



Biased Bayesian Learning with an Application to the Risk-Free Puzzle

Alexander Ludwig and Alexander Zimmer

ERSA working paper 390

November 2013

Economic Research Southern Africa (ERSA) is a research programme funded by the National Treasury of South Africa.

The views expressed are those of the author(s) and do not necessarily represent those of the funder, ERSa or the author's affiliated institution(s). ERSa shall not be liable to any person for inaccurate information or opinions contained herein.

Biased Bayesian learning with an application to the risk-free rate puzzle*

Alexander Ludwig[†] Alexander Zimmer[‡]

November 6, 2013

Abstract

Based on the axiomatic framework of Choquet decision theory, we develop a closed-form model of Bayesian learning with ambiguous beliefs about the mean of a normal distribution. In contrast to rational models of Bayesian learning the resulting Choquet Bayesian estimator results in a long-run bias that reflects the agent's ambiguity attitudes. By calibrating the standard equilibrium conditions of the consumption based asset pricing model we illustrate that our approach contributes towards a resolution of the risk-free rate puzzle. For a plausible parameterization we obtain a risk-free rate in the range of 3.5 – 5%. This is 1 – 2.5% closer to the empirical risk-free rate than according calibrations of the rational expectations model.

Keywords: Ambiguity, Non-additive probability measures; Bayesian learning; Truncated normal distribution; Risk-free rate puzzle

JEL Classification Numbers: C79, D83

*We thank Klaus Adam, Christian Hellwig and Peter Wakker for helpful comments and discussions.

[†]CMR, Center for Macroeconomic Research; University of Cologne; Albertus-Magnus-Platz; 50923 Köln; Email: ludwig@wiso.uni-koeln.de.

[‡]Corresponding author. Department of Economics, University of Pretoria, Private Bag X20, Hatfield 0028, South Africa. E-mail: alexander.zimper@up.ac.za

1 Introduction

Starting with the seminal contribution of Lucas (1978), the consumption based asset pricing model for a representative consumer economy has become the workhorse of the macroeconomic finance literature. As its main virtue the model derives the relationship between the consumer’s uncertainty with respect to future consumption growth and equilibrium asset prices. However, as first demonstrated by Mehra and Prescott (1985), predicted asset returns under the rational expectations hypothesis are, by some large margin, at odds with actually observed asset returns. Several authors—e.g., Cecchetti, Lam, and Mark (2000); Brav and Heaton (2002); Abel (2002); Giordani and Soderlind (2006); Jouini and Napp (2008)—have relaxed the rational expectation hypothesis in order to explain asset-pricing puzzles through “incorrect” subjective beliefs. Any attempt to relax the rational expectations hypothesis in this manner faces the following questions:

Question 1: Why should subjective beliefs significantly differ from objective probabilities?

Question 2: Even if there exist initial differences between subjective beliefs and objective probabilities, why do these differences not vanish through Bayesian learning from large data-samples?

Question 3: Even if any significant difference between subjective beliefs and objective probabilities is persistent, can plausible magnitudes of such differences explain observed asset-pricing puzzles?

We develop a closed-form model of biased Bayesian learning of the mean of a normal distribution which addresses the above questions on a sound decision-theoretic foundation. As our starting point we modify a standard closed-form model of Bayesian learning—according to which a decision maker has a normal distribution as prior over the unknown mean parameter of some normal distribution—by two more realistic assumptions. First, the prior is given as a truncated normal distribution. Second, information is given as interval- rather than point-observations. We then embed this benchmark model of rational Bayesian learning within a model of Choquet Bayesian learning. This allows us to capture an agent’s ambiguity attitudes through neo-additive capacities in the sense of Chateauneuf, Eichberger and Grant (2007). The agent of our model resolves this ambiguity by putting additional decision-weight on extreme (zero versus one) probabilities. As a consequence, the Choquet Bayesian estimator will generically result in an under- or over-estimation compared to the corresponding rational Bayesian estimator.

In the remainder of this introduction we sketch in more detail how our approach addresses Questions 1-3.

Addressing Question 1: Neo-additive capacities

We model subjective beliefs as non-additive probability measures such as neo-additive capacities. These arise as generalizations of subjective additive probability measures in Choquet expected utility (CEU) theory, which relaxes Savage's (1954) sure thing principle in order to accommodate for ambiguity attitudes as elicited in Ellsberg (1961) paradoxes (Schmeidler 1986, 1989; Gilboa 1987). CEU theory is formally equivalent to (cumulative) prospect theory (=PT 1992) (Tversky and Kahneman 1992; Wakker and Tversky 1993) whenever PT 1992 is restricted to gains (for a comprehensive treatment of this equivalence see Wakker 2010). PT 1992, in turn, extends the celebrated concept of original prospect theory by Kahneman and Tversky (1979) to the case of several possible gain values in a way that satisfies first-order stochastic dominance.

By focusing on neo-additive capacities in the sense of Chateauneuf et al. (2007) we model the possible difference between ambiguous subjective beliefs and objective probability measures in a parsimonious way. In addition to information about the sample mean the agent's estimator of the distribution's mean reflects ambiguity. This is expressed through the parameter $\delta \in [0, 1]$. Ambiguity is then resolved through the agent's ambiguity attitudes, resulting either in under- or overestimation. This is expressed through an "optimism" parameter $\lambda \in [0, 1]$. Our model of Choquet Bayesian learning thus considers two additional parameters that are rooted in axiomatic decision-theory. It nests the benchmark model of rational Bayesian learning ($\delta = 0$).

To work with neo-additive capacities is attractive for two reasons. First, neo-additive capacities reduce the potential complexity of non-additive probability measures in a very tractable way. Important empirical features, e.g., inversely S-shaped probability transformation functions are thereby portrayed (Chapter 11 in Wakker 2010).¹ Second, Bayesian learning with respect to neo-additive capacities is closed, i.e., the prior and posterior are conjugate distributions.

Addressing Question 2: Biased Bayesian learning

Standard models of consistent (=rational) Bayesian learning have been applied to several topics in economics and management sciences. For instance, early contributions by

¹To quote Peter Wakker (2010): "[...] the neo-additive functions are among the most promising candidates regarding the optimal tradeoff of parsimony and fit." (p. 209)

Cyert, DeGroot, and Holt (1978) and Tonks (1983) apply a Bayesian learning model—formulated within a normal distribution framework—to a firm’s decision problem to invest in different technological processes. Viscusi (1979) and Viscusi and O’Connor (1984) apply a Bayesian learning model—formulated within a Beta-distribution framework—to the risk-learning behavior of workers in injury-prone industries. Viscusi (1990, 1991) uses the same learning model to address the question of how far smokers underestimate their health risk.

Zimper (2009) and Zimper and Ludwig (2009) have developed Choquet Bayesian learning models as generalizations of Viscusi’s Beta-distribution approach to non-additive beliefs. In contrast, the present paper’s Choquet Bayesian learning model is based on a normal distribution framework similar to Cyert et al. (1978) and Tonks (1983). As one main difference, the rational benchmark model of our approach takes a truncated normal distribution as prior. As a second main difference, we consider updating conditional on imprecise information; that is, we formally define information as interval—rather than as point-observations. These assumptions turn out to be technically and conceptually convenient when we extend the rational benchmark model to our model of Choquet Bayesian learning. They also add more realistic appeal to the rational benchmark model. For example, a representative agent’s realistic prior over the mean of the annual consumption growth rate—which will be relevant to the risk-free rate puzzle discussed in this paper—should arguably assign positive probability mass only to some subset $(-\infty < a, b < \infty)$ and not to any subset of the $(-\infty, \infty)$ interval.

The Choquet Bayesian estimator of our model remains generically bounded away from the sample mean whenever the agent expresses the slightest initial ambiguity about his prior. This contrasts with consistency results for Bayesian estimators for additive probability spaces (Doob 1949). Consequently, our approach provides a sound decision-theoretic answer to Question 2 as to why an initial difference between subjective beliefs and objective probabilities may, actually, increase rather than decrease in the long-run.

Addressing Question 3: The risk-free rate puzzle

Three major asset pricing puzzles have emerged in the macro-economic asset pricing literature: the “equity premium puzzle”, the “risk-free rate puzzle” and the “excess-volatility puzzle” (Kocherlakota 1996; Campbell 1999). As an illustrative application of our theoretical approach we investigate in how far a sensible calibration of our learning model may contribute towards a resolution of the risk-free rate puzzle, originally formulated in Weil (1989). According to the risk-free rate puzzle, a realistically calibrated standard consumption based asset pricing model yields a return on risk-free assets of about 6% compared to a real-world actual risk-free rate of about 2%.

Under the assumption of myopic time-period agents—who ignore any dynamic inconsistencies between themselves and their future selves—we use the biased Choquet estimator for the mean of the consumption growth rate in the equilibrium condition of the standard consumption based asset pricing model. We argue that any resolution of the risk-free rate puzzle through our model requires low values of the optimism parameter in combination with high values of the ambiguity parameter. This qualitative relationship has intuitive appeal: Given that (unrealistically) high degrees of risk-aversion can resolve the original formulation of the risk-free rate puzzle by Mehra and Prescott (1985), it is not surprising that high degrees of ambiguity combined with a pessimistic resolution of this ambiguity can achieve the same result. We then establish that plausible behavioral parameter values result in risk-free rates that are in the range of 3.5–5%. Our approach can thus account for 1–2.5% of the risk-free rate puzzle. A full resolution of the puzzle through biased Bayesian learning of the mean alone appears however implausible.

The remainder of our analysis is structured as follows. Section 2 briefly discusses the related literature on learning models. In Section 3 we develop the rational benchmark case of Bayesian learning in the absence of ambiguity. Section 4 generalizes our rational learning model to a model of Choquet Bayesian learning which is generically biased. In Section 5 we apply our model to illustrate the partial resolution of the risk-free rate puzzle. Finally, Section 6 concludes. All formal proofs are relegated to the Mathematical Appendix. A Supplementary Appendix contains additional material.

2 Related literature

There exists a rich macroeconomic literature on the question in how far agents’ learning behavior will converge towards rational expectations. This literature can be roughly divided in boundedly rational and rational learning models. The boundedly rational approach describes agents’ learning behavior in terms of adaptive learning rules whereby the resulting learning process may or may not converge to the true underlying model (classical contributions include Evans and Honkapohja 2001, 2003; Bullard and Mitra 2002; Evans and McGough 2005; Branch and Evans 2007; for more recent contributions see Assenza and Berardi 2009; Tetlow and von zur Muehlen 2009 as well as references therein). In contrast, the rational approach is synonymous with (standard) Bayesian learning which—in line with Doob’s consistency theorem—implies convergence to the true value whenever the data is drawn from some i.i.d. process (Kasa 1999; Bigio 2010).

Our approach complements these learning models because we describe Bayesian learners whose learning process will remain bounded away from true values in the long

run. The key to the biased learning behavior of our model is thereby not the assumption of bounded rationality in the form of overly simplistic learning rules but rather Bayesian learning which additionally reflects the updating of existing ambiguity attitudes.

Within the rational approach, models with Bayesian updating that establish long-run biases of subjective beliefs are based on the assumption that data is not drawn from an i.i.d. process. The main assumption of Doob's consistency theorem does thereby not apply to these models. For example, Brennan and Xia (2001), Lewellen and Shanken (2002), and Pastor and Veronesi (2003) consider the case in which the mean of an exogenous dividend process may not be constant over time. Consequently, the agent can never fully learn the objective parameters of the underlying distribution because observed frequencies do not admit any conclusion about objective probabilities even in the long run. Along the same line, Weitzman (2007) considers a setup in which the variance of the consumption growth rate is a hidden parameter whereas the mean is known. In contrast to these approaches, the agent of our model observes data drawn from an i.i.d. process. He nevertheless fails to learn the true parameter value because his ambiguity attitudes rather increase than decrease in the long-run.

Marinacci (1999, 2002) studies a non-Bayesian learning environment with non-additive beliefs. He considers a decision maker who observes an experiment such that the outcomes are i.i.d. with respect to the decision maker's non-additive belief. In this setup, Marinacci derives laws of large numbers for non-additive probability measures to the effect that ambiguity does not necessarily vanish in the long-run. While Marinacci's approach can be regarded as a frequentist approach towards non-additive probabilities, our approach is purely Bayesian.

