

THE AMBIGUITY AVERSION LITERATURE: A CRITICAL ASSESSMENT

NABIL I. AL-NAJJAR AND JONATHAN WEINSTEIN

Northwestern University

We provide a critical assessment of the ambiguity aversion literature, which we characterize in terms of the view that Ellsberg choices are rational responses to ambiguity, to be explained by relaxing Savage's Sure-Thing principle and adding an ambiguity-aversion postulate. First, admitting Ellsberg choices as rational leads to behaviour, such as sensitivity to irrelevant sunk cost, or aversion to information, which most economists would consider absurd or irrational. Second, we argue that the mathematical objects referred to as "beliefs" in the ambiguity aversion literature have little to do with how an economist or game theorist understands and uses the concept. This is because of the lack of a useful notion of updating. Third, the anomaly of the Ellsberg choices can be explained simply and without tampering with the foundations of choice theory. These choices can arise when decision makers form heuristics that serve them well in real-life situations where odds are manipulable, and misapply them to experimental settings.

1. INTRODUCTION

Daniel Ellsberg's (1961) thought experiment spawned an extensive literature attempting to incorporate ambiguity-sensitive preferences in a subjective setting. The seminal works of Schmeidler (1989) and Gilboa and

We are grateful to Drew Fudenberg, Edi Karni, Bart Lipman, Marciano Siniscalchi, Costis Skiadas and Rakesh Vohra for detailed comments that substantially improved the paper. We also thank Luciano de Castro, Eddie Dekel, Daniel Gottlieb, Faruk Gul, Yoram Halevy, Peter Klibanoff, Stephen Morris, Sujoy Mukerji, Emre Ozdenoren, seminar participants at LSE, NYU, Penn, UCL and Washington for helpful conversations, and Aygun Dalkiran, Mallesh Pai and Pablo Schenone for excellent research assistance.

Schmeidler (1989) laid a formal foundation for this enterprise by modifying Savage's (1954) subjective expected utility model. Two decades after Gilboa and Schmeidler's contributions, there is an abundance of elegant formal models with considerable technical refinements and elaborate representations. There have also been some applications of these ideas to finance, contracting, auctions and macroeconomics.

This paper provides a critical assessment of the *ambiguity aversion literature*. We define this literature as the body of works that adopt the following three methodological positions: (1) Ellsberg choices are expressions of *rational* decision makers facing ambiguity; (2) ambiguity is to be modelled by relaxing the sure thing principle while keeping other aspects of Savage's subjective framework intact;¹ and (3) the decision maker's attitude towards ambiguity is a matter of taste.

The present paper questions the feasibility of coherently extending the Bayesian model along these lines. We will argue that doing so fundamentally contorts the concepts of beliefs and updating, and ends up creating more paradoxes and inconsistencies than it resolves. We build a case centred on three arguments:

- *Replacing one anomaly by other anomalies*: If we admit Ellsberg choices as rational, we must also admit choices most economists would consider absurd or irrational. Using simple examples, in Sections 2 and 3 we show that one must consider rational decision makers who base their decisions on irrelevant sunk cost; update their beliefs based on taste, and not just information; have the ability to deform their beliefs at will; or express an aversion to information. These examples highlight that adopting the ambiguity aversion literature's modelling approach comes at a substantial cost, a cost we believe most economists would be unwilling to bear.
- *Interpretation of beliefs*: The ambiguity aversion literature has generated exotic mathematical objects to be interpreted as "beliefs": sets of priors, capacities, second order probabilities, to name a few. In Section 4 we argue that these objects have little to do with how an economist understands and uses the concept of "beliefs". In economic models, beliefs only change in response to new information. In ambiguity aversion models, either belief updating is based on things other than information, or the decision maker anticipates reversals in how he interprets future evidence.

¹ Here, we will use sure thing principle and the substitution axiom interchangeably. Even more narrowly, what we are concerned with is the lack of probabilistic sophistication, i.e. the failure of Machina and Schmeidler's (1992) P4*.

- *Interpreting the Ellsberg experiment:* The all-consuming concern of the ambiguity aversion literature is the Ellsberg “paradox”.² In Section 5 we argue that this seemingly anomalous behaviour can be explained, without tampering with the foundations of choice theory, using standard tools of information economics and game theory. The approach based on standard tools offers insights into what causes Ellsberg choices, and how these choices may change with the environment. The ambiguity aversion literature, by contrast, accommodates experimental anomalies by relaxing foundational assumptions. Variations in behaviour are ascribed to inexplicable differences in tastes, while the increasingly permissive sets of axioms in this literature weaken the (already modest) modelling discipline of the Bayesian approach.

The problems we raise are neither new nor unknown to insiders of the ambiguity aversion literature. But where insiders see isolated curiosities and minor inconveniences, we see fundamental hurdles that put into question the entire enterprise. Our contribution is therefore a synthesis of seemingly disparate facts, ideas and examples into a compelling case questioning the methodological underpinnings of this literature. This paper is not a comprehensive survey. Rather, its aim is to help the reader cut through the complexity of a literature accessible mainly to insiders with a substantial knowledge of the requisite technical and axiomatic machinery. By illustrating the main ideas with simple examples, we hope to make it easier for an outsider, be it a theorist, or an economist concerned with applications in finance or macroeconomics, to fully appreciate the implications of these models.

Our arguments are not limited to specific functional forms or systems of axioms, but extend to *all* ambiguity models adhering to the three methodological premises outlined earlier: that preferences are consistent with Ellsberg choices, obtained by relaxing the sure thing principle, and incorporate a distaste for ambiguity. We do not comment on other approaches such as the incomplete preference approach of Bewley (1986, 2002), minimax regret of Savage (1951), or models inspired by classical statistics.³

We hasten to add that we do not deny the intuitive and rhetorical appeal of introducing ambiguity in decision making. Indeed, the questions

² A *paradox* is “an apparently true statement that leads to a contradiction”. Referring to choices in Ellsberg’s thought experiment as “paradoxical” implicitly confers on them an aura of rationality. Since we question the rationality of these choices, we prefer the more neutral term *anomaly*, which refers to “a deviation, irregularity, or an unexpected result”.

³ For instance, Al-Najjar (2009) where ambiguity arises from the difficulty of uniform learning in the sense of statistical learning theory.

motivating the ambiguity aversion literature – such as where do beliefs come from, or how to account for decision maker’s “model uncertainty” – are important and may, one day, be addressed in a satisfactory manner. We simply question the value of an approach that relegates these issues to matters of taste, while interpreting as “beliefs” probability-like objects that lack the tractability of the Bayesian model.

A leading interpretation of the ambiguity aversion literature is that the Ellsberg choices are *rational* responses by decision makers to a lack of reliable information that prevents them from forming beliefs with confidence.⁴ The rationality of the Ellsberg choices is not just a matter of semantics: it means that the ambiguity aversion literature does not view itself as a branch of behavioural economics, preoccupied with the study of biases and mistakes. Rather, this literature positions itself as an extension of the standard Bayesian paradigm, an extension made necessary by this paradigm’s unduly rigid conception of rational choice under uncertainty. The central task of the ambiguity aversion models, then, is to characterize ambiguity-sensitive behaviour in terms of normatively compelling axioms.

Are the Ellsberg choices “rational”? Since theories of decision making *define* what constitutes rational behaviour, the risk of circularity is obvious. So we adopt Gilboa and Schmeidler’s (2001: 17–18) criterion which we find intuitive and fairly neutral: “An action, or sequence of actions is rational for a decision maker if, when the decision maker is confronted with an analysis of the decisions involved, but with no additional information, she does not regret her choices.” Stated differently, a decision is rational if it is *immune to introspection*.⁵

To put this in perspective, consider Tversky and Kahneman’s (1983) classic experiments where subjects frequently judge an event as less likely than one of its subevents. This “conjunction fallacy” is irrational under this definition because, once the fallacy is explained, typical subjects will recognize their error and “feel embarrassed”. The rational vs. irrational distinction is *with* a difference: one would expect the forces of learning, introspection and incentives to make decision makers unlikely to repeat the same errors in the future. Charness, Karni, and Levin (2008) show, for example, that the likelihood of committing the conjunction fallacy drops significantly if subjects are either offered small monetary incentives or allowed to consult with each other.

To answer the question whether Ellsberg choices are “rational”, we confront a decision maker who expresses these choices with some of their implications in simple dynamic settings. If static Ellsberg choices are

⁴ For example, Epstein and Le Breton (1993) write that the “[Ellsberg] choices seem sensible at a normative level, since they correspond to an aversion to imprecise information”.

⁵ We thank Peter Klibanoff for suggesting this phrase.

indeed rational, and not just a behavioural bias or anomaly, then they ought to be immune to the decision maker's introspection not just about the static choices, but also about their dynamic implications. Our examples show that these choices may lead the decision maker to absurd consequences involving the most rudimentary of normative economic principles.

A striking example concerns how decision makers deal with irrelevant sunk cost. Consider a problem where the decision maker may or may not make an irreversible sunk expenditure. A piece of information is revealed, at which point he must make a further decision contingent on this information. Neither the information nor the payoffs in the contingent decision problem are affected by the sunk cost. In our example, a dynamically consistent, ambiguity-averse decision maker ought to condition his choices on whether the (now sunk) cost was incurred. Teachers of economics know that students are resistant to the idea of ignoring sunk cost. This, of course, does not make the incorporation of sunk cost any less flawed or irrational.

Our other examples provide additional evidence of the irrationality of ambiguity-sensitive decision makers, such as deforming their beliefs at will, expressing an aversion to information, or selecting dominated choices. The broader lesson is that the apparent reasonableness of Ellsberg choices in static settings is deceptive. These static choices do not confront the decision maker with some of the more interesting questions facing an economic actor, namely those involving dynamic choice and information.

