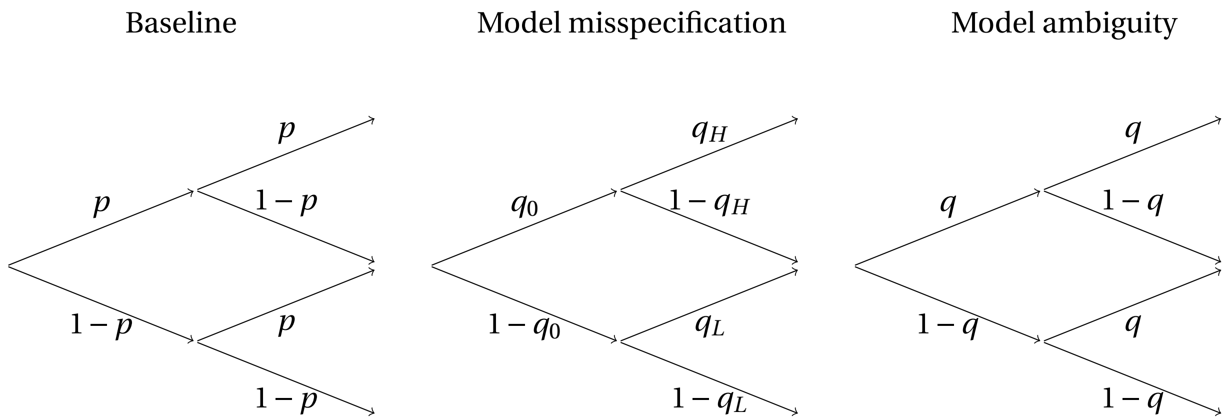

Figure 1: Binomial Tree

The analysis crucially hinges on the way in which the uncertainty sets are constructed, i.e. on the nature of the constraints imposed on the set of alternative measures.

1.1 A Single Entropy Constraint: Model Misspecification

Hansen and Sargent (2022) show that robustness against model misspecification can be obtained by allowing for the largest uncertainty set, in the sense of imposing the least structure, and by considering the worst-case alternative measure belonging to that uncertainty set.

For a given total amount K of relative entropy tolerated by the investor at $t = 0$, the largest uncertainty set is obtained by enforcing just a single intertemporal relative entropy constraint:

$$\begin{aligned} & q_0 \ln\left(\frac{q_0}{p}\right) + (1-q_0)\ln\left(\frac{1-q_0}{1-p}\right) \\ & + q_0 q_H \ln\left(\frac{q_H}{p}\right) + q_0(1-q_H)\ln\left(\frac{1-q_H}{1-p}\right) \\ & + (1-q_0) q_L \ln\left(\frac{q_L}{p}\right) + (1-q_0)(1-q_L)\ln\left(\frac{1-q_L}{1-p}\right) \leq K, \end{aligned} \tag{2}$$

where the first line on the left represents the expected relative entropy in the first period between $\mathbb{U}$ and $\mathbb{B}$, i.e., $\mathbb{E}^{\mathbb{B}}\left[\frac{d\mathbb{U}}{d\mathbb{B}} \log\left(\frac{d\mathbb{U}}{d\mathbb{B}}\right)\right]$, and the second and third line represent expected relative entropy in the second period.

In the language of Hansen and Sargent (2022), this uncertainty set includes unstructured alternative models, in contrast to the smaller uncertainty set that will be considered in the next subsection, which forms the set of structured alternative models. Intuitively, the wider uncertainty set with unstructured models leads the investor to explore deviations from the IID assumption, as seems natural when fearing *model misspecification*. On the other hand, when the investor fears *model ambiguity*, the set of structured models involves alternative models with distorted probabilities, but maintains the IID assumption.

Optimizing subject to this single entropy constraint corresponds to the "nonsequential constraint problem" (Hansen et al. (2006)). Nature is free to allocate entropy over time and across states, subject only to a single constraint. Epstein and Schneider (2003) are critical of this approach because it involves a nonrectangular set of priors or alternative measures. They advocate rectangularity as essential to rendering preferences recursive and hence dynamically consistent.[3]

The next subsection pursues a different formulation of the problem, with entropy constraints that impose rectangularity. However, first we solve the constraint problem with single entropy constraint (2) by introducing continuation entropy as a new state variable, as suggested by Hansen et al. (2006). This leads to a recursive formulation, time-consistent by construction. Intuitively, introducing the continuation entropy state variable avoids the potential for dynamic inconsistency described in footnote 3, because now the problem in state S is fully described by the state variables wealth W_1^S and continuation entropy K_S. Put differently, the continuation entropy state variable makes preferences conditional and de facto bakes in consequentialism.

Given our set-up and notation, the investor's objective function at time 0 is

$$\begin{aligned}\max_{\pi_0,\pi_H,\pi_L}\ \min_{q_0,q_H,q_L\in[0,1]} \frac{1}{1-\gamma}\Big(& q_0 q_H \left\{W_0\left(\pi_0(H-R_f)+R_f\right)\left(\pi_H(H-R_f)+R_f\right)\right\}^{1-\gamma} \\ & + q_0(1-q_H)\left\{W_0\left(\pi_0(H-R_f)+R_f\right)\left(\pi_H(L-R_f)+R_f\right)\right\}^{1-\gamma} \\ & + (1-q_0)q_L\left\{W_0\left(\pi_0(L-R_f)+R_f\right)\left(\pi_L(H-R_f)+R_f\right)\right\}^{1-\gamma} \\ & + (1-q_0)(1-q_L)\left\{W_0\left(\pi_0(L-R_f)+R_f\right)\left(\pi_L(L-R_f)+R_f\right)\right\}^{1-\gamma}\Big). \end{aligned} \tag{3}$$

For a recursive formulation we introduce the $t=1$ value functions at each node $S\in\{H,L\}$. The value V of the objective function depends on the state variables W_1^S and K_S. The minimizing agent has an entropy budget available at each state and in each period, and decides how much to use in the current period to distort the current probabilities and how much to carry over to each of the possible states next period. As a result of this decision at time $t=0$, the minimizing agent arrives in state S at time $t=1$ with continuation entropy K_S.

[3] The potential for dynamic inconsistency can easily be understood at the upper middle node where the minimizing agent may now deviate from the initially planned q_L so as to reallocate entropy from the counterfactual L state to state H in order to allow for a more extreme q_H than would otherwise be possible if q_L were fixed. Indeed, a prima facie feature of the single entropy constraint in (2) seems to be that "In contrast with recursive multiple-priors, *according to the robust control model behavior at any time-event pair may depend on what might have happened in unrealized parts of the tree.*" (Epstein and Schneider (2003), p. 19). This issue is closely related to the discussion in Knox (2003) and Chamberlain (2020) of consequentialism, the property that counterfactuals are ignored.

Consider state $S \in \{H, L\}$ at time $t = 1$. The value of the objective is

$$V_1(W_1^S, K_S) = \max_{\pi_S} \min_{q_S \in [0,1]} \frac{1}{1-\gamma}\Big(q_S \left\{W_1^S\left(\pi_S(H - R_f) + R_f\right)\right\}^{1-\gamma} + (1-q_S)\left\{W_1^S\left(\pi_S(L - R_f) + R_f\right)\right\}^{1-\gamma}\Big), \tag{4}$$

subject to

$$q_S \ln\left(\frac{q_S}{p}\right) + (1-q_S)\ln\left(\frac{1-q_S}{1-p}\right) \leq K_S. \tag{5}$$

We now construct the $t = 0$ objective recursively, based on the $t = 1$ conditional value functions.

$$V_0(W_0, K) = \max_{\pi_0} \min_{q_0, K_H, K_L} q_0 V_1(W_1^H, K_H) + (1-q_0)V_1(W_1^L, K_L) \tag{6}$$

$$\text{s.t. } q_0 \ln\left(\frac{q_0}{p}\right) + (1-q_0)\ln\left(\frac{1-q_0}{1-p}\right) + q_0 K_H + (1-q_0)K_L \leq K \tag{7}$$

Constraint problems (4)-(5) and (6)-(7) are converted into unconstrained ones by introducing Lagrange multipliers λ_0 on constraint (7), λ_H and λ_L on constraints (5) for $S = H$ and for $S = L$, respectively.

Lemma 1. *The constraint problems* (4)*-*(5) *and* (6)*-*(7) *are solved by the system of equations in Appendix A.1. The solution has the following properties:*

1. The Lagrange multipliers are equal $\lambda_0 = \lambda_H = \lambda_L$ *and homogeneous of degree* $1-\gamma$ *in initial wealth* W_0*.*

2. The optimal portfolio weights π_0*,* π_H *and* π_L *are independent of wealth, but state-dependent and not equal to each other for* $\gamma \neq 1$*.*

3. The optimal distorted probabilities q_0*,* q_H *and* q_L *are independent of wealth, but state-dependent and not equal to each other for* $\gamma \neq 1$*.*

The Lagrange multipliers are constant over time and across states as they put shadow prices on the same entropy constraint. Optimality requires that entropy be allocated to different states of the world such that shadow prices are equalized. Simultaneously, the Lagrange multipliers scale endogenously with initial wealth W_0, which safeguards homotheticity. Both findings are novel and crucial for understanding multiplier preferences and the debate about the penalty parameter in the robust control literature, as discussed in Section 3. A related result is that the portfolio weights are independent of wealth, due to homotheticity, but generally state-dependent. The reason for this is the state-dependent probability distortion, thereby deviating from the

IID assumption of the baseline measure. A subtle, but important finding therefore is that the state-dependence of the portfolio allocation reflects differences in probability distortions across states, not differences in wealth. Put differently, state-dependence of the portfolio weights is not to be confused with lack of homotheticity: the first-order conditions reveal that the problem is fully homothetic, by virtue of the endogenous scaling of the Lagrange multipliers with initial wealth.

In summary, the recursive solution to the single constraint entropy problem involves constant Lagrange multipliers that scale endogenously with initial wealth, and homothetic but state-dependent optimal portfolio weights. We now characterize the solution further in order to gain additional insight into the economic effect of model misspecification concerns.

Proposition 1. *Assume that $Hp + L(1-p) > R_f$. If $\gamma > 1$, the investor's pessimistic worst-case belief and the optimal portfolio weight in period 2 are both procyclical, i.e., $q_L < q_H < p$ and $\pi_L < \pi_H$. If $0 < \gamma < 1$, the investor's pessimistic worst-case belief and the optimal portfolio weight in period 2 are both countercyclical, i.e., $q_H < q_L < p$ and $\pi_H < \pi_L$. The log investor ($\gamma = 1$) has state-independent beliefs ($q_H = q_L = q_0 < p$) and invests myopically ($\pi_H = \pi_L = \pi_0$).*

Assuming a strictly positive equity premium under the baseline model, Proposition 1 shows that the worst-case probability distortion is pessimistic, and depends on the investor's risk preferences, since these drive optimal investment behavior. While returns under the baseline measure are IID, the worst-case measure induces serial correlation. The serial correlation is positive for a risk-averse investor ($\gamma > 1$) and negative for a risk-tolerant investor ($\gamma < 1$). For $\gamma > 1$, expected returns are less distorted downwards after a high return than following a low return, i.e., expected returns are higher in the good state of nature than in the bad state, with the opposite happening for the low risk-aversion investor ($\gamma < 1$). Put differently, the risk-averse investor fears persistence in returns, the risk-tolerant investor fears mean reversion.

The economic mechanism behind these results can be understood from the insights of Merton (1971) about intertemporal hedging. A risk-averse non-myopic investor holds more of the risky asset when returns and expected returns are negative correlated, because the expected improvement in investment opportunities following low returns makes the investment less risky over a longer horizon and provides an intertemporal hedge. Kim and Omberg (1996) offers a detailed analysis (see also Campbell and Viceira (2002) for additional insights and intuition). For CRRA they show that the $\gamma > 1$ investor hedges intertemporally, holding more of the risky asset than the mean-variance portfolio when returns and expected returns are negatively correlated (and less otherwise). In contrast, the $\gamma < 1$ investor speculates intertemporally and holds more of the risky asset than the mean-variance efficient allocation when returns and expected returns are positively correlated (and less otherwise), because this investor values the fact that an asset pays off highly precisely when investment opportunities improve.

Proposition 1 shows that the worst-case model misspecification for a $\gamma > 1$ investor, who has a preference for intertemporal hedging, is that intertemporal hedging is not possible because returns and expected returns are positively correlated. Similarly, the worst-case misspecification for a $0 < \gamma < 1$ investor, who likes to speculate intertemporally, is that returns and expected returns are negatively correlated, since this precludes intertemporal speculation.

The log investor is of course in between, and is indifferent between intertemporal hedging and intertemporal speculation. The serial correlation of returns is of no importance to this investor's optimal strategy, since he behaves myopically and invests based only on current investment opportunities, without any attempt at exploiting the dynamics of investment opportunities to either hedge intertemporally or to get a higher expected return via intertemporal speculation. Model misspecification is therefore irrelevant to the log investor, and we find that for $\gamma = 1$ the solution to the single-entropy constraint problem is the same as when there is no model misspecification concern, and only model ambiguity is at work. His portfolio weight is reduced, as with model ambiguity, but is state-independent.

This result is critical because virtually all papers applying robust control in finance and macroeconomics using multiplier preferences rely on log utility to get tractability. The first result in Lemma 1 gives one reason why log utility is highly tractable in the penalty approach: the endogenous Lagrange multiplier does not scale with wealth in this case, which explains why log multiplier preferences remain homothetic even when an exogenously fixed penalty parameter is imposed in lieu of an endogenous Lagrange multiplier. The second reason for their tractability is because log preferences lead to myopic behavior. Our findings clearly show that this feature of log utility is not innocuous at all and in fact eradicates the concern for model misspecification, reducing it to mere model ambiguity.

Figure 2 illustrates the findings of Proposition 1 for the relationship between the baseline model probability p and the distorted probabilities q_L and q_H along with the numerical solution for q_0, for values of the coefficient of relative risk aversion γ between 0.5 and 4. We see that for all values of γ, q_0 is between q_L and q_H. To appreciate the quantitative impact of these distorted probabilities, it is helpful to consider the corresponding optimal portfolio weights. For the return parameters assumed in Figure 2, a $\gamma = 2$ investor optimally allocates roughly 95% of wealth to the risky asset without robustness. When concerned about model misspecification the optimal weight for this investor shrinks to approximately 70% when $K = 0.01$ (top panel) and to roughly 55% when $K = 0.04$ (middle panel).

Figure 3 shows how the optimal solution depends on the total amount of entropy K, for a range of entropy from 0 to the maximum amount that is economically meaningful, namely such that the Sharpe ratio on the risky asset becomes zero in all states, so that the investor only holds the riskless asset. At the extremes of both zero and maximal entropy, the state-dependence

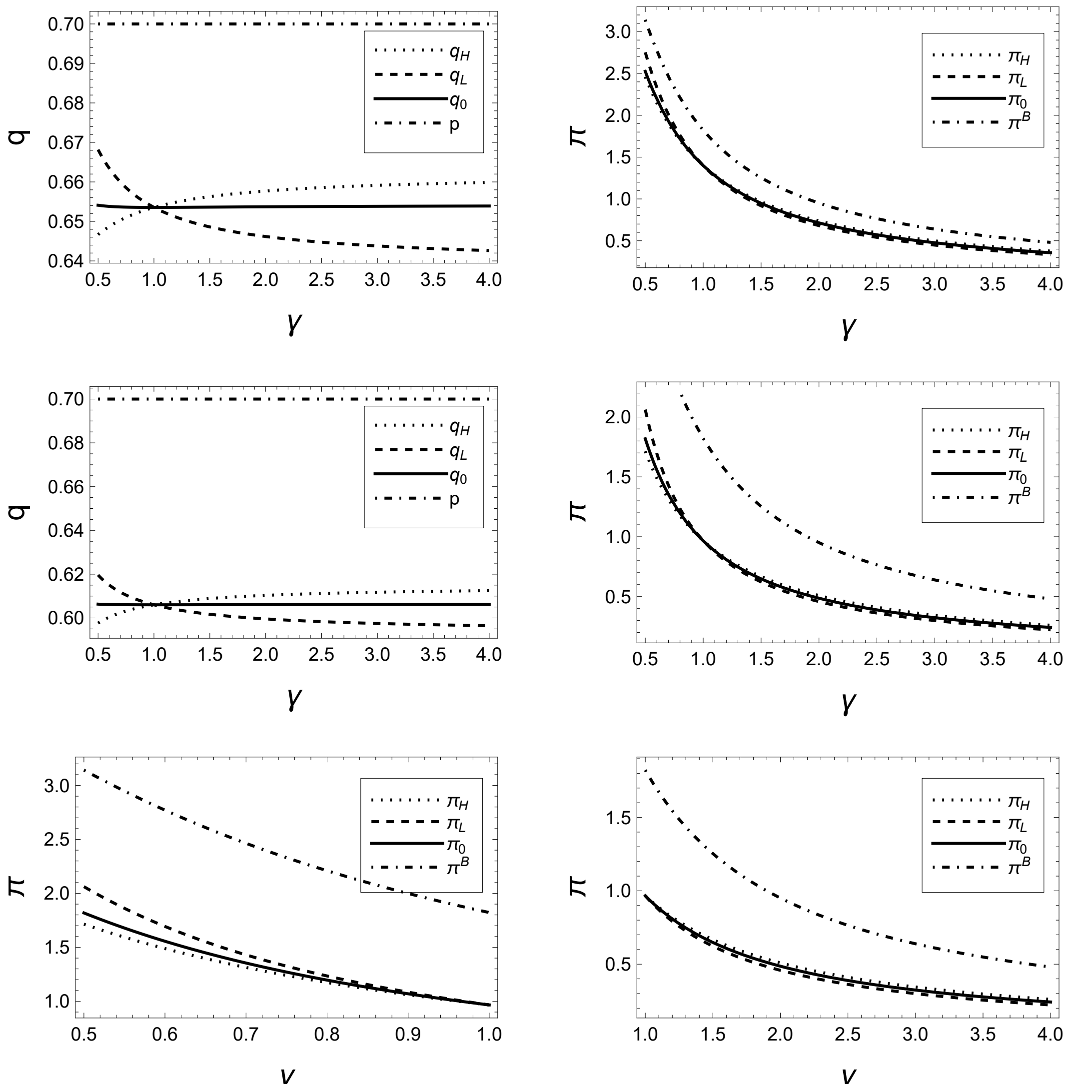

Figure 2:

Notes: This figure plots the optimally distorted probabilities q_0, q_L*, and* q_H*, and the optimal portfolio weights* π_0, π_L *and* π_H *corresponding to the solution of the problem in Lemma 1.* $\pi^{\mathbb{B}}$ *is the optimal portfolio weight corresponding to the baseline measure* $\mathbb{B}$*. The parameters are* $p = 0.7$*,* $H = 1.25$*,* $L = 0.8$*,* $R = 1.025$*. The problem is solved for coefficient of relative risk aversion* $\gamma \in [0.5, 4]$*. In the top panel* $K = 0.01$ *and in the middle panel* $K = 0.04$*. The bottom two panels split the middle right panel to two plots with* $\gamma \in (0.5, 1)$ *and* $\gamma \in (1, 4)$*, respectively.*

in distorted probabilities vanishes. This is intuitive: at zero entropy, the investor does not consider model misspecification or model ambiguity, and trusts the IID benchmark model. At the maximum amount of entropy, the Sharpe ratio is always zero, in which case there is no state-dependence either. We will find the same patterns in the continuous-time version of the model. We defer a discussion of how to interpret and calibrate the amount of entropy K to the continuous-time analysis. Note that the effect of state-dependence is very small here due to the very short investment horizon (only 2 periods). The continuous-time analysis will also address this by considering an infinite horizon.

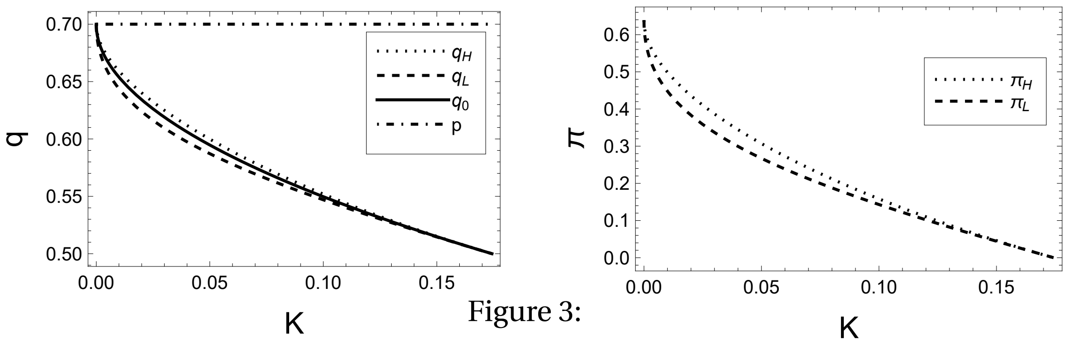

Figure 3:

Notes: This figure plots the optimally distorted probabilities q_0, q_L, and q_H, and the optimal portfolio weights π_L and π_H corresponding to the solution of the problem in Lemma 1. The parameters are $p = 0.7$, $H = 1.25$, $L = 0.8$, $R = 1.025$, and $\gamma = 3$. The problem is solved for a range of K, where the right end of both plots correspond to zero Sharpe ratio scenario.

Finally, we note that the endogenous worst-case beliefs of the risk-averse investor ($\gamma > 1$) can be viewed as extrapolative: despite returns being IID, the investor fears persistence when there is none under $\mathbb{B}$, and can be seen as extrapolating from recent experience when forming expectations for the future. Indeed, after experiencing a high return, expected returns increase and after the low return, expected returns shrink. Our findings are therefore similar to what Bidder and Dew-Becker (2016) and Bidder and Smith (2018) report in a macrofinance model, although they focus on expectations about consumption dynamics in an equilibrium setting rather than return expectations in a portfolio problem.

1.2 Two Pathwise Entropy Constraints: Model Ambiguity

Having discussed a single constraint in (2) and the allocation of continuation entropy in (6), we now present an alternative formulation where in the second period Nature can only spend whatever is left from the entropy budget after the first period. This leads to the following two pathwise entropy constraints:

$$q_H \ln\left(\frac{q_H}{p}\right) + (1-q_H)\ln\left(\frac{1-q_H}{1-p}\right) \leq K - q_0 \ln\left(\frac{q_0}{p}\right) - (1-q_0)\ln\left(\frac{1-q_0}{1-p}\right), \tag{8}$$

$$q_L \ln\left(\frac{q_L}{p}\right) + (1-q_L)\ln\left(\frac{1-q_L}{1-p}\right) \leq K - q_0 \ln\left(\frac{q_0}{p}\right) - (1-q_0)\ln\left(\frac{1-q_0}{1-p}\right). \tag{9}$$

Having two constraints rather than one obviously imposes more structure [4] and shrinks the uncertainty set against which the agent seeks robustness considerably. First, the two pathwise constraints (8) and (9) jointly imply (2). The resulting set of structured models considered by the investor is a strict subset of the set of unstructured models considered in the previous subsection. The pathwise constraints essentially impose that continuation entropy, which is the right-hand-side of (8) and (9), be state-independent. In other words, using the notation of the recursive formulation in the previous subsection, the constraints imply $K_H = K_L = K_1$. The decision-maker only needs to decide how to allocate entropy between today and tomorrow, but what is left over is now constant across states. This imposes strict rectangularity in the sense of Epstein and Schneider (2003). The pathwise entropy constraints can also be seen as putting a bound on the increment or change in entropy that is due to the probability distortion in the second period, i.e. the left-hand side of the constraints. Importantly, the bound in the pathwise constraint is state-independent, which is crucial for understanding the difference between model ambiguity and model misspecification in continuous time.

