



ECONOMIC RESEARCH
FEDERAL RESERVE BANK OF ST. LOUIS
WORKING PAPER SERIES

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

Authors	Massimo Guidolin, and Francesca Rinaldi
Working Paper Number	2010-028A
Creation Date	September 2010
Citable Link	https://doi.org/10.20955/wp.2010.028
Suggested Citation	Guidolin, M., Rinaldi, F., 2010; Ambiguity in Asset Pricing and Portfolio Choice: A Review of the Literature, Federal Reserve Bank of St. Louis Working Paper 2010-028. URL https://doi.org/10.20955/wp.2010.028

Federal Reserve Bank of St. Louis, Research Division, P.O. Box 442, St. Louis, MO 63166

The views expressed in this paper are those of the author(s) and do not necessarily reflect the views of the Federal Reserve System, the Board of Governors, or the regional Federal Reserve Banks. Federal Reserve Bank of St. Louis Working Papers are preliminary materials circulated to stimulate discussion and critical comment.

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

Massimo Guidolin*

Federal Reserve Bank of St. Louis and Manchester Business School

Francesca Rinaldi

Manchester Business School

Abstract

A growing body of empirical evidence 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 conduct a survey of the existing 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 empirical 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

Recent research in finance has investigated alternatives to standard models of investors' behavior. Traditional models assume that investors maximize (subjective) expected utility (EU); that agents are perfectly aware of their own preferences mapping from utility-relevant states of the world (such as individual consumption streams or wealth levels) into their perceived welfare; that investors' expectations are not systematically biased (rational expectations). However, a growing body of empirical evidence suggests that investors' behavior is not well described by this traditional paradigm, since actual choices are normatively questionable in the sense of being incompatible with the (S)EU 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; Daniel, Hirshleifer and Subrahmanyam, 1998; Hirshleifer, 2001; Shleifer, 2000; Shleifer and Vishny, 1998). 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., Guidolin and Timmermann, 2007; Lewellen and Shanken, 2002; Timmermann, 1993, 1996; Veronesi, 1999,

*Correspondence to: Research Division, P.O. Box 442, St. Louis, MO 63166, United States. E-mail: Massimo.Guidolin@stls.frb.org; phone: 314-444-8550.

2000). 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 instability of other types, like regimes). A third 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 (S)EU framework, in the sense that the underlying axioms are less demanding. In this area, particular attention has recently been dedicated to ambiguity and ambiguity aversion.¹ In this paper we conduct a systematic—albeit admittedly incomplete—review of the literature that has explored the implications of decision-making under ambiguity for financial market outcomes, such as portfolio choices and equilibrium asset prices.

Under (S)EU, if preferences satisfy certain axioms, there are numerical utilities and probabilities that represent acts (decisions under uncertainty) by a standard weighted sum of the utilities/outcomes deriving from acts in the possible states of the world, where the weights are (subjective) probabilities for each of the states. As innocuous as this basic principle may seem, there is a long, rich tradition of questioning whether it describes behavior adequately. Keynes (1921) was the first to draw a distinction between the implications of evidence—the likelihood judgments that evidence implies—and the “weight” that should be attached to this evidence, or the confidence in the assessed likelihoods. Keynes wondered whether a single probability number could express both dimensions of evidence. Knight (1921) distinguished *risk*, or known probability, from *uncertainty*. He suggested that economic returns could be earned for bearing uncertainty but not for bearing risk. However, the modern attack to (S)EU as a descriptive theory was made most directly by Ellsberg's (1961) paradox that we describe in Section 2.1. Ellsberg's (1961) thought-provoking article stirred a debate and between the 1970s and 1980s led researchers to assemble massive experimental evidence that indicates that people generally prefer the least ambiguous acts. This implies that the experimental subjects 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 (S)EU, the axiom by which a state with a consequence common to a pair of acts is irrelevant in determining preference between the acts.

Why is a survey of the literature on ambiguity in financial markets of any use? First, because the last decade has seen a tremendous growth in the number, breadth, and quality of papers that have exploited ambiguity in connection to research questions in financial economics. In fact, a quick scoring of our own references—as incomplete and deficient as they may be—reveals that out of a total of 195 references, 86 (i.e., a hefty 44%) have been either written (as indicated by the year of the working paper in our possession) or published on or after 2003.² Second, a thought-provoking paper by Al-Najjar and Weinstein (2009) has spurred a debate on whether ambiguity-based approaches to economic modelling would be as sensible as their growing popularity implies. Al-Najjar and Weinstein have been dismissive

¹Whether or not preferences reflecting ambiguity are “behavioral” is mostly a matter of tastes. However, many researchers that have used and critically commented ambiguity averse preferences (e.g., Backus et al., 2004) have noticed that most of these preferences do represent well-defined neo-classical preference orderings. Of course, this is not to be uncritically taken as a “good feature”. On the one hand, the strong theoretical foundations for ambiguity averse preferences allow a researcher to use all the tools of neo-classical economics, particularly optimization and welfare analysis. On the other hand, when these preferences come to ignore aspects of human behavior stressed in other social sciences, particularly sociology and social psychology, they may imply a loss of realism.

²Of course, this is only a very rough scoring system, but we doubt that the stunning percentage we have determined could be much affected by any other sensible assumption.

about the potential of ambiguity models in positive economic analyses. This potential may be fully appreciated without recourse to any presumption about the prescriptive validity of ambiguity-based preference models just by cataloguing and critically discussing the breadth and depth of applications of ambiguity models in finance. This is our main goal and it seems an appropriate time for it. Third, we have noticed that a fraction of the recent papers has 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.

Fourth, the background of our efforts and of the endeavour of 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, (S)EU-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 (S)EU/rational expectations (RE) paradigm. Indeed, the (S)EU hypothesis faces serious difficulties when confronted with asset markets data. Mehra and Prescott (1985) showed that for a standard rational (S)EU representative-agent 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 procyclically (Campbell and Shiller, 1988) and conditional expected equity premia move countercyclically (Fama and French (1989)). Excess returns are serially correlated, mean reverting (Fama and French, 1988b, Poterba and Summers, 1988) and forecastable (Fama and French, 1988a). 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.³ It is time to try and 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 uncovered a number of significant advances in our ability to rationalize important empirical features of asset returns and portfolio decisions as well as episodes of market failures, such as the market freezes that have characterized the recent "credit crunch." With regard to portfolio decisions, we have reviewed and summarized a substantial body of work that has concluded that ambiguity may imply violations of the classical two-fund separation theorem by which all investors should hold an identical risky portfolio with identical structure and can differ only in their relative allocations between the riskless asset and this risky mutual fund. This implies that professional investment advisors may play a key role in the identification of the ambiguity their clients suffer from. Many papers have proved that—even assuming constant investment opportunities over time (e.g., the absence of predictability)—ambiguity will affect the classical Merton-style expressions for optimal portfolio weights and that under a range of parameterizations we should expect ambiguity to decrease the optimal exposure to risky securities. The result that classical (S)EU portfolio weights—both in their myopic (mean-variance style) and in their hedging components—may be severely affected by ambiguity turns out to be robust to the choice of ambiguity averse

³These examples concern research on equity returns and portfolio diversification. Many more poorly understood stylized facts exist with reference to fixed income securities (such as the failure of the expectations hypothesis for risk-free bond yields, after adjusting for time-varying risk premia) or derivative securities (such as the fact that the volatilities implicit in derivative prices fail to predict subsequent, realized volatility). Excellent surveys of these literatures are in Campbell, Lo, and MacKinlay (1996), and Cochrane (2005).

preferences and to the parameterization of the problem. 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 national, country-specific stocks may result in optimal, 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 carrying an inflated, seemingly irrational weight in their portfolios.

In the asset pricing camp, there has been considerable interest in the fact that—when the economy is populated by both (S)EU and by ambiguity-averse investors—equilibrium pricing functions often become discontinuous around a limited participation region, a range of asset prices for which only (S)EU investors will trade. Moreover, in the participating equilibria—when both (S)EU and ambiguity-averse investors trade—it has been argued that changes in “off-equilibrium” (in the sense that these occur with small probabilities) potential outcomes have the potential to affect equilibrium prices. Therefore this literature has stressed the importance of regulation and policy interventions aimed at persuading (ambiguity averse) 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, it has been conjectured that the high risk premia that many papers have explained through ambiguity aversion may be more the result of ambiguity concerning idiosyncratic payoffs than ambiguity on the systematic ones. This may affect the way in which policy-makers fight the effects of ambiguity. 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 dynamic, Lucas-type endowment models, ambiguity has also been found to be a potential cause for asset price indeterminacy and therefore for endogenous volatility (of a “sunspot” type). However, when assumptions are introduced to avoid indeterminacy, then it has been shown that asset risk premia can be decomposed in two parts, a standard (S)EU risk premium component 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.

Researchers have not only busied themselves to show that simple ambiguity-related concerns may provide a solution to many of the puzzles that plague the classical consumption (C)CAPM model, but they have also scored 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 (S)EU, 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 in the sense that for any asset only its individual contribution to total market ambiguity will be compensated. 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 pricing equation, $1 = E[M_{t+1}R_{t+1}|\mathcal{F}_t]$ (where M_{t+1} denotes the stochastic discount factor) by introducing a multiplicative term M_{t+1}^u such that the Euler condition becomes:

$$1 = E[M_{t+1}^u M_{t+1} R_{t+1} | \mathcal{F}_t]. \quad (1)$$

Of course, under special parametric assumptions, much more can be said on the specific structure of the term M_{t+1}^u .

Finally, with reference to the recent financial crisis, papers old and new within the ambiguity literature applied to financial decisions have offered compelling explanations for recently observed phenomena such as trading break-downs (when markets freeze and trading stops altogether) and stubbornly 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. At prices below (above) the lower (upper) limit of the interval, the agent is willing to buy (sell) the asset, but when equilibrium forces fail to push the asset price outside the interval, there will be no willingness to trade. This is in sharp contrast with the standard (S)EU 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. A sufficient condition for ambiguity to induce market breakdowns 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. It is possible that it may not be ambiguity per se that causes the existence of 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 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 exist that 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). However, at least in our reading, none of these papers has the specific focus of our work. For instance, Camerer and Weber (1992), Etner, Jeleva and Tallon (2009), and Wakker (2008) have an explicit focus on defining ambiguity, ambiguity aversion, and how to best model such preferences, with a special focus issues of axiomatization of the resulting criteria and 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 ambiguity preferences.

Our claim that we somewhat differ from existing surveys, means in no way we are “better”. There are at least two

dimensions in which we feel any interested 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 (S)EU preferences. In this sense, we explain in greater depth what ambiguity means than the average finance paper can afford to do. However, that remains our benchmark—to understand enough to be able to appreciate the applied results. Second, we have devoted only occasional attention to the rich experimental literature on ambiguity (more generally, violations of the sure thing principle). This does not contain any implicit judgement on the relevance of such a literature. Camerer and Weber (1992) and Etner, Jeleva and Tallon (2009) review the experimental literature.

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 makers is not given probabilistic information about the external events that affect the outcome of a decision—as opposed to risk, which is “probabilized” uncertainty. We will concentrate on situations in which there is too little information to pin down probabilistic beliefs, as opposed to risky situations, in which objects of choice (gambles, lotteries) are formulated in terms of probability distributions. 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.

The paper has the following structure. Section 2 reviews a few key definitions relating to the concept of ambiguity, its difference from risk, and how aversion to ambiguity might be measured. The purpose of this Section is to allow all Readers to understand the rest of the paper. Section 3 reviews papers that have addressed issues related to optimal portfolio choice under ambiguity aversion. We start with simple, static models that introduce us to 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, to how it may best to introduce robustness in technical terms, and to large scale problems. We also explore the relationships between portfolio choice under incomplete information and ambiguity. 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 difficult but intriguing logical and technical issues, such as dynamic consistency and rational updating of ambiguous beliefs. We also review models that have not just priced stocks and short-term riskless bonds, but also other asset classes, such as equity options, the term structure of nominal, real riskless bonds, and foreign currencies. Section 5 is dedicated to recent advances in financial microstructure theory under ambiguity. Section 6 discusses policy and regulatory implications. Section 7 concludes and discusses the likely (hopeful) impact of ambiguity models on financial economists and policy-makers’ approaches to applied issues.⁴

⁴As a final note of caution, let us stress that our review has no ambition of being complete in the sense of discussing in adequate depth all the relevant papers. For instance, in this paper we say nothing about auctions under ambiguity whereas a number of interesting findings have been published (e.g., the failure of the classical Revenue Equivalence Theorem by which first and second price auctions are Pareto equivalent), see e.g., the seminal paper by Lo (1998) and the introduction to the relevant issues in Mukerji and Tallon (2003). Similarly, we have not covered the papers that have applied ambiguity to macroeconomic issues, unless they explicitly concern either asset allocation

2. Generalities and Definitions: What is Ambiguity?

Let's introduce some bits of notation that will become handy later on. 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 future uncertainty. Sets of states of nature, $E \subset \Omega$, are called events. The outcome space \mathcal{F} contains the possible, random results of any conceivable decision. The outcome space can be rather abstract: although in many applications we can take \mathcal{F} to be the set of real numbers (e.g., wealth), in principle it could be any relevant aspect of a decision problem.⁵ 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 of nature $s \in \Omega$ a possible consequence $f(s)$ (or f_s). $f \succsim g$ means that the decision maker weakly prefers decision f to decision g ; $f \sim g$ means that the decision maker is indifferent between f and g . Most of the time—we will clearly note when this is not the case—all preferences we consider are assumed to be complete (i.e., a decision maker is always able to rank decisions), reflexive ($f \succsim f$) and transitive (i.e., if a decision maker prefers f over g and g over h , then she also prefers f over h). We generally denote by U a standard VNM utility index, and we label by BM standard Brownian motions.

2.1. Early Literature: Ellsberg's Paradox

That a useful distinction might be drawn between standard (SEU) and more general decision-making models has been known to economists since the seminal work by Knight (1921). According to Knight's well-known distinction, there are two kinds of uncertainty: the first, sometimes called *risk*, corresponds to situations in which all relevant events are associated with a (objectively or subjectively) uniquely determined probability assignment; the second, often called *Knightian uncertainty*, corresponds to situations in which some events do not have an obvious probability assignment. Such a distinction is meaningful since, in reality, in most economic contexts where agents face uncertainties, no probabilities are actually given or easily computable.

The experimental relevance of the distinction between risk and uncertainty has been formally discussed by Ellsberg (1961), whose findings have induced researchers to elaborate a range of new preference classes to accommodate for Knightian uncertainty, now more commonly known as *ambiguity*. In the simplest formulation of the paradox, an individual is given the opportunity to bet on the draw of a ball from one of two urns: urn A has 50 red and 50 black balls, urn B has 100 balls which are from some unknown mix of red and black. First, subjects are offered a choice between two bets: \$1 if the ball drawn from urn A is red and nothing if it is black; or \$1 if the ball drawn from urn B is red and nothing if it is black. In experimental implementations of this setting, the first bet has been generally preferred over the second by a majority of the subjects. Therefore, if the agents have a prior on urn B , the predicted probability of red in urn B must be (strictly) *less* than 0.5. Next, subjects are offered a choice between two new bets: one pays out \$1 if the ball drawn from urn A is black and nothing if it is red; the other pays out \$1 if the ball drawn from urn B is black and nothing if it is red. Again the first bet has been generally preferred in experiments. Therefore, if decision makers have a prior on urn B , the predicted probability of black in urn B must be less than 0.5; equivalently,

or asset pricing. See Backus, Routledge and Zin (2004) for an introduction to the role of ambiguity in macroeconomic research.

⁵It will sometimes be 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 $1 - p$ ”.

the probability of red in urn B must be (strictly) *higher* than 0.5. This probability assessment is inconsistent since a unique prior cannot simultaneously assign to the event “red from urn B ” a probability that is strictly less and also strictly more than 0.5! Ellsberg’s (1961) original interpretation was that people are simply *averse to the ambiguity* about the odds for the “ambiguous” urn B : as a result, they would prefer to bet on events with known, rather than ambiguous, odds, consequently ranking bets on the unambiguous urn A over equivalent bets related to B .⁶

Another popular version of the Ellsberg’s paradox clearly shows that ambiguity might induce a violation of the standard independence axiom of (S)EU theory, according to which for any triplet of bets f , g and h , and for all $\alpha \in (0, 1)$, if $f \succeq g$, then $\alpha f + (1 - \alpha)h \succeq \alpha g + (1 - \alpha)h$, where $\alpha f + (1 - \alpha)h$ is the two-stage lottery that returns with probability α bet f , and with probability $(1 - \alpha)$ bet h .⁷ Specifically, a ball is drawn from one urn containing three balls, one which is certainly red and two that can be either white or green. The bets among which the decision maker is asked to express her preferences are as follows:

- f : win 100\$ if the green ball is drawn and 0 otherwise;
- g : win 100\$ if the red ball is drawn and 0 otherwise;
- f' : win 100\$ if the green or the white ball is drawn and 0 otherwise;
- g' : win 100\$ if the red or the white ball is drawn and 0 otherwise.

The experimental evidence shows that people generally prefer g to f , and f' to g' . This behavior represents a violation of expected utility since it is incompatible with any assignment of a probability distribution over random payoffs. As a matter of fact, the preference g over f shows that the event “a red ball is drawn” is considered more probable than the event “a green ball is drawn”, so that $p(R) > p(G)$.⁸ On the contrary, the preference for f' over g' shows that the event “a red or a white ball is drawn” is to be considered less probable than the event “a green or a white ball is drawn” and must reflect the subjective probability assignment $p(R \cup W) < p(G \cup W)$. This is not consistent with any probability assignment since the three events are mutually exclusive, so that (using the fact that $p(R) > p(G)$)

$$p(R \cup W) = p(R) + p(W) > p(G) + p(W) = p(G \cup W) \quad (2)$$

should hold (as the ball cannot be at the same time green and white, $G \cap W = \emptyset$).

These illustrations of Ellsberg’s paradox are important because they can be used to arrive to a formal definition of ambiguity and of aversion to ambiguity. Consider a decision maker 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 some unknown intermediary. Given that the agent is not familiar with the second coin, it is possible that she would consider “ambiguous” all the 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 seen as somewhat less desirable than bets that are “unambiguous”, such as a bet that pays \$1 if the first coin

⁶Interestingly, Ellsberg had not run carefully designed experiments in the modern sense of the term. His were mostly introspective experiments, supported by abundant casual evidence. However, starting with Becker and Brownson (1964), an experimental literature has developed confirming the introspective intuition originally described by Ellsberg (1961). A number of these experiments and of resulting stylized facts are described in Camerer and Weber (1992).

⁷In this example, h is the gamble: “win 100\$ if the green ball is drawn and 0 otherwise”; $f' = \alpha f + (1 - \alpha)h$ and $g' = \alpha g + (1 - \alpha)h$.

⁸ $P(R)$, $P(G)$, and $P(W)$ represents the subjective probability of a red, green or white ball being drawn from the urn.

lands with head up, or, equivalently, on the event $\{HH, HT\}$. If the decision maker 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 decision maker who is averse to ambiguity may prefer the equal-probability “mixture” of two ambiguous bets to either of the bets. Formally, Schmeidler (1989) called *ambiguity averse* a decision maker who prefers any mixture $\alpha f + (1 - \alpha)g$, $\alpha \in [0, 1]$ of two ambiguous bets f and g that she would otherwise (i.e., in the absence of ambiguity) consider indifferent, $f \sim g$: $\alpha f + (1 - \alpha)g \succsim f \forall f, g$ when $f \sim g$. This decision maker 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 .

An alternative perspective on ambiguity can be obtained by adopting as a benchmark a model that captures the way in which risk is commonly represented in the standard (S)EU framework. EU assumes that the probabilities of outcomes are known. If preferences follow a set of simple axioms, they can be represented by a real-valued utility function—preferred choices have higher utility numbers—and the utility of a choice is the expected utility of its possible outcomes, weighted by their probabilities, see Von Neumann and Morgenstern (1947). Under SEU, probabilities are not necessarily objectively known, so SEU applies more widely than EU.⁹ The SEU axioms show the conditions under which preferences can be represented by a numerical expected utility that uses subjective probabilities of states to weight consequence utilities (see Anscombe and Aumann, 1963): the stochastic consumption stream generated by the bet f is preferred to the one generated by the bet g , if and only if

$$(S)E_p[U(f)] \equiv \sum_{s=1}^{|\Omega|} p_s U(f_s) \geq \sum_{s=1}^{|\Omega|} p_s U(g_s) \equiv (S)E_p[U(g)], \quad (3)$$

where the p_s s are the subjective probabilities of each of the possible $|\Omega|$ states and $U(\cdot)$ is a Von Neumann-Morgenstern (VNM) utility index. The fact that axioms exist under which (3) obtains cannot be down-played: in Savage’s model, a decision maker behaves like a EU-subject who possesses a probability distribution over the states of the world, even if she is not actually aware of either probabilities of her own utility index. This implies that decisions under uncertainty can always be reduced to decisions under risk, with one caveat: beliefs are here a purely subjective construct, and can be interpreted as themselves part of the “preferences” of a decision maker.

To accommodate for some early experimental evidence of inconsistency of actual behavior with the key postulates (axioms) of rational choice under (S)EU, Quiggin (1982) had proposed early on a generalization of (3) based on the relaxation of the Von Neumann-Morgenstern/Savage independence axiom. Given the (subjective) probability distribution $\{p_j\}_{s=1}^{|\Omega|}$, Quiggin assumed the existence of a strictly increasing and continuous probability weighting function $v(\cdot)$ (also called *capacity*) which reflects the “sensitivity” of people towards their own uncertainty on the quantification of probability.¹⁰ Under Quiggin’s rank dependent utility (RDU), the risky bet f , with state-dependent payoffs $f_1 \geq$

⁹People are assumed to have subjective, or “personal” probabilities of the states, which may legitimately differ across people. In fact, de Finetti (1931) and Savage (1954) showed that probabilities can be defined in the absence of statistics, by relating them to observable choice. For example, using the so-called Laplace’s “principle of indifference”, $P(E) = 0.5$ can be derived from an observed indifference between receiving a prize under event E and receiving it under not- E (the complement of E).

¹⁰In general, a capacity v over Ω should satisfy the following 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) $v(\emptyset) = 0$ and $v(\Omega) = 1$ and (iv) $\sum_{s \in \Omega} v(s) \leq 1$. Probabilities are special cases of capacity functions that satisfy additivity: $v(A \cup B) = v(A) + v(B) - v(A \cap B)$.

$f_2 \geq \dots \geq f_{|\Omega|}$, is evaluated according to the functional:

$$RDU(f) = \sum_{s=1}^{|\Omega|} [v(p_1 + p_2 + \dots + p_s) - v(p_1 + p_2 + \dots + p_{s-1})] U(f_s), \quad (4)$$

Hence, a bet f is preferred to g ($f \succsim g$) if and only if $RDU(f) \geq RDU(g)$. It can be shown that, as long as the capacity function $v(\cdot)$ is convex, the preference functional in (4) is consistent with rationality and solves many of the experimental puzzles that had created early discomfort with the (S)EU framework (e.g., the Allais' paradox, see Allais, 1953). Under convexity, outcomes receive more weight when they are ranked worse, reflecting pessimism: convexity implies low evaluations of prospects relative to sure outcomes, enhancing risk aversion. For instance, when $n = 2$, if $U(f_2) > U(g_2)$ and $U(g_1) > U(f_1)$, provided that $v(p_2) > p_2$ and $v(p_1 + p_2) - v(p_2) < p_1$, RDU allows for the ordering $f \succsim g$, even if $E_p[U(f)] < E_p[U(g)]$.¹¹ It turns out that Quiggin's (1982) generalization of the standard (S)EU framework would play a considerable role in the development of the theory of rational decisions under ambiguity.

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 decision-maker 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 the use of Quiggin's weighting approach to assign non-additive probabilities, or capacities, to allow for the encoding of information that additive probabilities cannot represent. In the context of 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 $\{\text{neither } f_1 \text{ nor } f_2 \text{ occur}\}$, $\{f_1 \text{ or } f_2 \text{ occurs}\}$, $\{f_1 \text{ and } f_2 \text{ both occur}\}$, $\{f_1 \text{ occurs}\}$, $\{f_2 \text{ occurs}\}$, such that: i) $v(\cdot) \geq 0$ and $v(f_1) + v(f_2) \leq 1$; ii) $v(\text{neither } f_1 \text{ nor } f_2 \text{ occurs}) = 0$; iii) $v(f_1 \text{ or } f_2 \text{ or both occurs}) = 1$. Schmeidler's preference representation is based on the concept of *Choquet integral*, that in our simple example reduces to

$$\begin{aligned} CEU(f) &= \min_{\mu \in C(v)} [\mu U(f_1) + (1 - \mu) U(f_2)] \\ C(v) &= \{\mu \in [0, 1], \quad \mu \geq v(f_1), \quad 1 - \mu \geq v(f_2)\}, \end{aligned} \quad (5)$$

where $C(v)$ is the *core* of v .¹² 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, in the sense that—to protect herself against the possibility of mistakes—the agent considers the most unfavorable probability distribution. To express the degree of ambiguity that characterizes the capacity assignment

¹¹Such a ranking is in fact compatible with the idea that p_1 is so low that the two (higher) outcomes g_1 and f_1 can be approximately considered impossible. Hence the decision maker under-weights the probability p_1 , assigning lower weight $v(p_1 + p_2) - v(p_2) < p_1$, and makes her choice considering only the ranking between g_2 and f_2 (and consequently over-weighting the probability p_2 , as reflected in the capacity assignment $v(p_2) > p_2$).

¹²The core of a capacity v consists of all finitely additive probability measures μ that event-wise dominate v .

$(v(f_1), v(f_2))$, 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 given to the outcomes f_1 and f_2 .

In general, CEU preferences are of the *multiple-prior type*, in the sense that in practice, under CEU a rational decision-maker evaluates expected utility using a multi-valued set of priors, for instance as defined by $\mu \in C(v)$. Schmeidler (1989) proved that—when the capacity v defined on a state space Ω is convex (i.e., given two events A and B , $v(A \cup B) + v(A \cap B) \geq v(A) + v(B)$)—the Choquet integral of a utility index U with respect to v is equal to

$$\min_{\mu \in C(v)} \int U(f) d\mu, \quad (6)$$

while the ambiguity index is $\mathcal{A}(v) \equiv \max_{\mu \in C(v)} \mu(f) - \min_{\mu \in C(v)} \mu(f)$ (see Fishburn, 1993, and Mukerji, 1997). Hence ambiguity aversion coincides with and it is measurable as the convexity of the capacity function. Because the capacity $\mu \in C(v)$ needs not be additive, the objective in (6) cannot be computed using a standard Lebesgue integral.

Another interesting connection is between CEU preferences and RDU. Because a Choquet integral is evaluated as

$$\begin{aligned} \int U(f) d\mu &= U(f_1) + [U(f_2) - U(f_1)]\mu(s_2, s_3, \dots, s_{|\Omega|}) + \dots + [U(f_i) - U(f_{i-1})]\mu(s_i, s_{i+1}, \dots, s_{|\Omega|}) + \\ &\quad + \dots + [U(f_{|\Omega|}) - U(f_{|\Omega|-1})]\mu(s_{|\Omega|}) \end{aligned}$$

($|\Omega|$ is the total number of states and $f_{|\Omega|} > \dots > f_i > f_{i-1} > \dots > f_1$, the outcomes are ranked from worst to best), a decision maker evaluates f by considering first the lowest outcome and then adding the successive increments, weighted by her estimation of their occurrence, as captured by a capacity μ . Due to the non-additivity of the capacity, the weight of an outcome will depend on its place in the ranking. Clearly, CEU is a special case of RDU. Note that under an additive capacity, we get back to subjective expected utility, because, for $n < |\Omega|$, $U(f_n)[\mu(s_n, s_{n+1}, \dots, s_{|\Omega|}) - \mu(s_{n+1}, s_{n+2}, \dots, s_{|\Omega|})] = U(f_n) \Pr(s_n)$, so that $\int U(f) d\mu = \int U(f) d\Pr(f)$, a standard (S)EU Lebesgue integral.

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

$$f \succsim g \text{ if and only if } \min_{p \in \wp} E_p[U(f)] \geq \min_{p \in \wp} E_p[U(g)], \quad (7)$$

where $E_p[U(\cdot)]$ is a standard (S)EU-operator when the probability measure is $p \in \wp$. \wp is a convex set whose size can be interpreted as representing the level of perceived ambiguity.¹³ The intuition is that—because of ambiguity aversion—agents are considering as valid and relevant the prior which is most unfavorable. Given the multi-valued nature of the set \wp , preferences represented by (7) have became commonly known as *multiple prior preferences* (MPP). MPP can be shown to be equivalent to CEU when the capacity is convex, since weighting outcomes by sub-additive capacities expresses the same kind of pessimism as taking the minimum (S)EU over \wp (see e.g., Camerer and Weber, 1992). However, strictly speaking, neither approach is a special case of the other.¹⁴

¹³Already in 1950, Wald had developed the first max-min decision criterion based on the compelling intuition that any decision problem under unknown probabilities can be modeled as a two-player, zero-sum game. More specifically, the two players are the decision maker and a malevolent Nature whose choice set is a particular subspace \wp of the space of all probability distributions over outcomes. In 1951, Hurwicz had proposed a “generalized Bayes-minimax principle” that relates to Wald’s min-max decision rule. However, such min-max decision rules have been axiomatized in their rational foundations only by Schmeidler (1982) and Gilboa and Schmeidler (1989).

¹⁴As noted by Wakker (2006) and others, the (S)EU-like formula (7) is at first easier to understand than CEU. The flexibility of not having to specify precisely what “the” probability measure is, while usually perceived as an advantage, may however turn into a disadvantage when applying the model. One has then to specify exactly what “the” set of probability distributions \wp is, which may be more complex than specifying only one probability measure or a capacity. 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 needed.

It is worth noticing that Bewley (1986) had anticipated a few of the intuitions later developed by Gilboa and Schmeidler. 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 the idea 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.¹⁵ When lotteries are non 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., that has higher expected utility for all possible probability distributions), which is consistent with the experimental evidence of a status quo bias (e.g., Samuelson and Zeckhauser, 1988; Knetsch, 1989).

The 1990s have witnessed a number of extensions and critiques to the seminal definition of ambiguity provided by Schmeidler (1989) and Gilboa and 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 nor necessary nor sufficient to generate behaviors that are compatible with the intuition behind ambiguity aversion. Considering the three-ball version of the Ellsberg's paradox, Epstein noticed that even though CEU preferences characterized by the capacity

$$v(R) = 8/24, v(W) = v(G) = 7/24, v(W \text{ or } G) = 13/24, v(R \text{ or } G) = v(R \text{ or } W) = 1/2$$

lead to a Ellsberg's type ordering among bets, v is not convex. Viceversa, if the capacity is ξ , such that

$$\xi(R) = 1/12, \xi(W) = \xi(G) = 1/6, \xi(W \text{ or } G) = 1/3, \xi(R \text{ or } G) = \xi(R \text{ or } W) = 1/2,$$

then the ranking among bets under CEU is exactly opposite to the one showed by Ellsberg, and nevertheless ξ is convex. Wakker (2008) shows that in fact convexity of v is required to resolve the Allais', not Ellsberg's paradox.¹⁶

Motivated by these considerations, Epstein and Ghirardato and Marinacci proposed a comparative notion of ambiguity aversion that is closely related to the one of risk aversion.¹⁷ The key idea is to use a set of events which are exogenously known to be unambiguous as "benchmarks". Acts (bets) that only depend on these events are called unambiguous, since the probability distributions over their possible payoffs are objectively known.¹⁸ Assuming that

¹⁵Aumann (1962) showed that incomplete preferences can be represented by expected utilities over sets of probabilities. Then $f \sim g$ if and only if $E(U(f)|\pi) > E(U(g)|\pi)$ for all probability distributions π , otherwise f and g are "incomparable", as a result of preference incompleteness. Schmeidler (1989) has later argued that the most troubling axiom underlying (S)EU is not the infamous independence axiom but the more common assumption of completeness; the critical role of the independence axiom is to extend preferences from choices that seem obvious to those that do not, i.e., it just helps to deliver completeness.

¹⁶Epstein (1999) and Wakker (2001, 2006) have argued that the Ellsberg paradox does not speak of convexity of the capacity in an absolute sense, but only in a relative (within-person) sense, suggesting more convexity for unknown probabilities than for known probabilities. To obtain Ellsberg's paradox, it is then possible that the capacity is concave, and not convex, for both known and unknown probabilities, but is less concave (and thus more convex) for the unknown probabilities.

¹⁷In standard (S)EU theory, given a gamble f , a preference relation \succeq_2 is more risk averse than \succeq_1 if for any certain gamble \bar{x} (that is, a gamble that pays \bar{x} with probability 1), $\bar{x} \succeq_1 f \Rightarrow \bar{x} \succeq_2 f$. In particular, a preference relation \succeq is risk averse if it is more risk averse than any risk neutral preference ordering \succeq_{RN} . As usual, a preference relation is risk neutral if the corresponding utility index is linear.

¹⁸Formally, given the set Ω of the states of the world, an algebra of subsets of Ω called events, and a set \mathcal{F} of outcomes, an act f is a finite-valued measurable function $f : \Omega \rightarrow \mathcal{F}$. Given any $\bar{x} \in \mathcal{F}$, an act f is said to be constant , if $f(s) = \bar{x}$ for all $s \in \mathcal{F}$. Finally, given a set of unambiguous events \mathcal{U} , the act f is said to be unambiguous if for any pair of events $E, F \notin \mathcal{U}$, $f(s) = f(w) \forall s \in E, w \in F$. That is, on any ambiguous event, f yields the same outcome.

a collection \mathcal{U} of unambiguous bets is exogenously given (for Ghirardato and Marinacci such \mathcal{U} corresponds to the set of constant acts), the preference \succeq_2 is more ambiguity-averse than \succeq_1 if, for any $x \in \mathcal{U}$,

$$x \succeq_1 f \Rightarrow x \succeq_2 f, \quad (8)$$

meaning that if any unambiguous act is ever preferred to an ambiguous one according to preference \succeq_1 , the same ordering will be reflected by \succeq_2 which must then be even more averse to ambiguity. In particular, in Epstein's framework a preference relation \succeq is ambiguity-averse if it is more ambiguity-averse than any probabilistically sophisticated order \succeq_{PS} , that can be heuristically interpreted as an uncertainty neutral order,¹⁹ while in Ghirardato and Marinacci's setting the reference for ambiguity neutrality is represented by some benchmark (S)EU preference. Ghirardato and Marinacci also proved 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.

2.3. Robust Control and Variational Preferences

Well 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 φ 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 in it, they seek robustness against misspecifications. Usually, φ 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 poor quality of information used to select the approximating model. Therefore ϱ can be thought of as an ambiguity aversion index since it measures the fear of model misspecification: the lower is ϱ , the higher is the degree of ambiguity aversion. Hansen and Sargent and a number of co-authors have developed a class of preferences that may be represented as

$$f \succsim g \text{ if and only if } \min_{q \in \Delta(\Omega)} E_q[U(f) + \varrho R(q||p)] \geq \min_{q \in \Delta(\Omega)} E_q[U(g) + \varrho R(q||p)], \quad (9)$$

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 (9), 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—as measured by the Kullback–Leibler distance—to the approximating model, p .²⁰ The

¹⁹A preference relation \succeq_{PS} is probabilistically sophisticated if it ranks gambles only on the basis of a possible (compound) probability distribution assigned to the payoffs of the gamble.

²⁰In technical terms, the use of entropy imposes that alternative probability distributions have to be absolutely continuous with respect to the reference one and, in a diffusion setting, absolute continuity restricts the class of alternative models to those diffusions that only differ in terms of the drift function with respect to the approximating process. Increasing penalization of larger departures from the approximating model is consistent with the experimental results in Yates and Zukowski (1976), who find that increasing the range of possible experimental probabilities increases ambiguity aversion.

formulation in (9) is often referred to as the “penalty problem”. An alternative methodology used in the robust control literature consists instead 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 second approach, individual’s preferences can be represented by

$$f \succsim g \text{ if and only if } \min_{q \in C} E_q[U(f)] \geq \min_{q \in C} E_q[U(g)] \quad C = \{q \in \Delta(\Omega) : R(q||p) \leq \eta\} \quad (10)$$

which is a particular case of MPP. Hansen and Sargent (2001) showed that (10) and (9) are connected via the Lagrange multiplier theorem, since the preference orderings implied by the two approaches differ, but the two optimizations lead to identical decisions.

The Hansen-Sargent model (9) is highly suitable to applications as it can be embedded into intertemporal continuous-state-space frameworks, in particular let \mathbf{B}_t be a d -dimensional BM on a probability space (Ω, \mathcal{F}, p) , and x_t a state-dependent variable which is assumed to evolve according to $dx_t = \mu(c_t, x_t)dt + \boldsymbol{\sigma}(c_t, x_t)d\mathbf{B}_t$. The actions of a decision maker neutral to ambiguity are represented by a stochastic policy function $c_t(\cdot)$, which is the optimal solution to

$$\begin{aligned} & \sup_{c_t} E_p \left[\int_0^\infty \exp(-\delta t) U(c_t, x_t) dt \right] \\ & \text{s.t. } dx_t = \mu(c_t, x_t)dt + \boldsymbol{\sigma}(c_t, x_t)d\mathbf{B}_t, \end{aligned} \quad (11)$$

where δ is a standard subjective discount factor. 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 , since its Radon-Nikodym derivative, $z_t^\theta \equiv dq/dp$, is

$$z_t^\theta = \exp \left(-\frac{1}{2} \int_0^t \|\boldsymbol{\theta}_s\|^2 ds + \int_0^t \boldsymbol{\theta}_s d\hat{\mathbf{B}}_s \right), \quad (12)$$

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

$$R(q) \equiv \int_0^\infty \exp(-\delta s) E_q \left(-\frac{1}{2} \int_0^t \|\boldsymbol{\theta}_s\|^2 \right) ds. \quad (13)$$

Hence, a possible, alternative “twisted” model can be represented as:

$$d\mathbf{x}_t = \mu(\mathbf{x}_t, \mathbf{c}_t)dt + \boldsymbol{\sigma}(\mathbf{x}_t, \mathbf{c}_t)[\boldsymbol{\theta}(\mathbf{x}_t)dt + d\mathbf{B}_t], \quad (14)$$

and the optimal policy for an ambiguity-averse, Hansen-Sargent type agent can be derived by solving the program

$$\begin{aligned} & \sup_{c_t} \inf_q E_q \left\{ \int_0^\infty \exp(-\delta t) \left[U(c_t, x_t) - \frac{1}{2} \varrho E_q \left(\int_0^t \|\boldsymbol{\theta}_s\|^2 \right) \right] dt \right\} \\ & \text{s.t. } dx_t = \mu(c_t, x_t)dt + \boldsymbol{\sigma}(c_t, x_t)d\mathbf{B}_t, \end{aligned} \quad (15)$$

where the problem in (15) has the same structure as (11) with two key amendments: expectations are computed under the measure q and deviations from p are penalized by ϱ times the Kullback–Leibler entropy measure, $R(q)$.

