

Learning Under Ambiguity

LARRY G. EPSTEIN

Boston University

and

MARTIN SCHNEIDER

New York University, Federal Reserve Bank of Minneapolis and NBER

First version received March 2005; final version accepted February 2007 (Eds.)

This paper considers learning when the distinction between risk and ambiguity matters. It first describes thought experiments, dynamic variants of those provided by Ellsberg, that highlight a sense in which the Bayesian learning model is extreme—it models agents who are implausibly ambitious about what they can learn in complicated environments. The paper then provides a generalization of the Bayesian model that accommodates the intuitive choices in the thought experiments. In particular, the model allows decision-makers' confidence about the environment to change—along with beliefs—as they learn. A portfolio choice application compares the effect of changes in confidence under ambiguity vs. changes in estimation risk under Bayesian learning. The former is shown to induce a trend towards more stock market participation and investment even when the latter does not.

1. INTRODUCTION

Models of learning typically assume that agents assign (subjective) probabilities to all relevant uncertain events. As a result, they leave no role for *confidence* in probability assessments; in particular, the degree of confidence does not affect behaviour. Thus, for example, agents do not distinguish between *risky* situations, where the odds are objectively known, and *ambiguous* situations, where they may have little information and hence also little confidence regarding the true odds. The Ellsberg Paradox has shown that this distinction is behaviourally meaningful in a static context: people treat ambiguous bets differently from risky ones. Importantly, the lack of confidence reflected by choices in the Ellsberg Paradox cannot be rationalized by *any* probabilistic belief. In particular, it is inconsistent with the foundations of Bayesian models of learning.

This paper considers learning when the distinction between risk and ambiguity matters. We assume that decision-makers view data as being generated by the same memoryless mechanism in every period. This *a priori* view motivates the exchangeable Bayesian model of learning and is commonly imposed on learners in a wide variety of economic applications. We view our principal contribution as twofold. First, we describe examples or thought experiments, dynamic variants of those provided by Ellsberg, in which intuitive behaviour contradicts the (exchangeable) Bayesian learning model. As in the static setting addressed by the Ellsberg Paradox, an agent may lack confidence in her *initial* information about the environment. However, confidence may now change as the agent learns. The examples reveal also another sense in which the Bayesian model is extreme—it models agents who are very ambitious about what they can learn, even in very complicated environments.

Our second contribution is to provide a generalization of the Bayesian model that accommodates the intuitive choices in the thought experiments. The model describes how the confidence of ambiguity-averse decision-makers changes over time. Beliefs and confidence are jointly represented by an evolving *set* of conditional distributions over future states of the world. This

set may shrink over time as agents become more familiar with the environment, but it may also expand if new information is surprising relative to past experience. We argue also that it models agents who take the complexity of their environment seriously and, therefore, are less ambitious than Bayesians about how much can be learnt. Finally, we show that the model is tractable in economic settings by applying it to dynamic portfolio choice.

The Bayesian model of learning about a memoryless mechanism is summarized by a triple (Θ, μ_0, ℓ) , where Θ is a parameter space, μ_0 is a prior over parameters, and ℓ is a likelihood. The parameter space represents features of the data-generating mechanism that the decision-maker tries to learn. The prior μ_0 represents initial beliefs about parameters. For a given parameter value $\theta \in \Theta$, the data are an independent and identically distributed sequence of signals $\{s_t\}_{t=1}^{\infty}$, where the distribution of any signal s_t is described by the probability measure $\ell(\cdot | \theta)$ on the period state space $S_t = S$. The triple (Θ, μ_0, ℓ) is the decision-maker's theory of how data are generated. This theory incorporates both prior information (through μ_0) and a view of how the signals are related to the underlying parameter (through ℓ). Beliefs on (payoff-relevant) states are equivalently represented by a probability measure p on sequences of signals (*i.e.* on S^{∞}), or by the process $\{p_t\}$ of one-step-ahead conditionals of p . The dynamics of Bayesian learning can be summarized by

$$p_t(\cdot | s^t) = \int_{\Theta} \ell(\cdot | \theta) d\mu_t(\theta | s^t), \quad (1)$$

where $s^t = (s_1, \dots, s_t)$ denotes the sample history at t and where μ_t is the posterior belief about θ , derived via Bayes' Rule.

We generalize the Bayesian learning model within the decision-theoretic framework of *recursive multiple-priors* utility, a model of utility put forth in Epstein and Wang (1994) and axiomatized in Epstein and Schneider (2003a) that extends Gilboa and Schmeidler's (1989) model of decision-making under ambiguity to an inter-temporal setting. Under recursive multiple-priors utility, beliefs and confidence are determined by a process $\{\mathcal{P}_t\}$ giving for each time and history a *set* of one-step-ahead conditionals. Since these sets describe responses to data, they are the natural vehicle for modelling learning. The present paper introduces a structure for the process $\{\mathcal{P}_t\}$ designed to capture learning about a memoryless mechanism. The process is constructed from a triple of sets $(\Theta, \mathcal{M}_0, \mathcal{L})$. As in the Bayesian case, Θ is a parameter space. The *set* of priors \mathcal{M}_0 represents the agent's initial view of these parameters. When \mathcal{M}_0 is not a singleton, it also captures (lack of) confidence in the information upon which this initial view is based. Finally, a *set* of likelihoods \mathcal{L} represents the agent's *a priori* view of the connection between signals and the true parameter.

The multiplicity of likelihoods and the distinction between Θ and \mathcal{L} are central to the way we model an agent who has modest (or realistic) ambitions about what she can learn. The set Θ represents features of the environment that the agent views as constant over time and that she therefore expects to learn. However, in complicated environments she may be wary of a host of other poorly understood factors that also affect realized signals. They vary over time in a way that she does not understand well enough even to theorize about, and therefore she does not try to learn about them. The multiplicity of likelihoods in \mathcal{L} captures these factors. Under regularity conditions, we show that beliefs become concentrated on one parameter value in Θ . However, ambiguity need not vanish in the long run, since the time-varying unknown features remain impossible to know even after many observations. Instead, the agent moves towards a state of time-invariant ambiguity, where she has learnt all that she can.¹

1. Existing dynamic applications of ambiguity typically impose time-invariant ambiguity. Our model makes explicit that this can be justified as the outcome of a learning process. See Epstein and Schneider (2003b) for further properties of time invariant ambiguity.

As a concrete illustration, consider independent tosses of a sequence of coins. If the decision-maker perceives the coins to be identical, then after many tosses she will naturally become confident that the observed empirical frequency of heads is close to a “true” frequency of heads that is relevant for forecasting future tosses. Thus, she will eventually become confident enough to view the data as an i.i.d. process. This laboratory-style situation is captured by the Bayesian model.²

More generally, suppose that she believes the tosses to be independent, but that she has no reason to be sure that the coins are identical. For example, if she is told the same about each coin but very little (or nothing at all) about each, then she would plausibly *admit the possibility* that the coins are not identical. In particular, there is no longer a compelling reason why data in the future should be i.i.d. with frequency of heads equal to the empirical frequency of heads observed in the past. Indeed, in contrast to a Bayesian, she may not even be sure whether the empirical frequencies of the data will converge, let alone expect her learning process to settle down at a single i.i.d. process.³ This situation cannot be captured by the exchangeable Bayesian model.⁴ One can view our model as an attempt to capture learning in such complicated (or vaguely specified and poorly understood) environments.

To illustrate our model in an economic setting, we study portfolio choice and asset market participation by investors who are ambiguity averse and learn over time about asset returns. Selective participation in asset markets has been shown to be consistent with optimal *static (or myopic)* portfolio choice by ambiguity-averse investors (Dow and Werlang, 1992). What is new in the present paper is that we solve the—more realistic—*inter-temporal* problem of an investor who rebalances his portfolio in light of new information that affects both beliefs and confidence. A key property of the solution is that an increase in confidence—captured in our model by a posterior set that shrinks over time—induces a quantitatively significant trend towards stock market participation and investment. This is in contrast to Bayesian studies of estimation risk, which tend to find small effects of learning on investment.

Another novel feature of the optimal portfolio is that the investment horizon matters for asset allocation because investors hedge exposure to an unknown (ambiguous) factor they are learning about. This *hedging of ambiguity* is distinct from the familiar hedging demand driven by inter-temporal substitution effects stressed by Merton (1971). Indeed, we show that it arises even when preferences over risky pay-off streams are logarithmic, so that the traditional hedging demand is 0.

We are aware of only three formal treatments of learning under ambiguity. Marinacci (2002) studies repeated sampling with replacement from an Ellsberg urn and shows that ambiguity is resolved asymptotically. A primary difference from our model is that Marinacci assumes a single likelihood. The statistical model proposed by Walley (1991, pp. 457–472) differs in details from ours, but is in the same spirit; in particular, it also features multiple likelihoods. An important difference, however, is that our model is consistent with a coherent axiomatic theory of dynamic choice between consumption processes. Accordingly, it is readily applicable to economic settings. A similar remark applies to Huber (1973), who also points to the desirability of admitting multiple likelihoods and outlines one proposal for doing so.

The paper is organized as follows. Section 2 presents a sequence of thought experiments to motivate our model. Section 3 briefly reviews recursive multiple-priors utility and then introduces

2. Under regularity conditions, the posterior μ_t converges to a Dirac measure on Θ , almost surely under p , the Bayesian’s subjective probability on sequences of data; see Schervish (1995).

3. For the exchangeable Bayesian model, convergence of empirical frequencies is a consequence of the law of large numbers for exchangeable random variables; see Schervish (1995).

4. If the data are generated by a memoryless mechanism and information about each observation is *a priori* the same, the exchangeable model is the only relevant Bayesian model. More complicated Bayesian models that introduce either serial dependence or a dependence of the distribution on calendar time are not applicable.

the learning model. Section 4 establishes properties of learning in the short and long runs. The portfolio choice application is described in Section 5. Proofs for results in Section 4 are collected in the appendix; proofs for some results in Section 5 are relegated to a supplementary appendix to this paper (Epstein and Schneider, 2007b).

2. EXAMPLES

In this section, we consider three related learning scenarios. Each involves urns containing five balls that are either black or white; Figure 1 illustrates the composition of the urns. However, the scenarios differ in a way intended to illustrate alternative hypotheses about learning in economic environments and to motivate the main ingredients of our model. In particular, we highlight intuitive behaviour that we show later can be accommodated by our learning model but not by the Bayesian model.

2.1. Scenario 1: Pure risk

An urn contains five balls. The agent is told that two are black, two are white and the fifth is either black or white depending on the outcome of the toss of an unbiased coin—the “coin ball” is black if the coin toss produces heads, and it is white otherwise. She cannot see the outcome of the coin toss, but she does observe, every period, a black or white ball. Her *ex ante* view is that these data are generated by sampling with replacement from the given urn. This may be either because she is actually doing the sampling or because she has been told that this is how signals are generated. Alternatively, it may simply reflect her subjective view of the environment. In any case, the obvious version of the Bayesian learning model (1) seems appropriate. In particular, she can be confident of learning the colour of the coin ball.

2.2. Scenario 2: Ambiguous prior information about a fixed urn

The information provided about the urn is modified to there is at least one non-coin ball of each colour. Thus unlike Scenario 1, no objective probabilities are given for the composition of the urn—there is ambiguity about the number of black non-coin balls. However, signals are again observed in each period, and they are viewed as resulting from sampling with replacement from the given urn. Because the agent views the factors underlying signals, namely the colour composition of the urn, as fixed through time, she can hope to learn it. That is, she would try to learn the true composition of the ambiguous urn.

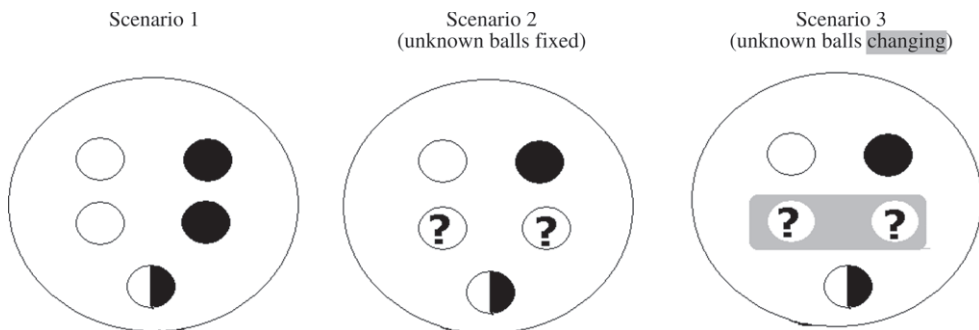


FIGURE 1
Urns for Scenarios 1–3

Compare this situation with the purely risky one, Scenario 1 above, assuming the two coin tosses are independent. In particular, consider the choice between betting on drawing black (white) from the risky urn as opposed to the ambiguous one. (Throughout, a bet on the colour of a ball drawn from an urn is understood to pay 1 until the ball has the desired colour and 0 otherwise.) The Ellsberg Paradox examines this choice prior to any sampling. As recalled in the Introduction, the intuitive behaviour pointed to by Ellsberg—the preference to bet on drawing black from the risky urn as opposed to the ambiguous one and a similar preference for white—is inconsistent with a single prior on the composition of the ambiguous urn.

Consider the ranking of bets after some sampling. It is plausible that *Ellsberg-type behaviour persists in the short run*: for any sufficiently short common sample history across the risky and ambiguous urns, a bet on the risky urn should remain preferable to an analogous bet on the ambiguous urn. At the same time, given the *ex ante* view that sampling is from an unchanging urn, it is intuitive that *learning should resolve ambiguity in the long run*. Asymptotically, the agent should behave as if she knew the fraction of black balls were equal to their empirical frequency. In the limit, for any common sample history across the two urns, she should be indifferent between betting on the next draw from the risky as opposed to the ambiguous urn. Both Scenarios 1 and 2 are thus conducive to inference and learning because there is one urn with fixed composition. Next, we examine learning in a more complex setting where signals are generated by a sequence of urns with time-varying composition.

