

Rochester Center for
Economic Research

Learning under Ambiguity

Epstein, Larry G., and Martin Schneider

Working Paper No. 497
March 2005

UNIVERSITY OF
ROCHESTER

LEARNING UNDER AMBIGUITY*

Larry G. Epstein

Martin Schneider

March 12, 2005

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 model of learning in complicated environments. The model allows decision-makers' confidence about the environment to change – along with beliefs – as they learn. Moreover, confidence need never become strong enough for decision-makers to identify a single stochastic process that is generating the data. The model is applied to dynamic portfolio choice. A calibrated example shows how learning about ambiguous returns leads to endogenous stock market participation costs that depend on past market performance. In addition, hedging of ambiguity provides a new reason why the investment horizon matters for portfolio choice.

1 INTRODUCTION

Models of learning typically assume that agents assign (subjective) probabilities to all relevant uncertain events. 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 and it provides a generalization of the Bayesian learning model that can accommodate it.

*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.

Our starting point is *recursive multiple-priors* utility, a model of utility axiomatized in Epstein and Schneider [13], that extends Gilboa and Schmeidler’s [19] 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 learning when data are viewed as being generated by the same memoryless mechanism in every period; this *a priori* view is commonly imposed on learners in a wide variety of economic applications. To illustrate the model in an economic setting, we solve the intertemporal portfolio choice problem of an investor who learns about stock returns that she perceives as ambiguous.

Learning under ambiguity has two interesting properties that cannot be captured by the standard Bayesian model of learning. First, confidence varies over time, together with beliefs. For example, in our portfolio choice model, beliefs and confidence are represented by an evolving *set* of conditional distributions over future stock returns. This set may shrink over time as agents become more familiar with the environment, but it may also expand if new information is surprising relative to past experience. Optimal portfolio choice then leads to endogenous market entry and exit rules that respond to past market performance.

The second important property of our model is that agents need not be certain ex ante that they will identify one “true” stochastic process in the long run. This possibility is relevant when agents take the complexity of the environment seriously and therefore are less ambitious than Bayesians about how much can be learned. To capture this attitude, our model allows for ambiguity that does not vanish in the long run. Agents in our model may learn only some features of the environment, while others remain forever ambiguous. They are aware of their own limitations and take them into account when making decisions. For example, a belief that the mean equity premium will never be known precisely, due to ongoing change in the economy for example, will lead them to be more cautious stock market investors. Asset markets are one example of an environment where such behavior appears relevant, but the structure is likely to be useful in other contexts as well.

The paper is organized as follows. The rest of this introduction outlines the benchmark Bayesian model of learning, then our generalization, and finally our application to portfolio choice. Section 2 presents a sequence of thought experiments to motivate our model. Section 3 briefly reviews recursive multiple-priors utility and then introduces the learning model. Section 4 establishes properties of learning in the short and long run. Section 5 applies our model to portfolio choice. A discussion of related literature follows in Section 6. The Appendix contains proofs.

1.1 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. For a given parameter value $\theta \in \Theta$, the data are an independent and identically distributed sequence of signals $\{s_t\}_{t=1}^\infty$, where the distribution of any signal s_t is described by the probability measure $\ell(\cdot|\theta)$ on the period state space $S_t = S$. The triple (Θ, μ_0, ℓ) is the decision-maker’s theory of how data are generated. This theory incorporates both prior information (through μ_0) and a view of how the signals are related to the underlying parameter (through ℓ). Beliefs on (payoff-relevant) states are equivalently represented by a probability measure p on *sequences* of signals (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

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

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

Foundations for this model are provided by the de Finetti Theorem: the representation (1) is possible if and only if the prior p on S^∞ is exchangeable. Nevertheless, there is a sense in which applications of the Bayesian model are still ad hoc. While the de Finetti theorem begins with a prior p and the assumption of exchangeability, its usefulness stems in large part from the reverse perspective as describing a cognitively plausible procedure for a decision-maker to arrive at a prior p (and to determine thereby if, indeed, exchangeability is acceptable).¹ In some settings, there is a natural pair (Θ, ℓ) and, when Θ is simple, forming a prior μ_0 over Θ is easier than forming one directly over S^∞ . Thus the representation (1) can be viewed as providing a procedure for the decision-maker to arrive at the measure p on S^∞ to use for decision-making. Moreover, to the extent that (Θ, ℓ) is simple and natural given her environment, the procedure is cognitively plausible.

We emphasize that cognitive plausibility is essential to support application of the Bayesian model: the typical application assumes a particular parametric specification of (Θ, μ_0, ℓ) , which amounts, via (1), to assuming a particular prior p . There is no axiomatic justification for this step - cognitive plausibility is the only supporting argument. However, outside of contrived settings, where the data generating mechanism is simple and transparent, the existence of a cognitively plausible specification of (Θ, μ_0, ℓ) is problematic. In particular, the Bayesian model presumes a degree of confidence on the part of the agent that seems unintuitive in complicated environments.

The examples in Section 2 illustrate two aspects of “lack of confidence”. On the one hand, a decision-maker may lack confidence in her *initial* information about the environment. In light of the Ellsberg Paradox, the Bayesian model is then not applicable. However, she may still expect to resolve initial ambiguity through learning and eventually become confident enough to view the data as a simple i.i.d. process. Alternatively, she may expect to *never* learn enough to confidently view the data generating process as an i.i.d. process. Indeed, in contrast to a Bayesian, she may not even be sure whether the

¹See, for example, Kreps [26, pp. 155-6].

empirical frequencies of the data will converge, let alone expect her learning process to settle down at a single i.i.d. process.²

As a concrete example, consider independent tosses of a sequence of coins. If the decision-maker perceives the coins to be identical, then after many tosses she will naturally become confident that the observed empirical frequency of heads is close to a “true” frequency of heads that is relevant for forecasting what happens in future tosses. This laboratory-style situation is captured by the Bayesian model.³ More generally, suppose that she believes the tosses to be independent, but that she has no reason to be sure that the coins are identical - for example, she might be told the same about each coin but very little (or nothing at all) about each. Then she would plausibly *admit the possibility* that the coins are not identical. In particular, there is no longer a compelling reason why data in the future should be i.i.d. with frequency of heads equal to the empirical frequency of heads observed in the past. This situation cannot be captured by the Bayesian model (1).⁴ One can view our model as an attempt to capture learning in such complicated environments.

1.2 Our Model

Generalize the Bayesian triple to $(\Theta, \mathcal{M}_0, \mathcal{L})$. As before, Θ is a parameter space. The *set* of priors \mathcal{M}_0 represents the agent’s initial view of these parameters. When \mathcal{M}_0 is not a singleton, it also captures (lack of) confidence in the information upon which this initial view is based. Finally, a *set* of likelihoods \mathcal{L} represents the agent’s *a priori* view of the connection between signals and the true parameter. These components are combined in a way described below, and paralleling the Bayesian model (1), to define the set of priors \mathcal{P} appearing in recursive multiple-priors utility. In particular, we specify sets of posteriors \mathcal{M}_t that update beliefs about the parameter, including the evolution of confidence. These sets of posteriors may shrink in response to a signal, if it confirms current beliefs and thus enhances confidence, or expand, if the signal is an outlier relative to current beliefs and reduces confidence.

The multiplicity of likelihoods and the distinction between Θ and \mathcal{L} are central to the way we model an agent who has modest (or realistic) ambitions about what she can learn. The set Θ represents features of the environment that the agent views as constant over time and that she therefore expects to learn. However, in complicated environments she may be wary of a host of other poorly understood or unknown factors that also affect realized signals. They vary over time in a way that she does not understand well

²For the exchangeable Bayesian model, convergence of empirical frequencies is a consequence of the LLN for exchangeable random variables; see Schervish [34].

³Under regularity conditions, the posterior μ_t converges to a Dirac measure on Θ , almost surely under p , the Bayesian’s subjective probability on sequences of data; see Schervish [34].

