Learning under Ambiguous Reversion

Hongseok Choi*

May 27, 2015

Abstract

While "it is now widely accepted that excess returns are predictable" (Lettau and Ludvigson, 2001, *Journal of Finance*), there also are authors finding otherwise, claiming that the predictive models are unstable or even spurious. This paper proposes a model of learning through which we can study the behavior of an investor under such ambiguous circumstances. The proposed model describes how observations are translated into a *set* of probability measures that represents the investor's view of the immediate future; and I explicitly characterize the set's evolution up to a system of differential equations that generalizes the Kalman-Bucy filter in the presence of ambiguity. The model of learning is then applied to the portfolio choice problem of a log investor; and learning under ambiguity is seen to have a significant effect on hedging demand: under a reasonable calibration, the optimal demand for the risky asset at zero instantaneous equity premium decreases, as the investor loses confidence, by half of wealth.

1 Introduction

There is a large body of empirical literature demonstrating stock return predictability. One of the most frequently considered predictors is the dividend-price ratio (Campbell and Shiller, 1988; Fama and French, 1988a). Since the dividend-price ratio is stationary, the predictability is related, in particular, to time-varying, mean-reverting expected returns (Fama and French, 1988b; Poterba and Summers, 1988).

Meanwhile, the form of time variation in investment opportunities is crucial to making investment decisions. When investment opportunities are constant, the optimal portfolio, too, is constant under standard utility functions (Merton, 1969; Samuelson, 1969). When investment opportunities are stochastic, on the other hand, the optimal portfolio has an additional component that reflects the investor's desire to hedge against adverse changes in the investment opportunities (Merton, 1973).

The facts outlined in the preceding paragraphs have led numerous authors to investigate portfolio choice under mean-reverting expected returns. To focus on continuous-time models,

*Republic of Korea Air Force Academy; Danjae-ro 635, Cheongju, Chungbuk 360-849, Republic of Korea; 82-10-5115-2037; aitch.choi@gmail.com. This paper is a revised version of the first two chapters of my doctoral dissertation at the University of Pennsylvania (Choi, 2012); and I owe my deepest gratitude to my advisor, Domenico Cuoco, for his guidance and support throughout the course of this research. I would also like to thank Philipp Illeditsch and Dirk Krueger; and the seminar participants at Seoul National University.

Kim and Omberg (1996) and Wachter (2002) characterized the optimal portfolio in closed form and showed, for example, that when relative risk aversion is greater than one and the equity premium is positive, the hedging demand is positive. Zohar (2001) and Brendle (2006) provided explicit solutions in the case of partial observation, in which the expected return, or the drift of the return process, is unobservable.[1]

However, claims of time variation in expected returns, or fundamentally of return predictability, also met with doubts. For example, Goetzmann and Jorion (1993) and Nelson and Kim (1993) found that small sample biases are significant, and once they are accounted for, there is much less evidence that the dividend-price ratio predicts stock returns. More recently, Welch and Goyal (2008) pointed out that standard predictors had performed poorly not only in-sample but especially out-of-sample. See, then, the apologetics in defense of predictability by Campbell and Thompson (2008) and Cochrane (2008)—the debate, after almost three decades, is yet to be resolved.

This paper proposes a framework through which we can study (among others) the behavior of an investor under such ambiguous circumstances. More specifically, the investor believes, as is the prevailing view of the financial economics profession, that mean-reverting expected returns is a plausible assumption. For concreteness, suppose that the only source of uncertainty in the market is the performance of a risky asset, and denote by R its cumulative return process. Under mean reversion, R satisfies the following system of stochastic differential equations (SDEs):

$$\mathrm{d}R(t) = x(t)\,\mathrm{d}t + \sigma_R\,\mathrm{d}w(t), \tag{1}$$

$$\mathrm{d}x(t) = \kappa(\bar{x} - x(t))\,\mathrm{d}t + \rho_w\,\mathrm{d}w(t) + \rho_v\,\mathrm{d}v(t), \tag{2}$$

where $0 \le t \le T < \infty$, w and v are independent Wiener processes, and the expected return process x is unobservable. The drift of x, $\kappa(\bar{x} - x(t))\,\mathrm{d}t$, drives mean reversion, a constant gravitation toward some fixed value $\bar{x}$. Now, facing at the same time nonnegligible evidence that questions its validity, the investor fails to have full confidence in mean reversion.

How, then, are we to represent a lack of confidence in mean reversion? In this paper, I ascribe it to model instability. Indeed, several authors have ascribed the failure to reject the null of random walk to the time-dependent nature of mean reversion, rather than interpreting it as a reason to embrace the random walk hypothesis. For example, Kim et al. (1991) and McQueen (1992) found that mean reversion is rather restricted to the Great Depression-World War II period; Welch and Goyal (2008) found that for many predictive models, statistical significance is based exclusively on the years of the Oil Shock of 1973-1975; and Lettau and Nieuwerburgh (2008) found that the steady-state value of the dividend-price ratio, or, precisely speaking, the value to which the dividend-price ratio is instantaneously tending, had shifted either once in the early 1990s or twice around 1954 and 1994, and showed that taking the structural breaks into account strengthens the evidence for predictability. These findings show not only that there is instability in the dynamics of expected returns, but also—notice the divergence in the timing of breaks—how difficult it is to infer the time

[1] These four papers assume that the expected return or the price of risk follows an Ornstein-Uhlenbeck process. This specification dates back to Merton (1971). Schroder and Skiadas (1999) and Liu (2007) explicitly characterize the optimal portfolio under more general return dynamics, up to a system of ordinary differential equations. See also Lakner (1998) and Xia (2001).

variation in the dynamics.[2]

Thus, I assume that the investor uses, in place of (2), the following SDE in making inferences:

$$\mathrm{d}x(t) = \kappa(\bar{x} + \kappa^{-1}\rho_v \eta(t) - x(t))\,\mathrm{d}t + \rho_w\,\mathrm{d}w(t) + \rho_v\,\mathrm{d}v(t) \tag{3}$$

where η is a function of time. Importantly, no assumption is made about η (besides the minimal technical one of square integrability). The presence of η clearly allows for the structural breaks suspected in the literature; meanwhile, the lack of assumptions on η reflects the fact that evidence is too inconclusive for the investor to confidently make one.

Here, it might seem that I have actually made an assumption about η, namely, it being a *deterministic* function of time; but I have not. When referring to η as a function of time, I am not highlighting η's time dependence but simply adopting the familiar parlance. η may very well vary over time, but the point is that we do not know if it does, and if it does, how. Thus, the more appropriate way to think of η is as a collection of parameters: $\eta = \{\eta(t) : t\}$. Viewed thus, $\eta(t)$ represents the perturbation that the investor suspects might have occurred at time t to the "steady state" $\bar{x}$ of the market; possibility of perturbations, however, is all he can be certain of, and he cannot form, a priori or a posteriori, a probability distribution over $\mathbb{R} \ni \eta(t)$, much less conceive a connection between $\eta(t)$ and $\eta(s)$, $t \neq s$.[3]

Thus, the investor in question is facing a kind of model uncertainty, or *ambiguity*; and the main contribution of this paper is to provide a model of learning under the ambiguity that conforms to the axiomatically founded *recursive multiple-priors utility*, thereby, most importantly, endogenizing the time variation in ambiguity.

Ambiguity, to be more specific, refers to the nonexistence of a probability measure that can rationalize given preferences. Standard models in economics, on the other hand, adopt Bayesianism, which presumes the existence of one. However, as exemplified by the divided views *even among experts* on the predictability of returns, there do exist situations where, due to limited knowledge and/or the complexity of the environment, agents are unable to settle on a single probabilistic model and tell the probabilities of all uncertain events precisely; see Ellsberg (1961) for more intuitive examples. In fact, one may go as far as to say that "it is sometimes *more* rational to admit that one does not have sufficient information for probabilistic beliefs than to pretend that one does" (Gilboa et al., 2012, emphasis in the original).

One way to model decision makers under ambiguity is Gilboa and Schmeidler's (1989) *multiple-priors utility*:

$$U(c) = \min_{P \in \mathcal{P}} \mathrm{E}^P\, u(c) \tag{4}$$

where $\mathcal{P}$ is called the *set of priors*. A straightforward interpretation is that the decision maker behaves as if he had multiple priors in mind and assessed each act under the respective worst-case scenario. Multiple-priors utility is thus sometimes criticized for its apparent extreme pessimism. Any critique of a behavioral model, however, should be directed at its axioms, and those of Gilboa and Schmeidler are standard (Anscombe and Aumann, 1963),

[2] To quote Merton (1980), "Estimating expected returns from time series of realized stock return data is very difficult. $\cdots$ Unless a significant portion of the variance of the market returns is caused by changes in the expected return on the market, it will be difficult to use the time series of realized market returns to distinguish among different models for expected return."

[3] In this respect, $\eta(t)$s are closely related to what Neyman and Scott (1948) called *incidental parameters*.

except for a weakening of the classical Independence axiom and the introduction of Uncertainty Aversion, which innocuously captures hedging behavior. Equally importantly, over the recent decades, multiple-priors models (static or dynamic) have been successful in explaining puzzling phenomena in financial markets: stock market nonparticipation (Dow and Werlang, 1992), excess volatility (Epstein and Wang, 1994; Illeditsch, 2011), excess equity premium (Chen and Epstein, 2002), and equity home bias (Epstein and Miao, 2003), to name a few.

The model of learning of this paper is then formulated in conformity with the dynamic version of multiple-priors utility known as *recursive multiple-priors utility*; it was axiomatized in discrete time by Epstein and Schneider (2003) and (nonaxiomatically) formulated in continuous time by Chen and Epstein (2002) in such a way that, most characteristically, the preferences are dynamically consistent:[4][5]

$$U(t,c) = u(c(t)) + \beta \min_{P_t^{+1} \in \mathcal{P}_t^{+1}} \mathrm{E}^{P_t^{+1}}\, U(t+1,c) \tag{5}$$

where $\mathcal{P}_t^{+1}$ is a set of probability measures on the period-$(t+1)$ information. (The present paper assumes continuous time, but in overviewing ideas, discrete-time representations are more convenient.) The natural interpretation of $\mathcal{P}_t^{+1}$ is that it models the agent's conditional beliefs about one-step-ahead uncertainties, larger sets signifying larger ambiguity; furthermore, as Epstein and Schneider emphasize, the axioms of recursive multiple-priors utility impose no restrictions on how observations are mapped to one-step-ahead beliefs, or the *one-step-ahead beliefs correspondence*. Therefore, modeling learning in the context of recursive multiple-priors utility comes down to defining the correspondence, especially the way the *size* of the set of one-step-ahead beliefs responds to observations.

Most applications of recursive multiple-priors utility, however, disregard learning and specify the time variation, or a lack thereof, in ambiguity exogenously. Most notably, Chen and Epstein (2002) proposed a time-invariant form of ambiguity called *independently and indistinguishably distributed (IID) ambiguity*,[6] which is to model "the agent after he has learned all that he can"; and ever since, it has been the dominant specification of ambiguity in the literature. See Hernández-Hernández and Schied (2006, 2007a,b), Schied (2008), Miao (2009), Routledge and Zin (2009), and Liu (2011, 2013) for applications of IID ambiguity to portfolio choice, and Epstein and Miao (2003), Trojani and Vanini (2004), and Gagliardini et al. (2009) for those to asset pricing. Time-varying ambiguity without a learning mechanism has been considered by Porchia (2005), Sbuelz and Trojani (2008), and Drechsler (2013).

Naturally, however, such lack of a learning mechanism is unsatisfactory: when does ambiguity resolve and when does it not? when it persists, what determines its long-run level and variations thereabout? and is the agent really going to face an IID ambiguity when he has learned all that he can? These questions do concern financial economics. For example, time variation in ambiguity can affect the optimal portfolio through hedging demand (Epstein and Schneider, 2007; Hernández-Hernández and Schied, 2007a); and in an economy populated

[4]See also Wang (2003), who provides related representation results in discrete time; and Epstein and Ji (2014), who extend Chen and Epstein (2002) to cases of nonequivalent priors.

[5]Since ambiguity aversion is a departure from Bayesianism, we are bound to lose along the way one or another of the desirable properties associated with the latter; and the loss may well be dynamic consistency. See Epstein and Schneider (2003, Section 4.1) and Siniscalchi (2011).

[6]The uppercase initialism IID is to be distinguished from i.i.d., independent and *identically* distributed.

by multiple-priors agents, the equilibrium equity premium includes a premium for bearing ambiguity as well as the risk premium (Chen and Epstein, 2002), and the specification of ambiguity thus has a direct effect on the level of and variations in the equity premium. In this paper, I provide answers to the above questions by deriving an endogenous dynamics of ambiguity and investigate its implications for portfolio choice.[7]

The endogenous dynamics of ambiguity is then derived by explicitly modeling the cognitive process of the agent, picturing him as a statistician. The theory of recursive multiple-priors utility is silent as to where the one-step-ahead beliefs come from in much the same way as Savage's is regarding the unique "a priori" measure; or, for that matter, as any other approaches that take preferences as the primitive. If we were to follow such positivist approaches, however, the most we can do with the weights appearing in the representations is to interpret them as subjective probabilities (or, perhaps preferably, none at all), whereas economic theories in fact typically proceed the other way around, deriving optimal choices from some sophisticated statistical model of the environment that the agent was declared at the outset to believe in. Of course, the logic is probably that if the agent somehow managed to cognitively construct such a statistical model, then it is reasonable to expect the formal model reconstructed from the resulting preferences to coincide with it. And in this paper, I make explicit the cognitive origin of the imprecise elicited probabilities.

Specifically, if each probability law induced by the SDEs (1) and (3) is viewed as a statistical model, or a *theory*, of the stock market, the investor is entertaining multiple theories (indexed by η) as in classical statistics. Having conceived these theories a priori,[8] the investor then computes the conditional one-step-ahead probability measures from a confidence set of theories: at each point in time, he (i) ranks the theories in order of plausibility based on the observations he has made; (ii) rules out the ones that are too implausible compared to the most plausible one; (iii) conditions the surviving ones individually on the observations; and, finally, (iv) restricts the conditioned measures to the one-step-ahead uncertainties. The exclusion, in particular, is only temporary; theories that are at one point disregarded for being too implausible can later turn out to be acceptably plausible.[9] Also, note that when the investor has full confidence in a single theory, the proposed learning mechanism reduces to the Bayesian one.

In defense of this nonaxiomatic incorporation of a non-Bayesian updating scheme, I emphasize that (i) in making decisions, the investor will still be conforming to all the axioms of recursive multiple-priors utility and (ii) what is described above is a standard procedure in (classical) statistics, that is, one that is actually being carried out everyday by many, if not the majority. With respect to the second point the present model may appear demanding

[7]Asset pricing implications are investigated in the third chapter of Choi (2012). There, considering an economy populated by a representative agent facing ambiguity in the dynamics of dividends, I show that (i) learning under ambiguity generates a declining trend in the equity premium, (ii) an improvement in the quality of a signal can increase the equity premium, and (iii) the relationship between the equity premium and the conditional variance of returns is unclear.

[8] This *a priori* is of course an *a posteriori* relative to a suppressed past. The investor has somehow arrived at the set of theories described above and stays with it for the time being; and the present paper describes his behavior during this period. It may be that the investor stays with the particular set of theories for the rest of his life; but it may as well be that at some point, he moves on to another set of theories or even to an entirely different mode of learning. The latter case is beyond the scope of this paper.

[9]See Footnote 8.

too much from "representative" investors. But, on the contrary, it is the "Bayesian dogma of precision" (Walley, 1991) that is unrealistically demanding, and it seems more reasonable to put the agents in economic models on a comparable footing with statisticians (Hansen and Sargent, 2001, 2008), who, after all, routinely consider multiple theories and report confidence sets.

A subtler point is that there are in fact two distinct types of probabilities involved here. To see this, note that recursive multiple-priors utility is embodying the *generalized Bayesian updating rule*, according to which the whole set of priors is Bayes-updated prior by prior. To elaborate, the recursive representation (5) can naturally be written in the static form (4) as well (with a suitable reinterpretation of u in (4)). In particular, at time t we have

$$U(t,c) = \min_{P_t \in \mathcal{P}_t} \mathrm{E}^{P_t} \sum_{s \geq t} \beta^{s-t} u(c(s))$$

where $\mathcal{P}_t$ is the generalized Bayesian update of $\mathcal{P}$. Thus, the measures associated with the agent's a priori theories are apparently not the measures $\mathcal{P}$ appearing in the representation of his preferences. To quote Gilboa and Marinacci (2013), "[The] set of priors need not necessarily reflect the individual's knowledge. Rather, information and personal taste jointly determine [the set of priors]." In this paper, I give the two types of a priori measures bland yet unmistakable names, the *theoretical* and the *preferential priors*, respectively.[10]

Thus, one way to summarize the description so far is that this paper portrays an investor/statistician presented with conflicting evidence regarding mean reversion, who, as a decision maker, behaves consistently; and the paper's focus is how his theoretical priors are translated into preferential priors.

The rest of the paper is organized as follows. Section 2 overviews the model of learning and its implications for portfolio choice. Section 3 defines and solves the model of learning. Section 4 applies it to portfolio choice. All proofs are collected in the appendix.

2 Overview

2.1 The Model of Learning

The full-fledged model that will actually be taken up in Section 3 considers the following partially observable system

$$\begin{aligned} \mathrm{d}y(t) &= (a(t,y) + b(t,y)x(t))\,\mathrm{d}t + \sigma(t,y)\,\mathrm{d}w(t), \\ \mathrm{d}x(t) &= \kappa(\bar{x} + \kappa^{-1}\rho_v \eta(t) - x(t))\,\mathrm{d}t + \rho_w\,\mathrm{d}w(t) + \rho_v\,\mathrm{d}v(t), \end{aligned} \tag{6}$$

where y is the observable process and x the unobservable process; all variables are allowed to be multidimensional; and a, b, and σ are nonanticipating path functionals. This is the

[10] The distinction is by no means novel; it is only that in Bayesian models, either the two notions of probability are mixed or the distinction is vacuous. According to Smets and Kennes (1994), for example, "There is a *credal level* where beliefs are entertained and a *pignistic level* [from *pignus*, a bet in Latin] where beliefs are used to make decisions[.] $\cdots$ The credal level precedes the pignistic level in that, at any time, beliefs are entertained (and updated) at the credal level; the pignistic level appears only when a decision needs to be made." (Italics in the original.)

general *conditionally Gaussian process* considered by Liptser and Shiryaev (1977, Chapter 12) with only the dynamics of the unobservable component x slightly restricted to fit the economic motivation of ambiguous reversion. In this overview, however, let us continue to work with the simplest special case where the observable process is the cumulative return process R of a risky asset; that is, (6) is replaced by

$$\mathrm{d}R(t) = x(t)\,\mathrm{d}t + \sigma_R\,\mathrm{d}w(t), \tag{7}$$

and all variables are scalar.

In the full-fledged model, not only η but $\bar{x}$, too, is assumed unknown. Then, were it not for η, a Bayesian agent would form a unique parameter prior[11] for $\bar{x}$, and his (unique) theory would be given by the parameter prior together with the (unique) log-likelihood function, $\ell_t(\bar{x})$. Following Epstein and Schneider (2007),[12] the present paper generalizes this Bayesian model of data generation by multiplying both the parameter prior and the likelihood; that is, due to ambiguity, the agent of this paper entertains multiple parameter priors and multiple likelihoods. Note that η is precisely what indexes the multiple likelihoods: $\ell_t(\bar{x}, \eta)$. Also, I assume that the parameter priors are all Dirac measures.[13] Then, each theory can be identified with $(\bar{x}, \eta)$, and the set of theories itself can be viewed as a semi-parametric model.

If we further assume that, for all theories, the initial distribution of x is Gaussian (with common mean and variance), then the conditional distribution of $x(t)$, conditional on the agent's observation of R up to time t, too, is Gaussian, and the conditional mean, $m^{\bar{x},\eta}(t)$, and the conditional variance, $\gamma(t)$, of $x(t)$ evolve according to the *Kalman-Bucy filter* (Liptser and Shiryaev, 1977, Chapter 10):

$$\begin{aligned} \mathrm{d}m^{\bar{x},\eta}(t) &= \kappa(\bar{x} + \kappa^{-1}\rho_v\eta(t) - m^{\bar{x},\eta}(t))\,\mathrm{d}t + (\rho_w + \gamma(t)/\sigma_R)\,\mathrm{d}\bar{w}^{\bar{x},\eta}(t), \\ \dot{\gamma}(t) &= \rho_w^2 + \rho_v^2 - 2\kappa\gamma(t) - (\rho_w + \gamma(t)/\sigma_R)^2, \end{aligned} \tag{8}$$

where the *innovation process* $\bar{w}^{\bar{x},\eta}$ is defined by

$$\mathrm{d}R(t) = m^{\bar{x},\eta}(t)\,\mathrm{d}t + \sigma_R\,\mathrm{d}\bar{w}^{\bar{x},\eta}(t).$$

Thus, in particular, the one-step-ahead conditionals[14] are given by

$$\mathrm{d}R(t)|\mathcal{G}_t \sim N(m^{\bar{x},\eta}(t)\,\mathrm{d}t, \sigma_R^2\,\mathrm{d}t)$$

where $\mathcal{G}_t$ denotes the agent's information at time t.

In other words, the ambiguity in question boils down to that in the conditional expectation of $x(t)$. Bayesian agents with unique theory $(\bar{x}, 0)$ would behave as if they were confident with the point estimate $m^{\bar{x},0}(t)$ of $x(t)$; the a posteriori dispersion of $x(t)$ is "second-order"

[11] To distinguish this a priori measure from the others, namely, the theoretical and the preferential priors, I call it the parameter prior. These three kinds of priors are all defined on different algebras.

[12] See also Epstein and Seo (2010).

[13] To quote Epstein and Schneider (2007), "Indeed, one may wonder whether there is a need for non-Dirac [parameter] priors at all."

[14] Since I work with continuous time, *one-step-ahead* is an abuse of language, which can be remedied, if need be, by introducing infinitesimal generators. See, for example, Anderson et al. (2003).

and is practically suppressed. Our agent, on the other hand, constructs a confidence set and behaves as if he responded to the uncertainty in estimation most cautiously.[15]

It remains to explain how the agent measures the plausibility of a theory. The natural criterion is the likelihood, but then, not surprisingly, the induced likelihood function for $m^{\bar{x},\eta}(t)$ is flat; each value of $m^{\bar{x},\eta}(t)$ can be supported equally well by some theory with a large η. In other words, the ambiguity is too large for there to be learning, if the agent assesses plausibility based on the likelihood alone.[16]

Indeed, "inductive inference based on objective criteria alone is bound to fail, while incorporating subjective criteria alongside objective ones can lead to successful learning," to quote Gilboa and Samuelson (2012). That is, "effective learning requires a willingness to sacrifice goodness-of-fit in return for enhanced subjective appeal" (ibid.).

Thus, I assume that the plausibility ranking, a binary relation "at least as plausible as," over the theories is represented by a penalized log-likelihood function. Specifically, the agent finds more appealing the "reference" or "simple" theories free of the poorly understood factors, and that subjective criterion is translated into a penalty on the log-likelihood proportional to the magnitude of η measured by the L^2-norm:

$$\ell_t^\lambda(\bar{x},\eta) \triangleq \ell_t(\bar{x},\eta) - \frac{\lambda}{2}\int_0^t |\eta(s)|^2 \,\mathrm{d}s$$

where $\lambda \in (0,\infty]$ measures the agent's a priori confidence about the reference theories. When $\lambda = \infty$, the set of theories reduces to $\{(\bar{x},0) : \bar{x}\}$ and the agent perceives no persistent source of ambiguity; when λ is small, the agent fits data with large ηs with little restraint. It is also worth noting that the L^2-norm of η is equal to the deviation of a theory $(\bar{x},\eta)$ from its simple counterpart $(\bar{x},0)$ measured by the Kullback-Leibler divergence.

Since the works of Good and Gaskins (1971) and Akaike (1973), penalizing the likelihood has been a standard method in statistics and information theory to strike a balance between the goodness of fit and the simplicity of the model; and the penalized log-likelihood representation of a plausibility ranking has recently been axiomatized by Gilboa and Schmeidler (2010). (See Remark 3.2 for a discussion on alternatives to the L^2-penalty.)

Finally, a theory $(\bar{x},\eta)$ is not ruled out if and only if

$$\ell_t^\lambda(\bar{x},\eta) \geq \max_{\bar{x}',\eta'} \ell_t^\lambda(\bar{x}',\eta') - \alpha$$

where $\alpha \in [0,\infty)$ measures how conservative the agent is in model selection. When $\alpha = 0$, in particular, the agent keeps nothing but the most plausible theories.

To characterize the set of preferential priors, define first a process ϵ by

$$\mathrm{d}R(t) = m^{\bar{x}_t^*,\eta_t^*}(t)\,\mathrm{d}t + \sigma_R\,\mathrm{d}\epsilon(t)$$

[15] Given that our agent is uncertain about the parameter $\bar{x}$, a fair comparison would require that the Bayesian agents be given a diffuse parameter prior. But the form of the parameter prior is irrelevant to the point I am trying to make here, namely, unique versus nonunique one-step-ahead conditionals.

[16] What is happening here is over-fitting. A related, more intuitive example is the tossing of an ambiguous coin (see Walley's (1991) discussion of vacuous previsions). If all we assume about the probability $p+q_i$ of the ith toss coming up heads is that $p+q_i \in [0,1]$, then after seeing, for example, HTTHT, the maximum likelihood estimate of $\{p+q_i\}$ is 10010; and in this case, past observations are of no help in guessing p. See also Epstein and Schneider's (2007) portfolio choice example, especially the proof of Proposition S1 in their supplementary appendix.

where $(\bar{x}_t^*, \eta_t^*)$ denotes the most plausible theory at time t. Then, there exists a unique probability measure P^0 on $\mathcal{G}_T$ such that ϵ is a Wiener process under P^0 (Proposition 3.8). This P^0 serves as the "center" of the set of preferential priors.

Two fundamental observations are as follows. First, the one-step-ahead conditionals are given by

$$\mathrm{d}R(t)|\mathcal{G}_t \sim N(m\,\mathrm{d}t, \sigma_R^2\,\mathrm{d}t),\ \frac{1}{2}\delta(t)^{-1}(m - m^{\bar{x}_t^*,\eta_t^*}(t))^2 \leq \alpha$$

(Lemma 3.6), where δ is an (absolutely continuous) process to be specified shortly. The second observation is the much advertised filtering equations (Propositions 3.6 and 3.7): $m^{\bar{x}_t^*,\eta_t^*}(t)$, $\delta(t)$, and $\bar{x}_t^*$ evolve according to

$$\mathrm{d}m^{\bar{x}_t^*,\eta_t^*}(t) = \kappa(\bar{x}_t^* - m^{\bar{x}_t^*,\eta_t^*}(t))\,\mathrm{d}t + [\rho_w + (\gamma(t) + \delta(t))/\sigma_R]\,\mathrm{d}\epsilon(t),$$

$$\dot{\delta}(t) = (\rho_w + \gamma(t)/\sigma_R)^2 - 2\kappa\delta(t) - [\rho_w + (\gamma(t) + \delta(t))/\sigma_R]^2 + 2\sigma_{\bar{x}^*}(t) + \lambda^{-1}\rho_v^2,$$

and

$$\kappa\,\mathrm{d}\bar{x}_t^* = (\sigma_{\bar{x}^*}(t)/\sigma_R)\,\mathrm{d}\epsilon(t), \tag{9}$$

where for the details regarding $\sigma_{\bar{x}^*}$ the reader is referred to Section 3.

Without ambiguity (refer to (8) with $\eta \equiv 0$), the weight on the innovation is increasing in the a posteriori Bayesian uncertainty γ in the state known as the *estimation risk*; the less trustworthy the current estimate, the more weight given to the new evidence. With ambiguity, the weight on the innovation is augmented by the a posteriori Knightian uncertainty δ in the data-generating mechanism, or *estimation ambiguity*. That is, in revising his estimates, the agent takes into account not only the uncertainty within each theory but also the uncertainty over the theories.

Section 3.4 further discusses the filtering equations.

Section 3.4.1 shows that the ambiguity associated with $\bar{x}$ eventually resolves in the sense that the confidence set shrinks to a singleton as time goes to infinity, at any significance level. Interestingly, Lettau and Nieuwerburgh (2008) assumed the steady state, or, to be precise, the value to which the dividend-price ratio is instantaneously tending, to be a martingale; and here, the *estimate* $\bar{x}^*$ of the constant $\bar{x}$ turns out to be one under the agent's reference (preferential) prior P^0 (see (9)). $\bar{x}^*$ is furthermore L^2-bounded, and therefore, under P^0, it converges by Doob's martingale convergence theorem.[17]

Section 3.4.2 compares the filter with the classical conditionally Gaussian filter, a special case of which is the celebrated Kalman-Bucy filter. Along the way I answer the questions asked in the introduction. First of all, ambiguity persists if and only if the agent lacks confidence in mean reversion, that is, $\lambda < \infty$. And when it persists, I identify in Section 3.4.3 a necessary and sufficient condition for convergence to an IID ambiguity: $\sigma^{-1}b$ *converge to a constant vector*, which is trivially satisfied in the single-asset case considered above. Since in this case, γ as well as δ converges to a constant, if we further assume that $\bar{x}^*$, too, converges (or is known), then, the agent at $t \geq \infty$ is observationally equivalent to someone in a fully observable market

$$\mathrm{d}R(t) = m(t)\,\mathrm{d}t + \sigma_R\,\mathrm{d}\epsilon(t), \tag{10}$$

$$\mathrm{d}m(t) = \kappa(\bar{m} - m(t))\,\mathrm{d}t + \sigma_m\,\mathrm{d}\epsilon(t), \tag{11}$$

[17] Convergence under P^0 is not exactly what we are after, however. See Section 3.4.1.

with an IID ambiguity; that is, the present paper justifies the IID assumption on ambiguity made in the context of standard dynamics like (10)-(11) (see, for example, Trojani and Vanini (2004) and Liu (2013)) indeed as the limit of a learning process. The size of the asymptotic ambiguity, too, depends intuitively on model primitives: it is increasing in the agent's lack of confidence and decreasing in the variability of the unobservable process relative to the observable process.[18] In general, the level of ambiguity varies over time in response to variations in the volatility of the observable process (viz. signal precision) and other quantities.

Next I look into the observational equivalence of agents with $\alpha = 0$, or *maximum plausibility agents*, to Bayesians. Clearly, the former are observationally equivalent to the latter with unique preferential prior P^0. In Section 3.4.4, I find a theory that is consistent with P^0, thereby rendering the maximum plausibility agents Bayesians under partial observation. For them, specifically, the effect of learning under ambiguous reversion is the same as an increase in the volatility of the unobservable process.

As mentioned in the introduction, Epstein and Miao (2003) have used a recursive multiple-priors model to explain the equity home-bias puzzle: agents invest less in foreign assets because they find foreign economies more ambiguous. Epstein and Miao, however, *took as given* the differing degrees of ambiguity. In Section 3.4.5, finally, I show that the present model on the other hand can deliver heterogeneous ambiguity *endogenously* as an outcome of learning under asymmetric information.

2.2 Portfolio Choice

In Section 4, I apply the model of learning under ambiguous reversion to the consumption/portfolio choice problem of a log investor and contrast the optimal portfolio with that of Epstein and Schneider's (2007) investor.

Log investors are well-known to be myopic under risk. Introducing ambiguity to stock returns, Epstein and Schneider have shown that log investors hedge against low continuation utility when they learn under ambiguity. Nevertheless, the hedging occurs only when the estimated equity premium is not unambiguously different from zero, and when it is close to zero, hedging demand is not significant. When the estimated equity premium is exactly zero, in particular, hedging demand is zero and so is the total demand (see their Figure 3). This turns out to be because of Epstein and Schneider's constant investment opportunities (IID returns) assumption and the consequent symmetry of the model.

The present paper assumes, on the other hand, stochastic investment opportunities (returns are, albeit ambiguously, predictable), and it then turns out that the investor is non-myopic for a wide range of estimated equity premia and the hedging demand is significant, even for a small ambiguity in the equity premium. When the model is calibrated so that the implied ambiguity in the instantaneous equity premium equals 0.01 (annual, decimal),[19] the optimal portfolio is such that the investor is nonmyopic when the estimate of the in-

[18] The last observation, in particular, is in line with the remark by Merton quoted in Footnote 2.

[19] The financial market is calibrated based on Barberis (2000), the investor's discount factor is set at 0.03, and the investor is assume to be facing a 10-year investment horizon after having observed the market for 20 years. The values of λ and α are then chosen so that the implied ambiguity in the instantaneous equity premium equals 0.01.

stantaneous equity premium falls between -0.03 and 0.01; and when the estimate is zero, in particular, he sells short an amount of the risky asset worth about half of his wealth. With mean reversion in expected returns, the temporary estimate of zero is more likely to rise than fall, and this asymmetry in the market dynamics renders small negative estimates worse (in the sense of low continuation utility) than zero.

So one interpretation of the hedging demand is that it reflects the Mertonian (1973) desire to hedge against adverse changes in the state variables. But why do Bayesian log investors not act on the same adverse possibilities, and why does the introduction of ambiguity make them do so all of a sudden? Whereas the optimal portfolio is only numerically characterized in Epstein and Schneider (2007), the continuous-time framework of this paper affords us a degree of analytical tractability,[20] and this facilitates an alternative interpretation of the optimal portfolio.