To solve the nonsequential formulation of the robust constraint problem, we introduce Lagrange multipliers λ_H and λ_L on constraints (8) and (9), respectively. The following result summarizes the solution to problem (3) subject to pathwise constraints (8) and (9).

Lemma 2. *The solution of* (3) *subject to* (8) *and* (9) *is* $q_0 = q_H = q_L$ *and* $\pi_0 = \pi_H = \pi_L$*. The Lagrange multipliers* λ_H *and* λ_L *are homogeneous of degree* $1-\gamma$ *in initial wealth* W_0*, and* $K_1 = K/2$.

Lemma 2 shows that continuation entropy is not only state-independent but also constant over time. Therefore, the pathwise entropy constraint formulation is equivalent to the κ-ignorance specification in Chen and Epstein (2002). Pathwise entropy constraints greatly simplify the solution and induce the investor as well as nature to act myopically, as evident from the state-independent and constant portfolio weights and worst-case probabilities. To preserve the homotheticity of the problem, the endogenous Lagrange multipliers scale with initial wealth raised to the power $1-\gamma$, which justifies the exact scaling proposed in Maenhout (2004). These findings therefore nicely parallel the result in Knox (2002) that the homothetic version in Maenhout (2004) of the penalty approach is structurally identical to κ-ignorance.

It is striking how much additional structure is imposed by merely using the pathwise constraints instead of the single constraint. As we show below, in a continuous-time model with Brownian information, the pathwise constraints correspond to a constraint on the instantaneous expected change in entropy, which must be bounded and state-independent. We will

[4] In fact it imposes consequentialism (cf. footnote 3) by virtue of being pathwise, thus ignoring what happens in the counterfactual state. As noted before, introducing continuation entropy as a state variable to obtain recursivity in the previous subsection also de facto bakes in consequentialism.

show that the correspondence between model ambiguity and pathwise entropy constraints, i.e., rectangularity, also continues to hold in the continuous-time setting.

2 Entropy Constraints in a Continuous-Time Model

2.1 Continuation Entropy in Continuous Time

Let $(\Omega, (\mathscr{F}_t)_{t\in[0,\infty)}, \mathbb{B})$ be a standard probability space. The probability measure $\mathbb{B}$ is the agent's baseline model and $B^{\mathbb{B}}$ is a 1-dimensional Brownian motion under $\mathbb{B}$. The agent considers a family of alternative models $\mathbb{U}$, which is parameterized by a real-valued bounded belief distortion process u. For each given bounded u and $T > 0$, $\mathbb{U}$ is defined as

$$\frac{d\mathbb{U}}{d\mathbb{B}}\Big|_{\mathscr{F}_T} = Z_T,^{5} \quad \text{where} \quad Z_t = \exp\Big(-\int_0^t \tfrac{1}{2}u_s^2 ds - \int_0^t u_s dB_s^{\mathbb{B}}\Big), \quad t \in [0, T]. \tag{10}$$

For every random outcome $\omega \in \Omega$, the density $Z_T(\omega)$ describes the change of likelihood for this outcome under $\mathbb{U}$ compared to $\mathbb{B}$. Larger $Z(\omega)$ implies that ω is more likely to happen under the alternative model $\mathbb{U}$.

Let us first introduce an entropy discrepancy measure between the alternative model $\mathbb{U}$ and the baseline model $\mathbb{B}$. Following Hansen and Sargent (2022), introduce the finite horizon entropy discrepancy measure

$$\mathbb{E}^{\mathbb{B}}\big[Z_T \log Z_T\big] = \tfrac{1}{2}\mathbb{E}^{\mathbb{B}}\Big[\int_0^T Z_s u_s^2 ds\Big] = \tfrac{1}{2}\mathbb{E}^{\mathbb{U}}\Big[\int_0^T u_s^2 ds\Big].$$

Taking the long run average and using Abel integral averages, Hansen and Sargent (2022) show

$$\lim_{T\to\infty} \tfrac{1}{T}\mathbb{E}^{\mathbb{B}}\big[Z_T \log Z_T\big] = \lim_{T\to\infty} \tfrac{1}{2T}\mathbb{E}^{\mathbb{U}}\Big[\int_0^T u_s^2 ds\Big] = \lim_{\delta\to 0}\mathbb{E}^{\mathbb{U}}\Big[\int_0^\infty \delta e^{-\delta s}\tfrac{1}{2}u_s^2 ds\Big],$$

provided that these limits exists. In the previous equation, the scaling by δ makes the weights $\delta e^{-\delta s}$ integrate to one and produces an exponential average. Choosing δ as the subjective discount rate, we define a discrepancy measure between $\mathbb{U}$ and $\mathbb{B}$ as

$$\mathbb{E}^{\mathbb{U}}\Big[\int_0^\infty \delta e^{-\delta s}\tfrac{1}{2}u_s^2 ds\Big].$$

The entropy allocation problem in (4)-(5) and (6)-(7) motivate us to introduce *continuation*

[5] Because u is bounded, $\mathbb{E}^{\mathbb{B}}[Z_T] = 1$, for any $T > 0$, therefore $\mathbb{U}$ is a probability measure equivalent to $\mathbb{B}$ on $\mathscr{F}_T$ for any $T > 0$.

entropy

$$E_t = \mathbb{E}_t^{\mathbb{U}}\Big[\int_t^\infty \delta e^{-\delta(s-t)}\tfrac{1}{2}u_s^2 ds\Big], \quad \text{for } t \geq 0,\text{[6]} \tag{11}$$

where $\mathbb{E}_t^{\mathbb{U}}[\cdot]$ stands for the conditional expectation $\mathbb{E}^{\mathbb{U}}[\cdot|\mathscr{F}_t]$. Intuitively, nature starts continuation entropy E from an initial entropy budget E_0, and allocates instantaneous distortions u_t and continuation entropy E_t to future dates and states. It follows from the martingale representation theorem that there exists a unique process g such that E follows the dynamics

$$dE_t = \delta E_t dt - \tfrac{\delta}{2}u_t^2 dt + g_t dB_t^{\mathbb{U}}. \tag{12}$$

Therefore, controlling continuation entropy E is equivalent to controlling the instantaneous distortion u and the volatility g of continuation entropy.

The martingale term in (12) is crucial as it allows Nature to control the volatility of continuation entropy and as such enables a state-dependent allocation of entropy. This was instrumental in distinguishing between model ambiguity and model misspecification in the two-state example, and the same holds true in continuous time. Put differently, without the martingale term in (12), nature can only control the drift or the expected change of the continuation entropy, exactly as with the pathwise entropy constraints defining the set of structured models in Section 1.2.

2.2 An Infinite Horizon Optimal Consumption-Investment Problem

Consider a financial market with a risk-free asset and a risky asset. The risk-free rate is assumed to be $r > 0$ and the price S of the risky asset follows

$$\frac{dS_t}{S_t} = \mu dt + \sigma dB_t^{\mathbb{B}},$$

where the expected return μ and the volatility σ are constants. Consider an agent whose preference over consumption is modelled by a CRRA utility $U(c) = \frac{c^{1-\gamma}}{1-\gamma}$ with $\gamma > 0$ ($U(c) = \log(c)$ when $\gamma = 1$). This agent faces an infinite-horizon consumption-investment problem

$$\sup_{\pi,c}\inf_{u,g}\mathbb{E}^{\mathbb{U}}\Big[\int_0^\infty \delta e^{-\delta s}U(c_s)ds\Big], \tag{13}$$

subject to the budget constraint

$$dW_t = rW_t dt + \pi_t W_t\big((\mu - r)dt + \sigma dB_t^{\mathbb{B}}\big) - c_t dt, \tag{14}$$

[6]Because u is assumed to be bounded, continuation entropy E is bounded as well. Therefore it satisfies the transversatility condition $\lim_{t\to\infty} e^{-\delta t}\mathbb{E}^{\mathbb{U}}[E_t] = 0$.

and the dynamics of continuation entropy E in (12).[7]

We take the agent's wealth W and Nature's continuation entropy E as natural state variables for the problem (13) and define the value function V via

$$V(W_t, E_t) = \sup_{\pi,c} \inf_{u,g} \mathbb{E}_t^{\mathbb{U}}\Big[\int_t^{\infty} \delta e^{-\delta(s-t)} U(c_s)\, ds\Big]. \tag{15}$$

It follows from dynamic programming that V satisfies the following HJB equation:

$$\begin{aligned}\delta V = \sup_{\pi,\tilde{c}} \inf_{u,g} \Big\{ & \delta U(W\tilde{c}) + W\partial_W V\big[r + \pi(\mu - r - \sigma u) - \tilde{c}\big] + \partial_E V\big(\delta E - \tfrac{\delta}{2}|u|^2\big) \\ & + \tfrac{1}{2}W^2\pi^2\sigma^2\partial^2_{WW}V + \tfrac{1}{2}g^2\partial^2_{EE}V + W\pi\sigma g\partial^2_{EW}V\Big\},\end{aligned} \tag{16}$$

where $\tilde{c} = c/W$ is the consumption-wealth ratio. Because of the homothetic property of U, the value function V admits the following homothetic decomposition

$$V(E, W) = \frac{W^{1-\gamma}}{1-\gamma}\mathcal{V}(E), \tag{17}$$

for a positive function $\mathcal{V}$ depending on E only. The following results summarize the agent's optimal strategies and the worst-case belief.

Proposition 2. *When $\gamma > 1$, $\mathcal{V}$ in (17) satisfies*

$$\begin{aligned}\delta\mathcal{V} = \inf_{\pi,\tilde{c}} \sup_{u,g} \Big\{ & \delta\tilde{c}^{1-\gamma} + (1-\gamma)[r + \pi(\mu - r - \sigma u) - \tilde{c}]\mathcal{V} + \big(\delta E - \tfrac{\delta}{2}|u|^2\big)\partial_E\mathcal{V} \\ & - \tfrac{1}{2}\gamma(1-\gamma)\pi^2\sigma^2\mathcal{V} + \tfrac{1}{2}g^2\partial^2_{EE}\mathcal{V} + (1-\gamma)\pi\sigma g\partial_E\mathcal{V}\Big\},\end{aligned} \tag{18}$$

with the boundary conditions at $E = 0$ and $E = \overline{E} = \frac{(\mu-r)^2}{2\sigma^2}$:

$$\mathcal{V}(0) = \delta\Big\{\tfrac{1}{\gamma}\Big[\delta - (1-\gamma)r - (1-\gamma)\tfrac{(\mu-r)^2}{2\gamma\sigma^2}\Big]\Big\}^{-\gamma}, \tag{19}$$

$$\mathcal{V}(\overline{E}) = \delta\Big\{\tfrac{1}{\gamma}\Big[\delta - (1-\gamma)r\Big]\Big\}^{-\gamma}. \tag{20}$$

Suppose that $\partial_E\mathcal{V} > 0$ and $\partial^2_{EE}\mathcal{V} < 0$, the agent's optimal investment fraction π^* and the optimal

[7] Given a consumption process c, the utility process $U_t = \inf_{u,g}\mathbb{E}_t^{\mathbb{U}}\big[\int_t^{\infty}\delta e^{-\delta(s-t)}U(c_s)ds\big]$ can be formulated as an optimization problem for a forward-backward stochastic differential equation: $U_t = \inf_{u,g} U_t^{u,g}$ where the backward component $U^{u,g}$ follows $dU_t^{u,g} = \big(\delta U_t^{u,g} - \delta U(c_t)\big)dt + \phi_t^{u,g} dB_t^{\mathbb{U}}$ and the forward component E follows (12). When E is not a state variable, robust control problems with entropy penalties are typically formulated as optimization problems for backward stochastic differential equations, see eg. Skiadas (2003).

consumption-wealth ratio $\tilde{c}^*$ *satisfy*

$$\pi^* = \frac{\mu - r}{\gamma\sigma^2 - (1-\gamma)\sigma^2 \frac{\mathcal{V}}{\delta \partial_E \mathcal{V}} + (1-\gamma)\sigma^2 \frac{(\partial_E \mathcal{V})^2}{\mathcal{V} \partial^2_{EE} \mathcal{V}}}, \tag{21}$$

$$\tilde{c}^* = \delta^{\frac{1}{\gamma}} \mathcal{V}^{-\frac{1}{\gamma}}. \tag{22}$$

The optimal instantaneous distortion u^* *and the volatility of continuation entropy* g^* *are*

$$u^* = -\frac{(1-\gamma)\sigma}{\delta} \frac{\mathcal{V}}{\partial_E \mathcal{V}} \pi^*, \tag{23}$$

$$g^* = -(1-\gamma)\sigma \frac{\partial_E \mathcal{V}}{\partial^2_{EE} \mathcal{V}} \pi^*. \tag{24}$$

When $0 < \gamma < 1$, $\inf_{\pi,\tilde{c}} \sup_{u,g}$ *is replaced by* $\sup_{\pi,\tilde{c}} \inf_{u,g}$ *in* (18) *and the optimality conditions become* $\partial_E \mathcal{V} < 0$ *and* $\partial^2_{EE} \mathcal{V} > 0$.

To understand the optimal strategies, let us start from the instantaneous distortion u. The optimization of u in (16) is

$$\inf_u \Big\{ - W \partial_W V \pi \sigma u - \tfrac{\delta}{2} |u|^2 \partial_E V \Big\}. \tag{25}$$

This is the familiar trade-off between the marginal entropy cost of a higher distortion and the marginal benefit to the minimizing agent due to the marginal utility loss of the investor. Using the decomposition (17), we obtain the form of u^* in (23).

The optimization of g in (16) is

$$\inf_g \Big\{ \tfrac{1}{2} g^2 \partial^2_{EE} V + W \pi \sigma g \partial^2_{EW} V \Big\}, \tag{26}$$

where the objective function measures the impact of g, which controls the variance or state-dependence of the entropy state variable, on the variance of the value function. Intuitively, the optimal state-dependence of entropy depends on second-order derivatives of the value function. The optimality condition reflects the trade-off between the change in marginal entropy cost $\partial^2_{EE} V$ and the change in marginal utility due to a change in entropy $\partial^2_{EW} V$, reflecting the covariance between the state variables. Using the decomposition (17), we obtain the form of g^* in (24). It follows from (24) and the conditions on $\partial_E \mathcal{V}$ and $\partial^2_{EE} \mathcal{V}$ that $g^* < 0$ when $\gamma > 1$ and $g^* > 0$ when $0 < \gamma < 1$. This extends the results in Proposition 1 to a continuous-time setting, because $g^* < 0$ for $\gamma > 1$ means continuation entropy is countercyclical, i.e., increases after negative return shocks. Therefore, the investor is more pessimistic and reduces the expected return following negative return shocks. As in the binomial tree, we find that the risk-averse investor endogenously fears return persistence, as this is indeed the worst-case model misspecification for a risk-averse investor with $\gamma > 1$, because Mertonian intertemporal hedging becomes impossible in this case. This investor can be interpreted as (endogenously) extrapolating from recent return realizations, in line with empirical evidence obtained from survey data on expectations (see for instance Hong and Stein (1999), Barberis et al. (2015), Barberis (2018), Jin and Sui (2022), and Li and Liu (2019)).

Similarly, $g^* > 0$ for $0 < \gamma < 1$ implies procyclical continuation entropy, i.e., continuation entropy decreases after negative return shocks. Therefore, this investor becomes less pessimistic and raises the expected return following negative shocks. Exactly in line with the results in the discrete-time model, the risk-tolerant investor ($0 < \gamma < 1$) fears mean-reversion in returns, because it precludes intertemporal speculation. The endogenous beliefs of this investor can be viewed as contrarian.

The optimal investment strategy in (21) can be reformulated as

$$\pi^* = \frac{\mu - r}{\sigma^2 \gamma^{\text{eff}}(E)},$$

where γ^{eff} is the effective risk aversion

$$\gamma^{\text{eff}}(E) = \gamma - (1-\gamma)\frac{\mathcal{V}}{\delta \partial_E \mathcal{V}} + (1-\gamma)\frac{(\partial_E \mathcal{V})^2}{\mathcal{V}\partial^2_{EE}\mathcal{V}}. \tag{27}$$

On the right-hand side of (27), the first term is the usual relative risk aversion coefficient. We obtain two additional terms, due to concerns for model ambiguity and model misspecification, respectively. The second term in (27) is driven by the optimal distortion u^* and is therefore attributable to model ambiguity. The third term is proportional to g^*/u^* hence reflecting the relative concern of model misspecification versus model ambiguity. Due to the sign conditions for optimality on $\partial_E \mathcal{V}$ and $\partial^2_{EE}\mathcal{V}$, which are verified numerically in Section 4, both additional terms on the right-hand side of (27) are strictly positive, indicating that model ambiguity aversion as well as model misspecification fears unambiguously increase the investor's effective risk aversion. Therefore the optimal investment strategy π^* in (21) is always more conservative than Merton's mean-variance ratio $\pi^M = \frac{\mu - r}{\gamma\sigma^2}$, which is the optimal investment strategy for an agent without model misspecification or model ambiguity concerns. Moreover, effective risk aversion is state-dependent. It generates dynamic investment behavior depending on the value of the continuation entropy state variable, as we discuss in more detail in our numerical analysis in Section 4.

For the boundary conditions of $\mathcal{V}$, when $E = 0$ there is no continuation entropy left, then the future instantaneous belief distortion u must be zero. In this case, problem (13) is reduced to a standard Merton problem without model ambiguity, whose value is given in (19). When

$E = \overline{E}$, the instantaneous belief distortion reaches its maximum value $\overline{u} = \frac{\mu - r}{\sigma}$, in which case the expected excess return of the risky asset becomes zero under the subjective belief $\mathbb{U}$. In this case, the agent does not invest in the risky asset anymore and non-participation becomes optimal. Then $\overline{E}$ is the maximum continuation entropy value corresponding to this no-investment worst-case scenario and (13) is reduced to an optimal consumption but no investment problem, with value given in (20).

3 Model Misspecification versus Model Ambiguity

To gain a deeper understanding of the precise difference between model misspecification and model ambiguity in our setting, it is useful to first also examine the broader question of the relationship between the penalty or multiplier preferences of Hansen and Sargent (2001) and the continuation entropy approach of this paper.

3.1 Connection to Hansen and Sargent Multiplier Preferences

We follow closely the logic of Hansen et al. (2006) based on duality theory, but obtain additional results that give crisp insights into the difference between model ambiguity and model misspecification in dynamic portfolio and asset pricing problems.

For given fixed $\theta > 0$, define

$$\widehat{V}(W,\theta) = \inf_{E\geq 0} \Big\{ V(W,E) + E\theta \Big\}, \tag{28}$$

where $V(W,E)$ is the value function for the problem in (15). We call $\widehat{V}$ the dual value function.

Plugging the optimal g^* into (16) and using the duality relation (28), we derive the following equation for $\widehat{V}$ (see Appendix A.5):

$$\delta\widehat{V} = \sup_{\pi,c}\inf_{u} \Big\{ \delta U(W\tilde{c}) + W\partial_W \widehat{V}\big[r + \pi(\mu - r - \sigma u) - \tilde{c}\big] + \tfrac{1}{2}W^2\pi^2\sigma^2\partial^2_{WW}\widehat{V} + \delta\tfrac{\theta}{2}|u|^2 \Big\}. \tag{29}$$

Standard dynamic programming arguments can be used to show that this equation satisfied by the dual value function $\widehat{V}$ is the HJB for the following optimal consumption-investment problem with multiplier preferences, as in for example Anderson et al. (2003):

$$\widehat{V}(W,\theta) = \sup_{\pi,c}\inf_{u} \mathbb{E}^{\mathbb{U}}\Big[\int_0^\infty e^{-\delta t}\delta\big(U(c_t) + \tfrac{\theta}{2}u_t^2\big)dt \Big], \tag{30}$$

subject to

$$dW_t = W_t\big[r + \pi_t(\mu - r - \sigma u_t)\big]dt + W_t\pi_t\sigma dB_t^{\mathbb{U}} - c_t dt.$$

In this problem, θ is fixed and does not depend on W. Hence $\widehat{V}$ is not homothetic in W, except for log utility ($\gamma = 1$), which we discuss in Section 3.4. As emphasized in Hansen et al. (2006), multiplier preferences have the appeal of parsimony owing to the fixed penalty parameter θ, which avoids carrying along the entropy state variable necessitated by our constraint approach. At the same time, because the penalty parameter θ is fixed, the problem loses its homotheticity, and as a result wealth now needs to be tracked explicitly as a state variable affecting optimal consumption and portfolio decisions.

3.2 Model Misspecification

Coming back to the value function V, by virtue of the duality relationship between V and $\widehat{V}$,

$$V(W,E) = \sup_{\theta\geq 0}\left\{\widehat{V}(W,\theta) - E\theta\right\}. \tag{31}$$

Therefore, $\partial_E V(W,E) = -\theta^*$. The decomposition of V in (17) then yields

$$\theta^* = -\frac{W^{1-\gamma}}{1-\gamma}\partial_E \mathcal{V}(E). \tag{32}$$

This implies that θ^* is homogeneous of degree $1-\gamma$ in W and also depends on the state E. Evidently, compared to the fixed penalty parameter θ in the Hansen-Sargent multiplier preferences of the previous subsection, the constraint approach seems substantially less parsimonious. We now show that optimality combined with preservation of homotheticity nonetheless results in parsimonious and tractable solutions for our approach.

To this end, we now evaluate θ^* along the path of the state process (W_t, E_t). Define

$$\theta_t^* = \theta^*(W_t, E_t) = -\partial_E V(W_t, E_t), \tag{33}$$

where W_t and E_t follow the optimal strategies π^*, u^*, and g^* of Proposition 2. It follows from Itô's formula that the volatility of θ_t^* is

$$-g^*\partial^2_{EE}V - W\pi^*\sigma\partial^2_{EW}V, \tag{34}$$

which is zero from the first-order condition in (26). Similarly, the drift of θ_t^* can be obtained from Itô's lemma as

$$\begin{aligned}&-\left(\delta E - \tfrac{\delta}{2}|u^*|^2\right)\partial^2_{EE}V - W\left[r + \pi^*(\mu - r - \sigma u^*) - \tilde{c}^*\right]\partial^2_{EW}V\\&-\tfrac{1}{2}W^2(\pi^*)^2\sigma^2\partial^3_{EWW}V - \tfrac{1}{2}(g^*)^2\partial^3_{EEE}V - W\pi^*\sigma g^*\partial^3_{EEW}V,\end{aligned} \tag{35}$$

which can be shown to equal zero. To see this, it suffices to differentiate the HJB equation (16) with respect to E, to use the envelope theorem, and to evaluate at the optimal choices to obtain that

$$0 = \left(\delta E - \tfrac{\delta}{2}|u^*|^2\right)\partial^2_{EE}V + W\left[r + \pi^*(\mu - r - \sigma u^*) - \tilde{c}^*\right]\partial^2_{EW}V \\ + \tfrac{1}{2}W^2(\pi^*)^2\sigma^2\partial^3_{EWW}V + \tfrac{1}{2}(g^*)^2\partial^3_{EEE}V + W\pi^*\sigma g^*\partial^2_{EEW}V. \tag{36}$$

The following result summarizes the discussion above.

