

# Limit Orders and Knightian Uncertainty\*

Michael Greinecker<sup>†</sup>      Christoph Kuzmics<sup>‡</sup>

May 26, 2020

## Abstract

Ambiguity averse decision-makers can behave in financial portfolio problems in ways that cannot be rationalized as subjective expected utility maximization. Indeed, [Dow and da Costa Werlang, *Econometrica* 1992] show that an ambiguity-averse decision-maker might abstain from trading an asset for a wide interval of prices; something no subjective expected utility maximizer can. Dow and da Costa Werlang assume that decision-makers know the price of an asset when trading. We show that when markets operate via limit orders instead, all investment behavior of an ambiguity-averse decision-maker is observationally equivalent to the behavior of a subjective expected utility maximizer with the same risk preferences; ambiguity aversion has no additional explanatory power.

**Keywords:** Knightian uncertainty, ambiguity aversion, subjective expected utility, financial market participation, strict dominance

**JEL Classification:** D81, G11, C72

---

\*Examples similar to our motivating example have been studied by Dominique Paper and Peter Habiger in their respective master theses written under the supervision of Christoph Kuzmics. We are grateful to both as well as to Patrick Beissner, Harald Uhlig, Frank Riedel, Jan Werner, and Michael Zierhut, as well as to seminar audiences at the University of Bielefeld, Caltech, UC Irvine, University of Minnesota, Vienna University of Business and Economics, and participants of the VII Hurwicz workshop in Warsaw for valuable comments and suggestions.

<sup>†</sup>University of Graz, Department of Economics, email: [michael.greinecker@uni-graz.at](mailto:michael.greinecker@uni-graz.at)

<sup>‡</sup>University of Graz, Department of Economics, email: [christoph.kuzmics@uni-graz.at](mailto:christoph.kuzmics@uni-graz.at)

# 1 Introduction

We show that any portfolio choice that cannot be explained by subjective expected utility maximization can also not be explained by ambiguity aversion when decision-makers have access to limit orders. This is in sharp contrast to a situation where decision-makers trade at given known prices, a setting first studied by [Dow and da Costa Werlang \(1992\)](#). In our model, a decision-maker faces uncertainty over the joint distribution of the price of an asset at the point of purchase and its final value. They can trade via limit orders that trade contingent on prices. Our results show that every choice that cannot be explained as the behavior of a subjective expected utility maximizer with the same Bernoulli utility function for some probabilistic belief must be strictly dominated by the choice of some limit order. This implies that any ambiguity averse decision maker, whose preferences satisfy the von Neumann Morgenstern axioms over the set of constant acts and a monotonicity axiom, can only choose a limit order that could also be chosen by a subjective expected utility maximizer with the same risk attitudes. Ambiguity aversion has no additional explanatory power.

For the standard Bayesian paradigm, all uncertainty can be quantified by a single probability distribution, and a rational decision-maker maximizes their expected utility with respect to this distribution. Under the Bayesian paradigm, both the probability distribution and the utility function to be maximized can be derived from personal preferences, and rational decision-makers are taken to be subjective expected utility maximizers.

The Bayesian paradigm has strong limitations as a positive theory in financial settings. Indeed, it cannot even explain the most fundamental fact of household finance in a frictionless portfolio model: The majority of households do not participate in the stock market at all.<sup>1</sup> But in standard portfolio theory, subjective expected utility maximizers participate in the stock market generically. A risk-neutral investor buys an asset if the expected value is higher than the price, (short) sells it if the expected value is lower than the price, and only abstains from trading if the expected value exactly equals the price. Since expectations are formed with respect to subjective beliefs, the latter case should rarely happen. As [Arrow \(1970\)](#) pointed out, a risk-averse investor behaves locally like a risk-neutral investor and does, therefore, also trade some amount of the asset unless the expected value exactly equals the price. Essentially, everyone should trade, at least a little bit.

---

<sup>1</sup>See, for example, [Haliassos and Bertaut \(1995\)](#) or [Campbell \(2006\)](#).

A popular explanation for this nonparticipation puzzle, due to [Dow and da Costa Werlang \(1992\)](#), gives up on the Bayesian paradigm and allows that not all people lump all forms of uncertainty together into a single, all-encompassing, probability distribution. People differentiate between calculable risk and fundamental unmeasurable uncertainty, and they are wary of the latter; people are ambiguity-averse. [Ellsberg \(1961\)](#) came up with a series of thought experiments in which one could clearly separate calculable probability in terms of draws from an urn with known composition and fundamental uncertainty in terms of an unknown urn composition. These thought experiments have been turned into actual lab experiments that confirmed that people do not behave as if all uncertainty was calculable risk; see, for example, the early survey by [Camerer and Weber \(1992\)](#) or the recent survey by [Trautmann and Van De Kuilen \(2015\)](#). Many experimental subjects are, indeed, wary of ambiguity.

[Dow and da Costa Werlang \(1992\)](#) show that the maxmin expected utility model of [Gilboa and Schmeidler \(1989\)](#) of ambiguity aversion provides a simple explanation for why people do not participate in the stock market. In the maxmin expected utility model, a decision-maker has a whole family of probability distributions and evaluates each action with respect to the lowest expected utility possible under any of these probability distributions. Even if a decision-maker's behavior is locally risk-neutral, it is not locally ambiguity neutral. The different probabilities used for evaluating buying and selling an asset drive a wedge between prices at which buying and selling, respectively, is optimal. There can be a wide interval of prices at which not trading is the unique optimal action.<sup>2</sup>

This argument has not just become the starting point for many models of ambiguity aversion in finance,<sup>3</sup> it has also become the leading example of the economic significance of ambiguity aversion outside laboratory settings in surveys such as those by [Gilboa, Postlewaite, and Schmeidler \(2008\)](#) and [Gilboa and Marinacci \(2013\)](#).<sup>4</sup>

---

<sup>2</sup>[Antoniou, Harris, and Zhang \(2015\)](#) provide empirical support from field data; ambiguity predicts abstention.

<sup>3</sup>See, for example, the extensive surveys of ambiguity aversion in finance by [Guidolin and Rinaldi \(2013\)](#) and [Epstein and Schneider \(2010\)](#).

<sup>4</sup>The mechanism [Dow and da Costa Werlang \(1992\)](#) identified has been used in various forms in much of finance research since. [Easley and O'Hara \(2009\)](#) study its impact on financial regulation. The portfolio inertia implicit in the model of [Dow and da Costa Werlang](#) also appears in the representative agent dynamic asset pricing model of [Epstein and Wang \(1994\)](#). In their model, supporting prices are robustly not unique and, therefore, equilibria locally indeterminate. [Rigotti and Shannon \(2012\)](#) show, however, that equilibria in heterogeneous agent general equilibrium models with both risk and uncertainty are generically determinate. But indeterminacy

Ambiguity aversion has also been used to explain quantitative empirical puzzles. In many models, subjective expected utility maximization can only explain the size of effects if one assumes agents to be overly risk-averse. For example, the famous equity premium puzzle of Mehra and Prescott (1985) concerns the overly high implied risk-aversion of investors that is needed when one attributes the historical difference between the returns of equity and Treasury bonds to an equilibrium risk premium in a parametric general equilibrium model. Several authors have argued that there is an additional ambiguity premium that can be studied from market data; see for example Brenner and Izhakian (2018), Izhakian (2020), and Collard, Mukerji, Sheppard, and Tallon (2018).

