NBER WORKING PAPER SERIES

MODELING UNCERTAINTY AS AMBIGUITY: A REVIEW

Cosmin L. Ilut
Martin Schneider

Working Paper 29915
http://www.nber.org/papers/w29915

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
April 2022

We thank Rudi Bachmann, Federico Bennett, Lautaro Chittaro, Florian Fiaux, Robert Garlick, Monika Piazzesi, Giorgio Topa, Rosen Valchev and Wilbert van der Klaauw for helpful comments. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Modeling Uncertainty as Ambiguity: a Review
Cosmin L. Ilut and Martin Schneider
NBER Working Paper No. 29915
April 2022
JEL No. D8,E2,E3,E4,G1

ABSTRACT

We survey literature on ambiguity with an emphasis on recent applications in macroeconomics and finance. Like risk, ambiguity leads to cautious behavior and uncertainty premia in asset markets. Unlike risk, ambiguity can generate first order welfare losses. As a result, precautionary behavior and ambiguity premia obtain even when agents have linear utility and are reflected in linear approximations to model dynamics. Quantitative work exploits this insight to estimate models that jointly match the dynamics of asset prices and macro aggregates. In micro data, inertia and inaction due to ambiguity help understand patterns such as non-participation in asset markets, price rigidities and simple contracts. Learning under ambiguity generates asymmetric responses to news that help connect higher moments in micro and macro data. Survey evidence is increasingly used to provide direct evidence on ambiguity averse behavior, as well as to discipline quantitative models.

Cosmin L. Ilut
Department of Economics
Duke University
223 Social Sciences Building
Box 90097
Durham, NC 27708
and NBER
cosmin.ilut@duke.edu

Martin Schneider
Department of Economics
Stanford University
579 Jane Stanford Way
Stanford, CA 94305-6072
and NBER
schneidr@stanford.edu

1 Introduction

To introduce ambiguity, consider a simple thought experiment. Suppose we are uncertain about whether sea level rise will exceed one foot by the year 2050. Alternative models of climate change generate very different probabilities of this event. How should we value insurance in the face of such model uncertainty? A Bayesian with standard expected utility preferences assigns probabilities to all relevant events: he forms a prior over the different models and arrives at a prediction by averaging probability forecasts. Suppose for concreteness that this exercise delivers a conditional probability of one third.

Now suppose additional evidence leads all climate change models to *agree* on a probability of one third for the event. In other words, scientific progress eliminates model uncertainty. For the Bayesian, this new information does not affect beliefs or behavior: model averaging still delivers a probability of one third, and hence the same valuation of insurance. However, a complete lack of response is not obviously intuitive: new information should make us more confident in our assessment of the event, and this should be reflected in behavior as well. The expected utility model does not allow for a response since model uncertainty and uncertainty about events described by models are both treated as risk, that is, uncertainty with known odds. Reduction of compound lotteries then implies the same probability of one third in both scenarios.

The term ambiguity refers to uncertainty when the odds are not known. Ambiguity averse agents are not confident enough to assign probabilities to events about which they have little information, such as what is the right model of climate change. The Ellsberg (1961) Paradox, reviewed below, illustrates that the distinction between ambiguity and risk is behaviorally meaningful; agents make choices that are inconsistent with the Bayesian model because they want to avoid ambiguity. Axiomatic models of rational choice consistent with the Ellsberg Paradox were developed in the 1980s. The key innovation was to represent preferences using sets of beliefs (Gilboa and Schmeidler (1989)). Uncertainty about probability assessments can then be described by the size of the belief set, similarly to how model uncertainty is handled in a related literature on robust statistics (for example, Huber 2009).

When ambiguity averse agents evaluate a plan, they act *as if* they are relatively pessimistic about the plan, and more so the larger their belief set. Continuing our example, an ambiguity averse agent confronted with disagreement among climate change models has a large set of beliefs. An agent who is adversely exposed to sea level rise evaluates insurance as if sea level rise is relatively likely. In contrast, an agent who stands to gain from sea level rise behaves as if climate change has little effect. This endogenous adjustment of pessimism generates cautious behavior that is different in nature from what the Bayesian model can capture. In particular, news that all models agree collapses the set of beliefs to one, and behavior becomes less cautious.

Early applications of ambiguity aversion were mostly confined to finance, following the seminal papers on portfolio choice by Dow and Werlang (1992) and on asset pricing by Epstein and Wang (1994). The early 2000s then saw a surge of research on theoretical foundations for dynamic choice under ambiguity. Over the last decade, we have seen new applications in many fields of economics. Our goal in this chapter is to provide an overview of this work. We first compare models of ambiguity aversion and the standard Bayesian approach in four contexts: a consumption-savings problem, a portfolio choice problem, an asset pricing Euler equation, and a dynamic stochastic general equilibrium model of the business cycle. We then survey how the literature has used ambiguity aversion to understand aggregate fluctuations, asset pricing puzzles, decisions of heterogeneous agents in micro data, and optimal policy.

We emphasize three themes throughout the chapter. First, since ambiguity and risk are both concepts of uncertainty, models of risk aversion and ambiguity aversion naturally share many properties. Uncertainty of either flavor lowers welfare. Agents' desire to avoid uncertainty is a key force in asset markets: higher uncertainty can explain why some assets are cheaper than others, and why prices are lower in certain periods such as recessions. More generally, the desire to avoid uncertainty guides intertemporal decisions towards precautionary behavior: higher uncertainty leads to more savings, less hiring of workers, or less issuance of debt. Time variation in uncertainty can therefore generate economic fluctuations that coincide with drops in asset prices.

Second, a key difference between ambiguity and risk is that only the former can generate first order welfare effects of uncertainty. Intuitively, an ambiguity averse agent with a set of beliefs that differ in mean evaluates plans *as if* the worst case mean is low; uncertainty does not matter only for higher moments. There are two implications for applied work. One is technical: models with ambiguity-averse agents allow for precautionary savings, uncertainty premia and excess volatility of asset prices, but can still be characterized using standard *linear* approximation and estimation techniques that are common in macroeconomics. In particular, when ambiguity is assumed to move with the volatility of shocks, effects of stochastic volatility can also be studied with linear models.

First order welfare effects of ambiguity further generate inaction or inertia from uncertainty alone. Intuitively, an agent contemplating positive (negative) exposure to ambiguity evaluates plans based on low (high) mean payoffs and chooses zero exposure if both are sufficiently unfavorable. This feature leads to parsimonious explanations for a number of robust facts in micro data - for example, non-participation in asset markets that an investor is not familiar with, rigidity and memory in nominal prices, or lack of adoption of new technologies. We note that market frictions, technological rigidities or curvature in utility is not required for the argument. This is contrast to results on inaction due to risk: for example, "wait-and-see" effects require irreversibility of investment, which induces curvature in the objective function.

Our third theme is that models of ambiguity aversion offer a way to *quantify* uncertainty that is conceptually more appealing and more tractable than standard practice, yet allows for similar discipline in specifying preferences. To elaborate, the standard approach to quantify models in macroeconomics and finance (i) estimates stochastic processes for exogenous fundamentals, such as productivity shocks or dividends, and (ii) assumes rational expectations: the estimated processes not only describe physical laws of motion, but also beliefs used to calculate agents' expected utility. As a result, agents in the model do not share researchers' doubts about model specification or the parameter uncertainty revealed by estimation. With agents that confident about the world, it is not surprising that the standard approach has run into many puzzles.

Quantitative models with ambiguity aversion capture doubts about model specification and parameter uncertainty through belief sets: they avoid puzzles by making agents less confident. Additional degrees of freedom are pinned down using survey data as well as model consistency criteria that relate actual and worst case laws of motion. At the same time, two attractive features of standard rational expectations modeling are retained. In particular, belief sets are typically parametrized by an interval of means centered around econometricians' point estimates: as in standard models, agents do not make systematic forecast errors. Moreover, agents have structural knowledge of the economy: they understand how endogenous variables depend on shocks as well as policy parameters and form belief sets accordingly. As in standard models, policy evaluation thus is not subject to the Lucas critique.

The rest of the chapter is structured as follows. Section 2 introduces preferences in a two period setting and studies savings and portfolio choice. Section 3 considers intertemporal preferences and studies consumption-based asset pricing and business cycle models. Section 4 reviews work that quantifies ambiguity using survey data. Section 5 covers aggregate models, Section 6 collects work that makes predictions for micro data and Section 7 is about optimal policy. A conclusion flags open questions.

2 Static choice under uncertainty

We start with a two period setup and define preferences that reflect risk or ambiguity aversion. Our focus on applications leads us to discuss preferences only in terms of utility functions, as opposed to their axiomatic foundations. We refer the reader to Epstein and Schneider (2010) or Gilboa and Marinacci (2016) for a discussion of the main representation results.

2.1 Preferences

There are two dates 1 and 2 and one consumption good. At date 2, one of a finite number of states $\omega \in \Omega$ is realized. At date 1, an agent chooses not only date 1 consumption c, but also consumption $\tilde{c}(\omega)$ in each state ω at date 2. We refer to the pair $(c, \tilde{c})$ as a consumption *plan.* For all models we consider, preferences over consumption plans can be represented by

$$u(c) + \beta u(CE(\tilde{c})) \tag{1}$$

Here the function CE returns the certainty equivalent of uncertain date 2 consumption $\tilde{c}$, that is, the number $\bar{c}$ such that a plan with constant consumption $\bar{c}$ rather than $\tilde{c}(\omega)$ in all states ω at date 2 is indifferent to $(c, \tilde{c})$. The shape of CE reflects whether agents are averse to risk or ambiguity. The function u ranks certain consumption plans; throughout, it is smooth, strictly increasing and concave. The agent's willingness to substitute intertemporally depends on the curvature of u. It can be measured, for example, by the intertemporal elasticity of substitution (IES) $\sigma(c) = -u'(c)/cu''(c)$.

Risk aversion. Let the probability P on the state space Ω denote the agent's subjective belief. In what follows we use superscripts to distinguish between expectations formed under different beliefs, for example, $\mathbb{E}^P$ is the expectation under P. We define the certainty equivalent for the case of risk aversion by

$$v(CE(\tilde{c})) = \mathbb{E}^P[v(\tilde{c})] \tag{2}$$

where v is a strictly increasing and concave function.

Risk aversion is captured by curvature in v. Strict concavity of v implies $CE(\tilde{c}) < \mathbb{E}^P[\tilde{c}]$, that is, the certainty equivalent is valued less than mean consumption under the subjective belief. In this sense, the agent dislikes risky consumption plans. The degree of risk aversion can be measured, for example, by the coefficient of relative risk aversion $\gamma(c) = -cv''(c)/v'(c)$. In general, we allow $u \neq v$, so risk tolerance – the willingness to substitute across states – is distinct from the willingness to substitute over time, as in the large applied literature following Epstein and Zin (1989). In the special case $u = v$, we obtain the familiar time separable expected utility model $u(c)+\beta\mathbb{E}^{P}[u(\tilde{c})]$.

The Ellsberg Paradox. People prefer to know the odds of an event. In Ellsberg's thought experiment, such preference is revealed by the ranking of bets. To illustrate, consider bets that pay one dollar if an event occurs and zero otherwise. Compare two events: a risky event is known to occur with probability one half ("known odds"), whereas no information is given about the other, ambiguous, event ("unknown odds"). It then makes sense to prefer a bet on the risky event to a bet on the ambiguous event. This is because we can be confident that the first bet is indeed fair, whereas the second bet might not be. By symmetry, it also makes sense to prefer a bet *against* the risky event – a bet that pays one dollar when it does not occur and zero otherwise – to a bet against the ambiguous event.

This intuitive "Ellsberg-type" behavior is incompatible with a single subjective belief on the state space. Indeed, a decision maker whose choice under uncertainty is represented by (1)-(2) has in mind a single probability not only for the risky event but also for the ambiguous event. Strict preference for a bet on the risky event therefore reveals that the subjective probability of the ambiguous event is smaller than one half. At the same time, strict preference for a bet against the risky event implies a probability of the ambiguous event larger than one half, a contradiction.

