WORKING PAPER SERIES

Ambiguity in Asset Pricing and Portfolio Choice: A Review of the Literature

Massimo Guidolin, Francesca Rinaldi

Working Paper n. 417

This Version: October 2011

IGIER – Università Bocconi, Via Guglielmo Röntgen 1, 20136 Milano –Italy
http://www.igier.unibocconi.it

Ambiguity in Asset Pricing and Portfolio Choice: A Review of the Literature*

Massimo Guidolin† Bocconi University and IGIER

Francesca Rinaldi‡ Banque de France

Abstract

Empirical research suggests that investors' behavior is not well described by the traditional paradigm of (subjective) expected utility maximization under rational expectations. A literature has arisen that models agents whose choices are consistent with models that are less restrictive than the standard subjective expected utility framework. In this paper we survey the literature that has explored the implications of decision-making under ambiguity for financial market outcomes, such as portfolio choice and equilibrium asset prices. We conclude that the ambiguity literature has led to a number of significant advances in our ability to rationalize empirical features of asset returns and portfolio decisions, such as the failure of the two-fund separation theorem in portfolio decisions, the modest exposure to risky securities observed for a majority of investors, the home equity preference in international portfolio diversification, the excess volatility of asset returns, the equity premium and the risk-free rate puzzles, and the occurrence of trading break-downs.

JEL codes: G10, G18, D81.

Keywords: ambiguity, ambiguity-aversion, participation, liquidity, asset pricing.

1. Introduction

Traditional finance models assume that investors are perfectly aware of their own preferences mapping from utility-relevant states of the world into their perceived welfare, that they form rational expectations free of systematic biases, and therefore that they maximize their (subjective) expected utility (SEU). However, a growing body of empirical evidence suggests that investors' behavior is not well described by this traditional paradigm, since actual choices are often incompatible with the SEU predictions. One direction taken by the recent literature is behavioral finance, according to which the absolute rationality of investors' is replaced by any number of psychology-based alternatives, such as over-confidence, under-reaction, loss aversion, etc. (see e.g., Barberis and Thaler, 2003; Hirshleifer, 2001). Another strand of literature has focused instead on Bayesian model uncertainty (see, e.g., Pastor, 2000) or econometric learning and incomplete information in financial decisions and asset pricing (see e.g., Veronesi, 1999). This literature replaces rational expectations with beliefs updated through a rational learning rule, for instance Bayes' rule, in the light of the arrival of stochastic signals with non-zero correlation with relevant fundamentals. Under these approaches, investment decisions can differ from standard ones due to the difficulty of learning the true state given a complex generating process (for instance, subject to breaks or regimes). A third

*We would like to thank one anonymous referee for his constructive comments, Hening Liu and Massimo Marinacci for their encouragement and many stimulating conversations. Yu Man Tam provided invaluable research assistance while M. Guidolin was with the Federal Reserve Bank of St. Louis. All errors and inaccuracies remain of course our own.

†Correspondence to: Department of Finance, Via Roentgen, 1, Bocconi University, 20136 Milan, ITALY. E-mail: massimo.guidolin@unibocconi.it.

‡E-mail: Francesca.Rinaldi@banque-france.fr.

approach focusses on alternative frameworks of rational decision-making. This literature entertains agents whose choices are consistent with models that are less restrictive than the standard SEU framework, in the sense that the underlying axioms are less demanding. In this area, particular attention has recently been dedicated to ambiguity aversion. Under SEU, there are numerical utilities and probabilities that represent acts by a weighted sum of the utilities/outcomes in the possible states of the world, where the weights are (subjective) probabilities for each of the states (see e.g., Savage, 1954). Keynes was the first to distinguish *risk*, or known probability, from *uncertainty*. However, the modern attack to SEU as a descriptive theory was made most directly by Ellsberg's (1961) paradox that led researchers to collect massive experimental evidence that indicates that people generally prefer the least ambiguous acts and take their own confidence in estimates of subjective probability into account when making decisions. Such a pattern is inconsistent with Savage's *sure-thing principle* of SEU, the axiom by which a state with a consequence common to a pair of acts is irrelevant in determining preference between the acts.

In this paper we systematically review the literature that has explored the implications of decision-making under ambiguity for financial market outcomes, such as portfolio choices and equilibrium asset prices. Can a survey of the literature on ambiguity in financial markets be of any use? Our answer is obviously positive and this for at least three reasons. First, because the last decade has seen a tremendous growth in the number, breadth, and quality of papers that have applied ambiguity to financial research.[1] Second, because several recent papers have connected key events from the Great Financial Crisis of 2008-2009 to ambiguity in the form of poorly understood information and to investors' aversion to difficult-to-quantify uncertainty, as opposed to risk. Third, the background of the efforts by the scores of researchers involved in advancing our understanding of ambiguity is the existence in the field of financial economics of dozens of empirical "stylized facts" for which standard, SEU-based models have been incapable of providing a consistent rationalization. The ambiguity literature is currently perceived as one of the promising reactions to the generalized dissatisfaction with the empirical and predictive performance of the SEU/rational expectations (RE) paradigm. Indeed, the SEU hypothesis faces serious difficulties when confronted with asset markets data. Mehra and Prescott (1985) showed that for a standard rational SEU model to explain the high equity premium observed in the data, an implausibly high degree of risk aversion is needed, resulting in the equity premium puzzle. Weil (1989) showed that this high degree of risk aversion generates an implausibly high risk-free rate, creating a twin risk-free rate puzzle. In addition, a number of empirical studies document hard-to-understand links between aggregate asset markets and macroeconomics; for example, price-dividend ratios move pro-cyclically and conditional expected equity premia move countercyclically. Since French and Poterba (1991), researchers have become acutely aware that investors tend to forego valuable diversification opportunities by biasing their equity portfolios towards domestic stocks. It is time to collect a series of coherent thoughts on whether and how models that take ambiguity into account may solve these puzzles.

A survey paper cannot claim to have found any new theoretical or empirical results. Yet, we think that our scrutiny of the literature on ambiguity has scored a number of steps forward in our ability to rationalize important empirical features of asset returns and asset allocation decisions. With regard to portfolio decisions, we have summarized a substantial body of work that concludes that ambiguity may imply violations of the classical two-fund separation theorem by which all investors should hold an identical risky portfolio and can differ only in

[1] Al-Najjar and Weinstein (2009) have spurred a debate on whether ambiguity-based approaches to economic modelling would be as sensible as their growing popularity implies. Al-Najjar and Weinstein are dismissive of the potential of ambiguity models in positive economics. In our view, this potential may be appreciated without any presumption about their prescriptive validity by discussing the breadth of applications in finance. See also Mukerji's (2010) discussion.

their relative allocations between the riskless asset and this risky mutual fund. Many papers have proved that—even assuming constant investment opportunities—ambiguity will affect the classical (Merton's) expressions for optimal portfolio weights and that under a range of parameterizations we should expect ambiguity to decrease the optimal exposure to risky securities. A few papers have applied these results to investigate why investors seem to irrationally bias their portfolios away from foreign stocks—the so-called home country equity bias. When ambiguity about the joint distribution of domestic and foreign returns is high, small differences in ambiguity in favor of home securities may result in rational portfolios that are significantly under-diversified relative to standard mean-variance portfolios. In a similar vein, ambiguity models have been applied to study flight to familiarity phenomena, which occur when investors suddenly switch towards familiar assets.

In the asset pricing camp, there has been considerable interest in the fact that—when the economy is populated by both SEU and by ambiguity-averse investors—equilibrium pricing functions become discontinuous around a limited participation region, a range of asset prices for which only SEU investors trade. Moreover, in the participating equilibria—when both SEU and ambiguity-averse investors trade—it has been argued that changes in "off-equilibrium" potential outcomes could possibly affect equilibrium prices. Therefore this literature has stressed the importance of regulation and policy interventions aimed at persuading investors that any extremely negative scenarios may be safely ruled out. Another recurrent finding is that whether ambiguity may concern systematic or idiosyncratic "risks" may play a role in the equilibrium outcomes of financial markets. Because the potential for ambiguity aversion to induce limited participation equilibria depends on the fact that the spread between the highest and the lowest possible return of the idiosyncratic risk component is *larger* than the spread between the highest and the lowest possible return of the systematic component, the high realized risk premia that many papers have explained through ambiguity aversion may be more the result of ambiguity concerning idiosyncratic payoffs than on the systematic ones. Other papers have analyzed not only the effects of ambiguity, but also of its "dispersion" across heterogeneous investors. The insight is that a limited participation equilibrium may exist and that in this equilibrium the rate of participation, the average measure of ambiguity, and the equity premium all decrease as uncertainty dispersion increases, whereas under full participation, the equity premium does not depend on uncertainty dispersion. In Lucas-type endowment models, ambiguity has also been found to be a cause for asset price indeterminacy and endogenous ("sunspot") volatility. However, when assumptions are introduced to avoid indeterminacy, it has been shown that risk premia can be decomposed in two parts, a standard SEU premium that tends to be proportional to the covariance between asset returns and (appropriate functions of) the rate of growth of fundamentals or the market portfolio, and an ambiguity premium. This dual structure greatly helps in developing a unified and elegant solution to the equity premium and risk-free rate puzzles.

Recently researchers have also made progress in characterizing alternative asset pricing models that formally encompass aversion to ambiguity. For instance, particular forms of dynamic, recursive ambiguity originate two-factor extensions of the classical, single-factor CAPM in which the second priced risk factor simply captures the effects of ambiguity on the intertemporal marginal rate of substitution. Interestingly, in an ambiguity-adjusted CAPM, it has been proven that while under SEU, the law of large numbers implies that the variance of the market portfolio tends to zero as the number of assets becomes large, under ambiguity, the market portfolio does not become less uncertain as the number of assets increases; ambiguity may be priced in equilibrium because it is only partially diversifiable. Therefore, the basic insight of the CAPM—that only systematic risk may be compensated—would remain valid. Related papers have shown that in general terms, aversion to ambiguity modifies the standard stochastic discount factor (SDF) pricing equation, $1 = E[\mathcal{M}_{t+1}R_{t+1}|\mathcal{F}_t]$ (where $\mathcal{M}_{t+1}$ denotes a SDF) by introducing a multiplicative term $\mathcal{M}^u_{t+1}$.

Finally, ambiguity has offered compelling insights in phenomena recently observed during the financial crisis, such as trading break-downs and persistently high fixed income yields for (essentially) riskless assets. It is an old result from the early 1990s that under ambiguity averse preferences, there exists an interval of prices within which the agent neither buys nor sells short the risky asset, so that when equilibrium forces fail to push the asset price outside such an interval, there will be no willingness to trade. This is in sharp contrast with the standard SEU framework, where a risk neutral investor will buy (sell) a positive share of the risky asset if its price is higher (lower) than the expected value of its future payoffs. These early findings have been recently revived to show that when there is more heterogeneity, an equilibrium with no-trade is more difficult to establish and that bid-ask spreads may derive entirely from ambiguity and not only from asymmetric information, as commonly assumed.

Before moving on to the core of our review, let us mention that a few papers have reviewed the literature on ambiguity (Camerer and Weber, 1992; Epstein and Schneider, 2010; Etner, Jeleva and Tallon, 2009; Gilboa, Postlewaite and Schmeidler, 2008; Mukerji and Tallon, 2004; Wakker, 2008; Gilboa and Marinacci, 2011). However none of these papers has the focus of our work. For instance, Camerer and Weber (1992), Etner, Jeleva and Tallon (2009), and Wakker (2008) have focus on defining ambiguity aversion and on how to best model such preferences, with a keen interest in issues of axiomatization, while Gilboa and Marinacci (2011) focus on theoretical foundations of ambiguity-averse preferences. In our view, Epstein and Schneider (2010) is the most closely comparable survey, although their attention is more on the mapping between the "smoothness" (or lack thereof) of ambiguity averse preferences and their potential implications in finance applications—which is clearly a key aspect—than on the breadth of the implications of ambiguity aversion for portfolio choice and asset pricing. Because their potential to affect financial decisions is maximum within the known class of ambiguity preferences, Epstein and Schneider focus on applications of multiple-prior preferences, while our review encompasses a broader spectrum of preferences.[2]

We have one final note of caution on terminology. In the literature, ambiguity and uncertainty are not always distinguished, nor are they clearly defined. In this survey, we will use both terms equivalently. Uncertainty or ambiguity is meant to represent "non probabilized" uncertainty—situations in which the decision maker (henceforth, DM) is not given probabilistic information about the external events that affect the outcome of a decision—as opposed to risk, which is "probabilized" uncertainty. Another issue revolves around the difference between ambiguity aversion and a "concern" (preference) for robustness. As we shall explain in Section 2.4, the two notions are not completely equivalent and their origins can be usefully told apart. However, for our purposes we will treat the literature on robust financial decisions as a specific strand of the general ambiguity literature, and insist more on the points for which the two approaches are similar than on the differences. Hansen and Sargent (2007) is the authoritative reference on robustness.

Section 2 reviews a few key definitions relating to ambiguity, its difference from risk, and how aversion to ambiguity might be measured. The purpose of this Section is to present some of the numerous specifications for ambiguity averse preferences proposed in the literature. Needless to say, the different theoretical issues that characterize alternative models are important but also beyond the scope of this paper (see e.g., Gilboa and Marinacci,

[2]There are at least two dimensions in which we feel any Reader will find precious companions in the papers referenced above. First, our treatment of the decision-theoretic aspects of ambiguity is cursory at best because we only care for a Reader to appreciate how and why the preferences entertained in the finance literature differ from standard SEU preferences. Second, we have devoted no space to the rich experimental literature on ambiguity. This does not reflect any implicit judgement on the relevance of such a literature.

2011, for an in-depth treatment). Section 3 reviews papers that have addressed issues related to optimal portfolio choice. We start with simple, static models that introduce the idea that under ambiguity, the best trade may easily be no trade at all. We connect ambiguity models to classical mean-variance asset allocation. We then extend the analysis to robust asset allocation and to large scale problems. Section 4 is devoted to models of equilibrium asset prices under ambiguity. The main distinction here is between simple, static two-period models and dynamic models. The latter have forced researchers to deal with a number of intriguing logical issues, such as dynamic consistency and rational updating of ambiguous beliefs. Section 5 is dedicated to advances in financial microstructure theory. Section 6 concludes.

2. **Generalities and Definitions: What is Ambiguity?**

A decision problem is structured on a state space, an outcome space, and a preference relation. The state space Ω, whose elements are called states of nature, represents all the possible realizations of uncertainty. Sets of states of nature, $E \subset \Omega$, are called events. The outcome space $\mathcal{F}$ contains the random outcomes of decisions.[3] A preference relation $\succsim$ is defined over the mappings from Ω to $\mathcal{F}$. These mappings are called acts or decisions and they associate to each state $s \in \Omega$ a consequence $f(s)$ (or f_s). $f \succsim g$ means that the decision maker (henceforth, DM) weakly prefers decision f to g; $f \sim g$ means that the DM is indifferent between f and g. Most of the time—we will clearly stress when this is not the case—all preferences are assumed to be complete (i.e., a DM is always able to rank decisions), reflexive ($f \succsim f$) and transitive (i.e., if $f \succsim g$ and $g \succsim h$, then $f \succsim h$). We generally denote by U a standard VNM utility index; BM means "Brownian motion"; $p(E)$ denotes the (subjective) probability of event E occurring.

2.1. *Early Literature: Ellsberg's Paradox*

That a distinction might be drawn between standard SEU and more general decision models has been known since Knight (1921). According to Knight, there are two kinds of uncertainty: the first, called *risk*, corresponds to situations in which all relevant events are associated with a (objectively or subjectively) uniquely determined probability assignment; the second, called (*Knightian*) *uncertainty*, corresponds to situations in which some events do not have an obvious probability assignment. Such a distinction is meaningful since in most economic contexts where agents face uncertainties, no probabilities are given or easily computable. The experimental relevance of the distinction between risk and uncertainty has been formally discussed by Ellsberg (1961), whose findings have shown that agents are not always able to derive a unique probability distribution over the reference state space. After Ellsberg's seminal paper, uncertain environments have became better known as *ambiguous* and the general dislike for them as *ambiguity aversion*.

To formally define aversion to ambiguity, consider a DM who places bets that depend on the result of two coin flips. The first coin is well known, while the second one is provided by an unknown intermediary. Given that the agent is not familiar with the second coin, it is possible that she would consider ambiguous all bets whose payoff depended on the result of the second flip. For instance, a bet f that pays \$1 if the second coin lands with head up, or, equivalently, on the event $\{HH, TH\}$ (here H denotes the head outcome and T the tail outcome), can be

[3] It will be sometimes convenient to assume that $\mathcal{F}$ is a set of lotteries. Thus, the result from a decision could for instance be: "if state s realizes, get a lottery that yields some amount x with probability p and some amount y with probability $1-p$".

seen as somewhat less desirable than bets that are "unambiguous", such as a bet that pays \$1 if the first coin lands with head up, or, equivalently, on the event $\{HH, HT\}$. If the DM is given the possibility of buying "shares" of bets that rely on the second coin only (namely the bet that pays only if $\{HH, TH\}$ obtains and the one that pays if $\{HT, TT\}$ occurs), so that she is offered a bet that pays \$0.50 on $\{HH\}$ and \$0.50 on $\{HT\}$, she may prefer it to either of the two ambiguous bets. In fact, such a bet has the same contingent payoffs of a bet which pays \$0.50 if the first coin lands with head up, which is unambiguous. That is, a DM who is averse to ambiguity may prefer the equal-probability "mixture" of two ambiguous bets to either of the individual bets. This is a stark violation of the standard SEU paradigm. Formally, Schmeidler (1989) called *ambiguity averse* a DM who prefers any mixture $\alpha f + (1-\alpha)g$, $\alpha \in [0,1]$ of two ambiguous bets f and g between which she is indifferent to each of the individual bets: $\alpha f + (1-\alpha)g \succsim f \sim g \ \forall \ f,\ g$. This DM is averse to ambiguity because she benefits from the fact that the mixture induced by the weight $\alpha \in [0,1]$ reduces the overall ambiguity of the two ambiguous bets f and g.

2.2. *Choquet Expected Utility (CEU) and Multiple Prior Preferences (MPP)*

At the core of Ellsberg's paradox there is the awareness that—when she has too little information to form a single prior—a DM may plausibly consider a *set* of probability distributions and not a unique prior. Schmeidler (1989) formalized this intuition starting from the observation that the probability attached to an uncertain event may not reflect the heuristic amount of information that has led to that particular probability assignment. For example, when there are only two possible equiprobable events, they are usually given probability 1/2 each, independently of whether the available information is meager or abundant. Motivated by this consideration, Schmeidler suggested to assign non-additive probabilities, or capacities, to allow for the encoding of information that additive probabilities cannot represent.[4] With reference to a bet f with only two possible mutually-exclusive outcomes, say f_1 or f_2, a capacity v is any assignment to the events {neither f_1 nor f_2 occur}, {f_1 or f_2 or both occurs}, {f_1 and f_2 both occur}, {f_1 occurs}, {f_2 occurs}, such that: i) $v(\cdot) \geq 0$ and $v(f_1) + v(f_2) \leq 1$; ii) v(neither f_1 nor f_2 occurs) $= 0$; iii) v(f_1 or f_2 or both occurs) $= 1$.[5] Schmeidler's representation is based on the concept of *Choquet integral*, that in our example reduces to

$$CEU(f) = \min_{\mu \in C(v)} \left[\mu U(f_1) + (1-\mu) U(f_2)\right] \qquad C(v) = \{\mu \in [0,1], \quad \mu \geq v(f_1),\ 1-\mu \geq v(f_2)\}, \tag{1}$$

where $C(v)$ is the *core* of v.[6] In the example, the individual acts as if she were only able to establish for each outcome f_s $(s = 1,2)$ the minimal probability of occurrence $v(f_s)$. Because of the existence of ambiguity, she considers a multi-valued set of probability distributions uniquely defined by $C(v)$. Ambiguity aversion is reflected by the use of the *min* operator—the agent considers the most unfavorable probability distribution. To express the degree of ambiguity that characterizes the capacity assignment $(v(f_1), v(f_2))$, Schmeidler suggested the use of an index $\mathcal{A}(v, f) \equiv 1 - v(f_1) - v(f_2)$, that measures how much overall "faith" should be placed on the outcomes f_1 and f_2. Notice that, under an additive capacity, we get back to the standard SEU case. In general, CEU preferences are of the *multiple-prior type*, in the sense that in practice, under CEU a rational decision-maker evaluates expected

[4] In general, a capacity v over Ω satisfies the properties: (i) $v(\cdot) \in R^+$; (ii) for $E, F \in \mathcal{F}$ s.t. $E \subseteq F$ then $v(E) \leq v(F)$ (monotonicity); (iii) $v(\varnothing) = 0$ and $v(\Omega) = 1$ and (iv) $\sum_{s\in\Omega} v(s) \leq 1$. Probabilities are special cases of capacity functions that satisfy additivity: $v(A \cup B) = v(A) + v(B) - v(A \cap B)$.

[5] For the purpose of this survey, this simplified setting suffices to provide an elementary intuition. We refer a Reader to Gilboa and Marinacci (2011) for a general set of definitions and models.

[6] The core of a capacity v consists of all finitely additive probability measures μ that event-wise dominate v.

utility using a multi-valued set of priors, as defined by $\mu \in C(v)$.

Gilboa and Schmeidler (1989) further extended the CEU model by suggesting the following representation

$$f \succsim g \text{ if and only if } \min_{p\in\wp} E_p[U(f)] \geq \min_{p\in\wp} E_p[U(g)], \tag{2}$$

where $E_p[U(\cdot)]$ is a standard SEU-operator when the probability measure is $p \in \wp$. $\wp$ is a convex set whose "size" can be interpreted as representing the perception of ambiguity. The intuition is that agents are considering as relevant the most unfavorable prior. Given the multi-valued nature of $\wp$, preferences represented by (2) are commonly known as *multiple prior preferences* (MPP). MPP can be shown to be equivalent to CEU when the capacity is convex, because weighting outcomes by sub-additive capacities expresses the same kind of pessimism as taking the minimum SEU over $\wp$ (see e.g., Camerer and Weber, 1992).[7] However, strictly speaking, neither approach is a special case of the other. As noted by Wakker (2001), the distinction between probability measures that are either possible (contained in $\wp$) or impossible (not contained in $\wp$), on the one hand adds to the tractability of the model, but on the other hand cannot capture cognitive states where different probability measures are plausible to different degrees. Multiple priors are, however, well suited for general theoretical analyses where only general qualitative properties of the model are interesting.[8]

The 1990s have witnessed a number of extensions and critiques to the seminal definition of ambiguity provided by Schmeidler (1989). Epstein (1999) and Ghirardato and Marinacci (2002) have criticized Schmeidler's (1989) notion of ambiguity aversion based on the convexity of the capacity v, showing that convexity is neither necessary nor sufficient to generate behavior that is compatible with the intuition of ambiguity aversion, meaning that a non-convex capacity may also induce Ellsberg's type ordering among bets, and also that a convex capacity can generate under CEU a ranking which is exactly opposite to the one showed by Ellsberg. For example, Ghirardato and Marinacci (2002) show that a CEU preference ordering is ambiguity-averse if and only if its capacity v has a non-empty core, a strictly weaker property than convexity (see Giboa and Marinacci, 2011, for a discussion of the different notions of ambiguity aversion in the literature).

2.3. *Robust Control and Variational Preferences*

After the seminal papers by Gilboa and Schmeidler had gained popularity (see e.g., the early review by Camerer and Weber, 1992), Anderson, Hansen and Sargent (1998, 2003) and Hansen and Sargent (2001) noted that multiple-prior criteria also appear in the robust control theory used in engineering. Robust control theory specifies the set of probabilities $\wp$ by taking a single "approximating model" and statistically perturbing it. This reflects a situation wherein agents have a specific model of reference and, acknowledging the possibility of errors, seek robustness against misspecifications. Usually, $\wp$ is implicitly parameterized through some coefficient ϱ such that the higher is ϱ, the less importance is given to alternative models deviating from the approximating one. Hansen and Sargent (2001) pointed out that the concern for the possibility of model-misspecification may derive from ambiguity and the

[7] That is, given two events A and B, $v(A \cup B) + v(A \cap B) \geqq v(A) + v(B)$

[8] Bewley (1986) had anticipated a few of these MPP intuitions. Specifically, Bewley had developed a class of preferences characterized by the multi-valued nature of the set of priors and by uncertainty aversion, but while most of the ambiguity literature has stressed that it is the bad quality of the information that is responsible for violations of the (S)EU axioms, Bewley suggested that the lack of information renders difficult the ranking of alternative bets, determining preference incompleteness. The intuition of Bewley's framework is that one bet is preferred to another if and only if its expected value is higher under all probability distributions that may be employed to capture risk. When lotteries are not comparable, Bewley assumes choices are made by inertia: the current choice, or *status quo*, is only abandoned if a new choice appears that is certainly better, i.e., it has higher expected utility for all possible probability distributions.

poor quality of information used to select the approximating model. Therefore ϱ can be thought of as an ambiguity aversion index: the lower is ϱ, the higher is the degree of ambiguity aversion. Hansen and Sargent and a number of coauthors have developed a class of preferences that may be represented as

$$f \succsim g \text{ if and only if } \min_{q\in\Delta(\Omega)} E_q[U(f)+\varrho R(q||p)] \geq \min_{q\in\Delta(\Omega)} E_q[U(g)+\varrho R(q||p)], \tag{3}$$

where $\Delta(\Omega)$ is the standard simplex, p is the approximating, baseline probability distribution, and $R(\cdot||p)$ is the Kullback-Leibler divergence between any probability distribution over the reference state space and p. For instance, $R(\cdot||p)$ may be specialized to be the entropy of q relative to p, i.e., $R(q||p) \equiv \sum_{s=1}^{|\Omega|} q(s)\ln[q(s)/p(s)] = E_q[\ln(q/p)] \geq 0$ so that $q = p$ implies $R(q||p) = 0$, which is the expected difference in log-likelihoods between the reference and transformed models, with the expectation based on the latter. In (3), the term $\varrho R(q||p)$ acts as a penalty term: intuitively, agents consider a range of models q in alternative to p, but they assign a higher "weight" to models that are close in a statistical sense to the approximating model, p.

The formulation in (3) is referred to as the "penalty problem". An alternative methodology used in the robust control literature consists in formulating a "constrained problem", i.e., defining the set $\wp$ of alternative distributions by constraining the relative entropy of the elements $q \in \wp$ with respect to the reference model p to be lower than a parameter η. According to this approach, preferences are represented by

$$f \succsim g \text{ if and only if } \min_{q\in C} E_q[U(f)] \geq \min_{q\in C} E_q[U(g)] \qquad C = \{q \in \Delta(\Omega) : R(q||p) \leq \eta\}, \tag{4}$$

which is a particular case of MPP. Hansen and Sargent (2001) showed that (4) and (3) are connected via the Lagrange multiplier theorem and that the preference orderings implied by the two approaches differ, but the two optimizations lead to identical decisions.