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

enough even to theorize about and therefore she does not even try to learn about them. The multiplicity of likelihoods in \mathcal{L} captures these factors. Under regularity conditions, we show that beliefs become concentrated on one parameter value in Θ , but the time-varying unknown features remain impossible to know even after many observations. Thus ambiguity does not vanish in the long run. Instead, the agent moves towards a state of time-invariant ambiguity, where she has learned all that she can.⁵

Though recursive multiple-priors is an axiomatic model, we do not have axiomatic foundations for the specialization described here. But as argued above, there is a sense in which applications of the Bayesian model are ad hoc in spite of the axiomatic foundations provided by the de Finetti Theorem, and cognitive plausibility is a concern. We view enhanced cognitive plausibility as a justification for our model. As illustrated by the examples in Section 2, greater cognitive plausibility is afforded if we allow the decision-maker to have multiple priors on the statistical parameter space and, more importantly, multiple likelihoods.

Another feature of our model should be made clear at the outset: the decomposition between the learnable parameters Θ and the unlearnable features of the environment embodied by a nonsingleton \mathcal{L} is exogenous. We do not explain which features the agent tries to learn in any given setting. Note that an exogenous decomposition is imposed also by the Bayesian model, where the decomposition is extreme - \mathcal{L} is a singleton and the agent tries to identify a true i.i.d. process. Thus, while we leave important aspects unmodeled, we do extend the Bayesian model in the direction of permitting more plausible ambitions about what can be learned in complex settings. The Ellsberg-style examples and the portfolio choice application illustrate that natural decompositions may be suggested by descriptions of the environment.

1.3 Stock Market Participation

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.⁶ 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. The model delivers

⁵Existing 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 [14] for further properties of time invariant ambiguity.

⁶See Guiso and Haliassos [20] 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 [11]. The argument is discussed in detail in Section 5.

endogenous market exit and entry rules that depend on past market performance as well as on the planning horizon.

We view excess stock returns as a data series that is generated by a memoryless mechanism. We model investors who perceive the mechanism driving returns as ambiguous. They try to resolve ambiguity about the equity premium by looking at the distribution of past returns. However, they are aware that they have only a limited understanding of the data generating mechanism, and that there are factors influencing returns about which they cannot even theorize. As a result, they are not confident that they will ever resolve all ambiguity about the equity premium. We propose a way for investors to quantify their degree of ambiguity. We then calculate real time optimal portfolios for the post-war period. For U.S. investors with a moderate degree of ambiguity aversion, the model recommends that they hold virtually no stocks in the 1970's, entering the market only in the late 1980's.

We also 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. Her 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 [30]. Indeed, we show that it arises even when preferences over risky payoff streams are logarithmic, so that the traditional hedging demand is zero.

2 EXAMPLES

In this section, we consider three related learning scenarios. Each involves urns containing five balls that are either black or white. However, the scenarios differ in a way intended to illustrate alternative hypotheses about learning in economic environments, and to motivate the main ingredients of our model. Figure 1 illustrates the composition of the urns.

Scenario 1: Pure risk

An urn contains five balls. The agent is told that two are black, two are white and that the fifth is either black or white depending on the outcome of the toss of an unbiased coin - the 'coin ball' is black if the coin toss produces heads and it is white otherwise. She cannot see the outcome of the coin toss, but she does observe some possibly related signals, namely, in each period she sees a black or white ball. Suppose that she anticipates all this and that her ex ante view is that these data are generated by sampling with replacement from the given urn. This may be because she is actually doing the sampling, or because she has been told that this is how signals are generated. Alternatively, it may simply reflect her subjective view of the environment. In any case,

the obvious version of the Bayesian learning model (1) seems appropriate. In particular, she can be confident of learning the color of the coin ball.

It is useful to compare the preceding with the following alternative scenario (referred to later as Scenario 1'): signals are generated by successive sampling from a *sequence* of urns, each containing black and white balls. Each urn is perceived as consisting of a coin ball and four other balls as above. The coin is tossed once at the start and determines the same color for coin-balls in all urns. However, non-coin balls are replaced every period with a new set of non-coin balls. Each set is determined by a new realization of an i.i.d. process that puts probability 1/2 each on 2 and 3 black non-coin balls.

At first sight, this scenario looks more complicated: the number of non-coin balls in any given urn is not known and cannot be learned because it changes over time. However, the fact that there is a known identical distribution over the number of black non-coin balls at all dates implies that Scenario 1' is effectively the same as Scenario 1: in both cases, the probability that a randomly drawn non-coin ball is black is 1/2 at all dates. This is all that matters for updating beliefs about what the next draw will be. Decisions that depend on the conditional probability of the next ball drawn, such as betting on the color of the latter, should thus be identical in the two scenarios. Similarly, in both cases beliefs about the total number of black balls in each urn will settle down at the long run empirical frequency of black draws.

The difference between the two alternative perceptions of signals - they are generated by sampling with replacement from a fixed urn on the one hand versus being due to sampling from different urns on the other - will play an important role below in the intuition for our model of learning under ambiguity. The fact that there is no meaningful difference between them here is a (restrictive) feature of the Bayesian model.

Scenario 2: Ambiguous priors and fixed urn

The information provided about the urn is modified to: there is at least one non-coin ball of each color. Thus unlike Scenario 1, no objective probabilities are given for the composition of the urn – there is ambiguity about the number of black non-coin balls. However, signals are again observed in each period and they are viewed as resulting from sampling with replacement from the given urn. Because the agent views the factors underlying signals, namely the color composition of the urn, as fixed through time, she can hope to learn it. That is, she would try to learn the true $\theta \in \Theta = \{\frac{1}{5}, \frac{2}{5}, \frac{3}{5}, \frac{4}{5}\}$, where the connection to signals is expressed by the single likelihood

$$\ell(B | \theta) = \theta.$$

Compare this situation with the purely risky one, Scenario 1 above, assuming the two coin tosses are independent. In particular, consider the choice between betting on drawing black (white) from the risky urn as opposed to the ambiguous one. (Throughout a *bet on the color of a ball* drawn from an urn is understood to pay one util if the ball has the desired color and zero otherwise.) The familiar Ellsberg Paradox examines this choice




Figure 1: Urns for Scenarios 1-3.

prior to any sampling. The 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 and hence also with any single prior on Θ . The intuition is that confidence about the odds differs across the two urns. This difference matters for behavior, but cannot be captured by probabilities. On the other hand, the multiple-priors model predicts the intuitive choices if the set \mathcal{M}_0 of priors on Θ admits the possibility that the number of black non-coin balls is less than 2 and similarly for white. Thus we are led to the triple $(\Theta, \mathcal{M}_0, \ell)$, which differs from the Bayesian model only in that \mathcal{M}_0 is a nonsingleton set.

Consider the ranking of bets after some sampling. Assuming that the agent's prior view of signals is correct, actual learning proceeds in each case as though sampling is from an unchanging urn. Thus it is intuitive that for the ambiguous urn, learning should resolve ambiguity in the long run. As the number of draws increases, ambiguity should diminish and asymptotically the agent should behave as if she knew the fraction of black balls were equal to their empirical frequency - in the limit, and for any common sample history across the two urns, she would be indifferent between betting on the next draw from the risky as opposed to the ambiguous urn. Formally, increasing confidence about the ambiguous urn is captured in our model because the set \mathcal{M}_t of posteriors shrinks with time.

Both Scenarios 1 and 2 are conducive to inference and learning because there is one urn with fixed composition. Next we examine learning in a more complex setting where signals are generated by a sequence of urns with time-varying composition. This added complexity was shown above to be of no consequence in the Bayesian model, but it will matter more generally.

Scenario 3: Unambiguous prior and changing urns

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

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

Consider first betting behavior after a small number of draws, assuming that the coin tosses in Scenarios 2 and 3 are independent. Before any sampling, one should be indifferent between bets in Scenarios 2 and 3, since the two settings are *a priori* identical. However, suppose that one black draw from each urn has been observed and consider bets on the next draw being black. It is intuitive that the agent prefer again to bet on the next urn in Scenario 2 because the black draw is stronger evidence there for the next draw also being black: in Scenario 3, the agent does not understand well the dynamics of the non-coin balls and thus is plausibly concerned that the next urn may contain more white balls than did the first. She thinks that it may also contain more black balls, but as in the static Ellsberg choice problem, the more pessimistic possibility outweighs the optimistic one. Thus she would rather bet in Scenario 2 where the number of black non-coin balls is perceived not to change between urns. Our model can accommodate this preference (see Section 4.1).