We note that the Ellsberg paradox not only reveals a problem with the standard expected utility model ($u=v$), but more broadly with any model that represents belief with a single probability, for example, the more general class (1)-(2) that distinguishes risk aversion from the willingness to smoothe over time. This is why models of ambiguity aversion turn to sets of probabilities in order to capture Ellsberg-type behavior. We also emphasize that the paradox can be understood as a *rational* response to a lack of confidence in probability assessments – it is not limited to inexperienced subjects in the lab. Relatedly, it is not crucial that the odds of the risky event are exactly known, it is enough that we are confident to think about it in terms of probabilities.

Multiple priors. Let $\mathcal{P}$ denote a set of beliefs P on Ω. For the *multiple priors* model, introduced by Gilboa and Schmeidler (1989), the certainty equivalent is

$$v(CE(\tilde{c}))=\min_{P\in\mathcal{P}}\mathbb{E}^{P}[v(\tilde{c})]. \tag{3}$$

Agents with multiple priors preferences deal with ambiguity by acting *as if* they were using the most pessimistic belief from $\mathcal{P}$, a belief that we denote generically by P^{*}. In the special case where the set contains only a single probability, we are back to the risk aversion case (2). More generally, the size of the belief set describes agents' lack of confidence in probability assessments: a larger set reflects more ambiguity.

The multiple priors model is consistent with Ellsberg-type behavior because the "worst-case belief" P^* endogenously varies with the consumption plan. To see this, fix a state space and consider a set of beliefs that agree on the probability of an event S^r, say, but disagree on an event S^a such that $\min_{P\in\mathcal{P}} P(S^a) < P(S^r) < \max_{P\in\mathcal{P}} P(S^a)$. Here S^r represents a risky event to which the agent can attach a probability, whereas S^a is an ambiguous event that may or may not be more likely than S^r. An agent then prefers a bet on S^r to a bet on S^a because he evaluates the bet on S^a using $\min_{P\in\mathcal{P}} P(S^a) < P(S^r)$. At the same time, he prefers a bet against S^r to a bet against S^a since he evaluates the latter using the worst-case probability $1 - \max_{P\in\mathcal{P}} P(S^a) < 1 - P(S^r)$.

The Gilboa-Schmeidler axioms imply a representation (3) with *some* (subjective) set of beliefs $\mathcal{P}$. In applications, the set $\mathcal{P}$ is thus a primitive chosen by the modeler to describe behavior under ambiguity, much like the subjective belief P above is chosen to describe behavior under risk. One extreme choice is to let $\mathcal{P}$ contain only measures P_ω that put probability one on a single state ω, that is, agents consider a set of event realizations and evaluate plans by their performance under the worst-case *realization*. However, nothing in the theory restricts attention to that extreme: elements of $\mathcal{P}$ need not be degenerate. In applications, it often makes sense to make the set relatively small in order to fit agents' willingness to pay for insurance against ambiguity. As the example above shows, even a very small set generates choice that differs from expected utility.

Other models of ambiguity aversion. Several other models of preference have been developed to account for Ellsberg-type behavior; a detailed survey is Machina and Siniscalchi (2013). The most prominent models in applied work are multiplier preferences, proposed by Anderson et al. (2003) and axiomatized by Strzalecki (2011) as well as the smooth ambiguity model (Neilson (1993), Nau (2006), Klibanoff et al. (2005)).[1] Like multiple priors, these models work with sets of beliefs, and agents who evaluate a plan behave *as if* they are endogenously more pessimistic about that plan because unfavorable beliefs matter more for utility. Much of the intuition we discuss for multiple priors thus applies to all widely-used models.

The models differ from multiple priors in how belief sets enter the certainty equivalent function – in particular, they make the certainty equivalent a smooth function of consumption plans. As a result, welfare effects of uncertainty are not first order in the sense we describe below. Applications with multiplier or smooth ambiguity preference thus do not typically

[1] Multiplier preferences are related to work on robust control in engineering. See Hansen and Sargent (2008) and Hansen and Sargent (2022) for an interesting recent extension that formalizes a concern for model misspecification. For a comprehensive survey of theoretical properties and applications of the smooth model, see Marinacci (2015).

use linear methods, and do not emphasize qualitative differences from risk aversion related to inaction or inertia. The common denominator of all models that we emphasize in our review of applications is that they capture agents' concern with model uncertainty, which allows for better quantitative performance compared to models with risk aversion and rational expectations.

2.2 Risk vs ambiguity: similarities and differences

The certainty equivalent functions for risk aversion (2) and ambiguity aversion (3) capture disutility from plans that vary across states. As a result, much of the standard intuition for behavior under uncertainty equally applies to both models. For example, a change in uncertainty that lowers the certainty equivalent of future consumption makes agents unhappy; they may respond by taking precautionary actions. The models differ, however, in what a change in uncertainty means. An increase in risk is typically captured by adding a mean-preserving spread in the distribution of date 2 consumption. An increase in ambiguity instead is an increase in the set of distributions used to evaluate date 2 consumption. This difference has important implications for modeling strategy and behavior.

Uncertainty without curvature. With risk aversion, utility costs of uncertainty are due to curvature in v. Indeed, if v is linear, a mean preserving spread in the distribution of $\tilde{c}$ does not affect utility, since CE only depends on the mean $\mathbb{E}^P[\tilde{c}]$. This observation has led modelers who want to avoid curvature in utility for other reasons, say to eliminate wealth effects for tractability, to abstract from aversion to uncertainty altogether. For example, many papers on search markets or contracting within firms work with linear $u = v$.

The multiple priors model, in contrast, captures uncertainty aversion even when $u = v$ is linear. Suppose the beliefs $P \in \mathcal{P}$ differ in means, and consider utility

$$c + \beta \min_{P \in \mathcal{P}} \mathbb{E}^P[\tilde{c}] \tag{4}$$

An increase in ambiguity – a larger set $\mathcal{P}$ – now lowers utility if it leads to a lower worst-case mean for the plan $\tilde{c}$. Incorporating ambiguity aversion thus allows tractable models without wealth effects that nevertheless feature compensation for uncertainty as part of contract design or prices in search markets.

Welfare costs of uncertainty. With risk aversion, the welfare costs of uncertainty are second order. The textbook result starts from a certain consumption plan $\tilde{c} = \bar{c}$ and adds a small mean preserving spread $\alpha\tilde{z}$, say, where $\mathbb{E}^P\tilde{z} = 0$. With concave v, welfare declines for any α, but the welfare loss scales with α^2, and so goes to zero faster than α. This property of risk preferences has far-reaching consequences for applications. It generates a strong appetite for risk taking: starting from certainty, a risky gamble is attractive even if it offers very small risk compensation, as long as that compensation declines linearly in α and hence more slowly than the welfare loss of uncertainty. More generally, it is a hallmark of quantitative models with risk aversion that agents are willing to tolerate high uncertainty in return for a small gain in the mean.

With multiple priors, in contrast, we can have first order welfare losses from uncertainty. Consider linear utility (4) with a set of beliefs that differ in means. Starting again at a certain consumption plan, add a random variable $\alpha\tilde{z}$ such that $\min_{P\in\mathcal{P}} \mathbb{E}^P \tilde{z} < 0 < \max_{P\in\mathcal{P}} \mathbb{E}^P \tilde{z}$, but all other moments are identical. The belief set that describes ambiguity is thus parametrized by an interval of means for $\tilde{z}$ that brackets zero: it contains beliefs with positive and negative means for $\tilde{z}$. Again utility declines for any α: the worst-case mean endogenously adjusts so its sign is opposite to that of α. The key difference is that the welfare loss from ambiguity is linear in α. In applications with multiple priors, mean and uncertainty can thus have the same order of magnitude.

Preference towards the timing of the resolution of uncertainty. It is common in applications with risk aversion to assume curvature in v to be much larger than in u so $\gamma(c)$ is much higher than the desire for consumption smoothing over time, measured by the inverse of the IES $\sigma(c)$. The assumption is particularly prominent in macroeconomics and finance in order to address asset pricing puzzles under rational expectations. However, it also implies that agents strictly prefer early resolution of uncertainty, that is, agents would pay to receive information even if they cannot adjust their consumption plans in response (see Epstein et al. (2014) for quantitative thought experiments). Standard expected utility agents with $u = v$ do not exhibit such behavior.

With multiple priors, modelers can avoid this tight connection between aversion to uncertainty and preference towards the timing of its resolution. In particular, multiple priors agents with $u = v$ are indifferent towards the timing of the resolution of uncertainty, just like standard expected utility maximizers.[2] For example, an agent with low IES but a large set of beliefs $\mathcal{P}$ also requires a lot of risk compensation, but would not pay up front for information unless he can use it in planning. Of course, if compensation for timing is desirable to match the data, then the multiple priors model can be augmented with $v \neq u$.

Uncertainty attitude versus beliefs. The two primitives that govern attitude towards risk – the belief P and the utility function v – permit a distinction between risk perception and

[2] Strzalecki (2013) shows that multiple priors is the only model of ambiguity aversion that allows for indifference to timing. Intuitively, this is because other models capture ambiguity via nonlinear functions that affect timing preference, much like nonlinearity of v does in the case of risk.

risk attitude, respectively. At the same time, both work through the certainty equivalent, implying that changes in risk or risk aversion often have similar implications - for example through intertemporal substitution behavior or required compensation.

Models of ambiguity aversion, in contrast, do not offer a distinction between attitude and beliefs. In particular, the set of beliefs $\mathcal{P}$ always reflects in part the agent's attitude towards ambiguity. Our Ellsberg paradox examples above illustrate this point: we used features of the set to describe events that are perceived as risky versus ambiguous. More generally, modelers often use the set of beliefs to encode different attitudes towards different sources of uncertainty. For example, perception of a source as ambiguous may be driven by a lack of information about it, which makes agents less confident in probability assessments.

Early quantitative work in economics relied heavily on the distinction between risk perception and risk attitude. The idea was to (i) obtain observable proxies for beliefs from survey data or statistical models and (ii) assume that a risk aversion coefficient is a "deep parameter" that can be transferred across settings; for example it can be measured in the laboratory and then used to calibrate models. As a result, low risk aversion coefficients measured in experiments were imposed on investors in financial markets.

More recently, the distinction between beliefs and attitude has arguably become less important. One reason is that the limits of assumption (ii) have become more clear. For example, Rabin (2000) provides a thought experiment to show that concavity as a source of risk aversion cannot deal well with plausible choices given small versus large risks. Moreover, the use of richer data sets and modern estimation techniques allows modellers to estimate preference parameters in many contexts. Of course, it is still useful to write parsimonious models such that behavior towards uncertainty depends on, say, a subjective conditional variance, a curvature parameters or the width of an interval of means for a belief set. But it is not crucial that any such parameter is transferable across settings.

2.3 Savings and portfolio choice

In this section we illustrate how first order welfare losses from ambiguity shape behavior in two familiar applications. In a consumption savings problem, precautionary savings obtain even when utility is quadratic, and can therefore be characterized using linear approximations. In a portfolio choice problem, uncertainty leads to robust non-participation in markets even in the absence of transaction costs.

Precautionary savings. Consider a two-period consumption-saving problem. An agent is endowed with date 1 wealth w and uncertain date 2 labor income $\tilde{y}$. Savings s earn a certain gross interest rate R between dates 1 and 2. The agent chooses a consumption plan $(c, \tilde{c})$ and savings s to maximize utility (1) subject to the budget constraints

$$\begin{aligned} c &= w - s; \qquad s \geq 0 \\ \tilde{c} &= Rs + \tilde{y}. \end{aligned}$$

We abstract from timing effects by assuming $u = v$.

For both risk aversion and multiple priors, the condition for an interior optimum for savings takes the form of a consumption Euler equation

$$u'(w - s) = \beta R \mathbb{E}^{P^*} \left[u'(Rs + \tilde{y}) \right] \tag{5}$$

The only difference between the models is the interpretation of P^*: with risk aversion, it is the unique subjective belief of the agent that is given as part of utility (2), whereas with multiple priors it is the worst-case prior that achieves the minimum when (3) is evaluated at the optimal consumption plan.