More closely related to our Bayesian learning model is the recursive multiple priors learning model of Epstein and Schneider (2007, 2008) which can be interpreted as a learning model with Bayesian overtones based on max min expected utility (MMEU) preferences of Gilboa and Schmeidler (1989).² There exist two major conceptual differences between the approach of Epstein and Schneider (2007, 2008) and our model. First, Epstein and Schneider go beyond the standard (i.e., based on behavioral axioms) Bayesian framework by introducing the notion of ambiguous information (in the form of multiple likelihoods). This implies that an agent might not express ex-ante ambiguity but rather deals ex-post with ambiguity by permanently receiving ambiguous signals. In

²Recall that the MMEU theory of Gilboa and Schmeidler (1989) formally coincides with CEU theory restricted to convex capacities (e.g., neo-additive capacities for which the optimism parameter equals zero) whenever the set of priors is given as the core of the convex capacity. Also notice that recent robust control applications to asset pricing puzzles where pessimism results from an agents' caution in responding to concerns about model misspecification can be rationalized by Gilboa and Schmeidler's (1989) MMEU theory, cf. Hansen and Sargent (2007) and the literature cited therein.

contrast, our model is axiomatically embedded within the standard Bayesian approach because our decision maker has a unique (neo-additive) prior defined over all possible future events whereby his (neo-additive) posteriors are generated through a Bayesian update rule. Our approach thus treats ambiguity as a psychological feature of the agent which impacts on his interpretation of (possibly imprecise but never ambiguous) information rather than as a property of information itself. Second, the agent of our learning model may express a wider range of ambiguity attitudes because—in contrast to Epstein and Schneider—these ambiguity attitudes are independent of the information that the agent receives.

To comprehend this later difference observe that the framework of Epstein and Schneider (2007, 2008) satisfies the reduction principle (cf. Siniscalchi 2011) according to which dynamic decision situations reduce to static decision situations within the Savage framework. Under the reduction principle the dynamic principles of consequentialism and dynamic consistency together³ imply Savage’s sure thing principle (Sarin and Wakker 1998; Ghirardato 2002). Under the Savage axioms, any dynamic decision theoretic model that universally satisfies consequentialism as well as dynamic consistency within a Bayesian framework is therefore unable to express ambiguity attitudes because it reduces to expected utility theory. Because Epstein and Schneider (2007, 2008) keep consequentialism, dynamic consistency as well as the possible expression of ambiguity attitudes, their model must restrict the domain of both dynamic principles. The Epstein and Schneider (2007, 2008) learning model is based on the recursive multiple priors approach of Epstein and Schneider (2003) which, in turn, goes back to Sarin and Wakker (1998, Theorem 2.1). Sarin and Wakker consider fixed decision trees. They show that updating with respect to the information structure of a given tree satisfies consequentialism and dynamic consistency whenever the set of multiple priors is given as the reduced family of probability measures⁴. Similarly, Epstein and Schneider (2003) fix some information structure and assume that Savage’s sure thing principle holds exactly at all observable events.

However, linking ambiguity attitudes to a given information structure is conceptionally questionable. For example, agents with different information partitions—e.g., a well informed versus a less informed agent—cannot typically have the same ambiguity attitudes within the Epstein and Schneider framework. Furthermore, an agent who will learn in the future every possible event cannot exhibit any ex-ante ambiguity attitudes

³Losely speaking, consequentialism refers to a situation where the agent “never looks back”, i.e., he does not care about how he arrived at a specific decision node and only looks forward to the future consequences of his actions. Dynamic consistency refers to a situation in which, from the perspective of any period, all future plans are consistent with all past plans.

⁴Rectangular priors in the terminology of Epstein and Schneider (2003, 2007).

in the Epstein-Schneider framework because he must be modeled as an expected utility decision maker. In contrast, there do not exist any such information driven restrictions on the ambiguity attitudes expressible in our learning model. A formal discussion of these aspects is given in our supplementary appendix.

3 The benchmark case: Rational Bayesian learning

This section develops a closed-form model of Bayesian learning of the mean of a normal distribution for the rational benchmark case according to which the agent's beliefs do not express any ambiguity. Formally, we consider a probability space $(\pi, \Omega, \mathcal{F})$ where π denotes a subjective additive probability measure such that the agent's prior over possible parameter values is given as a truncated normal distribution. In a first step (Subsection 3.1), we derive the resulting Bayesian estimator under the standard assumption that the agent updates his belief conditional on point-observations, which have zero prior probability in our framework. In a second step (Subsection 3.3), we formally introduce information partitions such that, first, the agent may exhibit a lack of memory and, second, the observed information has positive prior probability. The analytical characterization of the rational Bayesian estimator with respect to positive probability information (Proposition 1) constitutes this section's main result.

3.1 Preliminaries and a lemma

Construction of the state space Ω . For some numbers a, b such that $a < b$ denote by $\Theta = (a, b)$ the parameter-space that collects all parameter-values that the agent regards as possible. Denote by $X^\infty = \times_{i=1}^\infty \mathbb{R}$ the sample-space that collects all possible infinite sequences of observations. The state space of our model is defined as

$$\Omega = \Theta \times X^\infty, \tag{1}$$

with generic element $\omega = (\theta, x_1, x_2, \dots)$.

Construction of relevant σ -algebras. Endow Θ with the Euclidean metric and denote by $\mathcal{B}(\Theta)$ the Borel σ -algebra on Θ , i.e., the smallest σ -algebra containing all open subsets of the Euclidean interval (a, b) . Similarly, denote by $\mathcal{B}(\mathbb{R})$ the Borel σ -algebra on the Euclidean line \mathbb{R} . Denote by $\mathcal{T}(\Theta)$ the trivial σ -algebra in Θ , i.e., $\mathcal{T}(\Theta) = \{\emptyset, \Theta\}$, and denote by $\mathcal{T}(\mathbb{R})$ the trivial σ -algebra in \mathbb{R} , i.e., $\mathcal{T}(\mathbb{R}) = \{\emptyset, \mathbb{R}\}$.

- Define Σ_Θ as the product σ -algebra generated by $\mathcal{B}(\Theta), \mathcal{T}(\mathbb{R}), \mathcal{T}(\mathbb{R}), \dots$; formally written as

$$\Sigma_\Theta = \mathcal{B}(\Theta) \otimes_0 \mathcal{T}(\mathbb{R}) \otimes_1 \mathcal{T}(\mathbb{R}) \otimes_2 \mathcal{T}(\mathbb{R}) \dots$$

- Let $n \geq 1$ and define Σ_n as the following product σ -algebra

$$\Sigma_n = \mathcal{T}(\Theta) \otimes_0 \mathcal{B}(\mathbb{R}) \otimes_1 \dots \otimes_n \mathcal{B}(\mathbb{R}) \otimes_{n+1} \mathcal{T}(\mathbb{R}) \otimes_{n+2} \mathcal{T}(\mathbb{R}) \dots \quad (2)$$

Observe that the sequence of σ -algebras $\Sigma_1, \Sigma_2, \dots$ constitutes a filtration.

- Finally, we define the following σ -algebra

$$\mathcal{F} = \mathcal{B}(\Theta) \otimes_0 \mathcal{B}(\mathbb{R}) \otimes_1 \mathcal{B}(\mathbb{R}) \otimes_2 \mathcal{B}(\mathbb{R}) \dots,$$

which covers any event of interest in our framework; in particular, $\Sigma_\Theta, \Sigma_n \subset \mathcal{F}$.

Construction of the additive probability measure π . Define by $\tilde{\theta} : \Omega \rightarrow \Theta$ such that

$$\tilde{\theta}(\theta, x_1, x_2, \dots) = \theta$$

the Σ_Θ -measurable coordinate random variable that assigns to every state of the world the corresponding true parameter value. By assumption, the true parameter value lies in (a, b) implying that our agent always regards the true parameter value as possible. We assume that the agent's prior over $\tilde{\theta}$ is given as a *truncated normal distribution* with support on the open interval (a, b) such that, for all $A \in \mathcal{B}(\Theta)$,

$$\pi\left(\left\{\omega \in \Omega \mid \tilde{\theta}(\omega) \in A\right\}\right) = \int_{\theta \in A} \psi_{\tilde{\theta}}(\theta) d\theta$$

with

$$\psi_{\tilde{\theta}}(\theta) = \begin{cases} \frac{1}{F(b)-F(a)} \cdot \frac{1}{\sqrt{2\pi\tau^2}} \exp\left[-\frac{1}{2} \frac{(\theta-\mu)^2}{\tau^2}\right] & \text{for } \tau < \infty \\ \frac{1}{b-a} & \text{for } \tau = \infty, \end{cases} \quad (3)$$

for $\theta \in (a, b)$. $F(\cdot)$ denotes here the cumulative distribution function of the corresponding untruncated normal distribution with mean μ and variance τ^2 , i.e.,

$$F(z) = \int_{-\infty}^z \frac{1}{\sqrt{2\pi\tau^2}} \exp\left[-\frac{1}{2} \frac{(\theta-\mu)^2}{\tau^2}\right] d\theta.$$

Define the Σ_k -measurable coordinate random variable $X_k : \Omega \rightarrow \mathbb{R}$, with $k = 1, 2, \dots$, such that

$$X_k(\theta, x_1, x_2, \dots) = x_k.$$

We assume that, conditional on the parameter-value $\theta \in \Theta$, the X_k are independent and identically normally distributed with mean θ and variance σ^2 whereby σ^2 is known to the agent. That is, for all $\theta \in (a, b)$ and all $A \in \mathcal{B}(\mathbb{R})$,

$$\pi\left(\left\{\omega \in \Omega \mid \tilde{\theta}(\omega) = \theta, X_k(\omega) \in A\right\}\right) = \int_{x \in A} \psi_{X_k}(x \mid \theta) \cdot \psi_{\tilde{\theta}}(\theta) dx$$

with

$$\psi_{X_k}(x | \theta) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp \left[-\frac{1}{2} \frac{(x - \theta)^2}{\sigma^2} \right] \text{ for } x \in \mathbb{R}.$$

Lemma. *Consider some point-observation*

$$\{\omega | X_1(\omega) = x_1, \dots, X_n(\omega) = x_n\} \quad (4)$$

with corresponding sample mean

$$\bar{x}_n = \frac{1}{n} \sum_{k=1}^n x_k.$$

Under the above distributional assumptions the expected value of $\tilde{\theta}$ conditional on (4) is given as

$$E \left[\tilde{\theta}, \pi(\cdot | x_1, \dots, x_n) \right] = \mu_n - \rho_n \frac{f_n(b) - f_n(a)}{F_n(b) - F_n(a)} \quad (5)$$

such that

$$f_n(x) = \frac{1}{\sqrt{2\pi\rho_n^2}} \exp \left[-\frac{1}{2} \frac{(x - \mu_n)^2}{\rho_n^2} \right] \text{ for } x \in \mathbb{R} \quad (6)$$

with corresponding cumulative distribution function

$$F_n(z) = \int_{-\infty}^z f_n(x) dx \text{ for } z \in \mathbb{R} \quad (7)$$

where

$$\mu_n = \frac{\sigma^2/n}{\sigma^2/n + \tau^2} \cdot \mu + \frac{\tau^2}{\sigma^2/n + \tau^2} \cdot \bar{x}_n \quad (8)$$

and

$$\rho_n^2 = \frac{\tau^2 \sigma^2/n}{\tau^2 + \sigma^2/n}. \quad (9)$$

3.2 A consistency result

A sequence of posterior probability measures $\{\pi(\cdot | X_1, \dots, X_n)\}_{n \in \mathbb{N}}$ on Σ_Θ is called *consistent* at parameter-value $\theta_0 \in \Theta$ iff it weakly converges with $\pi(\cdot | \tilde{\theta}(\omega) = \theta_0)$ probability one to the degenerate probability measure d_{θ_0} on Σ_Θ such that, for all $A \in \Sigma_\Theta$,

$$d_{\theta_0}(A) = \begin{cases} 1 & \text{if } (\theta_0, \dots) \in A \\ 0 & \text{else;} \end{cases}$$

formally written as

$$\{\pi(\cdot | X_1, \dots, X_n)\} \rightarrow_w d_{\theta_0} \text{ a.s. } \pi(\cdot | \theta_0).$$

According to Doob's (1949) consistency theorem, the sequence $\{\pi(\cdot | X_1, \dots, X_n)\}_{n \in \mathbb{N}}$ is consistent at all parameter-values in some set $\Theta' \subseteq \Theta$ such that

$$\pi\left(\left\{\omega \in \Omega \mid \tilde{\theta}(\omega) \in \Theta'\right\}\right) = 1.$$

Applied to (5), Doob's consistency theorem therefore immediately establishes that the Bayesian estimator converges with probability one to **almost all** true parameter values $\mu^* \in (a, b)$ since, by weak convergence,

$$\begin{aligned} \lim_{n \rightarrow \infty} \int_{\Omega} \tilde{\theta}(\omega) d\pi(\cdot | X_1, \dots, X_n) &= \int_{\Omega} \tilde{\theta}(\omega) d\delta_{\mu^*} \text{ a.s. } \pi(\cdot | \mu^*) \\ &\Leftrightarrow \\ \lim_{n \rightarrow \infty} E\left[\tilde{\theta}, \pi(\cdot | X_1, \dots, X_n)\right] &= \mu^* \text{ a.s. } \pi(\cdot | \mu^*). \end{aligned}$$

The following corollary to the Lemma establishes the even stronger result that the Bayesian estimator converges with probability one to **all** true parameter values $\mu^* \in (a, b)$.