Suppose there are one risk-free asset (bond) and one risky asset (stock) with return dynamics (7). Then, the optimal demand for the stock has the form of a "shadow"—for it does not always fully realize—hedging demand bounded by the optimistic and pessimistic Bayesian demands, where the optimistic (pessimistic) Bayesian demand refers to the Bayesian demand when the estimated instantaneous premium equals the highest (lowest) in the confidence set. Regarding the shadow hedging demand, while the said interpretation still stands—it is proportional to the derivative of the value function and the negative of the covariation between return and state—I show that it is at the same time the position in the stock that eliminates the effect of misspecification on continuation utility. That is, the investor wishes to make the backward induction immune to misspecification by taking a particular position in the stock, but it may be disproportionate to the ambiguity actually present and cannot exceed the position dictated by the most optimistic or most pessimistic estimate. Note that this is in fact reminiscent of the other interpretation by Merton (1973) of his hedging demand that it minimizes the volatility of consumption.

Epstein and Schneider also note that when the estimated premium is not unambiguously different from zero, ambiguity-averse log investors exhibit contrarian behavior in the sense that they go long for negative premia and short for positive premia. Mean reversion affects this behavior in two ways. First, contrarian behavior has to be redefined more generally as decreasing one's stock holdings as the estimated premium increases. The shadow hedging demand is contrarian in this sense because the Gilboa-Schmeidler ambiguity aversion, with its "bounded reservations," does not fundamentally alter the quadratic (that is, faster than linear) structure of the Bayesian value function. And while we recover the symmetric contrarian policy (as a function of the instantaneous premium) of Epstein and Schneider when the long-run premium is zero, the interval of contrarian behavior is skewed toward negative instantaneous premia when the long-run premium is positive. Second, contrarian behavior is not robust. When returns and expected returns are negative correlated (as is the case in reality; see, for example, Barberis (2000)), it occurs only when ambiguity is sufficiently high.

Another important observation is that, for investors with low confidence, the less confident an investor is, the larger his hedging demand is (in absolute value). Intuitively, the estimation of the true equity premium is more difficult and unreliable for those investors who are less confident about their grasp of the environment; the consequent lack of confidence in the

[20] Closed-form solutions are still not available.

estimate combined with the (apparent) pessimism leads those investors, then, to try to transfer wealth even more to unfavorable states.

Lack of confidence can have significant effects on portfolio choice. When the estimated instantaneous premium is zero, for example, the difference between the demand that takes lack of confidence into account and the demand that does not can be as large as half of wealth. And the hedging of ambiguity persists in the present model,[21] because neither the current value of the hidden state nor the perturbations to the steady state can be learned.

2.3 Related Papers

As should be clear by now, the present paper is closely related to Epstein and Schneider (2007), in which the set of predictive measures is also constructed by a statistical test over multiple theories. The main differences are twofold. First, whereas Epstein and Schneider consider (non-Bayesian) exchangeable data-generating mechanisms, I consider data-generating mechanisms with serial dependence and possibly even with path dependence. Second, whereas Epstein and Schneider's model is set in discrete time, mine is set in continuous time. Thus, we can relate the present model to the abundance of continuous-time models with mean reversion; and exploit the analytical advantages continuous-time modeling affords. I also note that the continuous-time counterpart of Epstein and Schneider's portfolio choice example results in no learning because the likelihood function degenerates to infinity everywhere (see the supplementary appendix), and consequently, their discrete-time finding that learning resolves ambiguity does not immediately carry over to continuous time: learning under ambiguity in continuous time needs separate treatment.[22]

Another paper that considers the learning of a multiple-priors agent is Miao (2009). Specifically, he considers the consumption/portfolio choice problem of a multiple-priors investor in continuous time who partially observes stochastic investment opportunities. However, his notion of learning is fundamentally different from mine. Miao's investor obtains a benchmark predictive measure by updating a reference theory, and the set of predictive measures is given by a neighborhood of the benchmark with a fixed radius. Thus, learning and ambiguity do not interact. In fact, Miao's model is the limit of the present model as the investor gains confidence (Section 4.4.3).[23][24]

[21] Epstein and Schneider do not discuss whether it persists or not in their model.

[22] Campanale (2011) applies Epstein and Schneider's (2007) model in the context of life-cycle portfolio choice and Miao and Wang (2011) in the context of job matching. Epstein and Schneider's (2008) asset pricing model conforms to the formalism of Epstein and Schneider (2007) but their agents do not discard any of the models they a priori entertain.

[23] Liu (2011) considers the consumption/portfolio choice problem of a Miao investor when expected returns follow a Markov chain.

[24] See also Hansen and Sargent (2011) and Chen et al. (2014). The former paper considers learning in the context of robust control, and the latter, in the context of smooth ambiguity. In both papers, agents learn by updating a Bayesian model. Chen et al.'s agent, in particular, has a standard hierarchical Bayesian model; in a smooth ambiguity model, aversion to ambiguity is captured not by imprecise probabilities but by failures to reduce compound lotteries. Incidentally, Chen et al. (2014) is also motivated by ambiguous predictability.

3 The Model of Learning under Ambiguity

3.1 Preferences: Recursive Multiple-Priors

I assume that the agent has Chen and Epstein's (2002) recursive multiple-priors utility.

Specifically, time is continuous and varies over $[0, T]$, $T \in (0, \infty)$.

Let Ω denote the set of states of Nature and let a filtration $\mathbf{G} = \{\mathcal{G}_t\}$ on Ω represent the accrual of the agent's information. $\mathbf{G}$ is right-continuous.

There is a set $\mathcal{P}$ of equivalent probability measures on $(\Omega, \mathcal{G}_T)$, the set of priors, with the following properties. Let $P^0 \in \mathcal{P}$. There is an n_y-dimensional Wiener process $\epsilon = \{\epsilon(t), \mathcal{G}_t\}$ under P^0 that generates $\mathbf{G}$. (The notation $\epsilon = \{\epsilon(t), \mathcal{G}_t\}$ signifies that the process ϵ is adapted to the filtration $\mathbf{G}$. All vectors, including the gradient ∂f of a scalar function f, are column vectors. Hence, $\epsilon(t) = (\epsilon_1(t), \cdots, \epsilon_{n_y}(t))^\top$.) $\mathcal{G}_0$ contains all the P^0-null events in $\mathcal{G}_T$. Thus, in particular, $\mathbf{G}$ satisfies the usual conditions. Each prior is identified with the corresponding density generator $\xi = \{\xi(t), \mathcal{G}_t\}$ and is thus written $P^\xi \in \mathcal{P}$, where

$$\frac{\mathrm{d}P^\xi}{\mathrm{d}P^0} = \mathcal{E}^\xi(T) \tag{12}$$

and $\mathcal{E}^\xi$ denotes the Doléans-Dade exponential

$$\mathcal{E}^\xi(t) \triangleq \exp\left(\int_0^t \xi(s)\,\mathrm{d}\epsilon(s) - \frac{1}{2}\int_0^t |\xi(s)|^2\,\mathrm{d}s\right),\ 0 \le t \le T.$$

$\mathcal{P}$ is required to be *rectangular*, which means that there is a set-valued process $\Xi : [0, T] \times \Omega \to 2^{\mathbb{R}^{n_y}}$ such that the probability measure P^ξ defined by (12) is a prior if and only if ξ is a $\mathbf{G}$-progressive process and $\xi(t, \omega) \in \Xi(t, \omega)$ for Lebesgue$\times P^0$ almost every (t, ω). Since $\mathcal{P}$ consists of equivalent measures, "for Lebesgue$\times P^0$ almost every (t, ω)" is henceforth abbreviated without ambiguity to "a.e." Ξ is called the *one-step-ahead beliefs process* and is further required to be (i) uniformly bounded, (ii) compact-convex-valued, and (iii) "$\mathbf{G}$-progressive": (i) $\Xi(t, \omega) \subset K$ a.e. for some bounded $K \subset \mathbb{R}^{n_y}$, (ii) $\Xi(t, \omega)$ is compact-convex a.e., and (iii) the restriction of Ξ to $[0, t] \times \Omega$ is $\mathcal{B}[0, t] \otimes \mathcal{G}_t$-measurable[25] for all $t \in [0, T]$ where $\mathcal{B}X$ denotes the Borel σ-algebra of a topological space X. Incidentally, IID ambiguity refers to the situation where Ξ is constant, that is, $\Xi(t, \omega) = K$ a.e. for some compact-convex $K \subset \mathbb{R}^{n_y}$.

A scalar process $c = \{c(t), \mathcal{G}_t\}$ is a *consumption process* if it is progressive, positive, and integrable. Denote the set of consumption processes by $\mathcal{C}$. The agent's conditional preferences at time t, $t \in [0, T]$, are represented by

$$U^c(t, \omega) = \min_{P \in \mathcal{P}} U^{c,P}(t, \omega),\ c \in \mathcal{C} \tag{13}$$

where $U^{c,P} = \{U^{c,P}(s), \mathcal{G}_s\}$, the utility process under $P \in \mathcal{P}$, uniquely solves the backward stochastic differential equation (BSDE)

$$U^{c,P}(s) = \mathrm{E}^P\left(\int_s^T F(c(\tau), U^{c,P}(\tau))\,\mathrm{d}\tau \,\middle|\, \mathcal{G}_s\right),\ t \le s \le T.$$

[25] $\{(s, \omega) \in [0, t] \times \Omega : \Xi(s, \omega) \cap K' \neq \emptyset\} \in \mathcal{B}[0, t] \otimes \mathcal{G}_t$ for all closed $K' \subset K$. See Aliprantis and Border (1999), Sections 16.1, 16.2, and 17.1.

Here, F is the aggregator. See Chen and Epstein (2002), Section 2.5 for the conditions that the aggregator has to satisfy.

Rectangularity is an essential requirement for the Epstein-Schneider recursive representation and hence for their notion of dynamic consistency.[26] The utility process U^c defined by (13) satisfies

$$\mathrm{d}U^c(t) = \left(-F(c(t), U^c(t)) + \max_{\xi(t) \in \Xi(t)} \sigma_U^c(t)\xi(t) \right) \mathrm{d}t + \sigma_U^c(t)\,\mathrm{d}\epsilon(t),\ \ 0 \le t \le T \tag{14}$$

with terminal condition $U(T) = 0$, for some process $\sigma_U^c = \{\sigma_U^c(t), \mathcal{G}_t\}$. When (14) is viewed as a BSDE, the pair (U^c, σ_U^c) constitutes a solution.

To interpret (14), rewrite it, with a slight abuse of notation, as

$$\mathrm{d}U^c(t) = -F(c(t), U^c(t))\,\mathrm{d}t + \max_{\xi(t) \in \Xi(t)} \sigma_U^c(t)(\,\mathrm{d}\epsilon(t) + \xi(t)\,\mathrm{d}t)$$

and compare it with Duffie and Epstein's (1992) single-prior representation

$$\mathrm{d}U^c(t) = -F(c(t), U^c(t))\,\mathrm{d}t + \sigma_U^c(t)\,\mathrm{d}\epsilon(t).$$

Consider first the single-prior case. By the assumption that ϵ generates $\mathbf{G}$, all changes in the fundamentals to take place over the infinitesimal future are a function of $\mathrm{d}\epsilon(t)$; that is, the agent's conditional beliefs about the uncertainties to be resolved over the next instant are summarized by the unique distribution $N(0,\,\mathrm{d}t)$ for the one-step-ahead noise[27] $\mathrm{d}\epsilon(t)$. Now, the multiple-priors representation suggests the interpretation that the agent entertains a *set* of distributions $N(\xi(t),\,\mathrm{d}t)$, $\xi(t) \in \Xi(t)$; and a pessimist, he assesses each consumption process under the belief that the distribution that is the worst for the process is the case.

3.2 The Theories

Denote the observable process that generates the agent's information by $y = \{y(t), \mathcal{G}_t\}$. That is, $\mathbf{G}$ is the P^0-augmentation of the filtration generated by y. n_y denotes the dimension of y. Examples of y will be given shortly.

In this section, I define the theories that the agent entertains about how y is generated. They are given by a set of probability measures $Q \in \mathcal{Q}$ on a common measurable space.

3.2.1 The Reference Likelihood

Let there be a filtration $\mathbf{F} = \{\mathcal{F}_t\}$ on Ω and a probability measure $Q^{\bar{x},0}$ on $(\Omega, \mathcal{F}_T)$ where $\bar{x} \in \mathbb{R}^{n_x}$, $n_x \ge 1$. $\mathbf{F}$ satisfies the usual conditions with respect to $Q^{\bar{x},0}$. Let there also be two independent $Q^{\bar{x},0}$-Wiener processes $w = \{w(t), \mathcal{F}_t\}$ and $v^{\bar{x},0} = \{v^{\bar{x},0}(t), \mathcal{F}_t\}$, n_y-dimensional and n_x-dimensional, respectively. Under $Q^{\bar{x},0}$, y satisfies the following system of SDEs:

$$\begin{aligned} \mathrm{d}y(t) &= (a(t,y) + b(t,y)x(t))\,\mathrm{d}t + \sigma(t,y)\,\mathrm{d}w(t), \\ \mathrm{d}x(t) &= \kappa(\bar{x} - x(t))\,\mathrm{d}t + \rho_w\,\mathrm{d}w(t) + \rho_v\,\mathrm{d}v^{\bar{x},0}(t). \end{aligned}$$

[26] For a detailed discussion on rectangularity and its connection to dynamic consistency, see Epstein and Schneider (2003), Sections 3.1 and 5.2.

[27] See Footnote 14.

Here, $x = \{x(t), \mathcal{F}_t\}$ is an n_x-dimensional process that is unobservable to the agent; $a : [0,T]\times C([0,T],\mathbb{R}^{n_y}) \to \mathbb{R}^{n_y}$, $b : [0,T]\times C([0,T],\mathbb{R}^{n_y}) \to \mathbb{R}^{n_y\times n_x}$, and $\sigma : [0,T]\times C([0,T],\mathbb{R}^{n_y}) \to \mathbb{R}^{n_y\times n_y}$ are nonanticipating path functionals where $C([0,T],\mathbb{R}^{n_y})$ denotes the set of continuous functions from $[0,T]$ into $\mathbb{R}^{n_y}$;[28] κ is an $n_x \times n_x$ diagonal matrix with positive entries, ρ_w is an $n_x \times n_y$ matrix, and ρ_v is an $n_x \times n_x$ invertible matrix. Given that y is observed, the diffusion matrix process $\sigma\sigma^\top$, too, is observed via quadratic variation-covariation. The assumption that σ as a nonanticipating path functional depends only on y, or equivalently, σ as a process is adapted to $\mathbf{G}$, embodies the restriction that observing the diffusion matrix does not expand the agent's information. $x(0)$ is an $\mathcal{F}_0$-measurable random variable. The distribution of $x(0)$ conditional on $\mathcal{G}_0$ is normal with mean $m_0 \in \mathbb{R}^{n_x}$ and variance-covariance matrix $\gamma_0 \in \mathbb{R}^{n_x\times n_x}$. For simplicity, I assume $y(0)$ is nonrandom.

All the parameters and functionals are known but $\bar{x}$. While this assumption may seem unrealistic—if he knows so much, why not $\bar{x}$? or vice versa—I point out that (i) in many special cases considered in the literature (see the examples below), the functionals are simple, being constant or linear, and (ii) the restrictive form of ignorance is only a first step: the agent may well find κ ambiguous as well, for example.

Example 3.1 (Stock Returns with Constant Volatility)**.** *Suppose that the cumulative return process R of a stock satisfies*

$$\mathrm{d}R(t) = x(t)\,\mathrm{d}t + \sigma_R\,\mathrm{d}w(t)$$

and R is the only observable process. Then, $y = R$ with $a \equiv 0$, $b \equiv 1$, and $\sigma \equiv \sigma_R > 0$.

Example 3.2 (Stock Returns with Stochastic Volatility)**.** *Suppose that the conditional return variance follows a Cox-Ingersoll-Ross process (Heston, 1993), that is,*

$$\begin{aligned}\mathrm{d}R(t) &= x(t)\,\mathrm{d}t + \sqrt{A(t)}(\sqrt{1-r_{RA}^2}, r_{RA})\,\mathrm{d}w(t),\\ \mathrm{d}A(t) &= \nu(\bar{A} - A(t))\,\mathrm{d}t + \varsigma_A\sqrt{A(t)}(0,1)\,\mathrm{d}w(t),\end{aligned}$$

where $r_{RA} \in (-1,1)$, $\nu > 0$, $\bar{A} \in \mathbb{R}$, and $\varsigma_A > 0$. Suppose also that R and A are the only observable processes. Then, $y = (R, A)^\top$ with

$$a(t,y) = \begin{pmatrix} 0 \\ \nu(\bar{A} - A(t)) \end{pmatrix},\ b \equiv \begin{pmatrix} 1 \\ 0 \end{pmatrix},\ \text{and } \sigma(t,y) = \sqrt{A(t)}\begin{pmatrix} \sqrt{1-r_{RA}^2} & r_{RA} \\ 0 & \varsigma_A \end{pmatrix}.$$

Example 3.3 (Extra Signal)**.** *Suppose that the return volatility is constant as in Example 3.1, but now there is an extra signal about the hidden state x in addition to R (Detemple, 1986; Veronesi, 2000):*

$$\begin{aligned}\mathrm{d}R(t) &= x(t)\,\mathrm{d}t + \sigma_R(\sqrt{1-r_{RA}^2}, r_{RA})\,\mathrm{d}w(t),\\ \mathrm{d}A(t) &= x(t)\,\mathrm{d}t + \sigma_A(0,1)\,\mathrm{d}w(t),\end{aligned}$$

[28] Let ι be the canonical process on $C([0,T],\mathbb{R}^{n_y})$; that is, $\iota(t,f) = f(t)$, $0 \le t \le T$, $f \in C([0,T],\mathbb{R}^{n_y})$. Let $\mathcal{B}_t \triangleq \sigma(\iota(s) : 0 \le s \le t)$ and $\mathcal{B}_{t+} \triangleq \cap_{s>t}\mathcal{B}_s$ with $\mathcal{B}_{T+} \triangleq \mathcal{B}_T$. a, b, and σ are measurable and adapted to $\{\mathcal{B}_{t+}\}$.

where $r_{RA} \in (-1,1)$ *and* $\sigma_R, \sigma_A > 0$. *Then,* $y = (R, A)^\top$ *with*

$$a \equiv \begin{pmatrix} 0 \\ 0 \end{pmatrix}, \; b \equiv \begin{pmatrix} 1 \\ 1 \end{pmatrix}, \; \textit{and } \sigma \equiv \begin{pmatrix} \sigma_R\sqrt{1 - r_{RA}^2} & \sigma_R r_{RA} \\ 0 & \sigma_A \end{pmatrix}.$$

Similar examples can be constructed in a general equilibrium setting in which, for example, the aggregate consumption growth replaces the stock returns.

Now, the reference likelihood function of the unknown parameter $\bar{x}$ under full observation, or simply the *reference likelihood*, is defined by

$$L_{\mathrm{FO},T}(\bar{x}) \triangleq \frac{\mathrm{d}Q^{\bar{x},0}}{\mathrm{d}Q^{0,0}}.$$

3.2.2 Bayesian Benchmark

Let M be a probability measure on $(\mathbb{R}^{n_x}, \mathcal{B}\mathbb{R}^{n_x})$. Then, $(M, L_{\mathrm{FO},T})$ defines a Bayesian model of data generation, according to which $\bar{x}$ is drawn from M and conditional on $\bar{x}$ the law of (y, x) is given by $L_{\mathrm{FO},T}(\bar{x})$.

3.2.3 The Theories

Bayesian agents behave as if they knew the probabilities of all relevant events precisely. Consider in contrast an agent who lacks confidence in his understanding of the data-generating mechanism and finds both the parameter $\bar{x}$ and the reference likelihood $L_{\mathrm{FO},T}$ ambiguous.

Specifically, the agent's perception of ambiguity regarding $\bar{x}$ is expressed by multiplicity of parameter priors. For simplicity, I assume that the parameter priors are all Dirac measures:[29]

$$\mathcal{M} = \{\mathrm{Dirac}^{\bar{x}'} : \bar{x}' \in \mathbb{R}^{n_x}\}$$

where $\mathrm{Dirac}^{\bar{x}'}$ denotes the Dirac measure concentrated at $\bar{x}' \in \mathbb{R}^{n_x}$.

Similarly, the agent also entertains multiple likelihoods. Fix $\bar{x}$. Let there be a probability measure $Q^{\bar{x},\eta}$, $\eta \in L^2([0,T], \mathbb{R}^{n_x})$, on $(\Omega, \mathcal{F}_T)$ where $L^2([0,T], \mathbb{R}^{n_x})$ denotes the set of square-integrable, $\mathbb{R}^{n_x}$-valued functions. Let there also be an n_x-dimensional Wiener process $v^{\bar{x},\eta} = \{v^{\bar{x},\eta}(t), \mathcal{F}_t\}$ independent of w. Under $Q^{\bar{x},\eta}$, (y, x) satisfies the following system of SDEs:

$$\mathrm{d}y(t) = (a(t,y) + b(t,y)x(t))\,\mathrm{d}t + \sigma(t,y)\,\mathrm{d}w(t), \tag{15}$$

$$\mathrm{d}x(t) = \kappa(\bar{x} + \kappa^{-1}\rho_v\eta(t) - x(t))\,\mathrm{d}t + \rho_w\,\mathrm{d}w(t) + \rho_v\,\mathrm{d}v^{\bar{x},\eta}(t), \tag{16}$$

where, as before, $y(0) \in \mathbb{R}^{n_y}$ is nonrandom and the $\mathcal{F}_0$-measurable random variable $x(0)$ has the conditional distribution $x(0)|\mathcal{G}_0 \sim N(m_0, \gamma_0)$. The set of full-observation likelihoods is given by

$$\mathcal{L}_{\mathrm{FO},T} = \left\{\bar{x} \mapsto L_{\mathrm{FO},T}(\bar{x}, \eta) : \eta \in L^2([0,T], \mathbb{R}^{n_x})\right\},$$

$$L_{\mathrm{FO},T}(\bar{x}, \eta) \triangleq \frac{\mathrm{d}Q^{\bar{x},\eta}}{\mathrm{d}Q^{0,0}}.$$

[29]See Footnote 13.

(16) can alternatively be written as

$$\mathrm{d}x(t) = \kappa(\bar{x} - x(t))\,\mathrm{d}t + \rho_w\,\mathrm{d}w(t) + \rho_v(\,\mathrm{d}v^{\bar{x},\eta}(t) + \eta(t)\,\mathrm{d}t)$$

which shows that the ambiguity in question is equivalent to that in the noise $v^{\bar{x},\eta}$ specific to the state dynamics and not the other one w. Indeed, the agent of this paper is not mechanically taking into consideration all the theories that are close to a reference, in which case w, too, would be perturbed, but is rather questioning a particular aspect of market dynamics, namely, mean reversion.

Now I turn to the issue of the existence and uniqueness of a solution to the system of SDEs (15) and (16). $|\cdot|$ denotes the Euclidean norm for vectors and the Frobenius norm for matrices; that is, for a vector or matrix z, $|z| \triangleq \sqrt{\mathrm{tr}(zz^\top)}$. All numbered assumptions stand throughout the paper from their statement on, unless otherwise noted.

Assumption 3.1 (Sufficient Conditions for Unique Strong Existence)**.** *(i) b is uniformly bounded.*

(ii) For all $f \in C([0,T], \mathbb{R}^{n_y})$,

$$\int_0^T (|a(t,f)| + |\sigma(t,f)|^2)\,\mathrm{d}t < \infty.$$

(iii) a, b, and σ are locally Lipschitz. That is, for each N there is a K_N such that

$$\left(\sup_{s\le t}|f(s)|\right) \vee \left(\sup_{s\le t}|g(s)|\right) \le N \Rightarrow |\sigma(t,f) - \sigma(t,g)| \le K_N \sup_{s\le t}|f(s) - g(s)|$$

for all $t \in [0,T]$; and the same for a and b mutatis mutandis.

(iv) a and σ are linearly growing. That is, there is a K such that

$$|a(t,f)| + |\sigma(t,f)| \le K\left(1 + \sup_{s\le t}|f(s)|\right)$$

for all $(t,f) \in [0,T] \times C([0,T], \mathbb{R}^{n_y})$.

Proposition 3.1. *Strong existence and pathwise uniqueness hold for the system of SDEs* (15)*-*(16)*.*

Suppose $(w, v^{\bar{x},\eta})$ and (y,x) are defined on some filtered complete probability space. With a slight abuse of notation, $\sigma(t) \equiv \sigma(t,\omega) \equiv \sigma(t, y(\omega))$. With the notation $\sigma(t,\omega)$, σ can be considered a process (adapted to $\mathbf{G}$). Similar remarks apply to the other functionals. As is the custom, the qualification almost surely is suppressed unless necessary.

Assumption 3.2. *There is an $\varepsilon > 0$ such that*

$$z^\top \sigma(t)\sigma(t)^\top z \ge \varepsilon|z|^2 \text{ for all } z \in \mathbb{R}^{n_y} \text{ and all } t \in [0,T].$$

Remark 3.1. *The stochastic volatility model (Example 3.2) violates Assumption 3.2:*

$$\sigma(t)\sigma(t)^\top = A(t)\begin{pmatrix} 1 & r_{RA}\varsigma_A \\ \varsigma_A r_{RA} & \varsigma_A^2 \end{pmatrix}$$

and the last matrix that post-multiplies $A(t)$ satisfies Assumption 3.2,[30] *but $A(t)$ may get arbitrarily close to 0. In general, if Assumption 3.2 fails, then the subsequent findings based on the assumption hold up to the random time*

$$T \wedge \inf\left\{t > 0 : z^\top \sigma(s)\sigma(s)^\top z \geq N^{-1}|z|^2 \text{ for all } z \in \mathbb{R}^{n_y} \text{ and all } s \leq t\right\}.$$

Assumption 3.2 implies that $\sigma(t)$ has an inverse and $|\sigma(t)^{-1}z| \leq K^{-1/2}|z|$ for all $z \in \mathbb{R}^{n_y}$ and all $t \in [0,T]$; likewise, $\sigma(t)^\top$, too, has an inverse and $|(\sigma(t)^\top)^{-1}z| \leq K^{-1/2}|z|$ for all $z \in \mathbb{R}^{n_y}$ and all $t \in [0,T]$ (Karatzas and Shreve (1988), Problem 5.8.1). With the last observation, we can rewrite (15) and (16) as

$$\mathrm{d}w(t) = \sigma(t)^{-1}[\,\mathrm{d}y(t) - (a(t) + b(t)x(t))\,\mathrm{d}t], \tag{17}$$

$$\mathrm{d}v^{\bar{x},\eta}(t) = \rho_v^{-1}\left\{\mathrm{d}x(t) - \kappa(\bar{x} - x(t))\,\mathrm{d}t - \rho_w\sigma(t)^{-1}[\,\mathrm{d}y(t) - (a(t) + b(t)x(t))\,\mathrm{d}t]\right\} - \eta(t)\,\mathrm{d}t, \tag{18}$$

and use these SDEs to *define* w and $v^{\bar{x},\eta}$.

Let

$$\Omega \triangleq C([0,T],\mathbb{R}^{n_y}) \times C([0,T],\mathbb{R}^{n_x}),$$
$$\mathcal{F}^\circ \triangleq \mathcal{B}C([0,T],\mathbb{R}^{n_y}) \otimes \mathcal{B}C([0,T],\mathbb{R}^{n_x}),$$

let (y,x) be the identity map on Ω, and let

$$Q^{\bar{x},\eta} \triangleq \mathrm{law}(y,x)$$

be defined on $(\Omega,\mathcal{F}^\circ)$. Let $\mathbf{F} = \{\mathcal{F}_t\}$ be the augmented filtration generated by (y,x). Since $Q^{\bar{x},\eta}$, $(\bar{x},\eta) \in \mathbb{R}^{n_x} \times L^2([0,T],\mathbb{R}^{n_x})$, are equivalent, they all lead to the same augmentation. Finally, define σ by $\sigma(t,\omega) = \sigma(t,y(\omega))$; a and b similarly; and w and $v^{\bar{x},\eta}$ by (17) and (18). In particular, this construction (of weak solutions) explains why $v^{\bar{x},\eta}$ is superscripted.

In sum, the agent's theories of how the data y is generated can be identified with the probability measures

$$\mathcal{Q} \triangleq \left\{Q^{\bar{x},\eta} : (\bar{x},\eta) \in \mathbb{R}^{n_x} \times L^2([0,T],\mathbb{R}^{n_x})\right\}$$

[30] The question is if there is a (small) $\varepsilon > 0$ such that for all $z = (z_1, z_2)$,

$$\begin{aligned} 0 &\leq z_1^2 + 2z_1z_2\varsigma_A r_{RA} + z_2^2\varsigma_A^2 - \varepsilon(z_1^2 + z_2^2) \\ &= (1-\varepsilon)\left(\left(z_1 + \frac{z_2\varsigma_A r_{RA}}{1-\varepsilon}\right)^2 - \left(\frac{z_2\varsigma_A r_{RA}}{1-\varepsilon}\right)^2 + z_2^2\frac{\varsigma_A^2 - \varsigma_A}{1-\varepsilon}\right). \end{aligned}$$

The last inequality holds for all z if and only if

$$\varsigma_A^2(1 - \varepsilon - r_{RA}^2) \geq \varepsilon(1-\varepsilon).$$

Hence, the desired ε exists if and only if $r_{RA}^2 < 1$, which is assumed.

on the common measurable space $(\Omega, \mathcal{F}_T)$, all defined above. I call these probability measures *theoretical priors* to distinguish them from the probability measures $P \in \mathcal{P}$ that are part of the representation of the agent's preferences, or the *preferential priors*. Not only are these two types of priors conceptually distinct, but they *are* different; we will see (in Section 3.4.4) that $\mathcal{P} \nsubseteq \{Q|_{\mathcal{G}_T} : Q \in \mathcal{Q}\}$ where $Q|_{\mathcal{G}_T}$ denotes the restriction of Q to $\mathcal{G}_T$.

3.3 The Preferential Priors

3.3.1 Filtering

Recall the partially observable system

$$
\begin{aligned}
\mathrm{d}y(t) &= (a(t) + b(t)x(t))\,\mathrm{d}t + \sigma(t)\,\mathrm{d}w(t),\\
\mathrm{d}x(t) &= \kappa(\bar{x} - x(t))\,\mathrm{d}t + \rho_w\,\mathrm{d}w(t) + \rho_v(\,\mathrm{d}v^{\bar{x},\eta}(t) + \eta(t)\,\mathrm{d}t).
\end{aligned}
$$

I use the dot notation for the time derivatives: $\dot{f}(t) \equiv \mathrm{d}f(t)/\,\mathrm{d}t$.

Proposition 3.2. *The following standard results in Gaussian filtering hold:*
(i) (y, x) is conditionally Gaussian.
(ii) The conditional mean vector and variance-covariance matrix

$$
\begin{aligned}
m^{\bar{x},\eta}(t) &\triangleq \mathrm{E}^{Q^{\bar{x},\eta}}(x(t)|\mathcal{G}_t),\\
\gamma(t) &\triangleq \mathrm{E}^{Q^{\bar{x},\eta}}[(x(t) - m^{\bar{x},\eta}(t))(x(t) - m^{\bar{x},\eta}(t))^\top|\mathcal{G}_t],
\end{aligned}
$$

satisfy the system of differential equations

$$
\begin{aligned}
\mathrm{d}m^{\bar{x},\eta}(t) &= [\kappa(\bar{x} - m^{\bar{x},\eta}(t)) + \rho_v\eta(t)]\,\mathrm{d}t\\
&\quad + (\rho_w\sigma(t)^\top + \gamma(t)b(t)^\top)(\sigma(t)\sigma(t)^\top)^{-1}[\,\mathrm{d}y(t) - (a(t) + b(t)m^{\bar{x},\eta}(t))\,\mathrm{d}t]\\
&= (\kappa\bar{x} + \rho_v\eta(t) - \bar{\kappa}(t)m^{\bar{x},\eta}(t))\,\mathrm{d}t\\
&\quad + (\rho_w\sigma(t)^\top + \gamma(t)b(t)^\top)(\sigma(t)\sigma(t)^\top)^{-1}(\,\mathrm{d}y(t) - a(t)\,\mathrm{d}t),
\end{aligned}
\tag{19}
$$

$$
\begin{aligned}
\dot{\gamma}(t) &= \rho_w\rho_w^\top + \rho_v\rho_v^\top - \kappa\gamma(t) - \gamma(t)\kappa\\
&\quad - (\rho_w\sigma(t)^\top + \gamma(t)b(t)^\top(\sigma(t)\sigma(t)^\top)^{-1}(\rho_w\sigma(t)^\top + \gamma(t)b(t)^\top)^\top,
\end{aligned}
\tag{20}
$$

with initial conditions $m^{\bar{x},\eta}(0) = m_0$ and $\gamma(0) = \gamma_0$, where

$$
\bar{\kappa}(t) \triangleq \kappa + (\rho_w\sigma(t)^\top + \gamma(t)b(t)^\top)(\sigma(t)\sigma(t)^\top)^{-1}b(t).
$$

(iii) the process $\bar{w}^{\bar{x},\eta} = \{\bar{w}^{\bar{x},\eta}(t), \mathcal{G}_t\}$ defined by

$$
\bar{w}^{\bar{x},\eta}(t) \triangleq \int_0^t \sigma(s)^{-1}[\,\mathrm{d}y(s) - (a(s) + b(s)m^{\bar{x},\eta}(s))\,\mathrm{d}s],\ 0 \le t \le T
$$

is a Wiener process under $Q^{\bar{x},\eta}$ and generates $\mathbf{G}$.