We show that observable differences between ambiguity averse decision-makers and subjective expected utility maximizers break down completely when decision-makers can trade with common financial instruments. Dow and da Costa Werlang (1992) assume that decision-makers only trade at a given known price. If, in contrast, decision-makers can set price contingent orders—limit orders—before prices materialize, then for every behavior that can occur in standard models of ambiguity aversion, there exists a Bayesian probabilistic belief at which their choice maximizes subjective expected utility for the same Bernoulli utility function. Ambiguity aversion is not identifiable, and the ambiguity premium must be zero.

A wide variety of models of ambiguity aversion has been used in finance, and we make our argument robust to the choice of the underlying decision-theoretic model. Our argument applies to all those models of ambiguity aversion formulated in the framework of Anscombe and Aumann (1963), which combines both uncertainty and calculable risk, that satisfy a weak monotonicity requirement, and that apply expected utility theory to choice problems that only involve calculable risk. The weak monotonicity requirement is simply that no choice can be dominated by another single deterministic choice in each state of nature. The models of ambiguity aversion we allow for include, among others, the maxmin expected utility model of Gilboa and Schmeidler (1989), the Choquet expected utility model of Schmeidler (1989), the smooth ambiguity model of Klibanoff, Marinacci, and Mukerji (2005), the variational and multiplier preference models of Maccheroni, Marinacci, and Rustichini (2006) and Hansen and Sargent

---

is robust if preferences are incomplete as in the model of Bewley (2002), as shown by Rigotti and Shannon (2005), or if there is uncertainty in the price system as well, as shown by Beissner and Riedel (2019). Billot, Mukerji, and Tallon (2020) provide a survey of these models and of how they relate. Mukerji and Tallon (2001) obtain endogenous market incompleteness from ambiguity aversion.

(2001), confidence function preferences of Chateauneuf and Faro (2009), uncertainty aversion preferences of Cerreia-Vioglio, Maccheroni, Marinacci, and Montrucchio (2011), and the incomplete preference model of Bewley (2002).<sup>5</sup>

To make our argument robust to the choice of the decision-theoretic model, we make use of a dual characterization of subjective expected utility maximization. In all the models listed above, choices that can be explained by the model but not by subjective expected utility maximization must be strictly dominated by randomized choices, but not by deterministic choices. We can explain the logic of the last statement with a variant of one of the thought experiments of Ellsberg (1961): Consider two urns, each containing a hundred balls. One urn is unambiguous and is known to contain 49 white balls and 51 black balls. The other, ambiguous, urn is filled with 100 balls that are each white or black, but the composition is not known. A decision-maker has to choose between the following bets: In bet one, the decision-maker wins a prize if a ball drawn from the ambiguous urn is black. In bet two, the decision-maker wins the same prize if a ball drawn from the ambiguous urn is white. Finally, in bet three, the decision-maker wins the same prize if a ball drawn from the unambiguous urn is white. A decision-maker choosing bet three must be ambiguity-averse, for their winning chance is only 49/100, while for any probabilistic belief about the composition of the ambiguous urn, either bet one or two (or both) must have a winning chance of at least 1/2.

Note that none of the three bets dominates another in pairwise comparisons for every composition of the ambiguous urn. But as Raiffa (1961) pointed out, a lottery that chooses bets one and two with probability 1/2 each has a higher winning chance (of 1/2) than bet three independently of the composition of the ambiguous urn. Only those two bets that can be chosen by a subjective expected utility maximizer for some probabilistic belief over the urn composition are not dominated by a randomized bet.

As Kuzmics (2017) pointed out, this is true in general. A lemma of Pearce (1984) states that in a finite two-player game in normal form, a strategy is not dominated by any mixed strategy if and only if it is a best reply to a mixed strategy of the opponent. Related results can be traced back all the way to the complete class theorem of Wald (1947) in statistical decision theory. The lemma of Pearce, when translated to a decision problem in the setting of Anscombe and

---

<sup>5</sup>Our argument does not apply to some more behaviorally flavored models of ambiguity aversion such as those of Seo (2009), Saito (2015), and Ke and Zhang (2020), that violate the weak monotonicity condition. Such models have not been used in finance and lack normative appeal.

Aumann (1963) with finitely many states, says that every decision that does not maximize subjective expected utility with respect to some probabilistic belief over the states of nature must be strictly dominated by a randomized choice. The models of ambiguity aversion listed above do not allow for choices that are dominated by deterministic, nonrandomized choices. So the only choices that can be explained by ambiguity aversion and not by subjective expected utility maximization are those that are strictly dominated by randomized choices but not by nonrandomized choices. We show that in our setting, there are no such choices. We first illustrate this claim in a simple example that abstracts from various features of our general model, such as general risk aversion or informative prices, but allows for a very simple argument.

For the general model, we need to adapt our arguments slightly since we are working in an infinite-dimensional setting. Spaces of limit orders are infinite-dimensional function spaces, and the state space is allowed to be infinite. Nevertheless, we show in Proposition 1, using a generalization of Pearce's lemma due to Battigalli, Cerreia-Vioglio, Maccheroni, and Marinacci (2016), that a limit order is not dominated by a randomly chosen limit order if and only if it maximizes subjective expected utility with respect to some probabilistic belief. While Proposition 1 corresponds to a universal decision theoretic principle, Proposition 2, our main technical result, makes crucial use of our finance setting. It states that a limit order that is strictly dominated by a randomly chosen limit order, a mixed limit order in our language, must already be strictly dominated by a single (deterministic) limit order. Therefore, as we state in the main Theorem 1, everything that can be explained by ambiguity aversion could be explained by subjective expected utility maximization with the same Bernoulli utility function as well. We show what our results means for the problem of market participation, as studied by Dow and da Costa Werlang (1992), in Proposition 3.

For many models of ambiguity aversion, an even stronger conclusion holds than the one given in Theorem 1. Proposition 4 shows that behavior that can be explained by ambiguity aversion based on a model with a set of probability distributions to model beliefs can be explained as maximizing expected utility with respect to a probabilistic belief in the closed convex hull of the underlying set of probability distributions. Our argument does, therefore, not rely on decision makers having extreme probabilistic beliefs. Proposition 5 shows that the limit orders in our arguments can be implemented by realistic market portfolios.

Here is the structure of the rest of this paper: Section 2 provides a simple

example that shows the dramatic difference limit orders can make. Section 3 presents the environment we use and our main results. Section 4 shows what our main results imply for market participation. Section 5 discusses the scope of our results, the modeling choices we make, and how one can generalize some results. Section 6 collects all proofs, except for the proof of Theorem 1 which is a direct consequence of Propositions 1 and 2.

## 2 Motivating Example


We illustrate our central point in terms of the leading example of Dow and da Costa Werlang (1992). There is one asset whose future value can be either  $-1$  or  $1$ . A risk-neutral decision-maker is faced with a price of  $p$  and asked to buy or (short) sell at most one unit of the asset. The net amount the decision-maker can buy must, therefore, lie in the interval  $[-1, 1]$ . If the decision-maker maximizes subjective expected utility with respect to some probabilistic belief  $\beta$  on the value being  $1$ , buying one unit is uniquely optimal if  $p < \beta - (1 - \beta)$ , selling one unit is uniquely optimal if  $p > \beta - (1 - \beta)$ , and everything is optimal, including not trading at all, if  $p = \beta - (1 - \beta)$ .

