LEARNING UNDER AMBIGUITY*

Larry G. Epstein Martin Schneider

September 10, 2002

Abstract

This paper considers learning when the distinction between risk and ambiguity (Knightian uncertainty) matters. Working within the framework of recursive multiple-priors utility, the paper formulates a counterpart of the Bayesian model of learning about an uncertain parameter from conditionally i.i.d. signals. Ambiguous signals capture differences in information quality that cannot be captured by noisy signals. They may increase the volatility of conditional actions and they prevent ambiguity from vanishing in the limit.

Properties of the model are illustrated with two applications. First, in a dynamic portfolio choice model, stock market participation costs arise endogenously from preferences and depend on past market performance. Second, ambiguous news induce negative skewness of asset returns and may increase price volatility. Shocks that trigger a period of ambiguous news induce a price discount on impact and are likely to be followed by further negative price changes.

*Department of Economics, U. Rochester, Rochester, NY, 14627, lepn@troi.cc.rochester.edu, and Department of Economics, UCLA, California, 90095-1477, martin@econ.ucla.edu. Epstein acknowledges financial support from NSF (grant SES-9972442). We are grateful to Hal Cole, Mark Grinblatt, Massimo Marinacci, Monika Piazzesi, Bryan Routledge and seminar participants at Berkeley, Carnegie-Mellon, Davis, the Federal Reserve Bank of Atlanta, Iowa, the NBER Summer Institute and U. Rochester for helpful comments and suggestions.

1 INTRODUCTION

1.1 Motivation and Outline of the Model

Models of learning describe economic agents whose actions are based on out-of-sample forecasts subject to model uncertainty. In contrast to agents with rational expectations, who know the true data generating process at all times, learning agents behave more like economists themselves. This is appealing when thinking about complicated environments such as financial markets, where the true model is hard to come by. Thus asset prices should be driven by the dynamics of agents' out-of-sample forecasts and portfolio choice recommendations should take model uncertainty explicitly into account.

But how do (or should) agents process information to arrive at forecasts? Most existing models of rational learning have taken a Bayesian approach. This precludes a distinction between *risky* situations, where the odds are known, and *ambiguous* situations, where the odds are not known with precision. 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. Ambiguous signals are processed differently from noisy (or risky) ones, and ambiguity perceived about the environment changes over time and may never be fully resolved. We propose a model of learning under ambiguity and provide applications to portfolio choice and asset pricing to illustrate its properties in economic settings.

Our starting point is *(recursive) multiple-priors* utility, a model of utility axiomatized in Epstein and Schneider [18], that extends Gilboa and Schmeidler's [23] model of decision making under ambiguity to an intertemporal setting. Learning is completely determined by specification of an initial *set* of probability measures $\mathcal{P}$ consistent with regularity conditions. 'Measure-by-measure' Bayesian updating describes subsequent responses to observations. In the present paper, we provide further structure on $\mathcal{P}$ to capture an agent's *a priori* view that data are generated by the same memoryless mechanism every period. This is the same *a priori* view that motivates the Bayesian model of learning about an underlying parameter from conditionally i.i.d. signals.[1] In both models, there is a parameter space representing a feature of the environment that the agent tries to learn and a sequence of signals that provide information.

Our model captures three important aspects of learning about a memoryless mechanism that are missed by the Bayesian approach. First, it broadens the notion of information quality. For a *noisy* signal, the distribution conditional on the parameter is known. In this sense, the meaning of the signal is clear. However, in many settings it is plausible that the agent is wary of a host of poorly understood or unknown factors that underlie realized signals and that obscure their meaning. This type of low information quality

[1] One stylized environment where this *a priori* view is reasonable is data generated by sampling with replacement from an urn that contains balls of various colors in unknown proportions. Here the mechanism is 'the same every period' because the urn is always the same. It is memoryless because draws are independent.

can be captured via *ambiguous signals*, where the conditional distribution given the parameter is unknown (or nonunique). Second, ambiguity may change nonmonotonically over time. The set of conditional probabilities used for forecasting can shrink or expand with new data as signals either resolve or induce ambiguity. Third, agents facing an ambiguous memoryless mechanism need not expect to be able to identify a simple i.i.d. data generating process as the 'truth', not even in the long run. While their learning process does settle down, the presence of poorly understood or unknown factors may prevent ambiguity perceived about the environment from vanishing in the limit.

To illustrate, suppose the agent participates in a trial for a new drug. In the beginning, he is told that, with equal probability, he will receive either the medication or a placebo. Ex ante, a bet on receiving medication is risky - it is equivalent to a bet on a fair coin. Consider next the agent's view of the environment after the first day of the trial. He will consider a change in his health as an informative signal. For example, if his health improves, it is more likely that he is receiving medication. However, given the lack of information about how the drug works, the agent will typically not know the likelihood. Therefore, a bet on receiving treatment is no longer purely risky ex post: the signal induces ambiguity. To describe adequately the quality of this signal, the standard risk measure (precision) is inappropriate. Our model represents ambiguous signals by a *set* of likelihoods; the size of this set captures information quality.

Ambiguous signals affect what agents expect to learn in the long run. When signals are noisy, the agent views the memoryless mechanism as a sequence of experiments held under *identical* conditions. Thus he plausibly expects to learn a true i.i.d. process that generated the data. However, the *a priori* view that there is a memoryless mechanism applies also to situations where the hypothesis of 'identical conditions' is not intuitive. Consider again the drug trial example. Suppose the agent knows that any persistent changes to his condition must arise from the medication.[2] At the same time, he is aware of many other factors that affect his health. These factors are not only hard to describe, but also changing over time and in a way that he does not understand. A reasonable agent might (i) every day, view tomorrow's change in health as ambiguous, and (ii) given the lack of day-specific information about the other factors, perceive the degree of ambiguity to be the same every day. The mechanism is thus a series of *indistinguishable* experiments, where equally little is known about each of them.

For an agent faced with such a mechanism, it is plausible that the learning process 'settles down': beliefs should eventually change little with every new observation. However, ambiguity about the data may persist forever. For one thing, the next signal will always be viewed as ambiguous. In addition, ambiguity about the parameter need not vanish. In the example, ambiguity induced by the other factors might make it altogether impossible for the agent to be confident about whether or not he is receiving medication. Our model provides a flexible structure that allows the learning dynamics to depend on the agent.

[2] The thrust of the argument does not depend on this assumption, but the model would have to be more complicated if it were to be relaxed.

Ambiguous signals also affect the choice of actions. According to the multiple-priors model, an ambiguity averse agent evaluates an action using the (updated) probability measure from the set $\mathcal{P}$ that minimizes the payoff from that action.[3] This entails an asymmetric response to news: ambiguous signals that convey 'good news' for a given action are viewed as less precise and hence less reliable than signals conveying bad news. Ambiguity also affects the link between information quality and the volatility of actions. The typical result in the Bayesian case is that, with lower information quality, actions are less volatile since signals are given less weight. In contrast, if a decrease in information quality is captured by an increase in signal ambiguity, then actions may become more volatile. That is because a more ambiguous signal admits more possible interpretations and the volatility of the 'worst case' interpretation, and hence also the associated action, may increase.

The paper is organized as follows. The remainder of this introduction outlines our applications. Section 2 presents a sequence of thought experiments to argue that the Bayesian model is not a satisfactory model of learning. Section 3 briefly reviews recursive multiple-priors utility. Section 4 introduces the learning model. The next two sections contain the portfolio choice application and our model of asset pricing with ambiguous news. Section 7 discusses related literature. Some proofs are collected in an appendix.

1.2 Applications

As our first application, we study the portfolio choice and stock market participation decisions of an agent who is learning about mean stock returns. Among U.S. households, non-participation in the stock market is both widespread and variable over time. At this point, a satisfactory model of the cross section of holdings does not exist.[4] One problem is that, in the standard frictionless portfolio choice model based on expected utility, zero holdings of risky assets are almost never optimal. In contrast, with multiple-priors utility, zero holdings of ambiguous assets are often optimal. Intuitively, if the interval of expected excess returns contemplated by the agent contains zero, then the worst case payoff for both long and short positions is at most zero. This result was shown in a static setting by Dow and Werlang [16].

What is new in our analysis is that we consider a dynamic model with learning where agents hedge changes in the investment opportunity set that arise from changes in beliefs. We examine when non-participation is optimal in this setting. Learning endogenizes the interval of expected excess returns and links the participation decision explicitly to past market performance. Ambiguity averse investors tend to be momentum investors, buying (or entering the market) on good news and selling (or exiting) on bad news. We propose a way for an investor to determine his degree of ambiguity about returns and we calculate optimal portfolios in real time. For reasonable parameter values, the model recommends

[3] The min operator in the functional form is justified by axioms that formally capture ambiguity aversion. Section 2 illustrates how it helps to describe intuitive behavior, such as arises in the Ellsberg Paradox.

[4] See Guiso and Haliassos [24] for a survey of the existing empirical and theoretical literature.

that U.S. investors hold no stocks in the 1970s, entering the market only in the late 1980s.[5]

As a second application, we formulate a simple model of asset pricing in times of "ambiguous news". It is motivated by consideration of a shock that simultaneously (*i*) increases uncertainty about fundamentals and (*ii*) changes the nature of signals relevant for forecasting fundamentals. One example is the terrorist attack of September 11, 2001. This shock both increased uncertainty about future growth **and** shifted the focus to hitherto "unfamiliar" news about foreign policy and terrorism. The shock could also be idiosyncratic: for example, a merger announcement often increases uncertainty about future performance **and** changes the set of signals used for valuation. Since the shock increases uncertainty, it marks the start of a learning process that affects prices. Since the news are unfamiliar, it is natural to model this process as learning from signals with ambiguous precision. This approach gives rise to two important effects.

First, **before** seeing a signal, agents want to be compensated for the ambiguity that it will induce. The **prospect** of ambiguous news thus leads to a price discount for securities even if there is no change in expected payoffs or in risk premia. This requires reevaluation of conclusions commonly drawn from event studies. For example, a negative abnormal return after a merger announcement need not imply that the market views the merger as a bad idea. Instead it might simply reflect the market's discomfort in the face of the upcoming period of ambiguous news. Importantly, a discount due to low future **idiosyncratic** information quality cannot be captured by the Bayesian model: if lower information quality is captured by higher risk, then it should be diversified away.[6]

Second, **after** seeing a signal, agents react asymmetrically: bad news are taken more seriously than good news. The distribution of returns thus becomes more negatively skewed than the distribution of signals. As a result, the initial shock can have a drawn out negative effect on prices even if there is no long term negative change in fundamentals. This can help to explain price movements in the aftermath of 9/11. The initial drop in the stock market when it reopened on September 17 was followed by more losses over the following week, before a gradual rebound occurred. We calibrate a representative agent model with learning to compare Bayesian and multiple-priors accounts of this period. Our working hypothesis is that no long term structural change occurred.

A Bayesian model with known precision then has problems explaining the initial slide in prices. Roughly, if precision is high, the arrival of enough bad news to explain the first week is highly unlikely, while with low precision bad news will not be incorporated into prices in the first place. With ambiguous precision, bad news are taken especially seriously and hence a much less extreme sequence of signals suffices to account for prices

---

[5] The main point may be relevant more widely. 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.

[6] In a risky world, the law of large numbers says that a portfolio of many independent and securities has zero risk. In an ambiguous world, the return on a portfolio of many independent securities may still be ambiguous. See Marinacci [35] for a version of the law of large numbers under ambiguity.

in the first week. In this sense, the ambiguity aversion model outperforms the Bayesian model. In addition, it attributes most of the initial drop in prices to the prospect of ambiguous news. In contrast, the Bayesian model relies on large movements in expected growth to rationalize price movements.[7]

2 SOME STYLIZED EXAMPLES

In this section we present a series of examples to illustrate limitations of the Bayesian approach and to introduce our model informally.[8] In one way or another, they are all based on the Ellsberg Paradox, a version of which is reviewed below. While the Ellsberg Paradox applies also in static settings, we focus on concerns that are specific to the Bayesian model of learning and hence to dynamic settings. Three stylized examples below exhibit behavior that is intuitive, yet at odds with the Bayesian model. We use them to motivate the building blocks of our model and then we argue that the key principles behind the examples are applicable more widely in economic settings.

2.1 The Ellsberg Paradox

Consider two urns, each containing four balls that are either black or white. The agent is told that the first "risky" urn contains two balls of each color. For the second "ambiguous" urn, he is told only that it contains at least one ball of each color. Suppose one ball is to be drawn from each urn. In what follows, a **bet on the color of a ball** is understood to pay one dollar (or one util) if the ball has the desired color and zero otherwise. Intuitive behavior pointed to by Ellsberg is the preference to bet on drawing black from the risky urn as opposed to the ambiguous one, and a similar preference for white. This behavior is inconsistent with any single probability measure on the associated state space, but can be explained by the multiple-priors model.

Formally, let the state of the world be $s = (s^r, s^a) \in S = \{B, W\}^2$. Then s^a denotes the color of the ball drawn from the ambiguous urn and the indicator function $\mathtt{1}_{s^a}$ denotes the corresponding bet. Similarly for the other urn and bets. Given a set $\mathcal{P}$ of probability measures on S, the multiple-priors utility of the bet $\mathtt{1}_{s^a}$ is

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

Thus Ellsberg-type behavior is accommodated if $\mathcal{P}$ contains a probability measure that

[7] More precisely, a Bayesian story attributes movements in prices to changes either in expected future growth or in risk. If the relevant news are mainly about medium term effects of the attack, changes in risk, that is, in the conditional covariance between consumption and returns, are small.

[8] More formal treatments follow in Section 3, where recursive multiple-priors utility is reviewed, and in Section 4, where our learning model is described more precisely. Some readers may wish to proceed directly from this informal description to the applications. Others may wish to skip this section and proceed directly to the formalism.

assigns probability greater than $\frac{1}{2}$ to black in the ambiguous urn and another measure that assigns probability less than $\frac{1}{2}$.

2.2 Sampling from Ellsberg Urns: Ambiguous Prior Beliefs

A simple example of learning obtains if repeated sampling with replacement is permitted from the given ambiguous urn. Compare the choice of bets in light of different sequences of draws. Two properties of learning are intuitive. First, Ellsberg-type behavior should be exhibited in the short run, because a few draws cannot plausibly resolve the initial ambiguity. Second, in the long run, learning should resolve ambiguity, since the ambiguous urn remains unchanged. As the number of draws increases, ambiguity should diminish and asymptotically the agent should behave as if he knew the fraction of black balls was equal to their empirical frequency. While the Bayesian model cannot deliver the first property, a model of learning under ambiguity satisfying both properties can be constructed by augmenting multiple-priors utility with specific assumptions on beliefs along the lines of Marinacci [36].