It is the scrutiny of dynamic settings that reveals the extent to which a decision maker ought to view the Ellsberg choices as absurd and embarrassing. This scrutiny is appropriate since applications of ambiguity models to economics and finance involve information and dynamic choice. In fact, static models are frequently a stand-in for an incompletely modelled dynamic situation. As Epstein and Le Breton (1993: 2) write: "a satisfactory treatment of updating is a prerequisite for fruitful application of models of non-Bayesian beliefs [. . .] whether to intertemporal problems, game theory, or statistical theory". In a similar vein, Gilboa and Schmeidler (1993: 35) write: "the theoretical validity of any model of decision making under uncertainty is quite dubious if it cannot cope successfully with the dynamic aspect". To sum up, we take rationality to be whole: a decision-making paradigm should not selectively pick and choose when its behavioural implications are rational and when they are not.⁶

Some readers may view debates over whether the Ellsberg choices are rational as overly concerned with semantics. What matters, the argument

⁶ As Machina (1989) writes: "Whereas experimental psychologists can be satisfied as long as their models of individual behavior perform properly in the laboratory, economists are responsible for the logical implications of their behavioral models when embedded into social settings."

goes, is that these choices are empirically relevant, not whether they are behavioural anomalies or expressions of rational choice. The contribution of the ambiguity aversion literature, according to this interpretation, is to provide convenient functional forms to fit the data.

Ambiguity models routinely assume decision makers who both display the Ellsberg choices yet rationally perform tasks like solving for complex intertemporal saving/consumption plans, or calculating optimal portfolios. Admitting the static Ellsberg choices as just another behavioural anomaly would be in conflict with requiring these same irrational decision makers to pursue rationality to the fullest in every remaining aspect of the model. The issue of rational vs. behavioural interpretation of ambiguity models is discussed in greater detail in Section 6.

In Section 4 we investigate the root cause of the paradoxes that seem to accompany extensions of the ambiguity aversion literature to dynamic settings. A fundamental achievement of Savage's subjective expected utility theory is the decomposition of preferences into tastes and beliefs. Beliefs represent the part of the preference that is updated to incorporate new evidence and equilibrated in a strategic interaction. Central to an economist's or game theorist's use of beliefs is *fact-based updating*, i.e. that beliefs are updated based on facts, and *only* facts. Beliefs in Savage's theory are not merely what is left after extracting the taste component of the preference. Rather, we take Savage's notion of subjective beliefs and his celebrated separation of tastes from beliefs to be meaningful only in so far as updating is fact-based.

Like Savage's theory, ambiguity models separate taste from probability-like objects – such as capacities or sets of priors – and interpret these objects as “beliefs”. But this is where the similarity ends: updating these objects either leads to behaviour that is dynamically inconsistent, or fails to be based on facts alone. In a paper titled “Dynamically Consistent Beliefs must be Bayesian”, Epstein and Le Breton (1993: 5) write that “there does not exist a “satisfactory” decision theoretic foundation for *any* rule for updating vague [...] beliefs”. Our examples suggest that the subsequent attempts to provide a satisfactory theory to updating ambiguity preferences did not fare any better. This may explain why the use of ambiguity models in games, where updating is central, have been minimal, even though games seem like the most natural setting where ambiguity would arise.

In Section 5 we scrutinize the core empirical justifications of ambiguity models, namely their ability to account for Ellsberg choices in experiments. We note that these choices are equally consistent with behavioural models that view them as “mistakes”. One such explanation is that Ellsberg-style choices are the result of decision makers misapplying heuristics that serve them well in real-life situations where the odds may be

subject to manipulation by an opponent.⁷ This point, already made by Myerson (1991) and others, account for the Ellsberg anomaly in experiments by viewing decision makers as carrying over these heuristics to artificial experimental settings where probabilities are not manipulable. Behavioural explanations can account for Ellsberg choices in standard experiments simply and without creating additional anomalies. Moreover, by recognizing Ellsberg choices as mistakes, a behavioural approach also provides a more subtle understanding because it distinguishes between situations with manipulable probabilities, when ambiguity aversion is a rational strategic response, from non-manipulable situations when it is a mistake. The ambiguity aversion literature simply attributes ambiguity aversion in any setting to the decision-maker's taste.

A common argument is that ambiguity models should be viewed not as models of rational behaviour, but as *descriptive* accounts of an experimentally important phenomenon. In Section 6 we discuss this descriptive interpretation, and we find it wanting. It is not surprising that ambiguity models can rationalize a broader range of behaviour, since they also substantially increase the degrees of freedom available to the modeller. The important question is at what price does this improved fit come? As discussed earlier, the ambiguity aversion literature generates a plethora of new anomalies and paradoxes. We discuss in Section 5 the experimental findings of Halevy (2007) that Ellsberg choices are highly correlated with subjects not reducing objective compound lotteries, casting doubts on the value of an axiomatic approach to understanding these choices. What justifies requiring agents who, on the one hand, fail to reduce objective compound lotteries, commit the sunk cost fallacy, or display an aversion to information, to nevertheless follow rationally motivated updating rules or carry complicated optimal portfolio calculations?

2. TRADING OFF ANOMALIES

We focus on simple examples and draw conclusions that, with few exceptions, hold for all ambiguity models. We start by describing the Ellsberg anomaly and fixing some notation and baseline assumptions.

2.1 Rationalizing the Ellsberg anomaly

To fix ideas, we briefly describe the Ellsberg anomaly. Consider a decision problem with three states representing the colours of balls in an urn:

⁷ For an interesting alternative explanation that views Ellsberg choices as behavioural mistakes, see the recent paper by Gottlieb (2009) and the references therein.




FIGURE 1. The Ellsberg anomaly.

b (lack), r (ed), and y (ellow). A decision maker chooses among the following acts:

	\underline{b}	\underline{r}	\underline{y}
f_1 :	10	0	0
f_2 :	0	10	0
f_3 :	10	0	10
f_4 :	0	10	10

These acts are illustrated in Figure 1. The dotted lines indicate information sets (in this case, highlighting the fact that the colour of the ball is unknown). Thus, the act f_1 pays 10, if the ball is black, 0 if red or yellow, and so on. Let \succsim and $>$ denote the decision maker's weak and strict preference over acts, respectively.

The decision maker knows the urn contains 120 balls, of which 40 are black. The ratio of the other two colours is unknown. Ellsberg argued that a rational decision maker may display the following choices: $f_1 \succ f_2$ and $f_4 \succ f_3$.

Such preferences are inconsistent with choices made based on probabilities. For if there were a probability measure P underlying these choices, then $f_1 \succ f_2$ implies $P(b) > P(r)$, while $f_4 \succ f_3$ implies $P(r) + P(y) > P(b) + P(y)$; a contradiction.

There are several axiomatized representations with preferences consistent with the Ellsberg choices. Our arguments apply to all of the ambiguity aversion models, and the problems we identify are inherent in the approach, and not artefacts of a particular formalization. Nevertheless, we sometimes find it helpful to illustrate our points using a specific functional form.

The most popular of ambiguity models is Gilboa and Schmeidler's (1989) *minimax expected utility* (MEU) under which the preference is

represented by the utility function:⁸

$$V(f) = \min_{P \in C} E_P v(f(s))$$

where C is a compact and convex set of probability measures. When C is a singleton, this reduces to standard expected utility. Otherwise, C captures the idea that the decision maker is unsure about the probability to assign to each event and takes a pessimistic, ambiguity-averse, attitude to the evaluation of acts (hence the “ $\min_{P \in C}$ ”).

Example 1 (MEU example). *The Ellsberg choices above are consistent with an MEU representation with a set of priors $C \subset \Delta$ defined by: $P \in C$ iff*

- $P(b) = \frac{1}{3}$;
- $P(r), P(y) \in [0, \frac{2}{3}]$;
- $P(r) + P(y) = \frac{2}{3}$.

More generally, Cerreia *et al.* (2008) show that any preference that displays ambiguity aversion (plus other standard axioms) can be represented by a utility function:

$$(1) \quad V(f) = \inf_{P \in \mathcal{P}} [E_P v(f(s)) + c(E_P v(f(s)), P)],$$

where $c : \Delta \rightarrow \mathcal{R}$ is a non-negative cost function with additional properties that need not concern us here. The interpretation is that Nature chooses a distribution P to minimize the decision maker's expected payoff but must pay him a cost $c(E_P v(f(s)), P)$.

2.2 Assumptions and notation

Our critique of the ambiguity aversion literature is not confined to a particular model (e.g. MEU) but extends to any model of ambiguity averse preferences. We will therefore impose only minimal assumptions on preferences beyond consistency with the Ellsberg choices. In particular, we shall assume that the set of prizes is the set of real numbers, and that the preference \succsim is a complete order on the set of acts, denoted \mathcal{F} , satisfying continuity and dominance.⁹ These properties are natural, and orthogonal to the conceptual issues of modelling ambiguity. They are assumed in all of the axiomatic representations we are aware of.

For simplicity of exposition, we further assume a finite state space and that the decision maker is risk neutral.

⁸ Other models include Maccheroni, Marinacci and Rustichini (2006a) and Klibanoff, Marinacci and Mukerji (2005). These are covered by our critiques, discussed below.

⁹ Continuity means that the set of acts $\{f : f \succ g\}$ is open for any act g , while dominance means that, for any pair of acts f and g , $f(x) \geq g(x)$ for every x implies $f \succsim g$.

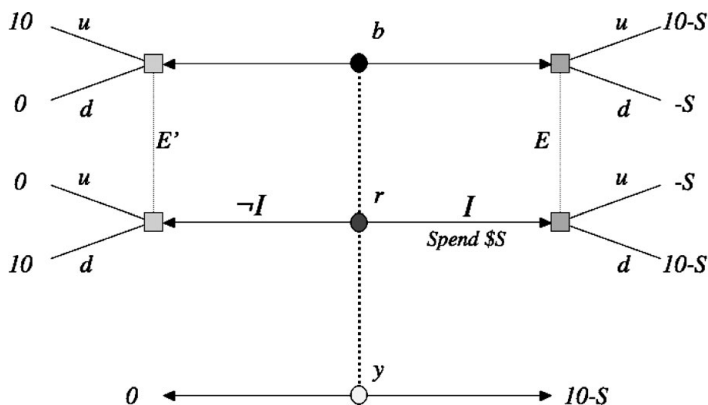


FIGURE 2. Sunk cost.

