Introduction: Legal Decision Making under Deep Uncertainty

David Weisbach

How do we decide when we lack the information needed to have well-formed judgments, when we face what is alternatively called Knightian uncertainty, deep uncertainty, or ambiguity? Many decisions must be made under these conditions. Keynes (1937, p. 214) in a famous passage gave a number of examples of deep uncertainty: "The sense in which I am using the term [uncertainty] is that in which the prospect of a European war is uncertain, or the price of copper and the rate of interest twenty years hence, or the obsolescence of a new invention, or the position of private wealth owners in the social system in 1970. About these matters there is no scientific basis on which to form any calculable probability whatever. We simply do not know." In all of these cases, society had to make choices in 1937 that were based on the likelihood of these events.

Writing 65 years later, Sunstein (2002) posited as a prototypical case the problem of regulating arsenic. Below a certain threshold, we do not know the dose-response curve for arsenic. Because the sample size needed to get reasonable estimates is too large, we are not going to know. Yet the Environmental Protection Agency (EPA) must decide on a permissible level of arsenic in the water. If arsenic turns out to be toxic at low levels, the savings from restrictive regulations would be large and the regulations would pass a cost-benefit test. If arsenic turns out to be mostly harmless or even beneficial at low levels, would restrictive regulations fail a cost-benefit test? How should the EPA decide what to do?

Another central case is the problem of climate change. As described in Greenstone, Kopits, and Wolverton (2013), in 2010 a group of agencies in the federal government convened to determine a unified value for the

DAVID WEISBACH is the Walter J. Blum Professor of Law and Economics at the University of Chicago Law School and a Senior Fellow at the Computation Institute of the University of Chicago and Argonne National Laboratory.

[*Journal of Legal Studies,* vol. 44 (June 2015)]

marginal harm from emissions of carbon dioxide, a value known as the social cost of carbon, or SCC. Estimating the SCC requires estimating the harm from emissions under a baseline scenario and comparing that to the harm from emissions under that scenario with 1 more ton of carbon emitted into the atmosphere. The SCC is the present-value difference in these two estimates. Carbon dioxide stays in the atmosphere for tens of thousands of years, and the harms are likely to be highly nonlinear. Estimating the harms from baseline emissions requires an accurate estimate of what the baseline emissions will be, which means understanding policy choices concerning emissions and energy choices for hundreds of years. These estimates must then be fed into a model of the climate to produce an estimate of the likely temperature change and other physical effects. These models have a high and seemingly irreducible level of uncertainty. Estimates of temperature changes must then be translated into social harms. Humans have never existed at the temperature levels we may face in the future, which means that we have little or no information about what the harms will be and how humans will adapt. All of these are then combined into a single model that is supposed to spit out a number to be used for cost-benefit analysis. The number is little more than a guess, yet agencies need to use some sort of number for decisions that affect emissions. How are we to come up with a value for the SCC?

Adrian Vermeule, in his article in this issue, gives a number of additional examples from decided cases. Among these are listing under the Endangered Species Act and nuclear waste storage. In particular, the Fish and Wildlife Service (FWS) needs to estimate the number of flat-tailed horned lizards to determine if the lizards are threatened under the Endangered Species Act. The method for estimating the population size, known as the scat count, has been discredited, and no other method currently exists. How is the FWS to determine whether the species is endangered? (See *Tucson Herpetological Society v. Salazar*, 566 F.3d 870 [9th Cir. 2009].)

The Nuclear Regulatory Commission (NRC) is charged with ensuring that spent nuclear fuel does not pose a threat to health or the environment. The half-lives of some of the materials are hundreds of thousands of years, far beyond the ability of the NRC to model. How should the NRC decide? (See *Baltimore Gas & Electric Co. v. Natural Resources Defense Council, Inc.*, 462 U.S. 87 [1983].)

Other examples abound. The National Nuclear Security Administration (NNSA) must certify the integrity of the nation's nuclear weapons.

One of the tasks is to determine the bomb blast radii of each weapon. The United States has not tested a nuclear weapon since 1992, and under the 1997 Comprehensive Test Ban Treaty, we will not be testing them in the near future. Aging weapons degrade: for example, the explosives that condense the fissile material change over time and the radioactive cores can become unreliable. The only way to determine the reliability of the stockpile is simulations. Even the most sophisticated uncertainties, however, have a large degree of uncertainty. How should the NNSA proceed if it must certify the integrity of the stockpile?