Lemma 3.1. *γ is uniformly bounded.*

Let $\varphi : [0,T] \times \Omega \to \mathbb{R}^{n_x \times n_x}$ be the solution of

$$\dot{\varphi}(t) = -\bar{\kappa}(t)\varphi(t),\ \varphi(0) = I_{n_x}$$

where I_{n_x} denotes the n_x-dimensional identity matrix. $\varphi(t)$ is invertible for all $t \geq 0$. Introduce the following notation: for functions f from $[0,T]$ into $\mathbb{R}^{n_x}$ or into $\mathbb{R}^{n_x \times n_x}$, Φ^f denotes the process defined by

$$\Phi^f(t) \triangleq \varphi(t) \int_0^t \varphi(s)^{-1} f(s)\,\mathrm{d}s,\ 0 \leq t \leq T.$$

Now

$$\begin{aligned} m^{\bar{x},\eta}(t) &= \varphi(t) \Big\{ m_0 + \int_0^t \varphi(s)^{-1} [(\kappa\bar{x} + \rho_v \eta(s))\,\mathrm{d}s \\ &\quad + (\rho_w \sigma(s)^\top + \gamma(s) b(s)^\top)(\sigma(s)\sigma(s)^\top)^{-1}(\,\mathrm{d}y(s) - a(t)\,\mathrm{d}t)] \Big\} \\ &= \varphi(t) m_0 + \Phi^{\kappa\bar{x}+\rho_v\eta}(t) \\ &\quad + \varphi(t) \int_0^t \varphi(s)^{-1}(\rho_w \sigma(s)^\top + \gamma(s) b(s)^\top)(\sigma(s)\sigma(s)^\top)^{-1}(\,\mathrm{d}y(s) - a(t)\,\mathrm{d}t). \end{aligned} \tag{21}$$

3.3.2 Likelihood of Theories

The log-likelihood of theories under partial observation, that is, given $\mathcal{G}_T$, is

$$\begin{aligned} \ell_T(\bar{x},\eta) &\triangleq \log \frac{\mathrm{d}(Q^{\bar{x},\eta}|_{\mathcal{G}_T})}{\mathrm{d}(Q^{0,0}|_{\mathcal{G}_T})} \\ &= \log \mathrm{E}^{Q^{0,0}} \left(\frac{\mathrm{d}Q^{\bar{x},\eta}}{\mathrm{d}Q^{0,0}} \middle| \mathcal{G}_T \right) \end{aligned} \tag{22}$$

where $Q^{\bar{x},\eta}|_{\mathcal{G}_T}$ denotes the restriction of $Q^{\bar{x},\eta}$ to $\mathcal{G}_T$. The choice of the reference, here $(\bar{x},\eta) = (0,0)$, is irrelevant to likelihood ratios.

Proposition 3.3.

$$\begin{aligned} \ell_T(\bar{x},\eta) &= \int_0^T (a(t) + b(t) m^{\bar{x},\eta}(t))^\top (\sigma(t)\sigma(t)^\top)^{-1}\,\mathrm{d}y(t) \\ &\quad - \frac{1}{2}\int_0^T (a(t) + b(t) m^{\bar{x},\eta}(t))^\top (\sigma(t)\sigma(t)^\top)^{-1} (a(t) + b(t) m^{\bar{x},\eta}(t))\,\mathrm{d}t \\ &\quad - \Big(\int_0^T (a(t) + b(t) m^{0,0}(t))^\top (\sigma(t)\sigma(t)^\top)^{-1}\,\mathrm{d}y(t) \\ &\quad - \frac{1}{2}\int_0^T (a(t) + b(t) m^{0,0}(t))^\top (\sigma(t)\sigma(t)^\top)^{-1} (a(t) + b(t) m^{0,0}(t))\,\mathrm{d}t \Big). \end{aligned}$$

The log-likelihood given $\mathcal{G}_t$, $t < T$, is obtained by replacing the arbitrary time horizon T with t:

$$\ell_t(\bar{x}, \eta) = \int_0^t (a(s) + b(s) m^{\bar{x},\eta}(s))^\top (\sigma(s)\sigma(s)^\top)^{-1} \, \mathrm{d}y(s) \\ - \frac{1}{2} \int_0^t (a(s) + b(s) m^{\bar{x},\eta}(s))^\top (\sigma(s)\sigma(s)^\top)^{-1} (a(s) + b(s) m^{\bar{x},\eta}(s)) \, \mathrm{d}s + f_t$$

where f_t is independent of $(\bar{x}, \eta)$.

Since $m^{\bar{x},\eta}(t)$ is linear in $\bar{x}$, $\ell_t(\bar{x}, \eta)$ is quadratic in $\bar{x}$. $\ell_t(\bar{x}, \eta)$ is also Gâteaux differentiable with respect to η and the derivative is linear in η:

Lemma 3.2. *The Gâteaux differential of $\ell_t(\bar{x}, \cdot)$ at $\eta \in L^2([0,T], \mathbb{R}^{n_x})$ in the direction $h \in L^2([0,T], \mathbb{R}^{n_x})$ is*

$$\int_0^t \Bigg((\varphi(s)^{-1} \rho_v)^\top \int_s^t \varphi(\tau)^\top b(\tau)^\top (\sigma(\tau)\sigma(\tau)^\top)^{-1} \\ \times \left[\mathrm{d}y(\tau) - (a(\tau) + b(\tau) m^{\bar{x},\eta}(\tau)) \, \mathrm{d}\tau \right] \Bigg)^\top h(s) \, \mathrm{d}s.$$

3.3.3 Learning

Recall the dynamics of the observable process

$$\mathrm{d}y(t) = (a(t) + b(t)x(t)) \, \mathrm{d}t + \sigma(t) \, \mathrm{d}w(t).$$

If the agent were a Bayesian with unique theoretical prior $Q^{\bar{x},0} \in \mathcal{Q}$,[31] then Bayesian updating would result in the filtered dynamics

$$\mathrm{d}y(t) = (a(t) + b(t) m^{\bar{x},0}(t)) \, \mathrm{d}t + \sigma(t) \, \mathrm{d}\bar{w}(t), \tag{23}$$

where $\bar{w}$, defined by (23), is a $(Q^{\bar{x},0}, \mathbf{G})$-Wiener process, and his time-t decisions would accordingly be based on the unique one-step-ahead conditional

$$\mathrm{d}y(t) | \mathcal{G}_t \sim N \left[(a(t) + b(t) m^{\bar{x},0}(t)) \, \mathrm{d}t, \sigma(t)\sigma(t)^\top \, \mathrm{d}t \right].$$

On the other hand, our agent entertains a set of theories, $\{Q^{\bar{x},\eta} : (\bar{x}, \eta) \in \mathbb{R}^{n_x} \times L^2([0,T], \mathbb{R}^{n_x})\}$, and rules out some of them in light of evidence. Hence, unless he rules out all but one theory, the agent will have *multiple* one-step-ahead conditionals of the form

$$\mathrm{d}y(t) | \mathcal{G}_t \sim N \left[(a(t) + b(t) m^{\bar{x},\eta}(t)) \, \mathrm{d}t, \sigma(t)\sigma(t)^\top \, \mathrm{d}t \right]$$

where $(\bar{x}, \eta)$ runs over a set. Note that the ambiguity in the data-generating mechanism boils down to that in the conditional expectation of $x(t)$.

[31] See Footnote 15.

Plausibility: Penalized Likelihood The ambiguity, however, is too large for there to be learning, if the agent assesses the plausibility of a theory based on the likelihood alone. To elaborate, define the log-likelihood induced by the transformation $(\bar{x}, \eta) \mapsto m^{\bar{x},\eta}(t)$ by

$$\ell_{t,m(t)}(m) \triangleq \sup_{(\bar{x},\eta)\in\mathbb{R}^{n_x}\times L^2([0,T],\mathbb{R}^{n_x})} \{\ell_t(\bar{x},\eta) : m^{\bar{x},\eta}(t) = m\}, \ m \in \mathbb{R}^{n_x}.$$

Then, $\ell_{t,m(t)}$ is constant, the constant value lying in $\mathbb{R}\cup\{\infty\}$; see the supplementary appendix for a proof. In other words, the conditional expectation of $x(t)$ is not identified. The reason is that each value of $m^{\bar{x},\eta}(t)$ can be supported equally well by some theory with a large η.[32]

Indeed, "inductive inference based on objective criteria alone is bound to fail, while incorporating subjective criteria alongside objective ones can lead to successful learning," to quote Gilboa and Samuelson (2012). In other words, "effective learning requires a willingness to sacrifice goodness-of-fit in return for enhanced subjective appeal" (ibid.).

Thus, I assume that the plausibility ranking, a binary relation "at least as plausible as," over the theories is represented by a penalized log-likelihood function. Specifically, the agent finds more appealing the "reference" or "simple" theories free of the poorly understood factors, and that subjective criterion is translated into a penalty on the log-likelihood proportional to the magnitude of η measured by the L^2-norm:

$$\ell_t^\lambda(\bar{x},\eta) \triangleq \ell_t(\bar{x},\eta) - \frac{\lambda}{2}\int_0^t |\eta(s)|^2 \, \mathrm{d}s$$

where $\lambda \in (0,\infty]$ measures the agent's a priori confidence about the reference likelihood. When $\lambda = \infty$, the set of theories reduces to $\{Q^{\bar{x},0} : \bar{x} \in \mathbb{R}^{n_x}\}$ and the agent perceives no persistent source of ambiguity; when λ is small, the agent fits data with large ηs with little restraint. It is also worth noting that the L^2-norm of η is equal to the deviation of a theory $Q^{\bar{x},\eta}$ from its simple counterpart $Q^{\bar{x},0}$ measured by the Kullback-Leibler divergence:

$$\begin{aligned} D_{\mathrm{KL}}(Q^{\bar{x},0} \| Q^{\bar{x},\eta}) &\triangleq \mathrm{E}^{Q^{\bar{x},0}} \log \frac{\mathrm{d}Q^{\bar{x},0}}{\mathrm{d}Q^{\bar{x},\eta}} \\ &= \frac{1}{2}\int_0^T |\eta(t)|^2 \, \mathrm{d}t. \end{aligned}$$

The idea of penalizing the likelihood was first discussed by Good and Gaskins (1971) in the context of nonparametric density estimation. Green (1987) extended it to semiparametric settings. In these non- or semi-parametric estimation problems, Sobolev norms of higher orders, as well as the L^2-norm, are favored; but for us, imposing smoothness on η would violate the assumption of symmetry. In the context of model selection, Akaike (1973) extended

[32] Precisely speaking, the supremum is not attained, that is, there does not exist a maximum likelihood estimate. Fix $\bar{x}$ and suppose there is a partial maximizer η of the likelihood $\ell_t(\bar{x},\eta)$, $0 < t \le T$, in $L^2([0,T],\mathbb{R}^{n_x})$. Then it must satisfy, from Lemma 3.2,

$$0 = (\varphi(s)^{-1}\rho_v)^\top \int_s^t \varphi(\tau)^\top b(\tau)^\top (\sigma(\tau)\sigma(\tau)^\top)^{-1} [\, \mathrm{d}y(\tau) - (a(\tau) + b(\tau)m^{\bar{x},\eta}(\tau))\, \mathrm{d}\tau], \ 0 \le s \le t$$

but the constancy of the left-hand side and the unbounded variation of the right-hand side are incompatible. It follows that for any given η, there is another η' with higher likelihood.

the maximum likelihood principle by proposing his celebrated criterion in the form of a penalized log-likelihood; and ever since, penalizing the likelihood has been a standard method in information theory to strike a balance between the goodness of fit and the simplicity of the model; see Konishi and Kitagawa (2008). The penalized log-likelihood representation of a plausibility ranking has recently been axiomatized by Gilboa and Schmeidler (2010).

In conclusion, a theory $Q^{\bar{x},\eta}$ is not ruled out if and only if

$$\ell_t^\lambda(\bar{x},\eta) \geq \max_{(\bar{x}',\eta')\in\mathbb{R}^{n_x}\times L^2([0,T],\mathbb{R}^{n_x})} \ell_t^\lambda(\bar{x}',\eta') - \alpha \tag{24}$$

where $0 \leq \alpha < \infty$. α measures how conservative the agent is in model selection; when $\alpha = 0$, in particular, the agent keeps nothing but the most plausible theories. And as shall be seen, the corresponding induced log-likelihood of the conditional expectation of $x(t)$

$$\ell_{t,m(t)}^\lambda(m) \triangleq \max_{(\bar{x},\eta)\in\mathbb{R}^{n_x}\times L^2([0,T],\mathbb{R}^{n_x})} \left\{\ell_t^\lambda(\bar{x},\eta) : m^{\bar{x},\eta}(t) = m\right\},\ m \in \mathbb{R}^{n_x}$$

has a nonzero curvature (Lemma 3.6).

Remark 3.2. *There are two prominent alternatives to the L^2-penalty.*

The first is Epstein and Schneider's (2007) L^∞-constraint: $\operatorname{ess\,sup}_{t\leq T}|\eta(t)| \leq \bar{\eta}$. This amounts to constraining instantaneous entropy rates point by point in time. While this is sensible when the agent is looking forward and fears misspecification of the infinitesimal future, in looking backward, it is not. What the agent tries to pin down here is the value of $m^{\bar{x},\eta}(t)$, and with this regard, $\eta(s)$, $s \leq t$, having large values for a short period of time has little significance.

*The other is an L^2-*constraint*: $\int_0^T |\eta(t)|^2\,\mathrm{d}t \leq \bar{\eta}_T$. Naturally, this is closely related to the L^2-penalty: First, the constraint is a penalty that is discontinuous. Second, the constraint is the dual of the penalty in Lagrange's theorem. The constant λ defines a shadow process $\bar{\eta}^\lambda = \{\bar{\eta}_t^\lambda\}$ that implies the same most plausible theories. And I note that the penalized likelihood ratio test with λ is more conservative than the constrained likelihood ratio test with $\bar{\eta}^\lambda$; that is,*

$$\ell_t(\bar{x},\eta) \geq \max_{(\bar{x}',\eta')\in\mathbb{R}^{n_x}\times L^2([0,T],\mathbb{R}^{n_x})} \ell_t(\bar{x}',\eta') - \alpha \text{ and } \frac{1}{2}\int_0^t |\eta(s)|^2\,\mathrm{d}s \leq \bar{\eta}_t^\lambda$$

implies (24)*.*

Compared to its penalty counterpart, however, the L^2-constraint has the following drawbacks. First, the sharp bounds seem to be at odds with the assumed a priori ignorance. Second, if, as is natural, the time-t bound $\bar{\eta}_t$ is lower than $\bar{\eta}_T$, $t < T$, then it implies that the agent has a time-varying parameter set; for example, he would deem $\eta(s) = \sqrt{2\bar{\eta}_T/t}(1,0,\cdots,0)^\top$, $s \leq t$, implausible at time t but plausible at time T.

Maximum Plausibility Estimation I will need the following facts to characterize the natural "center" of the set of preferential priors.

The *maximum plausibility estimate (MPE) of $(\bar{x},\eta)$ at time t* is defined as

$$(\bar{x}_t^*,\eta_t^*) \triangleq \operatorname*{arg\,max}_{(\bar{x},\eta)\in\mathbb{R}^{n_x}\times L^2([0,T],\mathbb{R}^{n_x})} \ell_t^\lambda(\bar{x},\eta).$$

The notion of the partial MPE of η given $\bar{x}$ will prove helpful:

$$\eta^*_{\bar{x},t} \triangleq \underset{\eta \in L^2([0,T],\mathbb{R}^{n_x})}{\arg\max} \ \ell^\lambda_t(\bar{x}, \eta).$$

Clearly, $\eta^*_t = \eta^*_{\bar{x}^*_t,t}$.

The first-order condition with respect to η (FOC(η)) is

$$\lambda\eta(s) = (\varphi(s)^{-1}\rho_v)^\top \int_s^t \varphi(\tau)^\top b(\tau)^\top (\sigma(\tau)\sigma(\tau)^\top)^{-1} \times [\,\mathrm{d}y(\tau) - (a(\tau) + b(\tau)m^{\bar{x},\eta}(\tau))\,\mathrm{d}\tau],\ 0 \le s \le t.$$

To write the solution of this integral equation, introduce the following notation. Let

$$\chi(s) \triangleq \begin{pmatrix} \rho_v\rho_v^\top \bar{\kappa}(s)^\top(\rho_v\rho_v^\top)^{-1} & \rho_v\rho_v^\top b(s)^\top(\sigma(s)\sigma(s)^\top)^{-1}b(s) \\ \lambda^{-1}I_{n_x} & -\bar{\kappa}(s) \end{pmatrix}$$

and let ψ be the matrix-valued process such that $\psi(0) = I_{2n_x}$ and

$$\dot{\psi}(s) = \chi(s)\psi(s),\ 0 \le s \le T.$$

$\psi(s)$ is invertible for all $s \ge 0$. Let $\iota_1 \triangleq (I_{n_x}, 0)^\top$, $\iota_2 \triangleq (0, I_{n_x})^\top$, and $A_{ij} \triangleq \iota_i^\top A \iota_j$ for a $2n_x \times 2n_x$ matrix A.

Lemma 3.3. *For all $t > 0$, (i) $\psi_{11}(t)$ is invertible and (ii) $\psi_{21}(t)\psi_{11}(t)^{-1}\rho_v\rho_v^\top$ is symmetric and positive definite.*

Let also

$$\Psi(s) \triangleq \psi(s)\int_0^s \psi(\tau)^{-1}\,\mathrm{d}\tau,\ 0 \le s \le T.$$

Proposition 3.4 (Partial MPE of η)**.**

$$\begin{aligned} \begin{pmatrix} \lambda\rho_v\eta^*_{\bar{x},t}(s) \\ \Phi^{\kappa\bar{x}+\rho_v\eta^*_{\bar{x},t}}(s) \end{pmatrix} &= \psi(s)\iota_1\psi_{11}(t)^{-1} \\ &\quad \times \left(\iota_1^\top\psi(t)\int_0^t \psi(\tau)^{-1}\iota_1\rho_v\rho_v^\top b(\tau)^\top(\sigma(\tau)^\top)^{-1}\,\mathrm{d}\bar{w}^{0,0}(\tau) - \Psi_{12}(t)\kappa\bar{x}\right) \\ &\quad - \psi(s)\int_0^s \psi(\tau)^{-1}\iota_1\rho_v\rho_v^\top b(\tau)^\top(\sigma(\tau)^\top)^{-1}\,\mathrm{d}\bar{w}^{0,0}(\tau) + \Psi(s)\iota_2\kappa\bar{x}. \end{aligned} \tag{25}$$

Hence, $m^{\bar{x},\eta^*_{\bar{x},t}}(s)$ is linear in $\bar{x}$ (recall (21)). Define $\theta(t)$ by

$$\theta(t) \triangleq \Psi_{22}(t) - \psi_{21}(t)\psi_{11}(t)^{-1}\Psi_{12}(t)$$

or

$$m^{\bar{x},\eta^*_{\bar{x},t}}(t) = m^{0,\eta^*_{0,t}}(t) + \theta(t)\kappa\bar{x}.$$

That is, $\theta(t)$ measures the sensitivity to $\kappa\bar{x}$ of $m^{\bar{x},\eta}(t)$ with η "profiled out." Let $\mathcal{I}_{\bar{x}}(t)$ denote the observed Fisher information about $\bar{x}$:

$$\mathcal{I}_{\bar{x}}(t) \triangleq -\left.\frac{\partial^2}{\partial(\kappa\bar{x})^2}\ell^\lambda_t(\bar{x}, \eta^*_{\bar{x},t})\right|_{\bar{x}=\bar{x}^*_t}.$$

Precisely speaking, $\mathcal{I}_{\bar{x}}(t)$ is the information about $\kappa\bar{x}$, but I adopt this slight abuse of terminology because κ is known and the parameter of interest is clearly $\bar{x}$.

Assumption 3.3. *(i) $n_x \leq n_y$.*

(ii) $b(t)$ is of full column rank (that is, n_x) for all $t \in [0, T]$.

Lemma 3.4.

$$\mathcal{I}_{\bar{x}}(t) = \int_0^t \theta(s)^\top b(s)^\top (\sigma(s)\sigma(s)^\top)^{-1} b(s)\theta(s)\,\mathrm{d}s$$

and is invertible for all $t > 0$.

FOC($\bar{x}$) is

$$0 = \int_0^t (b(s)\Phi^{I_{n_x}}(s)\kappa)^\top (\sigma(s)\sigma(s)^\top)^{-1}[\,\mathrm{d}y(s) - (a(s) + b(s)m^{\bar{x},\eta}(s))\,\mathrm{d}s].$$

Proposition 3.5 (MPE of $\bar{x}$)**.** *For $t > 0$,*

$$\kappa \bar{x}_t^* = \mathcal{I}_{\bar{x}}(t)^{-1} \int_0^t \Phi^{I_{n_x}}(s)^\top b(s)^\top (\sigma(s)^\top)^{-1}\,\mathrm{d}\bar{w}^{0,\eta_{0,t}^*}(s).$$

Remark 3.3. *Estimation is not defined at time 0, and consequently, neither is the time-0 decision making. This is natural. At time 0, the agent is in the state of sheer ignorance while once the observable process y starts to wiggle, information thereafter accrues continuously. The singularity at time 0 is not a problem because, as we will see, decision making is well-defined for all $t > 0$. Nevertheless, I assume purely for the brevity of exposition that* the agent's learning started prior to time 0 and all the statistics, including $\mathcal{I}_{\bar{x}}(0)$ and $\bar{x}_0^*$, have a definite, finite value at time 0. *The differential dynamics I am about to characterize determine their evolution from thenceforth. To maintain the convention that $\mathcal{G}_0$ is trivial, I assume that all the $\mathcal{G}_0$-measurable variables are nonrandom constants.*

The natural center of the time-t set of one-step-ahead conditionals is

$$\mathrm{d}y(t)|\mathcal{G}_t \sim N\left[(a(t) + b(t)m^{\bar{x}_t^*,\eta_t^*}(t))\,\mathrm{d}t, \sigma(t)\sigma(t)^\top\,\mathrm{d}t\right].$$

This observation motivates us to define a process $\epsilon = \{\epsilon(t), \mathcal{G}_t\}$ by

$$\mathrm{d}\epsilon(t) = \sigma(t)^{-1}[\,\mathrm{d}y(t) - (a(t) + b(t)m^{\bar{x}_t^*,\eta_t^*}(t))\,\mathrm{d}t],\ \ \epsilon(0) = 0. \tag{26}$$

To prove that there is a probability measure on $(\Omega, \mathcal{G}_T)$ under which ϵ is a Wiener process, I first observe the dynamics of the statistics.

Proposition 3.6 (Dynamics of the MPEs)**.**

$$\kappa\,\mathrm{d}\bar{x}_t^* = \sigma_{\bar{x}^*}(t)^\top b(t)^\top (\sigma(t)^\top)^{-1}\,\mathrm{d}\epsilon(t),$$
$$\mathrm{d}m^{\bar{x}_t^*,\eta_t^*}(t) = \kappa(\bar{x}_t^* - m^{\bar{x}_t^*,\eta_t^*}(t))\,\mathrm{d}t + [\rho_w\sigma(t)^\top + (\gamma(t) + \delta(t))b(t)^\top](\sigma(t)^\top)^{-1}\,\mathrm{d}\epsilon(t), \tag{27}$$

where

$$\sigma_{\bar{x}^*}(t) \triangleq \theta(t)\mathcal{I}_{\bar{x}}(t)^{-1},$$
$$\delta(t) \triangleq \psi_{21}(t)\psi_{11}(t)^{-1}\rho_v\rho_v^\top + \theta(t)\sigma_{\bar{x}^*}(t)^\top$$
$$= \psi_{21}(t)\psi_{11}(t)^{-1}\rho_v\rho_v^\top + \sigma_{\bar{x}^*}(t)\mathcal{I}_{\bar{x}}(t)\sigma_{\bar{x}^*}(t)^\top.$$

Note that δ is symmetric and positive definite. The following proposition closes the dynamics:

Proposition 3.7.

$$\dot{\theta}(t) = I_{n_x} - \{\kappa + [\rho_w\sigma(t)^\top + (\gamma(t) + \delta(t) - \theta(t)\sigma_{\bar{x}^*}(t)^\top)b(t)^\top](\sigma(t)\sigma(t)^\top)^{-1}b(t)\}\theta(t), \quad (28)$$

$$\dot{\sigma}_{\bar{x}^*}(t) = \mathcal{I}_{\bar{x}}(t)^{-1} - \{\kappa + [\rho_w\sigma(t)^\top + (\gamma(t) + \delta(t))b(t)^\top](\sigma(t)\sigma(t)^\top)^{-1}b(t)\}\sigma_{\bar{x}^*}(t),$$

$$\frac{\mathrm{d}}{\mathrm{d}t}(\mathcal{I}_{\bar{x}}(t)^{-1}) = -\sigma_{\bar{x}^*}(t)^\top b(t)^\top(\sigma(t)\sigma(t)^\top)^{-1}b(t)\sigma_{\bar{x}^*}(t),$$

$$\begin{aligned}\dot{\delta}(t) &= \sigma_{\bar{x}^*}(t) + \sigma_{\bar{x}^*}(t)^\top + \lambda^{-1}\rho_v\rho_v^\top \\ &+ (\rho_w\sigma(t)^\top + \gamma(t)b(t)^\top)(\sigma(t)\sigma(t)^\top)^{-1}(\rho_w\sigma(t)^\top + \gamma(t)b(t)^\top)^\top \\ &- \kappa\delta(t) - \delta(t)\kappa \\ &- [\rho_w\sigma(t)^\top + (\gamma(t) + \delta(t))b(t)^\top](\sigma(t)\sigma(t)^\top)^{-1}[\rho_w\sigma(t)^\top + (\gamma(t) + \delta(t))b(t)^\top]^\top.\end{aligned} \quad (29)$$

The Preferential Priors Make the following additional assumption:

Assumption 3.4. *θ, $\sigma_{\bar{x}^*}$ and δ are uniformly bounded.*

Here are simple example cases in which Assumption 3.4 holds:

Lemma 3.5. *Suppose either: (i) σ and b are deterministic or (ii) σ, ρ_w, ρ_v, and b are diagonal*[33] *and there is an $\varepsilon > 0$ such that $\bar{\kappa} = \kappa + (\rho_w\sigma^\top + \gamma b^\top)(\sigma\sigma^\top)^{-1}b \geq \varepsilon I_{n_x}$ a.e. Then Assumption 3.4 holds.*

Remark 3.4. *Given that σ, ρ_w, ρ_v, and b are diagonal, there trivially is an $\varepsilon > 0$ such that $\bar{\kappa} > \varepsilon I_{n_x}$ a.e. if $\rho_w = 0$.*

Proposition 3.8. *There is a unique probability measure on $(\Omega, \mathcal{G}_T)$, denoted by P^0, such that $P^0 \sim (Q^{0,0}|_{\mathcal{G}_T})$ and ϵ is a Wiener process under P^0. Also, ϵ generates $\mathbf{G}$.*

Observe that under P^ξ,

$$\mathrm{d}y(t)|\mathcal{G}_t \sim N\left[(a(t) + b(t)m^{\bar{x}_t^*, \eta_t^*}(t) + \sigma(t)\xi(t))\,\mathrm{d}t, \sigma(t)\sigma(t)^\top\,\mathrm{d}t\right].$$

Hence, the time-t set of one-step-ahead conditionals $\Xi(t)$ is defined by

$$\begin{aligned}&a(t) + b(t)m^{\bar{x}_t^*, \eta_t^*}(t) + \sigma(t)\Xi(t) \\ &\quad = \left\{\mu \in \mathbb{R}^{n_x} : \ell_t^\lambda(\bar{x}_t^*, \eta_t^*) - \max_{(\bar{x},\eta)\in\mathbb{R}^{n_x}\times L^2([0,T],\mathbb{R}^{n_x})}\{\ell_t^\lambda(\bar{x}, \eta) : a(t) + b(t)m^{\bar{x},\eta}(t) = \mu\} \leq \alpha\right\}\end{aligned}$$

where the maximum is defined to be $-\infty$ when there does not exist $(\bar{x}, \eta)$ satisfying the constraint. It turns out that $\delta(t)$ is the inverse of the observed Fisher information about $m^{\bar{x},\eta}(t)$:

[33] In case $n_x \neq n_y$, the $n_x \times n_y$ matrix ρ_w, for example, is diagonal if $\rho_w^{ij} = 0$ for all $i \neq j$.

Lemma 3.6.

$$\begin{aligned}\ell_t^\lambda(\bar{x}_t^*, \eta_t^*) - \max_{(\bar{x},\eta)\in\mathbb{R}^{n_x}\times L^2([0,T],\mathbb{R}^{n_x})} &\{\ell_t^\lambda(\bar{x},\eta) : m^{\bar{x},\eta}(t) = m\} \\ &= \ell_{t,m(t)}^\lambda(m^{\bar{x}_t^*,\eta_t^*}(t)) - \ell_{t,m(t)}^\lambda(m) \\ &= \frac{1}{2}(m - m^{\bar{x}_t^*,\eta_t^*}(t))^\top \delta(t)^{-1}(m - m^{\bar{x}_t,\eta_t^*}(t)),\ m \in \mathbb{R}^{n_x}.\end{aligned}$$

Proposition 3.9.

$$\sigma(t)\Xi(t) = b(t)\left\{\Delta m \in \mathbb{R}^{n_x} : \frac{1}{2}(\Delta m)^\top \delta(t)^{-1}\Delta m \le \alpha\right\},\ 0 \le t \le T. \tag{30}$$

The process $\Xi = \{\Xi(t), \mathcal{G}_t\}$ is uniformly bounded and compact-convex. If furthermore each of the processes b and σ^{-1} is left- or right-continuous, then Ξ is progressive.

Remark 3.5. *For $\xi(t) \in \Xi(t)$,*

$$\frac{1}{2}(\sigma(t)\xi(t))^\top (b(t)\delta(t)b(t)^\top)^+ \sigma(t)\xi(t) \le \alpha \tag{31}$$

where $(b(t)\delta(t)b(t)^\top)^+$ denotes the Moore-Penrose pseudoinverse:

$$(b(t)\delta(t)b(t)^\top)^+ = b(t)(b(t)^\top b(t))^{-1}\delta(t)^{-1}(b(t)^\top b(t))^{-1}b(t)^\top.$$

But the converse is not true, that is, (31) *does not imply $\xi(t) \in \Xi(t)$.*

With a slight abuse of notation, let $\xi \in \Xi$ mean that $\xi = \{\xi(t), \mathcal{G}_t\}$ is progressive and $\xi(t,\omega) \in \Xi(t,\omega)$ a.e. The set of preferential priors is given by

$$\mathcal{P} = \left\{P^\xi : P^\xi \text{ is a probability measure on } (\Omega, \mathcal{G}_T),\ \frac{\mathrm{d}P^\xi}{\mathrm{d}P^0} = \mathcal{E}^\xi(T),\ \xi \in \Xi\right\}.$$

3.4 Discussion

Assume throughout this section, with the exception of the last discussion (Section 3.4.5), $n_x = 1$. Still the setup is general enough to encompass all the examples given in Section 3.2.1.

3.4.1 Learning about $\bar{x}$

Proposition 3.10. *Suppose $b^\top(\sigma\sigma^\top)^{-1}b$ is uniformly bounded below. Then, the confidence interval for $\bar{x}$ shrinks to a point as time goes to infinity, for all critical values α.*

In other words, the ambiguity associated with $\bar{x}$ eventually resolves.

The question that naturally arises next is if $\bar{x}^*$ converges. But, since convergence under a probability measure does not imply convergence under another probability measure obtained by a Girsanov change of measure (see Karatzas and Shreve (1988), p. 193), to answer this question we need to take a stance on the true probability measure. Although my stance is that not only does the agent not know the true probability measure but he does not purport, either, to have identified a set of probability measures (theoretical priors) that includes it, if need be the natural candidate for the true probability measure is a theoretical prior $Q^{\bar{x},0} \in \mathcal{Q}$ of the agent (correct specification). It remains to be seen if $\bar{x}^*$ converges under $Q^{\bar{x},0}$.[34]

3.4.2 Comparison with the Classical Filter

The agent's learning process is summarized by a finite-dimensional filter (Propositions 3.6 and 3.7). The key components of the filter are $\{m^{\bar{x}^*_t,\eta^*_t}(t)\}$ and δ. $m^{\bar{x}^*_t,\eta^*_t}(t)$, in particular, is the agent's benchmark estimate of the conditional expectation of $x(t)$, and I accordingly defined the reference preferential prior P^0 to be the "concatenation" of the one-step-ahead conditionals computed using the benchmark estimates ((26) and Proposition 3.8). The set $\mathcal{P}$ of preferential priors is then given by a neighborhood of P^0, reflecting the a posteriori ambiguity in the conditional expectation of $x(t)$, or more fundamentally, in the data-generating mechanism. The set of the values of the conditional expectation of $x(t)$ that are sufficiently plausible is given by an interval centered at the benchmark estimate $m^{\bar{x}^*_t,\eta^*_t}(t)$ (Lemma 3.6 and Proposition 3.9). The length of the interval is proportional to the square root of $\delta(t)$. Hence, the latter, or $\delta(t)$ itself, is a measure of a posteriori ambiguity.

Remark 3.6. *In general, that is, when $n_x \in \mathbb{N}$, the set of alternative values of the conditional expectation of $x(t)$ is given by an n_x-dimensional hyper-ellipsoid centered at $m^{\bar{x}^*_t,\eta^*_t}(t)$. The lengths of the principal axes of the hyper-ellipsoid are proportional to the square-roots of the eigenvalues of $\delta(t)$.*

Therefore, of prime interest is how $m^{\bar{x}^*_t,\eta^*_t}(t)$ and $\delta(t)$ evolve. In what follows, I compare the filtering equations (27) and (29) with the classical conditionally Gaussian filter (Liptser and Shiryaev (1977), Chapter 12).