The Hansen-Sargent model (3) is highly suitable to applications as it can be embedded into intertemporal continuous-state-space frameworks. Let $\mathbf{B}_t$ be a d-dimensional BM on a probability space $(\Omega,\mathcal{F},p)$, and x_t a state-dependent variable which is assumed to evolve according to $dx_t = \mu(c_t,x_t)dt + \boldsymbol{\sigma}(c_t,x_t)d\mathbf{B}_t$. When ambiguity-aversion is introduced, the BM $\mathbf{B}_t$ is replaced by $\hat{\mathbf{B}}_t + \int_0^t \boldsymbol{\theta}_s ds$ where $\hat{\mathbf{B}}_t$ is another BM, and $\boldsymbol{\theta}_t$ is used to transform the probability distribution p into a new distribution q that is absolutely continuous with respect to p, because its Radon-Nikodym derivative, $z_t^{\theta} \equiv dq/dp$, is

$$z_t^{\theta} = \exp\left(-\frac{1}{2}\int_0^t \|\boldsymbol{\theta}_s\|^2 ds + \int_0^t \boldsymbol{\theta}_s d\hat{\mathbf{B}}_s\right), \tag{5}$$

so that the relative entropy between the measures q and p is given by:

$$R(q) \equiv \int_0^{\infty} \exp(-\delta s) E_q\left(-\frac{1}{2}\int_0^t \|\boldsymbol{\theta}_s\|^2\right) ds. \tag{6}$$

Hence, a possible, alternative "twisted" model can be represented as $d\mathbf{x}_t = \boldsymbol{\mu}(\mathbf{x}_t,\mathbf{c}_t)dt + \boldsymbol{\sigma}(\mathbf{x}_t,\mathbf{c}_t)[\boldsymbol{\theta}(\mathbf{x}_t)dt + d\mathbf{B}_t]$, and the optimal policy for an ambiguity-averse, Hansen-Sargent type agent can be derived by solving

$$\sup_{c_t}\inf_{q} E_q\left\{\int_0^{\infty} \exp(-\delta t)\left[U(c_t,x_t) - \frac{1}{2}\varrho E_q\left(\int_0^t \|\boldsymbol{\theta}_s\|^2\right)\right] dt\right\} \qquad \text{s.t. } dx_t = \mu(c_t,x_t)dt + \boldsymbol{\sigma}(c_t,x_t)d\mathbf{B}_t, \tag{7}$$

where the problem in (7) has the same structure as a standard SEU-optimization problem with two key amendments: expectations are computed under the measure q and deviations from p are penalized by ϱ times the Kullback-Leibler measure, $R(q)$. Hansen-Sargent robustness preferences (3) have been axiomatized only subsequently by Maccheroni, Marinacci, and Rustichini (2006) who have proven that robustness preferences are a specific sub-class of what they have labeled *variational preferences* (VP).[9] The tight relationship between VP and Hansen-Sargent robustness preferences is far from unique: VP have the advantage of nesting many of the known preferences structures that (may) capture ambiguity-aversion, including MPP. Another feature of VP is that they include as a special case Monotone Mean Variance Preferences (MMVP) that extend outside their domain of monotonicity the classical Mean Variance Preferences of Markowitz and Tobin.[10]

2.4. *Dynamic Representations of Ambiguity-Averse Preferences and Alternative Approaches*

Recently researchers in decisions have proposed ambiguity averse preferences within truly dynamic settings, which are important in financial applications. However, defining updating and ensuring dynamic consistency for CEU/MPP or similar models is a non-trivial task that has triggered extensive debated.[11,12] For instance, it is easy to assemble two-period extensions of Ellsberg's paradox in which an individual with Ellsberg-type preferences ends up displaying a dynamically inconsistent behavior as she fails to ex-post follow the path she has decided ex-ante, see e.g., the examples in Etner et al. (2009, pp. 32-33).

To see what the concern is, consider the following example adapted from Backus et al. (2004). Consider a simple model with three periods ($t = 0, 1, 2$) and in which—in each period—only two states are possible, $\Omega_t = \{1, 2\}$. Date 1 probabilities are $q(s_1 = 1) = q(s_1 = 2) = 1/2$ and they are not ambiguous. Date 2 (conditional) probabilities depend on s_1 and an autocorrelation parameter ϕ, for which the agent has point-wise priors on the values $+1$ and -1. As a result, the conditional probabilities of the four date 2 states are $q(s_2 = 1|s_1 = 1) = q(s_2 = 2|s_1 = 2) = (1 + \phi)/2$ and $q(s_2 = 2|s_1 = 1) = q(s_2 = 1|s_1 = 2) = (1 - \phi)/2$. The probabilities depend on whether s_1 and s_2 are the same or different and whether ϕ is $+1$ or -1. In this set up, consider the value of an asset that pays 1 unit of consumption good if $s_2 = 1$ and 0 otherwise, i.e., a state-contingent Arrow-Debreu security. Let $U(c) = c$, a risk-neutral linear utility index and set the subjective discount factor to 1. If any two recursive—over times and nodes of the tree—and date zero valuations of the asset differ, at least one set of preferences must be dynamically inconsistent. Consider recursive valuation first. At the node corresponding to the event $s_1 = 2$, the value of the asset is $(1 + \phi)/2$. Minimizing with respect to ϕ, as is implied by MPP preferences, implies $\phi = -1$ and a value of 0 for the Arrow-Debreu security. Similarly, the value at the node corresponding to $s_1 = 1$ is also 0, this time based on $\phi = +1$. The value at date zero is therefore 0 as well: there is no ambiguity, so the value is $(1/2)(0) + (1/2)(0) = 0$. Now consider a (naive) date zero problem based on the two-period probabilities of the four possible two-period paths: $(1 + \phi)/4$, $(1 - \phi)/4$, $(1 - \phi)/4$, and $(1 + \phi)/4$. Ambiguity on these probabilities is again represented by ϕ. Since the asset pays 1 if the first or third path occurs, its date zero value is $[(1 + \phi)/4] + [(1 - \phi)/4] = 1/2 > 0$, which is not ambiguous. The date 0 value (1/2) is clearly larger than the recursive value (0), so preferences are dynamically inconsistent. Here the problem is that the first, recursive valuation approach has allowed ϕ to differ across date 1 nodes, while the date 0, one-shot valuation does not. This may happen because giving the agent

[9] Strzalecki (2011) provides an axiomatization of Hansen and Sargent's preferences.

[10] Preferences characterized by functional representation $V(f) = \min_{p\in\Delta(\Omega)} E_p[f] + \varrho G(p||q)$ are a specific subclass of VP called Monotone Mean Variance Preferences (MMVP). Epstein (1999) has shown that in the domain of monotonicity of mean variance preferences the following equality holds: $E_p[f] - (2\rho)^{-1}Var(f) = \min_{p\in\Delta(\Omega)} E_p[f] + \varrho G(p||q)$, where $G(p||q)$ is the relative Gini concentration index. Section 4.1 describes applications of MMVP.

[11] For example, in the CEU model, conditionally on the realization of some event B, if one considers the Bayesian-updated core $C(v(\cdot|B))$ of a capacity v, this is in general different from $C(v')$, where v' is the lower envelope of $C(v(\cdot|B))$. Chateauneuf, Gajdas and Jaffray (2010) characterize a class of convex capacities for which $C(v(\cdot|B)) = core(v')$.

[12] As it is well known, dynamic consistency leads to Bayesian updating under (S)EU. Hanany and Klibanoff (2009) derive dynamically consistent updating rules for variational and KMM preferences.

access to date 1 information increases the amount of information but also increases the amount of ambiguity, which reduces the value of the asset. Therefore any resolution of this dynamic inconsistency problem must modify either the recursive or date 0 preferences.

Epstein and Schneider (2003) have proposed the latter approach. They show that if we expand the set of date zero probabilities in the right way, this will lead to the same preferences as under MPP, restoring dynamic consistency. In general, preferences depend on probabilities over complete paths, which in our example are associated with the four terminal nodes $\{s_0, s_1 = 2, s_2 = 2\}$, $\{s_0, s_1 = 2, s_2 = 1\}, \{s_0, s_1 = 1, s_2 = 2\}$, and $\{s_0, s_1 = 1, s_2 = 1\}$. Epstein and Schneider introduce a condition—called *rectangularity*—that instructs us to compute the set of probabilities recursively, one period at a time, starting at the end. At each step, we compute a set of probabilities for paths given our current history. In our example, the main effect of this approach is to eliminate any connection between the values of ϕ at the two date 1 nodes. The resulting date 0 probabilities are $(1 + \phi_1)/4$, $(1 - \phi_1)/4$, $(1 - \phi_2)/4$, and $(1 + \phi_2)/4$. The value of the asset is therefore $(1 + \phi_1)/4 + (1 - \phi_2)/4 = 1/2 + (\phi_1 - \phi_2)/4$. An ambiguity-averse investor will set $\phi_1 = -1$ and $\phi_2 = +1$ and the value is zero, the same value we computed recursively. Therefore expanding the date 0 set of probabilities in this way reconciles date 0 and recursive valuations and eliminates the dynamic inconsistency.

Chen and Epstein (2002) have extended these intuitions to develop a time-consistent, continuous-time intertemporal version of MPP that has been crucial in a number of asset pricing papers. In their model time varies over $[0, T]$ and uncertainty is represented by a probability space $(\Omega,\mathcal{F}, p)$. There is a single consumption good at each point in time and c_t denotes the consumption process. $\mathbf{B} = (\mathbf{B}_t)$ is a d-dimensional BM on $(\Omega,\mathcal{F}, p)$. A density generator is a d-dimensional process $\boldsymbol{\theta} = (\boldsymbol{\theta}_t)$, such that $\boldsymbol{\theta} \in \Theta$ and $E[\exp(-\frac{1}{2}\int_0^t \|\boldsymbol{\theta}_s\|^2 ds)] < \infty$. Under this assumption, $\boldsymbol{\theta}$ generates a set of probability measures, $\wp^\theta$, whose elements are defined through the Radon-Nikodym derivative, z_t^θ as in (5). Chen and Epstein generalized the notion of stochastic differential utility introduced by Duffie and Epstein (1992) allowing for ambiguity. Specifically, intertemporal MPP is represented as:

$$V_t \equiv \min_{q\in\wp^\theta} V_t^q = \min_{q\in\wp^\theta} E_q \left[\int_t^T \psi(c_s, V_s^q) ds | \mathcal{F}_t \right] \quad 0 \le t \le T. \tag{8}$$

where for each $q \in \wp^\theta$, V_t^q is a utility process and ψ is an aggregator as in Duffie and Epstein (1992). Under some regularity assumptions on the domain Θ, the ambiguity-aversion preferences in (8) are dynamically consistent, allowing for a recursive representation of the utility process.[13]

An alternative approach to the definition of ambiguity is the multi-stage approach. Notice that by construction any two-stage lottery is defined as a special bet with outcomes (at the end of the first stage) that can be thought of as further lotteries. Segal (1990) used two-stages lotteries to model ambiguous bets for which the first imaginary stage is a lottery over the possible values of the probability measure p that characterizes the second-stage lottery. Accordingly, in Ellsberg's experiment the DM would consider the urn with unknown composition as sampled from a set of urns: Ellsberg's behavior would be primarily related to a DM's attitude towards second-order risk, that is, risk aversion in the first stage. Ergin and Gul (2004) and Nau (2006) have extended the seminal intuitions by Segal and identified ambiguous bets with compound lotteries whose uncertainty derives from two independent issues. In their view, the violation of the SEU paradigm arises because agents are comparing bets that are based on different issues: some choices are based only on the issue of which ball is chosen (the knowledge of the composition of the

[13] Maccheroni, Marinacci, and Rustichini (2007), introduce dynamic VPs but also show that Hansen and Sargent's robustness preferences are dynamically consistent.

urn is not relevant), while other choices are based on both issues, the selection of the color of the ball and of the urn with its composition. Ambiguity-aversion translates into the preference for bets based only on the first issue.

Klibanoff, Marinacci and Mukerji (2005) have elaborated these multi-stage ideas by proposing that the ambiguity of a risky act f be characterized by a set $\wp = \{P_1, ..., P_n\}$ of subjectively plausible cumulative probability distributions for f. Letting f_j denote the random variable distributed as P_j, $j = 1, ..., n$, based on her subjective information, the DM associates a distribution $(q_1, ..., q_n)$ over $\wp$, where q_j is the subjective probability of P_j being the true distribution of f. The resulting preferences (KMM) have the representation[14]

$$f \succsim g \text{ if and only if } \sum_{j=1}^{n} q_j \zeta \left(\int U(f) dP_j \right) \geq \sum_{j=1}^{n} q_j \zeta \left(\int U(g) dP_j \right), \tag{9}$$

where $\zeta(\cdot)$ is an increasing real-valued function, whose shape describes the investor's attitude towards ambiguity. First, the DM evaluates the expected utility of f with respect to each P_j in $\wp$. Then she takes an expectation of distorted expected utilities. The role of ζ is crucial: if ζ were linear, the criterion reduces to SEU maximization. When ζ is not linear, qs and P_js cannot be combined. In this event, the DM takes the expected "ζ-utility" (with respect to q) of the expected "U-utility" (with respect to the Ps). A concave ζ will reflect ambiguity aversion, in the sense that it places a larger weight on bad expected "U-utility" realizations.

One important implication of this two-stage approach is that the DM is not forced to be so pessimistic as to select the act that maximizes the minimum expected utility as a consequence of the separation between ambiguity and the DM's attitude toward ambiguity. In this sense, KMM preferences may be interpreted as a "smooth" extension of Gilboa and Schmeidler's classical MPP. MPP is a limiting case of (9): up to ordinal equivalence, MPP is obtained in the limit as the degree of concavity of ζ increases without bound. As recently stressed by Maccheroni, Marinacci and Mukerji (2011), "The smooth ambiguity model allows us to explore implications of ambiguity aversion that do not have their source in preference kinks. Kinks are not implied by ambiguity averse (...) behavior (and, indeed, may be present without such behavior, see e.g. Segal and Spivak (1990)), yet they are what drive behavior in many applications of models like (...) Choquet expected utility (...) to economics and finance." Furthermore, an intuitive criterion for ambiguity aversion comparisons can be derived within the KMM framework: $\succsim_1^{KMM}$ is more ambiguity averse than $\succsim_2^{KMM}$ if and only if the associated ambiguity functions ζ_1 and ζ_2 are such that $-\zeta_1''/\zeta_1'$ is uniformly larger than $-\zeta_1''/\zeta_1'$, or, equivalently, if ζ_1 is more concave than ζ_2.[15] Contrary to MPP preferences, in the KMM model beliefs and ambiguity-aversion parameters are explicitly separated. For instance, it is the function $\zeta(\cdot)$ that captures aversion to ambiguity.[16,17]

2.5. *Learning and Ambiguity*

In the early stages of the literature, it was common to speculate (often erroneously) that ambiguity might be an asymptotically irrelevant phenomenon if agents were allowed to learn about their environment. For instance, it

[14]Kreps and Porteus' (1978) risk-sensitive preferences—that have received some attention in the asset pricing literature—have a functional representation similar to KMM. However Kreps and Porteus' preferences are not directly concerned with ambiguity, but are instead motivated by issues of time resolution of uncertainty in a classical SEU framework.

[15]Notice the similarity with the standard, SEU-style analysis of (local) risk aversion, as in Arrow and Pratt (1952).

[16]A few papers have debated whether KMM preferences may imply counterintuitive behaviors when the agent can bet directly on what the true model is. We refer the Reader to Epstein (2010) for two such examples and to Marinacci, Maccheroni and Mukerjie's (2011) rebuttal.

[17]For the purposes of this paper we do not explore any further alternative approaches towards ambiguity based on second order probabilities, that nonetheless constitute a rich and interesting literature.

was observed that in the classical Ellsberg paradox, any DM could benefit from repeated independent draws from the ambiguous urn. However, under non-additive probabilities, the notion of "independence" is not always well-defined, and, for particular experiments, it often coincides with the idea of "lack of known dependence" (see Dow and Werlang, 1994). Under this particular situation, uncertainty will fail to disappear because standard recursive drawing schemes might be applied without inducing any change in the agent's belief. In fact, models of ambiguity are known to cause a number of problems (see Etner et al., 2009) with updating algorithms, which interfere with the possibility "to learn ambiguity away".[18]

Epstein and Schneider (2007) have tackled these issues and formally introduced a learning model based on recursive MPP with three realistic features: first, people prefer risky bets to ambiguous ones with fixed composition in the short run, but not in the long run (so that Ellsberg-type behavior is observed in the short run only); second, risky bets are preferred to ambiguous bets with changing composition in both the short and the long run; finally, ambiguous bets with fixed compositions are always preferred to bets with changing composition. The proposed learning model allows DMs to express confidence about the changing environment and yet *ambiguity needs not vanish in the long run*: if some time-varying features of the model remain impossible to know even after many observations, then the agent moves towards a state of time-invariant ambiguity, where she has learnt all that she can.

Consider a setting characterized by a finite state space identical at all times, $\Omega_t = \Omega$, so that one state $s_t \in \Omega$ is observed in every period. At time t, the agent's information consists of the history $s^t = (s_1, ..., s_t)$. A standard Bayesian model of learning about a memoryless mechanism is summarized by a triplet (Ξ, μ_0, ℓ), where Ξ is a parameter space, μ_0 is a prior over it, and ℓ is the likelihood. Ξ collects features of the data-generating mechanism that the decision-maker is learning. For a given $\xi \in \Xi$, the data are an independent and identically distributed sequence of signals $\{s_t\}$. Beliefs are represented by a process $\{p_t\}$ of one-step-ahead conditionals that follow $p_t(\cdot|s^t) = \int_\Xi \ell(\cdot|\xi)\, d\mu_t(\xi|s^t)$, where μ_t is the posterior belief about ξ derived via Bayes' Rule. Ambiguity is introduced in the initial beliefs about parameters, represented by a set M_0 of probability measures on Ξ (and not by a unique μ_0), which reflects the decision-maker's (lack of) confidence in the prior information on which initial beliefs are based. A set of likelihoods L represents the agent's a priori view of the connection between signals and the true parameters. The multiplicity of likelihoods in L captures the complexity of the environment that prevents the agent to learn completely, since every parameter $\xi \in \Xi$ is associated with a set of probability measures, $L(\cdot|\xi)$. The dynamics of learning is summarized by a process of the one-step-ahead conditional beliefs

$$\wp_t(s^t) = \{p_t(\cdot) = \textstyle\int_\Xi \ell(\cdot|\xi) d\mu_t : \mu_t \in M_t^a(s^t), \ell \in L\}, \tag{10}$$

where $a \in (0, 1]$ is a parameter that governs the extent to which the decision-maker is willing to re-evaluate her views about how past data have been generated in the light of new sample information. To derive the updating rule that determines the set $\wp_t(s^t)$, once the sequence s^t has been observed, the elements of $M_t^a(s^t)$ have to be specified. To do so, consider a DM at time t looking back at the sample s^t and who entertains a number of different theories about how the sample has been generated. A theory is a pair (μ_0, ℓ^t), where μ_0 is a prior belief on Ξ and $\ell^t = (\ell_1, ..., \ell_t) \in L^t$ is a sequence of likelihoods. The decision-maker may re-evaluate her views about what sequence of likelihoods has been relevant for generating the data. To formalize re-evaluation, Epstein and Schneider suggest a two-step procedure. First, how well a theory (μ_0, ℓ^t) explains the data is captured by the data density evaluated at

[18] Pires (2002) analyses the effects of Bayesian updating on the relationship between conditional and unconditional MPP.

s^t, $\int \prod_{j=1}^{t} \ell_j(s_j|\xi)d\mu_0(\xi)$. Here conditional independence implies that the distribution given ξ is simply the product of the likelihood values ℓ_j, $j = 1, ..., t$. Prior information is taken into account by integrating out the parameters using the prior μ_0. The higher the likelihood of the data, the better is the observed sample s^t explained by the theory (μ, ℓ^t). Next, the posterior $\mu_t(\cdot; s^t, \mu_0, \ell^t)$ is derived from the theory (μ_0,ℓ^t) by Bayes' rule given the data s^t. $\mu_t(\cdot; s^t, \mu_0, \ell^t)$ can be calculated recursively, taking into account time variation in likelihoods:

$$d\mu_t(\cdot; s^t, \mu_0, \ell^t) = \frac{\ell_t(s_t|\cdot)}{\int_{\Xi} \ell_t(s_t|\xi^t)d\mu_{t-1}(\xi^t; s^{t-1}, \mu_0, \ell^{t-1})} d\mu_{t-1}(\cdot; s^{t-1}, \mu_0, \ell^{t-1}). \tag{11}$$

Recursive re-evaluation of the agent's beliefs takes the form of a likelihood-ratio test. The decision-maker discards all theories (μ_0, ℓ^t) that do not pass such a test against an alternative theory that puts maximum likelihood on the sample. Finally, posteriors $M_t^a(s^t)$ are formed only for theories that pass the test. This two-step process captures Epstein and Schneider's effort at integrating standard Bayesian learning within a MPP set up. Section 4 illustrates a few applications to key questions in finance.

3. **Optimal Consumption and Portfolio Decisions under Ambiguity**

3.1. *Trading Breakdowns and Limited Participation*

The seminal paper on the effects of ambiguity on portfolio choice is Dow and Werlang (1992), who studied a stylized two-period problem by considering a market with one ambiguous asset and one riskless bond. Their result is path-breaking: while under SEU, portfolio trades will occur *generically*, ambiguity may induce the existence of wide intervals for prices, for which the net demand of the asset is zero. Indeed, under CEU-preferences, there exists an interval of prices (and not a single value as under SEU) within which the agent neither buys nor sells short the risky asset. The intuition may be grasped from the following example. A risky asset pays off either $f_1 = 1$ or $f_2 = 3$, and the capacity v is such that $v(f_1) = 0.2$ and $v(f_2) = 0.4$. The investor is risk neutral, so that the expected payoff from *buying* a unit of the risky asset is given by $CE_v(\text{buy}) = 0.6 \times 1 + 0.4 \times 3 = 1.8$; this derives from the fact that in a max-min perspective, the weight is placed on a capacity of 0.4 in the good state and, as a result, of 0.6 in the bad state. On the other hand, the payoff from *selling* the asset is higher if the state underlying f_1 is realized; therefore $CE_v(\text{sell}) = 0.2 \times 1 + 0.8 \times 3 = 2.6$. Hence, for prices in the interval (1.8, 2.6), the investor would strictly prefer a zero position in the asset to either going short or long. This result is important to understand the existence of "market freezes", situations in which trading endogenously stops.

Easley and O'Hara (2009) also address the issue of under-participation in asset markets by proposing a simple portfolio choice problem in a standard normal-CARA framework under ambiguity. When agents' beliefs are compatible with a set of (normal) distributions characterized by parameters $\mu \in \{\mu_{\min}, ..., \mu_{\max}\}$ and $\sigma^2 \in \{\sigma^2_{\min}, ..., \sigma^2_{\max}\}$, if the price falls in the interval $[\mu_{\min}, \mu_{\max}]$, there will be no demand for the risky asset, because of ambiguity aversion. Hence, a region of possible beliefs for the expected payoff of risky assets may exist such that ambiguity averse investors abstain from trading and a limited participation equilibrium is possible. In a limited participation equilibrium, only SEU investors participate and trade the asset.

Guidolin and Rinaldi (2009) have extended the results in Easley and O'Hara (2009) to stress the distinction between two types of uncertainty: uncertainty that affects the entire market (*systematic* uncertainty) and uncertainty that just reflects circumstances peculiar to specific firms/stocks (*idiosyncratic* uncertainty). Guidolin and Rinaldi prove that a sufficient condition for ambiguity to induce market break-downs and limited participation equilibria is that the spread between the highest and the lowest possible return of the idiosyncratic risk component is *larger* than the spread between the highest and the lowest possible return of the systematic component.[19] Therefore it may not be ambiguity per se that causes limited participation equilibria, but instead the fact that markets tend to be characterized by much stronger ambiguity concerning idiosyncratic payoffs than ambiguity on the systematic ones.

These results on the relative importance of systematic vs. idiosyncratic risk in driving portfolio and limited participation results are consistent with Mukerji and Tallon (1999) who showed that, under CEU preferences characterized by a capacity v, if the range of variation of the idiosyncratic component in the asset's payoff is large relative to the range of variation of the systematic one, there exists an ambiguity level $\bar{\mathcal{A}}$, such that, if $\mathcal{A}(v) > \bar{\mathcal{A}}$, in equilibrium the investment level in each risky asset is zero for all investors, so that the market breaks down and agents are sub-optimally left exposed to aggregate risk. Instead, under SEU, the equilibrium allocation would approximate a complete market allocation, since usual diversification arguments imply that all idiosyncratic shocks wash away when the number of assets becomes arbitrary large. Under CEU preferences this result no longer holds and trading in well-diversified portfolios may collapse. The intuition is that the law of large number is still at work but agents' beliefs on the payoff of a well diversified portfolio converge to different values depending on their specific (short or long) position. Interestingly, it does not matter whether ambiguity concerns also aggregate, systematic, risk. Viceversa, the presence of an idiosyncratic component in assets' payoffs is necessary for trade to collapse, even if the financial market is already incomplete to begin with, and agents are ambiguity averse.[20]

3.2. *Mean-Variance Analysis Under Ambiguity*

The historical workhorse for the development of most modern finance theory has been the powerful mean-variance framework, in which SEU investors like expected portfolio returns and dislike variance—taken as a summary of risk—of portfolio returns (see e.g., Markowitz, 1959). Kogan and Wang (2003) have generalized the simple intuitions in Dow and Werlang (1992) and Easley and O'Hara (2009) to a general multivariate portfolio set-up, offering important insights on the implications of ambiguity (that they refer to as *model uncertainty*) for the cross-sectional properties of stock returns in a simple mean-variance framework. Consider a standard one-period representative agent economy, characterized by N risky assets and one riskless bond. Recall from Introductory Finance that under SEU and assuming that $\mathbf{R} \equiv [R_1, R_2, ..., R_N]'$ follows a joint multivariate normal distribution with known variance-covariance matrix $\mathbf{\Sigma}$ and known vector of means $\boldsymbol{\mu}$, a mean-variance investor will optimally hold the portfolio:[21]

$$\boldsymbol{\omega} = \frac{1}{\gamma}\mathbf{\Sigma}^{-1}(\boldsymbol{\mu} - r^f \boldsymbol{\iota}_N). \tag{12}$$

Kogan and Wang (2003) extend this result to the case in which the investor does not have perfect knowledge of the distribution of the returns $\mathbf{R}$. $\mathbf{R}$ still follows a multivariate normal distribution with known variance-covariance matrix $\mathbf{\Sigma}$, but the vector of mean returns, $\boldsymbol{\mu}$, is unknown. The agent displays MPP and some incomplete sources of information are available, so that she is able to estimate reference probabilistic models (hence, reference mean

[19] Under a few technical restrictions, they prove that—if the ambiguity concerns only the systematic risk component—then there will be always trade even when investors are ambiguity averse, which is the generic SEU outcome.