2.3. Scenario 3: Unambiguous prior and changing urns

Signals are generated by successive sampling from a *sequence* of urns, each containing black and white balls. The urns are perceived in the following way: each urn consists of a coin ball and four other balls as above. The coin is tossed once at the start and determines the same colour for coin balls in all urns. However, non-coin balls are replaced every period with a new set of non-coin balls. This task is performed every period by an administrator. It is also known that (i) a new administrator is appointed every period, (ii) no administrator knows the previous history of urns or draws, and (iii) the only restriction on any given administrator is that at least one non-coin ball of each colour be placed in the urn.

Ex ante, before any draws are observed, this environment looks the same as Scenario 2: there is one coin ball, there is at least one black and one white non-coin ball, and there are two non-coin balls about which there is no information. The new feature in Scenario 3 is that the non-coin balls change over time in a way that is not well understood. What might an agent now try to learn? Since the coin-ball is fixed across urns and underlies all signals, it still makes sense to try to learn its colour. At the same time, learning about the changing sequence of non-coin balls has become more difficult. Compared to Scenario 2, one would thus expect agents to perceive more uncertainty about the overall proportion of black balls. We now argue that this should lead to different behaviour in the two scenarios both in the short run and in the long run.

Begin again by considering betting behaviour in the short run, assuming that the coin tosses in Scenarios 2 and 3 are independent. Before any sampling, one should be indifferent between bets in Scenarios 2 and 3, since the two settings are *a priori* identical. However, *suppose that one black draw from each urn has been observed and consider bets on the next draw being black. It is intuitive that the agent prefer to bet on black in the next draw in Scenario 2 rather than on the next draw in Scenario 3*. This is because the observed black draw in Scenario 2 is stronger evidence for the next draw also being black than is the observed black draw in Scenario 3. In Scenario 3, the agent does not understand well the dynamics of the non-coin balls and thus is plausibly concerned that the next urn may contain more white balls than did the first. Of course, she thinks that the next urn may also contain more black balls, but as in the static Ellsberg

choice problem, the more pessimistic possibility outweighs the optimistic one. Thus she would rather bet in Scenario 2 where the number of black non-coin balls is perceived not to change between urns.

Turn to behaviour in the long run. Scenarios 2 and 3 differ in what qualifies as a reasonable *a priori* view about the convergence of empirical frequencies. Since the urn in Scenario 2 is essentially fixed, one should be confident that the fraction of black balls drawn converges to some limit. In contrast, nothing in the description of Scenario 3 indicates that such convergence will occur. The sequence of administrators suggests only that the number of black non-coin balls is independent across successive urns. Otherwise, the description of the environment is very much consistent with perpetual change. It is thus reasonable to expect *a priori* that nothing can be learnt about future non-coin balls from sampling. Given how little information is provided about any single urn, the number of non-coin balls should be perceived as ambiguous at all dates. In particular, *even after long identical samples from the two scenarios, one would still prefer to bet on the next draw being black from the urns in Scenario 2 rather than in Scenario 3.*

2.4. Desirable properties of a learning model

To sum up, we would like a model that captures the following intuitive choices between bets in Scenarios 1–3. First, bets on risky urns (Scenario 1) are preferred to bets on ambiguous urns with fixed composition (Scenario 2) in the short run, but not in the long run. In other words, with only ambiguous prior information, Ellsberg-type behaviour is observed in the short run, but not in the long run. Second, bets on risky urns (Scenario 1) are preferred to bets on ambiguous urns with changing composition (Scenario 3) in both the short and long runs. In other words, in a complicated environment with changing urns, Ellsberg-type behaviour persists. It is clear that the Bayesian model cannot generate these choices, since it never gives rise to Ellsberg-type behaviour. Finally, bets on ambiguous urns with fixed composition (Scenario 2) should be preferred to bets on ambiguous urns with changing composition (Scenario 3). This third choice emphasizes that the distinction between ambiguous prior information about a fixed urn and ambiguity perceived about changing urns is behaviourally important.

3. A MODEL OF LEARNING

3.1. Recursive multiple priors

We work with a finite period state space $S_t = S$, identical for all times. One element $s_t \in S$ is observed every period. At time t , an agent's information consists of the history $s^t = (s_1, \dots, s_t)$. There is an infinite horizon, so S^∞ is the full state space.⁵ The agent ranks consumption plans $c = (c_t)$, where c_t is a function of the history s^t . At any date $t = 0, 1, \dots$, given the history s^t , the agent's ordering is represented by a conditional utility function U_t , defined recursively by

$$U_t(c; s^t) = \min_{p \in \mathcal{P}_t(s^t)} E^p [u(c_t) + \beta U_{t+1}(c; s^t, s_{t+1})], \quad (2)$$

where β and u satisfy the usual properties. The set of probability measures $\mathcal{P}_t(s^t)$ models beliefs about the next observation s_{t+1} , given the history s^t . Such beliefs reflect ambiguity when $\mathcal{P}_t(s^t)$ is a non-singleton. We refer to $\{\mathcal{P}_t\}$ as the *process of conditional one-step-ahead beliefs*. The

5. In what follows, measures on S^∞ are understood to be defined on the product σ -algebra on S^∞ , and those on any S_t are understood to be defined on the power set of S_t . While our formalism is expressed for S finite, it can be justified also for suitable metric spaces S but we ignore the technical details needed to make the sequel rigorous more generally.

process of utility functions is determined by specification of $\{\mathcal{P}_t\}$, $u(\cdot)$, and β , which constitute the primitives of the functional form.

To clarify the connection to the Gilboa–Schmeidler model, it is helpful to rewrite utility using discounted sums. In a Bayesian model, the process of all conditional-one-step-ahead probability measures uniquely determines a probability measure over the full state space. Similarly, the process $\{\mathcal{P}_t\}$ determines a unique set of probability measures \mathcal{P} on S^∞ satisfying the regularity conditions specified in Epstein and Schneider (2003a).⁶ Thus, one obtains the following equivalent and explicit formula for utility:

$$U_t(c; s^t) = \min_{p \in \mathcal{P}} E^p \left[\sum_{s \geq t} \beta^{s-t} u(c_s) \mid s^t \right]. \quad (3)$$

This expression shows that each conditional ordering conforms to the multiple-priors model in Gilboa and Schmeidler (1989), with the set of priors for time t determined by updating the set \mathcal{P} measure-by-measure via Bayes' Rule.

Axiomatic foundations for recursive multiple-priors utility are provided in Epstein and Schneider (2003). The essential axioms are that conditional orderings (i) satisfy the Gilboa–Schmeidler axioms, (ii) are connected by dynamic consistency, and (iii) do not depend on unrealized parts of the decision tree: utility given the history s^t depends only on consumption in states of the world that can still occur. To ensure such dynamic behaviour in an application, it is sufficient to specify beliefs directly via a process of one-step-ahead conditionals $\{\mathcal{P}_t\}$. In the case of learning, this approach has additional appeal: because $\{\mathcal{P}_t\}$ describes how an agent's view of the next state of the world depends on history, it is a natural vehicle for modelling learning dynamics. The analysis in Epstein and Schneider (2003) restricts $\{\mathcal{P}_t\}$ only by technical regularity conditions. We now proceed to add further restrictions to capture how the agent responds to data.

3.2. Learning

Our model of learning applies to situations where a decision-maker holds the *a priori* view that data are generated by the same memoryless mechanism every period. This *a priori* view also motivates the Bayesian model of learning about an underlying parameter from conditionally i.i.d. signals.

As in the Bayesian model outlined in the introduction, our starting point is again (a finite period state space S and) a parameter space Θ that represents features of the data the decision-maker tries to learn. To accommodate ambiguity in initial beliefs about parameters, represent those beliefs by a set \mathcal{M}_0 of probability measures on Θ . The size of \mathcal{M}_0 reflects the decision-maker's (lack of) confidence in the prior information on which initial beliefs are based. A technically convenient assumption is that \mathcal{M}_0 is weakly compact;⁷ this permits us to refer to minima, rather than infima, over \mathcal{M}_0 .

Finally, we adopt a set of likelihoods \mathcal{L} —every parameter value $\theta \in \Theta$ is associated with a set of probability measures

$$\mathcal{L}(\cdot \mid \theta) = \{\ell(\cdot \mid \theta) : \ell \in \mathcal{L}\}.$$

Each $\ell : \Theta \rightarrow \Delta(S)$ is a likelihood function, so that $\theta \mapsto \ell(A \mid \theta)$ is assumed measurable for every $A \subset S$. Another convenient technical condition is that \mathcal{L} is compact when viewed as a

6. In the infinite horizon case, uniqueness obtains only if \mathcal{P} is assumed also to be *regular* in a sense defined in Epstein and Schneider (2003b), generalizing to sets of priors the standard notion of regularity for a single prior.

7. More precisely, measures in \mathcal{M}_0 are defined on the implicit and suppressed σ -algebra \mathcal{B} associated with Θ . Take the latter to be the power set when Θ is finite. The weak topology is that induced by bounded and \mathcal{B} -measurable real-valued functions on Θ . \mathcal{M}_0 is weakly compact iff it is weakly closed.

subset of $(\Delta(S))^\Theta$ (under the product topology). Finally, to avoid the problem of conditioning on zero probability events, assume that each $\ell(\cdot | \theta)$ has full support.

Turn to interpretation. The general idea is that the agent perceives the data-generating mechanism as having some unknown features represented by θ that are common across experiments. Because these features are perceived as common across time or experiments, she can try to learn about them. At the same time, she feels there are other factors underlying data, represented by \mathcal{L} , that are variable across experiments. It is known that the variable features are determined by a memoryless mechanism—this is why the set \mathcal{L} does not depend on history. It is also known that the relative importance of the variable features is the same every period, as in the urn example where there are always four non-coin balls and one coin ball. This is why the set \mathcal{L} does not depend on time. However, the factors modelled by the set \mathcal{L} are variable across time in a way that the agent does not understand beyond the limitation imposed by \mathcal{L} . In particular, at any point in time, *any* element of \mathcal{L} might be relevant for generating the next observation. Accordingly, while she can try to learn the true θ , she has decided that she will not try to (or is not able to) learn more.

Conditional independence implies that past signals s^t affect beliefs about future signals (such as s_{t+1}) only to the extent that they affect beliefs about the parameter. Let $\mathcal{M}_t(s^t)$, to be described below, denote the set of posterior beliefs about θ given that the sample s^t has been observed. The dynamics of learning can again be summarized by a process of one-step-ahead conditional beliefs. However, in contrast to the Bayesian case (1), there is now a (typically non-generate) *set* assigned to every history:

$$\mathcal{P}_t(s^t) = \left\{ p_t(\cdot) = \int_{\Theta} \ell(\cdot | \theta) d\mu_t(\theta) : \mu_t \in \mathcal{M}_t(s^t), \ell \in \mathcal{L} \right\}, \quad (4)$$

or, in convenient notation,

$$\mathcal{P}_t(s^t) = \int_{\Theta} \mathcal{L}(\cdot | \theta) d\mathcal{M}_t(\theta).$$

This process enters the specification of recursive multiple-priors preferences (2).⁸

3.2.1. Updating and reevaluation. To complete the description of the model, it remains to describe the evolution of the posterior beliefs \mathcal{M}_t . Imagine a decision-maker at time t looking back at the sample s^t . In general, she views both her prior information and the sequence of signals as ambiguous. As a result, she will typically entertain a number of different theories about how the sample was generated. A *theory* is a pair (μ_0, ℓ^t) , where μ_0 is a prior belief on Θ and $\ell^t = (\ell_1, \dots, \ell_t) \in \mathcal{L}^t$ is a *sequence* of likelihoods. The decision-maker contemplates different sequences ℓ^t because she is not confident that signals are identically distributed over time.

We allow for different attitude towards past and future signals. On the one hand, \mathcal{L} is the set of likelihoods possible in the future. Since the decision-maker has decided she cannot learn the true sequence of likelihoods, it is natural that beliefs *about the future* be based on the whole set \mathcal{L} as in (4). On the other hand, the decision-maker may *reevaluate*, with hindsight, her views about what sequence of likelihoods was relevant for generating data *in the past*. Such revision is possible because the agent learns more about θ and this might make certain theories more or less plausible. For example, some interpretation of the signals, reflected in a certain sequence

8. Given compactness of \mathcal{M}_0 and \mathcal{L} , one can show that $\mathcal{M}_t(s^t)$ defined below is also compact and subsequently that each $\mathcal{P}_t(s^t)$ is compact. This justifies the use of minima in (2).

$\ell^t = (\ell_1, \dots, \ell_t)$, or some prior experience, reflected in a certain prior $\mu_0 \in \mathcal{M}_0$, might appear not very relevant if it is part of a theory that does not explain the data well.