Let us begin with the unobservable process x:

$$\mathrm{d}x(t) = \kappa(\bar{x} - x(t))\,\mathrm{d}t + \rho_w\,\mathrm{d}w(t) + \rho_v(\,\mathrm{d}v^{\bar{x},\eta}(t) + \eta(t)\,\mathrm{d}t),$$
$$\frac{\mathrm{d}}{\mathrm{d}t}\mathrm{Var}(x(t)) = \underbrace{|\rho_w|^2 + \rho_v^2}_{\mathrm{Var}(\,\mathrm{d}x(t)|\mathcal{F}_t)/\,\mathrm{d}t} - 2\kappa\,\mathrm{Var}(x(t)).$$

The time-derivative of the unconditional variance of $x(t)$ is the conditional variance of $\mathrm{d}x(t)$ given $\mathcal{F}_t$ per unit time less the unconditional variance times the rate of reversion (times

---

[34] $\Delta\bar{x}^*(t) \triangleq \bar{x}^*_t - \bar{x}$ and $\Delta m^*(t) \triangleq m^{\bar{x}^*_t,\eta^*_t}(t) - m^{\bar{x},0}(t)$ satisfy

$$\kappa\,\mathrm{d}\Delta\bar{x}^* = \sigma_{\bar{x}^*}^\top b^\top(\sigma^\top)^{-1}(\,\mathrm{d}\bar{w}^{\bar{x},0} - \sigma^{-1}b\Delta m^*\,\mathrm{d}t)$$

$$\mathrm{d}\Delta m^* = \kappa(\Delta\bar{x}^* - \Delta m^*)\,\mathrm{d}t + \delta b^\top(\sigma^\top)^{-1}\,\mathrm{d}\bar{w}^{\bar{x},0} - (\rho_w\sigma^\top + (\gamma+\delta)b^\top)(\sigma\sigma^\top)^{-1}b\Delta m^*\,\mathrm{d}t$$

and $(\Delta\bar{x}^*, \Delta m^*)$ converges in L^2. But the difficulty is that $\sigma_{\bar{x}^*}$ is square-integrable but not integrable. It is not clear whether

$$\int_0^\infty \sigma_{\bar{x}^*}^\top b^\top(\sigma\sigma^\top)^{-1}b\Delta m^*\,\mathrm{d}t$$

is convergent or not. On the other hand, it is easy to see that $\bar{x}^*$ is an L^2-*bounded* continuous martingale under P^0, and therefore, under P^0, $\lim_{t\to\infty}\bar{x}^*_t$ exists by Doob's martingale convergence theorem (Rogers and Williams (1994), Theorem II.69.1).

two). This is intuitive. In particular, the unconditional variance is decreasing in the rate of reversion because as the reversion term becomes dominant, x stays closer to $\bar{x}$.

Recall next the classical conditionally Gaussian filter (19) and (20), slightly rephrased to facilitate the discussion:

$$\mathrm{d}m^{\bar{x},\eta}(t) = [\kappa(\bar{x} - m^{\bar{x},\eta}(t)) + \rho_v \eta(t)]\,\mathrm{d}t + \underbrace{[\rho_w + \underbrace{(\sigma(t)^{-1}b(t)\gamma(t))^\top}_{\text{Kalman gain}}]}_{\text{weight on the innovation}}\,\mathrm{d}\bar{w}^{\bar{x},\eta}(t),$$

$$\dot{\gamma}(t) = |\rho_w|^2 + \rho_v^2 - 2\kappa\gamma(t) - \underbrace{\left|\rho_w + (\sigma(t)^{-1}b(t)\gamma(t))^\top\right|^2}_{\text{weight on the innovation squared}}. \tag{32}$$

We revise $m^{\bar{x},\eta}(t)$ in consideration of two factors: (i) the estimation error and (ii) the variation in $x(t)$. First, the correction of the estimation error, $+(\sigma(t)^{-1}b(t)\gamma(t))^\top\,\mathrm{d}\bar{w}^{\bar{x},\eta}(t)$, is proportional to the innovation $\mathrm{d}\bar{w}^{\bar{x},\eta}(t)$. To understand, suppose $n_y = 1$ and $\sigma, b > 0$. Then, when, for example, the change in the observable variable exceeds what was expected, it is likely that the old estimate of the growth rate is an underestimation and it thus needs to be revised up. The multiplicative factor, or the Kalman gain, is increasing in the current uncertainty $\gamma(t)$ of $x(t)$ (the less trustworthy the current estimate, the more weight given to the new evidence) and is decreasing in the imprecision $\sigma(t)$ of the signal (the less informative the signal, the less weight given to the new evidence). Second, $m^{\bar{x},\eta}(t)$ as an estimate of $x(t)$ is also to be revised by $+\kappa(\bar{x} - m^{\bar{x},\eta}(t))\,\mathrm{d}t + \rho_w\,\mathrm{d}\bar{w}^{\bar{x},\eta}(t)$, to account for the corresponding (unobservable) changes $+\kappa(\bar{x} - x(t))\,\mathrm{d}t + \rho_w\,\mathrm{d}w(t)$ in $x(t)$.

$\dot{\gamma}(t)$ is given by the analogue of $(\mathrm{d}/\mathrm{d}t)\,\mathrm{Var}(x(t))$ less the weight on the innovation squared. The last term means that uncertainty resolves more quickly when the new evidence is taken more seriously. This is intuitive because if we consider the extreme case in which the weight is exactly zero, then it is equivalent to the case in which there temporarily is no signal ($\mathrm{d}\bar{w}^{\bar{x},\eta}(t) = 0$), in which case it is natural for the uncertainty to rise.

To be noted further is that while the motivation for the revision $+\rho_w\,\mathrm{d}\bar{w}^{\bar{x},\eta}(t)$ is to account for the state dynamics, in effect it acts as if it were a weight on the new evidence. And when the signs of ρ_w and b differ (assume $n_y = 1$), it can counteract the correction of the estimation error, and consequently, the uncertainty γ may depend nonmonotonically on the signal imprecision σ. To elaborate, suppose that the signs of ρ_w and b differ and that σ is time-varying and is currently sufficiently large that the signal is noninformative about $v^{\bar{x},\eta}$ but informative about w ($\lim_{t\to\infty}\gamma(t) = \rho_v^2/(2\kappa)$). If σ drops and the weight on the innovation becomes exactly zero, then by way of the earlier reasoning, uncertainty rises ($\lim_{t\to\infty}\gamma(t) = (\rho_w^2 + \rho_v^2)/(2\kappa)$). When ρ_w and b have the same sign, the filter always corrects the estimation error and uncertainty resolves.

Recall finally the dynamics (27) and (29) of $m^{\bar{x}_t^*,\eta_t^*}(t)$ and $\delta(t)$ under the reference preferential prior P^0, again slightly rephrased:

$$\mathrm{d}m^{\bar{x}_t^*,\eta_t^*}(t) = \kappa(\bar{x}_t^* - m^{\bar{x}_t^*,\eta_t^*}(t))\,\mathrm{d}t + \{\rho_w + [\sigma(t)^{-1}b(t)(\underbrace{\gamma(t)+\delta(t)}_{\text{estimation uncertainty}})]^\top\}\,\mathrm{d}\epsilon(t),$$

$$\dot{\delta}(t) = \underbrace{\left|\rho_w + (\sigma(t)^{-1}b(t)\gamma(t))^\top\right|^2}_{\mathrm{Var}(\,\mathrm{d}m^{\bar{x},\eta}(t)|\mathcal{G}_t)/\,\mathrm{d}t} - 2\kappa\delta(t) - \underbrace{\left|\rho_w + \left[\sigma(t)^{-1}b(t)(\gamma(t)+\delta(t))\right]^\top\right|^2}_{\text{weight on the innovation squared}} + \underbrace{2\theta(t)\mathcal{I}_{\bar{x}}(t)^{-1}}_{\text{ambiguity associated with } \bar{x}} + \underbrace{\lambda^{-1}\rho_v^2}_{\text{ambiguity associated with } \eta}.$$

As with the Bayesian estimate $m^{\bar{x},\eta}(t)$, the ambiguity-averse agent's estimate $m^{\bar{x}_t^*,\eta_t^*}(t)$, too, is revised in consideration of the estimation error and the variation in $x(t)$. But the difference is that now $\gamma(t)$ is replaced by the sum of $\gamma(t)$ and $\delta(t)$. $\gamma(t)$ is also known as the *estimation risk* in the literature (Kalymon, 1971; Barry, 1974; Klein and Bawa, 1976, 1977) and represents the Bayesian uncertainty under each theory that results because the agent cannot observe $x(t)$ and consequently has to estimate it. On the other hand, $\delta(t)$ represents the Knightian uncertainty that results because the agent does not know which theory (data-generating mechanism) is correct and consequently has to estimate it. Based on this parallelism, I call $\delta(t)$ the *estimation ambiguity* and the sum $\gamma(t)+\delta(t)$ the *estimation uncertainty*. When the estimated theory is imprecise (large $\delta(t)$), or the posterior distribution of $x(t)$ is diffuse under the theory (large $\gamma(t)$), or both, new evidence receives more weight. Note that this simple characterization of estimation uncertainty relies on the assumption that the prior variance of $x(0)$ is common to all theories; otherwise, different theories would result in different levels of estimation risk.

$\dot{\delta}(t)$ is given by an analogue of $\dot{\gamma}(t)$ plus terms accounting for the ambiguity in the data-generating mechanism. The first three terms reflect the fact that δ measures the imprecision in the estimation of $m^{\bar{x},\eta}$, as opposed to that in the estimation of x as does γ. To elaborate, the first term is the conditional variance of $\mathrm{d}m^{\bar{x},\eta}(t)$ (given $\mathcal{G}_t$) per unit time as opposed to that of $\mathrm{d}x(t)$ (given $\mathcal{F}_t$) per unit time; the parallelism between the second terms is obvious; and the third term is the weight on the innovation squared, exactly as in $\dot{\gamma}(t)$. Next, the fourth term captures the ambiguity in the estimate $m^{\bar{x}_t^*,\eta_t^*}(t)$ of $m^{\bar{x},\eta}(t)$ due to that in $\bar{x}$; recall that $\theta(t)$ measures the sensitivity to $\bar{x}$ of $m^{\bar{x},\eta^*_{\bar{x},t}}(t)$ and $\mathcal{I}_{\bar{x}}(t)^{-1}$ the imprecision of $\bar{x}_t^*$. The fifth and last term captures the ambiguity associated with η and sets the long-run level of δ as the fourth term vanishes as was observed in Section 3.4.1. As the ambiguity regarding $\bar{x}$ gets resolved, the agent's lack of confidence in mean reversion is what keeps the a posteriori ambiguity δ above zero: $\lim_{t\to\infty}\delta(t) = 0$ if and only if $\lambda = \infty$.

Similar remarks apply to the dual role of the revision $+\rho_w\,\mathrm{d}\epsilon(t)$. In particular, δ may depend nonmonotonically on the signal imprecision σ. But this time the dependence is subtler because the first term of $\dot{\delta}(t)$ also involves ρ_w.

3.4.3 Convergence to an IID Ambiguity

We have the following necessary and sufficient condition for convergence to an IID ambiguity:

Proposition 3.11. *Assume finite confidence $\lambda < \infty$ for nondegeneracy. Ξ converges to a constant subset of $\mathbb{R}^{n_y}$ if and only if $\sigma^{-1}b$ converges to a constant vector in $\mathbb{R}^{n_y}$.*

Suppose $\sigma^{-1}b$ converges to a constant vector and denote the latter again by $\sigma^{-1}b$. Suppose further $n_y = 1$ for simplicity. Then,

$$\lim_{t\to\infty} \bar{\xi}(t)^2 = 2\alpha \left(\sqrt{(\kappa + \sigma^{-1}b\rho_w)^2 + (1+\lambda^{-1})(\sigma^{-1}b\rho_v)^2} - \sqrt{(\kappa + \sigma^{-1}b\rho_w)^2 + (\sigma^{-1}b\rho_v)^2} \right)$$

where $\Xi(t) = [-\bar{\xi}(t), \bar{\xi}(t)]$. Note that $\bar{\xi}(\infty) \equiv \lim_{t\to\infty} \bar{\xi}(t)$ is nonzero if and only if λ is finite. Also, not surprisingly, it is increasing in α and decreasing in λ.

To see the dependence of $\bar{\xi}(\infty)$ on other parameters, let

$$Y(t) \triangleq \int_0^t \sigma(s)^{-1}\,\mathrm{d}y(s)$$

and assume $\sigma^{-1}a$ is deterministic. Then,

$$\mathrm{d}Y(t) = \sigma(t)^{-1}(a(t) + b(t)x(t))\,\mathrm{d}t + \mathrm{d}w(t)$$

and, denoting the asymptotic variability of Y by

$$V_Y \triangleq \lim_{t\to\infty} \frac{\mathrm{d}}{\mathrm{d}t} \mathrm{Var}(Y(t)),$$

we can rewrite $\bar{\xi}(\infty)^2$ as

$$\bar{\xi}(\infty)^2 = 2\alpha \left(\sqrt{\kappa^2 V_Y + \lambda^{-1}(\sigma^{-1}b\rho_v)^2} - \sqrt{\kappa^2 V_Y} \right).$$

Thus, $\bar{\xi}(\infty)$ is decreasing in κ and V_Y. This is intuitive. First, when κ is large, the unobservable process stays close to the attractor. Second, V_Y measures the variability of the unobservable process x relative to the measurement error w.[35]

Remark 3.7. *When $\sigma^{-1}b$ converges to a constant vector, γ as well as δ converges to a constant (see* (32)*). Thus, if we further assume that $\bar{x}^*$, too, converges (or is known), then the agent at $t \geq \infty$ is observationally equivalent to someone in a fully observable environment*

$$\mathrm{d}y(t) = (a(t) + b(t)m(t))\,\mathrm{d}t + \sigma(t)\,\mathrm{d}\epsilon(t), \tag{33}$$

$$\mathrm{d}m(t) = \kappa(\bar{m} - m(t))\,\mathrm{d}t + \sigma_m\,\mathrm{d}\epsilon(t), \tag{34}$$

with an IID ambiguity. That is, the present paper justifies the IID assumption on ambiguity made in the context of standard dynamics like (33)-(34) *(see, for example, Trojani and Vanini (2004) and Liu (2013)) indeed as the limit of a learning process.*

3.4.4 The Maximum Plausibility Agents

Agents with $\alpha = 0$ deserve a separate discussion; while a priori they lack confidence, a posteriori these agents manage to have full confidence in a single one-step-ahead conditional, the one that is implied by the most plausible theory of the moment; that is, they are observationally equivalent to Bayesians with unique preferential prior P_0. Refer to the agents with $\alpha = 0$ as the *maximum plausibility (MP) agents.*

In what follows, I find a theory that is consistent with P_0, thereby rendering the MP agents Bayesians under partial observation; and a byproduct of the discussion is an alternative interpretation of δ. Suppose for this discussion $\bar{x}$ is known (or has converged).

[35] See Footnote 2.

Observational Equivalence With $\bar{x}$ known, the observable process y satisfies (in reading what follows, recall the filtering equations in Section 3.4.2)

$$\mathrm{d}y(t) = (a(t) + b(t)m^{\bar{x},\eta_t^*}(t))\,\mathrm{d}t + \sigma(t)\,\mathrm{d}\epsilon(t), \tag{35}$$

$$\mathrm{d}m^{\bar{x},\eta_t^*}(t) = \kappa(\bar{x} - m^{\bar{x},\eta_t^*}(t))\,\mathrm{d}t + \{\rho_w + [\sigma(t)^{-1}b(t)(\gamma(t) + \delta(t))]^\top\}\,\mathrm{d}\epsilon(t). \tag{36}$$

If we define γ_λ by $\gamma_\lambda(t) \triangleq \gamma(t) + \delta(t)$, $t \geq 0$, then it satisfies[36]

$$\dot{\gamma}_\lambda(t) = |\rho_w|^2 + (1 + \lambda^{-1})\rho_v^2 - 2\kappa\gamma_\lambda(t) - \left|\rho_w + (\sigma(t)^{-1}b(t)\gamma_\lambda(t))^\top\right|^2.$$

This is the differential equation that γ satisfies, with ρ_v^2 replaced by $(1+\lambda^{-1})\rho_v^2$. That is, $\gamma_\lambda(t)$ is the posterior variance of $x(t)$ *under* $Q_\lambda^{\bar{x},0}$, where $Q_\lambda^{\bar{x},0}$, a probability measure on $(\Omega, \mathcal{F}^\circ)$, equals the law of (y, x) implied by the SDEs

$$\mathrm{d}y(t) = (a(t) + b(t)x(t))\,\mathrm{d}t + \sigma(t)\,\mathrm{d}w(t), \tag{37}$$

$$\mathrm{d}x(t) = \kappa(\bar{x} - x(t))\,\mathrm{d}t + \rho_w\,\mathrm{d}w(t) + \sqrt{1 + \lambda^{-1}}\rho_v\,\mathrm{d}v_\lambda^{\bar{x},0}(t). \tag{38}$$

Here, w and $v_\lambda^{\bar{x},0}$ are independent Wiener processes under $Q_\lambda^{\bar{x},0}$, and $v_\lambda^{\bar{x},0}$ in particular is *defined by* (38) as was $v^{\bar{x},\eta}$. (Recall the construction of weak solutions in Section 3.2.3.) Now note that for a Bayesian agent with theory (37)-(38), filtering yields

$$\mathrm{d}y(t) = (a(t) + b(t)m_\lambda^{\bar{x},0}(t))\,\mathrm{d}t + \sigma(t)\,\mathrm{d}\bar{w}_\lambda^{\bar{x},0}(t),$$

$$\mathrm{d}m_\lambda^{\bar{x},0}(t) = \kappa(\bar{x} - m_\lambda^{\bar{x},0}(t))\,\mathrm{d}t + [\rho_w + (\sigma(t)^{-1}b(t)\gamma_\lambda(t))^\top]\,\mathrm{d}\bar{w}_\lambda^{\bar{x},0}(t).$$

And compare, finally, the last pair of SDEs with (35)-(36) to see that an MP agent is observationally equivalent to a Bayesian agent with theory (37)-(38). In other words, the effect of learning under ambiguous reversion, when $\alpha = 0$, is observationally equivalent to an increase in the volatility of the unobservable process.

Interpretation of δ A simple rewriting of the definition of $\gamma_\lambda(t)$ reveals that

$$\delta(t) = \gamma_\lambda(t) - \gamma(t). \tag{39}$$

Thus, for an MP agent, the Knightian uncertainty is equal to the discrepancy between the Bayesian uncertainty implied by his behavior and that implied by his beliefs. When an MP agent is confident about mean reversion, that is, when $\lambda = \infty$, in particular, there will be no Knightian uncertainty as there will be no discrepancy, either, between his elicited and actual beliefs (note that $\gamma \equiv \gamma_\infty$).

3.4.5 Application: Assets with Differing Degrees of Ambiguity

As mentioned in the introduction, Epstein and Miao (2003) have used a recursive multiple-priors model to explain the equity home-bias puzzle: agents invest less in foreign assets because they find the dynamics of foreign economies more ambiguous.[37] Epstein and Miao,

[36]In recalling $\dot{\delta}(t)$, note that $\mathcal{I}_{\bar{x}}^{-1} \equiv 0$ since $\bar{x}$ is known.

[37]See also Uppal and Wang (2003), who address similar issues using a robust-control model.

however, *take as given* the differing degrees of ambiguity: specifically, there are two countries; fundamentals are driven by a two-dimensional Wiener process $\epsilon = (\epsilon_1, \epsilon_2)$; the investors of the countries (or simply the countries) have the common information structure generated by ϵ; but Country i, it is assumed, finds the drift of ϵ_j ambiguous, lying at all times between $\pm\bar{\xi}_i$, $\bar{\xi}_i > 0$, $i \neq j$.

The present model, on the other hand, can deliver heterogeneous ambiguity *endogenously* as an outcome of learning under asymmetric information.[38]

Suppose there are two countries and two assets (one in each country). The two countries are identical in the following aspects: (i) they both see the returns of both assets; (ii) they theorize about the returns in the same way: with R_i denoting the cumulative return process of the asset traded in Country i, or Asset i,

$$\mathrm{d}R_i(t) = x_i(t)\,\mathrm{d}t + \sigma_{R_i}\,\mathrm{d}w_i(t),$$
$$\mathrm{d}x_i(t) = \kappa_i(\bar{x}_i - x_i(t))\,\mathrm{d}t + \rho_{v,i}(\,\mathrm{d}v_i^{\bar{x},\eta}(t) + \eta_i(t)\,\mathrm{d}t),$$

where $i \in \{1, 2\}$ labels the countries (whether $\bar{x}_i$s are known or not is irrelevant because we will focus on what happens to ambiguity in the limit);[39] and (iii) they have the same level λ of confidence in mean reversion and the same level α of conservatism in model selection. But the two countries differ in that each has exclusive access to a country-specific signal: Country 1 sees A_1 but not A_2, and vice versa, where

$$\mathrm{d}A_i(t) = x_i(t)\,\mathrm{d}t + \sigma_{A_i}\,\mathrm{d}w_{i+2}(t),\ \ i = 1, 2.$$

Regarding the assumption of exclusive signals, Country i (Italy) may well, technically speaking, have access to all news circulating in Country j (Korea). But think of the signals as flows of noneconomic information concerning each market that is difficult to appreciate unless well-versed in the country's culture, history, and so on. Then, accessible as it may technically be, such information floating around Korea is likely to be absent in the minds of (representative) Italian investors at the time of their decision making.

Now, consider Country i, for which $y = (R_i, R_j, A_i)$ and $x = (x_i, x_j)$. Under a generic preferential prior P^ξ,

$$\mathrm{d}R_k(t) = (m_k^{\bar{x}_t^*,\eta_t^*}(t) + \Delta m_k(t))\,\mathrm{d}t + \sigma_{R_k}\,\mathrm{d}\epsilon_k^\xi(t)$$

where ϵ_k^ξ is a P^ξ-Wiener process and $\Delta m_k(t) = \sigma_{R_k}\xi_k(t)$, $k = i, j$. $\Delta m(t) = (\Delta m_i(t), \Delta m_j(t))$ is constrained by $\frac{1}{2}(\Delta m(t))^\top \delta(t)^{-1}\Delta m(t) \leq \alpha$, and so it only remains to compute $\delta(t)$. (Incidentally, note that the constraint is elliptic rather than rectangular as Epstein and Miao assume.)

Since $\sigma^{-1}b$ is constant and the coefficients of $\dot{\gamma}$ are all diagonal under the present assumptions, γ converges to a diagonal matrix; and so does δ (recall (39)). Write

$$\lim_{t\to\infty} \delta(t) = \begin{pmatrix} \delta_{\mathrm{d},i} & 0 \\ 0 & \delta_{\mathrm{f},i} \end{pmatrix}.$$

[38] I mention, however, that Epstein and Miao state, "The differing beliefs of the two individuals described below are *not* due to asymmetric information; they reflect differing prior views about the environment" (emphasis in the original). I am here providing a scenario, in contrast, where asymmetric information results in asymmetric priors, which are always posteriors relative to a suppressed past.

[39] In fact, Epstein and Miao consider a general-equilibrium model, in which context agents are to learn from observing the endowment process; but in this paper I try to stick to the partial-equilibrium perspective.

Thus, after enough time has passed the constraint becomes

$$\frac{1}{2}[\delta_{\mathrm{d},i}^{-1}(\Delta m_i(t))^2 + \delta_{\mathrm{f},i}^{-1}(\Delta m_j(t))^2] \leq \alpha;$$

and we can roughly interpret $\sqrt{2\alpha\delta_{\mathrm{d},i}}$ as the remaining ambiguity in the instantaneous expected return of the domestic asset and $\sqrt{2\alpha\delta_{\mathrm{f},i}}$ as that in the instantaneous expected return of the foreign asset. It is not difficult to compute[40]

$$\delta_{\mathrm{d},i} = (\sigma_{R_i}^{-2} + \sigma_{A_i}^{-2})^{-1}\left(\sqrt{\kappa_i^2 + (1+\lambda^{-1})\rho_{v,i}^2(\sigma_{R_i}^{-2} + \sigma_{A_i}^{-2})} - \sqrt{\kappa_i^2 + \rho_{v,i}^2(\sigma_{R_i}^{-2} + \sigma_{A_i}^{-2})}\right),$$
$$\delta_{\mathrm{f},i} = \sigma_{R_j}^2\left(\sqrt{\kappa_j^2 + (1+\lambda^{-1})\rho_{v,j}^2/\sigma_{R_j}^2} - \sqrt{\kappa_j^2 + \rho_{v,j}^2/\sigma_{R_j}^2}\right).$$

In particular, since Asset j is to Country i as Asset j is to Country j without the signal A_j, $\delta_{\mathrm{f},i} = \lim_{\sigma_{A_j}^2\to\infty}\delta_{\mathrm{d},j}$.

It is straightforward to check that $\delta_{\mathrm{d},i}$ is strictly increasing in $\sigma_{A_i}^2$ with $\lim_{\sigma_{A_i}^2\to 0}\delta_{\mathrm{d},i} = 0$.[41] Therefore, finally, (i) $\delta_{\mathrm{d},i} < \delta_{\mathrm{f},j} = \lim_{\sigma_{A_i}^2\to\infty}\delta_{\mathrm{d},i}$; that is, each asset appears more ambiguous to foreign investors than to domestic investors; and (ii) $\delta_{\mathrm{d},i} < \delta_{\mathrm{f},i}$ for sufficiently small $\sigma_{A_i}^2$; that is, to investors in either country, the foreign asset appears more ambiguous than the domestic one if only the domestic news is sufficiently informative. The second observation in particular means that even when an asset is riskier than another in the sense that it unambiguously has the higher volatility, it can nevertheless be less ambiguous to informed investors if the ambiguity in question is about expected returns.

4 Portfolio Choice

In Section 4, I apply the model of learning to the consumption/portfolio choice problem of a log investor. The investor finances his intertemporal consumption by trading one risk-free asset (bond) and a number of risky assets (stocks). He believes, as is the prevailing view of the financial economics profession, that mean reversion in stock returns, or more generally a form of predictability in them, is a plausible assumption; but facing at the same time nonnegligible evidence that questions its validity, he fails to have full confidence in it.

In Section 4.1, I explain the setup. Sections 4.2 and 4.3 characterize the optimal demand for stocks. In Section 4.4, I consider the special case in which there is a single stock and the stock return volatility is constant. This simplification allows us to establish certain analytical properties of the optimal policy. In Section 4.4.3, I numerically compute the optimal policy and discuss its behavior in comparison with the related models by Epstein and Schneider (2007) and Miao (2009).

[40] Since we are interested in the limit of δ, which is diagonal, we may assume that the initial values of γ and δ are diagonal to begin with. Then the systems (R_i, A_i, x_i) and (R_j, x_j) can be looked at separately.

[41] Denote by $\sigma_{\mathrm{eff},i} \triangleq (\sigma_{R_i}^{-2} + \sigma_{A_i}^{-2})^{-1/2} < \sigma_{R_i}$ the effective imprecision of R_i and A_i en masse as signals to the hidden state x_i. $\delta_{\mathrm{d},i}$ becomes

$$\delta_{\mathrm{d},i} = \sigma_{\mathrm{eff},i}^2\left(\sqrt{\kappa_i^2 + (1+\lambda^{-1})\rho_{v,i}^2/\sigma_{\mathrm{eff},i}^2} - \sqrt{\kappa_i^2 + \rho_{v,i}^2/\sigma_{\mathrm{eff},i}^2}\right).$$

$\delta_{\mathrm{d},i}$ is strictly increasing in $\sigma_{\mathrm{eff},i}^2$; and $\sigma_{\mathrm{eff},i}^2$ in $\sigma_{A_i}^2$.

4.1 The Setup

As with the previous section, time is continuous and varies over $[0,T]$, $T \in (0,\infty)$.

4.1.1 Securities Market Dynamics

There is a single consumption good in the economy, which is continuously consumed and serves as the numeraire. The investor finances his consumption by trading one risk-free asset (bond) and $n_R \geq 1$ risky assets (stocks).

The interest rate on the bond is constant at $r \in \mathbb{R}$.

Regarding how the stock returns are generated, on the other hand, the investor entertains multiple theories. Specifically, the theories take the form of probability measures on a common measurable space: Let there be a measurable space $(\Omega, \mathcal{F})$, a set $\mathcal{Q}$ of probability measures (theoretical priors) on $(\Omega, \mathcal{F})$, and a filtration $\mathbf{F} = \{\mathcal{F}_t\}$ of $\mathcal{F}$. The theoretical priors are equivalent and $\mathbf{F}$ satisfies the usual conditions with respect to the theoretical priors. Under the theoretical prior $Q^{\bar{x},\eta} \in \mathcal{Q}$, where $(\bar{x},\eta) \in \mathbb{R}^{n_x} \times L^2([0,T],\mathbb{R}^{n_x})$ and $n_x \geq 1$, the cumulative return process $R = \{R(t), \mathcal{F}_t\}$ of the stocks is given by part of the solution to the system of SDEs

$$\mathrm{d}R(t) = (a_R(t,R,A) + b_R(t,R,A)x(t))\,\mathrm{d}t + \sigma_R(t,R,A)\,\mathrm{d}w(t), \tag{40}$$
$$\mathrm{d}A(t) = (a_A(t,R,A) + b_A(t,R,A)x(t))\,\mathrm{d}t + \sigma_A(t,R,A)\,\mathrm{d}w(t), \tag{41}$$
$$\mathrm{d}x(t) = \kappa(\bar{x} - x(t))\,\mathrm{d}t + \rho_w\,\mathrm{d}w(t) + \rho_v(\,\mathrm{d}v^{\bar{x},\eta}(t) + \eta(t)\,\mathrm{d}t).$$

Here, $A = \{A(t), \mathcal{F}_t\}$ is an n_A-dimensional process, $n_A \geq 0$; $x = \{x(t), \mathcal{F}_t\}$ is an n_x-dimensional process; $n_x \leq n_R + n_A$; $w = \{w(t), \mathcal{F}_t\}$ and $v^{\bar{x},\eta} = \{v^{\bar{x},\eta}(t), \mathcal{F}_t\}$ are independent Wiener processes of dimension $n_R + n_A$ and n_x, respectively; a_R, b_R, σ_R, a_A, b_A, and σ_A are nonanticipating path functionals from $[0,T] \times C([0,T],\mathbb{R}^{n_R+n_A})$ into $\mathbb{R}^{n_R}$, $\mathbb{R}^{n_R \times n_x}$, $\mathbb{R}^{n_R \times (n_R+n_A)}$, $\mathbb{R}^{n_A}$, $\mathbb{R}^{n_A \times n_x}$, and $\mathbb{R}^{n_A \times (n_R+n_A)}$, respectively; κ is an $n_x \times n_x$ diagonal matrix with positive entries, ρ_w is an $n_x \times (n_R + n_A)$ matrix, and ρ_v is an $n_x \times n_x$ invertible matrix.

The investor observes R and A but not x. A in this context represents the observable macroeconomic variables in addition to the stock returns themselves that affect the stock returns; and x the latent state of the economy. Note that the characterization here of expected returns is more general than was considered in the introduction; they are "linear" in mean-reverting factors, rather than being mean-reverting factors themselves.

To conform to the notation of Section 3, let $y \triangleq (R^\top, A^\top)^\top$ and $n_y \triangleq n_R + n_A$. Then the dynamics (40)-(41) of the observable processes can be rewritten compactly as

$$\mathrm{d}y(t) = (a(t,y) + b(t,y)x(t))\,\mathrm{d}t + \sigma(t,y)\,\mathrm{d}w(t)$$

where the definitions of a, b, and σ are obvious. I continue to adopt the slightly abusive notation $f(t) \equiv f(t,\omega) \equiv f(t,y(\omega))$ for the path functionals f.

4.1.2 The Investor's Preferences

The investor has the Chen-Epstein recursive multiple-priors utility with log felicity. His conditional preferences at time $t \in [0,T]$ are represented by

$$\min_{\xi \in \Xi} \mathrm{E}^{P^\xi} \int_t^T e^{-\beta s} \log(c(s))\,\mathrm{d}s, \;\; c \in \mathcal{C}.$$

Under a generic preferential prior P^{ξ},

$$\mathrm{d}R(t) = (a_R(t) + b_R(t)m_t^*)\,\mathrm{d}t + \sigma_R(t)(\,\mathrm{d}\epsilon^{\xi}(t) + \xi(t)\,\mathrm{d}t)$$

where $m_t^* \equiv m^{\bar{x}_t^*,\eta_t^*}(t)$ is the maximum plausibility estimate of the conditional expectation of $x(t)$ given $\mathcal{G}_t$ and $\epsilon^{\xi} = \{\epsilon^{\xi}(t), \mathcal{G}_t\}$ is a P^{ξ}-Wiener process of dimension n_y. $\Xi(t)$ thus acquires a more specific interpretation as the ambiguity in the contemporaneous price of risk.

4.1.3 Trading Strategies and the Budget Constraint

A $(1+n_R)$-dimensional process (Π°, Π), $\Pi(t) = (\Pi_1(t), \ldots, \Pi_{n_R}(t))^{\top}$, is a *trading strategy* if $\mathbf{G}$-progressive and

$$\int_0^T (|\Pi^{\circ}(t)| + |\Pi(t)|^2)\,\mathrm{d}t < \infty.$$

Π° represents the amount of money invested in the bond and Π those invested in the stocks. A trading strategy (Π°, Π) *finances a consumption plan* $c \in \mathcal{C}$ if $\Pi^{\circ}(T) + \Pi(T)^{\top}\mathbf{1}_{n_R} \geq 0$ and

$$\mathrm{d}(\Pi^{\circ}(t) + \Pi(t)^{\top}\mathbf{1}_{n_R}) = \Pi^{\circ}(t)r\,\mathrm{d}t + \Pi(t)^{\top}\,\mathrm{d}R(t) - c(t)\,\mathrm{d}t$$

where $\mathbf{1}_{n_R}$ denotes the n_R-dimensional vector of ones. Denote the wealth process $\Pi^{\circ} + \Pi^{\top}\mathbf{1}_{n_R}$ by W. W satisfies

$$\mathrm{d}W(t) = (W(t) - \Pi(t)^{\top}\mathbf{1}_{n_R})r\,\mathrm{d}t + \Pi(t)^{\top}\,\mathrm{d}R(t) - c(t)\,\mathrm{d}t \tag{42}$$

with initial condition $W(0) = \Pi^{\circ}(0) + \Pi(0)^{\top}\mathbf{1}_{n_R}$. In fact, W is the unique strong solution to the last equation, and therefore, we can suppress Π° and identify a trading strategy with Π. A pair (Π, c) is *admissible for initial wealth* $W(0)$ if the corresponding wealth process $W^{\Pi,c,W(0)}$ is uniformly bounded below.