How does an agent expect to behave in the long run? A Bayesian might postulate a family of stochastic processes that describe the evolution of the non-coin balls over time. She would then hope to learn moments of that process, such as the fraction of black non-coin balls placed in the urn by the average administrator. However, any such approach makes sense only if the agent is confident that the fraction of black non-coin balls will converge to some limit. The description (or perception) of the environment provides no indication that this will happen. To the contrary, the ever-changing sequence of administrators suggests perpetual change.

Our model thus allows the agent to view the number of black non-coin balls as inde-

pendent across urns. Under this view, she does not expect to learn anything about future non-coin balls from sampling. Given how little information is provided about any single urn, it is then natural for her to perceive the number of non-coin balls as ambiguous *at all dates*. In particular, she expects that, after sufficiently long identical samples from the two scenarios, she will still prefer to bet on the next draw being black from the urns in Scenario 2 rather than in Scenario 3.

Multiple likelihoods to model changing urns

Our model of Scenario 3 builds on the Bayesian model in that it also starts from a parameter space – here $\Theta = \{B, W\}$, corresponding to the two possible colors of the coin-ball – and a prior μ_0 , here the $(\frac{1}{2}, \frac{1}{2})$ prior on Θ corresponding to the toss of the unbiased coin. In addition, ambiguity perceived about the changing non-coin balls motivates a *set of likelihoods* \mathcal{L} . The general idea is that of an agent who perceives the data generating mechanism as having some unknown features that are common across a set of experiments, represented by θ . At the same time, other unknown features are variable across experiments, represented by \mathcal{L} . It is known that the variable features are determined by a memoryless mechanism – this is why the set \mathcal{L} does not depend on history. It is also known that the relative importance of the variable features is the same every period, as in the urn example where there are always four non-coin balls and one coin ball. This is why the set \mathcal{L} does not depend on time.

In the urns example, the set of likelihoods might be specified as follows: let λ denote the number of black non-coin balls. Then the likelihood of drawing a black ball from that urn is given by

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

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

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

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

In the special case $\epsilon = 1$, the set \mathcal{L}_ϵ models an agent who attaches equal weight to all logically possible urn compositions $\lambda = 1, 2$, or 3 . More generally, (2) 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 cases $\lambda = 1, 3$. The parameter ϵ can be interpreted as the weight attached to the latter cases, as opposed to the focal likelihood. Indeed, \mathcal{L}_ϵ can be rewritten as⁷

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

⁷This 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 [16], for example.

We emphasize that the agent behaves as though she views any likelihood in \mathcal{L} as applying at any future time. In particular, she admits the possibility that the operative likelihood varies from period to period. Moreover, the way in which it can vary is constrained only by the requirement that every likelihood lies in \mathcal{L} . This view is perfectly consistent with her limited information about the urns. The description of the environment does not give any indication of patterns in the sequence of urns. In particular, past draws do not provide any information about the non-coin balls in future urns. The composition of the non-coin balls in the next urn is thus *always* ambiguous, even after a large number of draws.⁸

It follows that Ellsberg-type choices in betting on the color of the next draw should persist forever, in contrast to what we have observed for Scenario 2. For example, after a sufficiently long sample, an agent will prefer to bet on the next draw being black from the urns in Scenario 2 rather than in Scenario 3. The model thus rationalizes a preference for bets in Scenario 2 over those in Scenario 3 in the long run. Section 4.1 below will show that it also captures the intuitive preference for bets in Scenario 2 in the short run.

Updating

To motivate how we model behavior in the short run, turn to the agent’s view about θ after a finite number of draws. Intuitively, the fraction of black draws observed in the past is informative about the parameter: the more black draws are observed, the more confident she should become that the coin ball is black. How quickly confidence will change now also depends on what she believes, with hindsight, about the sequence of non-coin balls in past urns. Importantly, while it is not possible to learn about *future* non-coin balls, past draws are certainly informative about *past* non-coin balls.⁹ For example, if a black ball was drawn two periods ago, it is more likely, with hindsight, that the urn from which draws were made at that date contain three black non-coin balls.

More generally, any view one might have about the past sequence of non-coin balls can be evaluated in light of a sample $s^t = (s_1, \dots, s_t)$. In Scenario 3, such reevaluation of the non-coin balls is important for updating about the parameter θ . To illustrate, suppose 50 draws have been observed, and 30 of them have been black. Given a theory about the non-coin balls, these observations suggest that the coin ball is more likely to be black than white. At the same time, given a value of the parameter, the observations suggest that more than half of the past urns contained three black non-coin balls. Sensible updating takes both of these features into account. Our approach is based on selecting theories about the data described by a prior and a sequence of likelihoods, and then taking conditionals of theories that explain the given sample well.

⁸Scenario 3 is thus similar to Scenario 1’ above in that the agent is a priori certain that she will never learn more about future non-coin balls than what she already knows *a priori*. The difference is that initial beliefs in Scenario 1’ are captured by a single probability measure for the non-coin balls, while in Scenario 3 the non-coin balls are viewed as ambiguous.

⁹In this respect, Scenario 3 again resembles Scenario 1’. In Scenario 1’, one can infer something about past realizations of the i.i.d. non-coin ball process from past draws. At the same time, the conditional distribution of future non-coin balls is always the same.

More specifically, consider the following heuristic reasoning. There is a large set of “theories” about sequences of data that the agent might be willing to entertain. Every theory is described by a probability measure on S^∞ of the type

$$\int \prod_{j=1}^{\infty} \ell_j(\cdot|\theta) d\mu_0(\theta), \quad (3)$$

where $\mu_0 \in \mathcal{M}_0$ is a prior and $\{\ell_t\} \in \mathcal{L}^\infty$ is a sequence of likelihoods. Here every theory says that the data are serially independent conditional on θ , but many allow for nonstationarity, although there are also theories that say the data are conditionally i.i.d. We define the agent’s set of one-step-ahead beliefs after a sample s^t as conditionals (given s^t) of probabilities of type (3). To capture reevaluation, we allow only conditionals of those probability measures that do a “good enough job” in explaining the sample s^t , which is formalized using an appropriate likelihood ratio test. The stringency of the test becomes a parameter α that governs how quickly agents are willing to resolve ambiguity.

3 A MODEL OF LEARNING

3.1 Recursive Multiple-Priors

We work with a finite period state space $S_t = S$, identical for all times. One element $s_t \in S$ is observed every period. At time t , an agent’s information consists of the history $s^t = (s_1, \dots, s_t)$. There is an infinite horizon, so S^∞ is the full state space.¹⁰ The agent ranks consumption plans $c = (c_t)$, where c_t is a function of the history s^t . At any date $t = 0, 1, \dots$, given the history s^t , the agent’s ordering is represented by a conditional utility function U_t , defined recursively by

$$U_t(c; s^t) = \min_{p \in \mathcal{P}_t(s^t)} E^p [u(c_t) + \beta U_{t+1}(c; s^t, s_{t+1})], \quad (4)$$

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

To clarify the connection to the Gilboa-Schmeidler model, it is helpful to rewrite utility using discounted sums. In a Bayesian model, the 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^∞

¹⁰In what follows, measures on S^∞ are understood to be defined on the product σ -algebra on S^∞ , denoted by \mathcal{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.

satisfying the regularity conditions specified in Epstein and Schneider [13].¹¹ Thus one obtains the following equivalent and explicit formula for utility:

$$U_t(c; s^t) = \min_{p \in \mathcal{P}} E^p [\sum_{s \geq t} \beta^{s-t} u(c_s) \mid s^t]. \quad (5)$$

This expression shows that each conditional ordering conforms to the multiple-priors model in Gilboa and Schmeidler [19], 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 [13]. 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 [13] 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 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 [13] restricts $\{\mathcal{P}_t\}$ only by technical regularity conditions. We now proceed to add further restrictions to capture how the agent responds to data.