Corollary 1. *For all $\mu^* \in (a, b)$,*

$$\left\{E\left[\tilde{\theta}, \pi(\cdot | X_1, \dots, X_n)\right]\right\} \rightarrow \mu^* \text{ a.s. } \pi(\cdot | \mu^*).$$

3.3 The rational Bayesian estimator

According to our assumption of independent and identically normally distributed X_k , $k = 1, 2, \dots$, any point-observation $(X_1 = x_1, \dots, X_n = x_n)$ is a zero probability event so that the Bayesian estimator $E\left[\tilde{\theta}, \pi(\cdot | X_1, \dots, X_n)\right]$, as derived in the Lemma, is conditioned on zero-probability information only. In our opinion, however, a meaningful interpretation of Bayesian learning in the face of zero-probability information is highly problematic.

In this subsection we therefore introduce our concept of a rational Bayesian estimator which is defined for positive probability information. We thereby recover (5) as an approximation for this Bayesian estimator whenever the agent receives information with probability arbitrarily close to zero. We start with the construction of the period t agent's information structure whereby we allow for the possibilities of restricted memory or memory loss.

Construction of an information-partition and the corresponding σ -algebra. Fix some pair (t, n) such that $t \geq 1$ and $t - n \geq 0$. Consider a sequence of partitions Π_{t-n}, \dots, Π_t of the real-line such that $\Pi_i \subset \mathcal{B}(\mathbb{R})$, $i = t - n, \dots, t$. Denote by \mathcal{I}_t the period t information partition of Ω whereby \mathcal{I}_t is formally defined as the following collection of information cells

$$I_t = \{\Theta \times_0 \mathbb{R} \times_1 \dots \times_{t-n-1} \mathbb{R} \times_{t-n} A_{t-n} \times_{t-n+1} \dots \times_t A_t \times_{t+1} \mathbb{R} \times_{t+2} \dots \mid A_k \in \Pi_k, k = t - n, \dots, t\}.$$

- Given some period t information partition \mathcal{I}_t , define by $\Sigma(\mathcal{I}_t)$ the σ -algebra generated by \mathcal{I}_t . That is, $\Sigma(\mathcal{I}_t)$ is the smallest collection of subsets of Ω that is a σ -algebra and that contains all information cells in \mathcal{I}_t . Clearly, $\Sigma_n \subset \Sigma(\mathcal{I}_t)$ with Σ_n defined by (2).

The interpretation of \mathcal{I}_t is as follows: If $\omega \in \Omega$ is the true state of the world, the period t agent observes for the past n periods the information A_k , $k = t - n, \dots, t$, such that

$$\omega \in \Theta \times_0 \mathbb{R} \times_1 \dots \times_{t-n-1} \mathbb{R} \times_{t-n} A_{t-n} \times_{t-n+1} \dots \times_t A_t \times_{t+1} \mathbb{R} \times_{t+2} \dots$$

That is, whenever $t - n \neq 0$, the period t agent cannot observe any non-trivial information from the initial $1, \dots, t - n - 1$, periods. Consequently, the sequence of σ -algebras $\Sigma(\mathcal{I}_1), \Sigma(\mathcal{I}_2), \dots$ does not necessarily constitute a filtration thereby allowing for the possibility of agents who ignore previously available information.

The following proposition derives our rational benchmark Bayesian estimator conditional on “well-behaved” information, i.e., conditional on information cells with positive probability. More specifically, we consider information cells $I_t \in \mathcal{I}_t$ such that $\pi(A_k) > 0$ for all $k = t - n, \dots, t$, implying $\pi(I_t) > 0$ under the independence assumption. Define the set of possible values of the average \bar{X}_n that corresponds to information cell

$$I_t = \Theta \times_0 \mathbb{R} \times_1 \dots \times_{t-n-1} \mathbb{R} \times_{t-n} A_{t-n} \times_{t-n+1} \dots \times_t A_t \times_{t+1} \mathbb{R} \times_{t+2} \dots \quad (10)$$

as

$$\bar{I}_n = \left\{ \frac{1}{n} (x_{t-n} + \dots + x_t) \mid x_{t-n} \in A_{t-n}, \dots, x_t \in A_t \right\}.$$

Proposition 1. *Consider some information-cell $I_t \in \mathcal{I}_t$ such that $\pi(I_t) > 0$. The period t agent’s rational Bayesian estimator of the mean of X_k , $k \in \{1, 2, \dots\}$ conditional*

on I_t is given as

$$E \left[\tilde{\theta}, \pi(\cdot | I_t) \right] = \int_{y \in \bar{I}_n} E \left[\tilde{\theta}, \psi_{\tilde{\theta}}(\theta | y) \right] \cdot \frac{\psi_{\bar{x}_n}(y)}{\int_{y \in \bar{I}_n} \psi_{\bar{x}_n}(y) dy} dy \quad (11)$$

with $E \left[\tilde{\theta}, \psi_{\tilde{\theta}}(\theta | y) \right]$ given by (5) and

$$\psi_{\bar{x}_n}(y) = \frac{1}{F(b) - F(a)} \cdot \frac{1}{\sqrt{2\pi(\tau^2 + \sigma^2/n)}} \exp \left[-\frac{1}{2} \frac{(y - \mu)^2}{\tau^2 + \sigma^2/n} \right] \cdot \int_a^b \frac{1}{\sqrt{2\pi\rho_n^2}} \exp \left[-\frac{1}{2} \frac{(\theta - \mu_n)^2}{\rho_n^2} \right] d\theta. \quad (12)$$

Consider now the special case of an information cell (10), denoted $I_t(\varepsilon) \in \mathcal{I}_t$, such that $A_k = (x_k - \varepsilon, x_k + \varepsilon]$, $k = t - n, \dots, t$, for some $\varepsilon > 0$. That is, instead of learning the exact value of the random variable X_k the period t agent only learns that this value lies in the interval $(x_k - \varepsilon, x_k + \varepsilon]$ implying for the corresponding set of possible average values

$$\bar{I}_n(\varepsilon) = (\bar{x}_n - \varepsilon, \bar{x}_n + \varepsilon] \quad (13)$$

with

$$\bar{x}_n = \frac{1}{n} \sum_{k=t-n}^t x_k.$$

Corollary 2. *Consider some information cell $I_t(\varepsilon) \in \mathcal{I}_t$ as defined by (13). The period t agent's rational Bayesian estimator $E \left[\tilde{\theta}, \pi(\cdot | I_t) \right]$ can be approximated by (5) for sufficiently small ε . That is,*

$$\lim_{\varepsilon \rightarrow 0} E \left[\tilde{\theta}, \psi_{\tilde{\theta}}(\theta | \bar{I}_n(\varepsilon)) \right] = E \left[\tilde{\theta}, \psi_{\tilde{\theta}}(\theta | \bar{x}_n) \right].$$

With (11) we have constructed a closed-form model of a rational Bayesian estimator of the mean of a normal distribution under the assumption that the observed sample information has positive probability (e.g., belongs to some non-degenerate interval). Moreover, this Bayesian estimator converges to some value that is close to the true mean whenever, first, the sample information is close to point-observations and, second, the agent has no lack of memory (i.e., the sequence of σ -algebras $\Sigma(\mathcal{I}_1), \Sigma(\mathcal{I}_2), \dots$ constitutes a filtration). Such a consistency result does not apply to a Bayesian estimator defined as the Choquet expected value with respect to a non-additive probability measure. This will drive the results of the remainder of this paper.

4 Choquet Bayesian learning

4.1 Choquet decision theory and neo-additive capacities

We briefly recall basic elements of Choquet expected utility theory and neo-additive capacities. CEU theory was first axiomatized by Schmeidler (1986, 1989) within the Anscombe and Aumann (1963) framework, which assumes preferences over objective probability distributions. Subsequently, Gilboa (1987) as well as Sarin and Wakker (1992) have presented CEU axiomizations within the Savage (1954) framework, assuming a purely subjective notion of likelihood.

Consider a measurable space (Ω, \mathcal{F}) with \mathcal{F} denoting a σ -algebra on the state space Ω and a non-additive probability measure (=capacity) $\nu : \mathcal{F} \rightarrow [0, 1]$ satisfying

- (i) $\nu(\emptyset) = 0, \nu(\Omega) = 1$
- (ii) $A \subset B \Rightarrow \nu(A) \leq \nu(B)$ for all $A, B \in \mathcal{F}$.

The Choquet integral of a bounded function $f : \Omega \rightarrow \mathbb{R}$ with respect to capacity ν is defined as the following Riemann integral extended to domain Ω (Schmeidler 1986):

$$E[f, \nu(d\omega)] = \int_{-\infty}^0 (\nu(\{\omega \in \Omega \mid f(\omega) \geq z\}) - 1) dz + \int_0^{+\infty} \nu(\{\omega \in \Omega \mid f(\omega) \geq z\}) dz \quad (14)$$

whereby I simply write $E[f, \nu]$ for $E[f, \nu(d\omega)]$. For example, assume that f takes on m different values such that A_1, \dots, A_m is the unique partition of Ω with $f(\omega_1) > \dots > f(\omega_m)$ for $\omega_i \in A_i$. Then the Choquet expectation (14) becomes

$$E[f, \nu] = \sum_{i=1}^m f(\omega_i) \cdot [\nu(A_1 \cup \dots \cup A_i) - \nu(A_1 \cup \dots \cup A_{i-1})]. \quad (15)$$

We focus on non-additive probability measures that are defined as *neo-additive capacities* in the sense of Chateauneuf, Eichberger and Grant (2007). Recall that the set of *null events*, denoted \mathcal{N} , collects all events that the decision maker deems impossible.

Definition. Fix some set of null-events $\mathcal{N} \subset \mathcal{F}$ for the measurable space (Ω, \mathcal{F}) . The neo-additive capacity, ν , is defined, for some $\delta, \lambda \in [0, 1]$ by

$$\nu(A) = \delta \cdot \nu_\lambda(A) + (1 - \delta) \cdot \pi(A) \quad (16)$$

for all $A \in \mathcal{F}$ such that π is some additive probability measure satisfying

$$\pi(A) = \begin{cases} 0 & \text{if } A \in \mathcal{N} \\ 1 & \text{if } \Omega \setminus A \in \mathcal{N} \end{cases}$$

and the non-additive probability measure ν_λ is defined as follows

$$\nu_\lambda(A) = \begin{cases} 0 & \text{iff } A \in \mathcal{N} \\ \lambda & \text{else} \\ 1 & \text{iff } \Omega \setminus A \in \mathcal{N}. \end{cases}$$

For simplicity we restrict attention to sets of null-events \mathcal{N} such that $A \in \mathcal{N}$ if and only if $\pi(A) = 0$. Call A an *essential* event iff $0 < \pi(A) < 1$ and observe that the neo-additive capacity ν in (16) simplifies to

$$\nu(A) = \delta \cdot \lambda + (1 - \delta) \cdot \pi(A) \quad (17)$$

for essential A . Neo-additive capacities can thus be interpreted as non-additive beliefs that stand for deviations from additive beliefs such that a parameter δ (*degree of ambiguity*) measures the lack of confidence the agent has in some subjective additive probability measure π . The following observation extends a result (Lemma 3.1) of Chateauneuf, Eichberger and Grant (2007) for finite random variables to the more general case of random variables with a bounded range.

Observation 1. *Let $f : \Omega \rightarrow \mathbb{R}$ be an \mathcal{F} -measurable function with bounded range. The Choquet expected value (14) of f with respect to a neo-additive capacity (16) is then given by*

$$E[f, \nu] = \delta(\lambda \sup f + (1 - \lambda) \inf f) + (1 - \delta) E[f, \pi]. \quad (18)$$

According to Observation 1, the Choquet expected value of a random variable f with respect to a neo-additive capacity is a convex combination of the expected value of f with respect to some additive probability measure π and an ambiguity part. If there is no ambiguity, i.e., $\delta = 0$, then the Choquet expected value (18) reduces to the standard expected value of a random variable with respect to an additive probability measure. In case there is some ambiguity, however, the second parameter λ measures how much weight the decision maker puts on the least upper bound of the range of f . Conversely, $(1 - \lambda)$ is the weight he puts on the greatest lower bound. For high (resp. low) values of λ , the Choquet expected value $E[f, \nu]$ is greater (resp. less) than the corresponding additive expected values $E[f, \pi]$ whenever the agent expresses some ambiguity, i.e., $\delta > 0$. This formal feature gives rise to the possibility of over-estimation (resp. under-estimation) with respect to the rational benchmark case.