The market is dynamically incomplete if $n_A > 0$. Let

$$\zeta(t) \triangleq \sigma_R(t)^{\top}(\sigma_R(t)\sigma_R(t)^{\top})^{-1}(a_R(t) + b_R(t)m_t^* - r\mathbf{1}_{n_R}).$$

A consumption process $c \in \mathcal{C}$ can be financed by some trading strategy if and only if it satisfies the following static budget constraint:

$$\sup_{\nu \in \mathrm{Ker}(\sigma_R)} \mathrm{E}^{P^0} \int_0^T \mathcal{E}^{-(\zeta+\nu)}(t)e^{-rt}c(t)\,\mathrm{d}t \leq W(0) \tag{43}$$

where $\mathrm{Ker}(\sigma_R)$ denotes the set of processes ν such that $\sigma_R(t,\omega)\nu(t,\omega) = 0$ a.e. (He and Pearson, 1991; Karatzas et al., 1991; Cuoco, 1997).

Remark 4.1. *If the investor had full confidence in a simple theory $Q^{\bar{x},0} \in \mathcal{Q}$, then the present model would have as special cases the Bayesian learning models of Lakner (1998), Xia (2001), Zohar (2001), and Brendle (2006), in which the unobservable instantaneous expected return process follows an Ornstein-Uhlenbeck process. In other words, this section extends the latter models to a case of ambiguity.*

4.2 Optimal Consumption and Portfolio

Let $\mathcal{C}^2(u) \subset \mathcal{C}$ denote the set of consumption processes such that

$$\mathrm{E}^{P^0} \int_0^T [\log(c(t))]^2 \, \mathrm{d}t < \infty.$$

I define the investor's problem to be

$$\sup_{c \in \mathcal{C}^2(u)} \min_{\xi \in \Xi} \mathrm{E}^{P^\xi} \int_0^T e^{-\beta t} \log(c(t)) \, \mathrm{d}t \tag{44}$$

subject to the budget constraint (43). The objective function in (44) is finite for all $(c, \xi) \in \mathcal{C}^2(u) \times \Xi$ due to the definition of $\mathcal{C}^2(u)$ and the uniform boundedness of Ξ. Let $\mathcal{C}_{\text{budget}} \subset \mathcal{C}$ denote the set of consumption processes that satisfy the budget constraint.

Lemma 4.1. *The minimax theorem holds, that is,*

$$\sup_{c \in \mathcal{C}^2(u) \cap \mathcal{C}_{\text{budget}}} \min_{\xi \in \Xi} \mathrm{E}^{P^\xi} \int_0^T e^{-\beta t} \log(c(t)) \, \mathrm{d}t = \min_{\xi \in \Xi} \sup_{c \in \mathcal{C}^2(u) \cap \mathcal{C}_{\text{budget}}} \mathrm{E}^{P^\xi} \int_0^T e^{-\beta t} \log(c(t)) \, \mathrm{d}t.$$

Remark 4.2. *It is clear from the proof that the claim is true for any concave felicity, with the corresponding change to the definition of $\mathcal{C}^2(u)$.*

Proposition 4.1. *For a given $\xi \in \Xi$, the inner supremum*

$$\sup_{c \in \mathcal{C}^2(u) \cap \mathcal{C}_{\text{budget}}} \mathrm{E}^{P^\xi} \int_0^T e^{-\beta t} \log(c(t)) \, \mathrm{d}t \tag{45}$$

equals

$$\begin{aligned} &-\frac{1 - e^{-\beta T}}{\beta} \log\left(\frac{1 - e^{-\beta T}}{\beta}\right) + \frac{\beta - r}{\beta}\left(Te^{-\beta T} - \frac{1 - e^{-\beta T}}{\beta}\right) + \frac{1 - e^{-\beta T}}{\beta} \log W(0) \\ &\qquad + \mathrm{E}^{P^\xi} \int_0^T \frac{e^{-\beta t} - e^{-\beta T}}{\beta} \frac{1}{2} |\zeta(t) + \sigma_R(t)^\top (\sigma_R(t)\sigma_R(t)^\top)^{-1} \sigma_R(t)\xi(t)|^2 \, \mathrm{d}t. \end{aligned} \tag{46}$$

Let ξ^* denote the minimizer of the last expression:

$$\xi^* \triangleq \arg\min_{\xi \in \Xi} \mathrm{E}^{P^\xi} \int_0^T \frac{e^{-\beta t} - e^{-\beta T}}{\beta} \frac{1}{2} |\zeta(t) + \sigma_R(t)^\top (\sigma_R(t)\sigma_R(t)^\top)^{-1} \sigma_R(t)\xi(t)|^2 \, \mathrm{d}t. \tag{47}$$

The optimal consumption process is given by

$$c^*(t) = \beta W(0) e^{rt} \frac{e^{-\beta t}}{1 - e^{-\beta T}} \frac{\mathcal{E}^{\xi^*}(t)}{\mathcal{E}^{-(\zeta + \nu^*)}(t)}, \tag{48}$$

$$\nu^*(t) = [\sigma_R(t)^\top (\sigma_R(t)\sigma_R(t)^\top)^{-1} \sigma_R(t) - I_{n_y}] \xi^*(t). \tag{49}$$

Hence the key is to solve (47). Note for later reference that

$$\zeta(t) + \sigma_R(t)^\top(\sigma_R(t)\sigma_R(t)^\top)^{-1}\sigma_R(t)\xi(t) \\ = \sigma_R(t)^\top(\sigma_R(t)\sigma_R(t)^\top)^{-1}(a_R(t) + b_R(t)m_t^* - r\mathbf{1}_{n_R} + \sigma_R(t)\xi(t)).$$

To find the trading strategy that finances c^*, observe first that the wealth process corresponding to c^* is

$$\begin{aligned} W^*(t) &= \frac{1}{B(t)^{-1}\mathcal{E}^{-(\zeta+\nu^*)}(t)} \mathrm{E}^{P^0}\left(\int_t^T B(s)^{-1}\mathcal{E}^{-(\zeta+\nu^*)}(s)c^*(s)\,\mathrm{d}s\middle|\,\mathcal{G}_t\right) \\ &= W(0)e^{rt}\frac{e^{-\beta t} - e^{-\beta T}}{1 - e^{-\beta T}}\frac{\mathcal{E}^{\xi^*}(t)}{\mathcal{E}^{-(\zeta+\nu^*)}(t)}. \end{aligned} \tag{50}$$

Thus its differential is

$$\mathrm{d}W^*(t) = W^*(t)(\zeta(t) + \nu^*(t) + \xi^*(t))^\top\,\mathrm{d}\epsilon + \cdot\,\mathrm{d}t.$$

Comparing the last expression with (42) and recalling (49), we see that

$$\begin{aligned} \pi^*(t) &\triangleq \Pi^*(t)/W^*(t) \\ &= (\sigma_R(t)\sigma_R(t)^\top)^{-1}\left(a_R(t) + b_R(t)m_t^* - r\mathbf{1}_{n_R} + \sigma_R(t)\xi^*(t)\right) \end{aligned} \tag{51}$$

where π^* denotes the optimal *fraction* of wealth invested in the stock.

The optimal consumption plan c^* found above equals that of the Bayesian investor with unique prior P^{ξ^*}. Accordingly, π^* equals the stock demand of the same Bayesian investor, the term involving ξ^* accounting for the discrepancy between P^{ξ^*} and P^0. This observation also suggests that as is characteristic of Bayesian log investors, the optimal consumption is given by a fraction of wealth independent of other state variables, or precisely,

$$c^*(t) = \frac{\beta}{1 - e^{-\beta(T-t)}}W^*(t),$$

as can be verified from (48) and (50).

4.3 Markovian Characterization

Suppose the economy is Markovian, that is,

$$f(t, R, A) = f(t, R(t), A(t))$$

where $f = a$, b, or σ. Then the investor's information can be summarized by a finite number of Markovian variables.

Observe first that the Bayesian investor who has full confidence in a simple theory $Q^{\bar{x},0} \in \mathcal{Q}$ has the following as the state variables (see Proposition 3.2):

$$R(t),\ A(t),\ m^{\bar{x},0}(t),\ \text{and}\ \gamma(t).$$

Our investor also has these as state variables, with the obvious replacement of $m^{\bar{x},0}(t)$ by m^*_t, that is,

$$R(t),\ A(t),\ m^*_t,\ \text{and}\ \gamma(t), \tag{52}$$

and the following in addition:

$$\bar{x}^*_t,\ \sigma_{\bar{x}^*}(t),\ \mathcal{I}_{\bar{x}}(t)^{-1},\ \text{and}\ \delta(t). \tag{53}$$

See Propositions 3.6, 3.7, and 3.9. The first three of (53) originates from the estimation of $\bar{x}$; the last is needed to describe the set of one-step-ahead conditionals $\Xi(t)$. The standard control approach to the minimization (47) requires that $\Xi(t)$, $\zeta(t)$, and $\sigma_R(t)$ be functions of some (multidimensional) Markov process, and Propositions 3.6 and 3.7 confirm that the variables identified in (52) and (53) form a closed system of Markovian variables. Collect them in Z,

$$Z \triangleq (R, A, m^*, \gamma, \bar{x}^*, \sigma_{\bar{x}^*}, \mathcal{I}_{\bar{x}}^{-1}, \delta)^\top,$$

and write

$$\mathrm{d}Z(t) = \mu_Z(t, Z(t))\,\mathrm{d}t + \sigma_Z(t, Z(t))\,\mathrm{d}\epsilon(t). \tag{54}$$

Remark 4.3. *Some of the state variables identified above may be redundant. For example, if a, b, and σ are deterministic functions of time independent of R and A, then it suffices to take as state variables m^* and $\bar{x}^*$. See Section 4.4 below.*

Define the value function as

$$J(t,Z) \triangleq \min_{\xi\in\Xi} \mathrm{E}^{P^\xi}\left(\int_t^T \frac{e^{-\beta s}-e^{-\beta T}}{\beta} \times \frac{1}{2}|\zeta(s)+\sigma_R(s)^\top(\sigma_R(s)\sigma_R(s)^\top)^{-1}\sigma_R(s)\xi(s)|^2\,\mathrm{d}s \,\middle|\, Z(t)=Z\right)$$

subject to the state dynamics (54). Picking a particular $\xi\in\Xi$ is to say that $\epsilon^\xi = \{\epsilon(t), \mathcal{G}_t\}$ defined by $\mathrm{d}\epsilon^\xi(t) = \mathrm{d}\epsilon(t) - \xi(t)\,\mathrm{d}t$ is a Wiener process. Hence

$$J(t,Z) = \min_{\xi\in\Xi} \mathrm{E}^{P^0}\left(\int_t^T \frac{e^{-\beta s}-e^{-\beta T}}{\beta} \times \frac{1}{2}|\zeta^\xi(s)+\sigma^\xi_R(s)^\top(\sigma^\xi_R(s)\sigma^\xi_R(s)^\top)^{-1}\sigma^\xi_R(s)\xi(s)|^2\,\mathrm{d}s \,\middle|\, Z^\xi(t)=Z\right) \tag{55}$$

subject to

$$\mathrm{d}Z^\xi(t) = \mu_Z(t, Z^\xi(t))\,\mathrm{d}t + \sigma_Z(t, Z^\xi(t))(\,\mathrm{d}\epsilon(t) + \xi(t)\,\mathrm{d}t)$$

where $\sigma^\xi_R(s) \equiv \sigma_R(s, R^\xi(s), A^\xi(s))$. (55) is linear-quadratic in the control, although not in the state and hence not linear-quadratic in the classical sense. The corresponding Hamilton-Jacobi-Bellman (HJB) equation is

$$\begin{aligned}0 = \min_{\xi(t)\in\Xi(t,Z)} \Bigg(&\partial_t J(t,Z) + (\partial_Z J(t,Z))^\top(\mu_Z(t,Z)+\sigma_Z(t,Z)\xi(t)) \\ &+ \frac{1}{2}\operatorname{tr}[(\partial^2_Z J(t,Z))\sigma_Z(t,Z)\sigma_Z(t,Z)^\top] \\ &+ \frac{e^{-\beta t}-e^{-\beta T}}{\beta}\frac{1}{2}|\zeta(t,Z)+\sigma_R(t,Z)^\top(\sigma_R(t,Z)\sigma_R(t,Z)^\top)^{-1}\sigma_R(t,Z)\xi(t)|^2\Bigg)\end{aligned} \tag{56}$$

with boundary condition $J(T, Z) = 0$ for all Z. In general, (56) is of *degenerate* parabolic type and we can only say that the value function is a viscosity solution of (56). But see Section 4.4.1, where I consider a special case in which the value function is a unique classical solution to the HJB equation.

4.4 Examples

To gain intuitions, I consider in this section the special case in which there is a single stock, the stock return volatility is constant, and there are no other observable macroeconomic indicators that affect stock returns. That is, $n_R = 1$ and $n_A = 0$ so that $n_y = n_x = 1$ and $\sigma(t, y) = \sigma_R(t, y) = \sigma_R \in (0, \infty)$ for all (t, y). Assume furthermore $a_R \equiv 0$ and $b_R \equiv 1$. This setup is simple but rich enough to let us discuss key aspects of the optimal policy.

4.4.1 $\bar{x}$ Known

Suppose first that $\bar{x}$ is known.

Optimal Policy Revisited Under the aforementioned assumptions, the investor's problem is Markovian and his optimal stock demand can be written in a simple feedback form.

Recall Section 4.3 and note that (i) R and A are redundant as state variables because σ is constant, (ii) γ and δ are redundant because they are deterministic, and (iii) $\bar{x}^* \equiv \bar{x}$, $\sigma_{\bar{x}^*}$, and $\mathcal{I}_{\bar{x}}^{-1}$ are redundant because $\bar{x}$ is known. It thus suffices to take m^* as the sole state variable ($Z = m^*$). The controlled state dynamics is (see (27))

$$\begin{aligned} \mathrm{d}m_t^{*,\xi} &= \kappa(\bar{x} - m_t^{*,\xi})\,\mathrm{d}t + (\rho_w \sigma_R + \gamma(t) + \delta(t))\sigma_R^{-1}(\,\mathrm{d}\epsilon(t) + \xi(t)\,\mathrm{d}t) \\ &=: \mu_{m^*}(m_t^{*,\xi})\,\mathrm{d}t + \sigma_{m^*}(t)(\,\mathrm{d}\epsilon(t) + \xi(t)\,\mathrm{d}t). \end{aligned}$$

The price of risk under P^0 is simplified to

$$\zeta(m^*) = \frac{m^* - r}{\sigma_R}.$$

$\Xi(t)$ is given by an interval $[-\bar{\xi}(t), \bar{\xi}(t)]$ where

$$\bar{\xi}(t) \triangleq \frac{\sqrt{2\alpha\delta(t)}}{\sigma_R};$$

$\bar{\xi}(t)$ measures the magnitude of the ambiguity in the price of risk and is increasing in the investor's conservatism in model selection α and the estimation ambiguity $\delta(t)$. Also, it decreases monotonically and deterministically over time, converging to a constant, as is a property of δ. The estimated equity premium is $m^* - r$ and the ambiguity in the equity premium is $\sigma_R \bar{\xi}(t) = \sqrt{2\alpha\delta(t)}$. (Unless necessary, the [true or estimated] instantaneous equity premium will be referred to simply as the [true or estimated] equity premium.)

Next, the HJB equation (56) is simplified to

$$0 = \min_{\xi(t)\in\Xi(t)} \Bigg(\partial_t J(t,m^*) + (\partial_{m^*} J(t,m^*))(\mu_{m^*}(m^*) + \sigma_{m^*}(t)\xi(t)) \\ + \frac{1}{2}(\partial^2_{m^*} J(t,m^*))\sigma_{m^*}(t)^2 + \frac{e^{-\beta t} - e^{-\beta T}}{\beta}\frac{1}{2}(\zeta(m^*) + \xi(t))^2 \Bigg) \quad (57)$$

with boundary condition $J(T,m^*) = 0$ for all m^*. It is still not clear if (57) allows for an analytical solution, but we can now check some basic properties of the value function.[42] $C^{1,2}([0,T]\times\mathbb{R})$ denotes the set of real-valued functions f from $[0,T]\times\mathbb{R}$ such that $f(t,m^*)$ is continuously differentiable in t and twice continuously differentiable in m^*; and $C_p([0,T]\times\mathbb{R})$ the set of real-valued functions f from $[0,T]\times\mathbb{R}$ that are continuous and satisfy the polynomial growth condition:

$$|f(t,m^*)| \leq K(1 + |m^*|^n) \text{ for all } m^* \in \mathbb{R}$$

for some nonnegative constants K and n. Assume for the rest of Section 4.4.1,

Assumption 4.1. *$\sigma^2_{m^*} : [0,T] \to \mathbb{R}$ is bounded below away from zero.*

The assumption trivially holds if $\rho_w \geq 0$.

Proposition 4.2. *(i) The partial differential equation* (57) *with its boundary condition has a unique solution $K \in C^{1,2}([0,T]\times\mathbb{R}) \cap C_p([0,T]\times\mathbb{R})$.*

(ii) K is the value function, that is, $K = J$.

(iii)

$$\xi^*(t,m^*) = \max\left\{-\bar{\xi}(t), \min\left\{\bar{\xi}(t), \xi^U(t,m^*)\right\}\right\},$$
$$\xi^U(t,m^*) \triangleq -\zeta(m^*) - \frac{\beta e^{\beta t}}{1 - e^{-\beta(T-t)}}\sigma_{m^*}(t)\partial_{m^*} J(t,m^*).$$

Thus, in particular, the optimal control $\xi^* : [0,T]\times\mathbb{R} \to \mathbb{R}$ is continuous.

The expression for the optimal stock demand (51) becomes

$$\pi^*(t,m^*) = \frac{m^* - r + \sigma_R \xi^*(t,m^*)}{\sigma_R^2} \\ = \max\left\{\frac{m^* - r - \sigma_R\bar{\xi}(t)}{\sigma_R^2}, \min\left\{\frac{m^* - r + \sigma_R\bar{\xi}(t)}{\sigma_R^2}, -\frac{1}{\sigma_R^2}\frac{\beta e^{\beta t}}{1 - e^{-\beta(T-t)}}\sigma_R\sigma_{m^*}(t)\partial_{m^*} J(t,m^*)\right\}\right\}. \quad (58)$$

Lemma 4.2. *(i) $J(t,m^*)$ is convex in m^*.*

(ii)

$$\partial_{m^*} J(t,m^*) = \mathrm{E}^{P^0}\left(\int_t^T \frac{e^{-\beta s} - e^{-\beta T}}{\beta}\frac{e^{-\kappa(s-t)}}{\sigma_R}\frac{m^{*,\xi^*}_s - r + \sigma_R\xi^*(s)}{\sigma_R}\,\mathrm{d}s \,\middle|\, m^{*,\xi^*}_t = m^*\right).$$

[42] It is possible to formulate (57) as a free boundary problem and characterize the solution to a certain degree (cf. Davis and Norman (1990)), but there is little practical benefit and I do not pursue this direction.

From the convexity, that is, from the fact that $\partial_{m^*} J(t, m^*)$ is nondecreasing in m^*, it follows that

$$\frac{m^* - r - \sigma_R \bar{\xi}(t)}{\sigma_R^2} = -\frac{1}{\sigma_R^2} \frac{\beta e^{\beta t}}{1 - e^{-\beta(T-t)}} \sigma_R \sigma_{m^*}(t) \partial_{m^*} J(t, m^*)$$

and

$$\frac{m^* - r + \sigma_R \bar{\xi}(t)}{\sigma_R^2} = -\frac{1}{\sigma_R^2} \frac{\beta e^{\beta t}}{1 - e^{-\beta(T-t)}} \sigma_R \sigma_{m^*}(t) \partial_{m^*} J(t, m^*)$$

as equations in m^* each have a unique solution, $\overline{m^*}(t)$ and $\underline{m^*}(t) < \overline{m^*}(t)$, respectively. π^* can be rewritten as

$$\pi^*(t, m^*) = \begin{cases} \dfrac{m^* - r + \sigma_R \bar{\xi}(t)}{\sigma_R^2} & \text{if } m^* < \underline{m^*}(t) \\ \dfrac{m^* - r - \sigma_R \bar{\xi}(t)}{\sigma_R^2} & \text{if } m^* > \overline{m^*}(t) \\ -\dfrac{1}{\sigma_R^2} \dfrac{\beta e^{\beta t}}{1 - e^{-\beta(T-t)}} \sigma_R \sigma_{m^*}(t) \partial_{m^*} J(t, m^*) & \text{if } m^* \in [\underline{m^*}(t), \overline{m^*}(t)]. \end{cases}$$

Since ξ^* is bounded, the effect of ambiguity on $\partial_{m^*} J(t, m^*)$ is negligible for m^*s with a large absolute value. Combined with convexity, this implies that $m^* \mapsto J(t, m^*)$ is U-shaped. (Epstein and Schneider (2007) in p. 1296 make a similar observation from a numerical exercise.) As with his Bayesian counterpart with unique theoretical prior $Q^{\bar{x},0}$, our multiple-priors investor, too, is better off when the estimated equity premium is further away from zero, that is, when the stocks are (locally, in expected terms) more distinct from the bond. The U-shape implies that the optimal policy may have curvature in the central region $m^* \in [\underline{m^*}(t), \overline{m^*}(t)]$.

Compared to the Bayesian policy, our investor's stock demand is (i) shifted up by the ambiguity in the equity premium (divided by the return variance) when the estimated equity premium $m^* - r$ is sufficiently small (in the sense of $<$ on the real line), (ii) shifted down by the same amount when $m^* - r$ is sufficiently large, and (iii) proportional to the negative of the instantaneous covariation between the stock return and the state $(-\sigma_R \sigma_{m^*}(t))$ and the first derivative of the value function $(\partial_{m^*} J(t, m^*))$, when $m^* - r$ is intermediate. Clearly, the last case is reminiscent of Merton's (1973) hedging demand; it tells the investor to hold more of the stock if it pays in cases of low continuation utility. (But it is not exactly the same as Merton's hedging demand. His is such that the investor holds more of the assets that pay in cases of low *consumption*, or equivalently, high marginal utility.) I will have a deeper look at the quantity $-\sigma_R \sigma_{m^*}(t) \partial_{m^*} J(t, m^*)$ later, but to talk about hedging, first we have to clarify the myopic demand.

Myopic Demand The myopic demand is defined to be

$$\pi^*_{\text{myopic}}(t, m^*) \triangleq \left[\lim_{t \to T} \pi^*(t, m^*) \right]_{T=t}.$$

Proposition 4.3.

$$\pi^*_{\text{myopic}}(t, m^*) = \begin{cases} \dfrac{m^* - r + \sigma_R \bar{\xi}(t)}{\sigma_R^2} & \textit{if } m^* - r < -\sigma_R \bar{\xi}(t) \\ \dfrac{m^* - r - \sigma_R \bar{\xi}(t)}{\sigma_R^2} & \textit{if } m^* - r > +\sigma_R \bar{\xi}(t) \\ 0 & \textit{if } -\sigma_R \bar{\xi}(t) \le m^* - r \le \sigma_R \bar{\xi}(t). \end{cases}$$

The myopic demand is more conservative than that of the Bayesian investor with unique theoretical prior $Q^{\bar{x},0}$ in that in absolute values, the former is dominated by the latter:

$$|\pi^*_{\text{myopic}}(t, m^*)| \le \left| \frac{m^* - r}{\sigma_R^2} \right| \text{ for all } m^* \text{ and } < \left| \frac{m^* - r}{\sigma_R^2} \right| \text{ for all } m^* \ne r.$$

(I am comparing the feedback policies, considering m^* to be signifying the estimate of each investor. The actual values of m^* will differ between the two investors.) Furthermore, there is a range of estimated equity premia for which our investor neither buys nor sells short the stock. Say that the estimated equity premium is *unambiguously positive* if it is greater than the ambiguity in the equity premium, that is, if $m^* - r > \sigma_R \bar{\xi}(t)$; *unambiguously negative* if $m^* - r < -\sigma_R \bar{\xi}(t)$; and *not unambiguously distinct from zero*, otherwise. Then, the observation, rephrased, is that the multiple-priors investor, if myopic, does not participate in the stock market when his estimate of the equity premium is not unambiguously distinct from zero; and participates when it is unambiguously positive or negative but invests a smaller fraction of his wealth than the Bayesian counterpart with the same estimate would. See Dow and Werlang (1992), who first presented a nonparticipation result for ambiguity-averse (in the sense of Schmeidler (1989)) investors.

Hedging Demand Under risk, log investors do not hedge; under ambiguity, they do.

Recall the total demand (58) and let

$$\pi^{**}_{\text{hedging}}(t, m^*) \triangleq -\frac{1}{\sigma_R^2} \frac{\beta e^{\beta t}}{1 - e^{-\beta(T-t)}} \sigma_R \sigma_{m^*}(t) \partial_{m^*} J(t, m^*).$$

As noted earlier, $\pi^{**}_{\text{hedging}}$ reflects the investor's desire to hedge against adverse changes in the investment opportunities. Under ambiguity, an adverse change in the investment opportunities is a change in the state variables that is associated with a decrease in continuation utility. In the present case, if the (estimated) equity premium is sufficiently large that $\partial_{m^*} J(t, m^*) > 0$, then the investor would fear a decrease in the equity premium, that is, its becoming ambiguous, and want to transfer wealth to states with lower equity premia. And he could do this by holding more of the stock if it pays at times of lower equity premia and less of it if it does not.

However, the desire to hedge does not fully realize, and how much of it realizes depends on the magnitude of the ambiguity present. The total demand $\pi^*(t, m^*)$ is given by $\pi^{**}_{\text{hedging}}(t, m^*)$ confined between $(m^* - r \pm \sigma_R \bar{\xi}(t))/\sigma_R^2$, which collapse to the Bayesian demand $(m^* - r)/\sigma_R^2$ when no ambiguity is present. Hence I call $\pi^{**}_{\text{hedging}}$ the shadow hedging demand. Finally, based on the interpretation of $\pi^{**}_{\text{hedging}}$, the difference $\pi^* - \pi^*_{\text{myopic}}$ between the total demand and the myopic demand is called the hedging demand, although the intent is not fully realized:

$$\begin{aligned}\pi^*_{\text{hedging}}(t,m^*) &\triangleq \pi^*(t,m^*) - \pi^*_{\text{myopic}}(t,m^*)\\ &= \max\left\{\frac{m^*-r-\sigma_R\bar{\xi}(t)}{\sigma_R^2}, \min\left\{\frac{m^*-r+\sigma_R\bar{\xi}(t)}{\sigma_R^2}, \pi^{**}_{\text{hedging}}(t,m^*)\right\}\right\}\\ &\quad - \max\left\{\frac{m^*-r-\sigma_R\bar{\xi}(t)}{\sigma_R^2}, \min\left\{\frac{m^*-r+\sigma_R\bar{\xi}(t)}{\sigma_R^2}, 0\right\}\right\}.\end{aligned}$$

Long-horizon, multiple-priors log investors' nonmyopic behavior was first observed in discrete time by Epstein and Schneider (2007) and in continuous time by Hernández-Hernández and Schied (2007a).

In Comparison with Merton (1973) The shadow hedging demand $\pi^{**}_{\text{hedging}}$ is reminiscent of Merton's (1973), but not the same. The difference lies in what are adverse changes in the investment opportunities. Under risk, they are associated with low consumption; under ambiguity, with low continuation utility.

To draw further comparison between $\pi^{**}_{\text{hedging}}$ and Merton's hedging demand, recall that the latter is the position in the stock that minimizes the volatility of consumption. On the other hand, $\pi^{**}_{\text{hedging}}$ is the position in the stock that minimizes (to zero) the effect of misspecification on continuation utility. To elaborate, let

$$V(t,m^*,W) \triangleq \mathrm{E}^{P^{\xi^*}}\left(\int_t^T e^{-\beta s}\log(c^*(s))\,\mathrm{d}s \,\middle|\, m^*_t = m^*, W^{\pi^*,c^*}(t) = W\right).$$

As is characteristic of log investors, V additively separates to a part depending only on (t, W^*) and another depending only on (t, m^*), and I have been focusing on the latter denoted by J. Let further

$$f^{\xi}(t) \triangleq \frac{\mathrm{E}^{P^{\xi}}\left[\mathrm{d}V(t,m^*_t,W^{\pi,c}(t))\,|\, m^*_t = m^*, W^{\pi,c}(t) = W\right]}{\mathrm{d}t}$$

and observe that

$$\partial_{\xi(t)}(f^{\xi}(t) - f^0(t)) = W\pi(t)\sigma_R\partial_W V(t,m^*,W) + \sigma_{m^*}(t)\partial_{m^*}V(t,m^*,W).$$

From (46),

$$\partial_W V(t,m^*,W) = \frac{e^{-\beta t} - e^{-\beta T}}{\beta}\frac{1}{W} \text{ and } \partial_{m^*}V(t,m^*,W) = \partial_{m^*}J(t,m^*).$$

It follows that $|\partial_{\xi(t)}(f^{\xi}(t) - f^0(t))|$ attains its minimum (zero) at $\pi(t) = \pi^{**}_{\text{hedging}}(t,m^*)$.

4.4.2 $\bar{x}$ Unknown and Ambiguous

Suppose now that the investor does not know the value of $\bar{x}$ and entertains all the theoretical priors $\mathcal{Q} = \{Q^{\bar{x},\eta} : (\bar{x},\eta) \in \mathbb{R} \times L^2([0,T],\mathbb{R})\}$.

As before, R and A are redundant as state variables because σ is constant, and γ, δ, $\sigma_{\bar{x}^*}$, and $\mathcal{I}_{\bar{x}}^{-1}$ are redundant because they are deterministic. But now $\bar{x}^*$ needs to be taken as a state variable as well as m^*:

$$Z = (m^*, \bar{x}^*)$$

with dynamics (see Proposition 3.6)

$$\begin{aligned} \mathrm{d}Z^\xi(t) &= \begin{pmatrix} \kappa(\bar{x}_t^{*,\xi} - m_t^{*,\xi}) \\ 0 \end{pmatrix} \mathrm{d}t + \begin{pmatrix} \rho_w \sigma_R + \gamma(t) + \delta(t) \\ \kappa^{-1}\sigma_{\bar{x}^*}(t) \end{pmatrix} \sigma_R^{-1}(\,\mathrm{d}\epsilon(t) + \xi(t)\,\mathrm{d}t) \\ &= \mu_Z(Z^\xi(t))\,\mathrm{d}t + \sigma_Z(t)(\,\mathrm{d}\epsilon(t) + \xi(t)\,\mathrm{d}t). \end{aligned}$$

Since the diffusion matrix $\sigma_Z \sigma_Z^\top$ is degenerate, the value function J may not be differentiable. I assume nevertheless that $\partial_Z J(t, Z)$ exists everywhere and write

$$\pi^*(t, m^*, \bar{x}^*) = \max\left\{ \frac{m^* - r - \sigma_R \bar{\xi}(t)}{\sigma_R^2}, \min\left\{ \frac{m^* - r + \sigma_R \bar{\xi}(t)}{\sigma_R^2}, \pi^{**}_{\text{hedging}}(t, m^*, \bar{x}^*) \right\}\right\},$$

$$\begin{aligned} \pi^{**}_{\text{hedging}}(t, m^*, \bar{x}^*) \triangleq &-\frac{1}{\sigma_R^2} \frac{\beta e^{\beta t}}{1 - e^{-\beta(T-t)}} (\rho_w \sigma_R + \gamma(t) + \delta(t)) \partial_{m^*} J(t, m^*, \bar{x}^*) \\ &- \frac{1}{\sigma_R^2} \frac{\beta e^{\beta t}}{1 - e^{-\beta(T-t)}} \kappa^{-1} \sigma_{\bar{x}^*}(t) \partial_{\bar{x}^*} J(t, m^*, \bar{x}^*). \end{aligned}$$

I call the first term of $\pi^{**}_{\text{hedging}}$ the m^*-shadow hedging demand and the second the $\bar{x}^*$-shadow hedging demand.

4.4.3 Numerical Analysis

Continue to assume that the investor entertains all the theoretical priors $\mathcal{Q} = \{Q^{\bar{x},\eta} : (\bar{x}, \eta) \in \mathbb{R} \times L^2([0,T], \mathbb{R})\}$. In this section, I numerically compute the optimal stock demand $\pi^*(t, m^*, \bar{x}^*)$ and discuss its behavior.

The securities market model is calibrated based on Barberis (2000):[43]

$$\begin{aligned} \mathrm{d}R(t) &= x(t)\,\mathrm{d}t + 0.1428\,\mathrm{d}w(t), \\ \mathrm{d}x(t) &= 0.2743(\bar{x} - x(t))\,\mathrm{d}t - 0.0392\,\mathrm{d}w(t) + 0.0361\,\mathrm{d}v^{\bar{x},0}(t), \end{aligned}$$

and $r = 0.0432$ (all numbers are annual). The investor has observed 20 years of data and now faces a 10-year investment horizon. $\beta = 0.03$, $\lambda = \infty$, and $\alpha = 0.38$. These parameters translate into an ambiguity in the equity premium of 0.01. Also, $\sigma_Z(20) = (0.007, 0.009)^\top$.

