LEARNING UNDER AMBIGUITY*

Larry G. Epstein Martin Schneider

Abstract

This paper considers learning when the distinction between risk and ambiguity (Knightian uncertainty) matters. Working within the framework of recursive multiple-priors utility, the paper formulates a counterpart of the Bayesian model of learning about an uncertain parameter from conditionally i.i.d. signals. Ambiguous signals capture responses to information that cannot be captured by noisy signals. They induce nonmonotonic changes in agent confidence and prevent ambiguity from vanishing in the limit. In a dynamic portfolio choice model, learning about ambiguous returns leads to endogenous stock market participation costs that depend on past market performance. Hedging of ambiguity provides a new reason why the investment horizon matters for portfolio choice.

*Department of Economics, U. Rochester, Rochester, NY, 14627, lepn@troi.cc.rochester.edu, and Department of Economics, NYU, New York, 10003, ms1927@nyu.edu. Epstein acknowledges financial support from NSF (grant SES-9972442). We are grateful to Hal Cole, Mark Grinblatt, 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 helpful comments and suggestions.

1 INTRODUCTION

Models of learning typically assume that agents assign (subjective) probabilities to all relevant uncertain events.[1] These models leave no role for *confidence* in probability assessments; in particular, the degree of confidence does not affect behavior. 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 behaviorally meaningful in a static context: people treat ambiguous bets differently from risky ones. This paper argues that the distinction matters also for learning.

As a concrete example of the limitations of "probabilistic" learning, consider an agent who prepares for the possibility of an earthquake. To determine the likelihood that a major quake will occur in the next five years, he samples the opinions of several geologists. The experts disagree, but all provide numbers between 1% and 4%. If the agent thinks in terms of probabilities, he forms a (possibly weighted) average of the different estimates to arrive at a single probability of, say, 2%. Now suppose that there is news of a breakthrough in geological research that leads all experts to agree that the probability of a quake is in fact 2%. The agent will not react to this news at all. He will continue to believe that the probability of an earthquake is 2% and any decisions – how much insurance to buy, for example – will be made as before. However, the news is likely to make most people feel more confident about having the "right" estimate and will conceivably reduce their demand for insurance - after all, in light of the news they do not have to worry at all about high estimates near 4%. This illustrates our motivation for modeling learning under ambiguity - learning affects confidence and hence also behavior.

Our starting point is *recursive multiple-priors* utility, a model of utility axiomatized in Epstein and Schneider [15], that extends Gilboa and Schmeidler's [21] model of decision making under ambiguity to an intertemporal setting. Under recursive multiple-priors utility, learning is completely determined by a *set* of probability measures $\mathcal{P}$, and measure-by-measure Bayesian updating describes responses to data. The present paper introduces a tractable class of sets $\mathcal{P}$, designed to capture a decision-maker's *a priori* view that the data are independent, conditionally on an unknown (but fixed) parameter. This *a priori* view is commonly imposed on learners in a wide variety of economic applications.[2] We establish several distinctive properties of learning under ambiguity and provide an application to portfolio choice to illustrate it in an economic setting.

The dynamics of learning under ambiguity are generally richer than for probabilistic learning, since agent confidence changes along with beliefs. We focus in particular on learning from *ambiguous signals*, that is, situations where agents are uncertain about the

[1] This is true for models of rational learning, which typically follow a Bayesian approach, but also for models of adaptive learning.

[2] In a Bayesian framework, one would assume that signals are i.i.d. with a distribution that depends on an unknown parameter, and that the decision maker has a prior over the parameter. Equivalently, the decision-maker's subjective probability belief about the data is exchangeable.

distribution of the data, given the parameter. Ambiguous signals imply nonmonotonic changes in confidence – they can induce ambiguity, rather than resolve it. They also prevent ambiguity from vanishing in the limit: while the learning process always settles down, it may do so in a state of permanent ambiguity. Ambiguous signals play an important role in our intertemporal portfolio choice problem, where they are used to model stock returns. Changes in investor confidence then give rise to history-dependent market exit and entry rules even when there are no market frictions.

In our model, the set $\mathcal{P}$ is represented with the help of a tuple $(\Theta, \mathcal{M}_0, \mathcal{L})$. A parameter space Θ represents features of the environment that are reflected in every signal and that the agent therefore expects to learn. A set of priors $\mathcal{M}_0$ represents the agent's initial view of the 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 signals. When a signal is unambiguous, or noisy, the set $\mathcal{L}$ is a singleton – the distribution of the signal conditional on a parameter value is known. In this sense the meaning of a noisy signal is clear. However, in many settings it is plausible that the agent is wary of a host of poorly understood or unknown features of the environment that affect realized signals and obscure their meaning. Such situations are modeled via multiple likelihoods in $\mathcal{L}$: for ambiguous signals, the distribution conditional on a parameter is not unique.

Ambiguous signals may provide information that reduces confidence. In the context of the above earthquake example, imagine a further discovery, say of a new faultline, that again induces disagreement among geologists. Such a signal would take away confidence. To illustrate how a reduction in confidence is reflected in behavior, we describe a hypothetical choice problem of ranking bets about which there is either ambiguous or noisy information. We argue that this thought experiment reveals a failure of the Bayesian model that is similar to - and inspired by - the Ellsberg Paradox, but that arises specifically in the context of learning. Our model captures the evolution of confidence through changes in a set of posterior probabilities, denoted $\mathcal{M}_t$. Signals that are unambiguous, or that confirm current beliefs, shrink the set $\mathcal{M}_t$ as confidence grows. In contrast, signals that are perceived as ambiguous, or that are outliers relative to current beliefs, reduce confidence and expand $\mathcal{M}_t$.

To study long run outcomes, we consider learning from a sequence of *indistinguishable* experiments. The idea is that an agent perceives some features, modeled by parameters, as common across a set of experiments, and others, modeled by the set of likelihoods, as unknown and, in particular, variable across experiments in a way that he does not understand. However, the extent to which the unknown features matter is the same in every experiment, that is, the experiments are viewed as *a priori* indistinguishable. Formally, this is captured by a sequence of ambiguous signals described by the *same,* time-invariant, set of likelihoods. We show that in this case learning eventually settles down: in the long run, beliefs and confidence change little with every new observation.

We focus on representations of beliefs that are *learnable*: under regularity conditions, beliefs become concentrated on one parameter value. In other words, the agent is able to resolve ambiguity (in fact, any uncertainty) about those features of the environment that affect all experiments. However, with ambiguous signals, there are time-varying unknown features that remain impossible to know, even after many observations. Thus ambiguity does not vanish in the long run. Instead, the agent moves towards an environment of time-invariant ambiguity, where he has learned all that he can.[3]

We use our model to study portfolio choice and asset market participation by investors who are ambiguity averse and learn over time about asset returns. Standard models based on expected utility and rational expectations predict that investors should diversify broadly. In fact, the typical U.S. investor participates in only a few asset markets. For example, many households stay out of the stock market altogether, and those who do hold equity often pick only a few individual stocks. This selective participation in asset markets has been shown to be consistent with optimal *static (or myopic)* portfolio choice by ambiguity averse investors.[4] What is new in the present paper is that we solve the – more realistic – *intertemporal* problem of an investor who rebalances his portfolio in light of new information that affects both beliefs and confidence.

We develop two main points. First, we show that with time-varying ambiguity in returns, the investment horizon matters for asset allocation: the optimal myopic portfolio need not be optimal for a long-horizon investor. The reason is that if the investor perceives both returns in the short term and investment opportunities in the longer term to be affected by a common unknown (ambiguous) factor, then a long-horizon investor will have an implicit exposure to that factor. His optimal portfolio may try to hedge this exposure. This *hedging of ambiguity* is distinct from the familiar hedging demand driven by intertemporal substitution effects, stressed by Merton [31]. Indeed, we show that it arises even when preferences over risky payoff streams are logarithmic, so that traditional hedging demand is zero.

Second, the model delivers *endogenous market exit and entry rules* that depend on past market performance as well as on the planning horizon. We illustrate such rules in a calibrated example studying asset allocation between stocks and riskless bonds for U.S. investors in the post-war period. Stock returns are modeled as ambiguous signals - it is possible to learn about their distribution from new data, but they are affected also by unknown factors that change over time and that cannot be learned. We propose a way for an investor to determine his degree of ambiguity about returns and we calculate optimal portfolios in real time. For reasonable parameter values, the model recommends that U.S. investors hold virtually no stocks in the 1970s, entering the market only in the late 1980s.

The paper is organized as follows. Section 2 presents a sequence of thought experiments to argue that the Bayesian model is not a satisfactory model of learning. Section 3

[3]Existing applications of ambiguity to financial markets 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 [16] for further properties of time invariant ambiguity.

[4]See Guiso and Haliassos [22] for a survey of the existing empirical and theoretical literature on portfolio choice. Non-participation with multiple priors was first derived by Dow and Werlang [13]. The argument is discussed in detail in Section 6.

briefly reviews recursive multiple-priors utility. Section 4 introduces the learning model. Section 5 establishes properties of learning in the short and long run. Section 6 applies our setup to portfolio choice. Section 7 discusses related literature. The Appendix contains proofs.

2 EXAMPLES

In this section, we argue that changes in confidence and the distinction between ambiguous and noisy signals are important features of learning. We describe examples of decision problems where (i) a Bayesian would behave differently than an ambiguity averse learner and (ii) the behavior of the Bayesian learner appears counterintuitive. The examples are based on the Ellsberg Paradox, a version of which is reviewed below. However, while the Ellsberg Paradox refers to static or one-shot-choice, we focus on concerns that are specific to learning and hence to dynamic settings.

The Ellsberg Paradox

Consider two urns, each containing four balls that are either black or white. The agent is told that the first "risky" urn contains two balls of each color. For the second "ambiguous" urn, he is told only that it contains at least one ball of each color. One ball is to be drawn from each urn. In what follows, a *bet on the color of a ball* is understood to pay one dollar (or one util) if the ball has the desired color and zero otherwise. Intuitive behavior pointed to by Ellsberg is the preference to bet on drawing black from the risky urn as opposed to the ambiguous one, and a similar preference for white. This behavior is inconsistent with any single probability measure on the associated state space. The intuition is that confidence about the odds differs across the two urns. This difference matters for behavior, but cannot be captured by probabilities.

The multiple-priors model accommodates Ellsberg-type behavior. Formally, let the state of the world be $s = (s^r, s^a) \in S = \{B, W\} \times \{B, W\}$. Then s^r (s^a) denotes the color of the ball drawn from the risky (ambiguous) urn and the indicator function $\mathbf{1}_{s^r}$ $(\mathbf{1}_{s^a})$ denotes the corresponding bet. Given a set $\mathcal{P}$ of probability measures on S, the multiple-priors utility of the bet $\mathbf{1}_{s^a}$ is

$$U(\mathbf{1}_{s^a}) = \min_{p \in \mathcal{P}} \; p\left(s^a\right).$$

The model predicts the intuitive choices if $\mathcal{P}$ contains a probability measure that assigns probability greater than $\frac{1}{2}$ to black in the ambiguous urn and another measure that assigns probability less than $\frac{1}{2}$.