3.2 Learning

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

As in the Bayesian model outlined in the introduction, our starting point is again (a finite period state space S and) a parameter space Θ that represents features of the data the decision-maker tries to learn. To accommodate ambiguity in initial beliefs about

¹¹In the infinite horizon case, uniqueness obtains only if \mathcal{P} is assumed also to be *regular* in a sense defined in Epstein and Schneider [14], generalizing to sets of priors the standard notion of regularity for a single prior.

parameters, represent those beliefs by a *set* \mathcal{M}_0 of probability measures on Θ . The size of \mathcal{M}_0 reflects the decision-maker's (lack of) confidence in the prior information on which initial beliefs are based. A technically convenient assumption is that \mathcal{M}_0 is weakly compact;¹² this permits us to refer to minima, rather than infima, over \mathcal{M}_0 .

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

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

Each $\ell : \Theta \rightarrow \Delta(S)$ is a likelihood function, so that $\theta \mapsto \ell(A | \theta)$ is assumed measurable for every $A \subset S$. Another convenient technical condition is that \mathcal{L} is compact when viewed as a subset of $(\Delta(S))^\Theta$ (under the product topology). Finally, to avoid the problem of conditioning on zero probability events, assume that each $\ell(\cdot | \theta)$ has full support.

Turn to interpretation. If there is a single likelihood, that is $\mathcal{L} = \{\ell\}$, then any given θ determines a unique probability measure on S^∞ . More generally, a nonsingleton set of likelihoods reflects the decision-maker's view that θ alone does not uniquely determine the data generating mechanism. She feels there are other factors underlying data and this will influence her beliefs and preference. 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. Factors modeled by θ are perceived as common across time or experiments - thus the agent can try to learn about them. Those modeled by the set \mathcal{L} are variable across time in a way that she does not understand beyond the limitation imposed by \mathcal{L} . In particular, at any point in time, *any* element of \mathcal{L} might be relevant for generating the next observation. Accordingly, while she can try to learn the true θ , she has decided that she will not try to (or is not able to) learn more.

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

$$\mathcal{P}_t(s^t) = \left\{ p_t(\cdot) = \int_{\Theta} \ell(\cdot | \theta) d\mu_t(\theta) : \mu_t \in \mathcal{M}_t(s^t), \ell \in \mathcal{L} \right\}, \quad (6)$$

or, in convenient notation,

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

¹²More precisely, measures in \mathcal{M}_0 are defined on the implicit and suppressed σ -algebra \mathcal{B} associated with Θ . Take the latter to be the power set when Θ is finite. The weak topology is that induced by bounded and \mathcal{B} -measurable real-valued functions on Θ . \mathcal{M}_0 is weakly compact iff it is weakly closed.

This process enters the specification of recursive multiple priors preferences (4).¹³

Updating and reevaluation

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

We allow for different attitude towards past and future ambiguous signals. On the one hand, \mathcal{L} is the set of likelihoods possible in the future. Since the decision-maker has decided she cannot learn the true sequence of likelihoods, it is natural that beliefs *about the future* must be based on the whole set \mathcal{L} as in (6). On the other hand, the decision-maker may *reevaluate*, with hindsight, her views about what sequence of likelihoods was relevant for generating data *in the past*. Such revision is possible because the agent learns more about θ and this might make certain theories more or less plausible. For example, some interpretation of the signals, reflected in a certain sequence $\ell^t = (\ell_1, \dots, \ell_t)$, or some prior experience, reflected in a certain prior $\mu_0 \in \mathcal{M}_0$, might appear not very relevant if it is part of a theory that does not explain the data well.

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

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

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

$$d\mu_t(\cdot; s^t, \mu_0, \ell^t) = \frac{\ell_t(s_t | \cdot)}{\int_{\Theta} \ell_t(s_t | \theta') d\mu_{t-1}(\theta'; s^{t-1}, \mu_0, \ell^{t-1})} d\mu_{t-1}(\cdot; s^t, \mu_0, \ell^{t-1}). \quad (7)$$

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

¹³Given compactness of \mathcal{M}_0 and \mathcal{L} , one can show that $\mathcal{M}_t(s^t)$ defined below is also compact, and subsequently that each $\mathcal{P}_t(s^t)$ is compact. This justifies the use of minima in (4).

that puts maximum likelihood on the sample. Posteriors are formed only for theories that pass the test. Thus posteriors are given by

$$\mathcal{M}_t^\alpha(s^t) = \{\mu_t(s^t; \mu_0, \ell^t) : \mu_0 \in \mathcal{M}, \ell^t \in \mathcal{L}^t, \int \prod_{j=1}^t \ell_j(s_j|\theta) d\mu_0(\theta) \geq \alpha \max_{\substack{\tilde{\mu}_0 \in \mathcal{M}_0 \\ \tilde{\ell}^t \in \mathcal{L}^t}} \int \prod_{j=1}^t \tilde{\ell}_j(s_j|\theta) d\tilde{\mu}_0\}. \quad (8)$$

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

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

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

where \mathcal{M}_t^α is given by (8). The model reduces to the Bayesian model when both the set of priors and the set of likelihoods have only a single element.

Another important special case occurs if \mathcal{M}_0 consists of *several* Dirac measures on the parameter space in which case there is a simple interpretation of the updating rule: \mathcal{M}_t^α contains all $\tilde{\theta}$'s such that the hypothesis $\theta = \tilde{\theta}$ is not rejected by an asymptotic likelihood ratio test performed with the given sample, where the critical value of the $\chi^2(1)$ distribution is $-2 \log \alpha$. Since $\alpha > 0$, (Dirac measures over) parameter values are discarded or added to the set, and \mathcal{P}_t varies over time. The Dirac priors specification is convenient for applications – it will be used in our portfolio choice example below. Indeed, one may wonder whether there is a need for non-Dirac priors at all. However, more general priors provide a useful way to incorporate objective probabilities, as illustrated by the scenarios in Section 2.¹⁴

4 PROPERTIES

In this section we illustrate properties of learning under ambiguity in the context of the scenarios described above. We show that our model rationalizes the intuitive choices

¹⁴Another example is in Epstein and Schneider [15], 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.

described in Section 2, and we examine the dynamics of beliefs and confidence in the short term. Then we provide a result on convergence of the learning process in the long run.

4.1 Scenarios Revisited

We described in Section 2 how we would model the three scenarios. The intuitive Ellsberg-type behavior would obviously be captured in this way. Focus on the other choice described above - given that one black ball has been drawn, would you rather bet on the next draw being black in Scenario 2 or in Scenario 3? This question highlights one difference between learning in simple settings (Scenario 2) as opposed to complex ones (Scenario 3), and correspondingly demonstrates the key role played by multiple-likelihoods. We argued above that it is intuitive that one prefer to bet in Scenario 2 - the black draw is stronger evidence there for the next draw also being black. We now propose a natural specification of beliefs for both Scenarios for which these choices obtain.

Scenario 2: To define a representation $(\Theta, \mathcal{M}_0, \mathcal{L}, \alpha)$, take $\Theta = \{\frac{n}{5} : n = 1, 2, 3, 4\}$, the proportion of black balls (coin or non-coin), and $\ell(B | \theta) = \theta$. Specify the set \mathcal{M}_0 of priors on Θ as follows: the true overall proportion is θ iff there are 5θ black balls in total. Let

$$P \subset \Delta(\{1, 2, 3\})$$

denote beliefs about the number of black non-coin balls. Since the coin-ball and other colors are independent, each prior μ_0 should have the form

$$\mu_0^p(\theta) = \frac{1}{2}p(5\theta - 1) + \frac{1}{2}p(5\theta) \quad (9)$$

for some $p \in P$. Thus

$$\mathcal{M}_0 = \{\mu_0^p : p \in P\}.$$

Finally, fix a reevaluation parameter α .