The *updated* or *conditional preference* at E , denoted \succcurlyeq_E , is a preference relation on \mathcal{F} which we shall interpret as the decision-maker's preference over acts given the information that E occurred. Intuitively, \succcurlyeq_E should be related to the *ex ante*, or unconditional preference \succcurlyeq . Since we wish to illustrate our points under a variety of assumptions about the relationship between \succcurlyeq and \succcurlyeq_E , we do not impose any a priori restrictions on this relationship.

2.3 Sunk cost

Our first example shows that, in a dynamic setting, a decision maker with Ellsberg preferences may base his choice on irrelevant sunk costs. Unlike all subsequent examples, we assume here that the decision maker has the commitment power to be dynamically consistent.

In Figure 2, a decision maker must first decide whether to commit a sunk cost of S dollars. After making the decision to invest (denoted I) or not (denoted $-I$), he learns whether y did occur. If y does not occur, he must choose u or d , yielding the payoffs indicated in the figure. Compared to $-I$, investing here amounts to paying S dollars in exchange for improving the payoff when y occurs from 0 to 10. Aside from that, the information structure and available choices remain unchanged. The payoffs shown in the figure are net of the sunk cost.

The sequence of choices $(-I, u)$ and $(-I, d)$ correspond to f_1 and f_2 , while (I, u) and (I, d) correspond to the acts:

$$f_3 - S \quad \& \quad f_4 - S$$

which are the acts appearing in the Ellsberg example minus the sunk cost. The magnitude of the sunk cost is an exogenous parameter known to the decision maker. Because of dominance, the decision maker chooses I when

$S = 0$, and $-I$ when $S = 10$. By continuity, there must be a value \bar{S} at which he is indifferent between investing and not investing.

Suppose that \succsim is a preference that satisfies the following properties:

1. Ellsberg choices: $f_1 \succ f_2$ and $f_4 \succ f_3$.
2. Dynamic consistency: $f_1 \succ_{E'} f_2$ and $f_4 \succ_E f_3$.

Condition 2 is a weak form of dynamic consistency: it requires only that the *ex ante* optimal plan remains optimal after information is received.

To show that this ambiguity-sensitive decision maker takes irrelevant sunk cost into account, we consider two cases depending on whether his preference \succsim satisfies *additive invariance*: For any acts f and g and constant α ,

$$f \succ g \iff f + \alpha \succ g + \alpha.$$

To motivate this, recall our assumption of risk neutrality so acts are util-valued.¹⁰ A Bayesian decision maker with prior p always satisfies this condition, since $E_p(f + \alpha) = E_p f + \alpha$.

For general ambiguity preferences, the situation is more complicated, so we consider two cases. First, suppose that additive invariance fails. Then the decision maker's preference reverses as a consequence of the addition of a constant, so he ends up taking irrelevant sunk cost into account almost by definition. In this case, examples simpler than the one in Figure 2 would suffice to illustrate the point.¹¹

If \succsim satisfies additive invariance,¹² then the assumption that \succsim is consistent with Ellsberg choices (fact 1 above) implies that

$$f_4 - S \succ f_3 - S.$$

By dynamic consistency, we have:

$$f_1 \succ_{E'} f_2 \quad \text{yet} \quad f_4 - S \succ_E f_3 - S.$$

To see why this implies the incorporation of sunk cost, note that the choice problem at E differs from that at E' only by the fact that payoffs are scaled

¹⁰ The conditions would have to be slightly modified to cover the case of non-linear u .

¹¹ For example, if $f \succ g$ but $g - S \succ f - S$ then an objective lottery that picks a cost of either S or 0 with equal probability would reverse the comparison between the acts f and g even though this cost cannot be influenced by the decision maker.

¹² This is the case for some of the most important classes of preferences. Formally, the condition holds for any preference that admits a representation (1) with cost function that is additive in its two arguments. This includes, in particular, variational preferences, introduced by Maccheroni, Marinacci and Rustichini (2006a), which in turn include all MEU preferences and smooth ambiguity preferences (Klibanoff, Marinacci and Mukerji, 2005) with CARA ambiguity attitude.

down by a constant. Yet this constant, which reflects costs already sunk at an earlier stage, influences choice.

Should a rational decision maker be embarrassed by these choices? The ambiguity aversion literature is founded on the premise that the *static* Ellsberg choices, $f_1 \succ f_2$ and $f_4 \succ f_3$, are rational. Our view is that in testing for the rationality of a set of choices, one should confront the decision maker with the full set of implications of his preference, including what his preference would imply for his conditional choices as the parameters of the decision problem change. The sunk cost S is not relevant to what the decision maker should do at the information sets E and E' . However, as this cost crosses the threshold \bar{S} , and he changes his decision to invest, his second-period choice flips, although nothing about the second-period scenarios differs except for the sunk cost. How would a decision maker justify such choices as rational? Ignoring irrelevant sunk cost is one of the most basic lessons in economics education. From a normative point of view, at least by the standards of undergraduate textbooks, to behave differently in two otherwise identical situations is very embarrassing indeed.

We remind the reader that this argument is orthogonal to whether people are prone to making errors of judgment, such as the incorporation of irrelevant sunk cost in their decisions. The point is that the mere fact that errors are widespread does not alter their character as errors. Thus, the prevalence of sunk cost-related errors among undergraduates is generally not viewed as sufficient reason to revise the undergraduate curriculum. Commonplace errors of judgment, of course, provide fascinating and worthwhile topics of study in psychology and behavioural economics.

We conclude with a numerical example illustrating these points.

Example 2 (MEU example: sunk cost). Consider a decision maker with MEU preferences with the set of priors $C \subset \Delta$ introduced in Example 1.

The threshold for investing in this example turns out to be $\bar{S} = \frac{10}{3}$. For the purpose of illustration, assume that $S = 3$, so the decision maker invests. Then:

$$f_4 - 3 \succ f_3 - 3, \quad f_1 \succ f_2 \quad \text{and} \quad f_4 - 3 \succ f_1.$$

It is easy to calculate:

- $V(f_3 - 3) = 7 - 10 \max_{P \in C} P(r) = \frac{1}{3}$;
- $V(f_4 - 3) = \frac{11}{3}$;
- $V(f_1) = \frac{10}{3}$;
- $V(f_2) = 10 \min_{P \in C} \{P(r)\} = 0$;
- Therefore, $V(f_4 - 3) > V(f_3 - 3)$ and $V(f_1) > V(f_2)$.¹³

¹³ Calculations: $V(f_3 - 3) = \min_{P \in C} \{\frac{7}{3} - 3P(r) + (\frac{2}{3} - P(r))7\} = \min_{P \in C} \{7 - 10P(r)\}$;
 $V(f_4 - 3) = \min_{P \in C} \{-\frac{3}{3} + 7P(r) + (\frac{2}{3} - P(r))7\} = \frac{14}{3} - 1$; $V(f_1) = \min_{P \in C} \{\frac{10}{3} + 0P(r) + (\frac{2}{3} - P(r))0\}$; $V(f_2) = \min_{P \in C} \{0\frac{1}{3} + 10P(r) + (\frac{2}{3} - P(r))0\}$.

Thus, for $S = 3$, if the event E occurs, the decision maker chooses the continuation d . But for values $S > \frac{10}{3}$, if he reaches the event E' , he chooses u . As S changes, an outside observer will see the decision maker changing not only his investment behaviour, but also his conditional behaviour at the information sets E' and E . This occurs even though these information sets correspond to decision problems with identical information, sets of feasible action, and payoffs. The only difference is that sunk cost was incurred at E but not at E' .

2.4 Dynamic choice and fact-based updating

In the sunk cost example we assumed dynamic consistency, i.e. the decision maker's ability to commit, by fiat, to carry on his *ex ante* optimal plan. Of course, a major issue in dynamic choice is "to distinguish between an individual's *planned choices* for each decision node at the beginning of the decision problem [...] and his *actual choices* upon arriving at a given decision node" (Machina 1989: 1633).

The example in Figure 3 will illustrate that an ambiguity-sensitive decision maker who updates based on facts must display preference reversals, and thus faces an intra-personal conflict between his *ex ante* and *ex post* selves. Towards that end, we begin by formally defining what we mean by fact-based updating.

Two acts f and g agree on an event E , written $f \equiv_E g$, if they agree on the consequences they assign to states in E . The updated preference \succsim_E is *fact-based* if for all events E and acts f, g, f', g' :

$$[f \equiv_E f' \ \& \ g \equiv_E g'] \implies [f \succsim_E g \iff f' \succsim_E g'].$$

The requirement of fact-based updating has two parts.¹⁴ First, in comparing acts f and g given the event E , the conditional preference \succsim_E places no weight on the consequences of these acts at states that are now excluded by virtue of the knowledge that E occurred.¹⁵ Second, \succsim_E depends only on the information set E , and not on inconsequential aspects of the decision problem such as the *ex ante* optimal plan or feasibility constraints at various stages of the choice problem. This is reflected in the fact that \succsim_E is indexed only by E and nothing else.¹⁶

¹⁴ This condition also excludes those preferences, such as regret-based preferences, that depend on consequences that are no longer possible. If regret is modelled by adding a term into the final utilities, this issue goes away.

¹⁵ This part corresponds to the assumption of "null complements" in the literature (e.g., Hanany and Klibanoff, 2007: 285).

¹⁶ The reader may view the concept of fact-based updating as a way to formalize Machina's (1989) notion of consequentialism. On page 1641 he writes: "the *consequentialist* approach [...] consists of 'snipping' the decision tree at (that is, just before) the current choice node throwing the rest of the tree away, and recalculating by applying the *original* preference ordering [...] to alternative possible *continuations* of the tree".




FIGURE 3. Preference reversal.