Lars Hansen, the keynote speaker at the conference convened on this topic,[1] discussed macroeconomic policy. We have a model of the economy and a large amount of data that can be used to test the model. The data, however, are insufficient to distinguish this model from what might be thought of as a cloud of models surrounding it. That is, the data equally support an infinite number of different models, each with somewhat different implications for macroeconomic policy. How should the Federal Reserve or other macroeconomic policy makers proceed? Should they pick their favored model, engage in some sort of model averaging, or account for their fundamental uncertainty in some other way?

The problem of deep uncertainty is pervasive. In many contexts, we simply do not know what the consequences of our choices will be. Smart people can make guesses based on the best science, data, and models, but they cannot eliminate the uncertainty.

Decision making under conditions of uncertainty or ambiguity is not a new problem. Modern approaches tend to be based on the experiments proposed in Ellsberg (1961). Ellsberg considered two urns, one that contains an equal number, say 50 each, of red and black balls (the risky urn) and one that contains 100 balls that are either red or black but of unknown proportions (the ambiguous urn). There are four possible bets. You can choose between bets 1 and 2:

Bet 1. Draw a ball from the risky urn. If it is black, receive $10, and if it is red, receive nothing.

Bet 2. Draw a ball from the ambiguous urn, with the same payoffs as bet 1.

Many people prefer bet 1. If they choose 1 over 2, they implicitly are

1. The conference, “Developing Regulatory Policy in the Context of Deep Uncertainty: Legal, Economic, and Natural Science Perspectives,” was held April 26 and 27, 2013, at the University of Chicago Law School.

saying that they believe that the probability of a drawing a black ball from the ambiguous urn is less than 50 percent.

You can also choose between bets 3 and 4:

Bet 3. Draw a ball from the risky urn. If it is red, receive $10, and if it is black, receive nothing.

Bet 4. Draw a ball from the ambiguous urn, with the same payoffs as bet 3.

Ellsberg hypothesized that people who prefer bet 1 to bet 2 would also prefer bet 3 to bet 4. Choosing 3 over 4 implies that the probability of a drawing a red ball from the ambiguous urn is less than 50 percent. But the probabilities of drawing a black and drawing a red ball from the ambiguous urn cannot both be less than 50 percent because those are the only two possibilities. They must add to 1. Therefore, choosing both bet 1 and bet 3 creates a paradox. It would seem that these choices have to be inconsistent.

The standard explanation for this paradox is known as ambiguity aversion. Individuals choosing a set of payoffs will demand a higher premium for ambiguous choices than for merely risky ones. In the Ellsberg case, the payoffs are the same, so individuals prefer the risky bet. Only by increasing the payoffs for the ambiguous bet can individuals be made indifferent. The extent of the required additional payoff represents the size of an individual's ambiguity aversion.

Ambiguity aversion is contrary to the standard rational choice criteria for decision making set forth in Savage (1954). In particular, it violates what is known as the sure-thing principle, or the principle of irrelevant alternatives. The sure-thing principle states that events that do not affect payoffs do not affect choices. Savage used the example of the purchase of a piece of property. The purchaser is concerned about the effects of the next presidential election on the value of the property. Suppose he asks himself if he would buy the property if a Democrat wins, and the answer is yes. And suppose he considers whether he would buy the property if a Republican wins, and the answer is also yes. In this case, he should buy the property even though he does not know who will win the next election.

To see why choosing bets 1 and 3 violates the sure-thing principle, consider a variation of the bets described above. Suppose there is a single urn that contains yellow balls as well as red and black balls. There are 90 balls in the urn, 30 of which are yellow balls. The remaining 60 are either red or black, but you do not know how many of each. Consider two sets of bets. The first two are as follows, and you can choose which one you want to take:

Bet 5. Draw a yellow ball and receive $100.

Bet 6. Draw a red ball and receive $100.

Ellsberg hypothesizes that people prefer yellow because it has a known probability, $\frac{1}{3}$, while the chance of red is unknown. It is somewhere between 0 and $\frac{2}{3}$, but you cannot narrow it down more than that. You can also choose between the second two bets:

Bet 7. Draw yellow or black and receive $100.

Bet 8. Draw red or black and receive $100.

Many people, the claim is, prefer bet 8 because it has a known probability of $\frac{2}{3}$ while bet 7 is ambiguous. It has odds of somewhere between $\frac{1}{3}$ and 1, but you do not know more than that.

Note, however, that bets 7 and 8 are identical to bets 5 and 6 with the addition of the black ball. In both cases you get the benefit of the black ball. It is, therefore, an irrelevant alternative, which under the sure-thing principle, should have no effect. If you changed your choices between these two sets of bets, you have violated the sure-thing principle.