Scenario 3: Take $\Theta = \{B, W\}$, corresponding to the two possible colors of the coin ball, $\mu_0 = (\frac{1}{2}, \frac{1}{2})$, and define likelihoods as follows: let $P \subset \Delta(\{1, 2, 3\})$ denote beliefs about number of black non-coin balls, where P is the same set used in Scenario 2. Since the first urns in the two scenarios are identical, it is natural to use the same set P to describe beliefs about them. Moreover, though the urns differ along the sequence in Scenario 3, successive urns are indistinguishable, which explains why P can be used also to describe the second urn in the present scenario. Each p in P suggests one likelihood function via

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

and¹⁵

$$\mathcal{L} = \{\ell^p : p \in P\}.$$

¹⁵If $P_\epsilon = \{p \in \Delta(\{1, 2, 3\}) : \sum \lambda p(\lambda) \in [2 - \epsilon, 2 + \epsilon]\}$, for given ϵ , then the corresponding set of likelihoods is the set \mathcal{L}_ϵ described in Section 2.

The appendix shows that the intuitive choices discussed in Section 2 follow from these beliefs. In particular, it demonstrates that the minimal predictive probability for a black ball conditional on observing a black ball,

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

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

4.2 Inference from Small Samples

We use Scenario 3 here to illustrate the short run dynamics of beliefs (and confidence) that can be delivered by our model. Once again, take $\Theta = \{B, W\}$, $\mu_0 = (\frac{1}{2}, \frac{1}{2})$, and for the set of likelihoods take \mathcal{L}_ϵ from (2) with $\epsilon = 1$ - the agent weighs equally all the logically possible combinations of non-coin balls. The evolution of the posterior set \mathcal{M}_t^α in (8) shows how signals tend to induce ambiguity about parameters even where there is no ambiguity ex ante (singleton \mathcal{M}_0^α). This happens when the agent views some aspects of the data generating mechanism as too difficult to try to learn, as modeled by a nonsingleton set of likelihoods. More generally, \mathcal{M}_t^α can expand or shrink depending on the signal realization.

Posterior beliefs can be represented by the intervals

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

Figure 2 depicts the evolution of beliefs about coin-ball being black as more balls are drawn. The top panel shows the evolution of the posterior interval for a sequence of draws such that the number of black balls is $\frac{3t}{5}$, for $t = 5, 10, \dots$. Intervals are shown for $\alpha = .1$ and $\alpha = .001$, as well as for $\alpha = 0$, to illustrate what would happen without reevaluation. In addition, a single line is drawn for the case $\alpha = 1$, where the interval is degenerate. For example, after the first 5 signals, with 3 black balls drawn, the agent with $\alpha = .1$ 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 of past signals. For example, having seen 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




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.

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

4.3 Beliefs in the Long Run

Turn now to behavior after a large number of draws. As discussed above, our model describes agents who are not sure whether empirical frequencies will converge. Neverthe-

less, it is useful to derive limiting beliefs for the case when such convergence occurs: the limiting beliefs also approximately describe behavior after a large, but finite, number of draws. Therefore, they characterize behavior in the long run. The limiting result below holds with probability one under a true data-generating process, described by a probability P^* on (S^∞, \mathcal{S}) . For this process, we require only that the empirical frequencies of each of the finite number of states $s \in S$ converge, almost surely under P^* . In what follows, these limiting frequencies are described by a measure ϕ on S .

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

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

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

Theorem 1 *Suppose that Θ is finite and that:*

- (i) $\theta^* = \arg \max_{\theta} H(\theta)$ is a singleton;
- (ii) there exists κ such that for every μ_0 in \mathcal{M}_0^α ,

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

and $\mu_0(\theta^*) > 0$ for some μ_0 in \mathcal{M}_0^α .

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

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

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

Condition (i) 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. Condition (ii) is satisfied if every prior in \mathcal{M}_0^α assigns positive probability to θ^* (because \mathcal{M}_0^α is weakly compact). The weaker condition stated accommodates also the case where all priors are Dirac measures (including specifically the Dirac concentrated at θ^*), as well as the case of a single prior containing θ^* in its support. (In Scenario 3,

(ii) is satisfied for any set of priors where the probability of a black coin ball is bounded away from zero.)

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

The asserted uniform convergence is important in order that convergence of beliefs translate into long-run properties of preference. Thus it implies that for the process of one-step-ahead conditionals $\{\mathcal{P}_t\}$ given by (6),

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

we have

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

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

Long run behavior in Scenario 3

As a concrete example, consider long run beliefs about the urns in Scenario 3. Let ϕ_∞ denote the limiting frequency of black balls. Suppose also that beliefs are given by (2) with $\epsilon = 1$. 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 observes a large number of draws with a fraction of black balls above one 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 (8) in ensuring learnability of the representation. Suppose that $\phi_\infty = \frac{3}{5}$. The limiting frequency of $\frac{3}{5}$ black draws could be realized either because there is a black coin ball and on average one half of the non-coin balls were black, or because the coin ball is white and all of the urns contained 3 black non-coin balls. If α were equal to zero, both possibilities would be taken equally seriously and the limiting posterior set would contain Dirac measures that place probability one on either $\theta = B$ or $\theta = W$. This pattern is apparent from Figure 2. In contrast, with $\alpha > 0$, reevaluation eliminates the sequence where all urns contain three black non-coin balls as unlikely.¹⁶

Memoryless truths

The set of “true” processes P^* for which the theorem holds is large. Like a Bayesian who is sure the data are exchangeable, the agent described by the theorem is convinced that the data generating process is memoryless. As a result, learning is driven only by the empirical frequencies of the one-dimensional events, or subsets of S . Given any process for which those frequencies converge to a given limit, the agent will behave in the same way. For example, if the data were generated by a stationary Markov chain, the agent would not react to the resulting patterns of serial dependence, and act the same way as if the data were serially independent. Of course, in applications one would typically consider a “truth” that is memoryless. Like the exchangeable Bayesian model, our model is best applied to situations where the agent is *correct* in his a priori judgement that the data generating process is memoryless.

Importantly, a memoryless process for which empirical frequencies converge need not be i.i.d. There exists a large class of serially independent, but nonstationary, processes for which empirical frequencies converge. In fact, there is a large class of such processes that cannot be distinguished from an i.i.d. process with distribution ϕ on the basis of any finite sample. A concern with ongoing change may thus lead agents to never be confident that they have identified a “true” i.i.d. data generating process with distribution ϕ , even if they are convinced that the empirical frequencies converge.¹⁷

For a simple example, consider two measures ϕ_1 and ϕ_0 on S such that ϕ is a convex combination, that is, there is a number $\omega \in (0, 1)$ such that $\phi = \omega\phi_1 + (1 - \omega)\phi_0$. Now let $\tilde{\omega}_t$ be an i.i.d. process valued in $\{0, 1\}$ with $\Pr(\tilde{\omega}_t = 1) = \omega$. To construct a new nonstationary process on the state space S , first fix a realization $\omega = (\omega_t)$ of the process $(\tilde{\omega}_t)$. Then consider the serially independent process (s_t^ω) where the distribution of s_t^ω is $\phi_t = \omega_t\phi_1 + (1 - \omega_t)\phi_0$. For almost every realization ω , the empirical distribution of the nonstationary process (s_t^ω) will now converge to that of an i.i.d process with distribution ϕ .¹⁸ Without further information, agents who have observed a lot of i.i.d. data with

¹⁶With a stronger identification condition, then even if $\alpha = 0$, the posterior set converges to the singleton set consisting of the Dirac measure concentrated on θ^* . This case is discussed further in Epstein and Schneider [15].

¹⁷As discussed earlier, they may also reach this view because they are not sure whether the empirical frequencies converge in the first place.

¹⁸See Nielsen [32] for a formal statement and proof.

distribution ϕ cannot rule out the possibility that the true data generating process is nonstationary – it could be any of the nonstationary processes indexed by a typical realizations ω . Moreover, since ω_{t+1} is unknown for any finite t , agents may not be sure that the distribution of the next realization is ϕ – for all they know, it might be ϕ_0 or ϕ_1 . Non-vanishing ambiguity is one way to capture this concern with ongoing change.

5 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 [11] 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 [16] and Mukherji and Tallon [31] 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.¹⁹ Investors will participate in a market only if they are confident that they know how to make money in that particular market.²⁰

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

¹⁹For 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.

