

Dynamic portfolio choice under ambiguity and regime switching mean returns

Hening Liu

► To cite this version:

Hening Liu. Dynamic portfolio choice under ambiguity and regime switching mean returns. Journal of Economic Dynamics and Control, 2011, 35 (4), pp.623. 10.1016/j.jedc.2010.12.012 . hal-00781344

HAL Id: hal-00781344 https://hal.science/hal-00781344

Submitted on 26 Jan 2013

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers. L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Author's Accepted Manuscript

Dynamic portfolio choice under ambiguity and regime switching mean returns

Hening Liu

PII: S0165-1889(10)00283-6 DOI: doi:10.1016/j.jedc.2010.12.012 DYNCON 2521 Reference:



www.elsevier.com/locate/jedc

To appear in:

Journal of Economic Dynamics & Control

21 May 2008 Received date: 13 December 2010 Accepted date:

Cite this article as: Hening Liu, Dynamic portfolio choice under ambiguity and regime switching mean returns, Journal of Economic Dynamics & Control, doi:10.1016/j.jedc.2010.12.012

This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting galley proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

Dynamic Portfolio Choice under Ambiguity and Regime Switching Mean Returns

Hening Liu[†] University of Manchester

December 15, 2010

Abstract

I examine a continuous-time intertemporal consumption and portfolio choice problem under ambiguity, where expected returns of a risky asset follow a hidden Markov chain. Investors with Chen and Epstein's (2002) recursive multiple priors utility possess a set of priors for unobservable investment opportunities. The optimal consumption and portfolio policies are explicitly characterized in terms of the Malliavin derivatives and stochastic integrals. When the model is calibrated to U.S. stock market data, I find that continuous Bayesian revisions under incomplete information generate ambiguity-driven hedging demands that mitigate intertemporal hedging demands. In addition, ambiguity aversion magnifies the importance of hedging demands in the optimal portfolio policies. Out-of-sample experiments demonstrate the economic importance of accounting for ambiguity.

JEL: G11, D81, C61

Keywords: Hidden Markov model, Malliavin derivative, portfolio choice, recursive multiple priors

[†] I am grateful to Carl Chiarella (the editor), David Feldman, Massimo Guidolin, Frederik Lundtofte, Jianjun Miao, Martin Schneider and several anonymous referees for numerous valuable comments and suggestions. I am indebted to Evan Anderson for his guidance and support. All errors are my own. Corresponding address: Manchester Accounting and Finance, Manchester Business School, University of Manchester, Booth Street West, Manchester M15 6PB, UK. E-mail: Hening.Liu@mbs.ac.uk. Telephone: ++44 16127 54492. Fax: ++44 16127 54023.

1 Introduction

Since the seminal work of Merton (1971), a number of papers have examined dynamic portfolio choice when investment opportunities are time varying (e.g., Campbell and Viceira, 1999; Kim and Omberg, 1996; Schroder and Skiadas, 1999, 2003). Assuming fully observable investment opportunities, this bulk of literature finds that the intertemporal hedging demand, which arises due to stochastic variation in investment opportunities, is important in portfolio decisions. However, in reality, investment opportunities are partially observable as moments of probability distributions of investment opportunities are often unobservable and must be estimated from observed market signals. Dothan and Feldman (1986) and Detemple (1986) were the first to study asset prices under incomplete information in general equilibrium, followed by David (1997), Veronesi (1999), Ai (2009), among others.¹ Other papers, to name a few, Gennotte (1986), Brennan (1998), Lakner (1998) and Honda (2003), analyze dynamic portfolio choice under incomplete information. Feldman (2007) provides an elaborate review of this literature and related discussions. Recently, Bjök et al. (2010) obtain explicit representations of the optimal wealth and investment processes for a wide range of partially observable investment opportunity sets. This growing body of literature employs recursive-filtering methods to estimate unobservable moments of distributions of asset returns based on observed asset prices. The stochastic processes describing the dynamics of the estimated moments are then treated as perfectly known and optimal consumption and portfolio policies can be derived using techniques for solving complete information economies. All these papers assume that investors have complete confidence in the probability law governing the evolution of the estimated moments with no concerns regarding model uncertainty, which will be relaxed in this paper. The importance of model uncertainty (or ambiguity) has been largely recognized in both the asset pricing literature (e.g., Anderson et al., 2003; Chen and Epstein, 2002; Gagliardini et al., 2008; Hansen and Sargent, 2001; Leippold et al., 2008; Trojani and Vanini, 2004) and the portfolio choice literature (e.g., Campanale, 2009; Liu, 2010; Maenhout, 2004,2006; Uppal and Wang, 2003).

My aim in this paper is to examine the effects of ambiguity on intertemporal consumption and portfolio decisions in an incomplete information economy.² To this end, I follow Honda (2003)

¹ David (1997) investigates unobservable and regime switching investment opportunities in continuous time. Lundtofte (2008) examines expected life-time utility and hedging demands when endowments and their expected growth rate are imperfectly correlated.

² Under incomplete information, Cagetti et al. (2002) use the robust control approach and a hidden Markov model

and postulate that expected returns of a risky asset are unobservable and follow a hidden Markov chain. For the sake of analytical convenience, I assume that the hidden Markov chain has two different regimes.³ Investors update beliefs about the unobservable state according to the Bayes rule. Different from the literature assuming expected utility, I employ Chen and Epstein's (2002) recursive multiple priors utility (hereafter RMPU) to account for ambiguity.⁴ The present model therefore nests the expected utility model of Honda (2003) as a special case where there is no ambiguity. Ambiguity and ambiguity aversion are concepts used to describe a decision maker who is uncertain over which probability law describes the dynamics of state variables and is also averse to such uncertainty. In the present model, investors with RMPU endogenously choose the worstcase prior among a prescribed set of different priors inducing different posteriors. Investors take into account not only *incomplete information risk* resulting from time-varying precision of beliefs but also ambiguity about the probability law governing the dynamics of beliefs.⁵ In contrast to the i.i.d. case, the effect of ambiguity varies over time as investors engage in continuous Bayesian revisions under incomplete information. This way of modeling ambiguity builds on the work of Miao (2009) where intertemporal consumption and portfolio decisions are examined without considering specific investment opportunity sets.

I use the Malliavin calculus technique and Clark-Ocone formula to explicitly characterize the optimal consumption and portfolio policies in terms of the Malliavin derivatives and stochastic integrals. My solutions are based on the martingale method of Cox and Huang (1989).⁶ I calibrate the model to historical U.S. stock market data. Numerical calculations of the optimal portfolios are implemented through the Monte Carlo Malliavin derivative (MCMD) method developed by Detemple et al. (2003). Similar to others (e.g., Maenhout, 2004, 2006) I find that ambiguity lowers the total stock demand in all states of the economy. Moreover, under incomplete information, continuous Bayesian revisions interact with time-invariant ambiguity aversion to yield an ambiguity-driven hedging component that is state- and horizon- dependent. This component mitigates the hedging demand for stocks while magnifies the relative importance of

to examine asset pricing implications of ambiguity.

³ Guidolin and Timmermann (2007) analyze asset allocation decisions under multivariate regime switching asset returns. They use a four-regime model to characterize the joint distribution of stock and bond returns. Extending the present paper to the multivariate case would be interesting and is left for the future research.

⁴ See Gilboa and Schmeidler (1989) and Epstein and Schneider (2003) for axiomatic foundations for multiple priors utility and recursive multiple priors utility.

 $^{^5}$ I thank David Feldman for suggesting the terminology "incomplete information risk".

⁶ Karatzas and Xue (1991) use the martingale method to derive the optimal consumption and portfolio choice under incomplete information. Dybvig, Rogers and Back (1999) consider the application of the method to time-varying investment opportunities.

hedging demand in the optimal stock demand. Contrary to the well-known advice that investors with long horizons should invest aggressively in stocks, I find that the optimal stock demand is decreasing in the horizon, even when ambiguity is taken into consideration. Ambiguity also affects the economic value of learning about the unobservable state as opposed to assuming i.i.d. returns. Because investors worry that the model driving the dynamics of state beliefs is subject to model uncertainty, the economic value of accounting for regimes and filtering declines under ambiguity. For highly ambiguity-averse investors with very long horizons (e.g., 20 years), there is almost no economic value of accounting for regimes and filtering versus assuming i.i.d. returns.⁷ This is in contrast to the previous findings based on expected utility (Xia, 2001; Guidolin and Timmermann, 2007) that portfolio strategies taking into account unobservable states and filtering yield higher utility gains than the i.i.d strategy. Finally, out-of-sample experiments using the CRSP (Center for Research in Security Prices) data from 1996 to 2009 demonstrate that portfolio strategies accounted for ambiguity are superior to those ignoring ambiguity for investors with long horizons.

This paper differs from recent works examining implications of learning under ambiguity (Epstein and Schneider, 2007; Leippold, et al., 2008). These papers assume that information on the fundamental process is ambiguous and analyze the belief updating mechanism with multiple priors and likelihoods. Here, investors treat the model driving the dynamics of the filtered probabilities as ambiguous and have multiple beliefs with respect to the Bayesian estimated model. Schroder and Skiadas (2003) examine intertemporal consumption and portfolio policies for generalized recursive utility preferences that incorporate RMPU as a special case. They show that the optimal consumption and portfolio policies can be characterized up to the solution to a single constrained backward stochastic differential equation (BSDE). But they did not consider the role of incomplete information. Sbuelz and Trojani (2008) examine asset prices in a continuous-time exchange equilibrium with locally-constrained-entropy RMPU (LCE-RMPU). They exogenously posit that the local bound on the size of ambiguity is some function of time-varying state variables. Without deriving explicit solutions, they identify that the impacts of ambiguity on the optimal portfolio strategy are state-dependent in a non-standard way. Here, I show that even with a constant local bound on the size of ambiguity, the model can still generate state-dependence of the effects of ambiguity through endogenous Bayesian revisions. Furthermore, I explicitly show

⁷ I calibrate ambiguity (or ambiguity aversion) based on detection error probabilities. A high level of ambiguity results in a small detection error probability. See Section 3.2 for details.

the quantitative results on the state-dependence.

The rest of this paper is organized as follows. Section 2 presents the model and derives the optimal consumption and portfolio policies using the martingale method. Section 3 calibrates the model to U.S. stock market data and discusses the results. Section 4 concludes. Proofs are collected in the Appendix.