Example 1: Unambiguous Signals

A simple example of learning under ambiguity obtains if repeated sampling with replacement is permitted from the given ambiguous urn. Compare the choice of bets in light of different sequences of draws. Two properties of learning are intuitive. First, Ellsberg-type behavior should be exhibited in the *short run,* because a few draws cannot

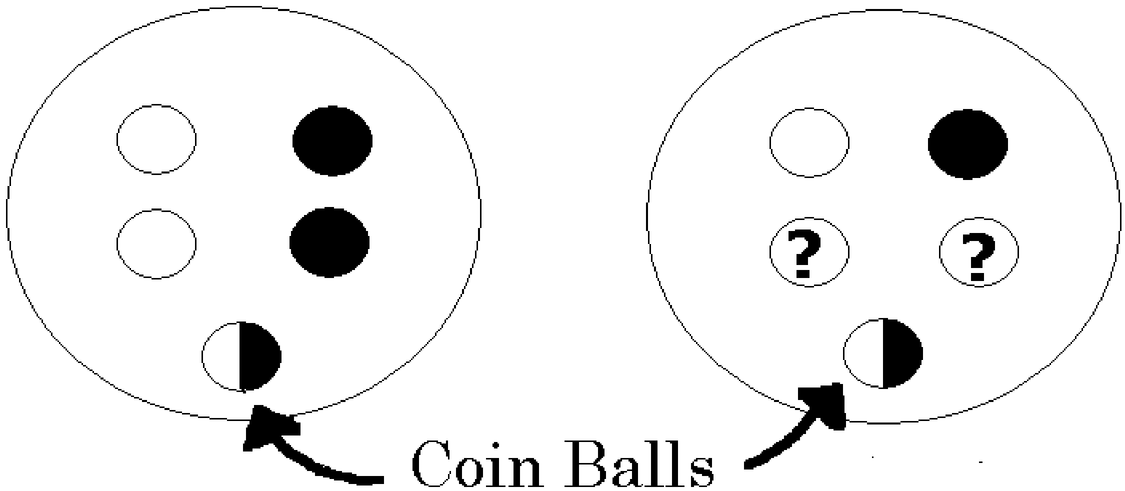

Figure 1: Experiment 2 – Risky urn on left, ambiguous urn on right.

plausibly resolve the initial ambiguity. Second, in the *long run,* learning should resolve ambiguity, since the ambiguous urn remains unchanged. As the number of draws increases, ambiguity should diminish and asymptotically the agent should behave as if he knew the fraction of black balls were equal to their empirical frequency. In other words, as one observes more draws, confidence in the odds increases and this is reflected in behavior.

The Bayesian model cannot deliver Ellsberg-type behavior in the short run. It therefore cannot adequately capture learning dynamics. However, a model of learning under ambiguity satisfying both properties can be constructed by augmenting multiple-priors utility with specific assumptions on beliefs. In particular, suppose that beliefs about the composition of the urn are represented by a set of probability measures ($\mathcal{M}_0$ in the formal model below). If this set is large enough, Ellsberg-type behavior is accommodated in the short run. As the set shrinks with time, confidence increases and ambiguity is resolved.

Example 2: Ambiguous Signals

Sampling with replacement from an Ellsberg urn is an example of learning from unambiguous signals. Indeed, conditional on the composition of the urn, the distribution of the signal is dictated by the physical environment and hence known. We now modify the example to introduce ambiguous signals. Suppose that one ball is added to each of the two urns (see Figure 1). Its color can be black or white, as determined by tossing a fair coin. The 'coin ball' is black if the coin toss produces heads and it is white otherwise. The two coin tosses that determine the color of the coin ball are independent across urns. Modify also the type of bet. Instead of betting on the next draw, the agent is now invited to bet *on the color of the coin ball.*

A priori, before any draw is observed, one should be indifferent between bets on the coin ball from either urn - all these bets amount to betting on a fair coin. Suppose now that one draw from each urn is observed and that both balls drawn are black. Neither draw gives perfect information about the coin ball in the urn from which it is drawn. Moreover, there is a difference between the information the draws provide about their respective urns. It is intuitive that one would prefer to bet on a black coin ball in the risky urn rather than in the ambiguous urn. The reasoning here could be something like "if I see a black ball from the risky urn, I know that the probability of the coin ball being black is exactly $\frac{3}{5}$. On the other hand, I'm not sure how to interpret the draw of a black ball from the ambiguous urn. There may be 3 black non-coin balls or there may be just 1 black non-coin ball. The posterior probability that the coin ball is black could be anywhere between $\frac{4}{7}$ and $\frac{2}{3}$. So I'd rather bet on the risky urn." By the same reasoning, if both drawn balls are white, one should prefer to bet on a white coin ball in the risky urn rather than in the ambiguous urn.

Could a Bayesian exhibit these choices? In principle, one can construct a subjective probabilistic belief about the composition of the ambiguous urn to rationalize these choices. However, any such belief must imply that the colors of the non-coin balls in the ambiguous urn depend on the color of the coin ball, directly contradicting the physical description of the experiment. To see this, let p^b (p^w) denote the conditional probability of drawing a black (white) ball from the non-coin balls in the ambiguous urn given that the coin ball in the ambiguous urn is black (white). Independence of the coin ball from the non-coin balls is respected only if $p^b = 1 - p^w$.

For the Bayesian, the probability of winning the bet on a black coin-ball in the ambiguous urn given that a black draw has been observed is

$$\pi = \frac{\frac{1}{2}\left(\frac{1}{5} + \frac{4}{5}p^b\right)}{\frac{1}{2}\left(\frac{1}{5} + \frac{4}{5}p^b\right) + \frac{1}{2}\frac{4}{5}\left(1 - p^w\right)} = \frac{1 + 4p^b}{5 + 4\left(p^b - p^w\right)},$$

and, with independence, we have $\pi = \frac{1+4p^b}{1+8p^b}$. The probability of winning the corresponding bet for the risky urn is simply $\frac{3}{5}$. The intuitive choice between bets on black coin balls given a black draw thus requires $\pi < \frac{3}{5}$, or $p^b > \frac{1}{2}$. By symmetry of the problem, the intuitive choice between bets on white coin balls given a white draw requires $p^w > \frac{1}{2}$. But then $p^b > 1 - p^w$, contradicting independence.

The above choices are intuitive, because the ambiguous signal (the draw from the ambiguous urn) appears to be of lower quality than the noisy signal (the draw from the risky urn). A perception of low quality arises because one cannot be completely confident about the true distribution of the ambiguous signal. The problem with the Bayesian model is that it cannot capture this dimension of information quality. The only way a Bayesian can rationalize the choices is by introducing systematic bias into the signal. This is not reasonable in the present context, because systematic bias was ruled out in the description of the experiment. We conclude that the Bayesian model cannot satisfactorily capture the difference between the two signals. The failure of the Bayesian model here is related to its failure in the Ellsberg paradox. However, a key difference is that here the prior belief about the color of the coin ball is *unambiguous*. Nevertheless, ambiguity in the signal induces conditional Ellsberg-type behavior. Our model below will accommodate this by permitting *multiple likelihoods.* Section 5 will revisit the current example and discuss learning from a sequence of draws.

3 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 $s^t = (s_1, ..., s_t)$. 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 $t = 0, 1, ...$, given the history s^t, the agent's ordering is represented by a conditional utility function U_t, defined recursively by

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

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 nonsingleton. We refer to $\{\mathcal{P}_t\}$ as the *process of conditional one-step-ahead beliefs.*

To clarify the connection to the Gilboa-Schmeidler model, it is helpful to rewrite utility using discounted sums. In a Bayesian model, the set of all conditional-one-step-ahead probabilities uniquely determine 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 [15].[6] Thus one obtains the following equivalent and explicit formula for utility:

$$U_t(c; s^t) = \min_{p \in \mathcal{P}} E^p \left[\Sigma_{s \geq t} \beta^{s-t} u(c_s) \mid s^t \right]. \tag{2}$$

This expression shows that each conditional ordering conforms to the multiple-priors model in Gilboa and Schmeidler [21], 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 [15]. The essential axioms are that (i) conditional orderings satisfy the Gilboa-Schmeidler axioms, and (ii) conditional orderings are connected by dynamic consistency. The analysis in [15] also clarifies the role of the set $\mathcal{P}$ in an intertemporal multiple-priors model. In particular, $\mathcal{P}$ should not be interpreted as the "set of time series models

[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.

[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 [16], generalizing to sets of priors the standard notion of regularity for a single prior.

that the agent contemplates". Indeed, the axioms imply restrictions on $\mathcal{P}$, although they do not impose structure on agents' beliefs. Instead, restrictions on $\mathcal{P}$ are needed to capture aspects of dynamic behavior, such as backward-induction reasoning implied by the dynamic consistency axiom. This observation is important for applications such as learning: if $\mathcal{P}$ is selected on the basis of statistical criteria alone, this might have unintended, or hard-to-understand, consequences for dynamic behavior.

Recursive multiple priors has some important features in common with the standard expected utility model. Decision making after a history s^t is not only dynamically consistent, but it also does not depend on unrealized parts of the decision tree. In other words, utility given the history s^t, depends only on consumption in states of the world that can still occur. To ensure such dynamic behavior 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 modeling learning dynamics. The analysis in [15] restricts $\{\mathcal{P}_t\}$ only by technical regularity conditions. We now proceed to add further restrictions to capture how the agent responds to data.

4 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.[7]

Bayesian Learning

The Bayesian model of learning about a memoryless mechanism can be 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, perhaps based on unmodeled prior information. 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 S. The pair (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 come about (through ℓ).

Beliefs are equivalently represented by a probability measure p on *sequences* of signals (that is, on S^∞), or by the process $\{p_t\}$ of one-step-ahead conditionals of p. The dynamics of Bayesian learning can be summarized by

[7] As an example of an environment where such a view is natural, consider data generated by sampling with replacement from an urn that contains balls of various colors in unknown proportions. Here the mechanism is the same every period because the urn is always the same. The mechanism is memoryless because draws are independent.

$$p_t\left(\cdot|s^t\right) \; = \int_\Theta \ell(\cdot \mid \theta)\, d\mu_t(\theta|s^t), \tag{3}$$

where μ_t is the posterior belief about θ, defined recursively using Bayes' Rule by

$$d\mu_t(\cdot) \; = \frac{\ell(s_t \mid \cdot)}{\int_\Theta \ell(s_t \mid \theta')\, d\mu_{t-1}(\theta'|s^t)}\; d\mu_{t-1}\left(\cdot|s^t\right). \tag{4}$$

Ambiguous Priors and Signals

Turn now to learning about ambiguous memoryless mechanisms. The starting point is again 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. To accommodate ambiguous signals, represent beliefs about the distribution of signals by a *set* of likelihoods $\mathcal{L}$. Every parameter value $\theta \in \Theta$ is associated with a set of probability measures $\mathcal{L}(\cdot \mid \theta) \; = \; \{\ell\,(\cdot \mid \theta) : \ell \in \mathcal{L}\}$. The size of this set reflects the decision-maker's (lack of) confidence in what an ambiguous signal means, given that the parameter is equal to θ. Signals are unambiguous only if there is a single likelihood, that is $\mathcal{L} = \{\ell\}$. Otherwise, the decision-maker feels unsure about how parameters are reflected in data.

Beliefs about the signal s_t are described by the same set $\mathcal{L}$ for every t – this captures the perception that the same mechanism is at work every period. Moreover, for a given parameter value $\theta \in \Theta$, signals are assumed to be independent over time – the mechanism is perceived to be memoryless. The agent perceives some factors, modeled by θ, as common across time or experiments, and others, modeled by the multiplicity of $\mathcal{L}$, as variable across time in a way that he 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, because θ is fixed over time, he can try to learn the true θ, but he has decided that he 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 $\mathcal{M}_t\left(s^t\right)$, 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 (3), there is now a (typically nondegenerate) *set* assigned to every history:

$$\mathcal{P}_t(s^t) = \left\{ p_t\left(\cdot\right) \; = \int_\Theta \ell(\cdot \mid \theta)\, d\mu_t(\theta) : \mu_t \in \mathcal{M}_t(s^t),\; \ell \in \mathcal{L} \right\}, \tag{5}$$

or, in convenient notation,

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

This process enters the specification of recursive multiple priors preferences (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, he views both his prior information and the sequence of signals as ambiguous. As a result, he will typically entertain a number of different theories about how the sample was generated. Adapting the notation used in the Bayesian case above, a theory is now summarized by a pair (μ_0, ℓ^t), where $\ell^t = (\ell_1, .., \ell_t) \in \mathcal{L}^t$ is a *sequence* of likelihoods. The decision-maker contemplates different sequences ℓ^t because he is not confident that ambiguous signals are identically distributed over time.