Rejection of the sure-thing principle seems, at least to many people, to be irrational. It is, moreover, relatively easy to induce what looks like irrational behavior in experimental settings. Perhaps the Ellsberg result is just another such behavior anomaly. The choice of bets on balls in urns reflects bad decision making, a failure to be rational. We learn nothing more from these experiments but that people are easily fooled.

Ellsberg argued, however, that the choice of bets is not irrational. If the behavior is irrational, people should alter their choices when the true nature of the problem is explained. They should be embarrassed to have made the mistake. When explained that their choices involve a clear error in the urn examples, however, Ellsberg claimed that many would choose to stick with their choices. As Ellsberg (1961, p. 656) noted: "[A]fter rethinking all their 'offending' decisions in the light of [Savage's] axioms, a number of people who are not only sophisticated but reasonable decide that they wish to persist in their choices. This includes people who previously felt a 'first order commitment' to the axioms, many of them surprised and some dismayed to find that they wished, in these situations, to violate the Sure-thing Principle. Since this group included L. J. Savage, when last tested by me (I have been reluctant to try him again), it seems to deserve respectful consideration."

Regardless of whether rational, Ellsberg behavior or ambiguity aversion has been used to explain a number of real-world phenomena. For example, French and Poterba (1991) suggested that ambiguity aversion can explain home bias in investment: investors might feel less comfortable assessing risks in foreign markets than in home markets. Rieger and Wang (2012) used ambiguity aversion to explain the equity risk premium. And Berger, Bleichrodt, and Eeckhoudt (2013) studied treatment decisions by doctors under conditions of ambiguity.

The Ellsberg result has led to a large literature in the economics of decision making. The goal is to find an attractive replacement for the sure-thing principle that produces Ellsberg-like behavior. Gilboa and Schmeidler (1989), in a canonical paper, introduced an approach known as maxmin expected utility. They supposed that we can put bounds on the likely outcomes, but within those bounds, we cannot determine the true probabilities. For example, we may not know the odds of a coin toss because the coin may be weighted, but we might be able to say that the odds heads are between 40 percent and 60 percent. Under maxmin expected utility, individuals choose the policy that maximizes the minimum possible utility within these bounds.

To illustrate,[2] suppose that we do not know the probability of a coin coming up heads because it might be weighted. We can, however, narrow down the possibility to be between 40 percent and 60 percent. Consider the following two bets on that coin:

Bet 9. Pay \$50 and receive \$110 if the coin comes up heads.

Bet 10. Pay \$50 and receive \$110 if the coin comes up tails.

One of those bets should be a good one. If the odds of heads are 50 percent, then each bet has an expected payment of \$55, an expected value of \$5. If the \$5 of expected earnings exceeds the cost of risk aversion, both bets are desirable. In the alternative, if the odds of heads is only 40 percent, the bet on tails is fantastic, and similarly if the odds of heads is 60 percent, the bet on heads is fantastic. No matter the odds, one or both bets should be a good one.

Gilboa and Schmeidler would, nevertheless, reject both bets. The worst case for the first bet is that the odds of heads is 40 percent, so you would expect to lose \$6. The worst case for the second bet is that the odds of heads is 60 percent, so you again would expect to lose \$6. You maximize the worst-case scenario by rejecting the bets so that your worst case is a payoff of \$0 rather than −\$6. Moreover, Gilboa and Schmeidler

2. Paraphrased from the very nice explanation in the blog Less Wrong (Nate Soares, Knightian Uncertainty and Ambiguity Aversion: Motivation, *LessWrong* [blog], July 21, 2014 [http://lesswrong.com/lw/kcl/knightian_uncertainty_and_ambiguity_aversion/]).

reject the claim that rejecting both bets is irrational. Rejecting both bets will systematically lose money, but it will also avoid very bad outcomes, and adopting a strategy to avoid very bad outcomes, they claim, is perfectly rational.

The Gilboa and Schmeidler approach focuses only on the worst case and, therefore, seems extremely averse to bad outcomes when faced with uncertainty. The so-called α-maxmin approach, introduced by Hurwicz (1951) and more recently by Ghirardato, Maccheroni, and Marinacci (2004), balances aversion to bad outcomes with an aversion to losing the benefit of good outcomes. We might think of it as avoiding regret, whether it is regret of getting a very bad outcome that could have been avoided with a different choice or the regret of failing to get a good outcome. It puts a weight α between 0 and 1 on the worst case—the Gilboa and Schmeidler case—and a weight of $(1 - \alpha)$ on the best case. The weight depends on how much a decision maker fears getting bad outcomes and how much the decision maker fears missing good outcomes.