²⁰This 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 [9, Ch. 43] for a discussion of this issue.

solving the recursive problem

$$\begin{aligned}
 V_t(W_t, s^t) &= \max_{\omega} \min_{p_t \in \mathcal{P}_t(s^t)} E^{p_t} [V_{t+1}(W_{t+1}, s^{t+1})] \\
 &\text{subject to} \\
 W_t &= (R^f + (R(s^t) - R^f)\omega(s^{t-1})) W_{t-1}, \quad t = 1, \dots, T. \quad (11)
 \end{aligned}$$

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 5.1, 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 [30], we focus on the case of log utility, where traditional hedging demand is zero. In Section 5.2, we study a calibrated example of asset allocation between stocks and riskless bonds for U.S. investors in the post-war period.

5.1 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 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 she 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 structure of ambiguity, 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

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

Consider next a *myopic* investor with nondegenerate beliefs process $\{\mathcal{P}_t\}$. Given the

history s^t , such an agent solves

$$\begin{aligned} & \max_{\omega} \min_{p_t(s^t) \in \mathcal{P}(s^t)} E^{p_t(s^t)} \left[\log \left(R^f + \left(\tilde{R}_{t+1} - R^f \right) \omega \right) \right] \\ &= \min_{p_t(s^t) \in \mathcal{P}(s^t)} E^{p_t(s^t)} \left[\log \left(R^f + \left(\tilde{R}_{t+1} - R^f \right) \omega^*(p_t(s^t)) \right) \right], \end{aligned} \quad (12)$$

where we have used the minmax theorem to exchange the order of optimization. Denote by $p_t^{myopic}(s^t)$ the minimizing (conditional-one-step-ahead) measure in (12). Then the optimal policy of a myopic agent is $\omega^*(p_t^{myopic}(s^t))$, the portfolio weights that are optimal for the corresponding Bayesian.

For the intertemporal problem (11), conjecture the value functions $V_t(W_t, s^t) = \log W_t + h_t(s^t)$, with $h_T = 0$. Again using the minmax theorem, as well as the budget constraint, we obtain

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

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

$$h_t(s^t) = \min_{p_t \in \mathcal{P}_t(s^t)} E^{p_t} \left[\log \left(R^f + \left(\tilde{R}_{t+1} - R^f \right) \omega(p_t(s^t)) \right) + h_{t+1}(s^{t+1}) \right]. \quad (13)$$

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

Comparison of (12) and (13) 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:

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

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

Independent vs. correlated ambiguity

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 - thus $R_t = R(x^{t-1}, y_t)$. Then define conditionals

$$\begin{aligned} \mathcal{P}_t(s^t) &= \{p_t^x \otimes p_t^y : p_t^x \in \mathcal{P}_t^x(x^t), p_t^y \in \mathcal{P}_t^y\} \\ &= \mathcal{P}_t^x(x^t) \otimes \mathcal{P}_t^y, \end{aligned} \tag{14}$$

where $\mathcal{P}_t^x(x^t)$ and \mathcal{P}_t^y are sets of probability measures on X and Y , respectively.²¹ We thus allow both the signal and the shock to be ambiguous.

With these beliefs, (13) becomes

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

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 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, she may be exposed to unknown factors that affect the distribution of h_{t+1} , even if she does not invest in the uncertain asset over the next period.²² 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 (13) 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 hedging is important in the context of learning. We present two examples where beliefs are given by representations $(\Theta, \mathcal{M}_0, \mathcal{L}, \alpha)$. In one case, ambiguity about returns and ambiguity about investment opportunities are independent, while in the other there is correlated ambiguity. The latter case is also relevant for the calibrated example of the next subsection.

²¹The 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 .

²²Exposure requires that (i) the realization s_{t+1} typically provides news about future opportunities, that is, it affects future belief sets $\mathcal{P}_{t+j}(s^{t+j})$ and (ii) that the news is payoff-relevant, so that the value function $h_{t+1}(s^{t+1})$ depends on s_{t+1} .

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

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

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 (7) and $\mathcal{M}_t^\alpha(s^t)$ depends only on x^t . Moreover, since ℓ^y does not depend on the parameter θ , the one-step-ahead conditionals $\mathcal{P}_t(s^t)$ will be a special case of (14). 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 correlated ambiguity, modify the above setup by setting $R_t = R(x_t)$, 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 $\mathcal{P}_t(s^t) = \{p_t(s^t; \bar{\theta}_0); \bar{\theta}_0 \in \bar{\Theta}_0\}$. 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 $E^{p_t}[R_{t+1}] = 0$ for some conditional p_t that achieves the minimum in (13). If distributions are chosen such that there is just one value $\bar{\theta}_0$ for which $E^{p_t(s^t; \bar{\theta}_0)}[R_{t+1}] = 0$, then this value will typically not achieve the minimum for the sum in (13). 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 structure of ambiguity. 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 [30] decomposition, and progress has been mostly through calibrated examples. We proceed to develop such an example in the next section.

5.2 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.²³ 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 $\theta \in \Theta \equiv [\bar{\lambda}, 1 - \bar{\lambda}]$, where $\bar{\lambda} < \frac{1}{2}$. However, she 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 $\ell(\cdot | \theta)$ such that

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

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 $\bar{\lambda} > 0$, returns are ambiguous signals: $\lambda_t \in [-\bar{\lambda}, \bar{\lambda}]$ 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^α 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

$$\mathcal{M}_t^\alpha(s^t) = \left\{ \theta \in \Theta : g(\theta; \phi_t) \geq \max_{\tilde{\theta} \in \Theta} g(\tilde{\theta}; \phi_t) + \frac{\log \alpha}{t} \right\}, \quad (15)$$

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

Using (15), 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^α collapses to the Dirac measure at θ^* , where

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

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

$$\Pr(s = 1) \in [\theta^* - \bar{\lambda}, \theta^* + \bar{\lambda}]. \quad (16)$$

²³Here we follow much of the finance literature and consider nominal returns.

The agent thus learns the true parameter value θ^* , in the sense that in the limit she behaves as if she 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 (11) can be rewritten using the fraction ϕ_t of high returns as a state variable:

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

subject to the transition equations

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

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 $\bar{\lambda}$. The parameter $\bar{\lambda}$ determines how much the agent thinks that she will learn in the long run. To determine a value, she 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. She 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 (16), the result is $\frac{1}{2} - \bar{\lambda}$. We present results for values of $\bar{\lambda}$ 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 will be reduced below 1% (to 88 basis points, to be exact) after a lot of data has been observed.

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 ($\bar{\lambda} = 0$), there is a simple interpretation of our updating rule: \mathcal{M}_t^α 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.




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.

5.3 Numerical Results

Our leading example is an investor in 1971:Q3 ($t = 0$), who plans over various horizons (up to 30 years). She looks back on data starting in 1927:Q3 and thus begins with a posterior set that is already based on 43 years of data. 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 $\bar{\lambda} = 0$. 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, she eventually goes 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, she takes ever larger long positions.

The optimal policy surface has also a flat piece at zero: when enough low returns are observed, the agent does 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 5.1, 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 she 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 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*. She is long in stocks if the sample equity premium is positive and short otherwise. She also reacts to high (low) returns by increasing (decreasing) her net position. If the absolute value of the equity premium is in an intermediate range, the agent is in a *contrarian region*. She is *short* in stocks for a positive sample equity premium and long otherwise. Moreover, she now reacts to high (low) returns by decreasing (increasing) her 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 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, she can exploit the expected equity premium by going long. On the other hand, she can insure herself 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.




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 $\bar{\lambda} = 0$.

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 hold stocks, although her 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.

The plot also shows that small amounts of signal ambiguity can significantly reduce the optimal stock position. The investor with $\bar{\lambda} = .01$ already holds essentially no stocks throughout most of the 1970s. An investor with $\bar{\lambda} = .02$ 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 she is long in stocks, changes in the ambiguity parameters α and $\bar{\lambda}$ 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




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.

look very different in states where the ambiguity averse investor moves in and out of the market.