Proposition 3. *$\theta^*_t = \theta^*(W_t, E_t)$ is constant when the state processes W_t and E_t follow the optimal strategies π^*, u^*, and g^* in Proposition 2.*

As we showed in the binomial tree case, allowing Nature to optimally control the volatility or state-dependence of continuation entropy is equivalent to equalizing the shadow price of continuation entropy across states and over time. Proposition 3 extends to the continuous-time setting the finding of Lemma 1 that the endogenous Lagrange multiplier on the entropy constraint is constant *conditional* on the state variables W_t and E_t. As is clear from (32), the Lagrange multiplier scales with wealth, exactly as in Lemma 1, and this preserves the homotheticity of the problem.

The fact that the drift of the Lagrange multiplier associated with the entropy constraint is zero can be seen as an application of the Inverse Euler Equation of dynamic principal-agent models due to Rogerson (1985) and Spear and Srivastava (1987). The line of reasoning used to prove this result above coincides with the proof of the version of the Inverse Euler Equation obtained in DeMarzo and Sannikov (2017). Hansen et al. (2006) pointed out in footnote 12 of the 2001 working paper version that there is a deep connection between the recursivity of the constraint formulation of robust control problems and dynamic contract theory. The continuation entropy state variable effectively plays the role of the continuation utility state variable introduced to satisfy the promise-keeping constraint, or of the agent's continuation payoff state variable used to impose the individual rationality constraint in DeMarzo and Sannikov (2017). While the Inverse Euler Equation produces that the Lagrange multiplier is a martingale, the result in Proposition 3 goes a step further, as the Lagrange multiplier is in fact constant, conditional on the state variables. It is precisely the first-order condition in (26), i.e. the optimal allocation of entropy across states of nature that achieves this. We note that the optimal choice of the sensitivity g^* of the entropy state variable to fundamental shocks is the mechanism that handles model misspecification in our analysis, in the exact same way that the sensitivity of the agent's promised utility to the fundamental shocks is the mechanism to handle moral hazard by providing incentives in dynamic principal-agent problems (Williams (2015))[8].

3.3 Model Ambiguity

We now zoom in on the difference between model misspecification and model ambiguity. Model ambiguity concerns lead an investor to consider adverse alternatives in the set of structured models. In the current set-up, this corresponds to the special case where Nature restricts the choice of g to be zero, which we will show is not without loss of generality.

Consider the problem

$$\widetilde{V}(W,E) = \sup_{\pi,c} \inf_{u} \mathbb{E}^{\mathbb{U}}\Big[\int_0^\infty e^{-\delta t}\delta U(c_t)\,dt\Big], \tag{37}$$

subject to (14) and

$$dE_t = \big(\delta E_t - \tfrac{\delta}{2}u_t^2\big)dt. \tag{38}$$

In contrast to (12), Nature restricts g to be zero. We denote the optimal strategies of (37) as $\tilde{\pi}^*$, $\tilde{c}^*$, and $\tilde{u}^*$. For a given initial entropy budget E, we first consider a stationary solution where the optimal instantaneous distortion $\tilde{u}^*$ always satisfies

$$E = \tfrac{1}{2}(\tilde{u}^*)^2. \tag{39}$$

Then the drift of the continuation entropy in (38) is constantly zero and continuation entropy remains at its initial value E. The notion of a stationary solution mirrors the finding in Lemma 2 that the optimal allocation of continuation entropy is constant over time in the two-period case.[9] At the end of this subsection we show that the stationary case is not only the natural one to consider, but also the more interesting one, because the non-stationary solution converges to the boundaries of either zero or maximal entropy, unless it happens to coincide with the stationary case to begin with.

Let us also consider a constrained problem where the instantaneous distortion is restricted as in Chen and Epstein (2002), in order to construct a rectangular set of priors. Consider the problem

$$\check{V}(W,\kappa) = \sup_{\pi,c} \inf_{u} \mathbb{E}^{\mathbb{U}}\Big[\int_0^\infty e^{-\delta t}\delta U(c_t)\,dt\Big], \tag{40}$$

[8] As Williams (2015) shows, in the First Best, i.e., the full information benchmark without hidden action, the principal is free to choose the sensitivity optimally, in the same way that Nature is able to freely choose g^* in our analysis, but this ceases to be the case with private information.

[9] The absence of discounting in the two-period case is an additional reason for the stationarity of its solution, as can also be seen from (38) for $\delta = 0$.

subject to (14) and

$$\tfrac{1}{2}u^2 \le \kappa, \tag{41}$$

for a given $\kappa > 0$. To rule out degenerate cases, we assume $\kappa < \frac{1}{2}\left(\frac{\mu-r}{\sigma}\right)^2$. If this condition is violated, $u = \frac{\mu-r}{\sigma}$ is an interior worst-case distortion, in which case the investor's optimal investment weight is identically zero.

Introducing the Lagrange multiplier associated with constraint (41), the dual formulation of the problem (40) is

$$\check{V}(W,\kappa) = \sup_{\lambda\ge 0}\sup_{\pi,c}\inf_{u}\mathbb{E}^{\mathbb{U}}\left[\int_0^\infty e^{-\delta t}\Big(\delta U(c_t) + \lambda\delta\big(\tfrac{1}{2}u_t^2 - \kappa\big)\Big)dt\right], \tag{42}$$

subject to (14). In (42), $\lambda\delta$ is the Lagrange multiplier associated with constraint (41). We denote the optimal strategies of (42) as $\check{\pi}^*, \check{c}^*$, $\check{u}^*$, and the optimal Lagrange multiplier as λ^*.

The following result connects these two formulations.

Proposition 4. *There is a one-to-one correspondence between the problem* (37) *and the problem* (40)*. In particular, for a given E, the stationary solution of* (37) *is*

$$\tilde{u}^* = \frac{\mu-r}{\sigma}\frac{\theta(E)}{\gamma+\theta(E)}, \tag{43}$$

$$\tilde{\pi}^* = \frac{\mu-r}{\sigma^2}\frac{1}{\gamma+\theta(E)}, \tag{44}$$

$$\tilde{c}^* = \frac{1}{\gamma}\left[\delta - (1-\gamma)r - \frac{(1-\gamma)\gamma}{2(\gamma+\theta(E))^2}\frac{(\mu-r)^2}{\sigma^2}\right], \tag{45}$$

where $\theta(E)$ satisfies

$$\frac{\mu-r}{\sigma}\frac{\theta(E)}{\gamma+\theta(E)} = \sqrt{2E}. \tag{46}$$

For a given κ, the solution of (40) *is given in* (43)*,* (44)*, and* (45) *with $\tilde{u}^*$, $\tilde{\pi}^*$, and $\tilde{c}^*$ on the left-hand side replaced by $\check{u}^*$, $\check{\pi}^*$, and $\check{c}^*$, respectively, and $\theta(E)$ on the right-hand side replaced by $\theta(\kappa)$ which satisfies*

$$\frac{\mu-r}{\sigma}\frac{\theta(\kappa)}{\gamma+\theta(\kappa)} = \sqrt{2\kappa}. \tag{47}$$

Moreover, $\widetilde{V}$, $\check{V}$, and λ^* are all homogeneous of degree $1-\gamma$ in W.

For a given E or κ, the optimal instantaneous distortion $\tilde{u}^*$ and the optimal investment weight $\tilde{\pi}^*$ also match those in Maenhout (2004) exactly. To illustrate the connection further, problem (42) can be recast as

$$\check{V}(W,\kappa) = \sup_{\lambda\ge 0}\Big\{\overline{V}(W,\lambda) - \lambda\delta\kappa\Big\}, \tag{48}$$

where $\overline{V}$ is

$$\overline{V}(W,\lambda) = \sup_{\pi,c} \inf_{u} \mathbb{E}^{\mathbb{U}}\Big[\int_0^\infty e^{-\delta t}\Big(\delta U(c_t) + \lambda \tfrac{\delta}{2} u_t^2\Big) dt\Big]. \tag{49}$$

Maenhout (2004) introduces the homothetic scaling, i.e., $\lambda \propto W_t^{1-\gamma}$, which is consistent with the homothetic scaling of the optimal Lagrange multiplier λ^* in (48). We therefore provide an alternative justification and rigorous proof for the homothetic scaling. This is also consistent with results in Knox (2002) that "homothetic robustness" of Maenhout (2004) is structurally equivalent to $\kappa-$ignorance in Chen and Epstein (2002).

In conclusion, Proposition 4 highlights that the homothetic version of Anderson et al. (2003) proposed in Maenhout (2004) and applications thereof mentioned in footnote 2 correspond to Nature restricting g to be zero, hence they deal with model ambiguity, not model misspecification. Importantly, this criticism does not apply to Anderson et al. (2003), as we showed in Section 3.1.

We briefly discuss the non-stationary solution. In this case, while the diffusion of entropy is set to zero, its drift in general is not. We refer to Appendix A.7 for the details. In summary, for the parameters we consider, we find that the non-stationary solution either converges to the boundaries of $E = 0$ (expected-utility) or $E = \overline{E}$ (maximal entropy), or coincides with the stationary solution, namely when the initial value of E is such that its drift equals zero.

3.4 Log Utility

We now explicitly cover the case of log utility, given its special role in the analysis and its importance in the literature on robustness. For log utility, the value function of problem (13) is additive in wealth and entropy, i.e., (17) changes to

$$V(E,W) = \log W + \mathcal{V}(E). \tag{50}$$

Clearly $\partial^2_{WE} V = 0$ in this case, which has several fundamental implications. First, it is immediate that the optimal g^* in (16) is zero for log utility. The fact that the optimal g^* is zero confirms the finding in Proposition 1 that with log utility the distinction between model ambiguity and model misspecification vanishes. In other words, in the special case of log utility, it is without loss of generality to set $g = 0$, since $g^* = 0$ anyway. But it does mean that this case does not speak to model misspecification in the setting that we consider. The reason the log investor is not affected by model misspecification concerns in the basic asset allocation problem underlying macrofinance, is because the log investor behaves myopically. Since this investor is indifferent between intertemporal speculation and intertemporal hedging, he makes no attempt to take the dynamics of investment opportunities into account when investing, and only considers current investment opportunities. The solution therefore coincides with the solution to a problem with model ambiguity only.

Second, we reconsider the dual problem formulated in (28). We have shown that $\widehat{V}$ is the value function for the optimal consumption-investment problem with Hansen-Sargent multiplier preferences in (30). Even though $\widehat{V}$ is not homothetic in W when $\gamma \neq 1$, it is for log utility. Therefore, the multiplier preferences with a fixed multiplier still maintain tractability when the utility is log. This is related to the result in Cerreia-Vioglio et al. (2011) that "a preference is both variational and homothetic if and only if is multiple priors". For log utility, Hansen-Sargent multiplier preferences are indeed homothetic and equivalent to multiple prior preferences in the Merton investment problem. This is also consistent with the result in Cerreia-Vioglio et al. (2022b) that variational preferences (which multiplier preferences are) are wealth constant relative ambiguity averse when logarithmic.

A third way to understand the economic significance of the assumption of log utility, builds on the insights of Section 3.2, namely equation (34). The only case where the volatility of θ_t^* is zero for $g^* = 0$ (model ambiguity concerns only) is when $\partial^2_{WE} V = 0$, i.e., for log utility. Put differently, unless $\partial^2_{WE} V = 0$, as is true for log utility, the volatility of the multiplier does not vanish when $g = 0$.

Finally, consider the duality relationship in (31). Combining (33) and (50), it follows that the optimal Lagrange multiplier θ^* is independent of wealth for log utility, consistent with the general finding of homogeneity of degree $1-\gamma$ in W.

4 Quantitative Model Implications

In this section, we present quantitative results for optimal consumption-investment problems with model ambiguity and model misspecification.

4.1 State-Dependent Optimal Strategies and Worst-Case Beliefs

We solve the ODE resulting from plugging the optimal controls into HJB (18) numerically with boundary conditions (19) and (20). We use the following estimates for the equity premium, risk-free rate and return volatility from the postwar US sample ($1947Q2-2001Q2$) reported in Campbell (2018): $\mu - r = 7.39\%$, $r = 0.72\%$ and $\sigma = 0.1598$.

The numerical results are presented in Figure 4 for discount rate $\delta = 0.02$ and risk aversion coefficient $\gamma = 3$. We see that small amounts of entropy reduce the optimal portfolio very significantly starting from the optimal Merton solution for the $E = 0$ boundary of the red dotted line. The reduction in the portfolio weight reflects the combined effect of a reduction in the average equity premium according to the worst-case belief distortion u^* and the optimal negative intertemporal hedging demands due to adverse dynamics of investment opportunities. The dynamics of the belief distortion and how the investor perceives investment opportunities are driven by the optimal volatility g^* of continuation entropy, which allows Nature to allocate entropy in a state-dependent fashion. The fact that $g^* < 0$ confirms the result in Propositions 1 and 2 that continuation entropy is countercyclical, so as to induce persistence in returns under the worst-case measure and a positive correlation between returns and expected returns for $\gamma > 1$. This is indeed a worst-case model misspecification for a risk-averse ($\gamma = 3$) investor, as it precludes intertemporal hedging.

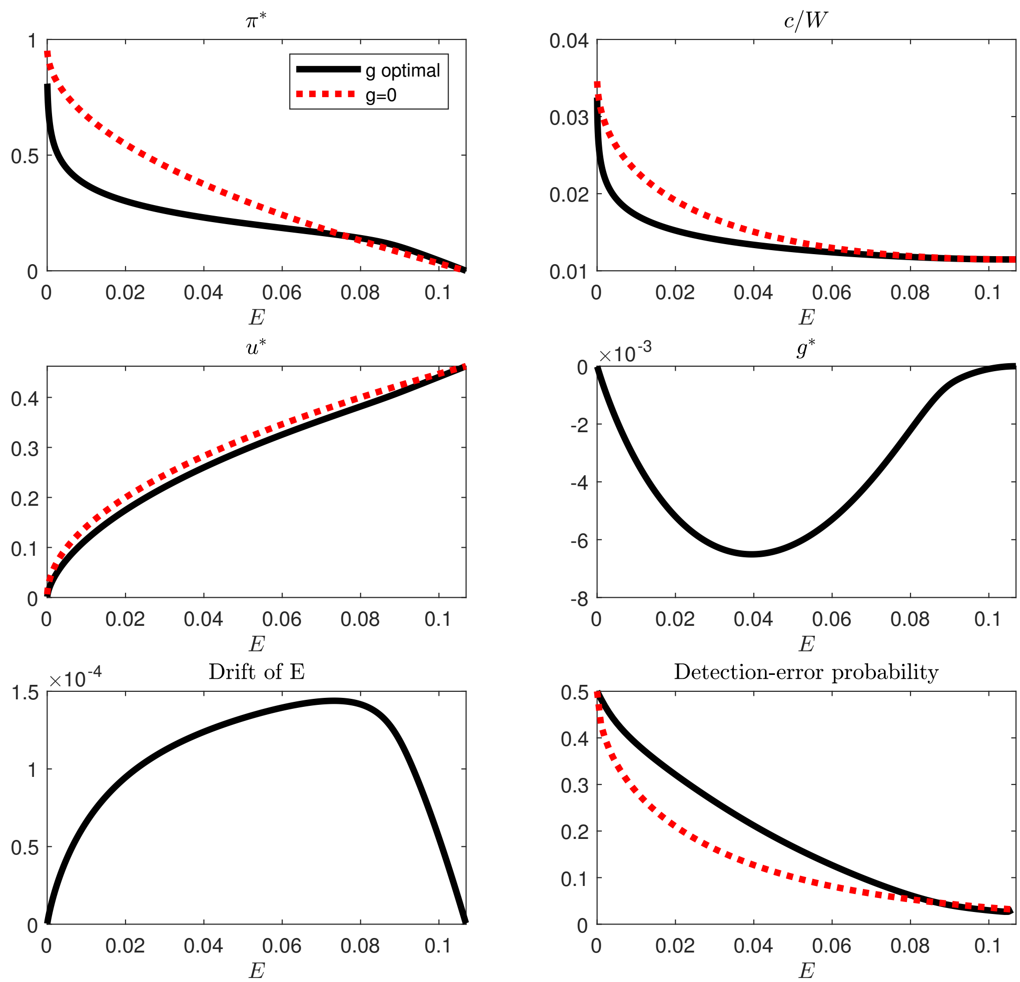

Figure 4:

Notes: This figure plots the optimal investment strategy π^*, the optimal consumption wealth ratio c/W, the worst-case belief distortion u^*, the entropy volatility g^*, the drift of entropy, and the detection-error probability with respect to the continuation entropy. The parameters are $\delta = 0.02$, $\gamma = 3$, $\mu - r = 0.0739$, $r = 0.0072$, and $\sigma = 0.1598$.

Because of the countercyclical dynamics of continuation entropy E, positive (negative) return shocks reduce (increase) E, and as a result optimal investment π^* and consumption wealth ratio c/W increase (decrease), meanwhile the belief distortion u^* decreases (increases), as the first three panels of Figure 4 show.

Higher values of entropy, corresponding to a stronger preference for robustness, lead to a lower consumption-wealth ratio and reflect a higher savings rate. This can be explained by two economic forces. First, the concern for model misspecification reduces the endogenous return on wealth for the agent. Since γ exceeds 1, the income effect dominates and the investor saves more. Second, robustness strengthens the precautionary savings motive. The reduced demand for the risky asset and the increased desire to save for precautionary reasons will of course help to explain a high equity premium and a low risk-free rate.

Not only the volatility, but also the endogenous drift of continuation entropy are important to understand the dynamics of the model. For the parameter values in Figure 4, the drift of continuation entropy is positive, so that as time passes by we expect u^* to drift up according to the increase in E. However, for the parameter values in Figure 4, the magnitude of continuation entropy volatility g^* is much larger than the magnitude of its drift, resulting in considerable randomness in the path of continuation entropy. Notice that both $E = 0$ and $E = \overline{E}$ are absorbing barriers. In both cases, the volatility g^* vanishes at the boundary. For $E = 0$, we have $u^* = 0$ so that the drift of E is zero as well. Similarly, the drift of E at $\overline{E}$ is $\overline{E} - \frac{1}{2}\overline{u}^2 = 0$. Therefore, whenever continuation entropy reaches either boundary, it stays there forever.

Our numerical analysis also illustrates the contrast between model ambiguity and model misspecification further. The difference in results is clear in Figure 4 where the red dotted lines graph the optimal portfolio weight π^*, the optimal consumption-wealth ratio $\tilde{c}^*$, and the optimal belief distortion u^* for $g = 0$. The parameters are otherwise identical to the ones used for the optimal g^* solution. Model ambiguity has a substantially smaller effect on the optimal portfolio share and optimal consumption than model misspecification does. It is also remarkable that a given value of E corresponds to a larger (more pessimistic) belief distortion u^* for the model ambiguity case, and yet has a smaller effect on portfolio choice and savings behavior. This highlights the importance of the state-dependence in u^* for the model misspecification case. For the model ambiguity case ($g = 0$), where the volatility of continuation entropy is zero, the drift vanishes as well, due to the stationary solution. Therefore, continuation entropy remains constant and all optimal strategies and worst-case belief distortions are constant as well. In sharp contrast, for the model misspecification (optimal g^*) case, continuation entropy is dynamic, inducing dynamic behavior for optimal strategies and worst-case belief. Belief distortions that are constant have less of an impact on behavior, even when the constant distortions are currently bigger. The state-dependence in u^* when entropy is allocated optimally makes it more potent in harming the agent, as Nature finds the state-contingent model perturbation that matters the most to the decision-maker. Furthermore, entropy drifts upwards, so that the agent fearing model misspecification anticipates larger distortions in the future, while the currently higher distortions with $g = 0$ stay constant and do not grow over time.

The last panel of Figure 4 plots the detection-error probabilities (see Appendix A.8), using a sample length T corresponding to the postwar data in Campbell (2018) of $T = 65$ years, for a given value of E. The model misspecification case with optimal g^* results in substantially higher detection-error probabilities than the case of model ambiguity ($g = 0$). Given 65 years of observations, detection-error probabilities can be used to find a reasonable initial value for the continuation entropy at time zero, i.e., E_0, so that it is challenging to distinguish between the worst-case belief and the objective data generating process. For $\gamma = 3$ a conservative detection-error probability of 20% allows $E_0 = 0.04$ when worried about model misspecification, but only $E_0 = 0.02$ when concerned about model ambiguity. This implies that the optimal portfolio weight at time 0 in the former case is less than half the optimal weight in the latter. The potential impact of model misspecification is therefore substantial. The same result obtains for the optimal consumption-wealth ratio.

We close this example by noting that the more direct mapping from entropy E_0 to detection-error probabilities is another advantage of the continuation entropy approach, as it avoids the extra steps that are otherwise required to infer the entropy that corresponds to a given worst-case distortion for a penalty parameter value θ.

Figure 5 plots the case where $\gamma = 5$ and $\delta = 0.02$. Relative to $\gamma = 3$ in Figure 4, all effects are more pronounced, due to the higher risk aversion. The state-dependence of entropy, driven by the optimal loading g^* on the Brownian motion is more significant and reflects the stronger preference for intertemporal hedging, and therefore greater fear of return persistence of the more risk-averse investor. Similarly, the stronger precautionary savings motive reduces the consumption-wealth ratio more relative to the no robustness case ($E = 0$ values on the red dotted lines).

Figure 6 shows the effect of a higher discount rate, namely $\delta = 0.05$ rather than $\delta = 0.02$ as in Figure 5. Not surprisingly, the less patient agent saves less and consumes more. Another consequence of the higher rate of time preference is the higher drift of the entropy state variable. As is clear from (12), when the drift of entropy is positive, a higher value of δ increases the drift. This effect will be important for understanding the simulated entropy dynamics in the next subsection, and especially in generating endogenous nonparticipation for sufficiently impatient and risk-averse investors.

Finally, Figure 7 presents the results for the risk-tolerant $\gamma = 0.8$ investor. Four findings stand out. First, the optimal consumption-wealth ratio increases relative to the Merton solution ($E = 0$).