In yet another approach, Hansen and Sargent (2001) adopted theories of robust control from engineering. They assumed that the decision maker has a best guess—they call it an approximating model—of the effects of a choice. It might be, for example, a model of the economy used by the Federal Reserve. The decision maker, however, is unsure that this guess is correct, and another model may in fact be the true one. The decision maker weights each possible model using a measure of its distance from the best guess and then chooses an outcome that maximizes the worst case from within this weighted cloud of models.

Numerous other decision criteria have been proposed. An important class decision criteria is inspired by the smooth-ambiguity approach of Klibanoff, Marinacci, and Mukerji (2005). This approach is able to characterize ambiguity aversion in a way that is similar to the way we characterize risk aversion: decision makers have a risk aversion parameter and an ambiguity aversion parameter that together combine to determine attitudes toward uncertain outcomes.

The papers in this issue are from a conference held at the University of Chicago involving people from a wide variety of disciplines, including lawyers, economists, decision theorists, psychologists, physicists, and climate scientists. Selfishly for the law and legal scholarship, a goal was to see whether there were techniques to borrow. Has anyone else figured this out, and can we use those solutions in law?

Recall that the motivation for much of the literature on ambiguity is the Ellsberg experiments. An initial question is whether the Ellsberg thought experiment is a robust model of actual behavior. At the conference, Stefan Trautmann presented a paper, now published as Trautmann and van de Kuilen (2016), that reviewed the findings.

In the 50 years since Ellsberg's paper, the result has been replicated in laboratory experiments numerous times. Using the urn experiments described above, the result is seen in students, nonstudents, non-Western subjects, children, and even monkeys. Notwithstanding these findings, Trautmann concluded that there are serious problems with the evidence for ambiguity aversion.

The extent and even the existence of ambiguity aversion varies widely with the elicitation method. For example, the elicitation method can be varied by giving the same subjects both the two-urn and three-urn versions of the experiment and compare the within-subject results. The comparison reveals that about 26 percent are risk averse but not ambiguity averse, more than 60 percent behave almost randomly, and only about 12 percent are ambiguity averse. Similarly, as discussed in Trautmann, Vieider, and Wakker (2008), if the elicitation method is designed to avoid peer evaluations and the fear of negative evaluations, ambiguity aversion disappears. And subjects are not ambiguity averse and may even seek ambiguity for low-likelihood events and for losses. In his presentation, Trautmann concluded that ambiguity aversion should not be taken as a universal phenomenon to be built into standard models of choice. Jumping from the Ellsberg results to full-blown models used for major legal or social choices may not be wise.

A second preliminary question before using the Ellsberg results for legal decision making is whether they represent desirable behavior. Even if individuals are robustly ambiguity averse, do we want our agents in the government or elsewhere to also be ambiguity averse? We might be better off if our delegated agents make decisions based on the usual rationality criteria. Do we want, for example, our agents making the sorts of choices described above with the Gilboa and Schmeidler bets, rejecting what seem like clearly winning strategies because of ambiguity aversion?

Ellsberg argued that ambiguity aversion was rational on the basis of the idea that agents, reflecting on the implications of their choices (such as nonadditive probabilities), would not change them. In an extended critique of this conclusion, Al-Najjar and Weinstein (2009) used the same methodology—asking what reflective agents would do—to argue that ambiguity aversion is not rational. Confronted with the implications of ambiguity aversion, they argued that a rational person would reject ambiguity-averse behavior. For example, ambiguity aversion implies a rejection of Bayes's rule. Bayes's rule is nothing more than a restatement of the definition of conditional probabilities. Rejecting Bayes's rule means rejecting the standard definition of conditional probabilities. Statements like "what is the probability that it will snow in Chicago conditional on it being winter" become difficult to analyze, and how one changes views if it is also an El Niño year may not follow what would seem like simple logic. Methods to solve this problem within the framework of ambiguity aversion lead to others, such as aversion to new information.

Nabil Al-Najjar's paper in this issue summarizes and extends the arguments in Al-Najjar and Weinstein (2009). He argues that to the extent that ambiguity aversion requires rejecting Savage's axioms, it should be seen as simply another facet of the numerous problems that humans have making decisions under uncertainty. They are an anomaly, to use the language of behavioral economics. And, therefore, we should demand that our agents, such as government agencies, follow standard Bayesian rationality. We should not want the people we hire to do things for us to replicate our mistakes.