To formalize reevaluation, we need two preliminary steps. First, how well a theory (μ_0, ℓ^t) explains the data is captured by the (unconditional) data density evaluated at s^t :

$$\int \prod_{j=1}^t \ell_j(s_j | \theta) d\mu_0(\theta).$$

Here conditional independence implies that the conditional distribution given θ is simply the product of the likelihoods ℓ_j . Prior information is taken into account by integrating out the parameter using the prior μ_0 . The higher the data density, the better is the observed sample s^t explained by the theory (μ_0, ℓ^t) . Second, let $\mu_t(\cdot; s^t, \mu_0, \ell^t)$ denote the posterior derived from the theory (μ_0, ℓ^t) by Bayes' Rule given the data s^t . This posterior can be calculated recursively by Bayes' Rule, taking into account time variation in likelihoods:

$$d\mu_t(\cdot; s^t, \mu_0, \ell^t) = \frac{\ell_t(s_t | \cdot)}{\int_{\Theta} \ell_t(s_t | \theta') d\mu_{t-1}(\theta'; s^{t-1}, \mu_0, \ell^{t-1})} d\mu_{t-1}(\cdot; s^{t-1}, \mu_0, \ell^{t-1}). \quad (5)$$

Reevaluation takes the form of a likelihood-ratio test. The decision-maker discards all theories (μ_0, ℓ^t) that do not pass a likelihood-ratio test against an alternative theory that puts maximum likelihood on the sample. Posteriors are formed only for theories that pass the test. Thus posteriors are given by

$$\mathcal{M}_t^\alpha(s^t) = \{ \mu_t(s^t; \mu_0, \ell^t) : \mu_0 \in \mathcal{M}, \ell^t \in \mathcal{L}^t, \quad (6)$$

$$\int \prod_{j=1}^t \ell_j(s_j | \theta) d\mu_0(\theta) \geq \alpha \max_{\substack{\tilde{\mu}_0 \in \mathcal{M}_0 \\ \tilde{\ell}^t \in \mathcal{L}^t}} \int \prod_{j=1}^t \tilde{\ell}_j(s_j | \theta) d\tilde{\mu}_0 \}.$$

Here α is a parameter, $0 < \alpha \leq 1$, that governs the extent to which the decision-maker is willing to *reevaluate her views about how past data were generated* in the light of new sample information. The likelihood-ratio test is more stringent and the set of posteriors smaller, the greater is α . In the extreme case $\alpha = 1$, only parameters that achieve the maximum likelihood are permitted. If the maximum likelihood estimator is unique, ambiguity about parameters is resolved as soon as the first signal is observed. More generally, we have that $\alpha > \alpha'$ implies $\mathcal{M}_t^\alpha \subset \mathcal{M}_t^{\alpha'}$. It is important that the test is done after every history. In particular, a theory that was disregarded at time t might look more plausible at a later time and posteriors based on it may again be taken into account.

In summary, our model of learning about an ambiguous memoryless mechanism is given by the tuple $(\Theta, \mathcal{M}_0, \mathcal{L}, \alpha)$. As described, the latter induces, or *represents*, the process $\{\mathcal{P}_t\}$ of one-step-ahead conditionals via

$$\mathcal{P}_t(s^t) = \int_{\Theta} \mathcal{L}(\cdot | \theta) d\mathcal{M}_t^\alpha(\theta),$$

where \mathcal{M}_t^α is given by (6). The model reduces to the Bayesian model when both the set of priors and the set of likelihoods have only a single element.

Another important special case occurs if \mathcal{M}_0 consists of *several* Dirac measures on the parameter space, in which case there is a simple interpretation of the updating rule: \mathcal{M}_t^α contains all $\tilde{\theta}$'s such that the hypothesis $\theta = \tilde{\theta}$ is not rejected by an asymptotic likelihood-ratio test performed with the given sample, where the critical value of the $\chi^2(1)$ distribution is $-2 \log \alpha$. Since $\alpha > 0$, (Dirac) measures over parameter values are discarded or added to the set, and \mathcal{P}_t varies over

time. The Dirac priors specification is convenient for applications—it will be used in our portfolio choice example below. Indeed, one may wonder whether there is a need for non-Dirac priors at all. However, more general priors provide a useful way to incorporate objective probabilities, as illustrated by the scenarios in Section 2.⁹

3.3. Scenarios revisited

We now describe how we would model the three scenarios of Section 2, and we show that our set-up can accommodate the intuitive behaviour described there. The results here clarify the role of multiple priors and multiple likelihoods, but are essentially independent of the revaluation parameter α . We return to the role of α in Section 4, where we consider the evolution of beliefs in the context of Scenario 3.

3.3.1. Scenarios 1 and 2. For both scenarios, specify $S = \{B, W\}$ and $\Theta = \{\frac{n}{5} : n = 1, 2, 3, 4\}$. The state s corresponds to the colour of a drawn ball, and θ indicates the proportion of black balls (coin or non-coin). For Scenario 1, assume that there is a single prior with $\mu_0(\theta) = \frac{1}{2}$ for $\theta = \frac{2}{5}$ or $\frac{3}{5}$ and define the single likelihood by $\ell(B | \theta) = \theta$. For Scenario 2, specify a representation $(\Theta, \mathcal{M}_0, \mathcal{L}, \alpha)$ by setting $\mathcal{L} = \{\ell\}$, fixing a revaluation parameter α and defining the set \mathcal{M}_0 of priors on Θ as follows. Let $P \subset \Delta(\{1, 2, 3\})$ denote beliefs about the number of black non-coin balls. Since the coin-ball and other colours are independent, each prior should have the form

$$\mu_0^p(\theta) = \frac{1}{2}p(5\theta - 1) + \frac{1}{2}p(5\theta), \quad (7)$$

for some $p \in P$. Thus let $\mathcal{M}_0 = \{\mu_0^p : p \in P\}$.

This set-up captures Ellsberg-type behaviour in the short run. Because the set of one-step-ahead probabilities for the next draw is non-degenerate in Scenario 2, a bet on black would be evaluated under a prior that puts the lowest possible conditional probability on a black draw, and similarly for white. As a result, either bet in Scenario 2 will be inferior to its counterpart in Scenario 1. In addition, keeping the history of draws the same across scenarios, any difference in utility would become smaller and smaller over time. With a single likelihood, it is easy to verify that the posteriors derived from \mathcal{M}_0 become more similar over time and eventually become concentrated on a single parameter value. Ambiguity is thus resolved in the long run.

3.3.2. Scenario 3. A natural parameter space here is $\Theta = \{B, W\}$, corresponding to the two possible colours of the coin-ball. There is a single prior μ_0 —the $(\frac{1}{2}, \frac{1}{2})$ prior on Θ corresponding to the toss of the unbiased coin. Ambiguity perceived about the changing non-coin balls is captured by a non-degenerate set of likelihoods, specified as follows. Let $P \subset \Delta(\{1, 2, 3\})$ denote beliefs about the number of black non-coin balls, where P is the same set used in Scenario 2. Since the first urns in the two scenarios are identical, it is natural to use the same set P to describe beliefs about them. Moreover, though the urns differ along the sequence in Scenario 3, successive urns are indistinguishable, which explains why P can be used also to describe the second urn in the present scenario. Each p in P suggests one likelihood function via

$$\ell^p(B | B) = \frac{\sum \lambda p(\lambda) + 1}{5}, \quad \ell^p(B | W) = \frac{\sum \lambda p(\lambda)}{5}. \quad (8)$$

9. Another example is in Epstein and Schneider (2007a), where a representation with a single prior, and $\alpha = 0$ is used to model the distinction between tangible (well-measured, probabilistic) and intangible (ambiguous) information.

Finally, let¹⁰

$$\mathcal{L} = \{\ell^p : p \in P\}. \quad (9)$$

Our model of Scenario 3 also leads to a non-degenerate set of one-step-ahead conditionals in the short run. The difference from Scenario 2 is that ambiguity enters through beliefs about signals, captured by multiple likelihoods, rather than through multiple priors on Θ . The multiplicity of likelihoods also implies directly that ambiguity never vanishes. Indeed, a bet on black in Scenario 3 will always be evaluated under a likelihood that puts the lowest possible conditional probability on a black draw, and similarly for white. Either bet in Scenario 3 will thus be inferior to its counterpart in Scenario 1, even in the long run.

Focus now on the other choice described above—given that one black ball has been drawn, would you rather bet on the next draw being black in Scenario 2 or in Scenario 3? This question highlights one difference between learning in simple settings (Scenario 2) as opposed to complex ones (Scenario 3), and thus *demonstrates the key role played by multiple likelihoods*. We argued above that it is intuitive that one prefer to bet in Scenario 2—the black draw is stronger evidence there for the next draw also being black. The appendix shows that our model predicts these intuitive choices. In particular, it demonstrates that the minimal predictive probability for a black ball conditional on observing a black ball,

$$\min_p \mathcal{P}_1(B) = \min_{\ell \in \mathcal{L}, \mu_1 \in \mathcal{M}_1^\alpha(B)} \int_{\Theta} \ell(B | \theta) d\mu_1(\theta),$$

is smaller under Scenario 3 than under Scenario 2. An ambiguity-averse agent who ranks bets according to the smallest probability of winning will thus prefer to bet on the urn from Scenario 2.

3.4. Discussion

Two further remarks are in order about the structure of our model and its relationship to the Bayesian model. Consider first the question of foundations. Though recursive multiple priors is an axiomatic model, we do not have axiomatic foundations for the specialization described here. For the Bayesian model, the de Finetti Theorem shows that the representation (1) is possible if and only if the prior p on S^∞ is exchangeable. We are missing a counterpart of the de Finetti Theorem for our setting. Nevertheless, without discounting the importance of this missing element, we would argue that our model can be justified on the less formal grounds of cognitive plausibility.

Begin with a perspective on the de Finetti Theorem. While de Finetti starts from a prior p on S^∞ and the assumption of exchangeability, the usefulness of his theorem stems in large part from the reverse perspective: it describes a cognitively plausible procedure for a decision-maker to arrive at a prior p and to determine thereby if, indeed, exchangeability is acceptable (see, for example, Kreps 1988, pp. 155–156). In some settings, there is a natural pair (Θ, ℓ) and, when

10. Given the symmetry of the environment, a natural candidate for P is

$$P_\varepsilon = \left\{ p \in \Delta(\{1, 2, 3\}) : \sum \lambda p(\lambda) \in [2 - \varepsilon, 2 + \varepsilon] \right\},$$

where ε is a parameter, $0 \leq \varepsilon \leq 1$. In the special case where $\varepsilon = 1$, the agent attaches equal weight to all *logically possible* urn compositions $\lambda = 1, 2$, or 3 . More generally, P_ε incorporates a *subjective* element into the specification. Just as subjective expected utility theory does not impose connections between the Bayesian prior and objective features of the environment, so too the set of likelihoods is subjective (varies with the agent) and is not uniquely determined by the facts. For example, for $\varepsilon < 1$, the agent attaches more weight to the “focal” likelihood corresponding to $\lambda = 2$ as opposed to the more extreme scenarios $\lambda = 1, 3$.

Θ is simple, forming a prior μ_0 over Θ is easier than forming one directly over S^∞ . Thus the representation (1) can be viewed as providing a procedure for the decision-maker to arrive at the measure p on S^∞ to use for decision-making. Moreover, to the extent that (Θ, ℓ) is simple and natural given the environment, the procedure is cognitively plausible.

In fact, cognitive plausibility is essential to support the typical application of the Bayesian model. After all, a Bayesian modeller must assume a particular parametric specification for (Θ, μ_0, ℓ) . This amounts, via (1), to assuming a *particular* prior p for which there is no axiomatic justification. Instead, cognitive plausibility is the only supporting argument. However, outside of contrived settings, where the data-generating mechanism is simple and transparent, the existence of a cognitively plausible specification of (Θ, μ_0, ℓ) is problematic. In particular, the Bayesian model presumes a degree of confidence on the part of the agent that seems non-intuitive in complicated environments. We view enhanced cognitive plausibility as a justification for our model. As illustrated by the examples and the way in which we model them, greater cognitive plausibility is afforded if we allow the decision-maker to have multiple priors on the parameter space Θ and, more importantly, multiple likelihoods.

A second feature of our model—and one that it shares with the Bayesian approach—is that the decomposition between learnable and unlearnable features of the environment is exogenous. Indeed, the learnable parameters Θ and the unlearnable features embodied by a non-singleton \mathcal{L} are given at the outset. We thus do not explain features that the agent tries to learn in any given setting. The same is true in the Bayesian case, where the decomposition is extreme— \mathcal{L} is a singleton and the agent tries to identify an i.i.d. process. Thus, while we leave important aspects unmodelled, we do extend the Bayesian model in the direction of permitting more plausible ambitions about what can be learnt in complex settings. The Ellsberg-style examples and the portfolio choice application illustrate that natural decompositions may be suggested by descriptions of the environment.

4. BELIEF DYNAMICS

In this section we illustrate further properties of learning under ambiguity. We first examine the dynamics of beliefs and confidence in the short term, using the example of Scenario 3. We then provide a general result on convergence of the learning process in the long run.

4.1. Inference from small samples

As in Section 3.3, take $\Theta = \{B, W\}$, $\mu_0 = \left(\frac{1}{2}, \frac{1}{2}\right)$, and assume further that the agent weighs equally all the logically possible combinations of non-coin balls, that is, \mathcal{L} is given by (8) with

$$P = \left\{ p \in \Delta(\{1, 2, 3\}) : \sum \lambda p(\lambda) \in [1, 3] \right\}.$$

The evolution of the posterior set \mathcal{M}_t^α in (6) shows how signals tend to induce ambiguity about parameters even where there is no ambiguity *ex ante* (singleton \mathcal{M}_0^α). This happens when the agent views some aspects of the data-generating mechanism as too difficult to try to learn, as modelled by a non-singleton set of likelihoods. More generally, \mathcal{M}_t^α can expand or shrink depending on the signal realization.