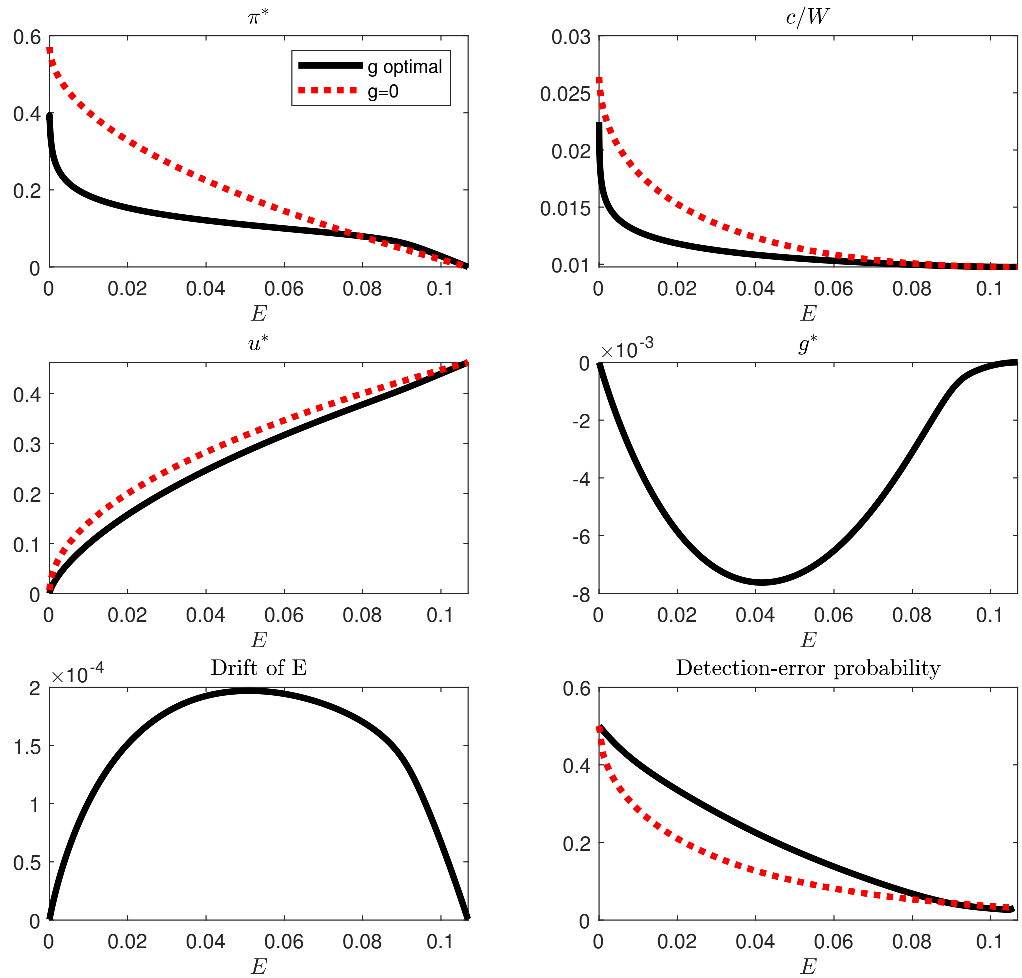

Figure 5:

Notes: This figure plots the optimal investment strategy π^*, the optimal consumption wealth ratio c/W, the worst-case belief distortion u^*, the entropy volatility g^*, the drift of entropy, and the detection-error probability with respect to the continuation entropy. The parameters are $\delta = 0.02$, $\gamma = 5$, $\mu - r = 0.0739$, $r = 0.0072$ and $\sigma = 0.1598$.

The reason is twofold. The substitution effect dominates the income effect for the risk-tolerant investor, so that he consumes more, not less, when the return on wealth is reduced. Moreover, the precautionary savings motive is relatively small and insufficient to trump the intertemporal substitution effect. A second important result in Figure 7 is that indeed $g^* > 0$ in this case. A positive diffusion coefficient for entropy results in mean reversion for stock returns under the worst-case model, rather than persistence, which would be required for the intertemporal speculation that is desired by a $\gamma < 1$ investor. The resulting negative speculative demands reduce the total optimal investment strategy below the Merton solution, with the reduction increasing

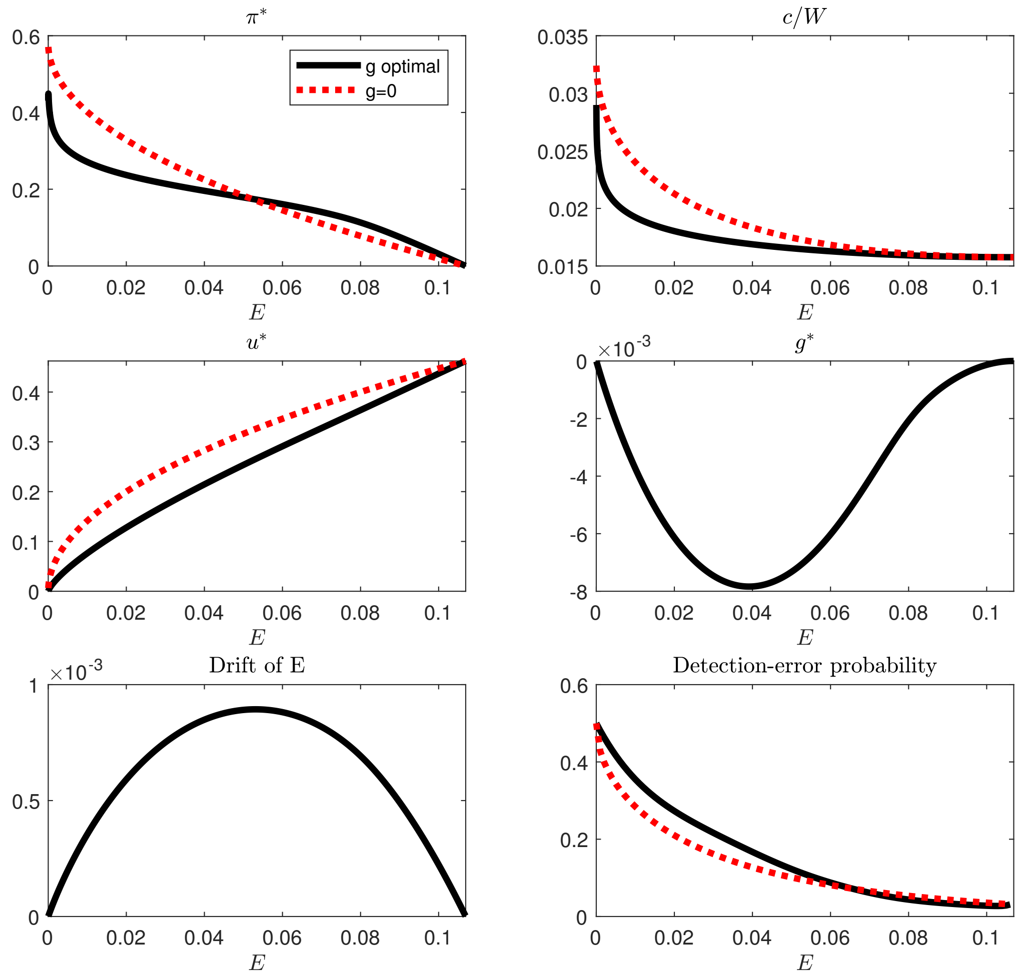

Figure 6:

Notes: This figure plots the optimal investment strategy π^**, the optimal consumption wealth ratio* c/W*, the worst-case belief distortion* u^**, the entropy volatility* g^**, the drift of entropy, and the detection-error probability with respect to the continuation entropy. The parameters are* $\delta = 0.05$*,* $\gamma = 5$*,* $\mu - r = 0.0739$*,* $r = 0.0072$ *and* $\sigma = 0.1598$.

in entropy. The lower optimal investment strategy also reflects a reduced myopic asset demand, due to the diminished equity premium according to the worst-case belief distortion u^*. A third finding in Figure 7 is that now the drift of entropy turns negative for higher values of entropy. Unlike in the previous figures for $\gamma > 1$, for $\gamma = 0.8$ the difference in optimal investment strategy between the case of model misspecification and model ambiguity seems smaller. This is not surprising given that $\gamma = 0.8$ is much closer to log than the other cases, where the distinction between model ambiguity and model misspecification blurs, as shown in Subsection 3.4.

In summary, the continuous-time model confirms and extends all results of the binomial tree

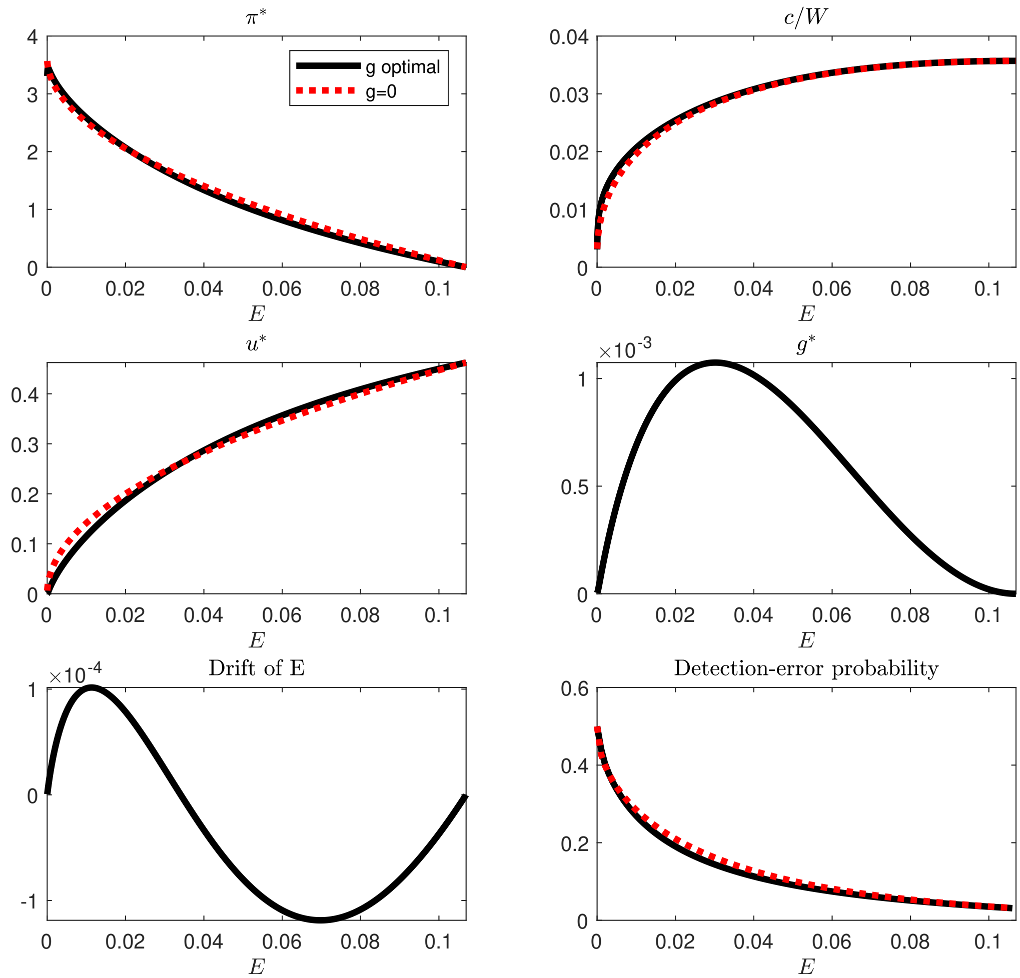

Figure 7:

Notes: This figure plots the optimal investment strategy π^**, the optimal consumption wealth ratio* c/W*, the worst-case belief distortion* u^**, the entropy volatility* g^**, the drift of entropy, and the detection-error probability with respect to the continuation entropy. The parameters are* $\delta = 0.03$*,* $\gamma = 0.8$*,* $\mu - r = 0.0739$*,* $r = 0.0072$ *and* $\sigma = 0.1598$.

intuition. The demand for the risky asset is substantially reduced due to the concerns for model misspecification, and reflects both a pessimistic distortion of current investment opportunities and a *dynamic perception* of investment opportunities. While the first effect also occurs with model ambiguity, the second effect is novel and unique to model misspecification. A risk-averse investor fears return persistence, as this means she cannot hedge intertemporally. A risk-tolerant investor on the other hand fears mean reversion in stock returns, since it rules out intertemporal speculation.

4.2 Simulated Entropy Dynamics

We now turn to the simulated dynamics of continuation entropy E implied by the model, for the parameter values considered in the previous subsection. We present the histogram for E from 10^5 simulated paths over 10 years and over 50 years under the pessimistic measure $\mathbb{U}$. The entropy state variable is initialized at $E_0 = \frac{1}{2}\overline{E}$, which corresponds to a very reasonable detection-error probability of roughly 20% according to Figure 4. For moderate risk aversion ($\gamma = 3$) and time discounting ($\delta = 0.02$), all paths in Figure 8 end up in the interior of the domain of E after 10 years, with optimal beliefs and portfolio and consumption decisions that are in between the extremes of no misspecification concerns ($E = 0$) and the upper boundary $E = \overline{E}$ where zero investment in the risky asset is optimal. However, when simulating paths for 50 years, Figure 8 shows that for $\gamma = 3$ roughly 10% of the paths end up at the absorbing barrier of $E = 0$ (expected utility), where the investor ceases to be concerned with model misspecification or model ambiguity and embraces expected-utility behavior with beliefs according to the benchmark model $\mathbb{B}$. This happens as a result of the investor experiencing a series of positive returns, which reduce the entropy state variable to zero and make the investor's pessimism fade away eventually. The behavior exhibited by this investor can be viewed as consistent with the evidence in Malmendier and Nagel (2011) of what they call the experience hypothesis, where risk-taking behavior is affected by the investor's personal experience in terms of the series of returns encountered by the investor. As an alternative to explanations in the finance literature based on constant-gain learning (see for instance Nagel and Xu (2022)), our model offers an intuitive explanation driven by pessimistic beliefs due to model misspecification concerns rather than by learning. There are no paths in Figure 8 that lead to $E = \overline{E}$ after 50 years, i.e., the other extreme where nonparticipation becomes optimal after experiencing a long series of negative returns. For different parameter values we will see that our model can generate nonparticipation endogenously, but first we show in more detail the extent to which the past experience of the investor affects her behavior, through the endogenous dynamics of the entropy state variable.

A first way to illustrate the experience hypothesis is to show the effect of early belief scarring for an investor who experiences 1 year of very negative returns. In particular, to mimic the experience of the "depression babies" of Malmendier and Nagel (2011), we subject the investor described so far to 1 year of negative return shocks, namely such that the average return experienced that year is 2 standard deviations below the mean, i.e., roughly −24%. After this disastrous year, the investor faces the same simulated dynamics as the 'non-depression' normal investor. The effect of the belief scarring is quite dramatic and long-lasting, in accordance with the experience hypothesis. Many more paths end up near the right-hand side limit of the state space than for the normal agent, and fewer end up at $E = 0$. Even after 50 years, the belief scarring has not disappeared and continues to affect beliefs and behavior.

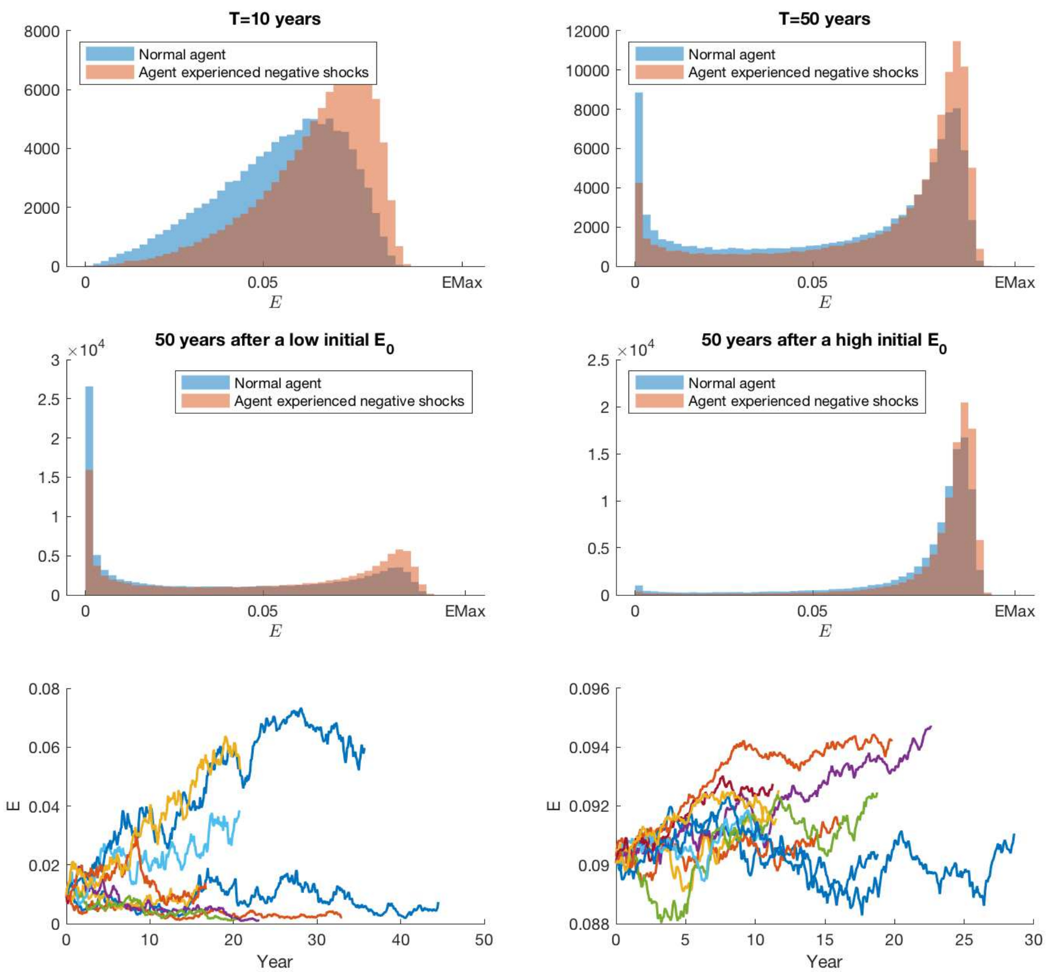

Figure 8:

Notes: This figure plots the histogram for the continuation entropy state variable E corresponding to 10^5 paths over 10 years and over 50 years in Monte Carlo simulations for the return shocks. The parameters are the same as in Figure 4 ($\delta = 0.02$ and $\gamma = 3$). The entropy state variable is initialized at $E_0 = \frac{1}{2}\overline{E}$, except in the middle panels, where $E_0 = 0.3\overline{E}$ (low initial E_0), and $E_0 = 0.7\overline{E}$ (high initial E_0). The agent who experienced negative shocks had a return on the risky asset that is 2 standard deviations below the mean in the first year of the simulation, but otherwise experienced the same return shocks as the normal agent in the next 50 years.

Second, the initial value for E also matters a great deal. We consider both the normal investor, as well as the 'depression baby', for $E_0 = 0.3\overline{E}$, and for $E_0 = 0.7\overline{E}$. The first case can be viewed as an investor who is less concerned with model misspecification and who starts off with less pessimism, the second case being a very pessimistic investor. One could interpret this as the effect of prior belief scarring. Compared to the standard initialization of $E_0 = \frac{1}{2}\overline{E}$ in the top panels, the behavior is again very significantly impacted, even though the histogram represents the state variable half a century after the initialization. Investor heterogeneity in initial model misspecification concerns and in preference for robustness leads to rich dynamics and a wide cross-section of investor behavior, even after 50 years.

In the bottom 2 panels we zoom in on individual sample paths to understand the dynamics of entropy near the boundaries. We report 10 randomly selected sample paths for E that reach a threshold of either $E = 0.01$ (near $E = 0$) or $E = 0.09$ (near $E = \overline{E}$). Whenever this threshold is reached, we report the path. Because the threshold is reached at different times, the paths have different length. The bottom left panel shows that the dynamics do allow very low entropy to escape the lower boundary of $E = 0$. Near the upper boundary on the other hand, where entropy is very high and the investor extremely pessimistic, it seems very difficult to overcome the pessimism, and the entropy state variable exhibits a lot of persistence. The beliefs of these investors seem to be stuck in a state of extreme pessimism. The strong persistence of extreme pessimism is also apparent from the middle right panel of Figure 8, which shows, by the Markov property, that when the state variable E reaches a high level $0.7\overline{E}$, the majority of paths for entropy still result in very high levels of entropy, even after 50 years.

Finally, it is noteworthy that our model is also consistent with the finding in Malmendier and Nagel (2011) that young investors with shorter life histories are more sensitive to recent returns than older investors. The mechanism behind this in our model is that we assume that young investor with short investment histories start their entropy state variable close to the initial value $E_0 = \frac{1}{2}\overline{E}$, while older investors are very spread out in terms of entropy state variable according to the histograms in Figure 8. The sensitivity of investors to recent returns is driven by g^*, which in Figure 4 clearly has larger magnitude in the middle of the state space (where young investors are likely to be) than near the edges (where old investors are more likely to be).

On the other hand, when $\gamma = 0.8$, both histograms in Figure 9 indicate stationarity of the distributions, without any case of either expected-utility behavior or nonparticipation. The reason for this finding is the mean-reverting drift of entropy for this case: as revealed in Figure 7, the drift of entropy is positive for low values of the entropy state variable and negative for high values, which drives the mean-reversion that explains the stationary distribution in this case. There is much less evidence of behavior consistent with the experience hypothesis in this case, and the effect of belief scarring is now very subdued, and essentially gone after 50 years.