We allow for different attitude towards past and future ambiguous signals. On the one hand, $\mathcal{L}$ is the set of likelihoods possible *in the future*. Since the decision-maker has decided he cannot learn the true sequence of likelihoods, it is natural that beliefs about the future must be based on the *whole* set $\mathcal{L}$ as in (5). On the other hand, the decision-maker may *reevaluate*, with hindsight, his 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 $\ell^t = (\ell_1, ..., \ell_t)$, 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:

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

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 $\mu_t\left(\cdot\, ; s^t, \mu_0, \ell^t\right)$ 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 adapting (4) to accommodate time variation in likelihoods:

$$d\mu_t\left(\cdot\, ; s^t, \mu_0, \ell^t\right) = \frac{\ell_t(s_t \mid \cdot)}{\int_\Theta \ell_t(s_t \mid \theta')\, d\mu_{t-1}(\theta'; s^{t-1}, \mu_0, \ell^{t-1})} d\mu_{t-1}(\cdot\, ; s^t, \mu_0, \ell^{t-1}). \tag{6}$$

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

$$\mathcal{M}_t^\alpha(s^t) = \{\mu_t\left(s^t;\mu_0,\ell^t\right):\ \mu_0 \in \mathcal{M},\ \ell^t \in \mathcal{L}^t, \qquad (7)$$
$$\int \Pi_{j=1}^t \ell_j\left(s_j|\theta\right) d\mu_0(\theta) \geq \alpha \max_{\substack{\tilde{\mu}_0 \in \mathcal{M}_0 \\ \tilde{\ell}^t \in \mathcal{L}^t}} \int \Pi_{j=1}^t \tilde{\ell}_j\left(s_j|\theta\right) d\tilde{\mu}_0\}.$$

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.

Special Cases

In general, our model of learning about an ambiguous memoryless mechanism is summarized 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

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

where $\mathcal{M}_t^\alpha$ is given by (7). The model reduces to the Bayesian model when both the set of priors and the set of likelihoods have only a single element.

An 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: $\mathcal{M}_t^\alpha$ 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 $\alpha > 0$, 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 – Example 2 in Section 2 is one concrete case.[8]

Learning from *unambiguous* signals is captured by a single likelihood, but a possibly nondegenerate set of priors. In this case, beliefs can be summarized alternatively by a set of *exchangeable* measures on S^∞, together with the reevaluation parameter α. Indeed, let $\mathcal{P}^e$ denote a set of exchangeable measures. By the de Finetti Theorem, every $p \in \mathcal{P}^e$

[8] Another example is in Epstein and Schneider (2004) 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.

can be represented by a tuple (Θ, μ, ℓ). Assume that all these representations share the parameter space Θ and likelihood ℓ, but that they differ in the priors, which make up a set $\mathcal{M}_0$. Then beliefs with representation $(\Theta, \mathcal{M}_0, \mathcal{L}, \alpha)$ can be written equivalently as[9]

$$\mathcal{P}_t\left(s^t\right) = \left\{p_t(\cdot|s^t) : p \in \mathcal{P}^e, \;\; p\left(s^t\right) \geq \alpha \max_{\tilde{p} \in \mathcal{P}^*} \tilde{p}\left(s^t\right)\right\}.$$

In particular, if $\alpha = 0$, then $\mathcal{P}_t\left(s^t\right)$ coincides with the set of all conditional one-step-ahead measures on S induced by $\mathcal{P}^*$ and the sample s^t.

5 PROPERTIES

In this section we illustrate properties of learning under ambiguity. Sections 5.1 and 5.2 consider the dynamics of beliefs and confidence in the short run, focusing on the cases of unambiguous and ambiguous signals, respectively. In particular, we show that the model rationalizes the intuitive choices in the examples of Section 2. Section 5.3 provides a result on convergence of the learning process in the long run.

5.1 Unambiguous Signals

To see how *unambiguous signals* resolve ambiguity, consider again Example 1 from Section 2. The natural period state space is $S = \{B, W\}$ – a state corresponds to the color of a drawn ball. Let $\theta \in \Theta = \{2, 3\}$ denote the number of black balls in the ambiguous urn. A prior over Θ can be described by the probability that $\theta = 3$. For simplicity, denote this probability directly by μ_0 and let the initial set of priors be given by an interval $\left[\underline{\mu}_0, \bar{\mu}_0\right]$. The objective description of the environment implies the single likelihood

$$\ell\left(B|\theta\right) = \frac{\theta + 1}{4}.$$

After t draws, n of which are black, the posterior probability that $\theta = 3$, based on the prior μ_0, is

$$\mu_t = \mu_t\left(n, \mu_0\right) = \frac{\mu_0\left(\frac{3}{4}\right)^n\left(\frac{1}{4}\right)^{t-n}}{\mu_0\left(\frac{3}{4}\right)^n\left(\frac{1}{4}\right)^{t-n} + \left(1-\mu_0\right)\left(\frac{1}{2}\right)^n\left(\frac{1}{2}\right)^{t-n}} = \frac{\mu_0}{\mu_0 + \left(1-\mu_0\right)2^t 3^{-n}}. \tag{8}$$

The effect of the sample on posteriors is captured by a single number, the number of black balls, which is a sufficient statistic in the Bayesian case. For $\alpha = 0$, the posterior set is conveniently described by the interval $\left[\mu_t(n, \underline{\mu}_0), \mu_t\left(n, \bar{\mu}_0\right)\right]$. The data density is simply the denominator in (8). The maximum likelihood theory must be either $\mu_0 = \bar{\mu}_0$ or $\mu_0 = \underline{\mu}_0$, depending on whether $2^t 3^{-n}$ is smaller or larger than one. In the knife edge

[9] Here $p\left(s^t\right)$ is short-hand for the marginal probability $p\left(\{s^t\} \times \Pi_{i=t+1}^{\infty} S_i\right)$, whereas p_t

case where $2^t 3^{-n} = 1$, all priors are equally likely. It follows that the posterior interval is always an interval with $\mu_t(n, \underline{\mu}_0)$ or $\mu_t(n, \bar{\mu}_0)$ as one of the bounds. The other bound will typically be determined by the likelihood ratio test in (7). As a general rule, the posterior interval will be larger the less informative is the sample, that is, the closer $2^t 3^{-n}$ is to one. It is natural that, with unambiguous signals, a single sufficient statistic can determine how informative is the sample.

The posterior set $\mathcal{M}_t$ is effectively the set of posteriors that would be obtained by a collection of Bayesian estimations that are performed with different priors μ_0, but with the same "model" ℓ. As more data become available, every posterior places less weight on the prior and more weight on the sample. Ambiguity about the parameter thus shrinks as the posteriors become more similar – the agent becomes more and more confident. Let the true data generating process be the i.i.d. measure implied by $\ell(.|\theta^*)$ for some true parameter θ^*. Under standard regularity conditions (and it is obvious in the present example), every posterior converges almost surely under the truth to δ_{θ^*}, the Dirac measure that puts probability one on the true parameter.[10] Thus ambiguity about the data generating process is resolved in the limit.

5.2 Ambiguous Signals

The evolution of the posterior set $\mathcal{M}_t$ in (7) shows how ambiguous signals tend to induce ambiguity in beliefs about parameters. Indeed, suppose there is only a single prior, $\mathcal{M}_0 = \{\mu_0\}$, but that there is a nondegenerate set of likelihoods. In this case, even though there is no ambiguity ex ante, the sets $\mathcal{M}_t$ will be nondegenerate. Lack of confidence when interpreting signals thus translates into lack of confidence in knowledge about parameters. This feature is at the heart of Example 2 of Section 2. We now demonstrate that a model with ambiguous signals can rationalize the intuitive choices discussed there.

Beliefs

To define a general class of beliefs about the urns of Example 2, it is natural to let the parameter be $\theta = (\theta^r, \theta^a) \in \Theta = \{B, W\} \times \{B, W\}$, where θ^r (θ^a) denotes the color of the coin ball in the risky (ambiguous) urn, and to assume the single prior μ_0, independent across urns and satisfying $\mu_0(\theta^a = B) = \mu_0(\theta^r = B) = \frac{1}{2}$. Let λ denote the number of black non-coin balls in the ambiguous urn. The likelihood of drawing a black ball from that urn is given by

$$\ell_\lambda(B|\theta^a) = \begin{cases} \frac{\lambda+1}{5} & \text{if } \theta^a = B \\ \frac{\lambda}{5} & \text{if } \theta^a = W. \end{cases}$$

Given the symmetry of the environment, a natural set of likelihoods is

$$\mathcal{L}_\epsilon = \{\ell_\lambda : \lambda \in [2-\epsilon, 2+\epsilon]\}, \tag{9}$$

[10]See Marinacci [30] for a formal result.

where ϵ is a parameter, $0 \leq \epsilon \leq 1$.

In the special case $\epsilon = 1$, the set $\mathcal{L}_\epsilon$ models an agent who attaches equal weight to all logically possible urn compositions $\lambda =$ 1, 2, or 3. More generally, (9) 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, the agent might attach more weight to the 'focal' likelihood corresponding to $\lambda = 2$ as opposed to the more extreme scenarios $\lambda = 1, 3$. The parameter ϵ can be interpreted as the weight attached to the latter scenarios, as opposed to the focal likelihood. Indeed, $\mathcal{L}_\epsilon$ can be rewritten as[11]

$$\mathcal{L}_\epsilon = \{(1-\epsilon)\, \ell_2 (\cdot \mid \theta) + \epsilon \ell_\lambda (\cdot \mid \theta) : \lambda = 1,\, 2,\, 3\}.$$

A Bayesian agent ($\epsilon = 0$) considers the focal case only, while in the intermediate range, ϵ models the importance of ambiguity in beliefs and preference.

To complete the description of beliefs, assume the obvious likelihood for the risky urn as well as independence across urns, and denote by $\mathcal{L}$ the set of likelihoods relevant for the Cartesian product space defined by the two urns.

Inference from the First Draw

Suppose one draw from each urn has been observed, and that both drawn balls are black. For any single likelihood ℓ_λ for the ambiguous urn, the posterior probability of a black coin-ball in the ambiguous urn is

$$\frac{\frac{1}{2}\frac{\lambda+1}{5}}{\frac{1}{2}\frac{\lambda+1}{5} + \frac{1}{2}\frac{\lambda}{5}} = \frac{\lambda+1}{2\lambda+1} \in \left[\frac{3+\varepsilon}{5+2\varepsilon}, \frac{3-\varepsilon}{5-2\varepsilon}\right].$$

For the risky urn, the posterior probability of a black coin ball corresponds to the case $\varepsilon = 0$, and therefore equals $\frac{3}{5}$. Since this is larger than the "worst-case" posterior for any $\varepsilon > 0$, a bet on a black coin ball in the risky urn will be strictly preferred. By symmetry, the same is true for bets on white coin balls when both drawn balls are white.

Inference from Small Samples

Figure 2 depicts the evolution of beliefs about the ambiguous urn as more balls are drawn. We set $\epsilon = 1$ so that the agent weighs equally all the logically possible combinations of non-coin balls. The top panel illustrates the aggregation of ambiguous signals. It shows the evolution of the posterior interval for a sequence of draws such that the number of black balls is $\frac{3t}{5}$, for $t = 5, 10, ...$ In particular, after the first 5 ambiguous signals, with 3 black balls drawn, the agent assigns a posterior probability between .4 and .8 to the coin ball being black.

What happens if the same sample is observed again? There are two effects. First, a larger batch of signals permits more possible interpretations. For example, having seen

[11] This is a form of the ϵ-contamination model employed in robust statistics (see Walley [37], for example). In economic modeling, it is used in Epstein and Wang [18], for example.

ten 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. At the same time, the story that all but one non-coin ball was always white is somewhat less believable if the sample is larger: reevaluation limits the scope for interpretation, and more so the more data are available. The evolution of confidence, measured by the size of the posterior interval, thus depends on how much agents reevaluate their views. For an agent with $\alpha = .001$, the posterior interval expands between $t = 5$ and $t = 20$. In this sense, a sample of ten or twenty ambiguous signals induces *more* ambiguity than a sample of five. However, reevaluation implies that large enough batches of ambiguous signals induce less ambiguity than smaller ones.

The lower panel of Figure 2 tracks the evolution of posterior intervals along a representative sample. 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 behavior is reminiscent of the evolution of the Bayesian posterior variance, which is also maximal if the fraction of black balls is one half.