Hansen-Sargent robustness preferences (9) have been axiomatized only subsequently by Maccheroni, Marinacci, and Rustichini (2006) who have formally proven that robustness preferences are in fact a specific sub-class of what they have labeled *variational preferences* (VP). VP are characterized by a grounded (with zero infimum value), convex, and

lower semi-continuous function ϕ defined on the standard simplex $\Delta(\Omega)$. ϕ is called ambiguity index, since different functional forms of ϕ determine different attitudes towards ambiguity. According to the VP criterion:

$$f \succsim g \text{ if and only if } \min_{p \in \Delta(\Omega)} E_p[U(f)] + \phi(p) \geq \min_{p \in \Delta(\Omega)} E_p[U(g)] + \phi(p). \quad (16)$$

In words, agents consider all possible probabilities in $\Delta(\Omega)$, giving weight $\phi(p)$ to each of them, and the minimization over $p \in \Delta(\Omega)$ reflects ambiguity aversion. The cautious attitude featured by VP agents can also be interpreted as the outcome of a two-player zero sum game between the decision maker and a malevolent Nature (with the goal of minimizing the agent's utility), that, for any selected act f , chooses a probability p , incurring a cost $\phi(p)$ which depends on the specific probability measure she has selected. Furthermore, two pairs $(U_0; \phi_0)$ and $(U; \phi)$ represent the same preferences if and only if there exist $a > 0$ and b such that $U = aU_0 + b$ and $\phi = a\phi_0$, which nicely extends the usual linearity property of VNM ordinal utility indices. VP are characterized by uncertainty aversion as defined by Schmeidler, and ambiguity aversion comparisons are based on the following criterion: \succsim_1^{VP} is more ambiguity averse than \succsim_2^{VP} if the associated ambiguity indexes ϕ_1 and ϕ_2 are such that $\phi_1(p) \leq \phi_2(p), \forall p \in \Delta(\Omega)$.

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. For instance, setting $\phi(p) = 1$ if $p \in \wp$ and $\phi(p) = 0$ if $p \notin \wp$, one obtains MPP. Obviously, noting that $R(p||q)$ is a convex function and setting $\phi(p) = R(p||q)$, one derives Hansen-Sargent (HS) robust preferences. 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.^{21,22}

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

With the exception of Hansen and Sargent continuous time modelling of robust preferences, a major concern with most of the advances reviewed in Sections 2.2 and 2.3 is that these are static settings. However, preference functionals that take ambiguity into account have also been proposed within truly dynamic settings, which are important in financial applications. However, because ambiguity leads to (and stems from) violations of the Sure-Thing Principle, defining updating and ensuring dynamic consistency for CEU/MPP or similar models is a non-trivial task and it has been extensively debated. For instance, it is easy to put together simple two-period extensions of Ellsberg's paradox in which an individual with Ellsberg-type preferences ends up displaying a dynamically inconsistent behavior as she will fail 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, from the set $\Omega = \{1, 2\}$. This is obviously a standard binomial tree with three dates and a total of four final possible nodes. Date

²¹Preferences 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:

$$E_p[f] - \frac{1}{2\rho} Var(f) = \min_{p \in \Delta(\Omega)} E_p[f] + \varrho G(p||q)$$

where $G(p||q)$ is the relative Gini concentration index. Section 4.1 describes applications of MMVP.

²²Because most models of ambiguity are based on relaxations of the independence axiom, Cerreia-Vioglio, Maccheroni, Marinacci and Montrucchio (2009) have developed a general class of *uncertainty averse* preferences that completely do away with the independence axiom.

1 probabilities are $q(s_1 = 1) = q(s_1 = 2) = 1/2$; 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 possible 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 . With these probabilities, consider the value of an asset that pays 1 unit of the consumption good if $s_2 = 1$ and 0 otherwise, i.e., a simple state-contingent Arrow-Debreu security. For convenience, let $U(c) = c$, a simple 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 standard MPP/max-min 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 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 in a MPP criterion, restoring dynamic consistency. In general, preferences depend on probabilities over complete paths, which in our example you might associate 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. In short, expanding the date 0 set of probabilities in this way reconciles date 0 and recursive valuations and resolves the dynamic inconsistency problem. One puzzling consequence of rectangularity is that it can induce ambiguity in events that have none to begin with. For instance, in our example it is obvious that $q(s_2 = 2) = q(s_2 = 2|s_1 = 2) + q(s_2 = 2|s_1 = 1) = (1 + \phi)/2 + (1 - \phi)/2 = 1/2$, for which there is no ambiguity. Yet, under Epstein and Schneider's rectangularity, $q(s_2 = 2) = [(1 + \phi_1)/4] + [(1 - \phi_2)/4] = 1/2 + (\phi_1 - \phi_2)/4$ and ambiguity seems to oddly persist. The apparent puzzle is resolved if we realize that the date-zero rectangular set does not represent date 0 ambiguity, but the date 0 probabilities needed to anticipate preferences over future ambiguity.

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 (Ω, \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 (Ω, \mathcal{F}, p) . A density generator is a d -dimensional processes $\boldsymbol{\theta} = (\boldsymbol{\theta}_t)$, such that $\boldsymbol{\theta} \in \Theta$ and $E[\exp(-\frac{1}{2} \int_0^t \|\boldsymbol{\theta}_s\|^2 ds)] < \infty$.²³ Under this assumption, $\boldsymbol{\theta}$ generates a set of probability measures, \wp^θ , whose elements are equivalent to p . In particular, \wp^θ is defined as:

$$\wp^\theta = \left\{ q : \frac{dq}{dp} = z_T^\theta, \quad z_t^\theta = \exp \left(-\frac{1}{2} \int_0^t \|\boldsymbol{\theta}_s\|^2 ds + \int_0^t \boldsymbol{\theta}_s d\mathbf{B}_s \right) \right\}. \quad (17)$$

Chen and Epstein generalized the notion of stochastic differential utility introduced by Duffie and Epstein (1992) allowing for ambiguity. Specifically, for each $q \in \wp^\theta$ they defined a utility process V_t^q as

$$V_t^q = E_q \left[\int_t^T \psi(c_s, V_s^q) ds | \mathcal{F}_t \right], \quad (18)$$

where the function ψ is an aggregator as in Duffie and Epstein (1992). With this definition for V_s^q ($s \geq t$), intertemporal MPP is represented as the lower envelope of V_t^q over $q \in \wp^\theta$:

$$V_t \equiv \min_{q \in \wp^\theta} V_t^q = \min_{q \in \wp^\theta} E_q \left[\int_t^T \psi(c_s, V_s^q) ds | \mathcal{F}_t \right] \quad 0 \leq t \leq T. \quad (19)$$

Under some regularity assumptions on the domain Θ , the ambiguity-aversion preferences in (19) are dynamically consistent, allowing for a recursive representation of the utility process. Whilst the important paper by Chen and Epstein (2002) extended in a dynamically consistent fashion the MPP introduced by Gilboa and Schmeidler, it was Maccheroni, Marinacci, and Rustichini (2007) who introduced dynamic VP. Recalling the previous interpretation of a two-players zero sum game against Nature, Maccheroni et al. show that a VP-decision maker is dynamically consistent if and only if she thinks that Nature is also dynamically consistent in its malevolence. As a corollary, it can be derived that dynamic MPP and Hansen and Sargent's robustness preferences are dynamically consistent.

An alternative approach to the definition of ambiguity is the multi-stage, multi-lottery approach. Notice that by construction any two-stage lottery is simply defined a special bet with outcomes (at the end of the first stage) that can be thought of as further lotteries. One early example of the connection between two-stage lotteries and ambiguity was offered by Segal (1990). Segal 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 decision maker considers the urn with unknown composition as sampled from a set of urns. Segal showed that agents who are (S)EU-maximizers in the second stage and display *anticipated utility preferences* (AUP) in the first stage, exhibit Ellsberg's behavior. AUP (see, e.g., Quiggin, 1982) is an extension of the expected utility framework based on the relaxation of the independence axiom, that allows for the representation of choices that are incompatible with standard theory.²⁴ In Segal's framework, this is precisely what is

²³Here, and only for simplicity, we treat Θ as exogenous. For more details on its construction, see Chen and Epstein (2002).

²⁴AUP permits the analysis of phenomena associated with the distortion of subjective probability, since it is based on a weighted sum of state-contingent utilities formed using decision weights. Each weight is derived from the entire distribution over the reference state space and not from the individual probabilities. Hence, when extreme events are overweighted, some intermediate outcome must be underweighted (even if it has same objective probability to occur).

needed to avoid the reduction of compound, multi-stage lotteries into simpler, single-stage lotteries (this would happen if agents were (S)EU-maximizer also in the first stage, because of the reduction of compound lotteries axiom).²⁵ In particular, Ellsberg's behavior would be primarily related to a decision maker's attitude towards second-order risk (informally defined as "ambiguity" by Segal, 1987), 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 independent issues. To motivate their approach, they used the three-ball version of the Ellsberg's paradox presented in Section 2.1. Specifically, they note that two types of uncertainty can be described: first there is the issue of which ball is chosen (red or non-red), second there is the issue of the color of the non-red balls. In their view, the violation of the expected utility paradigm arises because agents are comparing bets that are based on different issues: g and g' are based only on the issue of which ball is chosen (the knowledge of the composition of the urn is not relevant), while f and f' are based on both issues. Ergin and Gul associated to each possible bet a new compound two-stage lottery, where each stage is identified with the resolution of one issue. By ranking the compound lotteries, agents are implicitly ranking also the original ones. Specifically, ambiguity-aversion translates into the preference for bets based only on the first issue (red or non-red). The link with the traditional approach towards ambiguity is established by the fact that ambiguity aversion over original bets, implies risk aversion over compound lotteries.

Klibanoff, Marinacci and Mukerji (2005) have elaborated these multi-stage ideas by proposing that the ambiguity of a risky act f is characterized by a set $\wp = \{P_1, \dots, P_n\}$ of subjectively plausible cumulative probability distributions for f . Letting f_j denote the random variable distributed as P_j , $j = 1, \dots, n$, based on her subjective information, the decision maker associates a distribution (q_1, \dots, q_n) over \wp , where q_j is the subjective probability of P_j being the true distribution of f . The resulting preferences (call them KMM) have the representation²⁶

$$f \succsim g \text{ if and only if } \sum_{j=1}^n q_j \zeta \left(\int U(f) dP_j \right) \geq \sum_{j=1}^n q_j \zeta \left(\int U(g) dP_j \right), \quad (20)$$

where $\zeta(\cdot)$ is an increasing real-valued function, whose shape describes the investor's attitude towards ambiguity. First the decision maker evaluates the expected utility of f with respect to all the priors in \wp : each prior P_j is indexed by j so in the end, we get a set of expected utilities, each being indexed by j . Then, instead of taking the minimum of these expected utilities, as MPP would, take an expectation of distorted expected utilities. The role of ζ is crucial here: if ζ were linear, the criterion would simply reduce to (S)EU maximization with respect to the combination of the probabilities qs and possible distributions P_js . When ζ is not linear, one cannot combine qs and P_js to construct a reduced probability distribution. In this event, the decision maker takes the expected " ζ -utility" (with respect to q) of the expected " U -utility" (with respect to the P_s). A concave ζ will reflect ambiguity aversion, in the sense that it places a larger weight on bad expected " U -utility" realizations.²⁷

²⁵The reduction of compound lotteries axiom states that any two-stage lottery is equivalent to a simple lottery yielding the same prizes with the corresponding (and appropriately) computed total probabilities. This principle is key in the reductionist approach to ambiguity, as described in Camerer and Weber (1992) and first discussed by Marshack (1975). There is abundant evidence that the reduction axiom is systematically violated in experiments (e.g., Camerer and Ho, 1991).

²⁶Kreps and Porteus' (1978) risk-sensitive preferences—that have received considerable attention in the asset pricing literature—have a functional representation similar to KMM preferences. However Kreps and Porteus' preferences are not directly concerned with ambiguity, but are instead motivated by issues of time resolution of uncertainty in the classical (S)EU framework. The model extends of the standard (S)EU-framework since the agent is allowed to display a particular preference for an earlier (or later) resolution of uncertainty.

²⁷Ergin and Gul (2004), Nau (2006) and Neilson (2010) have interpreted (20) to stress that second stage events are from a different

One important implication of the two-stage approach is that the decision maker 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 decision maker'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 (20): up to ordinal equivalence, MPP is obtained in the limit as the degree of concavity of ζ increases without bound. Further, a very simple 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 $-\zeta''_2/\zeta'_2$, or, equivalently, if ζ_1 is more concave than ζ_2 .²⁸ Contrary to MPP and RDU preferences, in the KMM model beliefs and (ambiguity) attitude parameters are explicitly separated. For instance, it is the function $\zeta(\cdot)$ that captures aversion to (or a preference for) ambiguity. Moreover, in the smooth ambiguity representation, beliefs may be seen to have precisely the same connection to a decision maker's subjective information as in subjective expected utility representation and standard Bayesian theory routinely applied in economics. The possibility to obtain a clear and time-consistent separation between beliefs and aversion to ambiguity in KMM framework, clearly represents a key advantage over MPP, which has attached considerable attention in applied work in financial economics (see Sections 3.6 and 4.4). However, a number of recent papers have drawn the attention on the fact that KMM may often imply counterintuitive behaviors when the agent can bet directly on what the true model is.²⁹

Chateauneuf, Eichberger, and Grant (2002) have recently returned to work on the seminal intuition by Quiggin and Schmeidler that ambiguity is a phenomenon that must be captured as and related to the properties of capacities used in the representation of preferences. They suggest the use of *non-extreme outcome* (NEO) *additive capacities* to model optimistic and pessimistic attitudes towards uncertainty. After providing a weighting scheme for objective probabilities to derive neo-additive capacities, Chateauneuf et al. (2002) axiomatized a model based on Choquet expected utility with neo-additive capacities. For simplicity, consider an act f , whose payoff can be either f_1 or f_2 , with $f_1 > f_2$. The decision maker beliefs are such that she is able to derive a probability distribution for the two outcomes, say $p(f_1)$ and $p(f_2)$. The NEO additive weighting scheme is

$$\nu^{neo}(f_i) = \lambda + (1 - \lambda - \xi)p(f_i) \quad \lambda, \gamma \geq 0, \quad \lambda + \xi \leq 1, \quad (21)$$

where λ and ξ are parameters that capture pessimism and optimism, respectively. The above assignment can be re-interpreted as a convex combination of a probability and two capacities on two extreme outcomes: one is complete ignorance in everything but certain events (μ^0), and the other is complete confidence in everything but the null (impossible) event. Specifically, assuming the existence of a set N (C) of events that are considered virtually impossible (certain to occur), and letting for any event E , $\mu^0(E) = 1$ if $E \in C$ and 0 otherwise, and $\mu^1(E) = 0$ if $E \in N$ and 1 otherwise, then $\nu^{neo}(E) = \xi\mu^0(E) + \lambda\mu^1(E) + (1 - \lambda - \xi)p(E)$. The Choquet expected value of f under ν^{neo} is:

$$CEU^{\nu^{neo}}(f) = [\lambda + (1 - \lambda - \xi)p(f_1)]f_1 + [1 - \lambda - (1 - \lambda - \xi)p(f_1)]f_2.$$

From this expression it is easy to see that $\lambda = 0$ corresponds to pure pessimism, when the weight assigned to the source of ambiguity is greater than the first-stage events. A concave $\zeta(\cdot)$, for instance, suggests stronger preference for certainty, and more ambiguity aversion, for the first-stage uncertainty than for the second. Seo (2009) has recently circumvented this issue and proposed an axiomatization in which the domain on which decisions are defined is simply the intuitive state space (i.e., in Ellsberg example, the color of the ball drawn).

²⁸Notice the similarity with the standard, SEU-style analysis of (local) risk aversion, as in Arrow and Pratt (1952).

²⁹We refer the Reader to Epstein (2010) for an interesting example.

better outcome f_1 is $(1 - \xi)p(f_1) \leq p(f_1)$, while $\xi = 0$ corresponds to pure optimism since the weight assigned to the worse outcome f_2 is $(1 - \lambda)(1 - p(f_1)) \leq p(f_2)$. This representation can also be written as:

$$CEU^{\nu^{neo}}(f) = \xi \inf_{f_i} f_i + \lambda \sup_{f_i} f_i + (1 - \lambda - \xi)E_p[f] \quad (22)$$

Therefore, the payoff of any strategy is given by a weighted average of the expected payoff under the probability p , plus maximal and minimal payoffs over the reference state-space Ω . The additive part of $CEU(\nu^{neo})$, $E_p[f]$, can be interpreted as the agent's belief and $(1 - \lambda - \xi)$ as (her) degree of confidence in that belief which nicely connects with Keynes' (1921) distinction between the likelihood implied by the evidence and the weight to be assigned to it.

2.5. Learning and Ambiguity

In the early stages of the literature, it was common to stumble in the (as we shall see, sometimes erroneous) speculation that ambiguity might be an asymptotically irrelevant phenomenon if agents were allowed to learn about their stochastic environment. For instance, it was observed that in the classical Ellsberg paradox, it is obvious that any decision maker could benefit from repeated independent draws from the ambiguous urn. However, reality is often complex enough that even related events over time cannot be interpreted as independent, repeated draws from a stationary urn that never changes. More interestingly, there are also technical issues because, under non-additive probabilities, the notion of "independence" is not always well-defined. More specifically, in the context of the Ellsberg's experiment, successive independent random draws could correspond to two very different experiments (see Dow and Werlang, 1994):

1. Balls are drawn (with replacement) from an urn with an uncertain fraction of red and black balls;
2. One ball is drawn from a sequence of urns. The fraction of red and black balls in the urns is unknown.

While in the case of experiment 1, it is obvious that the agent's subjective beliefs would become less ambiguous over time, and eventually converge to a limit, under experiment 2 the notion of independence might be more properly described as "lack of known dependence". Under this particular situation, uncertainty will fail to disappear because standard recursive drawing schemes might be applied but without inducing any change in the agent's belief. For example, in the CEU-model, for a given capacity v over the state space, the Bayes' rule for updating v is given by

$$v(A|B) = \min_{q \in C(v)} \frac{q(A \cap B)}{q(B)} \quad q(B) > 0$$

but the core of $v(\cdot|B)$, $C(v(\cdot|B))$, is not necessarily "smaller" than $C(v)$, hence ambiguity is not progressively reduced by subsequent draws. Therefore, if the structure of the environment is sufficiently complex, ambiguity does not trivially disappear with the passage of time. In fact, models of ambiguity are known to cause a number of curious problems (see e.g., Gilboa, 1989, and Etner et al., 2009) to standard updating algorithms, which interfere with the possibility "to learn ambiguity away".³⁰ In particular, there are instances in which the standard Bayes rule fails, which makes modelling learning in the presence of ambiguity much harder than under (S)EU. For instance, Gilboa and Schmeidler

³⁰See Al Najjar and Weinstein (2009). The issues are not limited to updating algorithms. For instance, in a CEU framework it is easy to check that if probabilities are nonadditive, then maximizing $U(\cdot)$ is not necessarily the same as minimizing $-U(\cdot)$, as preference orders may differ when two different kinds of Choquet expectations are taken.

(1991) have studied how to update non-additive probabilities and have defined a family of pseudo-Bayesian updating rules. When probabilities are additive, each rule coincides with Bayes's rule, but in general this correspondence fails.

Epstein and Schneider (2007) have tackled these issues and formally introduced the effects of learning in a Epstein-Chen ambiguity set up, when the distinction between risk and ambiguity matters and the data are generated by the same memoryless mechanism in every period. Specifically, they contend that standard Bayesian learning models are extreme in the sense that they describe agents who are implausibly ambitious about what they can learn in complicated environments. On the opposite, Epstein and Schneider developed 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 decision makers 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.

In a setting characterized by a finite state space identical at all times, $\Omega_t = \Omega$, one state $s_t \in \Omega$ is observed in every period. At time t , the agent's information consists of the history $s^t = (s_1, \dots, s_t)$. A standard Bayesian model of learning about a memoryless mechanism is summarized by a triple (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 trying to learn, while μ_0 represents initial beliefs about them. For a given parameter value $\xi \in \Xi$, the data are an independent and identically distributed sequence of signals $\{s_t\}$, where the distribution of any signal s_t is described by the probability density $\ell(\cdot|\xi)$. Beliefs on (payoff-relevant) states are represented by a process $\{p_t\}$ of one-step-ahead conditionals, so that the dynamics of learning can be summarized by

$$p_t(\cdot|s^t) = \int_{\Xi} \ell(\cdot|\xi) d\mu_t(\xi|s^t), \quad (23)$$

where s^t represents history up to time t and μ_t is the posterior belief about ξ derived via Bayes' Rule. When agents are ambiguity-averse, 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). The size of M_0 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 complexity of the environment that prevents the agent to learn properly, since every parameter $\xi \in \Xi$ is associated with a set of probability measures, $L(\cdot|\xi) = \{\ell(\cdot|\xi) : \ell \in L\}$, where each $\ell : \Xi \rightarrow \Delta(\Omega)$ is a likelihood function. The dynamics of learning is summarized by a process of the one-step-ahead conditional beliefs

$$\wp_t(s^t) = \left\{ p_t(\cdot) = \int_{\Xi} \ell(\cdot|\xi) d\mu_t : \mu_t \in M_t^a(s^t), \ell \in L \right\}, \quad (24)$$

where a is a parameter ($0 < a \leq 1$) 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 fully derive the updating rule that determines the construction of 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 decision-maker at time t looking back at the sample s^t : in general, she

views both her prior information and the sequence of signals as ambiguous. As a result, she will typically entertain a number of different theories about how the sample has been generated: specifically, a theory is a pair (μ_0, ℓ^t) , where μ_0 is a prior belief on Ξ and $\ell^t = (\ell_1, \dots, \ell_t) \in L^t$ is a sequence of likelihoods. Attitude towards past and future signals are allowed to be different: on the one hand, L is the set of likelihoods possible in the future; on the other hand, 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 suggested the following two steps procedure. First, how well a theory (μ_0, ℓ^t) explains the data is captured by the (unconditional) data density evaluated at s^t as $\int \prod_{j=1}^t \ell_j(s_j|\xi) d\mu_0(\xi)$. Here conditional independence implies that the conditional distribution given ξ is simply the product of the likelihood values ℓ_j , $j = 1, \dots, t$. Prior information is taken into account by integrating out the parameter 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:

$$d\mu_t(\cdot; s^t, \mu_0, \ell^t) = \frac{\ell_t(s_t|\cdot)}{\int \ell_t(s_t|\xi^t) d\mu_{t-1}(\xi^t; s^{t-1}, \mu_0, \ell^{t-1})} d\mu_{t-1}(\cdot; s^{t-1}, \mu_0, \ell^{t-1}). \quad (25)$$

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. Posteriors are formed only for theories that pass the test and are given by:

$$\begin{aligned} M_t^a(s^t) &= \{\mu_t(s^t; \mu_0, \ell^t) : \mu_0 \in M, \ell^t \in L^t\}, \\ \int \prod_{j=1}^t \ell_j(s_j|\xi) d\mu_0(\xi) &\geq a \max_{\tilde{\mu}_0 \in M_0, \tilde{\ell}^t \in L^t} \int \prod_{j=1}^t \tilde{\ell}_j(s_j|\xi) d\tilde{\mu}_0 \end{aligned} \quad (26)$$

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 of this result to key questions in the theory of finance.³¹

3. Optimal Consumption and Portfolio Decisions under Ambiguity

3.1. Trading Breakdowns and Limited Participation under Ambiguity

The seminal paper on the effects of ambiguity on optimal portfolio choice is Dow and Werlang (1992), who studied a stylized two-period portfolio problem by considering a market with one ambiguous asset and one safe bond. Their result is path-breaking: while under (S)EU, portfolio trades will occur *generically*, which means that empirically trading will be widespread, ambiguity may induce the existence of wide and persistent intervals for prices, for which the net demand for the asset is zero. Of course, such intuition is very important to understand the existence of market freezes, situations in which trading endogenously stops.

It is instructive to take a closer look at Dow and Werlang's (1992) result. In a standard (S)EU framework, 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; therefore the threshold between optimal purchases and sales is represented by a single price (as

³¹Let us mention that alternative ambiguity preferences have often required less efforts. For instance, updating using Bayes' rule does not represent a problem in KMM dynamic preferences.

known since Arrow, 1965). On the contrary, under CEU-preferences, there exists an interval of prices (and not a single point as it would be under SEU) within which the agent neither buys nor sells short the risky asset. At prices below (above) the lower (upper) limit of the interval, the agent is willing to buy (sell) the asset, but when equilibrium forces fail to push the asset price outside the interval, there will be no willingness to trade. The intuition behind this finding 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 $CE_v(\text{sell}) = 0.2 \times 1 + 0.8 \times 3 = 2.6$; this derives from the fact that in a max-min perspective, the weight is placed on a capacity of 0.2 in the bad state and, as a result, of 0.8 in the good state. 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.

Given two capacities v and w , Dow and Werlang defined the capacity v to be at least as uncertainty averse as w if for all events $E \subset \Omega$, $\mathcal{A}(v, E) \geq \mathcal{A}(w, E)$. In particular, it is easy to show that the no-trade interval, or, equivalently, the bid-ask spread (see Section 5 for further elaboration on this interpretation), associated to v , $-CE_v[-f] - CE_v[f]$, is wider relative to the one derived under w .³²

Another simple attempt at using basic ambiguity-related intuitions to tackle the empirical evidence of under-participation in (risky) asset markets and of biases in portfolio decisions has been recently proposed by Easley and O'Hara (2009). They investigate the problem of portfolio diversification between a riskless and an uncertain asset whose payoff d_i is normally distributed, with mean μ and payoff variance σ^2 . In their model, the two assets' net supplies are 0 and x , respectively. The J investors in the market display constant absolute risk aversion (CARA) negative exponential utility, with unit risk aversion parameter. An ambiguity averse investor holds beliefs that are compatible with multi-valued sets of possible mean returns and variances, that is, $\mu \in \{\mu_{\min}, \dots, \mu_{\max}\}$ and $\sigma^2 \in \{\sigma_{\min}^2, \dots, \sigma_{\max}^2\}$. Due to ambiguity aversion, the optimal portfolio of this investor is chosen by maximizing the minimum expected utility over sets of possible parameters, which is compatible with static versions of MPP. Simple algebra shows that—given some current price p for the risky asset—an ambiguity averse investor's demand function for the risky asset is:

$$\omega^{AA} = \begin{cases} \frac{\mu_{\min} - p}{\sigma_{\max}^2} & \mu_{\min} > p \\ 0 & \mu_{\min} \leq p \leq \mu_{\max} \\ \frac{\mu_{\max} - p}{\sigma_{\max}^2} & \mu_{\max} < p \end{cases} . \quad (27)$$

Therefore the decision of ambiguity averse investors to participate in the market for the risky asset is determined by μ_{\min} and μ_{\max} , in the sense that when the price falls in the compact interval $[\mu_{\min}, \mu_{\max}]$ there is no demand for the risky asset by the ambiguity averse traders. 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 (S)EU investors participate and trade the asset. Of course, if we

³²Another interesting implication of Dow and Werlang (1992) concerns insurance choices. Under (S)EU, an agent who can buy actuarially fair insurance in any amount will always choose to be fully insured. Instead under CEU-type ambiguity, a range of premium costs will exist at which the agent is not willing to get fully insured, which is a surprising result that may reconcile the widely observed tendency by risk averse agents to under-insure with rational decisions under uncertainty. Camerer and Weber (1992) have noticed that starting in the 1980s American insurance firms began raising premiums dramatically (or refusing to sell insurance at all) for several classes of highly ambiguous risks, like environmental hazards or manufacturing defects.

were to assume that all investors in the market are identically ambiguity averse, then when the prices falls in $[\mu_{\min}, \mu_{\max}]$, the market will freeze and all trading will stop.³³

Guidolin and Rinaldi (2009) have extended the results in Easley and O'Hara (2009), using a similar model. However, Guidolin and Rinaldi 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*). The systematic component of the stock payoff is normally distributed, with mean μ_S and variance σ_S^2 . Also the idiosyncratic component is normally distributed, with parameters μ_I and σ_I^2 . Call d the total payoff on the stock. Because systematic and idiosyncratic risks are independent, $d \sim N(\mu_I + \mu_S, \sigma_I^2 + \sigma_S^2)$. If agents are (S)EU maximizers, they consider as plausible a *unique* probability distribution for the idiosyncratic component of the asset's payoff and $(\mu_I, \sigma_I^2)'$ is a sufficient statistic for the distribution of idiosyncratic risk. In that case, the demand function for the risky stock is as usual $\omega^{SEU} = (\mu_I + \mu_S - p) / (\sigma_I^2 + \sigma_S^2)$ (where the exogenous inflation rate in Guidolin and Rinaldi has been set to zero). The classical result by which the sign of ω^{SEU} depends on the relation between the price of the asset and its expected payoff— $\omega^{SEU} \leq 0$ if and only if $p \geq (\mu_I + \mu_S)$ —and a zero net demand represents a knife-edge case, obtains. As in Easley and O'Hara (2009), the ambiguity averse agents act as if they have a set of distributions on returns; one distribution for each possible value of the unknown parameters, the sets $\{\mu_1, \mu_2, \dots, \mu_P\}$ and $\{\sigma_1^2, \sigma_2^2, \dots, \sigma_Q^2\}$ with $P \geq 2$ and $Q \geq 2$. In Guidolin and Rinaldi, Easley and O'Hara's result (27) is replaced by

$$\omega^{AA} = \begin{cases} \frac{\mu_{\min} + \mu_S - p}{\sigma_{\max}^2 + \sigma_S^2} > 0 & \text{if } \mu_{\min} + \mu_S > p \\ 0 & \text{if } \mu_{\max} + \mu_S \geq p \geq \mu_{\min} + \mu_S \\ \frac{\mu_{\max} + \mu_S - p}{\sigma_{\max}^2 + \sigma_S^2} < 0 & \text{if } \mu_{\max} + \mu_S < p \end{cases}, \quad (28)$$

and also in this case an interval of prices, $[\mu_{\min} + \mu_S, \mu_{\max} + \mu_S]$, exists for which it is optimal not to trade the risky asset. Because $\sigma_I \leq \sigma_{\max}$, the position in the risky asset held by ambiguity averse agents is always smaller (in absolute value) than the one held by (S)EU agents; because $\mu_{\min} \leq \mu_I \leq \mu_{\max}$, when it is optimal for an ambiguity averse agent to buy the risky asset, so it is for a (S)EU agent. Similarly, when it is optimal for an ambiguity averse agents to sell the risky asset, so it is for a (S)EU agent. Guidolin and Rinaldi (2009) have also generalized their results to the case in which ambiguity averse investors also perceive ambiguity deriving from the uncertainty on the systematic component of the risky payoffs. They 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.³⁴ 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.

Guidolin and Rinaldi's results on the relative importance of systematic vs. idiosyncratic risk in driving portfolio and limited participation results are consistent with the analysis of Mukerji and Tallon (1999). Mukerji and Tallon have considered an economy populated by ambiguity averse agents with CEU preferences in which risky assets are affected by both systematic and idiosyncratic uncertainty.³⁵ In their model there are two periods: one for trading

³³The decision to participate does not depend on σ_{\max}^2 . In addition, portfolio allocations will depend on σ_{\max}^2 only, and not on any other element of the set $\{\sigma_{\min}^2, \dots, \sigma_{\max}^2\}$. Section 4 will return on Easley and O'Hara's paper to illustrate their equilibrium implications.

³⁴Under 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.

³⁵Mukerji and Tallon's findings are not a mere generalization of Dow and Werlang's (1992) no-trade result since simply “closing” Dow and Werlang's model to clear markets is not sufficient to obtain a no-trade equilibrium.

(time 0), and one for consumption (time 1). The financial market is characterized by N non-redundant risky assets and by a safe bond. The value of the payoff of each risky asset is affected by circumstances peculiar to that specific security (idiosyncratic uncertainty), and others that affect the overall market (systematic uncertainty). At time 1, uncertainty unfolds, and the overall market can be in a high (H) or a low (L) state, with known probabilities p and $1 - p$, respectively. Hence, the possible realizations of the systematic component are y^H and y^L , with $y^H > y^L$. At time 1, agents also receive a stochastic endowment which is affected by systematic uncertainty. The possible realizations of the idiosyncratic component in the assets' payoff are $y(1)$ and $y(0)$, with $y(1) > y(0)$. These realizations are ambiguous, however, a capacity $(v(1), v(0))$, with associated ambiguity level $\mathcal{A}(v)$, is exogenously given.

Assuming $(y^H - y^L) < (y(1) - y(0))$, Mukerji and Tallon (1999) prove that 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 0 for all investors, so that the market breaks down and agents are suboptimally left exposed to (endowment) aggregate risk. Instead, under (S)EU, the equilibrium allocation would approximate a complete market allocation, since usual diversification arguments would imply that all idiosyncrasies are washed away when the number of assets becomes arbitrary large. Under CEU preferences this result no longer holds and trading in the well-diversified portfolio collapses. The intuition is that the law of large number is still at work but agents' beliefs on the payoff of the well diversified portfolio converge to different values depending on their specific (short or long) position.³⁶ Hence, ambiguity leads to a collapse in the trade of financial assets whose payoff is greatly affected by idiosyncratic risk when 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, provided that the ambiguity level is sufficiently high. 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. Indeed, what is crucial in determining the collapse in trade is the fact that, for each agent, there is an interval of prices that supports the optimal decision not to invest in the risky asset. Removing idiosyncratic risk, such an interval collapses to a point.^{37,38}

3.2. Mean-Variance Analysis Under Ambiguity

The historical workhorse for the development of most modern finance theory has been the simple and yet powerful mean-variance framework, in which (S)EU 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, single-period mean-variance framework.

³⁶Laws of large numbers for ambiguous beliefs have been studied by, among others, Walley and Fine (1982), Dow and Werlang (1994), Marinacci (1996) and Marinacci (1999).

³⁷Rinaldi (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.

³⁸Mukerji and Tallon (2003) have subsequently proven that the seminal results in Dow and Werlang (1992) on the existence of no-trade intervals for asset prices is strictly caused by ambiguity aversion and does not derive from any specific functional representation of their preferences. In practice, they show that trading will break down if and only if the investors behave in ways consistent with Ellsberg's paradox and therefore must exhibit some form of ambiguity aversion.

Consider a standard one-period representative agent economy, characterized by N risky assets and one riskless bond. Recall from Introductory Finance that under (S)EU and assuming that $\mathbf{R} \equiv [R_1, R_2, \dots, R_N]'$ follows a joint multivariate normal distribution with known variance-covariance matrix Σ and known vector of means μ , a mean-variance investor will solve the problem

$$\max_{\omega} [\omega'(\mu - r^f \boldsymbol{\iota}_N) + (1 + r^f)] - \frac{1}{2}\gamma\omega'\Sigma\omega, \quad (29)$$

subject to $\omega'\boldsymbol{\iota}_N = 1$, where γ is a risk aversion coefficient. The well-known solution is:³⁹

$$\omega = \frac{1}{\gamma}\Sigma^{-1}(\mu - r^f \boldsymbol{\iota}_N). \quad (30)$$

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} ; specifically, \mathbf{R} follows a joint multivariate normal distribution with known variance-covariance matrix Σ , and *unknown* vector of mean returns, μ . The agent displays MPP and some incomplete sources of information are available, so that she is able at least to estimate reference probabilistic models (and hence reference mean returns) for the joint distributions of asset returns. Following the constraint approach (10), and assuming a unique source of information, so that the agent is able to derive only a reference joint normal distribution of asset returns, $\hat{\mathbf{f}} \sim N(\hat{\mu}, \Sigma)$, the set of effective priors $\wp(\hat{\mathbf{f}})$ is

$$\wp(\hat{\mathbf{f}}) = \left\{ q : E[z \ln z] \leq \eta \quad z \equiv \frac{dq}{d\hat{\mathbf{f}}} \right\}, \quad (31)$$

where η captures ambiguity aversion (a larger η means higher aversion), and $E[\cdot]$ is the standard expected value operator under the reference model $\hat{\mathbf{f}}$. The investment problem is

$$\max_{\omega} V(W, \wp(\hat{\mathbf{f}})) = \max_{\omega} \min_{q \in \wp(\hat{\mathbf{f}})} \{E_q[U(W)]\} \quad \text{s.t. } W = [\omega'(\mathbf{R} - r^f \boldsymbol{\iota}_N) + (1 + r^f)], \quad (32)$$

($\omega'\boldsymbol{\iota}_N = 1$) where ω is the vector of assets' holdings, W is final wealth, and 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 $\theta \equiv \mu - \hat{\mu}$ 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 problem in (32) can be re-written as

$$\max_{\omega} \min_{\theta \in \{\theta : \frac{1}{2}\theta'\Sigma^{-1}\theta \leq \eta\}} E(z(\mathbf{R})U(W)) \quad z(\mathbf{R}) \equiv \exp \left\{ \frac{1}{2}\theta'\Sigma^{-1}\theta - \theta'\Sigma^{-1}(\mathbf{R} - \mu + \theta) \right\}, \quad (33)$$

subject to a budget constraint, which is an interesting transformation of the constraint on the set of admissible models under the parameter η into a (multiplicative) factor that appears in the objective function. A measure of aversion to ambiguity implied by a given choice of portfolio weights ω can then be derived by solving:

$$\Delta(\omega) = \sup_{\theta \in \{\theta : \frac{1}{2}\theta'\Sigma^{-1}\theta \leq \eta\}} \omega'\theta. \quad (34)$$

Denoting by ω^* and $\theta(\omega^*) \equiv (\omega^*)'\mu - \hat{\mu}$ the optimal solutions to the problem, the following pricing condition holds in the absence of arbitrage opportunities,

$$E[U'(W - \Delta(\omega^*))(\mathbf{R} - r^f \boldsymbol{\iota}_N - \theta(\omega^*))] = 0 \quad (35)$$

³⁹In 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.