For $\gamma = 5$ and $\delta = 0.1$ in Figure 10 we find that a concern for model misspecification has the potential to endogenously explain nonparticipation, in a model with homotheticity and without relying on frictions. While only a handful of paths lead to nonparticipation for the 'normal investor', the investor having experienced a negative first year does in roughly 10% of the cases exit the equity market after 50 years. As discussed in Campbell (2018) and Gomes et al. (2021), one of the challenging stylized facts documented in recent work on household finance is limited

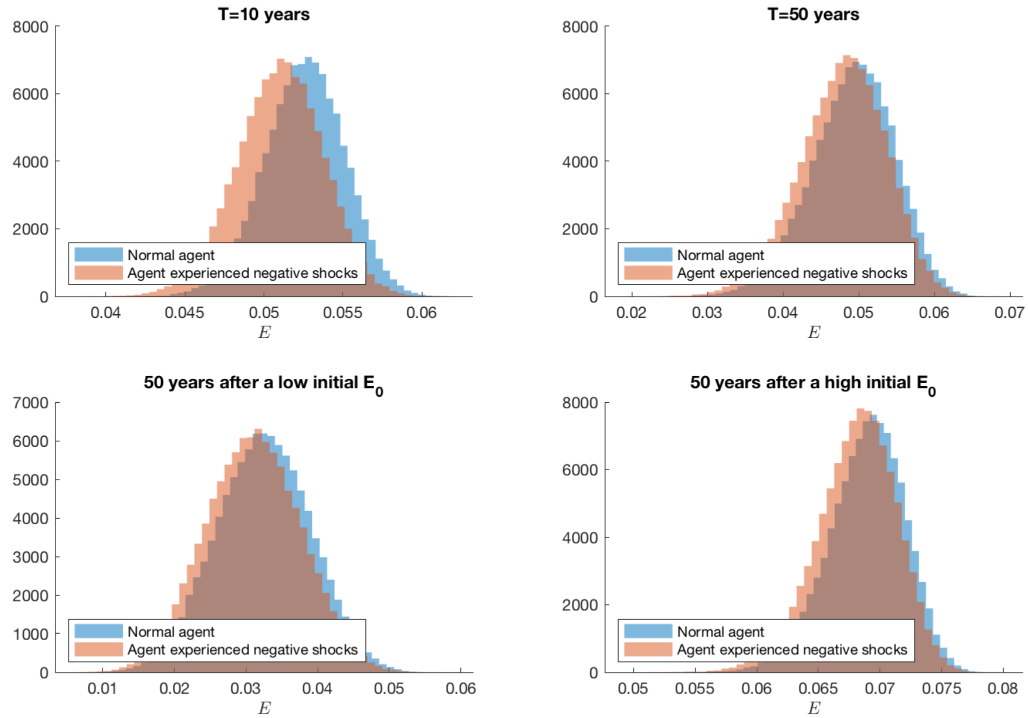

Figure 9:

Notes: This figure plots the histogram for the continuation entropy state variable E corresponding to 10^5 paths over 10 years and over 50 years in Monte Carlo simulations for the return shocks. The parameters are the same as in Figure 7 ($\delta = 0.03$ and $\gamma = 0.8$). The entropy state variable is initialized at $E_0 = \frac{1}{2}\overline{E}$, except in the bottom panels, where $E_0 = 0.3\overline{E}$ (low initial E_0), and $E_0 = 0.7\overline{E}$ (high initial E_0). The agent who experienced negative shocks had a return on the risky asset that is 2 standard deviations below the mean in the first year of the simulation, but otherwise experienced the same return shocks as the normal agent in the next 50 years.

stock-market participation, which is particularly hard to rationalize among wealthy households. Based on household survey data, Dimmock et al. (2016) find a significant negative relationship between stock market nonparticipation and their measure of ambiguity aversion.[10] In our model a concern for model ambiguity is not sufficient to generate the type of dynamics displayed in Figures 10, because the belief distortion, as well as corresponding optimal behavior, is constant in our setting with IID returns, so that nonparticipation cannot be generated endogenously as the result of a succession of negative stock returns. In contrast, the case of model misspecification can be interpreted as an investor who becomes increasingly pessimistic when experiencing a string of negative return shocks. As investor's continuation entropy increases due to a string of negative shocks, it may eventually reach the upper boundary $E = \overline{E}$, where the investor's beliefs

[10] The effect of risk aversion on nonparticipation in their data is insignificant, or at least not significant in a robust way.

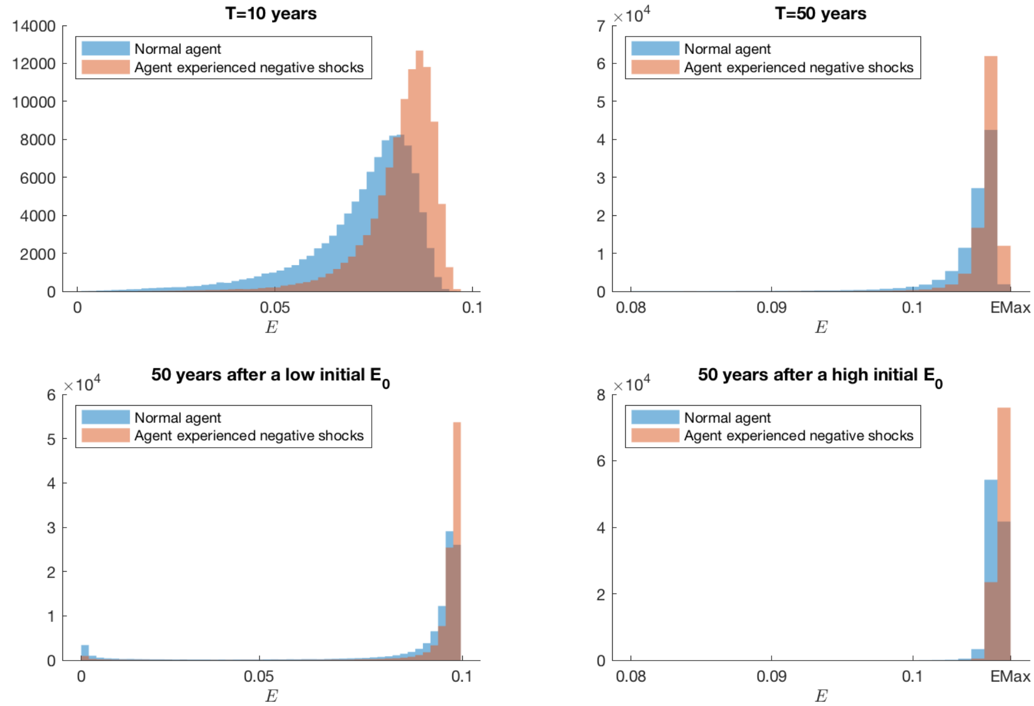

Figure 10:

Notes: This figure plots the histogram for the continuation entropy state variable E corresponding to 10^5 *paths over 10 years and over 50 years in Monte Carlo simulations for the return shocks. The parameters are* $\delta = 0.1$ *and* $\gamma = 5$*). The entropy state variable is initialized at* $E_0 = \frac{1}{2}\overline{E}$*, except in the bottom panels, where* $E_0 = 0.3\overline{E}$ *(low initial* E_0*), and* $E_0 = 0.7\overline{E}$ *(high initial* E_0*). The agent who experienced negative shocks had a return on the risky asset that is 2 standard deviations below the mean in the first year of the simulation, but otherwise experienced the same return shocks as the normal agent in the next* 50 *years.*

are sufficiently pessimistic to justify nonparticipation. Once continuation entropy reaches $\overline{E}$, it stays there forever. The mechanism can be viewed as a dynamic version of the model of Dow and Werlang (1992). It is important to point out that this only happens for sufficiently high risk aversion combined with sufficient impatience ($\delta = 0.1$). For more reasonable values of the rate of time preference δ the drift of entropy is not sufficiently large for high values of entropy to be able to reach the upper boundary $\overline{E}$. Economically speaking, it is not surprising that extreme parameter values as well as very pessimistic beliefs are required to endogenously generate nonparticipation.[11] In the final panel, where the entropy state variable is initialized at a high value of $E_0 = 0.7\overline{E}$, exit and nonparticipation is optimal for almost half the 'normal' investors and almost all 'depression' investors.

[11]For even higher risk aversion and impatience, nonparticipation can happen sooner. For example with $\gamma = 8$ and $\delta = 0.15$, 60 % of normal agents end up not participating after 30 years. The reason for this is that higher δ increases the positive drift of E.

5 General Equilibrium Asset Pricing

In this section, we examine the general equilibrium implications of model misspecification. We consider a Lucas economy in which the representative agent with Epstein-Zin preferences is concerned about model misspecification. We show that the equilibrium interest rate, expected stock return, and equity volatility are state-dependent. Moreover, model misspecification concerns can generate excess volatility.

5.1 Set-up of the Lucas Economy

Consider an economy with a single Lucas tree. We assume that its dividend follows the dynamics

$$\frac{dc_t}{c_t} = \mu^c dt + \sigma^c dB_t^{\mathbb{B}}, \tag{51}$$

where μ^c and σ^c are constants. The equilibrium price P^L of this Lucas tree is assumed to follow

$$\frac{dP_t^L + c_t dt}{P_t^L} = \mu_t^L dt + \sigma_t^L dB_t^{\mathbb{B}}, \tag{52}$$

whose expected return μ^L and volatility σ^L, together with the risk-free interest rate r will be determined in equilibrium.

Consider a representative agent whose preference over consumption streams is described by a continuous-time stochastic differential utility of the Kreps-Porteus and Epstein-Zin type. Given a discount rate δ, relative risk aversion $0 < \gamma \neq 1$, and EIS $0 < \psi \neq 1$, the Epstein-Zin aggregator F is $F(c,v) \equiv \delta \frac{c^{1-\frac{1}{\psi}}}{1-\frac{1}{\psi}} \left((1-\gamma)v\right)^{1-\frac{1}{\nu}} - \delta \nu v$, with $\nu = \frac{1-\gamma}{1-\frac{1}{\psi}}$ (see, e.g., Duffie and Epstein (1992)). The representative agent is concerned about model misspecification. The robust Epstein-Zin preference for a consumption stream c is

$$\mathcal{U}_t^c = \inf_{u,g} \mathbb{E}_t^{\mathbb{U}} \left[\int_t^{\infty} F(c_s, \mathcal{U}_s^c) ds \right], \tag{53}$$

where nature controls u and g subject to the continuation entropy dynamics in (12).

Given P^L and r, the representative agent invests and consumes with the optimal investment strategy π^* and the optimal consumption strategy c^*.

Definition 5.1. *$(r, \mu^L, \sigma^L, \pi^*, c^*)$ is an equilibrium if*

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

5.2 Equilibrium Results

The following result summarizes the equilibrium quantities.

Proposition 5. *Let $\mathcal{V}$ be a solution of* (144) *with boundary conditions* (146) *in the appendix. Then the equilibrium expected return, volatility, and risk-free interest rate are given by*

$$\mu^L - r = \gamma^{eff}(\sigma^L)^2, \tag{54}$$

$$\sigma^L = \Big[\frac{\mathcal{V}\partial^2_{EE}\mathcal{V}}{\mathcal{V}\partial^2_{EE}\mathcal{V} + (\psi-1)(\partial_E\mathcal{V})^2}\Big]\sigma^c, \tag{55}$$

$$\begin{aligned} r =& \delta + \frac{\mu^c}{\psi} - \frac{1+\psi}{2\psi}\gamma^{eff}(\sigma^L)^2 \\ &- \frac{1}{\nu}\frac{\partial_E\mathcal{V}}{\mathcal{V}}\Big[\frac{\delta}{2}(u^*)^2 - g^*(u^*+\sigma^L)\Big] - \frac{1}{2\nu}\Big[\Big(1+\frac{\psi}{\nu}\Big)\frac{(\partial_E\mathcal{V})^2}{\mathcal{V}^2} - \frac{\partial^2_{EE}\mathcal{V}}{\mathcal{V}}\Big](g^*)^2, \end{aligned} \tag{56}$$

where γ^{eff} is the effective risk aversion in (27). *When either $\gamma, \psi > 1$ or $0 < \gamma, \psi < 1$, there is excess volatility, i.e., $\sigma^L > \sigma^c$.*

The results indicate several important implications of model misspecification concerns in equilibrium asset pricing. The equilibrium equity premium is stochastic and state-dependent, reflecting the state-dependent effective risk aversion, which acts as a generalized price of risk. Recalling the equation for γ^{eff} from (27), the equilibrium equity premium in (54) has the following decomposition:

$$\mu^L - r = \gamma(\sigma^L)^2 + u^*\sigma^L + \tfrac{1-\gamma}{\delta}\tfrac{g^*}{u^*}(\sigma^L)^2. \tag{57}$$

This shows that the equilibrium equity premium is the sum of three components: the traditional *price of risk* $\gamma(\sigma^L)^2$, the *price of model ambiguity* $u^*\sigma^L$, and finally the *price of model misspecification* $\frac{1-\gamma}{\delta}\frac{g^*}{u^*}(\sigma^L)^2$, which is proportional to g^*/u^* the relative concern of model misspecification versus model ambiguity. Moreover, the second and third terms are positive, implying a unambiguously higher equity premium due to model ambiguity as well as model misspecification concerns.

Furthermore, for $\gamma > 1$ we expect the equilibrium Sharpe ratio to be countercyclical, reflecting $g^* < 0$. The subtle point is that the investor fears persistence in returns, and therefore invests cautiously. Equilibrium returns must adjust in order to entice the investor to hold the Lucas tree. As a result, equilibrium returns in fact mean-revert.

Turning to equilibrium volatility, the relationship in (55) can be rewritten, using (23) and (24), as

$$\sigma^L = \frac{1}{1+\frac{\psi-1}{\delta}\frac{g^*}{u^*}}\sigma^c. \tag{58}$$

As a result, $g^* \neq 0$, due to model misspecification, plays a key role in driving a wedge between return volatility σ^L and dividend (or consumption) volatility σ^c. This stands in sharp contrast to the case of model ambiguity, which yields $g^* = 0$ hence $\sigma^L = \sigma^c$. To generate excess volatility in our model, we need either $\gamma, \psi > 1$ or $0 < \gamma, \psi < 1$. In the former case, $\psi > 1$ and $g^* < 0$; in the later case, $\psi < 1$ and $g^* > 0$. In both cases, $\frac{(\psi-1)g^*}{u^*} < 0$, producing a factor exceeding unity in front of σ^c in (58), hence $\sigma^L > \sigma^c$. In addition to driving a wedge between return and consumption volatility, model misspecification also produces stochastic volatility, as the wedge is state-dependent and stochastic.

The equilibrium risk-free interest rate also reflects the additional and novel effects of model misspecification. The first two terms in (56) represent the effect on interest rates of impatience and of growth combined with intertemporal substitution, and are of course standard in any intertemporal macro model. The third term generalizes the usual precautionary savings motive, to reflect the higher effective risk aversion due to model ambiguity as well as model misspecification concerns. The second line in (56) reveals several additional terms, which are only present when $g^* \neq 0$ and therefore uniquely due to model misspecification worries.

Finally another insight from (57) is that the agent's pessimistic belief about the equilibrium Sharpe ratio can be characterized as follows:

$$\frac{\mu^L - r}{\sigma_L} - u^* = \left[\gamma + \frac{1-\gamma}{\delta}\frac{g^*}{u^*}\right](\sigma^L). \tag{59}$$

The gap between the baseline (objective) belief and the pessimistic belief about the equilibrium Sharpe ratio is u^*, i.e., the price of model ambiguity. A novel result is that the pessimistic assessment of the Sharpe ratio reflects the concern for model misspecification. Because $(1-\gamma)g^* > 0$ and $\sigma^L > \sigma^c$, the pessimistic assessment of the Sharpe ratio is always higher than $\gamma\sigma^c$, which is the assessment of the Sharpe ratio in the setting with model ambiguity only. Of crucial importance for conditional asset pricing are the dynamics of u^* and g^*. Based on the findings throughout the paper we expect u^* to be countercyclical for $\gamma > 1$, but the equilibrium dynamics $\frac{g^*}{u^*}$ are also expected to play a key role.

6 Conclusion

We show that model misspecification has profound effects on optimal asset allocation, dynamic savings behavior, and equilibrium asset pricing. We use a two-period binomial tree, the simplest possible dynamic model, to develop the analysis before extending to a general continuous-time Merton problem. For non-myopic investors, model misspecification concerns induce major changes in the optimal demand for the risky asset and in the consumption-savings policy function. When fearing that a baseline model of IID returns is misspecified, the robust solution to the Merton problem involves a substantially lower myopic demand, combined with negative intertemporal hedging demands for a risk-averse investor who fears return persistence, and negative speculative demands for a risk-tolerant investor who fears that stock returns might exhibit mean reversion. Our approach guarantees homotheticity and nests existing work on robust portfolio decisions, highlighting the difference between model misspecification and model ambiguity. Modelling entropy directly as an endogenous state variable also facilitates the calibration based on detection-error probabilities. For sufficiently impatient and risk-averse investors, our model can generate endogenous nonparticipation resulting from an increase in model uncertainty concerns and corresponding pessimism after experiencing a string of negative stock returns. Our model also generates belief scarring patterns, where investors who have experienced very negative returns early on in their investing life carry this experience with them for a long time. We are therefore able to offer an alternative explanation for the empirical evidence of the experience hypothesis (Malmendier and Nagel (2011)). For patient investors, a series of positive stock returns can result in the investor gradually overcoming uncertainty aversion and eventually embracing expected-utility behavior.

We further unpack the additional contribution of our analysis relative to existing work, by juxtaposing the effect of model misspecification on equilibrium asset pricing with the effect of model ambiguity.

A Appendix

This appendix contains the proofs of the lemmas and propositions, as well as more detail on the derivation of the detection-error probabilities.

A.1 Proof of Lemma 1

The Lagrangian of the problem (4) is

$$\begin{aligned}\mathscr{L}^S =& \frac{1}{1-\gamma}\Big(q_S \left\{W_1^S \left(\pi_S(H-R_f)+R_f\right)\right\}^{1-\gamma} \\ &+(1-q_S)\left\{W_1^S \left(\pi_S(L-R_f)+R_f\right)\right\}^{1-\gamma}\Big) \\ &+\lambda_S\Big(q_S \ln\left(\frac{q_S}{p}\right)+(1-q_S)\ln\left(\frac{1-q_S}{1-p}\right)-K_S\Big), \quad S\in\{H,L\}. \end{aligned} \tag{60}$$

Conditioning on the up node at $t=1$, the FOCs with respect to π_H, q_H, and λ_H are

$$\begin{aligned}\frac{\partial \mathscr{L}^H}{\partial \pi_H} =& q_H \left\{W_1^H \left(\pi_H(H-R_f)+R_f\right)\right\}^{-\gamma} W_1^H(H-R_f) \\ &+(1-q_H)\left\{W_1^H \left(\pi_H(L-R_f)+R_f\right)\right\}^{-\gamma} W_1^H(L-R_f)=0, \end{aligned} \tag{61}$$

$$\begin{aligned}\frac{\partial \mathscr{L}^H}{\partial q_H} =& \frac{1}{1-\gamma}\Big(\left\{W_1^H \left(\pi_H(H-R_f)+R_f\right)\right\}^{1-\gamma} \\ &-\left\{W_1^H \left(\pi_H(L-R_f)+R_f\right)\right\}^{1-\gamma}\Big) \\ &+\lambda_H\Big(1+\ln\left(\frac{q_H}{p}\right)-1-\ln\left(\frac{1-q_H}{1-p}\right)\Big)=0, \end{aligned} \tag{62}$$

$$\frac{\partial \mathscr{L}^H}{\partial \lambda_H} = q_H \ln\left(\frac{q_H}{p}\right)+(1-q_H)\ln\left(\frac{1-q_H}{1-p}\right)-K_H=0. \tag{63}$$

Conditioning on the down node at $t = 1$, the FOCs with respect to π_L, q_L, and λ_L are

$$\begin{aligned}
\frac{\partial \mathscr{L}^L}{\partial \pi_L} =& q_L \left\{W_1^L \left(\pi_L(H-R_f)+R_f\right)\right\}^{-\gamma} W_1^L(H-R_f) \\
&+ (1-q_L)\left\{W_1^L \left(\pi_L(L-R_f)+R_f\right)\right\}^{-\gamma} W_1^L(L-R_f) = 0, && (64)\\
\frac{\partial \mathscr{L}^L}{\partial q_L} =& \frac{1}{1-\gamma}\Big(\left\{W_1^L \left(\pi_L(H-R_f)+R_f\right)\right\}^{1-\gamma} \\
&- \left\{W_1^L \left(\pi_L(L-R_f)+R_f\right)\right\}^{1-\gamma}\Big) \\
&+ \lambda_L\left(\ln\left(\frac{q_L}{p}\right) - \ln\left(\frac{1-q_L}{1-p}\right)\right) = 0, && (65)\\
\frac{\partial \mathscr{L}^L}{\partial \lambda_L} =& q_L \ln\left(\frac{q_L}{p}\right) + (1-q_L)\ln\left(\frac{1-q_L}{1-p}\right) - K_L = 0. && (66)
\end{aligned}$$

The Lagrangian of the problem (6) is

$$\begin{aligned}
\mathscr{L}^0 =& q_0 V_1(W_1^H, K_H) + (1-q_0) V_1(W_1^L, K_L) \\
&+ \lambda_0\left(q_0 \ln\left(\frac{q_0}{p}\right) + (1-q_0)\ln\left(\frac{1-q_0}{1-p}\right) + q_0 K_H + (1-q_0) K_L - K\right), && (67)
\end{aligned}$$

where $W_1^H = W_0\left(\pi_0(H-R_f)+R_f\right)$ and $W_1^L = W_0\left(\pi_0(L-R_f)+R_f\right)$.

The FOCs with respect to q_0, K_H, and K_L are

$$\begin{aligned}
\frac{\partial \mathscr{L}^0}{\partial q_0} &= V_1(W_1^H, K_H) - V_1(W_1^L, K_L) + \lambda_0\left(\ln\left(\frac{q_0}{p}\right) - \ln\left(\frac{1-q_0}{1-p}\right) + K_H - K_L\right) = 0, && (68)\\
\frac{\partial \mathscr{L}^0}{\partial K_H} &= q_0 \frac{\partial V_1(W_1^H, K_H)}{\partial K_H} + q_0 \lambda_0 = 0, && (69)\\
\frac{\partial \mathscr{L}^0}{\partial K_L} &= (1-q_0) \frac{\partial V_1(W_1^L, K_L)}{\partial K_L} + (1-q_0)\lambda_0 = 0. && (70)
\end{aligned}$$

In (69), note that

$$\frac{\partial V_1(W_1^H, K_H)}{\partial K_H} = \frac{\partial \mathscr{L}^H}{\partial K_H} = -\lambda_H,$$

at the optimum. Therefore, (69) implies that $\lambda_H = \lambda_0$. Similarly, (70) implies $\lambda_L = \lambda_0$. Therefore $\lambda_0 = \lambda_L = \lambda_H$.