To see this, it is helpful to describe first the natural Bayesian specification. Let $\theta \in \Theta = \{2, 3\}$ denote the number of black balls in the ambiguous urn. The agent has prior μ_0 over Θ. If $\ell(\cdot|\theta)$ is the likelihood function suggested by the objective features of the setting described above, then after the t draws $s_1^t = (s_1, ..., s_t)$, beliefs about the next draw are given by

$$p_t(\cdot|s_1^t) = \int_\Theta \ell(\cdot|\theta)\ d\mu_t\left(\theta|s_1^t, \mu_0, \ell\right). \tag{1}$$

Here $\mu_t(\cdot|s_1^t, \mu_0, \ell)$ is the posterior over Θ determined by Bayes' Rule.

Ambiguity can be introduced by permitting a set $\mathcal{M}_0$ of priors on Θ. The counterpart of (1) is that conditional beliefs about the next draw are given by the set of 1-step-ahead conditionals

$$\mathcal{P}_t\left(s_1^t\right) = \left\{\int_\Theta \ell(\cdot|\theta)\ d\mu_t\left(\theta|s_1^t, \mu_0, \ell\right) : \mu_0 \in \mathcal{M}_0\right\}. \tag{2}$$

In other words, every prior in $\mathcal{M}_0$ is updated using the likelihood ℓ. After finitely many draws, there will still be a set of posteriors μ_t, and hence a nonsingleton set $\mathcal{P}_t$ of 1-step-ahead beliefs. This accommodates Ellsberg-type behavior in the short run. However, as time goes by, the influence of the priors diminishes and (under standard regularity conditions) $\mathcal{P}_t$ converges to a singleton.[9] Ambiguity is thus resolved in the long run.

2.3 Ambiguous Signals

The previous example is special because the physical environment suggests that the conditional distribution of signals is known. We now modify it to introduce ambiguous

[9] See Marinacci [36] for details.

signals. Suppose that one ball is added to each of the two urns. Its color can be black or white, as determined by tossing a fair coin: the 'coin ball' is black if the coin toss produces heads and white otherwise. The two coin tosses that determine the color of the coin ball are independent across urns. Modify also the type of bet. Instead of betting on the next draw, the agent is now invited to bet on the color of the coin ball.

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

The above rankings could be exhibited by a Bayesian agent who holds a single subjective probability belief about the composition of the ambiguous urn. However, his belief must imply that the colors of the non-coin balls in the ambiguous urn depend on the color of the coin ball.[10] Since such a belief does not respect independence of the coin ball, conclude that the Bayesian model cannot satisfactorily capture the difference between the two urns. A difference from the standard Ellsberg Paradox is that here the prior belief about the color of the coin ball is unambiguous. Nevertheless, ambiguity in the signal induces conditional Ellsberg-type behavior. Our model accommodates this by permitting multiple likelihoods.

It is natural to let the parameter be $\theta = (\theta^r, \theta^a) \in \Theta = \{B, W\}^2$, where θ^r (θ^a) denotes the color of the coin ball in the risky (ambiguous) urn, and to assume the single prior μ_0, independent across urns and satisfying $\mu_0(\theta^a = B) = \mu_0(\theta^r = B) = \frac{1}{2}$. If the ambiguous urn contains a black coin ball, then the likelihood of drawing a black ball

[10] To see this, let p^b (p^w) denote the conditional probability of drawing a black (white) ball from the non-coin balls given that the coin ball is black (white). The probability of winning the bet on a black coin-ball in the ambiguous urn after a black draw is

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

Independence of the coin ball is respected iff $p^b = 1 - p^w$, which implies $\pi = \frac{1+4p^b}{1+8p^b}$. Thus $\pi < \frac{3}{5}$, the probability of winning the corresponding bet for the risky urn, iff $p^b > \frac{1}{2}$. Argue similarly that $p^w > \frac{1}{2}$. Hence $p^b > 1 - p^w$, contradicting independence.

from that urn is given by

$$\ell_\lambda (B \mid B) = \tfrac{\lambda+1}{5},$$

where λ denotes the number of black non-coin balls; similarly for $\ell_\lambda (\cdot \mid W)$. Given the symmetry of the environment, a natural set of likelihoods is

$$\mathcal{L}_\epsilon = \{\ell_\lambda : \lambda \in [2-\epsilon, 2+\epsilon]\}, \tag{3}$$

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

In the special case where $\epsilon = 1$, the agent attaches equal weight to all logically possible urn compositions $\lambda = 1$, 2, or 3. More generally, (3) incorporates a subjective element into the specification. Just as subjective expected utility theory does not impose connections between the Bayesian prior and objective features of the environment, so too the set of likelihoods is subjective (varies with the agent) and is not uniquely determined by the facts. For example, the agent might attach more weight to the 'focal' likelihood corresponding to $\lambda = 2$ as opposed to the more extreme scenarios $\lambda = 1, 3$. The parameter ϵ can be interpreted as the weight attached to the latter scenarios, as opposed to the focal likelihood. Indeed, $\mathcal{L}_\epsilon$ can be rewritten as[11]

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

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

Implicit is that the agent does not understand the mechanism underlying λ well enough even to theorize about how its value is determined. Thus he considers the entire set $\mathcal{L}_\epsilon$ of likelihoods constructed as above (and assuming also the obvious likelihood for the risky urn and independence across urns). Denote by $\mathcal{L}$ the set of likelihoods relevant for the Cartesian product space defined by the two urns.

After one draw from each urn with outcome $s_1 = (s_1^r, s_1^a)$, updated beliefs about the next draw are given by the set of 1-step-ahead conditionals,

$$\mathcal{P}_1 (s_1) = \left\{ \int_\Theta \ell (\cdot|\theta)\ d\mu_1 (\theta | s_1, \mu_0, \ell) : \ell \in \mathcal{L} \right\}. \tag{4}$$

Thus the single prior has turned into a set of posteriors (compare with (2)). Conditional Ellsberg-type behavior is accommodated in this way.

One may view the draw of a ball as an experiment that is conducted to learn the color of the coin ball in an urn. For the risky urn, there is exactly one probability over draws for every θ. In other words, the agent is sure about the connection between the true parameter and the data generated by the experiment. This may be a tenable view of controlled laboratory experiments. It seems less plausible for 'experiments' observed

[11] This is a form of the ϵ-contamination model employed in robust statistics (see Walley [48], for example). In economic modeling, it is used in Epstein and Wang [20], for example.

in economic settings, such as in asset markets where returns are influenced by a large number of factors that are difficult to foresee and describe. Those factors obscure the link between the underlying feature of the environment that is of interest (such as the fundamental value of the security) and the data that are generated. Ambiguous signals can capture this intuition. Noisy signals cannot.

2.4 Sequences of Ambiguous Signals: Indistinguishable Experiments

Extend the preceding example to several draws. Assume that one ball from each urn is drawn every period and then replaced. In addition, assume that *all* non-coin balls from the ambiguous urn are replaced every period with a new set of non-coin balls, subject still to the restriction that one ball of each color is present. The only other information given to the agent about the replacement mechanism is that it is *memoryless*, that is, he should not expect a pattern in the sequence of sets of non-coin balls. Because the next urn always contains two new balls of unknown color, Ellsberg-type behavior is natural and it should persist *at any horizon*. In particular, the Bayesian model does not apply. However, our model with multiple-likelihoods captures intuitive features of this situation with a sequence of 'changing' ambiguous urns.

The fact that the set of likelihoods is the same for each time period captures the agent's perception that every new set of non-coin balls is equally ambiguous. This is reasonable because the description of the environment does not give any indication of patterns in the sequence of ambiguous urns. Using language of Walley [48], the ambiguous signals here are results of a sequence of independent and *indistinguishable* experiments. This is to be contrasted with experiments that are *identical*, that is, held under identical conditions, as illustrated by the example in Section 2.2 where the ambiguous urn was fixed through time.[12] Walley expresses the contrast also as one between "symmetry of evidence" and "evidence of symmetry". In the first case, there may be no reason to distinguish between experiments. However, particularly in real world settings, available evidence may be so meagre that there are no grounds for being confident that the experiments are identical - think of there being many underlying and poorly understood factors that influence data and that may vary over time. The Bayesian model is not able to distinguish between these two cases, because both must be captured by conditionally i.i.d. signals.

Since the Bayesian model views signals as results of identical experiments, a Bayesian learner is naturally very ambitious - he tries to, and eventually succeeds in, learning *all* events of interest in a given environment. Again, this is reasonable in some transparent environments, such as when sampling from the risky urn. For the ambiguous urn, and more generally in economic environments, a more modest stance seems plausible. There are often aspects of the environment that agents think are impossible to ever know. Accordingly, agents concentrate on trying to learn about a limited set of features.

[12] See Epstein and Schneider [19] for more on this distinction.

2.5 Beliefs in the Long Run

To see further desirable properties of a model of learning under ambiguity, contemplate again beliefs about the coin ball, but now conditional on a large number of draws. Consider a sequence of draws that leads to a limiting empirical frequency for black given by ϕ_∞. For ϕ_∞ sufficiently large (small), the agent should become confident that the coin ball is black (white). At issue are the meanings of 'sufficiently large' and 'sufficiently small.' For a Bayesian agent, who assigns a $\left(\frac{1}{2}, \frac{1}{2}\right)$ prior to the color composition of the non-coin balls in every ambiguous urn, precisely as for the risky urn, there is a unique cut-off value of $\frac{1}{2}$. He becomes certain asymptotically that the coin ball is black (white) according as $\phi_\infty > (<) \frac{1}{2}$.

An ambiguity averse agent, however, takes into account the possibility that many or all urns have more than 2 black non-coin balls and that this is what leads to more than half the draws being black. Ambiguity aversion means that he attaches *some* weight to this possibility, though not necessarily as much weight as he does to the 'focal' uniform color composition. But given *any positive weight* to the above noted possibility, then it is sensible to require stronger evidence than does the above Bayesian. Thus the ambiguity averse agent is eventually confident that the coin ball is black if (and only if) $\phi_\infty > \phi^{**}$ for some cut-off $\phi^{**} > \frac{1}{2}$. Similarly, confidence that the coin ball is white is sensible only if the evidence is strong in the sense that $\phi_\infty < \phi^*$ for some lower cut-off $\phi^* < \frac{1}{2}$.

This leaves the interval (ϕ^*, ϕ^{**}) where evidence is inconclusive. The size of this interval presumably depends on the weight attached to alternative scenarios and thus varies with the agent. (Because the Bayesian agent attaches zero weight, his interval is degenerate.) For ϕ_∞ in the inconclusive range, the agent might reason as follows: "If most of the urns have 3 white non-coin balls, then I would be certain that the coin ball is black and thus I would not be willing to pay any positive price for a bet on white. Similarly, for the symmetric scenario and a bet on black. Because both scenarios are possible and since I want to guard against the worst case, then I should not pay for a bet on either color." In other words, posterior beliefs about the color of the coin ball correspond to the entire interval $[0, 1]$, a situation that is sometimes referred to as *complete ignorance*.

Some agents might take this conservative stance. However, others might reason as follows given $\frac{1}{2} < \phi_\infty < \phi^{**}$: "I had incomplete *prior* information about the non-coin balls, but now that I look *back* at the large number of draws I have seen, it seems somewhat unlikely that in most cases there were 3 black non-coin balls. I conclude that the coin ball is probably black." In this way, reevaluation can eliminate the region where evidence is inconclusive even in the limit.

It seems reasonable to permit both types of behavior. To this end, our model features a parameter that measures the willingness to reevaluate views about how past data were generated (the conservative agent above was totally unwilling to reevaluate). Formally, adopt the same parameter space, prior and set of likelihoods $\mathcal{L}_\epsilon$ as in the previous example. The conservative agent's 1-step ahead beliefs are represented by the set

$$\mathcal{P}_t\left(s_1^t\right) = \left\{ \int_\Theta \ell_{t+1}\left(\cdot|\theta\right)\, d\mu_t\left(\theta|s_1^t, \mu_0, \ell_1^t\right) : \ell_\tau \in \mathcal{L}_\epsilon \text{ for } \tau \leq t+1 \right\}, \tag{5}$$

where μ_t now depends on a sequence $\ell_1^t = (\ell_1, ..., \ell_t)$ of likelihoods arbitrarily selected from $\mathcal{L}_\epsilon$. The set in (5) describes beliefs about the next ball to be drawn. Since ℓ_{t+1} ranges over all of $\mathcal{L}_\epsilon$ no matter how large is t, the model predicts that Ellsberg-type behavior towards bets on the next ball will persist forever.

We have seen that, if all sequences of likelihoods in $\mathcal{L}_\epsilon$ are permitted, the set of posteriors $\{\mu_t\}$ stays nondegenerate forever if the fraction of black balls converges to the appropriate interval. For the specifications of priors and likelihoods adopted, the cut-off values are (see Section 4.4)

$$\phi^* = \frac{\log\left(\frac{2+\epsilon}{3+\epsilon}\right)}{\log\left(\frac{4-\epsilon^2}{9-\epsilon^2}\right)} \quad \text{and} \quad \phi^{**} = \frac{\log\left(\frac{2-\epsilon}{3-\epsilon}\right)}{\log\left(\frac{4-\epsilon^2}{9-\epsilon^2}\right)}. \tag{6}$$

In particular, the corresponding interval shrinks monotonically about $\frac{1}{2}$ as ϵ declines to 0. Thus a modeler who feels that complete ignorance in the long run is reasonable at best for values of ϕ_∞ sufficiently close to $\frac{1}{2}$, can ensure this by taking ϵ suitably small. However, this comes at the cost of assuming that ambiguity aversion is small also in the short run.

To eliminate complete ignorance about the coin ball in the long run without unduly restricting ambiguity aversion in the short run, we model the less conservative attitude described above and permit the agent to discard certain sequences of likelihoods as implausible. A natural criterion is how well the prior and the sequence of likelihoods explains the observed history s_1^t, as measured by the probability

$$\Pr\left(s_1^t; \mu_0, \ell_1^t\right) = \int \prod_{j=1}^{t} \ell_j\left(s_j|\theta\right) d\mu_0\left(\theta\right).$$

In particular, assume that the agent retains only sequences of likelihoods such that

$$\Pr\left(s_1^t; \mu_0, \ell_1^t\right) \geq \alpha \max_{\mu_0,\, \ell_1^t} \Pr\left(s_1^t; \mu_0, \ell_1^t\right),$$

where $0 \leq \alpha \leq 1$ is a parameter. The special case $\alpha = 0$ captures the conservative agent. With $\alpha > 0$, an extreme sequence of likelihoods corresponding to "3 non-coin balls are black" is excluded once the empirical frequency of black balls is small enough. (Such a sequence may be readmitted subsequently if sufficiently many black balls are drawn.) Theorem 1 below implies that for any $\alpha > 0$, reevaluation leads to the cut-off values

$$\phi^* = \phi^{**} = \frac{1}{2}, \tag{7}$$

just as for the benchmark Bayesian agent. Thus limiting posteriors regarding the coin ball are identical for both agents. Nevertheless, these 2 agents behave differently, even in the limit, if we consider the ranking of bets on the color of the next ball (rather than on the color of the coin ball); only the ambiguity averse agent behaves in the intuitive Ellsbergian fashion. Moreover, they differ also in the short-run dynamics of their posteriors.