2 The Model

In this section, I first present a hidden Markov model with an unobservable state governing regimes of the expected return of a risky asset. A small investor filter his estimates of the unobservable state by observing past and current asset prices. In the absence of ambiguity, the investor relies on a non-linear recursive filter to extract the filtered probabilities that are updated according to the Bayes rule. This non-linear filter represents a reference model describing the dynamics of state beliefs. To model ambiguity, I use the κ -ignorance specification in Chen and Epstein (2002) and consider a constrained set of alternative models surrounding the reference model. This prescribed set of priors reflects the investor's lack of confidence in the reference model. I then describe the dynamic optimization problem under incomplete information and ambiguity aversion. Finally, I derive the optimal consumption and portfolio policies using the martingale method.

2.1 The investment opportunity set

I assume that there are two assets available for investment, a riskless short-term bond paying an instantaneous return r and a risky asset with the following price dynamics:

$$dS_t = S_t \mu_t dt + S_t \sigma_S dB_t.$$

The expected return μ_t follows a continuous-time Markov chain with two regimes μ_H and μ_L , where $\mu_H > \mu_L$. The infinitesimal generating matrix of the hidden Markov chain is

$$\Lambda \equiv \left(\begin{array}{cc} -\lambda & \lambda \\ \chi & -\chi \end{array} \right)$$

with $\lambda, \chi > 0$. The parameters $\mu_H, \mu_L, \lambda, \chi$ and σ_S are known constants. The transition probabilities over any interval of time can be constructed from the infinitesimal matrix via the exponential formula $T_s = \exp(s\Lambda)$.

I assume that at each point of time the investor can observe neither the expected return μ_t nor the Brownian motion B_t . Instead, he can only observe asset prices. Given an initial prior for the two regimes, the investor estimates the unobservable state, i.e., the probability of the current regime being μ_H , based on observed asset prices. As in previous works (e.g., Dothan and Feldman, 1986; Detemple, 1986; Feldman, 1989), I identify a σ -algebra equivalent economy where the state variable is the filtered probability of the unobservable state. This economy is a Markovian representation of the original economy.⁸ The optimal consumption and portfolio policies obtained in the Markovian-equivalent economy are also optimal in the original economy.

Define the filtered probability π_t as the posterior probability that the current regime is μ_H :

$$\pi_t = \Pr\left(\mu_t = \mu_H \mid \mathcal{F}_t^S\right)$$

with π_0 given, where $\{\mathcal{F}_t^S\}$ is the information filtration generated by the asset price process. It follows from Theorem 9.1 in Liptser and Shiryaev (2001) that π_t satisfies the stochastic differential equation (SDE)

$$d\pi_t = \left[\lambda - (\lambda + \chi) \,\pi_t\right] dt + \pi_t \left(1 - \pi_t\right) \frac{\mu_H - \mu_L}{\sigma_S} d\hat{B}_t \tag{1}$$

where \hat{B} is a standard Brownian motion with respect to the price filtration $\{\mathcal{F}_t^S\}$ and is defined by

$$\hat{B}_t = \int_0^t \frac{dS_\tau - S_\tau \hat{\mu}_\tau d\tau}{S_\tau \sigma_S}$$

The conditional expected return with respect to the price filtration, $\hat{\mu}_t$, is given by

$$\hat{\mu}_t = \mu_H \pi_t + \mu_L \left(1 - \pi_t \right).$$
(2)

In SDE (??), the precision of the conditional estimates depends on the filtered probability and is therefore stochastic. To explore the dynamics of the filtered probabilities driven by SDE (??), I

⁸ See Feldman (2007) for detailed discussions on the irrelevance of Separation Principle for solving optimization problems in dynamic incomplete information economies. See Detemple and Murthy (1994) for an early discussion in a general context.

rewrite (??) in terms of the original Brownian motion.⁹ Suppose the investor knows that during some time interval $[t_1, t_2]$, the true expected return is μ_t . It follows that π_t satisfies the following SDE

$$d\pi_t = \left[\lambda - (\lambda + \chi)\pi_t + \frac{\pi_t \left(\mu_H - \hat{\mu}_t\right)\left(\mu_t - \hat{\mu}_t\right)}{\sigma_S^2}\right]dt + \frac{\pi_t \left(\mu_H - \hat{\mu}_t\right)}{\sigma_S}dB.$$
(3)

If the true expected return is μ_H during the time interval, then the second term in the drift of (??) tends to trend π_t toward 1. When π_t is close to 1, both the second term in the drift and the diffusion term converge to zero. This makes the first term in the drift become dominant, and thus π_t is driven toward 0. Conversely, if the true expected return is μ_L during the time interval, the second term in the drift is negative and tends to trend π_t toward 0. When π_t is close to 0, the first term in the drift becomes significant and drives π_t toward 1. Thus, the drift term in (??) has the effect of mean reversion. In the special case of non-switching regimes ($\lambda = \chi = 0$), the mean-reverting effect vanishes and the precision of the filtered probabilities increases as returns are observed over time.

The diffusion term in SDE (??) quantifies incomplete information risk that stems from time variation in the precision of the filtered probabilities. When π_t takes values near 1/2, the investor is uncertain about the current estimate, which leads to a large instantaneous standard deviation of π_t and a large scope of incomplete information risk. In updating beliefs, the investor therefore relies heavily on innovations to returns. On the contrary, if the investor is fairly certain about the current estimate, that is, when π_t takes values near the boundaries of the interval [0, 1], the magnitude of incomplete information risk is small, and innovations to returns have little effect on revisions in beliefs.

2.2 Ambiguity and recursive multiple priors utility

2.2.1 The set of alternative models

Denote the planning horizon as T, the wealth process as $\{W_t\}_{t=0}^T$ and the consumption process as $\{C_t\}_{t=0}^T$. Given an initial wealth endowment $W_0 > 0$, the wealth dynamics can be described by the following SDE:

$$dW_t = [W_t \left(r + \alpha_t \left(\hat{\mu}_t - r\right)\right) - C_t] dt + W_t \alpha_t \sigma_S d\hat{B}_t \tag{4}$$

⁹ See David (1997) and Honda (2003) for more detailed discussions on the properties of this non-linear filter.

where α_t is the proportion of wealth invested in the risky asset. The investor makes consumption and portfolio decisions based on the filtered probabilities of the unobservable state. Thus far, the original incomplete information economy has been converted into a Markovian-equivalent economy with the state processes (??) and (??).

Define the state vector $Y_t \equiv (W_t, \pi_t)$. The reference model can be written as

$$dY_t = \boldsymbol{\mu} \left(Y_t \right) dt + \boldsymbol{\sigma} \left(Y_t \right) d\hat{B}_t$$
(5)

where μ and σ have the appropriate functional forms inherited from SDE (??) and (??). The reference model serves as a benchmark among all the candidate models that an ambiguity-averse investor is willing to consider. The investor doubts the full usefulness of the reference model and only deems it as an approximation of the dynamics of his beliefs. He therefore considers a constrained set of alternative models that are close to the reference model.

The set of alternative models on which RMPU is defined is constructed from $\{\mathcal{F}_t^S\}_{t=0}^T$ -adapted density generators defined by $\theta \equiv \{\theta_t\}_{t=0}^T \in \Theta$ satisfying $\sup |\theta_t| \leq \kappa$ with $\kappa \geq 0$. This specification is referred to as κ -ignorance in Chen and Epstein (2002). Each density generator θ delivers a local distortion to the reference model. Suppose \mathbb{P} is the subjective probability measure induced by the reference model, i.e., the probability measure with respect to the Brownian motion \hat{B} . Each density generator θ generates a martingale z^{θ} under \mathbb{P} :

$$z_t^{\theta} = \exp\left(-\frac{1}{2}\int_0^t \theta_s^2 ds - \int_0^t \theta_s d\hat{B}_s\right), \quad 0 \le t \le T.$$

The set of priors is denoted as \mathcal{P} and is defined in terms of the Radon-Nikodym derivatives of the alternative models with respect to the reference model:

$$\mathcal{P} \equiv \left\{ \mathbb{Q}^{\theta} : \theta \in \Theta, \frac{d\mathbb{Q}^{\theta}}{d\mathbb{P}} = z_T^{\theta} \right\}.$$

It is obvious that the size of \mathcal{P} increases with κ . A large κ means that the investor has little confidence in the reference model and is willing to consider a wide range of alternative models. In the special case $\kappa = 0$, all alternative models coincide with the reference model, and the set of priors collapses to a singleton $\{\mathbb{P}\}$. In this case, the investor has complete confidence in the reference model.

It follows from Girsanov's theorem that the distorted law of motion of the state vector Y implied by an alternative model \mathbb{Q}^{θ} is

$$dY_t = \boldsymbol{\mu}\left(Y_t\right)dt + \boldsymbol{\sigma}\left(Y_t\right)\left(d\hat{B}_t^{\mathbb{Q}^{\theta}} - \theta_t dt\right)$$
(6)

where $\hat{B}^{\mathbb{Q}^{\theta}}$ is a Brownian motion under \mathbb{Q}^{θ} . As a result, the alternative models are absolutely continuous with respect to the reference model. An ambiguity-averse investor is uncertain about whether \hat{B} is a Brownian motion with respect to his information filtration, and in this case ambiguity indeed concerns uncertainty induced by the drift distortion in the state processes. In particular, under \mathbb{Q}^{θ} , the distorted law of motion of π_t can be explicitly written as

$$d\pi_t = \left[\lambda_0 - \left(\lambda_0 + \lambda_1\right)\pi_t\right]dt + \pi_t \left(1 - \pi_t\right)\frac{\mu_H - \mu_L}{\sigma_S} \left(d\hat{B}_t^{\mathbb{Q}^{\theta}} - \theta_t dt\right).$$
(7)

It is worth noting that the drift distortion in (??) is a product of the term quantifying incomplete information risk and the size of ambiguity. If the precision of the filtered probabilities is low, which often occurs when estimating expected returns is difficult, the magnitude of the drift distortion is large, and thus, the investor has low confidence in the reference model. As shown below, this drift distortion generates an ambiguity-driven hedging demand. Under time-invariant ambiguity, that is, $\theta_t = \kappa$ for all $t \in [0, T]$, SDE (??) becomes

$$dY_{t} = \boldsymbol{\mu}\left(Y_{t}\right)dt + \boldsymbol{\sigma}\left(Y_{t}\right)\left(d\hat{B}_{t}^{\mathbb{Q}} - \kappa dt\right)$$

where \mathbb{Q} denotes \mathbb{Q}^{θ} with $\theta = \kappa$. In Section 2.3, I show that time-invariant ambiguity can be supported in the optimum.