[43] I annualized his monthly estimates (left panel of his Table II). His estimation is based on the monthly NYSE value-weighted returns as calculated by the CRSP, from June 1952 to December 1995.

Barberis assumes that excess stock returns are predicted by the dividend-price ratio, whereas the predictive variables of the present model, x, are unobservable. Hence, I calibrated the SDE for x so that the SDE for $m^{\bar{x},0}$ matches Barberis's estimation:

$$\mathrm{d}m^{\bar{x},0}(t) = 0.2743(\bar{x} - m^{\bar{x},0}(t))\,\mathrm{d}t - 0.0031\,\mathrm{d}\bar{w}^{\bar{x},0}(t)$$

where $-0.0031 = \lim_{t\to\infty}(\rho_w + \gamma(t)/\sigma_R)$. To be precise, Barberis finds, in accordance with other empirical works, excess stock returns and the dividend-price ratio to be highly negatively correlated (-0.9351), and I set R and $m^{\bar{x},0}$ to be perfectly negatively correlated.

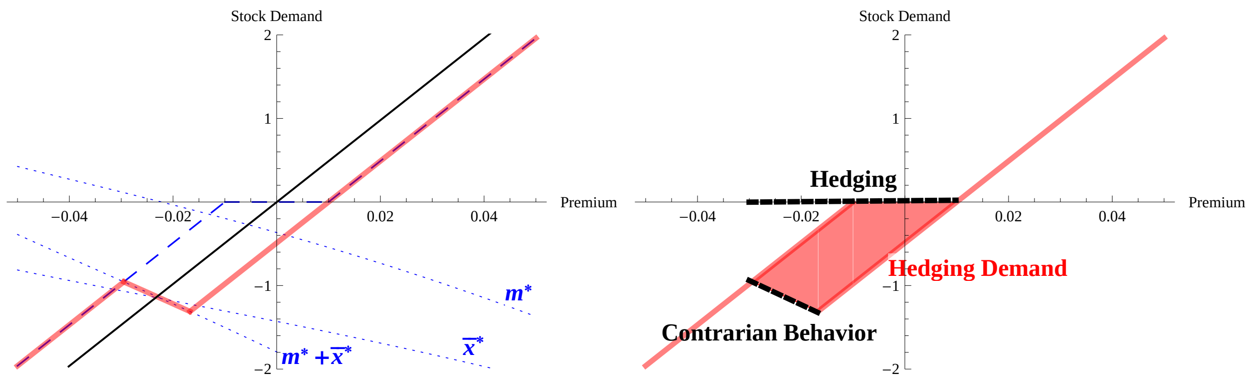

Figure 1: Optimal stock demand (fraction of wealth) as a function of the estimated instantaneous equity premium (annual, decimal). The investor has observed 20 years of data and now faces a 10-year investment horizon. $\beta = 0.03$, $\lambda = \infty$, $\alpha = 0.38$, and the estimated long-run equity premium $\bar{x}_t^* - r$ is fixed at 0.0458. *Left plot*: The *solid line passing through the origin* shows the Bayesian demand; the *dashed line*, the myopic demand; the *dotted lines*, the shadow hedging demands; finally, the *thick solid line* shows the total demand. *Right plot*: An analysis of the optimal stock demand.

Figure 1 shows the corresponding optimal stock demand as a function of the estimated instantaneous equity premium, with the estimated long-run equity premium fixed at 0.0458 (Barberis's estimate). In the left plot, the solid line passing through the origin shows the Bayesian demand; the kinked dashed line, the myopic demand; the dotted lines, the m^*-, $\bar{x}^*$-, and total shadow hedging demands; and finally, the thick solid line shows the total demand. As observed analytically, the total demand is given by the shadow hedging demand if the latter is moderate compared to the magnitude of the ambiguity present; otherwise, the investor behaves as if he were a Bayesian investor whose estimate of the equity premium is $m^* - r - \sigma_R \bar{\xi}(t)$ or $m^* - r + \sigma_R \bar{\xi}(t)$. The hedging demands are represented by a shaded region in the right plot. Note that the investor hedges for a range of estimated equity premia wider than dictated by the ambiguity in his estimate and the hedging demands are significant. For example, when the estimated equity premium is -0.01, the long-horizon investor facing a 10-year horizon sells short an amount of the stock worth about 100% of his wealth, whereas a myopic investor would take no position in the stock.

In Comparison with Epstein and Schneider (2007) To further analyze the optimal policy, it helps to contrast it with that of related models, and first I consider Epstein and Schneider (2007).

First, in Epstein and Schneider's model, a long-horizon multiple-priors investor still holds no stock when the estimated equity premium is zero. In Figure 1, on the other hand, π^* is negative around zero estimated premium. This is due to the asymmetry in the dynamics of the estimated premium $m^* - r$. When the true premium is constant and known, a log investor's value function is quadratic in it. Hence, in particular, it is symmetric at zero premium and is strictly increasing in the absolute value of the premium; that is, the investor is better off when the stock is (locally, in expected terms) more distinct from the bond. However, since in the present model $m^* - r$ is attracted to $\bar{x}^* - r$, the current value of which is positive, the value function rises in the right vicinity of zero estimated premium and rises

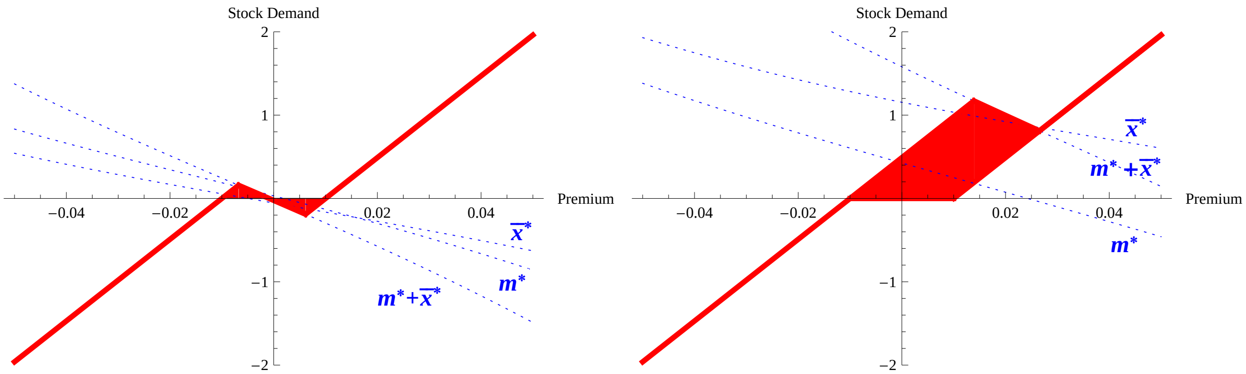

Figure 2: Optimal stock demand (fraction of wealth) as a function of the estimated instantaneous equity premium (annual, decimal). The parametrization is the same as Figure 1 except that the estimated long-run equity premium $\bar{x}_t^* - r$ is 0 for the left plot and -0.0458 for the right plot, as opposed to 0.0458.

more in the right than in the left because a negative $m^* - r$ will have to pass the minimum of the value function before reaching $\bar{x}^* - r$. Consequently, $\partial_{m^*} J(t, r, \bar{x}^*) > 0$. $\partial_{\bar{x}^*} J(t, r, \bar{x}^*) > 0$ for the obvious reason, and the negative demands around zero estimated premium follow.

When the desire to hedge fully realizes, it may give rise to contrarian behavior. Note from Figure 1 that when the estimated premium falls around -0.02, the investor exhibits contrarian behavior in the sense that he decreases his stock holdings as the estimated premium increases. In the absence of ambiguity (see, for example, Brendle (2006)), as the estimated premium improves, that is, as it moves toward the direction of increasing the continuation utility, the marginal indirect utility of such an improvement strictly increases and so does the desire to hedge. The introduction of ambiguity does not fundamentally alter this structure because the density generators are bounded by $\bar{\xi}$ and $\bar{\xi}$ is independent of the estimated premium. Epstein and Schneider make a similar observation that their investor is contrarian in the sense that when the estimated premium is not unambiguously distinct from zero, he goes long for negative premia and short for positive premia. This restricted form of contrarian behavior results from the symmetric structure of their model.

What, then, exactly is the dependence of the stock demand on the estimated *long-run* equity premium? The argument leading to $\partial_{m^*} J(t, r, \bar{x}^*) > 0$ suggests that if the current value of $\bar{x}^* - r$ is negative, $\partial_{m^*} J(t, r, \bar{x}^*) < 0$. This is indeed the case. See Figure 2; I changed the value of $\bar{x}_t^* - r$ from 0.0458 to 0 (left plot) and -0.0458 (right plot). When $\bar{x}_t^* - r = -0.0458$, both derivatives at zero instantaneous premium are negative and the corresponding hedging demand is positive. It is possible to show, following the proof of Lemma 4.2(i), that $J(t, m^*, \bar{x}^*)$ is convex in $(m^*, \bar{x}^*)$ and hence in particular in $\bar{x}^*$. Accordingly, the desire to hedge against low continuation utility results in contrarian behavior with respect to the long-run premium. Compare the monotonic dependence of the demand on the long-run premium with its nonmonotonic dependence on the instantaneous premium. Such a distinction is absent in Epstein and Schneider's model because they consider constant (in the sense of indistinguishability) investment opportunities.

The assumption of constant investment opportunities also implies that in Epstein and Schneider's model, hedging demands disappear as time goes to infinity. In contrast, in the

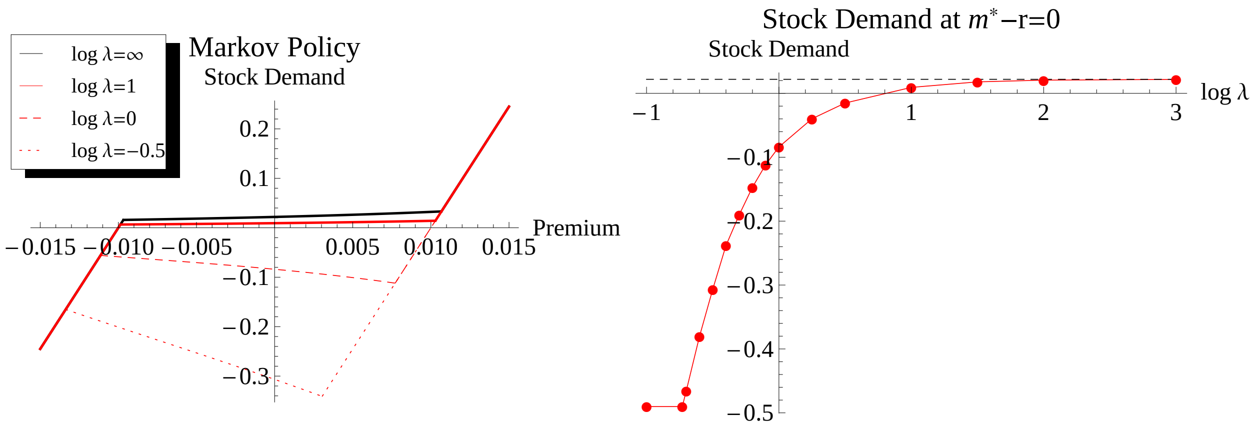

Figure 3: Confidence and optimal stock demand. The investor has learned all that he can ($t \to \infty$) and now faces a 10-year investment horizon. $\beta = 0.03$ and $\bar{x}^*_\infty - r = 0.0458$. α varies as λ does in such a way that the ambiguity in the instantaneous equity premium $\sqrt{2\alpha\delta(t=\infty;\lambda)}$ stays at 0.01. *Left plot*: Optimal stock demand (fraction of wealth) as a function of the estimated instantaneous equity premium (annual, decimal), for different levels of the investor's confidence λ in the reference likelihood. *Right plot*: The same demand at $m^* - r = 0$ as a function of λ.

present model, the desire to hedge against adverse changes in the estimate of the instantaneous premium, that is, the m^*-shadow hedging demand, persists.

In Comparison with Miao (2009) Miao (2009) also considers the consumption/portfolio choice problem of a multiple-priors investor in continuous time who partially observes stochastic investment opportunities. But his notion of learning is fundamentally different from mine.

To review Miao's model in the context of the present model, pick a theoretical prior $Q^{\bar{x},0}$, $\bar{x} \in \mathbb{R}$. A preferential prior P^ξ is characterized by the filtered stock return dynamics

$$\mathrm{d}R(t) = m^{\bar{x},0}(t)\,\mathrm{d}t + \sigma_R(\,\mathrm{d}\bar{w}^{\bar{x},0,\xi}(t) + \xi(t)\,\mathrm{d}t),\ \ |\xi(t)| \leq \bar{\xi}$$

where $\bar{w}^{\bar{x},0,\xi} = \{\bar{w}^{\bar{x},0,\xi}(t), \mathcal{G}_t\}$ is a P^ξ-Wiener process. That is, (i) the "center" of the set of one-step-ahead conditionals is obtained by the standard Bayesian learning under $Q^{\bar{x},0}$ and (ii) after the Bayesian learning, there remains an exogenous and time-invariant ambiguity. Thus, in particular, learning and ambiguity do not interact. In contrast, in the present model, the innovation receives a larger weight when the current estimate m^*_t is ambiguous, that is, when $\delta(t)$ is large.

In fact, Miao's model is the limit of the present model as $t, \lambda, \alpha \to \infty$ with the restriction $\sqrt{2\alpha\delta(t;\lambda)}/\sigma_R = \bar{\xi}$. Note that $t \to \infty$ is consistent with the IID ambiguity; $\lambda \to \infty$, that is, full confidence in the reference likelihood, with the Bayesian learning; and $\alpha \to \infty$ with the multiple one-step-ahead conditionals despite the full confidence.

In Figure 3, I plot the optimal stock demand corresponding to different levels of confidence λ. The investor has learned all that he can, meaning in particular that γ and δ have converged (assume that $\bar{x}^*$, too, has converged), and now faces a 10-year investment horizon. $\beta = 0.03$ as before and $\bar{x}^*_\infty - r = 0.0458$. α varies as λ does in such a way that the ambiguity in the instantaneous premium $\sqrt{2\alpha\delta(t=\infty;\lambda)}$ stays at 0.01. The left plot shows the Markov policies. The solid black line (top) in particular corresponds to full confidence and hence to the Miao demand. Note that it is increasing everywhere, that is, there is no region of contrarian behavior. This is because stock returns are negatively correlated with the state variable m^*: $\sigma_{m^*}(\infty) = \rho_w + \gamma(\infty)/\sigma_R = -0.003$.

More importantly, the stock demand monotonically decreases as the investor loses confidence. See the right plot, which shows the stock demand at $m^* - r = 0$ as λ varies. Intuitively, the estimation of the true premium is more difficult and unreliable for those investors who are less confident about their grasp of the environment; the consequent lack of confidence in the estimate combined with the (apparent) pessimism leads those investors, then, to try to transfer wealth even more to adverse states.

The effect of learning under ambiguity can be significant: the difference between Miao's prediction and mine can be as large as half of wealth, depending on the investor's confidence.

A Proofs

Proof of Proposition 3.1. The local version of the Itô existence-uniqueness result. See Rogers and Williams (1994), Theorem V.12.1. □

Proof of Proposition 3.2. Under Assumptions 3.1 and 3.2, the following theorems in Liptser and Shiryaev (1977) hold: (i) follows from Theorem 12.6; (ii), from Theorem 12.7; and (iii) from a multidimensional adaptation of Theorems 7.17 and 12.5. □

Proof of Lemma 3.1. Let $f(t) = e^{\kappa t}\gamma(t)e^{\kappa t}$. Then

$$\dot{f}(t) = e^{\kappa t}[\rho_w\rho_w^\top + \rho_v\rho_v^\top - (\rho_w\sigma(t)^\top + \gamma(t)b(t)^\top)(\sigma(t)\sigma(t)^\top)^{-1}(\rho_w\sigma(t)^\top + \gamma(t)b(t)^\top)^\top]e^{\kappa t}.$$

Since $(\rho_w\sigma^\top + \gamma b^\top)(\sigma\sigma^\top)^{-1}(\rho_w\sigma^\top + \gamma b^\top)^\top$ is symmetric and positive semidefinite,

$$\operatorname{tr} f(t) \le \operatorname{tr} f(0) + \int_0^t \operatorname{tr}(e^{\kappa t}(\rho_w\rho_w^\top + \rho_v\rho_v^\top)e^{\kappa t})\,\mathrm{d}t.$$

It follows that the sum of the variances is bounded: $\sup_{t\le T}\sum_i \gamma_{ii}(t) < \infty$. Since covariances are bounded by variances, the claim follows. □

Proof of Proposition 3.3.[44] Fix $(\bar{x}, \eta) \in \mathbb{R}^{n_x} \times L^2([0,T], \mathbb{R}^{n_y})$. Let $\bar{Q}^{\bar{x},\eta}$ be the measure under which j and $v^{\bar{x},\eta}$ are independent Wiener processes where j is defined by $\mathrm{d}j(t) = \sigma(t)^{-1}\,\mathrm{d}y(t)$, $j(0) = 0$. Then

$$\frac{\mathrm{d}Q^{\bar{x},\eta}}{\mathrm{d}\bar{Q}^{\bar{x},\eta}} = \Lambda(T),$$

$$\Lambda(t) \triangleq \exp\left(\int_0^t [\sigma(s)^{-1}(a(s) + b(s)x(s))]^\top\,\mathrm{d}j(s) - \frac{1}{2}\int_0^t |\sigma(s)^{-1}(a(s) + b(s)x(s))|^2\,\mathrm{d}s\right).$$

Let $\psi^{\bar{x},\eta}(t,\cdot)$ denote the unnormalized density of $x(t)$ given $\mathcal{G}_t$ under $Q^{\bar{x},\eta}$, defined by

$$\mathrm{E}^{\bar{Q}^{\bar{x},\eta}}[\Lambda(t)f(x(t))|\mathcal{G}_t] = \int_X f(x)\psi^{\bar{x},\eta}(t,x)\,\mathrm{d}x$$

[44] I thank Domenico Cuoco for this direct proof. Alternatively, we can differentiate (22) under the integral sign and re-construct the log-likelihood function back.

where $X \triangleq \mathbb{R}^{n_x}$, f denotes an arbitrary test function, and

$$\int_X \mathrm{d}x \equiv \int_{\mathbb{R}} \cdots \int_{\mathbb{R}} \mathrm{d}x_1 \cdots \mathrm{d}x_{n_x}.$$

Since (y, x) is conditionally Gaussian,

$$\psi^{\bar{x},\eta}(t,x) = \exp\left(u^{\bar{x},\eta}(t) - \frac{1}{2}(x - m^{\bar{x},\eta}(t))^\top \gamma(t)^{-1}(x - m^{\bar{x},\eta}(t))\right) \tag{59}$$

where $u^{\bar{x},\eta}(t)$ is independent of x. Now use Bayes' rule to see

$$\ell_T(\bar{x},\eta) = \log \mathrm{E}^{\bar{Q}^{\bar{x},\eta}}\left(\frac{\mathrm{d}Q^{\bar{x},\eta}}{\mathrm{d}\bar{Q}^{\bar{x},\eta}}\middle|\, \mathcal{G}_T\right) - \log \mathrm{E}^{\bar{Q}^{0,0}}\left(\frac{\mathrm{d}Q^{0,0}}{\mathrm{d}\bar{Q}^{0,0}}\middle|\, \mathcal{G}_T\right) + \log \mathrm{E}^{\bar{Q}^{0,0}}\left(\frac{\mathrm{d}\bar{Q}^{\bar{x},\eta}}{\mathrm{d}\bar{Q}^{0,0}}\middle|\, \mathcal{G}_T\right)$$

but the last term is 0 because under $\bar{Q}^{0,0}$, $v^{0,0}$ and j are independent. Thus

$$\begin{aligned}\ell_T(\bar{x},\eta) &= \log \int_X \psi^{\bar{x},\eta}(T,x)\,\mathrm{d}x - \log \int_X \psi^{0,0}(T,x)\,\mathrm{d}x \\ &= u^{\bar{x},\eta}(T) - u^{0,0}(T)\end{aligned}$$

and all boils down to computing $u^{\bar{x},\eta}(T)$.

To compute it, I compare the $\psi^{\bar{x},\eta}$ given in (59) with that as the solution to the Zakai equation:

Lemma A.1. *$\psi^{\bar{x},\eta}$ satisfies*

$$\begin{aligned}\mathrm{d}\psi^{\bar{x},\eta}(t,x) = \psi^{\bar{x},\eta}(t,x)[\sigma(t)^{-1}(a(t)+b(t)x)]^\top \mathrm{d}j(t) - \operatorname{div}[(\kappa(\bar{x}-x)+\rho_v\eta(t))\psi^{\bar{x},\eta}(t,x)]\,\mathrm{d}t \\ - \partial_x\psi^{\bar{x},\eta}(t,x)^\top \rho_w\,\mathrm{d}j(t) + \frac{1}{2}\operatorname{tr}[\partial_x^2\psi^{\bar{x},\eta}(t,x)(\rho_w\rho_w^\top + \rho_v\rho_v^\top)]\,\mathrm{d}t\end{aligned} \tag{60}$$

with initial condition $\psi^{\bar{x},\eta}(0,\cdot) \sim N(m_0,\gamma_0)$.

Proof. The derivation is standard; see, for example, Elliott and Krishnamurthy (1997). First, differentiating $\Lambda(t)f(x(t))$ and then re-integrating the resulting expression,

$$\begin{aligned}\Lambda(t)f(x(t)) = \Lambda(0)f(x(0)) &+ \int_0^t \Lambda(s)f(x(s))[\sigma(s)^{-1}(a(s)+b(s)x(s))]^\top \mathrm{d}j(s) \\ &+ \int_0^t \Lambda(s)\partial f(x(s))^\top\{[\kappa(\bar{x}-x(s))+\rho_v\eta(s)]\,\mathrm{d}s + \rho_w\,\mathrm{d}j(s) + \rho_v\,\mathrm{d}v^{\bar{x},\eta}(s)\} \\ &+ \frac{1}{2}\int_0^t \Lambda(s)\operatorname{tr}[\partial^2 f(x(s))(\rho_w\rho_w^\top + \rho_v\rho_v^\top)]\,\mathrm{d}s.\end{aligned}$$

Take the conditional expectation under $\bar{Q}^{\bar{x},\eta}$ given $\mathcal{G}_t$:

$$\begin{aligned}\int_X f(x)\psi^{\bar{x},\eta}(t,x)\,\mathrm{d}x = &\int_X f(x)\psi^{\bar{x},\eta}(0,x)\,\mathrm{d}x \\ &+ \int_0^t\int_X f(x)[\sigma(s)^{-1}(a(s)+b(s)x)]^\top\psi^{\bar{x},\eta}(s,x)\,\mathrm{d}x\,\mathrm{d}j(s) \\ &+ \int_0^t\int_X \partial f(x)^\top[\kappa(\bar{x}-x)+\rho_v\eta(s)]\psi^{\bar{x},\eta}(s,x)\,\mathrm{d}x\,\mathrm{d}s \\ &+ \int_0^t\int_X \partial f(x)^\top\rho_w\psi^{\bar{x},\eta}(s,x)\,\mathrm{d}x\,\mathrm{d}j(s) \\ &+ \frac{1}{2}\int_0^t\int_X \operatorname{tr}[\partial^2 f(x)(\rho_w\rho_w^\top + \rho_v\rho_v^\top)]\psi^{\bar{x},\eta}(s,x)\,\mathrm{d}x\,\mathrm{d}s.\end{aligned}$$

For the change in the order of the conditional expectation and the stochastic integral with respect to j, see Liptser and Shiryaev (1977), Theorem 5.14. Now, integration by parts with respect to x completes the derivation. $\square$

(Proof of the proposition continued.) From (59),

$$\begin{aligned}
\mathrm{d}\log\psi^{\bar{x},\eta}(t,x) = {} & \mathrm{d}u^{\bar{x},\eta}(t) + (x - m^{\bar{x},\eta}(t))^\top \gamma(t)^{-1}\,\mathrm{d}m^{\bar{x},\eta}(t) \\
& + \frac{1}{2}(x - m^{\bar{x},\eta}(t))^\top \gamma(t)^{-1}\dot{\gamma}(t)\gamma(t)^{-1}(x - m^{\bar{x},\eta}(t))\,\mathrm{d}t \\
& - \frac{1}{2}\operatorname{tr}[\gamma(t)^{-1}(\rho_w\sigma(t)^\top + \gamma(t)b(t)^\top)(\sigma(t)\sigma(t)^\top)^{-1}(\rho_w\sigma(t)^\top + \gamma(t)b(t)^\top)^\top]\,\mathrm{d}t
\end{aligned}$$

On the other hand, computing the spatial derivatives of $\psi^{\bar{x},\eta}$ using (59) and plugging them to (60), we obtain another expression for $\mathrm{d}\psi^{\bar{x},\eta}(t,x)$:

$$\begin{aligned}
\mathrm{d}\psi^{\bar{x},\eta}(t,x)/\psi^{\bar{x},\eta}(t,x) = {} & [\sigma(t)^{-1}(a(t) + b(t)x) + \rho_w^\top \gamma(t)^{-1}(x - m^{\bar{x},\eta}(t))]^\top\,\mathrm{d}j(t) \\
& + [(x - m^{\bar{x},\eta}(t))^\top \gamma(t)^{-1}(\kappa(\bar{x} - x) + \rho_v\eta(t)) + \operatorname{tr}\kappa]\,\mathrm{d}t \\
& + \frac{1}{2}(x - m^{\bar{x},\eta}(t))^\top \gamma(t)^{-1}(\rho_w\rho_w^\top + \rho_v\rho_v^\top)\gamma(t)^{-1}(x - m^{\bar{x},\eta}(t))\,\mathrm{d}t \\
& - \frac{1}{2}\operatorname{tr}(\gamma(t)^{-1}(\rho_w\rho_w^\top + \rho_v\rho_v^\top))\,\mathrm{d}t.
\end{aligned}$$

Then

$$\begin{aligned}
\mathrm{d}\log\psi^{\bar{x},\eta}(t,x) &= \frac{1}{\psi}\,\mathrm{d}\psi + \frac{1}{2}\left(-\frac{1}{\psi^2}\right)(\,\mathrm{d}\psi)^2 \\
&= [\sigma(t)^{-1}(a(t) + b(t)x) + \rho_w^\top \gamma(t)^{-1}(x - m^{\bar{x},\eta}(t))]^\top\,\mathrm{d}j(t) \\
&\quad + [(x - m^{\bar{x},\eta}(t))^\top \gamma(t)^{-1}(\kappa(\bar{x} - x) + \rho_v\eta(t)) + \operatorname{tr}\kappa]\,\mathrm{d}t \\
&\quad + \frac{1}{2}(x - m^{\bar{x},\eta}(t))^\top \gamma(t)^{-1}(\rho_w\rho_w^\top + \rho_v\rho_v^\top)\gamma(t)^{-1}(x - m^{\bar{x},\eta}(t))\,\mathrm{d}t \\
&\quad - \frac{1}{2}\operatorname{tr}(\gamma(t)^{-1}(\rho_w\rho_w^\top + \rho_v\rho_v^\top))\,\mathrm{d}t \\
&\quad - \frac{1}{2}|\sigma(t)^{-1}(a(t) + b(t)x) + \rho_w^\top \gamma(t)^{-1}(x - m^{\bar{x},\eta}(t))|^2\,\mathrm{d}t.
\end{aligned}$$

Equate the two expressions of $\mathrm{d}\log\psi^{\bar{x},\eta}(t,x)$ to see

$$\begin{aligned}
\mathrm{d}u^{\bar{x},\eta}(t) = {} & -\frac{1}{2}\operatorname{tr}(\gamma(t)^{-1}\dot{\gamma}(t))\,\mathrm{d}t + (a(t) + b(t)m^{\bar{x},\eta}(t))^\top(\sigma(t)\sigma(t)^\top)^{-1}\,\mathrm{d}y(t) \\
& - \frac{1}{2}(a(t) + b(t)m^{\bar{x},\eta}(t))^\top(\sigma(t)\sigma(t)^\top)^{-1}(a(t) + b(t)m^{\bar{x},\eta}(t))\,\mathrm{d}t.
\end{aligned}$$

Finally, note that $u^{\bar{x},\eta}(0) = u^{0,0}(0)$. $\square$

Proof of Lemma 3.2. Let $\varepsilon > 0$ and observe

$$\begin{aligned}
\ell_t(\bar{x}, \eta + \varepsilon h) - \ell_t(\bar{x}, \eta) = {} & \varepsilon\int_0^t \left(\int_0^s \varphi(\tau)^{-1}\rho_v h(\tau)\,\mathrm{d}\tau\right)^\top \varphi(s)^\top b(s)^\top(\sigma(s)\sigma(s)^\top)^{-1}\,\mathrm{d}y(s) \\
& - \varepsilon\int_0^t (a(s) + b(s)m^{\bar{x},\eta}(s))^\top(\sigma(s)\sigma(s)^\top)^{-1}b(s)\varphi(s)\int_0^s \varphi(\tau)^{-1}\rho_v h(\tau)\,\mathrm{d}\tau\,\mathrm{d}s + O(\varepsilon^2). \quad (61)
\end{aligned}$$

The first term can be rewritten, by integration by parts, as

$$\varepsilon\left(\int_0^t \varphi(\tau)^{-1}\rho_v h(\tau)\,\mathrm{d}\tau\right)^\top \int_0^t \varphi(s)^\top b(s)^\top(\sigma(s)\sigma(s)^\top)^{-1}\,\mathrm{d}y(s)$$
$$-\,\varepsilon\int_0^t (\varphi(s)^{-1}\rho_v h(s))^\top \int_0^s \varphi(\tau)^\top b(\tau)^\top(\sigma(\tau)\sigma(\tau)^\top)^{-1}\,\mathrm{d}y(\tau)\,\mathrm{d}s \tag{62}$$
$$=\varepsilon\int_0^t\left(\int_s^t \varphi(\tau)^\top b(\tau)^\top(\sigma(\tau)\sigma(\tau)^\top)^{-1}\,\mathrm{d}y(\tau)\right)^\top \varphi(s)^{-1}\rho_v h(s)\,\mathrm{d}s$$

and the second term, by changing the order of integration, as

$$\varepsilon\int_0^t\left(\int_s^t \varphi(\tau)^\top b(\tau)^\top(\sigma(\tau)\sigma(\tau)^\top)^{-1}(a(\tau)+b(\tau)m^{\bar{x},\eta}(\tau))\,\mathrm{d}\tau\right)^\top \varphi(s)^{-1}\rho_v h(s)\,\mathrm{d}s. \tag{63}$$

Now, plug (62) and (63) into (61), differentiate it with respect to ε, and set $\varepsilon = 0$. □

Proof of Lemma 3.3. Since $\psi_{11}(0) = I_{n_x}$, $\psi_{11}(t)$ is by continuity invertible up to a random time $\tau_{11} \in (0, T]$. Up to τ_{11}, $\psi_{21}(t)\psi_{11}(t)^{-1}\rho_v\rho_v^\top$ satisfies

$$\dot{f}(t) = \lambda^{-1}\rho_v\rho_v^\top - \bar{\kappa}(t)f(t) - f(t)\bar{\kappa}(t)^\top - f(t)b(t)^\top(\sigma(t)\sigma(t)^\top)^{-1}b(t)f(t),\ f(0) = 0. \tag{64}$$

(64) has a unique solution: suppose p and q solve (64), let $\Delta(t) = p(t) - q(t)$, and observe $\Delta(0) = \dot{\Delta}(0) = 0$. Thus, $\psi_{21}(t)\psi_{11}(t)^{-1}\rho_v\rho_v^\top$ is symmetric up to τ_{11}.

Consider the following hypothetical partially observable system

$$\mathrm{d}y(t) = b(t)x(t)\,\mathrm{d}t + \sigma(t)\,\mathrm{d}w(t),$$
$$\mathrm{d}x(t) = -\bar{\kappa}(t)x(t)\,\mathrm{d}t + \lambda^{-1/2}\rho_v\,\mathrm{d}v(t),\ x(0) \sim N(m_0, 0),$$

with the understanding $\bar{\kappa}(t) = \bar{\kappa}(t, y)$. By Liptser and Shiryaev (1977), Theorem 12.7, the conditional variance of $x(t)$ satisfies (64) and stays positive definite for $t > 0$. (The assumptions of the theorem are satisfied; in particular, $\bar{\kappa}$ is uniformly bounded by Lemma 3.1.) Hence, $\psi_{21}(t)\psi_{11}(t)^{-1}\rho_v\rho_v^\top$ is positive definite and consequently invertible up to τ_{11}.

Since $\psi_{21}(0) = 0$ and $\dot{\psi}_{21}(0) = \lambda^{-1}I_n$, $\psi_{21}(t)$, $t > 0$, too, is invertible up to a random time $\tau_{21} \in (0, T]$. By the last paragraph, $\tau_{21} \geq \tau_{11}$.