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 (35), and denoting by ω_m the composition of the market portfolio, Kogan and Wang show:

$$\begin{aligned}\mu - r^f \boldsymbol{\iota}_N &= \lambda \boldsymbol{\beta} + \lambda_u \boldsymbol{\beta}_u \quad \lambda_u = \Delta(\omega_m) \quad \boldsymbol{\beta} = \frac{1}{\sigma_m^2} \boldsymbol{\Sigma} \omega_m = \frac{1}{\sigma_m^2} \text{Cov}(\mathbf{R}, r_m) \\ \lambda &= \frac{E[U''(W - \Delta(\omega_m))]}{E[U'(W - \Delta(\omega_m))]} \sigma_m^2 = \frac{E[U''(W - \Delta(\omega_m))]}{E[U'(W - \Delta(\omega_m))]} \omega'_m \boldsymbol{\Sigma} \omega_m \\ \boldsymbol{\beta}_u &= \frac{1}{\sigma_m^2} \boldsymbol{\Sigma}_u(\omega_m) \boldsymbol{\theta}(\omega_m) \quad \text{where } \boldsymbol{\Sigma}_u(\omega_m) \text{ s.t. } \boldsymbol{\theta}(\omega_m) = [\boldsymbol{\Sigma}_u(\omega_m)] \omega_m.\end{aligned}\tag{36}$$

This elegant result shows that the risk premium on any risky asset can be written in a simple two-factor representation: (i) the standard, static CAPM component with structure $\lambda \boldsymbol{\beta}$, where λ is the market risk premium and $\boldsymbol{\beta}$ is the vector of standard betas measured with respect to returns on the market portfolio; (ii) a new component given by $\lambda_u \boldsymbol{\beta}_u$ 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 $\boldsymbol{\theta}(\omega_m) \equiv \omega'_m \boldsymbol{\mu} - \hat{\boldsymbol{\mu}}$. Ambiguity is only partially diversifiable in the sense that, in equilibrium, for any asset only its individual contribution to total market ambiguity will be compensated. Finally, since the investor bears both risk and what they call “Knighthian/Keynesian uncertainty” (ambiguity), two assets with the same beta with respect to the 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 standard (static) CAPM.

Garlappi, Uppal, and Wang (2007, henceforth GUW) have further extended the findings in Kogan and Wang (2003) modeling ambiguity through a MPP approach—as opposed to Kogan and Wang’s constrained robust approach (10)—and using a “confidence interval” framework that has become popular in the literature. GUW’s starting point is that in reality the parameters of the joint normal density characterizing the ambiguous asset returns \mathbf{R} have to be estimated. Assuming that a time series of length T of past asset returns \mathbf{h}_t is available, the conditional density distribution of returns $g(\mathbf{R}|\mathbf{h}_t)$ must be derived. Assuming that the returns depend on some unknown parameters $\boldsymbol{\theta}$, whose prior is $\pi(\boldsymbol{\theta})$, the predictive density $g(\mathbf{R}|\mathbf{h}_t)$ is

$$g(\mathbf{R}|\mathbf{h}_t) = \int g(\mathbf{R}|\boldsymbol{\theta}) p(\boldsymbol{\theta}|\mathbf{h}_t) d\boldsymbol{\theta},\tag{37}$$

where

$$p(\boldsymbol{\theta}|\mathbf{s}_t) = \frac{\prod_{t=1}^T g(\mathbf{s}_t|\boldsymbol{\theta}) \pi(\boldsymbol{\theta})}{\int \prod_{t=1}^T g(\mathbf{s}_t|\boldsymbol{\theta}) \pi(\boldsymbol{\theta}) d\boldsymbol{\theta}}.\tag{38}$$

If an investor solves the classical mean-variance problem

$$\max_{\omega} \omega' \int \mathbf{R} g(\mathbf{R}|\mathbf{h}_t) d\mathbf{R} - \frac{1}{2} \gamma \omega' \boldsymbol{\Sigma} \omega,\tag{39}$$

the resulting portfolio is of a Bayesian type and it will take into full account the existence of parameter uncertainty (see, e.g., Barberis, 2000). The portfolios in which only parameter uncertainty is considered often perform poorly out of sample, even in comparison to portfolios selected according to some simple ad hoc rules, see e.g., Garlappi, DeMiguel, and Uppal (2009). 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 appropriate statistical model, as identified here by the vector of expected returns μ . When MPP are used, the optimization takes the form

$$\begin{aligned} \max_{\omega} \min_{\mu} & \omega' \mu - \frac{1}{2} \gamma \omega' \Sigma \omega \\ \text{s.t. } & \mathbf{f}(\mu, \hat{\mu}, \Sigma) \leq \eta \quad \omega' 1 = 1, \end{aligned} \quad (40)$$

where \mathbf{f} is a vector-valued function, and $\hat{\mu}$ is the estimate of μ derived from the predictive density $g(\mathbf{R}|\mathbf{h}_t)$. GUW (2007) prove that the max-min problem above is equivalent to the simpler maximization problem

$$\max_{\omega} \omega' (\hat{\mu} - \mu^{adj}) - \frac{1}{2} \gamma \omega' \Sigma \omega, \quad (41)$$

where $\hat{\mu} - \mu^{adj}$ is the “adjusted estimated expected return” (the adjustment has the role to incorporate ambiguity), and define a vector μ^{adj} to satisfy:

$$\mu^{adj} = \left[sgn(\omega_1) \frac{\sigma_1}{\sqrt{T}} \sqrt{\eta_1} \ sgn(\omega_2) \frac{\sigma_2}{\sqrt{T}} \sqrt{\eta_2} \dots sgn(\omega_N) \frac{\sigma_N}{\sqrt{T}} \sqrt{\eta_N} \right]'. \quad (42)$$

The adjustment depends on the precision with which parameters are estimated, the length of the data series, and the investor’s aversion to ambiguity (η). Even though the adjustment in (42) may seem to make the optimization problem practically as simple as a classical mean variance program—with $\hat{\mu}$ replaced by $\hat{\mu} - \mu^{adj}$ —this impression is incorrect because of the complex definition of μ^{adj} : the vector of adjusted expected returns depends from the vector of portfolio weights ω . Therefore (41) has to be solved numerically because ω enters the problem in non-linear fashion.⁴⁰

Kogan and Wang (2003) and GUW (2007) 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 (he calls it Knightian uncertainty) concerns the underlying *asset pricing models* that are potentially generating asset returns, and has pursued the portfolio implications of such a framework. Notice that while the frameworks in Kogan and Wang (2003) and GUW (2007) are simple enough to be potentially “closed” to deliver asset pricing implications under ambiguity (as we have seen possible in Kogan and Wang’s case), Wang’s (2005) framework is instead of a partial equilibrium nature. In particular, Wang (2005) considered the structure of optimal 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, see Fama and French, 1992) and in which ambiguity aversion is captured through MPP. Assume there are M risky assets and K factor portfolios, with \mathbf{x}_{1t} and \mathbf{x}_{2t} their respective vectors of excess returns over the risk-free rate during period t , where also the K factor portfolios are investable. The vector time series of T observations $\mathbf{R} = \{\mathbf{r}_t\}_{t=1,\dots,T} = \{(\mathbf{r}'_{1t} \ \mathbf{r}'_{2t})'\}_{t=1,\dots,T}$ is assumed to follow a multivariate normal distribution with mean μ and variance Σ . The corresponding mean and the variances are derived from the regression (K -factor) model $\mathbf{x}_{1t} = \alpha + \mathbf{B}\mathbf{x}_{2t} + \mathbf{u}_t$, where α is the $M \times 1$ vector of Jensen’s alphas (which capture the portion

⁴⁰GUW (2007) present an interesting international asset allocation application in which the portfolio optimizer is called to choose among eight international equity indices—a problem with structure and scope typical in the literature. The portfolio weights for each strategy (that is, with and without accounting for ambiguity) are determined on each month using moments estimated from a rolling-window of 120 months. The resulting out-of-sample period spans the interval January 1980 to July 2001 and risk aversion is set to 1. The results show that the model with MPP exhibit uniformly higher ex-post means, lower volatility, and, consequently, substantially higher Sharpe ratios. In particular, the standard Bayesian model performs poorly because it puts too much weight on the estimated expected returns.

of excess returns which is not simply explained by exposure to risk), \mathbf{B} is the $M \times K$ matrix of factor loadings, and \mathbf{u}_t is the $M \times 1$ vector of residuals with constant covariance matrix equal to the $M \times M$ North-West submatrix of Σ . By construction, an exact K -factor model holds if and only if $\mathbf{x}_{1t} = \mathbf{B}\mathbf{x}_{2t} + \mathbf{u}_t$, which means that on average, excess returns can only be generated by exposures to the risks captured by \mathbf{x}_{2t} . According to the standard (S)EU approach, investors who do not believe in (trust) the asset pricing model will estimate the parameters without (with) the restriction $\boldsymbol{\alpha} = \mathbf{0}$. In the (S)EU/Bayesian framework, a fixed positive number ζ is assumed to control the variance of the prior distribution of the Jensen's alphas, so that $E[\mathbf{x}_{T+1}|\mathbf{R}, \zeta]$ and $Var[\mathbf{x}_{T+1}|\mathbf{R}, \zeta]$ are the mean and variance of the predictive distribution, and different prior beliefs in asset-pricing models (i.e. different ζ s) will lead to different predictive distributions. The investors' optimization problem is as usual

$$\max_{\boldsymbol{\omega}} E[U(W; \zeta)], \quad (43)$$

subject to a standard budget constraint. If the agent is ambiguity averse, instead, she will behave as if she had a set of priors parametrized in terms of $\zeta \in [0, \infty)$, so that her portfolio's choice problem is:

$$\max_{\boldsymbol{\omega}} \min_{\zeta \in [0, \infty)} E[U(W; \zeta)]. \quad (44)$$

Wang (2005) uses a shrinkage approach to evaluate the impact of prior beliefs on the predictive distributions in the standard Bayesian framework and, allowing for MPP, he also investigates the effects of model uncertainty. Specifically, the *shrinkage factor* λ is defined as

$$\lambda^{-1} \equiv 1 + \frac{T\zeta}{1 + \widehat{SR}^2}, \quad (45)$$

where \widehat{SR} is the highest Sharpe ratio on the efficient frontier spanned by the sample mean and variance of the factor portfolios. Wang shows that λ can be interpreted as the weight assigned to the restricted model in which $\boldsymbol{\alpha} = \mathbf{0}$: if $\lambda = 1$ ($\lambda = 0$), full weight is assigned to the restricted (unrestricted) model. Wang shows that the predictive mean $E[\mathbf{x}_{T+1}|\mathbf{R}, \lambda]$ will be a weighted average of the restricted and the unrestricted predicted means, while the predictive variance of excess returns will be a quadratic function of the shrinkage factor, λ ; the predictive covariance between the assets and the factors will be proportional to the weighted average of restricted and unrestricted estimates.

Using a shrinking approach, Wang shows that the optimization problem in (44) simplifies to

$$\max_{\boldsymbol{\omega}} \min_{\lambda \in [0, 1]} E[U(W; \lambda)], \quad (46)$$

and that the particular λ^* that solves the max-min problem is the optimal shrinkage factor to be used that takes into account the aversion to model ambiguity. Wang uses these results to analyze a simple portfolio problem that allocates wealth between the value-weighted market index portfolio of all NYSE, AMEX, and NASDAQ stocks and the riskless asset over the period July 1963 - December 1998. He finds 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.

The static portfolio problems in Garlappi, Uppal, and Wang (2007), Kogan and Wang (2003), and Wang (2005) all have clear connections to the classical mean-variance approach. Even though their way of capturing ambiguity is to some extent *ad-hoc*, Pflug and Wozabal (2007) have recently directly generalized Markowitz's mean-variance approach

to take ambiguity into account. As introduced in the seminal paper by Markowitz (1959), the basic portfolio selection problem consists in maximizing expected portfolio returns under a risk constraint. Pflug and Wozabal suggested to describe such a constraint through the concept of an *acceptability* function. Specifically, for any portfolio (defined by the weights ω), the acceptability $A(\omega)$ is a concave, monotonic, translation-invariant and positively homogeneous real valued function defined on the payoffs space, such that $A(\omega)$ is the maximal shift in the distribution of the return of the portfolio defined by ω that still keeps the return acceptable according to the mean-variance criterion. Hence the goal of a portfolio choice problem is to find an allocation that maximizes returns and leads to high acceptability. Assuming that p is a reference probabilistic model for assets returns, and b is a positive threshold constant for acceptability, the optimal portfolio under (S)EU can be computed as

$$\max_{\omega} E_p \left[\sum_{n=1}^N \omega_n R_n \right] \quad \text{s.t. } A \left(\sum_{n=1}^N \omega_n R_n \right) \geq b \quad \sum_{n=1}^N \omega_n = 1 \quad \omega_n \geq 0, \quad (47)$$

Standard theory often quantifies the robustness of the solution of the optimization problem through stress testing. The stress test consists in finding a plausible set of probabilistic models, \wp (called the ambiguity set), and in computing the expected value and the acceptability of the optimal portfolio under any model $q \in \wp$. To specify the ambiguity set \wp , Pflug and Wozabal assumed that a reference model p can be found by estimation from historic data and what they call “consistency” considerations, so that any other alternative model q should be not too distant from p , that is, $\wp = \{q : d(q, p) \leq \eta\}$, for a given η . The particular metric $d(q, p)$ used in the paper is the Kantorovich distance, which has the property that the expected return of any portfolio has a Lipschitz constant 1, that is, given a return \bar{R} , $|E_q[\bar{R}] - E_p[\bar{R}]| \leq d(q, p)$. The size of the ambiguity neighborhoods may be chosen to correspond to a probabilistic confidence region for the estimated basic model. A “robust” portfolio is then the solution to:

$$\max_{\omega} \min_{q \in \wp} E_q \left[\sum_{n=1}^N \omega_n R_n \right] \quad \text{s.t. } A \left(\sum_{n=1}^N \omega_n R_n \right) \geq b \quad \forall q \ni d(q, p) \leq \eta \quad \sum_{n=1}^N \omega_n = 1 \quad \omega_n \geq 0. \quad (48)$$

In numerical simulations based on six randomly selected Dow Jones stocks, Pflug and Wozabal’s robust portfolios turn out to be more diversified than standard, stressed-tested portfolios. The worst case expected return decreases as more weight is given to robustness, as captured by an increasing η ; however, the overall expected return only slightly drops. Further, for this data set, the price (loss of ex-post expected return) of model robustness remains very small while the portfolio composition changes dramatically with the increase in the size of \wp induced by increases in η .

3.3. Robust Portfolio Decisions

Although Pflug and Wozabal’s (2007) paper is couched in a mean-variance framework, their results target the notion of characterizing robust portfolios, which are portfolios suboptimal under the selected reference model, but that may perform well under similar models. This is obviously suggestive of the possibility of studying portfolio decisions within an explicit dynamic framework that incorporates the robustness ideas by Hansen, Sargent, and their co-authors.

One of the early examples of this strategy is Trojani and Vanini (2000), who have used the constraint framework (10) of Anderson et al. (2003) to analyze a robust version of the classical Merton’s (1969) 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 Brownian motion (GBM), with BM B_t and constant parameters μ

and σ . This model is the continuous time version of the random walk with drift model for stock prices postulated by much classical finance literature under the efficient market hypothesis. Letting ω_t be the proportion of current wealth W_t invested in the risky asset, the budget constraint can be expressed as

$$dW_t = \{[r^f + \omega_t(\mu - r^f)]W_t - c_t\}dt + W_t\omega_t\sigma dB_t, \quad (49)$$

where c_t is the consumption rate per infinitesimal period. For any given initial wealth-price pair (W_0, S_0) , it is possible to define the operator $\{T_t\}_{t \geq 0}$ such that:

$$T_t(\phi(W_0, S_0)) = E[\phi(W_t, S_t) | (W_0, S_0)]. \quad (50)$$

The operator process $\{T_t\}$ is important because it allows to define a continuum of alternative probabilistic models by ‘‘contamination’’ of some baseline, approximating model characterized by a non-negative random variable $\theta \geq 0$, so that the corresponding operators $(T_t^\theta)_{t \geq 0}$ 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 then given by:

$$\mathcal{I}_t(\theta) = T_t\left(\frac{\theta}{T_t(\theta)} \ln\left(\frac{\theta}{T_t(\theta)}\right)\right). \quad (51)$$

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 some parameter, η . Hence the infinite-horizon investor’s consumption/portfolio optimization problem is

$$\max_{c, \omega} \min_{\theta} E\left[\int_0^\infty \exp(-\delta s) U(\omega_s) ds \middle| (W_0, S_0)\right] \quad \text{s.t.: } \lim_{t \rightarrow 0} \frac{\mathcal{I}_t(\theta)}{t} \leq \eta, \quad (52)$$

subject to the postulated dynamic processes for the stock price and for wealth. The constrained minimization in θ transforms the optimization problem into the equivalent program

$$\max_{c, \omega} E\left[\int_0^\infty \exp(-\delta s) U(c_s) ds \middle| (W_0, S_0)\right], \quad (53)$$

subject to the distorted dynamics:

$$\begin{aligned} dS_t &= (\mu - \sqrt{2\eta}\sigma)S_t + \sigma S_t dB_t \\ dW_t &= \{[r_t^f + \omega_t(\mu - r^f - \sqrt{2\eta}\sigma)]W_t - c_t\}dt + W_t\sigma\omega_t dB_t. \end{aligned} \quad (54)$$

Trojani and Vanini prove that the optimal distortion has the simple structure $\sqrt{2\eta}\sigma$. Assuming an isoelastic utility index $U(c) = c^{1-\gamma}/(1-\gamma)$ (with $\gamma \in (0, 1)$), Trojani and Vanini derive the optimal robust solutions:

$$\begin{aligned} \omega^* &= \frac{\mu - r^f - \sqrt{2\eta}\sigma}{\sigma^2\gamma} < \frac{\mu - r^f}{\sigma^2\gamma} = \omega^{SEU} \\ \frac{c^*}{W} &= \left[\frac{\delta}{\gamma} - (1-\gamma) \left(\frac{(\mu - r^f - \sqrt{2\eta}\sigma)^2}{2\sigma^2\gamma^2} + \frac{r^f}{\gamma} \right) \right]^{-1/\gamma} > \left[\frac{\delta}{\gamma} - (1-\gamma) \left(\frac{1}{2}(\omega^{SEU})^2 + \frac{r^f}{\gamma} \right) \right]^{-1/\gamma} = \frac{c^{SEU}}{W}, \end{aligned} \quad (55)$$

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 (S)EU ($\eta = 0$), while the investor will consume a higher fraction of her wealth. Moreover, it is easy to prove that $\partial\omega^*/\partial\eta < 0$ and $\partial c^*/\partial\eta > 0$, 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 simply more risk averse than what can be captured by the parameter γ appearing in her utility index.

Maenhout (2004) has further explored the effects of model uncertainty on dynamic portfolio/consumption decisions by postulating and solving a problem similar to Trojani and Vanini's (2000) with the difference that the investor has a finite horizon T , γ is unconstrained (apart from $\gamma > 0$, capturing risk aversion), while the penalty framework (9) of Anderson et al. (2003) is used. Preferences are characterized by the isoelastic utility index $c_t^{1-\gamma}/(1-\gamma)$ and ambiguity aversion degree φ .⁴¹ Maenhout (2004) derives consumption-portfolio rules that perform reasonably well if there is some form of model misspecification. In particular, while the state equation for wealth under standard (S)EU is the usual (49), under HS-type preferences this process must be adjusted by adding a specific diffusion term which is endogenously chosen to minimize the sum of the expected intertemporal utility and of an entropy penalty. Anderson et al. (2003) have shown that this distorted wealth-process can be written as

$$\begin{aligned} dW_t &= \mu(W_t)dt + \sigma(W_t)[\xi(W_t)\sigma(W_t)dt]dB_t \\ \mu(W_t) &= W_t(r^f + \omega_t(\mu - r^f)) - c_t \quad \sigma(W_t) = \omega_t\sigma W_t \end{aligned} \quad (56)$$

where the drift adjustment ξ is the solution to:

$$\inf_{\xi} \mathcal{D}V + \xi(W_t)\sigma^2(W_t)V_W + \frac{1}{2\varphi}\xi^2(W_t)\sigma^2(W_t). \quad (57)$$

$\mathcal{D}V$ is the standard infinitesimal generator of the HJB equation (see, e.g., Merton, 1971) applied to the value function V . Hence, the first two terms in (57) represent the expected continuation payoff when the dynamics of wealth is described by (56). The penalty applied to models that are “too distant” from the reference one is reflected in the third term, the derivative of relative entropy.

Maenhout notices that a major drawback of HS preferences is the lack of homotheticity in case of a power-utility index, so that the resulting optimal portfolio is not independent from the wealth level.⁴² To overcome this problem, the ambiguity aversion coefficient φ is re-scaled at any time-state pair, and replaced by $\varphi_{s,t} \equiv \varphi / [(1-\gamma)V(s,t)]$.⁴³ Hence, the necessary conditions for an optimum are:

$$\begin{aligned} \omega^* &= \frac{1}{\varphi + \gamma} \frac{\mu - r^f}{\sigma^2} < \frac{\mu - r^f}{\sigma^2\gamma} = \omega^{SEU} \quad \frac{c_t^*}{W_t} = \frac{a}{[1 - \exp(-a(T-t))]} \\ \text{where } a &\equiv \frac{1}{\gamma} \left[\delta - (1-\gamma)r^f - \frac{1-\gamma}{2(\varphi + \gamma)} \left[\frac{\mu - r^f}{\sigma^2} \right]^2 \right]. \end{aligned} \quad (58)$$

As long as $\varphi > 0$, the optimal share of the risky asset ω^* will differ from the standard (S)EU (Merton) 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 of 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.⁴⁴

⁴¹In Maenhout (2004), φ is the reciprocal of ϱ in (9), hence higher ambiguity aversion is represented by higher values of φ .

⁴²Preferences are homothetic if $\forall x > 0$, $x \sim y \leftrightarrow \chi x \sim \chi y$.

⁴³Maenhout (2004) provides excellent intuition for this operation.

⁴⁴Further results from Maenhout (2004) will be presented in Section 4.3. Notice that whether or not the share of wealth to be optimally consumed is increased or decreased by ambiguity aversion depends on whether $\gamma \leq 1$. Another interesting property—which however entirely depends on the special utility index adopted in Manheout (2004) as well as in Trojani and Vanini (2000)—is that when $\gamma = 1$, the presence of ambiguity aversion ($\eta > 0$ in Trojani and Vanini's work, $\varphi > 0$ in Manheout's) will not affect consumption decisions.

3.4. Does Constraint vs. Penalty Matter, and How?

Section 3.3 has emphasized that adopting the HS-type constraint framework (10) or the penalty framework (9) may make a difference for the structure of optimal 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 have explicitly re-considered 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 affect in different ways the optimal portfolio rules. In both frameworks, the asset menu is composed of a pure discount bond with risk-free rate r^f and an ambiguous stock with ex-dividend price S . x represents the exogenous primitives of the economy and it influences 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 process for the state x . In Maenhout's (2004) model, the optimization problem is

$$\sup_{c,\omega} \inf_{\boldsymbol{\theta}} E^c \left[\int_0^\infty e^{-\delta t} \left[\frac{(cW)_t^{1-\gamma}}{1-\gamma} + \frac{1}{2\varphi(W_t, x_t)} \|\boldsymbol{\theta}_t\|^2 \right] dt \right], \quad (59)$$

where $\varphi > 0$ is the state- and wealth-dependent ambiguity aversion function. As in Maenhout (2004), $\varphi \rightarrow 0^+$ (from the right) implies that any perturbations brought about by $\boldsymbol{\theta}_t$ will carry enormous penalties, so that the standard (S)EU framework obtains; on the opposite, if $\varphi \rightarrow \infty$, infinite aversion to ambiguity results. The second, Anderson et al. (2003)-style setting is based on a robust control constraint

$$\sup_{c,\omega} \inf_{\boldsymbol{\theta}} E^c \left[\int_0^\infty e^{-\delta t} \frac{(cW_t)^{1-\gamma}}{1-\gamma} dt \right] \quad \text{s.t. } \frac{1}{2} \boldsymbol{\theta}' \boldsymbol{\theta} \leq \eta, \quad (60)$$

As usual, the larger η , the higher aversion to ambiguity. $\eta = 0$ implies (S)EU consumption and portfolio choices.

The value function $V(W, x)$ that affects $\varphi(W_t, x_t)$ and solves (59) can be expressed in terms of another function that in general 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 policy can be interpreted as the optimal consumption of a (S)EU-investor with elasticity of intertemporal substitution $1/\gamma$, while the investment policy ω^* corresponds to the optimal investment policy of a (S)EU-investor with relative risk aversion artificially increased to $\gamma + \varphi$. The optimal investment ω^* is the sum of the myopic demand (say $\omega_{myopic}^*(\varphi)$) and the hedging demand (say $\omega_{hedge}^*(\varphi)$). Comparing both demands to the standard ambiguity-free, (S)EU case (that is, with $\omega_{myopic}^*(0)$ and $\omega_{hedge}^*(0)$), one gets:

$$\begin{aligned} \omega_{myopic}^*(\varphi) &= \frac{\gamma}{\gamma + \varphi} \frac{1}{\gamma} \frac{\mu_S - r^f}{\sigma_S^2} = \frac{\gamma}{\gamma + \varphi} \omega_{myopic}^*(0) < \omega_{myopic}^*(0) \\ \omega_{hedge}^*(\varphi) &= \frac{\gamma}{\gamma + \varphi} \frac{1 - \gamma + \varphi}{1 - \gamma} \omega_{hedge}^*(0). \end{aligned} \quad (61)$$

Hence ambiguity aversion always 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, so that optimal intertemporal hedging is modified accordingly. The term $\gamma/(\gamma + \varphi) < 1$ measures the reduction in hedging demand that is made necessary by the lower,

baseline myopic portfolio rule. 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.

The value function $V(W, x)$ that solves (60) can be expressed in terms of a function $G(x)$ that cannot be computed in closed form too. 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 the portfolio rule corresponds to the optimal investment policy of a (S)EU investor with relative risk aversion artificially increased to $\gamma + \varphi$, in a constraint-style problem the optimal portfolio rule can be re-interpreted as the optimal strategy of an investor with an effective risk aversion coefficient 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 . In fact, the myopic and hedging demand functions ($\omega_{myopic}^*(\eta)$ and $\omega_{hedge}^*(\eta)$) are related to those of the ambiguity free case ($\omega_{myopic}^*(0)$ and $\omega_{hedge}^*(0)$) as follows:

$$\begin{aligned}\omega_{myopic}^*(\eta) &= \omega_{myopic}^*(0) \frac{\gamma}{\gamma + \sqrt{2\eta/G(x)}} < \omega_{myopic}^*(0) \\ \omega_{hedge}^*(\eta) &= \omega_{hedge}^*(0) \frac{\gamma}{\gamma + \sqrt{2\eta/G(x)}} \frac{1 - \gamma + \sqrt{2\eta/G(x)}}{1 - \gamma},\end{aligned}\quad (62)$$

and while it is unequivocal that ambiguity aversion has the power to structurally reduce 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?

The papers by Trojani and Vanini's (2000, 2004) and Maenhout (2004) drew the attention of researchers on the possibility to derive closed-form solutions for portfolio/consumption problems under ambiguity, at least when the aversion to uncertainty is captured in the form of a preference for robustness. However, these papers did propose and solve stylized problems in which the portfolio decision is between one risky asset and a riskless bond. Uppal and Wang (2003) showed instead that it is possible to use ambiguity aversion to solve realistic portfolio problems and to shed light on unresolved questions in empirical asset allocation. In particular, Uppal and Wang (2003) 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*, which refers to the tendency of investors to bias their portfolio decisions in favor of assets issued in their countries (or even regions and cities, see for instance Coval and Moskowitz, 1999) and away from optimal allocation weights implied by standard (S)EU models (e.g., the widely invoked mean-variance framework); see e.g., Lewis (1999) for a review of the literature.

Uppal and Wang's approach is simple. When investors are (S)EU-maximizers, their recursive utility in a dynamic problem may (under appropriate conditions) be represented as

$$V_t(c_t, \omega_t) = U(c_t) + \beta E^p [V_{t+1}(c_{t+1}, \omega_{t+1})], \quad (63)$$

where p is a reference probability measure, $V_t(c_t, \omega_t)$ is the value function, and β is the subjective discount factor. When investors have a preference for robustness, in the case of a single source of misspecification, the representation

(63) is replaced by

$$V_t(c_t, \omega_t) = U(c_t) + \beta \inf_z \{ \phi(E_t^z [V_{t+1}(c_{t+1}, \omega_{t+1})]) \varrho \Psi(q(z)) + E^z [V_{t+1}(c_{t+1}, \omega_{t+1})] \}, \quad (64)$$

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 , $\Psi(q(z))$ 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 $\varrho \rightarrow \infty$ ($\rightarrow 0^+$, i.e., from the right) ambiguity declines (increases); in fact a $\varrho \rightarrow \infty$ imposes an infinite penalty to any divergence (as measured by relative entropy) of z from p , and as such it forces the decision-maker to use the benchmark p , as postulated under (S)EU.

Uppal and Wang prove that when $\varrho \rightarrow \infty$, the optimal risky portfolio is identical to the standard Merton's portfolio; when $\varrho \rightarrow 0$, 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, invest in N risky stocks, and borrow or lend 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

$$\begin{aligned} d\mathbf{R}_t &= \boldsymbol{\mu}_R(\mathbf{R}_t, \mathbf{Y}_t)dt + \boldsymbol{\sigma}_R(\mathbf{R}_t, \mathbf{Y}_t)d\mathbf{B}_t \\ d\mathbf{Y}_t &= \boldsymbol{\mu}_Y(\mathbf{Y}_t)dt + \boldsymbol{\sigma}_Y(\mathbf{Y}_t)d\mathbf{B}_t, \end{aligned} \quad (65)$$

with \mathbf{B}_t a K -dimensional BM. As usual, the dynamics of wealth, for given decisions ω and c_t , is given by:

$$dW_t = W_t \left[r_t^f + \omega'_t(\boldsymbol{\mu}_R - r_t^f \boldsymbol{\iota}_N) - \frac{c_t}{W_t} \right] dt + W_t \omega'_t \boldsymbol{\sigma}_R d\mathbf{B}_t. \quad (66)$$

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 (S)EU Bellman equation augmented by the terms $\frac{\Psi(V)}{2} \mathbf{v}' \boldsymbol{\Phi} \mathbf{v} + \nabla V_W W \omega' \mathbf{v}_R + \nabla V_Y \mathbf{v}_Y + \nabla V_R \mathbf{v}_R$, where $\boldsymbol{\Phi}_{nn} \equiv \sum_{i=1}^N \varrho_i [\sigma_{J_i}, \sigma'_{J_i}]_n$ and $\mathbf{v} \equiv [v'_Y \ v'_R]'$. The term $0.5\Psi(V)\mathbf{v}' \boldsymbol{\Phi} \mathbf{v}$ is the penalty function for the robustness-driven distortions, while the next three terms are drift adjustments due to the change of measure. From the HJB, the optimal portfolio rule is

$$\omega^* = -\frac{1}{W \nabla^2 V_{WW}} [\boldsymbol{\sigma}_R \boldsymbol{\sigma}'_R]^{-1} [\nabla V_W (\boldsymbol{\mu}_R - r^f + \mathbf{v}_R^*) + \boldsymbol{\sigma}_R \boldsymbol{\sigma}'_Y \nabla^2 V_{WY} + \boldsymbol{\sigma}_R \boldsymbol{\sigma}'_R \nabla^2 V_{WR}], \quad (67)$$

where

$$\begin{bmatrix} \mathbf{v}_R^* \\ \mathbf{v}_Y^* \end{bmatrix} = -\frac{1}{\Psi(V)} \boldsymbol{\Phi}^{-1} \begin{bmatrix} \nabla V_W + \nabla V_R \\ \nabla V_Y \end{bmatrix} \quad (68)$$

Without the term \mathbf{v}_R^* , the portfolio rule reduces to the standard Merton's (1971) formula, indeed \mathbf{v}_R^* is the drift adjustment to $\boldsymbol{\mu}_R$ induced by ambiguity. Similarly, when the coefficients ϱ_i tends to infinity for $i = 1, 2, \dots, N$, $\boldsymbol{\Phi}^{-1}$ goes to \mathbf{O} and we obtain the familiar Merton's result.

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, but when the investor's risk aversion is α is replaced by (magnified to) $\alpha(1 + 1/\varrho)$:

$$\omega^* = \frac{1}{\alpha(1 + 1/\varrho)} \boldsymbol{\Sigma}_R^{-1} (\boldsymbol{\mu}_R - r^f \boldsymbol{\iota}_N) < \frac{1}{\alpha(1 + 1/\varrho)} \boldsymbol{\Sigma}_R^{-1} (\boldsymbol{\mu}_R - r^f \boldsymbol{\iota}_N) = \omega^{SEU}. \quad (69)$$

This means that—at least in the case in which ϱ is common across assets—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 $(1 + 1/\varrho) > 1$. Consistently with our earlier remarks, $\varrho \rightarrow \infty$ implies that this factor goes to 1, while $\varrho \rightarrow 0$ implies that the factor diverges and drives ω^* to a vector of zeros.

Uppal and Wang calibrated their model to the data on international equity returns from French and Poterba's (1991) paper on the empirics of the home country bias and compared the resulting optimal assets allocations to the standard Bayesian 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 puzzle: 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.⁴⁵

Although only in the context of a simple two-period (heterogeneous) two-agent economy, Epstein (2007) has shown that the intuition in Uppal and Wang (2003) extends to other frameworks that can be used to capture ambiguity, such as MPP. In particular, the macroeconomics literature has shown that another puzzle is associated with the home country bias in portfolio choice: individual country GDP (or consumption) growth rates tend to be insufficiently correlated and they seem to massively respond to idiosyncratic, country-specific shocks, reflecting relatively modest and sub-optimal risk-sharing (see Lewis, 1999).⁴⁶ Epstein considers two countries and two periods ($t - 1$ and t) populated by investors with log-utility indices, and uncertainty that has resolution at time t so that for each country the realized state can be either -1 or 1 . Countries display MPP, i.e., the agent in country $i = 1, 2$ assigns to its own country-specific state space a unique probability distribution $[1/2 \ 1/2]'$, while—with respect to the foreign country—she is only confident that the probability of each possible outcome lies in the interval $[(1 - \kappa^i)/2, (1 + \kappa^i)/2]$, where $0 < \kappa^i < 1$. The set of effective priors considered by country i is denoted by \wp^i and the aggregate stochastic endowment by $[Y_{t-1} \ Y_t]'$, with expected growth rate $e^{\mu^Y} - 1$. Under (S)EU, individual consumption c_t^i is an increasing linear function of Y_t , and therefore it is perfectly correlated across countries. Further, if preferences are homothetic, in equilibrium consumption growth rates are equal. Under MPP, country i 's utility function is given by:

$$V^i(c_{t-1}^i, c_t^i) = \ln c_{t-1}^i + \beta \min_{q^i \in \wp^i} E_{q^i} [\ln c_t^i]. \quad (70)$$

Solving the planning problem, Epstein finds that individual country consumption will be non-negatively correlated with shocks in both countries, and that consumption growth rates will not be perfectly correlated across countries.

Uppal and Wang's (2003) intuition that small amounts of heterogeneity in the ambiguity investors perceive “across” different asset types may cause important effects has been recently further explored by Boyle, Garlappi, Uppal, and Wang (2009) who study the role of ambiguity (or assets’ “familiarity” to a decision maker) in determining portfolio underdiversification and “flight to familiarity” episodes. Consider a mean-variance (or equivalently, CARA/Gaussian)

⁴⁵Epstein and Miao (2003) is another paper that has exploited ambiguity to discuss the home country bias in international equity portfolios. We will review Epstein and Miao's contributions in Section 4.2.

⁴⁶In a complete markets benchmark, investors worldwide should use asset markets to diversify away all idiosyncratic risks, including their national, business cycle risks. As a result, consumption growth processes across countries should become highly (in the limit, perfectly) correlated. Of course, if investors fail to internationally diversify their portfolios will display a home-country bias and shocks to their national markets (economies) will be strongly reflected in their consumption decision, with a limited amount of international risk-sharing.

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 Σ is the implied covariance matrix, with generic element off the main diagonal given by $\rho\sigma^2$. Using the framework developed by Garlappi et al. (2007), Boyle et al. (2009) write the optimization problem as

$$\max_{\omega} \min_{\mu} \omega' \mu - \frac{\gamma}{2} \omega' \Sigma \omega \quad \text{s.t.: } \frac{(\mu_i - \hat{\mu}_i)^2}{\sigma_{\hat{\mu}_i}^2} \leq \eta_i \quad \omega' \boldsymbol{1}_N = 1, \quad (71)$$

where μ is a N -dimensional vector that collects the individual expected returns, μ_i , $i = 1, \dots, N$. 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 $i = 1, \dots, N$ obtained by using a return time series of length T , and $\sigma_{\hat{\mu}_i}^2$ the variance of $\hat{\mu}_i$ —it is possible to define the confidence interval $\{T(\mu_i - \hat{\mu}_i)^2 / \sigma_{\hat{\mu}_i}^2 \leq \eta_i\}$ 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 decision maker, and the remaining $N - 1$ assets are all equally less familiar (i.e., $\eta_1 = \eta_F$ and $\eta_2 = \dots = \eta_U$, $\eta_F < \eta_U$), 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

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

and $\omega_1 > 1/N > \omega_i$, $i = 2, \dots, N$. 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

$$\omega_1 = 1 \quad \omega_i = 0 \quad i \geq 2,$$

and a 100% investment 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 regarding the return distributions, as captured by the term $\sqrt{\eta_U} - \sqrt{\eta_F}$. When the number of assets becomes arbitrarily large, it can be shown that 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 $N \rightarrow \infty$, while the existence of gains from diversification implies that an investor should hold only an infinitesimal amount in each of the remaining unfamiliar assets. Furthermore, Boyle et al. (2009) show that in both the first and second cases, portfolio variance will exceed what would be obtained in absence of ambiguity, since the investor's portfolio is less diversified. The additional risk in case 1 increases with the difference $\sqrt{\eta_U} - \sqrt{\eta_F}$ and the extent of average idiosyncratic risk characterizing the risky assets. Finally, an increase in correlation leads to an increase in the holding of the familiar asset (the so-called "flight to familiarity") since the unfamiliar assets become decreasingly useful for diversification purposes.