Suppose now that the decision-maker behaves according to the maxmin expected utility model of Gilboa and Schmeidler (1989) and that the two beliefs  $\beta_l < \beta_h$  form the extreme points of their set of beliefs. Now buying one unit is optimal if  $p \leq \beta_l - (1 - \beta_l)$ , selling one unit is optimal if  $p \geq \beta_h - (1 - \beta_h)$ , and not trading is uniquely optimal when  $p$  lies in the interval  $(\beta_l - (1 - \beta_l), \beta_h - (1 - \beta_h))$ . Ambiguity aversion explains why the decision-maker completely refrains from trading. This is the core of the argument of Dow and da Costa Werlang (1992). Implicitly, the decision-maker knows the price of the asset when trading, and we can only observe whether trade happens at the realized price.

We now consider a situation in which the decision-maker trades via (generalized) limit orders before prices realize. Such a decision-maker must have a model of how prices relate to final values. For the sake of this example, and only for the sake of this example, we follow Dow and da Costa Werlang (1992) and assume that prices do not tell us anything about the future value of the asset. Prices are distributed according to a distribution that has full support and no mass points. Finding the optimal choice for a subjective expected utility maximizer is straightforward. Their chosen limit order must maximize the

conditional expected utility given the price for almost every price. So they must choose a limit order that buys one unit if  $p < \beta - (1 - \beta)$  and sells one unit if  $p > \beta - (1 - \beta)$ , for almost every  $p$ . The limit order simply carries out what the decision-maker would have done if they knew the price of the asset.

The situation for a decision-maker in the maxmin expected utility model is different. Consider the limit order that buys one unit for prices below  $p_l = \beta_l - (1 - \beta_l)$ , sells one unit for prices above  $p_h = \beta_h - (1 - \beta_h)$ , and does not trade at prices in between. The resulting limit order is not optimal and, actually, strictly dominated by a simple threshold limit order that prescribes selling above a given price  $p_m$  and buying below this price. To see this, take  $p_m$  to be the unique point between  $p_l$  and  $p_h$  such that the probability of prices being in the intervals  $[p_l, p_m]$  and  $[p_m, p_h]$  is the same; in other words,  $p_m$  is the median price conditional on the price falling in the interval  $[p_l, p_h]$ .



Any payoff difference between the two limit orders comes from how they behave in the interval  $[p_l, p_h]$ . The original limit order generates an expected payoff of zero on this interval under both beliefs. The new threshold limit order does, conditional on the price being in the interval  $[p_l, p_h]$ , buy with probability 1/2 at the comparatively low prices between  $p_l$  and  $p_m$  and sell with probability 1/2 at the comparatively high price between  $p_m$  and  $p_h$ . This leads to an additional positive expected payoff that is independent of the beliefs; the original limit order is strictly dominated. To see this more explicitly, consider the possibility that the decision-maker can randomize, analogously to the [Raiffa \(1961\)](#) hedge in the [Ellsberg \(1961\)](#) examples, and buy and sell with probability 1/2 each for each price in  $[p_l, p_h]$ . The expected surplus would still be zero in the interval independent of the prior. But the threshold limit order we constructed also buys and sells with probability 1/2 in the interval but buys when prices are low and sells when prices are high. This generates an additional expected payoff and does not require any randomization. It can actually be shown as a straightforward corollary of our main results below that every limit order that is not a



threshold limit order is strictly dominated in the simple setting of this example. Since a threshold limit order could be chosen by a subjective expected utility maximizer whose probabilistic belief over final values has an expectation that coincides with the threshold, ambiguity aversion has no additional explanatory power.

It is tempting to explain the difference limit orders make by peculiarities of the example: Prices are uninformative; the decision-maker faces fundamental uncertainty about final values, but somehow not about prices arising in between; the decision-maker is risk-neutral. None of that matters. As we show in the next section in a setting in which the decision-maker faces ambiguity about the joint distribution of prices and final values, every choice of a limit order that cannot be rationalized as maximizing expected utility with respect to some probabilistic belief must be strictly dominated by a single, deterministic limit order. None of the standard theories of ambiguity aversion allows for choices that are dominated by deterministic choices, so they are unable to explain anything that cannot already be explained by subjective expected utility maximization. The price of the increased generality is that our proof is not constructive.

### 3 Environment and Main Result

The decision-maker faces uncertainty over which probabilistic model best describes the relationship between the price of an asset and its final value. We think of these as possible objectively correct models; all residual uncertainty conditional on the true model is objective, quantifiable risk. We follow [Keynes \(1937\)](#) here, who considered only those matters fundamentally uncertain for which “there is no scientific basis on which to form any calculable probability whatever.” Uncertainty within scientific models is taken to be objective uncertainty.

There is a compact metrizable space  $Y$  of models and for each model a joint density over prices and final values. This family of densities can be represented by a single nonnegative measurable function  $h : \mathbb{R} \times \mathbb{R} \times Y \rightarrow \mathbb{R}$  continuous in  $Y$  such that

$$\int \int h(p, x, y) \, dp \, dx = 1$$

for all  $y \in Y$ .

The decision-maker’s ultimate payoff from any net-gains from investing is

given by a continuous Bernoulli utility function  $u : \mathbb{R} \rightarrow \mathbb{R}$ .

**Remark 1.** We have assumed that the Bernoulli utility function is defined on the whole real line. This rules out certain Bernoulli utility functions such as logarithmic ones, and requires the Bernoulli utility function to be unbounded if it should be increasing and weakly concave. One could easily change the framework to have Bernoulli utility functions on domains that are bounded below. Later, the decision-maker is allowed to short-sell and, therefore, able to make large losses. If we want to have Bernoulli utility functions on a restricted domain, we would need an assumption that guarantees that final wealth levels stay in the domain of the Bernoulli utility function for all losses that can occur.

The net amount the decision-maker is allowed to invest is restricted to lie in an interval  $[b, t]$  with  $b < 0 < t$ . The bounds represent short-selling and debt constraints, respectively. To guarantee that expected utilities are defined and finite, we assume that there is an integrable function  $d : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$  satisfying

$$|u(tx - tp)|h(p, x, y) + |u(bx - bp)|h(p, x, y) \leq d(p, x)$$

for all  $(p, x) \in \mathbb{R}^2$  and all  $y \in Y$ .

**Remark 2.** This is a joint assumption on the Bernoulli utility function and the possible joint distributions over prices and payoffs. For bounded  $u$ , it holds automatically. If  $u$  is increasing and weakly concave, then the less concave the Bernoulli utility function is, the more restrictive is the assumption on the joint distribution over prices and payoffs. For the extreme case of a risk-neutral agent, it amounts to a uniform integrability condition on net-gains. Even then, the condition is fairly weak and mainly requires that the tails of all distributions vanish sufficiently fast in a uniform way. For example,  $h$  could be the density of a bivariate normal distribution or bivariate t-distribution with degrees of freedom  $2 + \epsilon$  (with  $\epsilon > 0$  arbitrarily small), and  $Y$  a compact set of pairs of means and invertible covariance matrices.<sup>6</sup> A possible choice of the bounding function  $d$  for all these cases would be a scaled-up joint density of two independent t-distributions with degrees of freedom strictly between 1 and  $1 + \epsilon$ .