Assuming that updating is fact-based does not rule out (or in) ambiguity-sensitive behaviour. To illustrate this point, consider the following extension of our earlier MEU example:

Example 3 (MEU example: Fact-based updating). *The decision maker is told that the event $E = \{b, r\}$ occurred – i.e. the ball is not yellow. Define \succ_E to be the MEU preference corresponding to the set C_E of Bayesian updates of the original set C . That is, $P_E \in C_E$ iff:*

- $P_E(b) \in [\frac{1}{3}, 1]$; and
- $P_E(b) + P_E(r) = 1$.

It is clear that \succ_E is fact-based since it is completely insensitive to the consequence of an act at the now excluded event $E^c = \{y\}$.

2.5 Intra-personal conflicts and preference reversals

In Figure 3, a decision maker with unconditional preference \succ consistent with the Ellsberg choices, initially chooses either L or R . Once this choice is made, he learns whether or not state y occurred. If y did not occur, the decision maker would have a further decision to make at E or E' (depending on whether he had chosen R or L). The Ellsberg choices correspond to choosing (R, d) over (R, u) and (L, u) over (L, d) .

If the updated preferences are fact-based, then \succ_E and $\succ_{E'}$ must coincide. If $u \succ_E d$, then upon choosing R , the decision maker would reverse his *ex ante* choice from d to u . If, on the other hand, $d \succ_E u$, then

upon choosing L , the decision maker would reverse his *ex ante* choice from u to d .

Such reversals have several problematic implications which we list in the next section. To set the stage for that discussion, it is useful to keep in mind that an act f in a dynamic setting should be identified with the set of contingent plans that, once implemented, yield the state-contingent payoffs implied by f . Thus, the act $(10, 0, 10)$ is identified with the contingent plan (R, u) and so on. The decision maker's *ex ante* preference \succsim is his evaluation of contingent plans to which he can commit in advance. Thus, \succsim does not reflect the dynamic nature of the decision problem where the decision maker is given the option to deviate from the *ex ante* plan.

When the decision maker is dynamically consistent he does not deviate from the *ex ante* plan, so the *ex ante* preference \succsim determines his full dynamic choice. But when preference reversals may be an issue, the choice made at the event E is "controlled" by the conditional preference \succsim_E , which may dictate a different course of action. To predict actual choices we must supplement \succsim with assumptions about whether the decision maker takes into account the possibility of future reversals.

The natural language to describe this intra-personal conflict is game theory. For our purposes, we consider a two-stage game with an *ex ante* player self whose preference is \succsim , and for each information set E , an *ex post* self with conditional preference \succsim_E . To avoid cumbersome notation, we avoid explicit listing of the information sets and denote this game by $(\succsim, \{\succsim_E\})$.

3. THE AMBIGUITY AVERSION LITERATURE'S ATTEMPTS TO DEAL WITH INTRA-PERSONAL CONFLICTS

This section describes the various attempts in the ambiguity aversion literature to solving the intra-personal game $(\succsim, \{\succsim_E\})$ between the *ex ante* and *ex post* selves. We illustrate each approach with a simple example then comment on the new set of anomalies that emerge.

3.1 Naiveté and dominated choices

The first approach to resolve the intra-personal conflict is to assume that the *ex ante* self selects a contingent plan according to \succsim without anticipating that the *ex post* selves, who make their choices according to \succsim_E , may reverse the initial plan. This behaviour is called *naive*. Its polar opposite, sophisticated behaviour, is examined in the next subsection.¹⁷

The next example illustrates the well-known fact that naive behaviour may lead to strictly dominated choices, i.e. choices that yield a lower payoff

¹⁷ Examples include prior-by-prior updating, as in Pacheco Pires (2002), and maximum likelihood updating in Gilboa and Schmeidler (1993).




FIGURE 4. Dominated choice.

in each and every state. Consider the choice problem in Figure 4. Note that $(R, d) = f_4 - \epsilon$ and $(R, u) = f_3 - \epsilon$. Assume that the decision-maker's conditional preference is such that $d \succ_E u$. Then $f_3 \succ_E f_4$ and thus, for any small enough $\epsilon > 0$, we also have $f_3 - \epsilon \succ_E f_4 - \epsilon$. On the other hand, Ellsberg preferences imply $f_4 \succ f_3$ and, by continuity, $f_4 - \epsilon \succ f_3$. Collecting all these facts, we conclude that the decision maker implements the plan that leads to payoffs $f_3 - \epsilon$.

A decision maker who ends up with the payoffs corresponding $f_3 - \epsilon$ when f_3 is available should be very embarrassed indeed: he has just selected an act that yields uniformly lower payoffs in every state. Whether or not people choose dominated acts in experiments, games or markets is not the point; rather what seems indisputable is that one should not call such choice rational.

To illustrate these points, we continue our MEU example with prior-by-prior updating. Assume that the set of beliefs at E is the set C_E identified in Example 3; the MEU choice given C_E is, of course, u .

Example 4 (MEU example: naive updating). *The decision maker's ex ante preference \succ has the MEU representation in Example 1 with a set of probabilities C .*

This decision maker faces the problem depicted in Figure 4. The \succ -optimal choice is (R, d) . His prior-by-prior Bayesian update upon learning that E occurred is the set of probabilities C_E calculated in Example 3:

- $P_E(b) \in [\frac{1}{3}, 1]$;
- $P_E(r) \in [0, \frac{2}{3}]$; and
- $P_E(b) + P_E(r) = 1$.

Applying the MEU criterion at the event E and the set of probabilities C_E , we can calculate the conditional distribution at E that minimizes the expected payoff for each choice:

Choice u : put the maximum weight on state r , hence

- $P_E(r) = \frac{2}{3}$;
- expected payoff = $\frac{10-3\epsilon}{3}$.

Choice d : put the maximum weight on state b , hence

- $P_E(b) = 1$;
- expected payoff = $-\epsilon$.

This leads to choosing u and reversing the original plan, resulting in the dominated act $f_3 - \epsilon$.

3.2 Sophistication and information aversion

The behaviour in Section 3.1 is naive in that the decision maker at the *ex ante* stage does not anticipate the subsequent reversal at the event E . The polar opposite of naiveté is *sophistication*, an approach advocated in Siniscalchi (2006).

The idea, roughly, is to solve the game $(\succ, \{\succ_E\})$ by backward induction. That is, the *ex ante* self anticipates that if the event E were to occur, the future course of action will be determined by the *ex post* self with preference \succ_E . In Example 4, the sophisticated *ex ante* self would have chosen L , since it fully anticipates that a choice of (R, d) will be reversed by \succ_E leading to the dominated act $f_3 - \epsilon$.¹⁸

While it escapes the paradox of choosing dominated acts and lends itself to backward induction, sophistication raises a fresh set of problems and paradoxes. The most important of these is *aversion to information*.¹⁹ To illustrate, consider the decision problem in Figure 5. The interpretation of the choices and outcomes following R are familiar from our earlier examples. A choice of L , on the other hand, is simply a commitment not to learn whether or not E occurred. Note that the payoffs at the right-hand branches are identical to those on the left.

Assume that the *ex post* self \succ_E picks u as in Example 4. Then a sophisticated *ex ante* self with Ellsberg preferences would choose (L, d) .

¹⁸ Formally, the sophisticated choice of a contingent plan takes into account the constraint that the subsequent choice must be optimal with respect to \succ_E , while the naive choice ignores this constraint.

¹⁹ Wakker (1988) argued that aversion to information is typical of non-expected utility models.




FIGURE 5. Dynamic consistency and value of information.

That is, he prefers not to look at the information (whether or not E occurred) as a commitment device to help him implement the *ex ante* choice.

Is a strict preference to remain ignorant rational? A statistician who finds himself choosing not to look at all available data should feel rather embarrassed. By the same token, a decision maker willing to pay not to see new evidence should find, upon introspection, such behaviour untenable.

Aversion to information under ambiguity is especially intriguing. A common justification for introducing ambiguity to begin with is to model the lack of reliable information. One would therefore expect information to be at least as valuable under ambiguity as under risk, if not more so. Aversion to information emerges here for the sole reason of providing the *ex ante* self a commitment device in the intra-personal conflict. The problem of aversion to information is one of the problems that prompted Epstein and Le Breton (1993: 3) to write: "From a normative point of view, it is difficult to imagine adopting or recommending a dynamically inconsistent updating rule for use in statistical decision problems."

A possible response to the above arguments is that the desire for commitment under ambiguity is analogous to the desire for commitment in games and under temptation preferences. We believe the analogy flawed because commitment to ignorance here lacks the motivation that justifies commitments in the contexts of games and temptation.

- In the case of games, the desire for commitment is motivated by its potential to influence the behaviour of an opponent. Game theory clarifies how this desire depends on the structure of the game (payoffs,

information, order of moves, and so on). A rational player need not be embarrassed for deciding to make an irreversible commitment.

- In the case of temptation preferences, as in the classic work of Strotz (1956), the source of temptation is psychological urges that have an independent motivation. For example, addiction to cigarettes or alcohol is, presumably, founded in the physiology of the brain and thus represents an objective and independently motivated constraint. One cannot wish or reason these urges away, any more than one can wish away other objective constraints. In our view, an individual who chooses to make commitments in anticipation of his urges has no reason to be embarrassed.

The desire for commitment under ambiguity lacks such motivations. A subjective ambiguity representation captures the decision-maker's model of his environment. While introspection is unlikely to eliminate physiologically induced urges, or force an opponent to change his behaviour, the subjective decision model is an entirely different matter. It is a mental construct the decision maker created to help him coherently think about the uncertainty he faces, interpret information and make decisions. The decision maker can change his model if, upon introspection, he finds it wanting or inadequate.