3.6. Smooth Recursive Preferences

One further approach to modelling ambiguity in asset allocation decisions has been recently pursued by Ju and Miao (2008) who capture ambiguity aversion through a generalization of KMM's recursive smooth preferences. In Ju and Miao's static setting the representative agent ranks acts according to the functional

$$v^{-1} \left(\int_{\wp} v \left(U^{-1} \left(\int U(c) dp \right) \right) dq(p) \right), \quad (72)$$

where $v(\cdot)$ is an increasing function, and \wp is a non-singleton set of probability distributions over the reference state space. The agent holds a prior q over \wp and attitudes towards ambiguity are characterized by the shape of $v(\cdot)$. Generalizing the previous, static representation to an infinite horizon, dynamic framework, the recursive representation of preferences in Ju and Miao (2008) is

$$V_t(c) = W(c_t, \mathbf{R}_t(V_{t+1}(c))) \quad \mathbf{R}_t(\cdot) \equiv v^{-1} \left(\int_{\wp_{x_t}} v \left(U^{-1} \left(\int U(\cdot) dp \right) \right) dq_t(p) \right), \quad (73)$$

where x_t is the time- t realization of a state variable x that determines the set \wp_{x_t} of effective measures.

In this framework, Chen, Ju and Miao (2008) consider 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 instance, the de-meanned dividend yield as in Fama and French, 1988) for excess stock returns:

$$\begin{aligned} \text{Model 1(IID): } r_{e,t+1} - r_{f,t+1} &= r + \varepsilon_{1,t+1} \quad \varepsilon_{1,t+1} \sim \mathcal{N}(0, \sigma_1^2) \\ \text{Model 2 (VAR): } \begin{cases} r_{e,t+1} - r_{f,t+1} = r + bx_t + \varepsilon_{2,t+1} \\ x_{t+1} = \rho x_t + \varepsilon_{3,t+1} \end{cases} & \quad [\varepsilon_{2,t+1}, \varepsilon_{3,t+1}]' \sim \mathcal{N}(0, \Omega) \quad \Omega = \begin{bmatrix} \sigma_2^2 & \sigma_{23} \\ \sigma_{23} & \sigma_3^2 \end{bmatrix}. \end{aligned} \quad (74)$$

The investor is uncertain as to which model is correctly specified, even when she can learn on the relative merits of the two processes by observing past data. Time is finite and discrete ($t = 0, 1, \dots, T$), and the investor's only source of income is her financial wealth (W_t). In each period t she 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. Using the updating rule in (73), 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 ($x_t < 0$).

Chen et al. (2008) calibrate their model using annual data from the U.S. stock market over the period 1926–2005. 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 a 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. For example, when the predictive variable takes a sufficiently high value, the VAR strategy advises the investor to allocate all her wealth to the stock. By contrast, the robust strategy advises the investor to sharply decreases her stock allocation to a level below that implied by the IID strategy, even when there is a small prior probability that the IID model is the true model of the stock return.

3.7. Incomplete Information vs. Ambiguity: Comparative Analysis and Interaction Effects

Before the advent of research on ambiguity and optimal asset allocation decisions one of most fruitful directions of investigation had focussed on solving dynamic portfolio choice/ consumption problems under incomplete information, when the investor cannot observe one or more features (either parameters or state variables) of the underlying environment and can only rationally learn/filter this information out of past data on asset prices and/or fundamentals (e.g., dividends). Among many others, seminal papers are Feldman (1986), Gennotte (1986), and David (1997). 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 as well as of ambiguity on optimal portfolio weights 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 remarkable steps towards addressing these questions.

Miao (2009) has studied optimal consumption/portfolio choices under incomplete information and ambiguity, assuming that agents display dynamic MPP as in Chen and Epstein (2001), with their κ -ignorance specification (see Section 4.2 for details). In his framework, time is continuous, there is a single good and c_t denotes consumption. $\mathbf{B} = \{\mathbf{B}_t\}$ is a d -dimensional BM on (Ω, \mathcal{F}, p) . The decision maker’s available information is represented by a sub-

filtration $\{\mathcal{F}_t^S\}$, where $\mathcal{F}_t^S \subset \mathcal{F}_t$, which is generated by $d' > d$ -dimensional BM $\hat{\mathbf{B}} = \{\hat{\mathbf{B}}_t\}$. In particular, there are $d+1$ securities consisting of one riskless bond and d non-dividend-paying stocks. The price of the riskless bond is given by $S_0^t = e^{r_f t}$, while the vector return process of the stocks, \mathbf{R} (where $R_t^i = dS_t^i/S_t^i$), follows a multivariate GBM with constant drift and diffusion $\mu^{\mathbf{R}}$ and $\sigma^{\mathbf{R}}$. The price process \mathbf{S} is what generates the observable filtration $\{\mathcal{F}_t^S\}$, and the consumption $\{c_t\}$ and portfolio $\{\omega_t\}$ processes must be adapted to $\{\mathcal{F}_t^S\}$. $\mathcal{F}_t^S \subset \mathcal{F}_t$ captures the presence of incomplete information. A recursive multiple-priors utility process $\{V_t(c)\}$ is defined by the information structure $\{\mathcal{F}_t^S\}$, the BM $\hat{\mathbf{B}} = \{\hat{\mathbf{B}}_t\}$, a set of priors \wp^θ on $\{\mathcal{F}_t^S\}$, a discount rate $\beta > 0$, and a utility index $u(\cdot)$. Without a concern for ambiguity, the BM $\hat{\mathbf{B}} = \{\hat{\mathbf{B}}_t\}$ can be defined as

$$\hat{\mathbf{B}}_t = \int_0^t (\sigma^{\mathbf{R}})^{-1} [d\mathbf{R}_s - \hat{\mu}_s^{\mathbf{R}} ds], \quad (75)$$

where $\hat{\mu}_t^{\mathbf{R}} = E[\mu^{\mathbf{R}} | \mathcal{F}_t^S]$ is a measurable version of the conditional expectation of $\mu^{\mathbf{R}}$ with respect to \mathcal{F}_t^S . Hence, the agent's perceived return dynamics is $\mathbf{R}_t = \hat{\mu}_t^{\mathbf{R}} dt + \sigma^{\mathbf{R}} d\hat{\mathbf{B}}_t$, so that the budget constraint becomes

$$dW_t = ([r^f + \omega'_t(\hat{\mu}_t^{\mathbf{R}} - r^f \iota_d)] W_t - c_t) dt + W_t \omega'_t \sigma^{\mathbf{R}} d\hat{\mathbf{B}}_t. \quad (76)$$

The perceived (as opposed to the actual unknown vector η) market price of risk is assumed to be:

$$\hat{\eta}_t = (\sigma^{\mathbf{R}})^{-1} (\hat{\mu}_t^{\mathbf{R}} - r_t^f \iota_d) \quad (77)$$

Hence, under (S)EU, the investor solves the maximization problem under the constraint (76). When ambiguity is taken into account, the κ -ignorance specification, leads to a density generator (vector) process θ that induces a distorted probability measure q^θ . Specifically, denoting by θ^* the generator that corresponds to the optimal consumption process c^* , Miao shows that the agent's perceived returns dynamics is

$$d\mathbf{R}_t = (\hat{\mu}_t^{\mathbf{R}} - \sigma^{\mathbf{R}} \theta_t^*) dt + \sigma^{\mathbf{R}} d\hat{\mathbf{B}}_t, \quad (78)$$

so that there are two factors influencing the deviations of the agent's perceived mean returns from their true values:

$$\mu^{\mathbf{R}} - (\hat{\mu}_t^{\mathbf{R}} - \sigma^{\mathbf{R}} \theta_t^*) = (\mu^{\mathbf{R}} - E[\mu^{\mathbf{R}} | \mathcal{F}_t^S]) + \sigma^{\mathbf{R}} \theta_t^*. \quad (79)$$

The first term represents what Miao's defines *estimation* (incomplete information) *risk*, and the second term reflects ambiguity. Because these terms are time-varying, investment opportunities change over time and two separate hedging motives will arise, namely for ambiguity and estimation risk.

Under power utility $U(c) = c^{1-\gamma}/(1-\gamma)$, the optimal portfolio process is

$$\omega_t^* = \frac{1}{\gamma} (\sigma^{\mathbf{R}} (\sigma^{\mathbf{R}})')^{-1} (\hat{\mu}_t^{\mathbf{R}} - r^f \iota_d) - \frac{1}{\gamma} ((\sigma^{\mathbf{R}})')^{-1} \theta_t^* + ((\sigma^{\mathbf{R}})')^{-1} \sigma_t^{\mathbf{H}}, \quad (80)$$

where $\{\mathbf{H}_t, \sigma_t^{\mathbf{H}}\}$ is the unique solution to the stochastic differential equation $d\mathbf{H}_t/\mathbf{H}_t = \mu_t^{\mathbf{H}} dt + \sigma_t^{\mathbf{H}} d\hat{\mathbf{B}}_t$ ($\mathbf{H}_T = \mathbf{0}$), where $\mu_t^{\mathbf{H}}$ is a complex function of the preference parameters, θ_t^* , $\hat{\eta}_t$, and $\sigma_t^{\mathbf{H}}$ is the matrix of diffusion coefficients of the optimal consumption process. The last two terms in the expression for ω_t^* in (80) are the hedging demands, which are visibly affected by both ambiguity (represented by θ_t^*) and by the estimation risk (represented by $\hat{\eta}_t$) caused by incomplete information:

$$-\frac{1}{\gamma} ((\sigma^{\mathbf{R}})')^{-1} \theta_t^* + ((\sigma^{\mathbf{R}})')^{-1} \sigma_t^{\mathbf{H}}. \quad (81)$$

Ambiguity distorts mean returns at time t by $\sigma^R \theta_t^*$ under the adjusted belief. As $\gamma \rightarrow 1$, 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 $\sigma_t^H = \mathbf{0}$, and the agent behaves myopically so that there is no hedging demand against future changes of investment opportunities. However, in this case $\gamma^{-1} ((\sigma^R)')^{-1} \theta_t^*$ simply converges to $((\sigma^R)')^{-1} \theta_t^*$ so that an ambiguity hedging demand will be present also in the logarithmic case, contrary to what found under incomplete information. Miao also shows that under ambiguity, but with complete information, the optimal portfolio reduces to

$$\omega_t^{**} = \frac{1}{\gamma} (\sigma^R (\sigma^R)')^{-1} (\mu_t^R - r^f \iota_d) - \frac{1}{\gamma} ((\sigma^R)')^{-1} \kappa, \quad (82)$$

where κ is the vector of security-specific κ -ignorance parameters. (82) is interesting because it indicates that under complete information, ambiguity does not induce any hedging demand, in the sense that the terms that is not in the classical Merton's mean-variance style solution, $\gamma^{-1} ((\sigma^R)')^{-1} \kappa$, fails to be time-varying and to depend from the state of the economy. It is therefore the interaction between incomplete information and ambiguity that induces the hedging demand $-\gamma^{-1} ((\sigma^R)')^{-1} \theta_t^*$ in (80): 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. Of course, it remains to be seen whether this intriguing result may generalize beyond the case of MPP with κ -ignorance.

A related paper is Liu (2009), who examines a continuous-time portfolio problem for an investor with recursive MPP and κ -ignorance (see Section 4.2 for this specification of MPP), 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 continuous 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 paying an instantaneous return r^f and of a risky asset with Ito dynamics. 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} ($i, j = 0, 1$).⁴⁷ 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 2 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 distorted law of motion of the estimate π_t is

$$d\pi_t = [\lambda_0 - (\lambda_0 + \lambda_1) \pi_t] dt + \pi_t(1 - \pi_t) \frac{\mu_H - \mu_L}{\sigma_t} (dB_t^{q^\theta} - \theta dt), \quad (83)$$

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

⁴⁷The infinitesimal generating matrix governs the dynamics of the Markov switching between states of high drift (expect returns) in the stock process and states of low drift.

Liu (2009) explicitly characterizes the optimal consumption and portfolio rules. Specifically, 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

$$\omega_t = \frac{(\pi_t \mu_H + (1 - \pi_t) \mu_L) - r^f}{\gamma \sigma_t^2} - \frac{\kappa}{\gamma \sigma_t^2} + \text{hedge}^{IIR} + \text{hedge}^{ambiguity}, \quad (84)$$

where γ is the standard constant relative risk aversion coefficient (the expressions of the terms hedge^{IIR} , $\text{hedge}^{ambiguity}$ are omitted to save space). The first term 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 term reflects ambiguity on the myopic component, which is a function of the magnitude of the 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 the size of ambiguity is equivalent to a decline in the effective market price of risk:

$$\frac{(\pi_t \mu_H + (1 - \pi_t) \mu_L) - r^f}{\gamma \sigma_t^2} - \frac{\kappa}{\gamma \sigma_t^2} = \frac{[(\pi_t \mu_H + (1 - \pi_t) \mu_L) - \kappa] - r^f}{\gamma \sigma_t^2} = \frac{\hat{E}_\pi^\kappa[\mu] - r^f}{\gamma \sigma_t^2}. \quad (85)$$

However, under incomplete information, this equivalence fails because ambiguity also has an impact on the intertemporal hedging demand. The term $(\text{hedge}^{IIR} + \text{hedge}^{ambiguity})$ quantifies 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 those derived in a standard (S)EU portfolio problem under Markov switching (see e.g., Guidolin and Timmermann, 2007, in discrete time and Honda, 2003, in continuous time) 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 $\text{hedge}^{ambiguity}$ 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 $\text{hedge}^{ambiguity}$ is purely driven by ambiguity. Whilst the hedging term hedge^{IIR} will also exist in a (S)EU framework with incomplete information such as Honda's (2003), $\text{hedge}^{ambiguity}$ 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 $\gamma > 1$ ($\gamma < 1$), 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 $\text{hedge}^{ambiguity}$ 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.

Liu (2009) has extended the ambiguity literature to consider continuous time models in which the process for the uncertain asset returns is non-linear, for instance according to a Markov switching model. 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 (S)EU framework of Chacko and Viceira (2005) 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, which is the same result as in Chacko and Viceira. 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 and this demand can be decomposed into three components, namely, myopic demand, hedging demand, and ambiguity demand. The same calibration of Chacko and Viceira (based on monthly excess stock returns on the CRSP value-weighted portfolio over the T-Bill rate from January 1926 to December 2000) shows that ambiguity hedging is larger than the intertemporal hedging demand.

As we have seen, Chen, Ju and Miao (2008) have systematically compared portfolio shares and realized recursive performance from Bayesian vs. their robust strategy. A number of authors have examined the differences between (estimation) risk in a Bayesian framework and uncertainty, when ambiguity is taken into account. This is an important issue because a strand of the empirical finance literature (see e.g., Pastor, 2000) had powerfully drawn the attention of academics and practitioners on the effects of *model uncertainty*—where “model” may stand both for a range of asset pricing models and for alternative statistical representations of the law of motion of asset returns—on optimal portfolio choice in Bayesian frameworks. One may suspect that a typical (S)EU/Bayesian approach to model uncertainty would generate implications that are at least qualitatively similar to ambiguity approaches to uncertainty. It turns out that this intuition is incorrect even to a first approximation, because the two approaches differ both in their methodological underpinnings and in their practical implications. For instance, Uppal and Wang (2003) perform a thorough comparison of their ambiguity-averse dynamic portfolio model that accounts for model misspecifications with the traditional (S)EU Bayesian approach that incorporates the effects of estimation risk (e.g., Barberis, 2000). In the Bayesian approach, model misspecification can only be captured in the form of parameter uncertainty. Suppose that a model of the probability law for asset returns is estimated and there exists one parameter which cannot be estimated precisely, say θ . Let $p(\mathbf{X}_t; \theta)$ be the probability distribution function for the state vector \mathbf{X}_t . Given that θ is unknown, in a typical Bayesian framework the issue faced by an investor is how to take estimation uncertainty into account. For instance, suppose $F(\theta)$ is a prior on the unknown parameter. Under standard Bayesian (S)EU preferences, the expected utility from a given consumption stream c_t is:

$$V_t(c_t, \omega_t) = U(c_t) + \beta E^F \{E^{p(\mathbf{X}_t; \theta)}[V_{t+1}(c_{t+1}, \omega_{t+1})]\}. \quad (86)$$

In contrast, in Uppal and Wang’s (2003) framework, one need not restrict model misspecification to uncertainty regarding a particular parameter and need not assume that model misspecification, as a subjective matter, can be represented by a probability distribution, which leads to consumption/portfolio problem

$$V_t(c_t, \omega_t) = U(c_t) + \beta \inf_z \{\phi(E_t^z [V_{t+1}(c_{t+1}, \omega_{t+1})]) \varrho R(q(z)) + E^z [V_{t+1}(c_{t+1}, \omega_{t+1})]\}, \quad (87)$$

where ϱ , $\phi(\cdot)$, z , and $R(q(z))$ are defined in Section 3.5. A number of papers—here we have reported implications from Chen, Ju and Miao (2008)—have stressed that the qualitative and quantitative asset allocation implications of Bayesian frameworks of model uncertainty and of models of ambiguity aversion may be strikingly different.

3.8. Learning, Ambiguity and Portfolio Choice

As we have argued in Section 2.5, one of the frontiers of research on the effects of ambiguity on financial decisions and prices consists of the interaction between learning dynamics and ambiguity. While the papers reviewed in Section 3.7 are mostly concerned with the interaction between the need to filter from the data some unknown features of the environment and a concern for ambiguity, a recent literature has also investigated how learning may affect recursive, dynamic portfolio decisions. In this new strand of the literature, the path has been opened by Epstein and Schneider (2007), who have used their learning model to introduce and 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 (i.e., the perception of ambiguity).

Consider an investor who believes that the equity (market) 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 (i.e., in the absence of either regime switching a' la Veronesi, 2002, and Whitelaw, 2002, or of breaks a' la Guidolin, 2006) learning produces effects that quickly dissipate over time. For instance, Guidolin's (2005) Bayesian dynamic model of international portfolio diversification can be directly contrasted to Uppal and Wang's (2003) application of ambiguity to solve the home country bias: Guidolin obtains strong but highly transient effects from Bayesian 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 adequately 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 that she derives from her final wealth W_{t+T} , where her utility index is logarithmic, $U(W_{t+T}) = \ln W_{t+T}$. The 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, the investor is assumed to have a prior over the probability of the high state, so that the posterior mean of p is \hat{p}_t . In the limit as $k \rightarrow \infty$ (if continuous rebalancing is possible), under (S)EU—here simply identified with Bayes' rule—Epstein and Schneider derive that the optimal portfolio allocation at any date t is the mean-variance allocation

$$\omega_t^{Bayes} = \frac{\bar{R}_t + \frac{1}{2}\sigma^2 - r^f}{\sigma^2}, \quad (88)$$

where \bar{R}_t is the monthly equity premium estimated as of time t . Allowing for ambiguity, beliefs are represented by $(\Xi, M_0, L_k, \varsigma)$, where Ξ is a parameter space, M_0 is a set of probability measures on Ξ (not necessarily unique), and L_k is a set of likelihoods that represent the agent's a priori view of the connection between the signals and parameters. The size of M_0 reflects the decision-maker's (lack of) confidence in the prior information. Differently from the (S)EU case, when M_0 does not degenerate in a unique prior, the set of likelihoods L_k is directly dependent on k , and there is an additional parameter (ς) that captures the investor's updating speed. Under ambiguity, the mean monthly log return is perceived to be equal to $\xi + v_t$ (for $\xi \in \Xi$), where ξ is a fixed parameter that can be learned, while the investor assumes that she will never fully learn (observe) v_t . The set of priors is consequently defined as:

$$L_k = \left\{ \ell(\cdot|\xi) : \ell^k(\text{high}|\xi) = \frac{1}{2} + \frac{\xi + v}{2\sigma/\sqrt{k}}, |v| < \bar{v} \right\}. \quad (89)$$

$\bar{v} > 0$ represents the presence of ambiguity, and each v_t parametrizes a different likelihood ℓ_t^k . Specifically, \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 $\bar{v} \rightarrow \infty$ no learning is anticipated. The posterior set for the realized sequence s^t is $M_{t,k}^a(s^t) = [\bar{R}_t - t^{-(1/2)}\sigma b_a, \bar{R}_t + t^{-(1/2)}\sigma b_a]$, where $b_a^2 = -2 \ln a$. From the definition of $M_{t,k}^a(s^t)$, it is immediate to compute the lowest conditional mean log return for $k \rightarrow \infty$, that is, $\bar{R}_t - \bar{v} - t^{-(1/2)}\sigma b_a$.

Epstein and Schneider (2007) prove 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 $k \rightarrow \infty$ is:

$$\lim_{k \rightarrow \infty} \omega_{t,k,1/k}^*(\bar{R}_t) = \max\{\omega_t^{Bayes} - \sigma^{-2}(\bar{v} + t^{-(1/2)}\sigma b_a), 0\} + \min\{\omega_t^{Bayes} + \sigma^{-2}(\bar{v} + t^{-(1/2)}\sigma b_a), 0\} \quad (90)$$

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 (90). 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, in the sense that 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 (90) reflects short selling when the equity premium is unambiguously negative. Interestingly, non-participation in the stock market—here, 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 (90) 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 $t \rightarrow \infty$. This result is clearly reminiscent of the early findings in Dow and Werlang (1992), in which states could exist in which an investor would simply abstain from trading.

The second important 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 be characterized by 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 nonparticipation, 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

It is often possible to “close” models of portfolio and consumption decisions to generate asset pricing models, i.e., to yield implications for the mappings between the state of the economy and equilibrium (no-arbitrage) asset prices. To review papers that have performed this operation is our goal in this Section. Our opening statement alerts a Reader that a portion of the models and results that we are about to review will overlap to frameworks and papers that we have covered in Section 3. In these cases we will try to economize on the details concerning the assumptions and solutions of each model and aim at immediately jumping to their asset pricing implications, counting on the good-will of a Reader to review the relevant material from Section 3.

4.1. Static Models

The simplest possible effort 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, simple algebra shows that—given some current price p for the risky asset—the AA investors’ demand for the risky asset is characterized by a non-participation interval, see (27). This outcome reduces the risk-sharing abilities of the market and it may consequently induce an increase in the equilibrium risk premium, since states exist in which the whole risk must now be born by the (S)EU-agents. Indeed, when the market is characterized by limited participation—which means that demand and supply are equated by a price that falls in the interval $[\mu_{\min}, \mu_{\max}]$ —such an equilibrium price is

$$p^{SEU} = \hat{\mu} - \frac{\hat{\sigma}^2}{1 - \alpha} x, \quad (91)$$

where the relevant parameters are simply the estimates by the (S)EU investors. Of course, this is the classical CARA/Gaussian result, with the only minor difference of the $(1 - \alpha)$ weighting. On the opposite, when both groups of investors are in the market, the equilibrium price is:

$$p^{AA} = \frac{\alpha \hat{\sigma}^2 \mu_{\min} + (1 - \alpha) \sigma_{\max}^2 \hat{\mu}}{(1 - \alpha) \sigma_{\max}^2 + \alpha \hat{\sigma}^2} - \frac{(\sigma_{\max}^2 + \hat{\sigma}^2)}{(1 - \alpha) \sigma_{\max}^2 + \alpha \hat{\sigma}^2} x. \quad (92)$$

Clearly, this is a weighted average (the weights are $(1 - \alpha) \sigma_{\max}^2$ and $\alpha \hat{\sigma}^2$) of expected payoffs minus a weighted average of perceived variances times the risky security's supply. 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 might shift from a nonparticipating to a participating one, as in (27) p^{SEU} may eventually move below the threshold μ_{\min} . Moreover, changes in the set of variances considered by AA investors have clearly no effect if the economy is a nonparticipating equilibrium. Viceversa, if the economy is in a participating equilibrium, increases in σ_{\max}^2 will reduce the price because the demand of AA investors shrinks when σ_{\max}^2 increases. If the market is in a limited participation equilibrium, increases in μ_{\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 μ_{\min} increases p^{AA} since ω^{AA} increases. Finally, an increase in μ_{\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, because AA agents attach great importance to worst case models.

Guidolin and Rinaldi (2009) have used their Easley and O'Hara (2009)-type model to solve for equilibrium prices when ambiguity may concern either the systematic risk or the idiosyncratic risks, or both. Also in their case, it is possible that the equilibrium price may fall in the region in which no AA agent expresses a non-zero demand for the risky asset, in which case the equilibrium price p^{SEU} is:

$$p^{SEU} = (\mu_I + \mu_S) - \frac{\sigma_I^2 + \sigma_S^2}{(1 - \alpha)} x. \quad (93)$$

This is a limited participation equilibrium. If both types of agents participate, the equilibrium price p^{AA} is:

$$p^{AA} = \frac{[\alpha \mu_{\min} (\sigma_I^2 + \sigma_S^2) + \mu_S (\alpha \sigma_I^2 + (1 - \alpha) \sigma_{\max}^2 + \sigma_S^2) + (1 - \alpha) \mu_I (\sigma_{\max}^2 + \sigma_S^2)]}{\sigma_S^2 + \alpha \sigma_I^2 + (1 - \alpha) \sigma_{\max}^2} - \frac{(\sigma_I^2 + \sigma_S^2) (\sigma_{\max}^2 + \sigma_S^2)}{\sigma_S^2 + \alpha \sigma_I^2 + (1 - \alpha) \sigma_{\max}^2} x. \quad (94)$$

This is a participation equilibrium. Interestingly, among the parameters that affect the (S)EU-only price, we also find α , the proportion of AA investors. Therefore, even when the AA agents are not trading, their mere existence will affect the equilibrium stock price. Moreover, when both types participate, both the single-prior ($\mu_I, \mu_S, \sigma_I^2, \sigma_S^2$) and the multiple-prior ($\mu_{\min}, \sigma_{\max}^2$) parameters enter the expression for p^{AA} .

Guidolin and Rinaldi (2009) have worked out expressions for (real) expected returns. However, in the presence of ambiguity, there are two notions of risk premium that can be considered. One is an objective notion and corresponds to an aggregate, market viewpoint that is based on true expectations for risky payoffs. This is also the risk premium anticipated by an external observer that understands the structure of the model and solves for the equilibrium. Naturally, this is the notion of risk premium relevant to an econometrician interested in understanding the data. The other is a subjective notion and corresponds to the expectation—obviously different across (S)EU 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 (S)EU and AA investors. Starting with the first, objective notion, Guidolin and Rinaldi show that real expected returns are

$$E^{SEU} [1 + R] = \frac{(1 - \alpha) (\mu_I + \mu_S)}{(1 - \alpha) (\mu_I + \mu_S) - x (\sigma_I^2 + \sigma_S^2)} \quad (95)$$

in the limited participation equilibrium and

$$E^{AA}[1+R] = \frac{(\mu_I + \mu_S)A}{[\alpha\mu_{\min}B + (1-\alpha)\mu_I C + \mu_S A] - xBC}, \quad (96)$$

where $A \equiv \alpha\sigma_I^2 + (1-\alpha)\sigma_{\max}^2 + \sigma_S^2$, $B \equiv \sigma_I^2 + \sigma_S^2$, and $C \equiv \sigma_{\max}^2 + \sigma_S^2$ in the equilibrium in which both types of investors trade. Required returns on the risky asset are increasing in the proportion α of AA agents because, given total supply, a smaller fraction of (S)EU-only investors will have to absorb the entire supply. (95) and (96) 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 to be drawn. However, these are different from the risk premia which are perceived by both categories of investors. The (S)EU 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 (S)EU-maximizers will be surprised by the level of the average realized real excess returns. In particular, (S)EU agents—in this sense they are akin to most in the economics profession—will systematically find that the realized real excess returns are “too high” in the light of the underlying economic environment, as they perceive it. Of course, this is due to the existence of an additional source of uncertainty—ambiguity, indeed—that is priced in equilibrium and of which the (S)EU investors are not aware, while AA investors are. The AA investors perceive instead a real expected return of:

$$\tilde{E}^{AA}[1+R] = \frac{(1-\alpha)(\mu_{\min} + \mu_S)}{(1-\alpha)(\mu_{\min} + \mu_S) - x(\sigma_{\max}^2 + \sigma_S^2)}. \quad (97)$$

Guidolin and Rinaldi show that

$$\tilde{E}^{AA}[1+R] > E^{AA}[1+R] > E^{SEU}[1+R] = \tilde{E}^{SEU}[1+R], \quad (98)$$

i.e., AA agents demand a risk premium which is higher than what (S)EU 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 (S)EU investors believe that the risky asset is under-priced, and thus that it offers abnormally high excess returns and ambiguity-averse traders believe that the asset is over-priced, and thus that it offers abnormally low excess returns in equilibrium, both types of traders will hold portfolios which differ from the market portfolio in their attempt to take advantage of the perceived mispricing. Finally, because Guidolin and Rinaldi (2009) argue that the ability of ambiguity aversion to induce participation equilibria depends on the fact that the spread between the highest and the lowest possible realization of idiosyncratic risk is *larger* than the spread between the highest and the lowest possible realization of systematic risk, the high risk premia may be more the result of ambiguity concerning idiosyncratic payoffs than ambiguity on the systematic ones.

Also Cao, Wang, and Zhang (2005) investigate the role of model uncertainty in determining limited participation and high equity premia. Their contribution consists of the adoption of a set up for ambiguity that gives an explicit role to a measure of *uncertainty dispersion*, i.e., the heterogeneity in the perceived ambiguity across market participants. Their setting is a heterogeneous agent economy, in which 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 risk aversion coefficient α . As in Kogan and Wang (2003), the set of priors φ considered by each agent can be characterized in terms of confidence regions: specifically, each agent considers as 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 the agents. Each investor solves the simple mean-variance problem

$$\max_{\omega_i} \min_{\theta: \theta^2 \leq \eta_i^2} \omega_i (\hat{\mu} - \theta - p) - \frac{\alpha \sigma^2}{2} \omega_i^2, \quad (99)$$

with typical, step-wise solution:

$$\omega_i^* = \begin{cases} \frac{1}{\alpha \sigma^2} (\hat{\mu} - \sigma \eta_i - p) & \text{if } \hat{\mu} - \sigma \eta_i > p \\ 0 & \text{if } \hat{\mu} - \sigma \eta_i \leq p \leq \hat{\mu} + \sigma \eta_i \\ \frac{1}{\alpha \sigma^2} (\hat{\mu} + \sigma \eta_i - p) & \text{if } \hat{\mu} + \sigma \eta_i < p \end{cases} . \quad (100)$$

Also in this case, for prices in the interval $[\hat{\mu} - \sigma \eta_i, \hat{\mu} + \sigma \eta_i]$ agent i will not trade. In case of full participation, the equilibrium asset pricing equation that derives from the usual clearing condition is:

$$\hat{\mu} - p = \alpha x \sigma^2 + \sigma \bar{\eta}, \quad (101)$$

where $\hat{\mu} - p$ may be interpreted as a risk premium and $\bar{\eta}$ is the average level of uncertainty perceived in the market. The term $\alpha x \sigma^2$ is the standard risk premium, while the second term $\sigma \bar{\eta}$ is the uncertainty premium. Instead, when some agents optimally decide to avoid trading, the pricing equation becomes:

$$\hat{\mu} - p = 2\sigma \sqrt{\alpha x \delta \sigma} + (\bar{\eta} \sigma - \delta) . \quad (102)$$

Cao et al. (2005) 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) is another mean-variance style paper that can be “closed” to generate interesting asset pricing implications. In the setting described in Section 3.2, there are N agents each of which is familiar with only one particular asset and (equally) unfamiliar with respect to the others, in particular agent n is familiar with asset n only, and her degree of familiarity with respect to asset i is η_i^n . Investors are assumed to be symmetric in their attitude towards ambiguity, so that for any two agents m and n , $\eta_n^n = \eta_m^m = \eta_F$ and $\eta_{i \neq n}^n = \eta_{i \neq m}^m = \eta_U$. If $\hat{\mu}/\sigma_{\hat{\mu}} > \sqrt{\eta_F} + \rho(1 - \rho)^{-1}(\sqrt{\eta_U} - \sqrt{\eta_F})$, ambiguity is relatively low with respect to the reward-to-risk ratio of each asset and each individual holds both the familiar asset and the portfolio made by all the unfamiliar ones. In this case of low ambiguity (relative to the reward-to-risk ratio), the equilibrium excess return of each risky asset is given by:

$$\hat{\mu} = \frac{1}{N} [\gamma \rho \sigma^2 (N - 1) + \sigma_{\hat{\mu}} \sqrt{\eta_U} (N - 1) + \sigma_{\hat{\mu}} \sqrt{\eta_F}] . \quad (103)$$

Hence, the equilibrium risk premium has three components: the conventional risk premium, $N^{-1}(N - 1)\gamma \rho \sigma^2$ which tends to $\gamma \sigma^2$ as $N \rightarrow \infty$ (here γ is the coefficient of risk aversion); a second term which is due to the concentration of

the portfolio in the familiar asset, the ambiguity premium relative to the unfamiliar assets ($N^{-1}(N-1)\sigma_{\hat{\mu}}\sqrt{\eta_U}$); an ambiguity premium relative to the familiar asset ($N^{-1}\sigma_{\hat{\mu}}\sqrt{\eta_F}$). When N becomes arbitrarily large, each agent holds all the assets only if $\sqrt{\eta_U} < \sqrt{\eta_F} + \gamma\sigma^2/\sigma_{\hat{\mu}}$ and the excess return of each risky asset is

$$\lim_{N \rightarrow \infty} \hat{\mu} = \gamma\rho\sigma^2 + \sigma_{\hat{\mu}}\sqrt{\eta_U}. \quad (104)$$

Therefore 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.

If $\sqrt{\eta_F} < \hat{\mu}/\sigma_{\hat{\mu}} < \sqrt{\eta_U} + \rho(1-\rho)^{-1}(\sqrt{\eta_U} - \sqrt{\eta_F})$ each agent considers unfamiliar assets as extremely ambiguous and therefore she holds only the familiar one. Hence, the equilibrium excess return does not reflect any ambiguity premium related to the unfamiliar assets, which implies $\hat{\mu} = \gamma\sigma^2 + \sigma_{\hat{\mu}}\sqrt{\eta_F}$, even for finite N . Finally, under the condition $0 < \hat{\mu}/\sigma_{\hat{\mu}} < \sqrt{\eta_F}$, the ambiguity level is so high for both familiar and unfamiliar assets that nobody is willing to trade, there are no equilibrium prices as the market collapses.

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. Furthermore, Maccheroni et al. (2009) have shown that MMVP representation is the minimal (and so the most cautious) monotone extension of the MVP functional and its best possible monotone approximation. Denoting by \mathbf{R} the N -dimensional vector of risky returns, and by R_f the gross return of a risk-free bond, the optimal portfolio of a MVP-agent with risk aversion γ is:

$$\boldsymbol{\omega}_{MVP}^* = \frac{1}{\gamma} \{Cov[\mathbf{R}]\}^{-1} (E[\mathbf{R} - R_f \boldsymbol{\iota}_N]). \quad (105)$$

For a MMVP-investor with uncertainty aversion ϱ , the optimal allocation rule is instead⁴⁸

$$\boldsymbol{\omega}_{MMVP}^* = \frac{1}{\varrho \Pr(W(\boldsymbol{\omega}_{MMVP}^*) \leq \psi^*)} \{Cov[\mathbf{R}]\}^{-1} E[\mathbf{R} - R_f \boldsymbol{\iota}_N | W(\boldsymbol{\omega}_{MMVP}^*) \leq \psi^*], \quad (106)$$

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 the non-linear system:

$$\begin{cases} \varrho \Pr(W(\boldsymbol{\omega}) \leq \psi) Cov[\mathbf{R}|W(\boldsymbol{\omega}) \leq \psi^*] \boldsymbol{\omega} = E[\mathbf{R} - R_f \boldsymbol{\iota}_N | W(\boldsymbol{\omega}) \leq \psi] \\ E[-\min(0, W(\boldsymbol{\omega}) - \psi)] = 1/\varrho \end{cases}. \quad (107)$$

While the first equation has a purely technical meaning, the second guarantees that the optimal choice of the portfolio will deliver a wealth process constrained to the monotonicity domain. 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

⁴⁸Over the domain of monotonicity of MVP, γ and ϱ simply coincide. However, whenever MMVP differ from standard MVP, it is better to distinguish them and to label ϱ as an uncertainty-aversion parameter.

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 $|\omega_{MMVP}^*| \geq |\omega_{MVP}^*|$ 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.⁴⁹

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

$$E[R_i] - R_f = \beta_i (E[R_m] - R_f) \quad \beta_i = \frac{Cov[R_i, m]}{Cov[R_m, m]}, \quad (108)$$

where R_m is the return of the market portfolio and m is the stochastic discount factor given by

$$m = -\varrho_m \min [0, R_m - \psi_m], \quad (109)$$

where the pair (ϱ_m, ψ_m) is a solution to a nonlinear system involving market quantities only. The monotone CAPM satisfies the standard two-fund separation property, since the optimal portfolios of risky assets 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 (because of monotonicity of MMVP), and it has the same empirical tractability of the standard CAPM since the variables R_m, R_i, R_f, M can all be fully computed from market data, while the betas can be estimated.