[20] Rinaldi (2009) has shown that these results are crucially related to the non-differentiability of the functional representation of CEU preferences. In particular, ambiguity is in general not sufficient to generate the no-trade result. However, for non-differentiable VP preferences a similar (more general) no-trade result holds.

[21] In what follows $\boldsymbol{\iota}_N$ is a $N \times 1$ vector of ones, so that $r^f \boldsymbol{\iota}_N$ is a $N \times 1$ vector that repeats r^f everywhere.

returns). Following the constraint approach (4), and assuming the agent is able to derive only a reference joint normal distribution of asset returns, $\hat{\mathbf{f}} \sim N(\hat{\boldsymbol{\mu}}, \boldsymbol{\Sigma})$, the set of priors $\wp(\hat{\mathbf{f}})$ is

$$\wp(\hat{\mathbf{f}}) = \left\{ q : E[z \ln z] \leq \eta \quad z \equiv \frac{dq}{d\hat{\mathbf{f}}} \right\}, \tag{13}$$

where η captures ambiguity aversion (a larger η means higher aversion), and $E[\cdot]$ is the expected value operator under the reference model $\hat{\mathbf{f}}$. In practice, the set $\wp(\hat{\mathbf{f}})$ constrains the statistical models for the vector process $\mathbf{R}$ to be "not too distant" from the benchmark $\hat{\mathbf{f}}$, with maximum distance given by η. Letting $\boldsymbol{\theta} \equiv \boldsymbol{\mu} - \hat{\boldsymbol{\mu}}$ be the divergence between one of the possible mean vectors under MPP and the vector of expected returns under the benchmark model, Kogan and Wang prove that the investment problem reduces to

$$\max_{\boldsymbol{\omega}} \min_{\boldsymbol{\theta} \in \left\{\boldsymbol{\theta} : \frac{1}{2}\boldsymbol{\theta}'\boldsymbol{\Sigma}^{-1}\boldsymbol{\theta} \leq \eta\right\}} E(z(\mathbf{R})U(W)) \quad z(\mathbf{R}) \equiv \exp\left\{ \frac{1}{2}\boldsymbol{\theta}'\boldsymbol{\Sigma}^{-1}\boldsymbol{\theta} - \boldsymbol{\theta}'\boldsymbol{\Sigma}^{-1}(\mathbf{R} - \boldsymbol{\mu} + \boldsymbol{\theta}) \right\}, \tag{14}$$

subject to the usual budget constraint, which is a transformation of the constraint on the set of admissible models under the parameter η into a (multiplicative) factor that appears in the objective function. Denoting by $\boldsymbol{\omega}^*$and $\boldsymbol{\theta}(\boldsymbol{\omega}^*) \equiv (\boldsymbol{\omega}^*)'(\boldsymbol{\mu} - \hat{\boldsymbol{\mu}})$ the optimal solutions to the problem, the following pricing condition holds in the absence of arbitrage opportunities,

$$E[U'(W - \Delta(\boldsymbol{\omega}^*))(\mathbf{R} - r^f \boldsymbol{\iota}_N - \boldsymbol{\theta}(\boldsymbol{\omega}^*))] = 0 \tag{15}$$

which modifies that standard Euler condition to account for the existence of ambiguity. Even though a closed-form solution for optimal portfolio weights fails to exist, using (15), and denoting by $\boldsymbol{\omega}_m$ the composition of the market portfolio, Kogan and Wang show that a basic two-factor beta representation holds:

$$\begin{aligned} \boldsymbol{\mu} - r^f \boldsymbol{\iota}_N &= \lambda\boldsymbol{\beta} + \lambda_u \boldsymbol{\beta}_u \qquad \lambda_u = \Delta(\boldsymbol{\omega}_m) \qquad \boldsymbol{\beta} = \frac{1}{\sigma_m^2}\boldsymbol{\Sigma}\boldsymbol{\omega}_m \qquad \lambda = \frac{E\left[U''(W - \Delta(\boldsymbol{\omega}_m))\right]}{E\left[U'(W - \Delta(\boldsymbol{\omega}_m))\right]}\boldsymbol{\omega}_m'\boldsymbol{\Sigma}\boldsymbol{\omega}_m \\ \boldsymbol{\beta}_u &= \frac{1}{\sigma_m^2}\boldsymbol{\Sigma}_u(\boldsymbol{\omega}_m)\boldsymbol{\theta}(\boldsymbol{\omega}_m) \qquad \text{where } \boldsymbol{\Sigma}_u(\boldsymbol{\omega}_m) \;\; s.t.\; \boldsymbol{\theta}(\boldsymbol{\omega}_m) = [\boldsymbol{\Sigma}_u(\boldsymbol{\omega}_m)]\,\boldsymbol{\omega}_m. \end{aligned} \tag{16}$$

This elegant result shows that the risk premium can be written as a two-factor representation: (i) the static CAPM component with structure $\lambda\boldsymbol{\beta}$, where λ is the market risk premium and $\boldsymbol{\beta}$ is the vector of betas measured with respect to returns on the market portfolio; (ii) a new component $\lambda_u\boldsymbol{\beta}_u$ where λ_u is the risk premium on ambiguity and $\boldsymbol{\beta}_u$ can be interpreted as a vector of betas the measure the exposure to the ambiguity contained in the market portfolio and quantified by $\boldsymbol{\theta}(\boldsymbol{\omega}_m) \equiv \boldsymbol{\omega}_m'(\boldsymbol{\mu} - \hat{\boldsymbol{\mu}})$. Ambiguity is only partially diversifiable in the sense that, in equilibrium and for all assets, only their individual contribution to total market ambiguity will be compensated. Finally, since the investor bears both risk and what they call "Knigthian/Keynesian uncertainty" (ambiguity), two assets with the same beta with respect to market risk may still have considerably different equilibrium expected returns, which Kogan and Wang interpret as highly realistic in the light of the voluminous literature on the empirical shortcomings of the static CAPM.

Garlappi, Uppal, and Wang (2007, GUW) have extended Kogan and Wang (2003) and modelled ambiguity through a MPP approach, using a "confidence interval" framework that is popular in the literature. GUW's starting point is that in reality the parameters of the joint conditional density characterizing the ambiguous asset returns $\mathbf{R}$ have to be estimated and are themselves random. If an investor solves the classical mean-variance problem taking parameter uncertainty into account, the resulting portfolio will be of a Bayesian type (see, e.g., Barberis, 2000). However, the portfolios in which only parameter uncertainty is considered often perform poorly out of sample, even in comparison to portfolios selected according to ad hoc rules. One reason for this result is that the vector of expected asset returns $\boldsymbol{\mu}$ is hard to estimate with any precision. This induces GUW to introduce ambiguity on the statistical model, as identified by the vector of expected returns $\boldsymbol{\mu}$. Under MPP, GUW (2007) prove that the optimization problem is equivalent to the simpler problem

$$\max_{\boldsymbol{\omega}} \boldsymbol{\omega}'(\hat{\boldsymbol{\mu}} - \boldsymbol{\mu}^{adj}) - \frac{1}{2}\gamma\boldsymbol{\omega}'\boldsymbol{\Sigma}\boldsymbol{\omega}, \tag{17}$$

where $\hat{\boldsymbol{\mu}} - \boldsymbol{\mu}^{adj}$ is the "adjusted estimated expected return". The adjustment $\boldsymbol{\mu}^{adj}$ has the role to incorporate ambiguity, and it depends on the precision with which parameters are estimated, the length of the data series, and the investor's aversion to ambiguity. As $\boldsymbol{\mu}^{adj}$ depends on $\boldsymbol{\omega}$, the problem has to be solved numerically.

Kogan and Wang (2003) and GUW are mean-variance papers under MPP in which ambiguity concerns some features of simple *statistical models*—such as a multivariate Gaussian distribution for asset returns—that are typical in finance. Wang (2005) has extended this line of work to the case in which ambiguity concerns the *asset pricing models* that are potentially generating returns, and has pursued the portfolio implications of such a framework. In particular, Wang considers the structure of equity portfolios when there is ambiguity on whether returns are generated by a Fama-French three-factor model (in which portfolios mimicking the exposure to size and book-to-market "risks" act as factors additional to the CAPM market portfolio) and in which ambiguity aversion is captured through MPP. His results show that, when aversion to model uncertainty is incorporated, the optimal portfolio is very different from the market portfolio prescribed by the CAPM, and also very different from the portfolio based on the unrestricted Bayesian sample estimate.

Key insights on how the mean-variance framework extends to the case of KMM preferences has been offered by Maccheroni, Marinacci and Ruffino (2010, MMR). They derive the analogue of the classic approximation, $CE(W + h, \mathbb{P}) \cong W + E_{\mathbb{P}}[h] - 0.5\lambda_u(W)\sigma^2_{\mathbb{P}}(h)$, for the certainty equivalent under KMM preferences

$$CE(W + h) \cong W + E_{\bar{\mathbb{Q}}}[h] - \frac{1}{2}\lambda_u(W)\sigma^2_{\bar{\mathbb{Q}}}(h) - \frac{1}{2}[\lambda_\zeta(W) - \lambda_u(W)]\sigma^2_\mu(E_{\mathbb{Q}}[h]), \tag{18}$$

where $\lambda_u(W) \equiv -u''(W)/u'(W)$ is local absolute risk aversion, h is a (self-financing) investment, $\mathbb{P}$ is the probabilistic model the agent uses to describe the stochastic nature of the problem, $CE(W + h)$ is the random variable that associates $CE(W + h, \mathbb{Q})$ to each model $\mathbb{Q}$ in the set of possible probability measures (Δ) that capture the presence of model uncertainty, the KMM function $\zeta(\cdot)$ measures the attitude toward model uncertainty, and μ is the investor's prior probability on the space of possible models Δ. $\bar{\mathbb{Q}}$ is the reduced probability $\int \mathbb{Q}d\mu(\mathbb{Q})$ induced by the prior $\mathbb{Q}$; $E[h] : \Delta \to \mathbb{R}$ is a random variable that associates the expected value $E_{\mathbb{Q}}[h]$ to each possible model $\mathbb{Q}$. The variance of $E_{\mathbb{Q}}[h]$, $\sigma^2_\mu(E[h])$, along with the difference $\lambda_\zeta(W) - \lambda_u(W)$ in uncertainty attitudes, determines an ambiguity premium that is novel relative to the classical result. Because the Arrow-Pratt approximation represents the microeconomic the foundation for the classical mean-variance model, (18) allows MMR to extend the standard mean-variance model to a framework of choice based on the maximization of the objective

$$V(W) = E_{\bar{\mathbb{Q}}}[W] - \frac{\lambda}{2}\sigma^2_{\bar{\mathbb{Q}}}(W) - \frac{\theta}{2}\sigma^2_\mu(E_{\mathbb{Q}}[W]), \tag{19}$$

where $\lambda \equiv \lambda_u(W)$, $\theta \equiv \lambda_\zeta(W) - \lambda_u(W)$, which represent negative attitudes towards risk and ambiguity, respectively. This augmented mean-variance model is determined by the three parameters λ, θ, and μ, as opposed to the two parameters λ and $\bar{\mathbb{Q}} = \mathbb{P}$ of the standard mean-variance model; μ can be considered an information parameter because higher values of $\sigma^2_\mu(E_{\mathbb{Q}}[W])$ correspond to poorer information on prospect's outcomes and on models. In particular, in a simple problem with a risk-free asset, a risky asset, and an ambiguous asset, they show that portfolio rebalancing in response to higher model uncertainty only depends on a special ambiguous asset's "alpha measure" that measures the excess return performance of the ambiguous asset that cannot be captured by its beta on the risky asset. Following the standard practice of considering the risky asset as a benchmark, this special alpha measure captures the residual performance of the ambiguous asset that cannot be explained in terms of risk only. Thus, an ambiguity averse agent that observes a positive alpha attributes the otherwise unexplained, augmented return to an unmeasurable increase in uncertainty (ambiguity) that drives him away from the ambiguous asset. Analogously, a negative is alpha associated with a diminution of uncertainty that, in turn, makes the ambiguous asset more desirable.

3.3. *Robust Portfolio Decisions*

The intuitive desire for robust portfolios, i.e., portfolios that may be sub-optimal under the selected reference model, but that may perform well under similar models, is obviously suggestive of the possibility of studying asset allocation decisions within an explicit dynamic framework that incorporates the robustness ideas by Hansen, Sargent, and their coauthors. One of the early examples of this strategy is Trojani and Vanini (2002), who have used the constraint framework (4) of Anderson et al. (2003) to analyze a robust version of the classical Merton's (1971) model in continuous time. In their set up, a representative agent can invest either in a riskless asset, with risk-free rate r^f, or in a risky asset, whose price dynamics is described by a simple Geometric BM with constant parameters μ and σ. For any given initial wealth-price pair (W_0, S_0), let the operator $\{T_t\}_{t\geq 0}$ be such that $T_t(\phi(W_0, S_0)) = E[\phi(W_t, S_t) \,|\, (W_0, S_0)]$. The operator process $\{T_t\}$ allows us to define a continuum of probabilistic models by "contamination" of a baseline, approximating model characterized by a non-negative random variable $\theta \geq 0$, so that the corresponding operators $\left(T_t^\theta\right)_{t\geq 0}$ are $T_t^\theta(\phi) = T_t(\theta\phi)/T_t(\theta)$. The relative entropy of the θ-model with respect to the original one is given by:

$$\mathcal{I}_t(\theta) = T_t\left(\frac{\theta}{T_t(\theta)} \ln\left(\frac{\theta}{T_t(\theta)}\right)\right). \tag{20}$$

As typical in the robust control literature, a preference for robustness is induced by choosing from a set of admissible θ-models the one that minimizes expected utility, where the admissible set is defined by constraining the relative entropy computed from the reference model to be lower than η. Hence, letting ω_t be the weight assigned to the risky asset, the infinite-horizon investor's optimization problem is

$$\max_{c,\omega} \min_{\theta} E\left[\int_0^\infty \exp(-\delta s)\, U(\omega_s) ds \middle| (W_0, S_0)\right] \qquad \text{s.t.: } \lim_{t\to 0} \frac{\mathcal{I}_t(\theta)}{t} \leq \eta, \tag{21}$$

subject to the standard dynamic processes for the stock price and for wealth. The constrained minimization in θ transforms the optimization problem into an equivalent program in which the drift of S_t is simply distorted by $-\sqrt{2\eta}\sigma$. Assuming positive risk premia and an isoelastic utility index $U(c) = c^{1-\gamma}/(1-\gamma)$ (with $\gamma \in (0,1)$), Trojani and Vanini derive the optimal robust solutions:

$$\begin{aligned}
\omega^* &= \frac{\mu - r^f - \sqrt{2\eta}\sigma}{\sigma^2\gamma} < \frac{\mu - r^f}{\sigma^2\gamma} = \omega^{SEU} \\
\frac{c^*}{W} &= \left[\frac{\delta}{\gamma} - (1\text{-}\gamma)\left(\frac{(\mu - r^f - \sqrt{2\eta}\sigma)^2}{2\sigma^2\gamma^2} + \frac{r^f}{\gamma}\right)\right]^{-1/\gamma} > \left[\frac{\delta}{\gamma} - (1\text{-}\gamma)\left(\frac{1}{2}(\omega^{SEU})^2 + \frac{r^f}{\gamma}\right)\right]^{-1/\gamma} = \frac{c^{SEU}}{W},
\end{aligned} \tag{22}$$

i.e., under ambiguity aversion (here, preference for robustness) and $\gamma \in (0,1)$, the share of the risky asset will be lower than the share that would be optimal under SEU ($\eta = 0$), while the investor will consume a higher fraction of her wealth. Moreover, $\partial\omega^*/\partial\eta < 0$ so that additional caution appears to be induced by aversion to ambiguity; in a sense, it is as if $\eta > 0$ helps mimicking the optimal portfolio selection of an investor that is more risk averse than what can be capture by the parameter γ in her utility index.

Maenhout (2004) has further explored the effects of model uncertainty on dynamic portfolio/consumption decisions by solving a problem similar to Trojani and Vanini's (2002) with the difference that the investor has a finite horizon T, $\gamma > 0$ is unconstrained, while the penalty framework (3) of Anderson et al. (2003) is used to derive consumption-portfolio rules that perform reasonably well under model misspecification. In particular, preferences are characterized by the isoelastic utility index $c_t^{1-\gamma}/(1-\gamma)$ and ambiguity aversion index φ.[22] Maenhout notices that a major drawback of Hansen-Sargent (HS) preferences is the lack of homotheticity in case of a power-utility index, so that the resulting optimal portfolio is not independent of wealth. Preferences are homothetic if $\forall\chi > 0$, $x \sim y \iff \chi x \sim \chi y$. To overcome this problem, the ambiguity aversion coefficient φ is re-scaled at any time-state pair, and replaced by $\varphi_{s,t} \equiv \varphi/\left[(1-\gamma)V(s,t)\right]$, where V is the value function defined as usual. Hence, the necessary conditions for an optimum are:

$$\omega^* = \frac{1}{\varphi+\gamma}\frac{\mu - r^f}{\sigma^2} < \frac{\mu - r^f}{\sigma^2\gamma} = \omega^{SEU} \qquad \frac{c_t^*}{W_t} = \frac{a}{1-e^{-a(T-t)}} \qquad a \equiv \frac{1}{\gamma}\left[\delta - (1\text{-}\gamma)r^f - \frac{1\text{-}\gamma}{2(\varphi+\gamma)}\left[\frac{\mu - r^f}{\sigma^2}\right]^2\right]. \tag{23}$$

As long as $\varphi > 0$, the optimal share of the risky asset ω^* will differ from the standard SEU demand function, a result qualitatively similar to Trojani and Vanini's, even though the specifics of the expression for ω^* are different because of the different (finite horizon) nature the problem and the way in which robustness is modelled. The simulations in Maenhout (2004) stress that a preference for robustness may dramatically decrease the portfolio demand for equities and lead to an additional hedging-type demand, so that the equilibrium equity premium increases.[23]

3.4. *Does Constraint vs. Penalty Matter?*

Section 3.3 has emphasized that adopting the HS-type constraint framework (4) or the penalty framework (3) may make a difference for investment and consumption rules. Trojani and Vanini (2004) have carefully investigated these differences and compared the effects of ambiguity aversion on asset pricing in the two robustness-frameworks. They re-consider the model in Maenhout (2004) under the constraint problem in Anderson et al. (2003) noting that two frameworks differ in the intertemporal structure of the representative agent's attitude towards ambiguity. Trojani and Vanini's key insight is that the specific way in which ambiguity aversion is modelled has no major qualitative implications for optimal consumption choices but affects in different ways the portfolio rules. In both frameworks, the asset menu is composed of a riskless discount bond that yields r^f and an ambiguous stock with ex-dividend price S. x represents the exogenous primitives that affect the opportunity set of the agent. Ambiguity aversion is modelled in the form of a contamination vector, $\boldsymbol{\theta}_t \equiv [\theta_t^x \ \theta_t^S]'$ that perturbs a BM B driving the state x. The problem is

$$\sup_{c,\omega}\inf_{\boldsymbol{\theta}} E^c\left[\int_0^\infty e^{-\delta t}\left[\frac{(cW)_t^{1-\gamma}}{1-\gamma} + \frac{1}{2\varphi(W_t,x_t)}\|\boldsymbol{\theta}_t\|^2\right]dt\right], \tag{24}$$

where $\varphi > 0$ is the state- and wealth-dependent ambiguity aversion function. As in Maenhout (2004), $\varphi \to 0^+$ (from the right) implies that any perturbations brought about by $\boldsymbol{\theta}_t$ will carry enormous penalties, so that the standard

[22] In Maenhout (2004), φ is the reciprocal of ϱ in (3), hence higher ambiguity aversion is represented by higher values of φ.

[23] Considering a mean-reverting risk premium, Liu (2010) shows that ambiguity aversion can instead increase the demand for the risky asset when the elasticity of intertemporal substitution is close to 1.

SEU framework obtains; on the opposite, if $\varphi \to \infty$, there is infinite aversion to ambiguity. The second, Anderson et al. (2003)-style setting is based on a robust control constraint

$$\sup_{c,\omega} \inf_{\boldsymbol{\theta}} E^c \left[\int_0^\infty e^{-\delta t} \frac{(cW_t)^{1-\gamma}}{1-\gamma} dt \right] \qquad \text{s.t. } \frac{1}{2}\boldsymbol{\theta}'\boldsymbol{\theta} \leq \eta, \tag{25}$$

The larger η, the higher aversion to ambiguity. $\eta = 0$ implies SEU consumption and portfolio choices.

The function $\varphi(W_t, x_t)$ that solves (24) cannot be computed in closed form. However, Trojani and Vanini show that optimal consumption is not affected by ambiguity aversion while optimal investment is. In particular, the consumption rule can be interpreted as the optimal consumption of a SEU-investor with elasticity of intertemporal substitution $1/\gamma$, while the investment policy ω^* corresponds to the investment policy of a SEU-investor with relative risk aversion artificially increased to $\gamma + \varphi$. The optimal investment ω^* is the sum of the myopic demand and the hedging demand. Comparing both demands to the standard ambiguity-free, SEU case (with $\omega^*_{myopic}(0)$ and $\omega^*_{hedge}(0)$), one gets:

$$\omega^*_{myopic}(\varphi) = \frac{\gamma}{\gamma+\varphi}\frac{1}{\gamma}\frac{\mu_S - r^f}{\sigma_S^2} = \frac{\gamma}{\gamma+\varphi}\omega^*_{myopic}(0) < \omega^*_{myopic}(0) \qquad \omega^*_{hedge}(\varphi) = \frac{\gamma}{\gamma+\varphi}\frac{1-\gamma+\varphi}{1-\gamma}\omega^*_{hedge}(0). \tag{26}$$

Ambiguity aversion reduces myopic portfolio exposures because of the increased effective risk aversion parameter $\gamma + \varphi$. The term $(1-\gamma+\varphi)/(1-\gamma) > 1$ in the hedging demand arises because, under the worst case probability, the dynamics of x is different from the reference one. The term $\gamma/(\gamma+\varphi) < 1$ measures the reduction in hedging demand required by the lower baseline myopic allocation.[24]

The value function that solves (25) can be expressed as a function $G(x)$ that cannot be computed in closed form. Again, optimal consumption is not affected by ambiguity, while the optimal investment policy is. However, the linearized portfolio rules are not identical. Trojani and Vanini obtain that in a constraint-style problem the portfolio rule can be re-interpreted as the optimal strategy of a SEU investor with an effective risk aversion of $\gamma + \sqrt{2\eta/G(x)}$, which is made state-dependent via the function $G(x)$. This is a key insight: in a constraint-style framework, ambiguity aversion affects optimal portfolio choices in a non-uniform way over the relevant support of x. The myopic and hedging demand functions are related to those of the ambiguity free case ($\omega^*_{myopic}(0)$ and $\omega^*_{hedge}(0)$) as follows:

$$\omega^*_{myopic}(\eta) = \omega^*_{myopic}(0)\frac{\gamma}{\gamma+\sqrt{2\eta/G(x)}} < \omega^*_{myopic}(0) \quad \omega^*_{hedge}(\eta) = \omega^*_{hedge}(0)\frac{\gamma}{\gamma+\sqrt{2\eta/G(x)}}\frac{1-\gamma+\sqrt{2\eta/G(x)}}{1-\gamma}, \tag{27}$$

and while ambiguity aversion structurally reduces myopic demands, the net effect on hedging demands depends on the balance of the terms $\gamma/(\gamma+\sqrt{2\eta/G(x)}) < 1$ and $(1-\gamma+\sqrt{2\eta/G(x)})/(1-\gamma) > 1$.

3.5. *Applications to Large-Scale Problems: Can Ambiguity Explain the Home Country Bias?*

Uppal and Wang (2003) showed that it is possible to use ambiguity aversion to solve more realistic portfolio problems than those considered in Trojani and Vanini (2002, 2004) and Maenhout (2004) and to shed light on unresolved questions in empirical asset allocation. In particular, Uppal and Wang solve an intertemporal portfolio choice problem characterized by investors with a preference for HS robustness who display different levels of ambiguity across different assets, and derive implications for the so-called *home country bias puzzle*, the tendency of investors

[24] Whether total hedging demand under ambiguity is increased or decreased by taking ambiguity into account will eventually depend on the relative strength of the two effects, i.e., on whether the term $[\gamma(1-\gamma+\varphi)]/[(1-\gamma)(\gamma+\varphi)] > 1$ or not.

to bias their portfolio decisions in favor of assets issued in their countries and away from optimal allocation weights implied by standard SEU models (e.g., the mean-variance framework); see e.g., Lewis (1999) for a review of the literature. Uppal and Wang's approach is simple. When investors have a preference for robustness, the standard (dynamic) recursive SEU representation is replaced by

$$V_t(c_t, \boldsymbol{\omega}_t) = U(c_t) + \beta \inf_z \left\{ \phi\left(E_t^z\left[V_{t+1}(c_{t+1}, \boldsymbol{\omega}_{t+1})\right]\right) \varrho \Psi\left(q\left(z\right)\right) + E^z\left[V_{t+1}(c_{t+1}, \boldsymbol{\omega}_{t+1})\right] \right\}, \tag{28}$$

where $\varrho \geq 0$ is an ambiguity aversion parameter, $\phi(\cdot)$ is a normalizing factor, z is the Radon-Nikodym derivative with respect to the reference probabilistic model p, $\Psi\left(q\left(z\right)\right)$ is the relative entropy of the distribution $q(z)$ vs. p, and $q(z)$ is such that $dq(z) = zdp$. As in Section 2.3, as $\varrho \to \infty$ ($\to 0^+$, i.e. zero from the right) ambiguity declines (increases).

Uppal and Wang prove that when $\varrho \to \infty$, the optimal risky portfolio is identical to the standard Merton's portfolio; when $\varrho \to 0$, the investment in the risky asset drops to zero. This means that when ambiguity diverges, pervasive and non-diversifiable model uncertainty will advise an investor to avoid all security risks. More generally, when $\varrho > 0$ but is finite, the total investment in the risky assets is less than what it would be in absence of ambiguity, which can be interpreted as a multivariate generalization of Maenhout (2004). Uppal and Wang consider the standard case in which the investor can consume a single good, invests in N risky stocks, and borrows or lends at an exogenously given riskless rate $r_t^f = r(\mathbf{Y}_t)$, where $\mathbf{Y}_t$ is the state of the economy. The return process of the N stocks is

$$d\mathbf{R}_t = \boldsymbol{\mu}_R(\mathbf{R}_t, \mathbf{Y}_t)dt + \boldsymbol{\sigma}_R(\mathbf{R}_t, \mathbf{Y}_t)d\mathbf{B}_t \qquad d\mathbf{Y}_t = \boldsymbol{\mu}_Y(\mathbf{Y}_t)dt + \boldsymbol{\sigma}_Y(\mathbf{Y}_t)d\mathbf{B}_t, \tag{29}$$

with $\mathbf{B}_t$ a K-dimensional BM. Uppal and Wang write the investors' indirect utility function as $V(W_t, \mathbf{R}_t, \mathbf{Y}_t, t)$ and use an appropriate drift adjustment to specify the HJB equation for the investor's utility maximization problem, which is the standard SEU Bellman equation augmented by the penalty function for the robustness-driven distortions and three drift adjustments due to the change of measure.