4.2 Full Bayesian updating of neo-additive capacities

Recall that a Savage act f is an \mathcal{F} -measurable function that maps the state space Ω into an arbitrary set of consequences X , i.e., $f : \Omega \rightarrow X$. Whenever $\omega \in \Omega$ is the (unique) true state of the world and the agent has chosen Savage act f , he ends up with the consequence $f(\omega)$ after the uncertainty is resolved. Given two Savage acts f, g , let $I, \neg I \in \mathcal{F}$ denote two complementary (essential) events and define the Savage act $f_I g : \Omega \rightarrow X$ as follows:

$$f_I g(\omega) = \begin{cases} f(\omega) & \text{for } \omega \in I \\ g(\omega) & \text{for } \omega \in \neg I. \end{cases} \quad (19)$$

Ex ante preferences over Savage acts, denoted \succeq , are interpreted as the decision maker's preferences before he receives any information. In contrast, ex post preferences over Savage acts, denoted \succeq_I , are interpreted as preferences conditional on I , i.e., after the decision maker has observed the information $I \in \mathcal{F}$.

A Bayesian update rule specifies how the ex ante preference ordering \succeq over Savage acts determines, for all essential $I \in \mathcal{F}$, the ex post preference orderings \succeq_I , which describe the agent's preferences over Savage acts after having received information I . Consider, for example, Bayesian update rules such that for every essential $I \in \mathcal{F}$ and every pair of Savage acts f, g there is some Savage act h satisfying

$$f_I h \succeq g_I h \text{ implies } f_I h' \succeq_I g_I h'' \text{ for all } h', h''. \quad (20)$$

Because the ex post preferences are independent of consequences that happened outside of observation I , such update rules satisfy the principle of consequentialism but violate dynamic consistency whenever the sure-thing principle is violated. In contrast to the unique Bayesian update rule for additive probability measures, there exists a multitude of perceivable Bayesian update rules satisfying (20). In the present paper we focus on the so-called *full* (or *generalized*) Bayesian update rule for which (20) becomes: for all acts f, g ,

$$f \succeq_I g \text{ if and only if } f_I h \succeq g_I h \quad (21)$$

such that h is the *conditional certainty equivalent* of g , i.e., h is the constant act such that $g \sim_I h$. The corresponding definition of conditional non-additive probability measures for Choquet decision makers is then given as follows (Eichberger, Grant, and Kelsey 2006):

$$\nu(A | I) = \frac{\nu(A \cap I)}{\nu(A \cap I) + 1 - \nu(A \cup \neg I)} \quad (22)$$

for $A, I \in \mathcal{F}$ such that I is essential. For the special case of conditional neo-additive capacities we obtain:

Observation 2. *Suppose that $0 < \pi(I) < 1$. Then an application of the generalized Bayesian update rule (22) to a prior neo-additive capacity (17) results in the posterior neo-additive capacity*

$$\nu(A | I) = \delta_I \cdot \lambda + (1 - \delta_I) \cdot \pi(A | I) \quad (23)$$

such that

$$\delta_I = \frac{\delta}{\delta + (1 - \delta) \cdot \pi(I)} \quad (24)$$

and

$$\pi(A | I) = \frac{\pi(A \cap I)}{\pi(I)}. \quad (25)$$

The full Bayesian update rule is popular in the literature because it avoids the extreme updating behavior of the so-called optimistic and pessimistic update rules (Gilboa and Schmeidler 1993). Furthermore, Cohen, Gilboa, Jaffray, and Schmeidler (2000) experimentally investigate the question whether the pessimistic—in their terms: the maximum likelihood (=ML)—update rule or the full Bayesian (=FB) update rule is consistent with ambiguity averse subjects' choice behavior. Their experimental findings establish an approximate ratio of 2:1 in favor for the full Bayesian update rule.

4.3 The Choquet Bayesian estimator of the mean of a normal distribution

This subsection goes back to the learning situation described in Section 3. However, instead of the additive probability space $(\pi, \Omega, \mathcal{F})$ we now consider the neo-additive probability space $(\nu(\cdot | \cdot), \Omega, \mathcal{F})$ such that, for all $A, I_t \in \mathcal{F}$ with $0 < \pi(A), \pi(I_t) < 1$, $\nu(A | I_t)$ is given by (23). The following proposition follows immediately from Observation 2 combined with Proposition 1.

Proposition 2. *Consider some information-cell $I_t \in \mathcal{I}_t$ such that $0 < \pi(I_t) < 1$. The period t agent's Choquet Bayesian estimator of the mean of X_k , $k \in \{1, 2, \dots\}$ conditional on I_t is given as*

$$\begin{aligned} & E \left[\tilde{\theta}, \nu(\cdot | I_t) \right] \\ &= \delta_{I_t} \cdot (\lambda \cdot b + (1 - \lambda) \cdot a) + (1 - \delta_{I_t}) \cdot E \left[\tilde{\theta}, \pi(\cdot | I_t) \right] \end{aligned} \quad (26)$$

with

$$\delta_{I_t} = \frac{\delta}{\delta + (1 - \delta) \cdot \pi(I_t)} \quad (27)$$

and $E \left[\tilde{\theta}, \pi (\cdot | I_t) \right]$ given by (11).

Observe that the Choquet Bayesian estimator (26) generically deviates from the rational Bayesian estimator (11) whenever $\delta_{I_t} > 0$. More precisely, whenever the agent expresses some ambiguity, i.e., $\delta > 0$, the Choquet Bayesian estimator results as a weighted average between the rational Bayesian estimator and the agent's attitudes towards resolving his ambiguity as expressed by the optimism parameter λ . If we think of the Choquet learning model of Proposition 2 as a class of learning models indexed by ambiguity parameter $\delta \in [0, 1]$, then rational Bayesian learning—as characterized in Proposition 1—corresponds to the Lebesgue measure zero special case for which the agent does not express any ambiguity, i.e., $\delta = 0$.

Furthermore, observe that the Choquet Bayesian estimator (26) puts the more weight δ_{I_t} on the ambiguity part the smaller the value of $\pi (I_t)$, i.e., the smaller the probability of observing information I_t about sample data. In other words, the more surprised an agent is when he observes new information, the less confident will he become about his additive belief. In contrast to the consistency result for rational Bayesian estimation, Choquet Bayesian estimation will therefore generically result in a long-run bias between the Choquet Bayesian estimator $E \left[\tilde{\theta}, \nu (\cdot | I_t) \right]$ and the rational Bayesian estimator $E \left[\tilde{\theta}, \pi (\cdot | I_t) \right]$. The following corollary characterizes the long-run Choquet Bayesian estimator for the relevant situation in which the additive prior probability attached to any specific sample observation goes towards zero as the sample-size becomes large.

Corollary 3. *Consider some sequence of information cells I_1, I_2, \dots such that*

$$\lim_{t \rightarrow \infty} \pi (I_t) = 0 \text{ a.s. } \pi.$$

Then

$$\lim_{n \rightarrow \infty} E \left[\tilde{\theta}, \nu (\cdot | I_t) \right] = \lambda \cdot b + (1 - \lambda) \cdot a \text{ a.s. } \pi.$$

That is, this long-run Choquet Bayesian estimator ignores any statistical information as conveyed by the limit of the sample mean.

Similarly, the Choquet Bayesian estimator $E \left[\tilde{\theta}, \nu (\cdot | I_t (\varepsilon)) \right]$ ignores any statistical information coming from the sample the more precise the information is, i.e., the closer $I_t (\varepsilon)$ resembles some point-observation.

Corollary 4. *Consider some information cell $I_t(\varepsilon) \in \mathcal{I}_t$ as defined by (13). The agent's Choquet Bayesian estimator $E[\tilde{\theta}, \nu(\cdot | I_t(\varepsilon))]$ can be approximated by*

$$\lambda \cdot b + (1 - \lambda) \cdot a$$

for sufficiently small ε . That is, for all $\omega \in \Omega$,

$$\lim_{\varepsilon \rightarrow 0} E[\tilde{\theta}, \nu(\cdot | I_t(\varepsilon))] = \lambda \cdot b + (1 - \lambda) \cdot a.$$

5 An illustrative application: The risk-free rate puzzle

In this section, we apply our model of Choquet Bayesian learning to the risk-free rate puzzle (Weil 1989). To this end we adopt the original utility structure of Mehra and Prescott (1985). The representative period t agent maximizes his expected utility from today's and future consumption under the assumptions that, first, the consumption growth rates are independent and identically normally distributed and, second, the agent's von Neumann Morgenstern utility function is additively time-separable and each period's utility of consumption is of the CRRA (constant relative risk aversion) form.

In contrast to the rational expectations hypothesis—which stipulates that subjective and objective distributions must coincide—the period t agent of our model bases his portfolio optimization decision on a subjective normal distribution over the consumption growth rate. The mean results from some process of Choquet Bayesian learning as derived in the previous section. That is, ambiguity in terms of non-additive beliefs enters the portfolio optimization problem only indirectly through the agent's learning behavior about distributional parameters. We further assume that the period t agent is unaware of any ambiguity-driven bias in his learning process to the effect that he—incorrectly—assumes that his Choquet estimator (11) is an unbiased estimator to which the law of iterated expectations applies. As a consequence, the period t agent of our model is myopic in the sense that he ignores any dynamic inconsistency that may result from the biased learning process but rather maximizes his expected utility over uncertain future consumption streams as if he solves a dynamically consistent maximization problem.⁵

⁵For an asset pricing model under ambiguity that considers dynamically inconsistent but sequentially sophisticated (i.e., non-myopic) agents see Zimper (2011).

5.1 The risk-free rate puzzle revisited

The period $t + 1$ equilibrium return, denoted R_{t+1} , of any given asset in a Lucas (1978) fruit-tree economy is characterized by the following Euler equation

$$1 = E [R_{t+1} \cdot M_{t+1}, \pi(\cdot | I_t)]. \quad (28)$$

where

$$M_{t+1} = \beta \cdot \frac{u'(C_{t+1})}{u'(C_t)}$$

is the *so-called stochastic discount factor*. $\beta < 1$ is the agent's time-discount factor, and C_t is consumption in period t .

For sake of simplicity, we assume that the per period utility function features constant relative risk aversion (CRRA),

$$u(C_t) = \frac{C_t^{1-\gamma} - 1}{1-\gamma}. \quad (29)$$

where γ is the coefficient of relative risk aversion. However, as pointed out by Weil (1989), the inter-temporal elasticity of substitution (IES)—given by $\phi = \frac{1}{\gamma}$ in case of CRRA preferences—is primarily relevant for the risk-free rate whereas the equity premium is almost exclusively driven by risk aversion. We therefore interpret (the inverse of) γ as measuring the IES. Our supplementary appendix provides a sensitivity analysis using recursive Epstein-Zin-Weil preferences (Epstein and Zin 1989; Weil 1989) that allow us to disentangle both motives. There, we show that the influence of risk aversion per se on our results is negligible. This finding directly relates to the equity premium puzzle, cf., also our concluding discussion.

With this parametric assumption, the constant equilibrium gross-return of any risk-free asset in (28) becomes

$$R_{t+1} = E [M_{t+1}, \pi(X_{t+1} | I_t)]^{-1} \quad (30)$$

where

$$M_{t+1} = \beta \cdot \exp(-\gamma X_{t+1})$$

and the random variable $X_{t+1} = \ln C_{t+1} - \ln C_t$ denotes consumption growth.

Suppose now that the probability distribution of the consumption growth rate is given as the normal distribution $N(\mu, \sigma^2)$ implying that

$$R_{t+1} = E [M, N(\mu, \sigma^2)]^{-1}. \quad (31)$$

Or, equivalently stated in logarithmic terms with $r = \ln R_{t+1}$,

$$r = -\ln \beta + \gamma \mu - \frac{1}{2} \gamma^2 \sigma^2. \quad (32)$$

Under the rational expectations hypothesis, the quantitative implications of (32) are strongly at odds with the data, as first demonstrated by Mehra and Prescott (1985). To see this let us calibrate the model with a discount factor of $\beta = 0.98$, an IES of $\phi = 0.5$ —corresponding to $\gamma = 2$ —which is the conventional estimate of the IES since Hall (1988). In a next step, we take as data updated versions of those studied, e.g., by Shiller (1981) and Mehra and Prescott (1985).⁶ The data are annual and we focus on the postwar period 1950 to 2004. The risk-free rate is computed as government bond yields and consumption is real per capita consumption of non-durables and services. Both series are inflation-adjusted by the annual consumer price index (CPI). The resulting moments are $\mu = 2.13\%$, and $\sigma = 1.08\%$. Substituting—under the rational expectations hypothesis—these moments in (32), we obtain a risk-free rate of about 5.7%. This exceeds our point estimate of the risk-free rate of 2.19% by about 3.6 percentage points.

5.2 Partial resolution of the puzzle

To address the risk-free rate puzzle through our model of biased Bayesian learning, we consider a representative period t agent who learns the mean of the economy’s consumption growth rate by observing sample-data from n previous periods with $n \leq t$. That is, we fix some period t —e.g., the “present”—and assume that the period t agent bases his estimator on information cells given as (10). In contrast to the rational expectations hypothesis we thus consider the subjective normal distribution $N(\mu^{bias}, \sigma^2)$ such that the subjective mean μ^{bias} of the consumption growth rate is given as the Choquet estimator derived in Proposition 2. As a consequence, the risk free rate of our approach, denoted r^{bias} , is characterized by the following equilibrium condition

$$r^{bias} = -\ln \beta + \psi \cdot E \left[\tilde{\theta}, \nu(\cdot | I_t) \right] - \frac{1}{2} \cdot \gamma^2 \cdot \sigma^2 \quad (33)$$

with the Choquet estimator $E \left[\tilde{\theta}, \nu(\cdot | I_t) \right]$ given by (26).