2.2.2 The RMPU preferences

The investor has a time preference rate ρ and a constant relative risk aversion (CRRA) utility function of the following form

$$u\left(C\right) = \begin{cases} \frac{C^{1-\gamma}}{1-\gamma}, & \gamma > 0, \neq 1\\ & \log C, & \gamma = 1 \end{cases}$$

With RMPU, the investor maximizes expected utility under the worst-case alternative model. Among all the candidate models induced by the allowed set of priors, the worst-case model delivers the minimum of expected utility for a consumption process and terminal wealth. Put formally, the value function is

$$V_0(C, W_T) = \max_{C_t, \alpha_t \mathbb{Q}^{\theta} \in \mathcal{P}} \mathbb{E}^{\mathbb{Q}^{\theta}} \left[\int_0^T e^{-\rho t} u(C_t) dt + e^{-\rho T} u(W_T) \right]$$
(8)

subject to the state dynamics (??) and (??), or compactly, the reference model (??). The minimization operator in (??) captures the concern that an alternative model may adversely affect the continuation value. The worst-case model is endogenously selected from the set of priors \mathcal{P} that contains a prescribed family of alternative models generated from locally distorting the reference model. The multiplicity of \mathcal{P} captures the investor's lack of confidence, and the minimization operator reflects the aversion to such ambiguity. Since each probability measure \mathbb{Q}^{θ} is associated with a density generator θ , the minimization operator in (??) is taken with respect to the process θ .

Chen and Epstein (2002) derive the following backward stochastic differential equation (BSDE) representation of the utility process V_t

$$dV_{t} = \left[-u\left(C_{t}\right) + \rho V_{t} + \max_{\theta \in \Theta} \theta_{t} \sigma_{t}^{V}\right] dt + \sigma_{t}^{V} d\hat{B}_{t}, \quad V_{T} = u\left(W_{T}\right).$$

where the diffusion term σ_t^V is endogenous and is a part of the complete solution to the above BSDE.¹⁰ For κ -ignorance, Chen and Epstein (2002) show that the endogenous density generator is given by

$$\max_{\theta \in \Theta} \theta_t \sigma_t^V = \theta_t^* \sigma_t^V = \kappa |\sigma_t^V|, \text{ with } \theta_t^* = \kappa \times \operatorname{sgn}\left(\sigma_t^V\right)$$
(9)

where $\operatorname{sgn}(\sigma_t^V) = |\sigma_t^V|/\sigma_t^V$ if $\sigma_t^V \neq 0$ and $\operatorname{sgn}(\sigma_t^V) = 0$ otherwise. The term $\kappa |\sigma_t^V|$ is entirely attributed to ambiguity aversion rather than risk aversion. The parameter κ can also be interpreted as an ambiguity aversion parameter. It is worth noting that the worst-case probability law associated with the density generator θ^* depends on the sign of the diffusion part of the utility process. The diffusion term of V_t , as a part of the complete solution to (??), relies on the optimal

¹⁰ Chen and Epstein (2002) show that the set of priors formulated by κ -ignorance satisfies a property called "rectangularity". Rectangularity and other technical conditions ensure that the utility process under multiple priors is recursive, and hence, time consistent in the sense of Johnsen and Donaldson (1985).

consumption path, which is in turn a function of the endogenous density generator θ^* (see the proof of Proposition 1). Thus, unlike other studies (e.g., Trojani and Vanini, 2002; Gagliardini et al., 2008) where one can explicitly derive the worst-case drift distortion and then tackle the maximization problem under the resulting worst-case model, here solutions to the optimal consumption and portfolio policies together with the endogenous worst-case model must be jointly determined from solving the optimization problem. In Section 2.3, I employ the guess-and-verify method to show that $\sigma_t^V > 0$ and $\theta_t^* = \kappa$ for all $t \in [0, T]$ can be supported together in the optimum. Thus, the size of ambiguity is time-invariant throughout the horizon and the optimal consumption and portfolio policies can be derived conditioning on $\theta^* = \kappa$.

2.3 Optimal consumption and portfolio choice

Define the conditional market price of risk $\hat{\nu}_t$ as

$$\hat{\nu}_t = \frac{\hat{\mu}_t - r}{\sigma_S}.$$

Since π_t is progressively measurable and always bounded between 0 and 1, the conditional market price of risk is also progressively measurable and bounded. Because $\hat{\nu}$ is a bounded process, Novikov's condition holds; that is, $\mathbb{E}^{\mathbb{P}}\left(\exp\left\{\frac{1}{2}\int_{0}^{T}\hat{\nu}_{t}^{2}dt\right\}\right) < \infty$. In addition, the market is complete in the Markovian-equivalent economy. Thus, there exists a unique equivalent martingale measure $\tilde{\mathbb{P}}$ given by $d\tilde{\mathbb{P}} = \zeta_{T}d\mathbb{P}$ where

$$\zeta_t = \exp\left(-\frac{1}{2}\int_0^t \hat{\nu}_s^2 ds - \int_0^t \hat{\nu}_s d\hat{B}_s\right)$$

with $\zeta_0 = 1$. Under $\tilde{\mathbb{P}}$, the process $\tilde{B}_t = \hat{B}_t + \int_0^t \hat{\nu}_s ds$ is a Brownian motion. The state price density, denoted as ξ , is given by

$$\xi_t = \exp\left(-\int_0^t r ds - \frac{1}{2}\int_0^t \hat{\nu}_s^2 ds - \int_0^t \hat{\nu}_s d\hat{B}_s\right).$$

Boundedness of the conditional market price of risk implies that all moments of the state price density are finite, which is a sufficient condition ensuring the applicability of the martingale method for time-varying investment opportunities (see Dybvig, Rogers and Back, 1999). This condition is referred to as the "DRB condition" in Korn and Kraft (2004). As a result, the

optimization problem is well-defined for general coefficients of relative risk aversion, and the pitfalls raised by Korn and Kraft (2004) can be naturally avoided.

Hereafter, I assume the following condition holds for the primitive parameters describing investment opportunities and preferences

$$\frac{\mu_L - r}{\sigma_S} - \kappa \ge 0. \tag{10}$$

This assumption guarantees that in all states of the economy, the conditional market price of risk accounted for ambiguity, or *effective* conditional market price of risk, is strictly positive unless π_t exactly reaches 0. This assumption plays a key role in proving that time-invariant ambiguity can be supported in the optimum (see the proof of Proposition 1).

The static variational problem corresponding to the optimization problem (??) is formulated by standard arguments as follows:

$$\max_{C_t, W_T \mathbb{Q}^{\theta} \in \mathcal{P}} \mathbb{E}^{\mathbb{Q}^{\theta}} \left[\int_0^T e^{-\rho t} u\left(C_t\right) dt + e^{-\rho T} u\left(W_T\right) \right]$$
(11)

s.t.
$$W_0 = \mathbb{E}^{\mathbb{P}}\left[\int_0^T \xi_t C_t dt + \xi_T W_T\right].$$
 (12)

Suppose \mathbb{Q}^{θ^*} solves the inner minimization problem. The Lagrangian of $(\ref{eq:problem})$ subject to $(\ref{eq:problem})$ is given by

$$\mathcal{L} = \mathbb{E}^{\mathbb{Q}^{\theta^*}} \left[\int_0^T e^{-\rho t} u\left(C_t\right) dt + e^{-\rho T} u\left(W_T\right) \right] + y \left\{ W_0 - \mathbb{E}^{\mathbb{P}} \left[\int_0^T \xi_t C_t dt + \xi_T W_T \right] \right\}$$
(13)

where θ^* is given in (??) and y is a scalar Lagrange multiplier. Different from the standard martingale formulation under expected utility, the endogenous probability law under which the expectation of discounted future utility is taken deviates from the probability law governing the static budget constraint. This deviation captures the investor's distrust in the reference model that delivers a particular consumption process. To solve (??), a change of measure is applied to convert the worst-case probability measure \mathbb{Q}^{θ^*} to the reference measure \mathbb{P} using the Radon-Nikodym derivative between the two measures. The first-order conditions (shown in the Appendix) are derived under the reference measure. The solutions to the optimal consumption and portfolio policies together with the endogenous probability law \mathbb{Q}^{θ^*} are presented in the

following proposition.

Proposition 1 Suppose $\gamma \neq 1$ and the following condition holds¹¹:

$$\mathbb{E}_{t}^{\mathbb{P}}\left[\int_{t}^{T} e^{-rs} C_{s}^{*} \int_{t}^{s} \left(D_{t} \hat{\nu}_{\tau} d\tilde{B}_{\tau}\right) ds + e^{-rT} W_{T}^{*} \int_{t}^{T} D_{t} \hat{\nu}_{s} d\tilde{B}_{s}\right] > 0 \quad for \ all \ t \in [0, T]$$
(14)

where C_s^* and W_T^* are given in (ii). The following solutions characterize an optimum:

- (i) The endogenous probability law is given by \mathbb{Q}^{θ^*} with $\theta_t^* = \kappa$ for all $t \in [0,T]$
- (ii) The optimal consumption C_t^* and terminal wealth W_T^* are

$$C_t^* = \left(\frac{e^{-\rho t} z_t^{\kappa}}{y\xi_t}\right)^{\frac{1}{\gamma}} \quad and \ W_T^* = \left(\frac{e^{-\rho T} z_T^{\kappa}}{y\xi_T}\right)^{\frac{1}{\gamma}} \tag{15}$$

where the Lagrange multiplier y satisfies

$$y = \left(\mathbb{E}^{\mathbb{P}}\left[\int_{0}^{T} (\xi_{t})^{\frac{\gamma-1}{\gamma}} (e^{-\rho t} z_{t}^{\kappa})^{\frac{1}{\gamma}} dt + (\xi_{T})^{\frac{\gamma-1}{\gamma}} (e^{-\rho T} z_{T}^{\kappa})^{\frac{1}{\gamma}}\right] / W_{0}\right)^{\gamma}$$
(16)

(iii) The optimal portfolio policy, α_t^* , is

$$\alpha_t^* = \frac{\hat{\mu}_t - r}{\gamma \sigma_S^2} - \frac{\kappa}{\gamma \sigma_S} + \frac{1 - \gamma}{\gamma} \frac{e^{rt}}{\sigma_S W_t^*} \mathbb{E}_t^{\tilde{\mathbb{P}}} \left[\int_t^T e^{-rs} C_s^* \int_t^s \left(D_t \hat{\nu}_\tau d\tilde{B}_\tau \right) ds + e^{-rT} W_T^* \int_t^T D_t \hat{\nu}_s d\tilde{B}_s \right]$$
(17)

where the conditional expectation is taken under the equivalent martingale measure $\tilde{\mathbb{P}}$. The Malliavin derivative of the conditional market price of risk is given by $D_t \hat{\nu}_s = \frac{(\mu_H - \mu_L)}{\sigma_S} D_t \pi_s$ where the Malliavin derivative of π , $D_t \pi_s$, satisfies the following SDE:

$$d(D_t\pi_s) = -(\lambda_0 + \lambda_1)D_t\pi_s ds + \frac{\mu_H - \mu_L}{\sigma}(1 - 2\pi_s)D_t\pi_s d\hat{B}_s$$

subject to the boundary condition $\lim_{s\to t} D_t \pi_s = \pi_t (1 - \pi_t) \frac{\mu_H - \mu_L}{\sigma_S}$.