How does an increase in uncertainty affect savings? With risk aversion, we have the textbook result that income uncertainty increases savings if and only if utility exhibits prudence, that is, $u''' > 0$. In particular, the tractable case of quadratic utility rules out precautionary savings. Moreover, a linear approximation to (5), say around the point of certain labor income, cannot be used to study precautionary savings, since it would reflect at most second derivatives of u. As a result, the standard practice of solving general equilibrium models by linearization abstracts from precautionary savings effects.

With multiple priors, in contrast, precautionary savings requires only curvature in utility. To see this, assume income ambiguity $\tilde{y} = \bar{y} + \alpha \tilde{z}$, where $\bar{y}$ is a constant and $\tilde{z}$ is a random variable such that the interval of means brackets zero but all other moments are the same. The worst-case belief P^* is then the belief that implies the lowest – and hence a negative – mean for $\tilde{z}$. Regardless of the savings choice s and the parameter α, this belief minimizes expected utility. We then obtain the impact of uncertainty on savings by applying the implicit function theorem to (5):

$$\left. \frac{ds}{d\alpha} \right|_{\alpha=0} = \frac{1}{\Delta} R u''(Rs + \bar{y}) \, \mathbb{E}^{P^*} [\tilde{z}] > 0,$$

where $\Delta > 0$ from the second order condition for optimality. Precautionary saving occurs as long as the agent is not willing to perfectly substitute consumption over time.

The example further illustrates two principles that are relevant also for understanding the quantitative applications we review below. First, for any given choice problem, there is an observational equivalence between ambiguity aversion and pessimism. Solving a decision problem with multiple priors involves first finding the worst-case belief, and then characterizing the solution given that belief, much like for the standard model. This feature is convenient for computation: we can often find the worst-case belief up front, as we did above, and then solve the problem with familiar tools. Importantly, the observational equivalence between ambiguity aversion and dogmatic pessimism (with a fixed prior) only holds for a given choice problem; it does not mean that we can simply replace the ambiguity aversion model by a model of dogmatic pessimism.

To see how observational equivalence does not extend beyond a given choice problem, compare responses to a change in the environment for an ambiguity-averse agent and an agent with a fixed prior. As a stark example, suppose the government introduces a transfer program that pays $-2\alpha\tilde{z}$ at date 2. In other words, the program subsidizes previously bad states and taxes good states, so much that it reverses their roles. Such a policy increases utility for the agent with the fixed prior: he now receives more income in the (previously bad) states he finds more likely. In contrast, for the ambiguity averse agent utility is unchanged: ambiguity about income is described by the same family of distributions as before, so the worst-case mean changes endogenously to put more weight on previously good, but now bad states. The agent's disutility stems from ambiguity, not mistaken belief, so a transfer program that leaves ambiguity unchanged is not beneficial.

A second principle is that curvature in the objective function matters for the effect of uncertainty. Indeed, the higher the curvature in u, or the more reluctant the agent is to substitute consumption over time, the more he responds to uncertainty by accumulating precautionary savings. Intuitively, the worry about a downward-sloping consumption path propels action. An agent with linear u still worries about uncertain income – utility is lower – but does not do anything about it. In many dynamic models, parameters that induce costs of uneven plans, such as adjustment costs, similarly strengthen precautionary behavior.

Portfolio inertia. To illustrate portfolio choice, we allow for two assets, but for simplicity abstract from the savings margin. Consider a mean-variance investor who maximizes end-of-period wealth. There is a bond with a constant risk-free rate of return as well as an asset with uncertain payoff that earns an excess return over the risk-free rate of $\tilde{R} = \bar{R} + \tilde{z}$. Assume that the belief set is such that the interval of means for $\tilde{z}$ brackets zero, and all other moments are the same. Let θ denote a zero-cost portfolio that is short in the bond and long in the uncertain asset. There are no borrowing or short sale constraints, so the agent solves

$$\max_{\theta} \min_{P\in\mathcal{P}} \left\{ \mathbb{E}^{P}\left[\tilde{R}\right]\theta - \frac{1}{2}\gamma var^{P}\left(\tilde{R}\right)\theta^{2} \right\}. \tag{6}$$

Without ambiguity, the optimal portfolio is $\theta = \mathbb{E}^P\left[\tilde{R}\right]/\gamma var^P\left(\tilde{R}\right)$. The agent participates in the market whenever the (unique) subjective expected excess return is not exactly zero. As long as there is a small compensation for risk taking, the risk averse investor will take on risk, since the welfare loss is second order. Moreover, any change in the expected excess return alters the optimal portfolio.

With ambiguity, in contrast, non-participation ($\theta = 0$) is optimal whenever $\bar{R}$ is close enough to zero so $\min_{P\in\mathcal{P}} \mathbb{E}^P\tilde{R} < 0 < \max_{P\in\mathcal{P}} \mathbb{E}^P\tilde{R}$, a result due to Dow and Werlang (1992).[3] Indeed, if the agent contemplates going long (short) in the uncertain asset, the worst-case expected excess return is lower (higher) than zero, so the payoff (6) is negative and $\theta = 0$ is strictly better. The *endogenous* switch to a pessimistic belief that depends on the action is crucial here. If instead $\bar{R}$ is sufficiently larger (smaller) than zero, then the solution is the same as under risk, but using the lowest (highest) mean of $\tilde{z}$. Put differently, the optimal portfolio policy as a function of $\bar{R}$ exhibits an inaction region, where the agent is out of the market and does not respond to changes in $\bar{R}$.

3 Dynamic choice and equilibrium

Dynamic applications require conditional preferences at every node of a decision tree. We thus describe briefly recursive versions of the models from Section 2.1 . We continue to assume that time is discrete, but now use Ω to denote a period state space: one element $\omega \in \Omega$ is realized every period, and the history of states up to date t is $\omega^t = (\omega_1, ..., \omega_t)$. Consumption plans are now sequences of random variables $C = (C_t)_{t=0}^{\infty}$, where $C_t : \Omega^t \to \Re$ maps histories up to date t into consumption. We write $U_t(C;\omega^t)$ for utility from the consumption plan C conditional on the history ω^t.

For any plan C, conditional utilities solve the difference equation

$$u\left(U_t\left(C;\omega^t\right)\right) = u\left(C_t\right) + \beta u\left(CE\left(U_{t+1}\left(C;\omega^{t+1}\right)\right)\right), \tag{7}$$

where CE can be one of the certainty equivalent functions from Section 2.1. The probabilities involved in writing those functions are now understood to be one-step-ahead conditionals – at each node ω^t in the decision tree, they provide probabilities over the next state ω_{t+1}. When both u and v are power functions and belief sets are singletons, the utility process is common recursive utility under risk (Epstein and Zin 1989). Recursive versions of multiple priors were axiomatized by Epstein and Schneider (2003) with $u = v$ and Hayashi (2005) for $u \neq v$.

[3] Bossaerts et al. (2010) and Asparouhova et al. (2015) confirm such behavior in experimental studies with laboratory asset markets.

We can summarize the primitives of recursive utility as functional forms for u and CE, a discount factor β, as well as an entire stochastic process of sets of conditional probabilities $\mathcal{P}_t(\omega^t)$ that describe beliefs about the next state ω_{t+1} . In the case of risk aversion only, all sets are singletons. In general, the dynamics of conditional beliefs can reflect changes in uncertainty, for example due to learning or uncertainty shocks. With risk aversion, such changes are captured only by higher moments under the singleton belief. With ambiguity aversion, fluctuations in the size of the belief set provide an additional dimension.[4]

Papers that provide axiomatic foundations for recursive utility consider families of conditional preference orderings, one for each history ω^t. Each conditional ordering satisfies the axioms justifying an atemporal model from section 2, suitably modified for multi-period plans. Moreover, conditional preferences at different histories are connected by dynamic consistency. The setup of the model further implies consequentialism, that is, utility from a consumption plan at history ω^t depends only on the consumption promised in future histories that can still occur after ω^t has been realized.

3.1 Asset pricing

We now use the first order condition from the dynamic model (7) to study asset pricing. Consider an investor who has access, after every history ω^t, to a full set of contingent claims on states $\omega \in \Omega$ that can occur next period. In this section, we are no longer interested only in choice, but also in equilibrium asset price dynamics measured by an econometrician. We thus denote by P^0 the true probability over sequences of states and assume that its one-step-ahead conditionals $P_t^0(\omega_{t+1}|\omega^t)$ have full support. For any one-step-ahead belief P_t, we can then define the random variable dP_t/dP_t^0 that makes a "change of measure" from the true Data Generating Process (DGP) to the belief: we have $dP_t/dP_t^0(\omega^t, \omega_{t+1}) = P_t(\omega_{t+1}|\omega^t) / P_t^0(\omega_{t+1}|\omega^t)$.

We write the price of a contingent claim that trades at date t after history ω^t and pays off one unit of the good if state ω_{t+1} occurs at date $t+1$ as $M_{t+1}(\omega_t, \omega_{t+1}) P_t^0(\omega_{t+1}|\omega^t)$, so the pricing kernel M_{t+1} represents one-step-ahead state prices normalized by (true) probabilities. We denote by $\theta_{t+1}(\omega^t, \omega_{t+1})$ the number of claims on state ω_{t+1} at $t+1$ purchased after history ω^t at date t. The agent's objective is to choose stochastic processes (C_t, θ_t) to maximize utility defined in (7) subject to the budget constraints

$$C_t + E_t^{P^0}[M_{t+1}\theta_{t+1}] = Y_t + \theta_t,$$

[4]Epstein and Schneider (2007, 2008) study learning under ambiguity. In particular, they propose functional forms such that the set of one-step-ahead conditionals can shrink as new data make agents more confident, but also expand when news is ambiguous.

where Y_t is an exogenous income process and the conditional expectation represents expenditure on contingent claims.

In a complete market, absence of arbitrage implies that the unique pricing kernel M_{t+1} contains all information needed to price assets: the date t price Q_t of any asset with state-dependent payoff Π_{t+1} at date $t+1$ is given by

$$Q_t = \mathbb{E}_t^{P^0}\left[M_{t+1}\Pi_{t+1}\right] = Q_t^b \mathbb{E}_t^{P^0}\left[\Pi_{t+1}\right] + cov_t^{P^0}\left(M_{t+1}, \Pi_{t+1}\right) \tag{8}$$

where $Q_t^b = E^{P^0}\left[M_{t+1}\right]$ is the price of a risk-free bond and the covariance reflects the measured uncertainty premium – the difference between price and expected present value under the statistical model used by the econometrician.[5]

Since the econometrician recovers P^0 from the empirical distribution of prices, payoffs and consumption, we can view M_{t+1} and C_{t+1} as the data the model needs to explain. For a quantitatively successful asset pricing model, we want M_{t+1} to vary across states ω_{t+1}, so average premia can differ a lot across assets with different payoff profiles – for example, equity on average pays a larger premium than bonds. We also want M_{t+1} to vary with time t, so premia are volatile; for example, uncertainty generates low prices, and high expected excess returns for many uncertain assets in recessions.

Investor optimization places restrictions on the joint distribution of M_{t+1} and C_{t+1}. Let J_t denote the stochastic process of utility at the optimal consumption plan. We can use dynamic programming to show that optimal portfolio choice implies

$$M_{t+1} = \beta \frac{u'\left(C_{t+1}\right)}{u'\left(C_t\right)} \frac{u'\left(CE\left(J_{t+1}\right)\right)}{v'\left(CE\left(J_{t+1}\right)\right)} \frac{v'\left(J_{t+1}\right)}{u'\left(J_{t+1}\right)} \frac{dP_t^*}{dP_t^0}. \tag{9}$$

where P_t^* is again the worst-case belief, which equals the unique subjective belief in the case of risk aversion.