The FOC with respect to π_0 is

$$\frac{\partial \mathscr{L}^0}{\partial \pi_0} = q_0 \frac{\partial V_1(W_1^H, K_H)}{\partial \pi_0} + (1-q_0) \frac{\partial V_1(W_1^L, K_L)}{\partial \pi_0}, \tag{71}$$

where

$$\frac{\partial V_1(W_1^H, K_H)}{\partial \pi_0} = q_H \left\{ W_0 \left(\pi_0(H-R_f)+R_f\right)\left(\pi_H(H-R_f)+R_f\right)\right\}^{-\gamma} W_0(H-R_f)\left(\pi_H(H-R_f)+R_f\right)$$
$$+ (1-q_H)\left\{ W_0 \left(\pi_0(H-R_f)+R_f\right)\left(\pi_H(L-R_f)+R_f\right)\right\}^{-\gamma} W_0(H-R_f)\left(\pi_H(L-R_f)+R_f\right), \tag{72}$$

$$\frac{\partial V_1(W_1^L, K_L)}{\partial \pi_0} = q_L \left\{ W_0 \left(\pi_0(L-R_f)+R_f\right)\left(\pi_L(H-R_f)+R_f\right)\right\}^{-\gamma} W_0(L-R_f)\left(\pi_L(H-R_f)+R_f\right)$$
$$+ (1-q_L)\left\{ W_0 \left(\pi_0(L-R_f)+R_f\right)\left(\pi_L(L-R_f)+R_f\right)\right\}^{-\gamma} W_0(L-R_f)\left(\pi_L(L-R_f)+R_f\right). \tag{73}$$

The FOC with respect to λ_0 is

$$\frac{\partial \mathscr{L}^0}{\partial \lambda_0} = \ln\left(\frac{q_0}{p}\right) - \ln\left(\frac{1-q_0}{1-p}\right) + K_H - K_L - K = 0. \tag{74}$$

To determine $q_H, K_H, \pi_H, q_L, K_L, \pi_L, q_0, \pi_0$, we observe that

$$(61) \Rightarrow q_H \left\{\left(\pi_H(H-R_f)+R_f\right)\right\}^{-\gamma}(H-R_f)$$
$$+ (1-q_H)\left\{\left(\pi_H(L-R_f)+R_f\right)\right\}^{-\gamma}(L-R_f) = 0, \tag{75}$$

$$(64) \Rightarrow q_L \left\{\left(\pi_L(H-R_f)+R_f\right)\right\}^{-\gamma}(H-R_f)$$
$$+ (1-q_L)\left\{\left(\pi_L(L-R_f)+R_f\right)\right\}^{-\gamma}(L-R_f) = 0, \tag{76}$$

$$(68) \Rightarrow \frac{1}{1-\gamma}\Big(q_H \left\{W_1^H\left(\pi_H(H-R_f)+R_f\right)\right\}^{1-\gamma} + (1-q_H)\left\{W_1^H\left(\pi_H(L-R_f)+R_f\right)\right\}^{1-\gamma}\Big)$$
$$- \frac{1}{1-\gamma}\Big(q_L \left\{W_1^L\left(\pi_L(H-R_f)+R_f\right)\right\}^{1-\gamma} + (1-q_L)\left\{W_1^L\left(\pi_L(L-R_f)+R_f\right)\right\}^{1-\gamma}\Big)$$
$$+ \lambda_0\left(\ln\left(\frac{q_0}{p}\right) - \ln\left(\frac{1-q_0}{1-p}\right) + K_H - K_L\right) = 0, \tag{77}$$

$$(71), (72), (73) \Rightarrow q_0\Big(q_H \left\{ W_0 \left(\pi_0(H-R_f)+R_f\right)\left(\pi_H(H-R_f)+R_f\right)\right\}^{-\gamma} W_0(H-R_f)\left(\pi_H(H-R_f)+R_f\right)$$
$$+ (1-q_H)\left\{ W_0 \left(\pi_0(H-R_f)+R_f\right)\left(\pi_H(L-R_f)+R_f\right)\right\}^{-\gamma} W_0(H-R_f)\left(\pi_H(L-R_f)+R_f\right)\Big)$$
$$+ (1-q_0)\Big(q_L \left\{ W_0 \left(\pi_0(L-R_f)+R_f\right)\left(\pi_L(H-R_f)+R_f\right)\right\}^{-\gamma} W_0(L-R_f)\left(\pi_L(H-R_f)+R_f\right),$$
$$+ (1-q_L)\left\{ W_0 \left(\pi_0(L-R_f)+R_f\right)\left(\pi_L(L-R_f)+R_f\right)\right\}^{-\gamma} W_0(L-R_f)\left(\pi_L(L-R_f)+R_f\right)\Big)$$
$$= 0. \tag{78}$$

Combining the previous 4 equations with (62), (63), (65), (66), and (74), we have 9 unknowns and 9 equations.

From (63) and (66), we can solve K_H and K_L. From (75) and (76), we obtain

$$\pi_H = R_f \frac{(1-q_H)^{-\frac{1}{\gamma}}(R_f - L)^{-\frac{1}{\gamma}} - q_H^{-\frac{1}{\gamma}}(H - Rf)^{-\frac{1}{\gamma}}}{(1-q_H)^{-\frac{1}{\gamma}}(R_f - L)^{1-\frac{1}{\gamma}} + q_H^{-\frac{1}{\gamma}}(H - Rf)^{1-\frac{1}{\gamma}}}, \tag{79}$$

$$\pi_L = R_f \frac{(1-q_L)^{-\frac{1}{\gamma}}(R_f - L)^{-\frac{1}{\gamma}} - q_L^{-\frac{1}{\gamma}}(H - Rf)^{-\frac{1}{\gamma}}}{(1-q_L)^{-\frac{1}{\gamma}}(R_f - L)^{1-\frac{1}{\gamma}} + q_L^{-\frac{1}{\gamma}}(H - Rf)^{1-\frac{1}{\gamma}}}. \tag{80}$$

Therefore, we have 5 unknowns $q_H, q_L, q_0, \pi_0, \lambda_0$ and 5 equations.

Dividing (78) by $W_0^{1-\gamma}$, we can rewrite it to

$$\begin{aligned}
&q_0\Big(q_H\left\{\left(\pi_0(H-R_f)+R_f\right)\left(\pi_H(H-R_f)+R_f\right)\right\}^{-\gamma}(H-R_f)\left(\pi_H(H-R_f)+R_f\right)\\
&+(1-q_H)\left\{\left(\pi_0(H-R_f)+R_f\right)\left(\pi_H(L-R_f)+R_f\right)\right\}^{-\gamma}(H-R_f)\left(\pi_H(L-R_f)+R_f\right)\Big)\\
&+(1-q_0)\Big(q_L\left\{\left(\pi_0(L-R_f)+R_f\right)\left(\pi_L(H-R_f)+R_f\right)\right\}^{-\gamma}(L-R_f)\left(\pi_L(H-R_f)+R_f\right)\\
&+(1-q_L)\left\{\left(\pi_0(L-R_f)+R_f\right)\left(\pi_L(L-R_f)+R_f\right)\right\}^{-\gamma}(L-R_f)\left(\pi_L(L-R_f)+R_f\right)\Big)=0
\end{aligned} \tag{81}$$

The only W_0 that remains is in the W_1 terms in (62), (65), and (77). Therefore $\lambda_0 = \lambda_L = \lambda_H$ is proportional in $W_0^{1-\gamma}$. After factoring $W_0^{1-\gamma}$ from (62), (65), and (77), the 5 equations for q_H, q_L, q_0, π_0, and λ_0 are independent of W_0, but still state dependent. In summary, these 5 equations are

$$\begin{aligned}
&\frac{1}{1-\gamma}\Big(\left\{W_1^H\left(\pi_H(H-R_f)+R_f\right)\right\}^{1-\gamma}-\left\{W_1^H\left(\pi_H(L-R_f)+R_f\right)\right\}^{1-\gamma}\Big)\\
&+\lambda_0\left(\ln\left(\frac{q_H}{p}\right)-\ln\left(\frac{1-q_H}{1-p}\right)\right)=0,
\end{aligned} \tag{82}$$

$$\begin{aligned}
&\frac{1}{1-\gamma}\Big(\left\{W_1^L\left(\pi_L(H-R_f)+R_f\right)\right\}^{1-\gamma}-\left\{W_1^L\left(\pi_L(L-R_f)+R_f\right)\right\}^{1-\gamma}\Big)\\
&+\lambda_0\left(\ln\left(\frac{q_L}{p}\right)-\ln\left(\frac{1-q_L}{1-p}\right)\right)=0,
\end{aligned} \tag{83}$$

$$
\begin{aligned}
&\frac{1}{1-\gamma}\Big(q_H\left\{W_1^H\left(\pi_H(H-R_f)+R_f\right)\right\}^{1-\gamma}+(1-q_H)\left\{W_1^H\left(\pi_H(L-R_f)+R_f\right)\right\}^{1-\gamma}\Big)\\
&-\frac{1}{1-\gamma}\Big(q_L\left\{W_1^L\left(\pi_L(H-R_f)+R_f\right)\right\}^{1-\gamma}+(1-q_L)\left\{W_1^L\left(\pi_L(L-R_f)+R_f\right)\right\}^{1-\gamma}\Big)\\
&+\lambda_0\Big(\ln\left(\frac{q_0}{p}\right)-\ln\left(\frac{1-q_0}{1-p}\right)+q_H\ln\left(\frac{q_H}{p}\right)+(1-q_H)\ln\left(\frac{1-q_H}{1-p}\right)-q_L\ln\left(\frac{q_L}{p}\right)-(1-q_L)\ln\left(\frac{1-q_L}{1-p}\right)\Big)=0,
\end{aligned}
\tag{84}
$$

$$
\begin{aligned}
&q_0\Big(q_H\left\{\left(\pi_0(H-R_f)+R_f\right)\left(\pi_H(H-R_f)+R_f\right)\right\}^{-\gamma}(H-R_f)\left(\pi_H(H-R_f)+R_f\right)\\
&+(1-q_H)\left\{\left(\pi_0(H-R_f)+R_f\right)\left(\pi_H(L-R_f)+R_f\right)\right\}^{-\gamma}(H-R_f)\left(\pi_H(L-R_f)+R_f\right)\Big)\\
&+(1-q_0)\Big(q_L\left\{\left(\pi_0(L-R_f)+R_f\right)\left(\pi_L(H-R_f)+R_f\right)\right\}^{-\gamma}(L-R_f)\left(\pi_L(H-R_f)+R_f\right)\\
&+(1-q_L)\left\{\left(\pi_0(L-R_f)+R_f\right)\left(\pi_L(L-R_f)+R_f\right)\right\}^{-\gamma}(L-R_f)\left(\pi_L(L-R_f)+R_f\right)\Big)=0,
\end{aligned}
\tag{85}
$$

$$
\begin{aligned}
&q_0\ln\left(\frac{q_0}{p}\right)+(1-q_0)\ln\left(\frac{1-q_0}{1-p}\right)\\
&+q_0\left(q_H\ln\left(\frac{q_H}{p}\right)+(1-q_H)\ln\left(\frac{1-q_H}{1-p}\right)\right)\\
&+(1-q_0)\left(q_L\ln\left(\frac{q_L}{p}\right)+(1-q_L)\ln\left(\frac{1-q_L}{1-p}\right)\right)=K,
\end{aligned}
\tag{86}
$$

where $W_1^H = W_0(\pi_0(H-R_f)+R_f)$, $W_1^L = W_0(\pi_0(L-R_f)+R_f)$, and π_H, π_L are given in (79) and (80).

A.2 Proof of Proposition 1

First, we observe that

$$
\pi_H > \pi_L \iff q_H > q_L. \tag{87}
$$

If we replace both ‘ $>$′ in (87) by ′ $\geq$′, the previous equivalence also holds. To see this why (87) holds, observe that $q_H > q_L$ is equivalent to

$$
\left(\frac{\pi_H(H-R_f)+R_f}{\pi_H(L-R_f)+R_f}\right)^{-\gamma}=\frac{R_f-L}{H-R_f}\frac{1-q_H}{q_H}<\frac{R_f-L}{H-R_f}\frac{1-q_L}{q_L}=\left(\frac{\pi_L(H-R_f)+R_f}{\pi_L(L-R_f)+R_f}\right)^{-\gamma},
$$

where the left and right equalities follow from equations (61) and (64). Then the previous inequality is equivalent to

$$
\frac{\pi_H(H-R_f)+R_f}{\pi_H(L-R_f)+R_f}>\frac{\pi_L(H-R_f)+R_f}{\pi_L(L-R_f)+R_f},
$$

which is further equivalent to $\pi_H > \pi_L$.

To prove the statement of the proposition, we first show $q_H, q_L < p$. Suppose that $q_H \geq p$. Then $\lambda_H\left(\ln\frac{q_H}{p} - \ln\frac{1-q_H}{1-p}\right) \geq 0$. When $\gamma > 1$, the previous inequality and equation (62) imply $(\pi_H(H-R_f)+R_f)^{1-\gamma} \geq (\pi_H(L-R_f)+R_f)^{1-\gamma}$, which implies $\frac{R_f-L}{H-R_f}\frac{1-q_H}{q_H} \geq 1$ thanks to equation (61). However, the previous inequality implies $Hq_H + L(1-q_H) \leq R_f$, then $R_f \geq L + (H-L)q_H \geq L + (H-L)p = Hp + L(1-p)$, where the second inequality follows from the assumption $q_H \geq p$. However, the previous inequality contradicts with the assumption $Hp + L(1-p) > R_f$. Hence we conclude $q_H < p$. When $0 < \gamma < 1$, the inequality $\lambda_0\left(\ln\frac{q_H}{p} - \ln\frac{1-q_H}{1-p}\right) \geq 0$ implies $(\pi_H(H-R_f) + R_f)^{1-\gamma} \leq (\pi_H(L-R_f)+R_f)^{1-\gamma}$ thanks to equation (62) and $0 < \gamma < 1$. The previous inequality also implies $\frac{R_f-L}{H-R_f}\frac{1-q_H}{q_H} \geq 1$ from equation (61). Then the remaining proof is the same as the $\gamma > 1$ case. In both cases, we have shown that $q_H < p$. The proof of $q_L < p$ is similar.

Next, we prove $\pi_H, \pi_L > 0$. Because $q_H < p$, equation (62) implies $\pi_H(H-R_f) + R_f > \pi_H(L-R_f) + R_f$ in both $\gamma > 1$ and $0 < \gamma < 1$ cases. Then the previous inequality implies $\pi_H > 0$.

Let us now prove $q_L < q_H$ when $\gamma > 1$. Suppose that $q_L \geq q_H$, then (87) implies $\pi_L \geq \pi_H$. Then the previous inequality, $\gamma > 1$, and $\pi_H, \pi_L > 0$ yield

$$\begin{aligned} 0 >& \left(\pi_H(H-R_f)+R_f\right)^{1-\gamma} - \left(\pi_H(L-R_f)+R_f\right)^{1-\gamma} \\ \geq& \left(\pi_L(H-R_f)+R_f\right)^{1-\gamma} - \left(\pi_L(L-R_f)+R_f\right)^{1-\gamma}. \end{aligned} \tag{88}$$

Due to the positive risk premium, $\pi_0 > 0$, therefore $W_1^H > W_1^L$. As a result, the second inequality in (88) and $\gamma > 1$ imply

$$\begin{aligned} &\frac{1}{\gamma-1}(W_1^H)^{1-\gamma}\left[\left(\pi_H(H-R_f)+R_f\right)^{1-\gamma} - \left(\pi_H(L-R_f)+R_f\right)^{1-\gamma}\right] \\ >& \frac{1}{\gamma-1}(W_1^L)^{1-\gamma}\left[\left(\pi_L(H-R_f)+R_f\right)^{1-\gamma} - \left(\pi_L(L-R_f)+R_f\right)^{1-\gamma}\right]. \end{aligned} \tag{89}$$

The previous inequality, $\lambda_0 > 0$, and equations (62), (65) combined imply

$$\ln\frac{q_H}{p} - \ln\frac{1-q_H}{1-p} > \ln\frac{q_L}{p} - \ln\frac{1-q_L}{1-p}. \tag{90}$$

Recall that both q_L and q_H are less than p. The previous inequality implies $q_L < q_H$ and contradicts with the assumption $q_L \geq q_H$.

When $0 < \gamma < 1$, suppose that $q_L \leq q_H$, then (87) implies $\pi_L \leq \pi_H$. Then the previous inequality, $0 < \gamma < 1$, and $\pi_H, \pi_L > 0$ yield

$$\begin{aligned} 0 <& \left(\pi_L(H-R_f)+R_f\right)^{1-\gamma} - \left(\pi_L(L-R_f)+R_f\right)^{1-\gamma} \\ \leq& \left(\pi_H(H-R_f)+R_f\right)^{1-\gamma} - \left(\pi_H(L-R_f)+R_f\right)^{1-\gamma}. \end{aligned} \tag{91}$$

Because $W_1^H > W_1^L$ and $0 < \gamma < 1$, the inequalities in (89) and (90) reversed. Because $q_H, q_L < p$, then $q_H < q_L$ which contradicts with the assumption $q_L \le q_H$. Therefore, $q_H < q_L < p$ when $0 < \gamma < 1$.

A.3 Proof of Lemma 2

The Lagrangian associated to problem (3) with the constraints (8) and (9) is

$$\begin{aligned}
\mathscr{L} =& \frac{1}{1-\gamma}\Big(q_0 q_H \left\{W_0\left(\pi_0(H-R_f)+R_f\right)\left(\pi_H(H-R_f)+R_f\right)\right\}^{1-\gamma} \\
&+ q_0(1-q_H)\left\{W_0\left(\pi_0(H-R_f)+R_f\right)\left(\pi_H(L-R_f)+R_f\right)\right\}^{1-\gamma} \\
&+ (1-q_0)q_L\left\{W_0\left(\pi_0(L-R_f)+R_f\right)\left(\pi_L(H-R_f)+R_f\right)\right\}^{1-\gamma} \\
&+ (1-q_0)(1-q_L)\left\{W_0\left(\pi_0(L-R_f)+R_f\right)\left(\pi_L(L-R_f)+R_f\right)\right\}^{1-\gamma}\Big) \\
&+ \lambda_H\Big(q_H \ln\left(\frac{q_H}{p}\right) + (1-q_H)\ln\left(\frac{1-q_H}{1-p}\right) + q_0\ln\left(\frac{q_0}{p}\right) + (1-q_0)\ln\left(\frac{1-q_0}{1-p}\right) - K\Big) \\
&+ \lambda_L\Big(q_L \ln\left(\frac{q_L}{p}\right) + (1-q_L)\ln\left(\frac{1-q_L}{1-p}\right) + q_0\ln\left(\frac{q_0}{p}\right) + (1-q_0)\ln\left(\frac{1-q_0}{1-p}\right) - K\Big).
\end{aligned} \tag{92}$$

The FOCs with respect to q_0, q_H, q_L are

$$\begin{aligned}
\frac{\partial \mathscr{L}}{\partial q_0} =& \frac{1}{1-\gamma}\Big(q_H\left\{W_0\left(\pi_0(H-R_f)+R_f\right)\left(\pi_H(H-R_f)+R_f\right)\right\}^{1-\gamma} \\
&+ (1-q_H)\left\{W_0\left(\pi_0(H-R_f)+R_f\right)\left(\pi_H(L-R_f)+R_f\right)\right\}^{1-\gamma} \\
&- q_L\left\{W_0\left(\pi_0(L-R_f)+R_f\right)\left(\pi_L(H-R_f)+R_f\right)\right\}^{1-\gamma} \\
&- (1-q_L)\left\{W_0\left(\pi_0(L-R_f)+R_f\right)\left(\pi_L(L-R_f)+R_f\right)\right\}^{1-\gamma}\Big) \\
&+ (\lambda_H+\lambda_L)\Big(1+\ln\left(\frac{q_0}{p}\right) - 1 - \ln\left(\frac{1-q_0}{1-p}\right)\Big) = 0,
\end{aligned} \tag{93}$$

$$\begin{aligned}
\frac{\partial \mathscr{L}}{\partial q_H} =& \frac{1}{1-\gamma}\Big(q_0\left\{W_0\left(\pi_0(H-R_f)+R_f\right)\left(\pi_H(H-R_f)+R_f\right)\right\}^{1-\gamma} \\
&- q_0\left\{W_0\left(\pi_0(H-R_f)+R_f\right)\left(\pi_H(L-R_f)+R_f\right)\right\}^{1-\gamma}\Big) \\
&+ \lambda_H\Big(1+\ln\left(\frac{q_H}{p}\right) - 1 - \ln\left(\frac{1-q_H}{1-p}\right)\Big) = 0, \quad \text{and}
\end{aligned} \tag{94}$$

$$\begin{aligned}
\frac{\partial \mathscr{L}}{\partial q_L} =& \frac{1}{1-\gamma}\Big((1-q_0)\left\{W_0\left(\pi_0(L-R_f)+R_f\right)\left(\pi_L(H-R_f)+R_f\right)\right\}^{1-\gamma} \\
&- (1-q_0)\left\{W_0\left(\pi_0(L-R_f)+R_f\right)\left(\pi_L(L-R_f)+R_f\right)\right\}^{1-\gamma}\Big) \\
&+ \lambda_L\Big(1+\ln\left(\frac{q_L}{p}\right) - 1 - \ln\left(\frac{1-q_L}{1-p}\right)\Big) = 0.
\end{aligned} \tag{95}$$

The FOCs with respect to π_0, π_H, and π_L are

$$\begin{aligned}\frac{\partial \mathscr{L}}{\partial \pi_0} =& \Big(q_0 q_H \left\{ W_0 \left(\pi_0(H-R_f)+R_f\right)\left(\pi_H(H-R_f)+R_f\right)\right\}^{-\gamma} W_0(H-R_f)\left(\pi_H(H-R_f)+R_f\right) \\ &+ q_0(1-q_H) \left\{ W_0 \left(\pi_0(H-R_f)+R_f\right)\left(\pi_H(L-R_f)+R_f\right)\right\}^{-\gamma} W_0(H-R_f)\left(\pi_H(L-R_f)+R_f\right) \\ &+ (1-q_0) q_L \left\{ W_0 \left(\pi_0(L-R_f)+R_f\right)\left(\pi_L(H-R_f)+R_f\right)\right\}^{-\gamma} W_0(L-R_f)\left(\pi_L(H-R_f)+R_f\right) \\ &+ (1-q_0)(1-q_L) \left\{ W_0 \left(\pi_0(L-R_f)+R_f\right)\left(\pi_L(L-R_f)+R_f\right)\right\}^{-\gamma} W_0(L-R_f)\left(\pi_L(L-R_f)+R_f\right)\Big) = 0,\end{aligned} \tag{96}$$

$$\begin{aligned}\frac{\partial \mathscr{L}}{\partial \pi_H} =& \Big(q_0 q_H \left\{ W_0 \left(\pi_0(H-R_f)+R_f\right)\left(\pi_H(H-R_f)+R_f\right)\right\}^{-\gamma} W_0 \left(\pi_0(H-R_f)+R_f\right)(H-R_f) \\ &+ q_0(1-q_H) \left\{ W_0 \left(\pi_0(H-R_f)+R_f\right)\left(\pi_H(L-R_f)+R_f\right)\right\}^{-\gamma} W_0 \left(\pi_0(H-R_f)+R_f\right)(L-R_f) = 0,\end{aligned} \tag{97}$$

$$\begin{aligned}\frac{\partial \mathscr{L}}{\partial \pi_L} =& (1-q_0) q_L \left\{ W_0 \left(\pi_0(L-R_f)+R_f\right)\left(\pi_L(H-R_f)+R_f\right)\right\}^{-\gamma} W_0 \left(\pi_0(L-R_f)+R_f\right)(H-R_f) \\ &+ (1-q_0)(1-q_L) \left\{ W_0 \left(\pi_0(L-R_f)+R_f\right)\left(\pi_L(L-R_f)+R_f\right)\right\}^{-\gamma} W_0 \left(\pi_0(L-R_f)+R_f\right)(L-R_f) = 0.\end{aligned} \tag{98}$$

The two constraints (8) and (9) imply that $q_H = q_L$ at optimum. We can conclude from (97) and (98) that $\pi_H = \pi_L$.

These insights, together with (94) and (95), lead to the conclusion that λ_H and λ_L are proportional to $W_0^{1-\gamma}$. We introduce the notation q_1 for both q_H and q_L and π_1 for π_H and π_L, then we can rewrite (96) as

$$\begin{aligned}&\Big(q_0 q_1 \left\{ \left(\pi_0(H-R_f)+R_f\right)\left(\pi_1(H-R_f)+R_f\right)\right\}^{-\gamma} (H-R_f)\left(\pi_1(H-R_f)+R_f\right) \\ &+ q_0(1-q_1) \left\{ \left(\pi_0(H-R_f)+R_f\right)\left(\pi_1(L-R_f)+R_f\right)\right\}^{-\gamma} (H-R_f)\left(\pi_1(L-R_f)+R_f\right) \\ &+ (1-q_0) q_1 \left\{ \left(\pi_0(L-R_f)+R_f\right)\left(\pi_1(H-R_f)+R_f\right)\right\}^{-\gamma} (L-R_f)\left(\pi_1(H-R_f)+R_f\right) \\ &+ (1-q_0)(1-q_1) \left\{ \left(\pi_0(L-R_f)+R_f\right)\left(\pi_1(L-R_f)+R_f\right)\right\}^{-\gamma} (L-R_f)\left(\pi_1(L-R_f)+R_f\right)\Big) = 0\end{aligned} \tag{99}$$

If we plug in (97) – coinciding with the first and third term – and (98) – coinciding with the second and fourth term – then (99) is satisfied if $\pi_0 = \pi_1$.