Uppal and Wang's key result is that the optimal portfolio rule differs from the standard Merton's (1971) formula because of ambiguity aversion (indeed when the coefficients ϱ_i tends to infinity for $i = 1, 2, ..., N$, we obtain the familiar Merton's result) and the consequent drift adjustment to $\boldsymbol{\mu}_R$. Moreover, a number of intriguing results can be derived when ambiguity is the same across all assets, $\varrho_i = \varrho$ $\forall i$, all moments are constant, there is no predictability, and the utility index is negative exponential with (absolute) risk aversion coefficient $\alpha > 0$. In this case, the optimal multivariate portfolio shares will be identical to those implied by an optimal Merton's portfolio, whose the investor's risk aversion α has been replaced by (magnified to) $\alpha(1 + 1/\varrho)$. This means that ambiguity generally reduces the demand of all risky assets, and that this is isomorphic to a generalized increase in the degree of aversion to risk, as measured by the factor $(1+1/\varrho) > 1$. They also calibrate their model to data on international equity returns from French and Poterba's (1991) paper on the empirics of the home country bias and compared the resulting assets allocations to the standard mean-variance style Merton portfolios. The calibration shows that when the overall ambiguity about the joint distribution of returns is high, then small differences in ambiguity for the marginal return distributions of national, country-specific stock return series will result in a portfolio that is significantly under-diversified relative to the standard mean-variance portfolio. This is a powerful explanation of the home country bias: modest heterogeneity in the ambiguity perceived by investors across different national asset markets may induce massive and yet optimal portfolio biases which are consistent with the data.

Uppal and Wang's (2003) intuition that small amounts of heterogeneity in the ambiguity investors perceive "across" assets may cause important effects has been further explored by Boyle, Garlappi, Uppal, and Wang (2009, henceforth BGUW) who study the role of ambiguity (or assets' "familiarity") in determining portfolio under-diversification and "flight to familiarity" episodes. Consider a mean-variance portfolio problem in which the asset menu is composed of N identical risky assets and one riskless asset. Each asset has (unknown) expected excess return μ_i, and common volatility σ. ρ is a common correlation coefficient across assets and $\mathbf{\Sigma}$ is the implied covariance matrix. Using the framework developed by Garlappi et al. (2007), BGUW write the optimization problem as

$$\max_{\boldsymbol{\omega}} \ \min_{\boldsymbol{\mu}} \boldsymbol{\omega}'\boldsymbol{\mu} - \frac{\gamma}{2}\boldsymbol{\omega}'\mathbf{\Sigma}\boldsymbol{\omega} \qquad \text{s.t.: } \frac{(\mu_i - \hat{\mu}_i)^2}{\sigma^2_{\hat{\mu}_i}} \leq \eta_i \quad \boldsymbol{\omega}'\boldsymbol{\iota}_N = 1, \tag{30}$$

where $\boldsymbol{\mu}$ is a N-dimensional vector that collects expected returns, μ_i, $i = 1, ..., N$. Under this specification, the ambiguity problem can be interpreted in terms of classical statistical analysis because—letting $\hat{\mu}_i$ be the estimated value of the mean return of asset $i = 1, ..., N$ obtained by using a return time series of length T, and $\sigma^2_{\hat{\mu}_i}$ the variance of $\hat{\mu}_i$—it is possible to define the confidence interval $\{T(\mu_i - \hat{\mu}_i)^2/\sigma^2_{\hat{\mu}_i} \leq \eta_i\}$ for expected returns. Hence, $\sqrt{\eta_i}$ is the critical value determining the size of the confidence interval that can be interpreted as a measure of the amount of ambiguity about the estimate of expected returns. A larger $\sqrt{\eta_i}$ determines a larger confidence interval and a larger set of possible distributions to which the true returns may belong. Without loss of generality, if we assume that the first asset is more familiar to the DM, and the remaining $N-1$ assets are all equally less familiar (i.e., $\eta_1 = \eta_F$ and $\eta_2 = ... = \eta_U$, $\eta_F < \eta_U$), the structure of the portfolio allocations will depend on a condition involving the quantities $\hat{\mu}_i/\sigma_{\hat{\mu}_i}$ and $\sqrt{\eta_F} + (1-\rho)^{-1}(\sqrt{\eta_U} - \sqrt{\eta_F})$, which can be interpreted as a scaled measure of "excess" unfamiliarity". To simplify the notation, let's assume that $\hat{\mu}_i/\sigma_{\hat{\mu}_i} = \hat{\mu}/\sigma_{\hat{\mu}}$. When $\hat{\mu}/\sigma_{\hat{\mu}} > \sqrt{\eta_F} + (1-\rho)^{-1}(\sqrt{\eta_U} - \sqrt{\eta_F})$ (that is, ambiguity is relatively low when compared to the reward-to-risk ratio), then

$$\begin{aligned}
\omega_1 &= \frac{1}{N} + \sigma_{\hat{\mu}} \frac{\left(\sqrt{\eta_U} - \sqrt{\eta_F}\right)\left[1 - \frac{1}{N} + \frac{\rho}{1-\rho}(N-1)\right]}{N\left[\hat{\mu}_1 - \sigma_{\hat{\mu}}\sqrt{\eta_F} - \sigma_{\hat{\mu}}(\sqrt{\eta_U} - \sqrt{\eta_F})\left(1 - \frac{1}{N}\right)\right]} > \frac{1}{N} \\
\omega_i &= \frac{1}{N} - \sigma_{\hat{\mu}} \frac{\left(\sqrt{\eta_U} - \sqrt{\eta_F}\right)\left[\frac{1}{1-\rho} - \left(1 - \frac{1}{N}\right)\right]}{N\left[\hat{\mu}_i - \sigma_{\hat{\mu}}\sqrt{\eta_F} - \sigma_{\hat{\mu}}\left(\sqrt{\eta_U} - \sqrt{\eta_F}\right)\left(1 - \frac{1}{N}\right)\right]} < \frac{1}{N} \quad i \geq 2,
\end{aligned} \tag{31}$$

and $\omega_1 > 1/N > \omega_i$, $i = 2, ..., N$. If $\sqrt{\eta_F} < \hat{\mu}/\sigma_{\hat{\mu}} \leq \sqrt{\eta_F} + (1-\rho)^{-1}(\sqrt{\eta_U} - \sqrt{\eta_F})$ (that is, ambiguity is relatively low for asset 1 and relatively high for all other assets), then a 100% share goes into the familiar asset. Finally, if $0 < \hat{\mu}/\sigma_{\hat{\mu}} \leq \sqrt{\eta_F}$ (that is, ambiguity is relatively high for all the assets), then $\omega_1 = 0 = \omega_i$ $(i \geq 2)$ and 100% of wealth ought to be invested in the riskless asset. In this case, ambiguity is so high across the board (including the case of familiar assets) that complete non-participation arises. Interestingly, in the first case the weight on the familiar asset exceeds that of each of the other assets. So, the investor holds familiar assets (as first advocated by Keynes) but balances this investment by holding also a portfolio of all the other assets (as advocated by Markowitz) which remains biased toward more familiar assets. The relative investment between the familiar asset and the unfamiliar portfolio depends on the relative ambiguity, as captured by the term $\sqrt{\eta_U} - \sqrt{\eta_F}$. When the number of assets becomes arbitrarily large, the weight for the more familiar asset approaches a positive constant that depends on $\sqrt{\eta_U} - \sqrt{\eta_F}$. In contrast, the weights for each of the other less-familiar assets approach zero as the number of assets increases, while the total weight in these $(N-1)$ unfamiliar assets approaches a positive fraction. Hence, familiarity of a specific asset implies that its holdings do not decrease to 0, even if $N \to \infty$, while the existence of gains from diversification implies that an investor should hold only an infinitesimal amount in each of the remaining unfamiliar assets.

3.6. *Smooth Recursive Preferences*

Chen, Ju and Miao (2011, CJM) capture ambiguity aversion through a generalization of KMM's recursive smooth preferences, in a discrete time-infinite horizon setting. Suppose that x_t is the time-t realization of a state variable x that determines the set of effective measures to be considered. The asset menu is composed of a risky asset and a risk-free bond with returns $r_{e,t+1}$ and $r_{f,t+1}$, respectively. Two possibly misspecified models of stock returns are considered. The first is characterized by identically and independently distributed (IID) returns and implies the absence of predictability; the second model is of a vector autoregressive (VAR) type, with a single zero-mean predictor variable for excess stock returns:

$$\text{Model 1(IID): } r_{e,t+1} - r_{f,t+1} = r + \varepsilon_{1,t+1} \quad \varepsilon_{1,t+1} \sim \mathcal{N}\left(0, \sigma_1^2\right) \tag{32}$$

$$\text{Model 2 (VAR): } \begin{cases} r_{e,t+1} - r_{f,t+1} = r + bx_t + \varepsilon_{2,t+1} \\ x_{t+1} = \rho x_t + \varepsilon_{3,t+1} \end{cases} \quad [\varepsilon_{2,t+1}, \varepsilon_{3,t+1}]' \sim \mathcal{N}(0, \mathbf{\Omega}) \quad \mathbf{\Omega} = \begin{bmatrix} \sigma_2^2 & \sigma_{23} \\ \sigma_{23} & \sigma_3^2 \end{bmatrix}.$$

The investor is uncertain as to which model is correctly specified, even when she can learn the relative merits of the two processes from past data. In each period t the investor decides how much to consume and how much to invest in the risky market portfolio, and at time T she will consume her total wealth. Beliefs evolve as follows: if the predictive variable is above average ($x_t > 0$), the VAR model predicts above average excess returns. Assuming that the volatility of returns are similar in the two models, a high realized return will be more likely in the VAR model than in the IID model. Thus, the investor revises downward her belief about the IID model. The opposite will occur if the predictive variable is below average ($x_t < 0$).

CJM calibrate their model using annual data from the U.S. stock market; preference parameters are set through hypothetical experiments a' la Ellsberg. The ambiguity-averse investor's optimal investment strategy is intuitively a robust strategy, and the calibration compares it to an IID strategy (i.e., when the investor is sure that the IID model is the data generating process), a VAR strategy, and a Bayesian strategy (which ignores ambiguity imposing a linear $v(\cdot)$). The results show that the robust stock allocation depends on the investment horizon, the beliefs about the model for stock returns, and the predictor. Compared to the Bayesian strategy, the robust portfolio rule is more conservative, inducing a lower rate of participation in the stock market. Indeed, both Bayesian and robust portfolio weights can be decomposed into a myopic demand and an intertemporal hedging demand. The former depends on the one-period ahead expected return, which is a weighted average of the expected excess returns from the two competing models 1 and 2. The hedging demand can be further decomposed into two components: the first is associated with variations in investment opportunities captured by the predictive variable; the second is the demand that hedges model uncertainty. While high (low) realized returns lead the Bayesian investor to shift her posterior beliefs towards (away from) the VAR model when the predictive variable is large and positive (small and negative), which generally makes his hedging demand negative (positive), the ambiguity-averse investor makes investment decisions using endogenously distorted beliefs, and not the actual predictive distribution. For a given non-degenerate prior, the distortion in beliefs is large when model uncertainty is high, so that both the myopic and hedging demands implied by the robust strategy can be quite different from those implied by the Bayesian strategy.

3.7. *Incomplete Information vs. Ambiguity*

Recently, a few researchers have started to introduce ambiguity in incomplete information problems with at least two questions in mind. First, when and how will it be possible to characterize the effects of both incomplete information and ambiguity on asset allocation and consumption decisions. Second, whether there can be any "interaction" or compounding effects between hedging demands caused by incomplete information and by ambiguity aversion. A few papers have made important steps addressing these questions.

Miao (2009) has studied consumption/portfolio choices under incomplete information and ambiguity, assuming that agents display dynamic κ-ignorance specification MPP as in Chen and Epstein (2001, see Section 4.2). In his framework, the κ-ignorance specification leads to a density generator (vector) process $\boldsymbol{\theta}$ that induces a distorted probability measure q^{θ}. Specifically, denoting by $\boldsymbol{\theta}^*$ the generator that corresponds to the optimal consumption process c^*, Miao shows that the agent's perceived return dynamics is

$$d\mathbf{R}_t = \left(\hat{\boldsymbol{\mu}}_t^R - \boldsymbol{\sigma}^R\boldsymbol{\theta}_t^*\right)dt + \boldsymbol{\sigma}^R d\hat{\mathbf{B}}_t, \tag{33}$$

where $\boldsymbol{\mu}^R$ and $\boldsymbol{\sigma}^R$ are the parameters of the multivariate GBM for stocks. Hence two factors influence the deviations of the agent's perceived mean returns from their true values:

$$\boldsymbol{\mu}^R - \left(\hat{\boldsymbol{\mu}}_t^R - \boldsymbol{\sigma}^R\boldsymbol{\theta}_t^*\right) = \left(\boldsymbol{\mu}^R - E\left[\boldsymbol{\mu}^R|\mathcal{F}_t^S\right]\right) + \boldsymbol{\sigma}^R\boldsymbol{\theta}_t^*. \tag{34}$$

The two terms above reflect to separate hedging motives: the first corresponds to what Miao's defines *estimation* (incomplete information) *risk*, and the second to ambiguity.

Under power utility, the optimal portfolio process is determined to be

$$\boldsymbol{\omega}_t^* = \frac{1}{\gamma}\left(\boldsymbol{\sigma}^R(\boldsymbol{\sigma}^R)'\right)^{-1}\left(\hat{\boldsymbol{\mu}}_t^R - r^f\boldsymbol{\iota}_d\right) - \frac{1}{\gamma}\left((\boldsymbol{\sigma}^R)'\right)^{-1}\boldsymbol{\theta}_t^* + \left((\boldsymbol{\sigma}^R)'\right)^{-1}\boldsymbol{\sigma}_t^H, \tag{35}$$

where $\{\mathbf{H}_t,\ \boldsymbol{\sigma}_t^H\}$ is the unique solution to the stochastic differential equation $d\mathbf{H}_t/\mathbf{H}_t = \boldsymbol{\mu}_t^H dt + \boldsymbol{\sigma}_t^H d\hat{\mathbf{B}}_t$ (with terminal condition $\mathbf{H}_T = \mathbf{0}$), $\boldsymbol{\mu}_t^H$ is a complex function of the preference parameters (and of $\boldsymbol{\theta}_t^*$, $\hat{\boldsymbol{\eta}}_t$), and $\boldsymbol{\sigma}_t^H$ is the matrix of diffusion coefficients of the optimal consumption process. The last two terms in (35) are hedging demands, which are affected by both ambiguity ($\boldsymbol{\theta}_t^*$) and estimation risk ($\hat{\boldsymbol{\eta}}_t$):

$$-\frac{1}{\gamma}\left((\boldsymbol{\sigma}^R)'\right)^{-1}\boldsymbol{\theta}_t^* + \left((\boldsymbol{\sigma}^R)'\right)^{-1}\boldsymbol{\sigma}_t^H. \tag{36}$$

Ambiguity distorts mean returns at time t by $\boldsymbol{\sigma}^R\boldsymbol{\theta}_t^*$ under the adjusted belief. As $\gamma \to 1$, the optimal portfolio processes converges to the one that is optimal in the logarithmic case. In particular, when $\gamma = 1$, it can be shown that $\boldsymbol{\sigma}_t^H = \mathbf{0}$, and the agent behaves myopically so that there is no hedging demand. However, in this case $\gamma^{-1}\left((\boldsymbol{\sigma}^R)'\right)^{-1}\boldsymbol{\theta}_t^*$ converges to $\left((\boldsymbol{\sigma}^R)'\right)^{-1}\boldsymbol{\theta}_t^*$ so that an ambiguity hedging demand will be present also in the logarithmic case, contrary to what is commonly found under incomplete information. Miao also shows that under ambiguity, but with complete information, the portfolio reduces to

$$\boldsymbol{\omega}_t^{**} = \frac{1}{\gamma}\left(\boldsymbol{\sigma}^R(\boldsymbol{\sigma}^R)'\right)^{-1}\left(\boldsymbol{\mu}_t^R - r^f\boldsymbol{\iota}_d\right) - \frac{1}{\gamma}\left((\boldsymbol{\sigma}^R)'\right)^{-1}\boldsymbol{\kappa}, \tag{37}$$

where $\boldsymbol{\kappa}$ is the vector of security-specific κ-ignorance parameters. (37) indicates that under complete information, ambiguity does not induce any hedging demand; the term that is not in the classical Merton's mean-variance style solution, $\gamma^{-1}\left((\boldsymbol{\sigma}^R)'\right)^{-1}\boldsymbol{\kappa}$, fails to depend on the state of the economy. It is therefore the interaction between incomplete information and ambiguity that induces the hedging demand $-\gamma^{-1}\left((\boldsymbol{\sigma}^R)'\right)^{-1}\boldsymbol{\theta}_t^*$ in (35): unless the incomplete information hedging demand disappears because the investor is rationally myopic and does not care for estimation risk ($\gamma = 1$), under MPP with κ-ignorance, one needs both incomplete information and ambiguity for the latter to generate a non-zero hedging demand.

A related paper is Liu's (2011), who examines a continuous-time portfolio problem for an investor with recursive MPP and κ-ignorance, when the expected return of a risky asset is unobservable and follows a hidden Markov chain. The investor takes into account incomplete information risk (resulting from time-varying precision of the conditional estimates of the unobservable state) and the ambiguity on the process that governs the dynamics of filtered probabilities of the underlying state. In Liu's paper, the investor treats filtered probabilities as ambiguous and has multiple beliefs with respect to the states resulting from Bayesian updating. She obtains the conditional estimates of the unobservable state by observing past and current asset prices and employs a non-linear recursive filter to extract regime probabilities that are updated according to Bayes rule and used to represent the approximating reference model, p. Alternative models (as parameterized by a law of motion q^θ) are obtained as a perturbation of this benchmark. The asset menu is composed of a riskless bond that yields r^f and of a risky asset. The drift process μ_t follows a continuous-time Markov chain with two states, $\mu_H > \mu_L$, and infinitesimal generating matrix with generic element λ_{ij} $(i, j = 0, 1)$.[25] The investor can observe neither the expected return μ_t, nor the BM B. Instead, she can only observe the asset price S. Given an initial prior π_0 over the two regimes, the investor estimates the unobservable state, that is, the probability of the current state being in the high-mean-return regime, based on the observed asset prices, $\pi_t \equiv \Pr(\mu_t = \mu_H|\mathcal{F}_t)$. Under the alternative model q^θ, the law of motion of π_t is

$$d\pi_t = \left[\lambda_0 - (\lambda_0 + \lambda_1)\,\pi_t\right]dt + \pi_t(1 - \pi_t)\frac{\mu_H - \mu_L}{\sigma_t}(dB_t^{q^\theta} - \theta dt), \tag{38}$$

where $B_t^{q^\theta}$ is the BM under q^θ, the distorted probability induced by θ. In the κ-ignorance MPP specification, the set of density generators is characterized through a coefficient κ, a measure of ambiguity aversion.

Liu (2011) characterizes optimal consumption and portfolio rules. Assuming that the lower bound of the conditional market price of risk adjusted for ambiguity is always nonnegative, that is, $(\mu_L - r^f)/\sigma_t - \kappa \geq 0$, under power utility, the optimal demand may be decomposed as the sum of four separate terms. The first component is the standard myopic demand for the risky asset under Markov switching, which is instantaneously mean-variance efficient and depends on the estimate of the unobservable state. The second one reflects ambiguity on the myopic demand, which is a function of perceived ambiguity, κ. Under the assumption that the conditional market price of risk must be non-negative, a larger κ implies that the investor allocates a smaller proportion of wealth to the risky asset when she behaves myopically by ignoring time variation of the conditional estimates. Together, the first two terms form the *ambiguity-adjusted myopic demand*. When returns are IID, and expected returns are fully observable, the optimal portfolio policy is given by the ambiguity-adjusted myopic component only. Interestingly, this result gives rise to a form of observational equivalence, in the sense that an increase in ambiguity is equivalent to a decline in the effective market price of risk. However, under incomplete information, this equivalence fails because ambiguity has an impact on the intertemporal hedging demand. The third and the fourth term ($hedge^{IIR}$ and $hedge^{ambiguity}$) quantify the intertemporal hedging demand, which insures the investor against future time variation of the conditional estimates of the unobservable state. A notable difference between these solutions and

[25] The infinitesimal generating matrix governs the dyanmics of the Markov switching between states of high drift (expect returns) in the stock price process and states of low drift.

those derived in a standard SEU portfolio problem under Markov switching (see e.g., Honda, 2003) is that the intertemporal hedging demand is driven, not only by incomplete information risk, but also by ambiguity. The term $hedge^{IIR}$ is obtained setting $\kappa = 0$, while $hedge^{ambiguity}$ accounts for the difference between $hedge^{IIR}$ and the total hedging demand. In this way, $hedge^{IIR}$ is solely attributed to intertemporal hedging of incomplete information risk, while $hedge^{ambiguity}$ is purely driven by ambiguity. Whilst the term $hedge^{IIR}$ will also exist in a SEU framework with incomplete information such as Honda's (2003), $hedge^{ambiguity}$ will not.

As far as consumption is concerned, two competing effects are generated by ambiguity: on the one hand, the investor will lower her consumption because a given level of wealth can deliver a smaller flow of consumption due to the perceived deterioration in investment opportunities. On the other hand, the agent is less interested in investing in the risky asset and tends to increase consumption. When risk aversion $\gamma > 1$ ($\gamma < 1$), the former (latter) effect dominates, and ambiguity aversion decreases (increases) the consumption-to-wealth ratio. Further, the myopic and the hedging demands are both reduced. Additionally, as the aversion to ambiguity (κ) increases, the size of $hedge^{ambiguity}$ increases, this term turns out to be non-monotonic and displays a mild hump-shape with respect to the estimate of the probability of the unobservable Markov state.

While Liu (2011) has examined continuous time models in which the process for the uncertain asset returns is Markov switching, Faria, Correia-da-Silvaz and Ribeirox (2010) have investigated another type of non-linearity, when the volatility of the uncertain asset is stochastic. Faria et al. extend the SEU framework common to the applied finance literature to allow for MPP and study the effects of ambiguity about stochastic variance or, more precisely, on its expected value. Their findings show that if the investor is able to continuously update her portfolio as a function of the observed instantaneous variance, then her optimal allocation will not be affected by ambiguity. Otherwise ambiguity on the variance becomes relevant, even if the agent can observe the instantaneous variance. Assuming that the investor uses her expectation about future variance, Faria et al. prove that ambiguity reduces the demand for the risky asset.

3.8. *Learning, Ambiguity and Portfolio Choice*

As already commented, one of the frontiers of research on ambiguity consists of the interaction between learning and ambiguity. Epstein and Schneider (2007) have used their learning model to solve an intertemporal asset allocation problem under ambiguity in which the investor can rebalance her portfolio in the light of new information that affects both beliefs and confidence. Consider an investor who believes that the equity risk premium is fixed, but unknown, so that it must be estimated from past returns. Intuitively, one would expect the investor to become more confident the longer is the observed series of returns used to derive the estimate. As a result, a given estimate should lead to higher portfolio weights on stocks, the longer is the available data. Similarly, one might expect the weights on stocks to be higher as the posterior variance of the equity premium declines. Bayesian analysis has tried to capture these intuitions by incorporating estimation risk into portfolio problems. However, learning seems to have moderate effects on investment decisions. In stationary environments (e.g., in the absence of regimes a' la Veronesi, 1999) learning produces effects that quickly dissipate over time. For instance, Guidolin's (2005) Bayesian model of international portfolio diversification can be contrasted to Uppal and Wang's (2003) application to solve the home country bias: Guidolin obtains strong but highly transient effects from learning under histories of different length, while Uppal and Wang obtain first-order and stationary effects in models with small heterogeneity across domestic vs. foreign perceived ambiguity. Epstein and Schneider (2007) argue that a Bayesian model often generates counter-intuitive results because a declining posterior variance as the sample size expands fails to capture changes in confidence, in the sense that even with expanding data sets it is plausible that an investor's confidence may remain time-varying.

In Epstein and Schneider's (2007) model there are k trading dates in a fixed period (say, a month), and the state space is simply $\Omega = \{high, low\}$, where the high state has a probability $p \in (0,1)$. Correspondingly, the log-return realizations for a risky asset can be either $R(high) = \sigma/\sqrt{k}$ or $R(low) = -\sigma/\sqrt{k}$, and the log risk free rate is r^f/k. The investor wants to maximize the utility she derives from her final wealth W_{t+T}, where her utility index is logarithmic. Her portfolio can be rebalanced $k(T-t)$ times between dates t and T. At any rebalancing date, the investor takes into account that she will learn in the future; therefore, the optimal portfolio weights at date t, $\omega^*_{t,T,k}$, depend on the calendar date, on the investment horizon, and on the number of future portfolio adjustments. In a standard Bayesian framework, Epstein and Schneider derive that the optimal portfolio allocation at any date t is the mean-variance allocation

$$\omega_t^{Bayes} = \frac{\bar{R}_t + \frac{1}{2}\sigma^2 - r^f}{\sigma^2}, \tag{39}$$

where $\bar{R}_t$ is the monthly equity premium estimated as of time t. Allowing for ambiguity, they prove instead that learning under ambiguity affects portfolio choice in two ways: directly through the optimal weights, and indirectly through the length of the investment horizon. To see this, consider first a myopic investor ($T = 1/k$). The limit of her optimal weights on stocks as $k \to \infty$ is:

$$\lim_{k\to\infty} \omega^*_{t,k,1/k}(\bar{R}_t) = \max\{\omega_t^{Bayes} - \sigma^{-2}(\bar{v} + t^{-(1/2)}\sigma b_a), 0\} + \min\{\omega_t^{Bayes} + \sigma^{-2}(\bar{v} + t^{-(1/2)}\sigma b_a), 0\}, \tag{40}$$

where $\bar{v}$ is a measure of the amount of information that the investor is expecting to learn in the future: if $\bar{v} = 0$, already today she is confident that she will completely learn the equity premium; if $\bar{v} \to \infty$ no learning is anticipated, and b_a is a parameter determined as a result of the strength of the updating. Hence, a myopic ambiguity-averse investor goes long in stocks only if the equity premium is unambiguously positive. The optimal position is then given by the first term in (40). This implies that an ambiguity-averse investor who goes long behaves as if she were a Bayesian who perceives the lowest conditional mean log return $\bar{R}_t - \bar{v} - t^{-(1/2)}\sigma b_a$. The optimal weight depends on the sample through the Bayesian position ω_t^{Bayes}, since, conditionally on participation, the optimal response to news is the same in the two models. However, ambiguity also introduces a trend component; the optimal position increases as confidence grows and the standard error $t^{-\frac{1}{2}}\sigma$ shrinks to zero: $\omega^*_{t,k,1/k}$ approaches the Bayesian position ω_t^{Bayes}(from below) in the long run, but, unless $\bar{v} = 0$ (that is, without ambiguity), it remains smaller. In this respect, Epstein and Schneider's conclusion is identical to the conclusion reported by many papers before: aversion to ambiguity—under log-preference—unequivocally reduces the demand for the risk asset. However, Epstein and Schneider explicitly link this finding to the strength of the learning process, both past (via the sample size t) and perspective (via the parameters $\bar{v}$). The second term in (40) reflects short selling when the equity premium is unambiguously negative. Interestingly, non-participation in the stock market—a zero demand for the risky asset—is optimal if the maximum likelihood estimate of the equity premium is small in absolute value so that both terms in the *max* and *min* operators in (40) are 0, that is, if $|\bar{R}_t + \frac{1}{2}\sigma^2 - r^f| < \bar{v} + t^{-(1/2)}\sigma b_a$. In particular, an investor who is not confident that the equity premium can be learnt ($\bar{v} > 0$) need not to be "in the market" even after having seen a large amount of data, i.e., even when $t \to \infty$.