A full resolution of the risk-free rate puzzle through our model of biased Bayesian learning means that the equilibrium condition (33) results in a risk-free rate of 2.19%. One perceivable approach to calibration is to accordingly let the data be informative about the structural model parameters. Following this approach, we would fix r^f to the empirical estimate and invert (33) to compute the Choquet Bayesian estimator

$$E \left[\tilde{\theta}, \nu(\cdot | I_t) \right] = \frac{1}{\gamma} \left(r^f + \ln \beta + \frac{1}{2} \cdot \gamma^2 \cdot \sigma^2 \right)$$

for appropriate choices of β, γ and σ .

⁶We take the data from Robert Shiller’s website <http://www.econ.yale.edu/shiller/data/chapt26.xls>.

To provide a numerical example, suppose that we feed into the model the same structural model parameters as for rational expectations, hence $\beta = 0.98$, $\gamma = 2$ and $\sigma = 0.0108$. We then get $E[\tilde{\theta}, \nu(\cdot | I_t)] = 0.001$. Instead of considering the objective mean of the consumption-growth rate, given by 2.13%, the period t -agent of our model would have to base his utility maximization problem on a subjective mean close to zero.

With the new⁷ free parameters of our model—mainly δ_{I_t} and λ (see below)—it is mathematically easy to generate a wide range of possible parameter constellations such that our model would fully “explain” the risk-free rate puzzle. The implied degrees of pessimism ($1 - \lambda$) and ambiguity would, however, be too high to be regarded as reasonable. We therefore choose a reasonable parametrization where key structural model parameters are not determined within the model. We then ask for the scope of our model to contribute towards a (partial) resolution of the risk-free rate puzzle.

5.2.1 Calibration of the Choquet model

For calibration of our model, notice, first, that our model significantly alters beliefs and updating of beliefs relative to rational expectations. Parameters used under rational expectations are therefore generally inappropriate. Second, our model gives rise to two additional parameters, the degree of optimism λ and the limit degree of ambiguity, δ_{I_t} .

As in our calibration of the rational expectations model, we exogenously feed in the point estimates of μ , σ and r^f . With regard to μ and the range of consumption realizations that the agent regards as reasonable, we introduce two rather innocuous assumptions.

Assumption (A1). We set zero as the largest lower bound for the possible mean of the consumption growth rate. In addition, we fix an interval of possible values for the mean of the consumption-growth rate that is centered around the true value $\mu = 0.00213$ resulting in the following support of the agent’s prior

$$\Theta \equiv (a, b) = (0, 2 \cdot \mu) = (0, 0.0426).$$

Assumption (A2). We set the agent’s rational Bayesian estimator for the consumption-growth rate, as derived in Proposition 1, equal to the true value $\mu = 0.0213$, i.e.,

$$E[\tilde{\theta}, \pi(\cdot | I_t)] = \mu = 0.0213.$$

⁷“New” in the sense that these parameters are relevant to our Choquet Bayesian learning model but become obsolete in the standard Bayesian learning model with additive beliefs.

By assumption (A1), we suppose that the agent only regards a relatively narrow range of mean consumption growth rates as possible. By assumption (A2), the rational part $E[\tilde{\theta}, \pi(\cdot | I_t)]$ of the Choquet estimator $E[\tilde{\theta}, \nu(\cdot | I_t)]$ behaves as if the agent has already learnt the true parameter value. As a consequence of assumption (A2), any difference between our equilibrium condition (33) and the rational expectations equilibrium condition (32) is therefore driven by the difference between our biased Bayesian learning model and the rational benchmark case. Combining assumptions (A1) and (A2), the Choquet Bayesian estimator (26) simplifies to

$$\begin{aligned} E[\tilde{\theta}, \nu(\cdot | I_t)] &= \delta_{I_t} \cdot \lambda \cdot b + (1 - \delta_{I_t}) \cdot \mu \\ &= (1 + (2 \cdot \lambda - 1) \cdot \delta_{I_t}) \cdot \mu. \end{aligned}$$

It remains to determine reasonable parameter values for the free structural model parameters β, γ, λ and δ_{I_t} . We are not aware of any experiments on simple capacities or the related probability weighting functions arising in PT 1992 in an inter-temporal setting. We therefore fix β and $\phi = \frac{1}{\gamma}$ to the above standard value from the rational expectations literature, hence $\beta = 0.98$ and $\gamma = 2$.

As to λ and δ_{I_t} we proceed as follows. According to observation 2, Choquet integration with simple capacities effectively transforms objective probabilities into a linear function with two discontinuous jumps at the endpoints, i.e., $\nu(A | I) = (\delta_{I_t}) \cdot \lambda + (1 - \delta_{I_t}) \cdot \pi(A | I) = 0$ for $\pi(A | I) \in (0, 1)$ and $\nu(A | I) = 0$ for $\pi(A | I) = 0$ as well as $\nu(A | I) = 1$ for $\pi(A | I) = 1$. For $\delta \in (0, 1)$ and $\lambda \in (0, 1)$ this gives higher decision weight to outcomes with low objective probability (and lower weight to outcomes with high objective probability) than under rational expectations, just as in PT 1992. Experiments based on PT 1992 typically use probability weighting functions with a single parameter, e.g.,

$$\tilde{\nu}(A)(\xi) = \frac{\pi(A)^\xi}{(\pi(A)^\xi + (1 - \pi(A))^\xi)^{\frac{1}{\xi}}} \quad (34)$$

cf., e.g., Tversky and Kahnemann (1992) as well as Wu and Gonzalez (1996). Standard estimates of ξ , as, e.g., reported by Wu and Gonzalez (1996), are in the range of $[0.3, 0.9]$. This gives rise to equivalent qualitative predictions as simple capacities. The point estimate of Tversky and Kahnemann (1992) is 0.61 which is well in this range. As our baseline value we use the point estimate reported by Tversky and Kahnemann. As sensitivity analysis, we consider the standard range of ξ , $\mathcal{G}^\xi = [0.3, 0.9]$.

In order to determine λ and δ_{I_t} from these values, we interpret the experimental evidence as reflecting stationary posterior beliefs. Given $\tilde{\nu}(A)(\xi)$ we can then identify λ

and δ_{I_t} from observation 2 by running an ordinary least-squares regression of $\tilde{\nu}(A)(\xi)$ on a constant and $\pi(A)$ for a pre-specified grid of $\pi(A) \in (0, 1)$. The regression equation is hence

$$\tilde{\nu}(A)(\xi) = \beta_0 + \beta_1\pi(A) + \epsilon$$

and we identify, according to observation 2, the respective degrees of optimism and ambiguity by $\beta_0 = \delta_{I_t}\lambda$ and $\beta_1 = 1 - \delta_{I_t}$. For $\xi = 0.61$ this results in $\lambda \approx 0.3$ and $\delta_{I_t} \approx 0.3$.

Figure 1 displays the linear map of objective probabilities—denoted as π —to their respective subjective counterparts—denoted by ν —under rational expectations—for which $\xi = 1$ —as a blue solid line. The black dashed dotted line is the PT 1992 probability weighting function derived from the Tversky and Kahnemann (1992) point estimate of $\xi = 0.61$. The linear red dashed line is the approximation with simple capacities resulting from our




Figure 1: Probability Weighting Function and Approximation with Simple Capacities

Figure 2 repeats the exercise for the entire range of $\xi \in \mathcal{G}^\xi$. As the probability weighting function becomes more strongly S-shaped—i.e., as ξ decreases—the implied

degree of ambiguity increases and the implied degree of pessimism decreases. Overall, we find that the implied degree of ambiguity is positively related to the degree of pessimism.




Figure 2: Implied Ambiguity and Optimism

Wu and Gonzalez (1996) derive these estimates of ξ from static experiments in which they combine (34) with a CRRA utility function as in (29). Consistently with numerous other experimental studies, Wu and Gonzalez report that estimated coefficients of relative risk aversion (RRA) are considerably less than one. The reason is that S-shaped probability functions preserve risk aversion. Estimates of risk aversion based on models with objective probabilities are consequently biased upward due to omitted variables. One may therefore argue that we should calibrate our model with simultaneously low values of the IES and RRA. As stated previously, we address this aspect in our supplementary appendix where we show that the effect of RRA per se is low.

5.2.2 Results

For our benchmark calibration—i.e., for $\beta = 0.98$, $\gamma = 2$, $\lambda \approx 0.3$ and $\delta_{I_t} \approx 0.3$ —our model predicts a biased mean consumption growth rate of 0.016 and a corresponding

risk-free rate of 5.2 percent. This is 0.5 percentage points lower than the corresponding rational expectations model predicts. Figure 3 repeats the exercise for all parameter combinations $\mathcal{G}^\xi \in [0.3, 0.9]$. The risk-free rate decreases quite strongly when ξ is decreased. The minimum risk-free rate predicted by our model for $\xi = 0.3$ —corresponding to $\delta_{I_t} \approx 0.7$ and $\lambda \approx 0.1$ —is roughly 3.5 percent. We therefore conclude that our model can contribute towards a resolution of the risk-free rate puzzle by up to 2 percentage points.




Figure 3: Risk-free rate under biased beliefs

6 Concluding remarks

This paper develops a closed-form model of rational Bayesian learning in which an agent learns the mean of a normal distribution by updating his truncated normal prior through observations from an independent and identically normally distributed data process. We further extend this model by allowing for ambiguity attitudes as expressed by neo-additive capacities. The resulting Choquet estimator of the mean of the normal distribution is generically biased, also in the long-run. We argue that our learning model thus gives a decision-theoretically rigorous answer to the question why subjective beliefs of Bayesian learners may persistently differ from their objective counterparts.

As an illustrative application of our learning model we calibrate the standard equilibrium conditions of the consumption based asset pricing model with normally distributed consumption growth. For plausible values of the decision-theoretic parameters which we derive from estimates arising in cumulative prospect theory we obtain a risk-free rate which can be considerably lower. The lowest value predicted in our model for a reasonable degree of pessimism is 3.5%. This is about two percentage points lower than the risk-free rate predicted under the rational expectations counterpart of the model. Although our approach can thus contribute towards a partial resolution of the risk-free rate puzzle, we also conclude that biased Bayesian learning alone cannot resolve the risk-free rate puzzle entirely.

Our work gives rise to two avenues of future research. First, in order to derive our illustrative findings, we have embedded our Choquet learning model of the mean consumption growth into a consumption based asset pricing model. In this experiment, we have treated as constant the variance of consumption growth and took the objective data as point estimate. We plan to extend our framework to a general asset pricing model with at least two assets which incorporates uncertainty and learning with respect to the mean and the variance. We will then ask whether such an extended framework will give rise to a reasonable risk-free rate and equity premium. Second, our discussion of calibration of our model has highlighted the importance of choosing reasonable parameters in models with psychologically biased beliefs. Future research has to design experiments in an inter-temporal setting that allow us to disentangle inter-temporal parameters (the discount rate and the inter-temporal elasticity of substitution) and parameters measuring risk (relative risk aversion, degree of optimism) and ambiguity. Such experiments will provide a more solid ground for calibration.