Suppose $\tau_{11} < T$. There are two cases to consider. First, $\tau_{11} = \tau_{21}$. This contradicts the invertibility of ψ. Second, $\tau_{11} < \tau_{21}$. Then $\psi_{11}(t)^{-1}$ will explode as $t \uparrow \tau_{11}$, which is impossible because $\psi_{21}(t)$ is invertible. To be concrete, let g be the solution of $\dot{g}(t) = g(t)\bar{\kappa}(t)$, $g(0) = I_n$, and let $h(t) \triangleq g(t)\psi_{21}(t)\psi_{11}(t)^{-1}\rho_v\rho_v^\top g(t)^\top$. Observe

$$\operatorname{tr} h(t) \leq \int_0^T \operatorname{tr}(g(s)\lambda^{-1}\rho_v\rho_v^\top g(s)^\top)\,\mathrm{d}s,\ t \leq \tau_{11},$$

Given that g and ψ_{21} are invertible, the left-hand side should explode as $t \uparrow \tau_{11}$ but the right-hand side is finite. Hence, $\tau_{11} = \tau_{21} = T$. (Note. T is arbitrary.) □

Proof of Proposition 3.4. Multiply ρ_v to FOC(η) to have

$$\lambda\rho_v\eta(s) = \rho_v\rho_v^\top(\varphi(s)^{-1})^\top \int_s^t \varphi(\tau)^\top b(\tau)^\top(\sigma(\tau)\sigma(\tau)^\top)^{-1}[\,\mathrm{d}y(\tau) - (a(\tau)+b(\tau)m^{\bar{x},\eta}(\tau))\,\mathrm{d}\tau].$$

Differentiate this with respect to s to see

$$\begin{aligned} \mathrm{d}(\lambda\rho_v\eta(s)) &= \rho_v\rho_v^\top\bar\kappa(s)^\top(\rho_v\rho_v^\top)^{-1}\lambda\rho_v\eta(s)\,\mathrm{d}s \\ &\quad - \rho_v\rho_v^\top b(s)^\top(\sigma(s)^\top)^{-1}\,\mathrm{d}\bar w^{0,0}(s) + \rho_v\rho_v^\top b(s)^\top(\sigma(s)\sigma(s)^\top)^{-1}b(s)\Phi^{\kappa\bar x+\eta}(s)\,\mathrm{d}s. \end{aligned}$$

Observe in turn that

$$\mathrm{d}\Phi^{\kappa\bar x+\rho_v\eta}(s) = -\bar\kappa(s)\Phi^{\kappa\bar x+\rho_v\eta}(s)\,\mathrm{d}s + (\kappa\bar x+\rho_v\eta(s))\,\mathrm{d}s$$

and that consequently we have a linear system of differential equations in $\lambda\rho_v\eta$ and $\Phi^{\kappa\bar x+\rho_v\eta}$. Written in the matrix form, the system is

$$\mathrm{d}\begin{pmatrix}\lambda\rho_v\eta(s)\\ \Phi^{\kappa\bar x+\rho_v\eta}(s)\end{pmatrix} = \chi(s)\begin{pmatrix}\lambda\rho_v\eta(s)\\ \Phi^{\kappa\bar x+\rho_v\eta}(s)\end{pmatrix}\mathrm{d}s + \begin{pmatrix}-\rho_v\rho_v^\top b(s)^\top(\sigma(s)^\top)^{-1}\,\mathrm{d}\bar w^{0,0}(s)\\ \kappa\bar x\,\mathrm{d}s\end{pmatrix}$$

It follows

$$\begin{aligned} \begin{pmatrix}\lambda\rho_v\eta(s)\\ \Phi^{\kappa\bar x+\rho_v\eta}(s)\end{pmatrix} &= \psi(s)\left(\iota_1\lambda\rho_v\eta(0) + \int_0^s\psi(\tau)^{-1}\begin{pmatrix}-\rho_v\rho_v^\top b(\tau)^\top(\sigma(\tau)^\top)^{-1}\,\mathrm{d}\bar w^{0,0}(\tau)\\ \kappa\bar x\,\mathrm{d}\tau\end{pmatrix}\right) \\ &= \psi(s)\iota_1\lambda\rho_v\eta(0) \\ &\quad -\psi(s)\int_0^s\psi(\tau)^{-1}\iota_1\rho_v\rho_v^\top b(\tau)^\top(\sigma(\tau)^\top)^{-1}\,\mathrm{d}\bar w^{0,0}(\tau) + \Psi(s)\iota_2\kappa\bar x. \end{aligned}$$

Finally, observe

$$\begin{aligned} \lambda\rho_v\eta(t) &= 0 \\ &= \psi_{11}(t)\lambda\rho_v\eta(0) - \iota_1^\top\psi(t)\int_0^t\psi(\tau)^{-1}\iota_1\rho_v\rho_v^\top b(\tau)^\top(\sigma(\tau)^\top)^{-1}\,\mathrm{d}\bar w^{0,0}(\tau) + \Psi_{12}(t)\kappa\bar x. \quad\square \end{aligned}$$

Lemma A.2. *θ is the unique solution of*

$$\dot f(t) = I_{n_x} - \bar\kappa^\lambda(t)f(t),\ \ f(0)=0 \tag{65}$$

where

$$\begin{aligned} \bar\kappa^\lambda(t) &\triangleq \bar\kappa(t) + \psi_{21}(t)\psi_{11}(t)^{-1}\rho_v\rho_v^\top b(t)^\top(\sigma(t)\sigma(t)^\top)^{-1}b(t) \\ &= \kappa + [\rho_w\sigma(t)^\top + (\gamma(t)+\psi_{21}(t)\psi_{11}(t)^{-1}\rho_v\rho_v^\top)b(t)^\top](\sigma(t)\sigma(t)^\top)^{-1}b(t). \end{aligned}$$

Proof. (65) follows from the direct differentiation of the definition. Uniqueness is standard. □

Define

$$p(s,t) \triangleq \Psi(s)\iota_2 - \psi(s)\iota_1\psi_{11}(t)^{-1}\Psi_{12}(t),\ \ s\le t\le T.$$

Then

$$\begin{pmatrix}\lambda\rho_v\eta^*_{\bar x,t}(s)\\ \Phi^{\kappa\bar x+\rho_v\eta^*_{\bar x,t}}(s)\end{pmatrix} = \begin{pmatrix}\lambda\rho_v\eta^*_{0,t}(s)\\ \Phi^{0+\rho_v\eta^*_{0,t}}(s)\end{pmatrix} + p(s,t)\kappa\bar x,\ \ s\le t\le T$$

and $\iota_2^\top p(t,t)=\theta(t)$. Also,

$$\frac{\partial}{\partial t}p(s,t) = -\psi(s)\iota_1\psi_{11}(t)^{-1}\rho_v\rho_v^\top b(t)^\top(\sigma(t)\sigma(t)^\top)^{-1}b(t)\theta(t).$$

Proof of Lemma 3.4. Let

$$M(s) \triangleq \begin{pmatrix} \lambda^{-1}(\rho_v\rho_v^\top)^{-1} & 0 \\ 0 & b(s)^\top(\sigma(s)\sigma(s)^\top)^{-1}b(s) \end{pmatrix}$$

and observe

$$\mathcal{I}_{\bar{x}}(t) = \int_0^t p(s,t)^\top M(s)p(s,t)\,\mathrm{d}s.$$

Thus

$$\begin{aligned}\frac{\mathrm{d}}{\mathrm{d}t}\mathcal{I}_{\bar{x}}(t) &= \theta(t)^\top b(t)^\top(\sigma(t)\sigma(t)^\top)^{-1}b(t)\theta(t) \\ &\quad - 2\underbrace{\int_0^t p(s,t)^\top M(s)\psi(s)\iota_1\,\mathrm{d}s\,\psi_{11}(t)^{-1}\rho_v\rho_v^\top}_{=:f(t)} b(t)^\top(\sigma(t)\sigma(t)^\top)^{-1}b(t)\theta(t),\end{aligned}$$

$$\begin{aligned}\dot{f} &= -f(\bar{\kappa}+\psi_{21}\psi_{11}^{-1}\rho_v\rho_v^\top b^\top(\sigma\sigma^\top)^{-1}b)^\top + \theta^\top b^\top(\sigma\sigma^\top)^{-1}b\Bigg(\psi_{21}\psi_{11}^{-1}\rho_v\rho_v^\top \\ &\quad - \underbrace{\int_0^t(\psi(s)\iota_1\psi_{11}(t)^{-1}\rho_v\rho_v^\top)^\top M(s)\psi(s)\iota_1\psi_{11}(t)^{-1}\rho_v\rho_v^\top\,\mathrm{d}s}_{=:g(t)}\Bigg),\end{aligned} \tag{66}$$

$$\begin{aligned}\dot{g} &= \lambda^{-1}\rho_v\rho_v^\top - \bar{\kappa}g - g\bar{\kappa}^\top - \psi_{21}\psi_{11}^{-1}\rho_v\rho_v^\top b^\top(\sigma\sigma^\top)^{-1}bg \\ &\quad + (\psi_{21}\psi_{11}^{-1}\rho_v\rho_v^\top - g)b^\top(\sigma\sigma^\top)^{-1}b\psi_{21}\psi_{11}^{-1}\rho_v\rho_v^\top,\end{aligned}$$

with $f(0) = g(0) = 0$, where I have suppressed t unless needed. $g = \psi_{21}\psi_{11}^{-1}\rho_v\rho_v^\top$ is the unique solution to the last equation. In turn, $f = 0$ is the unique solution to (66).

Suppose $\mathcal{I}_{\bar{x}}(t)$ is singular for some $t > 0$. Since it is symmetric and positive semidefinite, there must be a nonzero $z \in \mathbb{R}^{n_x}$ such that

$$\int_0^t z^\top\theta(s)^\top b(s)^\top(\sigma(s)\sigma(s)^\top)^{-1}b(s)\theta(s)z\,\mathrm{d}s = 0$$

or $\sigma(s)^{-1}b(s)\theta(s)z = 0$ for Lebesgue almost every $s \le t$ or $\theta(s)z = 0$ for all $s \le t$. Multiply z to (65) to see

$$\frac{\mathrm{d}}{\mathrm{d}s}(\theta(s)z) = 0 = z,\ \ s \le t$$

which is absurd. □

Lemma A.3.

$$\Phi^{I_{n_x}}(t)^\top - \int_0^t \Phi^{I_{n_x}}(s)^\top b(s)^\top(\sigma(s)\sigma(s)^\top)^{-1}b(s)\psi_{21}(s)\,\mathrm{d}s\,\psi_{11}(t)^{-1}\rho_v\rho_v^\top = \theta(t)^\top \tag{67}$$

and

$$\int_0^t \Phi^{I_{n_x}}(s)^\top b(s)^\top(\sigma(s)\sigma(s)^\top)^{-1}b(s)\iota_2^\top p(s,t)\,\mathrm{d}s = \mathcal{I}_{\bar{x}}(t). \tag{68}$$

Proof. (67): Denote the left-hand side by $f(t)$. Then

$$\begin{aligned}\dot{f} &= -(\Phi^{I_{n_x}})^\top \bar{\kappa}^\top + I_{n_x} - (\Phi^{I_{n_x}})^\top b^\top (\sigma\sigma^\top)^{-1} b\psi_{21}\psi_{11}^{-1}\rho_v\rho_v^\top \\ &+ \int_0^t \Phi^{I_{n_x}}(s)^\top b(s)^\top (\sigma(s)\sigma(s)^\top)^{-1} b(s)\psi_{21}(s)\,\mathrm{d}s\,\psi_{11}^{-1}\rho_v\rho_v^\top (\bar{\kappa}^\top + b^\top(\sigma\sigma^\top)^{-1} b\psi_{21}\psi_{11}^{-1}\rho_v\rho_v^\top). \quad (69)\end{aligned}$$

But by Lemma 3.3(ii),

$$\begin{aligned}\bar{\kappa}^\top + b^\top(\sigma\sigma^\top)^{-1} b\psi_{21}\psi_{11}^{-1}\rho_v\rho_v^\top &= (\bar{\kappa} + \psi_{21}\psi_{11}^{-1}\rho_v\rho_v^\top b^\top (\sigma\sigma^\top)^{-1} b)^\top \\ &= (\bar{\kappa}^\lambda)^\top\end{aligned}$$

and with this, (69) can be rewritten as $\dot{f} = I_{n_x} - f(\bar{\kappa}^\lambda)^\top$, which is also satisfied by $\theta^\top$. Since $f(0) = \theta(0)^\top = 0$, it follows that $f(t) = \theta(t)^\top$.

(68): Denote the left-hand side by $g(t)$. Then

$$\dot{g}(t) = \Phi^{I_{n_x}}(t)^\top b(t)^\top (\sigma(t)\sigma(t)^\top)^{-1} b(t)\theta(t) + \int_0^t \Phi^{I_{n_x}}(s)^\top b(s)^\top (\sigma(s)\sigma(s)^\top)^{-1} b(s) \frac{\partial}{\partial t}\iota_2^\top p(s,t)\,\mathrm{d}s$$

and that

$$\frac{\partial}{\partial t}\iota_2^\top p(s,t) = -\psi_{21}(s)\psi_{11}(t)^{-1}\rho_v\rho_v^\top b(t)^\top (\sigma(t)\sigma(t)^\top)^{-1} b(t)\theta(t).$$

By (67),

$$\dot{g}(t) = \theta(t)^\top b(t)^\top (\sigma(t)\sigma(t)^\top)^{-1} b(t)\theta(t).$$

Note finally that $g(0) = 0$. □

Proof of Proposition 3.5. From FOC($\bar{x}$),

$$\int_0^t \Phi^{I_{n_x}}(s)^\top b(s)^\top (\sigma(s)\sigma(s)^\top)^{-1} b(s)\iota_2^\top p(s,t)\,\mathrm{d}s\,\kappa\bar{x} = \int_0^t \Phi^{I_{n_x}}(s)^\top b(s)^\top (\sigma(s)^\top)^{-1}\,\mathrm{d}\bar{w}^{0,\eta^*_{0,t}}(s).$$

Recall (68). □

Proof of Proposition 3.6. (i) Differentiating FOC($\bar{x}$) with respect to t, we see

$$0 = \Phi^{I_{n_x}}(t)^\top b(t)^\top (\sigma(t)^\top)^{-1}\,\mathrm{d}\epsilon(t) - \int_0^t \Phi^{I_{n_x}}(s)^\top b(s)^\top (\sigma(s)\sigma(s)^\top)^{-1} b(s)\,\mathrm{d}_t\Phi^{\kappa\bar{x}^*_t + \rho_v\eta^*_t}(s)\,\mathrm{d}s.$$

Direct computation shows

$$\mathrm{d}_t\Phi^{\kappa\bar{x}^*_t + \rho_v\eta^*_t}(s) = \iota_2^\top p(s,t)\kappa\,\mathrm{d}\bar{x}^*_t + \psi_{21}(s)\psi_{11}(t)^{-1}\rho_v\rho_v^\top b(t)^\top (\sigma(t)^\top)^{-1}\,\mathrm{d}\epsilon(t). \quad (70)$$

Hence,

$$\begin{aligned}&\int_0^t \Phi^{I_{n_x}}(s)^\top b(s)^\top (\sigma(s)\sigma(s)^\top)^{-1} b(s)\iota_2^\top p(s,t)\,\mathrm{d}s\,\kappa\,\mathrm{d}\bar{x}^*_t \\ &\qquad = \left(\Phi^{I_{n_x}}(t)^\top - \int_0^t \Phi^{I_{n_x}}(s)^\top b(s)^\top (\sigma(s)\sigma(s)^\top)^{-1} b(s)\psi_{21}(s)\,\mathrm{d}s\,\psi_{11}(t)^{-1}\rho_v\rho_v^\top\right) \\ &\qquad\quad \times b(t)^\top (\sigma(t)^\top)^{-1}\,\mathrm{d}\epsilon(t).\end{aligned}$$

Use Lemma A.3.

(ii) Observe

$$\begin{aligned}\mathrm{d}m^{\bar{x}_t^*,\eta_t^*}(t) &= \mathrm{d}m^{\bar{x},\eta}(t)|_{\bar{x}=\bar{x}_t^*,\eta=\eta_t^*} + \Phi^{\kappa\,\mathrm{d}\bar{x}_t^*+\rho_v\,\mathrm{d}\eta_t^*}(t)\\ &= \kappa(\bar{x}_t^* - m^{\bar{x}_t^*,\eta_t^*}(t))\,\mathrm{d}t + (\rho_w\sigma(t)^\top + \gamma(t)b(t)^\top)(\sigma(t)^\top)^{-1}\,\mathrm{d}\epsilon(t) + \Phi^{\kappa\,\mathrm{d}\bar{x}_t^*+\rho_v\,\mathrm{d}\eta_t^*}(t).\end{aligned}$$

$\mathrm{d}\Phi^{\kappa\bar{x}_t^*+\rho_v\eta_t^*}(t)$ is, if computed from the definition,

$$\mathrm{d}\Phi^{\kappa\bar{x}_t^*+\rho_v\eta_t^*}(t) = -\bar{\kappa}(t)\Phi^{\kappa\bar{x}_t^*+\rho_v\eta_t^*}(t)\,\mathrm{d}t + \kappa\bar{x}_t^*\,\mathrm{d}t + \Phi^{\kappa\,\mathrm{d}\bar{x}_t^*+\rho_v\,\mathrm{d}\eta_t^*}(t)$$

and is, if computed from the solution (25) (recall (70)),

$$\begin{aligned}\mathrm{d}\Phi^{\kappa\bar{x}_t^*+\rho_v\eta_t^*}(t) = -\bar{\kappa}(t)\Phi^{\kappa\bar{x}_t^*+\rho_v\eta_t^*}(t)\,\mathrm{d}t + \kappa\bar{x}_t^*\,\mathrm{d}t&\\ + \psi_{21}(t)\psi_{11}(t)^{-1}\rho_v\rho_v^\top b(t)^\top(\sigma(t)^\top)^{-1}\,\mathrm{d}\epsilon(t) + \theta(t)\kappa\,\mathrm{d}\bar{x}_t^*.&\end{aligned}$$

Comparing the last two equations, we see

$$\Phi^{\kappa\,\mathrm{d}\bar{x}_t^*+\rho_v\,\mathrm{d}\eta_t^*}(t) = (\psi_{21}(t)\psi_{11}(t)^{-1}\rho_v\rho_v^\top + \theta(t)\sigma_{\bar{x}^*}(t)^\top)b(t)^\top(\sigma(t)^\top)^{-1}\,\mathrm{d}\epsilon(t). \quad \square$$

Proof of Proposition 3.7. All follow from direct differentiation. □

Proof of Lemma 3.5. (i) If σ and b are deterministic, then θ, $\sigma_{\bar{x}^*}$, and δ, too, are deterministic. Since the latter are continuous, boundedness follows.

(ii) Suppose σ, ρ_w, ρ_v, and b are diagonal; it then suffices to consider the scalar case. Suppose also $\bar{\kappa} \geq \varepsilon$ a.e. for some $\varepsilon > 0$.

Since $\delta - \theta\sigma_{\bar{x}^*} = \psi_{21}\psi_{11}^{-1}\rho_v^2 > 0$,

$$\dot{\theta}(t) < 1 - \varepsilon\theta(t) \text{ for all } t \geq 0.$$

Consider $\theta^\dagger$ defined by $\dot{\theta}^\dagger(t) = 1 - \varepsilon\theta^\dagger(t)$ and $\theta^\dagger(0) = \theta(0)$. $\theta(t) \leq \theta^\dagger(t)$ for all $t \geq 0$ because $\theta(t) = \theta^\dagger(t)$ implies $\dot{\theta}(t) < \dot{\theta}^\dagger(t)$. Now, $\theta^\dagger$ monotonically converges to ε^{-1}, and thus, θ is uniformly bounded by $\varepsilon^{-1} \vee \theta(0)$. Next, since $\mathcal{I}_{\bar{x}}^{-1}$ is decreasing,

$$\dot{\sigma}_{\bar{x}^*}(t) < \mathcal{I}_{\bar{x}}(0)^{-1} - \varepsilon\sigma_{\bar{x}^*}(t)$$

and $\sigma_{\bar{x}^*}$ is uniformly bounded by $(\varepsilon\mathcal{I}_{\bar{x}}(0))^{-1} \vee \sigma_{\bar{x}^*}(0)$. Note, finally, that

$$\dot{\delta}(t) < 2[(\varepsilon\mathcal{I}_{\bar{x}}(0))^{-1} \vee \sigma_{\bar{x}^*}(0)] + \lambda^{-1}\rho_v^2 - 2\varepsilon\delta(t). \quad \square$$

Remark A.1 (Sharpening of the Bound $\theta \leq \varepsilon^{-1} \vee \theta(0)$)**.** *First, $\theta \leq \varepsilon^{-1}$ because θ starts from 0 (at some suppressed time prior to time 0). Second,*

$$\begin{aligned}\theta &\leq \inf\{\varepsilon^{-1} : \bar{\kappa} \geq \varepsilon \text{ a.e.}\}\\ &= \frac{1}{\operatorname{ess\,inf}\bar{\kappa}}.\end{aligned}$$

Proof of Proposition 3.8. Since

$$\mathrm{d}\epsilon(t) = \mathrm{d}\bar{w}^{0,0}(t) - \sigma(t)^{-1}b(t)(m^{\bar{x}_t^*,\eta_t^*}(t) - m^{0,0}(t))\,\mathrm{d}t,$$

the question is whether

$$\mathcal{E}^{\sigma^{-1}b\Delta}(t) = \exp\left(\int_0^t \sigma(s)^{-1}b(s)\Delta(s)\,\mathrm{d}\bar{w}^{0,0}(t) - \frac{1}{2}\int_0^t |\sigma(s)^{-1}b(s)\Delta(s)|^2\,\mathrm{d}s\right),\ 0 \leq t \leq T$$

is a martingale under $Q^{0,0}|_{\mathcal{G}_T}$, where

$$\Delta(t) \triangleq m^{\bar{x}_t^*,\eta_t^*}(t) - m^{0,0}(t),\ 0 \le t \le T.$$

Observe

$$\begin{aligned}\mathrm{d}\Delta(t) &= \kappa(\bar{x}_t^* - \Delta(t))\,\mathrm{d}t + \delta(t)b(t)^\top(\sigma(t)^\top)^{-1}\,\mathrm{d}\bar{w}^{0,0}(t)\\ &\qquad - [\rho_w\sigma(t)^\top + (\gamma(t)+\delta(t))b(t)^\top](\sigma(t)\sigma(t)^\top)^{-1}b(t)\Delta(t)\,\mathrm{d}t.\end{aligned}$$

Hence, $(\Delta, \bar{x}^*)$ satisfies a linear SDE with uniformly bounded volatility. Thus, by a multidimensional adaptation of Liptser and Shiryaev (1977), Theorem 4.7, there is an $h > 0$ such that $\sup_{t\le T}\mathrm{E}^{Q^{0,0}}\exp(h|\Delta(t)|^2) < \infty$; in turn, by the uniform boundedness of $\sigma^{-1}b$, there is an $h' > 0$ such that $\sup_{t\le T}\mathrm{E}^{Q^{0,0}}\exp(h'|\sigma(t)^{-1}b(t)\Delta(t)|^2) < \infty$. Now, by Liptser and Shiryaev (1977), Section 6.2.3, Example 3, Novikov's condition holds and $\mathcal{E}^{\sigma^{-1}b\Delta}$ is a martingale. Define P^0 by $\mathrm{d}P^0/\,\mathrm{d}(Q^{0,0}|_{\mathcal{G}_T}) = \mathcal{E}^{\sigma^{-1}b\Delta}(T)$.

Denote by $\bar{\mathbf{F}}^\epsilon$ the augmented filtration generated by ϵ. From the definition of ϵ, we have $\bar{\mathcal{F}}_t^\epsilon \subseteq \mathcal{G}_t$, $0 \le t \le T$. For the other direction, observe the SDE that $(y, m^{\bar{x}_\cdot^*,\eta_\cdot^*}(\cdot), \bar{x}^*)$ satisfies with a, b, σ, γ, δ, and $\sigma_{\bar{x}^*}$ replaced by their respective nonanticipating path functionals in y. The drift is locally Lipschitz and linearly growing, and the volatility is linearly growing (Assumptions 3.1, 3.2, and 3.4). Hence, if in addition $(\gamma+\delta)b^\top(\sigma^\top)^{-1}$ and $\sigma_{\bar{x}^*}^\top b^\top(\sigma^\top)^{-1}$ are locally Lipschitz, then $(y, m^{\bar{x}_\cdot^*,\eta_\cdot^*}(\cdot), \bar{x}^*)$ is the unique strong solution to the SDE by Itô's existence and uniqueness theorem (Rogers and Williams (1994), Theorem V.12.1), and it would follow that $\bar{\mathcal{F}}_t^\epsilon \supseteq \mathcal{G}_t$, $0 \le t \le T$, or $\bar{\mathbf{F}}^\epsilon = \mathbf{G}$. (Recall that I assume all $\mathcal{G}_0$-measurable variables to be nonrandom constants.)

Since γ, δ, $b^\top$, $(\sigma^\top)^{-1}$, and $\sigma_{\bar{x}^*}^\top$ are uniformly bounded, it suffices to show that each of them is locally Lipschitz: suppose p and q are matrix-valued path functionals on $[0,T]\times C([0,T],\mathbb{R}^{n_y})$. Then

$$\begin{aligned}|p(t,f)q(t,f) - p(t,g)q(t,g)| &= |pq - pq' + pq' - p'q'|,\ p \equiv p(t,f) \text{ and } p' \equiv p(t,g)\\ &\le |p||q-q'| + |q'||p-p'|\end{aligned}$$

by the triangle and Cauchy-Schwarz inequalities.

γ is locally Lipschitz by the proof of Liptser and Shiryaev (1977), Theorem 12.5. b is so by assumption (Assumption 3.1). To see σ^{-1} is locally Lipschitz, observe that

$$\begin{aligned}|\sigma(t,f)^{-1} - \sigma(t,g)^{-1}| &= |\sigma(t,f)^{-1}(\sigma(t,g)-\sigma(t,f))\sigma(t,g)^{-1}|\\ &\le |\sigma(t,f)^{-1}||\sigma(t,g)^{-1}||\sigma(t,g)-\sigma(t,f)|.\end{aligned}$$

It remains to show that δ and $\sigma_{\bar{x}^*}$ are locally Lipschitz. Let $N > 0$, $t \in [0,T]$, and let $f, g \in C([0,T],\mathbb{R}^{n_y})$ be such that

$$\left(\sup_{s\le t}|f(s)|\right) \vee \left(\sup_{s\le t}|g(s)|\right) \le N.$$

Consider $\mathcal{I}_{\bar{x}}^{-1}$. Since

$$\frac{\mathrm{d}}{\mathrm{d}t}(\mathcal{I}_{\bar{x}}(t)^{-1}) = -\sigma_{\bar{x}^*}(t)^\top b(t)^\top(\sigma(t)\sigma(t)^\top)^{-1}b(t)\sigma_{\bar{x}^*}(t),$$

for $s \le t$,

$$|\mathcal{I}_{\bar{x}}(s,f)^{-1} - \mathcal{I}_{\bar{x}}(s,g)^{-1}| \le K_1\int_0^s |\sigma_{\bar{x}^*}(\tau,f) - \sigma_{\bar{x}^*}(\tau,g)|\,\mathrm{d}\tau + K_2\sup_{s\le t}|f(s)-g(s)|$$

where I use the same symbols for the path functionals and K_i are positive constants that do not depend on s or t. Proceeding similarly for $\sigma_{\bar{x}^*}$, and using the last inequality,

$$|\sigma_{\bar{x}^*}(s,f)-\sigma_{\bar{x}^*}(s,g)| \le K_3 \int_0^s |\sigma_{\bar{x}^*}(\tau,f)-\sigma_{\bar{x}^*}(\tau,g)|\,\mathrm{d}\tau \\ + K_4 \int_0^s |\delta(\tau,f)-\delta(\tau,g)|\,\mathrm{d}\tau + K_5 \sup_{s\le t}|f(s)-g(s)|.$$

In turn,

$$|\delta(s,f)-\delta(s,g)| \le K_6 \int_0^s |\delta(\tau,f)-\delta(\tau,g)|\,\mathrm{d}\tau + K_7 \sup_{s\le t}|f(s)-g(s)|.$$

By Gronwall's lemma,

$$|\delta(s,f)-\delta(s,g)| \le e^{K_6 s} K_7 \sup_{s\le t}|f(s)-g(s)|,\ \ s \le t$$

or

$$|\delta(t,f)-\delta(t,g)| \le e^{K_6 T} K_7 \sup_{s\le t}|f(s)-g(s)| =: K_8 \sup_{s\le t}|f(s)-g(s)|$$

where K_8 does not depend on t. Hence δ is locally Lipschitz. In turn, so is $\sigma_{\bar{x}^*}$. □

Proof of Lemma 3.6. Let

$$f_t(\Delta\bar{x},\Delta\eta) \triangleq \ell_t^\lambda(\bar{x}_t^*,\eta_t^*) - \ell_t^\lambda(\bar{x}_t^*+\Delta\bar{x},\eta_t^*+\Delta\eta) \ge 0 \\ = \frac{1}{2}\int_0^t \Phi^{\kappa\Delta\bar{x}+\rho_v\Delta\eta}(s)^\top b(s)^\top(\sigma(s)\sigma(s)^\top)^{-1} b(s)\Phi^{\kappa\Delta\bar{x}+\rho_v\Delta\eta}(s)\,\mathrm{d}t + \frac{\lambda}{2}\int_0^t |\Delta\eta(s)|^2\,\mathrm{d}s$$

where I have recalled the FOCs. We are to find

$$\min_{\Delta\bar{x},\Delta\eta}\left\{f_t(\Delta\bar{x},\Delta\eta) : \Phi^{\kappa\Delta\bar{x}+\rho_v\Delta\eta}(t) = \Delta m\right\}$$

where $\Delta m \equiv m - m^{\bar{x}_t^*,\eta_t^*}(t)$. Note that there always is a Δx such that $(\Delta x, \Delta\eta = 0)$ satisfies the constraint; $\Phi^\kappa(t)$ is invertible. Write the Lagrangian as

$$f_t(\Delta\bar{x},\Delta\eta) - \Lambda^\top(\Phi^{\kappa\Delta\bar{x}+\rho_v\Delta\eta}(t) - \Delta m);$$

the dependence of Λ on t is suppressed. FOC($\Delta\eta$) is

$$0 = (\varphi(s)^{-1}\rho_v)^\top \int_s^t \varphi(\tau)^\top b(\tau)^\top(\sigma(\tau)\sigma(\tau)^\top)^{-1} b(\tau)\Phi^{\kappa\Delta\bar{x}+\rho_v\Delta\eta}(\tau)\,\mathrm{d}\tau \\ + \lambda\Delta\eta(s) - (\Lambda^\top\varphi(t)\varphi(s)^{-1}\rho_v)^\top$$

or, multiplied by ρ_v and differentiated with respect to s,

$$\frac{\mathrm{d}}{\mathrm{d}s}(\lambda\rho_v\Delta\eta(s)) = \rho_v\rho_v^\top\bar{\kappa}(s)^\top(\rho_v\rho_v^\top)^{-1}\lambda\rho_v\Delta\eta(s) + \rho_v\rho_v^\top b(s)^\top(\sigma(s)\sigma(s)^\top)^{-1}b(s)\Phi^{\kappa\Delta\bar{x}+\rho_v\Delta\eta}(s).$$

Proceeding similarly to the proof of Proposition 3.4,

$$\begin{pmatrix} \lambda\rho_v\Delta\eta(s) \\ \Phi^{\kappa\Delta\bar{x}+\rho_v\Delta\eta}(s) \end{pmatrix} = \psi(s)\iota_1\lambda\rho_v\Delta\eta(0) + \Psi(s)\iota_2\kappa\Delta\bar{x}.$$

Let $s = t$ to obtain $\lambda\rho_v\Delta\eta(t) = \psi_{11}(t)\lambda\rho_v\Delta\eta(0) + \Psi_{12}(t)\kappa\Delta\bar{x}$. From FOC($\eta$),

$$0 = \lambda\Delta\eta(t) - (\Lambda^\top\rho_v)^\top.$$

Thus

$$\begin{pmatrix} \lambda\rho_v\Delta\eta(s) \\ \Phi^{\kappa\Delta\bar{x}+\rho_v\Delta\eta}(s) \end{pmatrix} = p(s,t)\kappa\Delta\bar{x} + \psi(s)\iota_1\psi_{11}(t)^{-1}\rho_v\rho_v^\top\Lambda. \tag{71}$$

Now, FOC($\Delta\bar{x}$) is

$$0 = \left(\int_0^t \Phi^{\kappa\Delta\bar{x}+\rho_v\Delta\eta}(s)^\top b(s)^\top(\sigma(s)\sigma(s)^\top)^{-1}b(s)\Phi^{I_{n_x}}(s)\kappa\,\mathrm{d}s\right)^\top - (\Lambda^\top\Phi^{I_{n_x}}(t)\kappa)^\top.$$

Substitute $\Phi^{\kappa\Delta\bar{x}+\rho_v\Delta\eta}(s)$ with that in (71) and use Lemma A.3 to see $\kappa\Delta\bar{x} = \sigma_{\bar{x}^*}(t)^\top\Lambda$. Plug this back to (71) and set $s = t$; the constraint is $\Phi^{\kappa\Delta\bar{x}+\rho_v\Delta\eta}(t) = \delta(t)\Lambda = \Delta m$ or

$$\Lambda = \delta(t)^{-1}\Delta m.$$

Thus

$$\begin{pmatrix} \lambda\rho_v\Delta\eta(s) \\ \Phi^{\kappa\Delta\bar{x}+\rho_v\Delta\eta}(s) \end{pmatrix} = (p(s,t)\sigma_{\bar{x}^*}(t)^\top + \psi(s)\iota_1\psi_{11}(t)^{-1}\rho_v\rho_v^\top)\delta(t)^{-1}\Delta m.$$

Observe

$$f_t(\Delta\bar{x},\Delta\eta) = \frac{1}{2}\int_0^t \begin{pmatrix} \lambda\rho_v\Delta\eta(s) \\ \Phi^{\kappa\Delta\bar{x}+\rho_v\Delta\eta}(s) \end{pmatrix}^\top M(s) \begin{pmatrix} \lambda\rho_v\Delta\eta(s) \\ \Phi^{\kappa\Delta\bar{x}+\rho_v\Delta\eta}(s) \end{pmatrix} \mathrm{d}s$$

where

$$M(s) = \begin{pmatrix} \lambda^{-1}(\rho_v\rho_v^\top)^{-1} & 0 \\ 0 & b(s)^\top(\sigma(s)\sigma(s)^\top)^{-1}b(s) \end{pmatrix}$$

as defined in the proof of Lemma 3.4. Proceeding similarly to that proof, we prove

$$f_t(\Delta\bar{x},\Delta\eta) = \frac{1}{2}(\Delta m)^\top\delta(t)^{-1}\Delta m. \quad \square$$

Proof of Proposition 3.9. Suppose first $\xi(t) \in \Xi(t)$. Then $b(t)m^{\bar{x}_t^*,\eta_t^*}(t) + \sigma(t)\xi(t) = b(t)m^{\bar{x},\eta}(t)$ for some theory $(\bar{x},\eta)$ and the theory passes the penalized likelihood ratio test. By Lemma 3.6,

$$\frac{1}{2}(m^{\bar{x},\eta}(t) - m^{\bar{x}_t^*,\eta_t^*}(t))^\top\delta(t)^{-1}(m^{\bar{x},\eta}(t) - m^{\bar{x}_t^*,\eta_t^*}(t)) \le \ell_t^\lambda(\bar{x}_t^*,\eta_t^*) - \ell_t^\lambda(\bar{x},\eta) \le \alpha.$$

Suppose next $\sigma(t)\xi(t) = b(t)\Delta m$, $\Delta m \in \mathbb{R}^{n_x}$, and $2^{-1}(\Delta m)^\top\delta(t)^{-1}\Delta m \le \alpha$. Let $\Delta\bar{x} \triangleq \Phi^\kappa(t)^{-1}\Delta m$. Then

$$b(t)m^{\bar{x}_t^*,\eta_t^*}(t) + \sigma(t)\xi(t) = b(t)m^{\bar{x}_t^*+\Delta\bar{x},\eta_t^*}(t).$$

There is a theory $(\bar{x},\eta)$ such that it passes the penalized likelihood ratio test and

$$b(t)m^{\bar{x},\eta}(t) = b(t)m^{\bar{x}_t^*+\Delta\bar{x},\eta_t^*}(t),$$

because

$$\begin{aligned}
&\ell_t^\lambda(\bar{x}_t^*,\eta_t^*) - \max_{\bar{x},\eta}\{\ell_t^\lambda(\bar{x},\eta) : b(t)m^{\bar{x},\eta}(t) = b(t)m^{\bar{x}_t^*+\Delta\bar{x},\eta_t^*}(t)\} \\
&\qquad = \ell_t^\lambda(\bar{x}_t^*,\eta_t^*) - \max_{\bar{x},\eta}\{\ell_t^\lambda(\bar{x},\eta) : m^{\bar{x},\eta}(t) = m^{\bar{x}_t^*+\Delta\bar{x},\eta_t^*}(t)\} \\
&\qquad = \frac{1}{2}(\Delta m)^\top\delta(t)^{-1}\Delta m \le \alpha
\end{aligned}$$

where the second equality follows from Lemma 3.6.