The second effect of learning under ambiguity is a new inter-temporal hedging motive that induces more participation as the investment horizon becomes longer. For MPP investors, the myopic and long horizon positions coincide if the equity premium is either high or low enough to induce participation of the myopic investor. However, for intermediate estimates of the premium, the myopic investor stays out of the stock market, while the long horizon investor takes "contrarian" positions and she goes short (long) in stocks for positive (negative) equity premia. For the relatively experienced investor that benefits from a large number of observations t, horizon effects will be small. However, for an inexperienced investor (who is likely to have a wide non-participation region), there can be sizeable differences between the optimal myopic and long horizon weights. Intuitively, agents with a low empirical estimate of the equity premium know that a further low return realizations may push them towards non-participation, and hence a low return on wealth. To insure against this outcome, they short the asset. As a result, the optimal intertemporal portfolio policy involves dynamic "exit and entry" rules: recursive updating shifts the interval of equity premia, and such shifts can make agents move in and out of the market, in terms of her participation decisions.

4. **Equilibrium Asset Prices under Ambiguity**

We now review papers that have "closed" models of portfolio and consumption decisions to generate equilibrium (no-arbitrage) asset pricing implications. When the models that we are about to survey will overlap with frameworks that we have covered in Section 3, we will try to economize on the details concerning assumptions and solutions and immediately jump to the asset pricing implications.

4.1. *Static Models*

The simplest effort aimed at developing equilibrium asset pricing implications of ambiguity is the two-period paper by Easley and O'Hara (2009). As discussed in Section 3.1, ambiguity reduces the risk-sharing capabilities of markets and may consequently induce an increase in the equilibrium risk premium, because states exist in which the whole risk must be absorbed by SEU agents. Indeed, when the market is characterized by limited participation—i.e., demand and supply are equated by a price that falls in the interval $[\mu_{\min},\ \mu_{\max}]$—such an equilibrium price is $p^{SEU} = \hat{\mu} - \hat{\sigma}^2 (1-\alpha)^{-1} x$ where the relevant parameters are estimates by the SEU investors and x is the supply of the stock. Of course, this is the classical CARA/Gaussian result, but among the parameters that affect the SEU-only price, we also find the proportion of ambiguity averse (AA) investors, α. Therefore, even when the AA agents are not trading, their mere existence affects the equilibrium price. On the opposite, when both groups of investors are in the market, the equilibrium price is:

$$p^{AA} = \frac{\alpha\hat{\sigma}^2\mu_{\min} + (1-\alpha)\,\sigma^2_{\max}\hat{\mu}}{(1-\alpha)\,\sigma^2_{\max} + \alpha\hat{\sigma}^2} - \frac{(\sigma^2_{\max} + \hat{\sigma}^2)}{(1-\alpha)\,\sigma^2_{\max} + \alpha\hat{\sigma}^2}x. \tag{41}$$

An increase in the fraction of AA investors α will decrease both p^{SEU} and p^{AA}. In a non-participating equilibrium, an increase in α does not necessarily induce full participation, however, because in the face of a shrinking number of active investors, the price must decrease to compensate the lower number of traders for the higher per-capita risk that they need to absorb to clear the market. However, if p^{SEU} falls enough, the equilibrium may shift from a nonparticipating to a participating one, if p^{SEU} eventually moves below the threshold $\mu_{\min}$. Moreover, changes in the set of variances considered by AA investors have no effect if the economy is in a nonparticipating equilibrium. Viceversa, if the economy is in a participating equilibrium, increases in $\sigma^2_{\max}$ will reduce the price because the demand of AA investors shrinks when $\sigma^2_{\max}$ increases. If the market is in a limited participation equilibrium, increases in $\mu_{\min}$ have no affect on p^{SEU} as long as AA investors are not induced to trade. However, if AA traders participate in the market, increasing $\mu_{\min}$ increases p^{AA} since the risky investment of AA traders increases. Finally, an increase in $\mu_{\min}$ can cause the market to switch from a nonparticipating to a participating equilibrium. Hence, changing the perception of extreme events may have large effects on equilibrium outcomes.

Guidolin and Rinaldi (2009) have used their Easley and O'Hara-type model to derive expressions for (real) expected returns. Under ambiguity, there are two notions of risk premium that can be considered. One is an objective notion and corresponds to an aggregate, market viewpoint based on true expectations. This is also the risk premium anticipated by an external observer that understands the structure of the model and solves for the equilibrium. The other is a subjective notion and corresponds to the expectation—obviously different across SEU and AA investors—of the premium that each individual investor will form before (or without) understanding the overall structure of the model and the outcomes generated by the interaction of SEU and AA investors. Starting with the first, objective notion ($E^{SEU}[1+R]$ in the limited, and $E^{AA}[1+R]$ in the full participation equilibrium), Guidolin and Rinaldi show that required returns on the risky asset are increasing in the proportion α of AA agents because, given total supply, a smaller fraction of SEU-only investors will have to absorb the entire supply. The expressions for the risk premia are derived assuming knowledge of the objective market outcome and—by the law of large numbers—correspond to the average, realized real excess returns on the risky asset if a long sequence of systematic and idiosyncratic shocks were drawn. However, this is different from the risk premium perceived by both categories of investors. The SEU investors perceive an ex-ante risk premium of $\tilde{E}^{SEU}[1+R] = E^{SEU}[1+R]$. This means that every time a participation equilibrium outcome obtains, the SEU-maximizers will be surprised by the "excessively high" average realized real excess returns. Of course, this is due to the existence of an additional source of uncertainty—ambiguity—that is priced in equilibrium and of which the SEU investors are not aware, while AA investors are. The AA investors perceive instead a different expected return $\tilde{E}^{AA}[1+R]$, and Guidolin and Rinaldi show that

$$\tilde{E}^{AA}[1+R] > E^{AA}[1+R] > E^{SEU}[1+R] = \tilde{E}^{SEU}[1+R], \tag{42}$$

i.e., AA agents demand a risk premium which is higher than what SEU agents' expect to receive. Therefore AA investors are always negatively surprised by the realized real excess returns, even when the investors themselves participate in the market. Because SEU investors believe that the risky asset is under-priced, while ambiguity-averse traders believe that it is over-priced, both types of traders will hold portfolios that differ from the market portfolio in their attempt to take advantage of the perceived mispricing.

Cao, Wang, and Zhang (2005, CWZ) adopt a set up that gives an explicit role to a measure of *uncertainty dispersion*, i.e., the heterogeneity in the perceived ambiguity across market participants. In their setting, investment decisions are made at time 0 to maximize utility of consumption at time 1. Agents can trade in a riskless bond and in a risky stock priced at p, whose per capita initial endowment of shares is $x > 0$. The payoff of the stock is normally distributed with mean μ and variance σ^2. σ^2 is precisely known, while μ is not. Agents display MPP (with heterogeneous degrees of aversion to uncertainty), characterized by a CARA utility index with coefficient α. The set of priors $\wp$ considered by each agent can be characterized in terms of confidence regions: each agent considers possible all mean returns $(\hat{\mu} + \theta)$, where $\hat{\mu}$ is an approximating estimate common to all agents, and θ is such that $\theta^2 \leq \sigma^2\eta_i^2$ for some parameter η_i that measures the uncertainty perceived by investor i. η is assumed to be uniformly distributed among investors on the interval $[\bar{\eta} - \delta, \bar{\eta} - \delta]$, for given $\bar{\eta}$, δ such that $\bar{\eta} \geq \delta$. δ can be interpreted as a measure of uncertainty aversion dispersion among agents. Each investor solves the mean-variance problem

$$\max_{\omega_i} \min_{\theta:\theta^2 \leq \eta_i^2} \omega_i \left(\hat{\mu} - \theta - p\right) - \frac{\alpha\sigma^2}{2}\omega_i^2, \tag{43}$$

with typical, step-wise solution characterized by a no-trade interval $[\hat{\mu} - \sigma\eta_i, \hat{\mu} + \sigma\eta_i]$. In the case of full participation, the risk premium is the sum of the standard risk premium $\alpha x\sigma^2$, and of the uncertainty premium $\sigma\bar{\eta}$. Instead, when some agents optimally decide to avoid trading, the pricing equation is:

$$\hat{\mu} - p = 2\sigma\sqrt{\alpha x\delta\sigma} + \left(\bar{\eta}\sigma - \delta\right). \tag{44}$$

CWZ show that in the limited participation equilibrium the rate of participation, the average measure of ambiguity, and the equity premium all decrease as uncertainty dispersion δ increases, whereas under full participation, the equity premium does not depend on uncertainty dispersion. Indeed, in the equilibrium with limited participation there are fewer active investors that have to bear all the market risk, so they demand a higher risk premium. However, only investors with relatively low uncertainty want to trade and they are willing to accept a lower uncertainty premium. Hence, the net effect on the equity premium depends on which force dominates.

Boyle, Garlappi, Uppal, and Wang (2009, BGUW) is another mean-variance paper that can be "closed" to generate interesting asset pricing implications in terms of decomposition of the risk premium. Assuming that each agent is familiar with only one particular asset and (equally) unfamiliar with respect to all the others available, BGUW show that if ambiguity is relatively low with respect to the reward-to-risk ratio of each asset, then each individual will hold both the familiar asset and the portfolio made by all the unfamiliar ones, as seen in Section 3.5. In this case of low ambiguity (relative to the reward-to-risk ratio), the equilibrium risk premium has three components: the conventional risk premium, a second term due to the concentration of the portfolio in the familiar asset, which is the ambiguity premium relative to the unfamiliar assets, and an ambiguity premium relative to the familiar asset. When the number of agents (and assets) becomes arbitrarily large, the premium for portfolio concentration in the familiar asset vanishes since the number of unfamiliar assets is large enough to compensate the effect of the ambiguity-induced portfolio bias towards the familiar asset. When the unfamiliar assets are considered extremely ambiguous, each investor holds only his own familiar one. Hence, the equilibrium excess return does not reflect any ambiguity premium related to the unfamiliar assets. Finally, when the ambiguity level is so high for both familiar and unfamiliar assets, nobody is willing to trade and there are no equilibrium prices.

Maccheroni, Marinacci, Rustichini and Taboga (2009) have derived an interesting version of the CAPM under ambiguity. They solve a Markowitz-style portfolio problem for an ambiguity averse agent that displays Monotone Mean Variance Preferences (MMVP). MMVP are a specific subclass of VP that expands Mean Variance Preferences (MVP) outside their domain of monotonicity. Indeed, while classical MVP are theoretically questionable because they may fail to be monotone, violating rationality and opening the door to the existence of unexploited arbitrage opportunities, MMVP are monotone everywhere and always agree with MVP where the latter are economically meaningful (i.e., monotone), but may depart from MVP when this is required to bar the existence of arbitrage opportunities. Denoting by $\mathbf{R}$ the N-dimensional vector of gross risky returns, by R_f the gross return of a risk-free bond, and by $\boldsymbol{\omega}^*_{MVP}$ the optimal portfolio of a MVP-agent with risk aversion γ, for a MMVP-investor with uncertainty aversion ϱ, the optimal allocation rule is:[26]

$$\boldsymbol{\omega}^*_{MMVP} = \frac{1}{\varrho \Pr\left(W(\boldsymbol{\omega}^*_{MMVP}) \leq \psi^*\right)} \{Cov\left[\mathbf{R}\right]\}^{-1} E\left[\mathbf{R} - R_f \boldsymbol{\iota}_N | W(\boldsymbol{\omega}^*_{MMVP}) \leq \psi^*\right], \tag{45}$$

where $W(\boldsymbol{\omega}^*_{MMVP})$ is future wealth expressed as a function of the selected portfolio, $W(\boldsymbol{\omega}) = (\mathbf{R} - R_f \boldsymbol{\iota}_N)\boldsymbol{\omega} + R_f$, and ψ^* is a constant determined jointly with $\boldsymbol{\omega}_{MMVP}$ by solving a non-linear system of two equations that involve the conditional moments of the returns' distribution. The use of conditional moments instead of unconditional ones implies that a MMVP agent ignores the part of the distribution where wealth is higher than the threshold parameter ψ^*, meaning that she will not take into account those high payoff states that contribute to increase the mean return, but give an even greater contribution to increase the variance and that imply an absurdly negative marginal utility of wealth. Maccheroni et al. (2009) prove that $|\boldsymbol{\omega}^*_{MMVP}| \geq |\boldsymbol{\omega}^*_{MVP}|$ (element-wise) so that MMVP-optimal portfolios are always more leveraged than MVP-ones, simply because some favorable investment opportunities are discarded by MVP-agents (because of preferences' non-monotonicity) and exploited by MMVP-investors.

Under MMVP, Maccheroni et al. also derive a monotone version of the standard CAPM. Specifically, for each asset the following pricing equation holds,

$$E\left[R_i\right] - R_f = \beta_i \left(E\left[R_m\right] - R_f\right) \quad \beta_i = \frac{Cov\left[R_i, \mathcal{M}\right]}{Cov\left[R_m, \mathcal{M}\right]}, \tag{46}$$

where R_m is the return of the market portfolio and $\mathcal{M}$ is the SDF, $\mathcal{M} = -\varrho_m \min[0, R_m - \psi_m]$, where the pair (ϱ_m, ψ_m) is a solution to a nonlinear system involving market quantities only. A monotone CAPM satisfies the standard two-fund separation property: the optimal portfolios held by agents with different degrees of uncertainty aversion are all proportional to each other, and, in equilibrium, the only difference between two agents is the amount of wealth invested in the risk-free asset. Furthermore, the monotone CAPM is truly arbitrage-free and it has the same empirical tractability of the standard CAPM. It remains unexplored whether the monotone CAPM may help empirical researchers in solving the many puzzles surrounding the empirical performance of the standard CAPM.

4.2. *Dynamic Asset Pricing Models*

One of the seminal contributions in the literature on intertemporal asset pricing models under ambiguity is Epstein and Wang (1994), in which MPP are generalized to a multiperiod discrete time framework. Epstein and Wang consider the case in which at each date t a representative agent observes some realization $s_t \in \Omega$. Beliefs about the evolution of the process $\{s_t\}$ are drawn from a time-homogeneous Markov chain. In standard models, this would imply a unique Markov probability kernel for the conditional probabilities of time $t+1$ events. However, under ambiguity, beliefs conditional on s_t are too vague to be represented by a unique prior and are therefore modeled as a probability kernel correspondence $\wp$, which is a (multi-valued) function such that, for each s_t, $\wp(s_t)$ is the set of probability measures representing beliefs. The recursive specification of utility under MPP preferences in an infinite horizon framework is

$$V_t\left(c, s^t\right) = u\left(c_t\left(s^t\right)\right) + \beta \inf_{q \in \wp(s_t)} \int V_{t+1}\left(c, s^t\right) dq, \tag{47}$$

where s^t denotes history up to time t, i.e. $s^t = \{s_0, s_1, ..., s_t\}$. At each date the agent receives a stochastic endowment $e = \{e_t, \mathbf{0}\}$ and can trade in N securities, each of which is in zero net supply and has dividend and price processes $d_i = \{d_{it}\}$ and $p_i = \{p_{it}\}$, respectively.

[26] Over the domain of monotonicity of MVP, γ and ϱ simply coincide. However, whenever MMVP differ from standard MVP, it is justified to distinguish them and to interpret ϱ as an uncertainty-aversion parameter.

An equilibrium is characterized by a process $\mathbf{p} = \{\mathbf{p}_t\}$, such that the endowment corresponds to the optimal consumption-investment plan for all times and states, and markets clear. The max-min criterion implicit in MPP leads to non-differentiability of the representation V. Nevertheless, one-sided Gateaux derivatives (a generalization of the standard concept of directional derivative) of V at c exist and are well-defined. Assuming that $\{\mathbf{p}_t\}$ is an equilibrium, any variation from the optimal policy $\{e_t, \mathbf{0}\}$ should leave the agent worse off. This consideration leads to the necessary condition:

$$\beta \max_{q\in\wp(s_t)} \int \frac{u'(e_{t+1})}{u'(e_t)} (\mathbf{p}_{t+1} + \mathbf{d}_{t+1})\, dq \geq \mathbf{p}_t \geq \beta \min_{q\in\wp(s_t)} \int \frac{u'(e_{t+1})}{u'(e_t)} (\mathbf{p}_{t+1} + \mathbf{d}_{t+1})\, dq. \tag{48}$$

Epstein and Wang prove that this double inequality is also sufficient for an equilibrium: under recursive MPP the optimality conditions take the form of inequalities because $V(\cdot; s)$ is generally non-differentiable, unless the correspondence $\wp : \Omega \to \wp(\Omega)$ is a probability kernel (i.e., there is no ambiguity). In general, these Euler "inequalities" are satisfied by a non-singleton set of prices, so that the equilibrium is not unique when for some t, the two one-sided Gateaux derivatives are different. Intuitively, we would expect a link between indeterminacy of assets prices and asset price volatility. This intuition can be easily formalized in the special case of "IID beliefs", that is, when $\wp$ is independent from s_t. Under this restriction, if a security has a time-homogeneous dividend process and the price is determinate, the latter must be necessarily constant across time and states; consequently, any fluctuation in price can only be a reflection of indeterminacy. Hence the equilibrium indeterminacy generated by ambiguity aversion may leave room for sunspots or Keynesian animal spirits to determine the equilibrium process, leading to higher volatility in asset prices than what is predicted in the standard Lucas model. Epstein and Wang (1995) have shown a connection between ambiguity and markets' extreme phenomena, such as crashes and booms: ambiguity aversion can generate equilibrium prices that are discontinuous processes of the state variables, so that even small variations in the market fundamentals are responsible for sudden significant changes in security prices.

Chen and Epstein (2002) have extended these results by building a dynamic, recursive MPP consumption-CAPM (CCAPM). While their model is more general than Epstein and Wang's (1994), the convenience of a continuous time framework allows them to obtain a few special cases of great interest. As in Section 2.4, under recursive MPP, the utility functional V is the solution to the differential equation

$$dV_t = \left[-\psi(c_t, V_t) + \max_{\theta\in\Theta} \boldsymbol{\theta}_t' \boldsymbol{\sigma}_t^V \right] dt + \boldsymbol{\sigma}_t^V d\mathbf{B}_t \tag{49}$$

where $V_T = 0$, $\mathbf{B}$ is a N-dimensional BM, and $\boldsymbol{\sigma}_t^V$ is a $N \times N$ covariance matrix. The aggregator function $\psi(c, v)$ is $\psi(c, v) \equiv [c^b - \varkappa(av)^{b/a}]/b(av)^{(b-a)/a}$, for $\varkappa \geq 0$, and parameters $b, a \neq 0$, with $a \leq 1$. There is a single consumption good, a riskless asset with return process r_t^f and N risky securities, one for each component of the BM $\mathbf{B}$. The returns $\mathbf{R}_t$ of the risky securities are described by a standard Ito process with parameters $\boldsymbol{\mu}_t$ and $\boldsymbol{\sigma}_t$. The set of density generators is derived by fixing a vector of parameters $\boldsymbol{\kappa}$ in $\mathbb{R}^N$ and let $\boldsymbol{\Theta}_t(\cdot) = \{\mathbf{y} \in \mathbb{R}^N\colon |y_i| \leq \kappa_i \text{ for all } i\}$, so that κ_i can be thought of as a measure of ambiguity aversion relative to asset i. Market completeness delivers a strictly positive state price process π_t, with a standard, vector of market prices of the N sources of uncertainty $\boldsymbol{\eta}_t$. Letting the optimal consumption c_t be an Ito process with parameters $(\mu_t^c, \boldsymbol{\sigma}_t^c)$, and writing the wealth process in terms of the market portfolio as $dW_t/W_t = b^M dt + \boldsymbol{\sigma}^M \cdot d\mathbf{B}_t$, the market price of risk and uncertainty is:

$$\boldsymbol{\eta}_t = b^{-1}[a(1-b)\boldsymbol{\sigma}_t^c + (b-a)\boldsymbol{\sigma}_t^M] + \boldsymbol{\theta}_t^*. \tag{50}$$

If the risk-free rate and the market price of uncertainty are deterministic constants, the optimal consumption process is geometric, therefore, $\boldsymbol{\sigma}_t^M = \boldsymbol{\sigma}_t^c$ and $\boldsymbol{\eta}_t = (1-a)\boldsymbol{\sigma}_t^c + \boldsymbol{\theta}_t^*$, where $\boldsymbol{\theta}_t^*$ is the solution of the maximization problem (49).

Assuming that ambiguity aversion is small in the sense that $0 \leq \kappa_i < |\eta_t^i|$ for all i, then $\sigma_t^{c,i} > (<)\ 0$ if $\eta_t^i > (<)$ 0, and optimal portfolio weights differ from the standard mean-variance expression only because of a correction $\boldsymbol{\kappa} \otimes sign(\boldsymbol{\eta}_t)$.[27] Hence, the optimal portfolio is not instantaneously mean-variance efficient and the mutual fund separation theorem holds if and only if $\boldsymbol{\kappa}$ is common to all agents. Otherwise, different investors that perceive heterogeneous κ-ambiguity, will have different optimal portfolio weights. Moreover, although the composition of the risky portfolio is independent of the risk aversion parameter, it depends on preferences through $\boldsymbol{\kappa}$.

Chen and Epstein show that in this model the equilibrium risk premium is the sum of a standard SEU risk premium, $b^{-1}[a(1-b)\boldsymbol{\sigma}_t\boldsymbol{\sigma}_t^h + (b-a)\boldsymbol{\sigma}_t\boldsymbol{\sigma}_t^M]$, and of an ambiguity premium, $\boldsymbol{\sigma}_t\boldsymbol{\theta}_t^*$, with ith component $\boldsymbol{\sigma}_t^i\boldsymbol{\theta}_t^* = -Cov_t[dR_t^i, dz_t^{\theta^*}/z_t^{\theta^*}]$ (here $\boldsymbol{\sigma}_t^i$ is the i-th row of the matrix $\boldsymbol{\sigma}_t$). Therefore, the premium is positive if the asset's return has negative co-variation with $dz_t^{\theta^*}/z_t^{\theta^*}$, where $z_t^{\theta^*}$ is the Radon–Nikodym derivative of the worst-case model with respect to the approximating one. If the consumption process follows a Geometric BM with constant drift and diffusion vector, Chen and Epstein show that their results strengthen to deliver additional insights. For instance, the vector of market prices of risk and uncertainty simplifies to $\boldsymbol{\eta}_t = (1-a)\boldsymbol{\sigma}_t^c + \boldsymbol{\kappa} \otimes sign(\boldsymbol{\sigma}_t^c)$. Crucially, the market price of uncertainty can be large even if consumption volatilities are small because the second term depends only on the sign of these volatilities and not on their magnitudes. Furthermore, one can prove that the vector of risk premia is:

$$\boldsymbol{\mu}_t - r_t^f \boldsymbol{\iota}_N = (1-a)\boldsymbol{\sigma}_t^i\boldsymbol{\sigma}_t^c + \boldsymbol{\kappa}(\boldsymbol{\sigma}_t^i \otimes sign(\boldsymbol{\sigma}_t^c)). \tag{51}$$

The ambiguity premium (the second term) for asset i is large if $\sigma_t^{i,j} \otimes sign(\sigma_t^{c,j})$ is large, and it is positive for components j of the driving Brownian process $\mathbf{B}_t$ that are very ambiguous, with large κ_js. Since the equilibrium risk premia depend on the endowment process only via the signs of the covariances $\sigma_t^{c,j}$, $j = 1, ..., N$, large ambiguity premia can occur even if consumption is relatively smooth. The excess return on the market portfolio (i.e., a portfolio exactly mimicking the consumption process itself) is given by:

$$\mu_t^M - r_t^f = (1-a)(\boldsymbol{\sigma}_t^c)'\boldsymbol{\sigma}_t^c + \boldsymbol{\kappa}'|\boldsymbol{\sigma}_t^c|. \tag{52}$$

The equity premium can be decomposed in the sum of the standard risk premium term $(1-a)(\boldsymbol{\sigma}_t^c)'\boldsymbol{\sigma}_t^c$ and of the ambiguity premium $\boldsymbol{\kappa}'|\boldsymbol{\sigma}_t^c|$. The ambiguity premium for the market portfolio vanishes as σ^c approaches zero. However, because it is a first-order function of volatility, it dominates the risk premium (a quadratic function) for small volatilities. Finally, the equilibrium risk-free rate will satisfy

$$r_t^f - \beta = (1-b)\left[\mu_t^c - \frac{(1-a)(2-b)}{2(1-b)}(\boldsymbol{\sigma}_t^c)'\boldsymbol{\sigma}_t^c - \boldsymbol{\kappa}'|\boldsymbol{\sigma}_t^c|\right], \tag{53}$$

and it will be decreasing in both ambiguity and risk aversion. An ambiguity CCAPM has the possibility to induce relatively low equilibrium (real) riskless rates under moderate degrees of risk and ambiguity aversions, while increasing the equity risk premium thanks to the contribution of the ambiguity premium, $\boldsymbol{\kappa}'|\boldsymbol{\sigma}_t^c|$.