Mathematical appendix

Proof of the Lemma. Note that under our distributional assumptions

$$\bar{X}_n = \frac{1}{n} (X_1 + \dots + X_n)$$

is a sufficient statistic for $\tilde{\theta}$, implying (cf. Lemma 1 in Berger 1985) that

$$\pi(\cdot \mid X_1, \dots, X_n) = \pi(\cdot \mid \bar{X}_n).$$

Consequently, for any point-observation (4),

$$\begin{aligned} E \left[\tilde{\theta}, \pi(\cdot \mid x_1, \dots, x_n) \right] &= E \left[\tilde{\theta}, \pi(\cdot \mid \bar{x}_n) \right] \\ &= E \left[\tilde{\theta}, \psi_{\tilde{\theta}}(\theta \mid \bar{x}_n) \right]. \end{aligned}$$

To complete the proof of the Lemma it therefore remains to derive an analytical expression for

$$E \left[\tilde{\theta}, \psi_{\tilde{\theta}}(\theta \mid \bar{x}_n) \right].$$

According to our distributional assumptions, \bar{X}_n is—conditional on parameter θ —normally distributed such that, for all $\theta \in (a, b)$ and all $A \in \mathcal{B}(\mathbb{R})$,

$$\pi \left(\left\{ \omega \in \Omega \mid \tilde{\theta}(\omega) = \theta, \bar{X}_n(\omega) \in A \right\} \right) = \int_{x \in A} \psi_{\bar{X}_n}(x \mid \theta) \cdot \psi_{\tilde{\theta}}(\theta) dx$$

with

$$\psi_{\bar{X}_n}(x \mid \theta) = \frac{1}{\sqrt{2\pi\sigma^2/n}} \exp \left[-\frac{1}{2} \frac{(x - \theta)^2}{\sigma^2/n} \right] \text{ for } x \in \mathbb{R}.$$

Let ψ denote the joint density function of $\tilde{\theta}$ and \bar{X}_n . Then, by Bayes' rule, the posterior density function of $\tilde{\theta}$ conditional on $\bar{x}_n \in \bar{X}_n(\Omega)$ is given by

$$\begin{aligned} \psi_{\tilde{\theta}}(\theta \mid \bar{x}_n) &= \frac{\psi(\theta, \bar{x}_n)}{\psi_{\bar{X}_n}(\bar{x}_n)} \\ &= \frac{\psi_{\bar{X}_n}(\bar{x}_n \mid \theta) \cdot \psi_{\tilde{\theta}}(\theta)}{\int_{\mu-a}^{\mu+a} \psi_{\bar{X}_n}(\bar{x}_n \mid \theta) \cdot \psi_{\tilde{\theta}}(\theta) d\theta}. \end{aligned}$$

Standard calculations (Berger 1985, p.127) show that, for $\theta \in (a, b)$,

$$\begin{aligned} \psi(\theta, \bar{x}_n) &= \frac{1}{F(b) - F(a)} \\ &\cdot \frac{1}{\sqrt{2\pi\rho_n^2}} \exp \left[-\frac{1}{2} \frac{(\theta - \mu_n^*)^2}{\rho_n^2} \right] \cdot \frac{1}{\sqrt{2\pi(\tau^2 + \sigma^2/n)}} \exp \left[-\frac{1}{2} \frac{(\bar{x}_n - \mu)^2}{\tau^2 + \sigma^2/n} \right] \end{aligned}$$

where

$$\mu_n = \frac{\sigma^2/n}{\sigma^2/n + \tau^2} \cdot \mu + \frac{\tau^2}{\sigma^2/n + \tau^2} \cdot \bar{x}_n$$

and

$$\rho_n^2 = \frac{\tau^2 \sigma^2/n}{\tau^2 + \sigma^2/n}.$$

Consequently,

$$\begin{aligned} \psi_{\bar{X}_n}(\bar{x}_n) &= \int_a^b \psi(\theta, \bar{x}_n) d\theta \\ &= \frac{1}{F(b) - F(a)} \cdot \\ &\quad \frac{1}{\sqrt{2\pi(\tau^2 + \sigma^2/n)}} \exp\left[-\frac{1}{2} \frac{(\bar{x}_n - \mu)^2}{\tau^2 + \sigma^2/n}\right] \cdot \int_{\mu-a}^{\mu+a} \frac{1}{\sqrt{2\pi\rho_n^2}} \exp\left[-\frac{1}{2} \frac{(\theta - \mu_n)^2}{\rho_n^2}\right] d\theta, \end{aligned} \quad (35)$$

so that, by an application of Bayes' rule,

$$\begin{aligned} \psi_{\tilde{\theta}}(\theta | \bar{x}_n) &= \frac{\psi(\theta, \bar{x}_n)}{\psi_{\bar{X}_n}(\bar{x}_n)} \\ &= \frac{\frac{1}{\sqrt{2\pi\rho_n^2}} \exp\left[-\frac{1}{2} \frac{(\theta - \mu_n)^2}{\rho_n^2}\right]}{\int_a^b \frac{1}{\sqrt{2\pi\rho_n^2}} \exp\left[-\frac{1}{2} \frac{(\theta - \mu_n)^2}{\rho_n^2}\right] d\theta} \\ &= \frac{1}{F_n(b) - F_n(a)} \frac{1}{\sqrt{2\pi\rho_n^2}} \exp\left[-\frac{1}{2} \frac{(\theta - \mu_n)^2}{\rho_n^2}\right], \end{aligned}$$

by the definition of F_n . Observe now that the conditional density function $\psi_{\tilde{\theta}}(\theta | \bar{x}_n)$ corresponds to a truncated normal distribution of $\tilde{\theta}$ with support on (a, b) . Conditional on \bar{x}_n , $\tilde{\theta}$ has therefore the expected value

$$\begin{aligned} E\left[\tilde{\theta}, \psi_{\tilde{\theta}}(\theta | \bar{x}_n)\right] &= \int_a^b \theta \cdot \psi_{\tilde{\theta}}(\theta | \bar{x}_n) d\theta \\ &= \mu_n - \rho_n \frac{f_n(b) - f_n(a)}{F_n(b) - F_n(a)}. \end{aligned}$$

This proves the Lemma. \square

Proof of Corollary 1. Observe at first that

$$\lim_{n \rightarrow \infty} \rho_n = \lim_{n \rightarrow \infty} \frac{\tau^2 \sigma^2}{n \cdot \tau^2 + \sigma^2} = 0. \quad (36)$$

By the strong law of large numbers (Theorem 22.1 in Billingsley 1995), $\{\bar{X}_n\}_{n \in \mathbb{N}}$ converges with $\pi(\cdot | \mu^*)$ probability one to μ^* as n gets large. Since all probability mass

gets thus concentrated around μ^* , it holds for any point-observation (4) that

$$\lim_{n \rightarrow \infty} f_n(b) = 0 \quad (37)$$

$$\lim_{n \rightarrow \infty} f_n(a) = 0 \quad (38)$$

$$\lim_{n \rightarrow \infty} F_n(b) = 1 \quad (39)$$

$$\lim_{n \rightarrow \infty} F_n(a) = 0 \quad (40)$$

with $\pi(\cdot | \mu^*)$ probability one whenever $\mu^* \in (a, b)$. Collecting equations (36) - (40) and substituting into the limit of (5) gives

$$\begin{aligned} \lim_{n \rightarrow \infty} E \left[\tilde{\theta}, \pi(\cdot | X_1, \dots, X_n) \right] &= \lim_{n \rightarrow \infty} \left(\frac{\sigma^2/n}{\sigma^2/n + \tau^2} \cdot \mu + \frac{\tau^2}{\sigma^2/n + \tau^2} \cdot \bar{X}_n \right) \\ &\quad - \lim_{n \rightarrow \infty} \rho_n \frac{f_n(b) - f_n(a)}{F_n(b) - F_n(a)} \\ &= \mu^* - 0 \text{ a.s. } \pi(\cdot | \mu^*), \end{aligned}$$

which proves Corollary 1. \square

Proof of Proposition 1. Since $\Sigma_n \subset \Sigma(\mathcal{I}_t)$, we have, by the law of iterated expectations, that

$$\begin{aligned} E \left[\tilde{\theta}, \pi(\cdot | I_t) \right] &= E \left[\tilde{\theta}, \psi_{\tilde{\theta}}(\theta | \bar{I}_n) \right] \\ &= E \left[E \left[\tilde{\theta}, \psi_{\tilde{\theta}}(\theta | y) \right], \pi(y | \bar{I}_n) \right] \\ &= \int_{y \in \bar{I}_n} E \left[\tilde{\theta}, \psi_{\tilde{\theta}}(\theta | y) \right] \cdot \frac{\psi_{\bar{X}_n}(y)}{\int_{y \in \bar{I}_n} \psi_{\bar{X}_n}(y) dy} dy. \end{aligned}$$

The analytical expression for the density $\psi_{\bar{X}_n}$ has already been derived as (35) in the proof of the Lemma. \square

Proof of Corollary 2. Observe that

$$\begin{aligned}
& \lim_{\varepsilon \rightarrow 0} E \left[\tilde{\theta}, \psi_{\tilde{\theta}}(\theta \mid \bar{I}_n(\varepsilon)) \right] \\
&= \lim_{\varepsilon \rightarrow 0} E \left[E \left[\tilde{\theta}, \psi_{\tilde{\theta}}(\theta \mid y) \right], \pi(y \mid \bar{I}_n(\varepsilon)) \right] \\
&= \lim_{\varepsilon \rightarrow 0} \int_{\bar{x}_n - \varepsilon}^{\bar{x}_n + \varepsilon} E \left[\tilde{\theta}, \psi_{\tilde{\theta}}(\theta \mid y) \right] \cdot \frac{\psi_{\bar{X}_n}(y)}{\int_{\bar{x}_n - \varepsilon}^{\bar{x}_n + \varepsilon} \psi_{\bar{X}_n}(y) dy} dy \\
&= \lim_{\varepsilon \rightarrow 0} \frac{E \left[\tilde{\theta}, \psi_{\tilde{\theta}}(\theta \mid \bar{x}_n + \varepsilon) \right] \cdot \psi_{\bar{X}_n}(\bar{x}_n + \varepsilon) + E \left[\tilde{\theta}, \psi_{\tilde{\theta}}(\theta \mid \bar{x}_n - \varepsilon) \right] \cdot \psi_{\bar{X}_n}(\bar{x}_n - \varepsilon)}{\psi_{\bar{X}_n}(\bar{x}_n + \varepsilon) + \psi_{\bar{X}_n}(\bar{x}_n - \varepsilon)} \\
&= E \left[\tilde{\theta}, \psi_{\tilde{\theta}}(\theta \mid \bar{x}_n) \right],
\end{aligned}$$

whereby the fourth line follows from L'Hopital's rule combined with Leibniz's integral rule. \square

Proof of Observation 1: By an argument in Schmeidler (1986), it suffices to restrict attention to a non-negative function f so that

$$E[f, \nu] = \int_0^{+\infty} \nu(\{\omega \in \Omega \mid f(\omega) \geq z\}) dz,$$

which is equivalent to

$$E[f, \nu] = \int_{\inf f}^{\sup f} \nu(\{\omega \in \Omega \mid f(\omega) \geq z\}) dz \quad (41)$$

since f bounded. We consider a partition P_n , $n = 1, 2, \dots$, of Ω with members

$$A_n^k = \{\omega \in \Omega \mid a_{k,n} < f(\omega) \leq b_{k,n}\} \text{ for } k = 1, \dots, 2^n$$

such that

$$\begin{aligned}
a_{k,n} &= [\sup f - \inf f] \cdot \frac{(k-1)}{2^n} + \inf f \\
b_{k,n} &= [\sup f - \inf f] \cdot \frac{k}{2^n} + \inf f.
\end{aligned}$$

Define the step functions $a_n : \Omega \rightarrow \mathbb{R}$ and $b_n : \Omega \rightarrow \mathbb{R}$ such that, for $\omega \in A_n^k$, $k = 1, \dots, 2^n$,

$$\begin{aligned}
a_n(\omega) &= a_{k,n} \\
b_n(\omega) &= b_{k,n}.
\end{aligned}$$

Obviously,

$$E[a_n, \nu] \leq E[f, \nu] \leq E[b_n, \nu]$$

for all n and

$$\lim_{n \rightarrow \infty} E[b_n, \nu] - E[a_n, \nu] = 0.$$

That is, $E[a_n, \nu]$ and $E[b_n, \nu]$ converge to $E[f, \nu]$ for $n \rightarrow \infty$. Furthermore, observe that

$$\begin{aligned} \inf a_n &= \inf f \text{ for all } n, \text{ and} \\ \sup b_n &= \sup f \text{ for all } n. \end{aligned}$$

Since $\lim_{n \rightarrow \infty} \inf b_n = \lim_{n \rightarrow \infty} \inf a_n$ and $E[b_n, \mu]$ is continuous in n , we have

$$\begin{aligned} \lim_{n \rightarrow \infty} E[b_n, \nu] &= \delta \left(\lambda \lim_{n \rightarrow \infty} \sup b_n + (1 - \lambda) \lim_{n \rightarrow \infty} \inf b_n \right) + (1 - \delta) \lim_{n \rightarrow \infty} E[b_n, \pi] \\ &= \delta (\lambda \sup f + (1 - \lambda) \inf f) + (1 - \delta) E[f, \pi]. \end{aligned}$$

In order to prove Observation 1, it therefore remains to be shown that, for all n ,

$$E[b_n, \nu] = \delta (\lambda \sup b_n + (1 - \lambda) \inf b_n) + (1 - \delta) E[b_n, \pi].$$