In the benchmark case with time separable expected utility $u = v$ and rational expectations $P^* = P^0$, the pricing kernel is simply equal to the intertemporal marginal rate of substitution $\beta u'\left(C_{t+1}\right)/u'\left(C_t\right)$. This model is soundly rejected by a large body of evidence. The other factors in (9) illustrate alternative approaches to generate larger premia due to preference for early resolution of uncertainty ($u \neq v$) or ambiguity aversion ($P^* \neq P^0$). We note that while J_{t+1} is not directly observable, it is related to consumption through (7), so observable counterparts can be constructed from consumption data, and possibly from wealth data using the budget constraint.

[5]Dividing by the price Q_t, the covariance term is the conditional expected excess return $E_t[\Pi_{t+1}]/Q_t Q_t^b$, a measure of uncertainty premia that one would get for example by regressing excess returns on conditioning information.

Intuitively, M_{t+1} indicates how bad a state is; the price of an asset is higher (and its premium lower) when it pays off relatively more in bad states, that is, it provides insurance. According to (9) a state can be bad for three reasons. First, future consumption is low – the standard effect from high marginal utility of consumption $u'(C_{t+1})$. Second, if the agent prefers early resolution of uncertainty – so there is more curvature in v than in u, then a state can be bad because future continuation utility is low, so the ratio $v'(J_{t+1})/u'(J_{t+1})$ is high. Importantly, utility can be low even if consumption at $t+1$ is high, say, because of bad news about consumption further in the future, a theme of the long-run risk literature. Finally, if the agent is ambiguity averse, then a state can be bad because it is very ambiguous, so its worst-case probability is much larger than its actual probability. Again here the future beyond next period matters – the worst-case can reflect beliefs about consumption in the long run, and it can do so even if $u = v$.

While compensation for ambiguity can be written as a covariance, much like compensation for risk, we emphasize that it generates premia even with risk neutrality. Indeed, with linear $u = v$, we have

$$cov_t^{P^0}(M_{t+1}, \Pi_{t+1}) = Q_t^b \left(\mathbb{E}_t^{P^*}[\Pi_{t+1}] - \mathbb{E}_t^{P^0}[\Pi_{t+1}] \right) \tag{10}$$

so the asset price in equation (8) becomes the worst-case present value $Q_t = Q_t^b \mathbb{E}_t^{P^*}[\Pi_{t+1}]$. The measured average excess return or premium then reflects only the difference between worst-case expected payoffs and expected payoffs measured by the econometrician.

We comment on two other properties of the pricing kernel that are important for quantitative performance. The first is that state prices – that is, cumulative products of one-step-ahead prices M_{t+1} – have a martingale component. As shown by Alvarez and Jermann (2005), a model without such a component has the property that very long horizon bonds have high premia relative to all other assets. As a result, standard business cycle models with trend stationary consumption and power utility are at odds with asset price data. Borovička et al. (2016) show that preference for early resolution introduces a martingale component. From (9), the cumulative product of changes of measure is also a martingale. This explains why multiple priors models can do a good job accounting for uncertainty premia even when consumption is trend stationary and there is no preference for early resolution.

A second property was highlighted by Ai and Bansal (2018), who build on the empirical work of Lucca and Moench (2015). It is a striking stylized fact that most of the equity premium in the US is earned through price jumps around macroeconomic announcement days, in particular monetary policy announcements. Ai and Bansal characterize the class of preferences that is consistent with price movements on pure news days, when nothing special happens to consumption. Both recursive utility with preference for early resolution and multiple priors belong to this class. The reason is apparent from (9): in both cases, the revelation of news about the far future, as opposed to only consumption, matter for the pricing kernel.

3.2 Business cycle models with uncertainty shocks

We consider a stylized model that can be solved in closed form, yet captures the main effects of uncertainty in quantitative models reviewed below. There are two goods: output of a numeraire consumption good Y_t is made from labor N_t according to the linear technology

$$Y_t = Z_t N_{t-1},$$

where Z_t is exogenous total factor productivity. The fruit of date $t-1$ labor effort thus only become available at date t. Since labor is an intertemporal decision, it depends on uncertainty. The resource constraint for the economy is $C_t = Y_t$.

To allow for flexible labor supply, we extend utility (7) to multiple goods. A consumption plan is a stochastic process for consumption and labor (C_t, N_t). Utility is defined recursively by

$$U_t = \frac{1}{1-\gamma} C_t^{1-\gamma} - \theta N_t + \beta \min_{P \in \mathcal{P}_t} \mathbb{E}^P \left[U_{t+1} \right]$$

where $\mathcal{P}_t$ is the process of one-step-ahead conditional belief sets, a singleton in the risk case.

Under the true data generating process, log TFP $z_t = \log Z_t$ is serially independent and normally distributed:

$$z_{t+1} = \mu_t - \frac{1}{2}\sigma_u^2 + u_{t+1} \tag{11}$$

Here u is an iid sequence of shocks, normally distributed with mean zero and variance σ_u^2. The sequence μ is deterministic and unknown to agents.

Agents treat the unknown component μ_t as ambiguous. We parameterize their one-step-ahead set of beliefs at date t by a set of means $\mu_t^P \in [-a_t, a_t]$. The stochastic process a_t captures agents' perceived ambiguity about TFP; it evolves as

$$a_{t+1} = (1-\rho_a)\,\bar{a} + \rho_a a_t + \varepsilon_{t+1}^a \tag{12}$$

with long run mean $\bar{a} > 0$ and $0 < \rho_a < 1$. Periods of low $a_t < \bar{a}$ represent unusually low ambiguity about future productivity, whereas $a_t > \bar{a}$ describes periods of high uncertainty. We further assume that ε_t^a is independent of u_t.

Equilibrium and planner problem. We assume that labor and consumption goods are traded in competitive spot markets and there are complete financial markets. The welfare theorems then hold and we can characterize competitive equilibrium allocations by solving the social planner problem: we maximize utility subject to the production function and the resource constraint. Using output as the only endogenous state variable, the Bellman equation is

$$V(Y,a) = \max_{N} \left\{ U(Y,N) + \beta \min_{\mu^P \in [-a,a]} \mathbb{E}^P \left[V\left(e^{\tilde{z}} N, \tilde{a}\right) \right] \right\}$$

where the conditional distribution of $\tilde{z}$ under belief P is given by (11) with $\mu_t = \mu^p$.

It is natural to conjecture that the value function is increasing in output. The worst-case belief P^* then has mean $\mu^{P^*} = -a$. Using normality of the shocks, we obtain a closed-form solution for log hours

$$n_t = \bar{n} - (1/\gamma - 1)\left(a_t + \frac{1}{2}\gamma\sigma_u^2\right) \tag{13}$$

where $\bar{n}$ is constant. Substituting back into the Bellman equation verifies the conjecture that the value function is increasing in output.

Business cycle dynamics. According to (13), uncertainty reduces hours – and hence log output $y_t = z_t + n_{t-1}$ whether it is ambiguity, as measured by a, or risk, as measured by the product of the quantity of risk σ_u^2 and risk aversion γ. Intuitively, an increase in uncertainty has wealth and substitution effects. On the one hand, higher uncertainty lowers the certainty equivalent of future production. Other things equal, the resulting wealth effect leads the planner to reduce consumption of leisure and increase hiring. However, higher uncertainty also lowers the uncertainty-adjusted return on labor. Other things equal, the resulting substitution effect leads the planner to reduce hiring. The net effect depends on the curvature in felicity from consumption, determined by γ. With a strong enough substitution effect (low enough γ), an increase in uncertainty lowers hiring.

For $1/\gamma > 1$, an increase in ambiguity a generates a recession, even when productivity is not unusually low. Hours are below steady state only because the marginal product of labor is more uncertain, so the planner finds it optimal not to make people work. Conversely, an unusual increase in confidence – a drop of a_t below its long run mean – generates a boom in which employment and output are unusually high, but productivity is not. In other words, a phase of low realizations of a_t will look to an observer like a wave of optimism, where output and employment surge despite average productivity realizations. In the time series, allocations thus look like there is a countercyclical "labor wedge" between the marginal product of labor and the marginal rate of substitution of consumption for labor, a key feature of aggregate data.

Bounding ambiguity by observed volatility. How much variation in ambiguity is plausible?

Ilut and Schneider (2014) propose a model-consistency criterion to bound the support of a. The idea is that even the most extreme forecasts implied by potential worst-case beliefs should be "good enough" in the long run. The concrete suggestion is to bound the support of ambiguity by $a \leq 2\sqrt{Var(z)}$, where $Var(z)$ is the observed long-run variance of z_{t+1} in (11). In other words, agents do not entertain forecasts outside a 95% confidence interval, centered around the long-run mean of z_{t+1}, given its observed variation. The criterion thus relaxes the rational expectations model-consistency criterion, which would collapse that interval of entertained beliefs to one number, but allows for more ambiguity when the data are more volatile.

A more general solution method. In contrast to our stylized example, quantitative business cycle models include further state variables such as capital and thus require numerical tools. An advantage of modeling uncertainty as ambiguity is that popular first-order perturbation methods do not lose uncertainty effects such as precautionary savings or asset premia, in contrast to what happens with risk. Ilut and Schneider (2014) leverage this insight to propose a tractable four-step approach that generalizes what we did in this section with the stylized model to a broad class of models with stochastic ambiguity about means of shocks.

The first step conjectures the worst-case belief. When ambiguity is about shocks that increase surplus, such as TFP above, the natural conjecture is that the worst-case mean is as low as possible. The second step solves the model with worst-case beliefs using standard methods for expectational difference equations, for example linearization around the deterministic steady state (e.g. Sims (2002)). It delivers a steady state and law of motion used by the agent for planning. The third step verifies the conjecture: since the agent optimizes as if the worst-case dynamics were correct, we can check the (equilibrium) value function. For example, a conjecture of a low worst-case mean for some shock is verified if the value function is increasing in that shock.

Finally, with agents' perceived law of motion in hand, the last step describes model dynamics as measured by an econometrician. Here we take into account that exogenous variables follow the true DGP, not the worst-case dynamics perceived by agents. In particular, one-step-ahead conditional means of ambiguous shocks under the true DGP are more favorable – the difference is the gap between true and worst-case means. In our example above, agents plan *as if* TFP is always below trend, but realized TFP in fact has mean zero. Moreover, endogenous variables respond to the true exogenous shocks, but according to agents' pessimistic equilibrium response. In our example, hours move with ambiguity only because of agents' pessimistic perception of future TFP.

If step two characterizes the law of motion by linearization, then ambiguity about conditional means matters for the coefficients of the approximating linear law of motion.

We have seen in Section 2.3 how precautionary savings under such ambiguity requires only curvature in utility, not a positive third derivative as with risk. It is therefore reflected in coefficients of a linearized consumption Euler equation. We have also seen in Section 3.1 how uncertainty premia in asset prices occur due to differences in worst-case versus true payoffs: while investors behave as if payoffs are low, so prices are below the present value of payoffs, actual payoffs are high, which generates high average excess returns from the perspective of the econometrician. With ambiguity about means, this effect is again present in the linear law of motion.

Volatility shocks and identification. The example law of motion (12) assumes that ambiguity is not connected to movements in actual TFP, the exogenous fundamental. Uncertainty about the future only responds to intangible information, such as signals about technological progress. Theory does not require this property. An alternative, proposed by Bianchi et al. (2018), makes ambiguity move with the volatility of fundamental shocks. In more turbulent times – that is, when larger shocks are expected – agents also find it harder to settle on a forecast of the future. Qualitatively, the model of Bianchi et al. (2018) thus works like a rational expectations model with stochastic volatility. The differences is that subjective uncertainty is captured by ambiguity. The estimation, in fact, allows ambiguity to move both with volatility and without it; the relative share of each type of movement is identified from the comovement of asset premia – which depend on subjective uncertainty – and conditional heteroskedasticity in fundamentals.

4 Quantifying ambiguity using survey data

In this section, we describe work that relates survey measures of ambiguity to experience and choice in the cross section of economic agents. We focus on large scale surveys, and do not cover the vast literature on experimental evidence reviewed for example by Trautmann and Van De Kuilen (2015). We emphasize three takeaways. First, ambiguity aversion is a widespread phenomenon among both households and decision makers in firms. Second, variation in ambiguity is large, both in the cross section of agents and in the time series for a given agent. It is at least partly accounted for by information, but not by sophistication. Third, measures of perceived ambiguity relate in sensible ways to observed choice. We conclude that survey evidence is consistent with rational behavior under ambiguity and hence a promising source for disciplining belief sets in structural models – in a final subsection we provide examples below for how this has been done.