Finally, multiplying $\frac{q_1}{q_0}$ on both sides of (94) and $\frac{1-q_1}{1-q_0}$ on both sides of (95), we obtain

$$\frac{1}{1-\gamma}\Big(q_1\left\{W_0\left(\pi_0(H-R_f)+R_f\right)\left(\pi_H(H-R_f)+R_f\right)\right\}^{1-\gamma}$$
$$-q_1\left\{W_0\left(\pi_0(H-R_f)+R_f\right)\left(\pi_H(L-R_f)+R_f\right)\right\}^{1-\gamma}\Big)$$
$$+\lambda_H\frac{q_1}{q_0}\left(\ln\left(\frac{q_H}{p}\right)-\ln\left(\frac{1-q_H}{1-p}\right)\right)=0,$$
$$\frac{1}{1-\gamma}\Big((1-q_1)\left\{W_0\left(\pi_0(L-R_f)+R_f\right)\left(\pi_L(H-R_f)+R_f\right)\right\}^{1-\gamma}$$
$$-(1-q_1)\left\{W_0\left(\pi_0(L-R_f)+R_f\right)\left(\pi_L(L-R_f)+R_f\right)\right\}^{1-\gamma}\Big)$$
$$+\lambda_L\frac{1-q_1}{1-q_0}\left(\ln\left(\frac{q_L}{p}\right)-\ln\left(\frac{1-q_L}{1-p}\right)\right)=0.$$

Now combining the previous two equations with (93), we obtain

$$-\left(\lambda_H\frac{q_1}{q_0}+\lambda_L\frac{1-q_1}{1-q_0}\right)\left(\ln\left(\frac{q_1}{p}\right)-\ln\left(\frac{1-q_1}{1-p}\right)\right)+(\lambda_H+\lambda_L)\left(\ln\left(\frac{q_0}{p}\right)-\ln\left(\frac{1-q_0}{1-p}\right)\right)=0,$$

which implies $q_0=q_1$. This, combined with (8) and (9), implies that $K_1=K/2$.

A.4 Proof of Proposition 2

We present the proof for the $\gamma>1$ case. The proof for the $0<\gamma<1$ case is similar. Plugging the decomposition (17) into (16) and cancelling $W^{1-\gamma}/(1-\gamma)$ on both sides, we obtain the equation (18). Because $1-\gamma<0$, $\sup_{\pi,\tilde{c}}\inf_{u,g}$ in (16) is transformed into $\inf_{\pi,\tilde{c}}\sup_{u,g}$ in (18) when $W^{1-\gamma}/(1-\gamma)$ is cancelled on both sides of (16).

For the left boundary condition of $\mathcal{V}$, as we discuss in the main text, both u^* and g^* are zero when $E=0$. In this case, the problem (13) is reduced to a standard Merton's problem without model ambiguity. The equation (18) is transformed into

$$\delta\mathcal{V}=\inf_{\pi,\tilde{c}}\Big\{\delta\tilde{c}^{1-\gamma}+(1-\gamma)[r+\pi(\mu-r)-\tilde{c}]\mathcal{V}+\frac{1}{2}(-\gamma)(1-\gamma)\pi^2\sigma^2\mathcal{V}\Big\},$$

where the optimal π^* is $\pi^*=\frac{\mu-r}{\gamma\sigma^2}$ and $\tilde{c}^*$ satisfies (22). The previous equation admits an explicit solution given in (19).

When $E=\overline{E}$, u^* reaches its maximum value $\overline{u}=\frac{\mu-r}{\sigma}$ and g^* is chosen to be zero so that E stays lower than $\overline{E}$. In this case, (18) is reduced to

$$\delta\mathcal{V}=\inf_{\tilde{c}}\Big\{\delta\tilde{c}^{1-\gamma}+(1-\gamma)[r-\tilde{c}]\mathcal{V}\Big\}, \tag{100}$$

which has an explicit solution (20).

When $\partial_E \mathcal{V} > 0$, the first order condition of u in (18) yields (23). When $\partial^2_{EE}\mathcal{V} < 0$, the first order condition of g in (18) implies (24). Plugging (23) and (24) into (18), the optimization of π in (18) is

$$(1-\gamma)(\mu-r)\mathcal{V}\pi + \frac{(1-\gamma)^2\sigma^2}{2\delta}\frac{\mathcal{V}^2}{\partial_E\mathcal{V}}\pi^2 - \frac{(1-\gamma)^2\sigma^2}{2}\frac{(\partial_E\mathcal{V})^2}{\partial^2_{EE}\mathcal{V}}\pi^2 - \frac{\gamma(1-\gamma)}{2}\sigma^2\mathcal{V}\pi^2.$$

When the second-order condition

$$\frac{(1-\gamma)^2}{\delta}\frac{\mathcal{V}^2}{\partial_E\mathcal{V}} - (1-\gamma)^2\frac{(\partial_E\mathcal{V})^2}{\partial^2_{EE}\mathcal{V}} - \gamma(1-\gamma)\mathcal{V} > 0 \tag{101}$$

is satisfied, then the optimal investment strategy π^* is given in (21). Note that the previous second-order condition holds thanks to $\gamma > 1$, $\mathcal{V} > 0$, and the assumption $\partial_E\mathcal{V} > 0$ and $\partial^2_{EE}\mathcal{V} < 0$.

A.5 Derivation of equation (29)

For the optimized value $E^* = E^*(W,\theta)$, we have from (28) that

$$\begin{aligned} V(W,E^*) =& \widehat{V}(W,\theta) - E^*\theta, \\ \partial_E V(W,E^*) =& -\theta, \quad \partial_W V(W,E^*) = \partial_W \widehat{V}(W,\theta), \\ \partial^2_{WW}\widehat{V}(W,\theta) =& \partial^2_{WW}V(W,E^*) + \partial^2_{WE}V(W,E^*)\partial_W E^* = \partial^2_{WW}V(W,E^*) - \frac{\left(\partial^2_{WE}V(W,E^*)\right)^2}{\partial^2_{EE}V(W,E^*)}. \end{aligned} \tag{102}$$

Meanwhile, note that

$$\begin{aligned} \inf_g\left\{\tfrac{1}{2}W^2\pi^2\sigma^2\partial^2_{WW}V + \tfrac{1}{2}g^2\partial^2_{EE}V + W\pi\sigma g\partial^2_{EW}V\right\} =& \tfrac{1}{2}W^2\pi^2\sigma^2\left\{\partial^2_{WW}V - \frac{(\partial^2_{WE}V)^2}{\partial^2_{EE}V}\right\} \\ =& \tfrac{1}{2}W^2\pi^2\sigma^2\partial^2_{WW}\widehat{V}, \end{aligned} \tag{103}$$

where the minimizer is

$$g = -W\pi\sigma\frac{\partial^2_{WE}V}{\partial^2_{EE}V}. \tag{104}$$

Plugging (102) and (103) into the HJB equation (16) where $E = E^*$, we obtain (29).

A.6 Proof of Proposition 4

It follows from dynamic programming that the value function $\widetilde{V}$ in (37) satisfies the following HJB equation

$$\delta\widetilde{V} = \sup_{\pi,\tilde{c}}\inf_{u}\Big\{\delta U(W\tilde{c}) + W\partial_W\widetilde{V}\big[r + \pi(\mu - r - \sigma u) - \tilde{c}\big] + \partial_E\widetilde{V}\big(\delta E - \tfrac{\delta}{2}|u|^2\big) + \tfrac{1}{2}W^2\pi^2\sigma^2\partial^2_{WW}\widetilde{V}\Big\}, \tag{105}$$

where $\tilde{c} = c/W$. Similar to (17), $\widetilde{V}$ admits the following homothetic decomposition

$$\widetilde{V}(E,W) = \frac{W^{1-\gamma}}{1-\gamma}\widetilde{\mathcal{V}}(E). \tag{106}$$

Plugging this decomposition into (105), we obtain the following equation for $\widetilde{\mathcal{V}}$:

$$\delta\widetilde{\mathcal{V}} = \inf_{\pi,\tilde{c}}\sup_{u}\Big\{\delta\tilde{c}^{1-\gamma} + (1-\gamma)\widetilde{\mathcal{V}}\big[r + \pi(\mu - r - \sigma u) - \tilde{c}\big] + \partial_E\widetilde{\mathcal{V}}\big(\delta E - \tfrac{\delta}{2}|u|^2\big) - \tfrac{1}{2}\gamma(1-\gamma)\widetilde{\mathcal{V}}\pi^2\sigma^2\Big\}, \tag{107}$$

when $\gamma > 1$, and $\inf_{\pi,\tilde{c}}\sup_u$ is replaced by $\sup_{\pi,\tilde{c}}\inf_u$ when $0 < \gamma < 1$.

The first-order condition of u in (107) yields

$$\tilde{u}^* = -\frac{(1-\gamma)\sigma\widetilde{\mathcal{V}}}{\delta\partial_E\widetilde{\mathcal{V}}}\pi^*. \tag{108}$$

The first-order condition in π and the previous equation combined yield

$$\tilde{\pi}^* = \frac{\mu - r}{\sigma^2\big(\gamma - \frac{(1-\gamma)\widetilde{\mathcal{V}}}{\delta\partial_E\widetilde{\mathcal{V}}}\big)}. \tag{109}$$

The first-order condition in $\tilde{c}$ yields

$$\tilde{c}^* = \delta^{\frac{1}{\gamma}}\widetilde{\mathcal{V}}^{-\frac{1}{\gamma}}. \tag{110}$$

Plugging these three first-order conditions into (107), we obtain

$$\tilde{c}^* = \frac{1}{\gamma}\Big[\delta - (1-\gamma)r - \frac{1-\gamma}{2\big(\gamma - \frac{(1-\gamma)\widetilde{\mathcal{V}}}{\delta\partial_E\widetilde{\mathcal{V}}}\big)}\frac{(\mu - r)^2}{\sigma^2} - E\frac{\delta\partial_E\widetilde{\mathcal{V}}}{\widetilde{\mathcal{V}}}\Big]. \tag{111}$$

For a stationary solution, E remains a constant. Therefore $\widetilde{\mathcal{V}}$ and $\partial_E\widetilde{\mathcal{V}}$ only take single value $\widetilde{\mathcal{V}}(E)$ and $\partial_E\widetilde{\mathcal{V}}(E)$, respectively. Introduce

$$\theta(E) = -\frac{(1-\gamma)\widetilde{\mathcal{V}}}{\delta\partial_E\widetilde{\mathcal{V}}}. \tag{112}$$

We obtain (43) and (44) from (108) and (109). Combining (39) and (43), we obtain (46). Finally (45) follows from (46) and (111).

The proof for the problem (40) is similar. We sketch it below. It follows from dynamic programming that $\check{V}$ in (40) satisfies the following HJB equation

$$\delta\check{V} = \sup_{\lambda,\pi,\check{c}}\inf_{u}\Big\{\delta U(W\check{c}) + W\partial_W\check{V}\big[r+\pi(\mu-r-\sigma u)-\check{c}\big] + \tfrac{1}{2}W^2\pi^2\sigma^2\partial^2_{WW}\check{V} + \lambda\delta\big(\tfrac{1}{2}|u|^2-\kappa\big)\Big\}. \tag{113}$$

Define $\check{\mathcal{V}}$ via $\check{V} = \frac{W^{1-\gamma}}{1-\gamma}\check{\mathcal{V}}$. The first order conditions of π and u in (113) yield

$$\check{\pi}^* = \frac{\mu-r}{\sigma^2}\frac{1}{\gamma+\frac{W^{1-\gamma}\check{\mathcal{V}}}{\lambda^*\delta}}, \tag{114}$$

$$\check{u}^* = \frac{\mu-r}{\sigma}\frac{\frac{W^{1-\gamma}\check{\mathcal{V}}}{\lambda^*\delta}}{\gamma+\frac{W^{1-\gamma}\check{\mathcal{V}}}{\lambda^*\delta}}. \tag{115}$$

Because we consider $\kappa < \frac{1}{2}\big(\frac{\mu-r}{\sigma}\big)^2$, the constraint (41) is binding. Otherwise, $\lambda^* = 0$ and $\check{u}^* = \frac{\mu-r}{\sigma}$ from (115), which violates the constraint $\frac{1}{2}|\check{u}^*| \leq \kappa$.

Denote $\theta = \frac{W^{1-\gamma}\check{\mathcal{V}}}{\lambda^*\delta}$. Combining (115) and $\frac{1}{2}|\check{u}^*|^2 = \kappa$, we obtain $\theta = \theta(\kappa)$ which satisfies (47). This implies

$$\lambda^* = \frac{W^{1-\gamma}\check{\mathcal{V}}}{\delta\theta(\kappa)},$$

hence λ^* and also $\check{V}$ are homothetic in W. Finally, plugging (114) and (115) back into the HJB equation (113), we obtain

$$\check{c}^* = \frac{1}{\gamma}\Big[\delta-(1-\gamma)r-\frac{(1-\gamma)\gamma}{2(\gamma+\theta(\kappa))^2}\frac{(\mu-r)^2}{\sigma^2}\Big].$$

A.7 Non-Stationary solution with model ambiguity

Consider the problem (37) subject to the budget constraint (14) and the dynamics of continuation entropy E in (38). Here we do not restrict the drift of E to be zero. In this case, the function $\widetilde{\mathcal{V}}$, defined via (106), satisfies the HJB equation (107) with the same boundary conditions (19) and (20). The first order conditions in u, π, and $\tilde{c}$ from (107) are (108), (109), and (110). Plugging these first order conditions into (107), we obtain

$$\delta = \gamma\delta^{\frac{1}{\gamma}}\widetilde{\mathcal{V}}^{-\frac{1}{\gamma}} + r(1-\gamma) + E\frac{\delta\partial_E\widetilde{\mathcal{V}}}{\widetilde{\mathcal{V}}} + \frac{1}{2}(1-\gamma)\frac{(\mu-r)^2}{\sigma^2}\frac{1}{\gamma-\frac{(1-\gamma)\widetilde{\mathcal{V}}}{\delta\partial_E\widetilde{\mathcal{V}}}}. \tag{116}$$

To simplify notation, let us introduce

$$\begin{aligned} A &= \frac{\delta \partial_E \widetilde{V}}{\widetilde{V}} \\ c &= \delta - r(1-\gamma) - \gamma \delta^{\frac{1}{\gamma}} \widetilde{V}^{-\frac{1}{\gamma}} \\ b &= (1-\gamma)E - \frac{1}{2}(1-\gamma)\frac{(\mu-r)^2}{\sigma^2} + \gamma c. \end{aligned}$$

Then equation (116) is transformed to

$$c = AE + \frac{1}{2}(1-\gamma)\frac{(\mu-r)^2}{\sigma^2}\frac{A}{\gamma A - (1-\gamma)}. \tag{117}$$

For given value of $\widetilde{V}$, the previous equation is quadratic in A

$$\gamma E A^2 - bA + (1-\gamma)c = 0. \tag{118}$$

When $\gamma > 1$, we have $\partial_E \widetilde{V} > 0$, hence we look for a positive solution of the previous quadratic equation.

$$A = \frac{b \pm \sqrt{b^2 - 4E\gamma(1-\gamma)c}}{2\gamma E}. \tag{119}$$

Using (119) and the boundary conditions (19) and (20), we can iteratively solve $\widetilde{V}$ and $\partial_E \widetilde{V}$ for $E \in (0, \overline{E})$. Using the same parameters as in Figure 4, a numeric example of the non-stationary solution is shown in Figure 11. The lower right panel shows that the drift of entropy is negative for small values of E and becomes positive when E crosses a threshold E^* (around 0.017). Therefore, if E starts from $E_0 < E^*$, the negative drift of E leads the entropy to zero and the investor eventually becomes a Merton investor; if E starts from $E_0 > E^*$, the positive drift of E leads the entropy to $\overline{E}$ and the investor eventually stops holding the risky asset; only when $E_0 = E^*$, is the drift of E zero, and does E remain constant. In this case, the non-stationary solution agrees with the stationary solution. The consumption-wealth ratios are similar between the non-stationary solution and the stationary solution, even though the ratio in the non-stationary case is marginally smaller.

A.8 Detection-error probabilities

A central idea in Anderson et al. (2003) is that to discipline the amount of entropy to a reasonable quantity, the agent is only willing to consider alternative models that are statistically difficult to distinguish from the baseline model. To quantify the difficulty of distinguishing alternative models from the baseline model, Anderson et al. (2003) introduce the detection-error probability

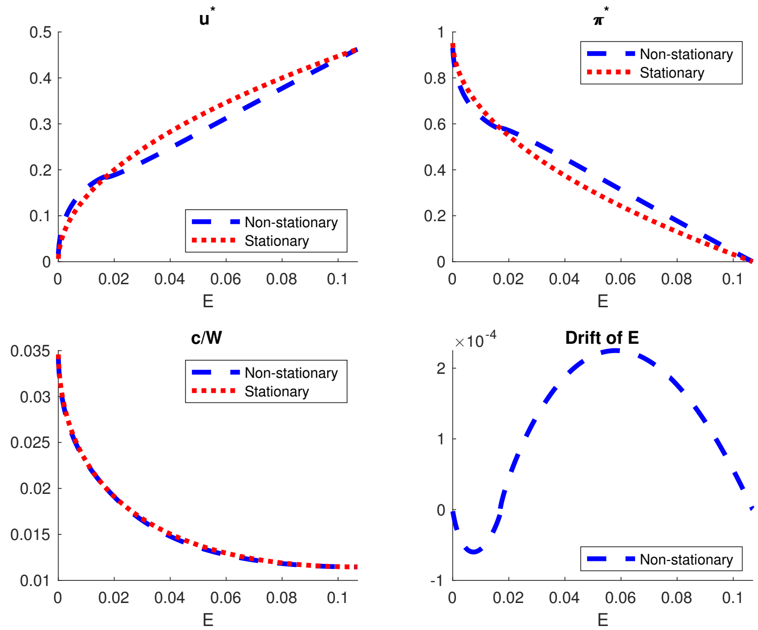

Figure 11:

Notes: This figure plots the worst-case belief distortion u^*, the optimal investment strategy π^*, the optimal consumption wealth ratio c/W, and the drift of entropy for both the non-stationary and the stationary solutions of the ambiguity-averse problem. The parameters are $\delta = 0.02$, $\gamma = 3$, $\mu - r = 0.0739$, $r = 0.0072$, and $\sigma = 0.1598$.

for a given time horizon T (see also Maenhout (2006)):

$$\psi(E) = \tfrac{1}{2}\mathbb{P}^{\mathbb{B}}(\log Z_T > 0) + \tfrac{1}{2}\mathbb{P}^{\mathbb{U}}(\log Z_T < 0), \tag{120}$$

where Z_T is defined in (10) with $u_t = u^*(E_t)$ therein. The quantity $\mathbb{P}^{\mathbb{B}}(\log Z_T > 0)$ is the probability that the agent mistakenly identifies the alternative model $\mathbb{U}$ rather than the baseline model $\mathbb{B}$ and the quantity $\mathbb{P}^{\mathbb{U}}(\log Z_T < 0)$ is the probability of misidentifying $\mathbb{B}$ rather than $\mathbb{U}$.[12] Both these two probabilities contribute equal weight in the definition (120). Anderson et al. (2003) propose that a reasonable initial value of the entropy budget at time zero is the value of E whose associated detection-error probability $\psi(E)$ is less than 10%.

To identify the detection-error probability $\psi(E)$, we define

$$\psi_{\mathbb{B}}(t,E,Z) = \mathbb{P}^{\mathbb{B}}(\log Z_T > 0 \,|\, E_t = E, Z_t = Z) \quad \text{and} \quad \psi_{\mathbb{U}}(t,E,Z) = \mathbb{P}^{\mathbb{U}}(\log Z_T < 0 \,|\, E_t = E, Z_t = Z). \tag{121}$$

Then the detection-error probability in (120) is identified as $\psi(E) = \frac{1}{2}\psi_{\mathbb{B}}(0,E,1) + \frac{1}{2}\psi_{\mathbb{U}}(0,E,1)$. We characterize $\psi_{\mathbb{B}}$ and $\psi_{\mathbb{U}}$ as solutions to two partial differential equations with two spatial

[12]When $Z \equiv 1$, i.e., $\mathbb{B}$ and $\mathbb{U}$ are the same, we set both probabilities to 0.5 so that distinguishing $\mathbb{B}$ and $\mathbb{U}$ is a random guess.

variables in Proposition 6 later in this section and numerically solve them to obtain the detection-error probability and the range of reasonable values for the initial entropy budget.

In the case of model ambiguity only, i.e., when $g = 0$, $u = \sqrt{2E}$, and the detection-error probability has the closed-form expression

$$\psi^{g=0}(0,E,0) = \frac{1}{2}\psi_{\mathbb{B}}^{g=0}(0,E,0) + \frac{1}{2}\psi_{\mathbb{U}}^{g=0}(0,E,0) = \frac{1}{2}\Phi\Big(-\frac{1}{2}u\sqrt{T}\Big) + \frac{1}{2}\Big(1-\Phi\Big(\frac{1}{2}u\sqrt{T}\Big)\Big). \tag{122}$$

To identify the detection-error probability $\psi(E)$, we introduce $x_t = \log Z_t$. It follows the dynamics

$$dx_t = -\tfrac{1}{2}|u_t^*|^2 dt - u_t^* dB_t^{\mathbb{B}} = \frac{1}{2}|u_t^*|^2 dt - u_t^* dB_t^{\mathbb{U}}, \tag{123}$$

where u^* is given in (23). The detection-error probabilities in (121) are

$$\psi_{\mathbb{B}}(t,E,x) = \mathbb{P}^{\mathbb{B}}(x_T > 0 \,|\, E_t = E, x_t = x) \quad \text{and} \quad \psi_{\mathbb{U}}(t,E,x) = \mathbb{P}^{\mathbb{U}}(x_T < 0 \,|\, E_t = E, x_t = x).$$

The following result characterizes $\psi_{\mathbb{B}}$ and $\psi_{\mathbb{U}}$ via two partial differential equations with two spatial variables.