Following the work of Halevy and Feltkamp (2005), Al-Najjar goes on to argue that we can produce what looks like ambiguity-aversion behavior within a standard Bayesian framework, which means that we can explain the results of Ellsberg-type experiments without rejecting the sure-thing principle. Consider the Ellsberg two-color urns, but now suppose that instead of a single bet on the color of the drawn ball, the urn is sampled twice (with the sample ball returned to the urn after the first draw).[3]

For the risky urn, with half red and half black balls, there is a 25 percent chance of getting two reds, a 25 percent chance of two blacks, and a 50 percent chance of a red and a black, in either order. Suppose for the ambiguous urn that the decision maker has a uniform prior belief, giving equal odds to any possible combination of black and red balls. For any probability p of drawing a red ball, the likelihood of drawing two reds is $p2$, drawing two blacks is $(1 - p)2$, and one red and one black is $2p(1 - p)$. As Bayesians, we should add up these probabilities over all possible

3. In most experiments, there is only a single draw. Al-Najjar as well as Helevy and Feltkamp (2005) argue that most experience is with the equivalent of multiple draws and that subjects simply replicate this multiple-draw behavior in the lab even though the lab has only a single draw.

values of p (weighted by the likelihood of that value). With a uniform prior belief, the probability of getting either two blacks or two reds is roughly $\frac{1}{3}$, and getting a black and red in either order is $\frac{1}{3}$.

As can be seen, the bet on the risky urn is strictly better than the bet on the ambiguous urn. Anyone who is risk averse would prefer the risky urn because it has the same expected value with less risk. And this is true without ambiguity aversion. It arises in a strictly Bayesian framework. Therefore, argues Al-Najjar, rejecting the standard Savagian framework not only leads to undesirable behaviors but is also unnecessary to explain the experimental results.

A common, and sensible, response to uncertainty is to try to gather more information, reducing or cabining the extent of uncertainty. The observation that more information is valuable has led to the massive literature on real options and the timing of choices. It has also led to a large literature on the use of experiments by legal institutions, changing laws as the experiments reveal information. Charles Manski's paper in this issue is in this tradition, supporting experimental strategies to reduce uncertainty.

Manski considers regulatory approvals, such as permitting, under conditions of uncertainty. He argues that agencies should randomize approvals. There are two reasons. The first is what he calls adaptive diversification, which uses diversification both to gain the usual benefits of diversification and as a form of experimentation.

Suppose, to start, that the decision of an individual or firm to apply for a permit is exogenous, in that it does not depend on the agency's approval process. In this situation, with risk but not uncertainty, the benefits of diversification are well known. Manski extends the arguments for diversification to uncertainty under the assumption that the agency, facing uncertainty, will want to minimize the maximum possible regret. (Regret is loss in welfare from a choice compared to the choice that would have turned out to be best with full information.) To diversify, the agency will randomize approvals within groups of otherwise observationally identical applicants. This randomization then generates data that can be used as if it were an experiment, allowing the agency to learn and improve its approach. The combination, diversification and using the resulting information to adapt, is what Manski calls adaptive diversification.

Now suppose that the decision to apply for a permit is endogenous, in that the choice depends on the approval process chosen by the agency. In this case, if there is risk but no uncertainty, Manski shows that agencies can use randomization to sort good applications from bad ones. That is, by setting a random approval rate (which also depends on the observed characteristics of the application), the agency can discourage nonmeritorious applications while encouraging meritorious ones.

The more difficult case is where uncertainty and endogenous applications are combined. This problem is formidable. Manski is able to show with some additional assumptions that adaptive diversification still is a desirable strategy.

Kip Viscusi and Richard Zeckhauser's paper in this issue also considers the possibility of learning or adaptation for permit approvals, in their case in the context of the Food and Drug Administration's (FDA's) approval process for pharmaceuticals. They show that the FDA has a strong bias against approving drugs when there is uncertainty about their effects, exhibiting some version of either ambiguity aversion or a bias against losses as compared to gains. As a result, the FDA has a strong bias against approving drugs that are potentially beneficial if there are uncertain harmful effects. Viscusi and Zeckhauser believe that the social losses from this approach are large.