5.3 Beliefs in the Long Run

To discuss what is learnable in the long run, we need to define a true data generating process. Take this process to be i.i.d. corresponding to the 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 by

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

The following result (proven in the appendix) summarizes the behavior of the posterior set in the long run.

Theorem 1 *Suppose that*

(i) Θ *is finite and for every* $\theta \in \Theta$, $\mu_0(\theta) > 0$ *for some* μ_0 *in* $\mathcal{M}_0^\alpha$*; and*

(ii) $\theta^* = \arg\max_\theta H(\theta)$ *is a singleton.*

Then every sequence of posteriors from $\mathcal{M}_t^\alpha$ *converges to the Dirac measure* δ_{θ^*}*, almost surely under the i.i.d. measure described by* ϕ*.*

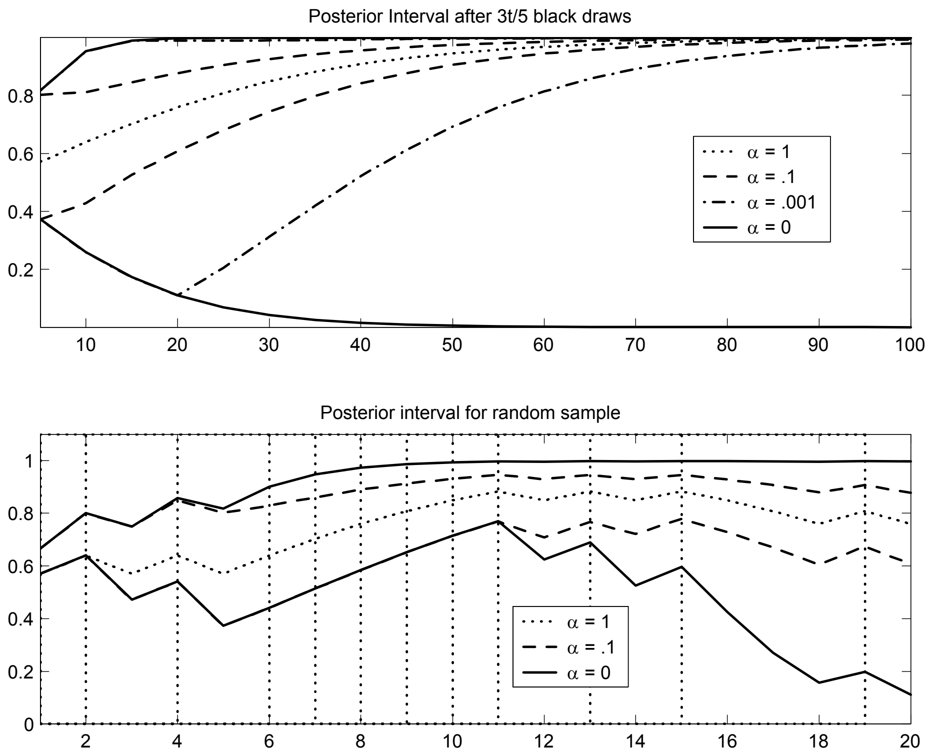

Figure 2: The Posterior Interval is the range of posterior probability that the coin ball is black, $\mu_t(B)$. In the top panel, the sample is selected to keep fraction of black balls constant. In the bottom panel, vertical lines indicate black balls drawn.

Condition (ii) is an identification condition: it says that there is at least one sequence of likelihoods (that is, the maximum likelihood sequence), such that the sample with empirical frequency measure ϕ can be used to distinguish θ^* from any other parameter value. If this condition holds, then in the long run only the maximum likelihood sequence is a permissible scenario 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.

As a concrete example, consider long run beliefs about the ambiguous urn of Example 2. Let ϕ_∞ denote the probability under the truth that a black ball is drawn. Suppose also that beliefs are given by (9), with $\epsilon = 1$: the agent views all possible combinations of black and white non-coin balls as equally likely. 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 reevaluates his views ($\alpha > 0$) and observes a large number of draws with a fraction of black balls above half believes it very likely that the color 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 (7) in ensuring learnability of the representation. Suppose that $\phi_\infty \to \frac{3}{5}$. The limiting frequency of $\frac{3}{5}$ black draws could be realized either because there is a black coin ball and on average one half of the non-coin balls were black, or because the coin ball is white, but it happens to be the case that all of the urns contained 3 black non-coin balls. If $\alpha = 0$, both possibilities are taken equally seriously and the limiting posterior set contains Dirac measures that place probability one on either $\theta = B$ or $\theta = W$. The convergence behavior is apparent from Figure 2. In contrast, for $\alpha > 0$, reevaluation eliminates the sequence where all urns contain three black non-coin balls as unlikely.

With a stronger identification condition, the posterior set converges to a set containing single Dirac measure concentrated on θ^*, even if $\alpha = 0$. Indeed, we show in the appendix that $\mathcal{M}_t^0 \longrightarrow \{\delta_{\theta^*}\}$, provided that for all $\theta \neq \theta^*$,

$$\sum_{s\in S} \phi(s) \max_{\ell \in \mathcal{L}} \log \frac{\ell(s|\theta)}{\ell(s|\theta^*)} < 0. \tag{11}$$

The latter condition is stronger than condition (ii) in the Theorem: it requires that θ^* is the most likely parameter under *all* sequences of likelihoods (ℓ_s), not simply under the maximum likelihood sequence. In the example, it reduces to

$$\phi_\infty \max_{\lambda_B} \log \frac{\lambda_B}{1+\lambda_B} + (1-\phi_\infty) \max_{\lambda_W} \log \frac{5-\lambda_W}{4-\lambda_W} < 0,$$

which satisfied if ϕ_∞ is greater than $\log(1/2) / \log(3/8)$.

6 DYNAMIC PORTFOLIO CHOICE

Selective asset market participation is puzzling in light of standard models based on expected utility and rational expectations. In contrast, it is consistent with optimal *static* portfolio choice by investors who are averse to ambiguity in stock returns. Such investors will take a nonzero (positive or negative) position in an asset only if it unambiguously promises an expected gain. Dow and Werlang (1992) consider an investor who allocates wealth between a riskless asset and one other asset with ambiguous return. They show that if the range of premia on the ambiguous asset contains zero, it is optimal to invest 100% in the riskless asset. The intuition is that, when the range of premia contains zero, the worst case expected excess return on both long and short positions is not positive.

In practice, most investors do not make one-shot portfolio choices, but have long investment horizons during which they repeatedly rebalance their portfolios. Here we consider an *intertemporal* problem, where the investor updates his beliefs about future returns. For simplicity, we follow Dow and Werlang by restricting attention to the case of a single uncertain asset. However, the effects we emphasize are relevant more widely. Indeed, work by Epstein and Wang [18] and Mukherji and Tallon [32] has clarified that the basic nonparticipation result extends to *selective* participation when many assets are available. The key observation is that the multiple-priors model allows confidence to vary across different sources of uncertainty.[12] Investors will participate in a market only if they are confident that they know how to make money in that particular market. [13]

Suppose the investor cares about wealth T periods from now and plans to rebalance his portfolio once every period. There is a riskless asset with constant gross interest rate R^f, as well as an asset with uncertain gross return $\tilde{R}_t = R(s^t)$ that depends on the history s^t observed by the investor. Signals may consist of more than current excess returns. Beliefs are given by a process of one-step-ahead conditionals $\{\mathcal{P}_t\}$. At date t, given information s^t, the investor selects a portfolio weight $\omega(s^t)$ for the risky asset by solving the recursive problem

$$\begin{aligned} V_t\left(W_t, s^t\right) &= \max_{\omega} \min_{p_t \in \mathcal{P}_t(s^t)} E^{p_t}\left[V_{t+1}\left(W_{t+1}, s^{t+1}\right)\right] \\ &\quad \text{subject to} \\ W_t &= \left(R^f + (R\left(s^t\right) - R^f)\omega\left(s^{t-1}\right)\right) W_{t-1}, \qquad t = 1, ..., T. \end{aligned} \tag{12}$$

This problem differs from existing work on portfolio choice with multiple-priors because beliefs are time-varying. We also emphasize the role of the investor's planning horizon T. In Section 6.1, we discuss the nature of intertemporal hedging with recursive multiple-priors. To distinguish intertemporal hedging due to time-varying ambiguity from the intertemporal substitution effects stressed by Merton (1973), we focus on the case of log utility, where traditional hedging demand is zero. In Section 6.2, we study a calibrated example of asset allocation between stocks and riskless bonds for U.S. investors in the post-war period.

6.1 Intertemporal Hedging and Participation

If future investment opportunities and future returns are affected by common unknown (ambiguous) factors, then investors who have a longer planning horizon, and therefore

[12] For example, suppose that the set of excess return distributions for each available asset contains a component with unknown mean that is uncorrelated with the return on other assets. The range of mean excess returns for a particular asset now reflects investor confidence towards that asset only. If the range is a large interval around zero, that is, investor confidence is low, it is optimal to allocate zero wealth to the asset.

[13] This intuition fits well with evidence from surveys among practitioners, who participate in a market only if they "have a view" about price movements of the asset in question. See Chew [11], Ch. 43 for a discussion on this issue.

care about future investment opportunities, will deal differently with ambiguity than short term investors. For example, if expected returns in some market are perceived to be highly ambiguous, the optimal myopic policy may be non-participation, to avoid exposure to ambiguity. However, if returns are also a signal of future expected returns, then a long-horizon investor is already exposed to the unknown (ambiguous) factor that is present in returns even if he does not currently participate in the market. It may then be better to take a position to offset existing exposure. We now derive intertemporal hedging demand due to ambiguity for the log utility case. We then discuss how the importance of hedging depends on the number of unknown factors, and what this implies for the emergence of dynamic participation rules.

Hedging Ambiguity

In the expected utility case, it is well known that log investors act myopically; even if the conditional distribution of future investment opportunities changes over time, the income and substitution effects of these changes are exactly offsetting. As a result, the investor's asset demand depends only on current investment opportunities, captured by the conditional one-period-ahead distribution of returns. The optimal weight after history s^t and given the beliefs process $\{p_t\}$ is simply

$$\omega^*\left(p_t\left(s^t\right)\right) := \arg\max_{\omega} E^{p_t\left(s^t\right)}\left[\log\left(R^f + \left(\tilde{R}_{t+1} - R^f\right)\omega\right)\right].$$

Consider next a *myopic* investor with nondegenerate beliefs process $\{\mathcal{P}_t\}$. Given the history s^t, such an agent solves

$$\begin{aligned} & \max_{\omega} \min_{p_t(s^t)\in\mathcal{P}(s^t)} E^{p_t\left(s^t\right)}\left[\log\left(R^f + \left(\tilde{R}_{t+1} - R^f\right)\omega\right)\right] \\ = & \min_{p_t(s^t)\in\mathcal{P}(s^t)} E^{p_t\left(s^t\right)}\left[\log\left(R^f + \left(\tilde{R}_{t+1} - R^f\right)\omega^*\left(p_t\left(s^t\right)\right)\right)\right], \end{aligned} \tag{13}$$

where we have used the minmax theorem to exchange the order of optimization. Denote by $p_t^{myopic}\left(s^t\right)$ the minimizing (conditional-one-step-ahead) measure in (13). Then the optimal policy of a myopic agent is $\omega^*\left(p_t^{myopic}\left(s^t\right)\right)$, the portfolio weights that are optimal for the corresponding Bayesian.

For the intertemporal problem (12), conjecture the value functions $V_t\left(W_t, s^t\right) = \log W_t + h_t\left(s^t\right)$, with $h_T = 0$. Again using the minmax theorem, as well as the budget constraint, we obtain

$$\begin{aligned} V_t\left(W_t, s^t\right) &= \min_{p_t(s^t)\in\mathcal{P}_t(s^t)} \max_{\omega} E^{p_t}\left[V_{t+1}\left(W_{t+1}, s^{t+1}\right)\right] \\ &= \min_{p_t(s^t)\in\mathcal{P}_t(s^t)} \left\{\max_{\omega} E^{p_t}\left[\log\left(R^f + \left(\tilde{R}_{t+1} - R^f\right)\omega\right)\right] + E^{p_t}\left[h_{t+1}\left(s^{t+1}\right)\right]\right\} + \log W_t. \end{aligned}$$

The first term depends only on time and s^t, verifying the conjecture for V_t. Denote by $p_t^*\left(s^t\right)$ a minimizing measure in

$$h_t\left(s^t\right) = \min_{p_t\in\mathcal{P}_t(s^t)} E^{p_t}\left[\log\left(R^f + \left(\tilde{R}_{t+1} - R^f\right)\omega\left(p_t\left(s^t\right)\right)\right) + h_{t+1}\left(s^{t+1}\right)\right]. \tag{14}$$

Then $\omega^*\left(p_t^*\left(s^t\right)\right)$ is an optimal policy for the intertemporal problem.