4.2. Dynamic Asset Pricing Models

Although rich insights have often been gained from simple, two-period models such as those reviewed so far, a number of researchers have strived to generalize truly intertemporal asset pricing models under ambiguity aversion. In particular, the standard Lucas's (1978) endowment economy has attracted intense efforts. One of the seminal contributions is Epstein and Wang (1994), in which MPP are consistently 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 about the next period state. The recursive specification of utility under MPP preferences in an infinite horizon framework is

$$V_t(c, s^t) = u(c_t(s^t)) + \beta \inf_{m \in \wp(s_t)} \int V_{t+1}(c, s^t) dm, \quad (110)$$

where s^t denotes history up to time t , i.e. $s^t = \{s_0, s_1, \dots, s_t\}$. 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 0-net supply and has dividend and price processes $d_i = \{d_{it}\}$ and $p_i = \{p_{it}\}$, respectively.

⁴⁹The notation $|\omega_{MMVP}^*| \geq |\omega_{MVP}^*|$ is meant to indicate that each component of the MMVP-portfolio is greater in absolute value than the corresponding component in the MVP-portfolio.

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 in any direction $\mathbf{h} \equiv [h_1 \ h_2 \ 0 \ 0 \ \dots]'$ are defined as:

$$\begin{aligned}\frac{d}{d\xi} V_t(c + \xi\mathbf{h}, s^t) \Big|_{0+} &= u'(c_t(s^t)) h_1 + \beta \min_{m \in \wp(s_t)} \int u'(c_{t+1}(s^{t+1})) h_2 dm \\ \frac{d}{d\xi} V_t(c + \xi\mathbf{h}, s^t) \Big|_{0-} &= u'(c_t(s^t)) h_1 + \beta \max_{m \in \wp(s_t)} \int u'(c_{t+1}(s^{t+1})) h_2 dm,\end{aligned}\quad (111)$$

Assuming that $\{\mathbf{p}_t\}$ is an equilibrium, any variation from the optimal policy $\{e_t, \mathbf{0}\}$ should leave the agent worse off. In particular, this consideration leads to the following necessary condition for an equilibrium price process \mathbf{p} :

$$\beta \max_{m \in \wp(s_t)} \int \frac{u'(e_{t+1})}{u'(e_t)} (\mathbf{p}_{t+1} + \mathbf{d}_{t+1}) dm \geq \mathbf{p}_t \geq \beta \min_{m \in \wp(s_t)} \int \frac{u'(e_{t+1})}{u'(e_t)} (\mathbf{p}_{t+1} + \mathbf{d}_{t+1}) dm. \quad (112)$$

Epstein and Wang prove that this double inequality is also sufficient for an equilibrium. Therefore, under recursive MPP the optimality conditions take the form of inequalities, just because $V(\cdot; s)$ is generally non-differentiable, unless the correspondence $\wp : \Omega \rightarrow \wp(\Omega)$ is a probability kernel (unless there is no ambiguity).⁵⁰ 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 :

$$\frac{d}{d\xi} V_t(e + \xi\mathbf{p}) \Big|_{0+} \neq \frac{d}{d\xi} V_t(e + \xi\mathbf{p}) \Big|_{0-}. \quad (113)$$

Intuitively, we would expect a link between indeterminacy of assets prices and intertemporal 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, and yet providing an avenue to approach the excessive volatility puzzle alternative to earlier attempts based on either un-modelled time-varying risk premia or learning (see e.g., Grossman and Shiller, 1982, or Timmermann, 1993).^{51,52}

Chen and Epstein (2002) have further generalized 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

⁵⁰In the standard (S)EU case, $\wp(s_t)$ is made of a unique probability distribution, say μ , so that $\max_{m \in \wp(s_t)} \int \frac{u'(e_{t+1})}{u'(e_t)} (\mathbf{p}_{t+1} + \mathbf{d}_{t+1}) dm = \min_{m \in \wp(s_t)} \int \frac{u'(e_{t+1})}{u'(e_t)} (\mathbf{p}_{t+1} + \mathbf{d}_{t+1}) dm = \int \frac{u'(e_{t+1})}{u'(e_t)} (\mathbf{p}_{t+1} + \mathbf{d}_{t+1}) d\mu$. Therefore the inequalities in the main text reduce to the standard vector of Euler equations, $\mathbf{p}_t = \beta \int \frac{u'(e_{t+1})}{u'(e_t)} (\mathbf{p}_{t+1} + \mathbf{d}_{t+1}) d\mu$.

⁵¹Epstein and Wang (1995) have extended these results and demonstrated a connection between ambiguity and markets’ extreme phenomena, such as crashes and booms. Indeed, 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.

⁵²Dana (2003) has analyzed the effects of ambiguity on equilibria determinacy. In the standard (S)EU-framework, agents’ consumption rules depend only on the aggregate endowment and are increasing functions of it; hence, the equilibrium is unique and so is the stochastic discount factor. Similarly, in a CEU/ambiguity framework, if there are no states of the world that yield the same aggregate endowment and agents have the same capacity assignments, then they also share a unique effective probability over the reference state space, and hence equilibria are determinate, so that the stochastic discount factor is unique. However, if non-aggregate (idiosyncratic) uncertainty exists and if the intersection of agents’ set of priors is not empty, then a multiplicity of equilibria may obtain due to idiosyncratic risk.

time framework has also allowed them to obtain a few special cases of great interest. As we have seen in Section 2.4, under recursive MPP, the utility functional V is the solution to the differential equation

$$dV_t = \left[-\psi(c_t, V_t) + \max_{\theta \in \Theta} \theta'_t \sigma_t^V \right] dt + \sigma_t^V d\mathbf{B}_t \quad (114)$$

where $V_T = 0$, \mathbf{B} is a d -dimensional BM, and σ_t^V is a $d \times d$ covariance matrix. In Chen and Epstein's application, the aggregator function $\psi(c, v)$ is taken to be $\psi(c, v) \equiv [c^b - \beta(av)^{b/a}]/b(av)^{(b-a)/a}$, for $\beta \geq 0$, 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 d risky securities, one for each component of the BM \mathbf{B}_t . The returns \mathbf{R}_t of the risky securities are described by the Ito process:

$$d\mathbf{R}_t = \boldsymbol{\mu}_t dt + \boldsymbol{\sigma}_t d\mathbf{B}_t. \quad (115)$$

The set of density generators is derived as follows. Fix a vector of parameters κ in \mathbb{R}^d and let $\Theta_t(\cdot) = \{\mathbf{y} \in \mathbb{R}^d : |y_i| \leq \kappa_i \text{ for all } i\}$, 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 ,

$$d\pi_t = \pi_t r_t^f dt + \pi_t \boldsymbol{\eta}_t d\mathbf{B}_t \quad \boldsymbol{\eta}_t = \boldsymbol{\sigma}_t^{-1} (\boldsymbol{\mu}_t - r_t^f \boldsymbol{\iota}_d) \quad (116)$$

where $\boldsymbol{\eta}_t$ is the $d \times 1$ vector of market prices of the d sources of uncertainty. The wealth dynamics is described by:

$$dW_t = ([r_t^f + \boldsymbol{\omega}'_t (\boldsymbol{\mu}_t - r_t^f \boldsymbol{\iota}_d)] W_t - c_t) dt + W_t \boldsymbol{\omega}'_t \boldsymbol{\sigma}_t d\mathbf{B}_t. \quad (117)$$

Letting the optimal consumption c_t be an Ito process with parameters (μ_t^c, σ_t^c) , and rewriting 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 becomes:

$$\boldsymbol{\eta}_t = b^{-1} [a(1-b)\boldsymbol{\sigma}_t^c + (b-a)\boldsymbol{\sigma}_t^M] + \boldsymbol{\theta}_t^*. \quad (118)$$

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 (114). Next, assuming that ambiguity aversion is small in the sense that $0 \leq \kappa_i < |\eta_t^i|$ for all i , then

$$\begin{aligned} \sigma_t^{c,i} &> (<) 0 \text{ if } \eta_t^i > (<) 0 \\ \boldsymbol{\omega}_t^* &= (1-a)^{-1} (\boldsymbol{\sigma}_t^T)^{-1} (\boldsymbol{\eta}_t - \boldsymbol{\kappa} \otimes \text{sign}(\boldsymbol{\eta}_t)) = (1-a)^{-1} (\boldsymbol{\sigma}_t^T)^{-1} ((\boldsymbol{\mu}_t - r_t^f \boldsymbol{\iota}_d) - \boldsymbol{\kappa} \otimes \text{sign}(\boldsymbol{\eta}_t)) \end{aligned} \quad (119)$$

This result implies that optimal portfolio weights differ from the standard mean-variance expression,

$$\boldsymbol{\omega}_t^{SEU} = (1-a)^{-1} (\boldsymbol{\sigma}_t^T)^{-1} (\boldsymbol{\mu}_t - r_t^f \boldsymbol{\iota}_d), \quad (120)$$

because of the correction implied by the term $\boldsymbol{\kappa} \otimes \text{sign}(\boldsymbol{\eta}_t)$, where

$$\boldsymbol{\kappa} \otimes \text{sign}(\boldsymbol{\eta}_t) = \begin{cases} \boldsymbol{\kappa}_i \frac{|\eta_t^i|}{\eta_t^i} & \text{if } \eta_t^i \neq 0 \\ \mathbf{0} & \text{otherwise} \end{cases}. \quad (121)$$

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 κ .

Chen and Epstein show that the equilibrium risk premia are:

$$\mu_t - r_t^f \iota_d = \sigma_t \eta_t = b^{-1}[a(1-b)\sigma_t \sigma_t^c + (b-a)\sigma_t \sigma_t^M] + \sigma_t \theta_t^*. \quad (122)$$

The right-hand side can be decomposed as the sum of a standard SEU risk premium component, $b^{-1}[a(1-b)\sigma_t \sigma_t^h + (b-a)\sigma_t \sigma_t^M]$, and of an ambiguity premium, $\sigma_t \theta_t^*$, such that

$$\sigma_t^i \theta_t^* = -\text{Cov}_t[dR_t^i, dz_t^{\theta^*}/z_t^{\theta^*}] \quad i = 1, 2, \dots, n, \quad (123)$$

where σ_t^i is the i -th row of the matrix σ_t . Therefore, the premium is positive if the asset's return has negative covariation 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 GBM with constant drift and diffusion vector, $dc_t = \mu^c dt + \sigma^c dB_t$, 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 $\eta_t = (1-a)\sigma_t^c + \kappa \otimes \text{sign}(\sigma_t^c)$. 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:

$$\mu_t - r_t^f \iota_d = (1-a)\sigma_t^i \sigma_t^c + \kappa(\sigma_t^i \otimes \text{sign}(\sigma_t^c)). \quad (124)$$

The ambiguity premium (represented by the second term) for asset i is large if $\sigma_t^{i,j} \otimes \text{sign}(\sigma_t^{c,j})$ is large, and it is positive for components j of the driving Brownian process B_t that are very ambiguous, in the sense of having large κ_j s. Since the equilibrium risk premia depend on the endowment process only via the signs of the covariances $\sigma_t^{c,j}$, $j = 1, \dots, d$, large ambiguity premia can occur even if consumption is relatively smooth. In particular, the excess return on the market portfolio (i.e., a portfolio exactly mimicking the consumption process itself, a tradable portfolio in complete markets) is given by:⁵³

$$\mu_t^M - r_t^f = (1-a)(\sigma_t^c)' \sigma_t^c + \kappa' |\sigma_t^c|. \quad (125)$$

The equity premium can be decomposed as the sum of the standard risk premium term $(1-a)(\sigma_t^c)' \sigma_t^c$ and of the ambiguity premium $\kappa' |\sigma_t^c|$. 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

$$r_t^f - \beta = (1-b) \left[\mu_t^c - \frac{(1-a)(2-b)}{2(1-b)} (\sigma_t^c)' \sigma_t^c - \kappa' |\sigma_t^c| \right], \quad (126)$$

and it will be decreasing in both ambiguity and risk aversion. Therefore, 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 separate contribution of the ambiguity premium component, $\kappa' |\sigma_t^c|$ (or more generally, $\sigma_t \theta_t^*$).

⁵³Denoting by $[\sigma_t^c]_{ij}$ and $[\sigma_t^c]_i$ the ij-component of $|\sigma_t^c|$ and σ_t^c respectively, by definition:

$$[\sigma_t^c]_{ij} = \begin{cases} [\sigma_t^c]_{ij} & \text{if } [\sigma_t^c]_{ij} \geq 0 \\ -[\sigma_t^c]_{ij} & \text{if } [\sigma_t^c]_{ij} < 0 \end{cases}$$

Epstein and Miao (2003) have used the recursive 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. 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 , respectively. Time varies over $[0, T]$. The cumulative dividends of the three securities are given by:

$$\mathbf{D}'_t = \begin{cases} \left[0 \int_0^t Y_s^1 ds \int_0^t Y_s^2 ds \right] & 0 \leq t < T \\ \left[1 \int_0^T Y_s^1 ds \int_0^T Y_s^2 ds \right] & t = T \end{cases}. \quad (127)$$

$B^1 = \{B_t^1\}$ and $B^2 = \{B_t^2\}$ are two BMs and \mathcal{F}_t is a filtration generated by them. The securities' prices are represented at each date by the Ito process $S_t = (S_t^0, S_t^1, S_t^2)'$, so that

$$d(\mathbf{S}_t + \mathbf{D}_t) = \boldsymbol{\mu}_t^G dt + \boldsymbol{\sigma}_t^G d\mathbf{B}_t, \quad (128)$$

is an Ito Process where $\mathbf{B}_t \equiv [B_t^1 \ B_t^2]'$. The riskless rate satisfies $r_t^f dt = dS_t^0/S_t^0$, while the return process for the ambiguous stocks is $d\mathbf{R}_t = \mathbf{b}_t dt + \boldsymbol{\sigma}_t d\mathbf{B}_t$ (with $\mathbf{R}_t \equiv [R_t^1, R_t^2]'$), where \mathbf{b}_t is the vector of expected stock returns, a quantity to be determined in equilibrium. The variable $\boldsymbol{\eta}_t = \boldsymbol{\sigma}_t^{-1}(\mathbf{b}_t - r_t^f \boldsymbol{\nu}_2)$ can be interpreted as the market price for uncertainty, including risk, and it is possible to define a state price process:

$$-d\pi_t/\pi_t = r_t^f dt + \boldsymbol{\eta}_t \cdot d\mathbf{B}_t, \quad \pi_0 = 1. \quad (129)$$

The aggregate endowment $Y_t \equiv Y_t^1 + Y_t^2$ is assumed to follow the GBM

$$dY/Y_t = \mu^Y dt + (\boldsymbol{\sigma}^Y)' d\mathbf{B}_t \quad \boldsymbol{\sigma}^Y \equiv [\boldsymbol{\sigma}_1^Y \ \boldsymbol{\sigma}_2^Y]', \quad (130)$$

and to not exhaust aggregate output, in the sense that $Y_t = Y_t^1 + Y_t^2 + \Phi_t$, where the non-traded quantity Φ_t is assumed to be equally divided between the two countries. Country-specific portfolios at any date are denoted by $\boldsymbol{\omega}_t^1$ and $\boldsymbol{\omega}_t^2$, respectively, while the initial holdings are assumed to be $\boldsymbol{\omega}_0^1 = [0 \ \boldsymbol{\omega}_{1,0}^1 \ \boldsymbol{\omega}_{2,0}^1]'$ and $\boldsymbol{\omega}_0^2 = [0 \ \boldsymbol{\omega}_{1,0}^2 \ \boldsymbol{\omega}_{2,0}^2]'$. To characterize the set of effective priors considered by each agent, while allowing for heterogeneity in ambiguity aversion, consider the vectors $\boldsymbol{\kappa}^1 = [0 \ \boldsymbol{\kappa}_1]'$ and $\boldsymbol{\kappa}^2 = [\boldsymbol{\kappa}_2 \ 0]'$, and for each agent define the set of density generators

$$\Theta^i = \left\{ \boldsymbol{\theta} = (\boldsymbol{\theta}_t): \sup_t |\theta_t^j| \leq \kappa_j^i \quad j = 1, 2 \right\} \quad i = 1, 2. \quad (131)$$

Θ^1 and Θ^2 determine the set of effective distributions considered by agent 1 and agent 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 $\boldsymbol{\kappa}^1$ ($\boldsymbol{\kappa}^2$).

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 the constraint

$$E \left[\int_0^T \pi_s \left(c_s^i - \frac{1}{2} \Phi_s \right) ds \right] \leq \boldsymbol{\omega}_{i,0}^i E \left[\int_0^T \pi_s Y_s^i ds \right] + \boldsymbol{\omega}_{j,0}^i E \left[\int_0^T \pi_s Y_s^j ds \right] \quad i = 1, 2 \quad i \neq j = 1, 2 \quad (132)$$

and the markets for contingent consumption clear, $c^1 + c^2 = Y$.⁵⁴ To characterize the equilibrium it is useful to define the shadow price of the non-traded endowment $\bar{S}_t = \pi_t^{-1} E \left[\int_t^T \pi_s \Phi_s ds \mid \mathcal{F}_t \right]$ and the total wealth process for each country $\bar{X}_t^i = S_t \omega_t^i + \frac{1}{2} \bar{S}_t$. The dynamics of the shadow price of the non-traded endowment can be characterized as

$$d\bar{S}_t = \bar{\mu}_t dt + \bar{\sigma}'_t d\mathbf{B}_t$$

for some $\bar{\mu}_t$, $\bar{\sigma}_t$. The equilibrium is given by the quantities:

$$c_t^1 = \frac{1}{1 + \lambda \varsigma_t} Y_t \quad c_t^2 = \frac{\lambda \varsigma_t}{1 + \lambda \varsigma_t} Y_t \quad \pi_t = \frac{c_0^1 e^{-\beta t} z_t^{\theta^{*1}}}{c_t^1} \quad (133)$$

where

$$\varsigma_t = \left[\frac{dq^{\theta^{*2}}/dp}{dq^{\theta^{*2}}/dp} \Big| \mathcal{F}_t \right] = \exp \left\{ \frac{1}{2} \left((\kappa_1)^2 - (\kappa_2)^2 \right) t + \kappa_1 B_t^2 - \kappa_2 B_t^1 \right\} \quad (134)$$

measures the difference in ambiguity-adjusted beliefs and hence can be called the “disagreement process”, while

$$\lambda = \frac{E \left[\int_0^T e^{-\beta t} z_t^{\theta^{*1}} (1 - \omega_{1,0}^1 Y_t^1 / Y_t - \omega_{2,0}^1 Y_t^2 / Y_t - \frac{1}{2} \Phi_t / Y_t) dt \right]}{E \left[\int_0^T e^{-\beta t} z_t^{\theta^{*2}} (\omega_{1,0}^1 Y_t^1 / Y_t + \omega_{2,0}^1 Y_t^2 / Y_t + \frac{1}{2} \Phi_t / Y_t) \right]} \quad (135)$$

measures the distribution of initial endowments in favor of country 2. As usual, $z_t^{\theta^{*1}}$ denotes the Radon–Nikodym derivative of the alternative measure with respect to the original one induced by the generator θ^{*1} . Defining the equilibrium prices of the two risky securities in the standard way (i.e., using no-arbitrage intertemporal conditions),

$$S_t^i = \frac{1}{\pi_t} E \left[\int_t^T \pi_s Y_s^i ds \mid \mathcal{F}_t \right] \quad (136)$$

and the equilibrium price for the bond as $S_t^0 = (1/\pi_t) E [\pi_T \mid \mathcal{F}_t]$, the Arrow-Debreu equilibrium can be implemented by a Radner equilibrium (i.e., a sequence of temporary equilibria) and it is identical to the one of an heterogenous ambiguity-free economy in which countries’ beliefs are represented by the two priors q^1 and q^2 , such that:

$$\frac{dq^i}{d\pi} = \exp \left\{ -\frac{1}{2} (\kappa_i)^2 T - \kappa_i B_T^j \right\} \quad i = 1, 2 \quad i \neq j = 1, 2. \quad (137)$$

The expression for the asset $i = 1, 2$ risk premium

$$\begin{aligned} b_t^i - r_t^f &= (\boldsymbol{\sigma}_t^i)' \boldsymbol{\sigma}^Y + \left(\kappa_j \frac{c_t^j}{Y_t} \sigma_t^{ii} + \kappa_i \frac{c_t^i}{Y_t} \sigma_t^{ij} \right) \quad i = 1, 2 \quad i \neq j = 1, 2 \\ r_t^f &= \beta + \mu^Y + (\boldsymbol{\sigma}^Y)' \boldsymbol{\sigma}^Y - [\kappa_2 \sigma_1^Y - \kappa_1 \sigma_2^Y] \quad \boldsymbol{\eta}_t = \boldsymbol{\sigma}^Y + \begin{bmatrix} \kappa_2 c_t^2 / Y_t \\ \kappa_1 c_t^1 / Y_t \end{bmatrix} \end{aligned} \quad (138)$$

shows that—besides the classical covariance (with the endowment process) term, $(\boldsymbol{\sigma}_t^i)' \boldsymbol{\sigma}^Y$ —a contribution, as measured by

$$\kappa_j \frac{c_t^j}{Y_t} \sigma_t^{ii} + \kappa_i \frac{c_t^i}{Y_t} \sigma_t^{ij}, \quad (139)$$

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

⁵⁴The (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 $\boldsymbol{\theta}_t^{*2} = (\kappa_2, 0)'$.

the general framework by Chen and Epstein (2002), $b_t^i - r_t^f$ has neither to shrink towards zero when $(\sigma_t^i)' \sigma^Y \rightarrow 0$ 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 κ s) is likely (or guaranteed, when $\sigma_t^{12} > 0$) to increase the equity risk premia, the factor $-\left[\kappa_2' \sigma_1^Y - \kappa_1' \sigma_2^Y\right]$ is also likely to convert such growing ambiguity in lower real interest rates, which is a key to solving the equity premium and risk-free rate puzzles.

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 , a number of interesting insights can be derived. For instance, it is well known that in standard (S)EU 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 $Cov_t(d\bar{c}_t^i/c_t^i - dY_t/Y_t, dB_t^i) > 0$, meaning that the country-specific growth rate is positively correlated with internal 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 . In general, ambiguity also reduces the riskless rate; only when the two countries face the same amount of ambiguity, r_t^f is constant and independent on the initial distribution of aggregate endowment. Ambiguity increases the market price for uncertainty, and, interpreting the i -th component of η_t as the domestic price of uncertainty for country i , it can be shown that η_t^i is decreasing in c_t^i/Y_t (the domestic share of consumption) in support to the empirical findings in Campbell (1998) that risk premia are higher in countries with higher saving rates. Finally, rewriting the return of asset 1 in terms of the Brownian driving processes that are appropriate for country 1 and 2, respectively, one gets:

$$\begin{aligned} dR_t^1 &= (b_t^1 - \kappa_1 \sigma_t^{12}) dt + \sigma_t^{11} dB_t^1 + \sigma_t^{12} d(B_t^2 + \kappa_1 t) \quad [\text{country 1 perspective}] \\ dR_t^1 &= (b_t^1 - \kappa_2 \sigma_t^{11}) dt + \sigma_t^{12} dB_t^2 + \sigma_t^{11} d(B_t^1 + \kappa_2 t) \quad [\text{country 2 perspective}]. \end{aligned} \quad (140)$$

Therefore, country 1 will impute a higher expected return to security 1 with respect to country 2 if and only if $\kappa_2 \sigma_t^{11} > \kappa_1 \sigma_t^{12}$, which holds only if country 1 considers its own security less ambiguous than country 2 is (the opposite holds for security 2). Hence country 1, which is less ambiguity averse with respect to its own security, attaches to it a higher expected return reflecting a “preference” for it over the foreign one, in accordance with the empirical findings (e.g., by French and Poterba, 1991) that show that large distortions in expected return beliefs are needed to support the observed home country bias.

Recently, 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 constant relative risk aversion (CRRA) representative agent, who can invest in a single, risky security. Although their analysis is performed at a level of generality that is inferior to Epstein and Wang (1994) and Chen and Epstein (2002), their findings are insightful. Assume that—according to some probabilistic reference model p —the security's dividend process $\{D_t\}$ follows the simple GBM

$$dD_t = D\mu dt + D\sigma_D dB_D, \quad (141)$$

where $\mu \in \{\mu_1, \mu_2, \dots, \mu_n\}$, and B_D is a p -BM. To allow for the possibility of model misspecification, for each μ in the

parameter set, the investor entertains a range of alternative specifications for the dynamics of D_t . In particular, for each $\mu \in \{\mu_1, \dots, \mu_n\}$, she considers the drift-distorted model:

$$dD(t) = D(t)(\mu + \sigma_D \theta) dt + D(t)\sigma_D dB_D(t). \quad (142)$$

Each distortion θ generates a new probability measure q^θ , under which the signal is assumed to be unbiased. For any μ in the parameter set, not all the distortions are possible: the agent constrains herself to consider only distortions for which $\frac{1}{2}\theta^2 \leq \eta(\mu_i)$, for some $\eta(\mu_i)$. In particular, for each μ , she considers the drift set:

$$\wp(\mu) = \left\{ \mu + \sigma_D \theta : \frac{1}{2}\theta^2 \leq \eta(\mu_i) \right\} \quad \forall \mu \in \{\mu_1, \dots, \mu_n\}. \quad (143)$$

The size of each neighborhood $\wp(\mu)$ describes the degree of ambiguity associated with any possible reference model for the drift μ , in the sense that, the broader is $\wp(\mu)$, the more ambiguous the signals about a specific dividend drift $\mu + \sigma_D \theta$ are. The investor considers some probability distribution $\pi(t) = \{\pi_1, \dots, \pi_n\}$ on the space of sets $\{\wp(\mu_1), \dots, \wp(\mu_n)\}$. Since the investor is basing her beliefs on the whole set of likelihoods implied by $\wp(\mu)$, she considers a whole class of dividend-drift prediction processes, given by:

$$\left\{ \sum_{i=1}^n (\mu_i + \theta \sigma_D) \pi_i(t) : \mu + \sigma_D \theta \in \wp(\mu), \mu \in \{\mu_1, \dots, \mu_n\} \right\}. \quad (144)$$

This set represents the investor's ambiguity about the true dividend-drift process at time t , conditionally on the available information \mathcal{F}_t generated by dividends and signals. The representative investor's program may be written as

$$\max_{c, \phi} \inf_{\theta} E \left[\int_0^\infty e^{-\delta s} U(c, s) ds \middle| \mathcal{F}_0 \right] \quad (145)$$

subject to the distorted-dividend and the standard wealth dynamics,

$$dW = W_t \left[\omega \frac{dS + Ddt}{S} + (1 - \omega) r^f dt \right] - cdt. \quad (146)$$

If γ is the risk aversion coefficient, and setting $\mu_i^* \equiv \delta + (\gamma - 1)\mu_i + 0.5\gamma(1 - \gamma)\mu_i\sigma_D^2$, Leippold et al. show that the solution of the infimum problem in (145) is $\theta^* = \sqrt{2\eta(\mu_i)}$, so that the equilibrium price and consumption choices are:

$$S = D \sum_{i=1}^n \pi_i \left[\mu_i + (1 - \gamma) \sqrt{2\eta(\mu_i)\sigma_D^2} \right]^{-1} \quad r^f = \delta - \frac{1}{2}\gamma(1 + \gamma)\sigma_D^2 + \gamma \left[\sum_{i=1}^n (\mu_i + \theta^* \sigma_D) \pi_i(t) \right] \quad (147)$$

Because the factors $[\mu_i + (1 - \gamma) \sqrt{2\eta(\mu_i)\sigma_D^2}]^{-1}$ ($i = 1, \dots, n$) are monotone decreasing in $\eta(\mu_i)$ if and only if $(1 - \gamma) \geq 0$, the price of any ambiguous state in $\wp(\mu_i)$ is a increasing (decreasing) function of the degree of ambiguity $\eta(\mu_i)$ if and only if $\gamma \leq 1$. Therefore ambiguity will decrease stock prices if $\gamma > 1$, which is normally considered the most plausible case.⁵⁵ In addition, ambiguity aversion implies lower equilibrium interest rates, regardless of risk aversion. Therefore, a low coefficient of risk aversion accommodates for both a higher equity premium and a lower interest rate, so that both the equity premium and the interest rate puzzles might find a rational explanation for moderate amounts of

⁵⁵However, the intuition according to which $\gamma > 1$ would be the sensible calibration overwhelmingly derives from an empirical literature in which estimation and/or calibration have been performed for (S)EU models. Therefore our claim in the text has to be taken with caution. The possibility to obtain growing effects on equilibrium asset returns for $\gamma < 1$ (or even as $\gamma \rightarrow 0^+$) is a well-established fact in the incomplete information and learning literatures, see e.g., Guidolin (2006) and David (2008).

ambiguity as well as relative risk aversion. In fact, if we call V^* the optimal value function of the problem solved by the investor, the equilibrium equity premium will be expressed as

$$\gamma(\sigma_D + V^*) + \left[\sum_{i=1}^n \sqrt{2\eta(\mu_i)\sigma_D^2} \pi_i(t) \right] + \left[\sum_{i=1}^n \sqrt{2\eta(\mu_i)\sigma_D^2} \pi_i(t) \right] kV^* \quad (148)$$

where $k \equiv k_D^2 + k_e^2$, and k_D , k_e are coefficients in the dynamics of the worst-case model. The first term in the equity premium expression, $\gamma(\sigma_D + V^*)$, is a standard risk exposure and also incorporates the premium attached to the worst-case state, while the sum of the second and the third terms represents ambiguity premium.

4.3. Robustness

Another, crucial strand of the asset pricing literature has adopted the robustness modelling approach proposed by Hansen, Sargent, and their co-authors. One of the first papers in this camp is the discrete-time, linear-quadratic permanent income model used by Hansen, Sargent, and Tallarini (1999) to study consumption/saving rules. Specifically, Hansen et al. 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 decision makers. To understand the connection between risk-sensitive preferences and robust preferences, consider a simple control problem whose state transition equation is $\mathbf{x}_{t+1} = \mathbf{Ax}_t + \mathbf{Bc}_t + \mathbf{Cw}_{t+1}$, where \mathbf{x}_t is a state vector, \mathbf{c}_t is the control vector (e.g., consumption and portfolio shares in a portfolio problem) 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

$$V_t = U(\mathbf{c}_t, \mathbf{x}_t) + \beta \mathcal{R}_t(V_{t+1}), \quad (149)$$

where $U(\mathbf{c}_t, \mathbf{x}_t)$ is a one period utility index and

$$\mathcal{R}_t(V_{t+1}) = \frac{2}{\sigma} \ln E \left[\exp \left(\frac{\sigma V_{t+1}}{2} \right) \middle| \mathcal{F}_t \right] \quad (150)$$

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 VNM specification. Moreover, while under (S)EU the transition equation is represented by $\mathbf{x}_{t+1} = \mathbf{Ax}_t + \mathbf{Bc}_t + \mathbf{Cw}_{t+1}$, in a robust control problem the transition equation is a distorted law of motion,

$$\mathbf{x}_{t+1} = \mathbf{Ax}_t + \mathbf{Bc}_t + \mathbf{C}(\mathbf{w}_{t+1} + \boldsymbol{\theta}_t), \quad (151)$$

where $\boldsymbol{\theta}_t$ is used to distort the mean of the innovation \mathbf{w} and it is chosen to minimize U_0 satisfying two constraints:⁵⁶

$$\hat{E} \left[\sum \beta^j \boldsymbol{\theta}'_{t+j} \boldsymbol{\theta}_{t+j} \right] \leq \eta_t \text{ and } \eta_{t+1} = \beta^{-1} (\eta_t - \boldsymbol{\theta}'_t \boldsymbol{\theta}_t) \text{ for a given } \eta_0. \quad (152)$$

⁵⁶ $\hat{E}[\cdot]$ 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. However, at least in a discrete time set up, the distortion would in this case also extend to the covariance matrix of the innovations, if given a chance. However, Hansen, Sargent, Turmuhambetova, and Williams (2006) have proven that in a continuous time diffusion setting, the optimal distortion will not involve the the volatility matrix because this would be infinitely costly in terms of entropy.

Hansen and Sargent (1998) have shown that—letting the Lagrange multiplier for the second constraint be equal 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. The relationship between the two value functions and the corresponding decision rules establishes how the risk-sensitive preferences specification induces the same behavior that would occur without the risk-sensitivity adjustment, but with the concern for robustness and the pessimism induced by the distortion operated by $\boldsymbol{\theta}$. Hence the parameters β and σ can be interpreted as indicating a desire for robustness.

Hansen, Sargent, and Tallarini (1999) apply this framework to produce asset pricing results. They prove that the concern for robustness modifies the standard pricing equation, $1 = E[M_{t+1}R_{t+1}|\mathcal{F}_t]$ (where M_{t+1} denotes the stochastic discount factor) by introducing a multiplicative term M_{t+1}^u :

$$M_{t+1}^u = \frac{\exp [-(\mathbf{w}_{t+1} - \boldsymbol{\theta}_t)'(\mathbf{w}_{t+1} - \boldsymbol{\theta}_t)]}{\exp (-\mathbf{w}'_{t+1}\mathbf{w}_{t+1}/2)}, \quad (153)$$

which means that M_{t+1}^u is the density ratio of the distorted relative to the reference probability distribution. Additionally, it turns out that

$$\text{St. dev}(M_{t+1}^u|\mathcal{F}_t) = [\exp(\boldsymbol{\theta}'_t\boldsymbol{\theta}_t) - 1]^{1/2} \simeq \|\boldsymbol{\theta}_t\|, \quad (154)$$

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) have further extended these results going beyond the linear-quadratic framework and allowed for more general return and transition functions. In particular, the technique employed by Anderson et al. (2003) 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 decision maker (we refer the Reader to 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

$$\max_{\{c_t\}} \min_{\{\boldsymbol{\theta}_t\}} E \int_0^\infty \exp(-\delta t) \left[U(\mathbf{x}_t, c_t) + \frac{\varrho}{2} \boldsymbol{\theta}'_t \boldsymbol{\theta}_t \right] dt, \quad (155)$$

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 (155) 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 stochastic discount factor which is similar—but obviously more general—than the one in Hansen et al. (1999). The stochastic discount factor (SDF) between dates t and $t+\tau$ is the positive random variable

$$M_{t,t+\tau} = \exp \left(- \int_t^{t+\tau} [\eta(\mathbf{x}_s)ds + \lambda'(\mathbf{x}_s)d\mathbf{B}_s] \right), \quad (156)$$

for some appropriate choices of the vectors $\eta(\mathbf{x}_t)$ and $\lambda(\mathbf{x}_t)$. Equivalently, denoting by $Mu(\mathbf{x}_t)$ the logarithm of the marginal utility process for the numeraire consumption good, in the standard (S)EU framework

$$M_{t,t+\tau} = e^{-\delta s} \exp [Mu(\mathbf{x}_{t+\tau}) - Mu(\mathbf{x}_t)], \quad (157)$$

which means that the SDF is a (discounted) transformation of the rate of growth of marginal utility of consumption. Further, using the stochastic integral representation of the evolution of the marginal utilities,

$$\begin{aligned} Mu(\mathbf{x}_{t+\tau}) &= Mu(\mathbf{x}_t) + \int_t^{t+\tau} \mu_M(\mathbf{x}_s) ds + \int_t^{t+\tau} \lambda_M(\mathbf{x}_s) d\mathbf{B}_s \text{ where} \\ \mu_M(\mathbf{x}) &= \left(\frac{\partial Mu}{\partial \mathbf{x}} \right)' \mu^* + \frac{1}{2} \text{tr} \left(\frac{\partial^2 Mu}{\partial \mathbf{x} \partial \mathbf{x}'} \Sigma^* \right) du \\ \lambda_M(\mathbf{x}) &= (\sigma^*)' \frac{\partial Mu}{\partial \mathbf{x}} \quad \Sigma^* = \sigma^*(\mathbf{x}_t)(\sigma^*(\mathbf{x}_t))', \end{aligned} \quad (158)$$

it is easy to see that $\eta(\mathbf{x}_s) = \delta - \mu_M(\mathbf{x}_s)$ and $\lambda(\mathbf{x}_s) = -\lambda_M(\mathbf{x}_s)$, while the risk-free rate is $\delta - \mu_M(\mathbf{x}_s) - \frac{1}{2}\lambda_M(\mathbf{x}_s)'\lambda_M(\mathbf{x}_s)$.

When we take into account also the concern for model misspecification, the term $M_{t,t+\tau}$ is multiplied by an *adjustment for model uncertainty*:

$$M_{t,t+\tau}^u \equiv \frac{\exp(-V^*(\mathbf{x}_{t+\tau})/\varrho)}{E[\exp(-V^*(\mathbf{x}_{t+\tau})/\varrho)|\mathbf{x}_t]}. \quad (159)$$

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

$$M_{t,t+\tau}^u = \exp \left(\int_t^{t+\tau} \left[-\frac{1}{2} (\theta^*(\mathbf{x}_s))' \theta^*(\mathbf{x}_s) + (\theta^*(\mathbf{x}_s))' d\mathbf{B}_s \right] \right) \quad \theta^*(\mathbf{x}_s) \equiv -\frac{1}{\varrho} (\sigma^*(\mathbf{x}))' \frac{\partial V^*(\mathbf{x})}{\partial \mathbf{x}}. \quad (160)$$

Hence, in the presence of ambiguity, the implied risk free rate is

$$\delta - \mu_M(\mathbf{x}_s) - \frac{1}{2} [\lambda_M(\mathbf{x}_s)' \lambda_M(\mathbf{x}_s) + \theta^*(\mathbf{x}_s)' \theta^*(\mathbf{x}_s)],$$

where the last term represents the robustness adjustment.

Maenhout (2004) has further progressed on the analysis of equilibrium asset prices under ambiguity within the robustness framework and homotheticity. 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 been sometimes 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.⁵⁷ 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 (155). Hence, the empirical evidence of substantial pessimism when forming portfolios 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, and Kocherlakota, 1993).