Hence (30).

Since δ is uniformly bounded, so are its eigenvalues; hence, the eigenvalues of δ^{-1} are uniformly bounded below away from 0. It follows that the right-hand side of (30) is uniformly bounded; so is Ξ by Assumption 3.2. Compact-convexity is clear. Finally, progressive measurability is proved as that of a single-valued, left- or right-continuous adapted process is proved: Suppose b and σ^{-1} are both right-continuous. Let $\{s_i^\nu : i\}$ denote the νth dyadic partition of $[0,t]$, $t \leq T$, and define δ_ν^{-1} by $\delta_\nu^{-1}(s) \triangleq \delta(s_{i+1}^\nu)^{-1}$ for $s_i^\nu < s \leq s_{i+1}^\nu$ and $\delta_\nu^{-1}(0) \triangleq \delta(0)^{-1}$; define b_ν and σ_ν^{-1} in the same way. Let F be a closed subset of $\mathbb{R}^{n_y}$ and observe that the weak inverse (Aliprantis and Border (1999), Section 16.1) is

$$\left\{(s,\omega) \in [0,t] \times \Omega : \sigma_\nu^{-1}(s,\omega)b(s,\omega)\left\{\Delta m \in \mathbb{R}^{n_x} : \frac{1}{2}(\Delta m)^\top \delta_\nu^{-1}(s,\omega)\Delta m \leq \alpha\right\} \cap F \neq \emptyset\right\}$$

which is trivially $\mathcal{B}[0,t] \otimes \mathcal{G}_t$-measurable. Now, note that δ^{-1} is differentiable and hence a fortiori continuous. We have $\delta_\infty^{-1} = \delta^{-1}$ as well as $b_\infty = b$ and $\sigma_\infty^{-1} = \sigma^{-1}$. Finally,

$$\begin{aligned}
&\left\{(s,\omega) : \sigma(s,\omega)^{-1}b(s,\omega)\left\{\Delta m : \frac{1}{2}(\Delta m)^\top \delta(s,\omega)^{-1}\Delta m \leq \alpha\right\} \cap F \neq \emptyset\right\} \\
&= \left\{(s,\omega) : \cap_{\mu=1}^\infty \cup_{\nu=\mu}^\infty \sigma_\nu^{-1}(s,\omega)b_\nu(s,\omega)\left\{\Delta m : \frac{1}{2}(\Delta m)^\top \delta_\nu^{-1}(s,\omega)\Delta m \leq \alpha\right\} \cap F \neq \emptyset\right\} \\
&= \cap_{\mu=1}^\infty \cup_{\nu=\mu}^\infty \left\{(s,\omega) : \sigma_\nu^{-1}(s,\omega)b_\nu(s,\omega)\left\{\Delta m : \frac{1}{2}(\Delta m)^\top \delta_\nu^{-1}(s,\omega)\Delta m \leq \alpha\right\} \cap F \neq \emptyset\right\} \\
&\in \mathcal{B}[0,t] \otimes \mathcal{G}_t. \quad \square
\end{aligned}$$

Proof of Proposition 3.10. The implied confidence set for $\bar{x}$ is the projection onto $\mathbb{R} \ni \bar{x}$ of the set of theories that pass the penalized likelihood ratio test. It equals the likelihood ratio set based on the profile penalized likelihood, which is

$$\left\{\bar{x} \in \mathbb{R} : \ell_t^\lambda(\bar{x}_t^*, \eta_{\bar{x}_t^*,t}^*) - \ell_t^\lambda(\bar{x}, \eta_{\bar{x},t}^*) \leq \alpha\right\} = \bar{x}_t^* + \left\{\Delta\bar{x} \in \mathbb{R} : \frac{1}{2}\mathcal{I}_{\bar{x}}(t)(\kappa\Delta\bar{x})^2 \leq \alpha\right\}.$$

Now, let $\varepsilon > 0$ be a lower bound of $b^\top(\sigma\sigma^\top)^{-1}b$. Observe from the dynamics of θ (28) and the boundedness of the statistics γ, θ, $\sigma_{\bar{x}^*}$, and δ (Lemma 3.1 and Assumption 3.4) that θ is bounded below as well: $\theta(t) \geq \underline{\theta}$, $t \geq 0$. (Keep in mind the convention that learning began prior to the decision making at time 0.) It follows

$$\begin{aligned}
\mathcal{I}_{\bar{x}}(t) &= \mathcal{I}_{\bar{x}}(0) + \int_0^t b(s)^\top(\sigma(s)\sigma(s)^\top)^{-1}b(s)\theta(s)^2\,\mathrm{d}s \\
&\geq \mathcal{I}_{\bar{x}}(0) + \varepsilon\underline{\theta}^2 t \to \infty \text{ as } t \to \infty. \quad \square
\end{aligned}$$

Proof of Proposition 3.11. The claim follows from Propositions 3.7 and 3.9. $\square$

Proof of Lemma 4.1. Let $\mathcal{M} \triangleq \{\mathcal{E}^\xi : \xi \in \Xi\}$. Define $f : \mathcal{M} \times (\mathcal{C}^2(u) \cap \mathcal{C}_{\text{budget}})$ by

$$f(M,c) \triangleq \mathrm{E}^{P^0}\int_0^T e^{-\beta t}M(t)\log(c(t))\,\mathrm{d}t.$$

The claim is

$$\sup_{c \in \mathcal{C}^2(u) \cap \mathcal{C}_{\text{budget}}} \min_{M \in \mathcal{M}} f(M,c) = \min_{M \in \mathcal{M}} \sup_{c \in \mathcal{C}^2(u) \cap \mathcal{C}_{\text{budget}}} f(M,c).$$

I apply the Kneser-Fan minimax theorem (Fan (1953), Theorem 2). The conclusion follows once the following three assumptions are checked.

(i) *$\mathcal{M}$ is a compact Hausdorff space.* Let $L^2([0,T]\times\Omega) \equiv L^2([0,T]\times\Omega, \mathcal{B}[0,T]\otimes\mathcal{G}_T, \text{Lebesgue}\times P^0)$ be the set of processes h such that

$$\|h\| \triangleq \left(\mathrm{E}^{P^0}\int_0^T h(t)^2\,\mathrm{d}t\right)^{1/2} < \infty.$$

$L^2([0,T]\times\Omega)$ is a reflexive Banach space with the norm $\|\cdot\|$ defined above. By design, $\mathcal{M}\subset L^2([0,T]\times\Omega)$. Let $K\geq 0$ be such that $\Xi(t)\in[-K,K]^{n_y}$, $t\geq 0$. (K may be state-dependent. See Section 4.3 and Remark 3.3.) For all $M\in\mathcal{M}$,

$$\|M\|^2 \leq \mathrm{E}^{P^0}\int_0^T \mathcal{E}^{(2\xi)}(t)e^{n_yK^2T}\,\mathrm{d}t = Te^{n_yK^2T}$$

and $\mathcal{M}$ is norm-bounded. $\mathcal{M}$ is norm-closed by Lemma B.1 of Cuoco and Cvitanić (1998) and is convex by (the proof of) Theorem 2.1(c) of Chen and Epstein (2002); thus, it is weakly closed. By Alaoglu's theorem, then, $\mathcal{M}$ is weakly compact. The weak topology of a normed space is Hausdorff and so is a subspace.

(ii) *For every $c\in\mathcal{C}^2(u)\cap\mathcal{C}_{\text{budget}}$, $f(M,c)$ is lower semicontinuous on $\mathcal{M}$.* Let span($\mathcal{M}$) be the linear span of $\mathcal{M}$ over $\mathbb{R}$; span($\mathcal{M}$) $\subset L^2([0,T]\times\Omega)$ is a normed space. For each $c\in\mathcal{C}^2(u)\cap\mathcal{C}_{\text{budget}}$, the map $\widetilde{f}^c:\text{span}(\mathcal{M})\to\mathbb{R}$,

$$M\mapsto \mathrm{E}^{P^0}\int_0^T e^{-\beta t}M(t)\log(c(t))\,\mathrm{d}t$$

is linear; by Hölder's inequality, the norm of $\widetilde{f}^c$ is bounded by $\|\log c\|<\infty$. Then there exists an extension f^c of $\widetilde{f}^c$ such that the linear functional f^c defined on $L^2([0,T]\times\Omega)$ is continuous in the norm topology, and consequently, in the weak topology (Aliprantis and Border (1999), Lemma 6.13). Being a restriction of f^c to $\mathcal{M}\subset\text{span}(\mathcal{M})$, $f(\cdot,c)$ is continuous as well.

(iii) *f is convexlike on $\mathcal{M}$ and concavelike on $\mathcal{C}^2(u)\cap\mathcal{C}_{\text{budget}}$.* $\mathcal{M}$ and $\mathcal{C}^2(u)\cap\mathcal{C}_{\text{budget}}$ are both convex. It then suffices to note that $(M,c)\mapsto M\log c$ is convex-concave on $(0,\infty)^2$. □

Proof of Proposition 4.1. Apply the minimax theorem and write the dual of the inner supremization as

$$\inf_{\nu}\mathrm{E}^{P^0}\int_0^T \max_{c(t)}\left[\mathcal{E}^{\xi}(t)e^{-\beta t}\log(c(t)) - \Lambda\mathcal{E}^{-(\zeta+\nu)}(t)e^{-rt}c(t)\right]\,\mathrm{d}t \tag{72}$$

where $\Lambda>0$. The solution to the dual problem solves the primal problem as well (He and Pearson, 1991; Karatzas et al., 1991). $c^*(t)$ and Λ^* are standard.

Plugging c^* to (72), ignoring irrelevant terms, and exchanging the order of integration, we reach

$$\mathrm{E}^{P^{\xi}}\int_0^T \frac{e^{-\beta t}-e^{-\beta T}}{\beta}\frac{1}{2}\inf_{\nu(t)}|\zeta(t)+\nu(t)+\xi(t)|^2\,\mathrm{d}t.$$

Without ξ, the minimizing $\nu(t)$ is 0 because $\nu(t)\in\text{Ker}(\sigma_R(t))$. With ξ, on the other hand,

$$|\zeta(t)+\nu(t)+\xi(t)|^2 = |\zeta(t)+\xi(t)|^2+|\nu(t)|^2+2\xi(t)^{\top}\nu(t)$$

and under the constraint $\sigma_R(t)\nu(t)=0$, the unique minimizer is given by $\nu^*(t)=f(t)\xi(t)$ where

$$f(t)\triangleq\sigma_R(t)^{\top}(\sigma_R(t)\sigma_R(t)^{\top})^{-1}\sigma_R(t)-I_{n_y}.$$

Observe that $f = f^\top$ and $f^2 = -f$, and plug c^*, ν^*, and Λ^* to (45). □

Proof of Proposition 4.2. (i) follows from Theorem IV.4.3 of Fleming and Soner (1993). The assumptions of the theorem are (IV.3.5) and (IV.4.6) in their book. (IV.3.5) is the uniform parabolicity assumption, which is equivalent in the present case to Assumption 4.1. (IV.4.6) is a collection of regularity conditions that can be checked straightforwardly. (ii) and (iii) follow from Theorem IV.3.1 of the same book. □

Proof of Lemma 4.2. (i) Let

$$F(t,m^*,\xi) \triangleq \mathrm{E}^{P^0}\left[\int_t^T \frac{e^{-\beta s}-e^{-\beta T}}{\beta}\frac{1}{2}\left(\frac{m_s^{*,\xi}-r+\sigma_R\xi(s)}{\sigma_R}\right)^2 \mathrm{d}s\,\middle|\, m_t^{*,\xi}=m^*\right]$$

so that $J(t,m^*) = \min_\xi F(t,m^*,\xi)$. The convexity of $m^* \mapsto J(t,m^*)$ follows from that of $(m^*,\xi) \mapsto F(t,m^*,\xi)$ and of Ξ: Suppose $m^* = hm_1^* + (1-h)m_2^*$, $h \in [0,1]$, and let ξ_1^* and ξ_2^* be the respective minimizers. Then

$$\begin{aligned} J(t,m^*) &\le F(t, hm_1^* + (1-h)m_2^*, h\xi_1^* + (1-h)\xi_2^*) \\ &\le hJ(t,m_1^*) + (1-h)J(t,m_2^*). \end{aligned}$$

(ii) $\partial_{m^*}J(t,m^*)$ is obtained via the envelope theorem: If $\partial_{m^*}J(t,m^*)$ and $\partial_{m^*}F(t,m^*,\xi^*)$ exist, then $\partial_{m^*}J(t,m^*) = \partial_{m^*}F(t,m^*,\xi^*)$. (See Milgrom and Segal (2002), Theorem 1.) Observe

$$m_s^{*,\xi} = e^{-\kappa(s-t)}\left(m_t^{*,\xi} + \int_t^s e^{\kappa(\tau-t)}[\kappa\bar{x}\,\mathrm{d}\tau + \sigma_{m^*}(\tau)(\,\mathrm{d}\epsilon(\tau) + \xi(\tau)\,\mathrm{d}\tau)]\right)$$

and let

$$f(s,t,m^*,\xi) \triangleq \frac{e^{-\beta s}-e^{-\beta T}}{\beta}\frac{1}{2}\left(\frac{m_s^{*,\xi}-r+\sigma_R\xi(s)}{\sigma_R}\right)^2,\ m_t^{*,\xi} = m^*$$

so that

$$F(t,m^*,\xi) = \mathrm{E}^{P^0}\int_t^T f(s,t,m^*,\xi)\,\mathrm{d}s.$$

Now, it is straightforward to check the conditions for differentiating under the integral (Durrett (2005), Theorem A.9.1) and we have $\partial_{m^*}F(t,m^*,\xi) = \mathrm{E}^{P^0}\int_t^T \partial_{m^*}f(s,t,m^*,\xi)\,\mathrm{d}s$. □

Proof of Proposition 4.3. It suffices to show

$$\lim_{t\to T}\frac{\beta e^{\beta t}}{1-e^{-\beta(T-t)}}\sigma_R\sigma_{m^*}(t)\partial_{m^*}J(t,m^*) = 0.$$

Recall Lemma 4.2(ii) and let

$$K(t,m^*) \triangleq \sup_{s\in[t,T]}\mathrm{E}^{P^0}\left|\frac{e^{-\kappa(s-t)}}{\sigma_R}\frac{m_s^{*,\xi^*}-r+\sigma_R\xi^*(s)}{\sigma_R}\right|,\ m_t^{*,\xi^*} = m^*.$$

Then

$$\begin{aligned} \left|\frac{\beta e^{\beta t}}{1-e^{-\beta(T-t)}}\partial_{m^*}J(t,m^*)\right| &\le \frac{\beta e^{\beta t}}{1-e^{-\beta(T-t)}}\int_t^T \frac{e^{-\beta s}-e^{-\beta T}}{\beta}\,\mathrm{d}s\,K(t,m^*) \\ &= \left(\frac{1}{\beta} - \frac{T-t}{e^{\beta(T-t)}-1}\right)K(t,m^*). \end{aligned}$$

$\lim_{t\to T} K(t,m^*) < \infty$ because (i)

$$m_s^{*,\xi^*} = m_s^{*,0} + e^{-\kappa(s-t)} \int_t^s e^{\kappa(\tau-t)} \sigma_{m^*}(\tau)\xi^*(\tau)\,\mathrm{d}\tau \text{ where } m_t^{*,0} = m_t^*,$$

(ii)

$$K(t,m^*) \leq \frac{1}{\sigma_R^2}\left(\sup_{s\in[t,T]} \mathrm{E}^{P^0}|m_s^{*,0}| + \int_t^T e^{\kappa(\tau-t)}|\sigma_{m^*}(\tau)|\bar{\xi}(\tau)\,\mathrm{d}\tau + r + \sigma_R\bar{\xi}(t)\right),$$

and (iii) $\mathrm{E}^{P^0}|m_s^{*,0}| = g(s-t,m^*)$ for some function g continuous in $s-t$. Thus,

$$\lim_{t\to T}\left|\frac{\beta e^{\beta t}}{1-e^{-\beta(T-t)}}\partial_{m^*}J(t,m^*)\right| \leq 0\cdot\lim_{t\to T}K(t,m^*) = 0. \quad \square$$

References

Akaike, Hirotugu (1973), "Information theory and an extension of the maximum likelihood principle." *Proc. 2nd Inter. Symposium on Information Theory, Budapest.*, 267–281.

Aliprantis, Charalambos D. and Kim C. Border (1999), *Infinite Dimensional Analysis: A Hitchhiker's Guide*. Springer.

Anderson, Evan W., Lars Peter Hansen, and Thomas J. Sargent (2003), "A quartet of semigroups for model specification, robustness, prices of risk, and model detection." *Journal of the European Economic Association*, 1, 68–123.

Anscombe, F. J. and R. J. Aumann (1963), "A definition of subjective probability." *Annals of Mathematical Statistics*, 34, 199–205.

Barberis, Nicholas (2000), "Investing for the long run when returns are predictable." *Journal of Finance*, 55, 225–264.

Barry, Christopher B. (1974), "Portfolio analysis under uncertain means, variances, and covariances." *Journal of Finance*, 29, 515–522.

Brendle, Simon (2006), "Portfolio selection under incomplete information." *Stochastic Processes and their Applications*, 116, 701–723.

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

Campbell, John Y. and Robert J. Shiller (1988), "The dividend-price ratio and expectations of future dividends and discount factors." *Review of Financial Studies*, 1, 195–228.

Campbell, John Y. and Samuel B. Thompson (2008), "Predicting excess stock returns out of sample: Can anything beat the historical average?" *Review of Financial Studies*, 21, 1509–1531.

Chen, Hui, Nengjiu Ju, and Jianjun Miao (2014), "Dynamic asset allocation with ambiguous return predictability." *Review of Economic Dynamics*, 17, 799–823.

Chen, Zengjing and Larry G. Epstein (2002), "Ambiguity, risk, and asset returns in continuous time." *Econometrica*, 70, 1403–1443.

Choi, Hongseok (2012), *Essays on Learning under Ambiguity.* Ph.D. dissertation, University of Pennsylvania.

Cochrane, John H. (2008), "The dog that did not bark: A defense of return predictability." *Review of Financial Studies*, 21, 1533–1575.

Cuoco, Domenico (1997), "Optimal consumption and equilibrium prices with portfolio constraints and stochastic income." *Journal of Economic Theory*, 72, 33–73.

Cuoco, Domenico and Jakša Cvitanić (1998), "Optimal consumption choices for a 'large' investor." *Journal of Economic Dynamics & Control*, 22, 401–436.

Davis, M. H. A. and A. R. Norman (1990), "Portfolio selection with transaction costs." *Mathematics of Operations Research*, 15, 676–713.

Detemple, Jérôme B. (1986), "Asset pricing in a production economy with incomplete information." *Journal of Finance*, 61, 383–392.

Dow, James and Sergio Ribeiro da Costa Werlang (1992), "Uncertainty aversion, risk aversion, and the optimal choice of portfolio." *Econometrica*, 60, 197–204.

Drechsler, Itamar (2013), "Uncertainty, time-varying fear, and asset prices." *Journal of Finance*, 68, 1843–1889.

Duffie, Darrell and Larry G. Epstein (1992), "Stochastic differential utility." *Econometrica*, 60, 353–394.

Durrett, Richard (2005), *Probability: Theory and Examples.* Thomson.

Elliott, Robert J. and Vikram Krishnamurthy (1997), "Exact finite-dimensional filters for maximum likelihood parameter estimation of continuous-time linear Gaussian systems." *SIAM Journal on Control and Optimization*, 35, 1908–1923.

Ellsberg, Daniel (1961), "Risk, ambiguity, and the Savage axioms." *Quarterly Journal of Economics*, 75, 643–669.

Epstein, Larry G. and Shaolin Ji (2014), "Ambiguous volatility, possibility and utility in continuous time." *Journal of Mathematical Economics*, 50, 269–282.

Epstein, Larry G. and Jianjun Miao (2003), "A two-person dynamic equilibrium under ambiguity." *Journal of Economic Dynamics & Control*, 27, 1253–1288.

Epstein, Larry G. and Martin Schneider (2003), "Recursive multiple-priors." *Journal of Economic Theory*, 113, 1–31.

Epstein, Larry G. and Martin Schneider (2007), "Learning under ambiguity." *Review of Economic Studies*, 74, 1275–1303.

Epstein, Larry G. and Martin Schneider (2008), "Ambiguity, information quality, and asset pricing." *Journal of Finance*, 63, 197–228.

Epstein, Larry G. and Kyoungwon Seo (2010), "Symmetry of evidence without evidence of symmetry." *Theoretical Economics*, 5, 313–368.

Epstein, Larry G. and Tan Wang (1994), "Intertemporal asset pricing under Knightian uncertainty." *Econometrica*, 62, 283–322.

Fama, Eugene F. and Kenneth R. French (1988a), "Dividend yields and expected stock returns." *Journal of Financial Economics*, 22, 3–25.

Fama, Eugene F. and Kenneth R. French (1988b), "Permanent and temporary components of stock prices." *Journal of Political Economy*, 96, 246–273.

Fan, Ky (1953), "Minimax theorems." *Proceedings of the National Academyof Sciences*, 39, 42–47.

Fleming, Wendell H. and H. Mete Soner (1993), *Controlled Markov Processes and Viscosity Solutions*. Springer-Verlag.

Gagliardini, Patrick, Paolo Porchia, and Fabio Trojani (2009), "Ambiguity aversion and the term structure of interest rates." *Review of Financial Studies*, 22, 4157–4188.

Gilboa, Itzhak and Massimo Marinacci (2013), "Ambiguity and the Bayesian paradigm." In *Advances in Economics and Econometrics* (Daron Acemoglu, Manuel Arellano, and Eddie Dekel, eds.), volume I, 179–242, Cambridge University Press.

Gilboa, Itzhak, Andrew Postlewaite, and David Schmeidler (2012), "Rationality of belief or: Why Savage's axioms are neither necessary nor sufficient for rationality." *Synthese*, 187, 11–31.

Gilboa, Itzhak and Larry Samuelson (2012), "Subjectivity in inductive inference." *Theoretical Economics*, 7, 183–215.

Gilboa, Itzhak and David Schmeidler (1989), "Maxmin expected utility with non-unique prior." *Journal of Mathematical Economics*, 18, 141–153.

Gilboa, Itzhak and David Schmeidler (2010), "Simplicity and likelihood: An axiomatic approach." *Journal of Economic Theory*, 145, 1757–1775.

Goetzmann, William N. and Philippe Jorion (1993), "Testing the predictive power of dividend yields." *Journal of Finance*, 48, 663–679.

Good, Irving J. and Ray A. Gaskins (1971), "Nonparametric roughness penalties for probability densities." *Biometrika*, 58, 255–277.

Green, Peter J. (1987), "Penalized likelihood for general semi-parametric regression models." *International Statistical Review*, 55, 245–259.

Hansen, Lars Peter and Thomas J. Sargent (2001), "Acknowledging misspecification in macroeconomic theory." *Review of Economic Dynamics*, 4, 519–535.

Hansen, Lars Peter and Thomas J. Sargent (2008), *Robustness*. Princeton University Press.

Hansen, Lars Peter and Thomas J. Sargent (2011), "Robustness and ambiguity in continuous time." *Journal of Economic Theory*, 146, 1195–1223.

He, Hua and Neil D. Pearson (1991), "Consumption and portfolio policies with incomplete markets and short-sale constraints: The infinite dimensional case." *Journal of Economic Theory*, 54, 259–304.

Hernández-Hernández, Daniel and Alexander Schied (2006), "Robust utility maximization in a stochastic factor model." *Statistics & Decisions*, 24, 109–125.

Hernández-Hernández, Daniel and Alexander Schied (2007a), "A control approach to robust utility maximization with logarithmic utility and time-consistent penalties." *Stochastic Processes and their Applications*, 117, 980–1000.

Hernández-Hernández, Daniel and Alexander Schied (2007b), "Robust maximization of consumption with logarithmic utility." *Proceedings of the 2007 American Control Conference*, 1120–1123.

Heston, Steven L. (1993), "A closed-form solution for options with stochastic volatility with applications to bond and currency options." *Review of Financial Studies*, 6, 327–343.

Illeditsch, Philipp K. (2011), "Ambiguous information, portfolio inertia, and excess volatility." *Journal of Finance*, 66, 2213–2247.

Kalymon, Basil A. (1971), "Estimation risk and the portfolio selection model." *Journal of Financial and Quantitative Analysis*, 6, 559–582.

Karatzas, Ioannis, John P. Lehoczky, Steven E. Shreve, and Gan-Lin Xu (1991), "Martingale and duality methods for utility maximization in an incomplete market." *SIAM Journal on Control and Optimization*, 29, 702–730.

Karatzas, Ioannis and Steven E. Shreve (1988), *Brownian Motion and Stochastic Calculus*. Springer-Verlag.

Kim, Myung J., Charles R. Nelson, and Richard Startz (1991), "Mean reversion in stock prices? a reappraisal of the empirical evidence." *Review of Economic Studies*, 58, 515–528.

Kim, Tong Suk and Edward Omberg (1996), "Dynamic nonmyopic portfolio behavior." *Review of Financial Studies*, 9, 141–161.

Klein, Roger W. and Vijay S. Bawa (1976), "The effect of estimation risk on optimal portfolio choice." *Journal of Financial Economics*, 3, 215–231.

Klein, Roger W. and Vijay S. Bawa (1977), "The effect of limited information and estimation risk on optimal portfolio diversification." *Journal of Financial Economics*, 5, 89–111.

Konishi, Sadanori and Genshiro Kitagawa (2008), *Information Criteria and Statistical Modeling*. Springer.

Lakner, Peter (1998), "Optimal trading strategy for an investor: The case of partial information." *Stochastic Processes and their Applications*, 76, 77–97.

Lettau, Martin and Sydney Ludvigson (2001), "Consumption, aggregate wealth, and expected stock returns." *Journal of Finance*, 56, 815–849.

Lettau, Martin and Stijn Van Nieuwerburgh (2008), "Reconciling the return predictability evidence." *Review of Financial Studies*, 21, 1607–1652.

Liptser, Robert S. and Albert N. Shiryaev (1977), *Statistics of Random Processes (Volumes I and II)*. Springer-Verlag.

Liu, Hening (2011), "Dynamic portfolio choice under ambiguity and regime switching mean returns." *Journal of Economic Dynamics & Control*, 35, 623–640.

Liu, Hening (2013), "Optimal consumption and portfolio choice under ambiguity for a mean-reverting risk premium in complete markets." *Annals of Economics and Finance*, 14, 21–52.

Liu, Jun (2007), "Portfolio selection in stochastic environments." *Review of Financial Studies*, 20, 1–39.

McQueen, Grant (1992), "Long-horizon mean-reverting stock prices revisited." *Journal of Financial and Quantitative Analysis*, 27, 1–18.

Merton, Robert C. (1969), "Lifetime portfolio selection under uncertainty: The continuous-time case." *Review of Economics and Statistics*, 51, 247–257.

Merton, Robert C. (1971), "Optimum consumption and portfolio rules in a continuous time model." *Journal of Economic Theory*, 3, 373–413.

Merton, Robert C. (1973), "An intertemporal capital asset pricing model." *Econometrica*, 41, 867–887.

Merton, Robert C. (1980), "On estimating the expected return on the market: An exploratory investigation." *Journal of Financial Economics*, 8, 323–361.

Miao, Jianjun (2009), "Ambiguity, risk and portfolio choice under incomplete information." *Annals of Economics and Finance*, 10, 257–279.

Miao, Jianjun and Neng Wang (2011), "Risk, uncertainty, and option exercise." *Journal of Economic Dynamics & Control*, 35, 442–461.

Milgrom, Paul and Ilya Segal (2002), "Envelope theorems for arbitrary choice sets." *Econometrica*, 70, 583–601.

Nelson, Charles R. and Myung J. Kim (1993), "Predictable stock returns: The role of small sample bias." *Journal of Finance*, 48, 641–661.

Neyman, Jerzy and Elizabeth L. Scott (1948), "Consistent estimates based on partially consistent observations." *Econometrica*, 16, 1–32.

Porchia, Paolo (2005), *Continuous-Time Asset Pricing with Ambiguity Aversion*. Ph.D. dissertation, University of Lugano.

Poterba, James M. and Lawrence H. Summers (1988), "Mean reversion in stock prices: Evidence and implications." *Journal of Financial Economics*, 22, 27–59.

Rogers, L. C. G. and David Williams (1994), *Diffusions, Markov Processes, and Martingales, Volumes 1 and 2*. Cambridge University Press.

Routledge, Bryan R. and Stanley E. Zin (2009), "Model uncertainty and liquidity." *Review of Economic Dynamics*, 12, 543–566.

Samuelson, Paul A. (1969), "Lifetime portfolio selection by dynamic stochastic programming." *Review of Economics and Statistics*, 51, 239–246.

Sbuelz, Alessandro and Fabio Trojani (2008), "Asset prices with locally constrained-entropy recursive multiple-priors utility." *Journal of Economic Dynamics & Control*, 32, 3695–3717.

Schied, Alexander (2008), "Robust optimal control for a consumption-investment problem." *Mathematical Methods of Operations Research*, 67, 1–20.

Schmeidler, David (1989), "Subjective probability and expected utility without additivity." *Econometrica*, 57, 571–587.

Schroder, Mark and Costis Skiadas (1999), "Optimal consumption and portfolio selection with stochastic differential utility." *Journal of Economic Theory*, 89, 68–126.

Siniscalchi, Marciano (2011), "Dynamic choice under ambiguity." *Theoretical Economics*, 6, 379–421.

Smets, Philippe and Robert Kennes (1994), "The transferable belief model." *Artificial Intelligence*, 66, 191–234.

Trojani, Fabio and Paolo Vanini (2004), "Robustness and ambiguity aversion in general equilibrium." *Review of Finance*, 8, 279–324.

Uppal, Raman and Tan Wang (2003), "Model misspecification and underdiversification." *Journal of Finance*, 58, 2465–2486.

Veronesi, Pietro (2000), "How does information quality affect stock returns?" *Journal of Finance*, 55, 807–837.

Wachter, Jessica A. (2002), "Portfolio and consumption decisions under mean-reverting returns: An exact solution for complete markets." *Journal of Financial and Quantitative Analysis*, 37, 63–91.

Walley, Peter (1991), *Statistical Reasoning with Imprecise Probabilities*. Chapman and Hall.

Wang, Tan (2003), "Conditional preferences and updating." *Journal of Economic Theory*, 108, 286–321.

Welch, Ivo and Amit Goyal (2008), "A comprehensive look at the empirical performance of equity premium prediction." *Review of Financial Studies*, 21, 1455–1508.

Xia, Yihong (2001), "Learning about predictability: The effects of parameter uncertainty on dynamic asset allocation." *Journal of Finance*, 56, 205–246.

Zohar, Gady (2001), "A generalized Cameron-Martin formula with applications to partially observed dynamic portfolio optimization." *Mathematical Finance*, 11, 475–494.