Thus, once confronted with choices like those in Figure 5, a rational decision maker should feel embarrassed by his choices and respond by changing how he models his environment. Rational decision makers may commit to abstinence by flushing away cigarettes or alcohol in anticipation of their urges. This is of a very different nature than commitments to "flush away" their ability to reason about the uncertainty they face. In sum, the desire for commitment under ambiguity originates in the way the decision maker chooses to incorporate ambiguity, something that, upon introspection, he is free to change. For this reason, we consider the sophistication approach not only an unappealing normative recommendation, but also unlikely to successfully describe behaviour. Sophisticates who are able to plan for all future contingencies are unlikely to persist in ambiguity aversion when perceiving their dynamic inconsistency, especially in light of our argument in Section 5 that ambiguity aversion is a heuristic misapplied by the relatively unsophisticated.

3.3 Distorting the updating rules

A third approach to overcome the updating paradoxes is that proposed by Hanany and Klibanoff (2007, 2008). They propose that decision makers update their beliefs in whatever way necessary to make them adhere to the *ex ante* optimal plan. In the terminology of the intra-personal game $(\succ, \{\succ_E\})$, a solution is found by changing the way beliefs are updated

to ensure that \succ_E lines up with \succ . In the MEU case, the decision maker avoids reversals by simply tossing out the problematic priors that could cause reversals.²⁰

Since the fine details of their theory are orthogonal to our main point, we illustrate their approach in the context of the sunk cost example, Example 2. In that example, the decision-maker's *ex ante* preference dictates that he chooses d at E and u at E' (he chooses to invest in the first stage, but this is not our main concern here). When given the opportunity to revise his decision at E or E' , then prior-by-prior updating applied to the initial set of priors C would dictate that he chooses u at both information sets, upsetting the *ex ante* plan. The solution of Hanany and Klibanoff is, depending on the information set, to retain only those beliefs whose Bayesian updates do not reverse the *ex ante* choice.

Example 5 (MEU example: belief distortion). *In the sunk cost example (Example 2), assume $S = 3$ and update beliefs prior-by-prior to obtain:*

$$C_E = C_{E'} = \left\{ P : P(r) \in \left[0, \frac{2}{3} \right] \text{ and } P(y) = 0 \right\}.$$

*Applying the MEU criterion at E' with respect to the set of priors $C_{E'}$ leads to u , consistently with the *ex ante* preference.*

*However, following the same procedure at E leads to reversing the *ex ante* choice d . This occurs because of priors that put mass less than $\frac{1}{6}$ on r . So to implement dynamic consistency at E , we simply toss out these troublesome priors.*

Specifically, prune the original set C to a smaller set C^ , with $P \in C^*$ iff*

- $P(b) = \frac{1}{3}$;
- $P(r) \in [\frac{1}{6}, \frac{2}{3}]$; and
- $P(r) + P(y) = \frac{2}{3}$.

The set of tossed out priors, $C - C^$, is precisely those priors with $P(r) < \frac{1}{6}$.*

Now that priors causing reversals are deleted, apply prior-by-prior updating to the pruned set C^ to obtain:*

$$C_E^* = \left\{ P : P(r) \in \left[\frac{1}{3}, \frac{2}{3} \right] \text{ and } P(y) = 0 \right\}.$$

²⁰ Hanany and Klibanoff (2007) show that their procedure (in addition to other standard auxiliary axioms) characterizes a weak form of dynamic consistency under MEU. They further show that stronger forms of consistency are incompatible with ambiguity models. In Hanany and Klibanoff (2008), they characterize the distortions necessary to restore dynamic consistency in other classes of preferences.

Applying MEU to C_E^* leads, as expected, to d , consistently with the *ex ante* choice.²¹

To appreciate the extraordinary nature of this behaviour, the decision maker is supposed to contort his beliefs in such a way that at E' he believes the probability of the red ball is between $\frac{1}{3}$ and $\frac{2}{3}$, while at E he believes it is between 0 and $\frac{2}{3}$. And this despite the fact that E and E' correspond to the same event, hence represent identical information about the state space.

Should we expect a *rational* decision maker to behave in this manner? Since dynamic consistency is imposed by fiat, the decision maker will not suffer the embarrassment associated with the reversals discussed earlier. On the other hand, he must accept that his updating is not fact-based: in updating at the event E , he must take into account his payoffs and tastes at states that he now knows are no longer relevant.²²

How would a rational decision maker justify updating some priors but not others in Figure 2? Imagine confronting him with the following analysis:

You initially chose d at information set E because $\{r, y\}$ hedged against the ambiguity about their probability. Now that E has occurred, y and any hedging advantage it may have offered *ex ante* is no longer relevant. This, after all, is what we mean by "learning that E occurred". Why let a state irrelevant to your present situation affect your decision? And if the initial set of priors represented the extent of your uncertainty about the odds, how can you justify selectively tossing out some of these priors at E but not at E' ?

In summary, Hanany and Klibanoff's approach envisions decision makers who evade the updating paradoxes by distorting the way beliefs are updated in response to new information. There is no independent motivation for why a rational decision maker would ever engage in such distortions. Although belief distortion is entirely conceivable as psychological bias, we believe most will find it difficult to swallow as a criterion consistent with rational behaviour.²³

²¹ The decision maker is actually indifferent between d and u . One can slightly perturb C^* to break the indifference. More generally, Hanany and Klibanoff advocate pruning C in whatever way necessary to prevent reversals of the *ex ante* choices.

²² Another issue with the Hanany and Klibanoff updating is its circularity. If g denotes the *ex ante* optimal act, then the updated beliefs at E may depend on g . On the one hand, the purpose of updating beliefs is to determine the optimal act conditional on E . On the other hand, beliefs at E are derived from the optimal act g . Note that while Epstein and Le Breton (1993) allow the updated beliefs at E to also depend on g , no circularity arises in their case. This is because the updated beliefs depend only on what the act g prescribes outside of the event E .

²³ Another issue with the Hanany and Klibanoff updating is that the optimal *ex ante* plan is embedded in their definition of the update rule. Thus, the rule does not provide any guidance to a decision maker beyond the advice to form the optimal plan according to

3.4 Restricting information structures

A fourth approach to eliminate the updating paradoxes is to limit attention to decision trees (information structures) on which no reversals occur.²⁴ Anomalies like those appearing in Example 4 are ruled out by “disallowing” decision trees that cause the decision maker to reverse his *ex ante* plan. In the terminology of the intra-personal game $(\succ, \{\succ_E\})$, a solution is obtained by imposing a joint restriction on \succ and the set of events $\{E\}$ for which the *ex ante* plan can be revised. Information sets E for which \succ_E reverses the *ex ante* choices are not considered.

The fact that there are *examples* where coherent updating is possible does little, in our opinion, to move matters forward. A theory of rational behaviour should not have the freedom to consider some information structures, while ignoring other equally relevant ones just because they are problematic for the theory. Taking the point of view that “economists are responsible for the logical implications of their behavioral models”, a minimal standard a subjective theory should meet is answering: “What should the decision maker do given the information he has?”

To illustrate this approach, we use Epstein and Schneider’s (2003) model of dynamic choice by an MEU decision maker who updates prior-by-prior. Consider a decision maker in the setup of Example 1, so he has MEU preferences with the set of priors C consisting of all probability distributions that put $\frac{1}{3}$ probability on b . The decision maker faces a two-stage dynamic choice problem with information represented by a partition of the state space \mathcal{E} . In the interim stage, an event $E \in \mathcal{E}$ is revealed to the decision maker who gets an opportunity to change his *ex ante* plan.

Epstein and Schneider give a condition characterizing absence of reversals, *rectangularity*: The set of priors C is rectangular with respect to the partition \mathcal{E} if for all $P, Q \in C$, their “composition” R , defined by:

$$R(\omega) = P(E)Q(\omega | E), \quad \omega \in E \in \mathcal{E}$$

is also in C . That may be interpreted to say that the set of priors C has a recursive structure. See Epstein and Schneider (2003) for motivation and details.

In the structure $\mathcal{E} = \{\{b\}, \{r\}, \{y\}\}$, the colour of the ball is revealed. In this case, updating is trivial. So we focus on the case where, at an interim

his preferences and stick to it. The updating rule does little beyond rationalizing the *ex ante* choice, so the decision maker gains no advantage from using that rule. A Bayesian decision maker, by contrast, computes his updated beliefs and optimal action only for the contingency that actually occurs. For a decision tree with many branches, this may be exponentially simpler than finding the optimal *ex ante* plan.

²⁴ This approach is pursued by, among others, Sarin and Wakker (1998), Epstein and Schneider (2003), Maccheroni, Marinacci and Rustichini (2006b), and Klibanoff, Marinacci and Mukerji (2006).

stage, the decision-maker's information consists of a partition $\{E, E^c\}$, with E^c consisting of a single state. There are three possible such structures, depending on whether E^c is b , r or y . We check whether rectangularity holds in each case:

1. $E = \{r, y\}$: For any $P, Q \in C$ and $x \in \{r, y\}$, since $P(E) = Q(E) = \frac{2}{3}$, we have

$$Q(x | E)P(E) = Q(x | E)Q(E) = Q(x),$$

and rectangularity holds.

2. $E = \{b, r\}$: Take $P = (\frac{1}{3}, 0, \frac{2}{3})$ and $Q = (\frac{1}{3}, \frac{2}{3}, 0)$; then

$$Q(b | E)P(E) = \frac{1}{3} \cdot \frac{1}{3} = \frac{1}{9},$$

but there is no prior in C that assigns probability $\frac{1}{9}$ to b , so C is not rectangular with respect to the information structure $\{E, E^c\}$.

3. $E = \{b, y\}$: Rectangularity fails for similar reasons as above.

In summary, in our simple example, rectangularity holds under the structures $\{\{b\}, \{r\}, \{y\}\}$ and $\{\{b\}, \{r, y\}\}$. In the first, ambiguity is completely resolved, while in the second the initial ambiguity remains. Rectangularity is violated when there is "partial" resolution of ambiguity, as in cases 2 and 3 above.