Epstein and Miao (2003) have used the results in Chen and Epstein (2002) to study asset pricing and portfolio choice in a realistic international finance application. Under special assumptions, Epstein and Miao are able to describe in closed form the equilibrium for a pure-exchange, continuous-time economy with two heterogeneous agents (or countries) and complete markets. In the model, two agents display recursive MPP with logarithmic utility indices and can trade in a locally riskless bond earning the instantaneous interest rate r^f, and in two ambiguous securities, with (non-negative) dividend streams Y_t^1 and Y_t^2. To characterize the set of effective priors

[27] $\boldsymbol{\kappa} \otimes sign(\boldsymbol{\eta}_t) = \boldsymbol{\kappa}_i \left|\eta_t^i\right| / \eta_t^i$ if $\eta_t^i \neq 0$, while otherwise $\boldsymbol{\kappa} \otimes sign(\boldsymbol{\eta}_t) = \mathbf{0}$.

considered by each agent, while allowing for heterogeneity in ambiguity aversion, consider the vectors $\boldsymbol{\kappa}^1 = [0\ \kappa_1]'$ and $\boldsymbol{\kappa}^2 = [\kappa_2\ 0]'$, and for each agent define the set of density generators

$$\boldsymbol{\Theta}^i = \left\{ \boldsymbol{\theta} = (\boldsymbol{\theta}_t) \colon\ sup_t \left| \theta_t^j \right| \leq \kappa_j^i \quad j = 1, 2 \right\} \quad i = 1, 2. \tag{54}$$

$\boldsymbol{\Theta}^1$ and $\boldsymbol{\Theta}^2$ determine the set of distributions considered by agents 1 and 2, respectively. Intuitively, agent 1 (2) is assumed to be "more familiar" with the first (second) security, which in her view is only risky and not ambiguous, as shown by the null first (second) component in κ^1 (κ^2). An Arrow-Debreu equilibrium is defined as a consumption process for each country, $\{c_t^1\}$ and $\{c_t^2\}$, and a state price process $\{\pi_t\}$, such that each country maximizes her MPP-utility subject to a budget constraint and the goods (consumption) markets clear.[28] Using standard no-arbitrage intertemporal conditions, the Arrow-Debreu equilibrium can be implemented by a sequence of temporary (Radner) equilibria and is identical to the one of an ambiguity-free economy in which countries' beliefs are represented by the two priors q^1 and q^2, such that:

$$\frac{dq^i}{d\pi} = \exp\left\{ -\frac{1}{2} (\kappa_i)^2 T - \kappa_i B_T^j \right\} \quad i = 1, 2 \quad i \neq j = 1, 2. \tag{55}$$

Epstein and Miao (2003) show that each asset's risk premium is the sum of the classical covariance (with the endowment) component $\left(\boldsymbol{\sigma}_t^i\right)' \boldsymbol{\sigma}^Y$ and of an additional term that comes from MPP aversion to ambiguity. Such a contribution appears to be monotonically increasing in κ_j (the amount of ambiguity on the other asset), while the effect of an increase in κ_i will depend on the sign of σ_t^{ij}. However, as in the general framework by Chen and Epstein (2002), the risk premium has neither to shrink towards zero when $\left(\boldsymbol{\sigma}_t^i\right)' \sigma^Y \to 0$ nor to be zero when $\sigma^Y = 0$ (i.e., when there is no fundamental risk in the economy). Similarly, the expression for the equilibrium riskless rate shows that while an increase in the ambiguity measures (the κs) is likely (or guaranteed, when $\sigma_t^{12} > 0$) to increase the equity risk premia, growing ambiguity is also likely to translate into lower real interest rates, which is a key to solving the equity premium and risk-free rate puzzles.[29]

If we interpret each agent as a representative consumer from a specific country and, consequently, B^i as the domestic shock to the endowment in country $i = 1, 2$, Y^i can be thought of as the dividend process on the domestic security of country i, and a number of interesting insights can be derived. For instance, it is well known that in standard SEU models with log-utility preferences, each country's consumption level will be a deterministic function of aggregate endowment, so that consumption growth rates worldwide should display perfect, unit correlation. However, the empirical evidence points to the fact that consumption growth rates are only weakly correlated, the so-called consumption home bias. Interestingly, under recursive MPP, Epstein and Miao show that the assumption $0 \leq \kappa_1 < \sigma_1^Y$, $0 \leq \kappa_2 < \sigma_2^Y$ implies that the country-specific growth rate is positively correlated with domestic shocks. Furthermore, the country with higher mean growth rate also has the largest variance of consumption growth. The higher is the ambiguity faced by country j, the higher is the mean growth rate of country i. Finally, rewriting the return of asset 1 in terms of the BM driving processes that are appropriate for country 1 and 2, respectively,

[28] The (Arrow-Debreu) equilibrium exists under the assumption $0 \leq \kappa_1 < \sigma_1^Y$, $0 \leq \kappa_2 < \sigma_2^Y$, which can be read as imposing an upper bound on the tolerable amount of ambiguity for trading to take place. In fact, under the assumption of "small ambiguity" the two worst-case density generators are $\boldsymbol{\theta}_t^{*1} = (0, \kappa_1)'$ and $\boldsymbol{\theta}_t^{*2} = (\kappa_2, 0)'$.

[29] Leippold, Trojani, and Vanini (2008) have re-examined the effects of ambiguity on the equity premium and risk-free rate puzzles in a Lucas exchange economy with a CRRA representative agent. Although their analysis is performed at a level of generality that is inferior to Epstein and Wang (1994) and Chen and Epstein (2002), they confirm that ambiguity aversion implies lower equilibrium interest rates, regardless of risk aversion and a higher equity premium.

one gets:

$$dR_t^1 = \left(b_t^1 - \kappa_i \sigma_t^{1j \neq i}\right) dt + \sigma_t^{1j \neq i} dB_t^i + \sigma_t^{1j \neq i} d\left(B_t^{j \neq i} + \kappa_i t\right) \text{ [country } i \text{ perspective] } j, i = 1, 2 \; j \neq i.$$

where b_t^i is the expected return of stock i. Therefore, country i will impute a higher expected return to security i with respect to country j if and only if $\kappa_j \sigma_t^{1i} > \kappa_i \sigma_t^{1j}$, which holds only if country i considers its own security less ambiguous than country j is. Hence country i attaches to its stock a higher expected return reflecting a "preference" for it over the foreign one, as in the empirical findings of French and Poterba (1991).

4.3. *Robustness*

Another, crucial strand of the asset pricing literature has adopted the robustness modelling approach proposed by Hansen, Sargent, and coauthors. One of the first papers in this camp is the discrete-time, linear-quadratic permanent income model used by Hansen, Sargent, and Tallarini (1999, HST) to study consumption/saving rules. Specifically, HST have re-interpreted *risk-sensitive preferences* as embedding a desire for robustness against model's misspecification, and showed that large estimates (or calibrated "perceptions") of market-based measures of risk aversion may result from misspecifications that ignore the existence of a concern for small specification errors by robust DMs. To understand the connection between risk-sensitive preferences and robust preferences, consider a simple control problem whose state transition equation is $\mathbf{x}_{t+1} = \mathbf{A}\mathbf{x}_t + \mathbf{B}\mathbf{c}_t + \mathbf{C}\mathbf{w}_{t+1}$, where $\mathbf{x}_t$ is a state vector, $\mathbf{c}_t$ is the control vector (e.g., consumption and portfolio shares) and $\mathbf{w}_{t+1}$ is an IID Gaussian random vector with zero mean and identity covariance matrix. The risk sensitive control problem consists in choosing the policy rule $\mathbf{c}$ to maximize

$$V_t = U(\mathbf{c}_t, \mathbf{x}_t) + \beta \mathcal{R}_t(V_{t+1}), \tag{56}$$

where $U(\mathbf{c}_t, \mathbf{x}_t)$ is a one period utility index and $\mathcal{R}_t(V_{t+1}) = \sigma^{-1} 2 \ln E\left[\exp\left(\sigma V_{t+1}/2\right) \middle| \mathcal{F}_t\right]$ is a risk adjustment operator. Specifically, when $\sigma \neq 0$, $\mathcal{R}_t$ makes an additional adjustment with respect to the one induced by U, in the sense that negative values of σ correspond to higher aversion to risk with respect to a SEU specification. Moreover, in a robust control problem the transition equation is the distorted law of motion,

$$\mathbf{x}_{t+1} = \mathbf{A}\mathbf{x}_t + \mathbf{B}\mathbf{c}_t + \mathbf{C}(\mathbf{w}_{t+1} + \boldsymbol{\theta}_t), \tag{57}$$

where $\boldsymbol{\theta}_t$ distorts the mean of the innovation $\mathbf{w}$ and is chosen to minimize U_0 satisfying two constraints:[30]

$$\hat{E}\left[\sum \beta^j \boldsymbol{\theta}'_{t+j} \boldsymbol{\theta}_{t+j}\right] \leq \eta_t \text{ and } \eta_{t+1} = \beta^{-1}\left(\eta_t - \boldsymbol{\theta}'_t \boldsymbol{\theta}_t\right) \text{ for a given } \eta_0. \tag{58}$$

Hansen and Sargent (1998) show that—setting the Lagrange multiplier for the second constraint to $-1/\sigma$—the optimal value function that solves the risk-sensitive control problem and the one that solves the robust control problem both have a quadratic form in $\mathbf{x}$ with a common characteristic matrix. Furthermore, the optimal policy rule $\mathbf{c}^*$ is the same in both problems. Hence the parameters β and σ can be interpreted as indicating a desire for robustness. HST have applied this framework to asset pricing. They prove that the concern for robustness modifies

[30] $\hat{E}[\cdot]$ is the expectation operator conditional on the distorted law of motion. The assumption that the distortion only affects the mean of the shocks to the linear system is actually not as restrictive as it may seem. Hansen and Sargent (2007) show that if we were to allow a general misspecification within a Gaussian framework, the optimal distortion would consist of a Gaussian transition density with the same mean distortion as when the misspecification is constrained to affect the mean.

the SDF by introducing a multiplicative term

$$\mathcal{M}_{t+1}^{u}=\frac{\exp\left[-\left(\mathbf{w}_{t+1}-\boldsymbol{\theta}_{t}\right)'\left(\mathbf{w}_{t+1}-\boldsymbol{\theta}_{t}\right)\right]}{\exp\left(-\mathbf{w}_{t+1}'\mathbf{w}_{t+1}/2\right)}, \tag{59}$$

which means that $\mathcal{M}_{t+1}^{u}$ is the density ratio of the distorted law of motion relative to the reference probability distribution. Additionally, it turns out that $St.dev\left(\mathcal{M}_{t+1}^{u}|\mathcal{F}_{t}\right)=\left[\exp\left(\boldsymbol{\theta}_{t}'\boldsymbol{\theta}_{t}\right)-1\right]^{1/2}\simeq\|\boldsymbol{\theta}_{t}\|$, i.e., the price of risk is approximately equal to the time t specification error, as captured by the distortion $\boldsymbol{\theta}_{t}$.

Anderson, Hansen, and Sargent (2003, AHS) have extended these results going beyond the linear-quadratic framework and allowed for more general transition functions. In particular, the technique employed by AHS is a generalization of the standard Hamilton-Jacobi-Bellman (HJB) equation applied to the pricing equations that twist the probability distribution of the reference model. Such a twist is governed by a single parameter that measures the set of alternative specifications that affect the DM (see Section 2.3). An ambiguity averse agent is attracted by decision rules that will work well across a set of $\boldsymbol{\theta}$s that are not implausibly "large". Therefore the classical utility maximization problem is replaced by a new max-min problem in which a penalty term is introduced to restrain the choice of $\boldsymbol{\theta}$. The optimal policy c solves the problem

$$\max_{\{c_t\}}\ \min_{\{\boldsymbol{\theta}_t\}} E\int_{0}^{\infty}\exp\left(-\delta t\right)\left[U(\mathbf{x}_t,c_t)+\frac{\varrho}{2}\boldsymbol{\theta}_t'\boldsymbol{\theta}_t\right]dt, \tag{60}$$

where ϱ is a parameter that measures the concern for model misspecification. Of course, this is exactly the typical specification for Hansen-Sargent preferences that reflect a preference for robustness, in penalty form. Letting V^{*} be the value function that solves (60) and $c_t=c^{*}(\mathbf{x}_t)$ the optimal solution for log-consumption, we define $\mu^{*}(\mathbf{x}_t)$ and $\sigma^{*}(\mathbf{x}_t)$ the drift and diffusion of the process $\mathbf{x}_t$ when the optimally distorted control law $c^{*}(\mathbf{x}_t)$ is imposed. Interpreting the state variable $\mathbf{x}_t$ as a vector of assets' returns, Anderson et al. (2003) derive a decomposition of the SDF which is more general than in HST. In particular, the multiplicative adjustment for model uncertainty is:

$$\mathcal{M}_{t,t+\tau}^{u}\equiv\frac{\exp\left(-V^{*}(\mathbf{x}_{t+\tau})/\varrho\right)}{E\left[\exp\left(-V^{*}(\mathbf{x}_{t+\tau})/\varrho\right)|\mathbf{x}_t\right]}. \tag{61}$$

where $\mathcal{M}_{t,t+\tau}^{u}$ is an exponential martingale and can be represented as:

$$\mathcal{M}_{t,t+\tau}^{u}=\exp\left(\int_{t}^{t+\tau}\left[-\frac{1}{2}(\theta^{*}(\mathbf{x}_s))'\theta^{*}(\mathbf{x}_s)+(\theta^{*}(\mathbf{x}_s))'d\mathbf{B}_s\right]\right)\qquad\theta^{*}(\mathbf{x}_s)\equiv-\frac{1}{\varrho}(\sigma^{*}(\mathbf{x}))'\frac{\partial V^{*}(\mathbf{x})}{\partial\mathbf{x}}. \tag{62}$$

where θ^{*}, $\sigma^{*}$and V^{*}derives from the solution to (60).

Maenhout (2004) has further progressed on the analysis of equilibrium asset prices under ambiguity within the robustness framework. Besides making closed-form expressions possible, under homotheticity the equilibrium worst-case scenario can be characterized to provide guidance on the choice of reasonable values for the uncertainty aversion parameter, which has often been the Achille's heel of applied ambiguity research. In addition, homotheticity makes robustness observationally equivalent to the stochastic differential utility of Duffie and Epstein (1992), so that robustness can be interpreted as increasing risk aversion without changing the willingness to substitute intertemporally.[31] As a result, an investor with homothetic robust preferences will be observationally equivalent to a Duffie-Epstein investor with elasticity of intertemporal substitution equal to γ^{-1} and risk aversion $\varrho^{-1}+\gamma$, where γ is the standard coefficient of relative risk aversion and ϱ measures the concern for robustness (or, equivalently, ambiguity aversion) in (60). Hence, the empirical evidence of substantial pessimism when forming portfolios

[31] In a finite-horizon continuous time setting, Skiadas (2003) has proven that the functional representation developed by Hansen and Sargent in a number of papers admits a recursive representation and that this takes the form of the Stochastic Differential Utility (SDU) of Duffie and Epstein (1992), independently of the specific dynamics of the underlaying state-variable.

might be consistent with rational behavior in the presence of moderate amounts of uncertainty aversion rather than by relative risk aversion. Therefore cautious portfolios and high equity premia can be obtained by keeping γ at reasonable levels (e.g., below the threshold of 10 argued by Mehra and Prescott, 1985).

Using AHS's result that a concern for model uncertainty simply adds an endogenous drift $\theta(W_t)$ to the law of motion of the state variable W_t, in Maenhout's framework stock market returns simply follow a GBM with constant drift (expected return) μ and constant diffusion coefficient σ so that the state variable is simply represented by financial wealth. The drift adjustment $\theta(W_t)$ is endogenously chosen and depends on the parameter ϱ. The more a DM cares for robustness (a lower ϱ), the less faith she will have in the reference model and the larger will be the set of drift distortions considered when evaluating her continuation payoffs from any given consumption and portfolio rule. Assuming that the endowment process D_t follows a GBM with fixed parameters μ_c and σ_c, in an equilibrium the excess return on the market portfolio under the distorted model is given by

$$\frac{dS_t + D_t dt}{S_t} - r^f dt = (\gamma + \varrho^{-1})\sigma_{cs} dt + \sigma_c dB_t, \tag{63}$$

with $\sigma_{cs} \equiv cov(dc/c, dS/S)$ being the instantaneous covariance between the growth rate of consumption (equal to endowment in equilibrium, as this is a standard in Lucas model) and the rate of price increase of the market portfolio. Clearly, the equity premium in (63)—the drift of the excess return process—will exceed the standard expression for the equity premium in the standard CCAPM, $\gamma\sigma_{cs}$. The difference is the term $\varrho^{-1}\sigma_{cs} > 0$ which is a decreasing (increasing) function of the parameter ϱ (ambiguity aversion). The term $\gamma + \varrho^{-1}$ captures the fact that both market risk and model uncertainty are priced in equilibrium. If we define the rate of intertemporal substitution in consumption as ψ, the equilibrium risk-free rate is

$$r^f = \delta + \psi\mu_c - \frac{1}{2}(1+\psi)(\gamma + \varrho^{-1})\sigma_c^2, \tag{64}$$

which differs from the standard CCAPM expression, $\delta + \psi\mu_c - \frac{1}{2}(1+\psi)\gamma\sigma_c^2$. Clearly, a concern for ambiguity lowers the risk-free rate vs. the SEU case because $(1+\psi)\varrho^{-1}\sigma_c^2 > 0$. Therefore the desire for robustness drives down the risk-free rate through a precautionary savings channel, since the separation between intertemporal substitution and risk aversion ($\gamma^{-1} \neq \psi$) allows a high (but reasonable) value of γ, without counterfactually producing a high risk-free rate (which would occur with time-additive, CRRA preferences, when $\gamma^{-1} = \psi$). This confirms earlier results in a continuous time MPP framework by Chen and Epstein (2002).

Similarly to Maenhout, also Sbuelz and Trojani (2008) introduce a time-state varying constraint in the formulation (4), specializing the results on the effects of a preference for robustness (ambiguity aversion) to the power utility case. They assume that the exogenous primitives of the economy follow a Ito process with BM B and to limit their analysis to models that are statistically indistinguishable from the approximating one, the perturbation $\boldsymbol{\theta}$ to B is constrained as $0.5\boldsymbol{\theta}'\boldsymbol{\theta} \leq \eta f^2(x)$, where η is a non-negative constant and f is a function of the current state x of the economy, to allow for state- and time-dependent ambiguity. Sbuelz and Trojani show that ambiguity aversion has a direct impact only on the myopic and the hedging demands for equity. The impact of ambiguity on the equity demand is state-dependent in a non-standard way: when equity risk barely pays off and there is limited need for hedging, or equity is an inadequate hedging tool, the preference for robustness generates the *first order risk aversion* (FORA) effect and the desired equity holdings are small even if ambiguity aversion is low. Sbuelz and Trojani quantify the impact of ambiguity by calibrating to real data some special cases, assuming approximately unit risk aversion and tiny amounts of ambiguity aversion. Their simulation findings confirm that ambiguity aversion reduces the risk-free rate but only indirectly affects equity returns and worst-case equity premia. In fact, log utility and worst-case equity premia remain completely unaffected.

Papers like Anderson et al. (2003), Maenhout (2004), and Sbuelz and Trojani (2008) have all derived key insights for the equilibrium quantities and the empirical properties of the SDF. In this vein, Hansen (2007) has illustrated in more general terms how statistical ambiguity—both intended as agents' inability to fully characterize the probabilistic model of reference, and as the econometrician's inability to infer the model actually used by economic agents—may create small errors that have however a potential to explain (generate) large asset pricing anomalies. Hansen's (2007) novelty consists in establishing a link between ambiguity aversion and the concept of *Chernoff's entropy*. Chernoff's entropy is often used in statistics to test a particular model against a competitor and it is defined as the rate at which the logarithm of the probability of the test model being wrong drops to zero. Hence, lower values of the Chernoff rate indicate that the errors deriving from the choice of a particular model are relatively small, or, equivalently, that the model is statistically indistinguishable from the "true" one. Using a simple framework, Hansen shows that models characterized by low Chernoff rates generate extremely different Sharpe-ratios relative to a "real" data-driven model, so that statistical ambiguity may greatly contribute to explaining a range of asset pricing puzzles. In particular small errors in the estimate of the mean assets return, represented by small Chernoff rates, induce sizable variations of the Sharpe ratio and large pricing anomalies, e.g., a Chernoff rate of 1% per annum increases the maximum Sharpe ratio by 0.14.

Barillas, Hansen and Sargent (2009, BHS) have followed up on Hansen's intuitions and have studied the problem of using ambiguity to produce SDFs with properties that are consistent with high mean excess returns and low real risk-free rates from a different angle. The equity premium and risk-free rate puzzles have been often expressed in terms of violations of Hansen and Jagannathan's (1991, HJ) pricing bounds. The equity premium puzzle arises because to reconcile the HJ bound with empirical measures of the market price of risk, an unreasonably high relative risk aversion coefficient (γ) is needed. Moreover, high values of γ deliver high market prices of risk but also push the reciprocal of the risk-free rate away from the HJ bound, determining the risk-free rate puzzle. Noting the observational equivalence between Kreps and Porteus' (1978) risk-sensitive preferences and MPP, Barillas, Hansen and Sargent have reinterpreted the coefficient γ calibrated by Tallarini (2000) to match asset market data as a measure for the consumer's doubts about the model specification. Consequently, the risk premia commonly estimated from asset market data would be measures of the benefits of the reduction in the uncertainty on the unknown model specification, and not only of the aversion to stochastic, aggregate endowment fluctuations. To discuss the plausibility of the parameter for uncertainty aversion calibrated on asset market data, BHS use *detection error probabilities* and find that values of these probabilities that imply a non-overly cautious attitude toward uncertainty yields market prices of risk that are compatible with HJ bounds.[32]

[32]Given a reference model A, and letting B^{ϱ} be the distorted worst-case model associated with the ambiguity aversion parameter ϱ, consider a fixed sample of observations on a problem-specific state variable x_t, $t = 0, ..., T-1$. Letting L_{ij} be the likelihood of that sample for model j assuming that model i generates the data, define the log-likelihood ratio as $r_i \equiv \ln(L_{ii}/L_{ij})$ ($j \neq i$, $i = A, B^{\varrho}$). Therefore, if we assume that model A has generated the data, $p_A = Prob(mistake|A) = freq(r_A \leq 0)$. Similarly, $p_{B^{\varrho}} = Prob(mistake|B^{\varrho}) = freq(r_{B^{\varrho}} < 0)$. The detection error probability associated with the ambiguity aversion parameter ϱ is defined as $p(\varrho) = \frac{1}{2}(p_A + p_{B^{\varrho}})$.

4.4. *Smooth Ambiguity Preferences and the SDF*

There is another strand of the ambiguity literature that has tackled the typical asset pricing questions using ambiguity preferences of the KMM type. For instance, Ju and Miao (2011) have developed a pure exchange economy model in which aggregate consumption follows a hidden Markov regime switching process, and the agent is ambiguous about the hidden state process. Preferences are represented by a recursive dynamic generalization of KMM preferences, as in Section 3.5. Ju and Miao (2011) work with the recursive functional

$$V_t(c) = W(c_t, \Gamma_t(V_{t+1}(c))) \qquad \Gamma_t(\cdot) \equiv \upsilon^{-1}\left(\int_{\wp_{x_t}} \upsilon\left(U^{-1}\int U(\cdot)dp\right) dq_t(p)\right), \tag{65}$$

where x_t is the realization of a hidden state x at time t. x follows a K-state Markov chain, and each of the realizations can be interpreted as a possible state of the economy: different realizations corresponds to different probability measures in the set $\wp_{x_t}$. The risky asset pays dividends D_t and has gross return R^e_{t+1} over the interval $[t, t+1]$. There is also a risk-free bond whose return is R^f_{t+1}. In equilibrium, consumption and dividends will be the same and consumption growth is governed by a Markov switching process:

$$\ln(c_{t+1}/c_t) = \mu_{x_{t+1}} + \sigma\epsilon_{t+1} \quad \epsilon_t \sim IIDN(0,1). \tag{66}$$

Setting $U(c) = c^{1-\gamma}/(1-\gamma)$, $\upsilon(x) = x^{1-\eta}/(1-\eta)$, and $W(c,\Gamma) = \left[(1-\beta)c^{1-\psi} + \beta\Gamma^{1-\psi}\right]^{\frac{1}{1-\psi}}$ with $\psi > 0$, the DM displays ambiguity aversion if and only if $\eta > \gamma$. η can be interpreted as an ambiguity aversion parameter, γ is the standard risk aversion coefficient, and $1/\psi$ represents the elasticity of intertemporal substitution. Under this particular specification, assuming that x can take only two possible state values (1 or 2, with 1 being a recession and 2 an expansion), beliefs are updated using Bayes' Rule. By standard principles, the corresponding SDF is the intertemporal marginal rate of substitution, here

$$\mathcal{M}_{x_{t+1}} = \beta\left(\frac{c_t}{c_{t+1}}\right)^{-\psi}\left(\frac{V_{t+1}}{\boldsymbol{R}_t(V_{t+1})}\right)^{\psi-\gamma} E_{t,x_{t+1}}\left(\frac{V_{t+1}^{1-\gamma}}{\boldsymbol{R}_t(V_{t+1})}\right)^{-(\eta-\gamma)}. \tag{67}$$

The last multiplicative factor reflects ambiguity. In the case of ambiguity neutrality, $\eta = \gamma$, this term vanishes. When the representative agent is ambiguity averse, a recession ($\mu_{x=1}$) is associated with a high value of the SDF. Intuitively, the agent has a lower continuation value V_{t+1} in a recession state, causing the multiplicative adjustment to take a higher value. This pessimistic behavior reduces the stock price and raises the expected stock return; furthermore, it reduces the risk-free rate because the agent wants to save more for the future. Because stocks do poorly in recessions when agents have higher marginal utility, ambiguity aversion generates a high negative correlation between the SDF and stock returns that increases the equity premium while a high conditional mean for the SDF decreases the real interest rate.[33]

[33] A calibration shows that the model can match mean excess stock returns on the market portfolio, the mean risk-free rate, and the observed volatility of excess return and its relationship with the (conditional, state-dependent) equity premium. However this result may be driven more by the separation between risk aversion and the elasticity of intertemporal substitution than by ambiguity aversion. Indeed, without such a separation, the equity premium can be matched, but not its volatility. In this sense, an important contribution is Collard et al.'s (2011), who provide a good match of both the level and the volatility of assets returns and of the counter-cyclicality of the market premium. Collard et al. (2011) use an exponential ambiguity function as opposed to Ju and Miao's power specification. Moreover, apart from being uncertain about the mean of the dividends' distribution, the agent is also uncertain about the magnitude of the parameter determining the persistence of the temporary shock on such a mean. These features allow them to match data even without a separate elasticity of intertemporal substitution.

4.5. *Issues with Asset Pricing Applications of Models of Ambiguity*

Although the papers we have surveyed so far have yielded a number of insights, research exists that has cautioned against the dangers of careless applications of ambiguity in financial modelling. For instance, Gollier (2009) and Maccheroni, Marinacci and Ruffino (2010, MMR) have warned that generalizations of the qualitative (calibration) results may be problematic, because some of their implications for the equity premium may derive not from KMM-type ambiguity per se, but also from the specific parametric models for risk and uncertainty that have been adopted. Indeed, while it has been suggested that ambiguity aversion will usually end up make individuals more cautious, thereby offering a potential explanation for the equity premium puzzle, Gollier proves that an increase in ambiguity aversion does not necessarily imply a reduction in the demand for risky assets. In particular, there exist sufficient conditions on the set of priors $\wp$ under which, *ceteris paribus*, an increase in ambiguity aversion reduces the demand for the ambiguous asset, and raises the equity premium. The intuition for Gollier's findings is that the first-order condition for portfolio optimality can be interpreted as a standard Euler equation under SEU, with a distortion in the compounded probabilities. This distortion is pessimistic and, under standard SEU-theory, pessimistic deteriorations in beliefs are known to not always reduce the demand for risky assets and do not necessarily increase the equity premium. Therefore the general credence that ambiguity will always reduce the demand for risky assets, possibly induce limited participation, and a higher equity premium may have been over-stated. Similarly, in a simple problem with a risk-free asset, a risky asset, and an ambiguous asset, MMR show that their optimal portfolio is not systematically conservative on the share held in the ambiguous asset: in general, it is not true that greater ambiguity reduces the optimal demand for the ambiguous asset.