Posterior beliefs can be represented by the intervals

$$\left[\min_{\mu_t \in \mathcal{M}_t^\alpha} \mu_t(B), \max_{\mu_t \in \mathcal{M}_t^\alpha} \mu_t(B) \right].$$

Figure 2 depicts the evolution of beliefs about the coin-ball being black as more balls are drawn. The top panel shows the evolution of the posterior interval for a sequence of draws such

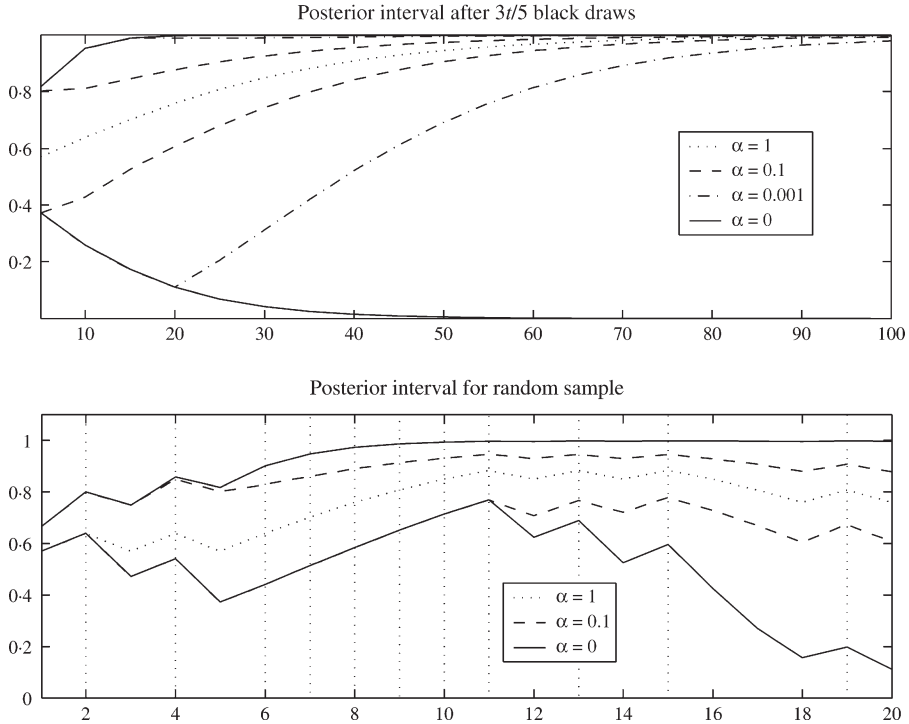


FIGURE 2

The posterior interval is the range of the posterior probability $\mu_t(B)$ that the coin ball is black. In the top panel, the sample is selected to keep the fraction of black balls constant. In the bottom panel, vertical lines indicate black balls drawn

that the number of black balls is $\frac{3t}{5}$, for $t = 5, 10, \dots$. Intervals are shown for $\alpha = 0.1$ and $\alpha = 0.001$, as well as for $\alpha = 0$, to illustrate what would happen without revaluation. In addition, a single line is drawn for the case $\alpha = 1$, where the interval is degenerate. For example, after the first five signals, with three black balls drawn, the agent with $\alpha = 0.1$ assigns a posterior probability between 0.4 and 0.8 to the coin ball being black.

What happens if the same sample is observed again? Our model captures two effects. First, a larger batch of signals permits more possible interpretations of past signals. For example, having seen 10 draws, the agent may believe that all six black draws came about *although* each time there were the “most adverse” conditions, that is, all but one non-coin ball was white. This interpretation strongly suggests that the coin ball itself is black. The argument also becomes stronger the more data are available—after only five draws, the appearance of three black balls under “adverse” conditions is not as remarkable.

Second, even though it is not possible to learn about *future* non-coin balls, past draws are informative about *past* non-coin balls, not just about the parameter. To illustrate, suppose 50 draws have been observed and that 30 of them have been black. Given a theory about the non-coin balls, these observations suggest that the coin ball is more likely to be black than white. At the same time, given a value of the parameter, the observations suggest that more than half of the past urns contained three black non-coin balls. Sensible updating takes the latter feature into account. In our model, this is accommodated by revaluation. The evolution of confidence, measured by the size of the posterior interval, then depends on how much agents reevaluate their views. For an agent with $\alpha = 0.001$, the posterior interval expands between $t = 5$ and $t = 20$. In

this sense, a sample of 10 or 20 signals induces *more* ambiguity than a sample of five. However, reevaluation implies that large enough batches of signals induce less ambiguity than smaller ones.

The lower panel of Figure 2 tracks the evolution of posterior intervals along one particular sample. Black balls were drawn at dates indicated by vertical lines, while white balls were drawn at the other dates. Taking the width of the interval as a measure, the extent of ambiguity is seen to respond to data. In particular, a phase of many black draws (periods 5–11, for example) shrinks the posterior interval, while an “outlier” (the white ball drawn in period 12) makes it expand again. This behaviour is reminiscent of the evolution of the Bayesian posterior variance, which is also maximal if the fraction of black balls is half.

4.2. Beliefs in the long run

Turn now to behaviour after a large number of draws. As discussed above, our model describes agents who are not sure whether empirical frequencies will converge. Nevertheless, it is useful to derive limiting beliefs for the case when such convergence occurs: the limiting beliefs also approximately describe behaviour after a large, but finite, number of draws. Therefore, they characterize behaviour in the long run. The limiting result below holds with probability 1 under a true data-generating process, described by a probability P^* on S^∞ . For this process, we require only that the empirical frequencies of each of the finite number of states $s \in S$ converge, almost surely under P^* . In what follows, these limiting frequencies are described by a measure ϕ on S .

By analogy with the Bayesian case, the natural candidate parameter value on which posteriors might become concentrated maximizes the data density of an infinite sample. With multiple likelihoods, any data density depends on the sequence of likelihoods that is used. In what follows, it is sufficient to focus on sequences such that the same likelihood is used whenever state s is realized. A likelihood sequence can then be represented by a collection $(\ell_s)_{s \in S}$. Accordingly, define the log data density after maximization over the likelihood sequence by¹¹

$$H(\theta) := \max_{(\ell_s)_{s \in S}} \sum_{s \in S} \phi(s) \log \ell_s(s | \theta). \quad (10)$$

The following result (proved in the appendix) summarizes the behaviour of the posterior set in the long run.

Theorem 1. *Suppose that Θ is finite and*

- (i) *for each state $s \in S$, the empirical frequency converges almost surely under P^* ;*
- (ii) *θ^* is the unique maximizer of $H(\theta)$, defined in (10);*
- (iii) *$\mu_0(\theta^*) > 0$ for some $\mu_0 \in \mathcal{M}_0$;*
- (iv) *there exists $\kappa > 0$ such that for every $\mu_0 \in \mathcal{M}_0$.*

$$\text{either } \mu_0(\theta^*) = 0 \quad \text{or} \quad \mu_0(\theta^*) \geq \kappa.$$

Then every sequence of posteriors from \mathcal{M}_t^α converges to the Dirac measure δ_{θ^} , almost surely under the true probability P^* , and the convergence is uniform, that is, there is a set $\Omega \subset S^\infty$ with $P^*(\Omega) = 1$ such that for every $s^\infty \in \Omega$,*

$$\mu_t(\theta^*) \rightarrow 1,$$

uniformly over all sequences of posteriors $\{\mu_t\}$ satisfying $\mu_t \in \mathcal{M}_t^\alpha$ for all t .

11. Our assumptions on \mathcal{L} —its compactness and that each $\ell(\cdot | \theta)$ has full support as well as finiteness of S and Θ , ensure that H is well defined.

Condition (ii) is an identification condition: it says that there is at least one sequence of likelihoods (*i.e.* the maximum likelihood sequence), such that the sample with empirical frequency measure ϕ can be used to distinguish θ^* from any other parameter value. Condition (iv) is satisfied if every prior in \mathcal{M}_0 assigns positive probability to θ^* (because \mathcal{M}_0 is weakly compact). The weaker condition stated accommodates also the case where all priors are Dirac measures (including specifically the Dirac measure concentrated at θ^*), as well as the case of a single prior containing θ^* in its support. (In Scenario 3, (iv) is satisfied for any set of priors where the probability of a black coin ball is bounded away from 0.)

Under conditions (ii) and (iv), and if Θ is finite, then in the long run only the maximum likelihood sequence is permissible, and the set of posteriors converges to a singleton. The agent thus resolves any ambiguity about factors that affect all signals, captured by θ . At the same time, ambiguity about future realizations s_t does not vanish. Instead, beliefs in the long run become close to $\mathcal{L}(\cdot | \theta^*)$. The learning process settles down in a state of time-invariant ambiguity.

The asserted uniform convergence is important in order that convergence of beliefs translate into long-run properties of preference. Thus it implies that for the process of one-step-ahead conditionals $\{\mathcal{P}_t\}$ given by (4),

$$\mathcal{P}_t(s^t) = \int_{\Theta} \mathcal{L}(\cdot | \theta) d\mathcal{M}_t^\alpha(\theta),$$

we have, almost surely under P^* ,

$$\begin{aligned} \min_{\mathcal{P}_t \in \mathcal{P}_t} \int f(s_{t+1}) d\mathcal{P}_t &= \min_{\mu_t \in \mathcal{M}_t^\alpha} \min_{\ell \in \mathcal{L}} \int_{\Theta} \left[\int_{S_{t+1}} f(s_{t+1}) d\ell(s_{t+1} | \theta) \right] d\mu_t(\theta) \\ &\rightarrow \min_{\ell \in \mathcal{L}} \int_{S_{t+1}} f(s_{t+1}) d\ell(s_{t+1} | \theta^*), \end{aligned}$$

for any $f : S_{t+1} \rightarrow \mathbb{R}^1$ describing a one-step-ahead prospect (in utility units).¹² In particular, this translates directly into the utility of consumption processes for which all uncertainty is resolved next period.

As a concrete example, consider long-run beliefs about the urns in Scenario 3. Let ϕ_∞ denote the limiting frequency of black balls. Suppose also that beliefs are given by (9). Maximizing the data density with respect to the likelihood sequence yields

$$\begin{aligned} H(\theta) &= \phi_\infty \max_{\lambda_B} \log \frac{1_{\{\theta=B\}} + \lambda_B}{5} + (1 - \phi_\infty) \max_{\lambda_W} \log \frac{5 - 1_{\{\theta=B\}} - \lambda_W}{5} \\ &= \phi_\infty \log \frac{1_{\{\theta=B\}} + 3}{5} + (1 - \phi_\infty) \log \frac{4 - 1_{\{\theta=B\}}}{5}. \end{aligned}$$

The first term captures all observed black balls and is therefore maximized by assuming $\lambda_B = 3$ black non-coin balls. Similarly, the likelihood of white draws is maximized by setting $\lambda_W = 1$. It follows that the identification condition is satisfied except in the knife-edge case $\phi_\infty = \frac{1}{2}$. Moreover, $\theta^* = B$ if and only if $\phi_\infty > \frac{1}{2}$. Thus the theorem implies that an agent who

12. The function $F(\ell, \mu) = \int_{\Theta} \left[\int_{S_{t+1}} f(s_{t+1}) d\ell(s_{t+1} | \theta) \right] d\mu(\theta)$ is jointly continuous on the compact set $\mathcal{L} \times \Delta(\Theta)$. Therefore, by the Maximum Theorem, $\min_{\ell \in \mathcal{L}} F(\ell, \mu)$ is (uniformly) continuous in μ . Finally, apply the (uniform over \mathcal{M}_t^α 's) convergence of posteriors to δ_{θ^*} proved in the theorem.

observes a large number of draws with a fraction of black balls above half believes it very likely that the colour of the coin ball is black. The role of α is only to regulate the speed of convergence to this limit. This dependence is also apparent from the dynamics of the posterior intervals in Figure 2.

The example also illustrates the role of the likelihood-ratio test (6) in ensuring learnability of the representation. Suppose that $\phi_\infty = \frac{3}{5}$. The limiting frequency of $\frac{3}{5}$ black draws could be realized either because there is a black coin ball and on average half of the non-coin balls were black, or because the coin ball is white and all of the urns contained three black non-coin balls. If α were equal to 0, both possibilities would be taken equally seriously and the limiting posterior set would contain Dirac measures that place probability 1 on either $\theta = B$ or $\theta = W$. This pattern is apparent from Figure 2. In contrast, with $\alpha > 0$, reevaluation eliminates the sequence where all urns contain three black non-coin balls as unlikely.

The set of “true” processes P^* for which the theorem holds is large. Like a Bayesian who is sure the data are exchangeable, the agent described by the theorem is convinced that the data-generating process is memoryless. As a result, learning is driven only by the empirical frequencies of the one-dimensional events, or subsets of S . Given any process for which those frequencies converge to a given limit, the agent will behave in the same way. Of course, in applications one would typically consider a “truth” that is memoryless. Like the exchangeable Bayesian model, our model is best applied to situations where the agent is *correct* in his *a priori* judgement that the data-generating process is memoryless. Importantly, a memoryless process for which empirical frequencies converge need not be i.i.d. There exists a large class of serially independent, but non-stationary, processes for which empirical frequencies converge.¹³ In fact, there is a large class of such processes that cannot be distinguished from an i.i.d. process with distribution ϕ on the basis of any finite sample. Thus even if they are convinced that the empirical frequencies converge, a concern with ongoing change may lead agents to never be confident that they have identified a “true” i.i.d. data-generating process with distribution ϕ .¹⁴

Incomplete learning can occur also in a Bayesian model. For example, if the true data-generating measure is not absolutely continuous with respect to an agent’s belief, violating the Blackwell and Dubins (1962) conditions, then beliefs may not converge to the truth.¹⁵ However, even then the Bayesian agent believes, and behaves as if, they will. In contrast, agents in our model are aware of the presence of hard-to-describe factors that prevent learning and their actions reflect the residual uncertainty.

Incomplete learning occurs also in models with persistent hidden state variables, such as regime switching models. Here learning about the state variable never ceases because agents are sure that the state variable is forever changing. The distribution of these regime changes is known *a priori*. Agents thus track a *known* data-generating process that is *not* memoryless. In contrast, our model applies to memoryless mechanisms and learning is about a fixed true parameter. Nevertheless, because of ambiguity, the agent reaches a state where no further learning is possible although the data-generating mechanism is not yet understood completely.

5. DYNAMIC PORTFOLIO CHOICE

Stock returns provide a natural economic example of an ambiguous memoryless mechanism. The assumption of i.i.d. stock returns is common in empirical work, since there is only weak evidence

13. See Nielsen (1996) for a broad class of examples.

14. As discussed earlier, they may also reach this view because they are not sure whether the empirical frequencies converge in the first place.