Comparison of (13) and (14) reveals that the supporting probabilities p_t^{myopic} and p_t^*, and hence the optimal policies in the myopic and intertemporal problems, need not agree. This motivates a decomposition of optimal portfolio demand ω_t in the intertemporal problem:

$$\omega_t = \underbrace{\omega^*\left(p_t^{myopic}\left(s^t\right)\right)}_{\text{myopic demand}} + \underbrace{\left(\omega^*\left(p_t^*\left(s^t\right)\right) - \omega^*\left(p_t^{myopic}\left(s^t\right)\right)\right)}_{\text{intertemporal hedging demand}}.$$

It is clear that the concept of intertemporal hedging demand identified here is unique to the case of ambiguity. With a singleton set $\mathcal{P} = \{p\}$, we have $p_t^{myopic}\left(s^t\right) = p_t^*\left(s^t\right) = p_t\left(s^t\right)$ and there is no intertemporal hedging.

Independent Unknown Factors

Intuitively, the intertemporal hedging demand should still be zero if returns are driven by factors that are completely unrelated to future investment opportunities. Under risk, this is captured by conditional independence of returns and the state variables that represent opportunities. The intuition carries through to the ambiguity case. However, independence is now required for sets of conditionals. In particular, if the set of conditional distributions for returns and state variables can be expressed as a Cartesian product of sets of independent marginal distributions, then myopic portfolio choice remains optimal. To see this, consider a simple example. Let $S = X \times Y$. Assume that, at date t, the investor observes $s_t = (x_t, y_t)$. Here x_t is a state variable that matters for future investment opportunities, but not for current returns, and y_t is a shock that affects only current returns, i.e., $R_t = R\left(x^{t-1}, y_t\right)$. Then define conditionals

$$\begin{aligned} \mathcal{P}_t\left(s^t\right) &= \left\{p_t^x \otimes p_t^y : p_t^x \in \mathcal{P}_t^x\left(x^t\right),\ p_t^y \in \mathcal{P}_t^y\right\} \\ &= \mathcal{P}_t^x\left(x^t\right) \otimes \mathcal{P}_t^y, \end{aligned} \tag{15}$$

where $\mathcal{P}_t^x\left(x^t\right)$ and $\mathcal{P}_t^y$ are sets of probability measures on X and Y, respectively.[14] We thus allow both the signal and the shock to be ambiguous.

With these beliefs, (14) becomes

$$h_t\left(x^t\right) = \min_{p_t^y \in \mathcal{P}_t^y} E^{p_t^y}\left[\log\left(R^f + \left(R\left(x^t, y_{t+1}\right) - R^f\right)\omega\left(p_t^y\right)\right)\right] + \min_{p_t^x \in \mathcal{P}_t^x\left(x^t\right)} E^{p_t^x}\left[h_{t+1}\left(x^{t+1}\right)\right],$$

and myopic behavior is clearly optimal. The key here is the separation into two minimization problems. This may no longer be possible if returns depend also on x_{t+1}. In that case, the set $\mathcal{P}_t^x$ may represent concern with *common factors* that affect both future opportunities and returns, or put alternatively, with "correlation" between the ambiguities perceived for future opportunities and returns.

To summarize, an investor who optimizes intertemporally is also concerned about future investment opportunities. As a result, he may be exposed to unknown factors

[14] The fact that $\mathcal{P}_t^y$ does not depend on history here is not restrictive since predictability of returns is permitted through the function R.

that affect the distribution of h_{t+1}, even if he does not invest in the uncertain asset over the next period.[15] However, it makes sense to react to such an exposure by changing portfolio strategy (and to deviate from the optimal myopic portfolio) only if the unknown factors that affect future opportunities have something to do with returns. Then selection of the worst case belief takes into account both terms in (14) and will typically end up at a different belief than in the myopic case, where only the first term matters.

Learning

The previous paragraph has shown that hedging demand depends on the nature of ambiguity. We now illustrate when it is important in the context of learning. We present two examples where beliefs are given by representations $(\Theta, \mathcal{M}_0, \mathcal{L}, \alpha)$. One has independent unknown common factors, and the other common factors (correlated ambiguity). The latter case is also relevant for the calibrated example of the next subsection.

To construct an example of learning with independent unknown factors, assume again $S = X \times Y$ and $R_t = R\left(x^{t-1}, y_t\right)$. Let beliefs have a representation with likelihood set

$$\mathcal{L} = \left\{\ell : \ell\left(s|\theta\right) = \ell^x\left(x|\theta\right)\ell^y\left(y\right);\ \ell^x \in \mathcal{L}_x,\ \ell^y \in \mathcal{L}_y\right\},$$

where $\mathcal{L}_x$ and $\mathcal{L}_y$ are sets of likelihoods on X and Y. The interpretation is that returns contain short term noise about which the investor does not expect to learn anything. Unknown factors that affect returns in the short run correspond to parameters that describe the set $\mathcal{L}_y$. At the same time, the investor can learn about conditional moments that depend on θ via the distribution of x. Unknown factors that affect learning are captured by $\mathcal{L}_x$ and are independent of those that affect returns in the short run. With this structure for $\mathcal{L}$, terms involving ℓ^y will cancel out of the updating equation (6) and $\mathcal{M}_t^\alpha\left(s^t\right)$ depends only on x^t. Moreover, since ℓ^y does not depend on the parameter θ, the one-step-ahead conditionals $\mathcal{P}_t\left(s^t\right)$ will be a special case of (15). Myopic portfolio choice is therefore optimal. This example illustrates that the presence of learning does not invalidate the basic intuition for non-participation.

To obtain an example of learning with common unknown factors, modify the above setup by setting $R_t = R\left(x_t\right)$, thus eliminating the shock y. For simplicity, let the signal x be unambiguous ($\mathcal{L}_x = \{\ell^x\}$) and assume that the set $\mathcal{M}_0$ is a one-parameter family of priors, parametrized by the prior mean of θ, say $\bar{\theta}_0 \in \bar{\Theta}_0$. Every set of one-step-ahead conditionals is then also a one-parameter family $0\mathcal{P}_t\left(s^t\right) = \left\{p_t\left(s^t;\bar{\theta}_0\right);\ \ \bar{\theta}_0 \in \bar{\Theta}_0\right\}$. The parameter $\bar{\theta}_0$ effectively describes a single common factor: since it affects all posteriors, it directly matters both for one-period-ahead returns and for later returns, and hence for future investment opportunities. When beliefs are parametrized this narrowly, one can construct examples where hedging demand will lead to participation at all times except in the final period. Indeed, non-participation requires that $E^{p_t}\left[R_{t+1}\right] = 0$ for a conditional p_t that achieves the minimum in (14). If distributions are chosen such that there is just one value $\bar{\theta}_0$ for which $E^{p_t\left(s^t;\bar{\theta}_0\right)}\left[R_{t+1}\right] = 0$, then this value will typically not

[15]Exposure requires that (i) the realization s_{t+1} typically provides news about future opportunities, that is, it affects future belief sets $\mathcal{P}_{t+j}\left(s^{t+j}\right)$ and (ii) that the news is payoff-relevant, so that the value function $h_{t+1}\left(s^{t+1}\right)$ depends on s_{t+1}.

achieve the minimum for the sum in (14). The minimizing value will in turn not set the mean excess return to zero.

It is difficult to make general statements beyond these examples. The importance of intertemporal hedging due to ambiguity depends on the precise factor structure. This is the case even when the analysis focuses on learning, as opposed to general time-variation in beliefs. In this respect, the situation is similar to hedging demand due to intertemporal substitution, where the literature contains few general results beyond Merton's (1973) decomposition, and progress has been mostly through calibrated examples. We proceed to develop such an example in the next section.

6.2 Learning and Stock Market Participation: A Calibrated Example

Suppose the investor allocates wealth to a riskless bond and a broad U.S. stock index once every quarter. To approximate the distribution of U.S. stock returns, assume that the state can take two values every period, $s_t \in \{0, 1\}$, and let $R(1) = 1.14$ and $R(0) = .92$. If s_t is i.i.d. with $\Pr\{s_t = 1\} = \frac{1}{2}$, then the mean and variance of NYSE returns from 1927:Q3 to 2001:Q2 are matched exactly. We fix the riskless rate at $R^f - 1 = .01$ per quarter.[16] The investor's utility function exhibits constant relative risk aversion with risk aversion parameter γ.

Beliefs are defined by a representation $(\Theta, \mathcal{M}_0, \mathcal{L}, \alpha)$. The investor thinks that something can be learned about the distribution of returns by looking at past data. This is captured by a parameter $\theta \in \Theta \equiv \left[\overline{\lambda}, 1 - \overline{\lambda}\right]$, where $\overline{\lambda} < \frac{1}{2}$. However, he also believes that there are many poorly understood factors driving returns. These are captured by multiple likelihoods, where the set $\mathcal{L}$ consists of all $\ell(\cdot \mid \theta)$ such that

$$\ell(1|\theta) = \theta + \lambda, \quad \text{for some } \lambda \in [-\overline{\lambda}, \overline{\lambda}].$$

Given our assumptions on Θ, $\ell(1|\theta)$ is between zero and one. The set of priors $\mathcal{M}_0$ on Θ is given by all the Dirac measures. For simplicity, we write $\theta \in \mathcal{M}_0$ if the Dirac measure on θ is included. If $\overline{\lambda} > 0$, returns are ambiguous signals: $\lambda_t \in [-\overline{\lambda}, \overline{\lambda}]$ parametrizes the likelihood ℓ_t. Since the set of priors consists of Dirac measures, reevaluation ($\alpha > 0$) is crucial for nontrivial updating; if $\alpha = 0$, then $\mathcal{M}_t = \mathcal{M}_0$ for all t.

Belief Dynamics

The above belief structure is convenient because the posterior set $\mathcal{M}_t^\alpha$ depends on the sample only through the fraction ϕ_t of high returns observed prior to t. More specifically, it is shown in the appendix that

$$\mathcal{M}_t^\alpha\left(s^t\right) = \left\{\theta \in \Theta : g(\theta; \phi_t) \geq \max_{\tilde{\theta} \in \Theta} g(\tilde{\theta}; \phi_t) + \frac{\log \alpha}{t}\right\}, \tag{16}$$

[16] Here we follow much of the finance literature and consider nominal returns.

where $g(\theta, \phi_t) = \phi_t \log\left(\theta + \overline{\lambda}\right) + (1 - \phi_t) \log\left(1 - \theta + \overline{\lambda}\right)$. The function $g(.;\phi)$ is strictly concave and has a maximum at $\theta = \phi + 2\overline{\lambda}\left(\phi - \frac{1}{2}\right)$.

Using (16), it is straightforward to determine the limiting behavior of the one-step-ahead beliefs $\mathcal{P}_t(s^t)$ as t becomes large. Suppose that the empirical frequency of high returns converges to ϕ_∞. Then $\mathcal{M}_t^\alpha$ collapses to the single number θ^*, where

$$\theta^* = \begin{cases} \overline{\lambda} & \text{if } \phi_\infty < \frac{2\overline{\lambda}}{1+2\overline{\lambda}} \\ \phi_\infty + 2\overline{\lambda}\left(\phi_\infty - \frac{1}{2}\right) & \text{if } \phi_\infty \in \left[\frac{2\overline{\lambda}}{1+2\overline{\lambda}}, \frac{1}{1+2\overline{\lambda}}\right] \\ 1 - \overline{\lambda} & \text{if } \phi_\infty > \frac{1}{1+2\overline{\lambda}}. \end{cases}$$

Thus $\mathcal{P}_t(s^t)$ collapses to the set $\mathcal{L}(.|\theta^*)$, which consists of all probabilities on $S = \{0,1\}$ with

$$\Pr(s = 1) \in [\theta^* - \overline{\lambda}, \theta^* + \overline{\lambda}]. \tag{17}$$

The agent thus learns the true parameter value θ^*, in the sense that in the limit he behaves as if he had been told that it equals θ^*. If the realized empirical distribution is symmetric ($\phi_\infty = \frac{1}{2}$), then $\theta^* = \phi_\infty$. We use this fact below to calibrate the belief parameters.