Proof. See Appendix.

Corollary 1 Suppose $\pi_0 > 0$. The solutions in Proposition 1 nest the case of logarithmic utility with $\gamma = 1$.

¹¹ Ideally, I would like to be able to show analytically that this condition holds for the solution in Proposition 1 (ii). Unfortunately, I have been unable to do so, because the condition involves the Malliavin derivatives and stochastic integrals, which must be computed numerically. Nevertheless, for all calibration exercises in this paper, I have verified that this condition does hold on a fine grid of the state probability $\pi \in [0, 1]$ and the horizon $t \in [0, T]$ with T = 40 years.

Proof. See Appendix.

Under (??) and (??), I can show that the time-invariant density generator $\theta^* = \kappa$ delivers the endogenous worst-case alternative model (see Appendix). As a result, the investor's attitude toward multiple priors does not depend on the state of the economy.¹² This simplified feature of the model provides analytical convenience, namely, delivers explicit solutions to the optimal consumption and portfolio policies.

In the optimal portfolio formula (??), the first term quantifies the myopic demand for the risky asset, which is instantaneously mean-variance efficient and only depends on the current estimate of the unobservable state. Here, myopia represents "generalized myopia" defined in Feldman (1992).¹³ Under incomplete information, investors with logarithmic preferences update their portfolios as functions of the filtered probability (see Corollary 1). The second term reflects the effect of ambiguity on myopic demand. Together, the first two terms are called *ambiguity-adjusted myopic demand*. Under the i.i.d. assumption, Chen and Epstein (2002) show that the optimal portfolio formula is given by the first two terms where $\hat{\mu}_t$ is replaced by the known and constant expected return. Their results hence give rise to a form of observational equivalence; i.e., concerning the effects on the optimal consumption and portfolio choice, an increase in the size of ambiguity is observationally equivalent to a decline in the effective market price of risk with the same magnitude. Under incomplete information, this form of observational equivalence cannot sustain because ambiguity also affects the intertemporal hedging demand in a non-standard way.

The third term quantifies hedging demand, which is induced to hedge against the future time variation in the conditional estimates of the unobservable state. In the optimal portfolio formula, the Malliavin derivative $D_t \pi_s$ captures the effect of an innovation in the Brownian motion \hat{B} at time t on the state variable π at time s. A notable difference between the solution derived here and those derived in the expected utility framework in the absence of ambiguity (Brennan, 1998; Honda, 2003) is that hedging demand is driven not only by incomplete information risk

¹² One can imagine a case where the density generator θ^* switches between the two different regimes (κ and $-\kappa$) whenever the diffusion term of the utility process changes its sign, as implied by (??). However, the case of regimeswitching density generator is intractable for two reasons. First, if the worst-case model switches between different regimes, a state vector solution to the optimal control would be difficult to obtain. Second, the diffusion part of the utility process, which determines the instantaneous regime of the density generator, is endogenous and depends on the decision variables when investment opportunities are time-varying. This adds enormous difficulty to solving the optimization problem in the presence of jumps in the density generator.

¹³ In incomplete information economies, logarithmic investors have the same informational requirements as non-logarithmic investors do, due to the need of estimating unobservable state variables. Feldman (1992) encompasses in the definition of generalized myopia the concept that investors are myopic in an incomplete information economy if and only if they are myopic in the Markovian-equivalent complete information economy.

but also by ambiguity. In particular, ambiguity affects hedging demand through intermediate consumption decisions, as is obvious from (??) and (??). The hedging term in (??) can be further decomposed into two components, $hedge^{IIR}$ and $hedge^{ambiguity}$, and the optimal portfolio formula can be rewritten as

$$\alpha_t^* = \frac{\hat{\mu}_t - r}{\gamma \sigma_S^2} - \frac{\kappa}{\gamma \sigma_S} + hedge^{IIR} + hedge^{ambiguity}$$

where $hedge^{IIR}$ is obtained from the third term in (??) by setting $\kappa = 0$, and $hedge^{ambiguity}$ accounts for the difference between $hedge^{IIR}$ and hedging demand. As a result, $hedge^{IIR}$ is solely attributed to the intertemporal hedging of incomplete information risk, while $hedge^{ambiguity}$ is purely driven by ambiguity. The component $hedge^{IIR}$ exists under expected utility, but $hedge^{ambiguity}$ does not. Although the effect of ambiguity on myopic demand is state-independent, $hedge^{ambiguity}$ is state-dependent. The following section shows quantitative results on the state-dependence.

3 Calibration and Results

3.1 Data, estimation and MCMD

In the calibration exercise, the U.S. stock market is a proxy for the risky asset. The parameters of the stochastic process for stock returns and the two-regime hidden Markov model are estimated using quarterly U.S. stock market return data for the sample period 1947.1–1995.4. Data from 1996.1 to 2009.4 are taken as a genuine post-sample and used for out-of-sample experiments. All data are from Center for Research in Security Prices. Nominal returns are deflated to obtain real returns using the CPI data from Federal Reserve Economic Data St. Louis. The real risk-free rate is set at a level comparable to that in Campbell and Viceira (1999). I jointly estimate the discrete-time analogue of the return process and hidden Markov model using the EM algorithm of Hamilton (1989). Here, I set the low-regime mean return in the discrete-time model at such a level that Inequality (??) holds for a sizable κ . The allowed values for κ are empirically plausible as the corresponding detection error probabilities are economically meaningful, which implies that the scope of ambiguity is empirically significant.¹⁴ Then I estimate a Markov switching model to obtain maximum-likelihood estimates. The estimates are annualized to yield appropriate continuous-time estimates. I employ the following formula in Israel et al. (2001) to compute the

 $^{^{14}\,\}mathrm{See}$ Section 3.2 and 3.3 for further details.

infinitesimal generating matrix Λ given the estimated transition probability matrix

$$\Lambda = \sum_{i=1}^{\infty} \frac{(-1)^{i+1}}{i} (P - I)^{i}$$

where P is the transition probability matrix, I is the identity matrix, and the series is approximated with length 5 in the calibration exercise. Table 1 summarizes the estimation results.

[Insert Table 1 here]

I use the Monte Carlo Malliavin derivative (MCMD) method developed by Detemple, et al. (2003) to compute optimal portfolio policies. As shown by Detemple et al. (2003), simulating optimal portfolios in explicit form can improve efficiency and accuracy of numerical approximations relative to the standard stochastic programming with backward iterations. I run a random number generator to simulate a large number of sample paths and numerically evaluate the stochastic integrals in the optimal portfolio policy rule. The state probability π and its Malliavin derivative are simulated using a variance-stabilizing transformation to minimize approximation error, as suggested by Detemple et al. (2003). The reformulation allows us to adopt a change of variables that normalizes the volatility of the filtered probability process to a constant. This makes the calculation of the Malliavin derivative $D_t \pi_s$ be of the same complexity as the numerical solution of an ordinary differential equation.¹⁵ The number of Monte Carlo replications is set at N = 20,000.

3.2 Detection error probabilities

I use the detection error probabilities technique developed by Anderson et al. (2003) to calibrate the size of ambiguity. Anderson et al. (2003) relate the calibration of ambiguity to a model selection problem in a Bayesian context. Maenhout (2004) and Maenhout (2006) examine detection error probabilities, respectively, under i.i.d. returns and mean-reverting returns. When choosing between two potential data generating processes, a decision maker performs likelihood ratio tests for a given data set. The two models are difficult to distinguish from each other if the probability of mistakenly rejecting one model in favor of the other is high. This probability is given by the probability that the log-likelihood ratio is negative when the rejected model is

¹⁵ See Detemple et al. (2003) for further details. Detemple et al. (2005) extend the variance-stabilizing transformation to cases with multiple state variables and assets.

the true data generating process. In continuous time, the log-likelihood ratio is given by the log of the Radon-Nikodym derivative. Specifically, the log of the Radon-Nikodym derivative of the distorted probability measure Q with respect to the reference measure P is

$$\eta_{1,t} \equiv \log(\frac{dQ}{dP} \mid \mathcal{F}_t^S) = -\int_0^t \theta_s^* d\hat{B}_s - \frac{1}{2} \int_0^t (\theta_s^*)^2 ds$$

The log of the Radon-Nikodym derivative of the probability measure P with respect to Q is

$$\eta_{2,t} \equiv \log(\frac{dP}{dQ} \mid \mathcal{F}_t^S) = \int_0^t \theta_s^* d\hat{B}_s + \frac{1}{2} \int_0^t (\theta_s^*)^2 ds.$$

If Model P is true, the decision maker will mistakenly reject it in favor of Model Q based on a finite sample with size N when $\eta_{1,N} > 0$. Conversely, if Q is correct, it will be rejected erroneously when $\eta_{2,N} > 0$. Assuming an initial prior of 0.5 on each model, the detection error probability $\varepsilon_N(\theta)$, which is the time-0 conditional probability of choosing a wrong model based on a finite sample of length N, is defined as

$$\varepsilon_N(\theta) = 0.5 \Pr(\eta_{1,N} > 0 \mid P) + 0.5 \Pr(\eta_{2,N} > 0 \mid Q).$$

As θ^* increases, the two models are easier to distinguished statistically from each other, and the detection error probability becomes lower as a result. For κ -ignorance, it is straightforward to derive

$$\varepsilon_N(\kappa) = \Pr\left(Z < -\frac{\kappa}{2}\sqrt{N}\right)$$

where Z is from the standard normal distribution. Anderson et al. (2003) advocate 10% as the lower bound of detection error probabilities.