Methodology. We distinguish two approaches to eliciting ambiguity. One follows the experimental literature and infers properties of preferences from choice, for example by incorporating survey questions about (hypothetical) Ellsberg-type choice situations. It can in principle test axioms underlying specific models of ambiguity aversion. The second approach elicits stated "imprecise probabilities": it asks respondents about the likelihood of an event, but provides the option to answer with a range of probabilities. It has roots in psychology (for example, Wallsten et al. (1983)) and was first used in economics by Manski and Molinari (2010); it is also discussed in Chapters 9 and 26 of this handbook. Asking for imprecise probabilities has recently become more prominent since it makes it easy to introduce questions about natural (not artificial) events into large scale surveys. An important agenda for future research is to compare the two approaches.[6]

Interpretation of imprecise probabilities needs to take into account that decision theory does not directly relate belief sets to survey responses. Belief sets only serve as devices to represent preferences – agents act *as if* they maximize with a belief set in mind. In particular, theory does not predict that ambiguity averse agents make pessimistic forecasts, just like the expected utility model does not predict that risk averse agents should do so.[7] Nevertheless, we can make two sharp statements. First, Bayesian respondents should answer questions about likelihoods with single probabilities. The frequency of imprecise probability answers thus measures the prevalence of non-Bayesian perception of uncertainty. Second, more ambiguity averse agents behave *as if* their belief sets are larger. If we are willing to make the additional assumption that such agents also state wider probability ranges in surveys, then we can rank ambiguity across agents, or for the same agent over time.

Evidence on households. Dimmock et al. (2016) elicit Ellsberg-style choice behavior among households in the American Life Panel. They show that, consistent with the theory described in Section 2.3, ambiguity aversion is negatively related to stock market participation and the portfolio weight on stocks, but positively related to own-company stock ownership and other measures of under-diversification. Bianchi and Tallon (2019) match survey responses to portfolio data from a large French financial institution. Ambiguity averse investors not only under-diversify and hence bear more risk, but are also more active traders who earn higher average returns. Delavande et al. (2021b) add an incentivized choice module to the UK Household Longitudinal Study to assess various models of behavior under uncertainty. An interesting result is that women exhibit less ambiguity aversion, in contrast to earlier

[6] For recent methodological advances in eliciting ambiguity perceptions about natural events from choice, see Baillon et al. (2018) or Abdellaoui et al. (2021).

[7] Recent work has tested this hypothesis in survey data. Under both risk or ambiguity, equation (8) says that uncertainty-adjusted expected returns should be equated across assets. However, subjective expected returns in surveys differ substantially across assets, see for example Adam et al. (2021) for securities and Kindermann et al. (2021) for housing returns.

work on risk; it underscores the need for comprehensive measurement of uncertainty.

Households prefer to communicate their uncertainty via imprecise probabilities for a diverse set of events, especially for those where information is scarce. Giustinelli et al. (2021) use the US Health and Retirement Study to show that about one half of dementia-free older Americans perceive ambiguity about dementia and the need for long-term care and that ambiguity about the latter *conditional* on the hypothetical dementia state is substantially lower. Delavande et al. (2021b) elicit imprecise probabilities for a range of health and financial outcomes and documents widespread ambiguity. Delavande et al. (2021a) study ambiguity about Covid-19-related health outcomes and also ask about ambiguity conditional on hypothetical behavior to show that perceived returns from protective measures predict actual choice. We find such evidence on conditional beliefs particularly interesting to guide future modeling of learning and updating under ambiguity.

Evidence on firms. Bachmann et al. (2020) elicit imprecise probabilities about future sales growth from leading executives in German manufacturing firms, based on the ifo Business Survey. They present results from a quarterly panel over four years: even though sales growth is a short term, routine random variable for top executives, 25% of firms respond with imprecise probabilities in any given quarter, while 70% of firms do so at least once. Time variation in ambiguity is systematically related to credit spreads, familiar measures of uncertainty in financial markets. For example, the share of ambiguous responses spikes up in the 2015 Greek default crisis, especially for exporting firms. While there is considerable variation across firms, perception of ambiguity is unrelated to the use of statistical analysis inside the firm, and forecast mistakes are similar for precise and imprecise probability responses.

Quantifying models. Survey data have played a useful role in quantifying models of decision making under ambiguity. It is again helpful to distinguish two approaches. One is to jointly design the survey and the structural model to tightly connect survey answers to their model counterparts. An example is the pioneering study of Giustinelli and Pavoni (2017) who estimate a model of learning under ambiguity by Italian high school students who choose career tracks. The data follow students' beliefs and choices over several years. Ambiguity is again prevalent, as is lack of awareness of possible choices. Imprecise probabilities elicited from students provide moments that discipline the evolution of uncertainty. A main result is that the speed of learning depends importantly on parents' background.

An alternative approach is to use the *dispersion* in survey forecasts as a proxy for ambiguity. Ilut and Schneider (2014) draw on the Survey of Professional Forecasters to construct an observable counterpart for variations in ambiguity that drive the business cycle. The basic idea is that an ambiguity-averse representative household samples experts' opinions before making decisions, and stronger disagreement among experts leads to lower confidence in probability assessments. More formally, beliefs are parameterized by an interval for the one-quarter-ahead conditional mean of total factor productivity, where the interval is assumed to be monotonically related to the interdecile range of expert forecasts of output.[8] Forecast dispersion thus serves as an index of ambiguity – its volatility and comovement with other variables discipline the scope for uncertainty to matter.

5 Aggregate applications

In this section, we review macro-finance models where ambiguity serves as a driver of the business cycle, asset prices, or both.

Ambiguity-driven business cycles. Ilut and Schneider (2014) propose and estimate a standard New Keynesian (NK) model with multiple priors utility. Ambiguity is about productivity, and its magnitude and fluctuations are disciplined by the dispersion in survey forecasts about GDP. The main result is that confidence shocks, that is, exogenous fluctuations in ambiguity, drive most of the business cycle. In particular, they generate comovement of hours, consumption and investment without strong reactions in inflation. The model thus meets the challenge put forward by Angeletos et al. (2020): models of demand-driven business cycle should be consistent with stable inflation. Intuitively, worry about productivity is not only a force for weak demand as households worry about future income, but also a force for high prices, as firms worry about future cost. The latter effect counteracts the deflationary force of weak demand.

Other papers in this area differ in how ambiguity is captured and what source of uncertainty it is about. Bidder and Smith (2012) consider a model with multiplier preferences and stochastic volatility; Bhandari et al. (2019) allow for shocks to agents' concerns for model misspecification. Related work analyzes the smooth ambiguity model with stochastic volatility (Backus et al. (2015)), learning about TFP (Altug et al. (2020)), or dispersed information (Pei (2018)). Masolo and Monti (2020), Michelacci and Paciello (2020) and Kroner (2021) propose quantitative models that use multiple priors preferences to focus specifically on ambiguity about monetary policy.

Ilut and Saijo (2021) go beyond exogenous ambiguity shocks to show how ambiguity is a *propagation mechanism* for standard fundamental shocks. Perceived uncertainty is endogenously higher in recessions, caused either by supply or demand shocks, because firms

[8]The mapping from survey forecast dispersion about output in the data to model beliefs about total factor productivity is a parametric function that respects the relationship between the variables in the model.

accumulate less information about themselves when they produce less. This depresses aggregate activity, since uncertainty-adjusted returns to working, consumption and investing are all lower. In other words, higher ambiguity in recessions *jointly* generates labor, consumption and investment "wedges" with respect to standard marginal conditions. The labor wedge works like in the simple model of section 3.2. The consumption wedge occurs because precautionary savings pushes the risk-free rate below the growth rate of marginal utility. The investment wedge reflects equilibrium compensation for ambiguous returns on capital.

The model provides a unified explanation for comovement of major aggregates, as well as countercyclical correlated wedges, even though prices are fully flexible. The comovement result is surprising in light of Barro and King (1984) who showed that in a standard RBC model hours and consumption co-move *only* in response to shocks to TFP or the disutility of work. New Keynesian models overturn this impossibility result through countercyclical markups. In Ilut and Saijo (2021), the endogenous countercyclical labor wedge allows labor and consumption to both fall *even conditional on demand* shocks, as in the data.

Asset pricing under ambiguity. A large literature studies asset pricing in representative agent endowment economies. Many early contributions are reviewed in detail in Epstein and Schneider (2010) and Guidolin and Rinaldi (2013). Ambiguity has been shown to generate the level and time-variation in the equity premium (eg. Chen and Epstein (2002), Hansen (2007), Barillas et al. (2009) Hansen and Sargent (2010), Ju and Miao (2012), Chen et al. (2014), Bidder and Dew-Becker (2016), Collard et al. (2018), Gallant et al. (2019)), carry trade excess returns (Ilut (2012)), CDS spreads (Boyarchenko (2012)), sovereign debt spreads (Pouzo and Presno (2016)), index options (Drechsler (2013)), or the variance premium (Miao et al. (2019)). Recent work studies asset pricing in production economies, where time-variation in premia arises from learning about an unobserved technology state, for example Jahan-Parvar and Liu (2014)), a process which further allows for time-varying volatility in Liu and Zhang (2021).

Bianchi et al. (2018) build on the points of Section 3 to study the joint dynamics of not only macro aggregates and asset prices in a production economy, but consider asset supply: they show that time varying ambiguity about technology generates observed fluctuations in shareholder payout and capital structure. Firms in the model face an increasing marginal cost of debt that is traded off against a tax advantage, as well as adjustment costs to equity. When ambiguity about future cost is high, shareholders worry more about the cost of debt, scale back leverage and lower payout. This mechanism accounts not only for the procyclicality of payout and debt, but also for medium term swings.

Importantly, ambiguity is not only about TFP, but also about operating costs that affect earnings but do not scale with production. The latter allows the model to account for high investment in times of high uncertainty, such as in the 1970s, whereas uncertainty about TFP generates low investment in recessions. Moreover, it helps reconcile volatile uncertainty premia with stable interest rates. Intuitively, operating cost is a small share of overall consumption, but a sizable share of shareholder payout, then uncertainty over these operating costs increases equity premia but affects uncertainty about future consumption only weakly, and thus has a small effect on precautionary savings and bond prices.

6 Heterogeneity and micro-to-macro applications

In this section we study review work that extends beyond the representative agent frameworks discussed so far. In particular, we discuss treatments of heterogeneous perceptions of uncertainty as well as micro-to-macro applications that use ambiguity to better account for behavior of heterogeneous firms, consumers and investors.

6.1 Heterogeneous perceptions of uncertainty

Heterogeneous ambiguity averse agents often act *as if* they disagree. One potential reason is heterogeneity in preferences. Ilut et al. (2016) study a model where more ambiguity averse households have larger belief sets about labor income and dividends. Equilibrium is then observationally equivalent to a model with persistent disagreement. The setup raises a computational challenge: in steady state, agents act *as if* they are on a transition path to *different* worst-case steady states. Ilut et al. (2016) thus propose a first order perturbation approach that solves jointly for steady state and transition dynamics. When quantified using US data on income and wealth, the model matches negative risk-free real rates, financial leverage and aggregate uncertainty premia. A key force is precautionary savings, which is reflected in the linear approximation and varies with aggregate uncertainty shocks.

Heterogeneity in the exposure to shocks can also make agents act *as if* they disagree. Michelacci and Paciello (2020) study a New Keynesian model with lenders and borrowers who receive ambiguous news of low future interest rates, a redistribution from lenders to borrowers. Each agent responds *as if* the news was unfavorable: under the worst-case belief, only lenders find the announcement credible. The main quantitative result is that this effect significantly weakens the power of ECB forward guidance. In Saijo (2020), owners and non-owners of capital act as if they disagree about future capital taxes. While capital owners fear high taxes and substitute away from investment, non-owners fear low transfers from austerity and reduce consumption. Uncertainty generates a recession that is substantially stronger than in a representative agent benchmark.