3 RECURSIVE MULTIPLE-PRIORS

We work with a period state space S, identical for all times, so that the (full) state space is $S^\infty = \Pi_{t=1}^\infty S_t$, $S_t = S$ all $t \geq 1$. 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. The agent observes the realization s_t in S_t at time t. Thus her information at t consists of the history $s_1^t = (s_1, ..., s_t)$.[13]

The agent ranks (nonnegative real-valued and adapted) consumption processes $c = (c_t)$. At any time $t = 0, 1, ...,$ and given the history s_1^t, her ordering is represented by the conditional utility function V_t, defined recursively by

$$V_t(c) = \min_{Q \in \mathcal{P}_t(s_1^t)} E_Q \left[u(c_t) + \beta V_{t+1}(c) \right], \tag{8}$$

where: β and u satisfy the usual properties, $\mathcal{P}_t(s_1^t)$ is a set of 1-step-ahead measures conditional on the history s_1^t and where the dependence of $V_t(c)$ on s_1^t is suppressed in the notation. The set $\mathcal{P}_t(s_1^t)$ embodies beliefs about the next step (about S_{t+1}) given the history of observations s_1^t. Such beliefs reflect ambiguity when $\mathcal{P}_t(s_1^t)$ is a nonsingleton. We refer to $\{\mathcal{P}_t\}$ as the process of conditional 1-step-ahead beliefs.

Recursive multiple-priors utility was introduced in [20] without axiomatic foundations, but these have been provided recently in [18]. The latter also provides a reformulation of utility that makes clearer the connection to the Gilboa-Schmeidler model. The key to establishing this reformulation is to observe that the collection of all sets $\mathcal{P}_t(s_1^t)$, as one varies over times and histories, determines a unique set of priors $\mathcal{P}$ on S^∞ satisfying the regularity conditions specified in [18].[14] Thus one obtains the following equivalent and explicit formula for utility:

$$V_t(c) = \min_{Q \in \mathcal{P}} E_Q \left[\Sigma_{s \geq t} \beta^{s-t} u(c_s) \mid s_1^t \right]. \tag{9}$$

In particular, this expression shows that each conditional ordering conforms to the multiple-priors model [23], with the set of priors for time t determined by updating the set $\mathcal{P}$ prior-by-prior via Bayes' rule.

Besides its simple and (in our view) appealing axiomatic foundations, recursive multiple-priors utility is attractive also because dynamic consistency ensures that behavior is determined, via preference maximization, without the need to resort to auxiliary and invariably ad hoc assumptions about how intrapersonal conflicts are resolved. In addition, the model permits a natural way to describe "greater ambiguity aversion," namely through expansion of the set $\mathcal{P}$.

An essential feature of recursive multiple-priors utility is that the process of conditional 1-step-ahead beliefs $\{\mathcal{P}_t\}$ is restricted only by technical regularity conditions. This

[13] Measures on S^∞ are understood to be defined on the product σ-algebra on S^∞ and those on any S_t are understood to be defined on the power set of S_t.

[14] In the infinite horizon case, uniqueness obtains only if $\mathcal{P}$ is assumed also to be regular in a sense defined in Epstein and Schneider [19], generalizing to sets of priors the standard notion of regularity for a single prior.

is important for learning theory because, as a description of the way in which the agent's view of the next step or experiment depends on history, the process $\{\mathcal{P}_t\}$ is the natural vehicle for modeling learning. It remains only to add restrictions on this process and thus on how the agent responds to data. We proceed now to describe such restrictions.

4 THE MODEL

We present the model in two steps. Begin with processes $\{\mathcal{P}_t\}$ that permit what we call a "statistical representation". The general model follows.

4.1 A Special Case: Statistical Representations

Since we extend the Bayesian learning model based on (1), we first specify the latter more precisely. To do so, denote by $\Delta(S)$ the set of all probability measures on S. Given a measurable space $(\Theta, \mathcal{B})$, interpreted as the space of unknown parameters, refer to $\ell : \Theta \longrightarrow \Delta(S)$ as a **likelihood function** if

$$\theta \longmapsto \ell(A \mid \theta) \text{ is } \mathcal{B}\text{-measurable for every } A \in \mathcal{S}.$$

Beliefs of a subjective expected utility maximizer are represented by a probability measure p on S^∞. Any such p determines a process $\{p_t\}$ of 1-step-ahead conditional measures and, in turn, it can be uniquely reconstructed from this process by the usual rules of probability calculus. Thus beliefs may be represented equivalently by the process $\{p_t\}$.

Say that $(\Theta, \mathcal{B}, \mu_0, \ell)$ is a **statistical representation** for $\{p_t\}$, and thus indirectly for p, if: $(\Theta, \mathcal{B}, \mu_0)$ is a probability space and ℓ is a likelihood function satisfying (1), which is rewritten here for convenience:

$$p_t(\cdot) = \int_\Theta \ell(\cdot \mid \theta)\, d\mu_t(\theta).$$

As mentioned earlier, μ_t is the posterior belief about Θ; it is defined recursively by

$$d\mu_t(\cdot) = \frac{\ell(s_t \mid \cdot)}{\int_\Theta \ell(s_t \mid \theta')\, d\mu_{t-1}(\theta')}\, d\mu_{t-1}(\cdot)\,. \tag{10}$$

Refer to Θ as a **(statistical) parameter** space and to μ_0 as the **prior** on Θ. The representation suggests the interpretation that each $\theta \in \Theta$ describes those features of the environment about which the agent is uncertain at the start. Her uncertainty is modeled by the subjective prior μ_0. Note that the parameter θ is viewed as providing a **complete** description of the environment in the sense that it determines a **unique** conditional 1-step-ahead probability law $\ell(\cdot \mid \theta)$ on S. In particular, the likelihood embodies the decision-maker's view about how unknown parameters are reflected in finite samples of data.

Discussion in Section 2 pointed to the need for modifications of the Bayesian model. We would like to permit the decision-maker to feel unsure about how parameters are reflected in data, or equivalently, to feel uncertain about whether different trials of the experiment are identical. Thus each θ should capture only some aspects of the data generating mechanism, suggesting a need for multiple likelihoods. Second, there may be *a priori* ambiguity about the true θ, suggesting a need for multiple-priors over the parameter space.

To proceed more formally, for a given measurable space $(\Theta, \mathcal{B})$, say that

$$\mathcal{L} : \Theta \text{ Ã } \Delta(S)$$

is a *likelihood correspondence* if it admits a likelihood selection, by which we mean that there exists a likelihood function ℓ such that

$$\ell(\cdot \mid \theta) \in \mathcal{L}(\cdot \mid \theta) \text{ for every } \theta.$$

Abbreviate the latter by $\ell \in \mathcal{L}$. Thus every likelihood function can be identified with a degenerate likelihood correspondence. The interpretation is that the agent perceives some factors, modeled by θ, as common across time or experiments, and others, modeled by the multiplicity of $\mathcal{L}$, as variable across time in a way that 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, because θ is fixed over time, she can try to learn the true θ, but she has decided that she will not try to (or is not able to) learn more.

Initial beliefs about the parameter are given by $\mathcal{M}_0 \subset \Delta(\Theta)$, a set of measures on $(\Theta, \mathcal{B})$. The evolution of these beliefs is represented by a sequence of sets of posteriors $\mathcal{M}_t$. The formation of beliefs about s_{t+1} combines beliefs $\mathcal{M}_t$ about parameters with beliefs (or information) $\mathcal{L}$ about how these parameters impact data generation. A natural analogue of the Bayesian model admits all posteriors that can be delivered by some likelihood function and some prior, that is,

$$\mathcal{M}_t\left(s_1^t\right) = \left\{\mu : d\mu(\theta) = \frac{\ell(s_t \mid \theta) d\mu_{t-1}(\theta)}{\int_\Theta \ell(s_t \mid \theta')\, d\mu_{t-1}(\theta')},\ \mu_{t-1} \in \mathcal{M}_{t-1}\left(s_1^{t-1}\right),\ \ \ell \in \mathcal{L}\right\}. \tag{11}$$

Now consider a collection $\{\mathcal{P}_t\}$ of conditional 1-step-ahead beliefs embodying the beliefs of an agent who conforms to recursive multiple-priors utility. Say that $(\Theta, \mathcal{B}, \mathcal{M}_0, \mathcal{L})$ is a *statistical representation* for $\{\mathcal{P}_t\}$, and thus indirectly for the corresponding set of priors $\mathcal{P}$ on S^∞, if: $(\Theta, \mathcal{B})$ is a measurable space, $\mathcal{M}_0 \subset \Delta(\Theta)$ and $\mathcal{L}$ is a likelihood correspondence satisfying

$$\mathcal{P}_t(s_1^t) = \{p_t(\cdot) = \int_\Theta \ell(\cdot \mid \theta)\, d\mu_t(\theta) : \mu_t \in \mathcal{M}_t(s_1^t),\, \ell \in \mathcal{L}\},$$

or in more compact notation,

$$\mathcal{P}_t = \int_\Theta \mathcal{L}(\cdot \mid \theta)\, d\mathcal{M}_t(\theta). \tag{12}$$

The case where Θ is a singleton has $\mathcal{P}_t = \mathcal{L}$ for all times and histories, implying therefore that there is no learning. This 'IID' case is studied in detail in [19].[15]

4.2 The General Model

Beliefs that permit statistical representations are special in that they shut off one channel of learning that is unique to the ambiguous case. We have interpreted $\mathcal{L}$ as the set of likelihoods possible *in the future*. Since the agent has decided he cannot learn the true sequence of likelihoods, it is natural that beliefs about the future must be based on the *whole* set $\mathcal{L}$ as in (12). On the other hand, we might expect her to revise, 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 sequences $\ell_1^t = (\ell_1, ..., \ell_t)$ of likelihoods more or less plausible. Similarly, $\mathcal{M}_0$ describes ambiguity at time 0 about parameters to be learned, reflecting the result of prior unmodeled observations, or, more generally, prior experience in similar situations. With hindsight she might resolve some ambiguity about what generated these prior observations and update her view about the relevance for the current situation of prior experience.

Reevaluation of how patterns were reflected in past data and experience is important for beliefs about the future because it can lead to a different view of which priors μ_0, and hence also, which *posteriors* are relevant. For example, if a prior belief is such that the sample appears very unlikely under *every* conceivable likelihood sequence (together with the given prior), then it is intuitive that the agent will consider the prior to be not particularly relevant. She might even *disregard* the posteriors obtained from that prior. Whether she actually does so presumably depends on an aspect of her *a priori* view of the environment and on her willingness to reevaluate the extent of ambiguity with hindsight. We now present a more general learning model where this willingness is a parameter.

View any pair $(\mu_0, \ell_1^t) \in \mathcal{M}_0 \times \mathcal{L}^t$ of initial beliefs and a sequence of likelihoods as a *theory* at t about how the sample to that point was created. For every theory, the agent can compute the likelihood assigned to the sample s_1^t. If the sample appears very unlikely given some theory about its creation, the agent might plausibly disregard posteriors based on this theory. In addition, how 'unlikely' a sample appears is presumably judged not in absolute terms, but relative to other available theories.

The preceding motivates the following definition of updating behavior: For any α, a parameter in the unit interval, define

[15] In particular, we prove a strong LLN appropriate for recursive multiple-priors.

$$\mathcal{M}_t^\alpha(s_1^t) = \{\mu_t\left(s_1^t;\mu_0,\ell_1^t\right) \in \mathcal{M}_t : \int \Pi_{j=1}^t \ell_j\left(s_j|\theta\right) d\mu_0(\theta) \geq \alpha \max_{\substack{\tilde{\mu}_0 \in \mathcal{M}_0 \\ \tilde{\ell}_1^t \in \mathcal{L}^t}} \int \Pi_{j=1}^t \tilde{\ell}_j\left(s_j|\theta\right) d\tilde{\mu}_0\}, \tag{13}$$

where $\mu_t\left(s_1^t;\mu_0,\ell_1^t\right)$ denotes the posterior obtained by updating μ_0 via (10), given s_1^t and the 'working hypothesis' that the data were generated by ℓ_1^t. The implied 1-step-ahead beliefs about the state space are described by

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

which generalizes (12). Refer to $(\Theta, \mathcal{B}, \mathcal{M}_0, \mathcal{L}, \alpha)$ as an α-ML (α-maximum likelihood) representation. Thus any statistical representation is a 0-ML representation.

Essentially, the agent disregards posteriors based on theories that do not pass a likelihood-ratio test against the alternative theory that puts maximum likelihood on the sample. The parameter α governs the extent to which the agent is willing to reevaluate her views about how past data were generated in the light of new sample information. In other words, it reflects her a priori judgement about how sample evidence should be weighed against prior evidence encoded in $\mathcal{M}_0$. If $\alpha = 0$, then data receive minimal weight and $\mathcal{M}_t^0 = \mathcal{M}_t$, the set of posteriors from the statistical representation. At the other extreme $\alpha = 1$, only posteriors based on theories assigning maximal weight to the sample are allowed. In general, the more weight placed on the sample, the less ambiguity there is about θ and the future: $\alpha > \alpha'$ implies $\mathcal{M}_t^\alpha \subset \mathcal{M}_t^{\alpha'}$ and $\mathcal{P}_t^\alpha \subset \mathcal{P}_t^{\alpha'}$.

The Bayesian model is the special case where both the set of priors and the set of likelihoods have only a single element. In particular, if there is a single likelihood and a single prior that is a Dirac measure on one parameter value, then data are perceived to be *i.i.d.* Another interesting special case occurs if $\mathcal{M}_0$ consists of several Dirac measures on the parameter space in which case there is a simple interpretation of the updating rule: $\mathcal{M}_t^\alpha$ contains all θ^*'s 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 $-2log\,\alpha$. In particular, if $\alpha > 0$, then there is updating, parameter values may be discarded or added to the set, and $\mathcal{P}_t^\alpha$ varies over time. In the extreme case $\alpha = 0$, ambiguity is not reevaluated and no learning takes place.