What would constitute a rational choice at an information set like $E = \{b, r\}$? If the theory continues to prescribe prior-by-prior updating for all information sets, this will result in the dynamic inconsistency discussed earlier. The theory eliminates this problem by ignoring situations like these, even though there is nothing peculiar or unusual about them. As Machina (1989) put it: "economists are responsible for the logical implications of their behavioral models when embedded into social settings". A sensible theory of updating should not selectively limit its scope to those situations where its desired implications seem to hold, while remain silent about what happens if slight perturbations to the information structure are introduced.

4. BELIEFS IN SUBJECTIVE MODELS

Ambiguity models often propose probability-like objects as a way to generalize the concept of beliefs. We shall argue that calling these objects "beliefs" stretches the meaning of this concept so much so that it has little to do with what economists and game theorists understand and use in their models. We will also explain why it is so difficult for the ambiguity aversion literature to produce an adequate notion of beliefs.

Beliefs in the Savage model: We begin by reminding the reader that tastes and beliefs are treated as conceptually distinct aspects of the preferences, both in Savage's writings, and in their subsequent interpretations. Thus, Aumann (1987) writes:

That Bayesian decision theory à la Savage derives both utilities and probabilities from preferences does not imply that it does not discriminate conceptually between these two concepts.

[U]tilities directly express tastes, which are inherently personal. It would be silly to talk about "impersonal tastes," tastes that are "objective" or "unbiased". But it is not at all silly to talk about unbiased probability estimates, and even to strive to achieve them. On the contrary, people are often criticized for wishful thinking – for letting their preferences color their judgement. One cannot sensibly ask for expert advice on what one's tastes should be; but one may well ask for expert advice on probabilities.

In all models of choice considered in this paper, the decision-maker's preference \succsim can be represented by a functional $I(u(f(\cdot)))$ for some utility function u on consequences. Savage's classic theorem represents the functional I as a weighted sum of the utilities, but other representations, under different axioms, are of course possible. In MEU, for instance, I is a minimum over integrals.

So what justifies referring to the weights in Savage's theory as "beliefs"? Why aren't they subject to the same arbitrariness that is characteristic of tastes? And what should one minimally expect of alternative conceptions that aspire to be useful generalizations of our standard notion of beliefs?

In the general formulation $I(u(f(\cdot)))$ one may call the functional I beliefs, if one wishes, and declare that u and I achieve a separation of tastes from beliefs. This, however, is a separation only in a trivial, purely mathematical sense. In our view, whether a mathematical object like the functional I embodies a meaningful notion of beliefs is inseparably tied to whether there is a coherent theory describing how I changes to incorporate new information.

The role of Bayesian updating: Most would agree that Bayesian updating is not merely an interesting adjunct to the Savage model, but a central part of its interpretation. In evaluating purported generalizations of Savage's conception of beliefs, one has to spell out which aspects of Bayesian updating can be relaxed or dispensed with, and which are fundamental and "not optional".

To organize the discussion, the following sketch of the role of beliefs and updating in Savage's theory will be helpful:

- *Separation of taste from beliefs:* A preference \succsim is represented in terms of a utility function u and a probability P , reflecting the decision-maker's beliefs.

- *Dynamic choice*: In anticipation of being told that some event E has occurred, the decision maker can proceed in two conceptually distinct ways:
 1. *Dynamic consistency*: Define a conditional preference \succsim_E that is dynamically consistent with the *ex ante* preference \succsim ; or
 2. *Updating*: Wait until he discovers that the event E has occurred; treat his subjective belief P as a classical probability; update P based on facts using Bayes rule; and use the updated beliefs to derive a conditional preference \succsim_E^* .
- *Characterization*: In the Savage setting, \succsim_E and \succsim_E^* coincide. That is, Savage's theory delivers two equivalent characterizations of conditional choices. The decision maker can either: (1) use his prior to determine the complete optimal contingent plan *ex ante* and stick to it; or (2) Wait until information arrives, use it to update the prior beliefs, and determine the E -contingent choices based on the updated beliefs.

The practical value of this characterization is hard to over estimate; it is what underlies the success of the Bayesian methodology in games, dynamic optimization and the economics of information. In the discussion that follows our concern is primarily with the conceptual issues, however.

A first aspect of Savage's theory that seems essential to preserve is that there be a *substantive* separation of tastes from beliefs. As pointed out earlier, beliefs are not merely what is left after extracting the taste component of the preference. To call a functional I "beliefs" we must require that its transformation in response to new information is not contaminated by the decision-maker's taste. Updating must be based on facts, on what actually occurs rather than what the decision maker would like to occur. While this is doubtless violated by real people every day (as in, e.g. Aumann's reference to wishful thinking), it is clear that it is a useful normative standard for rational behaviour. Formally, this is our requirement that updating is fact-based and it amounts to saying that separation applies not just to prior beliefs, but to updating as well.

The second crucial aspect of Savage's theory is that prior beliefs and their updates must, in some sense, be inter-temporally coherent. Intuitively, this says that the decision maker has a coherent theory of his environment: the choices made by applying the theory *ex ante* in anticipation of future contingencies do not contradict the choices made based on the implications derived from the theory via updating.

Beliefs in ambiguity models: The ambiguity aversion literature puts a premium on representations that can express the functional I in an aesthetically appealing and tractable functional form. While not disputing their value, aesthetics and tractability do not justify interpreting a mathematical object as "beliefs".

Can the Bayesian methodology be sensibly extended to ambiguity models? As noted earlier, separating beliefs from taste and fact-based updating are easily accomplished within ambiguity aversion models.²⁵ Dynamic consistency for preferences can also be defined in a way analogous to the Bayesian model.²⁶

The essence of the reversals discussed in Section 2.5 is that it is not possible to develop a coherent notion of beliefs and fact-based updating where choices based on updated beliefs are dynamically consistent. As argued earlier, this difficulty is inherent in the approach. In the case of non-additive probabilities, Gilboa and Schmeidler (1994) and Mukerji (1997) show that the state space on which the ambiguity model is defined is, in fact, a reduced form of an underlying state space on which beliefs are additive. Ambiguity arises because of "missing states" – states that are relevant to the decision maker but overlooked by the reduced form. These authors argue that the incomplete specification of the state space raises significant problems for updating.

Classifying the various approaches to updating in the ambiguity aversion literature: In light of our earlier discussion, the various efforts to define belief updating can be classified as follows:

1. Naive updating and sophistication dispense with dynamic consistency entirely. Reversals occur under both approaches; the difference is in how the decision maker is supposed to deal with them. Under naive updating, the decision maker is fooled into thinking that he will not reverse, only to be ultimately proven wrong. Sophistication presents the opposite case: the decision maker anticipates that he will change his mind when he sees the new information, so he strictly prefers not to see it.
2. The approach of restricting information eliminates troublesome information sets where the decision maker would have reversed his *ex ante* choice had he been given the opportunity to do so. The only motivation for this procedure seems to be the need to prevent choices that would be problematic for the theory.
3. Hanany and Klibanoff (2007) advocate updating rules under which the decision-maker's response to information is not fact-based. Under this approach, the decision maker changes his interpretation of the

²⁵ Indeed most updating rules in that literature are fact-based. An example is the prior-by-prior updating; see, for instance, the updating in Example 3, Epstein and Schneider (2003), Pacheco Pires (2002). The exceptions are the Hanany and Klibanoff models, and Hansen and Sargent (2001b). For comments on the latter approach, see Epstein and Schneider (2003, Section 5).

²⁶ For example, Epstein and Le Breton (1993) and Epstein and Schneider (2003) both provide such definitions. Hanany and Klibanoff (2007) introduce a weaker notion of dynamic consistency.

information in whatever way necessary to ensure that he does not reverse.

Is this just semantics? The reader may object that we are too hung up on semantics, e.g. what the words “beliefs” or “rationality” mean. These are just definitions, so who cares? The point is that words are powerful. Applied theorists should know that what is meant by beliefs and rationality in the ambiguity aversion literature is very different from their customary use. It is not a natural generalization of the standard concept.

A standard argument in the ambiguity aversion literature is that decision makers hold multiple priors (non-additive beliefs, or other forms of “beliefs”) because they do not know the true probabilities. These decision makers hedge against ambiguity about the true probability by being ambiguity-averse. This begs the question: what does it mean to use the “wrong” prior in a subjective setting?

In its strictest interpretation, the subjectivist view claims, as de Finetti (1974) so memorably put it, that “probabilities do not exist”. Under this strict subjectivist interpretation, there is no objective distribution which the subjective belief can match or fail to match. Being ambiguity averse would then amount to being cautious about things that “do not exist”.

A more nuanced view of probability is given by Borel:

Observe however that there are cases where it is legitimate to speak of the probability of an event: these are the cases where one refers to the probability which is common to the judgements of all the best informed persons, that is to say, the persons possessing all the information that it is humanly possible to possess at the time of the judgements [...] This surely captures exactly our intuition of what we mean by the true probability of an event.²⁷

Borel’s interpretation of true probabilities as relative is consistent with both the negative and positive results of the recent testing literature. Strong impossibility results prevent us from testing whether a single expert knows the “true” probabilities (see Sandroni 2003). But, as we showed in Al-Najjar and Weinstein (2008), we can test whether one expert knows more than another.

Under this interpretation, the usual motivation for the ambiguity literature – that one should be cautious if one does not know the “true” probabilities – makes sense just when there is another player who is better informed. We certainly should be cautious if uncertainty about the true probabilities really means that others may be better informed, and act contrary to our interest – but such caution should not be modelled as an issue of arbitrary taste in a one-person decision model. Rather, a more natural tool is game theory, which we turn to next.

²⁷ Quoted in Morris (1997).

5. EXPERIMENTAL EVIDENCE, GAMES, AND THE ELLSBERG ANOMALY

The core empirical justification of ambiguity models is their ability to account for Ellsberg choices in experiments. We argue here that these same experimental findings are equally consistent with other explanations, and thus lend no special support to the ambiguity models.