In contrast to Dow and da Costa Werlang (1992), we assume that the decision-maker acts before prices are known and chooses a limit order. A *limit*

---

<sup>6</sup>These examples include, therefore, bivariate distributions with existing means and variances. Such distributions need not have a kurtosis, and, thus, include bivariate distributions with heavy tails.

*order* is a measurable function from  $\mathbb{R}$  into the set  $A = [b, t]$ . More precisely, limit orders are taken to be equivalence classes of measurable functions from  $\mathbb{R}$  to  $A$  with two such functions being equivalent if they agree outside a set of Lebesgue measure zero. We denote the set of limit orders by  $L$  and endow it with the topology of convergence in measure<sup>7</sup> and its corresponding Borel  $\sigma$ -algebra. We embed  $L$  in the space of *mixed limit orders*  $\Delta(L)$ , the space of Borel probability measures on  $L$  by identifying the deterministic limit  $l$  order with the Dirac point mass  $\delta_l$  concentrated on  $l$ . A *deterministic limit order* is a mixed limit order of the form  $\delta_l$ .

**Remark 3.** We generally allow our decision-maker to choose mixed limit orders, but, as we see below, one can interpret them as purely ancillary mathematical constructs used in proofs.<sup>8</sup> Our main result, Theorem 1, does not refer to any mixed limit orders.

There is a payoff function  $v : \mathbb{R} \times A \times \mathbb{R} \rightarrow \mathbb{R}$  given by  $v(p, a, x) = u(a(x-p))$ . The expected payoff from  $\mu \in \Delta(L)$  if the true model is  $y$  is

$$V(\mu, y) = \int \int \int v(p, l(p), x) h(x, p, y) dp dx d\mu(l).$$

We say that  $\mu' \in \Delta(L)$  *strictly dominates*  $\mu \in \Delta(L)$  if  $V(\mu', y) > V(\mu, y)$  for all  $y \in Y$ . A mixed limit order that is not strictly dominated by another mixed limit order is *undominated*. A mixed limit order that is not strictly dominated by a nonrandomized limit order is *deterministically undominated*. The following proposition gives a version of the familiar characterization of undominated choices from statistical decision theory and game theory.

**Proposition 1.** *A mixed limit order  $\mu$  is undominated if and only if there exists a probability distribution  $\beta \in \Delta(Y)$  such that  $\mu$  maximizes  $\int V(\cdot, y) d\beta(y)$ .*

We are now ready for our central technical result.

**Proposition 2.** *A mixed limit order is undominated if and only if it is deterministically undominated.*

<sup>7</sup>Essentially, we use the topology of convergence in probability for any probability measure mutually absolutely continuous with respect to Lebesgue measure. The resulting topology does not depend on the specific choice of the probability measure and is separable and completely metrizable.

<sup>8</sup>We, thus, allow for randomization over “acts.” Though the original setting of Anscombe and Aumann (1963) explicitly allowed for randomization over acts, the commonly employed simplification due to Fishburn (1970) does not. It is the latter setting that has become the basis for most axiomatic treatments of ambiguity aversion.

The strategy of the proof of Proposition 2 is as follows: Suppose a mixed limit order  $\mu$  is dominated by a mixed limit order  $\mu'$ . Instead of choosing a deterministic limit order at random, as  $\mu'$  does, one could imagine randomizing conditionally on the price by “behavioral limit orders” that randomize over  $A$ . That this makes no difference follows from results in Balder (1981) or Ghossoub (1982). They are versions of the result of Kuhn (1953) that mixed and behavioral strategies induce the same distributions over plays in extensive form games of perfect recall, or, even closer, the result of Wald and Wolfowitz (1951) on the equivalence of these two forms of randomizing in statistical decision theory. There is a natural topology on the resulting “behavioral limit orders” under which payoffs are continuous and in which deterministic limit orders are dense. This corresponds to the denseness of controls in Warga (1972) or the denseness of pure strategies in Milgrom and Weber (1985). By this denseness, one can find a deterministic limit order that still dominates  $\mu$ . A deterministic limit order that closely approximates a nontrivial (non-deterministic) behavioral limit order must oscillate strongly.

In the remainder of this section we express the key consequence of Propositions 1 and 2 in purely decision theoretic terms. The decision problem of choosing a limit order can be naturally embedded within the usual Anscombe-Aumann framework. The state space is the set  $Y$  and the set of outcomes, the possible financial gains, is the set  $\mathbb{R}$ . Each limit order  $l$  induces an Anscombe-Aumann act,  $f_l : Y \rightarrow \Delta(\mathbb{R})$ , defined by

$$\int g(r) df_l(y)(r) = \int \int g(l(p)(x - p))h(p, x, y) dp dx,$$

for each bounded measurable function  $g : \mathbb{R} \rightarrow \mathbb{R}$ . That this defines  $f_l(y)$  follows from Dudley (2002, Theorem 4.5.2.).

Let  $\mathcal{L}$  be the set of all acts that are induced by some  $l \in L$ . We call a set of (measurable) acts  $\mathcal{F}$  *rich* if it includes  $\mathcal{L}$  as well as all constant acts with values of the form  $f_l(y)$  for some  $l \in L$  and  $y \in Y$ .

Let  $\succeq$  denote a (not necessarily complete or transitive) preference relation on  $\mathcal{F}$ . We say that  $\succeq$  is *compatible* with the Bernoulli utility function  $u$  if the restriction of  $\succeq$  to the set of constant acts in  $\mathcal{F}$  can be represented by the expectation of  $u$  with respect to (the values of) these constant acts. That  $\succeq$  is compatible with a Bernoulli utility function essentially amounts to  $\mathcal{F}$  being large enough to include the convex hull of its constant acts, and for the von Neumann

Morgenstern axioms to hold on the set of constant acts in  $\mathcal{F}$ .<sup>9</sup>

For any act  $f \in \mathcal{F}$  and any state  $y \in Y$  let  $f^y$  denote the constant act that satisfies  $f^y(y') = f(y)$  for all  $y' \in Y$ . We say that the preference relation  $\succeq$  (with strict part  $\succ$ ) is *monotone* if  $f \succ g$  whenever  $f^y \succ g^y$  for every state  $y \in Y$ .

With this in place we can state our main result.

**Theorem 1.** *Let  $\succeq$  be a monotone (not necessarily complete or transitive) preference relation on a rich set of acts  $\mathcal{F}$  that is compatible with the Bernoulli utility function  $u$ , and let  $l$  be a  $\succeq$ -maximal element in the set of limit orders. Then there exists a probability distribution  $\beta \in \Delta(Y)$  such that  $l$  maximizes  $\int \int u(m) df_l(y)(m) d\beta(y)$ .*

*Proof.* Since  $\succeq$  is monotone and compatible with  $u$ , the limit order  $l$  must be deterministically undominated. By Proposition 2,  $l$  is undominated. By Proposition 1, there exists  $\beta \in \Delta(Y)$  such that  $l$  maximizes

$$\begin{aligned} \int V(l, y) d\beta(y) &= \int \int \int v(p, l(p), x) h(x, p, y) dp dx d\beta(y) \\ &= \int \int \int u(l(p)(x - p)) h(x, p, y) dp dx d\beta(y) \\ &= \int \int u(m) df_l(y)(m) d\beta(y). \end{aligned}$$

□

Theorem 1 is our central result. It shows that every choice of a limit order of an ambiguity-averse decision-maker might as well be explained as the decision of a subjective expected utility maximizer with the same Bernoulli utility function for a suitably chosen probabilistic belief.