In addition to the above papers on interaction in markets, a growing literature studies contracting with heterogeneous agents. Carroll (2019) reviews work on mechanism design that introduces ambiguity in standard principal/agent problems. Recent work in corporate finance studies capital structure choice under ambiguity in with bankruptcy costs (Izhakian et al. (2021)), asymmetric information (Malenko and Tsoy (2020)), dynamic moral hazard (Miao and Rivera (2016), Dicks and Fulghieri (2021b)) or allocation of control rights (Garlappi et al. (2017)). In political economy, multiple priors utility has been used to model parties that perceive ambiguity about voter preferences (eg. Bade (2011)) as well as ambiguity-averse strategic voters, to explain selective abstention (Ghirardato and Katz (2006)) and show that abstention can be strong enough to prevent information aggregation (Ellis (2016)).

6.2 Inaction and inertia

Inaction due to first order effects of uncertainty helps understand a wide variety of phenomena. One important area is the behavior of nominal prices in micro data. Ilut et al. (2020) study price setting by ambiguity averse firms who learn about demand. The key idea is that uncertainty about the shape of a demand curve is local to every price point. A firm's uncertainty is therefore lowest near prices at which it has sold before. Formally, the model studies learning under ambiguity with sets of flexible Gaussian process distributions that allow many shapes of demand.

While local demand uncertainty could in principle be modelled as risk, in the presence of ambiguity aversion it implies price rigidity. For a firm that contemplates a price increase away from its current posted price, the worst-case is that the unknown demand function is very elastic at the higher price, so the firm loses a lot of sales from the price increase. For a price decrease, in contrast, the worst-case belief is a very inelastic demand, so the price cut does not stimulate sales much. This endogenous change in the worst-case belief forms a *kink* in the worst-case expected demand at previously posted prices, where the posterior uncertainty is the lowest.

Under ambiguity, firms thus behave *as if* they face time- and state-dependent costs of changing prices away from previously posted prices. Nominal prices stickiness emerges even without fixed ('menu') costs of changing prices. Moreover, in a quantitative evaluation based on micro-level pricing moments, the theory can account for facts that are difficult to explain with fixed costs: (i) price memory (the high likelihood of revisiting old price points), (ii) a decreasing price change hazard and (iii) the co-existence of both small and large price changes. Related work in industrial organization includes Bergemann and Schlag (2011) who analyze a static optimal pricing problem with multiple priors over the distribution of buyers' valuations, Handel and Misra (2015) who consider a model with maxmin regret, and Handel et al. (2013) who develop theoretical and econometric arguments to identify consumer preferences from discrete choice data.

Inaction due to ambiguity also motivates a recent literature on technology adoption. In development economics, Warnick et al. (2011) and Ross et al. (2012) find that ambiguity aversion, and not risk aversion, constrains adoption of new farming technologies. Other recent work on slow adoption includes Kala (2019), Bryan (2019), Dougherty et al. (2020) and Norton et al. (2020). Moreover, technology adoption gives rise to interesting strategic interactions. Mukerji and Tallon (2004b) reviews work on games with ambiguity about the others' strategies (e.g. Dow and Werlang (1994), Eichberger and Kelsey (2002)) or payoffs (e.g. Azrieli and Teper (2011)). A key insight is that when individual payoff functions depends on others' actions, equilibrium worst-case scenarios may reflect a fear of 'missing out', that is, not innovating or investing enough. In Beauchêne (2019), Bennett (2021), and Dicks and Fulghieri (2021a), various forms of payoff complementarity imply cautious individual decisions, but also equilibria that appear excessively *optimistic* from the perspective of an outside observer.

Non-participation in asset markets due to ambiguity has interesting equilibrium implications. Mukerji and Tallon (2001) and Mukerji and Tallon (2004a) model endogenously incomplete markets. In particular, ambiguity about relative prices can account for the observed lack of indexed debt. Condie and Ganguli (2011a,b, 2017)) study models with informational inefficiency: private information is not fully revealed in equilibrium due to non-participation of informed investors. *Informational inertia* means that equilibrium stock prices do not reflect public information in some parts of the state space (see for example Illeditsch (2011), Illeditsch et al. (2021) and supporting evidence in Ben-Rephael and Izhakian (2020))). We finally note that ambiguity changes tradeoffs in evaluating real options. For example, in the job search model without recall of Nishimura and Ozaki (2004), an increase in ambiguity leads an agent to resolve uncertainty sooner and thus *stop* looking for a job, in contrast to the effect of an increase in risk (for other optimal stopping problems under ambiguity see Riedel (2009), Miao and Wang (2011) and Li (2019)).

6.3 Ambiguous information and asymmetric decision rules

Ambiguity can be due not only to prior lack of information, but also to ambiguous signals that are hard to interpret. Epstein and Schneider (2007, 2008) propose a model of learning from ambiguous signals: agents update beliefs with multiple likelihoods. Epstein and Halevy (2021) clarify the distinction between multiple priors and likelihoods and provide experimental evidence. A prominent special case is ambiguous signal quality: multiple likelihoods differ in precision. It implies that posterior beliefs respond asymmetrically to news: agents act *as if* bad (good) news are relatively more (less) precise. Early work applied this effect to asset pricing (for example, Epstein and Schneider (2008), Illeditsch (2011) and Ilut (2012)), we focus here on recent macroeconomic applications.

Ilut et al. (2018) show how asymmetric concave decision rules that respond more to bad than good news endogenously generate (1) negative skewness in both the cross-section and time series even if firm-level shocks are not skewed, and (2) countercyclical cross-sectional and time series volatility – this is because a bad aggregate shock lowers the average firm-level shock and leads the typical firm to respond more. The cross-sectional dispersion of actions thus increases in bad times even if the volatility of firm-level shocks is unchanged. Furthermore, Ilut et al. (2018) use Census data on U.S. manufacturing establishments to document that the relationship between employment growth and innovations to profitability is usually concave and argue for the empirical relevance of this asymmetry-driven mechanism as a unified and endogenous link between micro and macro moments.

Baqaee (2020) shows how ambiguity about inflation news endogenously generates downward wage rigidity, as in the data. Workers worry that inflation might lower real wages. When they observe news of ambiguous quality about future prices, they are more sensitive to inflationary than to deflationary news, and this is reflected in the equilibrium wage. Yoo (2019) estimates a model of consumption with an ambiguous quality of TFP signals that can deliver quantitatively relevant negative skewness in aggregate consumption responses. Kroner (2021) builds a New Keynesian model with two sectors that differ in the perceived ambiguity about the relevance of forward guidance signals. The asymmetric response mechanism implies that, as in the data, high-uncertainty sectors are endogenously more responsive to contractionary forward guidance and less responsive to expansionary forward guidance.

7 Policy implications

In general, an optimal policy problem with ambiguity features (i) ambiguity in the policy maker's objective and (ii) ambiguity in agents' utility, which is implicitly part of the policy maker's constraint set. While both features can be present and interact in the same study, they give rise to different effects, so we discuss their effects in two subsections.

7.1 Ambiguous policy objectives

There is a long tradition of studying policy uncertainty as risk - a prominent theme is the Brainard principle: monetary policy should be less aggressive if key features of the economy are unknown. Manski (2011) and Manski (2013) argue that partial identification of policy effects point to ambiguity as a way to capture policy markers' uncertainty.

Whether cautious behavior under ambiguity leads to attenuated or more aggressive policy actions depends on the source and structure of ambiguity. Barlevy (2011) reviews literature on attenuation, which occurs when the policy maker is not confident in the *sign* of the net welfare gain from a policy, a version of the "Brainard principle". At the extreme, inaction due to ambiguity may lead the policy maker to not intervene at all, as in the social insurance model of Kocherlakota and Phelan (2009).

However, cautious optimal stabilization policy may also respond more aggressively to shocks. In Sargent (1999) and Coenen (2007)), persistent shocks create bigger losses. With ambiguity about persistence, caution calls for more aggressive action. In Giannoni (2002) and Leitemo and Söderström (2008)), aggressive policy stems from uncertainty about feedback effects between multiple policy objectives. Woodford (2010) and Adam and Woodford (2012)) show how history dependence in optimal policy becomes stronger when the policy-maker is uncertain about agents' model misspecification.

A related literature studies the performance of a given set of (usually simple) policies across a small but importantly *non-nested* set of standard models that have substantially different propagation mechanisms (see for example Levin et al. (2003), Levin and Williams (2003), Kuester and Wieland (2010) and the survey in Taylor and Williams (2010)). The overall policy performance is computed as a weighted average of the policy's utility outcomes conditional on each model. Robust policies are then those that do 'well enough' even when the model where that policy does the worst has most of the weight.

Finally, we briefly mention the growing interdisciplinary literature on ambiguity and environmental policy. Climate scientists report high uncertainty about key parameters such as the climate sensitivity to greenhouse emissions. Similarly,the economic costs and benefits of adopting environmental policies are highly uncertain, but are crucial for quantitative models (for example, Palmer and Stevens (2019)). For frameworks to capture such uncertainty as ambiguity, see Lange and Treich (2008), Asano (2010), Millner et al. (2010), Millner et al. (2013). Examples of dynamic, quantitative economic models of climate change with ambiguity are Lemoine and Traeger (2016), Lemoine and Rudik (2017), Barnett et al. (2020), Barnett et al. (2021) and Hansen (2021).

7.2 Optimal policy with ambiguity-averse agents

When policy changes, ambiguity averse agents' endogenously pessimistic beliefs adjust, as discussed in Section 2. The optimal policy problem thus changes, even if the policy maker itself is not ambiguity averse, as emphasized in Hansen and Sargent (2008) and Hansen and Sargent (2015).

Consider first fiscal policy. Karantounias (2013) studies optimal taxation when agents have multiplier preferences. Procyclical tax policy dampens the private agents' effective pessimism, which allows the government to issue debt at a higher equilibrium bond prices and reduce the welfare losses arising from distortionary taxes. Similar setups are studied in Karantounias (2020) and Ferriere and Karantounias (2019). Young (2012) and Bennett et al. (2022) study optimal macroprudential policy in a related setting. A key point there is that policies depend on whether the planner is ambiguity averse or not. In Ilut and Saijo (2021) the government spending multiplier is larger than under rational expectations because the generated increase in economic activity also endogenously raises agents' confidence.

For monetary policy, Ilut and Saijo (2021) show that an interest rate rule that reacts to the credit spread would significantly lower output variability because it stabilizes the variation in endogenous uncertainty and thus in the equilibrium worst-case beliefs. Benigno and Paciello (2014) show in a quantitative NK model that optimal policy under paternalism involves less inflation stabilization and alleviates inefficiencies through a lower average markup. In Baqaee (2020), although fluctuations are costly and asymmetric, the standard policy prescription of an inflationary bias is undone in equilibrium by the endogenously-formed household expectations.

A number of papers have explored policy with heterogeneous ambiguity averse agents. Lensman and Troshkin (2022) analyze optimal mechanisms for public good provision when agents are uncertain about the distribution of private valuations. They find that uncertainty can lead to simple implementations of efficient policies. Caballero and Krishnamurthy (2008) derive a role for a lender of last resort when bank portfolios are ambiguous. Easley and O'Hara (2009) and Brock and Manski (2011) study regulation of credit markets.

We finally comment on welfare costs. Quantitative business cycle models with risk typically exhibit low welfare costs, which has led some macroeconomists to emphasize policy aimed at economic growth (a first order effect) relative to stabilization. As we have seen above, however, ambiguity has *first order* welfare effects. As a result, Ilut and Schneider (2012) and Baqaee (2020)) find *welfare cost* of business cycles that are orders of magnitude larger than the standard Lucas (1987) number based only on risk (see also Barillas et al. (2009) for a discussion).

The same property can generate large welfare costs of *policy uncertainty*. In the linearized New Keynesian model of Saijo (2020), fiscal policy uncertainty leads to heterogeneous worst-case beliefs and altered aggregate dynamics. In Masolo and Monti (2020), ambiguity about monetary policy generates a low frequency drift in agents' worst-case beliefs about inflation. In Michelacci and Paciello (2020), ambiguity in monetary policy communication ignites concerns over redistribution.