15. As a simple example, if the parameter governing a memoryless mechanism were outside the support of the agent’s prior, the agent could obviously never learn the true parameter.

of return predictability. At the same time, the distribution of stock returns has proven hard to pin down; in particular, uncertainty about the size of the equity premium remains large. In this section, we solve the inter-temporal portfolio choice problem of an ambiguity-averse investor who learns about the equity premium and contrast it with that of a Bayesian investor. As emphasized above, the Bayesian model is extreme in that it does not allow for confidence in probability assessments; in particular, Bayesian learning cannot capture changes in confidence. In a portfolio choice context, this leads to counterintuitive predictions for optimal investment strategies.

Suppose an investor believes that the equity premium is fixed, but unknown, so that it must be estimated from past returns. Intuitively, one would expect the investor to become more confident in his estimate as more returns are observed. As a result, a given estimate should lead to a higher portfolio weight on stocks, the more data was used to form the estimate. Bayesian analysis has tried to capture the above intuition by incorporating estimation risk into portfolio problems. One might expect the weight on stocks to increase as the posterior variance of the equity premium declines through Bayesian learning. However, the effects of estimation risk on portfolio choice tend to be quantitatively small or even 0 (see, for example, the survey by Brandt, forthcoming). In fact, in a classic continuous time version of the problem, where an investor with log utility learns about the drift in stock prices, there is no trend in stock investment whatsoever: portfolio weights depend only on the posterior mean, not on the variance (Feldman, 1992). The Bayesian investor thus behaves exactly as if the estimate were known to be correct.

Below we revisit the classic log investor problem. We argue that the Bayesian model generates a counterintuitive result because a declining posterior variance does not adequately capture changes in confidence. In contrast, when returns are perceived to be ambiguous, an increase in investor confidence—captured in our model by a shrinking set of posterior beliefs—naturally induces a trend towards stock ownership. We show that this trend is quantitatively significant at plausible parameter values.

5.1. Investor problem

The standard binomial tree model of stock returns maps naturally into the finite state framework of Section 3.2. For calibration purposes, it is convenient to measure time in months. We assume that there are k trading dates in a month. Below, we will obtain simple formulae for optimal portfolio choice by taking the limit $k \rightarrow \infty$, which describes optimal choices when frequent portfolio rebalancing is possible. The state space is $S = \{hi, lo\}$. The return on stocks $R(s_t) = e^{r(s_t)}$ realized at date t is either high or low: we fix (log) return realizations $r(hi) = \sigma/\sqrt{k}$ and $r(lo) = -\sigma/\sqrt{k}$. In addition to stocks, the investor also has access to a riskless asset with constant per period interest rate $R^f = e^{r^f/k}$, where $r^f < \sigma$.

We compute the optimal fraction of wealth invested in stocks for investors who plan for T months starting in month t and who care about terminal wealth according to the utility function $V_T(W_{t+T}) = \log W_{t+T}$. Investors may rebalance their portfolio at all $k(T-t)$ trading dates between t and T . At time t , they know that rebalancing may be optimal in the future—for example, if learning changes beliefs and confidence—and thus take the prospect of future learning and rebalancing into account when choosing a portfolio at t . In general, the optimal portfolio weight on stocks chosen at date t , denoted $\omega_{t,T,k}^*$, will depend on calendar time t , the investment horizon T and on the number of trading dates per month k . In addition, it depends on the history of states that the investor has observed up to date t . Below we will isolate the effect of learning on portfolio choice by comparing optimal choices across agents choosing at different dates t and after different histories, holding fixed T and k .

Consider the investor's problem when beliefs are given by a general process of one-step-ahead conditionals $\{\mathcal{P}_\tau(s^\tau)\}$. Beliefs and the value function are defined for every trading date.

The history s^τ of state realizations up to trading date $\tau = t + j/k$ —the j -th trading date in month $t + 1$ —can be summarized by the fraction ϕ_τ of high states observed up to τ . The value function of the log investor takes the form $V_\tau(W_\tau, s^\tau) = h_\tau(\phi_\tau) + \log W_\tau$. The process $\{h_\tau\}$ satisfies $h_{t+T} = 0$ and

$$\begin{aligned}
 h_\tau(\phi_\tau) &= \max_{\omega_\tau} \min_{p_\tau \in \mathcal{P}_\tau(s^\tau)} E^{p_\tau} [\log(R^f + (R(s_{\tau+1/k}) - R^f)\omega_\tau) + h_{\tau+1/k}(\phi_{\tau+1/k})] \\
 &= \min_{p_\tau \in \mathcal{P}_\tau(s^\tau)} \max_{\omega_\tau} E^{p_\tau} [\log(R^f + (R(s_{\tau+1/k}) - R^f)\omega_\tau) + h_{\tau+1/k}(\phi_{\tau+1/k})], \quad (11)
 \end{aligned}$$

where we have used the minimax theorem to reverse the order of optimization.

For our numerical examples, we capture investors’ information set by a monthly sample of log real U.S. stock returns $\{r_\tau\}_{\tau=1}^t$.¹⁶ Agents in the model observe not only monthly returns, but binary returns $R(s_\tau)$ at every trading date. To model an agent’s history up to date t , we thus construct a sample of tk realizations of $R(s_\tau)$ such that the implied empirical distribution of monthly log returns is the same as that in the data. Let $\sigma = \frac{1}{\sqrt{12}} 19.0\%$, the sample standard deviation of monthly log returns from 1926 to 2004. The sample of binary returns up to some integer date τ is summarized by the fraction of high returns, defined by

$$\phi_\tau = \hat{\phi}_k(\bar{r}_\tau) := \frac{1}{2} + \frac{1}{2} \frac{\bar{r}_\tau}{\sigma \sqrt{k}} \quad (12)$$

where $\bar{r}_\tau := \frac{1}{\tau} \sum_{j=1}^\tau r_j$ is the mean of the monthly return sample. For given k , the sequence $\{\phi_\tau\}$ pins down a sequence of monthly log returns $\left\{ \sum_{j=1}^k \log R(s_{\tau+j/k}) \right\}$ that is identical to the data sample $\{r_\tau\}$. We also assume a constant riskless rate of 2% per year.

5.2. Bayesian benchmark

As a Bayesian benchmark, we assume that the investor has an improper beta prior over the probability p of the high state, so that the posterior mean of p after t months (or tk state realizations) is equal to ϕ_τ , the maximum likelihood estimator of p . The Bayesian’s probability of a high state next period is then also given by ϕ_τ . The optimal portfolio follows from solving (11) when $\mathcal{P}_\tau(s^\tau)$ is a singleton set containing only the measure that puts probability ϕ_τ on the high state. In what follows, we focus on beliefs and portfolios in integer periods t and use (12) to write beliefs directly as a function of k and the monthly sample mean \bar{r}_t .

The objective function on the R.H.S. of (11) now reduces to the sum of two terms: the expected pay-off of a myopic investor plus expected continuation utility. The first term depends on the portfolio weight ω_τ , but not on the investment horizon T , while the converse is true for the second term. As a result, the optimal portfolio weight is independent of the investment horizon. It is given by

$$\omega_{t,k,T}^*(\bar{r}_t) = \frac{e^{r^f/k} (\hat{\phi}_k(\bar{r}_t) (e^{\sigma/\sqrt{k}} - e^{r^f/k}) + (1 - \hat{\phi}_k(\bar{r}_t)) (e^{-\sigma/\sqrt{k}} - e^{r^f/k}))}{(e^{\sigma/\sqrt{k}} - e^{r^f/k}) (e^{r^f/k} - e^{-\sigma/\sqrt{k}})}. \quad (13)$$

As usual in binomial tree set-ups, simple formulae are obtained when taking the limit as k tends to infinity. For the Bayesian case, the supplementary appendix (Epstein and Schneider, 2007b) shows that, for all T ,

$$\lim_{k \rightarrow \infty} \omega_{t,k,T}^*(\bar{r}_t) = \frac{\bar{r}_t + \frac{1}{2} \sigma^2 - r^f}{\sigma^2} := \omega_t^{\text{bay}}. \quad (14)$$

16. We use the monthly NYSE-AMEX return from CRSP, net of monthly CPI inflation for all urban consumers, for 1926:1 to 2004:12.

Here the numerator is the (monthly) equity premium expected by the investor at his current estimate of mean log returns \bar{r}_t . This form of the optimal weight—the equity premium divided by the conditional variance of returns—is standard in continuous time models of portfolio choice with or without learning. In particular, suppose that cumulative stock returns V follow $dV(t)/V(t) = \theta dt + \sigma dB(t)$, where B is a standard Brownian motion. Feldman (1992) shows that the weight (14) is optimal for an investor with log utility who knows the volatility σ and who learns about the constant drift θ starting from an improper normal prior.

The benchmark clarifies the limitation of the Bayesian model in capturing confidence. In the binomial tree, the optimal portfolio weight of the log investor depends only on the estimate ϕ_τ , regardless of how many observations it is based on. Similarly, in the continuous time model considered by Feldman, the optimal portfolio weight depends only on the sample mean \bar{r}_t . This happens despite the fact that estimation risk is gradually reduced as the posterior variance—here σ^2/t —shrinks. Intuitively, estimation risk perceived by a Bayesian is second order relative to the risk in unexpected returns. In the continuous time limit, estimation risk vanishes altogether.¹⁷

Trends in stock investment can occur in the Bayesian case if the investment horizon T is long and the coefficient of relative risk aversion is different from 1.¹⁸ However, the effect of calendar time on the portfolio weight is then due to intertemporal hedging of future news, as emphasized by Brennan (1998). It does *not* occur because the resolution of risk makes risk averse agents more confident. Indeed, Brennan's results show that if the coefficient of relative risk aversion is less than 1, learning induces a *decreasing* trend in the optimal portfolio weight, although the investor is risk averse. Moreover, while high risk aversion can generate an increasing trend, the effect is quantitatively small in our setting, as we show below.

5.3. Beliefs under ambiguity

Beliefs under ambiguity are given by a representation $(\Theta, \mathcal{M}_0, \mathcal{L}_k, \alpha)$. We make explicit the dependence of the set of likelihoods on k because, as in the Bayesian case, simple expressions for the portfolio weight can be obtained by taking the limit for $k \rightarrow \infty$. The ambiguity-averse investor perceives the mean monthly log return as $\theta + \lambda_t$, where $\theta \in \Theta := \mathbb{R}$ is fixed and can be learned, while λ_t is driven by many poorly understood factors affecting returns and can *never* be learnt. The set \mathcal{L}_k consists of all $\ell(\cdot | \theta)$ such that

$$\ell^k(hi | \theta) = \frac{1}{2} + \frac{1}{2} \frac{\theta + \lambda}{\sigma \sqrt{k}}, \quad \text{for some } \lambda \text{ with } |\lambda| < \bar{\lambda}. \quad (15)$$

If $\bar{\lambda} > 0$, returns are ambiguous even conditional on the parameter θ , and λ_t parametrizes the likelihood ℓ_t^k .

The set of priors \mathcal{M}_0 on Θ consists of Dirac measures. For simplicity, we write $\theta \in \mathcal{M}_0$ if the Dirac measure on θ is included in the set of priors, and we define

$$\mathcal{M}_0 = \{\theta : |\theta| \leq \bar{\lambda} + 1/\sigma\}.$$

17. In discrete time, estimation risk would induce a small trend effect. Maintaining a diffuse prior and normality, the one-period-ahead conditional variance of log returns at date t would be $\sigma^2 + \sigma^2/t$. With a period length of one month, the effect of estimation risk becomes a small fraction of the overall perceived variance after only a couple of years.

18. In continuous time, it follows from results in Rogers (2001) that the optimal weight on stocks for a Bayesian investor with power utility and coefficient of relative risk aversion $\gamma > 1$ is given by

$$\omega_t = \frac{\bar{r}_t + \frac{1}{2}\sigma^2 - r}{\gamma\sigma^2} \frac{\gamma t}{\gamma t + (\gamma - 1)T}.$$

The first fraction is the “myopic demand” that obtains for $T = 0$; it depends only on the posterior mean. For $T > 0$, the second fraction captures inter-temporal hedging. For positive equity premia, the weight increases relative to the myopic demand with calendar time t .

The condition ensures that the probability (15) remains between 0 and 1 for all $k \geq 1$. Standard results imply that if the state is i.i.d. with fixed θ and λ , then the cumulative stock return process for large k approximates a diffusion with volatility σ and drift $\theta + \lambda$.

Beliefs depend on history only via the fraction of high returns $\phi_t = \hat{\phi}_k(\bar{r}_t)$. We focus again on integer periods t and write beliefs and optimal portfolios directly as functions of k and \bar{r}_t . In particular, we write $\theta \in \mathcal{M}_{t,k}^\alpha(\bar{r}_t)$ if the Dirac measure on θ is included in the posterior set at the end of month t after history $\hat{\phi}_k(\bar{r}_t)$. The supplementary appendix shows that the posterior set is an interval $\mathcal{M}_{t,k}^\alpha(\bar{r}_t) := [\underline{\theta}_k(\bar{r}_t), \bar{\theta}_k(\bar{r}_t)]$, with both bounds strictly increasing in \bar{r}_t , and that

$$\left[\lim_{k \rightarrow \infty} \underline{\theta}_k(\bar{r}_t), \lim_{k \rightarrow \infty} \bar{\theta}_k(\bar{r}_t) \right] = \left[\bar{r}_t - t^{-(1/2)} \sigma b_\alpha, \bar{r}_t + t^{-(1/2)} \sigma b_\alpha \right], \quad (16)$$

where $b_\alpha = \sqrt{-2 \log \alpha}$. The limiting posterior set is thus an interval centred around the sample mean \bar{r}_t that is wider the larger is the standard error $t^{-(1/2)} \sigma$ and the smaller is the parameter α ; in the Bayesian case, $\alpha = 1$ and $b_\alpha = 0$.