The specific alternative we propose is that subjects incorrectly extend heuristics that serve them well in real-world situations to experimental settings where these heuristics are inappropriate. As the literature cited below makes clear, many of the points we make can be traced, in one form or another, to points already made in the literature.²⁸ Our (modest) claim to novelty lies in providing a synthesis that is not only consistent with the Ellsberg choices, but has more predictive power, clarifies the updating paradoxes, and does so without having to tamper with foundational assumptions.

Here are the key steps of our argument:

- *Games*: In many real-world situations, individuals offered to bet on risky prospects would be wise to assume that the odds are *adversarially* manipulable. Whether it be betting on a horse in a horse race, buying a used car, or choosing a political strategy, we almost always find ourselves playing against opponents with the ability to change the odds. This, after all, is why we study game theory. Myerson (1991: 26) noted that commonly encountered real-world situations may contaminate subjects' behaviour in experiments. Calling these common situations "salient perturbations", he notes that "people usually offer to make bets only when they have some special information or beliefs. We can try to offer bets uninformatively, [...] but this is so unnatural that subjects may instead respond to the salient perturbation."

- *Heuristics*: It is eminently sensible for such individuals to adopt heuristics according to which they hedge against risks that can be manipulated. The behaviour implied by such a heuristic is consistent with Ellsberg choices.

In fact, one of the key lessons of the ambiguity aversion literature is that ambiguity-sensitive behaviour is observationally indistinguishable from the behaviour of a player in a game. This is already apparent in the classic MEU model. More generally, Cerreia *et al.* (2008) show that *all* ambiguity-averse preferences have (under mild auxiliary axioms)

²⁸ For example, Morris (1997) argued that Ellsberg choices can be explained without abandoning classical decision models. The difference is that he does not appeal to the role of heuristics as we do here.

a representation under which the decision maker behaves as if he thought he was in a game against adversarially determined odds.

- *Misapplying heuristics*: Subjects' behaviour in laboratory experiments can be affected by the heuristics they use to deal with the real-world situations they spend most of their time in. Indeed, this is the central theme of literature on heuristics and biases pioneered by Tversky and Kahneman (1974). In that paper they write: "people rely on a limited number of heuristic principles which reduce the complex tasks of assessing probabilities and predicting values to simpler judgmental operations. In general, these heuristics are quite useful, but sometimes they lead to severe and systematic errors."

The idea that agents constrained by a limited set of models may commit systematic errors in experiments (and decision making in general) is made formal by Samuelson (2001). He studies "decision makers characterized by a stock of models, or analogies, who respond to strategic interactions by applying what appear to be the most suitable models; balancing the gains from more sophisticated decision making against the cost of placing heavier demands on scarce reasoning resources". For such decision makers, "[i]nteractions that are infrequently encountered, relatively unimportant, or similar to other interactions may trigger seemingly inappropriate analogies, leading to behavioral anomalies". Samuelson discusses in detail how inappropriately triggered analogies can account for framing effects and other "anomalies" in experiments. Laboratory settings seem like good candidates of the infrequently encountered interactions he refers to.²⁹

We conclude that the basic experimental findings supporting Ellsberg choices cannot distinguish between two competing explanations:

1. Ambiguity models which explain these choices by appealing to taste (ambiguity aversion).
2. A model where subjects misapply heuristics that serve them well in real-world situations.³⁰

²⁹ A similar phenomenon appears in connection with the Allais paradox. List and Haigh (2005) write: "We find that both students and professionals exhibit some behavior consistent with the Allais paradox, but the data pattern does suggest that the trader population falls prey to the Allais paradox less frequently than the student population." They add that: "Indeed, according to some researchers, learning and familiarization with the decision tasks are required before true preferences settle on the genuine underlying form."

³⁰ Halevy and Feltkamp (2005) suggest another explanation of the misapplied heuristic variety: If more than one ball (a bundle) is to be drawn from an urn, even a subject with a Bayesian prior over the composition of an unknown urn will exhibit Ellsberg behaviour. In this case a lottery over the composition of the urn translates into additional risk, as compared with the urn having the average composition with certainty.

On the other hand, by linking Ellsberg choices to games and heuristics, the misapplied heuristic model offers a number of advantages not enjoyed by ambiguity models:

- First, this alternative explanation can account for the experimental findings without revising foundational assumptions.
- Second, the misapplied heuristic model provides a straightforward resolution for the updating paradoxes. Ozdenoren and Peck (2008) exhibit a number of alternative games against nature that a subject might perceive in an Ellsberg situation. The games vary according to the timing of moves by the player and the malevolent nature. They show that various updating rules emerge as backwards-induction outcomes depending on the game being played. For example, a sophisticated response to preference reversals arises if the subject considers it possible that the urn will be manipulated by nature at more than one turn.
- Third, the misapplied heuristic model can account for more recent experimental findings that would confound the explain-anomalies-by-taste approach. Misapplied heuristics readily explain, for instance, the recent findings of Halevy (2007). Most subjects in his experiments (80%) exhibit ambiguity aversion, in accordance with Ellsberg's thought experiment. But Halevy also find that most subjects (84%) also fail to reduce *objective* compound lotteries, in violation of standard decision theoretic models. Halevy finds that subjects' failure to multiply probabilities in order to reduce compound lotteries is highly correlated with whether they express ambiguity aversion. Of those subjects who understood basic probability enough to reduce objective compound lotteries, 96% were indifferent to ambiguity. On the other hand, 95% of those subjects who could not multiply objective probabilities expressed ambiguity aversion.

The problem is not unique to ambiguity. In a related context, List and Haigh (2005) find that "professional traders behave in accordance with the reduction principle (reducing compound lotteries to simple ones via the calculus of probabilities [. . .]), whereas students did not exhibit this tendency".

The misapplied heuristic explanation accounts for Halevy's (2007) findings simply and naturally in terms of subjects' heterogeneous abilities in judging whether a given heuristic is applicable in a given situation; 'smart' subjects are more discriminating in the analogies they draw.

Deviations from normative decision theory are to be expected and may well deserve behavioural explanations. One should not rush, however, to including such deviations as part of rational theory, as

this would be like “asking people’s opinion of $2 + 2$, obtaining an average of 4.31 and announcing this to be the sum. It would be better to teach them arithmetic.” (D. Lindley’s preface to de Finetti’s (1974) textbook).

An explanation based on game theory and heuristics can be of practical value in situations where ambiguity models are content to “explain” choices in terms of subjective sets of priors and distaste for ambiguity. In a non-technical paper on ambiguity-laden investment opportunities, Zeckhauser (2006) suggests that the typical investor is not cautious enough when facing gambles with severe adverse selection, but overly cautious when facing ambiguous situations where the other side is also ignorant. He thus offers the recommendation: “In a situation where probabilities may be hard for either side to assess, it may be sufficient to assess your knowledge relative to the party on the other side (perhaps the market).” The ambiguity aversion literature completely misses the distinction he makes, while game theory makes it clear this distinction is vital.³¹

To sum up, the misapplied-heuristic explanation provides a more nuanced understanding of ambiguity-averse choices than an approach that attribute them to inexplicable taste parameters. When manipulating the odds is impossible, ambiguity aversion is a mistake that does not warrant revision of the foundational decision-making paradigm. On the other hand, being averse to gambles when one does not know the probabilities involved is perfectly rational when the agent offering the gamble may have superior knowledge of the probabilities, or the ability to influence them. Few gambles are offered in a vacuum, and the fact that they are offered tends to be bad news!

If ambiguity aversion-like choices are, in fact, a reflection not of a decision problem, but of a game, then attributing this fundamentally strategic phenomenon to taste distorts our modelling efforts in at least two ways. First, by taking the amount of ambiguity aversion as an inexplicable primitive, this approach makes no effort to scrutinize the extent of manipulation, which should arise endogenously from the strategies and incentives of the other player(s). Second, when we are not in a game (as in Ellsberg experiments), it sanctifies a simple, understandable error by

³¹ Zeckhauser offers anecdotal evidence that the anomaly of over-cautiousness by most investors in ambiguous situations can lead to lucrative investment opportunities. In particular, Warren Buffett says he has frequently profited from such opportunities. Of course, Buffett may be in a better position than most to know when the odds are manipulable. This reinforces the importance of our strategic-based explanation of ambiguity aversion rather than one based on arbitrary preferences.

encoding it as part of the rational choice paradigm. By analogy, when students of microeconomics struggle with the concept of ignoring sunk costs, the role of theory is to clarify thought and help understand the sources of the fallacy. We do not develop a purportedly rational theory of sunk-cost-sensitive decision-makers.

6. CONCLUDING REMARKS: INTERPRETING AMBIGUITY MODELS

The wide appeal of the ambiguity aversion literature is understandable. The standard Bayesian paradigm stipulates decision makers who model uncertainty with a unique subjective prior. Since this paradigm offers no guidance as to how priors are formed, one is tempted to interpret experimental anomalies as an expression of agents' being unsure about the right prior. This concern about robustness, or model uncertainty, is then taken as justification to relax the seemingly unrealistic demands of a unique prior. For an exposition of this point of view, see Hansen and Sargent (2001a).^{32,33}

It is less clear whether the ambiguity aversion models are intended as rational or descriptive models. We are not wedded to either the rational or descriptive modelling approaches. A dogmatic commitment to one approach over the other is unhelpful since both provide indispensable tools and insights. What is *not* helpful is to be confused about which category a particular model falls into. Models based on rational behaviour represent a commitment to the full set of logical consequences of the rationality assumption. These models are quite different from, and judged by different criteria as, their descriptive counterparts that put a premium on the realism of assumptions and fitting empirical findings.

So should the ambiguity models be viewed as models of rational decision makers or descriptive accounts of behavioural biases and bounded rationality? It is easy to find references in that literature suggesting one interpretation or the other (or both). We find both interpretations questionable. First, as argued in this paper, once scrutinized based on their dynamic implications, ambiguity models lead to choices that most economists would view as irrational, and even absurd. This, in

³² Another example is Maccheroni, Marinacci and Rustichini (2006a) who write: "Under this hypothesis all agents share the same probability distribution on some relevant economic phenomenon and each agent has to be firmly convinced that the model he has adopted is the correct one. This is a strong requirement as agents can have different models, each of them being only an approximation of the underlying true model, and they may be aware of the possibility that their model is misspecified. A weakening of this requirement allows agents to entertain different priors on the economy."