## 4 Consequences for Market Participation

We show now what our results imply for market participation. A risk-averse subjective expected utility maximizer chooses, in general, a limit order that abstains from trading at a price  $p$  only if the subjective expected utility maximizer

<sup>9</sup>In fact, we need an additional continuity assumption since we need  $u$  to be continuous and integrable for all relevant distributions. This can be a subtle issue, since finance applications usually require unbounded Bernoulli utility functions. The appropriate continuity notion should be compatible with the allowed probability distributions; see Dillenberger and Vijay Krishna (2014) for an elegant approach.

believes that, on average, and over the models considered, the expected payoff conditional on the price equals the price. This is slightly more general than saying that each model (each distribution induced by some  $y \in Y$ ) satisfies the efficient market hypothesis at that price. At the heart of the argument is a generalization of the result of Arrow (1970) that a risk-averse expected utility maximizer behaves locally as if they were risk-neutral. This requires that the expected payoff of a risk-neutral decision-maker is well-defined, and we make an assumption to that effect.

For each  $\beta \in \Delta(Y)$ , we let  $\beta_p$  be the marginal with respect to the first coordinate of the probability measure on  $\mathbb{R} \times \mathbb{R} \times Y$  with density  $h$  under  $\lambda \otimes \lambda \otimes \beta$  with  $\lambda$  being Lebesgue measure. In order for the following definition to make sense, we assume that the function  $x \mapsto \sup_{y \in Y} |x - p| h(x, p, y)$  is integrable for almost all  $p$ . We call

$$M_\beta = \left\{ p \in \mathbb{R} \mid \int \int x h(p, x, y) dx d\beta(y) = p \right\}$$

the *martingale part* of  $\beta$ . The martingale part of  $\beta$  is the set of prices that equal the expected conditional value under this joint belief induced by  $\beta$ .

**Proposition 3.** *Let  $u$  be increasing, continuously differentiable at the origin, and assume that the function  $x \mapsto \sup_{y \in Y} |x - p| h(x, p, y)$  is integrable for almost all  $p$ . Let  $\mu$  be an undominated mixed limit order. Then there exists  $\beta \in \Delta(Y)$  such that  $\beta_p(l^{-1}(0) \setminus M_\beta) = 0$  for  $\mu$  almost all  $l$ .*

The assumption that  $u$  is continuously differentiable is satisfied for all but countably many income levels (recall that  $v$  is defined on net-gains) if the decision-maker is risk-averse; see Rockafellar (1970, Theorem 25.3.).

To understand Proposition 3, note that each belief  $\beta$  on  $Y$  induces a joint belief over prices, final values, and models. Proposition 3 says then that for each undominated mixed limit order, there is a belief  $\beta$  such that almost all deterministic limit orders, that the mixed limit order randomizes over, abstain from trading, if at all, only at prices in  $M_\beta$ .

## 5 Discussion

As stated before, Theorem 1 shows that every choice of a limit order of an ambiguity-averse decision-maker might as well be explained as the decision of

a subjective expected utility maximizer with the same Bernoulli utility function for a suitably chosen probabilistic belief.

Many models of ambiguity aversion represent uncertainty by a set  $\Pi$  of probability distributions over the states of nature. For such models, one can ask whether our rationalizing probabilistic belief may need to be more extreme than every member of  $\Pi$ . It does not. For sets of probability distributions to be interpretable as a collection of reasonable beliefs, it should be the case that a decision-maker will prefer one act over another whenever the former gives a higher expected utility with respect to every probability distribution in  $\Pi$ . Such a decision maker's preferences can be interpreted as a completion of the unanimity ordering of [Bewley \(2002\)](#). It, therefore, suffices to show that rationalizing probabilistic beliefs can be chosen to be no more extreme than every member of the set of probability distributions when preferences respect the unanimity ordering. We do so in the next proposition, which shows that behavior can then be rationalized as maximizing expected utility with respect to some belief in the closed convex hull of  $\Pi$ .

**Proposition 4.** *Assume that  $\Pi \subseteq \Delta(Y)$  is nonempty and weak\*-closed. Let  $\imath$  be a deterministic limit order and assume that there is no deterministic limit order  $\imath'$  such that  $\int V(\imath', y) d\beta(y) > \int V(\imath, y) d\beta(y)$  for all  $\beta \in \Pi$ . Then there exists a probability measure  $\beta$  on  $Y$  in the weak\*-closed convex hull of  $\Pi$  such that  $\imath$  maximizes  $\int V(\cdot, y) d\beta(y)$ .*

If the set of probability distributions on  $Y$  is already closed and convex, as it usually is in the corresponding representation results, the rationalizing belief can be chosen out of the set of probability distributions itself.

At the heart of Proposition 4 is a generalization of Proposition 1 that holds outside our setting and that might be of independent interest: If  $\Pi$  is nonempty and compact and a mixed act  $\mu$  is maximal under the unanimity ordering, then there exists a probabilistic belief  $\beta$  on  $Y$  in the closed convex hull of  $\Pi$  such that  $\mu$  maximizes expected utility with respect to  $\beta$ . To prove this, one simply applies a suitable version of the lemma of [Pearce \(1984\)](#) to a model in which the set of states of nature is replaced by  $\Pi$ . If one interprets elements of  $\Pi$  as states this way, being undominated is equivalent to being maximal in the unanimity order. So for a maximal mixed act, there exists a probabilistic belief over  $\Pi$  that rationalizes the choice of  $\mu$ . This probabilistic belief over members of  $\Pi$  averages out to the desired probabilistic belief  $\beta$  on  $Y$  in the closed convex hull of  $\Pi$ . To finish the proof of Proposition 4, one combines this generalization of Proposition

1 with arguments similar to those in the proof of Proposition 2.

In our model, all choice objects are limit orders. Not using limit orders is not an option, and one might worry that our results are only relevant for decision-makers who actually trade using limit orders. This is not the case. Suppose a decision-maker does not use a limit order and instead waits until prices are known and trades then at given prices. No matter how the decision-maker behaves in this dynamic setting, the resulting strategy will be equivalent to a limit order in the model. If this equivalent limit order is dominated, the decision-maker could do better by using an actual limit order, to begin with. A decision-maker who can reflect on their own behavior and waits in order to trade must, therefore, behave like a subjective expected utility maximizer. What matters is the option to use limit orders, not the actual use of limit orders.

One might also worry that the dominating deterministic limit orders shown to exist in Proposition 2 might be complicated measurable functions that do not correspond to anything that could be implemented in real financial markets. This is not the case. Market participants have access to so-called *stop-loss limit orders* that buy a fixed (possibly negative) amount for all prices within some interval. Positive linear combinations of such stop-loss limit orders represent, therefore, portfolios that are economically feasible. Each such portfolio corresponds to a linear combination of indicator functions of intervals. Negative weights represent stop-loss limit orders that (short) sell. We, therefore, define a *simple limit order* to be a linear combination of indicator functions of intervals. These are essentially step-functions, and standard arguments for approximating general measurable functions by step functions guarantee that we can find for each dominating deterministic limit order a dominating simple limit order.