Bellman Equation

The problem (12) can be rewritten using the fraction ϕ_t of high returns as a state variable:

$$\begin{aligned} V_t(W_t, \phi_t) &= \max_{\omega} \min_{p_t \in \mathcal{P}_t(s^t)} E^{p_t}\left[V_{t+1}\left(W_{t+1}, \phi_{t+1}\right)\right] \\ &= \max_{\omega_t} \min_{\substack{\lambda_t \in [-\overline{\lambda}, \overline{\lambda}] \\ \theta \in \mathcal{M}_t^\alpha}} \{\, (\theta + \lambda_t) V_{t+1}(W_{t+1}(1), \phi_{t+1}(1)) \\ &+ (1 - \theta - \lambda_t) V_{t+1}(W_{t+1}(0), \phi_{t+1}(0))) \,\}, \end{aligned}$$

subject to the transition equations

$$\begin{aligned} W_{t+1}(s_{t+1}) &= \left(R^f + \left(R(s_{t+1}) - R^f\right)\omega_t\right) W_t, \\ \phi_{t+1}(s_{t+1}) &= \frac{t\phi_t + s_{t+1}}{t+1}. \end{aligned}$$

Preference Parameters

We specify $\beta = .99$ and set the coefficient of relative risk aversion γ equal to 2. It remains to specify the belief parameters α and $\overline{\lambda}$. The parameter $\overline{\lambda}$ determines how much the agent thinks that he will learn in the long run. To determine a value, he could pose the following question: "Suppose I see a large amount of data and that the fraction of high returns is $\phi_\infty = \frac{1}{2}$. How would I compare a bet on a fair coin with a bet that next quarter's returns are above or below the median?" By the Ellsberg Paradox, we would expect the agent to prefer the fair bet. He could then try to quantify this preference by asking: "What is the probability of heads that would make me indifferent between betting on heads in a coin toss and betting on high stock returns?" In light of the range of limiting probabilities given in (17), the result is $\frac{1}{2} - \overline{\lambda}$. We present results for values of $\overline{\lambda}$ ranging from 0 to 2%. Even the upper bound of 2% leaves substantial scope for learning: in terms of mean returns, it implies that the investor believes that the range of equity premia be reduced below 1% (88 basis points, to be exact) in the limit.

The parameter α determines how fast the set of possible models shrinks. Here it can be motivated by reference to classical statistics. If signals are unambiguous ($\overline{\lambda} = 0$), there is a simple interpretation of our updating rule: $\mathcal{M}_t^{\alpha}$ contains all parameters θ^* such that the hypothesis $\theta = \theta^*$ 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 $-2\log\alpha$. For a 5% significance level, $\alpha = .14$, which is the value we use below.

6.3 Numerical Results

Our leading example is an investor in 1971:Q3 ($t = 0$), who plans over various horizons (up to 30 years). He looks back on data starting in 1927:Q3. We generate a discretized returns sample by letting $s_t = 1$ if the NYSE return was above the mean in quarter t and $s_t = 0$ otherwise. Figure 3 shows the optimal stock position as a fraction of wealth for a 5 year horizon if $\alpha = .14$ and $\overline{\lambda} = 0$. The axis to the right measures time in quarters t up to the planning horizon of 120 quarters. The axis to the left measures the number of high returns observed, $H = t\phi_t$. Thus only the surface above the region $H \leq t$ (that is, the region to the right to the diagonal $H = t$ in the plane) represents the optimal stock position.

The slope of the surface suggests that the agent is by and large a momentum investor. If low returns are observed (movements to the right, increasing t while keeping H fixed), the stock position is typically reduced. On the path above the time axis, which is taken if low returns are observed every period, agents eventually go short in stocks. In contrast, if high returns are observed (movements into the page, increasing H one for one with t), the stock position is typically increased. On the ridge above the diagonal $(H = t)$, which is taken if high returns are observed every period, investors take ever larger long positions. The optimal policy surface has also a flat piece at zero: when enough low returns are observed, agents do not participate in the market. The terminal period of our model corresponds to the static Dow-Werlang setup. In all earlier periods, participation and the size of positions is in part determined by intertemporal hedging.

Hedging Ambiguity

In the present example, both θ and λ capture unknown factors that matter not only for returns over the next period, but also for future investment opportunities. For the reasons discussed in Section 6.1, participation is thus more likely earlier in the investment period. In addition, hedging due to time-varying ambiguity implies that the agent follows contrarian, as opposed to momentum, strategies when he is not sure about the size of the equity premium. To illustrate this effect, Figure 4 focuses on an agent planning in 1928:Q3, who has only one year of previous data as prior information. While the effects

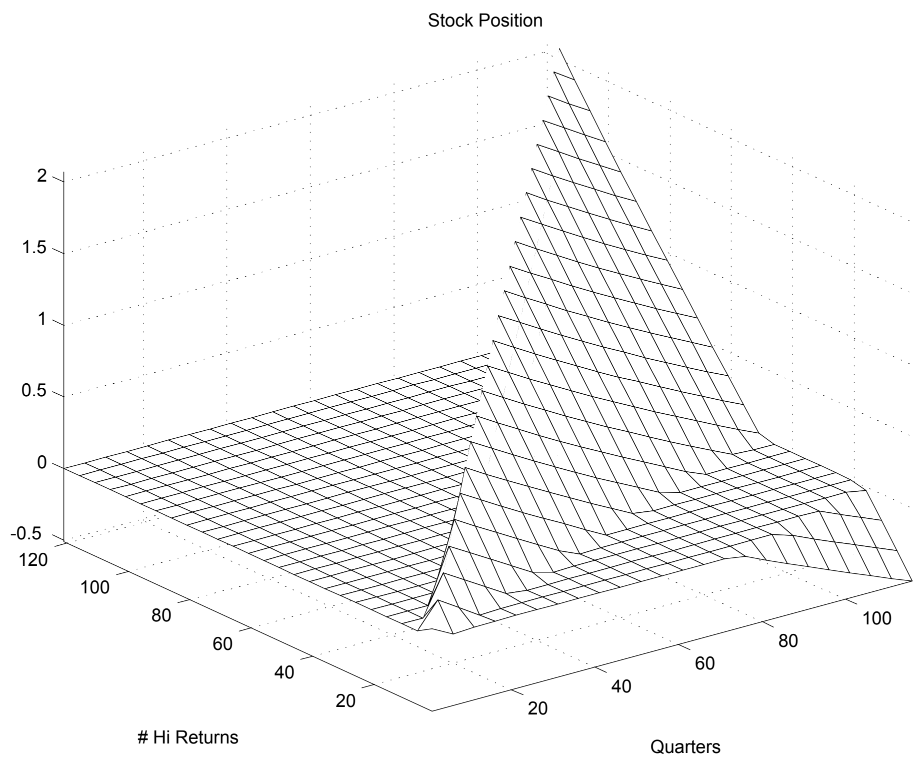

Figure 3: Optimal stock position for ambiguity-averse agents with $\alpha = .14$ and $\bar{\lambda} = 0$, for 30-year planning problem, beginning in 1971:Q3.

are qualitatively similar in later years (such as for our leading example), the effects of hedging are most pronounced if there is a large amount of prior uncertainty. The left hand panel of Figure 4 shows a representative 'section' of the optimal policy surface (for $t = 12$, here 1931:Q2). The right hand panel shows the *change* in the stock position at $t = 12$, as a function of the number of high returns, if the $13th$ observation was either a high or a low return.

Investment behavior falls into one of three regions. A *non-participation region* is reached if the absolute value of the sample equity premium is low, which occurs for ϕ_t slightly below $\frac{1}{2}$. If the equity premium has been either very high or very low, the agent is in a *momentum region.* He is long in stocks if the sample equity premium is positive and short otherwise. He also reacts to high (low) returns by increasing (decreasing) his net position. If the absolute value of the equity premium is in an intermediate range, the agent is in a *contrarian region.* He is *short* in stocks for a positive sample equity premium and long otherwise. Moreover, he now reacts to high (low) returns by decreasing (increasing) his net exposure. The contrarian region is also present in our leading example (investment starting in 1971:Q3), but is very small; this is why it is not discernible in Figure 3.

To understand why a contrarian region emerges, consider the dependence of continuation utility on ϕ_t. The term $h_t(\phi_t)$ is typically U-shaped in ϕ_t. Intuitively, the agent

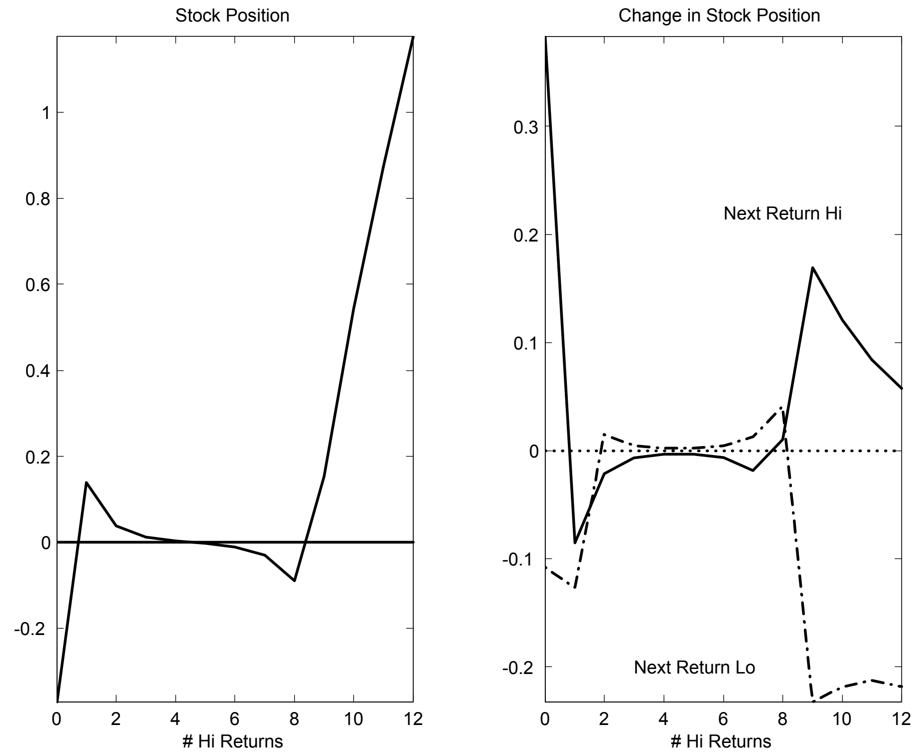

Figure 4: Stock position and responses to arrival of a high or low return after 12 quarters in 30-year planning problem beginning in 1928:Q3. Ambiguity averse investor with $\alpha = .14$ and $\bar{\lambda} = 0$.

prefers to be in a region where *either* the lowest possible expected return is much higher than the riskless rate *or* the highest possible return is far below the riskless rate, because in both cases there is an equity premium (positive or negative, respectively) that can be exploited. Suppose the lowest expected equity premium is positive. The agent must balance two reasons for investing in the stock market. On the one hand, he can exploit the expected equity premium by going long. On the other hand, he can insure himself against bad news about investment opportunities (low returns) by going short. Which effect is more important depends on how the size of the equity premium compares to the slope of the h_t's. If the equity premium is small in absolute value, the hedging effect dominates.

Participation over the last three decades

Figure 5 compares the positions that various investors would have chosen since 1971. An investor who believes that returns are unambiguous signals ($\bar{\lambda} = 0$) should always holds stocks, although his positions are quite small, barely reaching 30% even after the high returns of the 1990s. As a reference point, an agent with rational expectations who is sure that the equity premium is equal to its sample mean would hold 82% stocks every period. The Bayesian learner in Figure 5 lies between these two, increasing his position up to 80% toward the end of the sample.