The set of one-step-ahead beliefs $\mathcal{P}_{t,k}(\bar{r}_t)$ —defined as in (4)—contains all likelihoods of the type (15) for some $\theta \in \mathcal{M}_{t,k}^\alpha(\bar{r}_t)$ and $|\lambda| < \bar{\lambda}$. It can thus be summarized by an interval of probabilities for the high state next period, with bounds that are strictly increasing in \bar{r}_t because of (16). Consider now the lowest and highest conditional mean log returns per month. In the limit as $k \rightarrow \infty$,

$$\left[\lim_{k \rightarrow \infty} (\underline{\theta}_k(\bar{r}_t) - \bar{\lambda}), \lim_{k \rightarrow \infty} (\bar{\theta}_k(\bar{r}_t) + \bar{\lambda}) \right] = \left[\bar{r}_t - \bar{\lambda} - t^{-(1/2)} \sigma b_\alpha, \bar{r}_t + \bar{\lambda} + t^{-(1/2)} \sigma b_\alpha \right]. \quad (17)$$

In the long run as t becomes large, this interval shrinks towards $[\bar{r}_t - \bar{\lambda}, \bar{r}_t + \bar{\lambda}]$ so that the equity premium is eventually perceived to lie in an interval of width $2\bar{\lambda}$.

The explicit form of the posterior set and limiting beliefs suggests an intuitive way to select parameters $\bar{\lambda}$ and α for our numerical examples. The parameter $\bar{\lambda}$ determines what the investor expects to learn in the long run. The case $\bar{\lambda} = 0$ describes an investor who is sure today that she will eventually know what the equity premium is. More generally, the investor might believe that even far in the future there will be poorly understood factors that preclude full confidence in any estimate of the equity premium. Her lack of confidence can be quantified by the range of possible equity premia in the long run, which has width $2\bar{\lambda}$. In the calculations below, we consider both $\bar{\lambda} = 0$ and $\bar{\lambda} = 0.001$; the latter corresponds to a 2.4% range for the (annualized) expected equity premium.

Second, the parameter α that determines the investor's speed of updating can be interpreted in terms of classical statistics. The posterior set (16) contains all parameters θ^* such that the hypothesis $\theta = \theta^*$ is not rejected by an asymptotic likelihood-ratio test where the critical value of the $\chi^2(1)$ distribution is $b_\alpha^2 = -2 \log \alpha$. Every value of α thus corresponds to a significance level at which a theory must be rejected by the likelihood-ratio test in order to be excluded from the belief set. We consider significance levels of 5% and 2.5%, which leads to $\alpha = 0.15$ or ($b_\alpha = 1.96$) and $\alpha = 0.08$ (or $b_\alpha = 2.23$), respectively. The lower the significance level, the fewer theories are rejected for any given sample; in other words, confidence in understanding the environment is lower. If $\alpha = 1$ and $\bar{\lambda} = 0$, the posterior set contains only \bar{r}_t and $\mathcal{P}_t(s^t)$ contains only the Bayesian benchmark belief.

5.4. Optimal portfolios

Learning under ambiguity affects portfolio choice in two ways. The first is a direct effect of beliefs and confidence on the optimal weight, which is relevant at any investment horizon. Consider

a myopic investor ($T = 1/k$). The supplementary appendix shows that the limit of her optimal weight on stocks as k becomes large, $\lim_{k \rightarrow \infty} \omega_{t,k,1/k}^*(\bar{r}_t)$, is given by

$$\sigma^{-2} \max \left\{ \bar{r}_t + \frac{1}{2} \sigma^2 - r^f - (\bar{\lambda} + t^{-(1/2)} \sigma b_\alpha), 0 \right\} + \sigma^{-2} \min \left\{ \bar{r}_t + \frac{1}{2} \sigma^2 - r^f + \bar{\lambda} + t^{-(1/2)} \sigma b_\alpha, 0 \right\},$$

or

$$\lim_{k \rightarrow \infty} \omega_{t,k,1/k}^*(\bar{r}_t) = \max \left\{ \omega_t^{\text{bay}} - \sigma^{-2} \left(\bar{\lambda} + t^{-(1/2)} \sigma b_\alpha \right), 0 \right\} \quad (18)$$

$$+ \min \left\{ \omega_t^{\text{bay}} + \sigma^{-2} \left(\bar{\lambda} + t^{-(1/2)} \sigma b_\alpha \right), 0 \right\}.$$

A myopic ambiguity-averse investor goes long in stocks only if the equity premium is unambiguously positive. The optimal position is then given by the first term in (18). Comparison with (14) shows that an ambiguity-averse investor who goes long behaves *as if* he were a Bayesian whose perceives the lowest conditional mean log return in (17). The optimal weight depends on the sample through the Bayesian position ω_t^{bay} : conditional on participation, the optimal response to news is therefore the same in the two models.

However, ambiguity also introduces a trend component: conditional on participation, the optimal position increases as confidence grows and the standard error $t^{-\frac{1}{2}} \sigma$ falls. While it moves closer to the Bayesian position ω_t^{bay} in the long run, the optimal position remains forever smaller when there are multiple likelihoods ($\bar{\lambda} > 0$). The second term in (18) reflects short selling when the equity premium is unambiguously negative. Non-participation is optimal if the maximum likelihood (ML) estimate of the equity premium is small in absolute value so that both terms are 0, that is, if $|\bar{r}_t + \frac{1}{2} \sigma^2 - r^f| < \bar{\lambda} + t^{-(1/2)} \sigma b_\alpha$. In particular, an investor who is not confident that the equity premium can be learnt ($\bar{\lambda} > 0$) need not be “in the market” even after having seen a large amount of data.

The second effect of learning under ambiguity is a new inter-temporal hedging motive that induces more participation as the investment horizon becomes longer. Figure 3 compares myopic and long horizon investors by plotting optimal positions against the maximum likelihood estimate of the (annualized) equity premium $(\bar{r}_t + \frac{1}{2} \sigma^2 - r^f) \times 12$.¹⁹ We consider two long horizon investors, both with an investment horizon of $T = 10$ years: the dashed line is for an investor who has seen $t = 35$ years of data, while the dotted line is for an investor who has seen $t = 5$ years of data. Thick grey lines indicate optimal myopic positions, which are zero whenever they do not coincide with the long horizon optimal position. The solid line shows a Bayesian log investor, whose position depends neither on experience nor on the investment horizon.

For multiple-priors investors, the myopic and long horizon positions coincide if the equity premium is either high or low enough so that the myopic investor participates. However, for intermediate estimates of the premium, the myopic investor stays out of the stock market, while the long horizon investor takes “contrarian” positions: she goes short in stocks for positive equity premia, but long for negative premia. For the relatively experienced investor (dashed line), who has seen 35 years of data, horizon effects are small. However, for an inexperienced investor (dotted line), who has a wide non-participation region if she is myopic, there can be sizeable differences between the optimal myopic and the optimal long horizon weights. For example, if the first five years of data deliver a very bad sample equity premium of -15% , the long horizon investor does not stay out, but instead invests close to 30% of wealth in stocks.

Intuitively, contrarian behaviour can be optimal for long horizon investors because it provides a hedge against adverse outcomes of future learning under ambiguity. To see this, consider

19. Closed form solutions for optimal weights are not available with ambiguity aversion if the investor is not myopic. However, it is straightforward to compute the solution numerically using backward induction on (11). All calculations reported below assume $k = 30$.

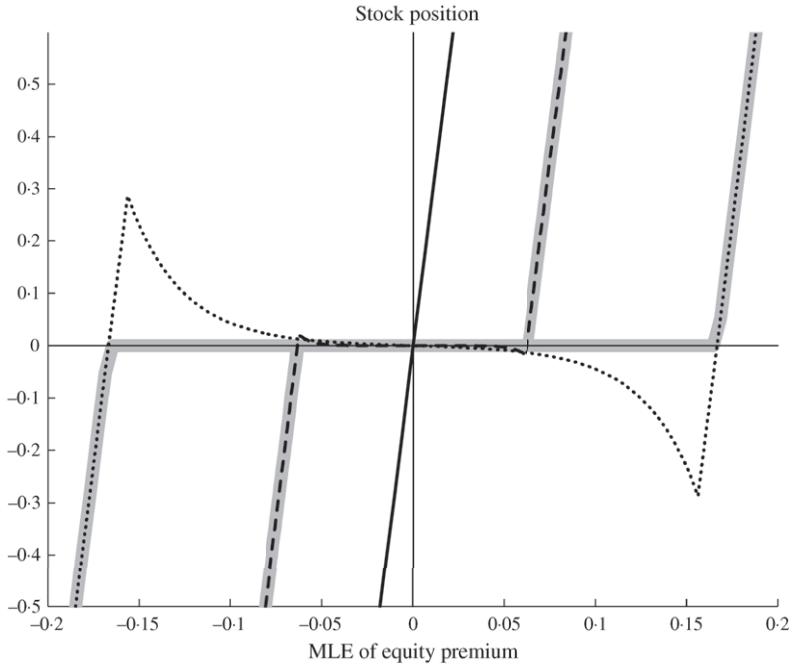


FIGURE 3

Optimal weight on stocks for log investors as a function of ML estimate of equity premium. *Dashed black line*: investor with $\alpha = 0.15$, $\lambda = 0$, sample size $t = 5$ years and horizon $T = 10$ years. *Dotted black line*: investor with the same belief and horizon, but sample size $t = 35$ years. Thick grey lines indicate the corresponding myopic positions ($T = 1/k$), which are 0 when they do not coincide with the positions for $T = 10$ years. *Solid black line*: Bayesian investor ($\alpha = 1$, $\lambda = 0$). All numerical calculations are for $k = 30$

first how returns affect the continuation utility of myopic and long horizon investors differently. To a myopic investor, stock returns over the next trading period matter only because they affect future wealth. In contrast, a learning non-myopic investor perceives returns also as a signal about the long-run equity premium. For example, high returns suggest that the long-run equity premium is higher than previously expected. The Bellman equation (11) illustrates how returns matter to a non-myopic investor in part because they affect future learning: continuation utility depends on $\phi_{t+1/k}$.

In particular, our numerical results show that continuation utility is U-shaped in ϕ , or, equivalently, in the sample equity premium: multiple-priors investors are better off, the further away the sample equity premium is from zero. This is because they believe that they can make a profit in the stock market only if the equity premium is unambiguously different from zero, either on the positive or on the negative side. Suppose now that the ML estimate of the equity premium is negative. A myopic investor will not take a long position, since the worst case mean return is negative. While this is also of concern to the non-myopic investor, the latter benefits in addition from the fact that the long position pays off precisely when continuation utility is low. Indeed, a long position pays off more when returns are high and thus when the estimate of the long-run equity premium moves closer to zero (from its initial negative value). In other words, the position provides “insurance against bad investment opportunities”. If the mean equity premium is not too negative, this hedging motive dominates and leads to contrarian investment, as in Figure 3.

The inter-temporal hedging motive derived here is reminiscent of hedging demand under risk stressed by Merton (1971), and, in the context of Bayesian learning, by Brennan (1998).

Nevertheless, it is unique to the case of ambiguity. Indeed, hedging demand under risk is zero for a log investor: separability of the objective in (11) implies that the optimal portfolio does not depend on the continuation value $h_{\tau+1/k}$. In contrast, with multiple conditionals, the minimizing probability $p_t^*(\phi_t)$ will in general depend on $h_{\tau+1/k}$ and hence on the investment horizon. The presence of ambiguity thus breaks the separability and generates hedging demand.

5.5. Confidence and participation over time

To illustrate the quantitative relevance of changes in confidence, Figure 4 compares optimal stock weights for myopic investors over the period 1960–2004. The top panel shows the position of the Bayesian log investor (grey, left scale). This position is highly levered and fluctuates around the sample mean of 176%. There is no upward trend—the Bayesian model actually predicts higher stock investment in the 1960’s than at the peak of the 1990’s boom. The black line in the top panel is the ambiguity-averse investor with $\bar{\lambda} = 0$ and $\alpha = 0.15$, which corresponds to a 5% significance level of the likelihood ratio test. Since this position is considerably smaller on average, a second (shifted) scale is marked on the right hand axis. While the fluctuations in the optimal weight are similar in magnitude to the Bayesian case—as predicted by (18)—the ambiguity-averse position trends up over time. In particular, it is higher in the 1990’s than in the 1960’s.