4.3 Informal Support

We do not yet know what axioms on preference must be added to those in [18] in order to deliver a statistical representation or the more general α-ML representation, and in that sense these are ad hoc. However, we argue next that formal axiomatic support is in an important sense limited even in the Bayesian model and that on informal grounds, our models are appealing in many settings. While the de Finetti Theorem delivers *some* statistical representation for any given exchangeable prior, in applications the modeler must choose a *particular* representation $(\Theta, \mathcal{B}, \mu_0, \ell)$ and there is no axiomatic foundation for any particular specification. Consequently, the latter is invariably ad hoc and is justified informally by what seems 'natural' or 'cognitively appealing' given the setting. Similarly, to apply our model, and even given a counterpart of the de Finetti Theorem, one would specify a representation $(\Theta, \mathcal{B}, \mathcal{M}_0, \mathcal{L}, \alpha)$ that seems natural. Our claim is that there are settings where a natural specification is easier to find in our framework than in the Bayesian one. The urns examples in Section 2 provide a number of such settings. Others are described in the applications below.

To be sure, we do not claim that the appeal of our model is universal. There are well known criticisms of the standard portrayal of an agent who thinks in terms of an exhaustive set of contingencies, conjures a prior over them and then applies Bayesian updating in response to new information. One could describe settings where our model is similarly problematic cognitively. We suggest only that it is suitable in a broader class of settings than the Bayesian model.

4.4 Properties

According to the exchangeable Bayesian model, the agent believes that any one of a set of i.i.d. processes is being observed, and learning gradually resolves uncertainty about the 'true' process. While a typical measure of 'estimation risk', such as the posterior variance, might fluctuate during the learning process, it eventually shrinks. The agent then behaves as if he had been told the distribution of the 'true' i.i.d. process. Ex ante, he is sure that he will actually learn it.

Learning under ambiguity introduces three additional aspects. First, new measures of uncertainty must be used to describe 'estimation uncertainty'. For example, measures of maximum distance between posterior distributions are an option. Second, according to reasonable measures, uncertainty need not be reduced by the learning process. The earlier observation that ambiguous signals can induce uncertainty applies also to sequences of such signals. Third, the agent eventually behaves as if he faced independently and *indistinguishably* distributed data. The learning process under ambiguity thus also 'settles down', in the sense that, given many observations, the agent's one-step-ahead view of the world changes little in response to new data. However, he might still perceive ambiguity. If this is the case, he will know ex ante that some ambiguity may never be resolved.

Beliefs in the Short Run

The short run properties of the learning process are illustrated in Figure 1, using the urn example of Section 2.3. We set $\epsilon = 1$: the agent weighs equally all the logically possible combinations of non-coin balls. The top panel illustrates the aggregation of ambiguous signals. It shows the evolution of the posterior interval for a sequence of draws such that the number of black balls is $\frac{3t}{5}$, for $t = 5, 10, ...$ In particular, after the first 5 ambiguous

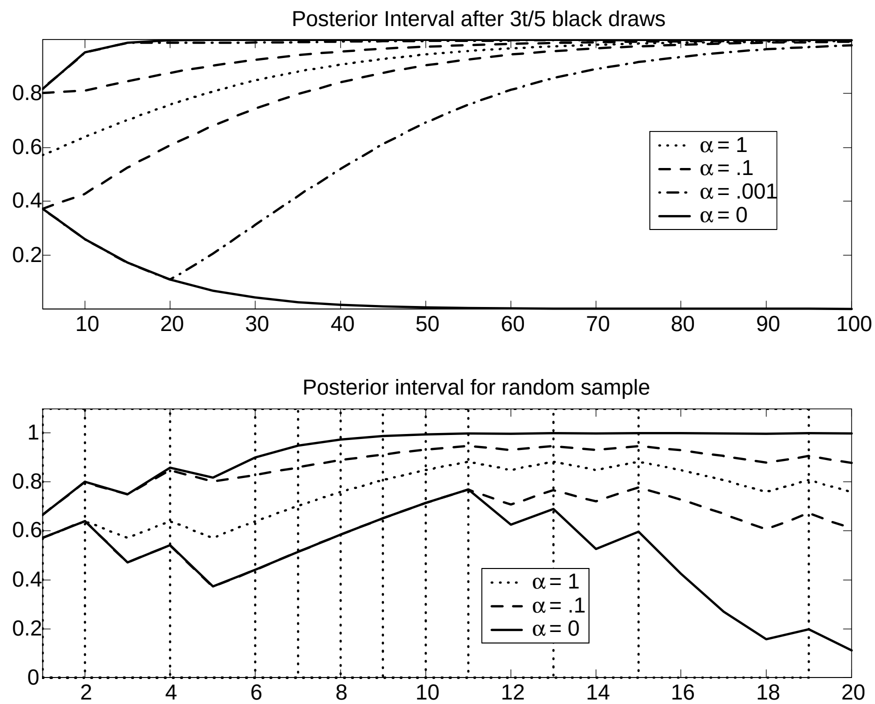

Figure 1: Posterior Interval is range of posterior probability that coin ball is black, $\mu_t(B)$. In top panel, sample is selected to keep fraction of black balls constant. In bottom panel, vertical lines indicate black balls drawn.

signals, with 3 black balls drawn, the agent assigns a posterior probability between .4 and .8 to the coin ball being black.

What happens if the same sample is observed again? There are two effects. First, a larger batch of signals permits more possible interpretations. For example, having seen 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 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 the posterior interval thus depends on how much agents reevaluate their views. For agents with $\alpha = 0$ and $\alpha = .001$, 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. If $\alpha = 0$, this is the only effect: the interval increases monotonically. An agent who does not reevaluate at all has an increasing stock of theories, none of which he is willing to disregard. However, any willingness to reevaluate ($\alpha > 0$) implies that large enough batches of ambiguous signals induce less ambiguity than smaller ones.

The lower panel of Figure 1 tracks the evolution of posterior intervals along a representative sample. Taking the width of the interval as a measure, the extent of ambiguity is seen to respond to data. In particular, a phase of many black draws (periods 5-11, for example) shrinks the posterior interval, while an 'outlier' (the white ball drawn in period 12) makes it expand again. This behavior is reminiscent of the evolution of the Bayesian posterior variance, which is also maximal if the fraction of black balls is one half.

Beliefs in the Long Run

To describe what happens in the long run, it is helpful to fix a 'true' data generating process. Assume that this process is i.i.d. and let the measure $\phi \in \Delta(S)$ denote the distribution of one component s_t. Equivalently, ϕ can be viewed as the empirical distribution of an infinite sample. 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. In our setup, any data density depends on the sequence of likelihoods that is used. In what follows, it is sufficient to focus on sequences such that the same likelihood is used whenever state s is realized. A likelihood sequence can then be represented by a collection $(\ell_s)_{s \in S}$.

Accordingly, define the log data density after maximization over the likelihood sequence by

$$H(\theta) := \max_{(\ell_s)_{s \in S}} \sum_{s \in S} \phi(s) \log \ell_s(s|\theta). \tag{14}$$

We now assume that $\theta^* = \arg\max_\theta H(\theta)$ is a singleton. This 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 ϕ can be used to discriminate θ^* from any other parameter value. The following proposition (proven in the appendix) summarizes the behavior of the posterior set in the long run. Write $\mathcal{M}_t \longrightarrow \{\delta_x\}$ if every sequence of posteriors from $\mathcal{M}_t$ converges to the Dirac measure δ_x, almost surely under the i.i.d. measure described by ϕ.

Theorem 1 Suppose that Θ is finite and that for every $\theta \in \Theta$, there is $\mu_0 \in \mathcal{M}_0$ with $\mu_0(\theta) > 0$. Suppose also that the above identification condition is satisfied.

(a) If for all $\theta \neq \theta^*$,

$$\sum_{s \in S} \phi(s) \max_{\ell \in \mathcal{L}} \log \frac{\ell(s|\theta)}{\ell(s|\theta^*)} < 0, \tag{15}$$

then $\mathcal{M}_t^0 \longrightarrow \{\delta_{\theta^*}\}$.

(b) If $\alpha > 0$, then $\mathcal{M}_t^\alpha \longrightarrow \{\delta_{\theta^*}\}$.

For a statistical representation $(\alpha = 0)$, the set of posteriors converges to δ_{θ^*} under condition (15). The latter requires that the agent cannot envision a sequence of likelihoods that makes any alternative parameter θ look more likely than θ^*, given the infinite sample with empirical frequencies ϕ.[16] If the agent reevaluates theories ($\alpha > 0$), and if the identification condition holds, then in the long run only the maximum likelihood sequence is a permissible scenario and the set of posteriors converges to a singleton.

The theorem is consistent with the earlier discussion (Sections 2.4-2.5) of learning in the long run in the example of ambiguous urns. In conformity with earlier notation, let ϕ_∞ denote the probability that a black ball is drawn "under the truth". Hence

$$\begin{aligned} H(\theta) &= \phi_\infty \max_{\lambda_1} \log \frac{1_{\{\theta=B\}} + \lambda_1}{5} + (1-\phi_\infty) \max_{\lambda_0} \log \frac{5 - 1_{\{\theta=B\}} - \lambda_0}{5} \\ &= \phi_\infty \log \frac{1_{\{\theta=B\}} + 3}{5} + (1-\phi_\infty) \log \frac{4 - 1_{\{\theta=B\}}}{5} \end{aligned}$$

It follows that the identification condition is satisfied except in the knife-edge case $\phi_\infty = \frac{1}{2}$. Moreover, $\theta^* = B$ if and only if $\phi_\infty > \frac{1}{2}$. Thus the theorem implies that an agent who reevaluates his views ($\alpha > 0$) and observes a large number of draws with a fraction of black balls above half "almost" believes that the color of the coin ball is black. The role of the size of α is only to regulate the speed of convergence to this limit.

An agent with $\alpha = 0$ acts more conservatively. Let $\phi_\infty > \frac{1}{2}$, so that $\theta^* = B$. Condition (15) reduces to

$$\phi_\infty \max_{\lambda_1} \log \frac{\lambda_1}{1+\lambda_1} + (1-\phi_\infty) \max_{\lambda_0} \log \frac{5-\lambda_0}{4-\lambda_0} < 0;$$

but this is satisfied only if ϕ_∞ is greater than a cutoff $\phi^{**} = \frac{log\,(1/2)}{log\,(3/8)}$.[17] As discussed in Section 2.5, "unlikely" sequences such as those with $\lambda_1 = \lambda_0 = 3$ are given weight by the conservative agent, but are eventually discarded if there is reevaluation.

5 DYNAMIC PORTFOLIO CHOICE

Portfolio choice is a natural application of our model. Given the substantial uncertainty about asset return processes, portfolio recommendations that ignore this uncertainty seem off the mark.

5.1 The Decision Problem

Consider an agent who invests for T quarters. He cares only about terminal wealth, but rebalances his portfolio every quarter. He can invest in a riskless asset with constant interest rate r and in stocks with uncertain returns $R(s_t)$. The state can take two values

[16] This is not implied by the identification condition (14), except in special cases such as the Bayesian case. The identification condition requires that θ^* is more likely than θ only when the sequence of likelihoods is actually the maximum likelihood sequence.

[17] The cutoff values in (6), for the case of general ϵ, can be proven similarly.

every period: $s_t \in \{0,1\}$. We choose parameters to approximate U.S. stock returns: $R(1) = 1.14$ and $R(0) = .92$. If s_t is iid with $\Pr\{s_t = 1\} = \frac{1}{2}$, the mean and variance of NYSE returns from 1927:Q3 to 2001:Q2 are matched exactly. We fix the riskless rate at $r = .01$ per quarter.[18]

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

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

The set of priors $\mathcal{M}_0$ on Θ is given by all the Dirac measures. For simplicity, we write $\theta \in \mathcal{M}_0$ if the Dirac measure on θ is included. If $\overline{\lambda} > 0$, returns are ambiguous signals. The likelihood parameter $\lambda_t \in [-\overline{\lambda}, \overline{\lambda}]$ represents movement in the poorly understood factors. 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

Before discussing the agent's optimization problem, we briefly describe the evolution of beliefs. The present example is convenient because the posterior set $\mathcal{M}_t^\alpha$ depends on the sample only through the fraction ϕ_t of high returns observed prior to t. More specifically, it is shown in the appendix that

$$\mathcal{M}_t^\alpha\left(s_1^t\right) = \left\{\theta \in \Theta : \; g(\theta;\phi_t) \geq \max_{\tilde{\theta}\in\Theta} g(\tilde{\theta};\phi_t) + \frac{\log \alpha}{t}\right\} \tag{16}$$

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

Using (16), it is straightforward to determine the limiting behavior of the 1-step-ahead beliefs $\mathcal{P}_t\left(s_1^t\right)$ as t becomes large. Suppose that the empirical frequency of high returns converges to ϕ_∞. Then $\mathcal{M}_t^\alpha$ collapses to the single number

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

Thus $\mathcal{P}_t\left(s_1^t\right)$ collapses to the set $\mathcal{L}(.|\theta^*)$, which consists of all probabilities on $S = \{0,1\}$ with

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

[18]Here 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 he behaves as if he had been told that it equals θ^*. If the realized empirical distribution is symmetric ($\phi_\infty = \frac{1}{2}$), then $\theta^* = \phi_\infty$. We use this fact below to calibrate the belief parameters.

Bellman Equation

The agent's optimization problem can be written using the fraction ϕ_t of high returns together with wealth w_t as state variables. The value functions J_t satisfy the terminal condition $J_{T+1}\left(w_{T+1}, \phi_{T+1}\right) = w_{T+1}$ and the Bellman equation

$$\begin{aligned} J_t\left(w_t, \phi_t\right) &= \max_{\Gamma_t} \min_{p \in \mathcal{P}_t\left(s_1^t\right)} \left(E_p\left(J_{t+1}\left(w_{t+1}, \phi_{t+1}\right)\right)^{1-\gamma}\right)^{\frac{1}{1-\gamma}} \\ &= \max_{\Gamma_t} \min_{\substack{\lambda_t \in [-\overline{\lambda}, \overline{\lambda}] \\ \theta \in \mathcal{M}_t^\alpha}} \{ \left(\theta + \lambda_t\right)\left(J_{t+1}(w_{t+1}\left(1\right), \phi_{t+1}\left(1\right))\right)^{1-\gamma} \\ &+ \left(1 - \theta - \lambda_t\right)\left(J_{t+1}(w_{t+1}\left(0\right), \phi_{t+1}\left(0\right))\right)^{1-\gamma} \}^{\frac{1}{1-\gamma}}, \end{aligned}$$

subject to the transition equations

$$\begin{aligned} w_{t+1}\left(s_{t+1}\right) &= \left(1 + r + \Gamma_t\left(R\left(s_{t+1}\right) - r\right)\right) w_t, \\ \phi_{t+1}\left(s_{t+1}\right) &= \frac{t\phi_t + s_{t+1}}{t+1}. \end{aligned}$$

Here Γ_t is the fraction of wealth w_t invested in the stock market. The value function has the form $J_t\left(w_t, \phi_t\right) = h_t\left(\phi_t\right) w_t$.