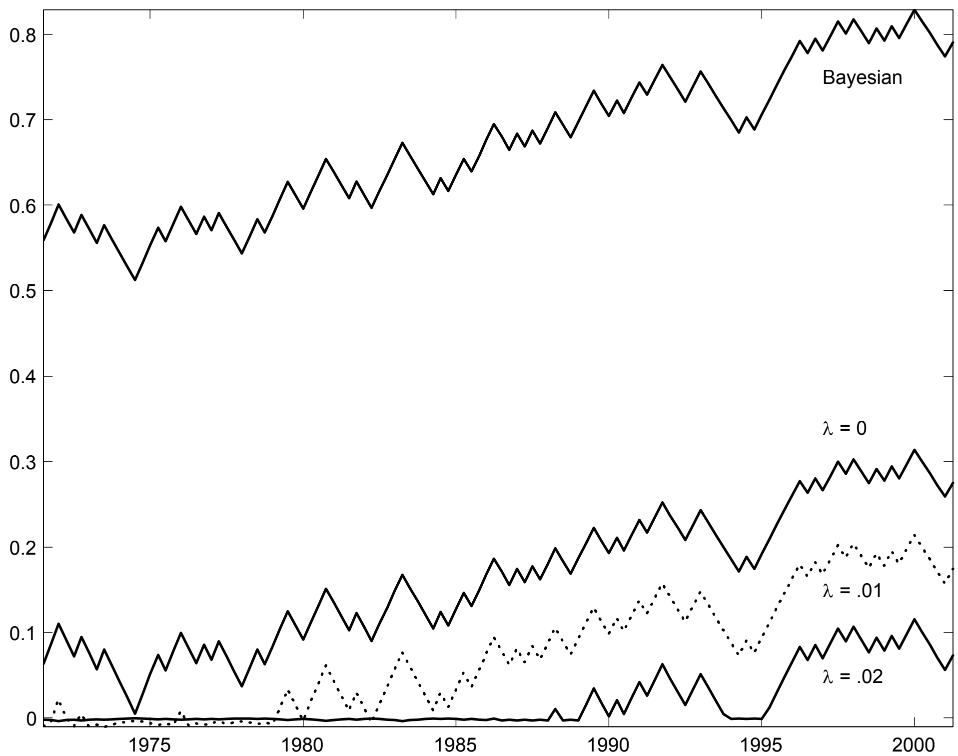

Figure 5: Optimal stock positions for ambiguity-averse agents with $\alpha = .14$ and different values of $\bar{\lambda}$ (here denoted λ) as well as for Bayesian agent with uniform prior. All planning problems are over 30 years beginning in 1971:Q3.

The plot also shows that small amounts of signal ambiguity can significantly reduce the optimal stock position. The investor with $\overline{\lambda} = .01$ already holds essentially no stocks throughout most of the 1970s. An investor with $\overline{\lambda} = .02$ does not go long in stocks until 1989. Both of these investors participate in the market in the 1970s, but spend most of the time in their contrarian region, where they take tiny short positions. As long as the investor remains in a region where he is long in stocks, changes in the ambiguity parameters α and $\overline{\lambda}$ for a given sample tend to affect the *level* of holdings, with a negligible effect on changes. Comparison of the Bayesian and ambiguity-averse solutions reveals that the supporting measure's means are essentially vertically shifted versions of each other. Of course, the Bayesian model cannot generate non-participation, so changes will look very different in states where the ambiguity averse investor moves in and out of the market.

7 RELATED LITERATURE

We are aware of only two formal treatments of learning under ambiguity. Marinacci [30] studies repeated sampling with replacement from an Ellsberg urn and shows that ambiguity is resolved asymptotically. This is a special case of our model in which signals are unambiguous. The statistical model proposed by Walley [37, pp. 457-72] differs in details from ours, but is in the same spirit; in particular, it also features what we call ambiguous signals. 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.[17]

With ambiguous signals, learning ceases without all ambiguity having been resolved. Our model thus proposes a way to model incomplete learning in complicated environments that is quite different from existing Bayesian approaches. One such approach starts from the assumption that the true data generating measure is not absolutely continuous with respect to an agent's belief.[18] This generates situations where beliefs do not converge to the truth even though agents believe, and behave as if, they will.[19] 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.

Our setup is also different in spirit from models with persistent hidden state variables, such as regime switching models. In these models, learning about the state variable never ceases because agents know that the state variable is forever changing. Agents thus track a *known* data generating process that is not memoryless. In contrast, our model applies to memoryless mechanisms. Accordingly, learning about the fixed true parameter does eventually cease, and agents know this. 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.

There exist a number of applications of multiple-priors utility or the related robust control model to portfolio choice or asset pricing. None of these is concerned with learning. Multiple-priors applications typically employ a *constant* set of one-step-ahead probabilities (Epstein and Miao [14], Routledge and Zin [33]). Similarly, existing robust control models (Hansen, Sargent, and Tallarini [23], Maenhout [28], Cagetti et al. [9]) do not allow the 'concern for robustness' to change in response to new observations. Neither is learning modeled in Uppal and Wang [35] that pursues a third approach to accommodating ambiguity or robustness.

Our paper contributes to a growing literature on learning and portfolio choice. Bawa, Brown, and Klein [7] and Kandel and Stambaugh [26] first explored the role of parameter uncertainty in a Bayesian framework.[20] Several authors have solved intertemporal portfolio choice problems with Bayesian learning.[21] The main results are conservative

[17] A similar remark applies to Huber [25], who also points to the desirability of admitting ambiguous signals and outlines one proposal for doing so.

[18] This violates the Blackwell-Dubins [8] conditions for convergence of beliefs to the truth. See Feldman [19] for an economic application.

[19] 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.

[20] There are alternatives to a Bayesian approach to the parameter uncertainty problem. See Ang and Bekaert [1] for a classical econometric strategy.

[21] Detemple [12], Gennotte [20], and Barberis [3] have considered the case of learning about mean

investment recommendations and optimal 'market timing' to hedge against changes in beliefs. While these effects reappear in our setup, our results are qualitatively different since multiple-priors preferences lead to non-participation.

Non-participation can be derived also from preferences with first-order risk aversion, as in Ang, Bekaert and Liu [2]. A key difference between first-order aversion and ambiguity aversion is that the latter allows different attitudes towards different sources of uncertainty. In applications, this is relevant for modeling participation in some assets markets and non-participation in others. Barberis and Huang [4, 5] generate selective participation by combining first-order risk aversion with narrow framing. In their model, agents derive utility directly from individual asset positions, so that risk aversion can differ by position. This is different in spirit from our model where ambiguity aversion differs by source of uncertainty – it differs by position only if different positions represent exposure to different sources of uncertainty. In addition, our learning process links ambiguity aversion directly to information: ambiguity aversion is smaller towards sources of uncertainty that are more familiar. This is consistent with evidence on investor behavior cited above.

Empirical work on stock market participation has considered agents with expected utility preferences that face a fixed per period participation cost. For example, Vissing-Jorgensen [36] estimates levels of per period fixed costs that would be required to rationalize observed participation rates in the U.S. Her approach exploits the tight link between wealth and participation predicted by fixed cost models. While modest fixed costs do help explain the lack of participation among poor households that have little financial assets in the first place, she also concludes that "it is not reasonable to claim that participation costs can reconcile the choices of all nonparticipants". This conclusion follows because participation is not as widespread among wealthy households as a fixed cost model would imply.

8 CONCLUSION

This paper has proposed a tractable framework for modelling learning under ambiguity. The key new property of learning is that confidence varies over time, together with beliefs. Ambiguous signals describe pieces of information that may reduce confidence. They also prevent ambiguity from vanishing in the long run. Our setup can thus be used to model decision making in complicated environments, where agents know that some features of the environment remain forever ambiguous, whereas others can be learnt over time. Asset markets are one example of such an environment, but the structure is likely to be useful in other contexts as well.

Our results on optimal portfolio choice suggest that learning under ambiguity could be a building block in a successful model of the cross section of asset holdings. While

returns, while Barberis [3] and Xia [38] have studied learning about predictability. See Ang and Bekaert [1] for portfolio choice in a regime-switching model.

more work is required to distinguish an ambiguity aversion story from one based on a technological participation cost, it is already clear that the two models have very different implications. For example, consider the issue of investing a social security fund in the stock market. If the participation cost is technological, then the government could reduce it by exploiting economies of scale. In contrast, if non-participation is due to ambiguity underlying preferences, then agents could not gain from being forced to invest.

Participation rules driven by learning from past returns may be relevant also for understanding the dynamics of selective participation in particular markets. For example, it has been noted that the decline in equity home bias ceased in the late 1990s although transaction costs were reduced further. Increases in the implicit participation cost induced by ambiguity after recent financial crises abroad provide a candidate explanation. Other examples include the increased participation by small investors in "hot" stocks or mutual funds after a period of high returns.

A APPENDIX

Proof of Theorem 1. For any sequence $s^\infty = (s_1, s_2, ...)$, 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 $\phi_t \to \phi$; this set has measure one under the truth. Fix a sequence $s^\infty \in \Omega$. For any likelihoods $(\ell_s)_{s\in S}$ and $\theta \in \Theta$, and for any probability measure λ on S, define

$$\widetilde{H}\left(\lambda, (\ell_s), \theta\right) = \sum_{s\in S} \lambda(s) \log \ell_s(s|\theta),$$

Below we take λ to be ϕ_t or ϕ.

Given μ_0 and the likelihood sequence ℓ^t, then the data density for the first t periods is

$$\Pr(s^t; \mu_0, \ell^t) = \sum_{\theta\in\Theta} \mu_0(\theta) \prod_{j=1}^{t} \ell_j(s_j|\theta).$$

In choosing a likelihood sequence that maximizes $\Pr(s^t; \mu_0, \ell^t)$, it is wlog to focus on sequences such that $\ell_j = \ell_k$ if $s_j = s_k$. Any such likelihood sequence can be identified with a collection $(\ell_s)_{s\in S}$ and we can write

$$\max_{\ell^t} \Pr(s^t; \mu_0, \ell^t) = \max_{(\ell_s)} \sum_{\theta\in\Theta} \mu_0(\theta)\ e^{t\widetilde{H}(\phi_t, (\ell_s), \theta)}.$$

By definition of H and the identification condition, there exists $\epsilon > 0$ such that

$$max_{(\ell_s)}\, \widetilde{H}\left(\phi, (\ell_s), \theta\right) \leq H(\theta^*) - \epsilon, \quad \text{for all } \theta \neq \theta^*.$$

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

$$max_{(\ell_s)}\, \widetilde{H}\left(\phi_t, (\ell_s), \theta\right) \leq max_{(\ell_s)}\, \widetilde{H}\left(\phi_t, (\ell_s), \theta^*\right) - \epsilon, \tag{18}$$

for all $\theta \neq \theta^*$ and $t > T$.

We claim that

$$\lim_{t\to\infty}\left(\max_{\mu_0,\ell^t}\Pr(s^t;\mu_0,\ell^t)\right)^{\frac{1}{t}} = e^{H(\theta^*)},$$

or equivalently, that

$$\left[\frac{\max_{\mu_0,\ell_1^t}\Pr(s^t;\mu_0,\ell^t)}{e^{t\max_{(\ell_s)}\widetilde{H}(\phi_t,(\ell_s),\theta^*)}}\right]^{\frac{1}{t}} \longrightarrow 1. \tag{19}$$

Rewrite the latter in the form

$$\left[max_{\mu_0,\ell^t}\sum_{\theta\in\Theta}\mu_0(\theta)\ e^{t\eta_t\left(\theta,\mu_0,\ell^t\right)}\right]^{\frac{1}{t}} \longrightarrow 1, \quad \text{where} \tag{20}$$

$$\eta_t\left(\theta,\mu_0,\ell^t\right) = \frac{1}{t}\Sigma_{j=1}^t log\ell_j(s_j \mid \theta) - \max_{(\ell_s)}\widetilde{H}(\phi_t,(\ell_s),\theta^*).$$

From (18), deduce that

$$\begin{aligned}
&max_{\mu_0,\ell^t}\mu_0(\theta^*)\,e^{\eta_t\left(\theta^*,\mu_0,\ell^t\right)}\\
&\leq max_{\mu_0,\ell^t}\sum_{\theta\in\Theta}\mu_0(\theta)\ e^{t\eta_t\left(\theta,\mu_0,\ell^t\right)}\\
&\leq max_{\mu_0,\ell^t}\mu_0(\theta^*)\ e^{t\eta_t\left(\theta^*,\mu_0,\ell^t\right)} + (1-\mu_0(\theta^*))\,e^{-\epsilon t},
\end{aligned}$$

for all $t > T$. But

$$\left[max_{\mu_0,\ell^t}\,\mu_0(\theta^*)\ e^{t\eta_t\left(\theta^*,\mu_0,\ell^t\right)}\right]^{\frac{1}{t}} \longrightarrow 1,$$

which proves (20).