**Proposition 5.** *A mixed limit order is deterministically undominated if and only if it is not dominated by a simple limit order.*

Finally, the formalism we use deserves some discussion. For the proof of Lemma 1, we need expected payoffs to be jointly continuous in mixed limit orders and beliefs over  $Y$ . Without some restriction on the dependence between the distributions of prices and payoffs, this is generally not possible. The approach we have taken is inspired by the existence results for Bayesian games of Milgrom and Weber (1985) and, in particular, Balder (1988), whose framework and results we rely on. That the joint distribution of prices and payoffs is



absolutely continuous with respect to a product measure corresponds to a diffuseness condition on types in Bayesian games. [Stinchcombe \(2011\)](#) discusses discontinuities that can arise without such an assumption. Given the mathematical machinery, it is straightforward to modify the setting so that Bernoulli utility functions are defined only for positive wealth levels, so that state-dependent payoffs are allowed for, or so that the decision-maker can choose more than one asset.

In general, there are problems with randomizing over measurable functions as [Aumann \(1961\)](#) showed.<sup>10</sup> The problem Aumann identifies is that the pointwise evaluation of measurable functions is, in general, not a jointly measurable mapping, no matter the  $\sigma$ -algebra one puts on the space of measurable functions. But limit orders are really equivalence classes of measurable functions that are not evaluated pointwise but by integration. The problems [Aumann \(1961\)](#) raises do, therefore, not affect our arguments.

## 6 Proofs

In what follows, we replace the product Lebesgue measure  $\lambda \otimes \lambda$  on  $\mathbb{R} \times \mathbb{R}$  by the product of two probability measures  $\pi \otimes \xi$  that are both mutually absolutely continuous with respect to  $\lambda$ . As we will see, all we really need is that  $\pi$  is atomless; no further property of Lebesgue measure is being used. The change of the underlying measures will not affect the validity of our assumptions. Let  $r_\pi$  be a nonnegative Radon-Nikodym derivative of  $\lambda$  with respect to  $\pi$  and  $r_\xi$  be a nonnegative Radon-Nikodym derivative of  $\lambda$  with respect to  $\xi$ . Let  $d' : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$  be given by  $d'(p, x) = d(p, x)r_\pi(p)r_\xi(x)$  and let  $h' : \mathbb{R} \times \mathbb{R} \times Y \rightarrow \mathbb{R}$  be given by  $h'(p, x, y) = h(p, x, y)r_\pi(p)r_\xi(x)$ . Then, by Fubini's theorem,

$$\begin{aligned} \int \int d(p, x) \, dp \, dx &= \int d(p, x) \, d\lambda \otimes \lambda(p, x) = \int r_\pi(p) \int d(p, x)r_\xi(x) \, d\xi(x) \, d\pi(p) \\ &= \int r_\pi(p)r_\xi(p)d(p, x) \, d\pi \otimes \xi = \int d' \, d\pi \otimes \xi(p, x), \end{aligned}$$

so  $d'$  is  $\pi \otimes \xi$ -integrable. Also, one can show by a similar argument that  $\int h'(p, x, y) \, d\pi \otimes \xi(p, x) = \int h(p, x, y) \, d\lambda \otimes \lambda(p, x)$  for all  $y \in Y$ . Finally, by multiplying both sides the original uniform integrability inequality by  $r_\pi(p)r_\xi(x)$ , we

<sup>10</sup>A simpler proof of Aumann's main result has been given by [Rao \(1971\)](#). For a textbook treatment of Rao's proof, see [Dudley \(2014, Section 5.2\)](#).

obtain

$$|u(tp - tx)|h'(p, x, y) + |u(bx - bp)|h'(p, x, y) \leq d'(p, x)$$

for all  $(p, x) \in \mathbb{R}^2$  and all  $y \in Y$ . So we can assume without loss of generality that our assumptions hold for the product of two probability measures.

*Proof of Proposition 1.* Let  $\Delta_\pi(\mathbb{R} \times A)$  be the space of Borel probability measures on  $\mathbb{R} \times A$  with  $\mathbb{R}$ -marginal  $\pi$ . For  $B \subseteq \mathbb{R} \times A$ , let  $1_B : \mathbb{R} \times A \rightarrow \{0, 1\}$  be the corresponding indicator function. We define  $\phi : \Delta(L) \rightarrow \Delta_\pi(\mathbb{R} \times A)$  by

$$\phi_\mu(B) = \iint 1_B(p, l(p)) d\pi(p) d\mu(l)$$

for each Borel set  $B \subseteq \mathbb{R} \times A$ . It follows from Balder (1981, Theorem 7.1) or the results in Ghossoub (1982, Section I) that  $\phi$  is a surjection. Moreover,

$$\begin{aligned} V(\mu, y) &= \iint v(p, l(p), x)h(p, x, y) d\pi \otimes \xi(p, x) d\mu(l) \\ &= \iint \int v(p, l(p), x)h(p, x, y) d\pi(p) d\mu(l) d\xi(x) \\ &= \iint v(p, a, x)h(p, x, y) d\phi_\mu(p, a) d\xi(x), \end{aligned}$$

so we can study undominated mixed limit orders in terms of  $\Delta_\pi(\mathbb{R} \times A)$ . We can identify  $\Delta_\pi(\mathbb{R} \times A)$  with a convex and compact subset of a locally convex Hausdorff topological vector space as in Balder (1988) by endowing  $\Delta_\pi(\mathbb{R} \times A)$  with the narrow topology on Young measures. It follows from the Scorza-Dragnoni Theorem, see, for example, Denkowski, Migórski, and Papageorgiou (2003, Theorem 2.5.19), that this topology coincides with the usual topology of weak convergence of measures.

We can then, abusing notation a bit, treat  $V$  as a continuous function  $V : \Delta_\pi(\mathbb{R} \times A) \times Y \rightarrow \mathbb{R}$ . We can also identify  $Y$  homeomorphically with a closed subset of the weak\*-compact set  $\Delta(Y)$  via the embedding  $y \mapsto \delta_y$ . We then extend  $V$  to a bilinear function  $V^* : \Delta_\pi(\mathbb{R} \times A) \times \Delta(Y) \rightarrow \mathbb{R}$  via integration. The function  $V^*$  is continuous by Balder (1988, Theorem 2.5). By Phelps (2001, Proposition 1.2),  $\Delta(Y)$  is, under the embedding, the closed convex hull of  $Y$ . So by Battigalli, Cerreia-Vioglio, Maccheroni, and Marinacci (2016, Lemma 1), an element  $\tau$  of  $\Delta_\pi(\mathbb{R} \times A)$  is undominated if and only if  $\tau \in \operatorname{argmax} V^*(\cdot, \beta)$  for some  $\beta \in \Delta(Y)$ .  $\square$

*Proof of Proposition 2.* One direction is trivial. For the other direction, assume

that  $\mu'$  strictly dominates  $\mu$ . By the Berge maximum theorem, the function  $\kappa \mapsto \min_y V(\kappa, y) - V(\mu, y)$  is continuous. By assumption, it achieves a strictly positive value at  $\mu'$ . To finish the proof, we make use of the fact that the set of deterministic limit orders, embedded via the function  $l \mapsto \phi_{\delta_l}$ , is dense in  $\Delta_\pi(\mathbb{R} \times A)$  when  $\pi$  is nonatomic. This denseness is familiar from the optimal control literature, the classic reference being [Warga \(1972, Theorem IV.2.6; 6\)](#). By this denseness, there exists a deterministic limit order  $l$  such that  $\min_y V(l, y) - V(\mu, y) > 0$ . Then  $\mu$  is dominated by the deterministic limit order  $l$ .  $\square$