5.2 Preference Parameters

We specify $\beta = .99$ and set the coefficient of relative risk aversion γ equal to 2. It remains to specify the belief parameters α and $\overline{\lambda}$. The parameter $\overline{\lambda}$ determines how much the agent thinks that he will learn in the long run. To determine a value, he could pose the following question: "Suppose I see a large amount of data and that the fraction of high returns is $\frac{1}{2}$. How would I compare a bet on a fair coin with a bet that next quarter's returns are above or below the median?" By the Ellsberg Paradox, we would expect the agent to prefer the fair bet. He could then try to quantify this preference by asking: "What is the probability of heads that would make me indifferent between betting on heads in a coin toss and betting on high stock returns?" In light of the range of limiting probabilities given in (17), the result is $\frac{1}{2} - \overline{\lambda}$. We present results for values of $\overline{\lambda}$ ranging from 0 to .02.

The parameter α determines how fast the set of possible models shrinks. Here it can be motivated by reference to classical statistics. If signals are unambiguous ($\overline{\lambda} = 0$), there is a simple interpretation of our updating rule: $\mathcal{M}_t^\alpha$ contains all parameters θ^* such that the hypothesis $\theta = \theta^*$ is not rejected by an asymptotic likelihood ratio test performed

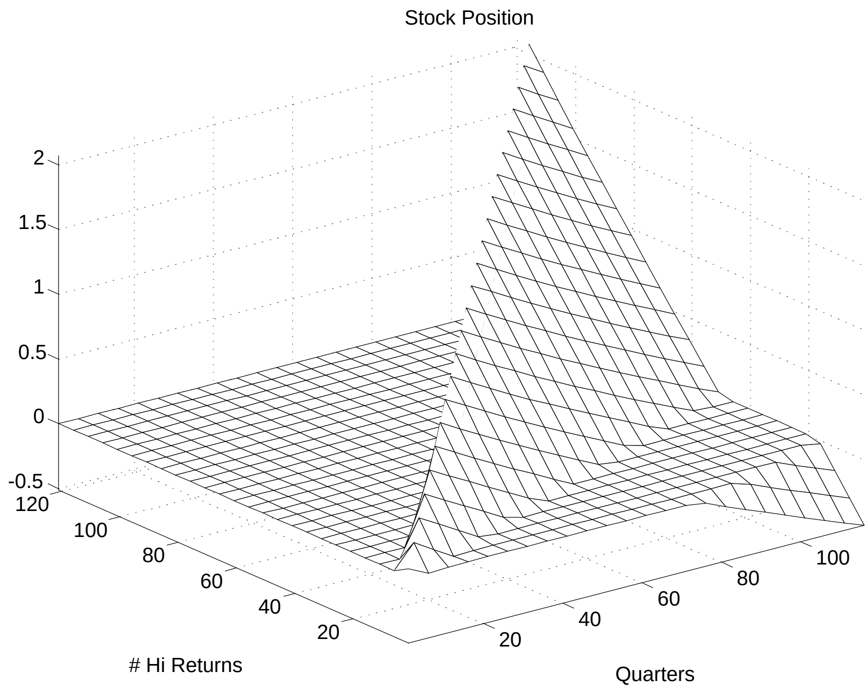

Figure 2: Optimal stock position for ambiguity averse agent with $\alpha = .14$ and $\lambda = 0$, for 30-year planning problem, beginning in 1971:Q3.

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.

5.3 Optimal Policies

Our leading example is an investor in 1971:Q3 ($t = 0$), who plans over various horizons (up to 30 years). He looks back on data starting in 1927:Q3. We generate a discretized returns sample by letting $s_t = 1$ if the NYSE return was above the mean in quarter t and $s_t = 0$ otherwise. Figure 2 shows the optimal stock position as a fraction of wealth for a 5 year horizon if $\alpha = .14$ and $\overline{\lambda} = 0$. The axis to the right measures time in quarters, the one to the left the number of high returns observed, $H = t\phi_t$. Thus only the surface above the region $H \leq t$ 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. The extreme is the path above the time axis which is taken if low returns are observed every period. In contrast, if high returns are observed (movements into the page, increasing H one for one with t), the stock position is typically increased. The extreme is now the ridge above the diagonal ($H = t$) which is taken if high returns are observed every period.

Non-Participation

The optimal policy surface has a flat piece at zero: when enough low returns are observed, agents do not participate in the market. Non-participation arises because the agent uses different 'worst case' probability measures to evaluate continuation utility given long or short positions. When the agent is long he employs the measure with the lowest possible expected return. In contrast, when short, as in the far right corner of the figure, his worst case scenario is provided by the measure with the highest expected return. In the non-participation region, continuation utility for any nonzero position is lower than what can be obtained by simply investing in the riskless asset.

For the static case, non-participation with multiple-priors utility was noted by Dow and Werlang [16]. The terminal period of our model corresponds to their setup. However, the emergence of a non-participation region in earlier periods is not obvious. In a dynamic model, agents invest in the stock market not only to exploit the expected equity premium, but also to hedge changes in the investment opportunity set. In our model, the latter is summarized by the set of beliefs about future returns, or, equivalently, by the state variable ϕ_t.

Hedging Changes in Ambiguity

The intertemporal hedging motive implies that (i) agents may follow contrarian, as opposed to momentum, strategies, and (ii) participation is more likely earlier in the investment period. The effects of hedging are most pronounced if there is a large amount of prior uncertainty. For this reason, Figure 3 illustrates them by focusing on an agent in 1928:Q3, who has only one year of previous data as prior information. The left hand panel shows a representative 'section' of the optimal policy surface (for $t = 12$). 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. The non-participation region is reached if the absolute value of the sample equity premium is low $\left(\phi_t = \frac{H}{t} \approx \frac{1}{2}\right)$. If the equity premium has been either very high or very low, the agent is in a **momentum region**. He is long in stocks if the sample equity premium is positive and short otherwise. He also reacts to high (low) returns by increasing (decreasing) his net position. If the absolute value of the equity premium is in an intermediate range, the agent is in a **contrarian region**. He is **short** in stocks for a positive sample equity premium and long otherwise. Moreover, he now reacts to high (low) returns by decreasing (increasing) his net exposure. The contrarian region is also present in our leading example (investment starting in 1971:Q3), but is very small; this is why it was not discernible in Figure 2.

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

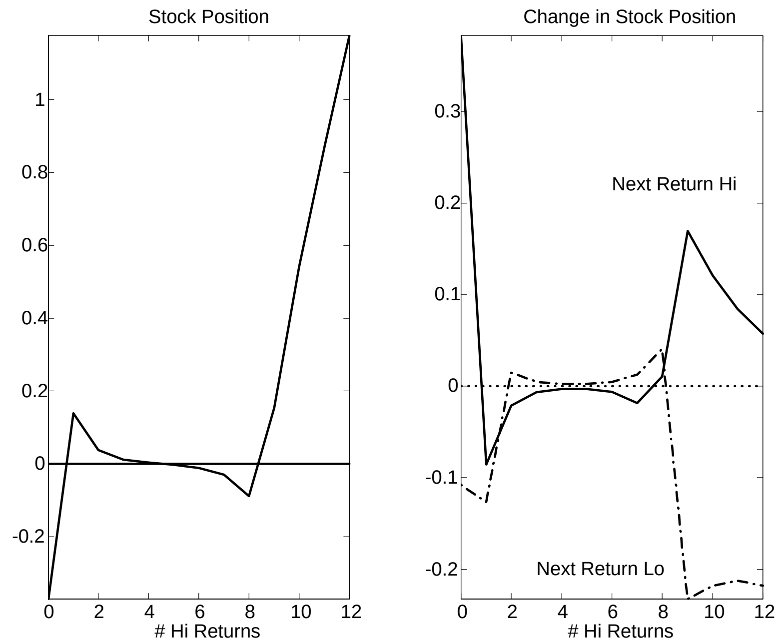

Figure 3: Stock position and responses to arrival of a high or low return after 12 quarters in 30-year planning problem beginning in 1928:Q3. Ambiguity averse investor with $\alpha = .14$ and $\overline{\lambda} = 0$.

balance two reasons for investing in the stock market. One the one hand, he can exploit the expected equity premium by going long. On the other hand, he can insure himself against bad news about investment opportunities (low returns) by going short. Which effect is more important depends on how the size of the equity premium compares to the slope of the h_t's. If the equity premium is small in absolute value, the hedging effect dominates.

Participation since 1971

Figure 4 compares the positions that various investors would have chosen over the last three decades. An investor who believes that returns are unambiguous signals ($\overline{\lambda} = 0$) should always holds stocks, although his positions are quite small, barely reaching 30% even after the high returns of the 1990s. As a reference point, a 'rational expectations' agent who is sure that the mean return is equal to the sample mean of 12% p.a. would hold 82% stocks every period. The Bayesian learner in Figure 4 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 $\overline{\lambda} = .01$ already holds essentially no stocks throughout most of the 1970s. An investor with $\overline{\lambda} = .02$ does not go long in stocks until 1989. Both of these investors participate in the market in the 1970s, but spend most of the time in their contrarian region, where they take tiny short positions. As long as the investor remains in a region where he is long in stocks, changes in the

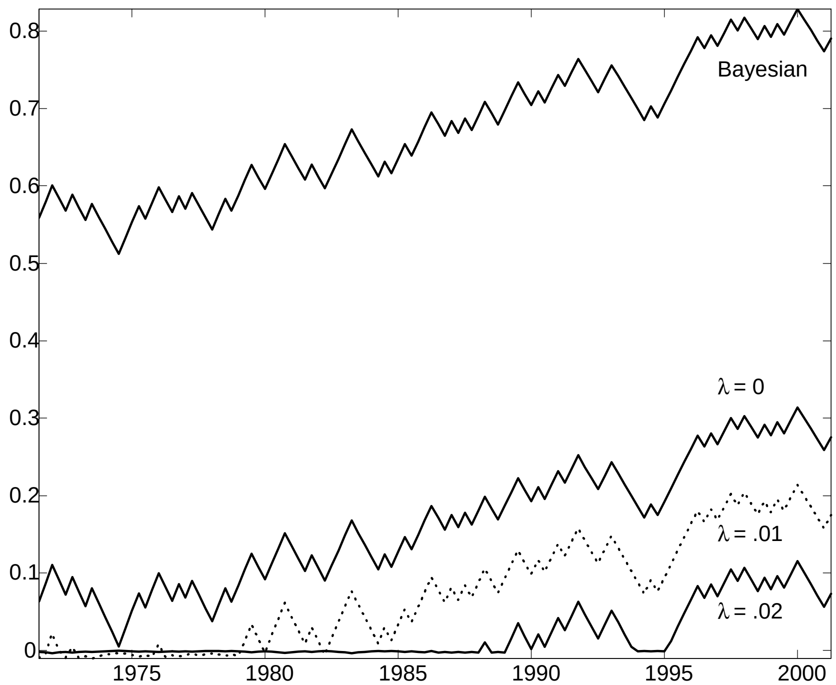

Figure 4: Optimal stock positions for ambiguity averse agents with $\alpha = .14$ and different values of $\overline{\lambda}$ (here denoted λ), as well as for a Bayesian agent with uniform prior. All planning problems are over 30 years beginning in 1971:Q3.

ambiguity parameters α and $\overline{\lambda}$ for a given sample tend to affect the level of holdings, with a negligible effect on changes. Comparison of our benchmark, our cases $\overline{\lambda} = .01$ and $\overline{\lambda} = .02$, and the Bayesian example reveals that the supporting measure's means are essentially vertically shifted versions of each other. Of course, the Bayesian model cannot generate non-participation, so changes will look very different in states where the ambiguity averse investor moves in and out of the market.

Overall, our results suggest that learning under ambiguity could be used as a building block of a successful model of the cross section of holdings. While more work is required to distinguish an ambiguity aversion story from a Bayesian model augmented with a technological participation cost, it is already clear that the two models have 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.

6 AMBIGUOUS NEWS AND ASSET PRICES

This section studies asset pricing in periods of ambiguous news. The first subsection below contains a stylized example to highlight the main effects discussed in the introduction: the asymmetric response to news, the discount due to future ambiguous news, and the link between information quality and volatility. These features are common to the aftermath of both systematic and idiosyncratic shocks. Examples of shocks that flag a period of ambiguous news include company announcements as well as unanticipated shocks to economic growth. In the second subsection, we calibrate a representative agent model of asset pricing with ambiguous news to the period after September 11. This exercise suggests that a model with ambiguous signals can help to account quantitatively for observed price movements.

6.1 A Simple Model of Ambiguous News

There are three periods. A risk neutral representative agent, who does not discount the future, can hold shares of a single asset that is available in unit supply. The asset pays θ units of consumption in period 3, where $\theta = 1$ with probability m and $\theta = 0$ with probability $1 - m$. In period 2, agents receive an ambiguous signal $s \sim \mathcal{N}\left(\theta, \lambda^{-1}\right)$ with $\lambda \in \left[\underline{\lambda}, \overline{\lambda}\right]$. An equilibrium is a price p_1 at which the asset trades in period 1 together with a price function $p_2(s)$ for period 2.

The agent updates beliefs in period 2 to arrive at a set of posterior probabilities for $\theta = 1$:

$$\mathcal{M}_1(s) = \left\{ \frac{m\phi\left(s; 1, \lambda^{-1}\right)}{m\phi\left(s; 1, \lambda^{-1}\right) + (1-m)\,\phi\left(s; 0, \lambda^{-1}\right)}; \quad \lambda \in \left[\underline{\lambda}, \overline{\lambda}\right] \right\}$$

where $\phi\left(s; \mu, \lambda^{-1}\right)$ is the normal density with mean μ and precision λ evaluated at s. It is helpful to rewrite the posterior probability given precision λ as the posterior mean

$$E\left[\theta | s, \lambda\right] = \left(1 + \frac{1-m}{m} \exp\left(\lambda\left(\frac{1}{2} - s\right)\right)\right)^{-1}. \tag{18}$$

The posterior mean is increasing in m and in s, the latter because high s is 'good news' about θ. It increases with the precision λ if and only if $s > \frac{1}{2}$: more precise good news increase the posterior mean, while more precise bad news decrease it.