Proposition 6. *$\psi_{\mathbb{B}}$ is the solution of the PDE*

$$\partial_t\psi_{\mathbb{B}} + \tfrac{1}{2}g^2\partial_{EE}^2\psi_{\mathbb{B}} + \tfrac{1}{2}u^2\partial_{xx}^2\psi_{\mathbb{B}} - ug\partial_{xE}^2\psi_{\mathbb{B}} + \big(\delta E + gu - \tfrac{\delta}{2}u^2\big)\partial_E\psi_{\mathbb{B}} - \tfrac{1}{2}u^2\partial_x\psi_{\mathbb{B}} = 0, \tag{124}$$

with the terminal condition

$$\psi_{\mathbb{B}}(T,E,x) = 1_{\{x>0\}} + 0.5 \times 1_{\{x=0\}}. \tag{125}$$

and boundary conditions

$$\psi_{\mathbb{B}}(t,E,-\infty) = 0, \tag{126}$$

$$\psi_{\mathbb{B}}(t,E,\infty) = 1, \tag{127}$$

$$\psi_{\mathbb{B}}(t,0,x) = 1_{\{x>0\}} + 0.5 \times 1_{\{x=0\}}, \tag{128}$$

$$\psi_{\mathbb{B}}(t,\overline{E},x) = \Phi\Bigg(\frac{x - \frac{(\mu-r)^2}{2\sigma^2}(T-t)}{\frac{\mu-r}{\sigma}\sqrt{T-t}}\Bigg). \tag{129}$$

$\psi_{\mathbb{U}}$ is the solution of the PDE

$$\partial_t\psi_{\mathbb{U}} + \tfrac{1}{2}g^2\partial_{EE}^2\psi_{\mathbb{U}} + \tfrac{1}{2}u^2\partial_{xx}^2\psi_{\mathbb{U}} - ug\partial_{xE}^2\psi_{\mathbb{U}} + \big(\delta E - \tfrac{\delta}{2}u^2\big)\partial_E\psi_{\mathbb{U}} + \tfrac{1}{2}u^2\partial_x\psi_{\mathbb{U}} = 0, \tag{130}$$

with the terminal condition

$$\psi_{\mathbb{U}}(T,E,x) = 1_{\{x<0\}} + 0.5 \times 1_{\{x=0\}}, \tag{131}$$

and boundary conditions

$$\psi_{\mathbb{U}}(t,E,-\infty) = 1, \tag{132}$$

$$\psi_{\mathbb{U}}(t,E,\infty) = 0, \tag{133}$$

$$\psi_{\mathbb{U}}(t,0,x) = 1_{\{x<0\}} + 0.5 \times 1_{\{x=0\}}, \tag{134}$$

$$\psi_{\mathbb{U}}(t,\overline{E},x) = \mathbb{P}^{\mathbb{B}}\Big(N > \frac{x+\frac{1}{2}u^2(T-t)}{u\sqrt{T-t}}\Big) = 1-\Phi\Big(\frac{x+\frac{1}{2}u^2(T-t)}{u\sqrt{T-t}}\Big). \tag{135}$$

Proof. Recall that E follows the dynamics

$$\begin{aligned} dE_t =& \delta E_t dt - \tfrac{\delta}{2}u_t^2 dt + g_t dB_t^{\mathbb{U}} \\ =& \delta E_t dt + \big(g_t u_t - \tfrac{\delta}{2}u_t^2\big)dt + g_t dB_t^{\mathbb{B}}. \end{aligned}$$

The differential equation (124) follows from the standard Feynman-Kac formula. For the boundary conditions, it is natural to set (126) and (127) given the definition of $\psi_{\mathbb{B}}$ in (121). When $E=0$, $u^*=0$, then x remains the same as its initial value. As a result, we choose (128) as the boundary condition. When $E=\overline{E}$, $u^* = \frac{\mu-r}{\sigma}$. Then $x_T = x_t - \frac{1}{2}(u^*)^2(T-t) - u^*(B_T^{\mathbb{B}} - B_t^{\mathbb{B}})$.

$$\psi_{\mathbb{B}}(t,\overline{E},x) = \mathbb{P}^{\mathbb{B}}\Big(x - \tfrac{1}{2}(u^*)^2(T-t) - u^*(B_T^{\mathbb{B}} - B_t^{\mathbb{B}}) > 0\Big) = \mathbb{P}^{\mathbb{B}}\Big(N < \frac{x-\frac{1}{2}(u^*)^2(T-t)}{u\sqrt{T-t}}\Big), \tag{136}$$

where N is a standard normal random variable.

The proof of the equation for $\psi_{\mathbb{U}}$ is similar.

□

A.9 Proof of Proposition 5

We conjecture μ^L, σ^L, and r as functions of the state variable r. We will prove this conjecture later. The value function V of the representative agent is

$$V(W_t,E_t) = \sup_{\pi,c}\inf_{u,g}\mathbb{E}_t^{\mathbb{U}}\Big[\int_t^\infty F\big(c_s, V(W_s,E_s)\big)ds\Big].$$

The homothetic property of the Epstein-Zin aggregator implies the decomposition (17) for a function $\mathcal{V}$ depending on E only. Following the same argument as in the proof of Proposition 2, when $\gamma > 1$, $\mathcal{V}$ satisfies

$$\delta v \mathcal{V} = \inf_{\pi,\tilde{c}} \sup_{u,g} \Big\{ \delta v \tilde{c}^{1-\frac{1}{\psi}} \mathcal{V}^{1-\frac{1}{v}} + (1-\gamma)[r + \pi(\mu^L - r - \sigma^L u) - \tilde{c}]\mathcal{V} + \big(\delta E - \tfrac{\delta}{2}|u|^2\big)\partial_E \mathcal{V} \\ - \tfrac{1}{2}\gamma(1-\gamma)\pi^2(\sigma^L)^2 \mathcal{V} + \tfrac{1}{2} g^2 \partial^2_{EE}\mathcal{V} + (1-\gamma)\pi\sigma^L g \partial_E \mathcal{V} \Big\}. \tag{137}$$

When $0 < \gamma < 1$, $\inf_{\pi,\tilde{c}} \sup_{u,g}$ is replaced by $\sup_{\pi,\tilde{c}} \inf_{u,g}$ The first order condition in $\tilde{c}$ yields

$$\tilde{c}^* = \delta^{\psi} \mathcal{V}^{-\frac{\psi}{v}}. \tag{138}$$

The optimal π^*, u^*, and g^* are the same as in (21), (23), and (24) with μ and σ therein replaced by μ^L and σ^L.

The portfolio choice (21) with $\mu = \mu^L$ and $\sigma = \sigma^L$ therein combined with capital market clearing $\pi^* = 1$ produces the CCAPM result (54)

From the consumption market clearing and (138),

$$c_t = \tilde{c}^*_t W_t = \delta^{\psi} \mathcal{V}^{-\frac{\psi}{v}} W_t$$

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

$$\frac{dc^*_t}{c^*_t} = -\frac{\psi}{v}\frac{\partial_E \mathcal{V}}{\mathcal{V}} dE + \frac{\psi}{2v}\Big[\Big(1+\frac{\psi}{v}\Big)\frac{(\partial_E \mathcal{V})^2}{\mathcal{V}^2} - \frac{\partial^2_{EE}\mathcal{V}}{\mathcal{V}}\Big](dE)^2 + \frac{dW}{W} - \frac{\psi}{v}\frac{\partial_E \mathcal{V}}{\mathcal{V}}\frac{dW}{W} dE. \tag{139}$$

Using the dynamics of wealth and entropy under $\mathbb{B}$, combined with capital market clearing $\pi^* = 1$ and the CCAPM result (54), we can determine all endogenous equilibrium outcomes. First, matching the volatility terms from the previous equation with the volatility of the Lucas tree dividend, we find

$$\sigma^L = \sigma^c + \frac{\psi}{v}\frac{\partial_E \mathcal{V}}{\mathcal{V}} g^*. \tag{140}$$

Using this in (24) with $\sigma = \sigma^L$ therein, again invoking financial market clearing $\pi^* = 1$ results in

$$g^* = -(1-\gamma)\Big[\frac{\partial^2_{EE}\mathcal{V}}{\partial_E \mathcal{V}} + (\psi - 1)\frac{\partial_E \mathcal{V}}{\mathcal{V}}\Big]^{-1}\sigma^c. \tag{141}$$

Combining the previous two equations, we confirm (55).

Collecting the drift terms in (139), using again $\pi^* = 1$ and the CCAPM result (54), we find

$$\mu^c = -\frac{\psi}{v}\frac{\partial_E \mathcal{V}}{\mathcal{V}}\big(\delta E - \tfrac{\delta}{2}(u^*)^2 + g^* u^*\big) + \frac{\psi}{2v}\Big[\Big(1+\frac{\psi}{v}\Big)\frac{(\partial_E \mathcal{V})^2}{\mathcal{V}^2} - \frac{\partial^2_{EE}\mathcal{V}}{\mathcal{V}}\Big](g^*)^2 \\ + r + \gamma^{\text{eff}}(\sigma^L)^2 - \tilde{c}^* - \frac{\psi}{v}\frac{\partial_E \mathcal{V}}{\mathcal{V}} g^* \sigma^L. \tag{142}$$

Financial market clearing $\pi^* = 1$ and the CCAPM relation (54) also allow us to rewrite the HJB equation (137) to find the optimal consumption-wealth ratio $\tilde{c}^*$ as

$$\tilde{c}^* = \psi\Big[\delta - \Big(1-\frac{1}{\psi}\Big)r - \frac{1}{2}\Big(1-\frac{1}{\psi}\Big)\gamma^{\text{eff}}(\sigma^L)^2 - \frac{\delta}{\nu}E\frac{\partial_E \mathcal{V}}{\mathcal{V}}\Big]. \tag{143}$$

The last two equations can now be combined to get the expression for the equilibrium risk-free interest rate in (56).

Plugging (56) into (143) and using (138) on the left-hand side, we obtain that $\mathcal{V}$ satisfies the equation

$$\begin{aligned}\frac{\delta^{\psi}\mathcal{V}^{-\frac{\psi}{\nu}}}{\psi-1} =& \frac{1}{\psi-1}\delta - \frac{\mu^c}{\psi} + \frac{1}{2\psi}\gamma^{\text{eff}}(\sigma^L)^2 \\ &+ \frac{1}{\nu}\frac{\partial_E \mathcal{V}}{\mathcal{V}}\Big[\frac{\delta}{2}(u^*)^2 - g^*(u^*+\sigma^L) - \frac{\delta\psi}{\psi-1}E\Big] \\ &+ \frac{1}{2\nu}\Big[\Big(1+\frac{\psi}{\nu}\Big)\frac{(\partial_E\mathcal{V})^2}{\mathcal{V}^2} - \frac{\partial^2_{EE}\mathcal{V}}{\mathcal{V}}\Big](g^*)^2.\end{aligned} \tag{144}$$

To specify the boundary conditions for (144), we specify the value of $\mathcal{V}$ when $E=0$ or $E=\overline{E}$ for some large value of $\overline{E}$. When $E=0$, $u^* = g^* = 0$. When $E=\overline{E}$, we specify the value of $\mathcal{V}$ by setting $g^* = 0$ there. It then follows from (43) and (46) that $u^*(\overline{E}) = \frac{\mu^L - r}{\sigma^L}\frac{\theta(\overline{E})}{\gamma+\theta(\overline{E})} = \sqrt{2\overline{E}}$, for some $\theta(\overline{E})$. Combining the previous equation with the CCAPM relation $\mu^L - r = (\sigma^L)^2(\gamma+\theta(\overline{E}))$, we obtain $\theta(\overline{E}) = \frac{\sqrt{2\overline{E}}}{\sigma^L}$. In both cases, $\mathcal{V}$ satisfies the following equation at the boundary

$$\delta\nu\mathcal{V} = \inf_{\pi,\tilde{c}}\Big\{\delta\nu\tilde{c}^{1-\frac{1}{\psi}}\mathcal{V}^{1-\frac{1}{\nu}} + (1-\gamma)\big[r + \pi(\mu^L - r - \sigma^L u^*) - \tilde{c}\big]\mathcal{V} - \frac{1}{2}\gamma(1-\gamma)\pi^2(\sigma^L)^2\mathcal{V}\Big\}. \tag{145}$$

The first order condition in $\tilde{c}$ yields $\tilde{c}^* = \delta^{\psi}\mathcal{V}^{-\frac{\psi}{\nu}}$. Following the same argument as in (140), we obtain $\sigma^L = \sigma^c$. The same argument leading to (54) and (56) yields

$$\begin{aligned}\mu^L - r =& \gamma^{\text{eff}}(\sigma^L)^2 \\ r =& \delta + \frac{\mu^c}{\psi} - \frac{1+\psi}{2\psi}\gamma^{\text{eff}}(\sigma^L)^2,\end{aligned}$$

where $\gamma^{\text{eff}} = \gamma$ when $E=0$ and $\gamma^{\text{eff}} = \gamma + \theta(\overline{E})$ when $E=\overline{E}$. Plugging the previous two equations and the expressions of $\tilde{c}^*$ and σ^L into (145), we obtain the boundary conditions

$$0 = \frac{1}{\psi-1}\delta^{\psi}\mathcal{V}^{-\frac{\psi}{\nu}} + \frac{\delta}{1-\psi} + \frac{\mu^c}{\psi} + \frac{\psi-1}{2\psi}\gamma^{\text{eff}}(\sigma^c)^2 - \sigma^c u^* - \frac{\gamma}{2}(\sigma^c)^2, \tag{146}$$

where $u^* = 0$ and $\gamma^{\text{eff}} = \gamma$ for $E = 0$ and $u^* = \sqrt{2\overline{E}}$ and $\gamma^{\text{eff}} = \gamma + \theta(\overline{E})$ for $E = \overline{E}$.

The claim about $\sigma^L > \sigma^c$ is proved in the discussion after (58).

References

Ai, H. (2010). Information quality and long-run risk: Asset pricing implications. *The Journal of Finance*, 65(4):1333–1367.

Ai, H. and Bansal, R. (2018). Risk preferences and the macroeconomic announcement premium. *Econometrica*, 86:1383–1430.

Ai, H., Bansal, R., and Han, L. (2022). Information acquisition and the pre-announcement drift. Working Paper, University of Wisconsin-Madison.

Aït-Sahalia, Y. and Matthys, F. (2019). Robust consumption and portfolio policies when asset prices can jump. *Journal of Economic Theory*, 179:1–56.

Anderson, E. W., Ghysels, E., and Juergens, J. L. (2009). The impact of risk and uncertainty on expected returns. *Journal of Financial Economics*, 94(2):233–263.

Anderson, E. W., Hansen, L. P., and Sargent, T. J. (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.

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., Greenwood, R., Jin, L., and Shleifer, A. (2015). X-capm: An extrapolative capital asset pricing model. *Journal of Financial Economics*, 115:1–24.

Bhamra, H. and Uppal, R. (2019). Does household finance matter? small financial errors with large social costs. *American Economic Review*, 109(3):1116–1154.

Bhandari, A., Borovička, J., and Ho, P. (2022). Survey data and subjective beliefs in business cycle models. Working Paper, University of Minnesota.

Bidder, R. and Dew-Becker, I. (2016). Long-run risk is the worst-case scenario. *American Economic Review*, 106(9):2494–2527.

Bidder, R. M. and Smith, M. E. (2018). Doubts and variability: A robust perspective on exotic consumption series. *Journal of Economic Theory*, 175:689–712.

Bordalo, P., Gennaioli, N., Ma, Y., and Shleifer, A. (2020). Overreaction in macroeconomic expectations. *American Economic Review*, 110(9):2748–2782.

Bordalo, P., Gennaioli, N., Porta, R. L., and Shleifer, A. (2019). Diagnostic expectations and stock returns. *Journal of Finance*, 74(6):2839–2874.

Bordalo, P., Gennaioli, N., and Shleifer, A. (2018). Diagnostic expectations and credit cycles. *Journal of Finance*, 73(1):199–227.

Branger, N. and Larsen, L. S. (2013). Robust portfolio choice with uncertainty about jump and diffusion risk. *Journal of Banking and Finance*, 37(12):5036–5047.

Branger, N., Larsen, L. S., and Munk, C. (2013). Robust portfolio choice with ambiguity and learning about return predictability. *Journal of Banking and Finance*, 37(5):1397–1411.

Campanale, C. (2011). Learning, ambiguity and life-cycle portfolio allocation. *Review of Economic Dynamics*, 14(2):339–367.

Campbell, J. Y. (2018). *Financial Decisions and Markets*. Princeton University Press.

Campbell, J. Y. and Cochrane, J. (1999). By force of habit: A consumption-based explanation of aggregate stock market behavior. *Journal of Political Economy*, 107(2):205–251.

Campbell, J. Y. and Viceira, L. M. (2002). *Strategic Asset Allocation*. Oxford University Press, 1 edition.

Cao, H., Wang, T., and Zhang, H. (2005). Model uncertainty, limited market participation, and asset prices. *Review of Financial Studies*, 18(4):1219–1251.

Cerreia-Vioglio, S., Hansen, L., Maccheroni, F., and Marinacci, M. (2022a). Making decisions under model misspecification. Working Paper, University of Chicago.

Cerreia-Vioglio, S., Maccheroni, F., and Marinacci, M. (2022b). Ambiguity aversion and wealth effects. *Journal of Economic Theory*, 199:104898.

Cerreia-Vioglio, S., Maccheroni, F., Marinacci, M., and Montrucchio, L. (2011). Uncertainty averse preferences. *Journal of Economic Theory*, 146:1275–1330.

Chamberlain, G. (2020). Robust decision theory and econometrics. *Annual Review of Economics*, 12:239–271.

Chen, Z. and Epstein, L. (2002). Ambiguity, risk, and asset returns in continuous time. *Econometrica*, 70(4):1403–1443.

Christensen, T. (2019). Dynamic models with robust decision makers: Identification and estimation. Working Paper, New York University.

Cogley, T. and Sargent, T. J. (2008). The market price of risk and the equity premium: A legacy of the great depression? *Journal of Monetary Economics*, 55(3):454–476.

Collin-Dufresne, P., Johannes, M., and Lochstoer, L. A. (2016). Parameter learning in general equilibrium: The asset pricing implications. *American Economic Review*, 106(3):664–698.

Constantinides, G. M. (1990). Habit formation: A resolution of the equity premium puzzle. *Journal of Political Economy*, 98:519–543.

DeMarzo, P. and Sannikov, Y. (2017). Learning, termination, and payout policy in dynamic incentive contracts. *Review of Economic Studies*, 84:182–236.

Detemple, J. B. and Zapatero, F. (1991). Asset prices in an exchange economy with habit formation. *Econometrica*, 59(6):1633–1657.

Dew-Becker, I. and Nathanson, C. G. (2019). Directed attention and nonparametric learning. *Journal of Economic Theory*, 181:461–496.

Dimmock, S., Kouwenberg, R., Mitchell, O., and Peijnenburg, K. (2016). Ambiguity aversion and household portfolio choice puzzles: Empirical evidence. *Journal of Financial Economics*, 119:559–577.

Dow, J. and Werlang, S. R. d. C. (1992). Uncertainty aversion, risk aversion, and the optimal choice of portfolio. *Econometrica*, 60(1):197–204.

Drechsler, I. (2013). Uncertainty, time-varying fear, and asset prices. *The Journal of Finance*, 68(5):1843–1889.

Duffie, D. and Epstein, L. G. (1992). Stochastic differential utility. *Econometrica*, 60(2):353–394.

Epstein, L. G. and Schneider, M. (2003). Recursive multiple-priors. *Journal of Economic Theory*, 113(1):1–31.

Gomes, F., Haliassos, M., and Ramadorai, T. (2021). Household finance. *Journal of Economic Literature*, 59(4):1219–1251.

Hansen, L. P. (2004). [exotic preferences for macroeconomists]: Comment. *NBER Macroeconomics Annual*, 19:391–405.

Hansen, L. P. and Miao, J. (2018). Aversion to ambiguity and model misspecification in dynamic stochastic environments. *Proceedings of the National Academy of Sciences*, 115(37):9163–9168.

Hansen, L. P. and Miao, J. (2022). Asset pricing under smooth ambiguity in continuous time. *Economic Theory*, 74:335–371.

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

Hansen, L. P. and Sargent, T. J. (2008). *Robustness*. Princeton University Press.

Hansen, L. P. and Sargent, T. J. (2022). Structured ambiguity and model misspecification. *Journal of Economic Theory*, 199:105–165.

Hansen, L. P., Sargent, T. J., Turmuhambetova, G., and Williams, N. (2006). Robust control and model misspecification. *Journal of Economic Theory*, 128(1):45–90.

Hong, H. and Stein, J. C. (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 Sui, P. (2022). Asset pricing with return extrapolation. *Journal of Financial Economics*, 145(2):273–295.

Jin, X., Luo, D., and Zeng, X. (2019). Tail risk and robust portfolio decisions. *Working Paper, Shanghai University*.

Kim, T. S. and Omberg, E. (1996). Dynamic nonmyopic portfolio behavior. *The Review of Financial Studies*, 9(1):141–161.

Knox, T. (2002). Learning how to invest when returns are uncertain. Working Paper, Harvard University.

Knox, T. (2003). Foundations for learning how to invest when returns are uncertain. Working Paper, University of Chicago.

Kozlowski, J., Veldkamp, L., and Venkateswaran, V. (2020). The tail that wags the economy: Beliefs and persistent stagnation. *Journal of Political Economy*, 128(8):2839–2879.

Li, K. and Liu, J. (2019). Extrapolative asset pricing. *Working Paper, University of California, San Diego*.

Liu, H. (2010). Robust consumption and portfolio choice for time varying investment opportunities. *Annals of Finance*, 6(4):435–454.

Liu, J., Pan, J., and Wang, T. (2005). An equilibrium model of rare-event premia and its implication for option smirks. *The Review of Financial Studies*, 18(1):131–164.

Luo, Y., Nie, J., and Wang, H. (2022). Ignorance, pervasive uncertainty, and household finance. *Journal of Economic Theory*, 199:105204.

Maccheroni, F., Marinacci, M., and Rustichini, A. (2006). Ambiguity aversion, robustness, and the variational representation of preferences. *Econometrica*, 74(6):1447–1498.

Maenhout, P. J. (2004). Robust portfolio rules and asset pricing. *The Review of Financial Studies*, 17(4):951–983.

Maenhout, P. J. (2006). Robust portfolio rules and detection-error probabilities for a mean-reverting risk premium. *Journal of Economic Theory*, 128(1):136–163.

Malmendier, U. and Nagel, S. (2011). Depression babies: Do macroeconomic experiences affect risk taking? *The Quarterly Journal of Economics*, 126(1):373–416.

Merton, R. C. (1971). Optimum consumption and portfolio rules in a continuous-time model. *Journal of Economic theory*, 3(4):373–413.

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

Peijnenburg, K. (2018). Life-cycle asset allocation with ambiguity aversion and learning. *Journal of Financial and Quantitative Analysis*, 53(5):1963–1994.

Rogerson, W. P. (1985). Repeated moral hazard. *Econometrica*, 53:69–76.

Skiadas, C. (2003). Robust control and recursive utility. *Finance and Stochastics*, 7(4):475–489.

Spear, S. and Srivastava, S. (1987). On repeated moral hazard with discounting. *Review of Economic Studies*, 54:599–617.

Strzalecki, T. (2011). Axiomatic foundations of multiplier preferences. *Econometrica*, 79(1):47–73.

Trojani, F. and Vanini, P. (2004). Robustness and ambiguity aversion in general equilibrium. *Review of Finance*, 8:279–324.

Uppal, R. and Wang, T. (2003). Model misspecification and underdiversification. *The Journal of Finance*, 58(6):2465–2486.

Williams, N. (2015). A solvable continuous time dynamic principal–agent model. *Journal of Economic Theory*, 159:989–1015.