Anderson et al. (2003) prove 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 , where 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, and

$$\mu(W_t) = W_t(r^f + \omega_t(\mu - r^f)) - c_t \text{ and } \sigma(W_t) = \omega_t \sigma W_t \quad (161)$$

⁵⁷In 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 underlying state-variable.

is short-hand notation for the drift and diffusion of the wealth process. The drift adjustment $\theta(W_t)$ is endogenously chosen and crucially depends on the parameter ϱ . The more a decision maker 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

$$\frac{dS_t + D_t dt}{S_t} - r^f dt = (\gamma + \varrho^{-1})\sigma_{cs} dt + \sigma_c dB_t, \quad (162)$$

with $\sigma_{cs} \equiv \text{cov}(dc/c, dS/S)$ 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 (162)—the drift of the excess return process—will exceed the standard expression for the equity premium in the CCAPM of Breeden (1979), $\gamma\sigma_{cs}$. 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

$$r^f = \delta + \psi\mu_c - \frac{1}{2}(1+\psi)(\gamma + \varrho^{-1})\sigma_c^2 \quad (163)$$

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 will lower the risk-free interest rate with respect to the one that derives from the standard CCAPM because $(1+\psi)\varrho^{-1}\sigma_c^2 > 0$. Therefore the desire for robustness drives down the equilibrium 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, constant relative risk aversion preferences, where $\gamma^{-1} = \psi$). In other words, model uncertainty is instrumental in avoiding Weil's (1989) risk-free rate puzzle.

As in Maenhout, Sbuelz and Trojani (2008) also introduce a time-state varying constraint in the formulation (10), specializing the results on the effects of a preference for robustness (ambiguity aversion) to the power utility case, obtaining further closed-form results. Assume that a state process x represents the exogenous primitives of the economy and that they follow the Ito process:

$$dx = \mu(x)dt + \sigma(x)dB^x. \quad (164)$$

The endowment and (cum dividend) stock price processes follow

$$\begin{aligned} dc &= \mu_c c(x)dt + \sigma_c c \left[\rho_c(x)dB^x + \sqrt{1 - \rho_c^2(x)}dB^c \right] \\ dS + cdt &= S\mu_S dt + S\sigma_S \left[\rho_S(x)dB^x + \sqrt{1 - \rho_S^2(x)}dB^c \right], \end{aligned} \quad (165)$$

where $\mathbf{B} \equiv [B^x \ B^c]'$ is a bivariate BM, $\rho_c(x)$ measures the correlation between shocks to the state and the endowment processes, and $\rho_S(x)$ captures the correlation between cum-dividend stock valuations and the state variable x . The cum-dividend return, $R \equiv (dS + cdt)/S$, displays a conditional expectation μ_S , a conditional volatility σ_S , and conditional correlation coefficient ρ_S to be determined in equilibrium. The dynamics described in (164)- (165) constitute

the reference model, $d\mathbf{Y} = \boldsymbol{\mu} dt + \boldsymbol{\Lambda} d\mathbf{B}$ (where $\mathbf{Y} \equiv [x \ W]'$) of this economy. As in Anderson et al. (2003) and Maenhout (2004), ambiguity aversion is modelled in the form of a contamination vector, $\boldsymbol{\theta} \equiv [\theta^x \ \theta^S]'$ that perturbs the BM \mathbf{B} so that, under the probabilistic measure induced by $\boldsymbol{\theta}$, the drift term of \mathbf{Y} is $\boldsymbol{\mu} + \boldsymbol{\Lambda}\boldsymbol{\theta}$:

$$d\mathbf{Y} = (\boldsymbol{\mu} + \boldsymbol{\Lambda}\boldsymbol{\theta})dt + \boldsymbol{\Lambda}d\mathbf{B}. \quad (166)$$

Following the robust control literature, to limit the concerns expressed by the representative investor to only the models that are statistically indistinguishable from the approximating, reference model, $\boldsymbol{\theta}$ 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, that is introduced to allow for state- and time-dependent ambiguity. Larger values of η represent higher degrees of ambiguity aversion. The agent optimization problem is

$$V(W, x) = \max_{c, \omega} \min_{\boldsymbol{\theta}} E_0^\theta \left[\int_0^\infty e^{-\delta t} \frac{(cW)^{1-\gamma}}{1-\gamma} dt \right], \quad (167)$$

subject to robustness constraint and the dynamics of \mathbf{Y} .

Sbuelz and Trojani show that a number of equilibrium quantities can be solved in (quasi-) closed form:⁵⁸

$$c^* = \left(\frac{\exp[(1-\gamma)g]}{\delta} \right)^{-\frac{1}{\gamma}} \quad \omega^* = \frac{1}{\gamma - \sqrt{\frac{2\eta}{G(\omega^*)} |f|}} \left[\frac{\mu_S - r^f}{\sigma_S^2} + \left(1 - \gamma - \sqrt{\frac{2\eta}{G(\alpha^*)} |f|} \right) \frac{\partial g}{\partial x} \frac{\rho_S \sigma(x)}{\sigma_S} \right]$$

Here $G(\omega^*) \equiv \sigma_S^2(\omega^*)^2 + \sigma^2(x)(g')^2 + 2\omega^* \rho_S \sigma_S \sigma(x)g'$ and g is a relatively complicated function of γ , $\sqrt{\eta}$, and x (we refer the interested Reader to Sbuelz and Trojani (2008) for the expressions for μ_S , σ_S^2 and ρ_S). Hence, ambiguity aversion has a direct impact only on the myopic and the hedging demands for equity. While in the absence of ambiguity ($\eta = 0$) the demand for the stock consists of the mean-variance style myopic term

$$\omega^{(S)EU} = \frac{1}{\gamma} \frac{\mu_S - r^f}{\sigma_S^2}, \quad (168)$$

when ambiguity is taken into account, the additional hedging term

$$\frac{\sqrt{\frac{2\eta}{G(\omega^*)} |f|}}{\gamma^2 - \gamma \sqrt{\frac{2\eta}{G(\omega^*)} |f|}} \frac{\mu_S - r^f}{\sigma_S^2} + \frac{1 - \gamma - \sqrt{\frac{2\eta}{G(\omega^*)} |f|}}{\gamma - \sqrt{\frac{2\eta}{G(\omega^*)} |f|}} \frac{\partial g}{\partial x} \frac{\rho_S \sigma(x)}{\sigma_S} \quad (169)$$

appears. The impact on the equity demand is state-dependent in a non-standard way: when equity risk barely pays off, and there is scarce need for hedging, or equity is an inadequate hedging tool, the preference for robustness generates the *first order risk aversion* (FORA) effect, in the sense that desired equity holdings are small even if ambiguity aversion is low. However ambiguity has no direct impact on the equilibrium equity price dynamics and it can only affect the parameters that enter the dynamic process of cum-dividend prices, i.e., equity expected returns, volatilities and correlations. In fact, under the worst case model that takes ambiguity aversion into account, we have

$$(\mu_S - r^f)_{\boldsymbol{\theta}} = \sigma_S^2 - (1 - \gamma) \left(\sigma_S^2 + \rho_S \sigma(x) \sigma_S \frac{\partial g}{\partial x} \right), \quad (170)$$

while under the reference model the equity premium is:

$$(\mu_S - r^f) = (\mu_S - r^f)_{\boldsymbol{\theta}} + \sqrt{\frac{2\eta}{\sigma_c^2 + \frac{2(1-\gamma)}{\gamma} \sigma(x) \rho_c \sigma_c \frac{\partial g}{\partial x} + \left(\frac{1-\gamma}{\gamma} \right)^2 \left(\frac{\partial g}{\partial x} \right)^2}} |f(x)| \left(\sigma_S^2 + \rho_S \sigma(x) \sigma_S \frac{\partial g}{\partial x} \right). \quad (171)$$

⁵⁸The ambiguous expression “quasi-closed form” refers to the fact that the expressions that follow contain terms g , $\frac{\partial g}{\partial x}$, and $G(\omega^*)$ that depend in possibly complicated, non-linear ways on γ , $\sqrt{\eta}$, and x .

Hence the equity premium under the reference model is directly affected by time varying ambiguity; specifically, it is higher in presence of model uncertainty if and only if the sum of the standard and the ambiguity-speculative demands for equity dominates over the corresponding intertemporal hedging demand ($\sigma_S^2 > |\rho_S \sigma(x) \sigma_S \partial g / \partial x|$) or if ambiguity-driven hedging implies a short position in equity ($\rho_S \sigma(x) \sigma_S \partial g / \partial x > 0$). By contrast, the worst-case model equity premium is indirectly affected by ambiguity only through g , ρ_S and σ_S . 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. The calibration shows that realistic amounts of ambiguity aversion give a substantial lift to the equity premium already at low levels of standard risk aversion. This happens without upward pressure on the risk-free rate, since the elasticity of intertemporal substitution is left unchanged.

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 generate large asset pricing anomalies. Hansen’s (2007) novelty consists in establishing a link between ambiguity aversion and the concept of Chernoff’s (1952)’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.

Assume that the N -dimensional random vector $\mu + \Lambda v$ represents the returns provided by N assets, where the vector v is normally distributed with mean zero and an identity covariance matrix. Therefore the matrix Λ determines the priced risk exposures of each asset to each of the shocks and the covariance matrix of asset returns is $\Sigma = \Lambda \Lambda'$. Under the standard normal/CARA framework, the risk factors are lognormal and they can be characterized as a transformation of the kernel

$$\exp\left(\omega' \mu + \omega' \Lambda v - \frac{1}{2} \omega' \Sigma \omega\right), \quad (172)$$

where ω represents a given portfolio choice. Denoting by $\exp(r^f)$ the risk-free return, Hansen assumes that the logarithm of the asset prices can be written as

$$\ln p(\omega) = \omega' \mu - r^f - \omega' \Lambda q, \quad (173)$$

for some N -dimensional vector q , whose components are the risk prices. Hence, q prices the exposure to shocks in v and it is the compensation for risk exposure in terms of logarithms of the means. Let’s assume that Λ is such that, whenever ω is a versor vector e_j (this is a vector that is zero everywhere but in its j -th position, where a 1 appears), then $p(e_j) = 1$ and the logarithm of the excess return is $\omega' \mu - r^f = \omega' \Lambda q$ (such that $\omega' 1 = 1$). Assuming that

investors maximize the logarithm of the Sharpe ratio by choosing ω , Hansen proves that:

$$\max_{\omega' \mathbf{1} = 1} \frac{\omega' \mu - r^f}{\sqrt{\omega' \Sigma \omega}} = \max_{\omega' \mathbf{1} = 1} \frac{\omega' \Lambda \mathbf{q}}{\sqrt{\omega' \Sigma \omega}} = |\mathbf{q}| = [(\mu - \iota_n r^f)' \Sigma^{-1} (\mu - \iota_n r^f)']^{1/2}. \quad (174)$$

In reality, the mean μ is not perfectly revealed to an econometrician or perhaps even to investors, that instead might be able to estimate an alternative mean $\hat{\mu}$. The two Sharpe ratios (associated to μ and $\hat{\mu}$) are related as follows:

$$[(\hat{\mu} - \mu)' \Sigma^{-1} (\hat{\mu} - \mu)]^{1/2} - [(\mu - \iota_n r^f)' \Sigma^{-1} (\mu - \iota_n r^f)']^{1/2} \leq [(\hat{\mu} - \iota_n r^f)' \Sigma^{-1} (\hat{\mu} - \iota_n r^f)']^{1/2}. \quad (175)$$

The first term $[(\hat{\mu} - \mu)' \Sigma^{-1} (\hat{\mu} - \mu)]^{1/2}$ is the square root of 8 times the Chernoff entropy of model $\hat{\mu}$ with respect to model μ . Therefore, Small Chernoff rates (that is, small errors in the estimate of $\hat{\mu}$) induce sizable variations of the Sharpe ratio and large pricing anomalies. For example, a Chernoff rate of 1% per annum increases the maximum Sharpe ratio by about 0.14.

Barillas, Hansen and Sargent (2009) have followed up on Hansen's (2007) intuitions and have attacked the problem of using ambiguity to produce SDFs with properties that are consistent with the observed (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 the literature in terms of violations of Hansen and Jagannathan's (1991) pricing bounds:

$$\frac{|E[R_{t+1}^e - R_t^f]|}{\sigma[R_{t+1}^e - R_t^f]} \leq \frac{\sigma[M_{t+1}]}{E[M_{t+1}]} \quad (176)$$

The ratio $\sigma[M_{t+1}] / E[M_{t+1}]$ is commonly called the market price of risk and it is the slope of the mean-standard deviation frontier for SDFs. 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.⁵⁹ To discuss the plausibility of the parameter for uncertainty aversion calibrated on asset market data, Barillas et al. 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.⁶⁰

⁵⁹A similar point has been made by Cagetti, Hansen, Sargent and Williams (2002) in the context of a stochastic growth model under uncertain growth, for which a high market price of uncertainty occurs not because of confidence in low growth but rather because of ambiguity about which growth regime rules the economy.

⁶⁰Given a reference probabilistic 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 , $t = 0, \dots, T - 1$. 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

$$r_i \equiv \log \frac{L_{ii}}{L_{ij}} \quad j \neq i \quad i = A, B^\rho.$$

Therefore, if we assume that model A has generated the data, $p_A = \text{Prob}(\text{mistake}|A) = \text{freq}(r_A \leq 0)$. Similarly, $p_{B^\rho} = \text{Prob}(\text{mistake}|B^\rho) = \text{freq}(r_{B^\rho} < 0)$. The detection error probability associated with the ambiguity aversion parameter ρ is defined as $p(\rho) = \frac{1}{2}(p_A + p_{B^\rho})$.

4.4. Smooth Ambiguity Preferences and the SDF

There is another strand of the ambiguity literature that has tackled the typical asset pricing questions—i.e., how to obtain plausible properties for the SDF in general equilibrium, how to rationalize high equity premia with low real interest rates along with high and strongly time-varying volatility—using ambiguity preferences of the KMM type. For instance, Ju and Miao (2008) have developed a pure exchange economy 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, see Section 3.5 for details. Ju and Miao (2008) work with the recursive functional

$$V_t(c) = W(c_t, \mathbf{R}_t(V_{t+1}(c))) \quad \mathbf{R}_t(\cdot) \equiv v^{-1} \left(\int_{\wp_{x_t}} v \left(U^{-1} \int U(\cdot) dp \right) dq_t(p) \right), \quad (177)$$

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:

$$\ln(c_{t+1}/c_t) = \mu_{x_{t+1}} + \sigma \epsilon_{t+1} \quad \epsilon_t \sim IIDN(0, 1). \quad (178)$$

Setting $U(c) = c^{1-\gamma}/(1-\gamma)$, $v(x) = x^{1-\eta}/(1-\eta)$, and

$$W(c, \mathbf{R}) = \left[(1-\beta)c^{1-\psi} + \beta \mathbf{R}^{1-\psi} \right]^{\frac{1}{1-\psi}} \quad \psi > 0 \text{ and } \beta \in (0, 1), \quad (179)$$

the decision maker 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 functional 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 that in this case reduces to

$$\Pr(x_{t+1} = 1 | s^t) = \frac{\lambda_{11} f(\ln(c_{t+1}/c_t), 1) q_t + \lambda_{21} f(\ln(c_{t+1}/c_t), 2) (1 - q_t)}{f(\ln(c_{t+1}/c_t), 1) q_t + f(\ln(c_{t+1}/c_t), 2) (1 - q_t)}, \quad (180)$$

where λ_{ij} is the transition probability of x , $f(\ln(c_{t+1}/c_t), s)$ is the probability density of log-consumption growth when the mean is $\mu_{x=s}$ ($s = 1, 2$), and the variance is σ . By standard principles, the corresponding pricing kernel is the intertemporal marginal rate of substitution, here

$$M_{x_{t+1}} = \beta \left(\frac{c_t}{c_{t+1}} \right)^{-\psi} \left(\frac{V_{t+1}}{\mathbf{R}_t(V_{t+1})} \right)^{\psi-\gamma} E_{t,x_{t+1}} \left(\frac{V_{t+1}^{1-\gamma}}{\mathbf{R}_t(V_{t+1})} \right)^{-(\eta-\gamma)}, \quad (181)$$

where $E_{t,x_{t+1}}$ denotes the expectation operator conditional on the history s^t and the state x_{t+1} . The last multiplicative factor in the definition of the SDF 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 pricing kernel. Intuitively, the agent has a lower continuation value V_{t+1} in a recession state, causing the adjustment factor in $M_{x_{t+1}}$ 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 put more weight on the pricing kernel, ambiguity aversion generates high negative correlation between the pricing kernel and stock return that increases the equity premium:

$$E_t [R_{e,t+1} - R_{f,t+1}] = -\frac{Cov_t(R_{e,t+1}, M_{x_{t+1}})}{E_t [M_{x_{t+1}}]} \quad R_{f,t+1} = \frac{1}{E_t [M_{x_{t+1}}]}, \quad (182)$$

while a high conditional mean for the SDF decreases the real interest rate. In Ju and Miao (2008), ambiguity aversion increases (in absolute value) the term $Cov_t(R_{e,t+1}, M_{x_{t+1}}) < 0$ and therefore—for given $E_t [M_{x_{t+1}}]$ —it magnifies the equity risk premium. 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.

4.5. Issues with Asset Pricing Applications of Models of Ambiguity

Although the papers we have reviewed yield a number of insights, a few contributions have cautioned against the dangers of careless applications of some of basic insights from models of ambiguity. For instance, Gollier (2009) has warned that generalizations of the qualitative (calibration) results in papers such as Ju and Miao (2008) 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. Gollier has investigated the effects of increases in ambiguity aversion on optimal portfolio allocation and equilibrium asset prices using KMM preferences as in Section 2.4. While it has been suggested that ambiguity aversion will inevitably make individuals more cautious, thereby offering a potential explanation for the equity premium puzzle, Gollier proves instead that an increase in ambiguity aversion does not necessarily imply a reduction in the demand for risky assets. In particular, there exist sufficient conditions under which, *ceteris paribus*, an increase in ambiguity aversion reduces the demand for the ambiguous asset, and raises the equity premium; however, absent those conditions not much can be said on the effects of ambiguity aversion on asset prices.

In particular Gollier shows that for particular sets of priors φ , the introduction of ambiguity aversion *increases* the investor's demand for the ambiguous asset. Indeed, the first-order condition for portfolio optimality can be interpreted as a standard Euler equation under (S)EU, with a distortion in the compounded probabilities. In particular, under ambiguity aversion, the agent computes the compounded distorted probability assigning more weight on those distributions yielding a smaller expected utility. This distortion is pessimistic and, under standard (S)EU-theory, pessimistic deteriorations in beliefs are known to not always reduce the demand for the risky asset 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 in the literature, because of the use of specific parametric set ups which are of course hard to avoid in practice, but also carry much more importance than it may have been commonly recognized.

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-(S)EU investors may deliver implications for asset prices that are substantially

different from those that a matching representative agent model would give.⁶¹ 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. In their framework, there are two periods, one for trading and one for consumption. Uncertainty is resolved at time 1, so that one of $|\Omega|$ possible states has realization. Each agent receives a fixed (known) endowment at time 0, and a further random endowment at time 1. Agents can trade N contingent claims, each of which is in zero net supply, and a risk-free bond. The utility index has a simple negative exponential form, while ambiguity is captured by MPP. Specifically, agents are assumed to have an reference probability \tilde{p} over the sample space Ω , so that the set of effective priors considered by each agent is

$$\wp = \left\{ q \in \Delta(\Omega) : \left| \frac{q_s}{\tilde{p}_s} - 1 \right| \leq \epsilon \text{ } s \in \Omega \right\} \quad (183)$$

for some specific $\epsilon \geq 0$, where $\epsilon = 0$ corresponds to the standard (S)EU-framework. In their calibration, ϵ is assumed to vary over $[0, 0.95]$. The calibration of the price of the market portfolio and of the risk-free asset 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 (S)EU and the other has $\epsilon = 0.95$. Similarly, when one of the agents is (S)EU, 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.

Trojani and Vanini (2000) have used their portfolio choice model with a concern for robustness described in Section 3.3 to study the effects of ambiguity in a heterogeneous, production general equilibrium economy characterized by two agents with power and logarithmic utility indexes, and ambiguity parameters $\varphi_1 > 0$, $\varphi_2 = 0$, $\eta_1 > 0$, $\eta_2 = 0$, respectively. This means that the first investor is an ambiguity-averse, power utility agent and the second investor an EU maximizer with unit relative risk aversion coefficient. Of course, their examples suffer from a tight parameterization and therefore are subject to the Gollier's (2009) critique. Yet, imposing specific parametric assumptions helps in developing results for the effects of heterogeneity in preferences on equilibrium asset prices. Trojani and Vanini (2000) analyze a robust version of the classical Merton's (1969) 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 GBM. The risky asset is a share of stock that gives ownership of the production technology with constant

⁶¹The analysis of the interaction between ambiguity and heterogeneity is one of the leading research themes for the future. In an experimental setting, Bossaerts, Ghirardato, Guarnaschelli and Zame (2010) study the effects of ambiguity on equilibrium asset prices. Their results show 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 prices for which these agents refrain from trading. Contrary to the standard (S)EU framework, investors are not infra-marginal, since ambiguity averse investors who decide to not trade still have an indirect effect on prices because the per-capita amount of risk that has to be borne by those who are left trading is affected.

return to scale. They find that in production economies, it is possible to find a threshold level for the cross-sectional wealth distribution $w = W_1 / (W_1 + W_2)$, below which only agent 2 trades the risky asset. In particular, if w falls below the threshold, agent 1 leaves the security market, so that prices reflect the preferences of the log-utility investor only. Such a threshold is a decreasing function of the volatility- (risk) to-ambiguity ratio, hence, when the asset's volatility is large relative to the amount of ambiguity, both agents tend to trade. Indeed, while risk aversion induces agents to trade in order to diversify risk, ambiguity prevents trade.

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) who can observe past prices of a risky asset (say “ a ”), dividends and an additional ambiguous signal that is informative about future dividends of the asset. 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. Hence, the market portfolio is a claim to the total dividend:

$$\frac{1}{n}d + \frac{n-1}{n}\tilde{d} = \frac{1}{n}m + \frac{n-1}{n}\tilde{m} + \varepsilon^a + \frac{1}{n}\varepsilon^i + \frac{n-1}{n}\tilde{\varepsilon}^i. \quad (184)$$

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 $\sigma_s^2 \in [\sigma_{*s}^2, \sigma_s^{*2}]$. There is a single normal prior for μ , the two-dimensional parameter $\mu \equiv [\varepsilon^a + \varepsilon^i \ \varepsilon^a]'$, 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 $\alpha^2\sigma_a^2 + \sigma_i^2 + \sigma_s^2$. The set of posteriors about μ at date 1 is calculated using standard rules for updating normal random variables. 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:

$$\xi(\sigma_s^2) = \frac{Cov(\varepsilon^a + \varepsilon^i, s)}{Var(s)} = \frac{Cov(\varepsilon^a + \varepsilon^i, \alpha\varepsilon^a + \varepsilon^i + \varepsilon^s)}{Var(\alpha\varepsilon^a + \varepsilon^i + \varepsilon^s)} = \frac{\alpha^2\sigma_a^2 + \sigma_i^2}{\alpha^2\sigma_a^2 + \sigma_i^2 + \sigma_s^2}. \quad (185)$$

Given s , the posterior density of μ is also normal. As σ_s^2 ranges over $[\sigma_{*s}^2, \sigma_s^{*2}]$, the coefficient $\xi(\sigma_s^2)$ also varies, tracing out a family of posteriors. The coefficient $\xi(\sigma_s^2)$ is commonly used to measure the information content of a signal relative to the volatility of the parameter, since it determines the fraction of prior variance in μ that is resolved by the signal. Hence, under ambiguity, $\xi(\sigma_{*s}^2)$ and $\xi(\sigma_s^{*2})$ provide bounds on the (relative) information content, while,

in the Bayesian case, $\xi(\sigma_{*s}^2) = \xi(\sigma_s^{*2})$, so that agents know how much information is conveyed by s . 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.

For simplicity, Epstein and Schneider focus on the case in which the representative agent is risk neutral and does not discount the future; additionally, she 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

$$p_1^{(S)EU}(s) = m + \varepsilon^a + \varepsilon^i = m + \xi(\sigma_s^2)s \quad p_0^{(S)EU} = E[p_1^{(S)EU}|s] = m, \quad (186)$$

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:

$$p_1^{AA}(s) = \min_{\sigma_s^2 \in [\sigma_{*s}^2, \sigma_s^{*2}]} E[d|s] = \begin{cases} m + \xi(\sigma_s^{*2})s & s \geq 0 \\ m + \xi(\sigma_{*s}^2)s & s < 0 \end{cases}. \quad (187)$$

Therefore $p_1^{AA}(s)$ is a piece-wise linear functions kinked at $s = 0$ and with lower slope over $s \geq 0$, differently from $p_1^{(S)EU}(s)$ which is simply 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} . This yields the expression:

$$p_0^{AA} = \min_{\sigma_s^2 \in [\sigma_{*s}^2, \sigma_s^{*2}]} E_s[p_1^{AA}|s] = m - (\xi(\sigma_{*s}^2) - \xi(\sigma_s^{*2})) \frac{1}{\sqrt{2\pi\xi(\sigma_s^{*2})}} \sqrt{\alpha^2\sigma_a^2 + \sigma_i^2}. \quad (188)$$

Therefore, p_0^{AA} exhibits an ambiguity premium,

$$p_0^{(S)EU} - p_0^{AA} = (\xi(\sigma_{*s}^2) - \xi(\sigma_s^{*2})) \frac{1}{\sqrt{2\pi\xi(\sigma_s^{*2})}} \sqrt{\alpha^2\sigma_a^2 + \sigma_i^2} > 0, \quad (189)$$

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 (and cheaper) already today, while, in the standard Bayesian framework, any feedback of future information quality on current utility is precluded. Moreover, while under (S)EU 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 the mean m , under ambiguity the market portfolio does not become less uncertain as the number of assets increases, since the ambiguity-averse investor acts as if the mean was lower, regardless of the number of assets.

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 just an ambiguous signal about the next innovation in dividends before that innovation is revealed. Because of risk neutrality, in equilibrium the riskless interest rate r^f equals $1/(1 + \beta)$, where β is the intertemporal discount rate. Assume a standard mean-reverting process for dividends, $d_t = \kappa\bar{d} + (1 - \kappa)d_{t-1} + u_t$, where \bar{d} is the unconditional mean dividend, $(1 - \kappa)$ measures mean reversion, and u_t is IID Gaussian. Every period an ambiguous signal about the value of next period's shock u_{t+1} is observed. For simplicity, the distinction between systematic and idiosyncratic shocks is omitted so that $s_t = u_{t+1} + \varepsilon_t^s$, where the variance of

ε_t^s is $\sigma_{s,t}^2 \in [\sigma_{*s}^2, \sigma_s^{*2}]$. In equilibrium, the price of the stock is:

$$p_t^{AA} = \min_{\sigma_{s,t}^2, \sigma_{s,t+1}^2 \in [\sigma_{*s}^2, \sigma_s^{*2}] \times [\sigma_{*s}^2, \sigma_s^{*2}]} \beta E_t[p_{t+1} + d_{t+1}]. \quad (190)$$

Focusing on stationary equilibria and conjecturing a price function with structure $p_t = \bar{Q} + Q_d d_t + Q_s \xi(\sigma_{s,t}^2) s_t$, Epstein and Schneider prove that:

$$p_t^{AA} = \frac{\bar{d}}{r^f} + \frac{1-\kappa}{r^f+\kappa}(d_t - \bar{d}) + \frac{1}{r^f+\kappa} \left(\begin{cases} \xi(\sigma_s^{*2}) s_t & s_t \geq 0 \\ \xi(\sigma_{*s}^2) s_t & s_t < 0 \end{cases} \right) - (\xi(\sigma_{*s}^2) - \xi(\sigma_s^{*2})) \frac{\sigma_u}{\sqrt{2\pi\xi(\sigma_s^{*2})}}, \quad (191)$$

which should be contrasted with the (S)EU Bayesian expression:

$$p_t^{(S)EU} = \frac{\bar{d}}{r^f} + \frac{1-\kappa}{r^f+\kappa}(d_t - \bar{d}). \quad (192)$$

This is the present discounted value of dividends: without ambiguity, prices are determined only by current interest rate and dividend levels. Hence the (time-varying) ambiguity premium is:

$$\frac{1}{r^f+\kappa} \left(\begin{cases} \xi(\sigma_s^{*2}) s_t & s_t \geq 0 \\ \xi(\sigma_{*s}^2) s_t & s_t < 0 \end{cases} \right) - (\xi(\sigma_{*s}^2) - \xi(\sigma_s^{*2})) \frac{\sigma_u}{\sqrt{2\pi\xi(\sigma_s^{*2})}}. \quad (193)$$

The first term captures the response to the current, ambiguous signal. As before, this response is asymmetric, so that bad news is incorporated into prices with a stronger effect. In addition, the strength of this reaction depends on the persistence of dividends: if κ is smaller (i.e., shocks to dividends 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. Next, Epstein and Schneider examine the properties of per-share excess returns $\hat{R}_{t+1}^{AA} \equiv [R_{t+1}^{AA} - (1+r^f)] \times p_t^{AA}$ under ambiguity:

$$\begin{aligned} \hat{R}_{t+1}^{AA} &= p_{t+1}^{AA} + d_{t+1} - p_t^{AA}(1+r^f) \\ &= \frac{1}{r^f+\kappa} \left(\begin{cases} \xi(\sigma_s^{*2}) s_{t+1} & s_{t+1} \geq 0 \\ \xi(\sigma_{*s}^2) s_{t+1} & s_{t+1} < 0 \end{cases} \right) + \frac{1+r^f}{r^f+\kappa} \left[u_{t+1} - \frac{1}{r^f+\kappa} \left(\begin{cases} \xi(\sigma_s^{*2}) s_t & s_t \geq 0 \\ \xi(\sigma_{*s}^2) s_t & s_t < 0 \end{cases} \right) \right] + \\ &\quad + (\xi(\sigma_{*s}^2) - \xi(\sigma_s^{*2})) \frac{\sigma_u}{\sqrt{2\pi\xi(\sigma_s^{*2})}}. \end{aligned} \quad (194)$$

The first term 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 s_{t+1} and u_{t+1} are. The last (constant) term represents a constant premium that investors obtain for enduring low information quality in the future. Without ambiguity we would simply have

$$\hat{R}_{t+1}^{(S)EU} = \frac{1}{r^f+\kappa} \xi(\sigma_s^2) s_{t+1} + \frac{1+r^f}{r^f+\kappa} (u_{t+1} - \xi(\sigma_s^2) s_t), \quad (195)$$

which shows that the ambiguity risk premium term $(\xi(\sigma_{*s}^2) - \xi(\sigma_s^{*2}))\sigma_u/\sqrt{2\pi\xi(\sigma_s^{*2})} > 0$ is uniquely contributing to raise the equity risk premium relative to the (S)EU case. The mean excess return under the true probability is

$$E^*[\hat{R}_{t+1}^{AA}] = (\xi(\sigma_{*s}^2) - \xi(\sigma_s^{*2})) \frac{\sigma_u}{\sqrt{2\pi\xi(\sigma_s^{*2})}} \left[1 + \frac{r^f}{r^f+\kappa} \sqrt{\frac{\xi(\sigma_s^{*2})}{\xi(\bar{\sigma}_s^2)}} \right] \sigma_u > 0, \quad (196)$$

which is to be contrasted with $E^*[\hat{R}_{t+1}^{(S)EU}] = 0$ which must obtain in a risk-neutral framework in which there is 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 $\xi(\sigma_{*s}^2) = \xi(\sigma_s^{*2}) = \xi(\bar{\sigma}_s^2)$, the derivatives of the volatility of equilibrium returns under ambiguity $Var^*[\hat{R}_{t+1}]$ 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 precisions 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 literature started by LeRoy and Porter (1981) and Shiller (1981). Finally, 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). In Epstein et al.'s model this means that large cap stocks are "in the news" more, so that traders have to elaborate more ambiguous information. In contrast, small stocks are not widely followed by the media and information mostly arrives in the form of the occasional earnings report, so that one would not expect a lot of negative skewness.

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.⁶² 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. Moreover, because of investors' heterogeneity, Illeditsch (2009) also proves interesting results for limited participation and market "freezes"—i.e., absence of active trading in response to large shocks (here, signal realizations)—in terms of the existence of an interval of prices such that investors do not participate in the market because of ambiguity.

4.7. Effects of Ambiguity on Other Asset Classes

Recent years have witnessed a flurry of papers that have taken applications of decision making under ambiguity to asset pricing well beyond the narrow boundaries of questions concerning equilibrium stock prices, interest rates, the equity premium and its relationship with volatility. Although this literature is in no way negligible for its applied

⁶²In 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 setting characterized by asymmetric information and both ambiguity averse and (S)EU-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 (S)EU-agents arise and under/over-reactions to announcements may occur.

contributions or the elegance of the ambiguity frameworks proposed—see e.g., the applications by Autore, Billingsley, and Schneller (2009), Benigno and Nistico (2009), Boyarchenko (2009), Kleshchelski and Vincent (2007), Miao and Wang (2006), and Uhlig (2009)—we will limit our treatment in this sub-section to four papers that—by necessity and besides their specific merits—have to be taken as mere examples.

Liu, Pan, and Wang (2002) have examined the equilibrium 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 to the 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. The rate of endowment flow Y_t is assumed to be a one-dimensional Markov process

$$dY_t = \mu Y_t dt + \sigma dB_t + (e^{M_t} - 1) Y_{t-} dN_t, \quad (197)$$

where B_t is a BM and N_t is a pure-jump process with intensity $\lambda > 0$. Given a jump arrival at time t , the jump amplitude is controlled by a normally distributed random variable with mean μ_J and standard deviation σ_J . Consequently, the mean percentage jump in Y is $k = \exp(\mu_J + \sigma_J^2/2) - 1 \leq 0$. Non-positivity is assumed to focus on worse case scenarios. Shares of ownership of the aggregate endowment Y_t are traded as stocks with price S_t while the risk-free rate is r_t^f . Assuming that the stock price has the form $S_t = A(t)Y_t$, where $A(t)$ is a deterministic function of t ($A(T) = 0$), under the reference measure p , in equilibrium the stock price follows the process:

$$dS_t = \left(\mu + \frac{A'(t)}{A(t)} \right) S_t dt + \sigma S_t dB_t + (e_t^M - 1) S_{t-} dN_t. \quad (198)$$

The alternative models that an ambiguity averse investor will entertain are defined by the probability measure q^z , where z is its Radon-Nikodym derivative with respect to p . This is a slight variation over Anderson et al.'s (2003) framework. In fact, here z simply changes the agent's probability assessment with respect to the jump component, without altering her view about the diffusive component. Letting φ be the entire collection of probabilities q^z , the agent's value function at time t is (with terminal condition $V_T = 0$)

$$V_t = \lim_{\Delta \rightarrow 0} \frac{c_t^{1-\gamma}}{1-\gamma} \Delta + e^{-\delta \Delta} \inf_{q^z \in \varphi} \left\{ \varrho \varphi(V_t) E_t^z \left[\psi \left(\ln \frac{z_t + \Delta}{z_t} \right) \right] + E_t^z [V_{t+\Delta}] \right\}, \quad (199)$$

where ϱ is the ambiguity aversion coefficient, $\varphi(V_t)$ is a normalizing factor, and $\psi(x) = x + \beta(e^x - 1)$. This specification implies that a preference for robustness affects the representative agent in two ways. On the one hand, to protect herself against model uncertainty, the investor evaluates her future prospect $E_t^z [V_{t+\Delta}]$ under a range of alternative measures q^z s in φ ; she focuses on jump models that provide worse prospects than the reference one, hence the minimization over φ . On the other hand, the reference probability measure p is the best representation of the data, so the investor penalizes her choice q^z according to how much it deviates from the reference p . In equilibrium, the equity premium can be decomposed into a diffusive risk premium $\gamma \sigma^2$, a jump risk premium $\lambda k - \bar{\lambda} \bar{k}$, and rare event premium $\bar{\lambda} \bar{k} - \lambda^{q^*} k^{q^*}$, where q^* is the solution to the minimization problem in (199) for the optimal consumption level and

$$\begin{aligned} \lambda^{q^*} &= \lambda \exp \left(-\gamma \mu_J + \frac{1}{2} \gamma^2 \sigma_J^2 + a^* - b^* \gamma \sigma_J^2 \right) & k^{q^*} &= (1+k) \exp((b^* - \gamma) \sigma_J^2) - 1 \\ \bar{\lambda} &= \lambda \exp \left(-\gamma \mu_J + \frac{1}{2} \gamma^2 \sigma_J^2 \right) & \bar{k} &= (1+k) \exp(-\gamma \sigma_J^2) - 1. \end{aligned}$$

The first premium depends on risk aversion only: therefore it approaches 0 when $\gamma \rightarrow 0$, and is positive for any risk-averse investor. 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 $\lambda^{q^*} k^{q^*} > 0$ for $\gamma \rightarrow 0$, while $\lambda^{q^*} k^{q^*} \rightarrow 0$ for $\varrho \rightarrow \infty$. Moreover, under the alternative specification q^* , the stock price follows the process

$$dS_t = (r^f - q) S_t dt + \sigma S_t dB_t^{q^*} + (c_t^z - 1) S_t dN_t - \lambda^{q^*} k^{q^*} dt, \quad (200)$$

where q is the dividend payout rate, $B_t^{q^*}$ is a standard BM under q^* , N_t is a Poisson process with intensity λ^{q^*} , and—given jump arrival at time t —the percentage jump amplitude is log-normally distributed with mean k^{q^*} . Liu, Pan, and Wang prove that in this framework, European-style options can be priced using a simple modification of the celebrated Black and Scholes (1973) formula ($BS(S_0, K, r_j, q, \sigma_j, \tau)$):

$$\begin{aligned} C_0 &= \exp[-\lambda^{q^*} (1 + k^{q^*}) \tau] \sum_{j=0}^{\infty} \frac{[\lambda^{q^*} (1 + k^{q^*}) \tau]^j}{j!} BS(S_0, K, r_j, q, \sigma_j, \tau) \\ &r^f - \lambda^{q^*} k^{q^*} + \frac{j \ln(1 + k^{q^*})}{\tau} \quad \sigma_j^2 = \sigma^2 + \frac{j \sigma_J^2}{\tau} \quad \text{for } j = 0, 1, \dots \end{aligned} \quad (201)$$

Liu, Pan, and Wang (2002) have actually verified that this option pricing formula may produce a sensible empirical performance. They use one-month S&P 500 European-style options (with a ratio between strike and spot prices varying from 0.9 to 1.1) to test the pricing implications of their formula. Under the assumption that $\varrho \rightarrow \infty$ (no 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 however very feeble when compared to the pronounced “smirk” patterns observed in actual markets. On the contrary, under the assumption that $\varrho < \infty$ (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 becomes even more pronounced for out-of-the-money puts, that are typically highly sensitive to adverse, rare events. This is another feature that echoes what is commonly observed in equity options markets.