In a similar vein, Chapman and Polkovnichenko (2009) have recently stressed that when building models that feature non-expected utility preferences, deriving a formal representative investor may be difficult and that even simple models with heterogeneous, non-SEU investors may deliver implications for asset prices that are substantially different from those that a matching representative agent model would imply.[34] In fact, Chapman and Polkovnichenko examine a range of two-date economies populated by heterogeneous agents displaying some common forms of non-expected utility preferences, including MPP, to show that the risk premium and the risk-free rate generated in the models are sensitive to ignoring heterogeneity because of potential limited participation effects a' la Dow and Werlang (1992). Chapman and Polkovnichenko compare the equilibrium prices derived in a model with two agents with different degrees of ambiguity aversion to the prices obtained with a single representative agent who has a wealth-weighted average of the ambiguity aversion parameters of the two-agent economy. The calibration shows that the single (weighted-average) agent model overstates the magnitude of the equity premium. The extent of this bias is always increasing in the difference in ambiguity between the two agents, with the largest difference occurring if one agent is SEU. Similarly, when one of the agents is SEU, the risk-free rate is overstated in the single-agent model (the overstatement is between 5.3% and 26.2% relative to the two-agent model), but this effect is weaker when both agents are ambiguity averse. Intuitively, more ambiguity averse agents may choose to not trade, which leaves the pricing of risk to the remaining agents. In these cases, it is incorrect to use single average agent pricing, since the risk premium is primarily determined by the agents who are most willing to bear risk, and it will differ considerably from the level implied by the "typical" preference parameters.

[34] Bossaerts, Ghirardato, Guarnaschelli and Zame (2010) have provided experimental evidence that attitudes toward ambiguity (characterized through CEU max-min preferences) are heterogeneous across the population, and sufficiently high degrees of ambiguity aversion generate open intervals of asset prices for which these agents refrain from trading.

4.6. *Ambiguity and Learning*

Epstein and Schneider (2008) have extended the results on asset pricing under ambiguity to derive insights on the effects of the quality of information. Financial market participants absorb on a daily basis a large amount of news (signals). Processing a signal involves qualitative judgments: news from a reliable source should lead to stronger portfolio changes than news from an obscure source. When their quality is difficult to evaluate, signals are treated as ambiguous. Epstein and Schneider have considered a simple, three-period asset pricing model populated by a representative agent with recursive MPP à la Epstein and Schneider (2007). Consider a market in which $1/N$ shares of a are traded, where each share is a claim to a dividend that follows the process $d = m + \varepsilon^a + \varepsilon^i$; additionally, a portfolio composed of all the other available assets is traded, with dividend process $\tilde{d} = \tilde{m} + \varepsilon^a + \tilde{\varepsilon}^i$. Here m and $\tilde{m}$ are the mean dividends of the two assets/portfolios, ε^a is an aggregate shock, and ε^i and $\tilde{\varepsilon}^i$ are idiosyncratic, asset-specific shocks. All shocks are mutually independent and normally distributed with mean zero. While trading in the assets occurs at time 1—when a signal s on a is revealed—at time 2 dividends are realized. The signal s has structure $s = \alpha\varepsilon^a + \varepsilon^i + \varepsilon^s$, where $\alpha \geq 0$ measures the degree to which the signal is specific to asset a. If $\alpha = 0$, then the news is 100% company-specific, that is, while it helps to forecast company cash flow d, the signal is not useful for forecasting the payoff $\tilde{d}$ of other assets. Because of ambiguity on signal quality, the variance of ε^s is only known to lie in the interval $\sigma_s^2 \in \left[\sigma_{*s}^2, \sigma_s^{*2}\right]$. There is a single normal prior for $\boldsymbol{\mu}$, the two-dimensional parameter $\boldsymbol{\mu} \equiv [\varepsilon^a + \varepsilon^i \;\; \varepsilon^a]'$, that agents try to infer from the signal s, and a set of normal likelihoods for s parameterized by σ_s^2. The set of one-step ahead beliefs about s at date 0 consists of normals with mean zero and variance $\alpha^2\sigma_a^2 + \sigma_i^2 + \sigma_s^2$. Denoting by $\xi(\sigma_s^2)$ the slope coefficient of the linear projection of $\varepsilon^a + \varepsilon^i$ on s, it is easy to check that:

$$\xi(\sigma_s^2) = \frac{Cov(\varepsilon^a + \varepsilon^i, s)}{Var(s)} = \frac{Cov(\varepsilon^a + \varepsilon^i, \alpha\varepsilon^a + \varepsilon^i + \varepsilon^s)}{Var(\alpha\varepsilon^a + \varepsilon^i + \varepsilon^s)} = \frac{\alpha^2\sigma_a^2 + \sigma_i^2}{\alpha^2\sigma_a^2 + \sigma_i^2 + \sigma_s^2}. \tag{68}$$

Given s, the posterior density of $\boldsymbol{\mu}$ is also normal. As σ_s^2 ranges over $\left[\sigma_{*s}^2, \sigma_s^{*2}\right]$, the coefficient $\xi(\sigma_s^2)$ also varies, tracing out a family of posteriors. The coefficient $\xi(\sigma_s^2)$ measures the information content of the signal relative to the volatility of the parameter, since it determines the fraction of prior variance in $\boldsymbol{\mu}$ that is resolved by the signal; under ambiguity, the larger is the difference $\xi(\sigma_{*s}^2) - \xi(\sigma_s^{*2})$, the less confident the investor is about the true information content.

Epstein and Schneider focus on the case in which the representative agent is risk neutral, she does not discount the future, and derives utility only from consumption at time 2. In the standard Bayesian case, this assumptions imply that the prices of asset a at dates 0 and 1 are given by $p_1^{SEU}(s) = m + \xi(\sigma_s^2)s$ and $p_0^{SEU} = E[p_1^{SEU}|s] = m$, respectively. At date 0, the expected present value is simply the prior mean dividend m, while at date 1 it is the posterior mean dividend, provided that the signal is informative ($\xi(\sigma_s^2) > 0$). The implied equilibrium prices are different under ambiguity. The price of the asset at time 1 is:

$$p_1^{AA}(s) = \min_{\sigma_s^2 \in [\sigma_{*s}^2, \sigma_s^{*2}]} E[d|s] = \begin{cases} m + \xi(\sigma_s^{*2})s & s \geq 0 \\ m + \xi(\sigma_{*s}^2)s & s < 0 \end{cases}. \tag{69}$$

$p_1^{AA}(s)$ is a piece-wise linear function kinked at $s = 0$ and with lower slope over $s \geq 0$, differently from $p_1^{SEU}(s)$ which is linear in $\xi(\sigma_s^2)$. The arrival of an ambiguous signal at date 1 is anticipated at date 0, and since the date 1 price is concave in the signal s, the date 0 conditional mean return is minimized by selecting the highest possible variance σ_s^{*2}. p_0^{AA} exhibits an ambiguity premium which is directly related to the extent of ambiguity measured by $\xi(\sigma_{*s}^2) - \xi(\sigma_s^{*2})$. Such a premium is increasing in the volatility of fundamentals, including the volatility of idiosyncratic risk. Ambiguous company-specific news induce a premium, whose size depends on total risk. Further, prices depend on the prospect of low information quality, in the sense that if it is known today that information on asset a will be difficult to interpret in the future, asset a will be less attractive (cheaper) today, whereas in a Bayesian framework, any feedback of future information quality on current utility is precluded. Moreover, while under SEU the law of large numbers implies that the variance of the market portfolio tends to zero as N becomes large, so that the value of any portfolio converges to m, under ambiguity the market portfolio does not become less uncertain as the number of assets increases.

This simple three-period model can be embedded into an infinite-horizon framework by chaining together a sequence of short learning episodes, in which the agent observes an ambiguous signal u_{t+1} about the next innovation in dividends before that innovation is revealed, and by assuming a mean-reverting process for dividends with mean reversion $(1-\kappa)$ and Gaussian shock u. In equilibrium, Epstein and Schneider prove that the (time-varying) ambiguity premium is the sum of two components. The first term captures the response to the ambiguous signal. As before, this response is asymmetric—bad news is incorporated into prices with a stronger effect. The strength of this reaction depends on the persistence of dividends: if κ is small (i.e., shocks are highly persistent), the effect of news on prices is stronger since the information becomes more important for later payoffs. The second term captures anticipation of future ambiguous news and, as before, it may reflect compensation for asset-specific shocks. The per-share excess returns,

$$\begin{aligned}\hat{R}^{AA}_{t+1} &= p^{AA}_{t+1} + d_{t+1} - p^{AA}_{t}(1+r^f) = \frac{1}{r^f+\kappa}\left(\begin{cases}\xi(\sigma_s^{*2})s_{t+1} & s_{t+1}\geq 0\\ \xi(\sigma_{*s}^{2})s_{t+1} & s_{t+1}<0\end{cases}\right) + \\ &+\frac{1+r^f}{r^f+\kappa}\left[u_{t+1} - \frac{1}{r^f+\kappa}\left(\begin{cases}\xi(\sigma_s^{*2})s_{t} & s_{t}\geq 0\\ \xi(\sigma_{*s}^{2})s_{t} & s_{t}<0\end{cases}\right)\right] + \left(\xi(\sigma_{*s}^{2}) - \xi(\sigma_s^{*2})\right)\frac{\sigma_u}{\sqrt{2\pi\xi(\sigma_s^{*2})}}. \qquad (70)\end{aligned}$$

are the sum of three terms, the first two of which would appear also in the absence of ambiguity. The first one is the response to the current, ambiguous signal. The second term reflects the realization of dividends, and it is proportional to the difference between the current innovation to dividends u_{t+1} and the response to last period's projection about that innovation. These two terms are independent of each other because so are the signal and the shock. Finally, the last (constant) positive term compensates investors for enduring low information quality in the future. The mean excess return under the true probability is positive which is to be contrasted with $E^*[\hat{R}^{SEU}_{t+1}]=0$ which obtains in a risk-neutral framework with no aversion to ambiguity. Moreover, the presence of ambiguous news induces an ambiguity premium which is increasing in σ_u, providing a theoretical foundation for idiosyncratic risk pricing. Epstein and Schneider show that at the point $\xi(\sigma_{*s}^2) = \xi(\sigma_s^{*2})$, the derivatives of the volatility of equilibrium returns under ambiguity with respect to both $\xi(\sigma_{*s}^2)$ and $\xi(\sigma_s^{*2})$ are positive. This is in sharp contrast to the Bayesian case, where price and return volatilities move in opposite directions. Further, the volatility of prices and returns can be much larger than the one of fundamentals, since, if the range of precision contemplated by ambiguity-averse agents is large, they will often attach more weight to a signal with respect to agents who know the true precision. Hence, ambiguous information can cause large price fluctuations. This finding relates to the excess price volatility puzzle (e.g., Shiller, 1981). Because ambiguity-averse market participants respond asymmetrically to news, the model has also the potential to produce skewed distributions for prices and returns, even though both dividends and noise have symmetric (normal) distributions. Interestingly, negative skewness should be more pronounced for stocks for which there is more ambiguous information. This feature helps explaining the existing evidence on skewness in the cross-section of stocks, see e.g., Harvey and Siddique (2000).

Illeditsch (2009) has extended the results in Epstein and Schneider (2008) in a number of directions: investors are allowed to be averse to both risk and ambiguity, they are heterogeneous in their aversion to risk and ambiguity and they can learn from ambiguous signals over time.[35] The model assumes that investors know the marginal distribution of fundamentals (the aggregate dividend) but are ambiguous about the conditional distribution of the signal given the dividend (the precision of the signal). The incorporation of both risk and ambiguity aversion turns out to be important: not only risk and ambiguity aversion generate very different qualitative implications for the equilibrium signal-to-price maps, but risk aversion also amplifies the effects of ambiguous information on the conditional distribution of stock returns, determining an interesting non-linear compounding effect.

4.7. *Other Asset Classes*

Although the literature on the effects of ambiguity on other asset classes is in no way negligible for its contributions or the elegance of the frameworks—see e.g., the applications by Boyarchenko (2009), Montesano (2008), Miao and Wang (2006), and Pritsker (2010)—we will limit to briefly discuss in this Section three papers that—besides their specific merits—have to be taken as mere examples.

Liu, Pan, and Wang (2002, LPW) have examined the option pricing implications of imprecise knowledge concerning rare events that affect the aggregate endowment using an approach that is inspired by HS robustness. They consider a pure exchange economy with one representative agent, a perishable good and a stock which is a claim on aggregate endowment. The latter is affected by two types of random shocks: one is diffusive in nature, capturing the daily fluctuations in fundamentals, while the other is a pure jump, representing rare events. While the probability law of both types of shocks can be estimated using time series data, the attainable precision for rare events is much lower than the one of diffusive shocks. LPW prove that the equilibrium equity premium can be decomposed into a diffusive risk premium, a jump risk premium, and rare event premium that depends on the worst-case distribution associated to the optimal consumption level. In particular the first term depends on risk aversion only, while the magnitude of the rare event premium also depends on the risk aversion parameter of the investor, but it is generated by uncertainty aversion, since the premium stays positive even when risk aversion goes to zero, while it drops to zero in the absence of ambiguity. LPW show that in this framework, European-style options can be priced using the celebrated Black and Scholes formula simply modified to take in to account ambiguity aversion through the worst-case distribution associated to the optimal consumption level. They use one-month S&P 500 European-style options to test their formula. Without aversion to uncertainty, at-the-money calls (with strike close to spot price) have an implied volatility of approximately 15.2% while out-of-the-money put options have a slightly higher implied volatility, as a result of risk aversion. Hence, the equilibrium model generates a "smile" curve—i.e., implied volatilities that are higher for options with strike prices very different from the current price of the underlying—that is feeble when compared to the pronounced "smirk" patterns observed in actual markets. On the contrary, under ambiguity aversion, at-the-money options have implied volatility of 15.5% and the model with uncertainty aversion predicts a premium of about 2% for options away from at-the-money. Furthermore, this implied volatility premium

[35] In Illeditsch (2009) investors share the same information but their interpretation varies because of heterogeneity in risk and ambiguity aversion. A novel avenue for research is the case in which information is asymmetric. For instance, using a model characterized by both ambiguity averse and SEU agents, Caskey (2009) has shown that mispricings may occur because ambiguity averse investors may prefer to reduce the perceived ambiguity using aggregate signals, even at the cost of losing some of the available information. Under these circumstances, equilibrium prices may fail to reflect all the available information, so that profit opportunities for SEU agents arise and under/over-reactions to announcements may occur.

becomes even more pronounced for out-of-the-money puts, that are typically highly sensitive to adverse, rare events.

Another field of asset pricing in which it appears natural to test the implications of ambiguity is the pricing of fixed income securities. For instance, Ulrich (2011) has investigated the effects of ambiguity on the inflation process for the term structure of nominal risk-free yields. Given their model of inflation dynamics, investors face inflation risk and require an inflation risk premium for bearing unanticipated shocks to inflation. However they may be uncertain about the true statistical process of inflation and require an ambiguity premium to protect themselves against any change or misspecification in their underlying inflation model. Ulrich builds and estimates a dynamic, equilibrium three-factor model for the nominal term structure which accounts for these two sources of inflation premia.[36] In Ulrich's model, a representative agent has time separable, logarithmic preferences over consumption holdings and real monetary holdings. The agent holds the capital stock and owns a linear technology which linearly produces an output good; he can invest a fraction of his wealth in this production technology and the rest in a nominal risk-free bond. The real return of the investment into the nominal bond is not known ex-ante, because inflation is stochastic. Inflation is affected by three random state variables: productivity, inflation volatility, and the inflation drift. The central bank is assumed to control an additional monetary shock which is orthogonal to the state shocks. The agent maximizes his life-time utility while taking into account that the real return on the nominal bond is risky and uncertain. Further, she expresses a preference for robustness by insuring herself against a worst case inflation distortion, *which the central bank chooses.* The entropy constraint specifies that the investor protects herself against conditional inflation drift distortions which do not exceed an upper bound η_t, which is allowed to be time-varying and stochastic. Ulrich proves that the optimal degree of inflation distortion that the ambiguity averse agent takes into account is $2\sqrt{\eta_t}$. Given this endogenous degree of robustness, the agent adjusts his probability measure and solves a standard maximization problem obtaining the consumption and money demand policy functions which turn out not to be affected by inflation ambiguity which instead affects the expected excess return of the inflation-sensitive nominal bond. The mirror image of consumption not being affected by inflation ambiguity is that the real interest rate and the market price of output risk are not affected by inflation ambiguity, either, so that they coincide with a standard Cox-Ingersoll-Ross (1985, CIR) economy. Since the endogenous process for inflation is affected by the ambiguity on monetary policy, Ulrich shows that the difference between real and nominal bond prices coincides with the inflation ambiguity premium, that is the reference model covariance between the nominal intertemporal marginal rate of substitution (IMRS) in consumption and the ambiguity kernel. Splitting up this term shows that in general, inflation ambiguity has two channels of impact on the nominal term structure. One is through co-variation of the ambiguity kernel with the IMRS and the second is through the co-variation of the ambiguity kernel with the price deflator. However, because the first covariance is zero because (by construction) the source of ambiguity affects only the price deflator and not the IMRS—because it does not affect consumption—the *inflation ambiguity premium* boils down to a simple covariance between the ambiguity kernel and the price deflator, and it equals the distance between the ambiguity-adjusted expected inflation rate and the "true" expected inflation rate.

Additionally, Ulrich proves an equivalence between the inflation ambiguity premium and the inflation variance

[36] Gagliardini, Porchia and Trojani (2008) have proposed a CIR-style real economy where the representative investor faces model uncertainty with regard to the production technology and only analyze the impact of ambiguity on the term structure of *real* (inflation-indexed) bonds. They model ambiguity as a preference for robustness on misspecifications of the process of the technology shock; the alternative probabilistic models are specified in terms of probability distributions q^{θ}, such that the vector state process is a BM under q^{θ} and the vector $\boldsymbol{\theta}$ satisfies $\boldsymbol{\theta}'\boldsymbol{\theta} \leq 2\eta$, for a constant ambiguity aversion parameter $\eta \geq 0$.

premium by assuming that the upper bound of inflation misspecifications moves with the volatility of inflation v_t and the stochastic technological innovation z_t

$$\eta(\mathbf{X}_t) = \frac{1}{2}\left(q_{a_1}\sqrt{v_t} + \frac{q_{a_2}}{\sqrt{v_t}} + q_{a_3}\frac{z_t}{\sqrt{v_t}}\right)^2, \tag{71}$$

for given q_{a_1}, q_{a_2}, and q_{a_3}. The assumed entropy constraint preserves the simple affine bond pricing structure that is known from models like CIR's. It also leads to the inflation ambiguity premium being proportional to the variance of inflation. Intuitively, when the volatility of inflation increases, it becomes difficult to estimate the drift of inflation precisely; the agent therefore doubts his underlying inflation model more if he observes more dispersed inflation realizations. The model is tested using data on US government bonds to ask whether there is evidence that investors demand an inflation ambiguity premium and whether this premium helps to explain movements in the nominal yield curve. Ulrich finds that the term structures of the inflation ambiguity premium are upward sloping. The inflation ambiguity premium is negative for short-maturity bonds and positive for long-maturity bonds. The term structure of inflation expectations is flat and the term structure of the inflation risk premium is downward sloping, high for short-maturity bonds and slightly negative for long-maturity bonds. The inflation risk premium, a measure of the conditional co-variation of inflation and consumption, is high for a horizon of up to two years, and quickly mean reverts to zero for longer maturities. Similarly to Gagliardini, Porchia and Trojani (2008), inflation ambiguity and ambiguity premia help to explain deviations from the expectation hypothesis. Finally, Ulrich finds that the estimated inflation ambiguity premium plays an important role in explaining the variance of nominal yields.

Ilut (2010) studies equilibrium exchange rate determination in a typical MPP framework in which an ambiguity-averse agent solves a signal extraction problem. Ilut uses this model to address one of the most important empirical puzzles in international finance, the uncovered interest rate parity (UIP) puzzle.[37] The only source of uncertainty in the environment is the domestic/foreign interest rate differential which is modelled as an exogenous stochastic process given by the sum of persistent and transitory components, both unobserved. Under MPP preferences, the agent simultaneously chooses a belief about the model parameter values and a decision about how many bonds to buy and sell. The bond decision maximizes expected utility subject to the chosen belief and a budget constraint. The belief is chosen so that, conditional on the agent's bond decision, expected utility is minimized subject to the constraint that the variance considered must belong to an exogenous finite set. The investor chooses this set so that, in equilibrium, the selected variance parameters are not implausible in a likelihood ratio sense. The dynamics is captured through overlapping generations of investors who each live two periods and derive utility from end-of-life wealth. There is one good for which purchasing power parity (PPP) holds. To capture the process of expectation formation, agents use the Kalman Filter which—given the Gaussian shocks and the linear set up—is the optimal filter. Under these assumptions, Ilut solves the model in closed-form. His findings show that, faced with uncertainty, agents choose to base their decisions on pessimistic beliefs so that, compared to the true but unknown underlying data generating process, they underestimate the unobservable state. This drives the differential between the interest rate paid by the bonds in the investment and funding currencies. In practice, because of ambiguity aversion, positive innovations are treated as temporary shocks, while negative innovations are considered persistent. This systematic underestimation is at the basis of the explanation for the UIP puzzle as it implies that agents perceive on average positive innovations in updating the estimate, which creates the possibility of further increases in the demand of

[37] There is a UIP puzzle because while according to a range of standard SEU models, high interest rate currencies should depreciate vis-a-vis low interest rate currencies, in practice the opposite tends to occur.

the investment currency and its consequent appreciation.

5. **Ambiguity in Market Microstructure Research**

A relatively new but promising area of financial economics in which models of ambiguity have found fruitful application is market microstructure, the microeconomic investigation of rational decisions by individual traders and intermediaries. Easley and O'Hara (2010) have developed a model in which illiquidity (defined as the absence of trades despite the existence of bid and ask quotes) arises from ambiguity aversion. Specifically, in their two-period (0 and 1), heterogenous agents framework, J investors display preferences (with exponential utility index and unit risk aversion parameter) a' la Bewley (1986) and, therefore, each of them is willing to trade if and only if the move from the status quo (i.e., the absence of trade) is expected to be utility-improving for every belief in their set of priors. The payoff of the asset is normally distributed with variance σ^2, but there is no general agreement on the mean of the distribution, so that each investor sets the mean to μ_j. At time 1, an unanticipated shock that reduces each agent's estimate μ_j occurs, so that the new expected value of the payoff is $\mu_j^1 = \mu_j\lambda$, $\lambda < 1$. However the magnitude of such a shock is unknown, so that agents are not able to derive a unique value for λ which is only known to belong to $[\lambda_l, \lambda^h]$. Each of the investors holds a set of beliefs about the future value of the risky asset and trades away from her initial portfolio only if the trade is beneficial according to every belief she considers. Denoting by ω_{0j} the optimal investment of agent j at time 0, Easley and O'Hara show that for prices in the interval $[\lambda_l\mu_j - \sigma^2\omega_{0j},\ \lambda^h\mu_j - \sigma^2\omega_{0j}]$ investor j will refrain from trading so that if the equilibrium price belongs to the set

$$\bigcap_{1,2,\ldots,J}[\lambda_l\mu_j - \sigma^2\omega_{0j},\ \lambda^h\mu_j - \sigma^2\omega_{0j}] \neq \oslash, \tag{72}$$

then no trades will occur and a market will freeze. Moreover, a sufficient condition for trade to collapse is $\{(\max_{1,2,\ldots,J}\mu_j)\ /\ (\min_{1,2,\ldots,J}\mu_j)\} < [(1-\lambda_l)/(1-\lambda^h)]$. The left-hand-side of this no-trade inequality is the ratio of the most optimistic trader's mean to the least optimistic trader's mean perception. Thus, when there is more heterogeneity, an equilibrium with no-trade is more difficult to establish because the left hand side will be large. In particular, when prior opinions are diverse, the portfolios that traders bring into period one are diverse, and thus there is more of an incentive for individual traders to move toward the mean portfolio in response to a decline in expectations of future prices. The right-hand side of the no-trade inequality is instead the ambiguity about the percentage decline in future mean values. The maximum (minimum) price in the no trade interval is the lowest (highest) price at which some trader is willing to sell (buy) the risky asset, hence it represents the ask (bid) price. Specifically, denoting by x supply at time 0, and by $\hat{\mu}$ the average (across agents) expected payoff at time 0, we have that:

$$ask = \min_j[\lambda^h\mu_j\text{-}\sigma^2\alpha_{0j}] = \hat{\mu} - \sigma^2 x - \max_j[(1\text{-}\lambda^h)\mu_j] \qquad bid = \max_j[\lambda_l\mu_j\text{-}\sigma^2\alpha_{0j}] = \hat{\mu} - \sigma^2 x - \min_j[(1\text{-}\lambda_l)\mu_j]. \tag{73}$$

Intuitively, the trader who is most optimistic (pessimistic) about the largest (smallest) possible decline in the value of the asset sets the bid (ask) price. Therefore in this model the bid-ask price difference is an ambiguity spread, in contrast to the standard asymmetric information spread. Indeed, in the presence of asymmetric information, the spread reflects the informational advantage that some traders have with respect to knowledge of the asset's true value, while, under their ambiguity averse model, no trader has an advantage. Yet, a bid-ask spread will exist because of the existence of ambiguity. Finally, the midpoint of the bid-ask spread is a reasonable approximation of the fair value of the asset, since the ask and the bid prices can be thought of overestimate and underestimate, respectively, of the true value of the asset.