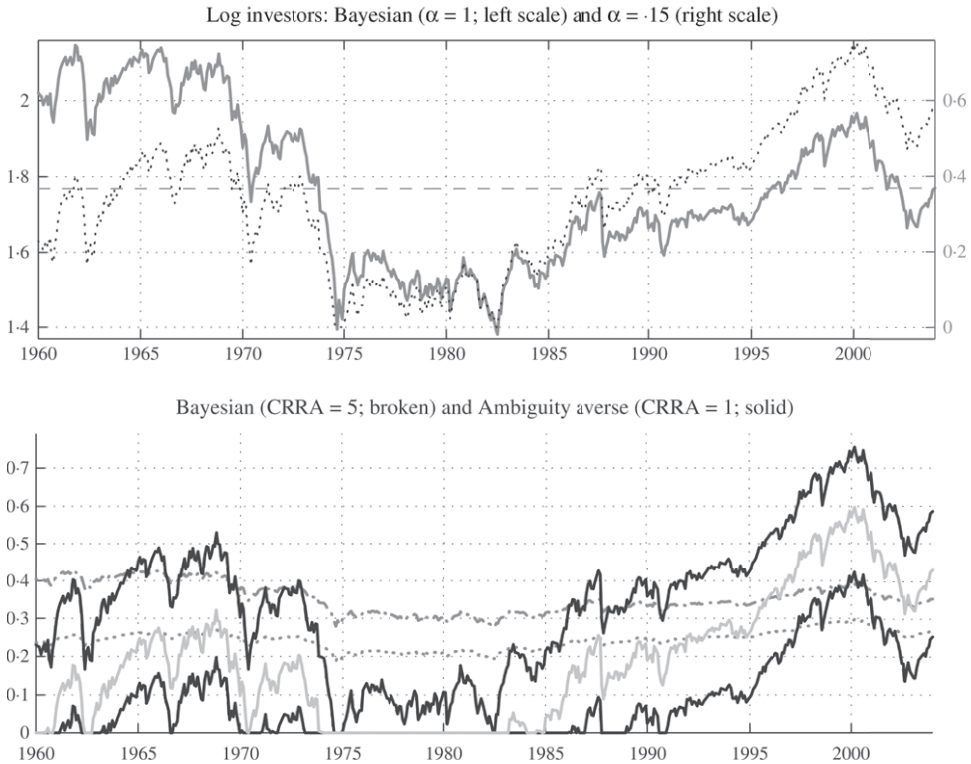


FIGURE 4

Optimal stock positions, 1960–2004, for investors who learn from monthly real NYSE stock returns since 1926. *Top panel:* log investor with $\alpha = 1$ (solid grey; left scale), and log investor with $\alpha = 0.15$ (dotted black; right scale). *Bottom panel:* broken lines are for Bayesian investors with coefficient of relative risk aversion (CRRA) = 5; $T = 1/k$ (dash-dot grey) and $T = 30$ years (dotted grey); solid lines are for ambiguity-averse log investors; $\alpha = 0.15$, $\bar{\lambda} = 0$ (top solid black line), $\alpha = 0.08$, $\bar{\lambda} = 0$ (solid grey line) and $\alpha = 0.15$, $\bar{\lambda} = 0.001$ (bottom solid black line)

The bottom panel of Figure 4 shows what happens as investors become more risk averse or more ambiguity averse. The upper black line replicates the case $\bar{\lambda} = 0$ and $\alpha = 0.15$ of the upper panel. The grey line corresponds to $\bar{\lambda} = 0$ and $\alpha = 0.08$ (a significance level of 2.5%), an investor whose confidence is not only lower, but also growing more slowly. The latter feature explains why the difference between the two lines is narrower in the 1990's than in the 1960's. In contrast, introducing multiple likelihoods shifts down the position in a parallel fashion: an example is the lower black line ($\alpha = 0.15, \bar{\lambda} = 0.001$). The latter two investors also go through lengthy periods of optimal non-participation. For comparison, the lower panel also reports optimal positions for Bayesian investors with a coefficient of relative risk aversion of 5. The solid grey line is the myopic Bayesian, while the dashed grey line assumes a horizon of $T = 30$ years. The point here is that once risk aversion is high enough for the Bayesian model to generate sensible portfolio allocations, these positions do not vary much over time. Moreover, even with a very long investment horizon, the trend in the position induced by inter-temporal hedging is small.

It is interesting to ask whether the positions of a multiple-priors investor could be replicated by a Bayesian with an informative prior. For any two points in the figure, this is always possible: we can reverse engineer a mean and variance for the prior normal distribution to match the posterior mean, and hence the optimal positions in equation (18), at any two calendar dates. For example, consider the investor with $\alpha = 0.15$ and $\bar{\lambda} = 0$, whose optimal weight on stocks grows from $\omega_t = 5.8\%$ in December 1976 to $\omega_t = 59.3\%$ in December 2004. A Bayesian with normal prior who chooses the same positions must start in 1926 from a prior mean of -6.41% p.a. and a prior standard deviation of 0.88%. In other words, she must have been essentially convinced that the equity premium is always negative. Since such extreme and dogmatic pessimism is needed to match the multiple-priors investor in 1976 and 2004, the optimal position is negative in all years before 1963. This is in sharp contrast to the multiple-priors investor, who invests more than 20% in stocks in the early 1960's.

APPENDIX

Comparison of predictive probabilities from Section 3.3.

We want to show that the minimal predictive probability for a black ball conditional on observing a black ball,

$$\min_P \mathcal{P}_1(B) = \min_{\ell \in \mathcal{L}, \mu_1 \in \mathcal{M}_1^\alpha(B)} \int_{\Theta} \ell(B | \theta) d\mu_1(\theta), \tag{A.1}$$

is smaller under Scenario 3 than under Scenario 2.

Consider first Scenario 2. The set \mathcal{M}_1^α of posteriors (after having observed one black draw) is constructed by updating only those priors μ_0^p according to which the black draw was sufficiently likely *ex ante* in the sense of the likelihood-ratio criterion in (6). The *ex ante* probability of a black draw under the prior μ_0^p is (using (7))

$$\int_{\Theta} \theta d\mu_0^p = \frac{1}{10} \left[2 \sum_{\lambda=1}^3 \lambda p(\lambda) + 1 \right]. \tag{A.2}$$

It follows that the agent retains only those priors μ_0^p corresponding to some p in $P^\alpha \subset P$, where

$$P^\alpha = \left\{ p \in P : \sum \lambda p(\lambda) + \frac{1}{2}(1 - \alpha) \geq \alpha \max_{p' \in P} \sum \lambda p'(\lambda) \right\}.$$

(Because all measures in P have support in $\{1, 2, 3\}$, it is easily seen that $P^\alpha = P$ for all $\alpha \leq \frac{3}{4}$. At the other extreme, if $\alpha = 1$, then P^α consists only of those measures in P that maximize $\sum \lambda p(\lambda)$, the expected number of black non-coin balls.) Thus \mathcal{M}_1^α consists of all measures μ_1^p , where for some $p \in P^\alpha$,

$$\mu_1^p(\theta) = \frac{\theta}{\int_{\Theta} \theta' d\mu_0} \mu_0^p(\theta) = \frac{\theta \left(\frac{1}{2} p(5\theta - 1) + \frac{1}{2} p(5\theta) \right)}{\frac{1}{10} \left[2 \sum \lambda p(\lambda) + 1 \right]}.$$

The set of predictive probabilities of a black ball on the next draw after having observed one black draw is given by

$$\mathcal{P}_1(B) = \left\{ \int_{\Theta} \theta d\mu_1^p : p \in P^\alpha \right\}.$$

Compute that

$$\begin{aligned} \int_{\Theta} \theta d\mu_1^p &= \frac{\sum_{\theta} \left[\theta^2 \left(\frac{1}{2} p(5\theta - 1) + \frac{1}{2} p(5\theta) \right) \right]}{\frac{1}{10} [2 \sum \lambda p(\lambda) + 1]} \\ &= \frac{1}{5} \frac{2 \sum \lambda^2 p(\lambda) + 2 \sum \lambda p(\lambda) + 1}{2 \sum \lambda p(\lambda) + 1} \geq \frac{1}{5} \frac{2 (\sum \lambda p(\lambda))^2 + 2 \sum \lambda p(\lambda) + 1}{2 \sum \lambda p(\lambda) + 1} \\ &= \frac{1}{5} \frac{\sum \lambda p(\lambda) + 1}{2 \sum \lambda p(\lambda) + 1} + \frac{1}{5} \sum \lambda p(\lambda), \end{aligned}$$

that is,

$$\int_{\Theta} \theta d\mu_1^p \geq \frac{1}{5} \frac{\sum \lambda p(\lambda) + 1}{2 \sum \lambda p(\lambda) + 1} + \frac{1}{5} \sum \lambda p(\lambda), \quad \text{for every } p \text{ in } P^\alpha.$$

Therefore, the minimum predictive probability satisfies

$$\min \left\{ \int_{\Theta} \theta d\mu_1^p : p \in P^\alpha \right\} \geq \frac{1}{5} \left[\min_{p \in P^\alpha} \frac{\sum \lambda p(\lambda) + 1}{2 \sum \lambda p(\lambda) + 1} + \min_{p \in P^\alpha} \sum \lambda p(\lambda) \right]. \tag{A.3}$$

This inequality is used below to draw comparisons with Scenario 3.

Turn now to Scenario 3. Note that

$$\ell^p(B | B) - \ell(B | W) = \frac{1}{5} \quad \text{for every } \ell^p.$$

After seeing one black draw, the agent reevaluates which likelihoods in \mathcal{L} could have been operative in the first draw. He does this by computing the *ex ante* probability of a black draw associated with each ℓ^p (and the fixed μ_0)—that probability is given by

$$\begin{aligned} \int_{\theta \in \{B, W\}} \ell^p(B | \theta) d\mu_0 &= \frac{1}{2} (\ell^p(B | B) + \ell^p(B | W)) \\ &= \frac{1}{10} \left(2 \sum \lambda p(\lambda) + 1 \right), \end{aligned}$$

which is identical to the *ex ante* probability computed in Scenario 2 (see (A.2)). It follows that the likelihood-ratio criterion leads to retention of only those likelihoods ℓ^p where p lies in P^α , the same set defined in the discussion of Scenario 2.

The set \mathcal{M}_1^α of posteriors (after having observed one black) consists of the measures

$$\mu_1^p(\theta) = \frac{\ell^p(B | \theta)}{\ell^p(B | B) + \ell^p(B | W)}, \quad \theta \in \{B, W\}, \text{ for some } p \text{ in } P^\alpha.$$

In particular,

$$\mu_1^p(B) = \frac{\ell^p(B | B)}{\ell^p(B | B) + \ell^p(B | W)} = \frac{\sum \lambda p(\lambda) + 1}{2 \sum \lambda p(\lambda) + 1}.$$

Finally, the set $\mathcal{P}_1(B)$ of predictive probabilities of a black on the next draw, after having observed one black draw, is

$$\mathcal{P}_1(B) = \left\{ \int_{\theta \in \{B, W\}} \ell(B | \theta) d\mu_1^p : \ell \in \mathcal{L}, p \in P^\alpha \right\}.$$

Therefore the minimum predictive probability equals

$$\begin{aligned} \min_{p \in P^\alpha, \ell \in \mathcal{L}} \left\{ \mu_1^p(B) [\ell(B|B) - \ell(B|W)] + \ell(B|W) \right\} &= \min_{p \in P^\alpha, \ell \in \mathcal{L}} \left\{ \mu_1^p(B) \frac{1}{5} + \ell(B|W) \right\} \\ &= \frac{1}{5} \min_{p \in P^\alpha} \mu_1^p(B) + \min_{\ell \in \mathcal{L}} \ell(B|W) = \frac{1}{5} \min_{p \in P^\alpha} \mu_1^p(B) + \min_{p' \in P} \frac{\sum_1^3 \lambda p'(\lambda)}{5} \\ &\leq \frac{1}{5} \min_{p \in P^\alpha} \mu_1^p(B) + \min_{p' \in P^\alpha} \frac{\sum_1^3 \lambda p'(\lambda)}{5} \quad (\text{replacing } P \text{ by its subset } P^\alpha) \\ &= \frac{1}{5} \min_{p \in P^\alpha} \frac{\sum \lambda p(\lambda) + 1}{2 \sum \lambda p(\lambda) + 1} + \min_{p' \in P^\alpha} \frac{\sum \lambda p'(\lambda)}{5}, \end{aligned}$$

which is no larger than the minimum predictive probability computed for Scenario 2 (see (A.3)). This proves that betting in Scenario 2 is preferable.

It is not difficult to see that there is indifference between the two scenarios iff there exists $\lambda^* \in \{1, 2, 3\}$ such that $P = \{\delta_{\lambda^*}\}$, that is, there is certainty that there are λ^* black non-coin balls. In that case, the set of likelihoods proposed for Scenario 3 collapses to a singleton. ||

Proof of Theorem 1. For any sequence $s^\infty = (s_1, s_2, \dots)$, denote by ϕ_t the empirical measure on S corresponding to the first t observations. We focus on the set Ω of sequences for which $\phi_t \rightarrow \phi$; this set has measure 1 under the truth. Fix a sequence $s^\infty \in \Omega$.

As in the text, use the notation $(\ell_s)_{s \in S}$ for likelihood sequences (ℓ_t) such that $\ell_j = \ell_k$ if $s_j = s_k$. Suppose that the relative frequencies of the states in a (finite or infinite) sample are represented by a probability measure λ on S —below we will take λ to be ϕ_t or ϕ . Given a likelihood sequence $(\ell_s)_{s \in S}$, the likelihood of the sample conditional on the parameter value θ is then

$$\tilde{H}(\lambda, (\ell_s), \theta) = \sum_{s \in S} \lambda(s) \log \ell_s(s | \theta).$$

We want to show that, for every sequence of posteriors $\{\mu_t\}$ with $\mu_t \in \mathcal{M}_t^\alpha$ for all t , the posterior probability of the parameter value θ^* converges (uniformly) to 1. By the definition of \mathcal{M}_t^α , for every posterior μ_t there exists an *admissible theory* $(\mu_0^{(t)}, \ell^{t,(t)})$, that is, a theory satisfying the conditions in (6), such that μ_t is the Bayesian update of $\mu_0^{(t)}$ along the sequence of likelihoods $\ell^{t,(t)} = (\ell_1^{(t)}, \ell_2^{(t)}, \dots, \ell_t^{(t)})$. Here the t 's in brackets indicate the place of the posterior in the given sequence of posteriors $\{\mu_t\}$ —they are needed to account for the fact that not all μ_t in the sequence $\{\mu_t\}$ need to be updates of the same initial prior μ_0 , or be updated along the same likelihood sequence. (As well, the sequence of updates of a given μ_0 along a given sequence of likelihoods need not be admissible at all dates.) Thus our objective is to show that

$$\lim_{t \rightarrow \infty} \mu_t(\theta^* | s^t, \mu_0^{(t)}, \ell^{t,(t)}) = 1, \tag{A.4}$$

uniformly in the sequences of admissible theories $(\mu_0^{(t)}, \ell^{t,(t)})$.