Another field of asset pricing in which it appears natural to test the implications of ambiguity is the pricing of fixed income securities, such as long-term coupon bonds, and derivatives that have bonds as their underlying assets. Gagliardini, Porchia and Trojani (2008) have extended the production economy of Cox, Ingersoll, and Ross (1985, henceforth CIR) and allowed for ambiguity aversion. They model ambiguity as a preference for robustness following Anderson et al. (2003) to capture a representative investor’s concern about misspecification of the process of the production technology shock. Gagliardini et al. adopt a setting in which the horizon is infinite and uncertainty is characterized in terms of a probability space (Ω, \mathcal{F}, p) and a $(k+1)$ -dimensional standard BM \mathbf{B}_t . p is the representative investor’s reference belief and alternative probabilistic models are specified in terms of probability distributions q^θ , such that the process $\mathbf{B}_t^\theta = \mathbf{B}_t + \int_0^t \boldsymbol{\theta}_s ds$ is a BM under q^θ and the vector $\boldsymbol{\theta}$ satisfies $\boldsymbol{\theta}' \boldsymbol{\theta} \leq 2\eta$, for a given (constant) ambiguity aversion parameter $\eta \geq 0$. The asset menu is composed of a locally riskless bond in zero net supply with return r_t^f , and K risky (real) bonds, also in zero net supply. Markets are complete and risk is characterized by the dynamics of K state variables with Ito process $d\mathbf{Y}_t = \boldsymbol{\mu}(\mathbf{Y}_t)dt + \boldsymbol{\sigma}(\mathbf{Y}_t)d\mathbf{B}_t$ that influence a linear technology Z , with

output rate

$$\frac{dZ_t}{Z_t} = \mu_Z(\mathbf{Y}_t)dt + \boldsymbol{\sigma}'_Z(\mathbf{Y}_t)d\mathbf{B}_t. \quad (202)$$

The K assets have a vector price process \mathbf{S} that satisfies the stochastic differential equation

$$d\mathbf{S}_t = \mathbf{I}_S \boldsymbol{\mu}_S(\mathbf{Y}_t)dt + \mathbf{I}_S \boldsymbol{\sigma}_S(\mathbf{Y}_t)d\mathbf{B}_t, \quad (203)$$

where $\mathbf{I}_S = \text{diag}[S_1, S_2, \dots, S_K]$. The representative agent's utility index is logarithmic, $U(c, t) = e^{-\delta t} \ln c$, where c denotes consumption. As usual, the optimization problem consists of choosing the optimal investment in the technology Z and in the K bonds ($\boldsymbol{\omega}_t \equiv [\omega_t^Z \ \boldsymbol{\omega}_t^S]'$) and optimal consumption levels, considering the worst case scenario induced by the perturbation $\boldsymbol{\theta}$:

$$\begin{aligned} & \max_{c, \boldsymbol{\omega}} \inf_{\boldsymbol{\theta}} E^{q^\theta} \left[\int_0^\infty e^{-\delta t} \ln(c, s) \middle| \mathcal{F}_0 \right] \\ \text{s.t. } & dW = W_t \left[\omega_t^Z (\mu_Z(\mathbf{Y}_t) - r_t^f) + (\boldsymbol{\omega}_t^S)' (\boldsymbol{\mu}_S(\mathbf{Y}_t) - r_t^f \boldsymbol{\iota}_k) + \left(r_t^f - \frac{c_t}{W_t} \right) dt \right] + \omega_t' \boldsymbol{\Sigma}(t) [d\mathbf{B}_t + \boldsymbol{\theta}_t dt]. \end{aligned} \quad (204)$$

Here W is the wealth process and $\boldsymbol{\Sigma}(t) \equiv [\boldsymbol{\sigma}'_Z(\mathbf{Y}_t), \boldsymbol{\sigma}_S(\mathbf{Y}_t)]'$. An equilibrium is a consumption process, a model contamination process, an instantaneous risk-free interest rate, and K asset return processes, such that these processes solve the optimization problem and optimal consumption is financed by a self-financing trading strategy for which $\omega_t^z = 1$ and $\boldsymbol{\omega}_t^S = \mathbf{0}$. Denoting by V the optimal value function and by $\nabla V_{\mathbf{Y}}$ the gradient of V with respect to the state vector \mathbf{Y} , the market prices of risk and ambiguity are given by:

$$\pi \equiv \sigma_Z(\mathbf{Y}_t) + \sqrt{2\eta} \frac{\boldsymbol{\sigma}'(\mathbf{Y}_t) \nabla V_{\mathbf{Y}} + \sigma_Z(\mathbf{Y}_t)}{\sqrt{(\boldsymbol{\sigma}'(\mathbf{Y}_t) \nabla V_{\mathbf{Y}} + \sigma_Z(\mathbf{Y}_t))' (\boldsymbol{\sigma}'(\mathbf{Y}_t) \nabla V_{\mathbf{Y}} + \sigma_Z(\mathbf{Y}_t))}}. \quad (205)$$

In this expression, the first term takes the standard CIR affine form in which compensation per unit of risk is identical to the volatility of the production technology. The second term is the market price of ambiguity, and its non-linear form implies a (real) yield curve that is not affine. The components of the product $\boldsymbol{\sigma}'(\mathbf{Y}_t) \nabla V_{\mathbf{Y}}$ can change sign over time. Therefore, bond excess returns can take both positive and negative values, consistently with the empirical evidence. Finally, since the component of bond excess returns due to ambiguity is proportional to volatility, while the component due to risk is proportional to variance, the former dominates the latter when aggregate risk in the economy is low. The dependence of the market price of ambiguity on the term $(\boldsymbol{\sigma}'(\mathbf{Y}_t) \nabla V_{\mathbf{Y}} + \sigma_Z(\mathbf{Y}_t))$ is due to non-myopic portfolio behavior: the investor hedges her portfolio against future changes in the worst-case opportunity set, which are in a one-to-one relation to the realizations of the worst-case drift distortion.

Given the expression for the market price of risk and ambiguity, any interest rate derivative can be priced by standard arbitrage arguments. Compared to a completely affine setting, ambiguity aversion alters the fundamental pricing equation through the different equilibrium interest rate process $r_t^f = \mu_Z(\mathbf{Y}_t) - \boldsymbol{\sigma}'_Z(\mathbf{Y}_t)\pi$ and the corresponding change of drift. A simple calibration exercise to U.S. Treasury yield data shows that the model is able to reproduce the deviations from the expectations hypothesis documented in the literature, without modifying in any substantial way the nonlinear mean-reversion dynamics of the short interest rate. In contrast to completely affine models, there is no apparent trade-off between fitting the first and second moments of the yield curve. These findings suggest that a small degree of ambiguity can have large implications for explaining the yield curve.

Ulrich (2008) 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. In bearing this inflation risk premium, investors take the distribution, especially the conditional volatility of inflation, as known. Additionally, 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. The first premium is determined by the product of risk aversion and (negative of) the covariance between inflation and consumption. The second premium is determined by the product of model uncertainty aversion and the volatility of inflation. His paper differs from Gagliardini et al. (2008) because the latter have proposed a real production 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, while Ulrich (2008) models ambiguity as concerning a monetary policy Taylor rule and characterizes the impact on the term structure of both real and nominal bonds.

In Ulrich's model, a representative agent has time separable and logarithmic preferences over consumption holdings c_t and real monetary holdings m_t . The agent holds a capital stock K_t and owns a linear technology A_t which produces an output good Y_t in a linear fashion, $Y_t = A_t K_t$. The exogenous growth rate of the technology follows the process

$$\begin{aligned} \frac{dA_t}{A_t} &= (\mu_A + v_A z_t) dt + \sigma_A \sqrt{z_t} dW_t^z \\ dz_t &= \kappa_z (\theta_z - z_t) dt + \sigma_z \sqrt{z_t} dW_t^z, \end{aligned} \quad (206)$$

with the stochastic technological innovation z_t driven by a stationary square-root process. The agent can invest ζ_P of his wealth in this production technology and the remaining wealth ζ_N 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: the productivity factor z_t , the inflation volatility factor v_t , and the inflation drift factor θ_t . These three state variables make up the state space \mathbf{X}_t of the model. The central bank is assumed to control an additional monetary shock which is denoted $W^{\hat{m}}$ and 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:

$$\max_{c_t, m_t, \zeta_P, \zeta_N} \min_{q^h(\mathbf{X}_t) \in h(\mathbf{X}_t)} E^{q^h(\mathbf{X}_t)} \left\{ \int_t^\infty e^{-\rho s} \left[\ln c_s + \frac{1}{20} \ln m_s \right] ds \middle| \mathcal{F}_t \right\} \quad (207)$$

subject to a standard budget constraint and the entropy bound $0.5h^2(\mathbf{X}_t) \leq \eta(\mathbf{X}_t)$. The agent has a preference for robustness because she insures herself against a worst case inflation distortion, $h(\mathbf{X}_t)$, which the central bank chooses. The entropy constraint specifies that the investor wants to protect herself against conditional inflation drift distortions which are smaller or equal to $\eta(\mathbf{X}_t)$. This upper bound is therefore allowed to be time-varying and stochastic. This is different from Gagliardini et al. (2008) who assume a constant bound. The general structure of the money supply (m_t^S) rule under the ambiguity-free measure p consists of a Taylor-type rule

$$\frac{dm_t^S}{m_t^S} = \omega_t dt + q_1 \left(\frac{dY_t}{Y_t} - \hat{y} dt \right) + q_2 \left(\frac{d\pi_t}{\pi_t} - \hat{\pi} dt \right) + \xi \sigma_m \sqrt{v_t} dW_t^v + \sqrt{1 - \xi^2} \sigma_m \sqrt{v_t} dW_t^{\hat{m}}, \quad (208)$$

where ξ is the correlation between money supply and volatility shocks, and \hat{y} ($\hat{\pi}$) are the growth (inflation) rate objectives of the central bank. Ulrich proves that the optimal degree of inflation distortion that the ambiguity

averse agent takes into account is given by $h^*(\mathbf{X}_t) = \sqrt{2\eta(\mathbf{X}_t)}$. Given this endogenous degree of robustness, the agent adjusts his probability measure and solves a standard maximization problem obtaining the policy functions $c_t^* = (20/21)\rho W_t$ and $m_t^* = (1/21)\rho W_t$. Inflation ambiguity does not affect the policy function for consumption and money demand. Instead, it 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, r_t , and the market price of output risk, $\lambda_{r,Y}(t)$, are not affected by inflation ambiguity. These values coincide with a standard CIR (1985) economy, $r_t = \mu_A + v_A z_t - \delta - \sigma_A^2 z_t$ and $\lambda_{r,Y}(t) = \sigma_A^2 \sqrt{z_t}$. The endogenous process for inflation is affected by the ambiguity on monetary policy

$$\frac{d\pi_t^*}{\pi_t^*} = \mu_\pi(\mathbf{X}_t)dt + \sigma_\pi(\mathbf{X}_t)[dW_t^z \ dW_t^v \ (dW_t^{\hat{m}, h^*(\mathbf{X}_t)} + h^*(\mathbf{X}_t)dt)]', \quad (209)$$

where the endogenous degree of inflation misspecification is given by $\sigma_p(\mathbf{X}_t)h^*(\mathbf{X}_t)$. At this point, it is possible to determine the nominal interest rate $R(t)$ and the nominal market prices of risk $\lambda_{R,Y}(t)$ that nominal assets have to pay. Using these quantities, the (S)EU nominal bond price in an economy that is *not* subject to inflation ambiguity is

$$\hat{N}_t(\tau) = B_t(\tau)E_t^p \left[\frac{\pi_t^*}{\pi_{t+\tau}^*} \right] + Cov_t^p \left[\frac{u_c(c_{t+\tau}^*, m_{t+\tau}^*)}{u_c(c_t^*, m_t^*)}, \frac{\pi_t^*}{\pi_{t+\tau}^*} \right], \quad (210)$$

where $B_t(\tau)$ is a real bond price. $\hat{N}_t(\tau)$ is the sum of the inflation-adjusted real bond price and the covariance between the intertemporal marginal rate of substitution in consumption and the (negative) of the inflation rate, the *inflation risk premium*. The difference between the nominal bond price and the product of real bond price and price deflator is entirely attributed to the inflation risk premium. The closed-form solution for the equilibrium price of an ambiguity-free nominal zero-coupon bond $\hat{N}_t(\tau)$ is given by a log-linear function of the state variables. Ulrich shows that the difference between real ($\hat{N}_t(\tau)$) and nominal ($N_t(\tau)$) bond prices coincides with the inflation ambiguity premium

$$\begin{aligned} \hat{N}_t(\tau) &= N_t(\tau) + Cov_t^p \left[\frac{u_c(c_{t+\tau}^*, m_{t+\tau}^*)}{u_c(c_t^*, m_t^*)} \frac{\pi_t^*}{\pi_{t+\tau}^*}, \frac{dq_{t,t+\tau}^{h^*(\mathbf{X}_t)}}{dp_{t,t+\tau}} \right], \\ \frac{dq_{t,t+\tau}^{h^*(\mathbf{X}_t)}}{dp_{t,t+\tau}} &= \exp \left[-\frac{1}{2} \int_t^{t+\tau} \|h^*(\mathbf{S}_u)\|^2 du + \int_t^{t+\tau} h^*(\mathbf{S}_u) dW_t^{\hat{m}}(u) \right], \end{aligned} \quad (211)$$

where the reference (p -) measure covariance is the ambiguity premium, the covariance between the nominal intertemporal marginal rate of substitution in consumption and the ambiguity kernel (the usual Radon-Nikodym derivative of the q -measure vs. the reference measure). Splitting up this term shows that in general, inflation ambiguity has two channels of impact on the nominal term structure. One is through covariation of the ambiguity kernel with the IMRS and the second is through the covariation of the ambiguity kernel with the price deflator. However, because the first covariance is zero because (by construction) the source of ambiguity, $W^{\hat{m}}$, 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:

$$IAP_t(\tau) = Cov_t^p \left[\frac{\pi_t^*}{\pi_{t+\tau}^*}, \frac{dq_{t,t+\tau}^{h^*(\mathbf{X}_t)}}{dp_{t,t+\tau}} \right] = E^{q^{h^*(\mathbf{X}_t)}} \left[\frac{\pi_t}{\pi_{t+\tau}} \right] - E^p \left[\frac{\pi_t}{\pi_{t+\tau}} \right]. \quad (212)$$

This means that the inflation ambiguity premium equals the distance between the ambiguity-adjusted expected inflation rate and the “true” expected inflation rate.

Ulrich proves an equivalence between the inflation ambiguity premium and the inflation variance premium by assuming that the upper bound of potential inflation misspecifications moves with the volatility of inflation,

$$\eta(\mathbf{X}_t) = \frac{1}{2} \left(q_{a_1} \sqrt{v_t} + \frac{q_{a_2}}{\sqrt{v_t}} + q_{a_3} \frac{z_t}{\sqrt{v_t}} \right)^2, \quad (213)$$

for given q_{a_1} , q_{a_2} , and q_{a_3} , when the market price of inflation ambiguity is $h^*(\mathbf{X}_t) = q_{a_1} \sqrt{v_t} + q_{a_2} (1/\sqrt{v_t}) + q_{a_3} (z_t/\sqrt{v_t})$. The assumed entropy constraint preserves the simple affine bond pricing structure that is known from models like CIR, Vasicek (1977), and others. 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 is a measure for the conditional covariation of inflation and consumption. This covariation is relatively high for a horizon of up to two years, and quickly mean reverts to zero for longer maturities as inflation shocks have no persistent effect on consumption. Similarly to Gagliardini et al. (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. Variations in nominal yields are mostly driven by changes in expected inflation and in the inflation ambiguity premium and not by changes in the consumption risk premium.

Ilut (2010) is an example of ambiguity paper addressing one of the most important outstanding empirical puzzles in international finance, the uncovered interest rate parity (UIP) puzzle. There is a UIP puzzle because while according to a range of standard, rational expectations (S)EU models, high interest rate currencies should depreciate vis-a-vis low interest rate currencies, in practice it has been widely observed that the opposite tends to occur (see e.g., Hansen and Hodrick, 1980; Fama, 1984; Engel, 1996). Although the standard approach to the puzzle has been to assume rational expectations but to model (often exogenously) time-varying risk premia, an empirical literature based on survey data from the foreign exchange market has found significant evidence against the RE assumption and challenged the idea that sensible patterns of time variation in risk premium may be sufficient to explain the puzzle. Ilut proposes instead a story built around the notion of ambiguity aversion: in equilibrium, agents invest in the higher interest rate currency (investment currency) by borrowing at the lower interest rate currency (funding currency) because they entertain the possibility that the data could have been generated by various sequences of time-varying signal to noise ratios.

Ilut (2010) proposes a typical MPP framework in which an ambiguity averse agent solves a signal extraction problem under ambiguity on the precision (variance) of the signals she receives, similarly to Epstein and Schneider (2007). The only source of uncertainty in the environment is the domestic/foreign interest rate differential, $r_t = i_t - i_t^*$, which is modelled as an exogenous stochastic process given by the sum of persistent and transitory components, both unobserved. Under ambiguity aversion with a max-min preference functional, 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 the 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, derive utility from end-of-life wealth. There is one good for which purchasing power parity (PPP) holds: $p_t = p_t^* + e_t$, where p_t is the log of price level of the good in the Home country and e_t is the log of the nominal exchange rate defined as the price of the Home currency per unit of foreign currency. Investors are born at time t , they are risk-neutral over end-of-life wealth, W_{t+1} , and face a convex cost of capital. Their preferences are

$$V_t = \max_{b_t} \min_{q \in \varphi} E_t^q [W_{t+1} - \frac{c}{2} b_t^2], \quad (214)$$

where b_t is the amount of foreign bonds purchased, and c parameterizes their cost of capital. The set φ comprises the alternative probability distributions entertained by the agents, who perceive the following state-space representation,

$$r_t = x_t + \sigma_{v,t} v_t \quad x_t = \rho x_{t-1} + \sigma_u u_t \quad (215)$$

where v_t and u_t are both Gaussian white noise, $\sigma_{v,t}$ are draws from a set Υ , and x_t is the hidden state. In practice, the investors' problem is:

$$V_t = \max_{b_t} \min_{\sigma_V^*(r^t) \in \sigma_V^t} E_t^q [W_{t+1} - \frac{c}{2} b_t^2]. \quad (216)$$

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 for the true data generating process. Under these assumptions, it is able to solve the model in closed-form,

$$e_t = \frac{E_t^{q^*}(e_{t+1}) - r_t}{1 + 0.5c},$$

where $E_t^{q^*}(e_{t+1})$ is the worst-case scenario expectation of future exchange rates. 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 driving 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 since it implies that agents perceive on average positive innovations in updating the estimate, which creates the possibility of further increases in the demand of the investment currency and its consequent gradual appreciation.

5. Ambiguity and Market Microstructure Research

A relatively new but promising sub-field 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, such as market makers. 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), heterogeneous agents framework, J investors display preferences (with exponential utility index and unit risk aversion parameter) à 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 the set of priors that represents her preferences. 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 such a 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 be in the interval $[\lambda_l, \lambda^h]$. Hence, each of the investors holds a set of beliefs about the future value of the risky asset and she trades away from the status quo (her current 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

$$\bigcap_{1,2,\dots,J} [\lambda_l \mu_j - \sigma^2 \omega_{0j}, \lambda^h \mu_j - \sigma^2 \omega_{0j}] \neq \emptyset, \quad (217)$$

then no trades will occur and a market will freeze. Moreover, a sufficient condition for trade to collapse is $\{\max_{1,2,\dots,J} [\mu_j] / \min_{1,2,\dots,J} [\mu_j]\} < [(1 - \lambda_l) / (1 - \lambda^h)]$. 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 relatively 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, Easley and O'Hara show that:

$$\begin{aligned} \text{ask} &= \min_j [\lambda^h \mu_j - \sigma^2 \alpha_{0j}] = \hat{\mu} - \sigma^2 x - \max_j [(1 - \lambda^h) \mu_j] \\ \text{bid} &= \max_j [\lambda_l \mu_j - \sigma^2 \alpha_{0j}] = \hat{\mu} - \sigma^2 x - \min_j [(1 - \lambda_l) \mu_j]. \end{aligned} \quad (218)$$

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 (see e.g., the literature reviewed in O'Hara, 1994). 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 and there is no learning from prices. 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. This may be clearly useful in valuation applications of the model, when only bid and ask quotes are observed, but no transactions occur. Clearly, fluctuating levels of ambiguity and its heterogeneity—as captured by the quantities $\max_{1,2,\dots,J} [\mu_j] / \min_{1,2,\dots,J} [\mu_j]$ and $(1 - \lambda_l)/(1 - \lambda^h)$ —will cause fluctuating levels of liquidity.

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 to conclude that—albeit this remains intuitively sensible—ambiguity will

decrease liquidity and/or be the cause of market break-downs. Their basic motivation is the identification of financial crises with the lack of liquidity that typically follows extreme market outcomes. Rutledge and Zin observe that these breakdowns in liquidity seem particularly acute in markets where traders heavily rely on complex and yet uncertain empirical models for the cash flows processes, as it happens for instance in derivative markets. This evidence suggests a link between ambiguity and liquidity, since severe reductions in the provision of liquidity seem to result from model uncertainty. 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 (whose 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, $\Pr(\delta_t = 1) = \Phi(b_t)$. Apart from the trading activity, 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.

For concreteness, let's examine a two-period framework. For given choices of a , b , and d , at time 1 the market maker solves the problem

$$\max_{\omega} \left\{ U(\vartheta_0 - \omega S_0) + \beta \min_{p \in \wp} E_p [U(\vartheta_1 + \omega S_1)] \right\}$$

$$\vartheta_0 = y_0 + \begin{cases} a & \text{if } d = -1 \\ 0 & \text{if } d = 0 \\ -b & \text{if } d = 1 \end{cases} \quad W_1 = y_1 + dD(S_1), \quad (219)$$

where β is a discount factor, y_0 is the exogenous income deriving from other trading operations, and \wp is the set of probability distributions under consideration. d is the dividend paid by the risky asset, and ϑ is the income (profit stream) of the market-maker in period t . Hence, $U(a, b, d)$ is the maximal (indirect) utility for given (a, b, d) . At time 0, optimal bid and ask prices are chosen by solving:

$$\max_{a, b} \{[1 - \Phi(a)]U_a + [\Phi(a) - \Phi(b)]U_b + \Phi(b)U_b\}$$

$$U_a = (a, b, -1) \quad U_b = (a, b, 1) \quad U_0 = (a, b, 0) \quad (220)$$

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 actually 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 $U_0 - U_{a=0}$ —where $U_{a=0}$ is the total indirect utility from a short position in the derivative when no period 0 ask price is received—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 produce markets that are less liquid.

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 ambiguous beliefs (represented by MPP) about the probability distribution of the asset payoff. 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. Interestingly, while in the learning model by Epstein and Schneider (2009) investors have unambiguous beliefs prior to the arrival of information and their posterior beliefs become ambiguous because of the signals, for the informed traders in Ozsoylev and Werner’s model, ambiguity disappears because of the arrival of the signal. 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 (S)EU 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. When comparing the standard (S)EU-case with the ambiguity framework, their results show that ambiguity leads to high liquidity risk, lower trading volume, and lower market depth.

Another phenomenon that has attracted the interest of microstructure researchers is herding, the tendency of traders to simultaneously implement similar strategies, which may cause severe market disruptions. For instance, Ford, Kelsey, and Pang (2005) have proposed a simple herding model where full rationality is relaxed by considering agents’ attitudes towards informational ambiguity introduced through neo-additive capacities (as in Chateauneuf et al., 2002). In their model, trades are assumed to occur sequentially and prices are determined endogenously at each trade. Neither market makers nor traders know exactly the value of the assets. In particular, market makers form their expectations on the basis of public information, while traders receive private information that enables them to update public beliefs.

Ford et al. show that a consistent signal reduces confidence less than an inconsistent one. Intuitively, although the arrival of enough public information will improve informational efficiency in decisions, the arrival of an inconsistent (and yet informative) signal on public disclosure may, paradoxically, make decisions worse. In particular, as argued in Bikhchandani, Hirshleifer, and Welch (1992), in the process of aggregating the information of fewer individuals, further additional information can encourage agents to fall into a cascade sooner, so there is no presumption that the signal will improve decisions. First, Ford, Kelsey, and Pang (2005) assume asymmetric ambiguity, meaning that only traders and not market makers display ambiguity aversion. This asymmetry is shown to rationally generate history-dependent behavior and contrarian behavior different from those described by the traditional (S)EU literature on herding. Next,

they allow for ambiguity averse market makers. Under this condition, ambiguity is responsible for herding. Further, herding only leads to short-run mispricing in financial markets since informational efficiency holds in the long-run.

6. Policy Implications

Ambiguity may be responsible for such far-reaching effects on optimal financial decisions and equilibrium asset prices, that it is natural to conjecture that a literature must have taken an active interest in the implications of ambiguity for macroeconomic policies, financial regulation, and market and security design. In this section we turn our attention to a few such examples, even though many Readers will have found already quite natural to stop and ponder the effects of ambiguity for the optimal design of policy and markets in the course of our earlier review.

Easley and O’Hara (2009, 2010) have discussed how a carefully designed market microstructure can reduce the negative effects of ambiguity. In their model with mixed populations of (S)EU and ambiguity-averse investors, Easley and O’Hara (2009) have shown that deposit insurance, NYSE and NASD suitability rules, Rule 16(b) of the Securities Exchange Act, and the Investment Act of 1940 can play a surprisingly effective role by inducing ambiguity-averse investors to participate in the banking and financial system.⁶³ Interestingly, Easley and O’Hara have stressed that these rules/provisions effectively limit themselves to rule out extremely negative events which have very low probability to occur, so that the rules become rarely binding. Nonetheless, their existence can induce higher participation by simply convincing investors to disregard extremely pessimistic scenarios.

Easley and O’Hara (2010) have used these intuitions to ask when will it be the case that regulated markets emerge in equilibrium over un-regulated (say, over-the-counter) markets. Within the framework proposed in Easley and O’Hara (2009), they consider two stock markets, *A* and *B*, which differ from each other in terms of the services they provide to listing firms and investors: Market *B* is simply a trading platform, while Market *A* is an exchange that provides a certification service that can reduce company-specific ambiguity by limiting which firms it allows to trade. Specifically, the exchange examines companies that apply to be listed and only agrees to list those that meet some minimum standards.⁶⁴ For ambiguity averse investors, this corresponds to a higher minimum expected return (μ_{\min}) and a lower maximum variance (σ_{\max}^2) of the stock’s distribution. Each firm must choose a market on which to list its stock. Listing in Market *A* implies a per-share fee c , which is deduced from the investors’ payoff. In this case the ambiguity averse investors’ demand function ω_A for the risky assets in market *A* will be given by (27) where μ_{\min} and μ_{\max} are replaced by $\mu_{\min} - c$ and $\mu_{\max} - c$, while the risky asset demand ω_B in the simple platform OTC *B* is exactly as in (27). Easley and O’Hara (2010) show that if Market *B* is characterized by limited participation—when only the (S)EU investors trade—the equilibrium price will be $p_B^{(S)EU} = \hat{\mu} - \hat{\sigma}^2 x / (1 - \alpha)$, while it will be

$$p_B^{AA} = \frac{\alpha \hat{\sigma}^2 \mu_{\min} + (1 - \alpha) \sigma_{\max}^2 \hat{\mu} - \sigma_{\max}^2 \hat{\sigma}^2 x}{(1 - \alpha) \sigma_{\max}^2 + \alpha \hat{\sigma}^2} \quad (221)$$

in case of full participation, where α is the percentage of investors that are ambiguity averse. In market *A*, these

⁶³Rule 16(b) requires an insider, defined as an officer, director, or 10% shareholder, to return any profits made on a trading position held for less than six months. The Investment Act of 1940 limits a mutual fund to holding no more than 5% of its assets in any one issue to enforce some minimal degree of diversification.

⁶⁴Practical examples are listing rules, delisting rules, trading halts, affirmative obligations of market makers, transparency requirements, price collars and daily limits, public comes first rules, clearing house rules and margin requirements, fast market rules, etc.

expressions are different: $p_A^{(S)EU} = \hat{\mu} - c - \hat{\sigma}^2 x / (1 - \alpha)$ if there is limited participation, and

$$p_A^{AA} = \frac{\alpha \hat{\sigma}^2 (\mu_{\max} - c) + (1 - \alpha) \sigma_{\max}^2 (\mu_{\min} - c) - \sigma_{\max}^2 \hat{\sigma}^2 x}{(1 - \alpha) \sigma_{\max}^2 + \alpha \hat{\sigma}^2} \quad (222)$$

in the case of full participation.

These expressions offer powerful insights. Because a firm will list its stock where the price is higher, and given that it is trivial to show that if Market A is in limited participation equilibrium so is Market B , in the case of limited participation no firm will choose to list on Market A . This means that when ambiguity averse traders choose to stay on the sidelines, the “bad” market will completely crowd out the “good” market, which is an obvious result, as (S)EU traders would have to pay the cost of a regulated market but would derive no benefits because they do not suffer from ambiguity aversion. This situation is very likely to occur when α is small so that nearly all investors are (S)EU maximizers, as to them the guarantees offered by Market A have no value. With few new potential investors that might be induced to trade by listing a stock on Market A , there is little or no gain to the stock price and consequently all firms will trade on Market B . On the contrary, when nearly all investors are ambiguity averse (as $\alpha \rightarrow 1$), then, if the stock is to be publicly traded at all, equilibrium requires a share price low enough to attract the ambiguity averse investors. Firms will list on Market A only if the increase in the minimum mean payoff that A offers is larger than the regulation cost c . For an intermediate composition of the population of traders (say α away from both 0 and 1), the listing decision will depend on the relative costs and benefits in each market. The greater (lower) α is, the more likely it is that listing on A (B) will dominate as the equilibrium outcome. Additionally, firms with moderate ambiguity (measured by small differences $\mu^n - \mu_{\min}^n$ and $\sigma_{\max}^n - \sigma^n$) will benefit from paying the listing fees to access Market A , provided that the certification is sufficient to attract ambiguity averse traders.⁶⁵

In a similar set up, Guidolin and Rinaldi (2009) have focussed on the policy options available to improve the equilibrium outcomes of asset markets populated by a fraction α of MPP-type traders. The first type of policy that may be pursued is to try and affect the proportions of ambiguity averse (AA) and (S)EU-maximizers which compose the economy, α and $1 - \alpha$.⁶⁶ Guidolin and Rinaldi’s results make it clear that the effect of such policies will depend on the initial configuration as represented by the equilibrium price function and the level of participation. For instance, suppose that initially both groups of investors participate in the market, so that the equilibrium price is p^{AA} . Even though the policy maker finds a way to instantaneously decrease the fraction of AA agents, α , because $\sigma_I^2 \leq \sigma_{\max}^2$ and $\mu_{\min} < \mu_I$, the numerator of the expression for p^{AA} increases; viceversa its denominator declines. Therefore, on the whole p^{AA} increases. However, as this happens it is possible for the condition $(\mu_{\min} + \mu_S)/(1 + \pi) > p^{AA}$ to start being violated and p^{AA} may enter the region $((\mu_{\min} + \mu_S)/(1 + \pi), (\mu_{\max} + \mu_S)/(1 + \pi))$. When this occurs, participation from the AA investors will cease and the pricing function will switch to $p^{(S)EU} > p^{AA}$. Therefore, an attempt to decrease α and simply increase the proportion of (S)EU-maximizers represents a policy approach with mixed outcomes. On the one hand, it is certain that such an intervention will increase equilibrium prices, which may

⁶⁵Viceversa, if ambiguity is low, investors may opt to participate even if the stock trades on Market B . Because if Market B is in full participation equilibrium so is Market A , firms will choose to list their stocks in Market B . For intermediate situations, listing on Market A is chosen provided that the certification is necessary and sufficient to attract ambiguity averse traders and $p_A^{AA} > p_B^{AA}$. Finally, there exists a threshold size x below (above) which firms choose to be listed in Market B (A).

⁶⁶Causing changes in α means that AA investors are selectively given (or taken away) enough information for her perception of ambiguity to disappear. We interpret the notion of “taking information away” as identical to “confusing” the investors on purpose, which—as counterintuitive as this may sound—in a MPP environment may produce some benefits.

be attractive for firms because this will lower the cost of capital. On the other hand, it may end up penalizing the overall participation to the market and reduce liquidity and trading volume.⁶⁷ The intuition is that if very few AA agents are left in the market, the equilibrium price will mostly reflect the fundamental risk assessments expressed by (S)EU investors, and the resulting price will end being “too high” for AA investors to participate.

Suppose instead that initially only (S)EU investors were present in the market. If α is reduced, the price increases and that makes the stock even less attractive to AA investors. As a result the limited participation equilibrium would not be affected and the risk premium would decline. However, if a policy-maker tried to increase α in such an environment, this will lead to a decrease in the equilibrium price, so that if α increases enough, the equilibrium price will fall below the threshold $(\mu_{\min} + \mu_S)/(1 + \pi)$, and both groups will be willing to trade in the stock market, which will actually increase the stock price. Guidolin and Rinaldi (2009) compute a threshold $\bar{\alpha}$ such that when α is raised sufficiently, then the equilibrium switches from limited participation by (S)EU-maximizers only to both agent types. These results reveal a clear asymmetry in the ability to use policy to affect equilibrium outcomes. While increasing α is always possible and it implies that while equilibrium prices decline, more participation may be ultimately attained, decreasing α has non-linear effects: a threshold exists, such that if α declines below such threshold and the equilibrium price increases enough, then the market may switch to an (S)EU-only participation regime and this may cause a sudden drop in trading and liquidity.

Guidolin and Rinaldi also discuss a policy that simply consists—similarly to Easley and O’Hara (2009)—of convincing the ambiguity averse investors that certain sensitive scenarios can be ruled out. For instance, if the policy makers manage to increase μ_{\min} and only (S)EU agents were initially trading, it is possible to switch to a new equilibrium with full participation when μ_{\min} is raised enough to exceed the threshold $\mu_I - \bar{z}(\sigma_I^2 + \sigma_S^2) / (1 - \alpha)$. In that case, the equilibrium price will jump down from $p^{(S)EU}$ to p^{AA} and the risk premium will increase. The implication is that starting from situations of (S)EU-only participation, it will take large jumps in $\bar{\mu}_{\min}$ for the equilibrium to be significantly affected; however, when this happens, one can also expect a detrimental effect on risky asset prices. This is surprising: even though the policy action consists of ruling out the worst possible scenarios increasing μ_{\min} , for an action of sufficient magnitude its eventual effect on equilibrium prices will be negative and the cost of enforcing a participation equilibrium will consist of a higher risk premium. This means that when a market has fallen into a state of disruption (a (S)EU-only equilibrium), bringing it back to higher liquidity and orderly functioning through a reduction in the amount of perceived ambiguity may actually go through further reductions in equilibrium prices. However, it is also possible to prove that when initially a full participation equilibrium rules, p^{AA} will increase continuously as μ_{\min} is increased while increasing μ_{\min} has no effect on the participation constraint, which will still be satisfied. In the presence of full participation, reducing the amount of ambiguity has only positive effects on prices and will reduce risk premia. It seems then that the key objective of any policy maker ought to prevent equilibria with limited participation to appear in the first instance.

One final implication of the analysis in Guidolin and Rinaldi (2009) is that—assuming policy makers have full control over the realized inflation rate (which is quite optimistic)—it is impossible for governments to try and inflate

⁶⁷Notice that there is no contradiction in this claim: starting from a situation in which 100% of the potential investors (including the AA fraction α_0) participate to the market trading stocks, by sufficiently decreasing α (say, to $\alpha_0 - \Delta\alpha$) the policy-maker may induce a switch to an equilibrium in which only $1 - \alpha_0 + \Delta\alpha < 1$ of the potential investors participate. Unless $\Delta\alpha = -\alpha_0$, this implies a loss in participation.

their way out of a situation characterized by market breakdowns. When the inflation rate allowed to surge, it always (i.e., independently of whether the initial equilibrium was characterized by limited participation) strengthens the segmentation and produces ambiguous effects on risky asset prices. Intuitively, this happens when inflation becomes so high that—given a perception of uncertainty as represented by σ_{\max}^2 —the real expected payoffs from the risky asset become insufficient to compete with the real, riskless rate of return guaranteed to AA investors by the money market account. Therefore, inflation as a policy tool seems either ineffective or perverse because it cannot relax participation constraints while it may depress real equilibrium prices. On the contrary, a substantial reduction of the inflation rate—below a threshold $\bar{\pi} = -1 + \bar{z} (\sigma_I^2 + \sigma_S^2) / [(1 - \alpha) (\mu_I - \mu_{\min})]$ —may lead to increased participation, but lower asset prices. Clearly, if the economy were to start from an equilibrium of full participation, there is no chance that the participation constraint may stop being satisfied: if both (S)EU and AA agents were initially trading in the market, this remains the case in a regime with lower inflation. Therefore low inflation has only virtues.