8 Concluding remarks

We mention two areas where we see the greatest need for future research. The first is to incorporate ambiguity into a wider variety of models with heterogeneous agents and aggregate uncertainty. Currently, most quantitative models fall into one of two categories. On the one hand, a very active literature on economic policy and inequality allows for rich heterogeneity consistent with micro data on income and wealth, but nevertheless assumes that there are no premia for aggregate uncertainty. On the other hand, models that feature such premia allow only for stylized heterogeneity. One reason is presumably that fast computational methods rely on linear approximations that cannot handle risk premia. As we have shown, ambiguity premia are captured easily using linear methods and are therefore straightforward to introduce into models with rich heterogeneity.

Another area where we look forward to both theoretical and quantitative advances is the endogenous evolution of belief sets. Our review of the applied literature has described a number of mechanisms by which changes in uncertainty move prices and quantities. In other words, the literature has made a lot of progress understanding how beliefs shape data. We feel that for the other direction, how data shape beliefs, existing quantitative work has only scratched the surface, in particular with respect to learning about a world that is constantly evolving due to structural change. We expect that the new literature that directly collects survey data on ambiguity can help distinguish between alternative feedback mechanisms from data to beliefs. A combination of theory and evidence should thus lead us to models of the joint dynamics of uncertainty and observables, with causal effects running in both directions.

References

Abdellaoui, M., H. Bleichrodt, E. Kemel, and O. l'Haridon (2021): "Measuring beliefs under ambiguity," *Operations Research*, 69, 599–612.

Adam, K., D. Matveev, and S. Nagel (2021): "Do survey expectations of stock returns reflect risk adjustments?" *Journal of Monetary Economics*, 117, 723–740.

Adam, K. and M. Woodford (2012): "Robustly optimal monetary policy in a micro-founded New Keynesian model," *Journal of Monetary Economics*, 59, 468–487.

Ai, H. and R. Bansal (2018): "Risk preferences and the macroeconomic announcement premium," *Econometrica*, 86, 1383–1430.

Altug, S., F. Collard, C. Cakmakli, S. Mukerji, and H. Özsöylev (2020): "Ambiguous business cycles: A quantitative assessment," *Review of Economic Dynamics*, 38, 220–237.

Alvarez, F. and U. J. Jermann (2005): "Using asset prices to measure the persistence of the marginal utility of wealth," *Econometrica*, 73, 1977–2016.

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.

Angeletos, G.-M., F. Collard, and H. Dellas (2020): "Business-cycle anatomy," *American Economic Review*, 110, 3030–70.

Asano, T. (2010): "Precautionary principle and the optimal timing of environmental policy under ambiguity," *Environmental and Resource Economics*, 47, 173–196.

Asparouhova, E., P. Bossaerts, J. Eguia, and W. Zame (2015): "Asset pricing and asymmetric reasoning," *Journal of Political Economy*, 123, 66–122.

Azrieli, Y. and R. Teper (2011): "Uncertainty aversion and equilibrium existence in games with incomplete information," *Games and Economic Behavior*, 73, 310–317.

Bachmann, R., K. Carstensen, S. Lautenbacher, and M. Schneider (2020): "Uncertainty is more than risk–survey evidence on knightian and bayesian firms," Working Paper.

Backus, D., A. Ferriere, and S. Zin (2015): "Risk and ambiguity in models of business cycles," *Journal of Monetary Economics*, 69, 42–63.

Bade, S. (2011): "Electoral competition with uncertainty averse parties," *Games and Economic Behavior*, 72, 12–29.

Baillon, A., Z. Huang, A. Selim, and P. P. Wakker (2018): "Measuring ambiguity attitudes for all (natural) events," *Econometrica*, 86, 1839–1858.

Baqaee, D. R. (2020): "Asymmetric inflation expectations, downward rigidity of wages, and asymmetric business cycles," *Journal of Monetary Economics*, 114, 174–193.

Barillas, F., L. P. Hansen, and T. J. Sargent (2009): "Doubts or variability?" *Journal of Economic Theory*, 144, 2388–2418.

Barlevy, G. (2011): "Robustness and macroeconomic policy," *Annual Review of Economics*, 3, 1–24.

Barnett, M., W. Brock, and L. P. Hansen (2020): "Pricing uncertainty induced by climate change," *The Review of Financial Studies*, 33, 1024–1066.

Barnett, M., W. Brock, L. P. Hansen, et al. (2021): "Climate Change Uncertainty Spillover in the Macroeconomy," NBER WP No 29064.

Barro, R. J. and R. G. King (1984): "Time-Separable Preferences and Intertemporal-Substitution Models of Business Cycles," *Quarterly Journal of Economics*, 99, 817–839.

Beauchêne, D. (2019): "Is ambiguity aversion bad for innovation?" *Journal of Economic Theory*, 183, 1154–1176.

Ben-Rephael, A. and Y. Y. Izhakian (2020): "Should I stay or should I go? trading behavior under ambiguity," Rutgers University, mimeo.

Benigno, P. and L. Paciello (2014): "Monetary policy, doubts and asset prices," *Journal of Monetary Economics*, 64, 85–98.

Bennett, F. (2021): "Strategic Complementarity and Ambiguity Aversion: A portfolio choice application," Duke, mimeo.

Bennett, F., G. Montamat, and F. Roch (2022): "Robust Optimal Macroprudential Policy," *mimeo.*

Bergemann, D. and K. Schlag (2011): "Robust monopoly pricing," *Journal of Economic Theory*, 146, 2527–2543.

Bhandari, A., J. Borovička, and P. Ho (2019): "Survey data and subjective beliefs in business cycle models," NYU mimeo.

Bianchi, F., C. Ilut, and M. Schneider (2018): "Uncertainty Shocks, Asset Supply and Pricing over the Business Cycle," *The Review of Economic Studies*, 85, 810–854.

Bianchi, M. and J.-M. Tallon (2019): "Ambiguity preferences and portfolio choices: Evidence from the field," *Management Science*, 65, 1486–1501.

Bidder, R. and I. Dew-Becker (2016): "Long-run risk is the worst-case scenario," *American Economic Review*, 106, 2494–2527.

Bidder, R. and M. E. Smith (2012): "Robust animal spirits," *Journal of Monetary Economics*, 59, 738–750.

Borovička, J., L. P. Hansen, and J. A. Scheinkman (2016): "Misspecified recovery," *The Journal of Finance*, 71, 2493–2544.

Bossaerts, P., P. Ghirardato, S. Guarnaschelli, and W. R. Zame (2010): "Ambiguity in asset markets: Theory and experiment," *The Review of Financial Studies*, 23, 1325–1359.

Boyarchenko, N. (2012): "Ambiguity shifts and the 2007–2008 financial crisis," *Journal of Monetary Economics*, 59, 493–507.

Brock, W. A. and C. F. Manski (2011): "Competitive lending with partial knowledge of loan repayment: some positive and normative analysis," *Journal of Money, Credit and Banking*, 43, 441–459.

Bryan, G. (2019): "Ambiguity aversion decreases the impact of partial insurance: Evidence from African farmers," *Journal of the European Economic Association*, 17, 1428–1469.

Caballero, R. J. and A. Krishnamurthy (2008): "Collective risk management in a flight to quality episode," *The Journal of Finance*, 63, 2195–2230.

Carroll, G. (2019): "Robustness in mechanism design and contracting," *Annual Review of Economics*, 11, 139–166.

Chen, H., N. Ju, and J. Miao (2014): "Dynamic asset allocation with ambiguous return predictability," *Review of Economic Dynamics*, 17, 799–823.

Chen, Z. and L. Epstein (2002): "Ambiguity, risk, and asset returns in continuous time," *Econometrica*, 70, 1403–1443.

Coenen, G. (2007): "Inflation persistence and robust monetary policy design," *Journal of Economic Dynamics and Control*, 31, 111–140.

Collard, F., S. Mukerji, K. Sheppard, and J.-M. Tallon (2018): "Ambiguity and the historical equity premium," *Quantitative Economics*, 9, 945–993.

Condie, S. and J. Ganguli (2017): "The pricing effects of ambiguous private information," *Journal of Economic Theory*, 172, 512–557.

Condie, S. and J. V. Ganguli (2011a): "Ambiguity and rational expectations equilibria," *The Review of Economic Studies*, 78, 821–845.

——— (2011b): "Informational efficiency with ambiguous information," *Economic Theory*, 48, 229–242.

Delavande, A., E. Del Bono, and A. Holford (2021a): "Perceived Ambiguity about COVID-related Health Outcomes and Protective Health Behaviors among Young Adults," Working Paper.

Delavande, A., J. Ganguli, and F. Mengel (2021b): "Uncertainty Attitudes, Subjective Expectations and Decisions under Uncertainty," Working Paper.

Dicks, D. and P. Fulghieri (2021a): "Uncertainty, Investor Sentiment, and Innovation," *The Review of Financial Studies*, 34, 1236–1279.

Dicks, D. L. and P. Fulghieri (2021b): "Uncertainty, Contracting, and Beliefs in Organizations," UNC Kenan-Flagler, mimeo.

Dimmock, S. G., R. Kouwenberg, O. S. Mitchell, and K. Peijnenburg (2016): "Ambiguity aversion and household portfolio choice puzzles: Empirical evidence," *Journal of Financial Economics*, 119, 559–577.

Dougherty, J. P., J. E. Flatnes, R. A. Gallenstein, M. J. Miranda, and A. G. Sam (2020): "Climate change and index insurance demand: Evidence from a framed field experiment in Tanzania," *Journal of Economic Behavior & Organization*, 175, 155–184.

Dow, J. and S. R. d. C. Werlang (1992): "Uncertainty aversion, risk aversion, and the optimal choice of portfolio," *Econometrica*, 197–204.

——— (1994): "Nash equilibrium under Knightian uncertainty: breaking down backward induction," *Journal of Economic Theory*, 64, 305–324.

Drechsler, I. (2013): "Uncertainty, time-varying fear, and asset prices," *The Journal of Finance*, 68, 1843–1889.

Easley, D. and M. O'Hara (2009): "Ambiguity and nonparticipation: The role of regulation," *The Review of Financial Studies*, 22, 1817–1843.

Eichberger, J. and D. Kelsey (2002): "Strategic complements, substitutes, and ambiguity: the implications for public goods," *Journal of Economic Theory*, 106, 436–466.

Ellis, A. (2016): "Condorcet meets Ellsberg," *Theoretical Economics*, 11, 865–895.

Ellsberg, D. (1961): "Risk, Ambiguity, and the Savage Axioms," *The Quarterly Journal of Economics*, 643–669.

Epstein, L. G., E. Farhi, and T. Strzalecki (2014): "How much would you pay to resolve long-run risk?" *American Economic Review*, 104, 2680–97.

Epstein, L. G. and Y. Halevy (2021): "Hard-to-Interpret Signals," University of Toronto, mimeo.

Epstein, L. G. and M. Schneider (2007): "Learning under ambiguity," *The Review of Economic Studies*, 74, 1275–1303.

——— (2008): "Ambiguity, information quality, and asset pricing," *The Journal of Finance*, 63, 197–228.

——— (2010): "Ambiguity and asset markets," *Annual Review of Financial Economics*, 2, 315–346.

Epstein, L. G. and T. Wang (1994): "Intertemporal Asset Pricing under Knightian Uncertainty," *Econometrica*, 62, 283–322.

Epstein, L. G. and S. E. Zin (1989): "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework," *Econometrica*, 937–969.

Ferriere, A. and A. G. Karantounias (2019): "Fiscal austerity in ambiguous times," *American Economic Journal: Macroeconomics*, 11, 89–131.

Gallant, A. R., M. R Jahan-Parvar, and H. Liu (2019): "Does smooth ambiguity matter for asset pricing?" *The Review of Financial Studies*, 32, 3617–3666.