They argue that instead of being averse to uncertainty, agencies should actually prefer uncertainty because it offers opportunities to learn. In particular, they show that a patient, or set of patients, who has to take a drug multiple times should prefer a drug with unknown risks to a drug with known risks. The reason is that the patient (or patients) can use outcomes from early treatments to choose later treatments in much the same way that learning takes place in the two-armed bandit problem.

For example, suppose that the FDA is considering two drugs, one with a known likelihood of success of 50 percent and one with an unknown likelihood of success, and consider a patient who has to take the drug least two times. With the known-likelihood drug, the patient will decide in each period whether taking the drug is desirable. Because there is no learning, the expected benefits are additive. For the drug with an unknown probability of success, after taking the drug the first time, the patient will update his prior beliefs about the likelihood of success and be able to make a choice in the second period that is better informed. Viscusi and Zeckhauser show that the expected value of the drug with the unknown risks is higher because of this potential for learning. They conclude that ambiguity aversion, however deep-seated, is irrational and that, moreover, in environments where learning can take place, we may even want to prefer rather than avoid choices with uncertain probabilities.

Regulatory environments are often sufficiently complex that agencies resort to computational models to perform their cost-benefit analyses. Even with sophisticated models, however, the results may be uncertain (in the sense of ambiguous as opposed to merely risky). One approach to uncertainty in this context is to build better models, so that their predictions are more reliable, reducing the extent of uncertainty. A common approach to building better models is to add more detail in the hopes that better fidelity to the modeled environment will produce better estimates of the effects of a policy.

Alan Sanstad, in his contribution, considers this approach in the context of energy models. He focuses on the Integrated Planning Model, which is used by the Environmental Protection Agency (EPA) for its cost-benefit analysis. The IPM is a large model: it has detail on more than 15,000 power plants, about 2 million decision variables, and on the order of 200,000 constraints. It was not developed casually. The EPA has a Council on Regulatory Environmental Modeling to improve its modeling practices. Its procedures have been reviewed by the National Research Council and the resulting recommendations taken into account. The EPA held workshops with invited experts to improve IPM's representation of important sectors such as its coal and natural gas supply assumptions. The EPA is pursuing best practices. If adding detail to models is going to work to reduce uncertainty, the EPA is a good test case.

Intuitively, greater detail should help improve the predictive accuracy or other useful outcomes of a model. A better representation of the modeled environment should produce more accurate outcomes. Unfortunately, it does not, Sanstad concludes. His conclusion, to a great extent, relies on 30 years of experience in energy modeling and the failure of those models to narrow the range of uncertainty in their predictions, notwithstanding massive efforts.

In a nutshell, here is his explanation for this failure: For a model to be reliable, it must be validated against data and the extent of model uncertainty quantified. There is, however, no apparent way to validate models of this sort. As a result, model builders rely on calibration. What this means is that they use a model structure that is based on economic theory and calibrate it to existing data. The model is then run into the future or with a chosen policy to evaluate possible outcomes. The results, however, are then a prediction of the theory, not a prediction of what will actually happen if the policy is imposed. While the outcomes are informative in the sense that they help us understand the implications of a given economic theory, they are not informative in the sense that they tell us what the actual effects of a policy will be.

Adding detail to a model that takes a calibrationist approach does not necessarily improve accuracy. For example, IPM has representations of 15,000 power plants, such as information about their capacity, thermal efficiency, and cost. Using economic theory, the model assumes that they will be run in a cost-minimizing manner, so that electricity is produced in the most efficient manner given the existing facilities. Because the operation of the plants does not conform to this theory, however, this assumption means that the output of the power plants in the model do not conform to the actual output of those plants. Running this model into the future with and without a policy will not necessarily tell us what the real-world effects of the policy will be.

Arguments such as this aside, the proof is in the pudding. Are these models able to make accurate assessments? Sanstad shows that they cannot: IPM itself produced highly inaccurate predictions of CO2 emissions from the electric power sector because it did not, and really could not, predict the strong shift to natural gas due to fracking. Of course, one might say, fracking was basically impossible to predict, but that is part of the point. Similarly, Stanford's Energy Modeling Forum does large-scale model comparisons in which the very best energy models are given the same data and told to run the same simulations. Sanstad considers an example in which the models were asked what carbon price would be needed to produce a given percentage reduction in US CO2 emissions. The answers varied by a factor of five, which effectively means that the models are not telling us anything.

If Manski, Viscusi and Zeckhauser, and Sanstad address ways that we might reduce uncertainty, the last three papers, by William Brock and Steven Durlauf, Vermeule, and Daniel Farber, address what we might do in the face of irresolvable uncertainty.