Easley and O’Hara (2009, 2010) and Guidolin and Rinaldi (2009) have insisted on examining either market design issues or the effectiveness of policy interventions that try and remedy to a financial crisis characterized by shrinking liquidity. A handful of other papers have used ambiguity to explain phenomena that lie at the intersection between finance and macroeconomics. For instance, Mukerji and Tallon (2000) have developed an ambiguity model to motivate the little practical use of indexation in financial securitization. In particular, they extend the results in Mukerji and Tallon (1999) by specifying what forms of idiosyncratic risk may lead to trading breakdowns in equilibrium. For an indexed asset, such an idiosyncrasy is represented by the presence of “irrelevant” goods (irrelevant in the sense that the agents trading the asset neither consume nor are endowed with them) in the indexation bundle. In a two-period competitive monetary general equilibrium model for goods, bonds and money market, agents display CEU-preferences. Nominal and indexed bonds are both available for trade and prices of all goods and bonds are determined endogenously. A first group of agents (say, agents of type *A*) trades on financial markets, while a second group (say, agents of type *B*) has no access to any financial markets, and, therefore, it simply consumes all the revenue from endowment. Three goods are available, x , y , and z . Agents of type *A* consume only goods x and z , while agents of type *B* consume only goods y and z . Agents *A* have real endowments composed of goods x and z . Agents *B* have real endowments composed of goods y and z . Uncertainty is realized at the beginning of the second period, agents 1’s endowment is uncontingent in both periods, while the endowment of agents 2 is contingent in the second period for good y . Money supply in period 0 is fixed at m^0 , but may take on two values in period 1, say m_1^l or m_1^h , with $m_1 = \lambda_1 m^0$, where $\lambda > 1$ is the inflation rate, that can be high or low. Finally, there are two financial assets traded at time 0: the first is a nominal bond that pays off one unit of money in all states; the second is an indexed bond, that pays off a bundle of goods in each state in period 1. Agents do not know the probability distribution over Ω ; nevertheless they are given a convex capacity v over it. Without ambiguity, for generic first period aggregate endowments, there is trade in the indexed bond whenever $\lambda \neq 1$. However, in the presence of ambiguity, there exists a bound $\bar{\lambda} > 1$, such that, if $\lambda < \bar{\lambda}$, there exists \bar{A} , $0 < \bar{A} < 1$, such that, if $\mathcal{A}(v) > \bar{A}$, in equilibrium the indexed bond is not traded while the nominal bond is. Hence, when the indexation bundle contains at least one good which is not consumed by any of the agents who actively trade in the financial market, and beliefs about the change in the price of these goods relative to the average price level are sufficiently ambiguous, trade of indexed bonds does not occur in equilibrium, provided that inflation is bounded above. Additionally, it does not matter whether ambiguity concerns also nominal bonds. These

implications are consistent with the fact that typically, indexed bonds are traded almost exclusively under extreme inflationary circumstances, as reflected by the necessary condition $\lambda < \bar{\lambda}$.

7. Discussion and Conclusions

In this paper we have reviewed and discussed a number of works that have brought models of decision making under ambiguity (sometimes also perceived as a preference for robust decisions) at the forefront of research in financial economics. A number of insights have emerged. Of course, a lot of work remains to be done. On the one hand, a few papers (e.g., Gollier, 2009, or Chapman and Polkovnichenko, 2009) have cautioned that we should hesitate before generalizing some of the implications of models of ambiguity outside the tight parametric set ups in which these results have been derived. This appears particularly important in the perspective of policy-makers and market regulators that are facing a complex reality that will hardly adhere to the restrictive but simplifying assumptions exploited in the literature to derive crisp results. Therefore, an obvious avenue for future research consists of generalizations—for instance using flexible or general models for the stochastic processes driving the state variables, and rich asset menus that may also shed light on the cross-sectional, cross-asset structure of risk premia and volatility dynamics—of the existing portfolio choice and asset pricing models. For instance, we are currently blessed with a range of insights derived from models in which a single state variables follows a Geometric Brownian motion and in which the asset menu is simply composed of a single risky asset and one riskless bond. There is an obvious need for extensions and generalizations. However, let us stress that it is comforting to notice that different Authors using different preference specifications that capture ambiguity aversion (e.g., MPP, smooth KMM preferences, but also robustness preferences) have all delivered results for asset allocation and asset pricing that are qualitatively similar. On the other hand, the literature appears to be still in a stage in which the need to approach technically complicated problems and to deal with deep logical issues (e.g., how beliefs should be updated in dynamic models) has 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 GE models, but have usually kept the asset menu fixed to two assets only. However, a few papers that have derived implications for either large-scale portfolio decisions or for cross-sectional asset pricing have had to focus on large asset menus, but often in simple static models or giving up the elegance (ultimately, necessity) of GE models. Similarly, most of the policy and regulatory implications from the literature emerge from simple, two-period models. One wonders how safe policy implementation may be when complicated intertemporal effects (if not dynamic games between investors and policy-makers) are ignored. Therefore it is legitimate to expect that in the future new and even more ambitious papers will pursue more complicated and interconnected questions using dynamic GE 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 and for which it is legitimate to expect intense efforts at extending what we currently understand. Even though it is often the case that the most important breakthroughs may happen exactly where they are not expected, in our minds it appears 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 (S)EU; Mele and Sangiorgi (2009) have studied the flip side of the same coin, i.e., on the incentives for information acquisition in markets under ambiguity. Although Section 3 has highlighted a range of implications for asset allocation decisions that are robust to the specifics of the parametric models proposed, it is difficult to deny a need to work on general frameworks for the process of asset returns—for instance, nesting within each other the cases of no predictability, of linear predictability, as well as of non-linear predictability (as in Garcia, Detemple and Rindisbacher, 2002)—able to derive general results for the effects of ambiguity aversion on myopic vs. hedging demands. Finally, while a limited number of researchers straddling the macroeconomics and finance camps (e.g., Backus et al., 2004, Hansen and Sargent, 2007, and more recently Ulrich, 2009) have made efforts at structuring their ambiguity models to allow easy connections to existing estimation methods for dynamic GE asset pricing models, an overwhelming majority of the papers we have reviewed still content themselves with solving models, and then simulating empirical outcomes from calibrated versions. Although this has been useful to allow researches from all backgrounds to understand what the qualitative implications and salient mechanisms of ambiguity aversion are, models of ambiguity offer rich sets of empirical restrictions—often (especially in the case of MPP, see Epstein and Schneider, 2010) qualitatively different, for instance of a non-linear type, from (S)EU—that should be tested giving us a possibility to reject the proposition that ambiguity aversion affects portfolio/consumption choices and are priced in equilibrium.

A different issue is whether and how the ever growing class of financial models we have reviewed in this paper has already impacted the way in which applied (financial) economists, policy-makers, and commentators perceive and 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 (S)EU perspective had been addressed by a few of the founding fathers of economics and decision theory alike, among them Arrow, Hurwicz, Keynes, Knight, Raiffa, and Wald (in no order whatsoever)—even though research on decision making under ambiguity has witnessed an acceleration during the 1990s, informal arguments connecting poorly understood phenomena to aversion to uncertainty have been appearing among both academics and practitioners at least 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 and have started impacting the way in which applied economists view markets. For instance, Epstein and Schneider (2008) have used their model to show that taking ambiguity into account may go a long way towards explaining price movements—especially large market slides—that would otherwise represent genuine puzzles in standard asset pricing models. They propose an intriguing explanation of the behavior of the stock prices around the 9/11 events, in 2001. The analysis is interesting because a standard Bayesian model with known signal quality would have enormous difficulties at explaining the initial slide in prices on 9/11 and the following days. The terrorist attack increased uncertainty about future economic growth and news about terrorism and foreign policy, which were less important in the minds of traders earlier on in the month of September, suddenly became crucial. It therefore became difficult to assess the importance to assign to information flows during that week. Under ambiguity the signal precision is unknown, and bad news is treated as extremely reliable, so that a much less extreme sequence of signals is sufficient to account for extremely low stock prices, as those registered in the first week after the attack.

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, we think that the reaction of a number of central banks and governments to the great financial crisis of 2008-2009 showcases features that may be as a minimum interpreted as being consistent with the typical actions that may effectively contrast the disruptive effects of ambiguity. For instance, with reference to papers such as Easley and O’Hara (2009) and Guidolin and Rinaldi (2009), a handful of policy measures may be interpreted as attempts at lowering μ_{\min} and/or σ^2_{\max} , i.e., to reduce the “amount” of ambiguity. For instance, in the Spring of 2008 the U.S. Treasury announced a temporary guarantee of the share prices of money market mutual funds and, beginning in October 2008, it used authority granted under the 2008 Emergency Economic Stabilization Act to purchase preferred shares in a large number of depository institutions. Similar policies have been enacted in Canada and in a number of European countries.

As far as market and security designs are concerned, it is clear that contracting cannot solve the participation problems induced by ambiguity aversion, because it would involve probability specifications for states of the world whose very existence would make complete contracting prohibitively expensive. A related difficulty concerns the enforcement of optimal contracts, since, even if complete contracting were feasible, enforcement would not be automatic. Similarly, disclosure cannot solve ambiguity-induced problems because of the difficulty in attaching a probability to unlikely states. Indeed, under ambiguity aversion, disclosure can even exacerbate problems by scaring away potential market participants. Further, the palliative role of arbitrage is limited because ambiguity aversion induces non-participation, and not merely mispricing. Interestingly, regulation may play an important role in ruling out aberrant outcomes caused by ambiguity aversion. As stressed by Easley and O’Hara (2009) with reference to exchanges’ rule books, there are a myriad of rules and requirements, some so arcane as to be rarely, if ever, actually binding in practice. Yet, such obscurity is perfectly consistent with the ambiguity-resolving role explained in Section 6.

References

- [1] Allais, M. (1953). “Rational Man’s Behavior in the Presence of Risk: Critique of the Postulates and Axioms of the American School.” *Econometrica*, 21, 503-46.
- [2] Al-Najjar, N.I., and J., Weinstein (2009). “The Ambiguity Aversion Literature: A Critical Assessment.” *Economics and Philosophy*, 25, 249-284.
- [3] 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.
- [4] Anderson, E.W., L.P., Hansen, and T.J., Sargent (1998). “Risk and Robustness in General Equilibrium.” Working paper, University of Chicago.
- [5] Anscombe, F., and R., Aumann (1963). “A Definition of Subjective Probability.” *Annals of Mathematical Statistics*, 34, 199-205.
- [6] Arrow, K.J. (1965). *Aspects of the Theory of Risk-Bearing*. Helsinki: Academic Bookstore.
- [7] Arrow, K.J. and L., Hurwicz (1972). “An Optimality Criterion for Decision Making Under Ignorance,” in *Uncertainty and Expectations in Economics*, by C., Carter and J., Ford (editors), pp. 1-11, Blackwell Press.

- [8] Aumann, R.J. (1962). "Utility Theory without the Completeness Axiom." *Econometrica*, 30, 445-462.
- [9] Autore D., R., Billingsley and M., Schneller (2009). "Information Uncertainty and Auditor Reputation." *Journal of Banking and Finance*, 33, 183-192.
- [10] Barberis, N. (2000). "Investing for the Long Run when Returns Are Predictable." *Journal of Finance*, 55, 225-264.
- [11] 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 (editors), volume 1, chapter 18, pp. 1053-1128, Elsevier.
- [12] Barberis, N., A., Shleifer and R., Vishny (1998). "A Model of Investor Sentiment." *Journal of Financial Economics*, 49, 307-345.
- [13] Barillas, F., L.P., Hansen, and T.J., Sargent (2009). "Doubts or Variability?" *Journal of Economic Theory*, 144, 2388-2418.
- [14] Becker, S., and F., Brownson (1964). "What Price Ambiguity? Or the Role of Ambiguity in Decision-Making." *Journal of Political Economy*, 72, 62-73.
- [15] Benigno P., and S., Nistico. (2009) "International Portfolio Allocation under Model Uncertainty." NBER Working Paper No. 14734.
- [16] Bewley, T. (1986). "Knightian Decision Theory, Part 1." Working Paper, Yale University.
- [17] Bikhchandani, S., D., Hirshleifer and I., Welch (1992). "A Theory of Fads, Fashion, Custom, and Cultural Change in Informational Cascades." *Journal of Political Economy*, 100, 992-1026.
- [18] Billot, A., A., Chateauneuf, I., Gilboa, and J.-M., Tallon (2000). "Sharing Beliefs: Between Agreeing and Disagreeing." *Econometrica*, 68, 685-694.
- [19] Black, F., and M., Scholes (1973). "The Pricing of Options and Corporate Liabilities." *Journal of Political Economy*, 81, 637-659.
- [20] 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.
- [21] Boyarchenko, N. (2009). "Ambiguity, Information Quality and Credit Risk." Working Paper, University of Chicago.
- [22] 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.
- [23] Breeden, D.T. (1979). "An Intertemporal Asset Pricing Model with Stochastic Consumption and Investment Opportunities." *Journal of Financial Economics*, 7, 265-296.
- [24] Cagetti, M., L.P., Hansen, T., Sargent, and N., Williams (2002). "Robustness and Pricing with Uncertain Growth." *Review of Financial Studies*, 15, 363-404.
- [25] Camerer, C., and M., Weber (1992). "Recent Developments in Modeling Preferences: Uncertainty and Ambiguity" *Journal of Risk and Uncertainty*, 5, 235-370.

- [26] Camerer, C., and T., Ho (1991). "Isolation Effects in Compound Lottery Reduction." Working paper, University of Pennsylvania.
- [27] Campbell, J.Y., and R.J., Shiller (1988). "Stock Prices, Earnings, and Expected Dividends." *Journal of Finance*, 43, 661-676.
- [28] Campbell, J.Y. (1998). "Asset Prices, Consumption and the Business Cycle." NBER Working Paper No. 6485.
- [29] Campbell, J.Y., A.W., Lo, and A.C., MacKinlay (1996). *The Econometrics of Financial Markets*. Princeton, NJ: Princeton University Press.
- [30] 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.
- [31] Caskey, J. (2009). "Information in Equity Markets with Ambiguity Averse Investors." *Review of Financial Studies*, 22, 3595-3627.
- [32] Cerreia-Vioglio, S., F., Maccheroni, M., Marinacci and L., Montrucchio (2008). "Risk Measures: Rationality and Diversification." Carlo Alberto Notebooks No. 100, Collegio Carlo Alberto, Turin.
- [33] Chacko, G., and L.M., Viceira (2005). "Dynamic Consumption and Portfolio Choice with Stochastic Volatility in Incomplete Markets." *Review of Financial Studies*, 18, 1369-1402.
- [34] Chapman, D.A., and V., Polkovnichenko (2009). "First-Order Risk Aversion, Heterogeneity, and Asset Market Outcomes." *Journal of Finance*, 64, 1863-1887.
- [35] Chateauneuf, A., and J., Faro (2009): "Ambiguity Through Confidence Functions." *Journal of Mathematical Economics*, 45, 535-558.
- [36] Chateauneuf, A., J., Eichberger and S., Grant (2003). "A Simple Axiomatization and Constructive Representation Proof for Choquet Expected Utility." *Economic Theory*, 22, 907-915.
- [37] Chateauneuf, A., J., Eichberger, and S., Grant (2007). "Choice Under Uncertainty with the Best and Worst in Mind: Neo-additive Capacities." *Journal of Economic Theory*, 137, 538-567.
- [38] Chateauneuf, A., R., Dana, and J.-M., Tallon (2000). "Optimal Risk Sharing Rules and Equilibria with Choquet-Expected-Utility." *Journal of Mathematical Economics*, 34, 191-214.
- [39] Chen, H., N., Ju and J., Miao (2009). "Dynamic Asset Allocation with Ambiguous Return Predictability." Working Paper, Boston University.
- [40] Chen, Z., and L.G., Epstein (2002). "Ambiguity, Risk, and Asset Returns in Continuous Time." *Econometrica*, 70, 1403-1443.
- [41] Chernoff, H. (1952). "A Measure of Asymptotic Efficiency for Tests of a Hypothesis Based on the Sum of Observations." *Annals of Mathematical Statistics*, 23, 493-507.
- [42] Cochrane, J.H. (2005). "Financial Markets and the Real Economy." NBER Working Paper No. 11193.
- [43] Cochrane, J.H. (2005). *Asset Pricing*. Princeton, NJ: Princeton University Press.

- [44] Condie S., and J., Ganguli (2009). "Ambiguity and Partially-Revealing Rational Expectations Equilibria." Working Paper, University of Cambridge.
- [45] Coval, J., and T., Moskowitz (1999). "Home Bias at Home: Local Equity Preference in Domestic Portfolios." *Journal of Finance*, 54, 2045-2073.
- [46] Cox, J.C., J.E., Ingersoll and S.A., Ross (1985). "A Theory of the Term Structure of Interest Rates" *Econometrica*, 53, 385-407.
- [47] Daniel, K., D., Hirshleifer and A., Subrahmanyam (1998). "Investor Psychology and Security Market Under- and Overreactions." *Journal of Finance*, 53, 1839-1885.
- [48] David, A. (1997). "Fluctuating Confidence in Stock Markets: Implication for Returns and Volatility." *Journal of Finance and Quantitative Analysis*, 32, 427-462.
- [49] David, A (2008). "Heterogeneous Beliefs, Speculation, and the Equity Premium." *Journal of Finance*, 63, 41-83.
- [50] Backus, D., B., Routledge and S., Zin (2004). "Exotic Preferences for Macroeconomists." *NBER Macroeconomics Annual*, 19, 319-390.
- [51] De Bondt, W., and R., Thaler (1985). "Does the Stock Market Overreact?" *Journal of Finance*, 40, 793-808.
- [52] de Finetti, B. (1931) "Sul Significato Soggettivo della Probabilità." *Fundamenta Mathematicae*, 17, 298-329.
- [53] DeMiguel, V., L., Garlappi and R., Uppal (2009). "Optimal Versus Naive Diversification: How Inefficient is the 1/N Portfolio Strategy?" *Review of Financial Studies*, 22, 1915-1953.
- [54] Detemple, J.B., R., Garcia and M., Rindisbacher (2003). "A Monte Carlo Method for Optimal Portfolios." *Journal of Finance*, 58, 401-446.
- [55] Dothan, U., and D., Feldman (1986). "Equilibrium Interest Rates and Multiperiod Bonds in a Partially Observable Economy." *Journal of Finance*, 61, 369-382.
- [56] Dow, J., and S.R., Werlang (1992). "Uncertainty Aversion, Risk Aversion, and the Optimal Choice of Portfolio." *Econometrica*, 60, 197-204.
- [57] Dow, J., and C., Werlang (1994) "Learning Under Knightian Uncertainty: The Laws of Large Numbers for Non-additive Probabilities." Working Paper, London Business School.
- [58] Duffie, D., and L.G., Epstein (1992). "Asset Pricing with Stochastic Differential Utility." *Review of Financial Studies*, 5, 411-436.
- [59] Duffie, D., and L.G., Epstein (1992). "Stochastic Differential Utility." *Econometrica*, 60, 353-394.
- [60] Easley, D., and M., O'Hara (2010). "Liquidity and Valuation in an Uncertain World." *Journal of Financial Economics*, 97, 1-11.
- [61] Easley, D., and M., O'Hara (2009) "Ambiguity and Nonparticipation: The Role of Regulation." *Review of Financial Studies*, 22, 1817-1843.
- [62] Easley, D., and M., O'Hara (2010) "Microstructure and Ambiguity." *Journal of Finance*, Forthcoming.
- [63] Ellsberg, D. (1961). "Risk, Ambiguity, and the Savage Axioms." *Quarterly Journal of Economics*, 75, 643-669.

- [64] Engel, C. (1996). “The forward discount anomaly and the risk premium: A survey of recent evidence” *Journal of Empirical Finance*, 3, 123-192.
- [65] Epstein, L.G., and T., Wang (1995). “Uncertainty, Risk-Neutral Measures and Security Price Booms and Crashes.” *Journal of Economic Theory*, 67, 40-80.
- [66] Epstein, L.G. (1999). “A Definition of Uncertainty Aversion.” *Review of Economic Studies*, 66, 579-608.
- [67] Epstein, L.G. (2001). “Sharing Ambiguity.” *American Economic Review*, 91, 45-50.
- [68] Epstein, L.G. and J., Miao (2003). “A Two-Person Dynamic Equilibrium under Ambiguity” *Journal of Economic Dynamics and Control*, 27, 1253-1288.
- [69] Epstein, L.G., and M., Schneider (2003). “Recursive Multiple Priors.” *Journal of Economic Theory*, 113, 1-31.
- [70] Epstein, L.G., and M., Schneider (2007). “Learning Under Ambiguity.” *Review of Economic Studies*, 74, 1275-1303.
- [71] Epstein, L.G., and M., Schneider (2008). “Ambiguity, Information Quality, and Asset Pricing.” *Journal of Finance*, 63, 197-228.
- [72] Epstein, L.G. and M. Schneider (2010). “Ambiguity and Asset Markets,” NBER Working Paper No. 16181.
- [73] Epstein, L.G., and T., Wang (1994). “Intertemporal Asset Pricing under Knightian Uncertainty.” *Econometrica*, 62, 283-322.
- [74] Epstein, L.G. (2010). “Three Paradoxes for the ‘Smooth Ambiguity’ Model of Preference.” Working Paper, Boston University.
- [75] Ergin, H., and F., Gul (2004). “A Subjective Theory of Compound Lotteries.” Working Paper, Princeton University.
- [76] 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.
- [77] Fama, E.F. (1984). “Forward and Spot Exchange Rates.” *Journal of Monetary Economics*, 14, 319-338.
- [78] Fama, E.F., and K.R., French (1988a). “Dividend Yields and Expected Stock Returns.” *Journal of Financial Economics*, 22, 3-25.
- [79] Fama, E.F., and K.R., French (1988b). “Permanent and Temporary Components of Stock Prices,” *Journal of Political Economy*, 96, 246-73.
- [80] Fama, E.F., and K.R., French (1989). “Business Conditions and Expected Returns on Stocks and Bonds.” *Journal of Financial Economics*, 25, 23-49.
- [81] Fama, E.F., and K.R., French (1992). “The Cross-Section of Expected Stock Returns.” *Journal of Finance*, 47, 427-65.
- [82] 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.

- [83] Fine, T. L., and P., Walley (1982). "Towards a Frequentist Theory of Upper and Lower Probability." *The Annals of Statistics*, 10, 741-761.
- [84] Fishburn, P. C. (1993). "The Axioms and Algebra of Ambiguity." *Theory and Decision*, 34, 119-137.
- [85] Ford, J. L., D., Kelsey and W., Pang (2005). "Ambiguity in Financial Markets: Herding and Contrarian Behaviour." Working Paper, University of Exeter.
- [86] French, K.R., and J.M., Poterba (1991) "Investor Diversification and International Equity Markets." *American Economic Review*, Papers and Proceedings, 81, 222-226.
- [87] Frydman, R. (1982). "Towards an Understanding of Market Processes, Individual Expectations, Market Behaviour, and Convergence of Rational Expectations Equilibrium." *American Economic Review*, 72, 652-668.
- [88] Gagliardini, P., P., Porchia and F., Trojani (2009). "Ambiguity Aversion and the Term Structure of Interest Rates." *Review of Financial Studies*, 22, 4157-4188.
- [89] Gajdos, T., T., Hayashi, J.-M., Tallon and J.-C., Vergnaud (2008) "Attitude Toward Imprecise Information." *Journal of Economic Theory*, 140, 23-56.
- [90] 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.
- [91] Gennotte, G. (1986). "Optimal Portfolio Choice Under Incomplete Information." *Journal of Finance*, 41, 733-746.
- [92] Ghirardato, P., and M., Marinacci (2002). "Ambiguity Made Precise: A Comparative Foundation." *Journal of Economic Theory*, 102, 251-289.
- [93] Ghirardato, P., F., Maccheroni, and M., Marinacci. (2004) "Differentiating Ambiguity and Ambiguity Attitude," *Journal of Economic Theory*, 118, 133-173.
- [94] Gilboa I., and D., Schmeidler (1989). "Maxmin Expected Utility with Non-Unique Prior." *Journal of Mathematical Economics*, 18, 141-153.
- [95] Gilboa, I. (1989). "Duality in Non-Additive Expected Utility Theory." *Annals of Operations Research*, 19, 405-414.
- [96] Gilboa, I., A. W., Postlewaite and D., Schmeidler (2008) "Probability and Uncertainty in Economic Modeling." *Journal of Economic Perspectives*, 22, 173-188.
- [97] Gilboa, I. and D., Schmeidler (1991). "Updating Ambiguous Beliefs." Working paper, Northwestern University.
- [98] Gollier, C. (2009). "Portfolio Choices and Asset prices: The Comparative Statics of Ambiguity Aversion." Working Paper.
- [99] Guidolin, M. (2005). "Home Bias and High Turnover in an Overlapping-generations Model with Learning." *Review of International Economics*, 13, 725-756.
- [100] Guidolin, M. (2006). "High Equity Premia and Crash Fears, Rational Foundations." *Economic Theory*, 28, 693-709.

- [101] Guidolin, M. (2006). "Pessimistic Beliefs under Rational Learning: Quantitative Implications for the Equity Premium Puzzle." *Journal of Economics and Business*, 58, 85-118.
- [102] 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.
- [103] Guidolin, M., and A., Timmermann (2007). "Asset Allocation under Multivariate Regime Switching." *Journal of Economic Dynamics and Control*, 31, 3503-3544.
- [104] Hansen, L.P. (2007). "Beliefs, Doubts and Learning: Valuing Macroeconomic Risk." *American Economic Review*, 97, 1-30.
- [105] Hansen, L.P., and R. J., Hodrick. (1980). "Forward Exchange Rates as Optimal Predictors of Future Spot Rates: An Econometric Analysis." *Journal of Political Economy*, 88, 829-853.
- [106] Hansen, L.P., and R., Jagannathan (1991). "Implications of Security Market Data for Models of Dynamic Economies." *Journal of Political Economy*, 99, 225-262.
- [107] Hansen, L.P., and T. J., Sargent (2001). "Robust Control and Model Uncertainty." *American Economic Review*, 91, 60-66.
- [108] Hansen, L P., T. Sargent, and T. Tallarini (1999). "Robust permanent income and pricing." *Review of Economic Studies*, 66, 873-907.
- [109] Hansen, L.P., and T.J., Sargent (2007). *Robustness*. Princeton University Press: Princeton, NJ.
- [110] Hansen, L.P., and T.J., Sargent (1998). "Alternative Representations of Discounted Robust Linear Quadratic Control." Working Paper, University of Chicago.
- [111] Hansen, L.P., T.J., Sargent, G., Turmuhambetova, and N., Williams (2006). "Robustness and Model Misspecification." *Journal of Economic Theory*, 128, 45-90.
- [112] Harvey, C.R., and A., Siddique (2000). "Conditional Skewness in Asset Pricing Tests." *Journal of Finance*, 55, 1263-1295.
- [113] Hirshleifer, D. (2001). "Investor Psychology and Asset Pricing", *Journal of Finance*, 56, 1533-1597.
- [114] Honda, T. (2003). "Optimal Portfolio Choice for Unobservable and Regime-Switching Mean Returns." *Journal of Economic Dynamics and Control*, 28, 45-78.
- [115] Hurwicz, L. (1951). "Optimality Criteria for Decision Making Under Ignorance." Discussion Paper, Statistics No. 370, Cowles Commission.
- [116] Illeditsch, P. (2009). "Ambiguous Information, Risk Aversion and Asset Pricing." Working Paper, University of Pennsylvania.
- [117] Ilut, C. L. (2010). "Ambiguity Aversion: Implications for the Uncovered Interest Rate Parity Puzzle." Working Paper, Duke University.
- [118] Ju, N., and J., Miao (2008). "Ambiguity, Learning, and Asset Returns." Working Paper, Hong Kong University of Science and Technology.

- [119] Kantorovich, L.V. (1939). *Mathematical Methods in the Organization and Planning of Production*.
- [120] Keynes, J.M. (1921). *A Treatise on Probability*. London: Macmillan.
- [121] Kleshchelski, I., and N. Vincent (2009). “Robust Equilibrium Yield Curves.” HEC Montreal, Working Paper.
- [122] Klibanoff, P., M., Marinacci, and S., Mukerji (2005). “A Smooth Model of Decision Making under Ambiguity.” *Econometrica*, 73, 1849-1892.
- [123] Klibanoff, P., M., Marinacci and S., Mukerji. (2009) “Recursive Smooth Ambiguity Preferences.” *Journal of Economic Theory*, 144, 930-976.
- [124] Knetsch, J. (1989). “The Endowment Effect and Evidence of Nonreversible Indifference Curves.” *American Economic Review*, 79, 1277-1284.
- [125] Knight, F.H. (1921). *Risk, Uncertainty, and Profit*. Boston and New York: Houghton Mifflin.
- [126] Kocherlakota, N. (1996). “The Equity Premium: It’s Still a Puzzle.” *Journal of Economic Literature*, 34, 42-71.
- [127] Kogan, L., and T., Wang (2003) “A Simple Theory of Asset Pricing under Model Uncertainty.” Massachusetts Institute of Technology, Working Paper.
- [128] Kogan, L., and R., Uppal. (2000) “Risk Aversion and Optimal Policies in Partial and General Equilibrium Economies.” Massachusetts Institute of Technology, Working Paper.
- [129] Kreps, D.M., and E.L., Porteus (1978). “Temporal Resolution of Uncertainty and Dynamic Choice Theory.” *Econometrica*, 46, 185-200.
- [130] Leippold M., F., Trojani, and P., Vanini (2008). “Learning and Asset Prices Under Ambiguous Information.” *Review of Financial Studies*, 21, 2565-2597.
- [131] LeRoy, S.F., and R.D., Porter (1981) “The Present-Value Relation: Tests Based on Implied Variance Bounds.” *Econometrica*, 49, 555-574.
- [132] Lewellen, J., and J., Shanken (2002). “Learning, Asset-Pricing Tests, and Market Efficiency.” *Journal of Finance*, 57, 1113-1145.
- [133] Lewis, K. (1999) “Trying to Explain Home Bias in Equities and Consumption.” *Journal of Economic Literature*, 37, 571-608.
- [134] Liu, H. (2010) “Robust Consumption and Portfolio Choice for Time Varying Investment Opportunities.” Working Paper, Manchester Business School.
- [135] 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.
- [136] Lo, K. (1998). “Sealed Bid Auctions with Uncertainty Averse Bidders.” *Economic Theory*, 12, 1-20.
- [137] Lucas, R.E. (1978). “Asset Prices in an Exchange Economy.” *Econometrica*, 46, 1429-1445.
- [138] Maccheroni, F., M., Marinacci and A., Rustichini (2006). “Ambiguity Aversion, Robustness, and the Variational Representation of Preferences.” *Econometrica*, 74, 1447-1498.

- [139] Maccheroni, F., M., Marinacci, A., Rustichini and M., Taboga (2009). “Portfolio Selection with Monotone Mean-Variance Preferences.” *Mathematical Finance*, 19, 487-521.
- [140] Maenhout, P.J. (2004). “Robust Portfolio Rules and Asset Pricing.” *Review of Financial Studies*, 17, 951-983.
- [141] Marinacci, M. (1996). “Limit Laws for Non-Additive Probabilites and Their Frequentist Interpretation.” Working Paper, University of Turin.
- [142] Marinacci, M. (1999). “Limit Laws for Non-Additive Probabilites and Their Frequentist Interpretation.” *Journal of Economic Theory*, 84, 145-195.
- [143] Markowitz, H.M. (1959). *Portfolio Selection: Efficient Diversification of Investments*. Wiley, New York.
- [144] Marschak, J. (1975). “Personal Probabilities of Probabilities.” *Theory and Decision*, 6, 121-153.
- [145] Mehra, R., and E.C., Prescott. (1985). “The Equity Premium: A Puzzle.” *Journal of Monetary Economics*, 15, 145-161.
- [146] Mele, A., and F., Sangiorgi. (2009) “Ambiguity, Information Acquisition and Price Swings in Asset Markets.” Working Paper, London School of Economics.
- [147] Merton, R.C. (1969). “Lifetime Portfolio Selection under Uncertainty: The Continuous-Time Case.” *Review of Economics and Statistics*, 51, 247-257.
- [148] Merton, R.C. (1971). “Optimum Consumption and Portfolio rules in a Continuous-Time Model.” *Journal of Economic Theory*, 3, 373-413.
- [149] Miao, J. (2009). “Ambiguity, Risk and Portfolio Choice under Incomplete Information.” *Annals of Economics and Finance*, 10, 257-279.
- [150] Miao, J., and N., Wang (2006). “Risk, Uncertainty, and Option Exercise.” Boston College, Working Paper.
- [151] Mukerji, S. (1998). “Ambiguity Aversion and Incompleteness of Contractual Form.” *American Economic Review*, 88, 1207-1231.
- [152] Mukerji, S., and J.-M., Tallon (2001) “Ambiguity Aversion and Incompleteness of Financial Markets.” *Review of Economic Studies*, 68, 883-904.
- [153] Mukerji, S., and J.-M., Tallon (2003): “Ellsberg’s 2-Color Experiment, Portfolio Inertia and Ambiguity.” *Journal of Mathematical Economics*, 39, 299-316.
- [154] 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.
- [155] Nau, R. (2006). “Uncertainty Aversion with Second-Order Utilities and Probabilities.” *Management Science*, 52, 136-145.
- [156] Neilson, W.S. (2010). “A Simplified Axiomatic Approach to Ambiguity Aversion.” *Journal of Risk and Uncertainty*, 41, 113-124.

- [157] Ozsoylev, H., and J., Werner. (2009). "Liquidity and Asset Prices in Rational Expectations Equilibrium with Ambiguous Information." Working Paper, University of Minnesota.
- [158] Pastor, L. (2000). "Portfolio Selection and Asset Pricing Models." *Journal of Finance* 55, 179-223.
- [159] Pflug, G., and D., Wozabal (2007). "Ambiguity in Portfolio Selection." *Quantitative Finance*, 7, 435-442.
- [160] Poterba, J.M., and L.H., Summers (1988). "Mean Reversion in Stock Prices: Evidence and Implications." *Journal of Financial Economics*, 22, 27-59.
- [161] Quiggin, J. (1982). "A Theory of Anticipated Utility." *Journal of Economic Behavior and Organization*, 3, 323-343.
- [162] Routledge, B.R., and S.E., Zin (2009). "Model Uncertainty and Liquidity." *Review of Economic Dynamics*, 12, 543-566.
- [163] Samuelson, W., and R., Zeckhauser (1988). "Status Quo Bias in Decision Making." *Journal of Risk and Uncertainty*, 1, 7-59.
- [164] Savage, L. (1954). *The Foundations of Statistics*. Wiley, New York.
- [165] 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.
- [166] Schmeidler, D. (1989). "Subjective Probability and Expected Utility without Additivity." *Econometrica*, 57, 571-587.
- [167] Segal, U. (1987). "The Ellsberg Paradox and Risk Aversion: an Anticipated Utility Approach." *International Economic Review*, 28, 175-202.
- [168] Segal, U. (1990). "Two-Stage Lotteries without the Reduction Axiom." *Econometrica*, 58, 349-377.
- [169] Seo, K. (2009). "Ambiguity and Second-Order Belief." *Econometrica*, 77, 1575-1605.
- [170] Shiller, R.J. (1981) "Do Stock Prices Move Too Much to be Justified by Subsequent Changes in Dividends?" *American Economic Review*, 71, 421-436.
- [171] Shleifer, A., and R., Vishny (1997). "The Limits of Arbitrage." *Journal of Finance*, 52, 35-55.
- [172] Shleifer, A. (2000). *Inefficient Markets: An Introduction to Behavioral Finance*, Oxford University Press.
- [173] Skiadas, C. (2003). "Robust Control and Recursive Utility." *Finance and Stochastics*, 7, 475-489.
- [174] Starmer, C. (2000). "Developments in Non-Expected Utility Theory: The Hunt for a Descriptive Theory of Choice Under Risk." *Journal of Economic Literature*, 38, 332-382.
- [175] Strzalecki, T. (2007). "Axiomatic Foundations of Multiplier Preferences." Working Paper, Harvard University.
- [176] Tallarini, T.D. (2000). "Risk-sensitive real business cycles." *Journal of Monetary Economics*, 45, 507-532.
- [177] Timmermann, A. (1996). "Excess Volatility and Predictability of Stock Prices in Autoregressive Dividend Models with Learning." *Review of Economic Studies*, 63, 523-557.
- [178] Timmermann, A. (1993). "How Learning in Financial Markets Generates Volatility and Predictability in Stock Prices." *Quarterly Journal of Economics*, 108, 1135-1145.

- [179] Townsend, R. (1978). "Market Anticipations Rational Expectations, and Bayesian Analysis." *International Economic Review*, 19, 481-494.
- [180] 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.
- [181] Trojani, F., and P., Vanini (2004). "Robustness and Ambiguity Aversion in General Equilibrium." *Review of Finance*, 8, 279-324.
- [182] Tversky, A., and D., Kahneman (1992). "Advances in Prospect Theory: Cumulative Representation of Uncertainty." *Journal of Risk and Uncertainty*, 5, 297-323.
- [183] Uhlig, H. (2009). "A Model of a Systemic Bank Run." Working Paper, University of Chicago.
- [184] Ulrich, M. (2008). "Inflation ambiguity and the term structure of arbitrage-free U.S. Government bonds.", New York University, Working Paper.
- [185] Uppal, R. and Wang, T. (2003). "Model Misspecification and Underdiversification." *Journal of Finance*, 58, 2465-2486.
- [186] Vasicek, O. (1977) "An Equilibrium Characterization of the Term Structure." *Journal of Financial Economics*, 5, 177-188.
- [187] Veronesi, P. (1999). "Stock Market Overreaction to Bad News in Good Times: A Rational Expectations Equilibrium Model." *Review of Financial Studies*, 5, 975-1007.
- [188] Veronesi, P. (2000). "How Does Information Quality Affect Stock Returns?" *Journal of Finance*, 55, 807-837.
- [189] von Neumann, J., and O., Morgenstern (1947). *Theory of Games and Economic Behavior*. Princeton, NJ: Princeton University Press.
- [190] Wakker, P.P. (2001). "Testing and Characterizing Properties of Nonadditive Measures through Violations of the Sure-Thing Principle." *Econometrica*, 69, 1039-1059.
- [191] Wakker, P.P. (2008). "Uncertainty." in The New Palgrave Dictionary of Economics, second edition, S., Durlauf and L., Blume (editors) Palgrave Macmillan.
- [192] Wald, A. (1950). *Statistical Decision Functions*. Wiley, New York.
- [193] Wang, Z. (2005). "A Shrinkage Approach to Model Uncertainty and Asset Allocation." *Review of Financial Studies*, 18, 673-705.
- [194] Weil, P. (1989) "The Equity Premium Puzzle and the Risk-Free Rate Puzzle." *Journal of Monetary Economics*, 24, 401-421.
- [195] Yates, F. and L. Zukowski (1976). "Characterization of Ambiguity in Decision Making." *Behavioral Science*, 21, 19-25.