Asymmetric Response and Price Discount

The worst state for the agent occurs if $\theta = 0$. Since he holds the asset in equilibrium, he evaluates the future in period 2 using the lowest element of $\mathcal{M}_1(s)$. He is willing to hold the asset at the price

$$\begin{aligned} p_2(s) &= \min_{\lambda \in \left[\underline{\lambda}, \overline{\lambda}\right]} E\left[\theta | s, \lambda\right] \\ &= \begin{cases} E\left[\theta | s, \underline{\lambda}\right] & \text{if } s \geq \frac{1}{2} \\ E\left[\theta | s, \overline{\lambda}\right] & \text{if } s < \frac{1}{2}. \end{cases} \end{aligned}$$

Ex post, the agent interprets bad news $\left(s < \frac{1}{2}\right)$ as very informative, whereas good news are viewed as imprecise. The price function is plotted in the left hand panel of Figure 5.

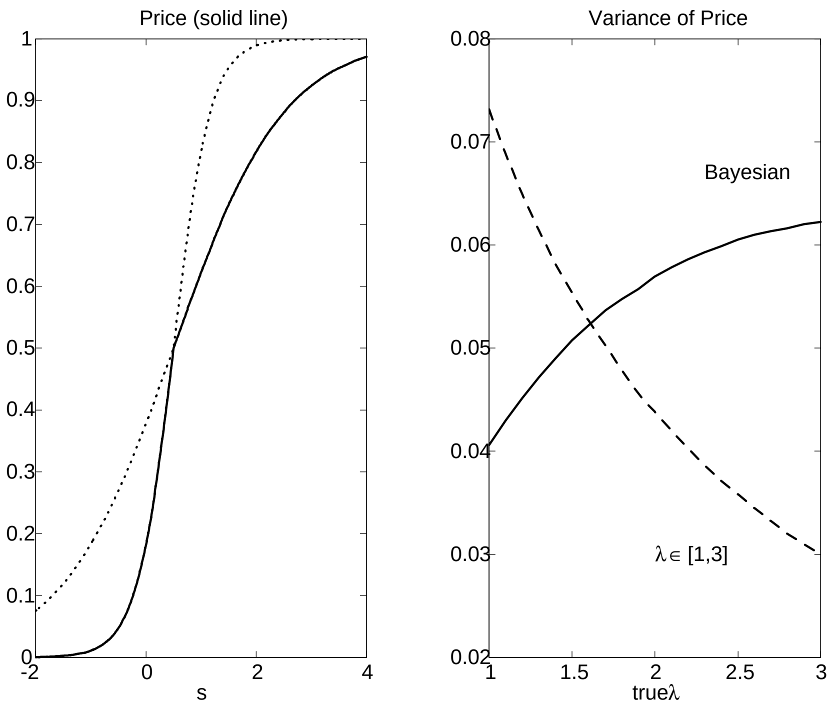

Figure 5: Price function $p_2(s)$ and variance of price p_2 as a function of true precision λ for ambiguous news model.

Its key feature is the kink at $\frac{1}{2}$. Ambiguity tends to make the distribution of price more negatively skewed than the distribution of signals.

In period 1, the price depends on the agent's ex ante perception of the signal precision. This minimizes the value of the asset in period 2. Thus,

$$\begin{aligned} p_1 &= \min_{\tilde{\lambda}\in\left[\underline{\lambda},\overline{\lambda}\right]} E\left[p_2\left(s\right)\right] \\ &= m \min_{\tilde{\lambda}\in\left[\underline{\lambda},\overline{\lambda}\right]} E\left[\frac{p_2\left(s\right)}{E\left[p_3|s,\tilde{\lambda}\right]}|s=1,\tilde{\lambda}\right] \\ &\leq m. \end{aligned} \tag{19}$$

If the signal is noisy, then the law of iterated expectations holds and the initial price is equal to m. If $\overline{\lambda} > \underline{\lambda}$, the price in the ambiguous signals model is strictly below the posterior mean. The prospect of ambiguous news induces a discount, as the agent wants to be compensated for the ambiguity to be endured.

Connection between Truth and Beliefs

To evaluate the equilibrium price behavior, one needs to specify the **true** distribution of the signal. We assume that the true precision λ is in the interval $\left[\underline{\lambda},\overline{\lambda}\right]$. The signal is thus $\mathcal{N}\left(\theta,\lambda^{-1}\right)$. For the Bayesian model ($\underline{\lambda}=\overline{\lambda}$), the truth is simply the likelihood at the true parameter. This is an intuitive restriction: the agent is not wrong about the meaning and precision of the signal. Similarly, in the ambiguous case, while the agent is uncertain about the precision, he is not entirely off target.

Information Quality and Volatility

The right hand panel of Figure 5 examines how price volatility, measured by the variance, changes with different measures of information quality. In the Bayesian case, precision itself is a natural measure. The solid line shows that price volatility increases monotonically with signal precision. Although the signal becomes less volatile as λ increases, it is also given less weight in the conditional expectation; on net, prices move around less. The dashed line presents the volatility that arises at the true λ if the agent is ambiguity averse with $\left[\underline{\lambda}, \overline{\lambda}\right] = [1, 3]$. For every true λ, the move from the dotted (Bayesian) line to the dashed line may thus be viewed as a comparative static decrease in information quality.

It is apparent that if a decrease in information quality arises through an increase in ambiguity, volatility can go up or down. In particular, volatility increases if the signal itself is relatively volatile. In this case, agents take bad news much more seriously than do Bayesians, while they treat good news similarly. The strong response to bad news increases volatility overall. In contrast, if the signal is relatively precise, then bad news are treated similarly by the Bayesian and the ambiguity averse agent, whereas the latter responds much less to good news. This decreases volatility overall.

Matching Prices after an Unanticipated Shock

After a shock that increases uncertainty, prices often fall not only on impact, but continue to slide for some time. For example, this was the case after 9/11, considered in more detail in the next subsection. We now examine when this is consistent with no long term change in fundamentals. Consider a price path with $p_1^* > p_2^*$. One can imagine that, before the shock, the price is $p_0^* = 1$. When the shock occurs, the price drops to p_1^*. A further drop to p_2^* is followed by a recovery to $p_3^* = \theta = 1$.

It is straightforward to back out the signal values (s^{bay} and s^{amb}, say) and the priors (m^{bay} and m^{amb}) required to match the above price path. For the Bayesian model, $m^{bay} = p_1^*$ and s^{bay} satisfies

$$E\left[\theta | s^{bay}, \lambda^*\right] = p_2^*.$$

In contrast, with an ambiguous signal, the prior and signal satisfy

$$\begin{aligned} \min_{\tilde{\lambda} \in \left[\underline{\lambda}, \overline{\lambda}\right]} E\left[p_2(s)\right] &= p_1^* \qquad \text{and} \\ E\left[\theta | s^{amb}, \overline{\lambda}\right] &= p_2^*. \end{aligned}$$

From (19), it follows that $m^{amb} > m^{bay}$. The ambiguous signal induces a discount for any given prior, so the prior that rationalizes a given price is higher in the ambiguous signal case. The two models thus provide different stories to explain the price path.

Under the Bayesian model, there was simply a large drop in the expected value of the asset. In contrast, in the ambiguous signal model a smaller revision of the expected value is amplified by the prospect of ambiguous news.

Which model is more plausible? This can be assessed by comparing the likelihood of the signal realizations s^{bay} and s^{amb} under the true density. Since $p_2^* < p_1^* = m^{bay} < m^{amb}$, (18) implies that both signal realizations must be smaller than $\frac{1}{2}$. Under the truth, which is normal with mean 1, the likelihood of the ambiguous model is higher if and only if $s^{amb} > s^{bay}$. This is true as long as p_2 is sufficiently small. Intuitively, a successful model here is one that can explain why prices fell so much without invoking a very low (and thus unlikely) signal. Since the ambiguity averse agent reacts strongly to bad news, the signal required to make him discount the asset to a rather low price need not be as extreme as that required for the Bayesian agent.

While this example clarifies the role of ambiguous signals in explaining prices after an increase in uncertainty, it has a number of drawbacks. First, the uncertainty is directly about changes in the asset value, rather than in economic growth. Second, it assumes risk neutrality. Finally, the recovery occurs by assumption within one period, as opposed to gradually over many periods, as it did after September 11. These assumptions are relaxed in the next subsection.

6.2 Asset Pricing after 9/11

This subsection constructs and calibrates a representative agent model to explain the movements in the S&P 500 index in September and early October of 2001. There is an infinitely-lived representative agent. A single Lucas tree yields dividends $Y_t = \exp\left(\sum_{j=1}^{t} \Delta y_j\right) Y_0$, with Y_0 given. According to the true data generating process, the growth rate of dividends is $\Delta y_t \sim i.i.\mathcal{N}\left(\theta^{hi}, \sigma^2\right)$ for all t. The agent knows that the mean growth rate is θ^{hi} from time 0 up to some given time $T+1$. However, he believes that with probability $1-\mu$, the mean growth rate drops permanently to θ^{lo} after $T+1$. Information about growth beyond $T+1$ is provided, at each date $t \leq T$, by a signal s_t that takes the values 1 or 0. Signals are serially independent and also independent of dividends before $T+1$; they satisfy $\Pr(s_t = 1) = \pi$. At time $T+1$, the long run mean growth rate is revealed.

The information structure captures the following scenario. First, there was no actual permanent structural change caused by the attack.[19] Second, agents were initially unsure if there would be such a change. Third, news reports were initially much more informative about the possibility of structural change than were dividend or consumption data. Of course, to the extent that dividend data were available, they may have provided some information. But initially, they are likely to have largely reflected decisions taken before

[19] The model can accommodate drops in dividends in September and stock price movements that reflect these drops. All that is required is that such movements come from the same distribution as movements before September 11.

the attack occurred, becoming more informative only with time. Our model captures this shift in relative informativeness in a stark way. We divide the time after the attack into two phases, a learning phase ($t \leq T$) where dividends are entirely uninformative about structural change, and a "new steady state" phase ($t > T$) where structural change actually materializes in dividends. In our calibration below, T corresponds to 26 days.[20] Finally, imposing a fixed T at which θ is revealed is not a strong restriction if beliefs are already close to the true θ at time T. We show a plot of our posterior means below.

The agent believes that signals are informative about growth, but views them as ambiguous. This feature is modeled via a set of likelihoods ℓ, where

$$\ell\left(s_t = 1|\theta^{hi}\right) = \ell\left(s_t = 0|\theta^{lo}\right) = \lambda \in [\underline{\lambda}, \overline{\lambda}], \tag{20}$$

with $\underline{\lambda} > \frac{1}{2}$. Beliefs about signals up to time T have a statistical representation $(\Theta, \mathcal{M}_0, \mathcal{L})$, where $\Theta = \left\{\theta^{hi}, \theta^{lo}\right\}$, $\mathcal{M}_0$ contains the single prior given by μ, and $\mathcal{L}$ is defined by (20). The special case $\underline{\lambda} = \overline{\lambda}$ is a Bayesian model. To ease notation, assume that signals continue to arrive after T, but that for $t > T$, $\ell\left(s_t = 1|\theta^{hi}\right) = 1 = \ell\left(s_t = 0|\theta^{lo}\right)$.

In terms of the notation of earlier sections, the state space is $S = \{0, 1\} \times \mathbb{R}$. Since Y is independent of s, the 1-step-ahead beliefs $\mathcal{P}_t\left(s_1^t, Y_1^t\right)$ for $t \leq T$ are given by the appropriate product of 1-step-ahead beliefs about s_{t+1} and the conditional probability law for Y_{t+1}. Preferences over consumption streams are then defined recursively by

$$V_t\left(c; s_1^t, Y_1^t\right) = \min_{p \in \mathcal{P}_t\left(s_1^t, Y_1^t\right)} \left(c_t^{1-\gamma} + \beta E_p\left[\left(V_{t+1}\left(c; s_1^{t+1}, Y_1^{t+1}\right)\right)^{1-\gamma}\right]\right)^{\frac{1}{1-\gamma}}, \tag{21}$$

where β and γ are the discount factor and the coefficient of relative risk aversion, respectively. Since only the signals are ambiguous, the minimization in (21) may be viewed as a choice over sequences $\left(\lambda_1^{t+1}\right)$ of precisions.

Connection between Truth and Beliefs

Discipline on beliefs is imposed in two ways. First, as above, assume that the true precision λ lies in $[\underline{\lambda}, \overline{\lambda}]$. This condition ensures that an agent's view of the world is not contradicted by the data. Suppose the agent looks back at the history of signals after he is told the true parameter at time T. If he is Bayesian ($\underline{\lambda} = \overline{\lambda}$), the distribution of the signals at the true parameter value is the same as the true distribution of the signals. In this sense, the agent has interpreted the signals correctly.[21] More generally, an ambiguity averse agent contemplates many 'theories' of how the signal history has been generated, each corresponding to a different sequence of precisions (λ_t). One might thus be concerned that theories that do not satisfy $\lambda_t = \lambda$ infintely often are contradicted

[20] The model could be extended to relax this strict division into phases. One might want to assume that both news reports and dividends are informative about structural change at all times. However, in such a setup, one would still like to let the informativeness of news reports decrease over time relative to that of dividends. It is plausible that the main effects of our setup would carry over to this more general environment.

[21] For example, he has not been "overconfident", interpreting every signal as more precise than it actually was.

by the data. However, this is not the case if $\lambda \in [\underline{\lambda}, \overline{\lambda}]$: there exists a large family of signal processes with **time varying precision** $\lambda_t \in [\underline{\lambda}, \overline{\lambda}]$ that cannot be distinguished from the true distribution on the basis of **any** finite sample.[22] While some of these processes will appear less likely than others in the short run, any of them is compatible with a sample that looks i.i.d. with precision λ. An agent who believes in the whole range $[\underline{\lambda}, \overline{\lambda}]$ need not, with hindsight, feel that he interpreted the signals incorrectly.

The second restriction is that agents would learn the true state θ^{hi} even if it were not revealed at $T+1$. This precludes an excessively pessimistic interpretation of news. A sufficient condition is that the posterior probability of θ^{hi}, $\mu_t\left(s_1^t, \lambda_1^t\right)$ say, converges to 1 if the truth equals the **lower bound** of the precision range:

$$\lim_{t\to\infty} \min_{\lambda_1^t} \ \mu_t\left(s_1^t, \lambda_1^t\right) = 1, \text{ a.s. for } s_t \text{ i.i.d. with } \Pr\left(s_t = 1\right) = \underline{\lambda}. \tag{22}$$

If $\overline{\lambda}$ were too large for given $\underline{\lambda}$, agents could interpret negative signals as very precise and never be convinced that the true state has occurred if the fraction of good signals is $\underline{\lambda}$. Thus the condition bounds $\overline{\lambda}$ for a given $\underline{\lambda}$.