3.3 Incomplete information risk, ambiguity and portfolio choice

I consider three alternative values for the ambiguity aversion parameter κ , $\kappa = 0$ (expected utility), 0.08 and 0.16. For the sample period 1947.1–1995.4, the detection error probabilities for $\kappa = 0.08$ and 0.16 are, respectively, 0.288 and 0.131, which are above the threshold advocated by Anderson et al. (2003). The coefficient of relative risk aversion, γ , is set at $\gamma = 4$. The investment horizon is 10 years unless explicitly stated. Figure 1 plots hedging demands as functions of the state probability when utility is defined over intermediate consumption and over terminal

wealth respectively. Under expected utility ($\kappa = 0$), the finding that hedging demands have a negative sign for $\gamma > 1$ is in line with the existing works (Brennan, 1998; Honda, 2003).¹⁶ For investors who are more risk averse than the logarithmic investor, their utility functions are unbounded from below but bounded from above. They therefore behave more conservatively than the logarithmic investor and are willing to hold a less risky portfolio than the myopic portfolio. Under incomplete information, continuous Bayesian revisions create a perfect positive correlation between innovations to returns and revisions in beliefs. Stocks tend to have low (high) returns when their mean returns are expected to be low (high). Because investors desire to smooth utility across states, a portfolio is deemed less risky if it can deliver wealth when investment opportunities are poor, that is, when mean return estimates are low. Thus, investors with $\gamma > 1$ optimally reduce the total stock demand relative to myopic demand. This implies that the intertemporal hedging demand has a negative sign. Moreover, similar to Honda (2003), I find that in the absence of ambiguity, hedging demand is high when π takes values near 1/2, that is, when incomplete information risk is sizable.

[Insert Figure 1 here]

Figure 2 shows that ambiguity mitigates hedging demand in all states of the economy and that this impact is state-dependent, where utility is assumed to be defined over intermediate consumption.¹⁷ The right panel in Figure 2 further separates hedging demand into the two hedging components, $hedge^{IIR}$ and $hedge^{ambiguity}$. The magnitude of $hedge^{ambiguity}$ depends on the ambiguity aversion parameter and incomplete information risk. A high ambiguity aversion parameter, κ , leads to a large $hedge^{ambiguity}$, all else being equal. A high instantaneous standard deviation of the filtered probabilities also results in a large $hedge^{ambiguity}$. If incomplete information risk shrinks toward a negligible level, $hedge^{ambiguity}$ vanishes. This state-dependence can be explained by the drift distortion in the worst-case alternative model (??). The drift distortion implied by the worst-case prior reinforces the effect of mean reversion in the filter (??) and drives the component $hedge^{ambiguity}$. Since hedging demand depends on the mean-reverting effect, the drift distortion in SDE (??) reduces hedging demand, as observed in Figure 2.¹⁸ Moreover, as the

¹⁶ I also compute hedging demands for various coefficients of relative risk aversion, assuming expected utility. Hedging demands have a negative sign for $\gamma > 1$ but a positive sign for $\gamma < 1$. For the sake of brevity, these results are not reported.

¹⁷ The results based on utility over terminal wealth are similar and thus not reported.

¹⁸I also perform the comparative statics experiments and find that mean reversion in the dynamics of state beliefs lowers hedging demand.

size of the drift distortion is large when π takes values near 1/2, hedge^{ambiguity} becomes large in magnitude. Figure 3 plots the optimal stock demand and the fraction of hedging demand in the optimal stock demand as functions of the state probability. Not surprisingly, ambiguity lowers the optimal stock demand in all states of the economy. In addition, ambiguity magnifies the relative importance of hedging demand in the optimal stock demand. This is because the effect of ambiguity on myopic demand is of the first order while its effect on hedging demand is of the second order.¹⁹ In particular, the impact of ambiguity on the ratio of hedging demand to the optimal stock demand turns out to be stronger when the current regime is more likely to be in the bad regime. As a result, ambiguity-averse investors who believe that the current state is poor prefer a conservative portfolio with its composition significantly steering toward hedging demand while away from myopic demand. Figure 4 shows the horizon effect on the optimal stock demand for a horizon ranging from 1 year to 20 years, where the state probability is set at its steady-state value. As the horizon increases, the optimal stock demand decreases because hedging demand is monotonically increasing in the horizon. This result holds true for all the values of κ being considered. Thus, the well-known advice that investors with long horizons should allocate a substantial proportion of wealth into stocks does not necessarily hold true in incomplete information economies.

> [Insert Figure 2 here] [Insert Figure 3 here] [Insert Figure 4 here]

3.4 Welfare implications

Two natural questions concerning the welfare implications of ambiguity are: (1) How economically important is it for an ambiguity-averse investor to learn about the unobservable state rather than treating investment opportunities as fully observable and i.i.d.? and (2) Are investors necessarily better off by accounting for ambiguity in portfolio choice from the out-of-sample perspective? To address the first question, I assess the economic importance of regimes under ambiguity aversion by computing the certainty equivalent wealth. I follow Xia (2001) and define the certainty equivalent wealth (CEW) of a portfolio s as the amount of wealth that makes an

¹⁹ Note that in the expression for the ambiguity-adjusted myopic demand, the ambiguity aversion parameter, κ , appears additively to the conditional market price of risk, which implies a first-order effect on myopic demand.

investor indifferent between receiving CEW with certainty at horizon T and having initial wealth \$1 to invest up to the horizon using strategy s:

$$e^{-\rho T} \frac{(CEW^s)^{1-\gamma}}{1-\gamma} = V_0(W_T \mid s) \quad W_0 = 1$$

where $V_0(W_T \mid s)$ is the maximized RMPU over terminal wealth when the optimal strategy s is adopted for a given investment opportunity set. The present value of CEW^s , $PCEW^s$, is defined as $PCEW^s = e^{-rT}CEW^s$. For instance, the symbol $PCEW^{IID}$ represents the present value of CEW assuming that asset returns follow a simple i.i.d. process and the corresponding optimal portfolio strategy is adopted. The economic value of regimes is quantified by taking the difference between $PCEW^{HMM}$ and $PCEW^{IID}$ for some ambiguity aversion parameter value, where $PCEW^{HMM}$ stands for PCEW under the hidden Markov model.

Figure 5 plots the percentage difference in *PCEWs* under the hidden Markov model versus i.i.d. returns as a function of the investment horizon ranging from 1 year to 20 years for the nonambiguity case ($\kappa = 0$) and two cases with ambiguity ($\kappa = 0.08$ and 0.16). The state probability π is set at its steady-state value. In the absence of ambiguity, the economic value of regimes is non-negligible and the welfare loss of ignoring regimes is increasing in the horizon. The utility cost measured by the difference in PCEWs reaches as high as 1.5% at long horizons. Nevertheless, the economic importance of regimes declines significantly when ambiguity is taken into account, and this effect is especially strong at long horizons. The decline in the economic value of regimes is due to the concern that the stochastic process driving the dynamics of state beliefs is subject to model uncertainty. This uncertainty reduces the scope of utility gains from accounting for regimes and filtering. For $\kappa = 0.16$, the economic value of regimes becomes even negative at a horizon as long as 20 years, suggesting that a highly ambiguity-averse investor may expect to be better off by choosing the i.i.d. strategy. Figure 6 displays the time series plot of *PCEWs* under different strategies for the sample period 1986.1–1995.4 and an investment horizon of 10 years. The state probability, starting from its steady-state value, is updated after each period's return is realized. The results indeed confirm the findings above. In the absence of ambiguity, the economic value of regimes is sizable: the *PCEW* under the hidden Markov model stays above the *PCEW* under i.i.d. returns in most cases during the sample period. Ambiguity, however, decreases the economic value of regimes by driving the *PCEW* under the hidden Markov model

toward the PCEW under i.i.d. returns.

To address the second question, I evaluate out-of-sample performance of different portfolio strategies for the genuine out-of-sample period 1996.1–2009.4 in a way similar to Guidolin and Timmermann (2007). However, unlike Guidolin and Timmermann (2007) who only consider the buy-and-hold strategy, I examine continuous-time strategies that take into account continuous updating and intermediate consumption.²⁰ I compare the performance of a two-regime hidden Markov model, a simple i.i.d. model, the hidden Markov model with ambiguity and the i.i.d. model with ambiguity. The models of asset returns will be re-estimated once the next period's return is realized, and the state probability will be updated based on the re-estimated parameter values. For instance, the hidden Markov model is initially estimated for 1947.1–1995.4 and the corresponding estimates and state probability (initially set at its steady-state value) are used to calculate the optimal stock allocation and portfolio return for 1996.1. Then the sample is extended to 1947.1–1996.1 and the model is re-estimated. The state probability is revised based on the estimates for 1947.1–1996.1 and portfolio optimization is repeated, and so forth. I consider three investment horizons, T = 1, 5 and 10 years.

I compute realized utility under different models to evaluate out-of-sample performance. The realized utility measure is defined as

$$U_t^T \equiv \frac{W_{t+T}^{1-\gamma}}{1-\gamma}, W_t = 1$$
$$W_{t+\tau} - W_{t+\tau-1} = \left[\left(1 - \psi_{t+\tau-1}^T \right) r_f + \psi_{t+\tau-1}^T r_{t+\tau-1,t+\tau} \right] W_{t+\tau-1}$$

where r_f is the quarterly risk-free rate, $\{r_{t+\tau-1,t+\tau}\}_{\tau=1}^T$ are the realized stock returns between t+1 and t+T and $\psi_{t+\tau}^T$ is the period- $(t+\tau)$ optimal stock allocation with the horizon T. The optimal portfolio weights are computed by solving the optimization problem (??) for a given κ and investment opportunity set. For each of the six models (or strategies, see Table 2) and each investment horizon, I obtain a time series $\{U_{\tau}^T\}$, $\tau = 1996.1, \dots, 2009.4 - T$ of realized utilities. Following Guidolin and Timmermann (2007), I use a block bootstrap with 50,000 independent trials to construct the empirical distribution of U_{τ}^T . Table 2 reports summary statistics for the empirical distribution of U_{τ}^T . Among the six portfolio strategies, the strategy $\{HMM, \kappa = 0.16\}$ produces the highest mean realized utility for all the three invest-

²⁰Out-of-sample experiments assuming no intermediate consumption produce similar findings.

ment horizons. Among all the horizons considered, the difference in the mean realized utilities between the strategy $\{HMM, \kappa = 0.16\}$ and other strategies is the highest for T = 10 years, indicating the importance of accounting for ambiguity in out-of-sample performance for investors with long horizons. The strategy $\{IID, \kappa = 0.16\}$ generates a slightly lower mean realized utility than $\{HMM, \kappa = 0.16\}$ for all the horizons, suggesting the importance of accounting for regimes in evaluating out-of-sample performance. Moreover, the strategy $\{HMM, \kappa = 0.16\}$ has the lowest standard deviations of realized utility levels for all the horizons, and its 5% and 10%confidence intervals of realized utility levels have the smallest size among all the strategies. These results indicate that ambiguity can reduce uncertainty in out-of-sample welfare levels. It is also worth noting that the 10% confidence intervals for the horizon T = 5 years do not even overlap under the strategy $\{HMM, \kappa = 0.16\}$ and $\{HMM, \kappa = 0\}$, which provides strong evidence to show that their performance are statistically distinguishable. Table 3 reports summary statistics for differences in realized utility levels between the benchmark strategy $\{HMM, \kappa = 0.16\}$ and all the other strategies. The results indicate that except for $\{IID, \kappa = 0.16\}$, the out-of-sample performance of $\{HMM, \kappa = 0.16\}$ is significantly better than other strategies for all the horizons. For the horizons T = 5 and 10 years, the out-of-sample performance of $\{HMM, \kappa = 0.16\}$ is sufficiently good to outperform $\{IID, \kappa = 0.16\}$.