*Proof of Proposition 3.* Let  $\mu$  be undominated. By [Proposition 1](#), there exists  $\beta \in \Delta(Y)$  such that  $\mu$  is a maximizer of  $\int V(\cdot, y) d\beta(y)$ . For this to be possible,  $\mu$  almost all  $l$  must be maximizers of  $\int V(\cdot, y) d\beta(y)$ . Such  $l$  must, by a standard argument, prescribe a conditional best response at  $\beta_P$ -almost every price  $p$ , so that

$$l(p) \in \operatorname{argmax} \iint v(p, \cdot, x) h(x, p, y) d\xi(x) d\beta(y).$$

Our general integrability condition implies that the function  $(x, y) \mapsto v(p, a, x) h(x, p, y)$  is integrable for all  $a$ . Together with the assumptions that  $x \mapsto \sup_{y \in Y} |x - p| h(x, p, y)$  is integrable for almost all  $p$  and that  $u$  is continuously differentiable at the origin, this implies that for all  $a$  in a neighborhood of the origin, the function  $(x, y) \mapsto u'(a)(x - p) h(x, p, y)$  is dominated by the function  $(x, y) \mapsto C|x - p| h(x, p, y)$  for  $C$  large enough. This dominating function is clearly integrable. Therefore, the function  $(x, y) \mapsto u'(a)(x - p) h(x, p, y)$  is integrable. [Klenke \(2014, Theorem 6.28\)](#) guarantees then that we can differentiate the integral by differentiating under the integral, so

$$\frac{\partial}{\partial a} \iint v(p, 0, x) h(x, p, y) d\xi(x) d\beta(y) = \iint u'(0)(x - p) h(x, p, y) d\xi(x) d\beta(y).$$

Since  $u$  is strictly increasing and weakly concave,  $u'(0) > 0$ . If  $\iint x \cdot h(p, x, y) d\xi(x) d\beta(y) > p$ , it is therefore strictly optimal to invest a strictly positive amount and if  $\iint x \cdot h(p, x, y) d\xi(x) d\beta(y) < p$ , it is strictly optimal to invest a strictly negative amount. In particular, investing 0 is only optimal if  $\iint x \cdot h(p, x, y) d\xi(x) d\beta(y) = p$ .  $\square$

*Proof of Proposition 4.* Let  $V^* : \Delta_\pi(\mathbb{R} \times A) \times \Delta(Y) \rightarrow \mathbb{R}$  be the bilinear continuous function introduced in the proof of [Proposition 1](#). An argument parallel to the proof of [Proposition 2](#) shows that there being no deterministic limit order  $l$  such

that  $\int V(l', y) d\beta(y) > \int V(l, y) d\beta(y)$  for all  $\beta \in \Pi$  implies that there is no  $\mu \in \Delta_\pi(\mathbb{R} \times A)$  such that  $V^*(\mu, \beta) > V^*(l, \beta)$  for all  $\beta \in \Pi$ . Applying Battigalli, Cerreia-Vioglio, Maccheroni, and Marinacci (2016, Lemma 1) to the restriction  $V^* : \Delta_\pi(\mathbb{R} \times A) \times \Pi \rightarrow \mathbb{R}$ , we obtain a probability measure  $\beta$  on  $Y$  in the closed convex hull of  $\Pi$  such that  $l$  maximizes  $V^*(\cdot, \beta) = \int V(\cdot, y) d\beta(y)$ .  $\square$

*Proof of Proposition 5.* As one sees from the proof of Proposition 2, it suffices to prove that the family of simple limit orders is dense in the space of limit orders in the topology of convergence in measure. We can metrize the topology of convergence in measure by the *Ky Fan metric*  $\alpha$  given by  $\alpha(l, l') = \inf\{\epsilon > 0 \mid \pi(|l - l'| > \epsilon) \leq \epsilon\}$ ; see Dudley (2002, Theorem 9.2.2.). By a standard argument, one can approximate each limit order arbitrarily well by a simple function  $f = \sum_{i=1}^m \lambda_i 1_{A_i}$  with the  $A_i$  disjoint. Since each  $A_i$  can be approximated from below by compact sets by Ulam's theorem, Dudley (2002, Theorem 7.1.4.), one can approximate  $f$  arbitrarily well by a simple function  $f' = \sum_{i=1}^m \lambda_i 1_{A'_i}$  with each  $A'_i$  a compact subset of  $A_i$ . The function  $f'$  is continuous on the compact set  $\bigcup_{i=1}^m A'_i$  and has, by the Tietze extension theorem, Dudley (2002, Theorem 2.6.4.), a continuous extension to all of  $\mathbb{R}$  whose range is contained in the convex hull of the range of  $f'$ . This way, one can approximate each limit order arbitrarily well by a continuous function. Clearly, one can approximate a continuous function arbitrarily well uniformly by a step function and thus a simple limit order on a compact set whose complement has arbitrarily small measure. The form of the Ky-Fan metric shows that this gives us the desired approximation.  $\square$

## References

- Francis J. Anscombe and Robert J. Aumann. A definition of subjective probability. *Annals of Mathematical Statistics*, 34:199–205, 1963.
- Constantinos Antoniou, Richard D.F. Harris, and Ruogu Zhang. Ambiguity aversion and stock market participation: An empirical analysis. *Journal of Banking & Finance*, 58:57 – 70, 2015.
- Kenneth J. Arrow. *Essays in the Theory of Risk-Bearing*. North-Holland Publishing Co., Amsterdam-London, 1970.
- Robert J. Aumann. Borel structures for function spaces. *Illinois Journal of Mathematics*, 5:614–630, 1961.

- Erik J. Balder. *Mathematical Foundations of Statistical Decision Theory: The Modern Viewpoint I*. Univ. Utrecht, Department of Mathematics, 1981.
- Erik J. Balder. Generalized equilibrium results for games with incomplete information. *Mathematics of Operations Research*, 13(2):265–276, 1988.
- Pierpaolo Battigalli, Simone Cerreia-Vioglio, Fabio Maccheroni, and Massimo Marinacci. A note on comparative ambiguity aversion and justifiability. *Econometrica*, 84(5):1903–1916, 2016.
- Patrick Beissner and Frank Riedel. Equilibria under Knightian price uncertainty. *Econometrica*, 87(1):37–64, 2019.
- Truman F. Bewley. Knightian decision theory. Part I. *Decisions in Economics and Finance*, 25(2):79–110, 2002.
- Antoine Billot, Sujoy Mukerji, and Jean-Marc Tallon. Market allocations under ambiguity: A survey. *Revue Économique*, 71(2):267–282, 2020.
- Menachem Brenner and Yehuda Izhakian. Asset pricing and ambiguity: Empirical evidence. *Journal of Financial Economics*, 130(3):503 – 531, 2018.
- Colin Camerer and Martin Weber. Recent developments in modelling preferences: Uncertainty and ambiguity. *Journal of Risk and Uncertainty*, 5(4):325–370, 1992.
- John Y. Campbell. Household finance. *The Journal of Finance*, 61(4):1553–1604, 2006.
- Simone Cerreia-Vioglio, Fabio Maccheroni, Massimo Marinacci, and Luigi Montrucchio. Uncertainty averse preferences. *Journal of Economic Theory*, 146:1275–1330, 2011.
- Alain Chateauneuf and José Heleno Faro. Ambiguity through confidence functions. *Journal of Mathematical Economics*, 45(9-10):535–558, 2009.
- Fabrice Collard, Sujoy Mukerji, Kevin Sheppard, and Jean-Marc Tallon. Ambiguity and the historical equity premium. *Quantitative Economics*, 9(2):945–993, 2018.
- Zdzisław Denkowski, Stanisław Migórski, and Nikolas S. Papageorgiou. *An Introduction to Nonlinear Analysis: Theory*. Kluwer Academic Publishers, Boston, MA, 2003.