Supporting Measure and Asset Prices

Following [20], equilibrium asset prices can be read off standard Euler equations once a (1-step-ahead) "supporting measure" that achieves the minimum in (21) has been determined. Suppose that the intertemporal elasticity of substitution is greater than one. It is then easy to show that continuation utility is always higher after good news ($s = 1$) than after bad news ($s = 0$). Thus the sequence of precisions $\left(\lambda_1^{*t+1}\right)$ that determines the supporting measure at time t and history s_1^t is chosen to **minimize** the probability of a high signal in $t+1$. For the past signals s_1^t, this requires maximizing the precision of bad news ($\lambda_j^* = \overline{\lambda}$ if $s_j = 0$) and minimizing the precision of good news ($\lambda_j^* = \underline{\lambda}$ if $s_j = 1$). For the future signal s_{t+1}, it requires maximizing (minimizing) the precision λ_{t+1}^* whenever news are more likely to be bad (good) next period, that is, whenever the posterior probability of θ^{hi} is smaller (larger) than $\frac{1}{2}$.

Let P_t denote the price of the Lucas tree. Since signals and dividends are independent for $t \leq T$, the price-dividend ratio $v_t = P_t/Y_t$ depends only on the sequence of signals. It satisfies the difference equation

$$v_t\left(s_1^t\right) = \hat{\beta} E_t^* \left[\left(1 + v_{t+1}\left(s_1^{t+1}\right)\right)\right]$$

where E_t^* denotes expectation with respect to the (1-step-ahead) supporting measure and where the new discount factor $\hat{\beta} = \beta e^{\left(1-\frac{1}{\sigma}\right)\theta^{hi} - \frac{1}{2}(1-\gamma)\left(1-\frac{1}{\sigma}\right)\sigma^2}$ is adjusted for dividend risk. Once θ has been revealed in period $T+1$, the price dividend ratio settles at a constant value.

[22] To construct such precision sequences, pick any ω such that $\omega\underline{\lambda} + (1-\omega)\overline{\lambda} = \lambda$. Let $\tilde{\lambda}_t$ be an i.i.d. process valued in $\{\underline{\lambda}, \overline{\lambda}\}$ with $\Pr\left(\tilde{\lambda}_t = \underline{\lambda}\right) = \omega$. For almost every realization (λ_t) of $\left(\tilde{\lambda}_t\right)$, the empirical distribution of the nonstationary signal process with precision sequence (λ_t) converges to the true distribution of the signals. See Nielsen [38] for a formal proof.

Calibration

To illustrate different models of learning, we calibrate the model to the aftermath of September 11, 2001. The stock market was closed in the week after the attack; the first trading day was Monday, September 17. Figure 6 plots the price-dividend ratio for the S&P 500 index for 19 trading days, starting 9/17, including the pre-attack value (9/10) as day 0. At the end of our window (Friday, October 5), the market had climbed, for the first time, to the pre-attack level. It subsequently remained between 68 and 73 for another three weeks (not shown). The discount rate is 4% p.a. and the coefficient of relative risk aversion is $\gamma = .5$. The average growth rate of dividends is fixed to match the price-dividend ratio, yielding a number of 5.2% p.a. This is clearly larger than the historical average, which reflects the high p/d ratio. The volatility of consumption is set at the historical value of 2% p.a. reported by Campbell [13] for postwar data. Finally, we assume that the potential permanent shock corresponds to a drop in consumption growth of .5 % p.a. In steady state, this would correspond to a price-dividend ratio of 61.

Having fixed these parameters, we infer, for every learning model, the sequence of signals that must have generated our price-dividend ratio sample if the model is correct. If the signals had a continuous distribution, this map would be exact. Here we assume that agents observe 20 signals per day. We then compute the model-implied price path that best matches the data. While the price distribution is still discrete, it is sufficiently fine to produce sensible results. A model is discarded if its 'pricing errors' are larger than .5 at any point in time. Finally, we compute the likelihood conditional on the first observation for each model, using the distribution of the fitted price paths. This is a useful criterion for comparing models, since the first observation is basically explained by the choice of the prior.

Numerical Results

To select a Bayesian model, we search over priors μ and precision parameters λ to maximize the likelihood. This yields an interior solution for both parameters. For example, for precision, the intuition is as follows: the path of posteriors is completely determined by the path of p/d ratios. Thus performance differences across Bayesian models depend on how likely the path of posteriors is under the truth. If precision is very large, then it is highly unlikely that there could have been enough bad news to explain the initial price decrease. In contrast, if precision is very small, then signals are so noisy that posteriors do not move much in response to any given news. Highly unlikely 'clusters' of first bad and then good news would be required to explain the price path. This tradeoff gives rise to an interior solution for precision.

To select a multiple-priors model, we need to specify both the true precision and the range of precisions the agent thinks possible. To sharpen the contrast with Bayesian models, we focus on models where ambiguity is large; we set $\overline{\lambda}$ slightly (.001) below the upper bound associated with the requirement (22). We also assume that the truth corresponds to $\underline{\lambda}$.[23] With these two conventions, we search over $\underline{\lambda}$ to find our favorite

[23] Strictly speaking, this polar case is not permitted by the restriction that the truth lie in the interior

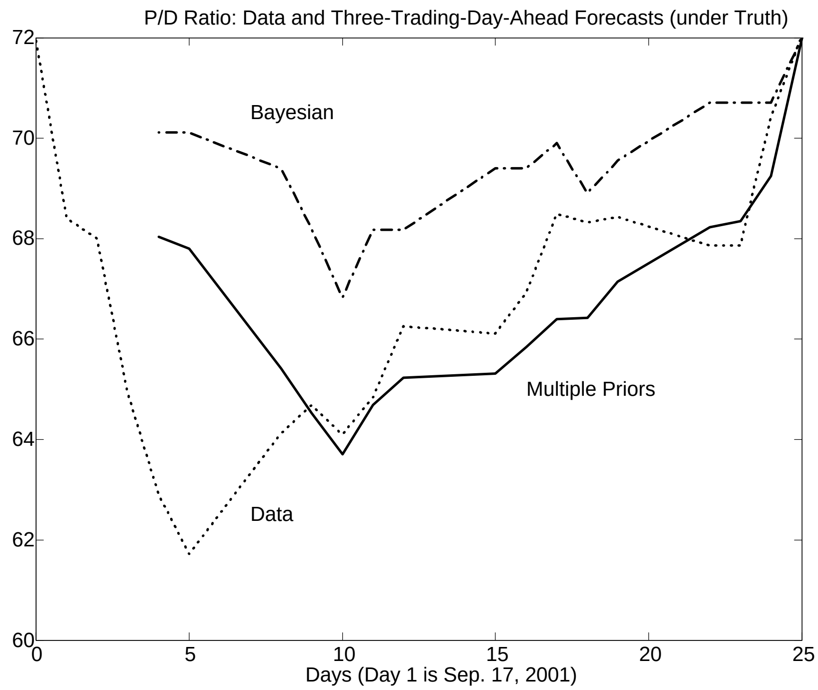

Figure 6: Data and In-Sample Forecasts for 9/11 Calibration.

multiple-priors model. This model is compared to the Bayesian model in Figures 6 and 7. The favorite multiple-priors model begins with a much higher prior probability, and the precision range for $\lambda \in [.58, .608]$ is higher than the precision for the best Bayesian model, $\lambda = .56$. The multiple-priors model (log likelihood $= -33.29$) outperforms the Bayesian model (log likelihood $= -36.82$). Figure 6 plots the 1-step-ahead conditional likelihoods to illustrate the source of the difference. The multiple-priors model is better able to explain the downturn in the week of September 17. The models do about the same during the recovery. Figure 6 plots, together with the data, three-trading-day-ahead in-sample forecasts. This shows that the Bayesian model predicts a much faster recovery than the multiple-priors model throughout the sample.

The result shows how the effects discussed in the previous subsection operate in a setting with many signals. The two models represent two very different accounts of market movements in September 2001. According to the Bayesian story, all price movements reflect changes in beliefs about future growth. In particular, the initial drop in prices arose because market participants expected a permanent drop in consumption of .2% (see Figure 7). During the first week, bad news increased the expected drop to almost .5%. In contrast, the 'ambiguity story' says that agents begin with a prior opinion that basically nothing has changed. However, they know that the next few weeks will be one of increased confusion and uncertainty. Anticipation of this lowers their willingness to pay for stocks. In particular, they know that future bad news will be interpreted (by

of the precision range. However, there is always an admissible model arbitrarily close to the model we compute.

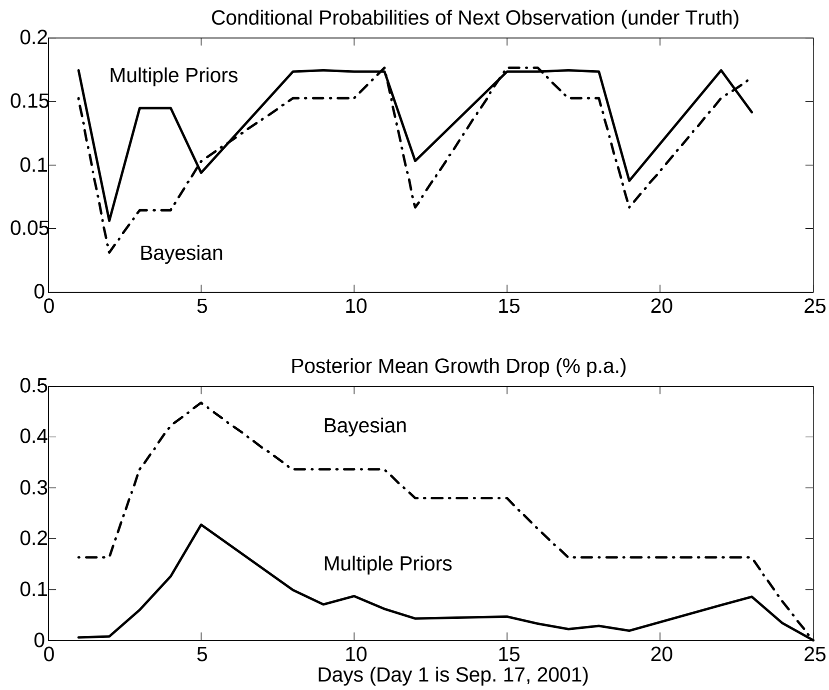

Figure 7: Conditional Probabilities and Posterior Mean for 9/11 Calibration.

future ambiguity averse market participants) as very precise, whereas future good news will be interpreted as noisy. This makes it more likely that the market will drop further in the short run than for the Bayesian model.

For representative agent asset pricing models with multiple-priors utility, there is always an observationally equivalent Bayesian model that yields the same equilibrium price. This begs the questions why one should not consider this Bayesian model directly.[24] Here, the reason is that this Bayesian model cannot be motivated by the same plausible a priori view of the environment as our ambiguity aversion model. We want to capture a scenario where signals are generated by a memoryless mechanism, and where precision does not depend on the state of the world: learning in good times is not expected to occur at a different speed than in bad times. An ambiguity aversion model with these features outperforms a Bayesian model with these features. Some other Bayesian model which does not have these features is not of interest. In addition, such a model would yield misleading comparative static predictions. The observationally equivalent model is much like a 'reduced form' which is not invariant to changes in the environment.

[24] This model would be an expected utility model with pessimistic beliefs, similar to the one in Abel [1].

7 RELATED LITERATURE

We are aware of only two formal treatments of learning under ambiguity. Marinacci [36] studies repeated sampling with replacement from an Ellsberg urn and shows that ambiguity is resolved asymptotically. This is a special case of our model in which signals are unambiguous. The statistical model proposed by Walley [48, pp. 457-72] differs in details from ours, but is in the same spirit; in particular, it also features ambiguous signals. An important difference, however, is that our model is consistent with a coherent axiomatic theory of dynamic choice between consumption processes. Accordingly, it is more readily applicable to economic settings.[25]

Our model proposes a way to model incomplete learning in complicated environments that is quite different from existing Bayesian approaches. One such approach starts from the assumption that the true data generating measure is not absolutely continuous with respect to an agent's belief.[26] This generates situations where beliefs do not converge to the truth even though agents believe, and behave as if, they will.[27] In contrast, agents in our model are aware of the presence of hard-to-describe factors that prevent learning and their actions reflect the residual uncertainty.

Our setup is also different in spirit from models with persistent hidden state variables, such as regime switching models. In these models, learning about the state variable never ceases because agents know that the state variable is forever changing. Agents thus track a *known* data generating process that is not memoryless. In contrast, our model applies to memoryless mechanisms. Accordingly, learning about the fixed true parameter does eventually cease. Nevertheless, because of ambiguity, the agent reaches a state where no further learning is possible although the true parameter is not yet known.

There exist a number of applications of multiple-priors utility or the related robust control model to asset pricing or portfolio choice. None of these is concerned with learning. Multiple-priors applications typically employ a *constant* set of one-step-ahead probabilities (Epstein and Miao [17], Routledge and Zin [39]). Similarly, existing robust control models (Hansen, Sargent, and Tallarini [26], Maenhout [33], Cagetti et al. [12]) do not allow the 'concern for robustness' to change in response to new observations. Neither is learning modeled in Uppal and Wang [44] that pursues a third approach to accommodating ambiguity or robustness.

Our paper contributes to a growing literature on learning and portfolio choice. Bawa, Brown, and Klein [9] and Kandel and Stambaugh [30] first explored the role of parameter uncertainty in a Bayesian framework.[28] Several authors have solved intertemporal

[25] A similar remark applies to Huber [29], who also points to the desirability of admitting ambiguous signals and outlines one proposal for doing so.

[26] This violates the Blackwell-Dubins [10] conditions for convergence of beliefs to the truth. See Feldman [21] for an economic application.

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

[28] There are alternatives to a Bayesian approach to the parameter uncertainty problem. See Ang and Bekaert [2] for a classical econometric strategy.

portfolio choice problems with Bayesian learning.[29] 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.

Non-participation can be derived also from preferences with first-order risk aversion, as in Ang, Bekaert and Liu [3]. The difference between first-order risk aversion and ambiguity aversion models is that, in the latter, the degree of deviation from expected utility behavior depends on the environment. For example, it can vary systematically with new information about the market, as in our model.[30] There is also a 'technological' approach, that derives non-participation from fixed participation costs.[31] Whether technology or preferences are behind observed no-participation is an important question for future research, in particular since the welfare implications of the two explanations are quite different, as discussed in Section 5.