4 Conclusion

I have examined a continuous-time intertemporal consumption and portfolio choice problem for an ambiguity-averse investor when expected returns of a risky asset are unobservable and follow a continuous-time Markov chain with two regimes. Investors lack confidence in the reference model that characterizes a Markovian equivalent representation of the original incomplete information economy. Facing an optimization problem with recursive multiple priors utility, investors consider the worst-case model among a family of alternative models surrounding the reference model. In calibrating the model to U.S. stock market data, I find that ambiguity with respect to the reference model lowers the optimal stock demand and hedging demand while magnifies the relative importance of hedging demand in the optimal portfolio policy. The economic value of regimes and filtering about the unobservable state declines under ambiguity. Finally, out-ofsample experiments confirm the economic importance of accounting for ambiguity in portfolio

decisions.

There are several ways to extend this paper. In a related work, Guidolin and Timmermann (2007) consider regime-switching predictability of both stock and bond returns from the dividend yield. One could extend their framework to incorporate ambiguity and ambiguity aversion. Such a new framework may deepen our understanding of the effects of ambiguity, multiple regimes and uncertain predictability on asset allocation decisions. Extensions along the lines of Ju and Miao (2010) and Chen et al. (2009) seem therefore very promising. In addition, one could also examine the effects of ambiguity on consumption and portfolio choice under incomplete information and recursive preferences, where the separation between risk aversion and the attitude Accepted manuscrit toward intertemporal substitution is allowed for.

Appendix

The Clark-Ocone Formula The space of random variables for which Malliavin derivatives are defined is called $\mathbb{D}^{1,2}$.²¹ Any random variable $F \in \mathbb{D}^{1,2}$ can be decomposed as

$$F = \mathbb{E}(F) + \int_{t}^{T} \mathbb{E}[D_{t}F \mid \mathcal{F}_{t}] dB_{t}$$

where \mathcal{F}_t represents the information filtration generated by the Brownian motion B up to time t.

Proof of Proposition 1 The proof employs the guess-and-verify method to show that in the optimum, the density generator θ is given by $\theta_t^* = \kappa$ for all $t \in [0, T]$ and to derive the optimal consumption and portfolio policies. First, I conjecture $\theta_t^* = \kappa$ for all $t \in [0, T]$ and proceed to derive the solution to the volatility of the utility process σ_t^V . Then I verify that $\sigma_t^V > 0$ holds for all $t \in [0, T]$, which is also a verification of $\theta_t^* = \kappa$ for all $t \in [0, T]$, in accordance with (??). The explicit solutions to the optimal consumption and portfolio choice conditioning on $\theta^* = \kappa$ follow naturally.

The first-order conditions of the Lagrangian problem (??) are

$$e^{-\rho t} z_t^{\kappa} (C_t^*)^{-\gamma} = y \xi_t \text{ and } e^{-\rho t} z_t^{\kappa} (W_T^*)^{-\gamma} = y \xi_T$$
 (18)

where z_t^{κ} is given by

$$z_t^{\kappa} = \exp\left(-\frac{1}{2}\int_0^t \kappa^2 ds - \int_0^t \kappa d\hat{B}_s\right).$$

The optimal consumption and terminal wealth in (??) can be obtained from the first-order conditions. Substituting (??) into the budget constraint (??) yields the Lagrange multiplier y in (??).

Applying Ito's lemma to the optimal consumption in (??) yields the following dynamics of the consumption process

$$\frac{dC_t^*}{C_t^*} = \mu_t^C dt + \sigma_t^C d\hat{B}_t$$

²¹Oksendal(1997) provides a concise introduction to Malliavin calculus. Interested readers can also refer to Nualart (1995) for a textbook treatment.

where

$$\begin{split} \mu^C_t &= \frac{1}{\gamma} \left(r - \rho \right) + \frac{1}{2} \left(1 + \gamma \right) \left(\sigma^C_t \right)^2 + \sigma^C_t \kappa \\ \sigma^C_t &= \frac{1}{\gamma} \left(\hat{\nu}_t - \kappa \right). \end{split}$$

By homogeneity, I conjecture that the utility process along the optimal path has the following form

$$V_t = \frac{\left(C_t^*\right)^{1-\gamma}}{1-\gamma} A_t$$

with $V_T = \frac{\left(W_T^*\right)^{1-\gamma}}{1-\gamma} A_T$, where A_t satisfies the BSDE:

satisfies the BSDE:

$$\frac{dA_t}{A_t} = \mu_t^A dt + \sigma_t^A d\hat{B}_t, \ A_T = 1.$$

$$V_t \text{ satisfies the BSDE:}$$

It follows from Ito's lemma that V_t satisfies the BSDE:

$$dV_t = \mu_t^V dt + \sigma_t^V d\hat{B}_t, \ V_T = \frac{(W_T^*)^{1-\gamma}}{1-\gamma}$$

with $\sigma_t^V = V_t \left[(1 - \gamma) \, \sigma_t^C + \sigma_t^A \right]$. The derivation of μ_t^V is irrelevant to the proof and thus omitted. To obtain σ_t^V and σ_t^A in explicit forms, I first derive the optimal portfolio policy α_t^* as a function of both σ_t^C and σ_t^A and then characterize α_t^* in terms of the Malliavin derivatives and stochastic integrals. The explicit expression of σ_t^A can be immediately obtained. Then, σ_t^V can be solved explicitly, and its sign can be readily determined.

Multiply both sides of the first-order condition for consumption in (??) by C_t^* and integrate over the product space $dt \otimes d\mathbb{P}$ to obtain

$$\mathbb{E}^{\mathbb{P}}\left[\int_{0}^{T} e^{-\rho t} \left(C_{t}^{*}\right)^{1-\gamma} z_{t}^{\kappa} dt\right] = y \mathbb{E}^{\mathbb{P}}\left[\int_{0}^{T} \xi_{t} C_{t}^{*} dt\right].$$

Similarly, one can obtain $\mathbb{E}^{\mathbb{P}}\left[e^{-\rho T} (W_T^*)^{1-\gamma} z_T^{\kappa}\right] = y\mathbb{E}^{\mathbb{P}}\left[\xi_T W_T^*\right]$. It follows from the complementary slackness condition (??), the equality $y = (C_0^*)^{-\gamma}$ and the definition of RMPU (??) that the following equality holds

$$W_0 = (1 - \gamma) (C_0^*)^{-\gamma} V_0.$$

In the same way, one can deduce for all $t \in [0, T]$

$$W_t^* = (1 - \gamma) \left(C_t^* \right)^{-\gamma} V_t.$$
(19)

Applying Ito's Lemma to (??) and matching the volatility term with that in the budget constraint (??) yield an expression for the optimal portfolio α_t^*

$$\alpha_t^* = \frac{\sigma_t^C + \sigma_t^A}{\sigma_S} = \frac{\hat{\mu}_t - r}{\gamma \sigma_S^2} - \frac{1}{\gamma} \frac{\kappa}{\sigma_S} + \frac{\sigma_t^A}{\sigma_S}.$$
(20)

The martingale representation theorem implies that wealth at time t, W_t , is given by

$$\xi_t W_t^* = \mathbb{E}_t^{\mathbb{P}} \left[\int_t^T \xi_s C_s^* ds + \xi_T W_T^* \right].$$
(21)

By Ito's lemma, the volatility of the left-hand side of (??) is $-\xi_t W_t^* \hat{\nu}_t + \xi_t W_t^* \alpha_t^* \sigma_s$. By the Clark-Ocone formula, the volatility of the right-hand side is given by $\mathbb{E}_t^{\mathbb{P}} \left[D_t \left(\int_t^T \xi_s C_s^* ds + \xi_T W_T^* \right) \right]$. The two volatilities must be equal, leading to the following equality

$$\xi_t W_t^* \alpha_t^* = \frac{1}{\sigma_S} \xi_t W_t^* \hat{\nu}_t + \frac{1}{\sigma_S} \mathbb{E}_t^{\mathbb{P}} \left[D_t \left(\int_t^T \xi_s C_s^* ds + \xi_T W_T^* \right) \right].$$
(22)

In (??), the Malliavin derivative $D_t \left(\int_t^T \xi_s C_s^* ds + \xi_T W_T^* \right)$ can be computed as follows. First, by linearity and exchangeability between the Malliavin derivative and the ordinary Lebesgue integral, we have

$$D_t \left(\int_t^T \xi_s C_s^* ds + \xi_T W_T^* \right) = \int_t^T D_t \left(\xi_s C_s^* \right) ds + D_t \left(\xi_T W_T^* \right).$$
(23)

The second term on the right-hand side of (??) can be written as $D_t (\xi_T W_T^*) = W_T^* D_t \xi_T + \xi_T D_t W_T^*$ where $D_t W_T^*$ is computed by the chain rule of Malliavin calculus:

$$D_t W_T^* = \frac{W_T^*}{\gamma} \left(\frac{z_T^{\kappa}}{\xi_T}\right)^{-1} D_t \left(\frac{z_T^{\kappa}}{\xi_T}\right) = \frac{W_T^*}{\gamma} \left(\frac{z_T^{\kappa}}{\xi_T}\right)^{-1} \left(\frac{1}{\xi_T} D_t z_T^{\kappa} - \frac{z_T^{\kappa}}{\xi_T^2} D_t \xi_T\right).$$

The term $D_t z_T^{\kappa}$ is further computed as

$$D_t z_T^{\kappa} = D_t \exp\left(-\frac{1}{2}\int_0^T \kappa^2 ds - \int_0^T \kappa d\hat{B}_s\right)$$
$$= z_T^{\kappa} D_t \left(-\frac{1}{2}\int_0^T \kappa^2 ds - \int_0^T \kappa d\hat{B}_s\right) = -\kappa z_T^{\kappa} \mathbf{1}_{t \le T}$$

where $1_{t \leq T}$ is an indicator function. Rearranging the terms and assuming $t \leq T$ yield

$$D_t\left(\xi_T W_T^*\right) = -\frac{\kappa}{\gamma} \xi_T W_T^* + \frac{\gamma - 1}{\gamma} W_T^* D_t \xi_T.$$
(24)