- David Dillenberger and R. Vijay Krishna. Expected utility without bounds—a simple proof. *Journal of Mathematical Economics*, 52:143–147, 2014.
- James Dow and Sérgio Ribeiro da Costa Werlang. Uncertainty aversion, risk aversion, and the optimal choice of portfolio. *Econometrica*, 60(1):197–204, 1992.
- Richard M. Dudley. *Real Analysis and Probability*, volume 74 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 2002. Revised reprint of the 1989 original.
- Richard M. Dudley. *Uniform Central Limit Theorems*, volume 142 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, New York, second edition, 2014.
- David Easley and Maureen O’Hara. Ambiguity and nonparticipation: The role of regulation. *The Review of Financial Studies*, 22(5):1817–1843, 2009.
- Daniel Ellsberg. Risk, ambiguity, and the Savage axioms. *The Quarterly Journal of Economics*, 75(4):643–669, 1961.
- Larry G. Epstein and Martin Schneider. Ambiguity and asset markets. *Annual Review of Financial Economics*, 2(1):315–346, 2010.
- Larry G. Epstein and Tan Wang. Intertemporal asset pricing under Knightian uncertainty. *Econometrica*, 62(2):283–322, 1994.
- Peter C. Fishburn. *Utility Theory for Decision Making*. Publications in Operations Research, No. 18. John Wiley & Sons, Inc., New York-London-Sydney, 1970.
- Nassif Ghoussoub. An integral representation of randomized probabilities and its applications. In *Seminar on Probability, XVI*, volume 920 of *Lecture Notes in Mathematics*, pages 519–543. Springer, Berlin-New York, 1982.
- Itzhak Gilboa and Massimo Marinacci. Ambiguity and the Bayesian paradigm. In Daron Acemoglu, Manuel Arellano, and Eddie Dekel, editors, *Advances in Economics and Econometrics: Tenth World Congress*, volume 1 of *Econometric Society Monographs*, pages 179—242. Cambridge University Press, 2013.
- Itzhak Gilboa and David Schmeidler. Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics*, 18(2):141–153, 1989.

- Itzhak Gilboa, Andrew W. Postlewaite, and David Schmeidler. Probability and uncertainty in economic modeling. *Journal of Economic Perspectives*, 22(3): 173–88, 2008.
- Massimo Guidolin and Francesca Rinaldi. Ambiguity in asset pricing and portfolio choice: A review of the literature. *Theory and Decision*, 74(2):183–217, 2013.
- Michael Haliassos and Carol C. Bertaut. Why do so few hold stocks? *The Economic Journal*, 105(432):1110–1129, 1995.
- Lars Peter Hansen and Thomas J. Sargent. Robust control and model uncertainty. *American Economic Review*, 91(2):60–66, 2001.
- Yehuda Izhakian. A theoretical foundation of ambiguity measurement. *Journal of Economic Theory*, 187:105001, 2020.
- Shaowei Ke and Qi Zhang. Randomization and ambiguity aversion. *Econometrica*, 2020. forthcoming.
- John M. Keynes. The general theory of employment. *The Quarterly Journal of Economics*, 51(2):209–223, 1937.
- Achim Klenke. *Probability Theory*. Universitext. Springer, London, second edition, 2014. A comprehensive course.
- Peter Klibanoff, Massimo Marinacci, and Sujoy Mukerji. A smooth model of decision making under ambiguity. *Econometrica*, 73(6):1849–1892, 2005.
- Harold W. Kuhn. Extensive games and the problem of information. In *Contributions to the theory of games, vol. 2*, Annals of Mathematics Studies, no. 28, pages 193–216. Princeton University Press, Princeton, N. J., 1953.
- Christoph Kuzmics. Abraham Wald’s complete class theorem and Knightian uncertainty. *Games and Economic Behavior*, 104:666–673, 2017.
- Fabio Maccheroni, Massimo Marinacci, and Aldo Rustichini. Ambiguity aversion, robustness, and the variational representation of preferences. *Econometrica*, 74(6):1447–1498, 2006.
- Rajnish Mehra and Edward C. Prescott. The equity premium: A puzzle. *Journal of Monetary Economics*, 15(2):145 – 161, 1985.

- Paul R. Milgrom and Robert J. Weber. Distributional strategies for games with incomplete information. *Mathematics of Operations Research*, 10(4):619–632, 1985.
- Sujoy Mukerji and Jean-Marc Tallon. Ambiguity aversion and incompleteness of financial markets. *The Review of Economic Studies*, 68(4):883–904, 2001.
- David G. Pearce. Rationalizable strategic behavior and the problem of perfection. *Econometrica*, 52:1029–50, 1984.
- Robert R. Phelps. *Lectures on Choquet's Theorem*, volume 1757 of *Lecture Notes in Mathematics*. Springer-Verlag, Berlin, second edition, 2001.
- Howard Raiffa. Risk, ambiguity, and the Savage axioms: Comment. *The Quarterly Journal of Economics*, 75(4):690–694, 1961.
- B. V. Rao. Borel structures for function spaces. *Colloquium Mathematicum*, 23: 33–38, 1971.
- Luca Rigotti and Chris Shannon. Uncertainty and risk in financial markets. *Econometrica*, 73(1):203–243, 2005.
- Luca Rigotti and Chris Shannon. Sharing risk and ambiguity. *Journal of Economic Theory*, 147(5):2028–2039, 2012.
- Ralph Tyrrell Rockafellar. *Convex Analysis*. Princeton Mathematical Series, No. 28. Princeton University Press, Princeton, N.J., 1970.
- Kota Saito. Preferences for flexibility and randomization under uncertainty. *American Economic Review*, 105(3):1246–1271, 2015.
- David Schmeidler. Subjective probability and expected utility without additivity. *Econometrica*, 57:571–587, 1989.
- Kyoungwon Seo. Ambiguity and second-order belief. *Econometrica*, 77:1575–1605, 2009.
- Maxwell B. Stinchcombe. Balance and discontinuities in infinite games with type-dependent strategies. *Journal of Economic Theory*, 146(2):656–671, 2011.
- Stefan T. Trautmann and Gijs Van De Kuilen. Ambiguity attitudes. *The Wiley Blackwell Handbook of Judgment and Decision Making*, 1:89–116, 2015.



Abraham Wald. An essentially complete class of admissible decision functions. *The Annals of Mathematical Statistics*, 18(4):549–555, 1947.

Abraham Wald and Jacob Wolfowitz. Two methods of randomization in statistics and the theory of games. *Annals of Mathematics. Second Series*, 53:581–586, 1951.

Jack Warga. *Optimal Control of Differential and Functional Equations*. Academic Press, New York-London, 1972.