Given any likelihood tuple ℓ^t and any parameter value θ , define

$$\eta_t(\theta, \ell^t) := \frac{1}{t} \sum_{j=1}^t \log \ell_j^t(s_j | \theta) - \max_{\ell_s} \tilde{H}(\phi_t, (\ell_s), \theta^*).$$

Here η_t is the log likelihood ratio between the likelihood of the sample under ℓ^t at the parameter value θ and the likelihood of the sample under the sequence (ℓ_s) that maximizes the likelihood of the sample given the parameter θ^* . A posterior μ_t derived from a theory $(\mu_0^{(t)}, \ell^{t,(t)})$ can be written as

$$\begin{aligned} \mu_t(\theta^* | s^t, \mu_0^{(t)}, \ell^{t,(t)}) &= \frac{\mu_0^{(t)}(\theta^*) \prod_{j=1}^t \ell_j^{(t)}(s_j | \theta^*)}{\sum_{\theta \in \Theta} \mu_0^{(t)}(\theta) \prod_{j=1}^t \ell_j^{(t)}(s_j | \theta)} = \frac{\mu_0^{(t)}(\theta^*) e^{t \eta_t(\theta^*, \ell^{t,(t)})}}{\sum_{\theta \in \Theta} \mu_0^{(t)}(\theta) e^{t \eta_t(\theta, \ell^{t,(t)})}} \\ &= \mu_0^{(t)}(\theta^*) \left(\mu_0^{(t)}(\theta^*) + \sum_{\theta \neq \theta^*} \mu_0^{(t)}(\theta) e^{t(\eta_t(\theta, \ell^{t,(t)}) - \eta_t(\theta^*, \ell^{t,(t)})} \right)^{-1}. \end{aligned}$$

Claim 1. For every $\varepsilon > 0$ there exists T such that $\eta_t(\theta, \ell^{t,(t)}) \leq -\varepsilon$, for all $t > T(\varepsilon)$, $\theta \neq \theta^*$ and for all sequences of likelihood tuples $(\ell^{t,(t)})$.

Claim 2. $\eta_t(\theta^*, \ell^{t,(t)}) \rightarrow 0$ uniformly in admissible theories $(\mu_0^{(t)}, \ell^{t,(t)})$.

Claim 3. There exists T such that $\mu_0^{(t)}(\theta^*) > 0$ for all $t > T$ and all admissible theories $(\mu_0^{(t)}, \ell^{t,(t)})$.

Claims 1 and 2 together imply that

$$\sum_{\theta \neq \theta^*} \mu_0^{(t)}(\theta) e^{t(\eta_t(\theta, \ell^{t,(t)}) - \eta_t(\theta^*, \ell^{t,(t)}))} \rightarrow 0$$

uniformly. Claim 3 and hypothesis (iv) imply that $\mu_0^{(t)}(\theta^*) > \kappa$ for large enough t , and (A.4) follows.

Proof of Claim 1. If ℓ^t is to be chosen to maximize $\frac{1}{t} \sum_{j=1}^t \log \ell_j(s_j | \theta)$, it is wlog to focus on sequences such that $\ell_j = \ell_k$ if $s_j = s_k$. Therefore, any likelihood tuple $\ell^{t,(t)}$ in the sequence satisfies

$$\frac{1}{t} \sum_{j=1}^t \log \ell_j^{(t)}(s_j | \theta) \leq \max_{(\ell_s)} \tilde{H}(\phi_t, (\ell_s), \theta). \tag{A.5}$$

By definition of H and the identification hypothesis (ii), there exists $\varepsilon > 0$ such that

$$\max_{(\ell_s)} \tilde{H}(\phi, (\ell_s), \theta) \leq H(\theta^*) - 2\varepsilon, \quad \text{for all } \theta \neq \theta^*.$$

Thus the Maximum Theorem implies that, for some sufficiently large T ,

$$\max_{(\ell_s)} \tilde{H}(\phi_t, (\ell_s), \theta) \leq \max_{(\ell_s)} \tilde{H}(\phi_t, (\ell_s), \theta^*) - \varepsilon, \tag{A.6}$$

for all $\theta \neq \theta^*$ and $t > T$. The claim now follows from (A.5).

Proof of Claim 2. By definition,

$$\eta_t(\theta^*, \ell^t) = \frac{1}{t} \sum_{j=1}^t \log \ell_j^t(s_j | \theta^*) - \max_{(\ell_s)} \sum_{s \in S} \phi_t(s) \log \ell_s(s | \theta^*).$$

By the definition of \mathcal{M}_t^α , every element of a sequence of admissible theories $(\mu_0^{(t)}, \ell^{t,(t)})$ satisfies

$$\left(\max_{\hat{\mu}_0, \hat{\ell}^t} \Pr(s^t; \hat{\mu}_0, \hat{\ell}^t) \right)^{1/t} \geq \left(\Pr(s^t; \mu_0^{(t)}, \ell^{t,(t)}) \right)^{1/t} \geq \alpha^{1/t} \left(\max_{\hat{\mu}_0, \hat{\ell}^t} \Pr(s^t; \hat{\mu}_0, \hat{\ell}^t) \right)^{1/t}, \tag{A.7}$$

where $\Pr(s^t; \mu_0^{(t)}, \ell^{t,(t)}) = \int \prod_{j=1}^t \ell_j^{(t)}(s_j | \theta) d\mu_0^{(t)}(\theta)$; $\Pr(s^t; \hat{\mu}_0, \hat{\ell}^t)$ is defined similarly. Consider the long-run behaviour of $\max_{\hat{\mu}_0, \hat{\ell}^t} \Pr(s^t; \hat{\mu}_0, \hat{\ell}^t)$, the data density under the maximum likelihood theory. We claim that

$$\left[\frac{\max_{\hat{\mu}_0, \hat{\ell}^t} \Pr(s^t; \hat{\mu}_0, \hat{\ell}^t)}{e^{t \max_{(\ell_s)} \tilde{H}(\phi_t, (\ell_s), \theta^*)}} \right]^{1/t} \rightarrow 1, \tag{A.8}$$

or equivalently,

$$\left[\max_{\hat{\mu}_0, \hat{\ell}^t} \sum_{\theta \in \Theta} \hat{\mu}_0(\theta) e^{t \eta_t(\theta, \hat{\ell}^t)} \right]^{1/t} \rightarrow 1. \tag{A.9}$$

Claim 1 implies that (for all $t > T$)

$$\max_{\hat{\mu}_0, \hat{\ell}^t} \hat{\mu}_0(\theta^*) e^{t \eta_t(\theta^*, \hat{\ell}^t)}$$

$$\begin{aligned} &\leq \max_{\hat{\mu}_0, \hat{\ell}^t} \sum_{\theta \in \Theta} \hat{\mu}_0(\theta) e^{t\eta_t(\theta, \hat{\ell}^t)} \\ &\leq \max_{\hat{\mu}_0, \hat{\ell}^t} \hat{\mu}_0(\theta^*) e^{t\eta_t(\theta^*, \hat{\ell}^t)} + (1 - \hat{\mu}_0(\theta^*)) e^{-\varepsilon t}. \end{aligned}$$

Hypothesis (iii) says that there is a prior μ_0 such that $\mu_0(\theta^*) > 0$. Verify that

$$\eta_t(\theta^*, \hat{\ell}^t) \leq 0 \text{ and } \max_{\hat{\ell}^t} \eta_t(\theta^*, \hat{\ell}^t) = 0. \quad (\text{A.10})$$

Then, by hypothesis (iv),

$$\kappa^{1/t} = \kappa^{1/t} \max_{\hat{\ell}^t} e^{\eta_t(\theta^*, \hat{\ell}^t)} \leq \left[\max_{\hat{\mu}_0, \hat{\ell}^t} \hat{\mu}_0(\theta^*) e^{t\eta_t(\theta^*, \hat{\ell}^t)} \right]^{1/t} \leq \left[\max_{\hat{\mu}_0} \hat{\mu}_0(\theta^*) \right]^{1/t} \leq 1,$$

which implies that

$$\left[\max_{\hat{\mu}_0, \hat{\ell}^t} \hat{\mu}_0(\theta^*) e^{t\eta_t(\theta^*, \hat{\ell}^t)} \right]^{1/t} \rightarrow 1,$$

and hence also (A.9).

Combining (A.8) and (A.7), it now follows that, for any sequence of admissible theories $(\mu_0^{(t)}, \ell^{t,(t)})$,

$$\left(\sum_{\theta \in \Theta} \mu_0^{(t)}(\theta) e^{t\eta_t(\theta, \ell^{t,(t)})} \right)^{1/t} = \left[\frac{\Pr(s^t; \mu_0^{(t)}, \ell^{t,(t)})}{e^{t \max_{\ell_s} \widetilde{H}(\phi_t, (\ell_s), \theta^*)}} \right]^{1/t} \rightarrow 1, \quad (\text{A.11})$$

and that the convergence is uniform.

Finally, $\eta_t(\theta^*, \ell^{t,(t)}) \leq 0$ by construction—see (A.10). Suppose that $\eta_t(\theta^*, \ell^{t,(t)}) < -\delta$ for some positive δ and infinitely many t . Then Claim 1 (with $\varepsilon = \delta$) implies that

$$\left[\frac{\Pr(s^t; \mu_0^{(t)}, \ell^{t,(t)})}{e^{t \max_{\ell_s} \widetilde{H}(\phi_t, (\ell_s), \theta^*)}} \right]^{1/t} = \left(\sum_{\theta \in \Theta} \mu_0^{(t)}(\theta) e^{t\eta_t(\theta, \ell^{t,(t)})} \right)^{1/t} \leq e^{-\delta} < 1,$$

for infinitely many t , contradicting (A.11). This proves Claim 2.

Proof of Claim 3. Suppose that for infinitely many t there exists an admissible theory $(\mu_0^{(t)}, \ell^{t,(t)})$ such that $\mu_0^{(t)}(\theta^*) = 0$. Claim 1 (with $\varepsilon = \delta$) now implies that, for infinitely many t ,

$$\left[\frac{\Pr(s^t; \mu_0^{(t)}, \ell^{t,(t)})}{e^{t \max_{\ell_s} \widetilde{H}(\phi_t, (\ell_s), \theta^*)}} \right]^{1/t} = \left(\sum_{\theta \neq \theta^*} \mu_0^{(t)}(\theta) e^{t\eta_t(\theta, \ell^{t,(t)})} \right)^{1/t} \leq e^{-\delta} < 1,$$

which contradicts (A.11). \parallel

Acknowledgements. Epstein acknowledges financial support from NSF (grant SES-9972442). We are grateful to Fabrizio Zilibotti (the editor), and three referees for helpful comments and suggestions. We also thank Hal Cole, Mark Grinblatt, Narayana Kocherlakota, Mark Machina, Massimo Marinacci, Monika Piazzesi, Bryan Routledge, Pietro Veronesi, and seminar participants at Berkeley, Carnegie-Mellon, Chicago, Davis, the Federal Reserve Bank of Atlanta, Iowa, the NBER Summer Institute, NYU, Northwestern, Rice, Rochester, Stanford, and Wharton for useful comments. This paper reflects the views of the authors and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System.

REFERENCES

- BLACKWELL, D. and DUBINS, L. (1962), "Merging of Opinions with Increasing Information", *Annals of Mathematical Statistics*, **33**, 882–886.
- BRANDT, M. "Portfolio Choice Problems", in Y. Ait-Sahalia and L. Hansen (eds.) *Handbook of Financial Econometrics* (forthcoming).
- BRENNAN, M. (1998), "The Role of Learning in Dynamic Portfolio Decisions", *European Finance Review*, **1**, 295–306.
- DOW, J. and WERLANG, S. (1992), "Uncertainty Aversion, Risk Aversion and the Optimal Choice of Portfolio", *Econometrica*, **60**, 197–204.
- EPSTEIN, L. G. and SCHNEIDER, M. (2003a), "Recursive Multiple-Priors", *Journal of Economic Theory*, **113**, 1–31.
- EPSTEIN, L. G. and SCHNEIDER, M. (2003b), "IID: Independently and Indistinguishably Distributed", *Journal of Economic Theory*, **113**, 32–50.
- EPSTEIN, L. G. and SCHNEIDER, M. (2007a), "Ambiguity, Information Quality and Asset Pricing", *Journal of Finance* (forthcoming).
- EPSTEIN, L. G. and SCHNEIDER, M. (2007b), "Supplementary Appendix to 'Learning under Ambiguity'" (Working Paper).
- EPSTEIN, L. G. and WANG, T. (1994), "Intertemporal Asset Pricing under Knightian Uncertainty", *Econometrica*, **62**, 283–322.
- FELDMAN, D. (1992), "Logarithmic Preferences, Myopic Decisions and Incomplete Information", *Journal of Financial and Quantitative Analysis*, **27**, 619–629.
- GILBOA, I. and SCHMEIDLER, D. (1989), "Maxmin Expected Utility with Nonunique Prior", *Journal of Mathematical Economics*, **18**, 141–153.
- HUBER, P. J. (1973), "The Use of Choquet Capacities in Statistics", *Bulletin of the International Statistical Institute*, **45**, 181–188.
- KREPS, D. M. (1988), *Notes on the Theory of Choice* (Boulder: Westview).
- MARINACCI, M. (2002), "Learning from Ambiguous Urns", *Statistical Papers*, **43**, 145–151.
- MERTON, R. C. (1971), "Optimum Consumption and Portfolio Rules in a Continuous Time Model", *Journal of Economic Theory*, **3**, 373–413.
- NIELSEN, C. K. (1996), "Rational Belief Structures and Rational Belief Equilibria", *Economic Theory*, **8** (3), 399–422.
- ROGERS, L. C. G. (2001), "The Relaxed Investor and Parameter Uncertainty", *Finance and Stochastics*, **5**, 134–151.
- SCHERVISH, M. J. (1995), *Theory of Statistics* (New York: Springer).
- WALLEY, P. (1991), *Statistical Reasoning with Imprecise Probabilities* (New York: Chapman and Hall).