Brock and Durlauf take the perspective of an expert giving advice to an agency head or other decision maker. Their key recommendation is the use of what they call value dispersion plots.

They consider four different levels of uncertainty. In the first, which we might think of as a standard modeling environment, the analyst has a model, a set of observable variables and parameters, unobserved variables, and a policy that is being modeled. If the unobserved variables can be estimated conditional on the observed variables, the analyst can then produce an expected outcome in the usual fashion. This case is a baseline, the sort of decision making with which everyone is familiar.

Turn now to their second case. Suppose that there are a number possible ways to model a problem and that we do not know which model is best. If analysts have reasonable informed views on the likelihood of each model being correct, they can use what is called Bayesian model averaging, which effectively weights outcomes from each model using its likelihood to produce an overall distribution of possible outcomes. The variance in reported outcomes will in general be larger because of model uncertainty, but the resulting probability distribution is still of the sort with which we are familiar.

In the third case, suppose that we do not have enough data to construct a posterior model probability so that all we have are uninformed prior beliefs about the usefulness of any given model. At this point, we face irresolvable uncertainty, and model averaging is not helpful. Brock and Durlauf consider and reject minmax expected utility and minimax regret approaches because they view the rejection of the sure-thing principle as irrational. Instead, they recommend that analysts provide policy makers with plots of possible outcomes, their value dispersion plots. These plots are histograms of possible outcomes from a policy, weighting the output of each possible model equally. (For an interesting example involving the model uncertainty regarding the deterrent effects of the death penalty, see Durlauf, Fu, and Navarro [2012].) The argument for this modest approach is that analysts should not attempt to convey more than they know or embed axiomatic assumptions about the proper decision criteria in the information that they present.

Finally, Brock and Durlauf consider environments where we face what they call radical uncertainty, where we truly lack information. They suggest here that value dispersion plots may still be useful. In addition, they suggest that analysts might have enough information to detect early warning signals of impeding bad outcomes.

Vermeule, in his contribution, considers how a reviewing court should evaluate a decision by an agency made under conditions of uncertainty. Among other examples, he considers the case of the flat-tailed horned lizard, discussed above. Suppose that the FWS does not have a valid method of estimating the population of flat-tailed horned lizards or whether the population has been increasing or declining. It must, nevertheless, decide whether to list the lizard as endangered. It cannot avoid a decision because not listing is also a decision. As Vermeule puts it, the FWS has no first-order principle for making a decision because it lacks information about the likelihood that the lizard is endangered, but it has a second-order reason, that is, a reason to make a choice from within the feasible set.

Suppose for simplicity that the population level can either be high or low. If the population level is high, the FWS should not list the lizard, and if it is low, it should. Also suppose, as actually happened, that the FWS does not list the lizard, which means that the FWS was acting as if the population level was high. A court could ask whether there are grounds for this choice or whether it was arbitrary. The FWS would not be able to point to first-order reasons for the decision because it does not have information about whether the population level was actually high or low. Under these conditions, choosing "high" seems arbitrary. A reviewing court rejected the FWS's decision on precisely these grounds.

Suppose, however, that the FWS chooses to list the lizard, acting as if the population level were low. This choice would also have been arbitrary using the same reasoning. Vermeule concludes that although the choice to list is arbitrary in a statistical sense, it should not be treated as arbitrary in the legal sense of a decision that is without support. The FWS had to make a choice. It lacked the necessary information. So it picked. Vermeule argues that reviewing courts should allow agencies to make arbitrary decisions under conditions of uncertainty. Doing so simply recognizes the epistemic uncertainty that agencies often face.

Farber, in the final paper in this issue, considers how the problem of uncertainty about future outcomes affects discount rates for policies with long-term effects. For example, determining what to do about climate change requires discounting future harms and future costs of mitigation over hundreds of years. We have little idea what the likely harms and costs will be over this sort of time scale. How do decide what to do? Farber focuses on a particular aspect of this choice, the discount rate. Because of the mathematics of discounting over long periods, small changes to the discount rate can have massive effects on what sort of polices we choose.

Following an argument developed in Weitzman (1998), Farber argues that uncertainty in future rates of growth should lead to a declining and ultimately very low discount rate. To see why, consider an estimate of the future effects of a policy that produces a future benefit of $100 million in 200 years and suppose that there are two possible discount rates, 6 percent and 2 percent, because of uncertainty about the economic environment. The resulting implied discount rate that should be used to evaluate the present value of the policy is not the simple average of 6 percent and percent, or 4 percent. Instead, the proper procedure is to take the present value of the project being evaluated in each scenario—the 2 percent and the 6 percent scenarios—and average those values. Taking this average gives an expected present value of just under $1 million. The implied discount rate—the rate that gives $100 million in 200 years a present value of $1 million—is 2.35 percent.