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

6 RELATED LITERATURE

We are aware of only two formal treatments of learning under ambiguity. Marinacci [29] studies repeated sampling with replacement from an Ellsberg urn and shows that ambiguity is resolved asymptotically. His setup corresponds to a version of our model in which there is a single likelihood and no reevaluation. 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 multiple likelihoods. An important difference, however, is that our model is consistent with a coherent axiomatic theory of dynamic choice between consumption processes. Accordingly, it is readily applicable to economic settings.²⁴

In our model, learning may cease without all ambiguity having been resolved, which is quite different from existing Bayesian approaches to modeling incomplete learning in complicated environments. One such approach starts from the assumption that the true data generating measure is not absolutely continuous with respect to an agent's belief.²⁵ This generates situations where beliefs do not converge to the truth even though agents believe, and behave as if, they will.²⁶ 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 are sure that the state variable is forever changing, and the distribution of these regime changes is known. Agents thus track a *known* data generating process that is not memoryless. In contrast, our model applies to memoryless mechanisms and learning is about a fixed true parameter. Nevertheless, because of ambiguity, the agent reaches a state where no further learning is possible although the data generating mechanism is not yet understood completely.

The role of ambiguity for explaining asset market participation was first studied by Dow and Werlang [11] in a static context. We build on their analysis, but consider a dynamic problem with hedging demand. 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. Non-participation can also 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 only the latter allows different attitudes towards different sources of uncertainty. In applications, this is relevant for modeling participation in some asset markets and non-participation in others.

²⁴A similar remark applies to Huber [23], who also points to the desirability of admitting multiple likelihoods and outlines one proposal for doing so.

²⁵This violates the Blackwell-Dubins [6] conditions for convergence of beliefs to the truth. See Feldman [17] for an economic application.

²⁶As 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.

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 to 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.

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 [12], Routledge and Zin [33]). Similarly, existing robust control models (Hansen, Sargent, and Tallarini [21], Maenhout [27], Cagetti et al. [7]) 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 also contributes to a growing literature on learning and portfolio choice. Bawa, Brown, and Klein [5] and Kandel and Stambaugh [25] first explored the role of parameter uncertainty in a Bayesian framework.²⁷ Several authors have solved intertemporal portfolio choice problems with Bayesian learning.²⁸ The main results are conservative 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.

A APPENDIX

A.1 Comparison of predictive probabilities from Section 4.1.

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

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

is smaller under Scenario 3 than under Scenario 2.

²⁷There are alternatives to a Bayesian approach to the parameter uncertainty problem. See Ang and Bekaert [1] for a classical econometric strategy.

²⁸Detemple [10] and Gennotte [18] have considered the case of learning about mean 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.

Consider first Scenario 2. The set \mathcal{M}_1^α of posteriors (after having observed one black draw) is constructed by updating only those priors μ_0^p according to which the black draw was sufficiently likely ex ante in the sense of the likelihood-ratio criterion in (8). The ex ante probability of a black draw under the prior μ_0^p is (using (9))

$$\int_{\Theta} \theta d\mu_0^p = \frac{1}{10} [2 \sum_{\lambda=1}^3 \lambda p(\lambda) + 1]. \quad (17)$$

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

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

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

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

for some p in P^α .

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

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

Compute that

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

that is,

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

Therefore, the minimum predictive probability satisfies

$$\min \left\{ \int_{\Theta} \theta d\mu_1^p : p \in P^\alpha \right\} \geq \frac{1}{5} \left[\min_{p \in P^\alpha} \frac{\sum \lambda p(\lambda) + 1}{2 \sum \lambda p(\lambda) + 1} + \min_{p \in P^\alpha} \sum \lambda p(\lambda) \right]. \quad (18)$$

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

Turn now to Scenario 3. Note that

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

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

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

which is identical to the ex ante probability computed in Scenario 2 (see (17)). It follows that the likelihood-ratio criterion leads to retention of only those likelihoods ℓ^p where p lies in P^α , the same set defined in the discussion of Scenario 2.

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

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

In particular,

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

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

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

Therefore the minimum predictive probability equals

$$\begin{aligned} &\min_{p \in P^\alpha, \ell \in \mathcal{L}} \{ \mu_1^p(B) [\ell(B | B) - \ell(B | W)] + \ell(B | W) \} = \\ &\quad \min_{p \in P^\alpha, \ell \in \mathcal{L}} \{ \mu_1^p(B) \frac{1}{5} + \ell(B | W) \} = \\ &\quad \frac{1}{5} \min_{p \in P^\alpha} \mu_1^p(B) + \min_{\ell \in \mathcal{L}} \ell(B | W) = \frac{1}{5} \min_{p \in P^\alpha} \mu_1^p(B) + \min_{p' \in P} \frac{\Sigma_1^3 \lambda p'(\lambda)}{5} \\ &\leq \frac{1}{5} \min_{p \in P^\alpha} \mu_1^p(B) + \min_{p' \in P^\alpha} \frac{\Sigma_1^3 \lambda p'(\lambda)}{5} \quad (\text{replacing } P \text{ by its subset } P^\alpha) \\ &= \frac{1}{5} \min_{p \in P^\alpha} \frac{\Sigma \lambda p(\lambda) + 1}{2 \Sigma \lambda p(\lambda) + 1} + \min_{p' \in P^\alpha} \frac{\Sigma \lambda p'(\lambda)}{5}, \end{aligned}$$

which is no larger than the minimum predictive probability computed for Scenario 2 (see (18)). This proves that betting in Scenario 2 is preferable.

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

A.2 Proof of Theorem 1

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

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

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

We want to show that, for every sequence of posteriors $\{\mu_t\}$ with $\mu_t \in \mathcal{M}_t^\alpha$ for all t , the posterior probability of the parameter value θ^* converges (uniformly) to one. By the definition of \mathcal{M}_t^α , for every posterior μ_t there exists an *admissible theory* $(\mu_0^{(t)}, \ell^{t,(t)})$, that is, a theory satisfying the conditions in (8), such that μ_t is the Bayesian update of $\mu_0^{(t)}$ along the sequence of likelihoods $\ell^{t,(t)} = (\ell_1^{(t)}, \ell_2^{(t)}, \dots, \ell_t^{(t)})$. Here the t 's in brackets indicate the place of the posterior in the given sequence of posteriors $\{\mu_t\}$ – they are needed to account for the fact that not all μ_t in the sequence $\{\mu_t\}$ need to be updates of the same initial prior μ_0 , or be updated along the same likelihood sequence. (As well, the sequence of updates of a given μ_0 along a given sequence of likelihoods need not be admissible at all dates.) Thus our objective is to show that

$$\lim_{t \rightarrow \infty} \mu_t(\theta^* | s^t, \mu_0^{(t)}, \ell^{t,(t)}) = 1, \quad (19)$$

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

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

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

Here η_t is the log likelihood ratio between the likelihood of the sample under ℓ^t at the parameter value θ and the likelihood of the sample under the sequence (ℓ_s) that maximizes the likelihood of the sample given the parameter θ^* . A posterior μ_t derived from a theory

$(\mu_0^{(t)}, \ell^{t,(t)})$ can be written as

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

Claim 1: For every $\epsilon > 0$ there exists T such that

$$\eta_t(\theta, \ell^{t,(t)}) \leq -\epsilon$$

for all $t > T(\epsilon)$, $\theta \neq \theta^*$ and for all sequences of likelihood tuples $(\ell^{t,(t)})$.

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

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

Claims 1 and 2 together imply that

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

uniformly. Claim 3 and hypothesis (ii) imply that $\mu_0^{(t)}(\theta^*) > \kappa$ for large enough t , and (19) follows.

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

$$\frac{1}{t} \sum_{j=1}^t \log \ell_j^{(t)}(s_j | \theta) \leq \max_{(\ell_s)} \tilde{H}(\phi_t, (\ell_s), \theta) \quad (20)$$

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

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

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

$$\max_{(\ell_s)} \widetilde{H}(\phi_t, (\ell_s), \theta) \leq \max_{(\ell_s)} \widetilde{H}(\phi_t, (\ell_s), \theta^*) - \epsilon, \quad (21)$$

for all $\theta \neq \theta^*$ and $t > T$. The claim now follows from (20).