Rutledge and Zin (2009) have examined the effects of ambiguity on liquidity and shown that as the framework of analysis grows more realistic, it is hard —albeit this remains intuitively sensible—to conclude that ambiguity will decrease liquidity and be the cause of market break-downs. They observe that liquidity breakdowns are particularly acute in markets where traders heavily rely on complex and yet uncertain empirical models for cash flows, as in the derivative markets. Therefore they propose a model in which a monopolistic ambiguity-averse market maker sets bid and ask prices for a European call option, and the bid-ask spread is used as a measure of liquidity. Ambiguity concerns the probability distribution of the underlying security (with price is S_t), and the market maker displays MPP. Time is discrete and, at any date, the market maker observes the orders from the traders and consequently sets her own strategy $\delta_t \in \{-1, 0, 1\}$ (respectively short, no-trade, or long strategy) and the size of her trade s. The demand for the derivative is summarized by the arrival of a random willingness-to-trade $\tilde{v}_t$ signal. If $\tilde{v}_t$ is greater (lower) than or equal to the posted ask (bid) price, a_t (b_t), then a buy (sell) order is received, and the market maker must set $\delta_t = -1$ ($\delta_t = 1$). If $\tilde{v}_t$ lies between the bid and ask prices, no trade takes place ($\delta_t = 0$). The willingness to trade is assumed to be an IID process with $\Phi(v) = \Pr(\tilde{v}_t < v)$. The bid and ask prices determine the likelihood of a trade in the derivative since $\Pr(\delta_t = -1) = [1 - \Phi(a_t)]$, $\Pr(\delta_t = 0) = [\Phi(a_t) - \Phi(b_t)]$, and, $\Pr(\delta_t = 1) = \Phi(b_t)$. The market maker also chooses optimal consumption (c_t) and investment in the underlying security (ω_t) levels at each date. Importantly, the investment in the risky asset after observing a trade in the derivative gives to the market maker the opportunity to partially hedge the realized position in the derivative market. In a two-period framework the market maker solves

$$\max_{\omega} \left\{ U\left(\vartheta_0 - \omega S_0\right) + \beta \min_{p \in \wp} E_p \left[U\left(\vartheta_1 + \omega S_1\right)\right] \right\} \qquad \vartheta_0 = y_0 + \begin{cases} a & \text{if } d = -1 \\ 0 & \text{if } d = 0 \\ -b & \text{if } d = 1 \end{cases}, \tag{74}$$

where d is the dividend paid by the risky asset, and ϑ is the income of the market-maker. Hence, $U(a, b, d)$ is the indirect utility for given (a, b, d). At time 0, bid and ask prices are chosen by solving

$$\max_{a,b} \left\{ [1 - \Phi(a)] U_a + [\Phi(a) - \Phi(b)] U_0 + \Phi(b) U_b \right\}, \tag{75}$$

where $U_a \equiv (a, b, -1)$, $U_b \equiv (a, b, 1)$, and $U_0 \equiv (a, b, 0)$. Choosing a high (low) value for the ask (bid) price generates a higher payoff should a high (low) value order arrive; however, it also lowers the probability of such a trade occurring, so that the period 0 income effect is offset by the period 1 income effect of the derivative's payoff. The first order conditions of the problem reveal that the presence of ambiguity lowers the baseline level of utility a market makers may attain in the absence of trades. For example, in the case where the optimal portfolio with no position in the derivative yields consumption that is close to riskless (and hence is not affected by ambiguity), the level of such a baseline utility (call it U_0) for an ambiguity averse and an ambiguity neutral market maker are approximately equal, while the ambiguity bid-ask spread is larger. However, if U_0 is affected by ambiguity, then it is possible that the ambiguity-averse market-maker posts a bid or ask that is more aggressive. This occurs when the derivative position "hedges ambiguity." In particular, the difference $U_0 - U_{a=0}$ may be smaller for an ambiguity-averse market-maker if the worst-case distribution in the U_0 case differs from the one in the $U_{a=0}$ case, so that the ambiguity-averse market-maker may post a more aggressive (lower) ask price. Hence, uncertainty aversion does not always lower liquidity.

Ozsoylev and Werner (2009) propose a different approach to connect ambiguity to liquidity and trading volume in a typical micro-structure model that emphasizes the interaction between uncertainty and private information. In Ozsoylev and Werner's static framework, all investors have MPP preferences for asset payoffs; while informed agents receive a private signal that resolves the ambiguity, arbitrageurs, instead, cannot observe such a signal and extract information from prices, so that their beliefs continue to be affected by ambiguity. Ozsoylev and Werner show that if the asset's supply is deterministic, the equilibrium is fully revealing, otherwise it is only partially revealing, which echoes standard SEU findings. However, there exists a range of values of the signal and the random asset supply such that the arbitrageurs choose not to trade. When there is ambiguity, under limited participation the sensitivity of prices to the signal and to the random supply is lower than under full participation. Using reciprocals of price sensitivities as measures of market depth, they show that limited participation also induces lower market depth.

6. **Discussion and Conclusions**

We have surveyed a body of work that has brought models of decision making under ambiguity at the forefront of research in financial economics. A number of insights have emerged but, as usual, considerable work remains to be done. An obvious avenue for future research consists of generalizations of the existing models. Cerreia-Vioglio et al. (2011) represents an important contribution in this direction as it provides a general preference representation that encompasses both smooth KMM and VP. Such an effort is crucial for a simple reason: although some of the asset pricing and portfolio choice predictions are qualitatively similar across different models, other implications either appear uniquely in some specific models, or predictions may even be heterogeneous across the range of existing frameworks. For example, the feature of kinked indifference curve of MPP generates some unique predictions such as market non-participation and multiplicity of pricing kernels; on the other hand, the realistic prediction of countercyclical SDFs and market price of ambiguity appears to be unique to the smooth ambiguity models (KMM). Therefore it is only developing frameworks that may nest the alternative ambiguity-averse preferences in current use that such alternative predictions may be empirically tested.[38] Additionally, the literature appears to be still in a stage in which the need to approach technically complicated problems and to deal with deep issues (e.g., how beliefs should be updated in dynamic models) has often advised to build models that are specialized to the objective at hand. For instance, papers on ambiguity and the equity premium and risk-free rate puzzles are technically sophisticated because they solve dynamic general equilibrium models, but have usually limited the asset menu to two assets only. We expect that ambitious papers will soon pursue more complicated and interconnected questions using dynamic models with rich asset menus.

Our review has also isolated a number of research topics on which only the first initial steps have been moved. Even though it is often the case that the most important breakthroughs may happen exactly where they are not expected, it is safe to say that considerable progress should be welcomed from papers that explore the interaction between ambiguity aversion and the presence of asymmetric information. For instance, Condie and Ganguli (2009) show that if an ambiguity averse investor has private information, then portfolio inertia a' la Dow and Werlang (1992) may prevent the revelation of information by prices in set ups in which we would observe substantial revelation under SEU. Moreover, while a limited number of researchers straddling the macroeconomics and finance camps (e.g., Backus et al., 2004, and Hansen and Sargent, 2007) have made efforts at structuring their models to allow

[38] For instance, Skiadas (2011) has developed a simple asset pricing model for general ambiguity averse, homothetic preferences.

using existing estimation methods for dynamic asset pricing models, a majority of the papers we have reviewed still content themselves with solving models and simulating empirical outcomes from calibrated versions. Although this has been useful to allow researches from all backgrounds to understand what the qualitative implications of ambiguity aversion are, models of ambiguity offer rich sets of empirical restrictions that should be tested giving us a possibility to reject the proposition that ambiguity aversion affects portfolio choices and is priced in equilibrium. Finally, a number of new papers have started to explore the effects of ambiguity aversion on portfolio and consumption choices over an investor's life-cycle, a field of application that currently abounds of puzzle and ill-understood phenomena. The first steps moved by Campanale (2009) and Peijnenburg (2011) show that the pay-offs of ambiguity modelling may be substantial.

A different issue is whether and how the ever growing class of financial models we have surveyed may affect the way in which applied (financial) economists interpret economic phenomena. We must admit that—because the topic of the existence of "risks" in economic decision making that are hardly quantifiable in a SEU perspective had been addressed by a few of the founding fathers of economics and decision theory alike (such as Arrow, Hurwicz, Keynes, Knight, Raiffa, and Wald, in no order whatsoever)—informal arguments connecting poorly understood phenomena to aversion to uncertainty have been popular since the 1930s. Yet, besides an informal appreciation that investment and pricing decisions may reflect an aversion to poor information, it is only recently that these arguments have been made formal. Similarly, given the existence of a handful of papers that have explored how policy makers may deal with ambiguity, it is natural to ask whether there is so far any evidence that these normative ideas have broken into the concerns of policy makers in practice. Even though it is hard to interpret the minds of policy makers, the reaction of a number of central banks and governments to the financial crisis of 2008-2009 may be interpreted as being consistent with the typical actions that may contrast the disruptive effects of ambiguity (see e.g., Caballero and Krishnamurthy, 2008, Easley and O'Hara, 2009, Guidolin and Rinaldi, 2009, and Pritsker, 2010).

References

[1] Al-Najjar, N.I., and J., Weinstein. (2009). "The Ambiguity Aversion Literature: A Critical Assessment." *Economics and Philosophy*, 25, 249-284.

[2] Anderson, E.W., L.P. Hansen and T.J., Sargent. (2003). "A Quartet of Semigroups for Model Specification, Robustness, Prices of Risk, and Model Detection." *Journal of the European Economic Association*, 1, 68-123.

[3] Anderson, E.W., L.P., Hansen, and T.J., Sargent. (1998). "Risk and Robustness in General Equilibrium." Working paper, University of Chicago.

[4] Backus, D., B., Routledge and S., Zin. (2004). "Exotic Preferences for Macroeconomists." *NBER Macroeconomics Annual*, 19, 319-390.

[5] Barberis, N. (2000). "Investing for the Long Run when Returns Are Predictable." *Journal of Finance,* 55, 225-264.

[6] Barberis, N., and R., Thaler. (2003). "A Survey of Behavioral Finance," in *The Handbook of the Economics of Finance* by G. M., Constantinides, M., Harris, and R., Stulz (eds.), volume 1, chapter 18, pp. 1053-1128, Elsevier.

[7] Barillas, F., L.P., Hansen, and T.J., Sargent. (2009). "Doubts or Variability?" *Journal of Economic Theory*, 144, 2388-2418.

[8] Bewley, T. (1986). "Knightian Decision Theory, Part 1." Working Paper, Yale University.

[9] Bossaerts, P., P., Ghirardato, S., Guarnaschelli and W.R., Zame. (2010). "Ambiguity in Asset Markets: Theory and Experiment." *Review of Financial Studies*, 23, 1325-1359.

[10] Boyarchenko, N. (2009). "Ambiguity, Information Quality and Credit Risk." Working Paper, University of Chicago.

[11] Boyle, P., L., Garlappi, R., Uppal, and T., Wang. (2009). "Keynes Meets Markowitz: the Trade-off Between Familiarity and Diversification." Working paper, London Business School.

[12] Caballero, R., and A., Krishnamurthy.(2008). "Collective Risk Management in a Flight to Quality Episode," *Journal of Finance*, 63, 2195-2236.

[13] Camerer, C., and M., Weber. (1992). "Recent Developments in Modeling Preferences: Uncertainty and Ambiguity" *Journal of Risk and Uncertainty,* 5, 235-370.

[14] Campanale, C. (2010). "Learning, Ambiguity, and Life-Cycle Portfolio Allocation." *Review of Economic Dynamics*. In press.

[15] Cao, H.H., T., Wang and H.H., Zhang. (2005). "Model Uncertainty, Limited Market Participation, and Asset Prices." *Review of Financial Studies*, 18, 1219-1251.

[16] Caskey, J. (2009). "Information in Equity Markets with Ambiguity Averse Investors." *Review of Financial Studies*, 22, 3595-3627.

[17] Cerreia-Vioglio, S., F., Maccheroni, M., Marinacci and L., Montrucchio. (2011). "Uncertainty Averse Preferences" *Journal of Economic Theory,* 146, 1275-1330.

[18] Chapman, D.A., and V., Polkovnichenko. (2009). "First-Order Risk Aversion, Heterogeneity, and Asset Market Outcomes." *Journal of Finance*, 64, 1863-1887.

[19] Chateauneuf, A., J., T. Gajdos, and J.-Y., Jaffray. (2010). "Regular updating." *Theory and Decision*, In press.

[20] Chen, H., N., Ju and J., Miao. (2011). "Dynamic Asset Allocation with Ambiguous Return Predictability." Working Paper, Boston University.

[21] Chen, Z., and L.G., Epstein. (2002). "Ambiguity, Risk, and Asset Returns in Continuous Time." *Econometrica*, 70, 1403-1443.

[22] Collard, S., S., Mukerji, K., Sheppard and J-M., Tallon. (2011). "Ambiguity and the Historical Equity Premium". Working paper, Oxford University.

[23] Condie S., and J., Ganguli. (2009). "Ambiguity and Partially-Revealing Rational Expectations Equilibria." Working Paper, University of Cambridge.

[24] Cox, J.C., J.E., Ingersoll and S.A., Ross. (1985). "A Theory of the Term Structure of Interest Rates" *Econometrica*, 53, 385-407.

[25] Dow, J., and S.R., Werlang. (1992). "Uncertainty Aversion, Risk Aversion, and the Optimal Choice of Portfolio." *Econometrica*, 60, 197-204.

[26] Duffie, D., and L.G., Epstein. (1992). "Stochastic Differential Utility." *Econometrica*, 60, 353-394.

[27] Easley, D., and M., O'Hara. (2010). "Liquidity and Valuation in an Uncertain World." *Journal of Financial Economics*, 97, 1-11.

[28] Easley, D., and M., O'Hara. (2009). "Ambiguity and Nonparticipation: The Role of Regulation." *Review of Financial Studies,* 22, 1817-1843.

[29] Easley, D., and M., O'Hara. (2010). "Microstructure and Ambiguity." *Journal of Finance*, In press.

[30] Ellsberg, D. (1961). "Risk, Ambiguity, and the Savage Axioms." *Quarterly Journal of Economics*, 75, 643-669.

[31] Epstein, L.G. (1999). "A Definition of Uncertainty Aversion." *Review of Economic Studies*, 66, 579-608.

[32] Epstein, L.G. (2010). "A Paradox for the 'Smooth Ambiguity' Model of Preference." *Econometrica,* 78, 2085-2099.

[33] Epstein, L.G. and J,. Miao. (2001). "A Two-Person Dynamic Equilibrium under Ambiguity" *Journal of Economic Dynamics and Control*, 27, 1253-1288.

[34] Epstein, L.G., and M., Schneider. (2003). "Recursive Multiple Priors." *Journal of Economic Theory,* 113, 1-31.

[35] Epstein, L.G., and M., Schneider. (2007). "Learning Under Ambiguity." *Review of Economic Studies,* 74, 1275-1303.

[36] Epstein, L.G., and M., Schneider. (2008). "Ambiguity, Information Quality, and Asset Pricing." *Journal of Finance*, 63, 197-228.

[37] Epstein, L.G. and M. Schneider. (2010). "Ambiguity and Asset Markets," NBER Working Paper No. 16181.

[38] Epstein, L.G., and T., Wang. (1994). "Intertemporal Asset Pricing under Knightian Uncertainty." *Econometrica*, 62, 283-322.

[39] Epstein, L.G., and T., Wang. (1995). "Uncertainty, Risk-Neutral Measures and Security Price Booms and Crashes." *Journal of Economic Theory,* 67, 40-80.

[40] Ergin, H., and F., Gul. (2004). "A Subjective Theory of Compound Lotteries." *Journal of Economic Theory,* 144, 899-929.

[41] Etner, J., M., Jeleva, J.-M., Tallon. (2009). "Decision Theory under Uncertainty." Documents de Travail du Centre d'Economie de la Sorbonne No. 2009.64.

[42] Faria, G., J., Correia-da-Silva, and C., Ribeiro. (2010). "Dynamic Consumption and Portfolio Choice with Ambiguity about Stochastic Volatility." Working Paper, Universidade de Porto.

[43] Fishburn, P. C. (1993). "The Axioms and Algebra of Ambiguity." *Theory and Decision*, 34, 119-137.

[44] French, K.R., and J.M., Poterba. (1991). "Investor Diversification and International Equity Markets." *American Economic Review*, Papers and Proceedings, 81, 222-226.

[45] Gagliardini, P., P., Porchia and F., Trojani. (2009). "Ambiguity Aversion and the Term Structure of Interest Rates." *Review of Financial Studies*, 22, 4157-4188.

[46] Garlappi, L., R., Uppal, and T., Wang. (2007). "Portfolio Selection with Parameter and Model Uncertainty: A Multi-Prior Approach." *Review of Financial Studies*, 20, 41-81.

[47] Ghirardato, P., and M., Marinacci. (2002). "Ambiguity Made Precise: A Comparative Foundation." *Journal of Economic Theory*, 102, 251-289.

[48] Ghirardato, P., F., Maccheroni, and M., Marinacci. (2004). "Differentiating Ambiguity and Ambiguity Attitude," *Journal of Economic Theory*, 118, 133-173.

[49] Gilboa I., F., Maccheroni, M., Marinacci., and D., Schmeidler. (2010). "Objective and Subjective Rationality in a Multiple Prior Model" *Econometrica*, 78, 755-770.

[50] Gilboa I., and M., Marinacci. (2011). "Ambiguity and the Bayesian Paradigm", in D. Acemoglu, M. Arellano, and E. Dekel (eds.), *Advances in Economics and Econometrics: Theory and Applications*, Tenth World Congress of the Econometric Society, Cambridge University Press, In press.

[51] Gilboa, I., A. W., Postlewaite and D., Schmeidler (2008). "Probability and Uncertainty in Economic Modeling." *Journal of Economic Perspectives*, 22, 173-188.

[52] Gilboa I., and D., Schmeidler. (1989). "Maxmin Expected Utility with Non-Unique Prior." *Journal of Mathematical Economics*, 18, 141-153.

[53] Gollier, C. (2009). "Portfolio Choices and Asset prices: The Comparative Statics of Ambiguity Aversion." Working Paper.

[54] Guidolin, M. (2005). "Home Bias and High Turnover in an Overlapping-generations Model with Learning." *Review of International Economics*, 13, 725-756.

[55] Guidolin, M. and F. Rinaldi. (2009). "A Simple Model of Trading and Pricing Risky Assets under Ambiguity: Any Lessons for Policy-Makers?" Working paper No. 2009-020A, Federal Reserve Bank of St. Louis.

[56] Hanany, E., P., Klibanoff. (2009) "Updating ambiguity averse preferences". *The B.E. Journal of Theoretical Economics*, 9, Article 37.

[57] Hansen, L.P. (2007). "Beliefs, Doubts and Learning: Valuing Macroeconomic Risk." *American Economic Review*, 97, 1-30.

[58] Hansen, L.P., and R., Jagannathan. (1991). "Implications of Security Market Data for Models of Dynamic Economies." *Journal of Political Economy*, 99, 225-262.

[59] Hansen, L.P., and T.J., Sargent. (1998). "Alternative Representations of Discounted Robust Linear Quadratic Control." Working Paper, University of Chicago.

[60] Hansen, L.P., and T. J., Sargent. (2001). "Robust Control and Model Uncertainty." *American Economic Review*, 91, 60-66.

[61] Hansen, L.P., and T.J., Sargent. (2007). *Robustness.* Princeton University Press: Princeton, NJ.

[62] Hansen, L P., T. Sargent, and T. Tallarini. (1999). "Robust permanent income and pricing." *Review of Economic Studies*, 66, 873-907.

[63] Harvey, C.R., and A., Siddique. (2000). "Conditional Skewness in Asset Pricing Tests." *Journal of Finance*, 55, 1263-1295.

[64] Hirshleifer, D. (2001) "Investor Psychology and Asset Pricing", *Journal of Finance,* 56, 1533-1597.

[65] Honda, T. (2003). "Optimal Portfolio Choice for Unobservable and Regime-Switching Mean Returns." *Journal of Economic Dynamics and Control*, 28, 45-78.

[66] Illeditsch, P. (2009). "Ambiguous Information, Risk Aversion and Asset Pricing." Working Paper, University of Pennsylvanya.

[67] Ilut, C. L. (2010). "Ambiguity Aversion: Implications for the Uncovered Interest Rate Parity Puzzle." Working Paper, Duke University.

[68] Ju, N., and J., Miao. (2011). "Ambiguity, Learning, and Asset Returns." *Econometrica*, In press.

[69] Klibanoff, P., M., Marinacci, and S., Mukerji. (2005). "A Smooth Model of Decision Making under Ambiguity." *Econometrica*, 73, 1849-1892.

[70] Klibanoff, P., M., Marinacci and S., Mukerji. (2009). "Recursive Smooth Ambiguity Preferences." *Journal of Economic Theory*, 144, 930-976.

[71] Klibanoff, P., M., Marinacci and S., Mukerji. (2011). "On the Smooth Ambiguity Model: A Reply." *Econometrica*, In press.

[72] Knight, F.H. (1921). *Risk, Uncertainty, and Profit.* Boston and New York: Houghton Mifflin.

[73] Kogan, L., and T., Wang (2003). "A Simple Theory of Asset Pricing under Model Uncertainty." Massachusetts Institute of Technology, Working Paper.

[74] Kogan, L., and R., Uppal. (2000). "Risk Aversion and Optimal Policies in Partial and General Equilibrium Economies." Massachusetts Institute of Technology, Working Paper.

[75] Kreps, D.M., and E.L., Porteus. (1978). "Temporal Resolution of Uncertainty and Dynamic Choice Theory." *Econometrica*, 46, 185-200.

[76] Leippold M., F., Trojani, and P., Vanini. (2008). "Learning and Asset Prices Under Ambiguous Information." *Review of Financial Studies*, 21, 2565-2597.

[77] Lewis, K. (1999). "Trying to Explain Home Bias in Equities and Consumption." *Journal of Economic Literature*, 37, 571-608.

[78] Liu, H. (2010). "Robust Consumption and Portfolio Choice for Time Varying Investment Opportunities." *Annals of Finance*, 6, 435-454.

[79] Liu, H. (2011). "Dynamic Portfolio Choice under Ambiguity and Regime Switching Mean Returns." *Journal of Economic Dynamics and Control*, 35, 623-640.

[80] Liu, J., J., Pan and T., Wang. (2005). "An Equilibrium Model of Rare-Event Premia and its Implication for Option Smirks." *Review of Financial Studies*, 18, 131-164.

[81] Maccheroni, M., M., Marinacci, and D., Ruffino. (2010). "Alpha as Ambiguity: Robust Mean-Variance Portfolio Analysis". Working Paper, Bocconi University.

[82] Maccheroni, F., M., Marinacci and A., Rustichini. (2006). "Ambiguity Aversion, Robustness, and the Variational Representation of Preferences." *Econometrica*, 74, 1447-1498.

[83] Maccheroni, F., M., Marinacci, A., Rustichini and M., Taboga. (2009). "Portfolio Selection with Monotone Mean-Variance Preferences." *Mathematical Finance*, 19, 487-521.

[84] Maenhout, P.J. (2004). "Robust Portfolio Rules and Asset Pricing." *Review of Financial Studies*, 17, 951-983.

[85] Markowitz, H.M. (1959). *Portfolio Selection: Efficient Diversification of Investments.* Wiley, New York.

[86] Mehra, R., and E.C., Prescott. (1985). "The Equity Premium: A Puzzle." *Journal of Monetary Economics*, 15, 145-161.

[87] Merton, R.C. (1971). "Optimum Consumption and Portfolio rules in a Continuous-Time Model." *Journal of Economic Theory*, 3, 373-413.

[88] Miao, J. (2009). "Ambiguity, Risk and Portfolio Choice under Incomplete Information." *Annals of Economics and Finance*, 10, 257-279.

[89] Miao, J., and N., Wang. (2011). "Risk, Uncertainty, and Option Exercise." *Journal of Economic Dynamics and Control*, 35, 442-461.

[90] Montesano, A. (2008). "Effects of Uncertainty Aversion on the Call Option Market." *Theory and Decision*, 65, 97-123.

[91] Mukerji, S. (2010). "Foundations of Ambiguity and Economic Modeling", *Economics and Philosophy*, 25, 297-302.

[92] Mukerji, S., and J.-M., Tallon. (2001). "Ambiguity Aversion and Incompleteness of Financial Markets." *Review of Economic Studies*, 68, 883-904.

[93] Mukerji, S., and J.-M., Tallon. (2004). "An Overview of Economic Applications of David Schmeidler's Models of Decision Making under Uncertainty," in *Uncertainty in Economic Theory*, by Itzhak Gilboa (editor)., pp. 283-302. Routledge, London.

[94] Nau, R. (2006). "Uncertainty Aversion with Second-Order Utilities and Probabilities." *Management Science*, 52, 136-145.

[95] Ozsoylev, H., and J., Werner. (2009). "Liquidity and Asset Prices in Rational Expectations Equilibrium with Ambiguous Information." Working Paper, University of Minnesota.

[96] Pastor, L. (2000). "Portfolio Selection and Asset Pricing Models." *Journal of Finance* 55, 179-223.

[97] Peijnenburg, K. (2011). "Life-Cycle Asset Allocation with Ambiguity Aversion and Learning", Working Paper, Bocconi University and IGIER.

[98] Pires, C.P. (2002). "A Rule For Updating Ambiguous Beliefs." *Theory and Decision*, 53, 2, 137-152.

[99] Pritsker, M. (2010). "Knightian Uncertainty and Interbank Lending." Working Paper, Board of Governors, Federal Reserve System.

[100] Rinaldi, F. (2009). "Endogenous Incompleteness of Financial Markets: the Role of Ambiguity and Ambiguity Aversion.." *Journal of Mathematical Economics*, 45, 2, 880-901.

[101] Routledge, B.R., and S.E., Zin. (2009). "Model Uncertainty and Liquidity." *Review of Economic Dynamics*, 12, 543-566.

[102] Savage, L. (1954). *The Foundations of Statistics.* Wiley, New York.

[103] 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.

[104] Schmeidler, D. (1989). "Subjective Probability and Expected Utility without Additivity." *Econometrica*, 57, 571-587.

[105] Segal, U. (1990). "Two-Stage Lotteries without the Reduction Axiom." *Econometrica*, 58, 349-377.

[106] Segal, U., and A., Spivak. (1990). "First Order versus Second Order Risk Aversion." *Journal of Economic Theory*, 51, 111-125.

[107] Skiadas, C. (2003). "Robust Control and Recursive Utility." *Finance and Stochastics*, 7, 475-489.

[108] Skiadas, C. (2011). "Scale Invariant Asset Pricing Theory: A General Discrete Framework with Ambiguity-Aversion Applications", Working Paper.

[109] Strzalecki, T. (2011), "Axiomatic Foundations of Multiplier Preferences", *Econometrica*, 79, 47-73.

[110] Tallarini, T.D. (2000). "Risk-sensitive real business cycles." *Journal of Monetary Economics*, 45, 507-532.

[111] 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.

[112] Trojani, F., and P., Vanini. (2004). "Robustness and Ambiguity Aversion in General Equilibrium." *Review of Finance*, 8, 279-324.

[113] Ulrich, M. (2011). "Inflation Ambiguity and the Term Structure of Arbitrage-Free U.S. Government Bonds.", *Journal of Monetary Economics*, In press.

[114] Uppal, R. and Wang, T. (2003). "Model Misspecification and Underdiversification." *Journal of Finance,* 58, 2465-2486.

[115] Veronesi, P. (1999). "Stock Market Overreaction to Bad News in Good Times: A Rational Expectations Equilibrium Model." *Review of Financial Studies*, 5, 975-1007.

[116] von Neumann, J., and O., Morgenstern. (1947). *Theory of Games and Economic Behavior*. Princenton, NJ: Princeton University Press.

[117] Wakker, P.P. (2001). "Testing and Characterizing Properties of Nonadditive Measures through Violations of the Sure-Thing Principle." *Econometrica,* 69, 1039-1059.

[118] Wakker, P.P. (2008). "Uncertainty." in The New Palgrave Dictionary of Economics, second edition, S., Durlauf and L., Blume (eds.). Palgrave Macmillan.

[119] Wang, Z. (2005). "A Shrinkage Approach to Model Uncertainty and Asset Allocation." *Review of Financial Studies*, 18, 673-705.

[120] Weil, P. (1989). "The Equity Premium Puzzle and the Risk-Free Rate Puzzle." *Journal of Monetary Economics*, 24, 401-421.