This effect, that lower discount rates count more in the averaging, gets bigger as the time periods get longer and as uncertainty increases. As a result, uncertainty in growth rates should lead us to value the future more, effectively as if we are buying insurance against the possibility of a bad future outcome.

This result may not change very much if there is ambiguity or uncertainty about what the future holds rather than just risk (as there surely is). Correctly averaging discount rates already pushes discount rates toward low values. That is, with risk rather than uncertainty, we already weight the bad cases more heavily. Aversion to ambiguity would not, under many calculations, cause us to weight those cases very much more. If this is true, then one of the biggest sources of uncertainty with respect to very long-term issues such as climate change—what economic growth will be in the distant future and what the implied discount rate will be—may not affect our choices very much. We will still have a very hard time estimating likely scenarios and putting weights on them, but the problems of choice under ambiguity rather than under risk would not be a significant complicating factor.

REFERENCES

Al-Najjar, Nabil I., and Jonathan Weinstein. 2009. The Ambiguity Aversion Literature: A Critical Assessment. *Economics and Philosophy* 25:249–84.

Berger, Loïc, Han Bleichrodt, and Louis Eeckhoudt. 2013. Treatment Decisions under Ambiguity. *Journal of Health Economics* 32:559–69.

Durlauf, Steven N., Chao Fu, and Salvador Navarro. 2012. Capital Punishment and Deterrence: Understanding Disparate Results. *Journal of Quantitative Criminology* 29:103–21.

Ellsberg, Daniel. 1961. Risk, Ambiguity, and the Savage Axioms. *Quarterly Journal of Economics* 75:643–69.

French, Kenneth R., and James M. Poterba. 1991. Investor Diversification and International Equity Markets. *American Economic Review* 81:222–26.

Ghirardato, Paolo, Fabio Maccheroni, and Massimo Marinacci. 2004. Differentiating Ambiguity and Ambiguity Attitude. *Journal of Economic Theory* 118:133–73.

Gilboa, Itzhak, and David Schmeidler. 1989. Maxmin Expected Utility with a Non-unique Prior. *Journal of Mathematical Economics* 18:141–53.

Greenstone, Michael, Elizabeth Kopits, and Ann Wolverton. 2013. Developing a Social Cost of Carbon for US Regulatory Analysis: A Methodology and Interpretation. *Review of Environmental Economics and Policy* 7:23–46.

Halevy, Yoram, and Vincent Feltkamp. 2005. A Bayesian Approach to Uncertainty Aversion. *Review of Economic Studies* 72:449–66.

Hansen, Lars Peter, and Thomas J. Sargent. 2001. Robust Control and Model Uncertainty. *American Economic Review* 91:60–66.

Hurwicz, Leonid. 1951. Some Specification Problems and Applications to Econometric Models. *Econometrica* 19:343–44.

Keynes, J. M. 1937. The General Theory of Employment. *Quarterly Journal of Economics* 51:209–23.

Klibanoff, Peter, Massimo Marinacci, and Sujoy Mukerji. 2005. A Smooth Model of Decision Making under Ambiguity. *Econometrica* 73:1849–92.

Rieger, Marc Oliver, and Mei Wang. 2012. Can Ambiguity Aversion Solve the Equity Premium Puzzle? Survey Evidence from International Data. *Finance Research Letters* 9:63–72.

Savage, Leonard. 1954. *The Foundations of Statistics*. New York: Wiley.

Sunstein, Cass R. 2002. *The Cost-Benefit State: The Future of Regulatory Protection*. Chicago: American Bar Association.

Trautmann, Stefan T., Ferdinand M. Vieider, and Peter P. Wakker. 2008. Causes of Ambiguity Aversion: Known versus Unknown Preferences. *Journal of Risk and Uncertainty* 36:225–43.

Trautmann, Stefan, and Gijs van de Kuilen. 2016. Ambiguity Attitudes. Pp. 89–116 in *The Wiley Blackwell Handbook of Judgment and Decision Making,* edited by Gideon Keren and George Wu. Malden, MA: John Wiley & Sons.

Weitzman, Martin L. 1998. Why the Far-Distant Future Should Be Discounted at Its Lowest Possible Rate. *Journal of Environmental Economics and Management* 36:201–8.