There is a large literature on asset pricing with Bayesian learning. One set of papers argues that learning can explain excess volatility and in-sample predictability of returns. In these applications, the learning process starts at the beginning of the sample and is usually reset periodically, for example due to regime shifts.[32] In contrast, our application focuses on a learning process triggered by an event that increases uncertainty. It is thus closer to a second group of papers that tries to explain post-event abnormal returns ("underreaction") through the gradual incorporation of information into prices.[33] Our setup may be viewed as a model of negative underreaction in periods of ambiguous news. In these periods, underreaction is likely even if there is no change in fundamentals that is gradually revealed. Moreover, the slide in prices is reversed in the long run as agents learn that fundamentals have not changed.

The mechanism that generates negative skewness in our model also differs from existing explanations. Veronesi [46] shows, in a Bayesian model with risk averse agents, that prices respond more to bad news in good times and conversely. This obtains because, in his setup, news that contradict the current belief increase the conditional variance of asset payoffs. Our result differs in two ways. First, since it does not rely on risk aversion, it is relevant also if uncertainty is idiosyncratic and investors are well diversified. Second, ambiguous signals entail an asymmetric response whether or not times are good. They

[29] Detemple [15], Gennotte [22], and Barberis [4] have considered the case of learning about mean returns, while Barberis [4] and Xia [49] have studied learning about predictability. See Ang and Bekaert [2] for portfolio choice in a regime-switching model.

[30] Another difference is that perceived ambiguity can differ across bets on different parts of the state space. Mukerji and Tallon [37] show how this differentiates multiple-priors from other preferences with 'kinked indifference curves'. In applications, this is relevant for modeling participation in some assets markets and non-participation in others.

[31] See, for example, Vissing-Jorgensen [47] or Haliassos and Michealides [25]). Non-participation of young households can also be explained by short-sale constraints, as in Storesletten, Telmer and Yaron [41].

[32] See Timmermann [42, 43], Bossaerts [5], and Lewellen and Shanken [32] for models of nonstationary transitions and Brandt, Xeng, and Zhang [6], Veronesi [46] and Brennan and Xia [11] for models with persistent hidden state variables.

[33] See, Brav and Heaton [7] for an overview and discussion of this literature.

thus induce **unconditional** negative skewness in returns. To explain the latter fact, some authors have proposed mechanisms for bad news to be more concentrated in time.[34] Such mechanisms could reinforce negative skewness in our setting, but they are not necessary for it to obtain.

A APPENDIX

Proof of (16). Write the likelihood of a sample s_1^t under some theory, here identified with a pair $\left(\theta, \lambda_1^t\right)$, as

$$L\left(s_1^t, \theta, \lambda_1^t\right) = \prod_{j=1}^{t} \left(\theta + \lambda_j\right)^{s_j} \left(1 - \theta - \lambda_j\right)^{1-s_j}. \tag{23}$$

Let $\tilde{\lambda}_1^t$ denote the sequence that maximizes (23) for fixed θ. This sequence is independent of θ and has $\tilde{\lambda}_j = \overline{\lambda}$ if $s_j = 1$ and $\tilde{\lambda}_j = -\overline{\lambda}$ if $s_j = 0$, for all $j \leq t$. It follows that $L\left(s_1^t, \theta, \tilde{\lambda}_1^t\right)$ depends on the sample only through the fraction ϕ_t of high returns observed. The set $\mathcal{M}_t^\alpha$ can be expressed in terms of $L\left(s_1^t, \theta, \tilde{\lambda}_1^t\right)$, because $\theta \in \mathcal{M}_t^\alpha$ if and only if the theory $\left(\theta, \tilde{\lambda}_1^t\right)$ passes the likelihood ratio criterion. Indeed, if $\theta \in \mathcal{M}_t^\alpha$, then there exists **some** λ_1^t such that the theory $\left(\theta, \lambda_1^t\right)$ passes the criterion. Thus $\left(\theta, \tilde{\lambda}_1^t\right)$ 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_1^t such that the theory $\left(\theta, \lambda_1^t\right)$ passes the criterion. Finally, one can use

$$g\left(\theta, \phi_t\right) = \frac{1}{t} \log L\left(s_1^t, \theta, \tilde{\lambda}_1^t\right)$$

to express the criterion in (16).

Proof of Theorem 1. (a) Refer to the i.i.d. measure with one-period distribution ϕ as the 'truth'. Every posterior in $\mathcal{M}_t^0\left(s_1^t\right)$ corresponds to some μ_0 and ℓ_1^t :

$$\begin{aligned} \mu_t\left(\theta^* | s_1^t, \mu_0, \ell_1^t\right) &= \frac{\mu_0\left(\theta^*\right) \prod_{j=1}^t \ell_j(s_j|\theta^*)}{\sum_{\theta \in \Theta} \mu_0\left(\theta\right) \prod_{j=1}^t \ell_j(s_j|\theta)} \\ &= \frac{1}{1 + \sum_{\theta \in \Theta} \frac{\mu_0(\theta)}{\mu_0(\theta^*)} \exp\left(t\left(\frac{1}{t} \sum_{j=1}^t \log \frac{\ell_j(s_j|\theta)}{\ell_j(s_j|\theta^*)}\right)\right)} \end{aligned} \tag{24}$$

By (15), then a.s. under the truth and for all $\theta \neq \theta^*$,

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

[34] See, for example, Hong and Stein [28] or Veldkamp [45].

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

where the last equality follows from the law of large numbers because the stochastic process $\max_\ell \log \frac{\ell(s_t|\theta)}{\ell(s_t|\theta^*)}$ is i.i.d. under the truth. It follows that the sum in the denominator of (24) converges to 0 a.s. under the truth and hence that $\mu_t(\theta^*|s_1^t, \mu_0, \ell_1^t) \rightarrow 1$.

(b) For any sequence s_1^∞, denote by ϕ_t the empirical measure on S corresponding to the first t observations. We focus on the set of sequences Ω for which $\phi_t \rightarrow \phi$; this set has measure one under the truth. Fix a sequence $s_1^\infty \in \Omega$. For any likelihoods $(\ell_s)_{s \in S}$ and $\theta \in \Theta$, and for any probability measure λ on S, define

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

Below we take λ to be ϕ_t or ϕ.

Given μ_0 and the likelihood sequence ℓ_1^t, then the data density for the first t periods is

$$\Pr(s_1^t; \mu_0, \ell_1^t) = \sum_{\theta \in \Theta} \mu_0(\theta) \prod_{j=1}^{t} \ell_j(s_j|\theta).$$

In choosing a likelihood sequence that maximizes $\Pr(s_1^t; \mu_0, \ell_1^t)$, it is wlog to focus on sequences such that $\ell_j = \ell_k$ if $s_j = s_k$. Any such likelihood sequence can be identified with a collection $(\ell_s)_{s \in S}$ and we can write

$$\max_{\ell_1^t} \Pr(s_1^t; \mu_0, \ell_1^t) = \max_{(\ell_s)} \sum_{\theta \in \Theta} \mu_0(\theta)\, e^{t\widetilde{H}(\phi_t, (\ell_s), \theta)}.$$

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

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

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

$$max_{(\ell_s)}\, \widetilde{H}(\phi_t, (\ell_s), \theta) \leq max_{(\ell_s)}\, \widetilde{H}(\phi_t, (\ell_s), \theta^*) - \epsilon, \tag{25}$$

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

We claim that

$$\lim_{t \to \infty} \left(\max_{\mu_0, \ell_1^t} \Pr(s_1^t; \mu_0, \ell_1^t) \right)^{\frac{1}{t}} = e^{H(\theta^*)},$$

or equivalently, that

$$\left[\frac{\max_{\mu_0, \ell_1^t} \Pr(s_1^t; \mu_0, \ell_1^t)}{e^{t \max_{(\ell_s)} \widetilde{H}(\phi_t, (\ell_s), \theta^*)}} \right]^{\frac{1}{t}} \longrightarrow 1. \tag{26}$$

Rewrite the latter in the form

$$\left[max_{\mu_0,\ell_1^t} \sum_{\theta\in\Theta} \mu_0\left(\theta\right)\ e^{t\eta_t\left(\theta,\mu_0,\ell_1^t\right)}\right]^{\frac{1}{t}} \longrightarrow 1, \quad \text{where} \tag{27}$$

$$\eta_t\left(\theta,\mu_0,\ell_1^t\right) = \frac{1}{t}\Sigma_{j=1}^t log\ell_j\left(s_j \mid \theta\right) - \max_{(\ell_s)} \widetilde{H}\left(\phi_t,\left(\ell_s\right),\theta^*\right).$$

From (25), deduce that

$$\begin{aligned}
& max_{\mu_0,\ell_1^t}\mu_0\left(\theta^*\right)e^{\eta_t\left(\theta^*,\mu_0,\ell_1^t\right)} \\
& \leq max_{\mu_0,\ell_1^t}\sum_{\theta\in\Theta}\mu_0\left(\theta\right)\ e^{t\eta_t\left(\theta,\mu_0,\ell_1^t\right)} \\
& \leq max_{\mu_0,\ell_1^t}\mu_0\left(\theta^*\right)\ e^{t\eta_t\left(\theta^*,\mu_0,\ell_1^t\right)} + \left(1-\mu_0\left(\theta^*\right)\right)e^{-\epsilon t},
\end{aligned}$$

for all $t > T$. But

$$\left[max_{\mu_0,\ell_1^t}\,\mu_0\left(\theta^*\right)\ e^{t\eta_t\left(\theta^*,\mu_0,\ell_1^t\right)}\right]^{\frac{1}{t}} \longrightarrow 1,$$

which proves (27).

Now consider any admissible 'theory' (μ_0, ℓ_1^t). By the definition of $\mathcal{M}_t^\alpha$, (μ_0, ℓ_1^t) must satisfy

$$\left(\max_{\mu_0,\ell_1^t}\Pr(s_1^t;\mu_0,\ell_1^t)\right)^{\frac{1}{t}} \geq \left(\Pr(s_1^t;\mu_0,\ell_1^t)\right)^{\frac{1}{t}} \geq \alpha^{\frac{1}{t}}\left(\max_{\mu_0,\ell_1^t}\Pr(s_1^t;\mu_0,\ell_1^t)\right)^{\frac{1}{t}}.$$

Thus (26) implies that

$$\left[\frac{\Pr(s_1^t;\mu_0,\ell_1^t)}{e^{t\max_{(\ell_s)}\ \widetilde{H}(\phi_t,(\ell_s),\theta^*)}}\right]^{\frac{1}{t}} \longrightarrow 1. \tag{28}$$

for **any** admissible theory.

From (24), the posterior derived from (μ_0, ℓ_1^t) satisfies

$$\begin{aligned}
\mu_t(\theta^*|s_1^t,\mu_0,\ell_1^t) &= \frac{\mu_0\left(\theta^*\right)e^{t\eta_t(\theta^*)}}{\sum_{\theta\in\Theta}\mu_0\left(\theta\right)\ e^{t\eta_t(\theta)}} \\
&= \mu_0\left(\theta^*\right)\left(\mu_0\left(\theta^*\right) + \sum_{\theta\neq\theta^*}\mu_0\left(\theta\right)e^{t(\eta_t(\theta)-\eta_t(\theta^*))}\right)^{-1};
\end{aligned}$$

here and below we suppress the dependence of η_t on (μ_0, ℓ_1^t) because the latter is fixed. Thus we are done if we can show that

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

This follows from two claims.

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

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

The claim follows from (25).

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

$$\left[\frac{\Pr(s_1^t; \mu_0, \ell_1^t)}{e^{t\max_{(\ell_s)}\ \widetilde{H}(\phi_t,(\ell_s),\theta^*)}}\right]^{\frac{1}{t}} = \left(\sum_{\theta\in\Theta}\mu_0(\theta)\, e^{\eta_t(\theta)}\right)^{\frac{1}{t}} < e^{-\delta} < 1$$

for all sufficiently large t, contradicting (28). ■

References

[1] A. Abel, An exploration of the effects of pessimism and doubt on asset returns, J. Econ. Dynam. Control, 26 (2002), 7-8, 1075-1092.

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

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

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

[5] P. Bossaerts, Learning-induced securities price volatility, Caltech, 2001.

[6] M. Brandt, Q. Zeng and L. Zhang, Equilibrium stock return dynamics under alternative rules of learning about hidden states, Wharton 2001.

[7] A. Brav and J. Heaton, Competing theories of financial anomalies, Rev. Finan. Studies 15 (2002), 575-606.

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

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

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

[11] M. J. Brennan and Y.Xia, Stock price volatility and the equity premium, J. Monet. Econ. 47 (2001), 249-283.

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

[13] J. Campbell, Asset Prices, consumption and the business cycle, in Handbook of Macroeconomics, J. Taylor and M. Woodford, eds., Amsterdam: North Holland, 1999.

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

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

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

[17] L.G. Epstein and J.J. Miao, A two-person dynamic equilibrium under ambiguity, J. Econ. Dyn. and Control, forthcoming.

[18] L.G. Epstein and M. Schneider, Recursive multiple-priors, U. Rochester, August 2001.

[19] L.G. Epstein and M. Schneider, IID: independently and indistinguishably distributed, U. Rochester, December 2001.

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

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

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

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

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

[25] M. Haliassos and A. Michaelides, Portfolio choice and liquidity constraints, Int. Econ. Rev., forthcoming.

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

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

[28] H. Hong and J. Stein, Differences of opinion, short-sales constraints and market crashes, Stanford 2002.

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

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

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

[32] J. Lewellen and J. Shanken, Learning, asset pricing tests and market efficiency, MIT, 2001.

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

[34] 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.

[35] M. Marinacci, Limit laws for non-additive probabilities, and their frequentist interpretation, J. Econ. Theory 84 (1999), 145-195.

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

[37] S. Mukerji and J. Tallon, Ellsberg's 2-color experiment, bid-ask behavior, and ambiguity, Oxford 2001.

[38] C.K. Nielsen, Rational belief structures and rational belief equilibria, Economic Theory 8 (1996) 3, 399-422.

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

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

[41] K. Storesletten, C. Telmer and A. Yaron, Asset pricing with idiosyncratic risk and overlapping generations, Wharton 2001.

[42] A. Timmermann, How learning in financial markets generates excess volatility and predictability in stock returns, Quart. J. Economics 108 (1993), 1135-1145.

[43] A. Timmermann, Excess volatility and predictability of stock returns in autoregressive dividend models with learning, Rev. Ec. Stud. (1996), 523-557.

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

[45] L. Veldkamp, Slow boom, sudden crash, INSEAD, 2001.

[46] P. Veronesi, Stock market overreaction to bad news in good times: a rational expectations equilibrium model, Rev. Finan. Stud. 12 (1999), 976-1007.

[47] A. Vissing-Jorgensen Towards an explanation of household portfolio choice heterogeneity: nonfinancial income and participation cost structures, U. Chicago, 2002.

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

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