Similarly, for $t \leq s$, one can derive

$$D_t\left(\xi_s C_s^*\right) = -\frac{\kappa}{\gamma} \xi_s C_s^* + \frac{\gamma - 1}{\gamma} W_s^* D_t \xi_s.$$
(25)

Substituting (??) and (??) into (??), rearranging terms and applying the equality (??) give us

$$\alpha_t^* = \frac{\hat{\nu}_t}{\sigma_S} - \frac{\kappa}{\gamma \sigma_S} + \frac{\gamma - 1}{\gamma} \frac{1}{\sigma_S W_t^*} \mathbb{E}_t^{\mathbb{P}} \left[\int_t^T C_s^* \frac{D_t \xi_s}{\xi_t} ds + W_T^* \frac{D_t \xi_T}{\xi_t} \right]$$

where $D_t \xi_s, t \leq s$ is computed as

$$D_t \xi_s = -\xi_s \left(\hat{\nu}_t + \int_t^s \left(d\hat{B}_\tau + \hat{\nu}_\tau d\tau \right) D_t \hat{\nu}_\tau \right).$$

The optimal portfolio α_t^* is given by

$$\alpha_t^* = \frac{\hat{\nu}_t - \kappa}{\gamma \sigma_S} + \frac{1 - \gamma}{\gamma} \frac{1}{\sigma_S X_t^*} \mathbb{E}_t^{\mathbb{P}} \left[\int_t^T \frac{\xi_s}{\xi_t} C_s^* \left(\int_t^s D_t \hat{\nu}_\tau \left(d\hat{B}_\tau + \hat{\nu}_\tau d\tau \right) \right) ds + \frac{\xi_T}{\xi_t} W_T^* \int_t^T \left(d\hat{B}_s + \hat{\nu}_s ds \right) D_t \hat{\nu}_s \right]$$

which can be rewritten using the equivalent martingale measure $\tilde{\mathbb{P}}$ as

$$\alpha_t^* = \frac{\hat{\mu}_t - r}{\gamma \sigma_S^2} - \frac{\kappa}{\gamma \sigma_S} + \frac{1 - \gamma}{\gamma} \frac{e^{rt}}{\sigma_S W_t^*} \mathbb{E}_t^{\tilde{\mathbb{P}}} \left[\int_t^T e^{-rs} C_s^* \int_t^s \left(D_t \hat{\nu}_\tau d\tilde{B}_\tau \right) ds + e^{-rT} W_T^* \int_t^T D_t \hat{\nu}_s d\tilde{B}_s \right]$$
(26)

where $\tilde{B}_t = \hat{B}_t + \int_0^t \hat{\nu}_s ds$. Comparing (??) to (??), we have

$$\sigma_t^A = \frac{1-\gamma}{\gamma} \frac{e^{rt}}{W_t^*} \mathbb{E}_t^{\tilde{\mathbb{P}}} \left[\int_t^T e^{-rs} C_s^* \int_t^s \left(D_t \hat{\nu}_\tau d\tilde{B}_\tau \right) ds + e^{-rT} W_T^* \int_t^T D_t \hat{\nu}_s d\tilde{B}_s \right]$$

Thus, σ_t^V is given by

$$\sigma_t^V = \frac{(1-\gamma)V_t}{\gamma} \left[(\hat{\nu}_t - \kappa) + \frac{e^{rt}}{W_t^*} \mathbb{E}_t^{\tilde{\mathbb{P}}} \left[\int_t^T e^{-rs} C_s^* \int_t^s \left(D_t \hat{\nu}_\tau d\tilde{B}_\tau \right) ds + e^{-rT} W_T^* \int_t^T D_t \hat{\nu}_s d\tilde{B}_s \right] \right].$$

Because $\mathbb{E}_{t}^{\tilde{\mathbb{P}}}\left[\int_{t}^{T} e^{-rs}C_{s}^{*}\int_{t}^{s}\left(D_{t}\hat{\nu}_{\tau}d\tilde{B}_{\tau}\right)ds + e^{-rT}W_{T}^{*}\int_{t}^{T}D_{t}\hat{\nu}_{s}d\tilde{B}_{s}\right] > 0$ and $W_{t}^{*} > 0$ for all $t \in [0,T]$ and the condition (??) implies that $\nu_{t} - \kappa \geq 0$ always holds true, σ_{t}^{V} is strictly positive for all $t \in [0,T]$ when either $\gamma > 1$ or $0 < \gamma < 1$ holds. As a result, one can verify by (??) that $\theta_{t}^{*} = \kappa$ for all $t \in [0,T]$ and the optimal consumption and portfolio choice are given by (??) and (??). This completes the proof. \Box

Proof of Corollary 1 For $\gamma = 1$, the utility process along the optimal path has the form $V_t = \log C_t^* + \log A_t$. By Ito's lemma, σ_t^V is given by $\sigma_t^V = \sigma_t^C + \sigma_t^A$, where $\sigma_t^C = \hat{\nu}_t - \kappa$. The optimal portfolio α_t^* can be expressed as

$$\alpha_t^* = \frac{\hat{\mu}_t - r}{\sigma_S^2} - \frac{\kappa}{\sigma_S} + \frac{\sigma_t^A}{\sigma_S}.$$

It follows from $\gamma = 1$ that

$$D_t\left(\xi_T W_T^*\right) = -\kappa \xi_T W_T^*$$
 and $D_t\left(\xi_s C_s^*\right) = -\kappa \xi_s C_s^*$

and that the optimal portfolio α_t^* is given by

$$\alpha_t^* = \frac{\hat{\mu}_t - r}{\sigma_S^2} - \frac{\kappa}{\sigma_S}$$

Thus, $\sigma_t^A = 0$ and $\sigma_t^V = \hat{\nu}_t - \kappa$. Since $\Pr(\pi_t > 0) = 1$ for $0 \le t \le T$ under the condition $\pi_0 > 0$ (see Lemma 9.3, Liptser and Shiryaev 2001), it follows that $\sigma_t^V > 0$ for $0 \le t \le T$. Thus, $\theta_t^* = \kappa$ for all $t \in [0, T]$ and the optimal consumption and portfolio policies are those given in (??) and (??) by setting $\gamma = 1$. This completes the proof. \Box

References

- Ai, H., 2009. Information about long-run risks: asset pricing implications. Journal of Finance forthcoming.
- [2] Anderson, E. W., L. P. Hansen, and T. J. Sargent (2003). A quartet of semigroups for model specification, robustness, price of risk, and model detection. Journal of the European Economic Association 1, 68-123.
- [3] Bjök, T., M. H. A., Davis, and C. Landé, 2010, Optimal investment under partial information. Mathematical Methods of Operations Research 71, 371-399.
- [4] Brennan, M. J., 1998. The role of learning in dynamic portfolio decisions. European Finance Review 1, 295-306.
- [5] Cagetti, M., L. P. Hansen, T. J. Sargent and N. Williams, 2002. Robustness and pricing with uncertain growth. Review of Financial Studies 15, 363-404.
- [6] Campbell, J. Y. and L. M. Viceira, 1999. Consumption and portfolio decisions when expected returns are time varying. Quarterly Journal of Economics 114, 433-495.
- [7] Chen, Z. and L. G. Epstein, 2002. Ambiguity, risk and asset returns in continuous time. Econometrica 4, 1403-1445.
- [8] Chen, H., N. Ju and J. Miao (2009). Dynamic asset allocation with ambiguous return predictability. Working paper, Boston University.
- [9] Campanale, C., 2009. Learning, ambiguity and life-cycle portfolio allocation. Review of Economic Dynamics, forthcoming.
- [10] Cox, J. C. and C. Huang, 1989. Optimal consumption and portfolio policies when asset prices follow a diffusion process. Journal of Economic Theory 49, 33-83.
- [11] David, A., 1997. Fluctuating confidence in stock market: implications for returns and volatility. Journal of Financial and Quantitative Analysis 32, 427-462.
- [12] Detemple, J., 1986. Asset pricing in a production economy with incomplete information. Journal of Finance 61, 383-392.

- [13] Detemple, J. and S. Murthy, 1994. Intertemporal asset pricing with heterogeneous beliefs. Journal of Economic Theory 62, 294-320.
- [14] Detemple, J., R. Garcia and M. Rindisbacher, 2003. A Monte-Carlo method for optimal portfolios. Journal of Finance 58, 401-446.
- [15] Detemple, J., R. Garcia and M. Rindisbacher, 2005. Representation formulas for Malliavin derivatives of diffusion processes. Finance and Stochastics 9, 349-367.
- [16] Dothan, M. U. and D. Feldman, 1986. Equilibrium interest rates and multiperiod bonds in a partially observable economy. Journal of Finance 41, 369-382.
- [17] Dybvig, P. H., L. C. G. Rogers, and K. Back, 1999. Portfolio turnpikes. Review of Financial Studies 12, 165-195.
- [18] Epstein, L. G. and J. Miao, 2003. A two-person dynamic equilibrium under ambiguity. Journal of Economic Dynamics and Control 27, 1253-1288.
- [19] Epstein, L. G. and M. Schneider, 2003. Recursive multiple priors. Journal of Economic Theory 113, 1-31.
- [20] Epstein, L. G. and M. Schneider, 2007. Learning under ambiguity. Review of Economic Studies 74, 1275-1303.
- [21] Feldman, D., 1989. The term structure of interest rates in a partially observable economy. Journal of Finance 44, 789C812.
- [22] Feldman, D., 1992. Logarithmic preferences, myopic decisions, and incomplete information. Journal of Financial and Quantitative Analysis 27, 619-629.
- [23] Feldman, D., 2007. Incomplete information equilibria: separation theorem and other myths. Annals of Operations Research 151, 119-149.
- [24] Gagliardini, P., P. Porchia, and F. Trojani. 2009. Ambiguity aversion and the term structure of interest rates. Review of Financial Studies 22, 4147-4188.
- [25] Gennotte, G., 1986. Optimal portfolio choice under incomplete information. Journal of Finance 41, 733-749.