Since b_n is a step function, (41) becomes

$$\begin{aligned} E[b_n, \nu] &= \sum_{A_n^k \in \mathcal{P}_n} \nu(A_n^{2^n} \cup \dots \cup A_n^k) \cdot (b_{k,n} - b_{k-1,n}) \\ &= \sum_{A_n^k \in \mathcal{P}_n} b_{k,n} \cdot [\nu(A_n^{2^n} \cup \dots \cup A_n^k) - \nu(A_n^{2^n} \cup \dots \cup A_n^{k-1})], \end{aligned}$$

implying for a neo-additive capacity

$$\begin{aligned} E[b_n, \nu] &= \sup b_n [\delta \lambda + (1 - \delta) \pi(A_n^{2^n})] + \sum_{k=2}^{2^n-1} b_{k,n} (1 - \delta) \pi(A_n^k) \\ &\quad + \inf b_n \left[1 - \delta \lambda - (1 - \delta) \sum_{k=2}^{2^n} \pi(A_n^k) \right] \\ &= \delta \lambda \sup b_n + (1 - \delta) \sum_{k=1}^{2^n} b_{k,n} \pi(A_n^k) + \inf b_n [\delta - \delta \lambda] \\ &= \delta (\lambda \sup b_n + (1 - \lambda) \inf b_n) + (1 - \delta) E[b_n, \pi]. \end{aligned}$$

□

Proof of Observation 2. Observe that

$$\begin{aligned}
\nu(A | I) &= \frac{\delta \cdot \lambda + (1 - \delta) \cdot \pi(A \cap I)}{\delta \cdot \lambda + (1 - \delta) \cdot \pi(A \cap I) + 1 - (\delta \cdot \lambda + (1 - \delta) \cdot \pi(A \cup \neg I))} \\
&= \frac{\delta \cdot \lambda}{\delta + (1 - \delta) \cdot \pi(I)} + \frac{(1 - \delta)}{\delta + (1 - \delta) \cdot \pi(I)} \pi(A \cap I) \\
&= \delta_I \cdot \lambda + (1 - \delta_I) \cdot \pi(A | I)
\end{aligned} \tag{42}$$

with δ_I given by (24). \square

Proof of Corollary 3. At first observe that, for all $\omega \in \Omega$ and all t ,

$$E \left[\tilde{\theta}, \pi(\cdot | I_t) \right] \in (\alpha, b);$$

that is, the possible values of $E \left[\tilde{\theta}, \pi(\cdot | \mathcal{I}_t) \right]$, $t = 1, 2, \dots$, are bounded in an Euclidean sense. To complete the proof just observe that

$$\begin{aligned}
\lim_{t \rightarrow \infty} \delta_{I_t} &= \lim_{t \rightarrow \infty} \frac{\delta}{\delta + (1 - \delta) \cdot \pi(I_t)} \\
&= 1
\end{aligned}$$

with probability one. \square

References

- Abel, A. B. (2002) “An Exploration of the Effects of Pessimism and Doubt on Asset Returns” *Journal of Economic Dynamics and Control* **26**, 1075-1092.
- Anscombe, F .J., and R. J. Aumann (1963) “A Definition of Subjective Probability” *Annals of American Statistics* **34**, 199-205.
- Assenza, T., and M. Berardi (2009) “Learning in a Credit Economy” *Journal of Economic Dynamics and Control* **33**, 1159-1169.
- Berger, J. O. (1985) *Statistical Decision Theory and Bayesian Analysis*, Springer: Berlin.
- Bigio, S. (2010) “Learning under Fear of Floating” *Journal of Economic Dynamics and Control* **34**, 1923-1950.
- Billingsley, P. (1995) *Probability and Measure*, John Wiley and & Sons, Inc.: New York, London, Sydney.
- Branch, William A., and George W. Evans, 2007, “Model Uncertainty and Endogenous Volatility,” *Review of Economic Dynamics*, 10, 207-237.
- Brav, A., and J. B. Heaton (2002) “Competing Theories of Financial Anomalies” *Review of Financial Studies* **15**, 575-606.
- Brennan, M. J., and Y. Xia (2001) “Stock Price Volatility and Equity Premium” *Journal of Monetary Economics* **47**, 249-283.
- Bullard, J., and K. Mitra (2002) “Learning about Monetary Policy Rules” *Journal of Monetary Economics* **49**, 1105-1129,
- Campbell, J. (1999) “Asset Prices, Consumption, and the Business Cycle” in: J. B. Taylor and M. Woodford (Eds.), *Handbook of Macroeconomics*, pp. 1231-1303, Amsterdam: North-Holland.
- Cecchetti, S. G., Lam, P., and N. C. Mark (2000) “Asset Pricing with Distorted Beliefs: Are Equity Returns Too Good to Be True?” *American Economic Review* **90**, 787-805.
- Chateauneuf, A., Eichberger, J., and S. Grant (2007) “Choice under Uncertainty with the Best and Worst in Mind: Neo-additive Capacities” *Journal of Economic Theory* **127**, 538-567.

- Cochrane, J. H. (2001) *Asset Pricing*, Princeton University Press: Princeton.
- Cohen, M., Gilboa, I., Jaffray, J.Y., and D. Schmeidler (2000) “An Experimental Study of Updating Ambiguous Beliefs” *Risk, Decision, and Policy* **5**, 123–133.
- Cyert, R. M., DeGroot, M. H., and C. A. Holt (1978) “Sequential Investment Decisions with Bayesian Learning” *Management Science* **24**, 712-718.
- Doob, J. L. (1949) “Application of the Theory of Martingales” in *Le Calcul des Probabilite's et ses Applications*, Colloques Internationaux du Centre National de la Recherche Scientifique Vol. 13. Centre National de la Recherche Scientifique, Paris, 23-27.
- Eichberger, J., Grant, S., and D. Kelsey (2006) “Updating Choquet Beliefs” *Journal of Mathematical Economics* **43**, 888-899.
- Ellsberg, D. (1961) “Risk, Ambiguity and the Savage Axioms” *Quarterly Journal of Economics* **75**, 643-669.
- Epstein, L.G. and Zin, S. (1989) “Substitution, risk aversion and the temporal behavior of consumption and asset returns: A theoretical framework ” *Econometrica* **57**, 937-969.
- Epstein, L. G., and M. Schneider (2003) “Recursive Multiple-Priors” *Journal of Economic Theory* **113**, 1-31.
- Epstein, L. G., and M. Schneider (2007) “Learning under Ambiguity” *Review of Economic Studies* **74**, 1275-1303.
- Epstein, L. G., and M. Schneider (2008) “Ambiguity, Information Quality, and Asset Pricing” *The Journal of Finance* **63**, 197-228.
- Evans, G., and S. Honkapohja (2001) *Learning and Expectations in Macroeconomics*. Princeton University Press: Princeton, NJ.
- Evans, G., and S. Honkapohja (2003) “Adaptive Learning and Monetary Policy Design” *Journal of Money, Credit and Banking* **35**, 1045-1072.
- Evans, G., and B. McGough (2005) “Monetary Policy, Indeterminacy and Learning” *Journal of Economic Dynamics and Control* **29**, 1809-1840.
- Ghirardato, P. (2002) “Revisiting Savage in a Conditional World” *Economic Theory* **20**, 83–92.

- Gilboa, I. (1987) “Expected Utility with Purely Subjective Non-Additive Probabilities” *Journal of Mathematical Economics* **16**, 65-88.
- Gilboa, I., and D. Schmeidler (1989) “Maxmin Expected Utility with Non-Unique Priors” *Journal of Mathematical Economics* **18**, 141-153.
- Gilboa, I., and D. Schmeidler (1993) “Updating Ambiguous Beliefs” *Journal of Economic Theory* **59**, 33-49.
- Giordani, P., and P. Soderlind (2006) “Is there Evidence of Pessimism and Doubt in Subjective Distributions? Implications for the Equity Premium Puzzle” *Journal of Economic Dynamics and Control* **30**, 1027–1043.
- Hall, R. E. (1988) “Intertemporal substitution in consumption” *Journal of Political Economy* **96**, 339-357.
- Hansen, L. P. and T. J. Sargent (2007). “Robustness”. Princeton University Press: Princeton.
- Jouini, E., and C. Napp (2008) “On Abel’s Concepts of Doubt and Pessimism” *Journal of Economic Dynamics and Control* **32**, 3682-3694.
- Kahneman, D., and A. Tversky (1979) “Prospect Theory: An Analysis of Decision under Risk” *Econometrica* **47**, 263-291.
- Kasa, K. (1999) “Will the FED ever Learn?” *Journal of Macroeconomics* **21**, 279–292.
- Kocherlakota, N. (1996) “The Equity Premium: It’s Still a Puzzle” *Journal of Economic Literature* **34**, 42-71.
- Lewellen, J., and J. Shanken (2000) “Learning, Asset-Pricing Tests, and Market Efficiency” *Journal of Finance* **57**, 1113-1145.
- Lucas, R.E., Jr. (1978) “Asset Prices in an Exchange Economy” *Econometrica* **46**, 1429-1446.
- Marinacci, M. (1999) “Limit Laws for Non-additive Probabilities and Their Frequentist Interpretation” *Journal of Economic Theory* **84**, 145-195.
- Marinacci, M. (2002) “Learning from Ambiguous Urns” *Statistical Papers* **43**, 145-151.
- Mehra, R., and E. C. Prescott (1985) “The Equity Premium: A Puzzle” *Journal of Monetary Economics* **15**, 145-161.

- Mehra, R., and E. C. Prescott (2003) “The Equity Premium in Retrospect” in: G.M. Constantinides, M. Harris and R.M. Stulz, eds., *Handbook of the Economics and Finance*. Amsterdam: Elsevier, 808-887.
- Pastor, L., and P. Veronesi (2003) “Stock Valuation and Learning about Profitability” *Journal of Finance* **58**, 1749-1789.
- Sarin, R., and P. P. Wakker (1998) “Dynamic Choice and NonExpected Utility” *Journal of Risk and Uncertainty* **17**, 87-119.
- Savage, L. J. (1954) *The Foundations of Statistics*, John Wiley and Sons, Inc.: New York, London, Sydney.
- Schmeidler, D. (1986) “Integral Representation without Additivity” *Proceedings of the American Mathematical Society* **97**, 255-261.
- Schmeidler, D. (1989) “Subjective Probability and Expected Utility without Additivity” *Econometrica* **57**, 571-587.
- Shiller, R. J., (1981) “Do Stock Prices Move Too Much to Be Justified by Subsequent Changes in Dividends?” *American Economic Review* **71**, 421-436.
- Siniscalchi, M. (2011) “Dynamic Choice under Ambiguity” *Theoretical Economics* **6**, 379-421.
- Tetlow, R. J., and P. von zur Muehlen (2009) “Robustifying Learnability” *Journal of Economic Dynamics and Control* **33**, 296-316.
- Tonks, I. (1983) “Bayesian Learning and the Optimal Investment Decision of the Firm” *The Economic Journal* **93**, 87-98.
- Tversky, A., and D. Kahneman (1992) “Advances in Prospect Theory: Cumulative Representations of Uncertainty” *Journal of Risk and Uncertainty* **5**, 297-323.
- Viscusi, W. K. (1979) *Employment Hazards: An Investigation of Market Performance*, Harvard University Press: Cambridge, Mass.
- Viscusi, W. K. (1990) “Do Smokers Underestimate Risks?” *The Journal of Political Economy* **98**, 1253-1269.
- Viscusi, W. K. (1991) “Age Variations in Risk Perceptions and Smoking Decisions” *The Review of Economics and Statistics* **73**, 577-588.

- Viscusi, W. K., and C. J. O'Connor (1984) "Adaptive Responses to Chemical Labeling: Are Workers Bayesian Decision Makers?" *The American Economic Review* **74**, 942-956.
- Wakker, P. P (2004) "On the Composition of Risk Preference and Belief" *Psychological Review* **111**, 236-241.
- Wakker, P. P. (2010) *Prospect Theory: For Risk and Ambiguity*, Cambridge University Press: Cambridge, UK.
- Wakker, P. P., and A. Tversky (1993) "An Axiomatization of Cumulative Prospect Theory" *Journal of Risk and Uncertainty* **7**, 147-176.
- Weil, P. (1989) "The Equity Premium Puzzle and the Risk-Free Rate Puzzle" *Journal of Monetary Economics* **24(2)**, 401-421.
- Weitzman, M. (2007) "Subjective Expectations and Asset-Return Puzzles" *American Economic Review*, **79(4)**, 1102-1130.
- Zimper, A. (2009) "Half Empty, Half Full and Why We Can Agree to Disagree Forever" *Journal of Economic Behavior and Organization*, **71**, 283-299.
- Zimper, A. (2011) "Asset Pricing in a Lucas Fruittree Economy with the Best and Worst in Mind" *Journal of Economic Dynamics and Control*, forthcoming.
- Zimper, A., and A. Ludwig (2009) "On Attitude Polarization under Bayesian Learning with Non-Additive Beliefs" *Journal of Risk and Uncertainty* **39**, 181-212.

Supplementary appendix

On dynamic consistency and consequentialism

The following discussion about the relationship between Savage's sure thing principle, dynamic consistency and consequentialism closely follows Zimper (2011). Throughout the text a Bayesian framework is considered that reduces in the standard way dynamic decision situations to a static Savage framework such that ex ante and ex post preferences are defined over Savage acts; that is, the reduction principle holds (Siniscalchi 2011). For such a framework Epstein and Schneider (2003) describe a decision theoretic model, i.e., the rectangular MMEU model, that combines some notion of ambiguity attitudes with dynamic consistency. Since the rectangular MMEU model also satisfies consequentialism, it follows from existing results in the decision theoretic literature (Ghirardato 2002, Epstein and Le Breton 1993) that this MMEU model must somehow violate universalism to the effect that both principles, consequentialism and dynamic consistency, cannot simultaneously hold at all events. In this supplementary appendix we make this argument more precise by showing in which specific sense the rectangular MMEU model violates universalism.