³³ For a critique of the use of ambiguity aversion-based ideas in macroeconomics, see Sims (2001).

our view, undermines the rationality interpretation of these models even in settings that do not involve dynamic choice, because rational models cannot selectively pick which logical implication of rationality to retain, and which to discard.

We also find the descriptive case for ambiguity models suspect. As discussed in Section 5, the anomalous experimental findings on which such a descriptive case is founded can be explained simply within standard theory. At a minimum, the experimental findings do not favour the explanations of the ambiguity models over others, such as those based on the misapplied heuristic idea discussed earlier. In fact, we have made the case that alternative models offer better insights since they can account more convincingly for updating paradoxes and recent experimental findings that would baffle the ambiguity interpretation.

A second problem with taking the descriptive interpretation of ambiguity aversion models seriously is the haphazard manner in which rationality is introduced. If these were truly descriptive models of a behavioural bias, then one would have to justify requiring agents who do not reduce objective compound lotteries (Halevy 2007) to behave in a dynamically consistent manner, to follow rationally motivated updating rules, and to carry complicated optimal portfolio calculations. In rational models, these are justified as consequences of the underlying rationality assumption. In descriptive models, by contrast, these implications of rationality are replaced by behavioural assumptions whose motivations are found in compelling stylized facts, empirical findings and experimental data. In the descriptive rendering of the ambiguity aversion idea, by contrast, we find no convincing reasons for imposing the full burden of rationality on otherwise behaviourally biased agents.³⁴

A third difficulty with a descriptive interpretation of the ambiguity aversion models is that, although they rationalize behaviour that is anomalous for standard theory, they do so only by substantially increasing the degrees of freedom available to the modeller.³⁵ Fitting empirical findings cannot be all we care about. For if it were, then we should embrace physical theories that dispense with conservation laws, or evolutionary theories that appeal to an intelligent designer. These permissive theories necessarily “explain” more observed phenomena by

³⁴ This methodological point was made, in the context of behavioural economics, by Fudenberg (2006): “Specifically, after modifying one or two of the standard assumptions, the modeler should consider whether the other assumptions are likely to be at least approximately correct in the situations the model is intended to describe, or whether the initial modifications suggest that other assumptions should be modified as well.”

³⁵ In the case of MEU, for example, the subjectivity of the single prior is replaced by a subjective set of priors. Once dynamics is introduced, even more degrees of freedom are added via the specification of the updating rules.

appealing to inexplicable free parameters about which the theory has little to say.

As an example, consider the well-known equity premium puzzle. Over the years there has been no shortage of explanations of this puzzle based on behavioural, institutional or other factors. Ambiguity models attempt to explain it in terms of investors' aversion to the ambiguity of stock returns. Why is this superior to other behavioural or ad hoc explanations that fit the data equally well? One answer is that the ambiguity approach is superior because it has firm decision-theoretic foundations. This presumably means not just the existence of *some* axioms that characterize behaviour, but that the underlying behaviour is compelling in a way not captured by the behavioural or ad hoc models. The goal of this paper is to cast doubt on such claims.

REFERENCES

- Al-Najjar, N. I. 2009. Decision makers as statisticians: diversity, ambiguity and learning. *Econometrica*, forthcoming.
- Al-Najjar, N. I. and J. Weinstein. 2008. Comparative testing of experts. *Econometrica* 76(3): 541–59.
- Aumann, R. J. 1987. Correlated equilibrium as an expression of Bayesian rationality. *Econometrica* 55(1): 1–18.
- Bewley, T. 1986. Knightian decision theory: Part I. Cowles Foundation Discussion Paper no. 807.
- Bewley, T. 2002. Knightian decision theory. Part I. *Decisions in Economics and Finance* 25(2): 79–110.
- Cerreia, S., F. Maccheroni, M. Marinacci and L. Montrucchio. 2008. Uncertainty averse preferences. Working paper; Collegio Carlo Alberto.
- Charness, G., E. Karni and D. Levin. 2008. On the conjunction fallacy in probability judgment: new experimental evidence. Johns Hopkins University.
- de Finetti, B. 1974. *Theory of Probability, Vol. 1–2*. New York: Wiley.
- Ellsberg, D. 1961. Risk, ambiguity, and the Savage axioms. *Quarterly Journal of Economics* 75(4): 643–69.
- Epstein, L. G. and M. Le Breton. 1993. Dynamically consistent beliefs must be Bayesian. *Journal of Economic Theory* 61(1): 1–22.
- Epstein, L. and M. Schneider. 2003. Recursive multiple-priors. *Journal of Economic Theory* 113, 1–31.
- Fudenberg, D. 2006. Advancing beyond advances in behavioral economics. *Journal of Economic Literature* 44(3): 694–711.
- Gilboa, I. and D. Schmeidler. 1989. Maxmin expected utility with nonunique prior. *Journal of Mathematics and Economics* 18(2): 141–53.
- Gilboa, I. and D. Schmeidler. 1993. Updating ambiguous beliefs. *Journal of Economic Theory* 59, 33–49.
- Gilboa, I. and D. Schmeidler. 1994. Additive representations of non-additive measures and the choquet integral. *Annals of Operations Research (Historical Archive)* 52(1): 43–65.
- Gilboa, I. and D. Schmeidler. 2001. *A Theory of Case-Based Decisions*. New York: Cambridge University Press.
- Gottlieb, D. 2009. *Imperfect Memory and Choice under Risk*. Boston, MA: MIT Press.

- Halevy, Y. 2007. Ellsberg revisited: an experimental study. *Econometrica* 75(2): 503–36.
- Halevy, Y. and V. Feltkamp. 2005. A Bayesian approach to uncertainty aversion. *Review of Economic Studies* 72(2): 449–66.
- Hanany, E. and P. Klibanoff. 2007. Updating preferences with multiple priors. *Theoretical Economics* 2, 261–98.
- Hanany, E. and P. Klibanoff. 2008. *Updating Ambiguity Averse Preferences*. Northwestern University.
- Hansen, L. and T. Sargent 2001a. Acknowledging misspecification in macroeconomic theory. *Review of Economic Dynamics* 4(3): 519–35.
- Hansen, L. and T. Sargent 2001b. Robust control and model uncertainty. *American Economic Review* 91(2): 60–6.
- Klibanoff, P., M. Marinacci, and S. Mukerji. 2005. A smooth model of decision making under ambiguity. *Econometrica* 73(6): 1849–92.
- Klibanoff, P., M. Marinacci, and S. Mukerji. 2006. Recursive smooth ambiguity preferences. *Carlo Alberto Notebooks*.
- List, J., and M. Haigh. 2005. A simple test of expected utility theory using professional traders. *Proceedings of the National Academy of Sciences* 102(3): 945–48.
- Maccheroni, F., M. Marinacci, and A. Rustichini. 2006a. Ambiguity aversion, malevolent nature, and the variational representation of preferences. *Econometrica* 74, 1447–98.
- Maccheroni, F., M. Marinacci, and A. Rustichini. 2006b. Dynamic variational preferences. *Journal of Economic Theory* 128, 4–44.
- Machina, M. 1989. Dynamic consistency and non-expected utility models of choice under uncertainty. *Journal of Economic Literature* 27(4): 1622–68.
- Machina, M. J., and D. Schmeidler. 1992. A more robust definition of subjective probability. *Econometrica* 60(4): 745–80.
- Morris, S. 1997. Risk, uncertainty and hidden information. *Theory and Decision* 42(3): 235–69.
- Mukerji, S. 1997. Understanding the nonadditive probability decision model. *Economic Theory* 9(1): 23–46.
- Myerson, R. 1991. *Game Theory: Analysis of Conflict*. Harvard: Harvard University Press.
- Ozdenoren, E. and J. Peck. 2008. Ambiguity aversion, games against nature, and dynamic consistency. *Games and Economic Behavior* 62(1): 106–15.
- Pacheco Pires, C. 2002. A rule for updating ambiguous beliefs. *Theory and Decision* 53(2): 137–52.
- Samuelson, L. 2001. Analogies, adaptation, and anomalies. *Journal of Economic Theory* 97(2): 320–66.
- Sandroni, A. 2003. The reproducible properties of correct forecasts. *International Journal of Game Theory* 32(1): 151–59.
- Sarin, R., and P. Wakker. 1998. Dynamic choice and unexpected utility. *Journal of Risk and Uncertainty* 17(2): 87–120.
- Savage, L. J. 1951. The theory of statistical decision. *Journal of the American Statistical Association* 46(253): 55–67.
- Savage, L. J. 1954. *The Foundations of Statistics*. New York: John Wiley & Sons Inc.
- Schmeidler, D. 1989. Subjective probability and expected utility without additivity. *Econometrica* 57(3): 571–87.
- Sims, C. 2001. Pitfalls of a minimax approach to model uncertainty. *American Economic Review, Papers and Proceedings* 91(2): 51–4.
- Siniscalchi, M. 2006. *Dynamic Choice under Ambiguity*. Northwestern University.
- Strotz, R. 1956. Myopia and inconsistency in dynamic utility maximization. *Review of Economic Studies* 23(3): 165–80.

- Tversky, A., and D. Kahneman. 1974. Judgment under uncertainty: heuristics and biases. *Science* 185(4157): 1124–31.
- Tversky, A., and D. Kahneman. 1983. Extensional versus intuitive reasoning: the conjunction fallacy in probability judgment. *Psychological Review* 90(4): 293–315.
- Wakker, P. 1988. Nonexpected utility as aversion of information. *Journal of Behavioral Decision Making* 1, 169–75.
- Zeckhauser, R. 2006. Investing in the unknown and unknowable. *Capitalism and Society* 1, 1–41.