Now consider any admissible 'theory' (μ_0,ℓ^t). By the definition of $\mathcal{M}_t^\alpha$, (μ_0,ℓ^t) must satisfy

$$\left(\max_{\mu_0,\ell^t}\Pr(s^t;\mu_0,\ell^t)\right)^{\frac{1}{t}} \geq \left(\Pr(s^t;\mu_0,\ell^t)\right)^{\frac{1}{t}} \geq \alpha^{\frac{1}{t}}\left(\max_{\mu_0,\ell^t}\Pr(s^t;\mu_0,\ell^t)\right)^{\frac{1}{t}}.$$

Thus (19) implies that

$$\left[\frac{\Pr(s^t;\mu_0,\ell^t)}{e^{t\max_{(\ell_s)}\widetilde{H}(\phi_t,(\ell_s),\theta^*)}}\right]^{\frac{1}{t}} \longrightarrow 1. \tag{21}$$

for *any* admissible theory.

The posterior derived from (μ_0,ℓ^t) satisfies

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

here and below we suppress the dependence of η_t on (μ_0, ℓ^t) because the latter is fixed. Thus we are done if we can show that

$$\sum_{\theta \neq \theta^*} \mu_0(\theta) e^{t(\eta_t(\theta) - \eta_t(\theta^*))} \to 0.$$

This follows from two claims.

Claim 1: For any $\epsilon > 0$ and all $\theta \neq \theta^*$, $\eta_t(\theta) \leq -\epsilon$ for all $t > T(\epsilon)$. To maximize $\frac{1}{t}\Sigma_{j=1}^t log \ell_j(s_j \mid \theta)$, it is wlog to focus on sequences such that $\ell_j = \ell_k$ if $s_j = s_k$. Therefore,

$$\frac{1}{t}\Sigma_{j=1}^t log \ell_j(s_j \mid \theta) \leq max_{(\ell_s)} \widetilde{H}(\phi_t, (\ell_s), \theta)$$

The claim follows from (18).

Claim 2: $\eta_t(\theta^*) \to 0$. By construction, $\eta_t(\theta^*) \leq 0$. Suppose that $\eta_t(\theta^*) < -\delta$ for some δ and all $t > T$. Then claim 1 (with $\epsilon = \delta$) implies that

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

for all sufficiently large t, contradicting (21). ■

Proof of (11). Refer to the i.i.d. measure with one-period distribution ϕ as the 'truth'. Every posterior in $\mathcal{M}_t^0(s^t)$ corresponds to some μ_0 and ℓ_1^t :

$$\mu_t\left(\theta^* | s^t, \mu_0, \ell_1^t\right) = \frac{\mu_0(\theta^*) \prod_{j=1}^t \ell_j(s_j|\theta^*)}{\sum_{\theta \in \Theta} \mu_0(\theta) \prod_{j=1}^t \ell_j(s_j|\theta)} \tag{22}$$

$$= \frac{1}{1 + \sum_{\theta \in \Theta} \frac{\mu_0(\theta)}{\mu_0(\theta^*)} \exp\left(t\left(\frac{1}{t}\sum_{j=1}^t \log \frac{\ell_j(s_j|\theta)}{\ell_j(s_j|\theta^*)}\right)\right)}$$

By (11), then a.s. under the truth and for all $\theta \neq \theta^*$,

$$\lim_{t \to \infty} \max_{\ell_1^t} \left(\frac{1}{t}\sum_{j=1}^t \log \frac{\ell_j(s_j|\theta)}{\ell_j(s_j|\theta^*)}\right) = \lim_{t \to \infty} \frac{1}{t}\sum_{j=1}^t \max_{\ell} \log \frac{\ell(s_j|\theta)}{\ell(s_j|\theta^*)}$$

$$= \sum_{s \in S} \phi(s) \max_{\ell} \log \frac{\ell(s|\theta)}{\ell(s|\theta^*)} < 0,$$

where the last equality follows from the law of large numbers because the stochastic process $\max_\ell \log \frac{\ell(s_t|\theta)}{\ell(s_t|\theta^*)}$ is i.i.d. under the truth. It follows that the sum in the denominator of (22) converges to 0 a.s. under the truth and hence that $\mu_t(\theta^* | s^t, \mu_0, \ell_1^t) \to 1$.

Proof of (16). Write the likelihood of a sample s^t under some theory, here identified with a pair (θ, λ^t), as

$$L\left(s^t, \theta, \lambda^t\right) = \prod_{j=1}^t (\theta + \lambda_j)^{s_j} (1 - \theta - \lambda_j)^{1 - s_j}. \tag{23}$$

Let $\tilde{\lambda}^t$ denote the sequence that maximizes (23) for fixed θ. This sequence is independent of θ and has $\tilde{\lambda}_j = \overline{\lambda}$ if $s_j = 1$ and $\tilde{\lambda}_j = -\overline{\lambda}$ if $s_j = 0$, for all $j \leq t$. It follows that $L\left(s^t, \theta, \tilde{\lambda}^t\right)$ depends on the sample only through the fraction ϕ_t of high returns observed. The set $\mathcal{M}_t^\alpha$ can be expressed in terms of $L\left(s^t, \theta, \tilde{\lambda}^t\right)$, because $\theta \in \mathcal{M}_t^\alpha$ if and only if the theory $\left(\theta, \tilde{\lambda}^t\right)$ passes the likelihood ratio criterion. Indeed, if $\theta \in \mathcal{M}_t^\alpha$, then there exists *some* λ^t such that the theory $\left(\theta, \lambda^t\right)$ passes the criterion. Thus $\left(\theta, \tilde{\lambda}^t\right)$ must also pass it, since its likelihood is at least as high. In contrast, if $\theta \notin \mathcal{M}_t^\alpha$, then there is *no* λ^t such that the theory $\left(\theta, \lambda^t\right)$ passes the criterion. Finally, one can use

$$g(\theta, \phi_t) = \frac{1}{t} \log L\left(s^t, \theta, \tilde{\lambda}^t\right)$$

to express the criterion in (16).

References

[1] A. Ang and G, Bekaert, International asset allocation with regime shifts, Columbia, 2001.

[2] A. Ang, G. Bekaert, and J. Liu, Why stocks may disappoint, Columbia, 2002.

[3] N. Barberis, Investing for the long run when returns are predictable, *J. Finance* 55 (2000), 225-264.

[4] N.Barberis and M. Huang, Mental Accounting, Loss Aversion and Individual Stock Returns, *J. Finance* 56 (2001), 1247-1292.

[5] N.Barberis and M. Huang, The Loss Aversion / Narrow Framing Approach to Stock Market Pricing and Participation Puzzles, U. Chicago, 2004.

[6] T.F. Bewley, Knightian decision theory and econometric inference, Cowles #868, 1988.

[7] V.S. Bawa, S.J. Brown and R.W. Klein, *Estimation Risk and Optimal Portfolio Choice*, Amsterdam: North Holland, 1979.

[8] D. Blackwell and L. Dubins, Merging of opinions with increasing information, *Annals of Math. Statistics* 33 (1962), 882-886.

[9] M. Cagetti, L.P. Hansen, T. J. Sargent and N. Williams, Robustness and pricing with uncertain growth, *Rev. Financ. Stud.* 15 (2002), 363-404.

[10] Z. Chen and L.G. Epstein, Ambiguity, risk and asset returns in continuous time, *Econometrica* 70 (2002), 1403-43.

[11] D. Chew, *The New Corporate Finance*, 2nd edition, McGraw Hill, 1999.

[12] J. Detemple, Asset pricing in a production economy with incomplete information, *J. Finance* 41 (1986), 383-390.

[13] J. Dow and S. Werlang, Uncertainty aversion, risk aversion and the optimal choice of portfolio, *Econometrica* 60 (1992), 197-204.

[14] L.G. Epstein and J.J. Miao, A two-person dynamic equilibrium under ambiguity, *J. Econ. Dyn. and Control* 27 (2003), 1253-1288.

[15] L.G. Epstein and M. Schneider, Recursive multiple-priors, *J. Econ. Theory* 113 (2003), 1-31.

[16] L.G. Epstein and M. Schneider, IID: independently and indistinguishably distributed, *J. Econ. Theory* 113 (2003), 32-50.

[17] L.G. Epstein and M. Schneider, Ambiguity, Information Quality, and Asset Pricing, 2004.

[18] L.G. Epstein and T. Wang, Intertemporal asset pricing under Knightian uncertainty, *Econometrica* 62 (1994), 283-322.

[19] M. Feldman, On the generic nonconvergence of Bayesian actions and beliefs, *Econ. Theory* 1 (1991), 301-321.

[20] G. Gennotte, Optimal portfolio choice under incomplete information, *J. Finance* 41 (1986), 733-746

[21] I. Gilboa and D. Schmeidler, Maxmin expected utility with nonunique prior, *J. Math. Econ.* 18 (1989), 141-153.

[22] L. Guiso, M. Haliassos, and T. Jappelli, *Household Portfolios,* MIT Press, 2002.

[23] L.P. Hansen, T.J. Sargent and T. Tallarini, Robust permanent income and pricing, *Rev. Ec. Stud.* 66 (1999), 873-907.

[24] L.P. Hansen, T.J. Sargent and N. Wang, Robust permanent income and pricing with filtering, *Macroeconomic Dynamics* 6 (2002), 40-84.

[25] P.J. Huber, The use of Choquet capacities in statistics, *Bull. Internat. Statist. Inst.* 45 (1973), 181-188.

[26] S. Kandel and R. Stambaugh, On the Predictability of Stock Returns: An Asset Allocation Perspective, *J. Finance,* 51 (1996), 385-424.

[27] D.M. Kreps, *Notes on the Theory of Choice*, Westview, 1988.

[28] P. Maenhout, Robust portfolio rules and asset pricing, INSEAD, 1999.

[29] R. Marimon, Learning from learning in economics, pp. 278-315 in *Advances in Economics and Econometrics: Theory and Applications,* D. Kreps and K. Wallis eds., Cambridge U. Press, 1997.

[30] M. Marinacci, Learning from ambiguous urns, *Statistical Papers* 43 (2002), 145-151.

[31] R. Merton, Optimal Consumption and Portfolio Rules in a Continuous Time Model, *J. Econ. Theory* 3 (1971), 373-413.

[32] S. Mukerji and J. Tallon, Ambiguity aversion and incompleteness of financial markets, *Rev. Econ. Stud.* 68 (2001), 883-904.

[33] B. Routledge and S. Zin, Model uncertainty and liquidity, NBER #8683, 2001.

[34] M.J. Schervish, *Theory of Statistics*, Springer, 1995.

[35] R. Uppal, and T. Wang, Model misspecification and under-diversification, MIT, 2001.

[36] A. Vissing-Jorgensen, Perspectives on behavioral finance: does "irrationality" disappear with waelth? Evidence from expectations and actions, in *NBER Macro Annual,* 2003.

[37] P. Walley, *Statistical Reasoning with Imprecise Probabilities*, Chapman and Hall, 1991.

[38] Y. Xia, Learning about predictability: the effect of parameter uncertainty on dynamic asset allocation, *J. Finance* 56 (2001), 205-246.