- [26] Gilboa, I. and D. Schmeidler, 1989. Maxmin expected utility with non-unique priors. Journal of Mathematical Economics 18, 141-153.
- [27] Guidolin, M. and A. Timmermann, 2007. Asset allocation under multivariate regime switching. Journal of Economic Dynamics and Control 31, 3503-3544.
- [28] Hansen, L. P., and T. J. Sargent, 2001. Robust control and model uncertainty. American Economic Review 91, 60-66.
- [29] Israel, R. B., J. S. Rosenthal, and J. Z. Wei, 2001. Finding generators from markov chains via empirical transition matrices with applications to credit ratings. Mathematical Finance 11, 245-265.
- [30] Johnsen, T. H., and J. B. Donaldson, 1985. The structure of intertemporal preferences under uncertainty and time consistent plans. Econornetrica 53, 1451-58.
- [31] Ju, N., and J. Miao (2010). Ambiguity, learning and asset returns. Working paper, Boston University.
- [32] Karatzas, I., X. Xue, 1991. A note on utility maximization under partial observations. Mathematical Finance 1, 57-70.
- [33] Kim, T. S., and E. Omberg, 1996. Dynamic nonmyopic portfolio behavior. Review of Financial Studies 9, 141-161.
- [34] Korn, R., and H. Kraft, 2004. On the stability of continuous-time portfolio problems with stochastic opportunity set. Mathematical Finance 14, 403-414.
- [35] Lakner, P., 1998. Optimal trading strategy for an investor: the case of partial information. Stochastic Processes and Their Applications 76, 77-97.
- [36] Leippold, M., F. Trojani and P. Vanini, 2008. Learning and asset prices under ambiguous information. Review of Financial Studies 21, 2565-2597.
- [37] Liu, H., 2010. Robust consumption and portfolio choice for time varying investment opportunities. Annals of Finance 6, 435-454.
- [38] Lundtofte, F., 2008. Expected life-time utility and hedging demands in a partially observable economy. European Economic Review 52, 1072-7096.

- [39] Maenhout, P., 2004. Robust portfolio rules and asset pricing. Review of Financial Studies 17, 951-983.
- [40] Maenhout, P., 2006. Robust portfolio rules and detection-error probabilities for a meanreverting risk premium. Journal of Economic Theory 128, 136-163.
- [41] Merton, R. C., 1971. Optimal consumption and portfolio rules in a continuous-time model. Journal of Economic Theory 3, 373-413.
- [42] Merton, R. C., 1980. On estimating the expected return on the market: an exploratory investigation. Journal of Financial Economics 8, 323-361.
- [43] Miao, J., 2009. Ambiguity, risk and portfolio choice under incomplete information. Annals of Economics and Finance 10, 257-279
- [44] Nualart, D., 1995. The Malliavin Calculus and Related Topics. Springer-Verlag, New York.
- [45] Ocone, D. L., I. Karatzas, 1991. A generalized Clark representation formula, with application to optimal portfolios. Stochastic and Stochastic Reports 34, 187-220.
- [46] Oksendal, B., 1997. An introduction to malliavin culculus with applications to economics. Teaching Notes, University of Oslo Norway.
- [47] Schroder M., and C. Skiadas, 1999. Optimal consumption and portfolio selection with stochastic differential utility. Journal of Economic Theory 89, 68-126.
- [48] Schroder M., and C. Skiadas, 2003. Optimal lifetime consumption-portfolio strategies under trading constraints and generalized recursive preferences. Stochastic Processes and their Applications 108, 155-202.
- [49] Sbuelz A., and F. Trojani, 2008. Asset prices with locally-constrained-entropy recursive multiple-priors utility. Journal of Economic Dynamics and Control 32, 3695-3717.
- [50] Trojani, F., and P. Vanini., 2002. A note on robustness in Merton's model of intertemporal consumption and portfolio choice. Journal of Economic Dynamics and Control 26, 423-435.
- [51] Trojani, F., and P. Vanini., 2004. Robustness and ambiguity aversion in general equilibrium. Review of Finance 8, 279-324.

- [52] Uppal, R., and T. Wang, 2003. Model misspecification and underdiversification. Journal of Finance 58, 2465-2486.
- [53] Veronesi, P., 1999. Stock market overreaction to bad news in good times: a rational expectation equilibrium model. Review of Financial Studies 12, 975-1007.
- [54] Xia, Y., 2001. Learning about predictability: the effect of parameter uncertainty on optimal dynamic asset allocation. Journal of Finance 56, 205-246.

Accepted manuscript



Figure 1: Optimal hedging demands plotted against the state probability: interim consumption and terminal wealth. The graphs plot the optimal hedging demands as functions of π for the horizon T = 10 years and $\kappa = 0, 0.08$ and 0.16.



Figure 2: Hedging demand components (with utility over interim consumption) plotted against the state probability: $\kappa = 0,0.08$ and 0.16. The investment horizon is 10 years. The left graph plots the optimal hedging demand. The right graph plots the hedging demand components, $hedge^{IIR}$ (the solid line), $hedge^{\kappa}$ with $\kappa = 0.08$ (the dash-dot line) and $hedge^{\kappa}$ with $\kappa = 0.16$ (the dashed line). $hedge^{IIR}$: hedging component due to incomplete information risk. $hedge^{\kappa}$ ($hedge^{ambiguity}$): hedging component due to ambiguity.



Figure 3: Optimal stock demand and the fraction of hedging demand in the optimal stock demand (with utility over interim consumption) plotted against the state probability for $\kappa = 0, 0.08$ and 0.16. The investment horizon is 10 years.



Figure 4: Horizon effect. The investment horizon ranges from T = 0 to 20 years. The state probability π is set at its steady-state value.

anges from 2.



Figure 5: The economic value of regimes. The graph plots the percentage difference in the present value of the certainty equivalent wealth (*PCEW*) between a hidden Markov model and an IID model as a function of the horizon (ranging from 1 year to 20 years) for $\kappa = 0,0.08$ and 0.16. The parameter estimates of the hidden Markov model are reported in Table 1. The state probability π is set at its steady-state value.



Figure 6: Time series plot of *PCEWs* under different models and preferences: 1986.1–1995.4. The model specifications include a hidden Markov model (HMM) with two regimes and an IID model. Three alternative preferences include $\kappa = 0$ (expected utility), $\kappa = 0.08$ and $\kappa = 0.16$.

Table 1: Parameter estimates of a two-regime Markov switching model for returns (1947.1–1995.4)

Parameter Descriptions	Notation	Values
	1101411011	0.1100
High mean return regime	μ_H	0.1129
Low mean return regime	μ_L	0.0283
Volatility of returns	σ	0.1569
Transition density parameter	χ	0.2589
Transition density parameter	λ	0.6018

This table reports the ML estimates of a two-regime hidden Markov model for stock returns where the low mean return regime, μ_L , is exogenously fixed. The estimates reported are continuous-time estimates that are derived from the discrete-time estimates, where the latter are obtained based on real quarterly log returns on the CRSP value-weighted market portfolio for the period 1947.1–1995.4. The parameter values are calculated by suitably annualizing the discrete-time estimates.

		HMM		IID		
	T = 1	T = 5	T = 10	T = 1	T=5	T = 10
Panel A: $\kappa = 0$						
Mean	-0.336	-0.276	-0.190	-0.337	-0.284	-0.204
Stdev	0.022	0.024	0.033	0.023	0.025	0.038
5% c.ilower	-0.381	-0.324	-0.256	-0.384	-0.333	-0.281
5% c.iupper	-0.294	-0.231	-0.128	-0.295	-0.236	-0.133
10% c.ilower	-0.373	-0.316	-0.245	-0.376	-0.325	-0.268
10% c.iupper	-0.301	-0.238	-0.137	-0.301	-0.243	-0.143
Panel B: $\kappa = 0.08$						
Mean	-0.321	-0.240	-0.147	-0.321	-0.245	-0.156
Stdev	0.018	0.018	0.022	0.018	0.019	0.025
5% c.ilower	-0.356	-0.276	-0.191	-0.358	-0.282	-0.206
5% c.iupper	-0.288	-0.206	-0.107	-0.288	-0.209	-0.110
10% c.ilower	-0.350	-0.270	-0.184	-0.352	-0.276	-0.198
10% c.iupper	-0.293	-0.211	-0.113	-0.293	-0.215	-0.117
Panel C: $\kappa = 0.16$						
Mean	-0.308	-0.211	-0.117	-0.309	-0.214	-0.122
Stdev	0.014	0.013	0.014	0.014	0.014	0.016
5% c.ilower	-0.336	-0.237	-0.145	-0.337	-0.241	-0.154
5% c.iupper	-0.283	-0.185	-0.091	-0.283	-0.187	-0.092
10% c.ilower	-0.331	-0.233	-0.141	-0.332	-0.237	-0.149
10% c.iupper	-0.287	-0.189	-0.095	-0.287	-0.192	-0.097

Table 2: Comparison of out-of-sample performance I

This table summarizes out-of-sample performance measures for portfolio strategies under different return specifications and preferences. The return specifications include a hidden Markov model (HMM) with two regimes and an IID model. The alternative preferences include $\kappa = 0$ (expected utility), $\kappa = 0.08$ and $\kappa = 0.16$. The out-of-sample performance measures are calculated based on realized power utility with $\gamma = 4$ and 50,000 bootstrap trials for three investment horizons: 1, 5, and 10 years and for the out-of-sample period 1996.1–2009.4. The results reported include the mean, standard deviation, 5% and 10% confidence intervals of the bootstrap trials of realized power utility.

_								
_		HMM				IID		
		T = 1	T = 5	T = 10	T = 1	T=5	T = 10	
	Panel A: $\kappa = 0$							
	Mean	-0.027	-0.065	-0.072	-0.029	-0.073	-0.086	
	Stdev	0.009	0.011	0.019	0.010	0.012	0.024	
Κ	<i>t</i> -stat	3.136	6.248	3.797	3.016	6.229	3.564	
	Panel B: $\kappa = 0.08$							
	Mean	-0.012	-0.029	-0.030	-0.013	-0.034	-0.039	
	Stdev	0.004	0.005	0.008	0.005	0.006	0.011	
	<i>t</i> -stat	3.056	6.267	3.918	2.819	6.120	3.515	
_	Panel C: $\kappa = 0.16$							
	Mean	n.a.	n.a.	n.a.	-0.001	-0.003	-0.005	
	Stdev	n.a.	n.a.	n.a.	0.001	0.001	0.002	
_	<i>t</i> -stat	n.a.	n.a.	n.a.	0.595	2.568	2.258	

Table 3: Comparison of out-of-sample performance II

This table summarizes out-of-sample performance measures for portfolio strategies under different return specifications and preferences. The return specifications include a hidden Markov model (HMM) with two regimes and an IID model. The alternative preferences include $\kappa = 0$ (expected utility), $\kappa = 0.08$ and $\kappa = 0.16$. The out-of-sample performance measures are calculated based on differences in realized power utility with $\gamma = 4$ between the strategy {HMM, $\kappa = 0.16$ } and all the other strategies. The number of bootstrap trials is 50,000. The investment horizons are 1, 5, and 10 years. The out-of-sample period is 1996.1–2009.4. The results reported include the mean, standard deviation, and *t*-stat of the bootstrap trials of differences in realized power utility.