Garlappi, L., R. Giammarino, and A. Lazrak (2017): "Ambiguity and the corporation: Group disagreement and underinvestment," *Journal of Financial Economics*, 125, 417–433.

Ghirardato, P. and J. N. Katz (2006): "Indecision theory: Weight of evidence and voting behavior," *Journal of Public Economic Theory*, 8, 379–399.

Giannoni, M. P. (2002): "Does model uncertainty justify caution? Robust optimal monetary policy in a forward-looking model," *Macroeconomic Dynamics*, 6, 111–144.

Gilboa, I. and M. Marinacci (2016): "Ambiguity and the Bayesian paradigm," in *Readings in formal epistemology*, Springer, 385–439.

Gilboa, I. and D. Schmeidler (1989): "Maxmin Expected Utility with Non-unique Prior," *Journal of Mathematical Economics*, 18, 141–153.

Giustinelli, P., C. F. Manski, and F. Molinari (2021): "Precise or Imprecise Probabilities? Evidence from Survey Response Related to Late-onset Dementia," *Journal of the European Economic Association*, forthcoming.

Giustinelli, P. and N. Pavoni (2017): "The evolution of awareness and belief ambiguity in the process of high school track choice," *Review of Economic Dynamics*, 25, 93–120.

Guidolin, M. and F. Rinaldi (2013): "Ambiguity in asset pricing and portfolio choice: A review of the literature," *Theory and Decision*, 74, 183–217.

Handel, B. R. and K. Misra (2015): "Robust new product pricing," *Marketing Science*, 34, 864–881.

Handel, B. R., K. Misra, and J. W. Roberts (2013): "Robust firm pricing with panel data," *Journal of Econometrics*, 174, 165–185.

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

——— (2021): "Central Banking Challenges Posed by Uncertain Climate Change and Natural Disasters," Becker Friedman Institute for Economics Working Paper No. 2021-64.

Hansen, L. P. and T. J. Sargent (2008): *Robustness*, Princeton University Press.

——— (2010): "Fragile Beliefs and the Price of Model Uncertainty," *Quantitative Economics*, 1, 129–162.

——— (2015): "Four types of ignorance," *Journal of Monetary Economics*, 69, 97–113.

——— (2022): "Structured ambiguity and model misspecification," *Journal of Economic Theory*, 199, forthcoming.

Illeditsch, P. K. (2011): "Ambiguous information, portfolio inertia, and excess volatility," *The Journal of Finance*, 66, 2213–2247.

Illeditsch, P. K., J. V. Ganguli, and S. Condie (2021): "Information inertia," *The Journal of Finance*, 76, 443–479.

Ilut, C. (2012): "Ambiguity Aversion: Implications for the Uncovered Interest Rate Parity Puzzle," *American Economic Journal: Macroeconomics*, 4, 33–65.

Ilut, C., M. Kehrig, and M. Schneider (2018): "Slow to hire, quick to fire: Employment dynamics with asymmetric responses to news," *Journal of Political Economy*, 126, 2011–2071.

Ilut, C., P. Krivenko, and M. Schneider (2016): "Uncertainty aversion and heterogeneous beliefs in linear models," Duke Univ., mimeo.

Ilut, C. and H. Saijo (2021): "Learning, confidence, and business cycles," *Journal of Monetary Economics*, 117, 354–376.

Ilut, C. and M. Schneider (2012): "Ambiguous business cycles," NBER Working Paper No. 17900.

——— (2014): "Ambiguous business cycles," *American Economic Review*, 104, 2368–99.

Ilut, C., R. Valchev, and N. Vincent (2020): "Paralyzed by fear: Rigid and discrete pricing under demand uncertainty," *Econometrica*, 88, 1899–1938.

Izhakian, Y., D. Yermack, and J. F. Zender (2021): "Ambiguity and the tradeoff theory of capital structure," *Management Science*.

Jahan-Parvar, M. R. and H. Liu (2014): "Ambiguity aversion and asset prices in production economies," *The Review of Financial Studies*, 27, 3060–3097.

Ju, N. and J. Miao (2012): "Ambiguity, learning, and asset returns," *Econometrica*, 80, 559–591.

Kala, N. (2019): "Learning, adaptation, and climate uncertainty: Evidence from Indian agriculture," MIT Sloan School of Management, mimeo.

Karantounias, A. G. (2013): "Managing pessimistic expectations and fiscal policy," *Theoretical Economics*, 8, 193–231.

——— (2020): "Doubts about the model and optimal policy," Federal Reserve Bank of Atlanta WP 2020-12.

Kindermann, F., J. Le Blanc, M. Piazzesi, and M. Schneider (2021): "Learning about housing cost: Survey evidence from the german house price boom," NBER Working Paper 28895.

Klibanoff, P., M. Marinacci, and S. Mukerji (2005): "A smooth model of decision making under ambiguity," *Econometrica*, 73, 1849–1892.

Kocherlakota, N. and C. Phelan (2009): "On the robustness of laissez-faire," *Journal of Economic Theory*, 144, 2372–2387.

Kroner, N. (2021): "Firm-Level Uncertainty and the Transmission of Forward Guidance to Investment," UT Austin, mimeo.

Kuester, K. and V. Wieland (2010): "Insurance policies for monetary policy in the euro area," *Journal of the European Economic Association*, 8, 872–912.

Lange, A. and N. Treich (2008): "Uncertainty, learning and ambiguity in economic models on climate policy: some classical results and new directions," *Climatic Change*, 89, 7–21.

Leitemo, K. and U. Söderström (2008): "Robust monetary policy in the New Keynesian framework," *Macroeconomic Dynamics*, 12, 126–135.

Lemoine, D. and I. Rudik (2017): "Managing climate change under uncertainty: Recursive integrated assessment at an inflection point," *Annual Review of Resource Economics*, 9, 117–142.

Lemoine, D. and C. P. Traeger (2016): "Ambiguous tipping points," *Journal of Economic Behavior & Organization*, 132, 5–18.

Lensman, T. and M. Troshkin (2022): "Implications of uncertainty for optimal policies," *Journal of Economic Theory*, 199.

Levin, A., V. Wieland, and J. C. Williams (2003): "The performance of forecast-based monetary policy rules under model uncertainty," *American Economic Review*, 93, 622–645.

Levin, A. T. and J. C. Williams (2003): "Robust Monetary Policy with Competing Reference Models," *Journal of Monetary Economics*, 50, 945–975.

Li, J. (2019): "The k-armed bandit problem with multiple priors," *Journal of Mathematical Economics*, 80, 22–38.

Liu, H. and Y. Zhang (2021): "Financial uncertainty with ambiguity and learning," *Management Science*.

Lucas, R. E. (1987): *Models of business cycles*, vol. 26, Basil Blackwell Oxford.

Lucca, D. O. and E. Moench (2015): "The pre-FOMC announcement drift," *The Journal of Finance*, 70, 329–371.

Malenko, A. and A. Tsoy (2020): "Asymmetric information and security design under Knightian uncertainty," Univ. of Toronto, mimeo.

Manski, C. F. (2011): "Choosing treatment policies under ambiguity," *Annual Review of Economics*, 3, 25–49.

——— (2013): *Public policy in an uncertain world*, Harvard University Press.

Manski, C. F. and F. Molinari (2010): "Rounding probabilistic expectations in surveys," *Journal of Business & Economic Statistics*, 28, 219–231.

Marinacci, M. (2015): "Model uncertainty," *Journal of the European Economic Association*, 13, 1022–1100.

Masolo, R. and F. Monti (2020): "Ambiguity, monetary policy and trend inflation," *Journal of the European Economic Association*, 19, 839–871.

Miao, J. and A. Rivera (2016): "Robust contracts in continuous time," *Econometrica*, 84, 1405–1440.

Miao, J. and N. Wang (2011): "Risk, uncertainty, and option exercise," *Journal of Economic Dynamics and Control*, 35, 442–461.

Miao, J., B. Wei, and H. Zhou (2019): "Ambiguity aversion and the variance premium," *Quarterly Journal of Finance*, 9, 1950003.

Michelacci, C. and L. Paciello (2020): "Ambiguous policy announcements," *The Review of Economic Studies*, 87, 2356–2398.

Millner, A., S. Dietz, and G. Heal (2010): "Ambiguity and climate policy," NBER Working Paper No. 16050.

——— (2013): "Scientific ambiguity and climate policy," *Environmental and Resource Economics*, 55, 21–46.

Mukerji, S. and J.-M. Tallon (2001): "Ambiguity aversion and incompleteness of financial markets," *The Review of Economic Studies*, 68, 883–904.

——— (2004a): "Ambiguity aversion and the absence of indexed debt," *Economic Theory*, 24, 665–685.

——— (2004b): "An overview of economic applications of David Schmeidler's models of decision making under uncertainty," *Uncertainty in economic theory*, 299–318.

Nau, R. (2006): "Uncertainty aversion with second-order probabilities and utilities." *Management Science*, 52, 136–145.

Neilson, W. S. (1993): "Ambiguity aversion: An axiomatic approach using second order probabilities," Mimeo.

Nishimura, K. G. and H. Ozaki (2004): "Search and Knightian uncertainty," *Journal of Economic Theory*, 119, 299–333.

Norton, B. P., J. B. Hoel, and H. Michelson (2020): "The demand for (fake?) fertilizer: Using an experimental auction to examine the role of beliefs on agricultural input demand in Tanzania," Agricultural and Applied Economics Association Annual Meeting.

Palmer, T. and B. Stevens (2019): "The scientific challenge of understanding and estimating climate change," *Proceedings of the National Academy of Sciences*, 116, 24390–24395.

Pei, G. (2018): "Ambiguity, Pessimism and Economic Fluctuations," University of Zurich mimeo.

Pouzo, D. and I. Presno (2016): "Sovereign default risk and uncertainty premia," *American Economic Journal: Macroeconomics*, 8, 230–66.

Rabin, M. (2000): "Risk Aversion and Expected-Utility Theory: A Calibration Theorem," *Econometrica*, 68, 1281–1292.

Riedel, F. (2009): "Optimal stopping with multiple priors," *Econometrica*, 77, 857–908.

Ross, N., P. Santos, and T. Capon (2012): "Risk, ambiguity and the adoption of new technologies: Experimental evidence from a developing economy," Conference, Brazil, International Association of Agricultural Economists.

Saijo, H. (2020): "Redistribution and fiscal uncertainty shocks," *International Economic Review*, 61, 1073–1095.

Sargent, T. (1999): "Comment on 'Policy Rules for Open Economies' by Laurence Ball," in *Monetary Policy Rules*, University of Chicago Press, vol. 31, 144–154.

Sims, C. A. (2002): "Solving Linear Rational Expectations Models," *Computational Economics*, 20, 1–20.

Strzalecki, T. (2011): "Axiomatic foundations of multiplier preferences," *Econometrica*, 79, 47–73.

——— (2013): "Temporal resolution of uncertainty and recursive models of ambiguity aversion," *Econometrica*, 81, 1039–1074.

Taylor, J. B. and J. C. Williams (2010): "Simple and robust rules for monetary policy," in *Handbook of monetary economics*, Elsevier, vol. 3, 829–859.

Trautmann, S. T. and G. Van De Kuilen (2015): "Ambiguity attitudes," *The Wiley Blackwell handbook of judgment and decision making*, 1, 89–116.

Wallsten, T. S., B. H. Forsyth, and D. V. Budescu (1983): "Stability and coherence of health experts' upper and lower subjective probabilities about dose—response functions," *Organizational Behavior and Human Performance*, 31, 277–302.

Warnick, J. C. E., J. Escobal, and S. C. Laszlo (2011): "Ambiguity aversion and portfolio choice in small-scale Peruvian farming," *The BE Journal of Economic Analysis & Policy*, 11.

Woodford, M. (2010): "Robustly optimal monetary policy with near-rational expectations," *American Economic Review*, 100, 274–303.

Yoo, D. (2019): "Ambiguous information, permanent income, and consumption fluctuations," *European Economic Review*, 119, 79–96.

Young, E. R. (2012): "Robust policymaking in the face of sudden stops," *Journal of Monetary Economics*, 59, 512–527.