To be more precise consider some state space Ω and some σ -algebra \mathcal{F} which contains all events that are relevant to the decision maker in the sense that he attaches probabilities to these events, (for example, we might simply follow Savage (1954) and assume that \mathcal{F} is given as the power-set 2^Ω). In a next step, let us follow Epstein and Schneider (2003) and fix some partition Π of Ω whereby $\sigma(\Pi) \subseteq \mathcal{F}$ denotes the σ -algebra generated by Π . The events in $\sigma(\Pi)$ are interpreted as the events that are ex post observable by the decision maker. In what follows we will define behavioral principles restricted to the events in $\sigma(\Pi)$ whereby the principle of *universalism* is satisfied if and only if $\sigma(\Pi) = \mathcal{F}$. Intuitively speaking, universalism means that the behavioral principles will hold at all relevant events, i.e., at all events included in the largest possible information structure $\sigma(\Pi) = \mathcal{F}$, to the effect that the resulting decision theoretic model will be invariant with respect to any particularly fixed information structure $\sigma(\Pi)$.

Recall that a Savage act f maps Ω into a set of consequences X , i.e., $f : \Omega \rightarrow X$. Further, consider any finite partition $\{B_1, \dots, B_n\}$ of Ω such that $B_i \in \mathcal{F}$ for $i = 1, \dots, n$. Given Savage acts f^1, \dots, f^n , denote by $f_{B_1}^1 \dots f_{B_n}^n : \Omega \rightarrow X$ the Savage act such that

$$f_{B_1}^1 \dots f_{B_n}^n(\omega) = \begin{cases} f^1(\omega) & \text{for } \omega \in B_1 \\ \dots & \dots \\ f^n(\omega) & \text{for } \omega \in B_n. \end{cases}$$

For complements $B, \neg B \in \mathcal{F}$ and acts f, g , let us simply write $f_B g$ instead of $f_B g_{\neg B}$. Ex ante preferences over Savage acts, denoted \succeq , are interpreted as the decision maker's preferences before he receives any information. In contrast, ex post preferences over Savage acts, denoted \succeq_B , are interpreted as preferences conditional on B , i.e., after the decision maker has observed the occurrence of some non-null event $B \in \mathcal{F}$. It is assumed that \succeq as well as \succeq_B impose a weak order on all Savage acts.

Definition: $\sigma(\Pi)$ -consequentialism. *For all Savage acts f, g, h, h', h'', h''' and all non-null events $B \in \sigma(\Pi)$, the following condition holds for ex post preferences:*

$$f_B h \succeq_B g_B h' \Leftrightarrow f_B h'' \succeq_B g_B h''' \quad (43)$$

The notion of consequentialism in the context of decision trees under risk goes back to Hammond (1989) and Machina (1989) whereby Burks (1977) already refers to the same concept as “invariance axiom” (see the formal definition at p. 268, Burks 1977). The above definition is closest to Machina's interpretation once the reduction principle is applied to decision trees:

“The philosophy behind this approach [=consequentialism; the author] is that the uncertainty that was involved in the rest of the tree, as represented by the probabilities at the snipped-off chance nodes and the planned choices at the snipped-off choice nodes, is now irrelevant and should be treated as if it never existed.” (p. 1641)

In other words, (43) states that it should not matter for ex post preferences conditional on observation B whatever consequences, governed either by h, h', h'' , or h''' , might have happened outside of B . Innocuous and straightforward as this dynamic principle may appear, Machina (1989) argues strongly against its appeal for non-EU decision making. Nevertheless, consequentialism is, e.g., satisfied for the CEU decision makers of our model as well as for the MMEU decision makers in Epstein and Schneider (2003).

Definition: $\sigma(\Pi)$ -dynamic consistency. *For all Savage acts f, g, h and all non-null events $B \in \sigma(\Pi)$, the following condition holds for the relationship between ex ante and ex post preferences:*

$$f_B h \succeq g_B h \Leftrightarrow f_B h \succeq_B g_B h \quad (44)$$

Definition: $\sigma(\Pi)$ -sure thing principle. For all Savage acts f, g, h, h' and all events $B \in \sigma(\Pi)$, the following condition holds for ex ante preferences:

$$f_B h \succeq g_B h \Rightarrow f_B h' \succeq g_B h'. \quad (45)$$

The following observation basically restates an implication of Lemma 2 in Ghirardato (2002) for our slightly different definitions of consequentialism and dynamic consistency.

Observation. $\sigma(\Pi)$ -consequentialism (43) combined with $\sigma(\Pi)$ -dynamic consistency (44) implies the $\sigma(\Pi)$ -sure thing principle (45) under reduction.

Proof: Consider any non-null $B \in \sigma(\Pi)$. Observe that

$$\begin{aligned} f_B h &\succeq g_B h \\ &\Rightarrow f_B h \succeq_B g_B h \text{ by (44)} \\ &\Rightarrow f_B h' \succeq_B g_B h' \text{ by (43)} \\ &\Rightarrow f_B h' \succeq g_B h' \text{ by (44),} \end{aligned}$$

which proves the claim. \square

Next consider the following version of dynamic consistency restricted to events in $\sigma(\Pi)$ as introduced in Epstein and Schneider (2003) for their rectangular MMEU model.

Definition: Epstein and Schneider (E-S) $\sigma(\Pi)$ -dynamic consistency. Given any finite partition B_1, \dots, B_n of Ω such that $B_i \in \sigma(\Pi)$ for $i = 1, \dots, n$. For all Savage acts f^1, \dots, f^n and g^1, \dots, g^n , the following condition holds:

$$f^1 \succeq_{B_1} g^1, \dots, f^n \succeq_{B_n} g^n \Rightarrow f_{B_1}^1 \dots f_{B_n}^n \succeq g_{B_1}^1 \dots g_{B_n}^n, \quad (46)$$

whereby \succeq is strict iff at least one \succeq_{B_i} is strict.

Observation. $\sigma(\Pi)$ -consequentialism (43) combined with E-S $\sigma(\Pi)$ -dynamic consistency (46) implies the $\sigma(\Pi)$ -sure thing principle (45) under reduction.

Proof: By E-S dynamic $\sigma(\Pi)$ -consistency, we have for any non-null $B \in \sigma(\Pi)$

$$f_B h \succeq_B g_B h, h_{B^c} f \succeq_{B^c} h_{B^c} g \Rightarrow f_B h \succeq g_B h. \quad (47)$$

Furthermore, observe that

$$\begin{aligned} h_{-B} f &\succeq_{-B} h_{-B} f \text{ by reflexivity of } \succeq_{-B} \\ &\Leftrightarrow h_{-B} f \succeq_{-B} h_{-B} g \text{ by (43),} \end{aligned}$$

so that

$$h_{-B}f \succeq_{-B} h_{-B}g$$

is always satisfied under our assumptions of weak order and consequentialism. Consequently, (47) is equivalent to

$$f_B h \succeq_B g_B h \Rightarrow f_B h \succeq g_B h,$$

which implies, by completeness,

$$f_B h \succeq_B g_B h \Leftarrow f_B h \succeq g_B h.$$

That is, under the assumptions of the observation, E-S $\sigma(\Pi)$ -dynamic consistency is equivalent to $\sigma(\Pi)$ -dynamic consistency (44), which proves the result. \square

By the above observation, the rectangular MMEU model of Epstein and Schneider (2003)—satisfying $\sigma(\Pi)$ -consequentialism as well as the reduction principle—satisfies the sure thing principle at all events in $\sigma(\Pi)$. As a consequence, a violation of the sure thing principle is restricted to events that are not observable by the decision maker to the effect that ex ante ambiguity attitudes are governed by the decision maker's ex post information structure.

Results for recursive preferences

The effect of risk aversion per se on our results is negligible. To show this we here adopt recursive Epstein-Zin-Weil preferences (Epstein and Zin 1989, 1991; Weil 1989). Wu and Gonzalez (1996) document that estimates of risk-aversion are below one, once non-standard subjective beliefs are considered. Our calibration will therefore consider a flexible parametrization of risk aversion, holding the IES constant.

Adopting recursive preferences, the period t utility function is given by

$$U_t = \left\{ c_t^{1-\frac{1}{\phi}} + \beta \left(E_t [U_{t+1}^{1-\gamma}] \right)^{\frac{1-\frac{1}{\phi}}{1-\gamma}} \right\}^{\frac{1}{1-\frac{1}{\phi}}} \quad (48)$$

where β is the discount factor, ϕ is the inter-temporal rate of substitution and γ is the measure of relative risk aversion.

We continue to assume that the representative period t agent maximizes expected utility under the assumption that the consumption growth rates are independent and identically normally distributed. Employing results of Kocherlakota (1990, 1996) for Epstein-Zin-Weil preferences and our i.i.d. distributional assumption gives

$$M_{t+1} = \beta \left(E_t [(X_{t+1})^{1-\gamma}] \right)^{\frac{\gamma-\frac{1}{\phi}}{1-\gamma}} E_t [(X_{t+1})^{-\gamma}], \quad (49)$$

where the random variable $X_{t+1} = \ln C_{t+1} - \ln C_t$ denotes consumption growth as before (details of the derivation are provided below).

Under i.i.d. normally distributed consumption growth with parameters (μ, σ^2) we get from the above that

$$r = \ln(R) = -\ln \beta + \frac{1}{\phi} \mu - \frac{1}{2} \sigma^2 \left(\gamma - \frac{1}{\phi} (1 - \gamma) \right).$$

The point estimate of γ reported in the pooled regressions of Gonzalez and Wu (1986) is $\gamma = 0.5$. In order to illustrate how the *indirect* effect of biased beliefs—via re-calibration of γ relative to the rational expectations benchmark—adds to their *direct* effect on expected mean consumption growth, we extend this range to include the range used under rational expectations which is usually $\gamma \in [1, 4]$. This consideration gives $\gamma \in \mathcal{G}^\gamma = [0.5, 4.0]$.

We here repeat our calculation of subsection 5.2 for a box of all parameter combinations $\mathcal{G}^\xi \otimes \mathcal{G}^\gamma = [0.5, 0.9] \otimes [0.5, 4]$ holding ϕ constant at 0.5. Results are shown in figure 4. Our conclusions from the main text remain unaltered. In particular, we observe that increasing γ more or less decreases the risk-free rate predicted under rational expectations and our Choquet model in parallel such that differences between the two models are roughly constant. Moreover, the overall effect of γ on the level of the risk-free rate is small.

First-order conditions for recursive preferences

We here repeat the steps of Kocherlakota (1990). By the chain rule, the partial derivatives of (48) with respect to c_t and c_{t+1} are given by

$$\begin{aligned} U_{tc_t} &= U_t^{\frac{1}{\phi}} c_t^{-\frac{1}{\phi}} \\ U_{tc_{t+1}} &= U_t^{\frac{1}{\phi}} \beta \left(E_t [U_{t+1}^{1-\gamma}] \right)^{\frac{\gamma - \frac{1}{\phi}}{1-\gamma}} E_t \left[U_{t+1}^{\frac{1}{\phi} - \gamma} c_{t+1}^{-\frac{1}{\phi}} \right]. \end{aligned}$$

The marginal rate of substitution, or the stochastic discount factor M_{t+1} , is therefore

$$M_{t+1} = \beta \left(E_t [U_{t+1}^{1-\gamma}] \right)^{\frac{\gamma - \frac{1}{\phi}}{1-\gamma}} E_t \left[U_{t+1}^{\frac{1}{\phi} - \gamma} \left(\frac{c_{t+1}}{c_t} \right)^{-\frac{1}{\phi}} \right].$$

Next, observe that homotheticity and the i.i.d. assumption imply

$$U_t = \kappa c_t$$

for all t . Using this in the above we get (49).




Figure 4: Risk-free rate under biased beliefs: Recursive preferences

Additional references

- Burks, A.W. (1977) *Chance, Cause, Reason: An Inquiry into the Nature of Scientific Evidence* The University of Chicago Press, Chicago
- Hammond, P.J. (1988) "Consequentialist Foundation for Expected Utility" *Theory and Decision* **25**, 25-78.
- Kocherlakota, N. (1990) "Disentangling the Coefficient of Relative Risk Aversion from the Elasticity of Inter-temporal Substitution: An Irrelevance Result " *Journal of Finance* **45(1)**, 175-190.
- Machina, M. (1989) "Dynamic Consistency and Non-Expected Utility Models of Choice under Uncertainty" *Journal of Economic Literature* **27**, 1622-1668.