Proof of Claim 2: By definition,

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

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

$$\left(\max_{\widehat{\mu}_0, \widehat{\ell}^t} \Pr(s^t; \widehat{\mu}_0, \widehat{\ell}^t) \right)^{\frac{1}{t}} \geq \left(\Pr(s^t; \mu_0^{(t)}, \ell^{t,(t)}) \right)^{\frac{1}{t}} \geq \alpha^{\frac{1}{t}} \left(\max_{\widehat{\mu}_0, \widehat{\ell}^t} \Pr(s^t; \widehat{\mu}_0, \widehat{\ell}^t) \right)^{\frac{1}{t}}, \quad (22)$$

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

$$\left[\frac{\max_{\widehat{\mu}_0, \widehat{\ell}^t} \Pr(s^t; \widehat{\mu}_0, \widehat{\ell}^t)}{e^{t \max_{(\ell_s)} \widetilde{H}(\phi_t, (\ell_s), \theta^*)}} \right]^{\frac{1}{t}} \longrightarrow 1, \quad (23)$$

or equivalently,

$$\left[\max_{\widehat{\mu}_0, \widehat{\ell}^t} \sum_{\theta \in \Theta} \widehat{\mu}_0(\theta) e^{t\eta_t(\theta, \widehat{\ell}^t)} \right]^{\frac{1}{t}} \longrightarrow 1. \quad (24)$$

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

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

We have assumed that there is a prior μ_0 such that $\mu_0(\theta^*) > 0$. Verify that

$$\eta_t(\theta^*, \widehat{\ell}^t) \leq 0 \text{ and } \max_{\widehat{\ell}^t} \eta_t(\theta^*, \widehat{\ell}^t) = 0. \quad (25)$$

Then, by hypothesis (ii),

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

which implies that

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

and hence also (24).

Combining (23) and (22), it now follows that, for any sequence of admissible theories $(\mu_0^{(t)}(t), \ell^{t,(t)}(t))$,

$$\left(\sum_{\theta \in \Theta} \mu_0^{(t)}(\theta) e^{t\eta_t(\theta, \ell^{t,(t)})} \right)^{\frac{1}{t}} = \left[\frac{\Pr(s^t; \mu_0^{(t)}, \ell^{t,(t)})}{e^{t \max(\ell_s) \widetilde{H}(\phi_t, (\ell_s), \theta^*)}} \right]^{\frac{1}{t}} \longrightarrow 1, \quad (26)$$

and that the convergence is uniform.

Finally, $\eta_t(\theta^*, \ell^{t,(t)}) \leq 0$ by construction - see (25). Suppose that $\eta_t(\theta^*, \ell^{t,(t)}) < -\delta$ for some positive δ and infinitely many t . Then Claim 1 (with $\epsilon = \delta$) implies that

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

for infinitely many t , contradicting (26). This proves Claim 2.

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

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

which contradicts (26). ■

A.3 Proof of (15)

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

$$L(s^t, \theta, \lambda^t) = \prod_{j=1}^t (\theta + \lambda_j)^{s_j} (1 - \theta - \lambda_j)^{1-s_j}. \quad (27)$$

Let $\tilde{\lambda}^t$ denote the sequence that maximizes (27) for fixed θ . This sequence is independent of θ and has $\tilde{\lambda}_j = \bar{\lambda}$ if $s_j = 1$ and $\tilde{\lambda}_j = -\bar{\lambda}$ if $s_j = 0$, for all $j \leq t$. It follows that $L(s^t, \theta, \tilde{\lambda}^t)$ depends on the sample only through the fraction ϕ_t of high returns

observed. The set \mathcal{M}_t^α 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 (θ, λ^t) 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 (θ, λ^t) 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 (15).■

References

- [1] A. Ang and G. Bekaert, International asset allocation with regime shifts, *Rev. Financ. Stud.* 15 (2002), 1137-1187.
- [2] A. Ang, G. Bekaert, and J. Liu, Why stocks may disappoint, *J. Finan. Econ.*, forthcoming.
- [3] N. Barberis, Investing for the long run when returns are predictable, *J. Finance* 55 (2000), 225-264.
- [4] T.F. Bewley, Knightian decision theory and econometric inference, Cowles #868, 1988.
- [5] V.S. Bawa, S.J. Brown and R.W. Klein, *Estimation Risk and Optimal Portfolio Choice*, Amsterdam: North Holland, 1979.
- [6] D. Blackwell and L. Dubins, Merging of opinions with increasing information, *Annals of Math. Statistics* 33 (1962), 882-886.
- [7] M. Cagetti, L.P. Hansen, T. J. Sargent and N. Williams, Robustness and pricing with uncertain growth, *Rev. Financ. Stud.* 15 (2002), 363-404.
- [8] Z. Chen and L.G. Epstein, Ambiguity, risk and asset returns in continuous time, *Econometrica* 70 (2002), 1403-43.
- [9] D. Chew, *The New Corporate Finance*, 2nd edition, McGraw Hill, 1999.
- [10] J. Detemple, Asset pricing in a production economy with incomplete information, *J. Finance* 41 (1986), 383-390.
- [11] J. Dow and S. Werlang, Uncertainty aversion, risk aversion and the optimal choice of portfolio, *Econometrica* 60 (1992), 197-204.

- [12] L.G. Epstein and J.J. Miao, A two-person dynamic equilibrium under ambiguity, *J. Econ. Dyn. and Control* 27 (2003), 1253-1288.
- [13] L.G. Epstein and M. Schneider, Recursive multiple-priors, *J. Econ. Theory* 113 (2003), 1-31.
- [14] L.G. Epstein and M. Schneider, IID: independently and indistinguishably distributed, *J. Econ. Theory* 113 (2003), 32-50.
- [15] L.G. Epstein and M. Schneider, Ambiguity, information quality and asset pricing, 2004.
- [16] L.G. Epstein and T. Wang, Intertemporal asset pricing under Knightian uncertainty, *Econometrica* 62 (1994), 283-322.
- [17] M. Feldman, On the generic nonconvergence of Bayesian actions and beliefs, *Econ. Theory* 1 (1991), 301-321.
- [18] G. Genotte, Optimal portfolio choice under incomplete information, *J. Finance* 41 (1986), 733-746
- [19] I. Gilboa and D. Schmeidler, Maxmin expected utility with nonunique prior, *J. Math. Econ.* 18 (1989), 141-153.
- [20] L. Guiso, M. Haliassos, and T. Jappelli, *Household Portfolios*, MIT Press, 2002.
- [21] L.P. Hansen, T.J. Sargent and T. Tallarini, Robust permanent income and pricing, *Rev. Ec. Stud.* 66 (1999), 873-907.
- [22] L.P. Hansen, T.J. Sargent and N. Wang, Robust permanent income and pricing with filtering, *Macroeconomic Dynamics* 6 (2002), 40-84.
- [23] P.J. Huber, The use of Choquet capacities in statistics, *Bull. Internat. Statist. Inst.* 45 (1973), 181-188.
- [24] M. Jackson, E. Kalai and R. Smorodinsky, Bayesian representation of stochastic processes under learning: de Finetti revisited, *Econometrica* 67 (1999), 875-893.
- [25] S. Kandel and R. Stambaugh, On the predictability of stock returns: an asset allocation perspective, *J. Finance*, 51 (1996), 385-424.
- [26] D.M. Kreps, *Notes on the Theory of Choice*, Westview, 1988.
- [27] P. Maenhout, Robust portfolio rules and asset pricing, *Rev. Financ. Stud.* 17 (2004), 951-983.
- [28] 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.

- [29] M. Marinacci, Learning from ambiguous urns, *Statistical Papers* 43 (2002), 145-151.
- [30] R. Merton, Optimal consumption and portfolio rules in a continuous time model, *J. Econ. Theory* 3 (1971), 373-413.
- [31] S. Mukerji and J. Tallon, Ambiguity aversion and incompleteness of financial markets, *Rev. Econ. Stud.* 68 (2001), 883-904.
- [32] C.K. Nielsen, Rational belief structures and rational belief equilibria, *Economic Theory* 8 (1996) 3, 399-422.
- [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, *J. Finance* 58 (2003), 2465-2486.
- [36] A. Vissing-Jorgensen, Perspectives on behavioral finance: does “irrationality” disappear with wealth? 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.